IMPERIAL COLLEGE LONDON

DEPARTMENT OF COMPUTING

Modelling Spatial Constraints on Brain Connectivity

Author:
Maxwell POPESCU

Supervisor:
Prof. Murray SHANAHAN

June 17, 2014

Acknowledgements

I would like to thank my supervisor, Professor Murray Shanahan for providing me with support and advice throughout the course of this project. I would also like to thank my second marker, Professor Abbas Edalat, for taking his time to provide me with feedback on my Interim Report. Finally, I am grateful to my family and friends for their continued support during my time at university.

Abstract

In the mammalian brain, there is a large volume of white matter dedicated to long-range wiring, separated from the neuron cell bodies and short-range wiring. In the avian brain, no such dedicated volume exists. Assuming that evolution has maximised brain functionality, what are the advantages of the mammalian arrangement from an engineering point of view?

To find possible answers to the question, this project considers the implications of changing the spatial arrangements of networks on a two-dimensional computer model. As a basis for the research, we test results from previous analytical studies. We find that, when a network has high local connectivity and sparse global connectivity, the mammalian arrangement can potentially reduce the total length of wiring. We then draw on percolation theory to consider the possibility that separating long-range fibers enables neuron cell bodies to be more densely packed.

Contents

Chapter 1

Introduction

1.1 Motivation

In vertebrates, the brain is formed of two types of tissue: grey matter and white matter. The grey matter consists of local neuronal networks, with non-myelinated axons forming the local connections. The white matter enables global communication across the brain by providing myelinated, long range connections. In the mammalian cortex, there is a clear segregation between the two types of tissue, with the grey matter circumferentially surrounding the white matter. This can be contrasted with the avian brain, which has no such large dedicated volume of white matter, and instead has white matter fascicles scattered throughout the whole forebrain [31].

From an evolutionary point of view, what is the rationale behind the segregation of grey and white matter? The exact reason is still unknown, but one study [24] reveals that a segregated cortical architecture potentially minimises total wiring volume, suggesting that the brain has had to adapt to the spatial constraints imposed by the cranium. Using the same model of connectivity adopted by this study, Wen and Chklovskii [31] instead sought to find whether minimisation of conduction delays - the time taken for a neuron to transmit a signal to another - could provide another explanation. Indeed, under certain parameters, segregating grey and white matter appears to lead to reductions in conduction delay.

The previous studies use an analytical approach, where equations for measures such as volume and delay are defined formally. The advantage of this is that we can abstract the extreme complexity of the brain and produce meaningful results that scale to any number of neurons. However, this of course comes with the cost of abstraction. For example, the volume taken up by wiring is not the same as the volume that is needed as an enclosure, since wiring will not be packed infinitely close together.

Therefore, it seems worthwhile to study the problem using a synthetic approach, where much smaller scale, but concrete, computer models are used. The applicability of previous studies could be tested on these models, while alternative explanations for grey and white matter segregation could be discovered.

1.2 Project Aim

This project sets out to seek possible explanations for the grey and white matter segregation that is observed in the brains of vertebrates, using a synthetic approach. Being a fundamentally open-ended, research-based project, the main aim is to establish an area of interest that could be used as a basis for further research.

1.3 Contributions

1. A concrete implementation of a visualised network model with predefined spatial constraints.

2. Results which show that in this model implementation and under certain circumstances, a segregated architecture can lead to reduction of wiring length while maintaining the same (and potentially higher) connectivity.

3. Results which shows how the density of network components affects the chances that a connection can be spatially embedded.

4. The proposal that grey and white matter segregation is possibly caused by the drive to increase neuron cell body density while maintaining very long-distance connections.

1.4 Report Structure

We begin the report by presenting the reader with a background in the human brain followed by an overview of related work and discussion of key terminology (Chapter 2). We then outline our methodology and provide details on the network model used in the project (Chapter 3). Some initial observations and minimisation of conduction delays is then discussed (Chapter 4). We explore percolation theory (Chapter 5) and use it as a basis for explaining some previous results (Chapter 6). The relevance of our findings to the brain

is then considered, and the challenges of the project are evaluated (Chapter 7). We then conclude the report, outlining some possible areas of further work (Chapter 8).

Chapter 2

Background

In this chapter, we first present the reader with a background of the human brain that they should familiarise themselves with before progressing through the rest of the report. We then examine previous studies which bear relevance to this project. Finally, we take a look into network theory and the characteristics that are typical of high functioning networks, and consider various models of networks that are used in the project.

2.1 The Human Brain

2.1.1 Neurons

Neurons are the basic data processing units of the brain, which connect with each other to form neural networks. Each neuron receives electrical inputs from thousands of other neurons (on average), and if the sum of impulses arriving simultaneously together is strong enough, the neuron will generate an electrical signal (an 'action potential') which becomes the input to other neurons in the network.

The structure of a typical neuron can be broken down into three distinct elements: A neuronal cell body (soma), dendrites, and an axon. The neuronal cell body is responsible for generating and processing electrical signals; the dendrites act to receive incoming signals and propagate them to the cell body; and the axon acts to conduct signals away from the cell body. Axons can be highly branched, meaning an action potential is received by many other neurons. The axon of one neuron and the dendrite of another are not physically in direct contact, but instead the electrical signal travels across a small gap known as a synapse.

Each connection between neurons therefore comprises of two parts. The dendrites are restricted to small regions around the cell body, whereas axons can be much longer. Axons

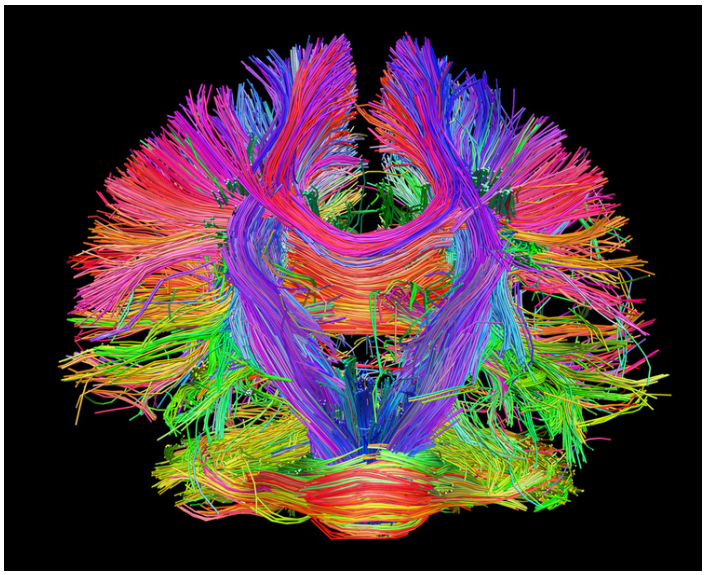

Fig. 2.1 White Matter fibers, HCP Dataset. Source: [23]

can be very short, but also very long, for example those that run outward to the body from the brain. In order to keep the conduction delays across long range axons reasonable, axons can be coated in a fatty substance known as myelin, which significantly increases the conduction velocity.

2.1.2 Grey and White Matter

Within the central nervous system (CNS), it is possible to distinguish regions of different shading with the naked eye on unstained cuts through the brain [25]. *Grey matter*, as the name suggests, is grey in appearance and consists primarily of neuronal cell bodies, unmyelinated axons, dendrites and glial cells. This region is where the neurons interact. The *white matter*, on the other hand, can be thought of as a cable system [25] which enables communication between different regions of grey matter. It contains axons and no other neuronal structures, many of which are myelinated, giving the region a white colouring.

The Human Connectome Project is a project that aims to map the brain's fibrous long distance connections using diffusion MRI [14]. As shown in Fig. 2.1, white matter tracts travel alongside each other in an orderly fashion, making long distance connections between different regions of the brain.

2.1.3 Axon Guidance

In neural development, the axons of neurons are found to reach target cells by following a series of intermediate targets, consisting of glial cells or intermediate guidepost cells [6]. An axon owes its ability to follow these guidance cues to its growth cone, which sniffs out the extracellular environment that instruct the axon which way to grow.

It is observed that developing axons prefer to grow along pre-existing axonal tracts [13]. Therefore, the first axons that grow in a largely axon free environment are termed pioneer axons, establishing the scaffolding that lays down the path for future follower axons [6].

2.1.4 Complexity of the Brain

The human brain is an extremely complex organ, with one estimate placing the number of neurons at 86 billion [3], and others with even larger numbers. The number of connections is even less precise, but one article suggests there are 125 trillion synapses in the cerebral cortex alone [12]. Furthermore, the brain is dynamic: many connections are established and pruned every day [21]. As a result, the availability of data on human neural connectivity is extremely limited.

2.2 Related Work

2.2.1 The Economy of Brain Network Organization

Ed Bullmore and Olaf Sporns [4] suggest that many aspects of the brain's organisation can be explained by a drive to minimise the high material and metabolic costs that it incurs. One major source of this cost comes in the form of wiring: It consumes volume, comes at metabolic expense, produces signal delay and attenuation and incurs a developmental overhead. It follows that reducing wiring cost confers an evolutionary advantage, and leads to a hypothesis referred to as the 'wiring economy principle', that claims wiring minimization influences brain organisation [7].

However, wiring minimisation is not the only criteria. If it was, the network would adopt a lattice-like topology where nodes would simply connect with their (spatially) nearest neighbours. This does not favour global integration of information processing, and therefore the efficiency is low. On the other hand, a random topology provides good integrative processing, but comes at too high a wiring cost, due to the large number of long-distance connections. Instead, the brain appears to adopt a middle ground, where there are clusters

of highly intra-connected nodes, but also high cost components which act as short-cuts between different brain regions.

Furthermore, a study by Marcus Kaiser and Claus C Hilgetag [18] demonstrates that neuronal components (node arrangements) are also not optimally placed for wiring minimisation. By searching for alternative node arrangements given existing datasets and spatial coordinates of the nodes, the authors were able to decrease total wiring among 95 primate cortical areas by 32% using a simulated annealing approach.

Overall, it appears that while wiring cost is an important consideration, the brain organisation involves an economic trade-off between minimising costs and enabling efficient connectivity across the brain.

2.2.2 Volume Efficiency of Cortical Architectures

Ruppin et al.[2] produced a possible explanation for grey and white matter segregation in terms of volume minimisation, by calculating the brain volume based on geometric characteristics. They considered two kinds of possible wiring organisations: First, the Homogeneously Dispersed (HD) organisation, where neurons and local connections (grey matter) are intermixed with long range connections (white matter); second, the Segregated (SG) organisation, where white matter and grey matter are found in two spatially distinct areas. The volume occupied by the two organisations are calculated as follows:

$$V(HD) = N[u + n_s f_s t_s(HD) + n_l f_l t_l(HD)]$$

$$V(SG) = N[u + n_s f_s t_s(SG) + n_l f_l t_l(SG)]$$

where N is the number of processors (neurons), u is the volume of a processor, (n_s, n_l) are the average number of (short and long range, respectively) connections per processor, (f_s, f_l) are average cross-sectional areas of connections, and (t_s, t_l) are average connection lengths.

For a given volume, Ruppin et al. notes that a SG organisation will have a lower value of t_s, since neurons have to be more densely packed. Therefore, $V(SG) < V(HD)$ if:

$$\frac{n_s}{n_l} > \frac{f_l(t_l(SG) - t_l(HD))}{f_s(t_s(HD) - t_s(SG))}$$

By modelling values for the average connectivity lengths in each type of organisation, the authors establish the parameters in order for one organisation to be more volume-efficient than the other.

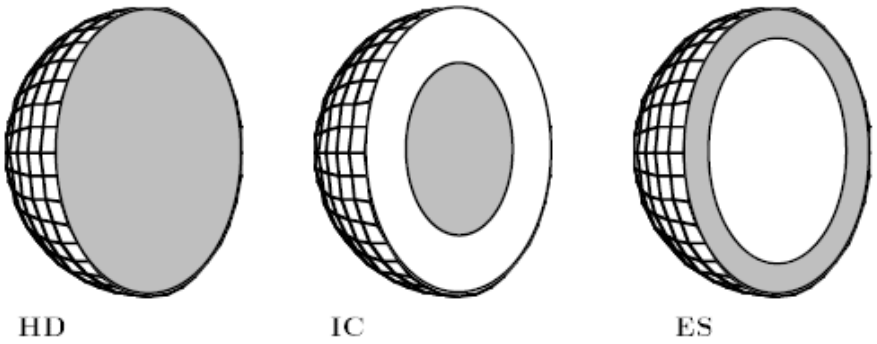

Fig. 2.2 Homogeneously Dispersed, Inner Core and External Sheet organisations. Source: [24]

An interesting point to take away from the study is how the SG organisation itself can be split into two categories, as shown in Fig. 2.2: The Inner Core (IC) organisation, where white matter surrounds the grey matter, and the External Sheet (ES) organisation where grey matter surrounds white matter. The latter is the organisation adopted in the mammalian cortex, and the paper argues that the ES organisation is more economical given available cortical data.

2.2.3 Wiring Efficiencies at different scales

Murre and Sturdy [22] also analysed grey and white matter segregation in terms of volume minimisation, but instead using a scaling approach. Their calculations suggest that depending on the scale being considered, different designs are more efficient. Most relevant to this project are their findings at the macro-structural level, where they suggest packing neurons on the surface of the exterior (grey matter) and packing fibers interiorly is most efficient in terms of volume.

In their model, it is assumed nodes are arranged neatly in a lattice, and they simplify the branching model by proposing a 'snaking' connection strategy, where each neuron snakes along all its target neurons, as shown in Fig. 2.3. They note that if neurons are not distributed neatly in the lattice, then optimization strategies could be used to reduce the path length of snaking along neurons. However, since this is the Travelling Salesmen Problem, which is NP-Complete and hence impossible for neurons to solve, they do not consider potentially suboptimal paths to be a problem for their results.

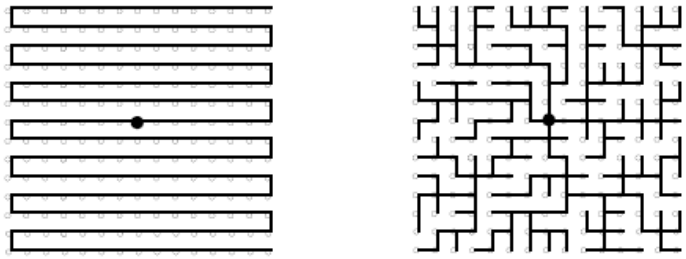

(a) (b)

Fig. 2.3 Branching can be reduced to a neuron 'snaking' along target neurons in a neatly arranged lattice. Source: [22]

2.2.4 Segregated Design and Conduction Delay minimisation

A more recent article by Wen and Chklovskii [31] analyses the advantages of grey and white matter segregation, not in terms of volume, but from the perspective of conduction delays. Using the same connectivity model used by Ruppin et al [24] and Murre and Sturdy [22], they determine parameters for which the segregated design has smaller conduction delays than the homogeneous design. Of particular interest was their appreciation for the different properties of local and global connections: The latter can generally be made to conduct faster since its conduction velocity scales linearly with wire diameter whereas the former can not. Some key points from the paper are:

1. Higher connectivity comes at the cost of longer conduction delays, since more wiring means increased volume and increased distances between neurons. This means connections are longer and hence conduction delays are higher.

2. Introducing global connections into local circuits causes swelling of the volume and local connections become longer, increasing local conduction delay.

3. Introducing local circuits to global connections will also cause a swelling of the volume, and hence global connections are longer. However, this can be compensated for by increasing wire diameter.

4. Therefore, from the perspective of reducing conduction delays, it can be advantageous to reduce the impact of global connections on local conduction delay. This can be achieved by separating it from the grey matter.

Given the studies by Ruppin et. al, Murre and Sturdy and Wen and Chklovskii [22, 24, 31], it seems the brain can be seen as efficient from two perspectives: Volume and conduction delay. However, the two have not yet been related to each other and a unified framework is yet to be achieved.

2.3 Networks

2.3.1 Network Terminology

- A network $G =< V, E >$ has a set V of nodes and a set $E \subseteq < V, V >$ of connections between nodes. E may also be represented by a two-dimensional connectivity matrix A, such that, for all $i, j \in V$:

$$A(i,j) = \begin{cases} 1 \text{ if } (j,i) \in E \\ 0 \text{ otherwise.} \end{cases}$$

- The degree k_i of a node i is the number of edges it participates in.
- The *path length* between nodes $i \in V$ and $j \in V$ is the number of edges in the shortest path between the two nodes.

2.3.2 Small-World Networks

A network G is said to be 'small-world' if the following properties hold:

1. Sparse: On average, each node is connected to only a very small fraction of the total number of nodes.
2. High clustering: If $A(x,y) = 1$ and $A(x,z) = 1$, then there is a high probability that $A(y,z) = 1$. The extent of clustering is characterised by the network's clustering coefficient, γ_g.
3. Low *mean path length*, λ_G: The *path length* averaged over every distinct pair of nodes in V is small.

We can quantify how small-world a network is using the small-world index, σ_G:

$$\sigma_G = \frac{\sigma_G / \gamma_{rand}}{\lambda_G / \lambda_{rand}}$$

where γ_{rand} and λ_{rand} are a random network's clustering coefficient and mean path length, respectively.

In nature, high connectivity and short conduction delays are competing requirements [31], but the small-world properties of the brain mean the brain can maintain high functionality by having fast global communication despite the sparse connectivity. Small-world networks are found to strike a good balance between local and global *efficiency*, as defined by Latora and Marchiori [20], which characterises how efficiently information can be propagated on a local and global scale.

2.3.3 Network Modularity

A modular network is one where its nodes can be partitioned into modules. Each module is highly intra-connected, but modules themselves are sparsely inter-connected. The modularity for a graph G with m edges can be measured using Q:

$$Q = \frac{1}{2m}\sum_{i,j}(A(i,j) - \frac{k_i k_j}{2m})\delta_{c_i c_j}$$

$$\text{where } \delta_{xy} = \begin{cases} 1 \text{ if } x = y \\ 0 \text{ otherwise.} \end{cases}$$

In this formula, the sum of $\frac{A_{i,j}}{2m}$ represents the actual fraction of intra-community edges whereas the fraction $\frac{k_i k_j}{2m}$ is what we would expect in a comparable random network. c_i denotes the community node i belongs to. Therefore, the value of Q represents gives an idea of how densely connected communities are compared to what we would expect given the degrees of nodes. Note that the community structure is not necessarily known in advance, in which case it is instead determined by the community structure that maximises Q.

2.3.4 Network Topologies

We outline a few network topologies which can be established using an established procedure:

Fully Connected

In a fully connected network, each node $v \in V$ is connected to every other node in the network. Assuming the network is undirected, it will contain $n(n-1)/2$ edges, and each node will have degree $n-1$.

Random

In a random network, every possible edge occurs independently with equal probability $0 < p < 1$. This is the model proposed by Edgar Gilbert [11].

Modular

Three of the previous studies mentioned [22, 24, 31] propose a 'modular' topology, where each node connects with its k-nearest neighbours, and makes an arbitrary long-distance connection to another local network in the brain. For the sake of simplicity, this long-distance connection is made with a randomly picked node.

Distance-based

We can produce a spatially embedded network that has a high small world index and medium to high modularity [27] based on the distances between different nodes with the following algorithm:

1. Create a network with n nodes, and assign each node random coordinates in the range 0 to 0.5, to ensure the distance between any two nodes lies in the range 0 to $\frac{\sqrt{2}}{2}$.
2. Consider every pair of distinct nodes and connect them with probability p:

$$p = e^{-hd}$$

where h is a constant and d is the Euclidean distance between the two nodes. This results in a connectivity pattern than favours mostly short connections while still producing some longer range connections. When $h = 0$, the graph is random, while increasing h produces relatively high local connectivity, as shown in Fig. 2.4.

2.3.5 Physical Models

The background papers present various physical 'packing strategies' of networks, which defines the placement of nodes and local connections (grey matter) and long-distance fibers. These abbreviations will be used in the rest of the report.

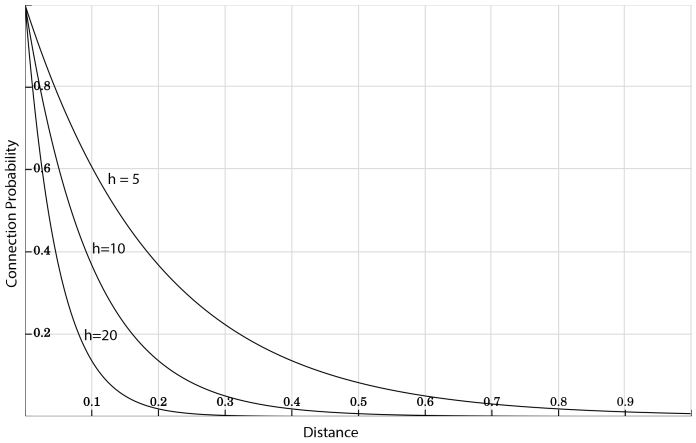

Fig. 2.4 Effect of varying h in the distance-based spatially embedded topology.

HD and SD

First, we define two high level packing strategies. In a homogeneous design (HD), local and global connections are uniformly and finely intermixed [31]. In a segregated design (SD), local connections are separated from global connections.

Forms of SD

SD itself can be subdivided into a number of different designs, depending on how the grey matter is positioned relative to the white matter:

- **Horizontally Sheeted**: We consider a laminar layer of grey matter, which exists at the top of the area with global connections running in the region below.
- **Inner Core (IC)**: As defined by Ruppin et al. [24], in IC the grey matter is found surrounded by an outer layer of white matter.
- **External Sheet (ES)**: ES is the mirror image of IC, where the grey matter is found on the external surface of the area.

Chapter 3

Modelling the Problem

3.1 Project Focus

Being an open-ended project, the initial problem was to establish its focus. The logical topology of a neural network is largely independent of its physical characteristics. Therefore, the mystery behind grey and white matter segregation appeared to be a physical concern. Two areas were considered: (a) The spatial efficiency of HD and SD, given certain parameters. (b) The implications of segregating grey and white matter on the difficulty of growth cone pathfinding. The latter proposition essentially asks: Does a more structured layout of nodes and wiring reduce the difficulty of finding their destination?

Although (b) is an interesting question, little is known about how neurons decide their target neurons, and whether or not the connections are even targetted, or based on chance encounters. There is a suggestion that this depends on the type of neuron in question [5]. This project does not concern itself with the difficulties a growth cone may encounter as it tries to establish connections, and simply distinguishes between connections which are possible or impossible to establish. It is therefore focused on (a).

3.2 Methodology

In order to gain an insight into the spatial efficiencies of HD and SD, two general approaches were considered. First, an analytical method that would represent the problem in the form of general formulae, and use mathematical reasoning to make deductions. Second, a synthetic method that would involve running computer simulations on concrete models to make inferences. The related work presented in the background was analytical in nature, producing

mathematically sound results at the cost of a high degree of abstraction.

The synthetic method was more attractive and thus chosen for the following reason: The existence of studies concerning the effects of segregating matter on wiring lengths, brain volume and conduction delays would provide a basis for testing the results on a concrete computer model; at the same time computer simulation could potentially become a source of inspiration for new theory by exposing complex behaviours otherwise lost through abstraction.

3.3 Establishing a Suitable Model

The first challenge was to establish a model that could be used to perform various experiments that would distinguish the properties of a homogeneous neural network architecture, where nodes and both local and global connections are intermixed, from a segregated architecture that sees global connections primarily be routed through a region separated from the nodes and local connections.

Given the inherent scale and complexity of the brain, creating a model which closely resembled its architecture was not only infeasible, but beyond the scope of a project that aims to produce speculative theories as a basis for further study. Instead, a simplified model that benefits from faster development and computational efficiency, while maintaining relevance to the network of the brain, is used. The key features of the model that were ultimately decided upon, and the reasoning behind them, are as follows:

Spatial Constraints

Since the research is concerned with the layout and spatial properties of the network, the nodes and edges of the graph must have some notion of location and occupy a quantity of finite space. In a brain, the space occupied by both neuron cell bodies and their connective components (axons, dendrites and synaptic clefts) is significant, so the model could not completely ignore either. However, realistic values of size are substituted in favour of artificial values. The nodes, which represent neuron cell bodies, are all of the same size. The edges of the graph represent all the connective components involved in a single connection, and have variable size.

2-dimensional Space

One major consideration was whether to work in a 2-dimensional space or a 3-dimensional space. The latter is more biologically realistic, but poses a significantly more difficult challenge. The increase in dimension would increase the computational cost of operations relevant to the project, such as collision detection and pathfinding, while visualisation of the model would require more development time and the use of more sophisticated software. By opting for a 2-dimensional model, these costs are greatly reduced, and findings remain relevant to 3-dimensional models.

Grid-based layout

A grid-based layout simplifies useful operations such as node or edge indexing and the higher degree of detail a real coordinate space would provide was deemed unnecessary for this project.

Edges can cross

The edges must consume some quantity of the finite space available. If edges were allowed to travel on top of one another, while some abstract quantity could still be assigned, the notion of connections becoming obstacles to one another would be lost completely. On the other extreme, if nodes are prohibited from crossing completely, a planar graph problem emerges, severely restricting connection possibilities. This would also be biologically unrealistic, so the middle ground is chosen, where nodes can cross but not overlap along the same direction.

Edges are limited to orthogonal movement

Movement on a square grid is limited to either completely orthogonal movement, or movement which permits diagonals. It was felt that limiting the model to orthogonal movement would still produce useful results, while reducing search spaces and avoiding a very minor problem introduced by allowing diagonal movement: When two nodes are next to each other diagonally, the gap between them is larger than the gap between two non-diagonally adjacent nodes. Should edges be allowed to move through this larger gap?

No branching connections

In biological neural networks, axons can branch, enabling each neuron to potentially establish connections with hundreds of other neurons. Initially, a branching design for the model was considered. As each node occupies a single square cell on the grid and edges travel orthogonally, each node can only have four terminals. Branching on outgoing edges is possible, but nodes cannot 'accept' more than four of these edges since it would have to involve the overlapping of separate edges. This cap on the maximum node degree would severely reduce the significance of branching, so such behaviour is not present in the model.

Undirected

The research is mostly concerned with the physical properties of the graph, as opposed to the logical topology. The question of whether edges are directed is of no significance in this regard, but the model assumes the graph is undirected when logical topology is considered.

3.4 Model Implementation

Initially, models were built using MATLAB which offers useful tools for network analysis, matrix manipulations which are suited to common operations on networks, and built-in visualization and graph plotting tools. While these models were useful for prototyping and preliminary analysis, it quickly became difficult to manage the complexity of them, and making low level changes was challenging. Furthermore, the network libraries were found to be less relevant to the project as the project became more focused on the spatial and physical characteristics of the network, as opposed to the logical characteristics.

The model that was eventually decided upon, as described in the previous section, is written in Java. An Object-Oriented approach was suited to the composition of components in the model, such as the grid, nodes and wiring. Visualization is done using Java Swing, graph plotting is done using JFreeChart [17], while analysis of the logical topology of networks are generally left to third-party applications that can import network graph data in standardised formats (e.g. CSV).

3.4.1 Low-Level Components

The three key low-level components of the model are as follows:

Grid The grid is initialised as a fixed size two-dimensional array of object type Cell, that represents the area available to the network. Each grid cell can contain objects of type CellObject. Network components are successively added to the grid, provided the operation does not violate any of the spatial constraints.

Node A Node is a CellObject which represents a neuron cell body. Each node in a network has a unique identifier, and contains relevant information such as its position and degree.

Edge When a logical connection is physically embedded onto the grid, it is composed of a series of 'edges' on the grid. Each edge travels from one cell to an adjacent cell.

3.4.2 High-Level Components

The bulk of the implementation is concerned with the manipulation of the low level components. The key components are as follows:

Network A Network object comprises of a set of nodes and edges. This is a logical network and is not necessarily embedded on the grid.

Network Generators In order to create networks, a series of building operations are applied. We start with node placement, using a specific type of packing strategy, and establish a connectivity matrix based on a chosen topology. Modifiers can be applied at the end, for example randomisation of certain connections in order to create surrogate networks.

Routing The Routing package contains classes concerned with physically embedding wiring, such as the A* pathfinder and AStarCell wrapper.

Simulators Implementations of the base class Simulator are high level controllers that aggregate various network creating and embedding operations based on the parameters set.

Utility Classes Network relevant operations, such as distance calculations, nearest neighbour searches and sorting operations are considered utility classes.

At an even higher-level are classes involved in visualisation, executing scripts and exporting network data to standard file formats.

3.4.3 Wiring using the A* Algorithm

In our model, wiring placement is significant because there is only a finite space available. If we are given a graph G with fixed vertex positions, and attempt to embed all the wiring, it is possible that not all connections can be established due to the spatial constraints. One possible optimization criteria would be to pack *as many* of these connections onto the embedding as possible. This kind of optimisation is akin to circuit board routing optimisation, for which almost all problems are NP-Hard [1]. Given that this makes it likely impossible for biological growth cones to consistently produce optimal wiring, sub-optimal wiring placement is a reasonable feature to assume in our model.

However, the model assumes that connections always take the shortest path when being embedded. This is achieved using the popular A* pathfinding algorithm [8], which is optimal in terms of the wiring length for the new connection, but may not lead to optimal wiring for the entire graph. The algorithm is superior to Dijkstra's algorithm since it uses an admissible heuristic to estimate the cost of the distance from a candidate node to the target node, making the search more informed. Other shortest path algorithms, such as the Floyd-Warshall algorithm, were also considered, but since the grid is modified after each embedding and we only consider one source node, A* was deemed to be the most applicable.

The A* algorithm works by considering three functions: f, g and h. At any given point, the A* algorithm will compute the cost associated with an adjacent node x by computing f:

$$f(x) = g(x) + h(x)$$

where $g(x)$ is the cost of reaching node x from starting node v_{start} and h_x is the estimated cost (using an admissible heuristic), of the path from node x to destination node v_{end}.

Since we use a grid based network model where connections can only travel along the x and y axis, the Manhattan distance is a suitable admissible heuristic.

A* Algorithm with Aesthetic Paths

When calculating the shortest path between two nodes using the A* algorithm, there are often many different paths across the grid with equal cost. Many of these paths will appear to zig-zag towards the target node (Fig. 3.1). These paths are not aesthetically pleasing, so a modification to the A* algorithm was made to ensure the path with the fewest turns is taken.

This is achieved by checking whether the expanded node involved a change of direction, and calculating the value for g as follows:

Algorithm 1 A* Pathfinding Algorithm

```
1:  function SHORTESTPATH(v_start, v_end)
2:      OPENSET ← v_start
3:      while OPENSET ≠ ∅ do
4:          remove x ∈ OPENSET with lowest value of f(x)
5:          if x == v_end then
6:              return path to x
7:          end if
8:          add x to CLOSEDSET
9:          for all n ∈ neighbour(x) do
10:             if n ∈ CLOSEDSET then
11:                 continue
12:             end if
13:             g_new ← g(x) + 1
14:             if n ∉ OPENSET || g_new < g(n) then
15:                 g(n) ← g_new
16:                 f(n) ← g(n) + h(n)
17:                 insert n into OPENSET
18:             end if
19:         end for
20:     end while
21:     return null
22: end function
```

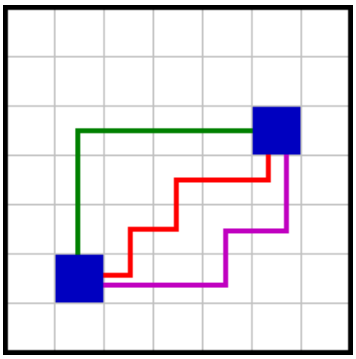

Fig. 3.1 Equal length shortest paths. Source: [16]

$$g_{new} = \begin{cases} g(current) + 1 + \varepsilon & \text{if direction changed} \\ g(current) + 1 & \text{otherwise.} \end{cases}$$

where ε is a value significantly smaller than the smallest weight, 1.

Although initially this choice was made primarily to aid visualisation, we later discovered that this feature significantly affected the number of connections that we could embed. This is discussed in chapter 5.

3.4.4 Summary

This chapter first looked at the thought process that went into how the problem would be approached. We decided to use a synthetic methodology which involved experimentation on a concrete 2D, spatially embedded model in order to make inferences. The rest of chapter looked at the design of the model, the reasoning behind the model specification, and touched on its implementation details.

Chapter 4

Minimising Conduction Delays

The first notable finding was the potential for SD to reduce conduction delays in our model. This area of study was inspired by the study conducted by Wen and Chklovskii [31]. This section describes the thought process and considers the parameters in which SD outperforms HD.

4.1 First Observations

4.1.1 The Effect of Increasing Density

The first question asked was "Given an embedding of an arbitrary graph G using the model presented, what are the effects of packing the same number of nodes into a smaller space?" This question arises when considering the brain could either embed G in such a way that neurons (nodes) are distributed across its entire volume, or instead place the same number of neurons into only a fraction of the available space.

The average distance between nodes decreases

The nodes are now closer together, so on average, the distance between any two nodes is reduced. On an NxN grid, the average distance grows linearly with the square of of N.

Long-range connectivity decreases but short-range connectivity remains high

Suppose we embed the nodes of G randomly, such that the probability of a node occupying any one of the squares is $\frac{n}{A}$ where n is the number of nodes in G and A is the number of squares on the grid. We then take all possible node pairings of the graph, sorted by the

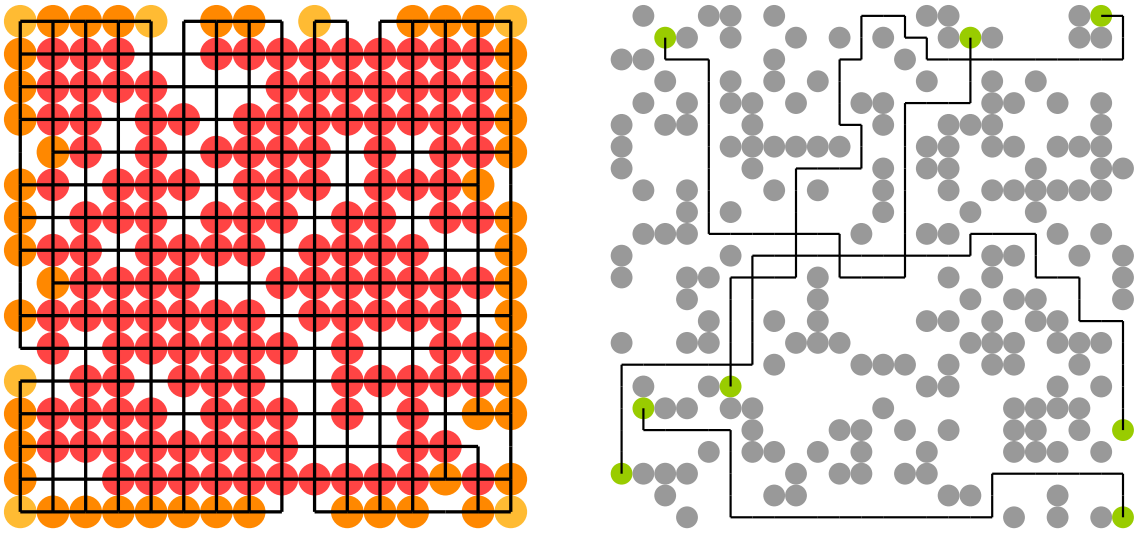

(a) 200 Nodes, 16x16 grid

(b) 200 Nodes, 24x24 grid

Fig. 4.1 Connectivity: Short connections (Left); Long connections (Right). For (b), d_{min} is the Manhattan distance between the grid's origin and midpoint. Note that the colour of nodes is indicative of node degree.

distance between the pair of nodes, and embed edges in two different ways: *(i)* Embed edges starting with the closest node pairs. *(ii)* Embed edges starting with the node pairs which are furthest apart, and only embed edges if the distance between the two nodes is greater than a threshold value d_{min}, which is scaled accordingly with the size of the grid. In order to achieve this a double-ended priority queue is used, with each element being a pair of nodes, ordered by the distances between two nodes in the pairs.

As demonstrated by Fig. 4.1, when nodes are connected with their closest nodes, many connections are possible. For graphs with 200 nodes, each with a maximum degree of 4, 400 connections is the theoretical maximum as each connection contributes a total degree increase of 2. When we instead try and make long range connections, connectivity is substantially reduced. Long range connections are less likely to find a free path to their target node due to the presence of nodes, while adding wiring to the embedding further contributes to the number of obstacles. A graph comparing the number of connections achieved as the area of the grid is varied is shown in Fig. 4.2.

4.1.2 Introducing White Matter: Random Connections

We now outline a preliminary analysis on the effect of changing the structure of the embedding by splitting it into two regions. Given an grid of length l and width w, we only permit

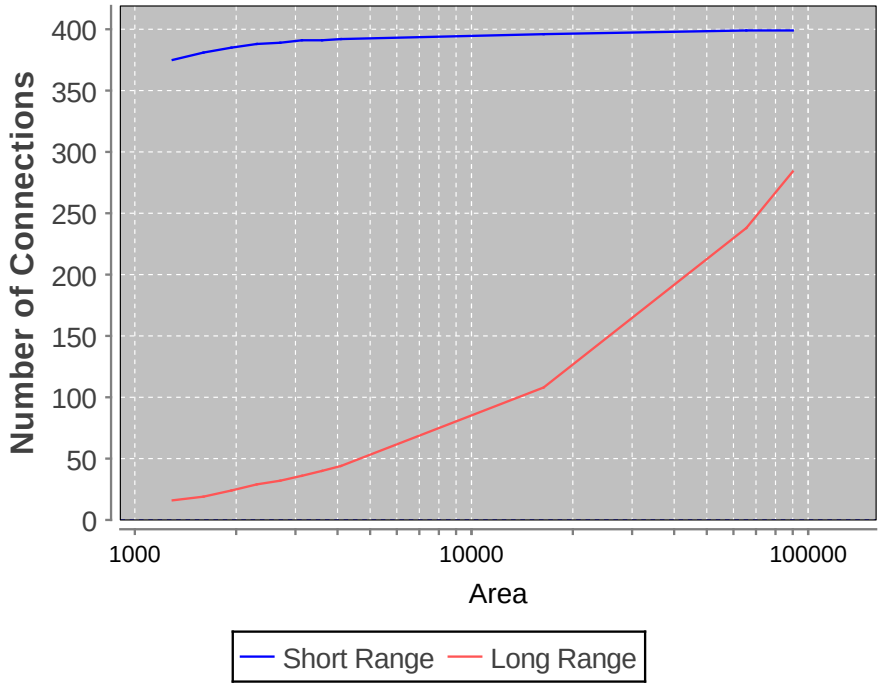

Fig. 4.2 How the maximum number of connections attained between 200 randomly placed nodes changes with total area, in terms of short and long range connections.

nodes to be added onto a bounded region of the grid, with length l_{grey} and width w_{grey}. This is representative of grey matter, and edges can now only begin or terminate in this region.

First, we consider what happens when we attempt full (all-to-all) connectivity. The number of possible node pair combinations is $n(n-1)/2$, which would require extremely large grid areas, and the maximum node degree of 4 makes this topology impossible for $n > 5$. Since this is infeasible, and the fact that this kind of topology is rare in nature, we instead consider random networks.

Random networks are created as follows. We consider all possible edges between nodes, and independently, with some probability $0 < p < 1$, add edges to the network. With this sample of randomly generated edges, we attempt to embed them onto the grid. It is important to note that in all cases, successfully embedding the entire sample of connections onto the grid is not likely, as specific connections are unlikely to find a path to its target node as the number of obstacles of the grid increases. Instead, we compare the different architectures by comparing the number of successful connections established, given a sample of random connections.

Fig. 4.4 shows how many connections, given a random sample, were successfully embedded onto a 100x100 grid with $n = 500$. Three architectures are considered: *(i)* A homo-

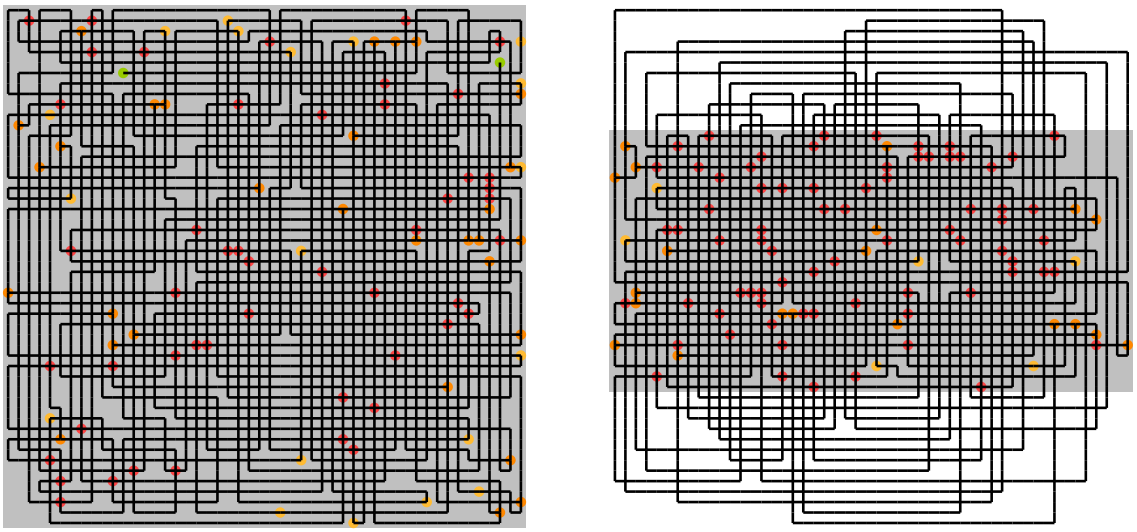

(a) Homogenous Design

(b) Segregated Design with grey matter centered

Fig. 4.3 Random Connectivity for 50x50 grid with 100 nodes

geneous design. *(ii)* A segregated design with a horizontal sheet of grey matter, with white matter below *(iii)* A segregated design with a horizontal sheet of grey matter, but this time placed in between two regions of white matter. Scaled down examples of *(i)* and *(iii)* are shown in Fig. 4.3a and Fig. 4.3b respectively.

The results show that the more white matter we introduce, the connectivity for random connections decreases. In all architectures, each successive connection made reduces the chance that the next connection will be successfully embedded. This is evident from the falling rate of change in the number of connections. In this regard, SD does not appear to result in any advantages over HD, but we do notice that when the horizontal sheet of grey matter is centered in the enclosure, connectivity is higher.

The Cherry Picking Problem

One problem with this experiment is that, as we increase the sample size of connections to embed, the graph that gets embedded becomes less and less random. This is a result of certain types of connection being 'cherry picked'. For example, if the random sample size is $n(n-1)/2$, it is the same as trying to embed a graph with all-to-all connectivity. In this scenario,

It is worth noting that as the number of connection attempts increases, the network that is embedded becomes less random as a result of 'cherry picking' certain types of connection. As Fig. 4.4 shows, the segregated designs slowly approach the curve of the homogeneous

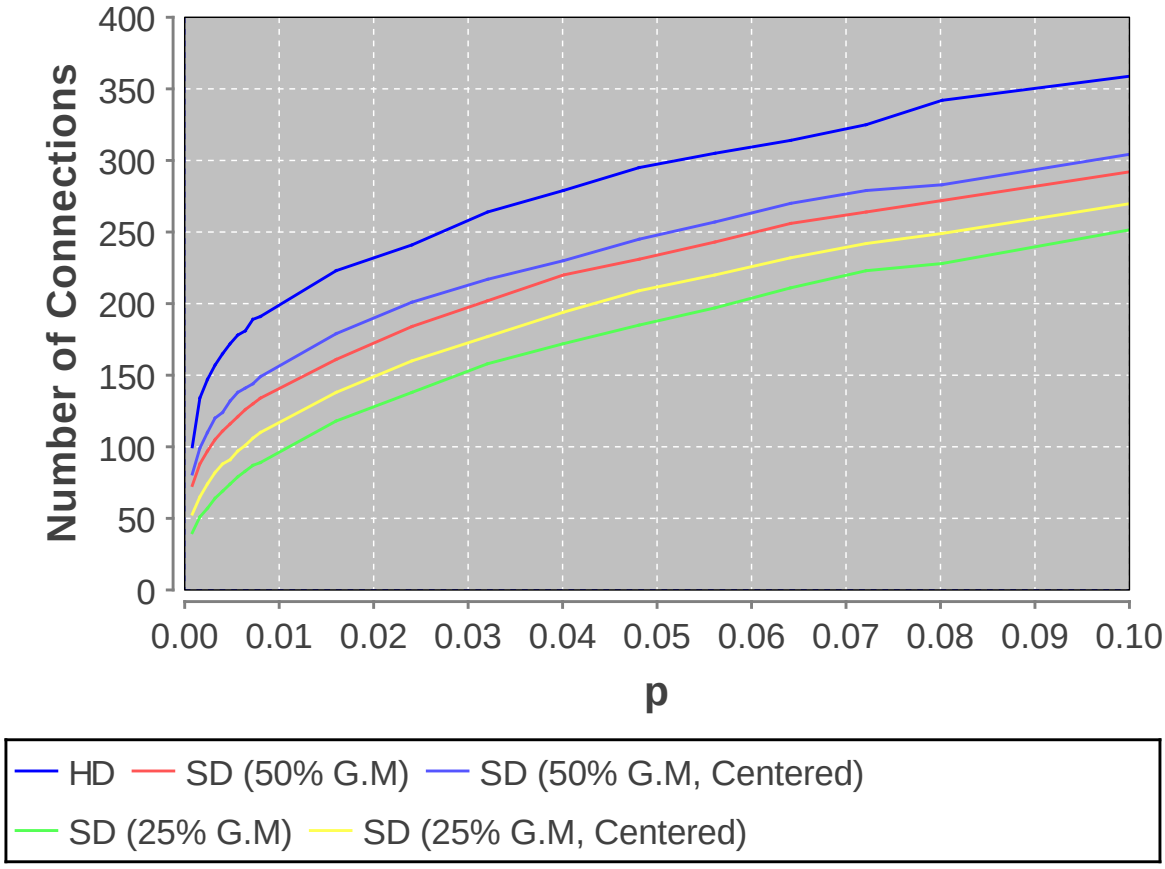

Fig. 4.4 Random connectivity performance between different architectures.

design, but this is mainly a result of selectively choosing shorter connections which were shown to be easier to connect. This growing bias is inherently difficult to solve, so we limit the results shown to $0 < p \leq 0.1$. Results for other values are included in Appendix A.

4.2 Minimising Conduction Delays

The previous section showed that random connectivity was significantly reduced when segregating the grey and white matter. However, the background research showed that the networks of mammalian brains do not have such a connectivity pattern, and instead it is the case that the vast majority of connections a neuron makes is with its nearest neighbours. This section looks to see if it is possible to generate embeddings with similar levels of connectivity, while benefiting from reduced conduction delays, as discussed in the study by Wen and Chklovskii [31].

4.2.1 Measuring Conduction Delays

In this model we assume the conduction delay across a connection increases with connection length. However, one key consideration to make when assessing conduction delays in the brain is to differentiate between **myelinated** and non-myelinated axons. In myelinated axons, conduction velocity scales linearly with the fiber's diameter. Wen and Chklovskii's [31] study highlights the importance of minimising the conduction delays in short range, unmyelinated axons, which do not share the same property. Fig. 4.5 shows the theoretical relationship between the two types of axon.

In our model we do not capture the notion of connection diameter, but we still classify connections can still distinguish between the two types of connection. We simplify the problem by assuming the non-myelinated connections are the bottleneck, and aim to reduce the short range conduction delays as much as possible. It is therefore important to bear in mind that this simplification is taking place in the model.

A computer network analogy

While the model aims to discern possible implications of grey and white matter segregation in the context of neural networks, the idea of distinguishing between different kinds of connection is present in other real-world networks. One analogy that could be used is a large scale computer network. Suppose we are trying to connect local computers using cost-efficient Ethernet cables, but also have a smaller number of costly Fiber Optic cables

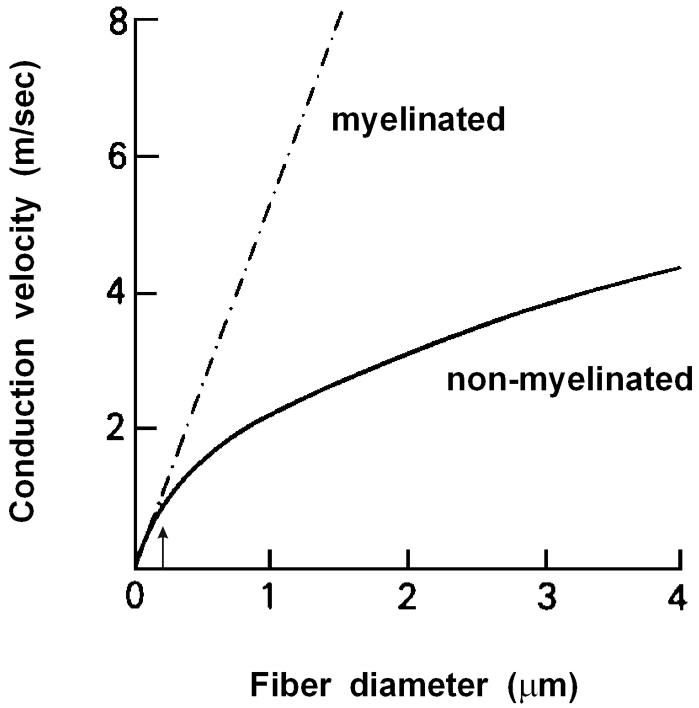

Fig. 4.5 Theoretical relationship between conduction velocity and axon diameter for small myelinated and non-myelinated axons. Source: [29]

available, which we assume can transport information fast enough, regardless of distance. What spatial arrangement of computers and cabling would favour the minimisation of the length of Ethernet cables?

4.2.2 Connection ordering

In order to minimise the conduction delays of these short range connections, the order in which we embed connections on the graph becomes important. As the grid becomes more congested, it is less likely that a connection will find a direct path to its target node, but instead have to wriggle around the existing wiring and take a longer route. Therefore, it is advantageous to begin by embedding the shortest connections first, given this optimization criteria. This principle will be applied when we compare conduction delays in homogeneous and segregated designs.

4.2.3 HD vs SD: Modular Connectivity Pattern

We consider how HD and SD compare in terms of conduction delay using the modular brain connectivity model presented in the background. In this topology, each node connects with

its k-nearest neighbours, and has a chance of making a long-range connection. Again, we consider HD and the two forms of SD which are horizontally sheeted.

Finding the k-Nearest neighbours

In order to efficiently locate the k-nearest neighbours of any given node, we maintain a K-d tree [28] in dimension 2, which organises nodes based on their grid coordinates. Building the tree is an O(n log n) operation, while searching for a nearest node is on average an O(log n) operation. The tree structure is ideal for organising the nodes in our model since each node has its own 2-d coordinates and no two nodes can occupy the same co-ordinates.

Experiment: Varying the minimum distance between modules

In this experiment, we attempt to embed a modular network with 2300 nodes and 200 hub nodes onto a grid of size 500x80, over HD and the two forms of horizontally sheeted design. We do not consider a square 200x200 grid, for which SD produced unremarkable results. Reasons for this are explained later. Hub nodes are nodes which are candidates for making long-distance connections, and are spread randomly like any other node. The variable in this experiment is the minimum distance between hub nodes we attempt to connect. The procedure is as follows:

1. Embed 2300 nodes, randomly placed, onto the grey matter of the grid.
2. Embed 200 hub nodes, randomly placed, onto the grey matter of the grid.
3. Each node attempts to connect with its k-nearest nodes.
4. Consider each pair of hub nodes, and if the **horizontal** distance between them exceeds a threshold value G_{min}, attempt to embed a connection between them.

There are a couple of points to note. First, hub nodes are not physically connected to the non-hub nodes. Initially, experiments were performed using a 'rewiring' strategy, where each node would consider making a global connection with some probability p. However, we are concerned with the physical constraints, and since each node only has 4 terminals, a rewiring strategy would restrict the long range connection to exiting from just one terminal, which is an unnecessary hurdle. Second, we only consider the horizontal distance between the two hub nodes to keep the experiment fair, as through segregation we reduce the vertical distance of grey matter, meaning otherwise the non-segregated layout would consider more

hub node pairings. We run this experiment in each design 50 times for each value of G_{min}, which makes 50 unit increments.

The results are as follows:

- The average short-range connection length decreases by a factor of around 0.82 when we go from HD to SD, which is slightly worse than the factor we would expect based on the change in average Manhattan distance between any two nodes, 0.75. This is likely due to relatively more congestion in the grey matter.
- Short range connectivity remains very similar between the different architectures. HD, SD and SD with centered grey matter achieved an average of 4136, 4090 and 4111 local connections, respectively. This gives a local connectivity factor of ≈ 0.99. The discrepancy between the figures for the two segregated designs is due to the boundary effect, which is discussed later.
- In the segregated designs, global connections are seen to leave the grey matter and travel along white matter, before re-entering the grey matter to reach their target node. It must be the case that either the grey matter offers no direct path due to congestion, or the path is not the shortest.
- When $G_{min} \geq 150$, both forms of SD achieve superior global connectivity over HD, with the centered form producing the highest levels. These results are shown in Fig. 4.6.
- The average **global** connection length is shorter in both forms of SD than HD when $G_{min} \geq 250$. This statistic is slightly more difficult to assess since the number of global connections is variable between the three designs. However, grey matter becomes more densely packed than white matter, and hence in HD projections are perturbed further from the minimal-cost straight line path. This is the area exclusion factor, analogous to volume exclusion found in a three-dimensional space [4]. The overhead of exiting and reentering grey matter can, as a result, be outweighed by the cost of avoiding obstacles within the grey matter.

4.2.4 White Matter Efficiency

When we change to SD and embed the local connections first, we guarantee reduced local conduction delay, while maintaining almost the same local connectivity. However, the relative global connectivity appears to depend on the lengths of global connections. When we

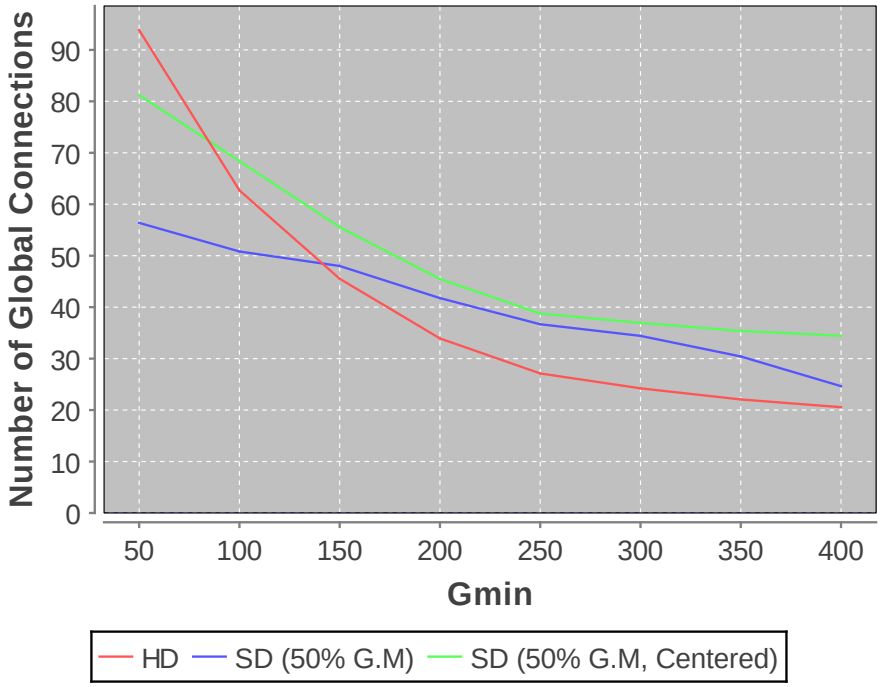

Fig. 4.6 Number of global connections embedded onto grid, given a minimum absolute distance between connected hub nodes

only consider connecting hub nodes which are separated by $G_{min} = 150$ units in the x-axis, SD was capable of embedding more global connections. As the value of G_{min} is further increased, the global connectivity gap between HD and SD increased further.

Intuitively, this seemed to be the result of white matter being used more efficiently. A global connection spends a certain fraction of its time leaving and re-entering grey matter, which becomes an overhead since it is not moving directly towards its target node. The relative impact of this overhead is reduced as the connection spends more time benefiting from the obstacle-free path provided by white matter. We define a measurement for the white matter efficiency of an embedding as follows:

$$\text{W.M. Efficiency} = \frac{\text{Total global connection length in W.M.}}{\text{Total global connection length}}$$

To test this hypothesis, we carry out an experiment similar to the previous experiment, but now place an upper bound, G_{max}, on the horizontal distance between hub nodes that are to be connected. By specifying the range of distances between the hub nodes we attempt to connect, we can explicitly influence the amount of time a global connection will spend in white matter, relative to grey matter. Fig. 4.8 shows how when we exclusively only permit relatively short global connections (50-100), the white matter efficiency is low. This

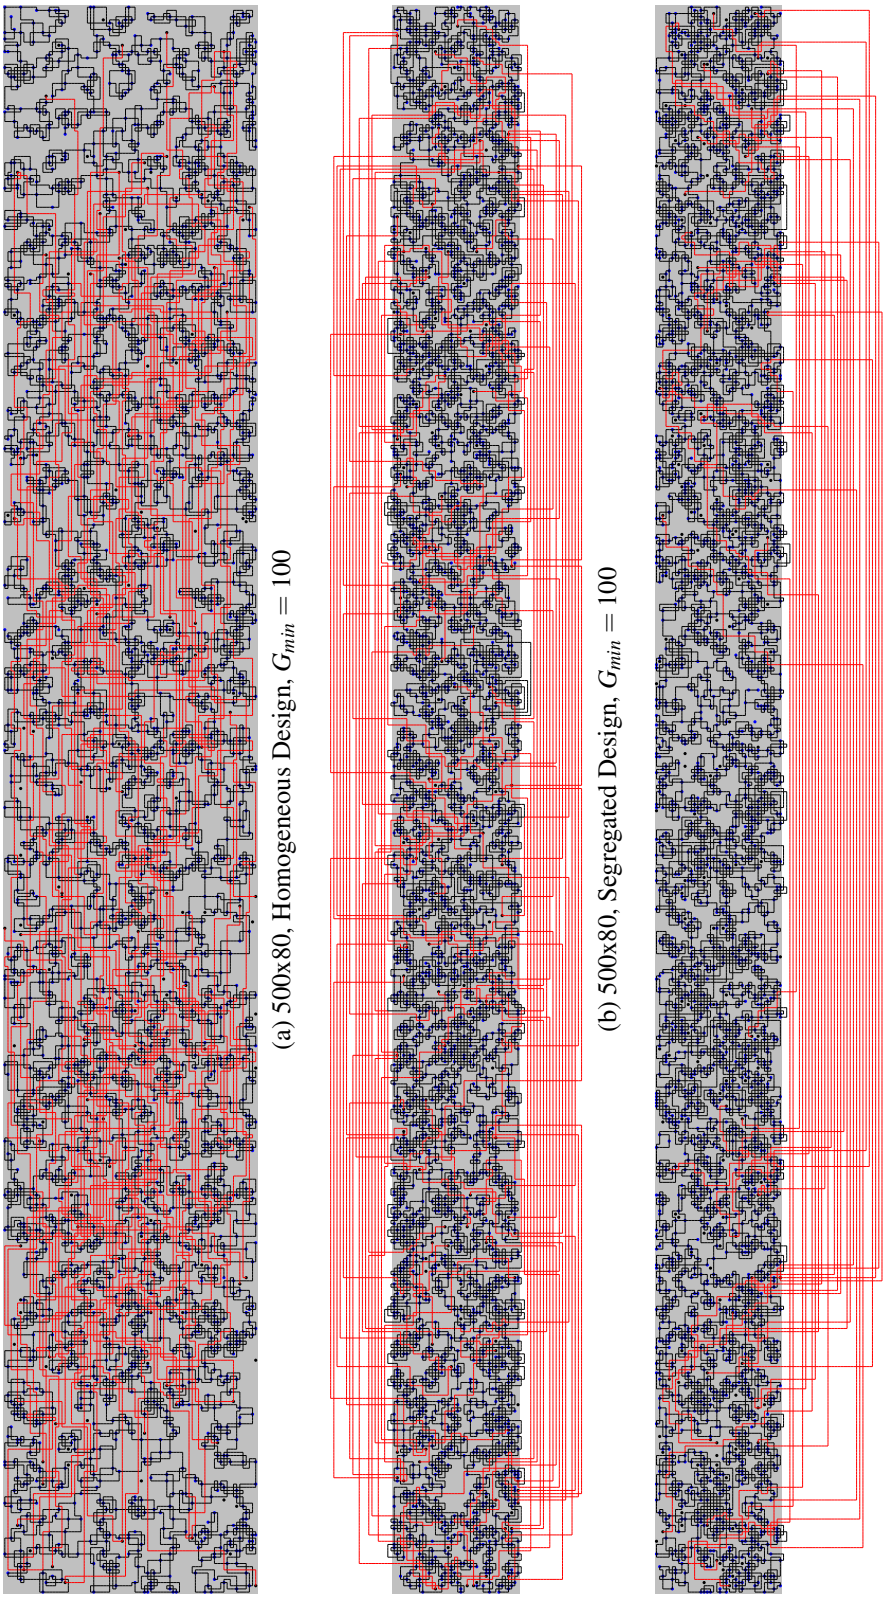

(a) 500x80, Homogeneous Design, $G_{min} = 100$

(b) 500x80, Segregated Design, $G_{min} = 100$

(c) 500x80, Segregated Design, $G_{min} = 300$

Fig. 4.7 Visualisation of the first experiment, which varied G_{min}, using three examples. 'Global' connections are shown in red.

is especially the case for when the grey matter is not vertically offset (centered), as the distance between a node and the boundary between grey and white matter is potentially doubled, meaning proportionally more time being spent in grey matter.

Maximising White Matter Efficiency

Since white matter efficiency is dictated by the proportion of length a connection covers in white matter, when there is a range of global connection lengths, the shorter global connections should be wired closest to the grey matter.

4.2.5 Varying White Matter Percentage

So far we have been concerned with varying the lengths of global connections, but have only considered SD with a fixed amount of white matter - exactly half of the available area. To examine the effects of varying the proportion of the area allocated to white matter, we perform the same experiment as before but fix G_{min} at 250. This means global connections span at least half the horizontal distance of the grid. The variable in this experiment is now the percentage of white matter, which we increase from 0% to 90%. Fig. 4.9 presents three sets of results: (a) The number of local connections achieved. (b) The average length of local connections. (c) The number of global connections achieved. Note that when the white matter percentage is zero, the model is equivalent to HD.

Results

- The number of local connections is fairly constant, but begins to dip slightly beyond 50%.
- The mean local connection length falls as the nodes become more densely packed, indicating shorter conduction delays. Beyond 50% white matter, the values vary significantly between the two forms of horizontally sheeted SD.
- Increasing white matter also increases the number of global connections being embedded, until a certain point.

The number of local connections embedded is slightly higher when the grey matter is centered as there is white matter above and below, meaning fewer nodes are adjacent to the external boundary which completely blocks wiring, unlike white matter.

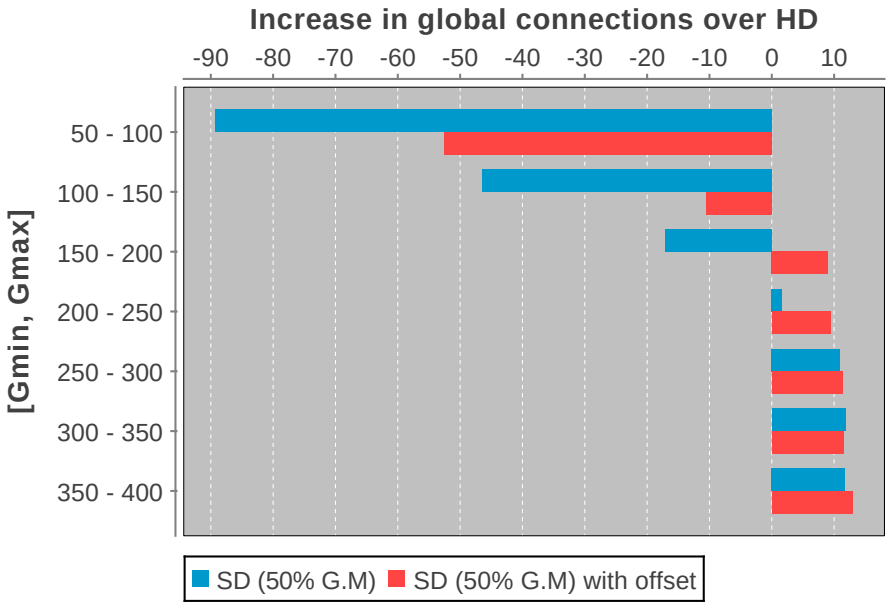

(a) HD achieves higher global connectivity when connected hub nodes are relatively close together. The opposite occurs when connected hub nodes are relatively far apart.

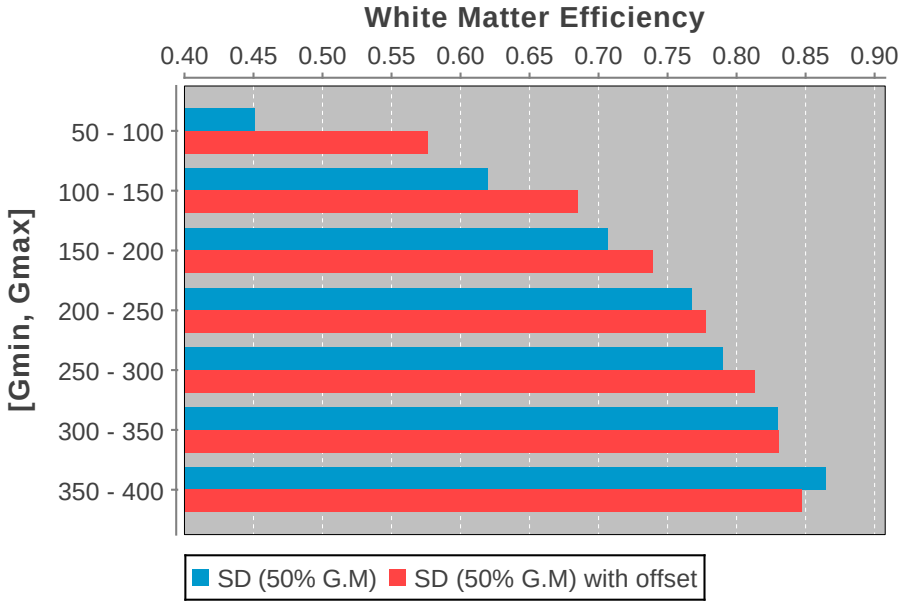

(b) When global connections are longer, the white matter efficiency is higher.

Fig. 4.8 Improved global connectivity is reflected by the white matter efficiency

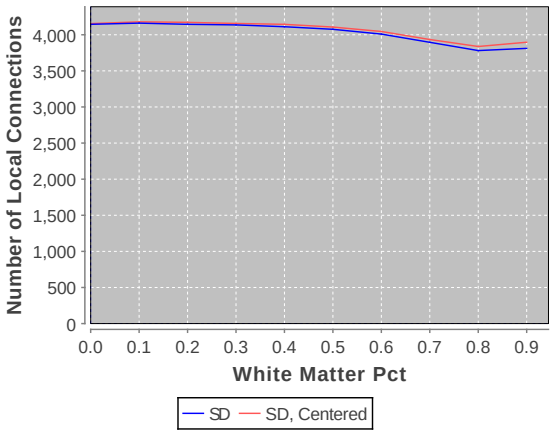

(a) Local connectivity

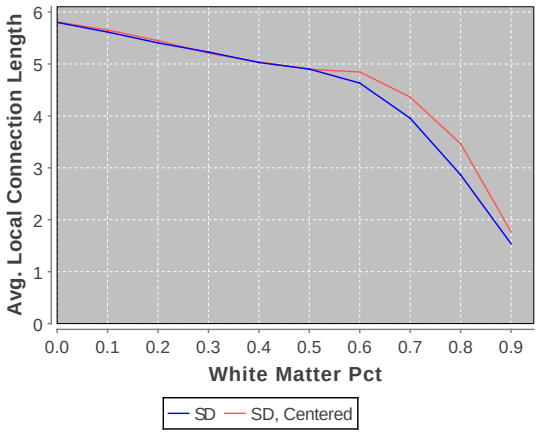

(b) Average length of local connections

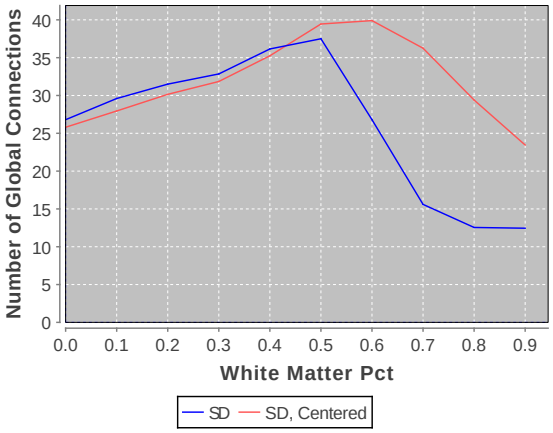

(c) Global connectivity

Fig. 4.9 Effects of varying the white matter percentage

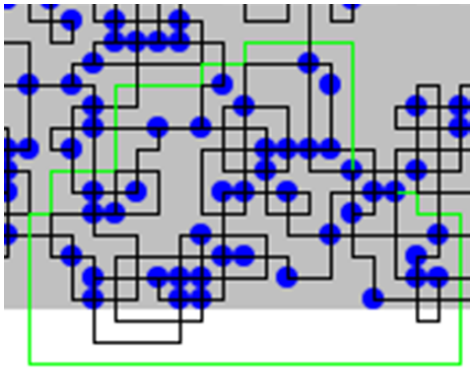

Fig. 4.10 At high densities, local connections can become indirect (shown in green).

We also notice that the local connection length begins to momentarily plateau for the case where grey matter is centered. This happens as the high density becomes an obstacle for further connections, making local connections compensate by taking long routes via the white matter, as shown in Fig. 4.10. The effect is less pronounced when the grey matter is in contact with the boundary, as instead fewer connections are made.

The most interesting result is the change in the number of global connections. Consistent with our first experiment where $G_{min} = 250$, we can embed more global connections as we gradually increase the amount of white matter. However, after 50% white matter coverage, things begin to change. When grey matter is centered, the number of global connections embedded continues to increase until 60% white matter, but in the other design there is a sharp fall. The density in both designs is the same, yet having that white matter above as well as below seems to yield remarkably better results. An explanation for this, in terms of percolation theory, is discussed in section 6.2.2.

A possible explanation for increased global connectivity over HD

In SD, the connection lengths are generally shorter. We found this to be the case in both local and global connections. Since wiring is taking up less space, more space is left over for further global connections.

4.3 Summary

This chapter demonstrated cases in which, for our model, a segregated design had clear advantages over a homogeneous one. We first observed that it is possible to densely pack nodes while maintaining high levels of connectivity, provided that we only connect nodes which are close together. We took this idea and applied it to a modular network, and found that we could pack in just as many local connections into a smaller area and benefit from shorter path lengths (and hence shorter conduction delays). The area left over from the compression of nodes was white matter, and we found this could be exploited to embed the global connections that would otherwise be impossible in the compressed grey matter. Further, we found we could potentially embed **more** global connections than a homogeneous design.

The segregated design required a few parameters for it to clearly outperform the homogeneous design. First, our optimization criteria was to minimize local conduction delay, and hence the local connections were embedded first. If our optimization criteria was to

minimize the global conduction delay, global connections would initially find shorter paths through the grey matter, meaning local connectivity would be severely impacted. Second, global connections needed to be sufficiently long to make the segregated design effective. This was a result of the relative overhead of leaving and re-entering grey matter, which we analysed in terms of white matter efficiency. The requirement of having sufficiently long global connections was one of the reasons using a 500x80 grid, but in chapter 6 we discover that a thin layer of grey matter is an essential feature for high global connectivity in SD.

Chapter 5

Percolation Theory

The previous chapter demonstrated a specific case where segregating nodes and local connections from global connections led to reduced conduction delays and higher global connectivity. The idea of white matter efficiency was introduced to better understand the conditions in which segregation becomes effective. This chapter aims to provide more general explanations for the differences between HD and SD, in terms of percolation theory, which realises that the grid will become disconnected as it becomes more densely packed.

5.1 Motivation

In chapter 4 we mentioned one possible explanation for why SD could enable higher global connectivity than HD under certain circumstances. The explanation was based on the idea that since connections become shorter, given a fixed number of connections, more free space would be left over under SD. While this is a significant factor, consider the following two randomly selected embeddings for a graph of size 500x80, 2300 nodes, 200 hub nodes, $G_{min} = 400$:

	Embedding 1 (HD)	Embedding 2 (SD)
Number of local connections	4174	4170
Avg. local connection length	5.77	4.88
Number of global connections	21	37
Avg. global connection length	562.47.	509.91
Total connection length	35909	39212

The connections embedded in HD occupy a smaller total length, yet no further global connections can be established. It appears that the placement of nodes and wiring is less

efficient than SD for global connectivity.

5.2 Global Connectivity

5.2.1 Site percolation

The nodes and wiring that gets embedded onto the graph become potential obstacles when trying to embed additional wiring. We start by abstracting away the details of our model and consider how connectivity is affected as cells on a grid become 'occupied'.

Percolation Trials

Our first experiment is as follows:

1. Take a grid of size *NxN* with all cells initialised to 'empty'.
2. Go through each cell in turn and with a probability $1 - p$, make the cell 'occupied'.
3. If there exists a path comprising only of empty cells going from the top of the grid to the bottom, count this trial as a 'success'.
4. Repeat this process r times, and calculate the probability of there being a path from top to bottom.

This problem is known as site percolation [9] and as a matter of convention we take p to be the probability that a cell is 'empty'. The results of this experiment are shown in Fig. 5.2. Two things can be observed: First, there is a clear phase transition as we increase p. Second, as we begin to consider larger lattices, the transition becomes more abrupt. When N is infinite, the probability of success becomes a step function, in which there is no connectivity if p is below a critical value, p_c, and certain connectivity if $p > p_c$. The exact value of p_c for a square lattice is unknown, but one numerical analysis puts it at 0.592746010 [15].

Checking for connectivity with Union-Find

In order to efficiently run percolation trials, we do not need to actually trace a path, but instead only check whether a path exists. This is achieved using a disjoint-set data structure, where the elements are the linear index of the cell position on the grid, and the union-find algorithm [26] that keeps 'connected' cells (cells between which an open path exists) in the same subset and allows for fast querying of whether two cells are connected.

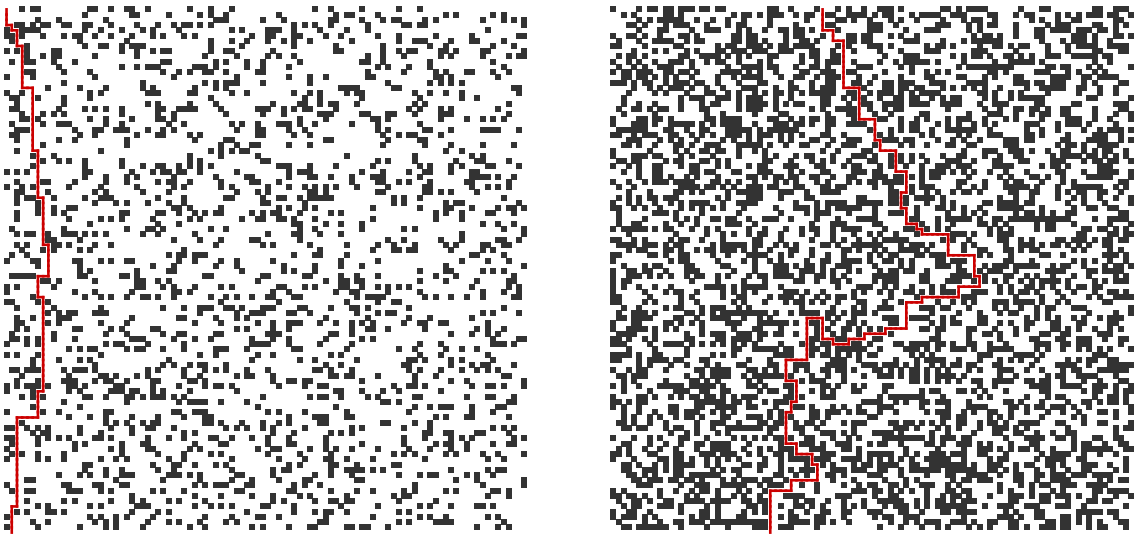

(a) p = 0.8 (b) p = 0.6

Fig. 5.1 Examples of successful percolation trials, $N = 100$

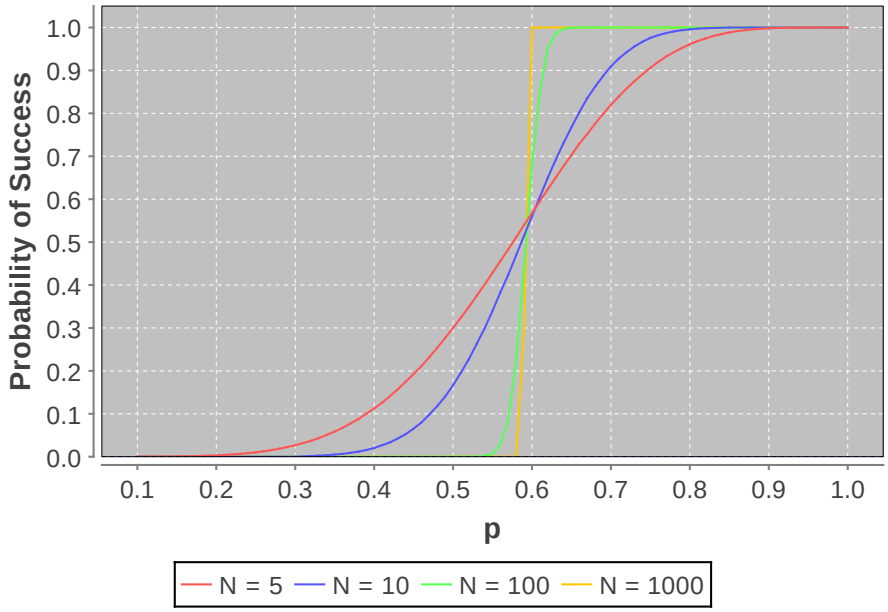

Fig. 5.2 The probability a path exists from top to bottom given a NxN grid with each cell being 'empty' with probability p (Nodes only).

5.2.2 Percolation with Wiring

The experiments performed on our model thus far have all initially started with a sparsely populated grid, with only nodes (neurons) occupying cells. Nodes were allocated randomly, and the value of p can be calculated as n/A, where n is the number of nodes and A is the total area of the grid. The sparse placement of nodes means $n/A >> p_c$, so global connectivity is almost certain. The more interesting question to ask is what happens as more and more wiring gets added to the grid.

Like nodes, wiring constitute obstacles on the grid, so it is likely that the principles of our first experiment will carry over. However, the properties of wiring is different from nodes, since they can be crossed orthogonally. We carry out the percolation trials as before, but now occupied cells can either contain a node or some wiring. To remain as true as possible to our model, we make the probability that a cell contains a node relatively small. Since wiring does not permit backtracking, all cells excluding those that contain a node will only have up to two outgoing edges.

To make a direct comparison with the abstract percolation model, we continue to vary a single variable, p, representing the probability a cell is empty. In the trials performed, we set the probability an occupied cell contains a node to be 0.1 and the probability it contains one or two edges as 0.6 and 0.3 respectively. The results are shown in Fig. 5.3.

The results in Fig. 5.3 show a similar phase transition, but this time a much greater amount of objects on the grid are tolerated before global connectivity is lost. We can vary the specific probabilities that a cell is occupied by a node/one edge/two edges, leading to shifts in the curve.

5.2.3 Measuring Percolation using Growth Cones

As expected, the phase transition of global connectivity still appears when wiring is introduced. However, wiring in these trials were added on a completely random basis, ignoring the fact that the presence of edges in one cell will generally mean adjacent cells also contain edges that follow on from these. The resultant embedding is hence unrealistic, as demonstrated by Fig. 5.4a.

Inspired by biological growth cones, which are at the head of an axon and responsible for seeking target neurons, we can measure percolation for our model in a more realistic way. Instead of using probabilities to determine node and wiring placement, we have each node send out a growth cone, which seeks out other nodes. We are not concerned with whether or not the growth cone makes a connection, but simply that it produces a realistic,

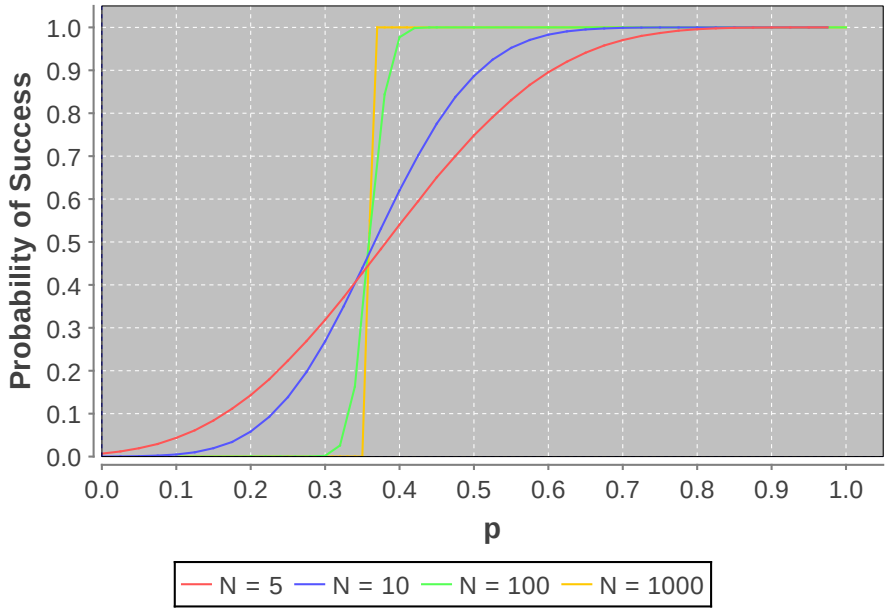

Fig. 5.3 The probability a path exists from top to bottom given a NxN grid with each cell being 'empty' with probability p (Nodes and Wiring).

contiguous piece of wiring.

In the following percolation trials that use growth cones, the variable is now coverage, c, given by:

$$c = \frac{\text{Total Length of wiring}}{\text{Grid Area}}$$

Coverage represents how much 'growth' has occurred between all the growth cones, giving us a sense of density, while being based on embeddings which are more realistic than those based purely on probabilities. On each iteration, each growth cone will attempt to grow their connection by one unit by reaching an adjacent cell. Assuming there is no obstacle ahead, they will preferentially continue travelling in the same direction with probability s, otherwise changing direction. As soon as the target coverage is met, the percolation trial is performed with the resulting embedding. The results of these trials for $s = 0.55$, $s = 0.65$ and nodes occupying 5% of cells are shown in Fig. 5.5.

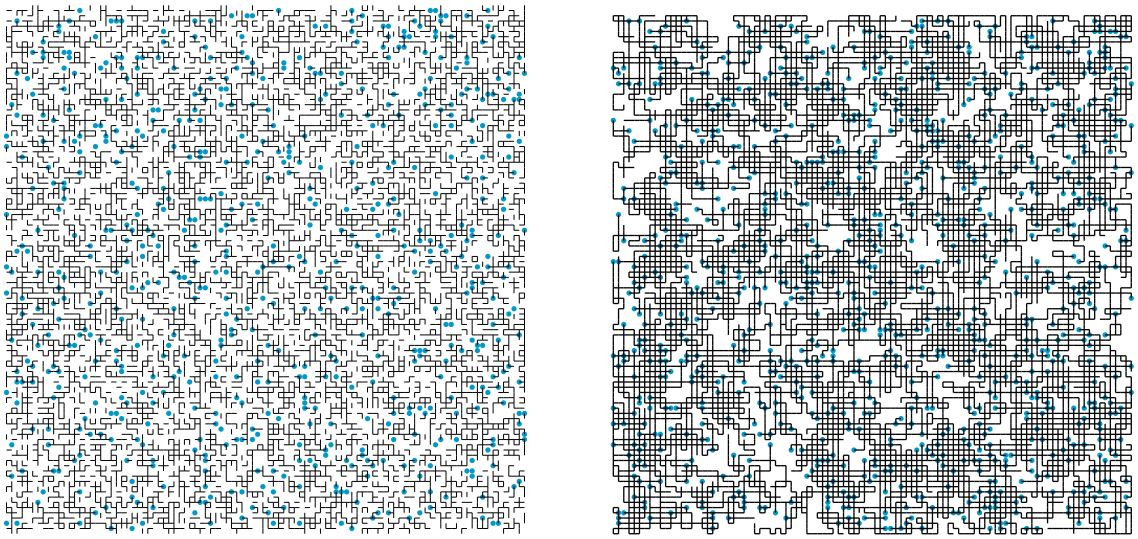

(a) $p = 0.35$ (b) $c = 1.35, s = 0.55$

Fig. 5.4 Unsuccessful percolation trials using two models: (a) Abstract wiring, (b) Growth cone based wiring

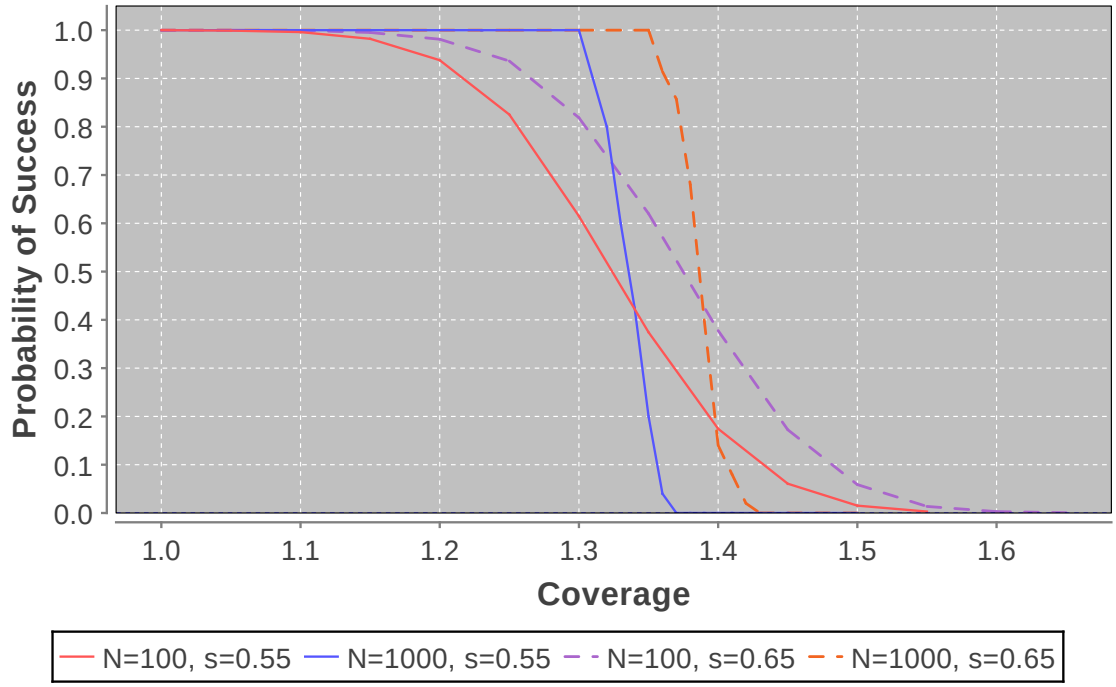

Fig. 5.5 Percolation Trials on Growth Cone based embeddings. s is a measure of how straight connections grow.

Discussion of results

Once again, a phase transition is observed. For both values of s, loss of connectivity is only seen to occur when $c > 1$. This means there is certainly some crossing involved in the embedding, as the total length of wiring exceeds the area of the grid.

A second observation is that as a growth cones follow more linear trajectories, global connectivity is maintained at higher levels of coverage. It is possible that, the more 'chaotic' wiring in the grid becomes, it becomes more difficult to find a path from top to bottom as there are more points at which we have to change direction, in order to obey the rule of not being able to travel on top of another piece of wiring. When connections are more ordered and found travelling in parallel straight lines, we guarantee that we can cross them all without having to change direction. In Appendix B, we revisit this point and show how the A* pathfinding algorithm achieves higher connectivity in embeddings when it travels in straighter lines.

5.2.4 Clustering

The percolation trials have given us an idea of how connectivity is affected as more and more wiring is embedded. Initially, the grid is sparsely populated and global connectivity is almost certain. The addition of wiring will eventually lead to a stage where a connection going from one side of the grid to the other is impossible. Another way of looking at this problem is in terms of clustering: A cluster is group of cells on the grid, in which for any one of its cells, a path exists to every other cell in the group. If we can reach one side of the grid to the other, it means there is at least one cluster which spans the entire length of the grid.

Fig. 5.6 shows examples of clustering on the abstract site percolation model we started with. Recall that for this model, the critical value p_c is 0.593 to 3.s.f. In (a) we see one cluster span from top to bottom. In (b), (c) and (d), where $p < p_c$, we are in the "subcritical" phase, and the clusters become progressively smaller. In fact, Menshikov [2] showed that in this phase, the probability that a specific point (for example, the origin) is contained in an open cluster of size r decays to zero exponentially in r. Put another way, let $S_{x,d}$ be the event where an open path from a cell x to another cell whose distance from x is d. Then let the probability that an event E holds when sites are 'empty' with probability p be $P(E)$. Menshikov's theorem implies that for a given $p < p_c$, there is a constant a such that:

$$P(S_{x,d}) \leq e^{-ad}$$

This is significant to our model because it suggests that as wiring is added and the embedding transitions into its subcritical phase, the probability of being able to establish a path between two nodes separated by a Manhattan distance d decays exponentially with d.

5.3 Summary

This chapter explored percolation theory, and its potential applications to our model. Site percolation deals with the issue of connectivity on a much more abstract level, by considering whether a cell is either occupied or empty, based on probability. Our network model cannot be reduced to a single probability value, since the embeddings are more complex, but we carried out similar trials on more concrete models which involved wiring. We consequently found that increasing the density of wiring in our model produced similar phenemona. Increasing the amount of wiring on the grid leads to the grid becoming disconnected, meaning certain connections become impossible to embed.

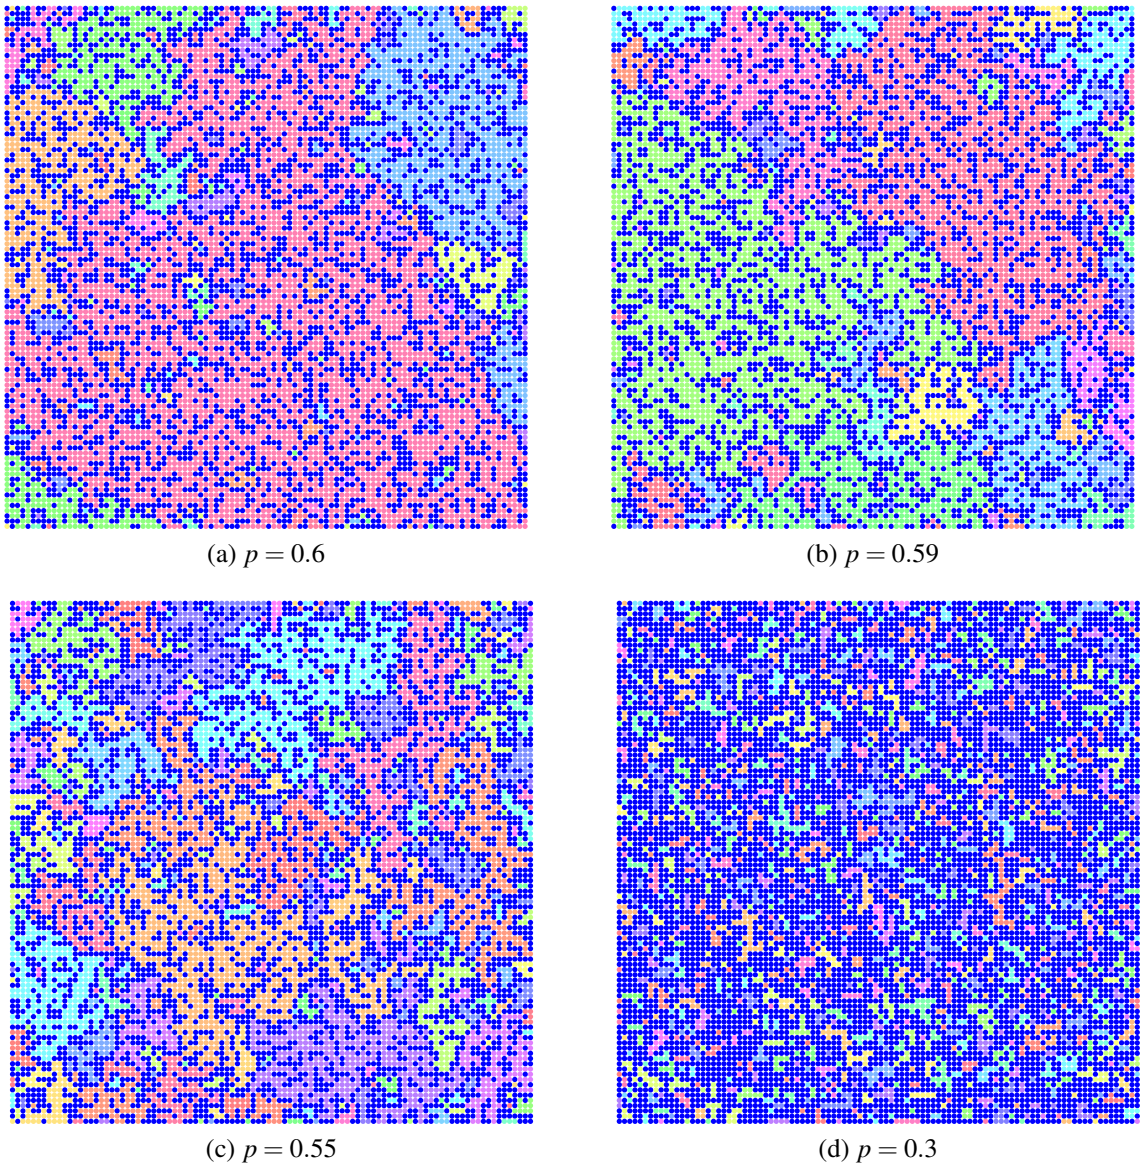

(a) $p = 0.6$ (b) $p = 0.59$

(c) $p = 0.55$ (d) $p = 0.3$

Fig. 5.6 Clustering on abstract embeddings with 'occupied' (shown in blue) and 'empty' cells.

Chapter 6

Applying Percolation Theory

At a fundamental level, SD guarantees, regions where nodes are completely absent. Results from chapter 4 suggest that this is potentially advantageous for global connectivity. This section provides additional explation for the results of the chapter in terms of of how clusters form in the respective designs.

6.1 Achieving Global Connections in dense embeddings

The previous chapter explained how longer range connections become increasingly difficult to establish as the coverage of the grid increases. We now look at how clusters form in graph embeddings that we create, and suggest how global connectivity can be facilitated by introducing white matter.

Consider an arbitrary graph embedding which is in its subcritical phase and therefore has finitely sized clusters, like the random network shown in Fig. 6.1.

Suppose we wish to make a connection from a node in a cluster on one end to a node in a cluster on the other end. This is not possible, since by the definition of clusters no such open path exists. However, if we now surround the entire grid with a new layer of completely empty cells, the connection is made possible: It first leaves its cluster of origin, moves onto the new layer of cells, before entering the cluster of its target node.

Landlocked Clusters

We can distinguish between clusters that are completely surrounded by other clusters and those which contain cells adjacent to the grid boundaries. We can call the former 'land-locked' clusters, and even if we introduce a new layer of cells, it is impossible for these

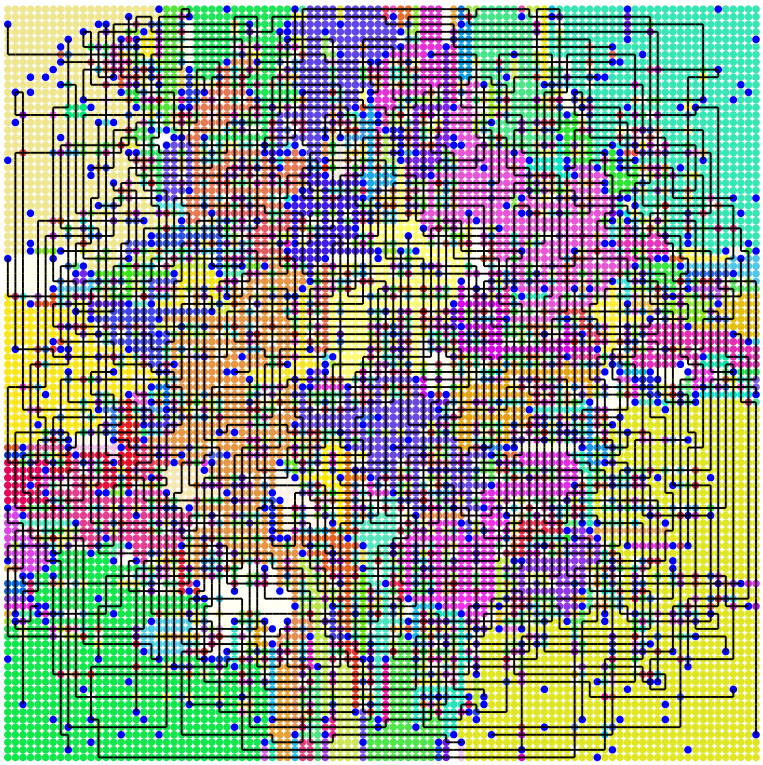

Fig. 6.1 A random network in the subcritical phase.

clusters to take advantage of them. Hence no direct connection can be made with other clusters. Note that: (a) Connections are already present in the embedding, so nodes in separate clusters may already be connected, but **additional** connectivity is impossible. (b) We are not considering logical connectivity, meaning indirect 'paths' that involve hops between nodes.

6.1.1 Boundary Length

In HD, we can reduce the density of the grid by spreading out nodes and connections evenly. However, if the grid was dense, the clusters will remain small and global connectivity will remain impossible. Instead, if we artificially add a layer of empty, contiguous cells, we can guarantee a connection between two non-landlocked cells.

There also exists a solution to landlocked clusters. In chapter 5 we only considered nodes being arranged on a square grid, with nodes occupying any cell with even probability. From Menshikov's Theorem, we know the probability that there is an open path from a cell to a boundary, given some density, is exponential in the distance between the cell and boundary, d. If we reduce the maximum distance between any node and a boundary, land-

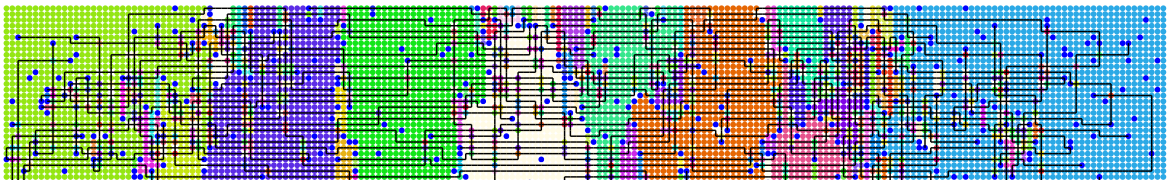

Fig. 6.2 Reducing landlocked clusters by reducing lattice width

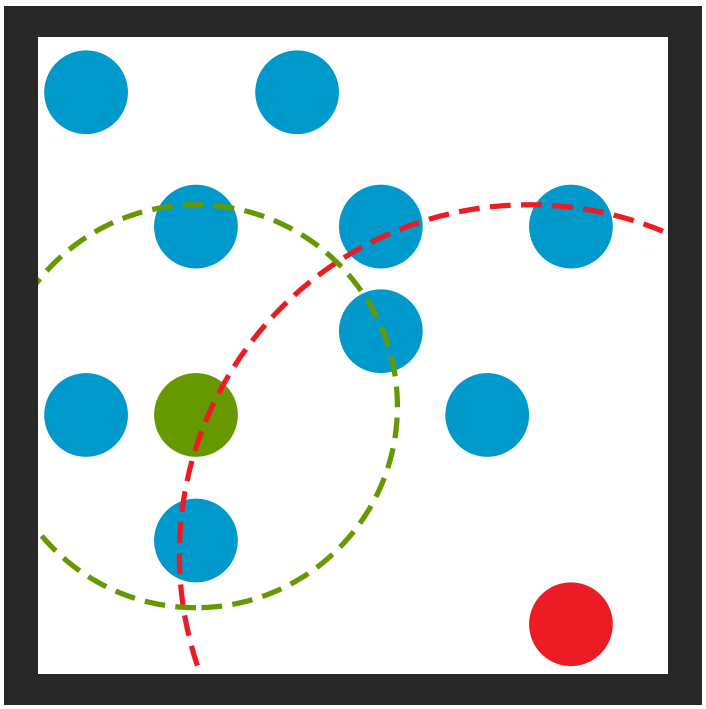

Fig. 6.3 Nodes closer to a boundary must expand their search radius to connect with k nodes.

locked clusters become far less common. Fig. 6.2 demonstrates this behaviour. Reducing this maximum distance from the boundary involves increasing the total length of boundary present in the embedding.

Cost of increasing boundary length

Modifying node arrangement to increase the total boundary length has its own side-effects. The cost depends on the connectivity pattern that is adopted by the graph. For example, if the graph has random connectivity, the average distance between any two nodes increases as the diameter of the rectangle is larger.

For modular connectivity, which is more biologically realistic, the average path length increases for a different reason. Each node tries to make a connection to the nearest k nodes, and therefore we can imagine it has a search radius that is large enough to include k nodes.

Since our model uses taxicab geometry, when no boundary exists (consider an infinite grid) and each cell contains a node with probability p, the 'circle' of radius r is the diamond given by:

$$|x| + |y| = r$$

This search radius contains $4r$ cells. We therefore expect $4pr$ nodes to be found in this search radius. However, nodes within distance r from the boundary are prone to the boundary effect, as their search radius must be higher. This means the outgoing connections from these nodes are, on average, longer. In our model, this also means potentially reduced local connectivity. Fig. 6.3 demonstrates this effect.

6.1.2 Varying Grey Matter layout

So far we have been comparing HD and SD with a 500x80 rectangle for the external enclosure. However, boundary length can be increased within any given space by laying out the grey matter in different ways. The advantages of having fewer landlocked nodes is potentially greater global connectivity, given the same density.

In order to compare different grey matter layouts on a two-dimensional grid, we perform the following experiment, **for each type of layout**:

1. Create a 200x200 grid and embed 2300 nodes.
2. Randomly place 100 hub nodes in vacant cells in the grey matter.
3. Use the k-Nearest connectivity pattern on the non-hub nodes, attempting to connect each node to its 4 closest neighbors.
4. Sort hub node pair combinations by the distance separating the two hub nodes, and create a new list which only contains the second half of the sorted combinations.
5. In a random order, attempt to embed connections between hub nodes in this new list.
6. Repeat for different amounts of white matter.

This is very similar to the experiments carried out in chapter 4. We only consider hub node pair distances which are in the longer half to ensure relatively long-range connections are being made. We do not use an absolute threshold length since different layouts mean the distances between hub nodes (which are situated in the grey matter) are skewed.

We consider three different layouts: horizontally sheeted, as used in chapter 4 (not centered) but on a square grid, and two of the other layouts described in the background: inner core (IC) and external sheet (ES). The results of this experiment are shown in Fig. 6.4. Visual examples of the layouts, post connection embedding, are shown in Fig. 6.5.

Results

- The number of local connections is very similar across all three layouts, with IC having the most and EC having the least.
- Local wiring becomes shorter as white matter is introduced and nodes are more densely packed. Beyond 60%, local connectivity begins to fall, leading to a sharper fall in the mean local connection length. (Discussed in previous chapter).
- ES demonstrates superior global connectivity to the other two layouts. Its performance in this respect peaks when there is 36% white matter.

The difference in local connectivity between the layouts can be explained by the amount of boundary. Nodes by the external boundary are less able to make the connections according to the k-Nearest pattern, as they only have a maximum node degree of 3. ES has the largest external boundary (800), while IC has the smallest (0). Furthermore, local connections are slightly longer in ES, due to the boundary effect, where nodes near the external boundary and grey-white matter boundaries have to expand their search radius.

The difference in global connectivity is also explained in terms of the boundary, but this time in terms of the distance of nodes from the boundary. Initially, horizontally sheeted and IC layouts establish connections within the grey matter (unless due to area exclusion, a path utilising white matter is found to be shorter). However, as the grey matter becomes more densely packed with wiring, they begin to take advantage of the white matter. From percolation theory, as the distance of a node from its nearest boundary increases, the probability of finding a cluster that is in contact with white matter decreases exponentially (technically, this would not be an enclosed cluster since it is the same cluster as the white matter, but imagine the white matter boundary as a barrier). Since proportionally more nodes are further away from the nearest boundary in horizontal split and IC, fewer nodes can take advantage of it. Fig. 6.6 shows points where nodes are furthest from the boundary in the respective layouts. Note how in ES these points are at the corners, explaining why the corners appear to have less global wiring in Fig. 6.5b

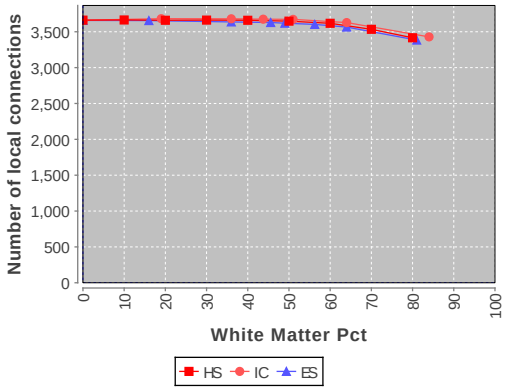

(a) Number of local connections embedded

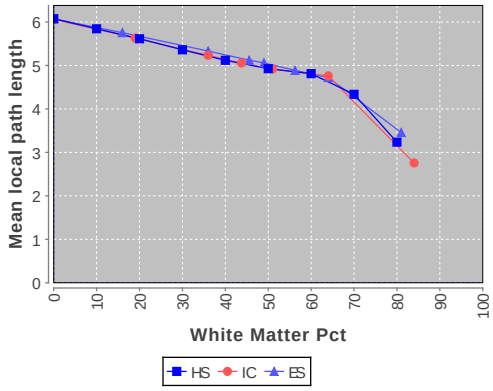

(b) Mean local connection length

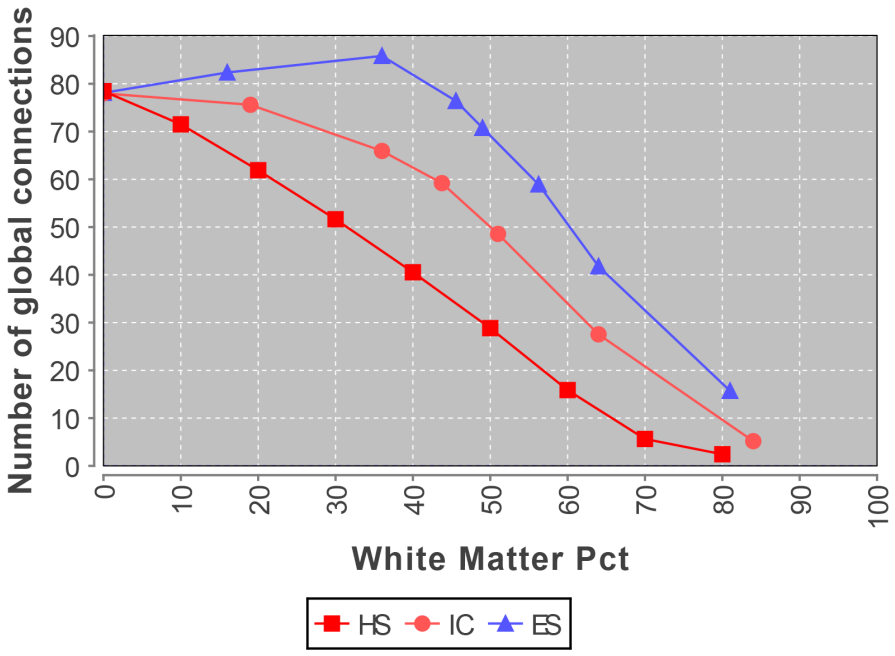

(c) Number of global connections embedded

Fig. 6.4 Results from experiments comparing horizontally sheeted, inner core (IC) and external sheet (ES) layouts

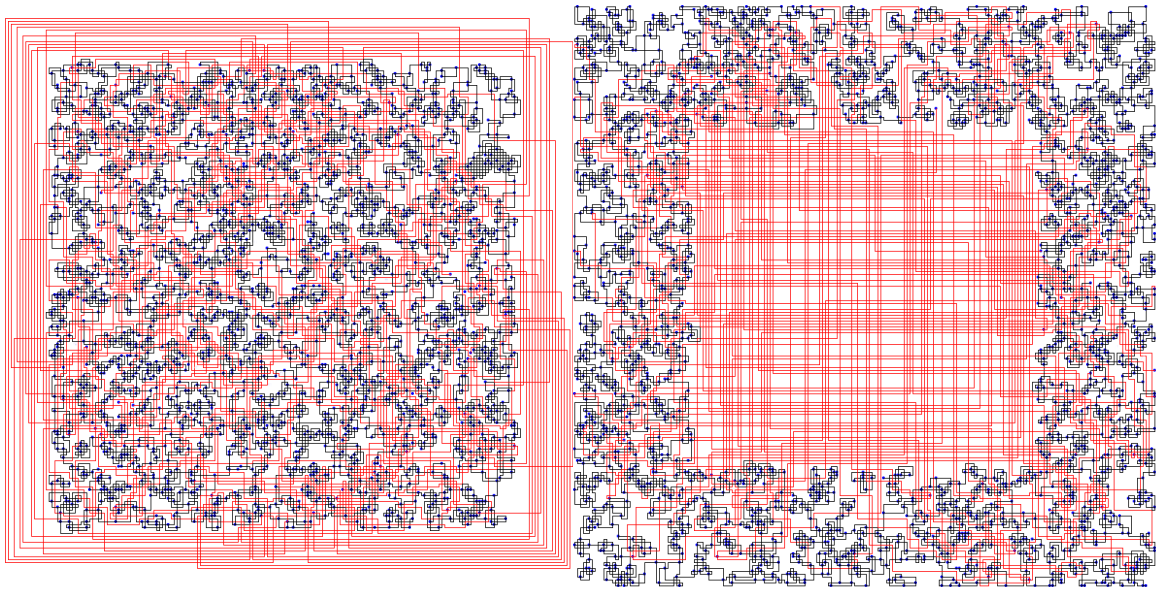

(a) IC (b) ES

Fig. 6.5 Examples of horizontally sheeted, inner core (IC) and external sheet (ES) layouts with 36% white matter

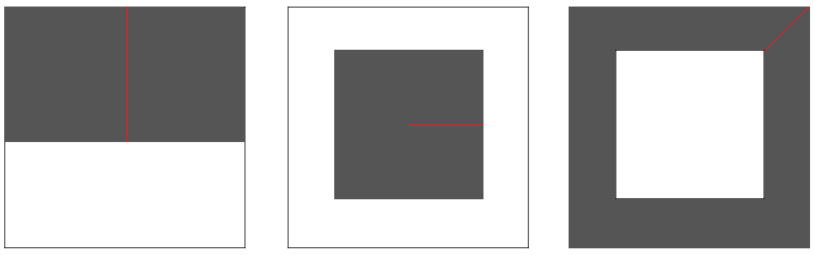

Fig. 6.6 Points where distance from nearest white matter boundary is maximal

6.2 Clusters in Modular Connectivity Networks

6.2.1 HD vs SD

We are now in a position to provide an explanation for why SD achieved a higher number of global connections than HD under certain circumstances in chapter 4. This is illustrated using a visual walkthough shown in Fig. 6.7, which demonstrates how clusters form as wiring is embedded.

We use the same experiment setup as before (500x80 grid, 2300 nodes and 200 hub nodes) and set $G_{min} = 250$, exactly half of the length of the grid. Initially, we embed all the local connections in order to optimise short conduction delays. We find that in both HD and SD, global connectivity remains, as indicated by the large red cluster (HD, 6.7a) and large green cluster (SD, 6.7c). We can also spot much smaller clusters forming, indicating regions where nodes have a limited range of connectivity.

After 25 connections have been established in this instance of HD, we detect a sudden loss of global connectivity. The 25th global connection that was embedded has left the graph with 2 large clusters, but not an infinitely large one. The situation is different in the instance of SD, where we have produced more small clusters, but a cluster of infinite size remains. Global connectivity is eventually lost in SD, albeit with more global connections being successfully embedded.

6.2.2 SD vs SD with centered grey matter

In chapter 4, we made a peculiar observation when varying the white matter percentage: When white matter covered 50% of the grid, both SD and SD with centered grey matter had embedded a similar number of global connections; but at 60% white matter coverage, a considerable gap had suddenly developed between the two (Fig. 4.9c).

A 10% increase in density of the grey matter causes a proportionally large decrease in the average cluster size of grey matter (we are considering the stage at which the connections have been embedded), of over 50%. This makes it significantly less likely a projection near the external boundary above will be able to access the white matter, hence causing the abrupt dip in the curve. When the grey matter is centered, on the other hand, the maximum distance from the boundary is halved and far more connections end up being embedded (Fig. 6.8c).

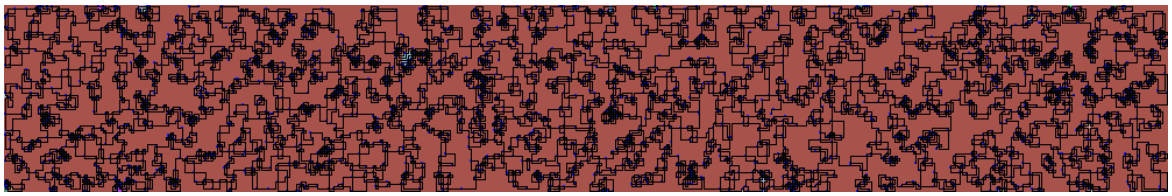

(a) HD, only local connections embedded

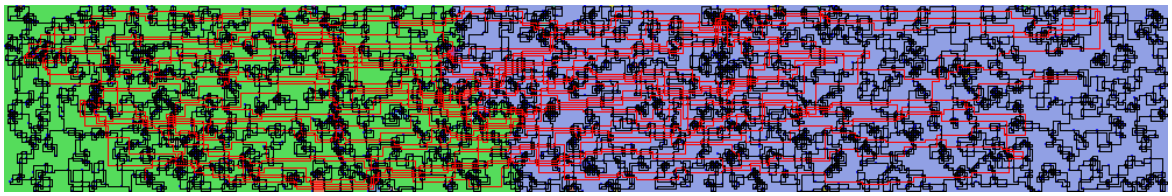

(b) HD, 25 global connections embedded

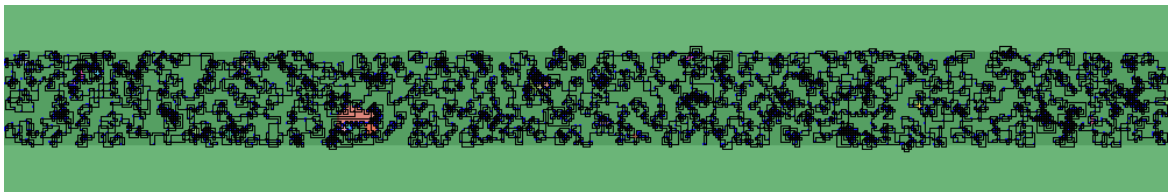

(c) SD, only local connections embedded

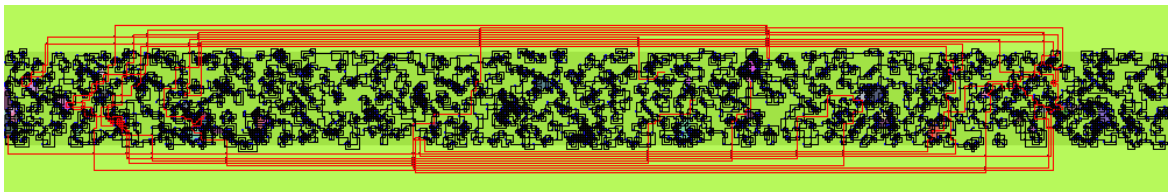

(d) SD, 25 global connections embedded

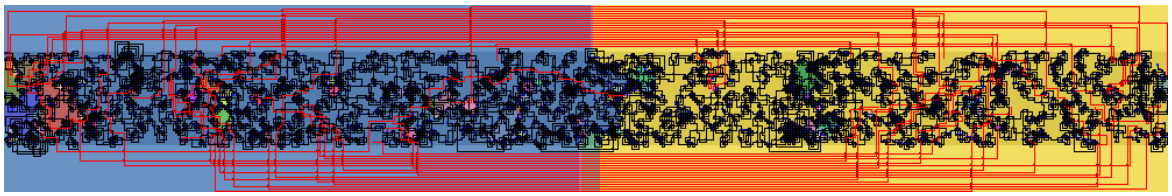

(e) SD, 37 global connections embedded

Fig. 6.7 Progressive cluster formation in HD and SD for modular connectivity

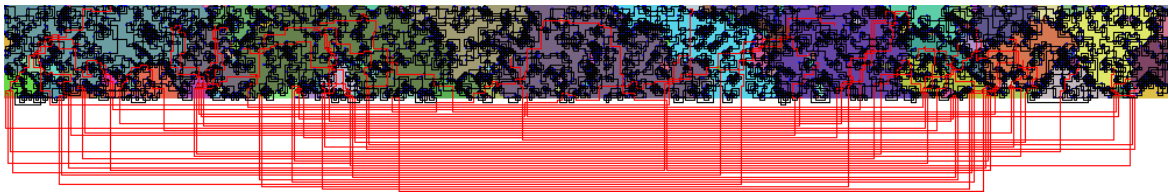

(a) SD, 50% white matter

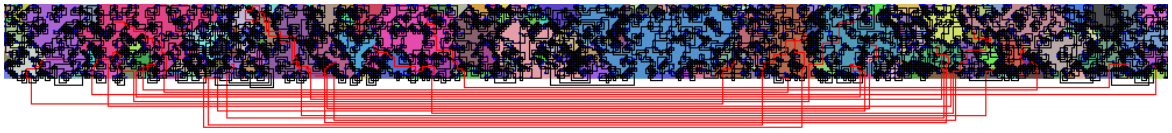

(b) SD, 40% white matter

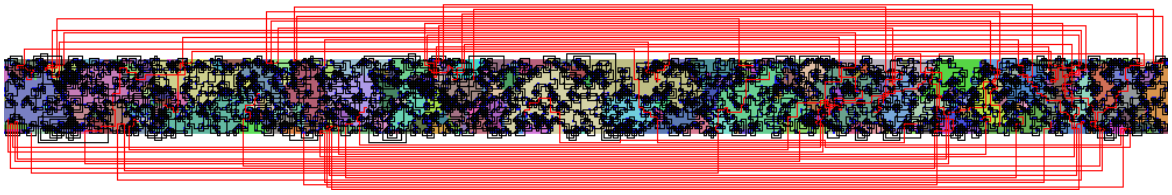

(c) SD (centered), 40% white matter

Fig. 6.8 Resulting clusters after an embedding trial with $G_{min} = 250$. Here the white matter boundary is treated as a barrier, and isn't included in any clusters.

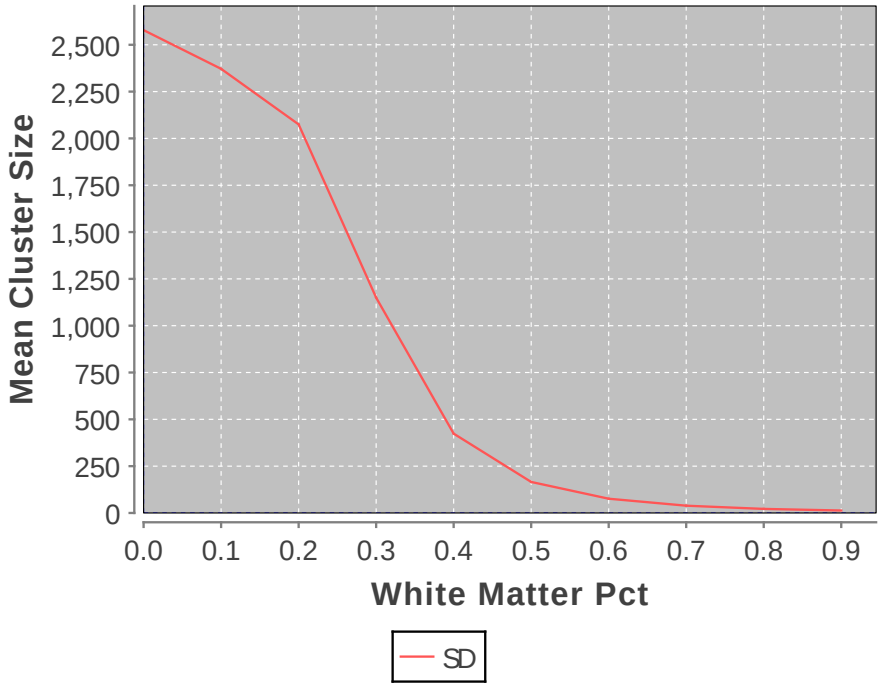

Fig. 6.9 Mean Cluster Sizes (Cluster sizes of size 1 are excluded from the calculation)

Mean cluster size

We can see how the mean cluster size of the grey matter changes as we introduce more white matter. We run the same experiment as we did in chapter 4: Create a 500x80 grid, embed 2300 nodes and 200 nodes and establish local and global connections. After we have exhausted the attempts to make global connections, we are left with our final embedding. At this point, the mean cluster size is taken, and only the cells in the grey matter are considered. The results of the experiment for different values of white matter percentage are shown in Fig. 6.9. Note that: (a) When white matter is low, we have big clusters left over, but global connectivity is lost (b) When nodes are closer together, a cluster of some size s is more valuable than one of the same size in an embedding where nodes are further apart,

6.3 Distance-Based Spatially Embedded Networks

In chapter 4, it appeared SD resulted in a loss of connectivity in networks with random connectivity, yet when the network had a modular topology, connectivity was potentially higher than HD. The difference between the two boils down to the wiring length distribution; the former involves embedding wiring of various lengths, whereas the latter predominantly

consists of very short wiring, with a handful of much longer connections.

In this section, we consider how HD and SD compare when we try to embed the distance-based spatially embedded networks in the manner described in the background. In principle, segregating grey and white matter should be suited to these, since it produces a large number of short range connections and relatively fewer long range connections.

6.3.1 Challenges

Mismatch between method and model

Unlike the modular topology, the method used to create distance-based spatially embedded networks presented a much bigger challenge. The problem is that the method generates far more connections than can possibly be embedded, as the algorithm will consider every node pair combination as a potential candidate. Consider the case where $h = 0$, i.e. the graph is completely random, then the method will generate $n(n-1)/2$ connections (assuming undirected connections) based off nodes on a grid, which is far beyond what can be embedded in the remaining space. As h increases, this value falls, but it would still not be possible to embed a large fraction of the connections, unless h is extremely high. This is not an issue if wiring and the degree of nodes are not constrained, as is the case in many models, but these spatial constraints are the very thing we are interested in.

Therefore, it was decided that we would place a limit on the number of node pairs generated by the spatially embedded algorithm. This, together with the fact that no node can exceed a degree of 4, potentially influence the modular characteristics of the network. However, we are more concerned with the connected node pair distance distribution, which would still be represented by the small sample of generated pairings.

Making a direct comparison

Since the spatially embedded algorithm generates a different connectivity matrix on every run, and the networks generated are based off the spatial location of the nodes themselves, it is a difficult challenge to compare HD and SD. If, for example, we were to compare a spatially embedded HD grid with a spatially embedded SD (ES) design (where nodes are placed around the surface of the exterior border), the logical topology would be too different to make a valid comparison.

Instead, we tried the following: First generate a random spatially embedded network (using the Manhattan Distance) from nodes which are dispersed across the entire grid area (HD). Then, reposition nodes by compressing them into a smaller region of grey matter

(ensuring relative positioning remains the same), leaving behind a region of white matter. The connectivity matrix is still the same. A more direct comparison of how the two designs are able to embed the same network is now possible.

6.3.2 The Experiment

1. Create a 500x80 grid and embed 2300 nodes randomly across the grid.
2. For each node, select another node at random and connect with probability $p = e^{-hd}$. Repeat this until each node is connected with another node.
3. Create a 500x80 grid and embed the same 2300 nodes in a (centered) grey matter region of size 500x40, by compressing nodes in the vertical axis.
4. For both grids, attempt to embed the connectivity matrix generated, in a random order.
5. Repeat for different values of h.

Results

First, we note that in terms of the total number of connections embedded, SD doesn't manage to attain the same number as HD, for all h. However, when we establish connections ordered by connection length, shortest first, this performance gap is reduced. This indicates that SD relies heavily on the long distance connections being routed through the white matter. HD, on the other hand, benefits relatively less from the order in which connections are embedded.

The number of connections embedded is fairly constant amongst all values of h. However, it appears that as the composition of short connections increases, the mean connection length becomes smaller than HD, provided they are embedded first. This is expected, as the nodes are now closer together. When h is low, the mean connection length is higher than SD, indicating that many connection lengths suffer from the overhead of exiting and reentering grey matter in order to make a connection.

Discussion

Compared to the modular connectivity topology, the case for SD is less convincing in distance-based spatially embedded graphs. When h is high and shortest connections are embedded first, it still has a substantially lower total number of connections. However, once

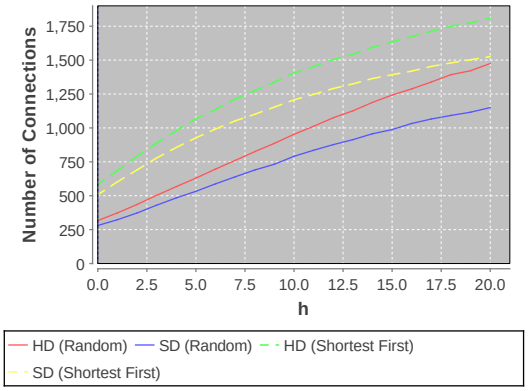

(a) Absolute difference in connectivity between HD and SD

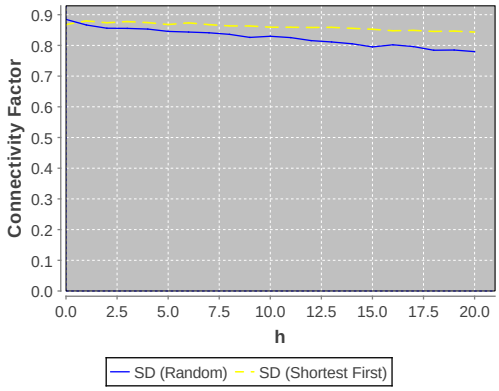

(b) The fraction of connections embedded by SD compared to HD

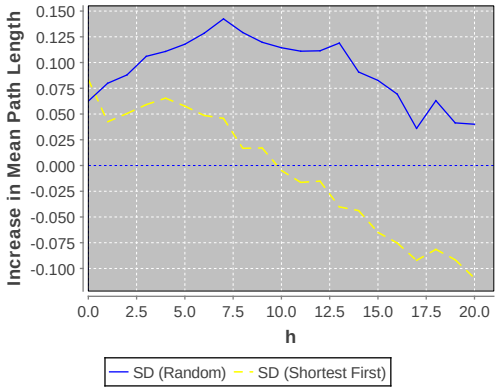

(c) Average wiring length increase in SD compared to SD

Fig. 6.10 Results comparing HD and SD when a distance-based topology is used. 'Shortest first' indicates that the connections in the connectivity matrix were embedded in order of length.

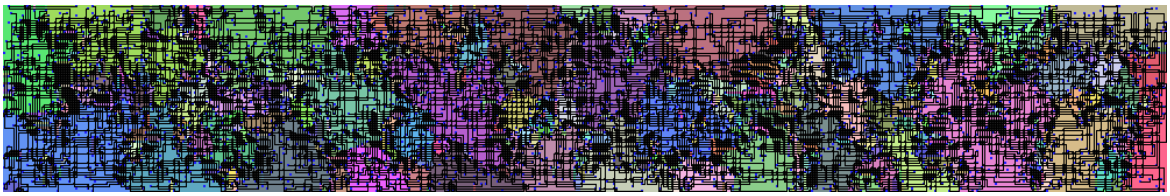

(a) HD, $h = 20$

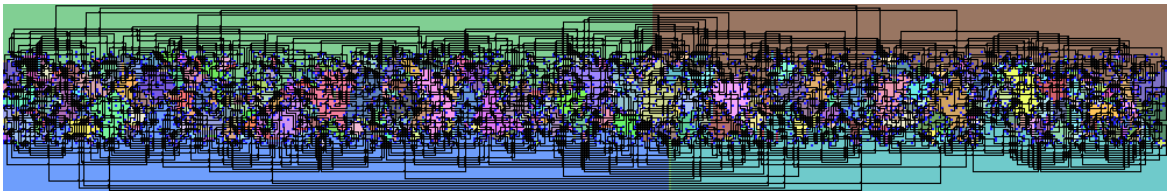

(b) SD, $h = 20$

Fig. 6.11 Resulting clusters after distance-based spatial embeddings in HD and SD. In SD, nodes maintain the same relative positioning, but are squeezed together vertically

again, analysing how clusters form in the respective designs becomes useful as it clarifies the underlying connective potential.

In SD, as we embed connections, the grey matter region becomes increasingly dense, while white matter is relatively sparsely populated. As a result, the cluster size distribution after the embedding is that of many small clusters, and a few very large clusters. In HD on the other hand, clusters form uniformly, with many 'medium' sized clusters. Fig. 6.12 shows the cluster distribution that results from the instances shown in Fig. 6.11.

If we check the difference between the networks that were embedded, we find that SD was embedding much longer connections than HD (Fig. 6.13).

The case for HD

Given the previous display of SD embedding longer-range connections, one might be led to believe it has better very long-range connective potential altogether. However, when we take the successfully embedded set of connections in SD, and attempt to embed it back in a new HD grid in the same order, we find HD can absorb the same set of connections. In other words, while SD was forced to embed the longer connections, HD simply became saturated with 'medium' range connections, meaning by the time it attempted to embed the long-range connections, the grid had no cluster big enough to accommodate them. This is not the case in certain modular networks, where HD could not absorb the same network that SD was able to embed.

Therefore, unlike the experiments carried out with the modular topology, where SD

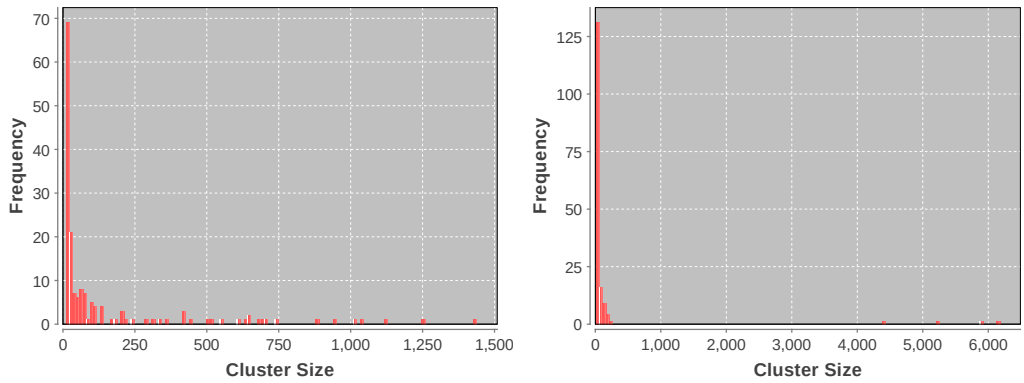

(a) HD, $h = 20$ (b) SD, $h = 20$

Fig. 6.12 Cluster size distribution after a distance-based spatial embeddings in HD and SD. Note the four large clusters in SD.

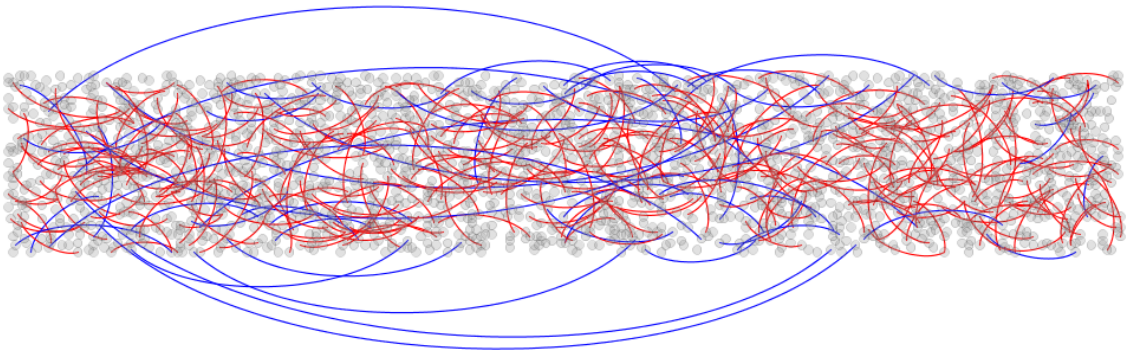

Fig. 6.13 Red edges indicate connections that were embedded in HD (6.11a) but NOT embedded in SD (6.11b). Blue edges indicate connections that were embedded in SD but NOT in HD. Visualised using Gephi [10].

performed outstandingly in terms of reduced conduction delays and increased global connectivity over HD, in this one the results are mixed. The short-range connection lengths may be reduced, but in terms of connectivity, HD is superior. While we tried to tweak various parameters for the distance-based method, we could not achieve a situation where the discrepancy between very short and very long connections was similar to that of the modular connectivity pattern.

Wiring Diameter

In our model all wiring is of the same diameter. However, if we imagine the diameter of wiring is doubled, then the maximum density that can be tolerated before long and thick connections lose connectivity is potentially much lower. It is possible that this would make SD relatively more effective than HD compared to the results in our model.

6.4 Summary

This chapter applied the concepts of clustering to our model, and used it to explain some observations previously made. Essentially, clustering analysis is useful because it allows derive a new graph based on the state of the cells on the grid, where an edge between any two nodes (representing grid cells) exist if there an open path between them on the grid. Clusters therefore represent disconnected components, giving us a sense of the connective potential that exists on the grid. We then compared HD and SD using the distance-based topology, where SD reduced the length of connections but had lower connective potential.

Chapter 7

Evaluation

This chapter first considers how the findings of the project bear relevance to the brains of mammals. Afterwards, we discuss the key challenges of the project.

7.1 Comparison with the Mammalian Brain

In the last three chapters, we established certain conditions that appeared favourable for SD. Here, we compare whether the brain's characteristics are in accordance with these conditions.

7.1.1 Width of Grey Matter

In chapter 6 we noted that given some density, the probability a neuron could find an open path to the grey-white matter boundary would fall exponentially with its distance from the boundary. Therefore, a thin layer of grey matter is most favourable for achieving long-distance connections in SD.

The vertebrate brain appears to have this characteristic. Grey matter is a relatively uniform and thin structure (Fig. 7.1a), around 2 to 4mm thick in the Cerebral cortex [19]. In order to pack in a vast number of neurons, cortical folding occurs, which explains the wrinkled appearance of the brain (Fig. 7.1b). This maximises the surface area between the grey and white matter regions, which is similar to maximising the boundary length in our two-dimensional model. It therefore could be that, in trying to densely pack neurons as closely together as possible (perhaps in order to reduce local inter-neuronal distances), the thickness of the cortex was minimised to maintain good global connectivity in light of dense grey matter.

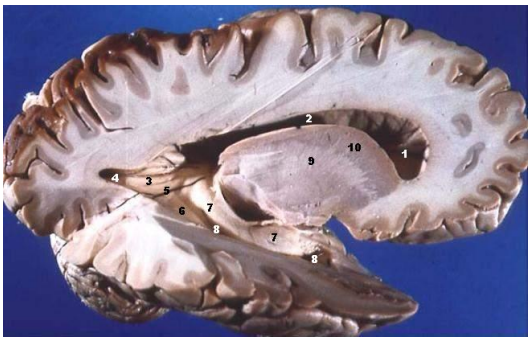

(a) Thin layer of grey matter. Source: [32]

(b) Cortical Folding. Source: [30]

Fig. 7.1

7.1.2 External Sheet Architecture

When we compared various ways of organising grey and white matter on a square grid, we considered an Inner Core architecture (EC) and its mirror-image, the external sheet (ES) architecture. We found that the reduced thickness of grey matter in ES enabled it to achieve a greater global connectivity than IC. The brain also has the grey matter exterior to the white matter (Fig. 7.1a).

7.1.3 Distribution of Wiring Lengths

SD was effective in embedding the modular topologies, provided the distance between connected hub nodes was large enough to make efficient use of the white matter, and resulted in higher global connectivity while reducing conduction delays. However, the distance-based spatial embeddings saw a trade off between conduction delays and connectivity. It is worth speculating that wiring length distribution in the mammalian brain does not closely resemble the distribution of the distance-based topology, since regions of the brain are highly connected with specific regions. Although no such data is available for large mammalian brains, in the C. elegans the wiring length distribution of global connections that would suit our SD model more than that created by the distance-based topology method (Fig. 7.2).

7.1.4 Favouring Linear Trajectories

In chapter 5 we saw that in our model, the fewer turns connections made, the more densely packed wiring could be before global connectivity was lost. We speculated that by giving wiring a more organised structure, future connections would be more likely to find large contiguous parallel connections that they can cross, whilst when connections are less linear

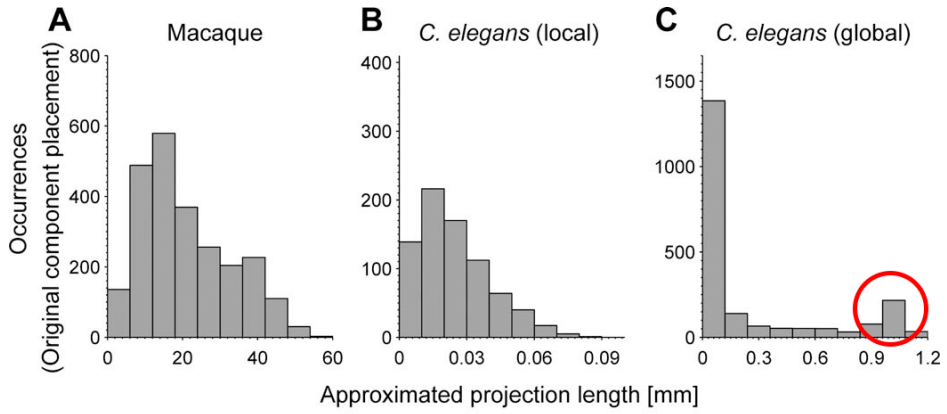

Fig. 7.2 Approximated global projection length distribution in C. elegans. Source: [31]

and more complex, these large crossable regions would disappear. The axons of pyramidal axons are regular and follow linear trajectories [5], while white fibers travel in bundles. In our model, this is advantageous for maximising connectivity.

7.1.5 Connection Ordering

The results in chapter 4 for modular networks were arrived at by establishing the shortest connections first, the motivation being that these were the connections that needed to be optimised for minimal length. Although this may not be the case in development of the brain, this was simply one way of achieving the embedding we did, which is where the long range connections make use of the white matter.

7.2 Main Challenges

One of the main challenges of the project was finding an area to focus on. Initially, a lot of background research on the brain was carried out, and we considered novel approaches to the problem such as the relative ease for growth cones to find targets. Unfortunately, the results were rather inconclusive, so we went back to the background studies for inspiration. By testing our model against some of the conclusions drawn by Wen and Chklovskii [31], we were led to the idea of the possible application of percolation theory to the problem.

Another challenge was dealing with the stochastic nature of our model. Since almost every aspect of the model is non-deterministic, making the discoveries was generally a trial and error process, making small changes to the model and observing the differences to make broader generalisations. A particular difficulty was reading behind the lines of experimen-

tal results. For example, in terms of absolute connectivity in the distance-based topology embedding, HD was found to be superior, but the results did not directly highlight the finer intricacies like the composition of connections being made.

Making comparisons between HD and SD as 'fair' as possible was challenging. Fundamentally, changing the physical packing strategy of network components could directly introduce a bias into our results. This was particularly the case when random and distance-based spatially embedded network topologies were considered. In the case of distance-based networks, changing the node placement and using the method to create the connectivity matrix would inevitably result in networks with different logical topologies. Our strategy to overcome this was to create the logical connectivity based off a HD layout and 'compress' nodes into a horizontally sheeted SD (maintaining relative distances). However, the same method could not be applied to IC and ES variants of SD, since relative positioning cannot be maintained.

Finally, although we have produced empirical results, we were not able to establish more formal results that could are more generalised. Our results are based on specific parameters, so there is still a void in understanding behaviours across a broader range of parameters. It was hoped that existing research from percolation theory could be applied, but unfortunately most of the research is concerned with finding critical values for different lattices, which is not particularly relevant to our model, which is less abstract. It did, however, become a source of inspiration that enabled us to speculate on the possibility that grey-white matter segregation exists to permit a higher density of components in grey matter whilst retaining longer distance connectivity.

Chapter 8

Conclusions and Future Work

8.1 Conclusions

In this project we created a concrete model that could be used to model spatial constraints on neural networks. Using this, we found that SD could potentially be a more spatially efficient network packing structure than a HD, depending on the topology of the network. This was the case in certain cases of modular networks and distance-based spatial embeddings. In modular networks, we also found that it could also potentially yield higher global connectivity, given an embedding in which local connections were already present.

We also discussed some of the underlying routing issues in terms of percolation theory. While the area taken up by wiring is important, the potential to successfully embed further connections also depends on the spatial arrangement of network components. It is possible that, in a relatively sparsely filled embedding, global connectivity is lost as over these long ranges there is almost no probability of an axon percolating. However, an axon has a good chance of making it through dense matter if the distance is short. Therefore,

There are therefore multiple factors in play, even in our simplified model. A segregated design can reduce the total length of wiring (through closer node cell bodies (short-range wiring) and reduced impact of area exclusion (long-range wiring)), provided the white matter is used efficiently. This can result in more free space being left over, which leaves more vacant cells for future connections. Reducing the amount of long-range wiring present in grey matter also means that it contributes less ‘coverage’ in the region, enabling the area to be more densely packed with other components. Finally, long-distance connections are capable of being embedded, despite the higher density within grey matter, due to the ability to exploit white matter.

8.2 Further Work

- Consider building an equivalent three-dimensional model of the system. Percolation theory also applies to three-dimensional cubic lattices, which is an extension of our grid based model in the third dimension.
- Build a more realistic model where distances are continuous and wiring is not limited to orthogonal movement.
- Consider the impact of changing wire diameter. White matter fibers are generally thicker [29] than unmyelinated ones, so the factor is significant.
- Consider dynamic networks where wiring is constantly changing. Although our grid was constantly being grown, it did not consider the effects of a dynamic rewiring process.
- Establish results which cover a broader range of scales and dimensions, by using algorithms which scale better.
- An optimisation-based approach could be applied to our model, where all the combinations are tested against optimisation criteria (wiring length, global connectivity, etc.) to arrive at a more quantitative conclusion.

References

[1] Abboud, N., Grötschel, M., and Koch, T. (2008). Mathematical methods for physical layout of printed circuit boards: an overview. *OR Spectrum*, 30:453–468.

[2] Austin, D. (2014). Percolation: Slipping through the cracks. http://www.ams.org/samplings/feature-column/fcarc-percolation.

[3] Azevedo, F. and et al. (2009). Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *The Journal of Comparative Neurology*, 513:532–541.

[4] Bullmore, E. and Sporns, O. (2012). The economy of brain network organization. *Nature Reviews Neuroscience*, doi:nrn3214.

[5] Callaway, E. (2004). Close encounters: How cortical neurons find and connect to their correct synaptic partners depends on the cell type. *Neuron*, 43(2):156–158.

[6] Chedotal, A. and Richards, L. (2010). Wiring the brain: The biology of neuronal guidance. *Cold Spring Harb Persepect Biol 2010*, 2:a001917.

[7] ChklovskiiLab (2014). Wiring economy. (Online). https://openwiki.janelia.org/wiki/display/chklovskiilab/Wiring+Economy.

[8] Dechter, R. and Pearl, J. (1985). Generalized best-first search strategies and the optimality of a*. *Journal of the Association for Computing Machinery*, 32:505–536.

[9] Gastner, M. (2014). Percolation theory. (Online). http://www2.imperial.ac.uk/ mgastner/percolation/percolation.html.

[10] Gephi (2014). An open source graph visualization and exploration platform. (Online). https://gephi.org/.

[11] Gilbert, E. (1959). Random graphs. *Annals of Mathematical Statistics*, 30(4):1141–1144.

[12] Goldman, B. (2010). New imaging method developed at stanford reveals stunning details of brain connections. (Online). http://med.stanford.edu/news/all-news/2010/11/new-imaging-method-developed-at-stanford-reveals-stunning-details-of-brain-connections.html.

[13] Goodman, C. and Shatz, C. (1993). Developmental mechanisms that generate precise patterns of neuronal connectivity. *Cell*, 72:7798.

[14] HCP (2014). Human connectome project: Official website. (Online). http://www.humanconnectomeproject.org/.

[15] Jacobsen, J. (2014). High-precision percolation thresholds and potts-model critical manifolds from graph polynomials. *Journal of Physics A: Mathematical and Theoretical*, 47(13):135001.

[16] Janko, A. and Janko, O. (2014). Rätsel. (Online). http://www.janko.at/Raetsel/Rules.htm.

[17] JFreeChart (2014). Java chart library. (Online). http:// www.jfree.org/jfreechart.

[18] Kaiser, M. and Hilgetag, C. (2006). Nonoptimal component placement, but short processing paths, due to long-distance projections in neural systems. *PLoS Comput Biol*, 2:e95.

[19] Kandel, E., Schwartz, J., and Jessell, T. (2000). *Principles of Neural Science*. McGraw-Hill, fourth edition. p. 324.

[20] Latora, V. and Marchiori, M. (2001). Efficient behavior of small-world networks. *PhysRevLett.87.198701*, 87(198701).

[21] Lieff, J. (2012). Neuronal connections and the mind, the connectome. (Online). http://jonlieffmd.com/blog/neuronal-connections-and-the-mind-the-connectome.

[22] Murre, J. and DPF, S. (1995). The connectivity of the brain: Multi-level quantitative analysis. *Biol. Cybern.*, 73:529–545.

[23] Oh, J. (2014). Mapping the human brain to understand the human mind. (Online). http://thinktank.uchicago.edu/blog/2014/1/8/mapping-the-human-brain-to-understand-the-human-mind.

[24] Ruppin, E., Schwartz, E., and Yeshurun, Y. (1994). Examining the volume-eciency of the cortical architecture in a multi-processor network model. *Biol. Cybern.*, 70:89–94.

[25] Schüz, A. (2008). Neuroanatomy. *Scholarpedia*, 3:3158.

[26] Sedgewick, R. and Wayne, K. (2011). *Principles of Neural Science*. Addison-Wesley, fourth edition. p.216.

[27] Shanahan, M. (2013). Imperial college computational neurodynamics course notes.

[28] Skiena, S. (2008). *The Algorithm Design Manual*. Springer, second edition. p.389.

[29] Swadlow, H. and Waxman, S. (2012). Axonal conduction delays. *Scholarpedia*, 7:1451.

[30] Wei, W. (2014). Picture of cortical surface. (Online). http://cheba.unsw.edu.au/project/morphology-cortical-surface-cortical-folding-pattern-and-sulcal-width.

[31] Wen, Q. and DB, C. (2005). Segregation of the brain into gray and white matter: A design minimizing conduction delays. *PLoS Comput Biol*, 1:e78.

[32] Wikipedia (2014). Dissected human brain, lateral view. (Online). http://en.wikipedia.org/wiki/White_matter.

Appendix A

Cherry Picking when Embedding Random Networks

In chapter 4, we noted that if we increase the sample size of random edges, the embedding itself becomes less random as the model has the luxury of picking from a vast selection of edges. Fig. A.1 shows how this effects the results: By the end of the experiment, the model which had the highest failure rate of embedding connections had managed to embed the most.

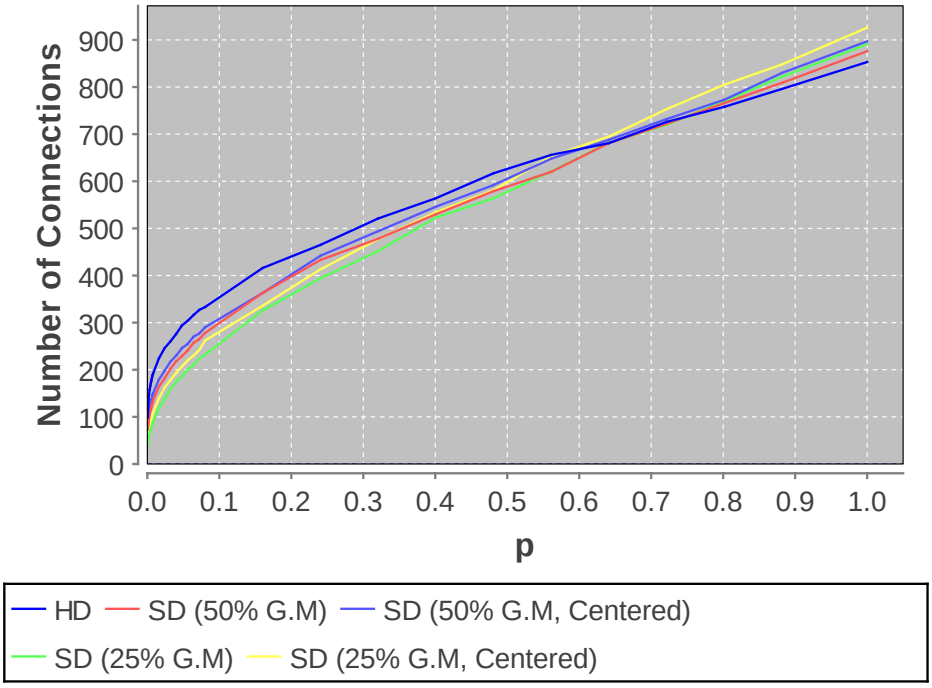

Fig. A.1 Random Connectivity on a 100x100 grid with 500 nodes.

Appendix B

Higher Connectivity when Paths are Straight

In Chapter 5 we performed percolation trials using growth cones. We noticed that the straighter connections were, the greater the amount of wiring 'coverage' that could be tolerated before an infinite cluster would disappear. In our model, minimising the number of turns made by connections produces higher levels of connectivity under all circumstances. Fig. B.1 shows how the A* algorithm adopted by the model (which guarantees the straightest path) produces more connections given a random network (100x100 grid, 500 nodes).

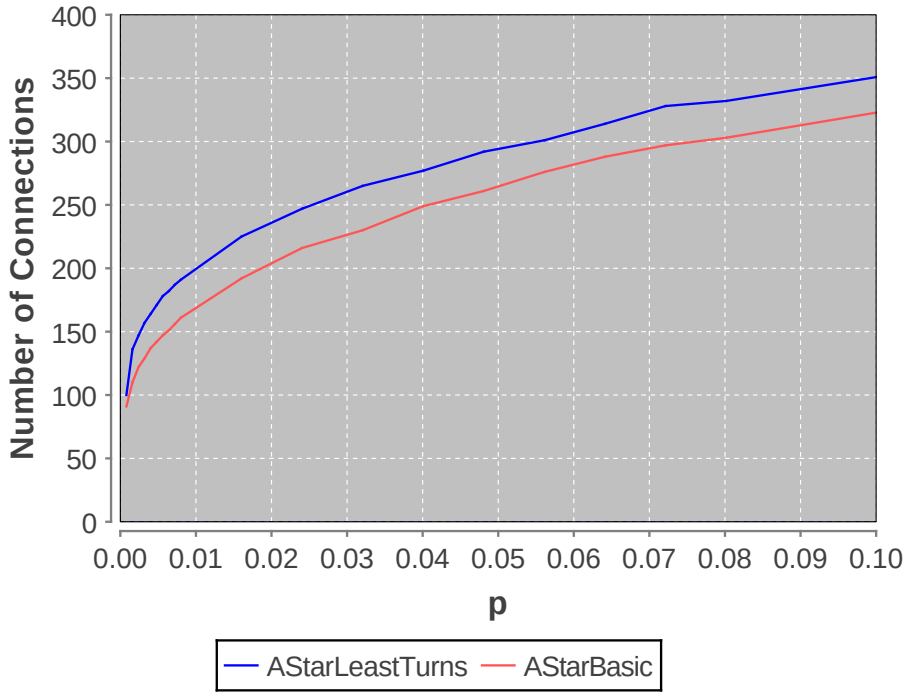

Fig. B.1 Comparing the two A* implementations on random networks

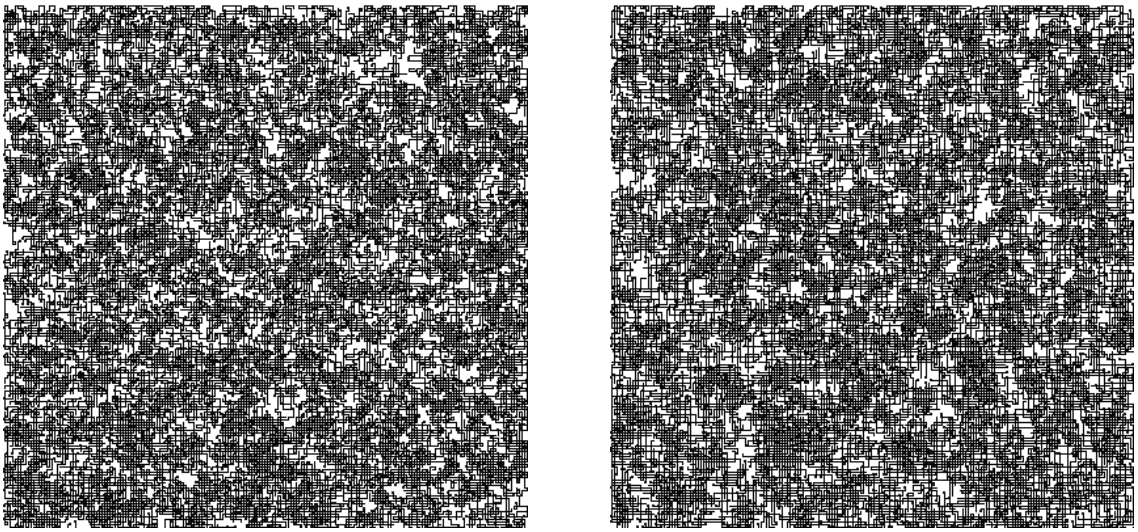

(a) $s = 0.55$ (Fail)

(b) $s = 0.65$ (Success)

Fig. B.2 Example growth cone percolation trials, $N = 200$, coverage = 1.35. s is a measure of how straight growth cones grow.