
Perturbing Hebbian Rules

Peter Dayan
CNL, The Salk Institute
PO Box 85800
San Diego CA 92186-5800, USA
dayan@helmholtz.sdsc.edu

Geoffrey Goodhill
COGS
University of Sussex, Falmer
Brighton BN1 9QN, UK
geoffg@cogs.susx.ac.uk

Abstract

Recently Linsker [2] and MacKay and Miller [3, 4] have analysed Hebbian correlational rules for synaptic development in the visual system, and Miller [5, 8] has studied such rules in the case of two populations of fibres (particularly two eyes). Miller's analysis has so far assumed that each of the two populations has exactly the same correlational structure. Relaxing this constraint by considering the effects of small perturbative correlations within and between eyes permits study of the stability of the solutions. We predict circumstances in which qualitative changes are seen, including the production of binocularly rather than monocularly driven units.

1 INTRODUCTION

Linsker [2] studied how a Hebbian correlational rule could predict the development of certain receptive field structures seen in the visual system. MacKay and Miller [3, 4] pointed out that the form of this learning rule meant that it could be analysed in terms of the eigenvectors of the matrix of time-averaged presynaptic correlations. Miller [5, 8, 7] independently studied a similar correlational rule for the case of two eyes (or more generally two populations), explaining how cells develop in V1 that are ultimately responsive to only one eye, despite starting off as responsive to both. This process is again driven by the eigenvectors and eigenvalues of the developmental equation, and Miller [7] relates Linsker's model to the two population case.

Miller's analysis so far assumes that the correlations of activity within each population are identical. This special case simplifies the analysis enabling the projections from the two eyes to be separated out into sum and difference variables. In general,

one would expect the correlations to differ slightly, and for correlations between the eyes to be not exactly zero. We analyse how such perturbations affect the eigenvectors and eigenvalues of the developmental equation, and are able to explain some of the results found empirically by Miller [6].

Further details on this analysis and on the relationship between Hebbian and non-Hebbian models of the development of ocular dominance and orientation selectivity can be found in Goodhill (1991).

2 THE EQUATION

MacKay and Miller [3, 4] study Linsker's [2] developmental equation in the form:

$$\dot{\mathbf{w}} = (\mathbf{Q} + k_2\mathbf{J})\mathbf{w} + k_1\mathbf{n}$$

where $\mathbf{w} = [w_i], i \in [1, n]$ are the weights from the units in one layer \mathcal{R} to a particular unit in the next layer \mathcal{S} , \mathbf{Q} is the covariance matrix of the activities of the units in layer \mathcal{R} , \mathbf{J} is the matrix $J_{ij} = 1, \forall i, j$, and \mathbf{n} is the 'DC' vector $n_i = 1, \forall i$.

The equivalent for two populations of cells is:

$$\begin{pmatrix} \dot{\mathbf{w}}_1 \\ \dot{\mathbf{w}}_2 \end{pmatrix} = \begin{pmatrix} \mathbf{Q}_1 + k_2\mathbf{J} & \mathbf{Q}_c + k_2\mathbf{J} \\ \mathbf{Q}_c + k_2\mathbf{J} & \mathbf{Q}_2 + k_2\mathbf{J} \end{pmatrix} \begin{pmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \end{pmatrix} + k_1 \begin{pmatrix} \mathbf{n} \\ \mathbf{n} \end{pmatrix}$$

where \mathbf{Q}_1 gives the covariance between cells within the first population, \mathbf{Q}_2 gives that between cells within the second, and \mathbf{Q}_c (assumed symmetric) gives the covariance between cells in the two populations. Define \mathbf{Q}_* as this full, two population, development matrix.

Miller studies the case in which $\mathbf{Q}_1 = \mathbf{Q}_2 = \mathbf{Q}$ and \mathbf{Q}_c is generally zero or slightly negative. Then the development of $\mathbf{w}_1 - \mathbf{w}_2$ (which Miller calls \mathbf{S}^D) and $\mathbf{w}_1 + \mathbf{w}_2$ (\mathbf{S}^S) separate; for $\mathbf{Q}_c = 0$, these go like:

$$\frac{\delta \mathbf{S}^D}{\delta t} = \mathbf{Q}\mathbf{S}^D \text{ and } \frac{\delta \mathbf{S}^S}{\delta t} = (\mathbf{Q} + 2k_2\mathbf{J})\mathbf{S}^S + 2k_1\mathbf{n}.$$

and, up to various forms of normalisation and/or weight saturation, the patterns of dominance between the two populations are determined by the initial value and the fastest growing components of \mathbf{S}^D . If upper and lower weight saturation limits are reached at roughly the same time (Berns, personal communication), the conventional assumption that the fastest growing eigenvectors of \mathbf{S}^D dominate the terminal state is borne out.

The starting condition Miller adopts has $\mathbf{w}_1 - \mathbf{w}_2 = \epsilon' \mathbf{a}$ and $\mathbf{w}_1 + \mathbf{w}_2 = \mathbf{b}$, where ϵ' is small, and \mathbf{a} and \mathbf{b} are $\mathcal{O}(1)$. Weights are constrained to be positive, and saturate at some upper limit. Also, additive normalisation is applied throughout development, which affects the growth of the \mathbf{S}^S (but not the \mathbf{S}^D) modes. As discussed by MacKay and Miller [3, 4], this is approximately accommodated in the $k_2\mathbf{J}$ component.

Mackay and Miller analyse the eigendecomposition of $\mathbf{Q} + k_2\mathbf{J}$ for general and radially symmetric covariance matrices \mathbf{Q} and all values of k_2 . It turns out that the eigendecomposition of \mathbf{Q}_* for the case $\mathbf{Q}_1 = \mathbf{Q}_2 = \mathbf{Q}$ and $\mathbf{Q}_c = 0$ (that studied by Miller) is given in table form by:

E-vector	E-value	Conditions	
$(\mathbf{x}_i, \mathbf{x}_i)$	λ_i	$Q\mathbf{x}_i = \lambda_i\mathbf{x}_i$	$\mathbf{n}\cdot\mathbf{x}_i = 0$
$(\mathbf{x}_i, -\mathbf{x}_i)$	λ_i	$Q\mathbf{x}_i = \lambda_i\mathbf{x}_i$	$\mathbf{n}\cdot\mathbf{x}_i = 0$
$(\mathbf{y}_i, -\mathbf{y}_i)$	μ_i	$Q\mathbf{y}_i = \mu_i\mathbf{y}_i$	$\mathbf{n}\cdot\mathbf{y}_i \neq 0$
$(\mathbf{z}_i, \mathbf{z}_i)$	ν_i	$(Q + 2k_2)\mathbf{z}_i = \nu_i\mathbf{z}_i$	$\mathbf{n}\cdot\mathbf{z}_i \neq 0$

Figure 1 shows the matrix and the two key $(\mathbf{y}, -\mathbf{y})$ and $(\mathbf{x}, -\mathbf{x})$ eigenvectors.

The details of the decomposition of Q_* in this table are slightly obscured by degeneracy in the eigendecomposition of $Q + k_2J$. Also, for clarity, we write $(\mathbf{x}_i, \mathbf{x}_i)$ for $(\mathbf{x}_i, \mathbf{x}_i)^\top$. A consequence of the first two rows in the table is that $(\eta\mathbf{x}_i, \theta\mathbf{x}_i)$ is an eigenvector for any η and θ ; this becomes important later.

That the development of S^D and S^S separates can be seen in the (\mathbf{u}, \mathbf{u}) and $(\mathbf{u}, -\mathbf{u})$ forms of the eigenvectors. In Miller's terms the onset of dominance of one of the two populations is seen in the $(\mathbf{u}, -\mathbf{u})$ eigenvectors – dominance requires that μ_j for the eigenvector whose elements are all of the same sign (one such exists for Miller's Q) is larger than the μ_i and the λ_i for all the other such eigenvectors. In particular, on pages 296-300 of [6], he shows various cases for which this does and one in which it does not happen. To understand how this comes about, we can treat the latter as a perturbed version of the former.

3 PERTURBATIONS

Consider the case in which there are small correlations between the projections and/or small differences between the correlations within each projection. For instance, one of Miller's examples indicates that small within-eye anti-correlations can prevent the onset of dominance. This can be perturbatively analysed by setting $Q_1 = Q + \epsilon E_1$, $Q_2 = Q + \epsilon E_2$ and $Q_c = \epsilon E_c$. Call the resulting matrix Q_*^ϵ .

Two questions are relevant. Firstly, are the eigenvectors stable to this perturbation, *ie* are there vectors \mathbf{a}_1 and \mathbf{a}_2 such that $(\mathbf{u}_1 + \epsilon\mathbf{a}_1, \mathbf{u}_2 + \epsilon\mathbf{a}_2)$ is an eigenvector of Q_*^ϵ if $(\mathbf{u}_1, \mathbf{u}_2)$ is an eigenvector of Q_* with eigenvalue ϕ ? Secondly, how do the eigenvalues change?

One way to calculate this is to consider the equation the perturbed eigenvector must satisfy:¹

$$Q_*^\epsilon \begin{pmatrix} \mathbf{u}_1 + \epsilon\mathbf{a}_1 \\ \mathbf{u}_2 + \epsilon\mathbf{a}_2 \end{pmatrix} = (\phi + \epsilon\psi) \begin{pmatrix} \mathbf{u}_1 + \epsilon\mathbf{a}_1 \\ \mathbf{u}_2 + \epsilon\mathbf{a}_2 \end{pmatrix}$$

and look for conditions on \mathbf{u}_1 and \mathbf{u}_2 and the values of \mathbf{a}_1 , \mathbf{a}_2 and ψ by equating the $\mathcal{O}(\epsilon)$ terms. We now consider a specific example. Using the notation of the table above, $(\mathbf{y}_i + \epsilon\mathbf{a}_1, -\mathbf{y}_i + \epsilon\mathbf{a}_2)$ is an eigenvector with eigenvalue $\mu_i + \epsilon\psi_i$ if

$$\begin{aligned} (Q - \mu_i I) \mathbf{a}_1 + k_2 J (\mathbf{a}_1 + \mathbf{a}_2) &= -(\epsilon_1 - \epsilon_c - \psi_i I) \mathbf{y}_i, \text{ and} \\ (Q - \mu_i I) \mathbf{a}_2 + k_2 J (\mathbf{a}_1 + \mathbf{a}_2) &= -(\epsilon_c - \epsilon_2 + \psi_i I) \mathbf{y}_i. \end{aligned}$$

Subtracting these two implies that

$$(Q - \mu_i I) (\mathbf{a}_1 - \mathbf{a}_2) = -(\epsilon_1 - 2\epsilon_c + \epsilon_2 - 2\psi_i I) \mathbf{y}_i.$$

¹This is a standard method for such linear systems, *eg* in quantum mechanics.

However, $\mathbf{y}_i^\top (\mathbf{Q} - \mu_i \mathbf{I}) = 0$, since \mathbf{Q} is symmetric and \mathbf{y}_i is an eigenvector with eigenvalue μ_i , so multiplying on the left by \mathbf{y}_i^\top , we require that

$$2\psi_i \mathbf{y}_i^\top \mathbf{y}_i = \mathbf{y}_i^\top (\mathbf{E}_1 - 2\mathbf{E}_c + \mathbf{E}_2) \mathbf{y}_i$$

which sets the value of ψ_i . Therefore $(\mathbf{y}_i, -\mathbf{y}_i)$ is stable in the required manner.

Similarly $(\mathbf{z}_i, \mathbf{z}_i)$ is stable too, with an equivalent perturbation to its eigenvalue. However the pair $(\mathbf{x}_i, \mathbf{x}_i)$ and $(\mathbf{x}_i, -\mathbf{x}_i)$ are not stable – the degeneracy from their having the same eigenvalue is broken, and two specific eigenvectors, $(\alpha_i \mathbf{x}_i, \beta_i \mathbf{x}_i)$ and $(-\beta_i \mathbf{x}_i, \alpha_i \mathbf{x}_i)$ are stable, for particular values α_i and β_i . This means that to first order, \mathbf{S}^D and \mathbf{S}^S no longer separate, and the full, two-population, matrix must be solved.

To model Miller's results, call $\mathbf{Q}_*^{\epsilon, m}$ the special case of \mathbf{Q}_*^ϵ for which $\mathbf{E}_1 = \mathbf{E}_2 = \mathbf{E}$ and $\mathbf{E}_c = 0$. Also, assume that the $\mathbf{x}_i, \mathbf{y}_i$ and \mathbf{z}_i are normalised, let $\mathbf{e}_1(\mathbf{u}) = \mathbf{u}^\top \mathbf{E}_1 \mathbf{u}$, etc, and define $\gamma(\mathbf{u}) = (\mathbf{e}_1(\mathbf{u}) - \mathbf{e}_2(\mathbf{u}))/2\mathbf{e}_c(\mathbf{u})$, for $\mathbf{e}_c(\mathbf{u}) \neq 0$, and $\gamma_i = \gamma(\mathbf{x}_i)$. Then we have

$$\beta_i/\alpha_i = -\gamma_i \pm \sqrt{1 + \gamma_i^2} \quad (1)$$

and the eigenvalues are:

E-vector	Eigenvalue for case:		
	\mathbf{Q}_*	$\mathbf{Q}_*^{\epsilon, m}$	\mathbf{Q}_*^ϵ
$(\alpha_i \mathbf{x}_i, \beta_i \mathbf{x}_i)$	λ_i	$\lambda_i + \epsilon \mathbf{e}_1(\mathbf{x}_i)$	$\lambda_i + \epsilon [\mathbf{e}_1(\mathbf{x}_i) + \mathbf{e}_2(\mathbf{x}_i) + \Xi_i /2]$
$(-\beta_i \mathbf{x}_i, \alpha_i \mathbf{x}_i)$	λ_i	$\lambda_i + \epsilon \mathbf{e}_1(\mathbf{x}_i)$	$\lambda_i - \epsilon [\mathbf{e}_1(\mathbf{x}_i) + \mathbf{e}_2(\mathbf{x}_i) + \Xi_i /2]$
$(\mathbf{y}_i, -\mathbf{y}_i)$	μ_i	$\mu_i + \epsilon \mathbf{e}_1(\mathbf{y}_i)$	$\mu_i + \epsilon [\mathbf{e}_1(\mathbf{y}_i) + \mathbf{e}_2(\mathbf{y}_i) - 2\mathbf{e}_c(\mathbf{y}_i) /2]$
$(\mathbf{z}_i, \mathbf{z}_i)$	γ_i	$\gamma_i + \epsilon \mathbf{e}_1(\mathbf{z}_i)$	$\gamma_i + \epsilon [\mathbf{e}_1(\mathbf{z}_i) + \mathbf{e}_2(\mathbf{z}_i) + 2\mathbf{e}_c(\mathbf{z}_i) /2]$

where $\Xi_i = \sqrt{[\mathbf{e}_1(\mathbf{x}_i) - \mathbf{e}_2(\mathbf{x}_i)]^2 + 4\mathbf{e}_c(\mathbf{x}_i)^2}$. For the case Miller treats, since $\mathbf{E}_1 = \mathbf{E}_2$, the degeneracy in the original solution is preserved, ie the perturbed versions of $(\mathbf{x}_i, \mathbf{x}_i)$ and $(\mathbf{x}_i, -\mathbf{x}_i)$ have the same eigenvalues. Therefore the \mathbf{S}^D and \mathbf{S}^S modes still separate.

This perturbed eigendecomposition suffices to show how small additional correlations affect the solutions. We will give three examples. The case mentioned above on page 299 of [6], shows how small same-eye anti-correlations within the radius of the arbor function cause a particular $(\mathbf{y}_i, -\mathbf{y}_i)$ eigenvector (i.e. one for which all the components of \mathbf{y}_i have the same sign) to change from growing faster than a $(\mathbf{x}_i, -\mathbf{x}_i)$ (for which some components of \mathbf{x}_i are positive and some negative to ensure that $\mathbf{n} \cdot \mathbf{x}_i = 0$) to growing slower than it, converting a monocular solution to a binocular one.

In our terms, this is the $\mathbf{Q}_*^{\epsilon, m}$ case, with \mathbf{E}_1 a negative matrix. Given the conditions on signs of their components, $\mathbf{e}_1(\mathbf{y}_i)$ is more negative than $\mathbf{e}_1(\mathbf{x}_i)$, and so the eigenvalue for the perturbed $(\mathbf{y}_i, -\mathbf{y}_i)$ would be expected to decrease more than that for the perturbed $(\mathbf{x}_i, -\mathbf{x}_i)$. This is exactly what is found. Different binocular eigensolutions are affected by different amounts, and it is typically a delicate issue as to which will ultimately prevail. Figure 2 shows a sample perturbed matrix for which dominance will not develop. If the change in the correlations is large ($\mathcal{O}(1)$), then the eigenfunctions can change shape (eg 1s becomes 2s in the notation of [4]). We do not address this here, since we are considering only changes of $\mathcal{O}(\epsilon)$.

Positive opposite-eye correlations can have exactly the same effect. This time $e_c(\mathbf{y}_i)$ is greater than $e_c(\mathbf{x}_i)$, and so, again, the eigenvalue for the perturbed $(\mathbf{y}_i, -\mathbf{y}_i)$ would be expected to decrease more than that for the perturbed $(\mathbf{x}_i, -\mathbf{x}_i)$. Figure 3 shows an example which is infelicitous for dominance.

The third case is for general perturbations in Q_*^ϵ . Now the mere signs of the components of the eigenvectors are not enough to predict which will be affected more. Figure 4 gives an example for which ocular dominance will still occur. Note that the $(\mathbf{x}_i, -\mathbf{x}_i)$ eigenvector is no longer stable, and has been replaced by one of the form $(\alpha_i \mathbf{x}_i, \beta_i \mathbf{x}_i)$.

If general perturbations of the same order of magnitude as the difference between w_1 and w_2 (ie $\epsilon' \simeq \epsilon$) are applied, the α_i and β_i terms complicate Miller's S^D analysis to first order. Let $w_1(0) - w_2(0) = \epsilon \mathbf{a}$ and apply Q_*^ϵ as an iteration matrix. $w_1(n) - w_2(n)$, the difference between the projections after n iterations has no $\mathcal{O}(1)$ component, but two sets of $\mathcal{O}(\epsilon)$ components; $\{2\mu_i^n (\mathbf{a} \cdot \mathbf{y}_i) \mathbf{y}_i\}$, and

$$\left\{ \begin{array}{l} \lambda_i^n [1 + \epsilon(\Upsilon_i + \Xi_i)/2\lambda_i]^n (\alpha_i \mathbf{x}_i \cdot \mathbf{w}_1(0) + \beta_i \mathbf{x}_i \cdot \mathbf{w}_2(0)) (\alpha_i - \beta_i) \mathbf{x}_i - \\ \lambda_i^n [1 + \epsilon(\Upsilon_i - \Xi_i)/2\lambda_i]^n (\alpha_i \mathbf{x}_i \cdot \mathbf{w}_2(0) - \beta_i \mathbf{x}_i \cdot \mathbf{w}_1(0)) (\alpha_i + \beta_i) \mathbf{x}_i \end{array} \right\}$$

where $\Upsilon_i = e_1(\mathbf{x}_i) + e_2(\mathbf{x}_i)$. Collecting the terms in this expression, and using equation 1, we derive

$$\left\{ \lambda_i^n \left[(\alpha_i^2 + \beta_i^2) \mathbf{x}_i \cdot \mathbf{a} + 2n \frac{\Xi_i}{\lambda_i} \Upsilon_i \alpha_i \beta_i \mathbf{x}_i \cdot \mathbf{b} \right] \mathbf{x}_i \right\}$$

where $\mathbf{b} = w_1(0) + w_2(0)$. The second part of this expression depends on n , and is substantial because $w_1(0) + w_2(0)$ is $\mathcal{O}(1)$. Such a term does not appear in the unperturbed system, and can bias the competition between the \mathbf{y}_i and the \mathbf{x}_i eigenvectors, in particular towards the binocular solutions. Again, its precise effects will be sensitive to the unperturbed eigenvalues.

4 CONCLUSIONS

Perturbation analysis applied to simple Hebbian correlational learning rules reveals the following:

- Introducing small anti-correlations within each eye causes a tendency toward binocularity. This agrees with the results of Miller.
- Introducing small positive correlations between the eyes (as will inevitably occur once they experience a natural environment) has the same effect.
- The overall eigensolution is not stable to small perturbations that make the correlational structure of the two eyes unequal. This also produces interesting effects on the growth rates of the eigenvectors concerned, given the initial conditions of approximately equivalent projections from both eyes.

Acknowledgements

We are very grateful to Ken Miller for helpful discussions, and to Christopher Longuet-Higgins for pointing us in the direction of perturbation analysis. Support

was from the SERC and a Nuffield Foundation Science travel grant to GG. GG is grateful to David Willshaw and the Centre for Cognitive Science for their hospitality. GG's current address is The Centre for Cognitive Science, University of Edinburgh, 2 Buccleuch Place, Edinburgh EH8 9LW, Scotland, and correspondence should be directed to him there.

References

- [1] Goodhill, GJ (1991). *Correlations, Competition and Optimality: Modelling the Development of Topography and Ocular Dominance*. PhD Thesis, Sussex University.
- [2] Linsker, R (1986). From basic network principles to neural architecture (series). *Proc. Nat. Acad. Sci., USA*, **83**, pp 7508-7512, 8390-8394, 8779-8783.
- [3] MacKay, DJC & Miller, KD (1990). Analysis of Linsker's simulations of Hebbian rules. *Neural Computation*, **2**, pp 169-182.
- [4] MacKay, DJC & Miller, KD (1990). Analysis of Linsker's application of Hebbian rules to linear networks. *Network*, **1**, pp 257-297.
- [5] Miller, KD (1989). *Correlation-based Mechanisms in Visual Cortex: Theoretical and Empirical Studies*. PhD Thesis, Stanford University Medical School.
- [6] Miller, KD (1990). Correlation-based mechanisms of neural development. In MA Gluck & DE Rumelhart, editors, *Neuroscience and Connectionist Theory*. Hillsborough, NJ: Lawrence Erlbaum.
- [7] Miller, KD (1990). Derivation of linear Hebbian equations from a nonlinear Hebbian model of synaptic plasticity. *Neural Computation*, **2**, pp 321-333.
- [8] Miller, KD, Keller, JB & Stryker, MP (1989). Ocular dominance column development: Analysis and simulation. *Science*, **245**, pp 605-615.

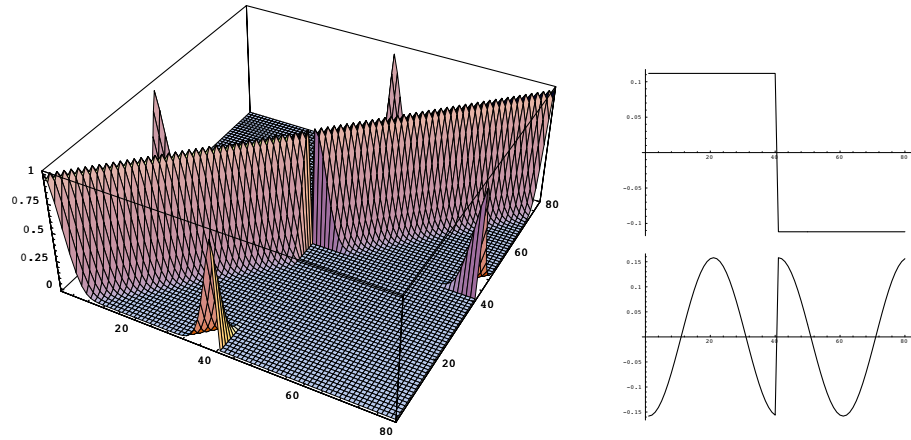


Figure 1: Unperturbed two-eye correlation matrix and $(y, -y)$, $(x, -x)$ eigenvectors. Eigenvalues are 7.1 and 6.4 respectively.

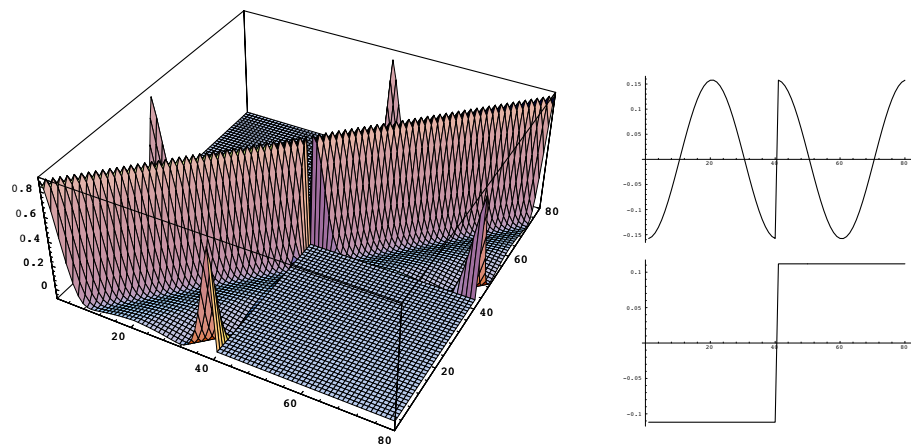


Figure 2: Same-eye anti-correlation matrix and eigenvectors. $(y, -y)$, $(x, -x)$ eigenvalues are 4.8 and 5.4 respectively, and so the order has swapped.




Figure 3: Opposite-eye positive correlation matrix and eigenvectors. Eigenvalues of $(y, -y)$, $(x, -x)$ are 4.8 and 5.4, so ocular dominance is again inhibited.




Figure 4: The effect of random perturbations to the matrix. Although the order is restored (eigenvalues are 7.1 and 6.4), note the $(\alpha x, \beta x)$ eigenvector.