# Dynamics and bifurcations in neural nets

Bard Ermentrout

March 19, 2002

## Introduction

A recurrent neural net, whether it is continuous or discrete in space and time defines a dynamical system. Thus, it is possible to apply the powerful qualitative and geometric tools of dynamical systems theory to understand the behavior of neural networks. These techniques are most useful when the behavior of interest is stationary in the sense that the inputs are at most time- or space- periodic. Thus, we can ask what kind of behavior we can expect over the long run for a given neural network. Such information is important both in artificial neural networks and biological neural nets. In the former, the final state of the neural network may represent the recognition of an input pattern, the segmentation of an image, or any number of machine computations. The stationary states of biological neural networks may correspond to cognitive decisions (e.g. binding via synchronous oscillations) or to pathological behavior such as seizures and hallucinations.

Another important issue that is addressed by dynamical systems theory is how the qualitative dynamics depends on parameters. The qualitative change of a dynamical system as a parameter is changed is the subject of bifurcation theory. The word "bifurcation" is derived from the Greek word for branching; we are concerned with the appearance and disappearance of branches of solutions to a given set of equations as some parameters vary. There are now a large number of very good general books on the mathematical theory behind dynamical systems and bifurcation. In this article we show how to use these techniques to understand the behavior of neural nets. A fundamental problem for both artificial and biologically motivated neural nets is to understand how the solutions depend on the parameters and the initial states of the network.

# Some basic definitions.

A **dynamical system** consists of a phase space, $X$, a time domain, $T$, and a function which describes the evolution of the phase space, $\phi(x,t)$. The function $\phi$ gives the value of an element in the phase space at time $t$ given that at $t = 0$ it was $x$. The two main motivating examples are differential equations and maps. In the former case, the time domain is the real line (continuous time) while in the latter it is the integers (discrete time). Consider the ordinary differential equation:

$$\frac{dx}{dt} = F(x) \quad x \in X, \quad t \in R. \tag{1}$$

Then we define $\phi(x_0, t)$ to be the solution to (1) with initial condition $x_0$. For example, if

$$\frac{dx}{dt} = x, \quad x(0) = x_0$$

then $x(t) = \phi(x_0, t) \equiv x_0 e^t$. (We have restricted our attention to *autonomous* systems in which there is no explicit time-dependence.) Consider, next, the iteration:

$$x(n+1) = F(x(n)) \quad x \in X, \quad n \in Z. \tag{2}$$

Then $\phi(x_0, n)$ is defined as the solution to (2) with initial conditions $x = x_0$. For example, if

$$x(n+1) = 2x(n), \quad x(0) = x_0$$

then $x(n) = \phi(x_0, n) = x_0 2^n$.

There is nothing that prevents us from considering infinite dimensional dynamical systems such as partial differential equations (eg see the article by Murray) or neural networks distributed in space. The set of states $\Gamma(x_0) = \{\phi(x_0, t) : t \in T, \phi(x_0, t)\text{defined}\}$ is called the **orbit** or **trajectory** through $x_0$. The orbit is a curve in state space for continuous systems and a sequence of points for discrete systems. If $\Gamma(x_0)$ consists of a single point in phase space, then we say that $x_0$ is a **fixed point** or **equilibrium** for the system. Fixed points are easily found by solving $F(x_0) = 0$ or $F(x_0) = x_0$ for continuous and discrete dynamical systems respectively. If $\phi(x_0, t + P) = \phi(x_0, t)$ for some nonzero value $P$ then the orbit is called **periodic** with period $P$. A set $S$ is **invariant** with respect to the dynamical system if $y \in S$ implies that $\phi(y, t) \in S$ for all $t \in T$. Thus, any orbit is an invariant set as is any fixed point or periodic solution. The partitioning of the state space into orbits is called the **phase portrait** of the dynamical system and is one of the goals of dynamical systems.

Another one of the key questions in dynamical systems is the issue of stability. We say that an invariant set $S_0$ is **stable** if for any $y$ close to $S_0$, $\phi(y,t)$ stays close to $S_0$ for all $t \in T$. An invariant set is **asymptotically stable** if for any $y$ near $S_0$, the distance between $\phi(y,t)$ and $S_0$ tends to zero as $t \to \infty$. The stability of a fixed point of a discrete or continuous dynamical system is easily determined by studying the eigenvalues of an associated linear operator or matrix. Suppose that $x_0$ is a fixed point. Let

$$A = DF(x_0)$$

be the matrix obtained by taking the partial derivatives of $F$ with respect to the state variables and evaluating it at the fixed point. The dynamical system obtained by replacing $F(x)$ with $Ax$ is called the **linearized system.** For (1) if all of the eigenvalues of $A$ have strictly negative real parts, then the fixed point is asymptotically stable (and all solutions to the linearized system decay to 0.) If any eigenvalue has a positive real part, then the fixed point is unstable. For (2) if all of the eigenvalues of $A$ lie inside the unit circle, then the fixed point is asymptotically stable. If any eigenvalue lies outside the unit circle, then the fixed point is unstable. As long as none of the eigenvalues have zero real part (respectively, lie on the unit circle), we say the fixed point of the differential equation (1) (respectively, map, (2)) is **hyperbolic.**

Eigenvalues that have negative real parts (lie in the unit circle) are called **stable eigenvalues**, those with positive real parts (lie outside the unit circle) are called **unstable eigenvalues** and those that have zero real parts (lie on the unit circle) are called **neutral eigenvalues** for the continuous time (discrete time) fixed point. The invariant set $W^s(x_0) = \{y \in X : \phi(y,t) \to x_0 \mathrm{as} t \to \infty\}$ (respectively $W^u(x_0)\{y \in X : \phi(y,t) \to x_0 \mathrm{as} t \to -\infty\}$) is called the **stable manifold** (respectively, **unstable manifold**) of the fixed point $x_0$. The stable (unstable) manifold is just the set of all points that tend to the fixed point as time increases (decreases) to infinity (negative infinity). The dimension of the stable (respectively, unstable) manifold is the number of stable (respectively unstable) eigenvalues of the linearized system. A fixed point that has both unstable and stable eigenvalues is called a **saddle point.** If the fixed point is asymptotically stable, then the stable manifold has the same dimension as the phase space and is then often called the **basin of attraction** for the fixed point. For neutral eigenvalues there is a **center manifold** which is invariant and has the same dimension as the number of neutral eigenvalues. The center manifold is extremely useful and important as we shall see since it allows one to study the behavior of high dimensional systems in a lower dimensional setting.

So far, the discussion of asymptotic behavior has been restricted to fixed points. A continuous dynamical system can often be reduced to a discrete one by introducing a Poincare map. Suppose the system has a periodic orbit. A **cross-section** for an $n-$dimensional continuous dynamical system is an $n-1-$dimensional hypersurface that is orthogonal to the tangent of the periodic orbit. (See figure 1.) A point on the surface that starts near the periodic orbit will be brought back to the surface at a later time. This produces a locally defined map from the surface back to itself which is then a discrete dynamical system. From Figure 1, it is clear that a fixed point of the Poincare map corresponds to a periodic solution of the original system. Thus, we can determine stability, stable, unstable, and center manifolds for periodic solutions just by studying the behavior of a Poincare map defined in some local neighborhood of it. An isolated periodic solution is called a **limit cycle** and its stability is determined by studying the stability of the fixed point for the associated Poincare map. A limit cycle solution is hyperbolic if the fixed point of the Poincare map is hyperbolic.

Invariant sets are not all as simple as periodic solutions and fixed points. In fact, they can be quite complicated. A stable invariant set which has irregular behavior (*e.g.,* it is not a simple curve or point) is called a **strange**

**attractor**. Similarly, the behavior on an invariant set can be quite complex as well. We say that an invariant set is **chaotic** if it displays sensitive dependence on initial conditions; that is, the orbits through two arbitrarily close points on the set diverge from each other exponentially. (For examples of chaotic behavior in neurons, see the articles by Aihara, Faure & Korn, and Glass.)

A dynamical system which has only hyperbolic fixed points and periodic orbits will maintain the same qualitative behavior if the parameters are varied slightly. Thus, qualitative changes are seen when the fixed points and periodic orbits become non-hyperbolic. This happens when eigenvalues cross the critical axis and thus a fixed point loses stability. This is one of the key ideas behind **bifurcation theory.** Roughly, we expect to see qualitative changes as a parameter varies when some of the fixed points or periodic orbits become non-hyperbolic. Bifurcation theory gives us a method of studying arbitrary dynamical systems near these critical values of parameters. The idea is that near a critical parameter value, there will be some neutral eigenvalues. This will imply that there is a nontrivial center manifold and in fact the local dynamics of the full system can be completely understood by studying the dynamics restricted to this center manifold. Bifurcation methods give a recipe for computing the form of the equations on this low-dimensional system. This means that one can study a possibly infinite-dimensional system by looking at the dynamics of possibly one-dimensional system!

The simplified dynamical systems that one obtains near critical values of the parameters are called **normal forms**. Thus, the understanding of the local behavior of the full system comes from studying the behavior of the relevant normal form.

## Local Bifurcations.

Local bifurcation theory allows one to study the behavior of discrete and continuous dynamical systems near fixed points. Since the behavior of periodic solutions to continuous systems reduces to the analysis of fixed points of the Poincare map, local bifurcation of maps enables us to analyze bifurcations of limit cycles in continuous systems.

### Continuous time systems.

There are two ways in which a fixed point of a continuous time dynamical system can become nonhyperbolic as a parameter varies: (i) an eigenvalue crosses zero or (ii) a pair of complex eigenvalues cross the imaginary axis. In the case of a zero eigenvalue, this signifies the appearance of new fixed points near the original one. In the most general setting, with no symmetries, a zero eigenvalue implies a **fold** or **turning point** bifurcation. In this bifurcation, the deviation from the fixed point obeys one-dimensional dynamics:

$$r' = ar^2 + c(\mu - \mu^*) \tag{3}$$

where $a, c$ are problem dependent parameters, $r \in R$ and $\mu$ is the parameter that is varied. The fixed points of this simple system (called the normal form) correspond to fixed points of the full system even if it is infinite dimensional. In Figure 2A, steady state solutions are shown for $a, c > 0$. As $\mu$ increases past $\mu^*$ two fixed points (a stable and unstable one) coalesce and disappear leaving no nearby fixed point. (The picture is essentially the same for other choices of $c, a$.) In cases with additional symmetry (such as the requirement that there is always at least one fixed point or the system has some symmetry) there are two additional common normal forms:

$$r' = ar^2 + cr(\mu - \mu^*) \tag{4}$$

$$r' = ar^3 + cr(\mu - \mu^*). \tag{5}$$

(4) leads to the **transcritical** bifurcation shown in Figure 2B. This is also called an **exchange of stability** bifurcation since the stability of the two fixed points switches at the point of bifurcation. As in the fold, the signs of $a, c$ are irrelevant to the picture. With additional symmetry, (5) occurs and this is called the **pitchfork** bifurcation. In this case, the sign of $a$ is important. If $a < 0$ then the bifurcation is **supercritical** and two new **stable** fixed points arise (see Figure 2C). If $a > 0$ then two new **unstable** fixed points occur and the bifurcation is called **subcritical.** (See Figure 2D.) In many biological and physical systems, subcritical pitchfork bifurcations "turn around" at a pair of fold points and restabilize (as in Figure 2E). This leads to **bistability** between several fixed points and to hysteresis.

Suppose that stability is lost when a pair of complex eigenvalues crosses the imaginary axis. Then the system can undergo what is called a **Hopf** or **Andronov** bifurcation. The dynamics are locally determined by the simple complex differential equation:

$$z' = az^2\bar{z} + cz(\mu - \mu^*) + i\omega z \qquad a, c, z \in C. \tag{6}$$

If the real part of $a$ is negative, then a stable periodic orbit will bifurcate from the fixed point with amplitude that is proportional to $\sqrt{\mu - \mu^*}$. The bifurcation is called **supercritical**. If the real part of $a$ is positive, then an

**unstable** periodic orbit bifurcates from the fixed point and the bifurcation is called **subcritical.** The bifurcation diagrams look like those of the pitchfork bifurcation (Fig 2C,D) where $r = |z|$. As with the pitchfork bifurcation, physical systems that have subcritical Hopf bifurcations often restabilize as in Figure 2E. Thus, there will be a range of the parameter for which there is a stable limit cycle and a stable fixed point (analogous to Fig 2E.)

## Discrete dynamics.

For a discrete dynamical system to undergo a local bifurcation, eigenvalues must cross the unit circle. There are three ways in which this can happen: (i) an eigenvalue is $+1$, (ii) an eigenvalue is -1, or (iii) a complex pair of eigenvalues lie on the unit circle. The first case of an eigenvalue of 1 is completely analogous to the case of a zero eigenvalue for continuous systems. The third case is similar to the Hopf bifurcation for continuous systems and is called a **Neimark-Sacker** bifurcation. However, the dynamics of the bifurcating solution can be complicated and all one can conclude is that there is a small invariant circle. The second case of a -1 eigenvalue is called a **flip** or **period-doubling** bifurcation. The dynamics is determined by the behavior of the simple one-dimensional map:

$$r_{n+1} = -r_n + c(\mu - \mu^*)r_n + ar_n^3. \quad (7)$$

If $a > 0$ then a stable period 2 fixed point appears; that is every other iterate of the map is the same. If $a < 0$ then an unstable period 2 fixed point occurs. Period-doubling bifurcations are very important as they often signal the onset of chaotic behavior. Indeed, often a period 2 point itself will become unstable through another period doubling bifurcation to a period 4 point. This continues as the parameter is changed and a whole **cascade** of period doublings occurs terminating in chaotic behavior.

## Periodic orbits in continuous systems.

Periodic orbits undergo bifurcations similar to those of discrete dynamical systems since their local behavior is reducible to a discrete system. Folds of the Poincare map correspond to the annihilation of a stable and unstable limit cycle; flips correspond to period-doubling of the limit cycle; Neimark-Sacker bifurcations correspond to the appearance of an invariant 2-torus.

In addition to these local bifurcations, there are two common **global** bifurcations in which the period of the orbit tends to infinity. A **heteroclinic** orbit, $\gamma(t)$ is a nonconstant orbit satisfying

$$\lim_{t \to \pm\infty} \gamma(t) = x^{\pm}$$

where $x^{\pm}$ are fixed points. If $x^+ = x^-$ we call $\gamma(t)$ a **homoclinic** orbit. In general, one cannot expect to find a homoclinic orbit; rather, as some parameter, $\mu$ changes, the homoclinic orbit appears at one value of that parameter, $\mu^*$. (See Figure 3A.) Thus, the appearance of a homoclinic is a bifurcation. If the fixed point $x^+ = x^-$ is a hyperbolic saddle then as the parameter moves away from criticality, a periodic orbit arises and the period of this orbit goes to infinity as the homoclinic orbit is approached. The period is given by

$$T_{hom} \sim K \log \frac{1}{|\mu - \mu^*|}.$$

If instead of a hyperbolic saddle point, the fixed point is a fold point, then a **saddle-node on a limit cycle** occurs (Figure 3B). Limit cycles occur as the parameter is moved from criticality with a period

$$T_{sn} \sim \frac{K}{\sqrt{|\mu - \mu^*|}}.$$

## Applications to continuous time neural nets.

So far, we have introduced definitions but not used them in the context of neural networks. There have been numerous papers that use bifurcation methods to analyze neural networks including the chapters on phase-plane methods, cooperative phenomena, bursting, and pattern formation. (See Izhikevich, 2000, for a long well-illustrated review.) The types of neural net models to which the theory has been applied generally have one of the following two forms:

$$\tau_j \frac{du_j}{dt} = -u_j + F_j(\sum_k W_{jk} u_k + I_j) \quad (8)$$

$$u_j(n+1) = \mu_j u_j + F_j(\sum_k W_{jk} u_k(n) + I_j). \quad (9)$$

The functions $F_j$ are generally monotone and increasing, $W_{jk}$ are the weights, $I_j$ are the inputs and $1/\tau_j, 1-\mu_j$ are decay rates. With proper limits, the sums in the above can go to integrals over space in order to represent a continuum model for neural nets (see the article by Cowan et al). In this case the interaction takes the form

$$\int_\Omega W(x,y)u(y)\ dy$$

where $\Omega$ is the spatial domain, a one- or two- dimensional region in space. The behavior of these models is generally impossible to analyze completely except in certain simple cases. For example, Hopfield (1984) shows that if the weights, $W_{jk}$ are symmetric then all solutions to (8) tend to fixed points. In the case of one- or two-dimensions, (8) can be completely analyzed (see the chapter on phase-plane methods, or Hoppensteadt and Izhikevich, chapt 2). The discrete system (9) can be completely characterized only in one-dimension.

## Local bifurcations of fixed points.

To see how bifurcation methods can be used in neural nets, consider a simple continuous neural network with no inputs and satisfying $\tau_j = 1, F_j(u) = F(u)$ with $F(0) = 0$, $F'(0) = \mu$, a gain parameter. Thus the trivial state of the network, $u_j = 0$ is a fixed point. To study stability and bifurcation, we linearize about $u_j = 0$ and obtain the matrix:

$$M_{jk} = -\delta_{jk} + \mu W_{jk}$$

where $\delta_{jk}$ is the Kronecker delta function. If $\lambda$ is an eigenvalue of $W_{jk}$, then $-1+\mu\lambda$ is an eigenvalue for $M$; thus if the gain is sufficiently small, the trivial fixed point is stable. If some of the eigenvalues of the weight matrix, $W$ have positive real parts, then for sufficient gain, some of the eigenvalues of $M$ will cross the imaginary axis and the rest state will lose stability. In particular, let

$\lambda_0$ be the eigenvalue with maximal real part and let $\Phi_0$ be the corresponding eigenvector. Since this is a continuous dynamical system, there are only two cases of interest. Suppose that $\lambda_0$ is real. Then if $\mu = \mu^* = 1/\lambda_0$ there will be a zero eigenvalue and a bifurcation to new fixed points. Since 0 is always a fixed point, the bifurcation will be transcritical or a pitchfork. In any case, the new solution will have a nonzero amplitude and be proportional to $\Phi_0$. This is the essence of pattern formation. The trivial state $u_j = 0$ loses stability and a new state that is coded in the weight matrix bifurcates from rest. Analysis of pattern formation in neural networks and for that matter, any pattern formation models all comes down to this calculation (see Murray, 1989). If the weight matrix is symmetric, then the eigenvalues are real and the initial bifurcation from the rest state will always be through a zero eigenvalue. If the weight matrix is non-negative, then the eigenvalue with largest modulus will be real and positive and the first bifurcating mode is proportional to the principal component of the weight matrix.

Suppose the weight matrix is not symmetric and not all the same sign; that is, there are "inhibitory" and "excitatory" interactions. Then, the eigenvalue with maximal real part could be a complex eigenvalue. This means that the first instability is through a Hopf bifurcation and a oscillatory pattern of activities can bifurcate. In situations in which there is symmetry or the network has a particular geometry, this type of bifurcation can lead to periodic wave trains. Recently Hoppensteadt and Izhikevich have proposed a general theory of neural networks which exploits the kind of local analysis that we have only sketched here.

## Beyond local bifurcations.

To go beyond the simple bifurcation analysis described above, it is necessary to turn to numerical methods. There are several numerical packages for the analysis of bifurcations in nonlinear dynamics. AUTO (Doedel, 1997) is among the best of them and works on many operating systems. To illustrate the concepts of section 3, I present a global numerical diagram for a six neuron model whose weight matrix was chosen randomly from a uniform distribution with mean 0 and standard deviation 0.5. (Note that this particular example was picked out due to its rich behavior.) The function $F(x) = \tanh(\mu x)$ was chosen for simplicity. Once the weight matrix is chosen, there is only one parameter, the gain, $\mu$. Figure 3 shows the norm of the solutions that are computed numerically as the gain is increased. The eigenvalues for the weight matrix are approximately $-0.83, 0.25 \pm 0.2i$, $-0.16, 0.12, -0.47$. Thus, for positive gains, there will be a Hopf bifurcation at approximately $\mu = 4$ and a pitchfork bifurcation at approximately $\mu = 8.33$. Both of these branches are shown in the diagram and labeled **H1** and **P** respectively. The pitchfork is subcritical and undergoes a fold bifurcation labeled **F1** and stabilizes. A subcritical Hopf bifurcation **H3** occurs on this branch which turns around at the limit-cycle fold bifurcation, **Flc1** giving rise to a stable periodic orbit. This orbit turns around again at **Flc2** and the resulting unstable limit cycle terminates at a homoclinic bifurcation, **Ho**. The curve of fixed points turns around at **F2** leaving a stable fixed point which persists for all higher values of $\mu$.

The fate of the periodic orbit that arises at **H1** is more interesting. This orbit is supercritical and leads to small amplitude stable periodic solutions. As $\mu$ increases, the periodic orbit (an hence the fixed point to the Poincare map) loses stability as an eigenvalue crosses $+1$. This results in a pitchfork bifurcation at **Plc**. The unstable periodic orbit appears to persist for all values of $\mu$ beyond the bifurcation point but never restabilizes. The pitchfork bifurcation of the Poincare map is super critical and results in a stable periodic orbit. This orbit becomes unstable at a flip or period-doubling bifurcation **Pd**; the branch eventually terminates at a Hopf bifurcation **H2** on the branch of fixed points that bifurcated at **P.** The periodic branch also undergoes a Neimark-Sacker bifurcation **NS** resulting in an unstable torus. The flip bifurcation is supercritical and leads to a new branch of periodic solutions with twice the period. This branch in turn undergoes a flip bifurcation and so on producing a period doubling cascade. The regime between $\mu = 5.5$ and $\mu = 7.5$ is very complicated. There are many bifurcations and the appearance of many stable and unstable periodic orbits as well as chaotic behavior that are not described on this plot. In order to depict this behavior we plot the Poincare map whose cross-section is defined as the hyper-plane where $u_4' = 0$. The $u_2$ component of this map is shown in Figure 5 as a function of $\mu$. In this diagram, each dot represents an iteration of the Poincare map, so that for a fixed value of $\mu$ a single dot implies a periodic solution, a pair is a period two orbit, and so on. The period doubling cascades are clear. A period three orbit at about $\mu = 6.85$ is seen to undergo a period doubling cascade as $\mu$ is decreased. Chaotic behavior terminates near a saddle point at $\mu \approx 7.57$ and disappears "instantly." This global bifurcation is called a **crisis** by Yorke.

## Conclusion.

The general behavior of recurrent neural networks as parameters vary remains an open problem. Dynamical systems methods and bifurcation theory provide a general approach toward analyzing these interesting systems. Pattern formation, spatio-temporal behavior, and complex dynamics are all aspects of recurrent neural networks that can be understood by these useful mathematical tools. Analytical methods along with the careful use of numerical methods allow one to globally characterize complex biological and artificial neural networks.

**Related reading:** Chaos in Neural Systems; Phaseplane Analysis of Neural Activity; Oscillatory and Bursting Properties of Neurons; Pattern Formation, Biological; Traveling Activity Waves.

References

E. Doedel, A. Champneys, T. Fairgrieve, Y. Kuznetsov, B. Sandstede, and XJ Wang, (1997) AUTO97: Continuation and bifurcation software for ordinary differential equations (with HomCont), Computer science, Concordia University, Montreal, Canada. (ftp://ftp.cs.concordia.ca/pub/doedel/auto)

*B. Ermentrout (1998) Neural networks as spatio-temporal pattern-forming systems, Reports on Progress in Physics 61:353-430.

J. Guckenheimer and P.J. Holmes (1983) Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields. Springer-Verlag, Heidelberg.

J. Hopfield, (1984) Neurons with graded response have collective computational properties like those of two-state neurons, PNAS USA 81:3088-3092.

*F. Hoppensteadt and E. Izhikevich, (1997) Weakly Connected Neural Networks, Springer-Verlag, New York.

E.M. Izhikevich E.M. (2000) Neural Excitability, Spiking, and Bursting. International Journal of Bifurcation and Chaos. 10:1171-1266.

*Y. A. Kuznetsov (1998) Elements of Applied Bifurcation Theory, Springer-Verlag, New York.

J. D. Murray (1989) Mathematical Biology, Springer-Verlag, Heidelberg.

S. Wiggins (1990) Introduction to Applied Nonlinear Dynamical Systems and Chaos, Springer-Verlag, New York.

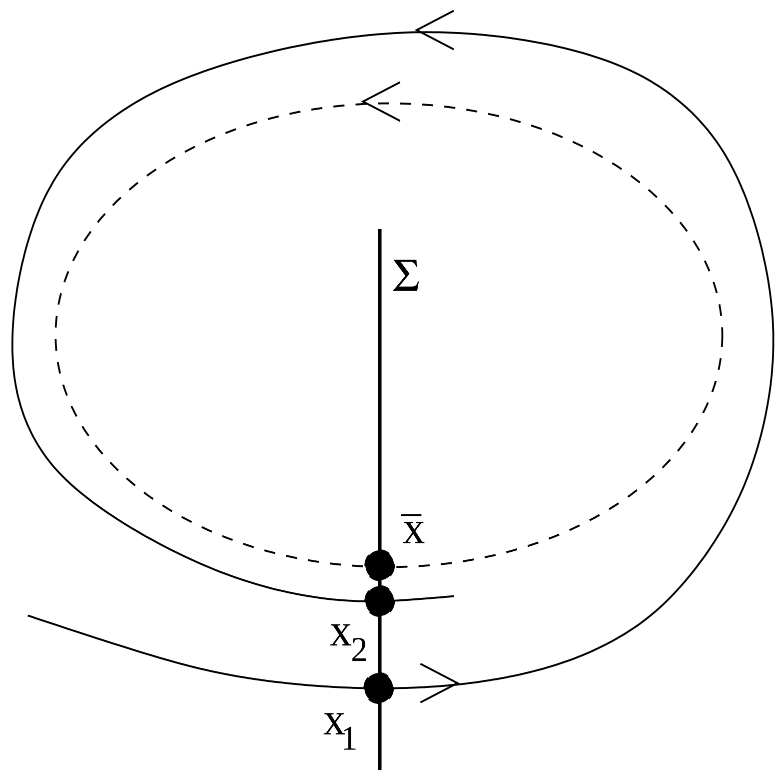


Figure 1: Construction of the Poincare map for a two-dimensional system.

# Figure legends

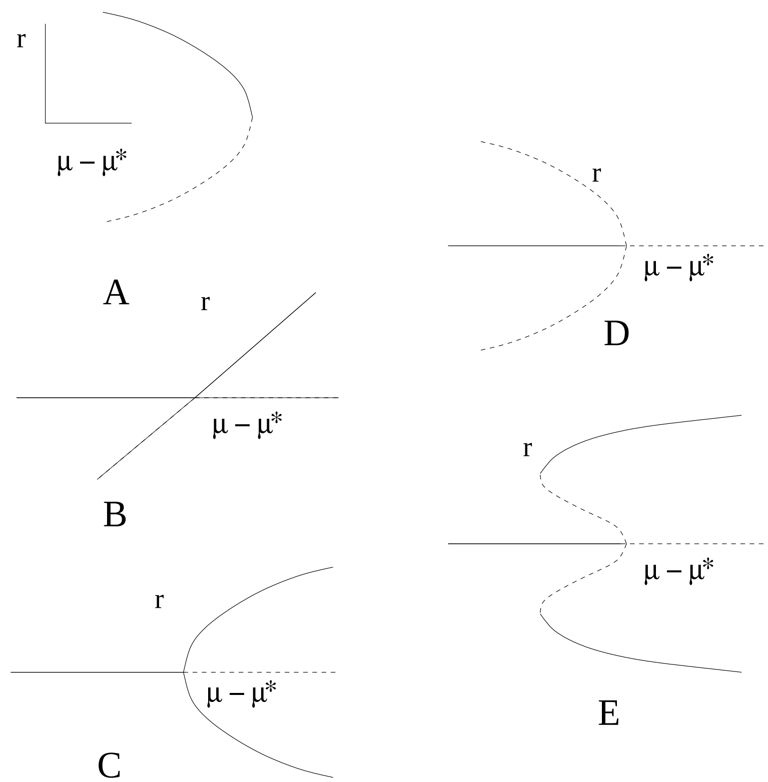


Figure 2: Bifurcation diagrams for fixed points of continuous time dynamical systems. Solid lines are stable, dashed unstable. See text for description.

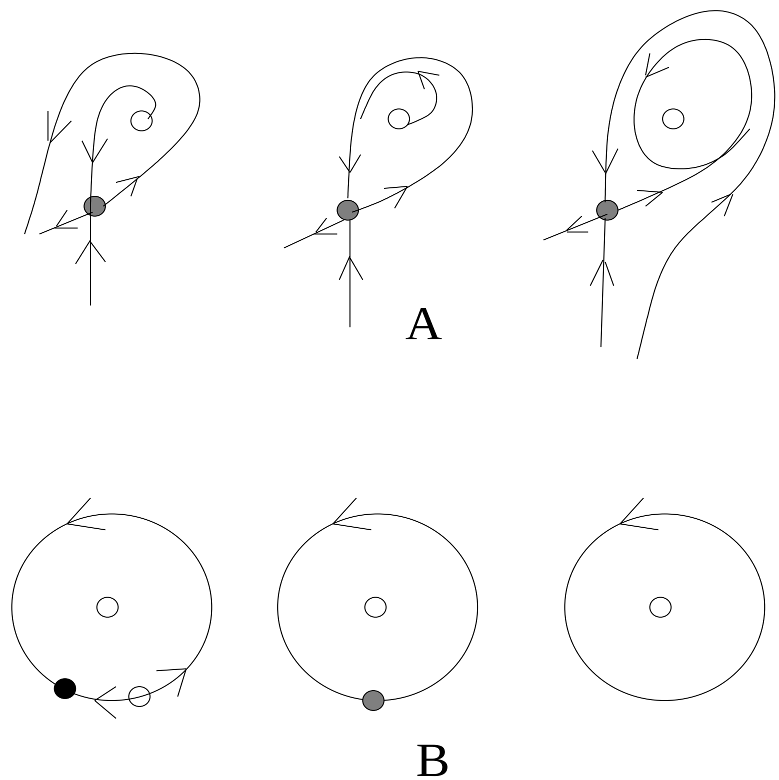


Figure 3: Two distinct types of homoclinic bifurcation. Bifurcation of a (A) homoclinic orbit from a hyperbolic saddle or a saddle-node on a limit cycle (B) leading to periodic orbits.

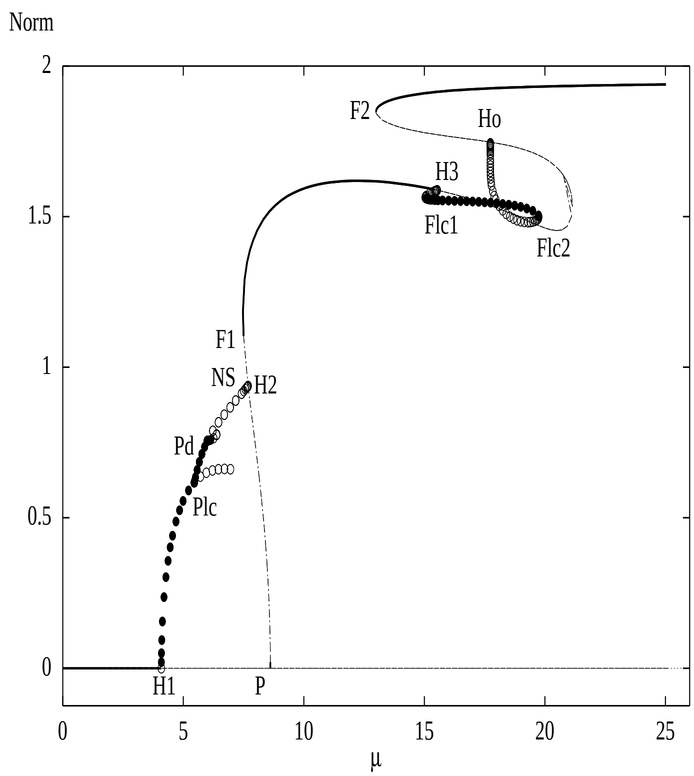

Figure 4: Numerically computed bifurcation diagram for a 6 neuron network with random weights chosen between -0.5 and 0.5.

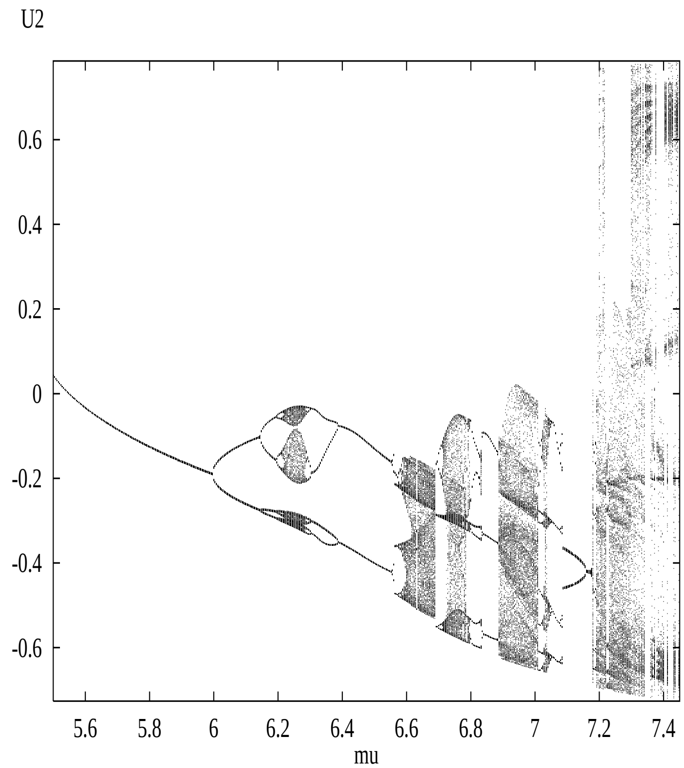


Figure 5: Poincare map of a region from Figure 3 showing complex dynamics.

# Appendix

The neural network used in Figs 3,4 has the form:

$$U_j' = -U_j + \tanh(\mu \sum_{k=1}^{6} W_{jk} U_k)$$

where

$$W = \begin{pmatrix} -0.42473 & 0.243325 & -0.267939 & 0.308063 & -0.0370201 & 0.394969 \\ -0.166832 & 0.474204 & -0.0151443 & 0.476774 & 0.211162 & 0.401305 \\ -0.427914 & 0.370044 & -0.0675567 & 0.276535 & 0.45988 & -0.457168 \\ -0.362131 & 0.033334 & -0.196538 & -0.037606 & -0.125548 & 0.143851 \\ 0.429334 & -0.306886 & 0.402954 & -0.166799 & -0.45518 & -0.0304156 \\ 0.294663 & -0.346348 & -0.138444 & 0.334973 & 0.13884 & -0.364227 \end{pmatrix}.$$