Winner-Take-All Neural Networks and Visual Search Tasks

Stephanie Lewkiewicz

Advisors: Daniel Martí and John Rinzel

1   Introduction

The research conducted and presented within this report is a mathematical exploration of a model of visual search. The fundamental question is one posed in Hugh R. Wilson's book *Spikes, Decisions, and Actions* [1]. He designs an experimental task where the subject is presented a two-dimensional image that contains a group of objects. All except one of the objects are identical. The task consists of identifying the outlier, or *target* object (Figure 1). The goal of this experiment is to test identification time, and how this time depends on the number of objects in the image. Thus, if there are N objects, then there are $N-1$ "distractors" to accompany the target. All distractors are identical to one another and slightly different from the target. When the target and distractors are almost indistinguishable, reaction times are very high on account of the increased difficulty in isolating the target [2].
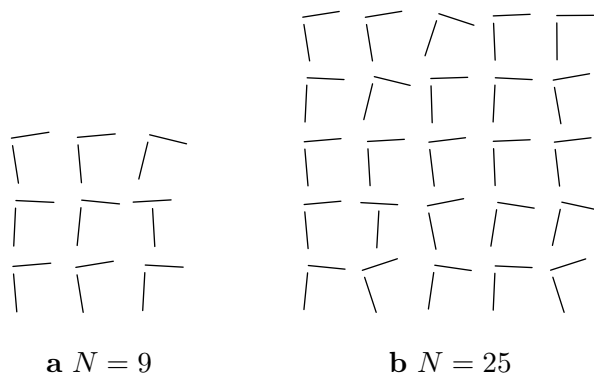

a $N = 9$ **b** $N = 25$

Figure 1: Search task. The subject has to find a gapped 'T' (the target) in a crowd of $N-1$ gapped 'L's (distractors). The time it takes to identify the target increases with the number of distractors.

In theory, a dynamical system that models this setup calls for the inclusion of N distinct equations, each of which defines the behavior of one population. However, we desire a geometric analysis of the system's behavior, and clearly visualizing more than three dimensions is, as we know, virtually impossible. Luckily, if we impose a minor restriction on the behavior of all the losing populations, we can circumvent this pictorial nightmare and perform the desired simple geometric analysis: merely assume that all the losing populations behave in exactly the same way. Doing so allows us to collapse $N-1$ equations down to one, as long as we take into account the exact number of total populations in that condensed expression. We pair this expression with that for the winning population's change in firing rate and have formed a set of two equations that together describe the activity of the full N-dimensional system:

If x denotes the firing rate of the sole winning population, and y the firing rate of any given losing population, the system evolves according to

$$\dot{x} = -x + \phi(E_x), \tag{1}$$

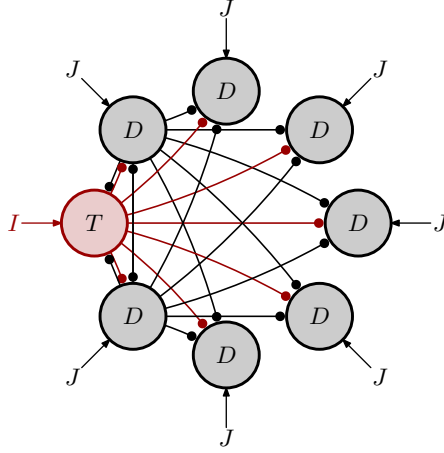$$\dot{y} = -y + \phi(E_y), \tag{2}$$

Figure 2: Architecture of the network. Black circles denote losing populations, and red the winning one. Dot-headed arrows represent inhibitory connections. For clarity we plot only the connections between the winning and two losing populations.

where the dot denotes derivative with respect to time and where ϕ is an input-output function, which transforms inputs into firing rates. We chose this function as a sigmoid curve

$$\phi(s) = \frac{1}{1 + \exp\left(-g[s - \theta]\right)}, \tag{3}$$

where $g > 0$ is a gain parameter and θ is a soft threshold.

This function is differentiable, monotonically increasing, and takes positive values between 0 and 1 for all reals, with limits $\lim_{s \to \infty} \phi(s) = 1$, $\lim_{s \to -\infty} \phi(s) = 0$. The inputs to the populations x and y, denoted by E_x and E_y respectively, are given by:

$$E_x = \alpha x - (N - 1)\beta y + I, \tag{4}$$
$$E_y = \alpha y - (N - 2)\beta y - \beta x + J, \tag{5}$$

where $\alpha > 0$ is the population's self-excitation, and $\beta > 0$ is the strength with which each population inhibits the others. The parameter I is the external input to the winning x-group, while J is the input to each losing y-group. Note that by assumption $I > J$. Throughout this paper the difference $I - J$ is kept at 0.1. Note also that while N is by definition a natural number, we will take N to be real for mathematical convenience. As is evident in the expressions, each population inhibits all others, self-excites, and receives some external stimulus. (see Fig. 2)

The two-dimensionality of the system lends itself nicely to analysis in the *xy-phase plane*. In particular, we observe the trajectory solutions of the system and the fixed points generated in the phase plane. The first are representative of the course of a decision from start to finish. As a point $(x^\star, y^\star)$ in the plane denotes a particular instantaneous firing rate for each of the populations in the network, the unique trajectory solution to the conditions $(x, y) = (x^\star, y^\star)$, $\cdot x$, and $\cdot y$ demarcates the path of the decision in terms of how firing rates of the populations change and finally settle.

Also of critical importance are the steady states of the system. The stable steady states into which trajectories are drawn indicate the final firing rates approached as a decision is made, and thus essentially represent the possible decisions themselves. (The relative magnitudes of the x and y coordinates indicate which groups win and lose, so if a trajectory heads for a fixed point with an x coordinate of norm much larger than its y component, then it's clear that x's object is being chosen.) Of greatest importance for our purposes are the saddle points of the system, which are steady states into which a trajectory is drawn before being thrust back

2

out again. (Mathematically, the system's change is negative in one direction, and simultaneously positive in the other.) These points hold the key to an elegant and predictive understanding of decision-making times, as the behavior of the system near a saddle point provides data about the speed of the trajectory along its journey. In order to grasp an estimate of the saddle point-predicted time courses, we study the eigenvalues of the total derivative matrix at the saddle. The search time is associated with the characteristic time a trajectory with initial condition at $(0,0)$, associated with a neutral activity state, in comparison with the time courses predicted by the saddle point.

In varying the parameters of which the expressions are a function, a set of several "canonical" phase planes is produced, each of which corresponds to a distinct possible outcome. These images summarize all potential destinations of a given trajectory in the phase plane. The sequences of phase planes we produce for study are taken over changes to the input, inhibition, and of course the most crucial parameter, N. In general, when input is low–effectively nonexistent–the system is mostly inactive with one fixed point near the origin, indicating low firing rates for all populations. As we increase input to both populations, stable and unstable steady states appear. These pairs of stable steady states are situated in such a way that they are highly asymmetrical and indicate large firing disparities between the two groups. They mirror one another in the sense that the x firing rate at one will be high and the y firing rate low, with roles reversed at the other. It is inevitable that when the input to the groups is not equal, the group(s) with a higher input will ultimately win, as one might anticipate. This arrangement epitomizes true competitive behavior in which either the target or the distractors can win, and one of them must–what is called *winner take all behavior*.

It was hypothesized that increases in the number of distractors would yield an increase in the search time— clearly a logical assumption, as the sorting mechanism has more objects through which to search. The goal of this project was to investigate this claim further, recreating the experiment with a similar but somewhat different model, in order to determine how changes to parameters other than N affect the time courses as N changes.

2 Qualitative description of the system

The rate of change of x and y, given by the right hand sides of Eqs. (1)–(2), consists in the sum of a linear term and a nonlinear function of the inputs (4)–(5). Given that a negative firing rate is physically absurd, in all plausible situations the negative linear terms in Eqs. (1)–(2) contribute to a decrease in the firing rate over time. The input-output function, on the other hand, serves to represent response of the neuronal population to the inputs. The bounds on the sigmoid confine all fixed points to the unit square, $0 \leq x, y \leq 1$, in two-dimensional xy-space.

when α, I, and J are all positive, increasing the value of these parameters pushes the value of $\phi(s)$ closer to one, thus increasing the populations's activity. This is natural considering that they represent self-excitation and external input, both of which logically stimulate the population. Larger positive values of β imply larger inhibition from other populations, and thus smaller values of $\phi(s)$, as s is pushed towards zero, helping to stifle the populations's activity.

Changes in the parameter α affect the curvature of the nullclines; as α increases, the nullcline is more bent. Small positive values of α yield a soft curve, whereas large positive values of α lead to a twisted curve that resembles an S stretched horizontally in both directions. The inhibition β also affects the degree of the nullcline's curvature. Large positive values of β create a very sharp nullcline. When β is large enough, the nullcline is so sharp that it resembles a step function. Indeed, it approaches a step function as β becomes infinitely large.

The input parameters I and J shift the nullclines without changing their shape. As the value of either I or J increases, the corresponding nullcline moves away from the origin in the positive direction. Changes in each of these parameters ultimately give rise to saddle-node bifurcations, which will be the primary method of creation and destruction of fixed points and will lead to substantial changes in the configuration of the phase plane. A bifurcation, in the general sense, is just that—the creation and destruction of fixed points or changes

to the stability of existent fixed points as a result of changes in parameter values. A saddle-node bifurcation occurs when a pair of fixed points, one unstable and one stable, suddenly appears and the points move away from one another as the parameter continues to change. However, changes in β also lend themselves to an interesting subcritical pitchfork bifurcation that creates a saddle point, instead of the traditional unstable fixed point formed when two unstable fixed points and one stable fixed point are formed and branch out from a single unstable fixed point.

3 Closed Expression for the Nullclines

The x and y-nullclines of the system are the curves in the phase-plane where, respectively, $\dot{x}$ and $\dot{y}$ are zero. Thus, instantaneous motion at a given point on the, say, x-nullcline, is entirely in the y-direction. (The intersection of two nullclines is a fixed point of the system.) In our system, this is equivalent to:

$$x = \phi\big(\alpha x - (N-1)\beta y + I\big),$$

and hence to

$$\phi^{-1}(x) = \alpha x - (N-1)\beta y + I, \quad \Rightarrow \quad y = -\frac{\phi^{-1}(x) - \alpha x - I}{(N-1)\beta}$$

where $\phi^{-1}(s)$ is the inverse of the sigmoid in Eq. (3), given by

$$\phi^{-1}(s) = \theta - \frac{1}{g}\ln\left(\frac{1}{s} - 1\right)$$

and defined in the domain $(0,1) \subset \mathbf{R}$. Likewise, the y-nullcline is defined by the relationship

$$x = -\frac{1}{\beta}\big(\phi^{-1}(y) - \alpha y + [N-2]\beta y - J\big).$$

4 The General System

Now we begin our study of the system in all its generality, allowing N to change. Bifurcation analysis will underlie all dissection of the system, with N the focal parameter of the bifurcation study. We will, in a sense, construct a two-parameter bifurcation analysis that is in some measurements continuous and others discrete. That is to say, we will start by fixing all the parameters except for our primary one, N, and then constructing a continuous bifurcation diagram as N ranges over some specified values. (Usually, this will begin with $N = 2$, because it's the logical choice; competition between one group and itself is a bit meaningless. The ending N-values chosen usually represent a point at which all "interesting" bifurcation activity has clearly ceased.) Then we will in turn change the non-N parameter values we had fixed in the beginning, usually by way of a linear translation (perhaps adding 0.1 to the previous value). We will compute the bifurcation diagram again for these new parameter values, and then repeat the procedure after changing those parameters once more, continuing until the parameters are just absurdly large and impractical. This way, we come to obtain a finite series of phase plane images, iterated on changing values of secondary parameters, each of which depicts a continuous bifurcation diagram of x vs N.

From this analysis we aim at understanding the effect of each parameter on the behavior of the system. We are asking the question: When parameter β is very small, and parameter J is simultaneously very large, what does the bifurcation diagram of N look like? (For example.) For what values of N is the behavior competitive? And what happens to this information after a sequence of small incremental changes to the other parameters?
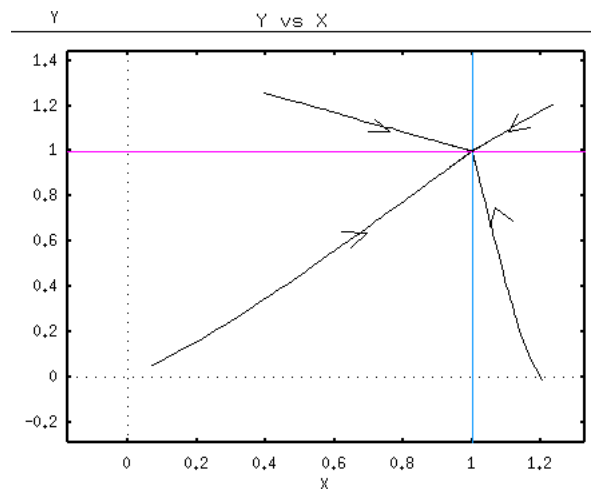
Figure 3: Phase plane in the absence of inhibition. The system has a single stable fixed point, regardless of the value of N. Cyan curve: x-nullcline. Magenta curve: y-nullcline. Parameters are $\beta = 0$, $I = 0.6$, $J = 0.5$, and $N = 3$.

4.1 No inhibition

First, we look to the simple cases, namely when one parameter is taken to be zero. The experimental analogue of this is that one form of interaction between populations is completely cut off. We begin by cutting off inhibition taking $\beta = 0$. Incidentally, this allows us to isolate the effect of the inputs. For all values of N, and nonzero values of I and J, the system has a single stable fixed point, or stable node, at $(x, y) = (1, 1)$, see Figure 3. Making β vanish removes any possible dependence of the system on the size N. As expected, the ability of one group to win is inextricably dependent on its capacity to inhibit the other(s) —no winners without losers. Strong inputs are not enough, competition requires inhibition.

4.2 No inputs

Now let us reverse the roles and consider the case in which both inputs are cut off, but the populations can inhibit each other. Although one form of interaction among the populations is removed, as with the previous case of no inhibition, this setup loses the simplicity and constancy of the last. Rather, it ends up assuming the phase plane complexity that will be seen later on when the parameter relationships themselves become more complicated. We also encounter here our first real bifurcations.

When inhibition is low ($\beta = 0.1$) and inputs are switched off ($I = J = 0$), the system begins has four stable fixed points, located at $(0, 0)$, $(0, 1)$, $(1, 0)$ and $(1, 1)$ (see Fig. 4a). This corresponds to a situation where the population x, favored by the inputs, and the populations receiving less inputs (firing each at rate y), can be simultaneously active, simultaneously inactive, x can win, or y can win. A bifurcation which might be described as a hybrid of the subcritical pitchfork and saddle node—two unstable fixed points coalesce with one stable, but annihilate instead of forming a different unstable steady state—brings the system to what will be a familiar configuration for small N and higher β at $N = 2.593$. (Fig. 4b) Thus, we are left with three stable steady states, corresponding to the possibilities of x winning, y winning, and suppression of both groups. (Fig. 4c, with $N = 3$.) Another saddle node bifurcation follows at $N = 3.5875$, removing the possibility that y might win. (Fig. 4d) As N becomes quite large, the curves are flattened but maintain the same fixed point arrangement. (Fig. 4e, with $N = 6$)

It turns out that the four-point configuration seen for small N only exists for $\beta = 0.1$. When β is increased to 0.2, we already find at $N = 2$ the configuration which was achieved after the first saddle node bifurcation with $\beta = 0.1$. (Images of these phase planes are virtually identical to the previous, so including them would be
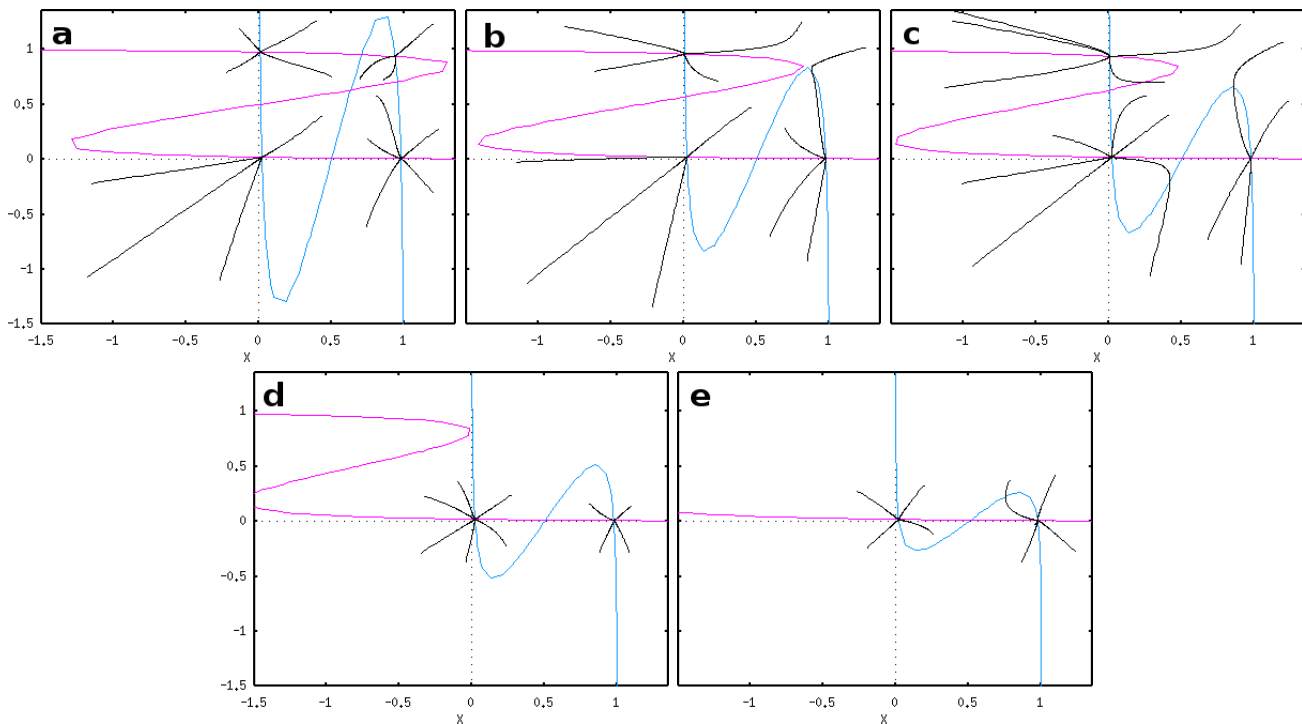
Figure 4: Change in the phase plane as N varies, in the absence of inputs.

unnecessarily excessive.) What was the second saddle node bifurcation has now become the first and only; for $\beta = 0.2$, this happens when $N = 2.8$. As β is increased incrementally, the same bifurcation pattern continues. The major change to occur alongside the increase in N is that the N-value at which the saddle-node bifurcation happens decreases. Recall that it began at 3.5875 with $\beta = 0.1$, and by the time β has reached 1, has made its way down to 2.15925, seemingly converging to 2. It's interesting to note that this change actually affects the discrete dynamics of the competition between *populations* minimally, if at all. For $0.2 \le \beta \le 1$, the phase plane is in the five-point configuration when $N = 2$ and the three-point configuration when $N = 3$; we need not concern ourselves with when, between $N = 2$ and 3, the changeover occurs. However, mathematically, the window of N values in which behavior is competitive increases, as the bifurcation to eliminate mutual inactivity takes place for continually smaller N. In conjunction with the results of the previous section with no inhibition, the criticality of inhibition's role in actually governing the competition among populations has become extremely evident. Without it, there was really no noteworthy behavior at all: No one won, no one lost, and the phase plane never changed. And now it has proven itself an even stronger force in its ability to produce phase planes and bifurcations in the complete absence of input that mimic those with input to the system. The value of the input is not diminished, as will be seen, as its presence alters the phase planes we see now in significant ways; however, this commentary on the inhibition provided by the previous two sections is truly remarkable.

This concept of input deprivation lends itself nicely to the question of whether it is possible to somehow overcome the x-victory encoded in the equations themselves by stimulating the y groups with input, but not x. As it turns out, the x group's imminent win is relatively immune to even this so y-favored skewness of the inputs. Invariably, as N is made large enough, x will always win except in instances of *extreme* stimulus to y. And even in those cases, the y victory is short-lived to a mutual fade of both groups's activity. For very small stimuli to y (take $J = 0.1$), the difference between the phase plane changes as N increases is for all intents and purposes identical to the phase plane changes in the previous case of no input whatsoever. Once the stimulus to y surpasses a certain threshold (near $J = 0.5$), there is found the presence of a somewhat lasting stable point in favor of y. (Figs. 5a, b, c) Instead of bifurcating away, it slowly creeps down towards the origin,
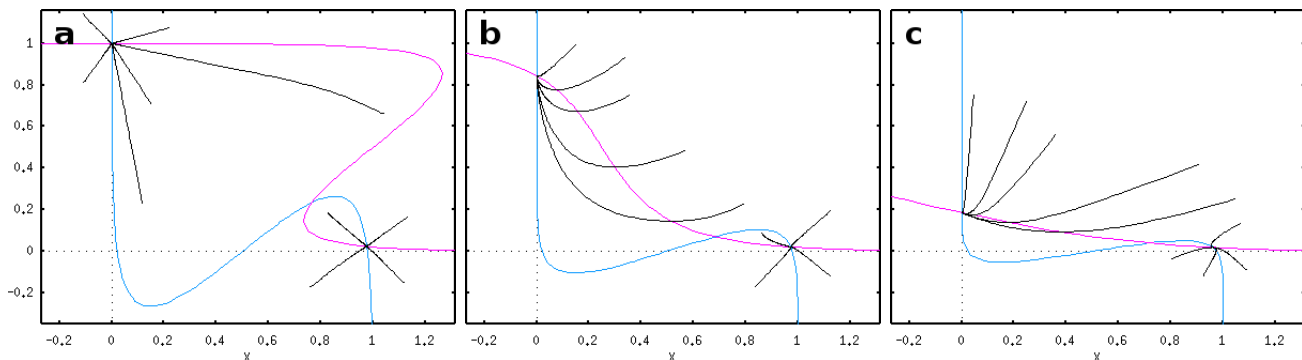
Figure 5: Change in the phase plane as N varies, in the absence of inputs. In a, b, and c $N = 2, 3.5, 6$, respectively.

preserving for longer the possibility for the y populations to not be silenced entirely, as they remain active although atrophying. For very strong stimuli to y, as one would expect, the system has only one stable steady state at $(0, 1)$. (Fig. 5d, with $N = 2$ and $J = 1$.) As N is made arbitrarily large, the fixed point travels towards the origin, signifying the dual repression of both populations rather than the common capitulation of y to x. So it seems that with absolutely no input to x and very high input to y, the best that the y-populations can do is to put up a fight until N itself becomes too large. Even as N approaches infinity, the y populations do maintain some part of that prior success in that they stop x from ever actually winning, instead settling for the compromise of both groups becoming weak and ceasing all activity. As y's diminished success comes when N is very high, we might be led to believe that this is a product of the intense inhibition imposed on any given group by the ever-increasing number of competitors.

4.3 Close Parameter Values

Now that we have thoroughly investigated the cases of absence of one channel for interaction among populations, we consider the circumstance in which all possible forms of communication are in play—the populations can inhibit each other, self-stimulate, and receive an external stimulus. Now the question remains of what the resultant behavior will look like as secondary restraints—those in addition to the new restraint that the inhibition and input parameters are strictly nonzero—are placed on the system. First we look to what is perhaps the most unpredictable case of input and inhibition parameters all being very close to one another. It gives a sense of commonality, of "average" behavior. Later we will look to the extreme cases of a large difference maintained between input and inhibition, and those are perhaps slightly more predictable given their similarity to the past two sections.

Under these circumstances, the system always bifurcates between one of three common phase plane configurations. The effect of changing N is on the amount of time the system spends in a given configuration, and thus how long it maintains the competitive state.

We will begin by showing a bifurcation diagram of N vs. x, taking the input and inhibition to be very small (Fig. 6). This bifurcation diagram is a sketch of the movement of the fixed points as the N parameter varies. Dark lines are stable fixed points, and broken lines are unstable fixed points. This bifurcation diagram for small input and inhibition is indicative of how the fixed points change location after increasing those parameter values as well. As usual, the difference is merely an issue of small numerical changes.

As is easily seen, there is one stable steady state which lingers permanently at $x = 1$, and two additional stable/unstable pairs that bifurcate away as N increases, one pair with x-value near 1, and the other near 0. From this diagram, we can construct the basic configuration of the phase plane for any N-value on the continuous spectrum from 2 to 4, without actually computing the phase plane image. Before the first bifurcation, the system has three stable fixed points at $(0, 1)$, $(1, 0)$, and $(1, 1)$. These fixed points correspond to y's win, x's
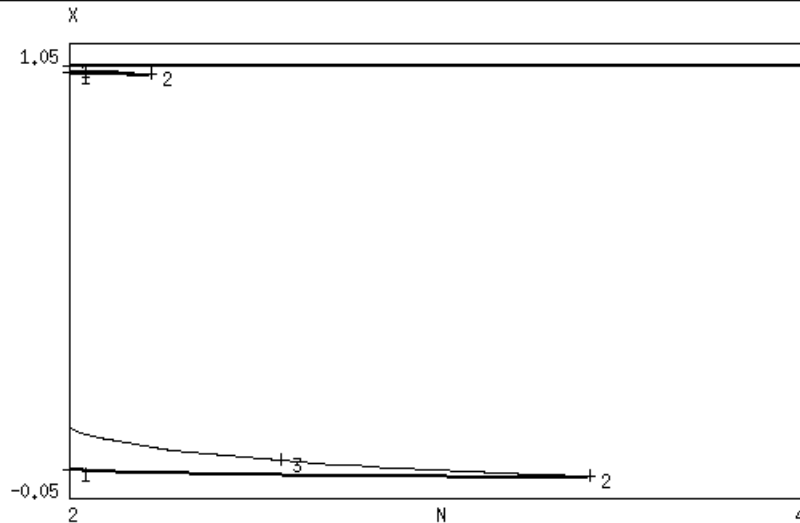
Figure 6: The bifurcation diagram N vs. x with other parameters $\beta = I = 0.2$ and $J = 0.1$. First saddle-node bifurcation at $N = 2.222$ and second at $N = 3.4159$.

win, and mutual activity, respectively (Fig. 8a). After the first saddle-node bifurcation (Fig. 8b), we are left with only two stable fixed points at $(0, 1)$ and $(1, 0)$ respectively, instigating competitive behavior. Then another bifurcation occurs (Fig. 8c), which removes two of the centrally-located unstable fixed points. This bifurcation is relatively minor, however, because it doesn't add or remove outcome possibilities, only warping the path of the trajectories. Now we have come to a rather canonical image of competition (Fig. 7d), with trajectories approaching the fixed points representing each respective group's win and the saddle—in this case, notice its location near $(0, 0.5)$—determining their paths. Soon after, a final bifurcation (Fig. 7e) removes the possibility for y to win, leaving a lone fixed point at $(1, 0)$ and—finally—inevitable victory for the predetermined x-group (Fig. 7f).

The bifurcation diagrams and phase plane images have all made reference to the continuous change in N, but what is the discrete analogue of these results as applications to the notion of population competition? We see that the initial possibility for any population to win or all to become equally active exists only for the N-value 2. Competition is only encountered with the N-value 3, and for all $N \geq 4$ the x population must win. In terms of the discrete analysis, the window of competition is very small—just a single value—when the input and inhibition are small themselves. This bodes the question of whether that window might expand or contract when we change the specific values of the input and inhibition while maintaining the idea that the input and inhibition are close to one another. The answer is yes, the window does change. The competition window expands as input and inhibition grow. To be exact, when the values of input and inhibition are simultaneously translated by 0.4, the new window of N-values for competition increases by exactly 1, i.e., we identify one additional integer N as conducive to competition. Thus, when we take $\beta = 0.6 = I$ and $J = 0.5$, the competitive window now includes both $N = 3$ and $N = 4$. The expansion is clearly on the right side and not the left. In a qualitative sense, this pattern of short-lived competitive behavior for weak parameter values and longer-lived competitive behavior for stronger ones indicates the system's dependence on energetic parameters—"fuel" in a sense—to combat its own fated x win.

4.4 Low Inhibition, High Input

Now we choose to skew the disparity between the values of input and inhibition, while keeping them all nonzero. This addresses the physical situation in which all channels of interaction among populations are open, but the input is slated as more active. Recalling from earlier that inhibition was determined to be the more change-
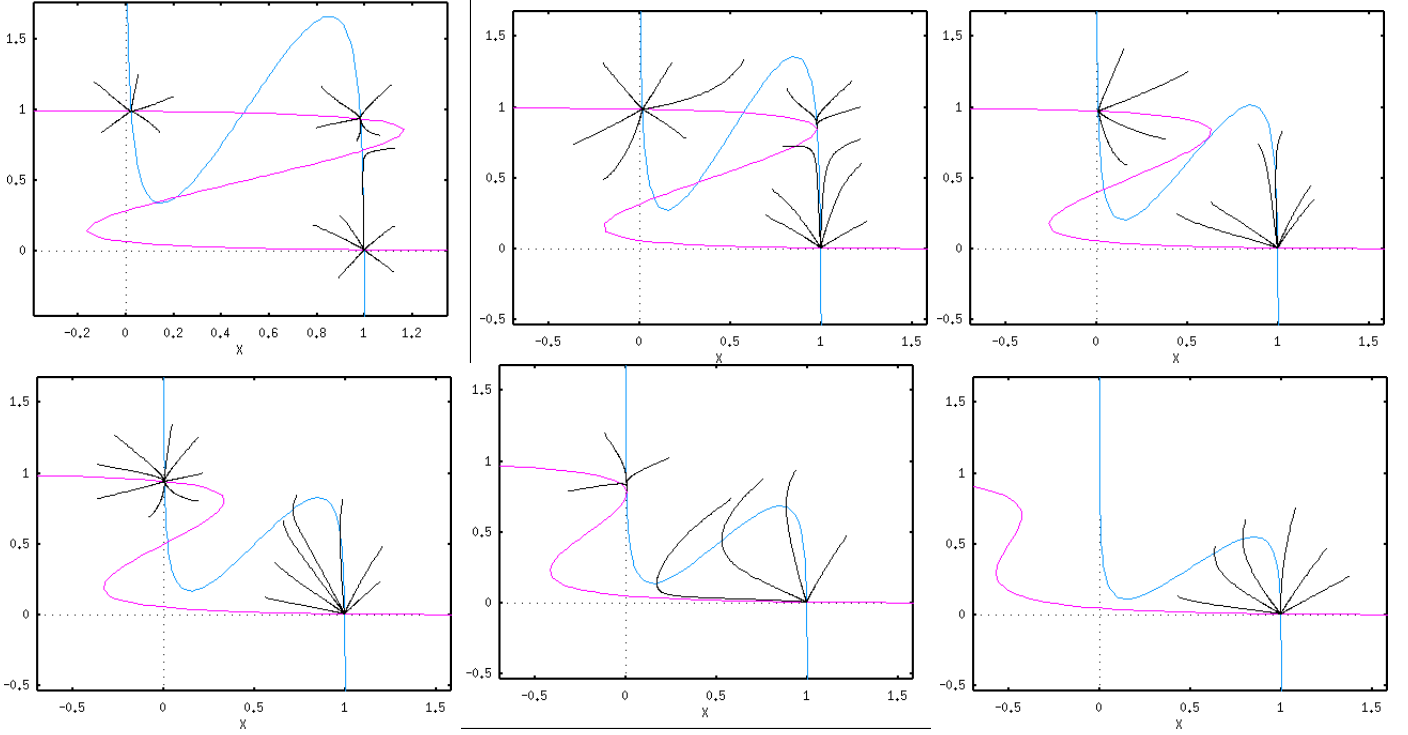
Figure 7: Phase plane images of bifurcation points and competition as N varies for the parameter conditions $\beta = 0.2 = I$ and $J = 0.1$.

effective parameter, this will test whether small inhibition, like totally absent inhibition, creates absurd system dynamics. As it turns out, there is a solid level of truth to this hypothesis: as earlier, we encounter instances of the same one stable steady state at $(1, 1)$ phase plane configuration. However, the inclusion of this nonzero inhibition has served to complicate the bifurcation diagram.

Just as last time, we will start by fixing the difference between the input and inhibition, and then setting small starting values for the smaller parameter (inhibition). While maintaining the usual difference of 0.1 between I and J, fix a constant difference of 0.5 between β and I (which implies a constant difference of 0.4 between β and J).

First observing the system at $N = 2$, the phase plane starts off with one stable steady state at $(1, 1)$ (Fig. 8a). The nullcline curves are very close to one another, and at $N = 5.72$ there is catalyzed a saddle-node bifurcation producing a second stable steady state at $(0, 1)$ and saddle between it and the original stable point. This will turn out to be a very large N-value for this initial bifurcation point. This of course produces the possibility for y to win, and ironically the system seems to allow for either the losing group to win or a state of mutual activity (Fig. 8b). As N increases, all three fixed points travel towards the x-axis, before a second saddle-node bifurcation annihilates the same fixed points that were just created. In the meantime, the y-components of the stable steady states are decreasing, corresponding to an atrophy of the losing groups' stamina. This second bifurcation happens at $N = 13.24$, which will also prove to be a very large value of the second saddle-node bifurcation. At the bifurcation point, the original stable steady state is on course towards $(1, 0)$, which it approaches as N goes to infinity, representing the eventual dominance of the predetermined x-group (Fig. 8c). Also, for small windows of very high values of N, the nullclines bifurcate again, producing circumstances under which both groups become inactive, with firing rates close to 0. This happened at the N-value of 14, which will be a peak, corresponding to the necessity of *many* distractors being added to the image before competition goes to the x-group.

This continuous analysis' discrete analogue is an interesting inverse of the last case with equal input and
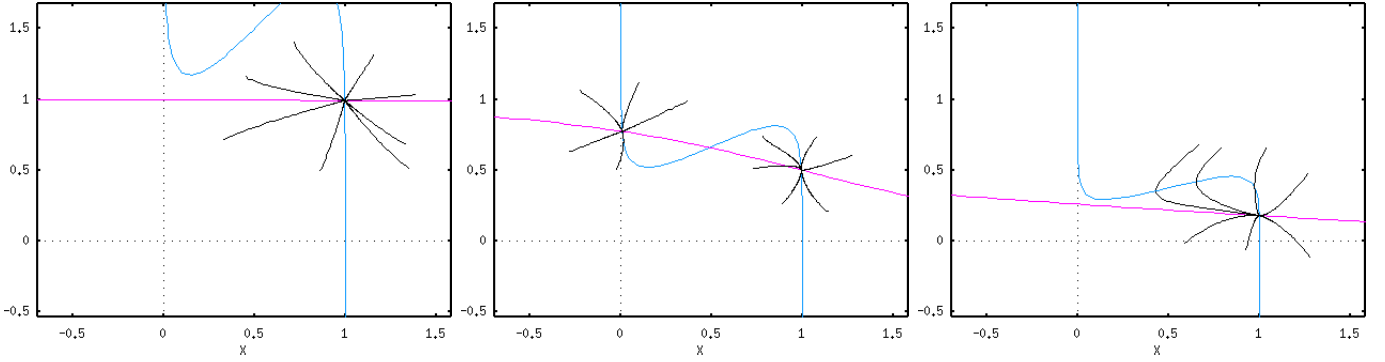
Figure 8: Phase plane sequence as N increases for large input ($I = 0.6$ and $J = 0.5$) and small inhibition ($\beta = 0.1$).

inhibition, which maintains the ability to be broken down into three major phase plane configurations that follow one another as N changes. This time, for small N, no one can win as both groups are firing equally. The pocket of competition seen before at intermediate N has been replaced by an ironic period of possibility for either y's win or mutual activity. N values once opportunistic for both groups are now conducive to "false" wins and bizarre behavior. However, the system does resolve itself as earlier, with an inevitable x win for large enough N.

This time, as we look to change the non-N parameters, we increase by increments of 0.1 because the parameter values are so widely spread that large jumps will throw the system conditions far outside the realm of feasibility quite quickly. Unlike the last case in which the N-window grew alongside the parameters, this time it is shrunk drastically and also translated toward the origin.

Thus, as input and inhibition are increased equally in a linear fashion, fewer and fewer distractor objects need be added to the composition of the image in order to guarantee domination by the x-group. Another consequence of these particular parameter increases is that the window of N-values near 2 during which neither group can win is also shortened, forcing one group to win most of the time. Another noteworthy "anomaly" of the window between saddle-node bifurcations is that—with the exception of low values within the window—the winning group does not entirely suppress the activity of the losing groups. In this N-neighborhood, the losing group maintains a solid firing rate, often near half that of the winning group. The fight between x and y is nuanced, on some level, for the first time. The fixed points that bifurcate in this case do so after lengthy movements that represent massive decreases in their firing rates (albeit ones that maintain the overall trends in dominance); in the past, the fixed points remained stationary and it was the nullclines that slid around under them, even through bifurcations. A group loses not by having its activity instantaneously eliminated, but by being slowly weakened to a point beneath some threshold—no longer a strict zero—at which the other group's firing rate is comparatively large enough to be deemed winner.

4.5 Low Input, High Inhibition

Now we flip the skewness of the relationship between the input and inhibition parameter values to make the inhibition large and the input small. Recalling the poignantly strong effect of inhibition on the system—in particular lobbying for x—we anticipate dramatic behavior. The most noteworthy distinction between this circumstance and the last is that the system has already entered into a state of competition by $N = 2$ (Fig. 8a), with the possibility of the y-groups' win being short-lived as annihilation of the fixed points which allow for it happens quickly thereafter. The more pertinent window of N-values in these configurations is not one of competition or anomalous behavior, rather that in which a win by x is guaranteed—the very situation we set out for Fig. 8b. For relatively high values of N, a saddle node bifurcation creates a steady state at the origin, representing suppression of both groups' activity, as always (Fig. 8c). This is maintained as N becomes
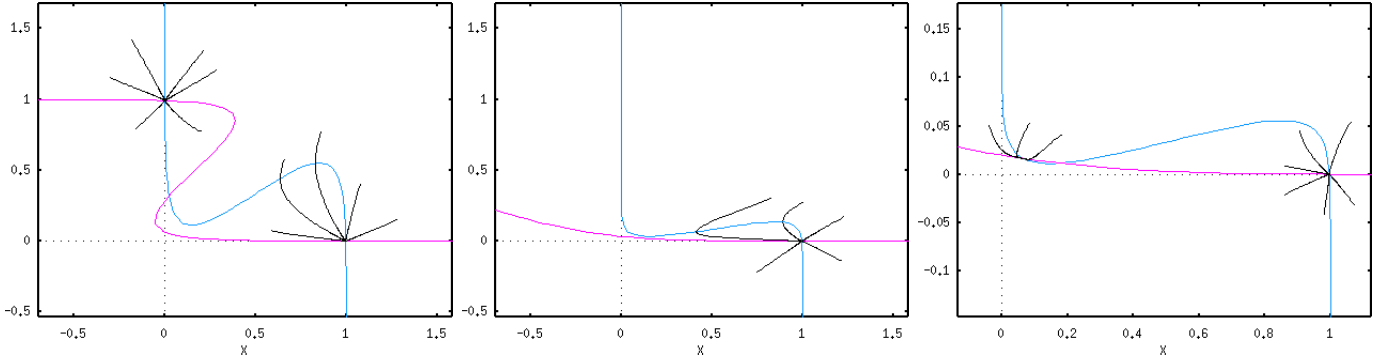
Figure 9: Phase planes for large input and low inhibition, $\beta = 0.6$, $I = 0.2$, $J = 0.1$. In (a) $N = 2$, (b) $N = 5$, (c) $N = 11$. First bifurcation occurs at $N = 2.48$, second at $N = 8.125$.

arbitrarily large.

As we increase the starting parameter conditions by units of 0.1, the central N-window, now a period of x-dominance, expands very quickly: In the circumstance above, the N-value at which begins the simultaneous possibilities of x winning and mutual inactivity is 9, but after only 4 parameter increases, it has moved to 38. This explosion of the N window is much greater than anything we have seen before. We see the powerful inhibition maintain its footing in this case given that for large N, the suppressive behavior is sustained. (The fixed point at $(0, 0)$ never vanishes for high N). It effectively does its job of rendering both populations—in particular x—less powerful (in some sense, this is a lesser victory for y), whereas large input was not enough to help y suppress x for very many N values. In other words, high inhibition was much stronger in achieving the theoretically "impossible" behavior than high input, which, to the contrary, fueled the pre-programmed outcome. This might be taken as a credit to inhibition, but even that is not entirely deserved because inhibition was only able to remove the predestined result, not actually enabling the actively-prevented result to transpire. We might have anticipated that external stimulus would translate nicely to inhibition—if a population is being stimulated, wouldn't it have lots of fuel to suppress the others? But the fact that high levels of input yield furtherance of the desired result suggests otherwise, that input heightens the mechanism already in place.

5   Saddle Points and their Eigenvalues

We have spent much time analyzing many common phase plane arrangements, and we will now focus in on one. We draw our attention to the competitive phase plane, that which, under certain parameter conditions, sports three fixed points: two stable fixed points near $(1, 0)$ and $(0, 1)$, and a saddle somewhere between them. The first two fixed points are of course the x-dominant and y-dominant, respectively, and the saddle the source of manifolds that steer a given trajectory into one of them. A *manifold*, in this context, is a curve in the phase plane assigned to a saddle point, which partitions the phase plane in such as way as to give us data on how the trajectories relate to the saddle. Each saddle is accompanied by two manifolds that intersect at the point itself, one called *stable* and the other *unstable*. If a trajectory is initialized on the stable manifold itself, it will be drawn directly into the saddle. And if a trajectory is drawn into the saddle, one can be sure it came from the stable manifold. Similarly, a trajectory is thrust directly away from the saddle if and only if it was initialized on the unstable manifold. The manifolds define boundaries of the *basins of attraction* of the two stable steady states that accompany the saddle in our system. The basin of attraction of a fixed point (x, y) is the set of all points in the phase plane at which a trajectory can be initialized and definitely approach (x, y). A trajectory that began interior to a basin of attraction can be thought of as gliding towards the saddle alongside the nearest stable manifold, then turning and gliding away alongside the nearest unstable manifold.

An analysis of the properties of the saddle point's total derivative provides crucial information about the

system given certain parameter conditions. As it turns out, the eigenvectors of the saddle point help to linearize the manifolds passing through it; one eigenvalue is tangent to the stable manifold, the other to the unstable manifold, at the saddle point. Their corresponding eigenvalues measure the magnitude and direction of a trajectory's motion with respect to the eigenvector, i.e. the manifold. Negative eigenvalues correspond to the stable manifold, and positive eigenvalues to the unstable manifold. It follows that if an eigenvalue relates to speed, we can use them to get a sense of the *time course*, or the amount of time it takes for a trajectory to "reach" a fixed point.

5.1 Numerical Computation of Manifold Structure

If we denote the system (1)–(2) generically as

$$\dot{x} = f(x, y),$$
$$\dot{y} = g(x, y),$$

the total derivative of the system is as follows

$$\begin{pmatrix} \frac{\partial f}{\partial x} & \frac{\partial f}{\partial y} \\ \frac{\partial g}{\partial x} & \frac{\partial g}{\partial y} \end{pmatrix} = \begin{pmatrix} -1 + \alpha\phi' & -\beta(N-1)\phi' \\ -\beta\phi' & -1 + (\alpha - (N-2)\beta)\phi' \end{pmatrix}, \tag{6}$$

where ϕ' is the derivative of the input-output transfer function evaluated at the fixed point under consideration. The eigenvalues of the total derivative matrix (6) determine the stability of the fixed point.

In this section, we will measure the eigenvalues numerically as discrete functions of N. A trajectory is initiated at the origin, and measured for time immediately upon reaching the ball of radius $0 < \epsilon \ll 1$ about the fixed point. Given the nature of the trajectory's decay, it never actually reaches the fixed point, so we can only estimate its travel time with the epsilon-ball method. To maintain experimental consistency, it is necessary to ensure that all trajectories initiated at the point $(0,0)$ approach the same fixed point, i.e., that the origin remains interior to the same basin of attraction for all N. This is precisely the ultimate purpose of using slightly higher inputs for the target than for the distractors.

In all cases, the saddle being observed is the "last saddle standing", i.e., the last to be annihilated by a saddle-node bifurcation before the system settles with one stable steady state. It is observed throughout the window of N-values during which it is the only saddle point among two stable fixed points. Therefore, the interval over which it is measured begins with the N-value corresponding to the annihilation of the phase plane's other saddle, and ends with the N-value representing its own destruction. We want a significant but not too large data set, and 20 measurements seem an adequate sample size. We determine the N-values at which to collect data by partitioning the interval into twenty equally-sized smaller intervals and using the boundaries.

In terms of the intuition of the implications of the system's linearization at the saddle point, eigenvalues of small norm correspond to low speeds and catalyze long time courses. On the other hand, eigenvalues of large norm correspond to high speeds and quick time courses. Thus, when the eigenvalues simultaneously approach zero, the speed decreases in both the direction of the stable and the unstable manifold, and it is clear in both theory and practice that the overall time course must decrease. Similarly, if the norms of both eigenvalues approach infinity, the speed alongside both manifolds increases and the time course drops rapidly. However, the system resolves to neither of these states. For all measured circumstances of input and inhibition, the positive eigenvalue and negative eigenvalue decrease simultaneously as N is varied. This lack of predictability necessitates a direct and precise analysis of the time courses themselves.

The particular parameter conditions under which the eigenvalues and time courses have been measured are the instance in which β, I, and J are nonzero and close in value, as well as the instance in which there is a large disparity between values of β, I, and J, weighted in favor of the inhibition. Below is displayed a sample, the graph of time course as a function of N (Figure 10) when inhibition is small and inputs moderate. For lower N-values in the range of 6 to 9, there is a generally monotonic growth to the graph, as predicted by the experimental hypothesis.
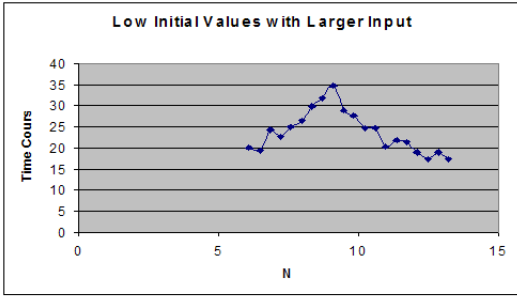
Figure 10: Graphs of search time vs. N. Small Values, larger input: $\beta = 0.1$, $I = 0.6$, $J = 0.5$.

6   Acknowledgements

References

[1] Wilson, H.R. (1998), *Spikes, decisions, and actions. The dynamical foundations of neuroscience*, Oxford.

[2] Treisman, A. (1982), Perceptual grouping and attention in visual search for features and for objects, *J Exp Psychol Hum Percept and Perform*, **8**(2), 194–214.