

Self-organisation can generate the discontinuities in the somatosensory map

Tom Stafford*, Stuart P. Wilson

Adaptive Behaviour Research Group, Department of Psychology, University of Sheffield, Sheffield S10 2TP, UK

Available online 15 November 2006

Abstract

The primary somatosensory cortex contains a topographic map of the body surface, with two notable discontinuities—the representation of the face is next to that of the hands, and that of the feet is next to the genitals. Farah [Why does the somatosensory homunculus have hands next to face and feet next to genitals? a hypothesis. *Neural Computation* 10(8) (1998) 1983–1985] has suggested that these discontinuities are due to the mechanisms of self-organisation which underlie cortical map development. The typical position of the foetus in the womb means that these two pairs of body parts will often touch and hence their representations will be simultaneously co-active, even though they are distal in terms of the body surface. We use the Kohonen self-organising map algorithm to provide an existence proof of the plausibility of Farah's hypotheses. We then use the model to test the viability of other possible causes of the known map structure and to explore the limitations of self-organisation for explaining the features of the somatosensory map. The model shows that (a) the Kohonen algorithm requires high frequencies of co-activation to introduce a selective discontinuity into the map, and (b) that higher frequency of separate activation of the critical patterns alone is not sufficient to generate the selective discontinuity and (c) the consistency of near-optimal map formation, and in particular the medial–lateral ordering, cannot be reliably generated by a simple Kohonen algorithm.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Model; Somatosensory map; Self-organisation

1. Introduction

The human somatosensory cortex contains a topographic map of the body surface, with more sensitive body parts having a larger amount of the cortical surface devoted to their representation. As is characteristic of cortical maps, adjacent body parts are, in general, represented by adjacent cortical areas. There are, however, two notable discontinuities in the somatosensory map. Firstly, the area representing the face adjoins the area representing the hands. Secondly, the genital area adjoins the feet area (although see [11], discussed below).

Farah [6] proposes that the arrangement of the somatosensory map, and specifically the existence of the discontinuities in the map, is due to the prenatal position of the foetus during map development. Specifically, she proposes that the arrangement is due to self-organisation

of the map under conditions where co-activation of those two pairs of body parts is more probable, because in the prenatal fetal position these are the parts most likely to touch each other.

We have investigated this hypothesis using self-organising Kohonen networks [12] which utilise a Hebb learning rule [9] to create self-organised topographic maps. Simply, we ask, using a grossly abstracted model of the mapping between body surface and cortical representation, is it possible for co-activation of distal body parts to create a selective discontinuity in the topographic map? How sensitive is map formation to the frequency of co-activation of these distal parts? The Kohonen algorithm is the core of a fundamental class of self-organising network: are there elements of the organisation of the somatosensory map which it does not explain?

Our motivation is not to capture all details of the physiology of cortical somatotopic map development, but to test the Farah [6] hypothesis. A working simulation can merely provide an existence proof that self-organising

*Corresponding author.

E-mail address: t.stafford@shef.ac.uk (T. Stafford).

processes have the potential to explain the discontinuities in the somatotopic map. Additionally, modelling provides a structured framework for investigating the interaction of the principles of self-organisation with the particular details of somatotopy.

A recent report has claimed, counter to the assumption of Farah [6] and all those following the convention established by Penfield and Rasmussen [15], that the location of the representation of the penis in the male is not such as to create a discontinuity in the somatotopic map [11]. This single study alone, using a single methodology, is not sufficient cause to abandon the conventional map—not least because of the known capacity of the somatosensory map to dynamically reorganise according to input and task-dependent/attentional factors [1].

This finding does, however, suggest a useful simplification of our modelling effort. Our purpose is to investigate how principles of self-organisation, as embodied by the Kohonen algorithm, can lead to map discontinuities, not to provide a full account of the structure of the somatotopic map. Given this, and given the high level of abstraction at which we are working, it is as appropriate to investigate how the principles of self-organisation can lead to just *one* major map discontinuity rather than two. Kell et al.'s [11] report suggests that only the discontinuity at the representation of the face and the hands should be the focus of our initial investigation.

2. Model specification

The model inputs come from a sheet of 12 by 12 input units. These represent the sensory inputs at the skin surface and are universally connected to a sheet of 12 by 12 output units, which represent the neurons in the primary somatosensory cortex. These dimensions are arbitrary and chosen merely to allow reasonable simulation run times.

The patterns that make up the basic training set are generated by selecting a random coordinate in the 12 by 12 input space and activating the corresponding input unit and those immediately surrounding it. All input patterns are 1 normalised.

The standard Kohonen algorithm [12,8], using a Hebb learning rule, and a linearly decreasing neighbourhood function is employed. A conscience mechanism [3] was used to finesse map formation ([10] discusses why map formation can be unreliable without such a mechanism). Sources of variability in the model outcome are limited to, firstly, the initial weights of the network which are small and random, and, secondly, the random selection of training patterns.

These dynamics, and the basic training set, lead to the formation of topographic maps which approach criteria for map optimality: *coverage universality* (all output units respond to some input), *coverage uniformity* (equal representational space is given to all inputs), *completeness* (all inputs are represented by some output unit) and map




Fig. 1. Fishnet plot showing typical mapping generated by the basic training set. Each unit in the output layer is represented by a dot, which is shown positioned according to which input it responds most strongly to. Lines connect those units which are immediately adjacent in the output layer.

continuity [2]. This can be seen in the plot shown in Fig. 1. Note how the plot highlights a quirk of this instantiation of the Kohonen algorithm whereby those inputs just inside the edges do not have units tuned to them.

3. Results

To investigate the Farah [6] hypothesis we investigated networks trained on altered training sets which incorporate training patterns representing the co-activation of face and hand inputs. Rows 12 and 6 on the sheet of input units are denoted the rows representing 'face' and 'hands' inputs respectively. For 'co-active' patterns, when the randomly selected training pattern has a centre of activity which falls on either of these rows, the corresponding unit in the other row, and its surrounding units, are also active. The probability of co-activation was systematically varied between 0 and 1 to gain a more nuanced understanding of how co-activation during self-organisation affects representation co-localisation (and thus the map necessarily containing the critical discontinuity). Fifty networks were trained at each level of co-presentation probability.

Co-localisation of the 'hands' and 'face' representations was measured in the following way. The component inputs of the critical patterns were presented and the maximally responding units for each input were recorded. Thus two sets of maximally responding units were identified, one set for each critical pattern. The shortest distance between any two maximally responding units, where one came from each of the two sets, was taken as the distance measure.

The simulations show that co-presentation of the critical patterns can lead to co-localisation in the resulting representational maps. Increased probability of co-presentation leads to an increased frequency of co-localisation (see Fig. 2), such that above an 0.85 probability of co-presentation the representations of the critical patterns are nearly always adjacent. Interestingly co-localisation of representations does not increase significantly just before this critical frequency.

As a control condition, we investigated the effect that increasing frequency of presentation of the critical patterns, separately rather than simultaneously, has on the frequency of co-localisation of their representations. This gives a baseline for the effect of mere stimulus preponderance against which to compare the first manipulation. A rationale for this is that the face and hands also command the largest representational area in the somatosensory map as well as being the site of the major map discontinuity. Are these two facts related? By increasing the frequency of the presentation of the face and hands inputs, which increases the number of output units that come to respond to these inputs, we can test whether stimulus preponderance alone can generate co-localisation, or whether co-activation in particular is critical.

In order to encourage the formation of the basic topographic map—it is nonsensical to look for discontinuities if there are no underlying elements of topography to start with—the absolute frequency of the presentation of the noncritical patterns was kept constant. This necessarily means that as the frequency of presentation of the critical patterns increases the absolute number of patterns the model is trained on increases.

More frequent presentation of the critical patterns separately does not result in more frequent co-localisation in the resulting representational maps (see Fig. 3). In fact, at first elevated stimulus preponderance increases, rather




Fig. 2. Co-presentation of patterns leads to their co-localisation in the somatotopic map, but only consistently at high frequencies. Standard error bars are shown.

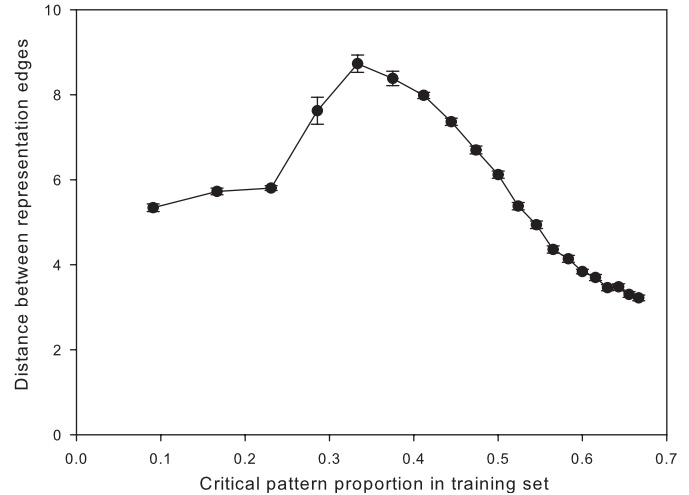


Fig. 3. Presenting the critical patterns more frequently, but not together, does not lead to map co-localisation, even at very high frequencies. Standard error bars are shown.

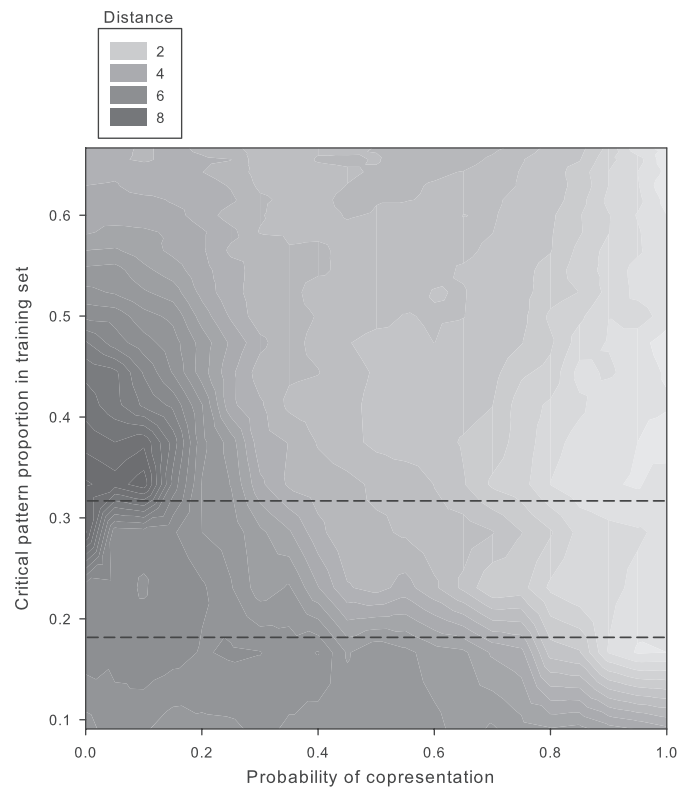


Fig. 4. Shaded contour plot showing average minimum distance between representations for the critical patterns as probability of co-presentation and frequency of presentation are varied in combination. Each point is the average of 20 simulations. Dashed lines show the critical pattern proportion values used for Fig. 5.

than decreases, the minimum distance between the representations of the critical patterns.

Fig. 4 shows the interaction of these two factors, using averages from 20 simulations at each point. Note that an increase in the frequency of stimulus preponderance reduces the minimum probability of co-activation that is

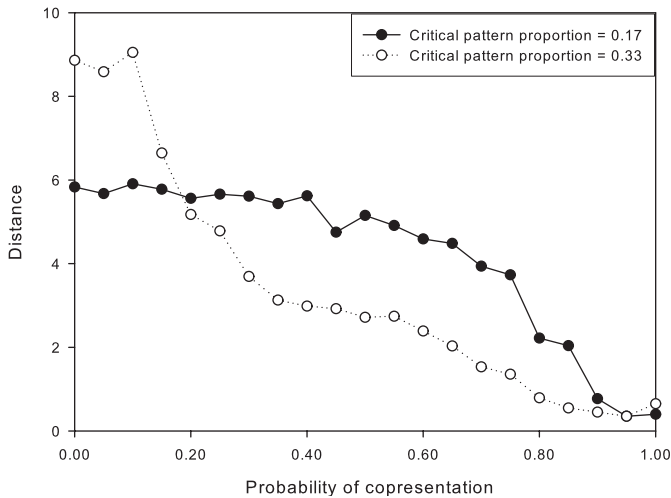


Fig. 5. The minimum required probability of co-activation is reduced as the critical patterns become more common in the training set. Each point is the average of 20 simulations.

required to generate co-localisation (shown as the lightest shading in Fig. 4). This can be seen clearly if selective data from Fig. 4 are redrawn. Fig. 5 shows line plots of the minimum average difference between the representations of the critical patterns compared against all levels of the probability of co-activation (as shown in Fig. 4), but for only two levels of stimulus preponderance. The first of these, a critical pattern proportion of 0.17, is the frequency in the basic training set (the critical patterns are 2 rows out of 12. $2/12 = 0.17$ to 2 d.p.). The second is a frequency of critical patterns approximately twice that of the basic training set.

4. Discussion

Our simulation results confirm that self-organisation, in combination with co-active inputs—as is plausible given the prenatal fetal position—can lead to selective discontinuities in the somatotopic map. Figs. 6 and 7 illustrate the change that co-presentation of the critical patterns introduces. Fig. 6 shows the distribution of the units which maximally respond to the critical patterns in a typical, well-ordered, map. Fig. 7 shows how co-activation of these patterns during training of the map results in the co-localisation of the units which respond maximally to them. This is an existence proof of the plausibility of the Farah [6] hypothesis, although it in no way demonstrates that other processes are not or could not lead to known map structure.

Further to this result, a higher frequency of co-presentation results in co-localisation in a higher proportion of resulting maps. Increasing the frequency at which the ‘face’ and ‘hand’ patterns are separately presented to the model does not produce co-localisation of representations. Instead, initially, the increase in the preponderance of the critical stimuli in the training set *increases* the minimum distance between their representations. This can




Fig. 6. A typical distribution of the maximally responding units for each of the inputs of the two critical patterns for a well-ordered topographic map (‘face’ shown by unfilled circles, ‘hands’ shown by filled circles).




Fig. 7. A typical distribution of the maximally responding units for each of the inputs of the two critical patterns when the map demonstrates co-localisation (‘face’ shown by unfilled circles, ‘hands’ shown by filled circles).

be understood in terms of the dynamics of the Kohonen algorithm which attempts to create output units trained to all inputs in the training set. As the critical patterns are presented more frequently they will compete for representational resources and this mutual competition will drive them to colonise more distal regions of the output space.

As the frequency of presentation increases further these patterns will come to dominate the representational map, distorting the underlying topography, so that some decrease in the distance measure can be seen (Fig. 3), but this is still not comparable to the decrease seen with pattern co-activation.

A moderate increase in the frequency of presentation of the critical patterns appears to increase the influence that the probability of co-activation has on co-localisation; the average minimum distance between representations falls more rapidly as the probability of co-presentation increases. This suggests that the fact that the face and hands command the largest representational areas in the somatotopic map may contribute to their co-localisation in the map, although their co-activation remains the driving force overall.

There are features of the topographic map which are not explained by the model. For example, the topographic map, although it contains significant local variation between individuals, is consistent in its medial–lateral ordering. Conversely, the model is orientation neutral and, although it will learn a topographic map of the input space, does not do this at a consistent orientation—in other words the top of the input space is equally likely to be represented at the bottom of the output space as it is at the top (compare Figs. 6 and 7, for example). Additionally, no simulation conditions explored here, including the basic condition in which no inputs have preponderance or are co-active, ensure near-perfect map formation on every simulation run.

To some extent this variability in map outcome is seen in biological maps (e.g. [14]), so its existence in the model should not overly concern us. However, given also the demonstrable optimality of some cortical maps [17,4], and the consistency of the medial–lateral ordering of the somatotopic map, there is sufficient reason to look for additional mechanisms which play a role in map formation and/or additional constraints on the orientation of the map.

It is known that some topographic organisation of somatosensory inputs begins before the primary somatosensory cortex, even as early as the spinal cord. There is strong evidence that pre-cortical organisation is not essential to the formation of an ordered somatosensory map [7], but it seems possible that pre-cortical map processes form part of the additional, and multiply redundant, mechanisms which allow optimal, consistent and developmentally robust cortical map formation. Of particular interest to the present topic, it has been shown that the pre-cortical segregation of innervation plays a role in the hand/face discontinuity of the somatosensory map [13].

Consistency of medial–lateral map ordering may be established by different areas of somatosensory cortex becoming plastic at different points in development. Such a ‘chronotopic constraint’ [5] would be a biologically plausible way of promoting consistent orientation of the

somatotopic map. The possibility of timing-dependent plasticity being sufficient to produce consistent map orientation is certainly amenable to investigation via simulation studies (see, for example, [16]).

Explaining other salient features of the human somatotopic map is not the sole motivation or benefit for extending the model. Incorporating additional biologically plausible mechanisms of self-organisation, such as these, into the model will formalise the possibilities for which biological mechanisms could be operating in map formation, as well as drawing out the strengths and limitations of the standard self-organisation algorithm as established by Kohonen [12].


Acknowledgments

Thanks are due to Kevin Gurney, Jim Stone and John Porrill for extensive advice on the nuances of self-organising map algorithms. This work was supported by EPSRC Grant EP/C516303/1 to Kevin Gurney.

References

- [1] C. Braun, U. Heinz, R. Schweizer, K. Wiech, N. Birbaumer, H. Topka, Dynamic organization of the somatosensory cortex induced by motor activity, *Brain* 124 (2001) 2259–2267.
- [2] M.A. Carreira-Perpinan, G.J. Goodhill, Are visual cortex maps optimized for coverage?, *Neural Comput.* 14 (7) (2002) 1545–1560.
- [3] D. Desieno, Adding a conscience to competitive learning, in *Proceedings of the IEEE International Conference on Neural Networks*, vol. I, 1988, pp. 117–124.
- [4] R. Durbin, G. Mitchison, A Dimension Reduction Framework For Understanding Cortical Maps, *Nature* 343 (6259) (1990) 644–647.
- [5] J. Elman, E. Bates, M. Johnson, A. Karmiloff-Smith, D. Parisi, K. Plunkett, *Rethinking Innateness: A Connectionist Perspective on Development*, MIT Press, Cambridge, MA, 1996.
- [6] M.J. Farah, Why does the somatosensory homunculus have hands next to face and feet next to genitals? A hypothesis, *Neural Comput.* 10 (8) (1998) 1983–1985.
- [7] S.L. Florence, N. Jain, M.W. Pospichal, P.D. Beck, D.L. Sly, J.H. Kaas, Central reorganization of sensory pathways following peripheral nerve regeneration in fetal monkeys, *Nature* 381 (6557) (1996) 69–71.
- [8] K. Gurney, *An Introduction to Neural Networks*, UCL Press, London, 1997.
- [9] D.O. Hebb, *The Organization of Behaviour*, Wiley, New York, 1949.
- [10] R. Hecht-Nielsen, *Neurocomputing*, Addison-Wesley, Reading, MA, 1990.
- [11] C.A. Kell, K. von Kriegstein, A. Rosler, A. Kleinschmidt, H. Laufs, The sensory cortical representation of the human penis: revisiting somatotopy in the male homunculus, *J. Neurosci.* 25 (25) (2005) 5984–5987.
- [12] T. Kohonen, *Self-organization and Associative Memory*, Springer, Berlin, 1984.
- [13] P.R. Manger, T.M. Woods, A. Munoz, E.G. Jones, Hand/face border as a limiting boundary in the body representation in monkey somatosensory cortex, *J. Neurosci.* 17 (16) (1997) 6338–6351.
- [14] M.M. Merzenich, R.J. Nelson, J.H. Kaas, M.P. Stryker, W.M. Jenkins, J.M. Zook, M.S. Cynader, A. Schoppmann, Variability in hand surface representations in areas 3b and 1 in adult owl and squirrel monkeys, *J. Comp. Neurol.* 258 (2) (1987) 281–296.
- [15] W. Penfield, T. Rasmussen, *The Cerebral Cortex of Man: A Clinical Study of Localization of Function*, Macmillan, New York, 1950.

- [16] J. Shrager, M.H. Johnson, Dynamic plasticity influences the emergence of function in a simple cortical array, *Neural Networks* 9 (7) (1996) 1119–1129.
- [17] N.V. Swindale, D. Shoham, A. Grinvald, T. Bonhoeffer, M. Hubener, Visual cortex maps are optimized for uniform coverage, *Nat. Neurosci.* 3 (8) (2000) 822–826.



Tom Stafford (<http://tomstafford.staff.shef.ac.uk/>) is a lecturer in the Psychology Department at the University of Sheffield. His main interest is the cognitive neuroscience of action selection. In particular he is interested in computational and experimental psychological approaches to the study of simple decisions. At the time of writing he is trying to persuade Reed Elsevier to cease their involvement with the international trade in arms and torture equipment. Reed Elsevier is the parent company of Elsevier, the publisher of *Neurocomputing*.



Stuart P. Wilson has a B.Sc. in Psychology and Cognitive Neuroscience from the University of Sheffield (2006), and is currently undertaking an M.Sc. in Informatics at the University of Edinburgh (specialising in Bioinformatics and Neural Computation). His particular interest is self-organisation in cortical map development.