SYMPOSIUM: TOWARD A MORE NEUROTROPIC TRANSFER FUNCTION

Chaired by Frederick J. Bremner, *Trinity University*

Multiphasic neuronal transfer functions for representing temporal structure

ADAM PHILIP KING and C. R. GALLISTEL
University of California, Los Angeles, California

Neural network models of timing have struggled to account for animal timing capabilities using the accepted connectionist assumptions, in most cases without postulating the existence of explicit neuronal time-keeping mechanisms. Current ethological and physiological data, however, suggest that cellular oscillators form the foundation for animals' temporal capacities. We propose that these oscillators could be used as temporal filters that capture the temporal structure of the animal's experience. A model is presented that accounts for a number of the salient features of the feeding anticipatory response with only a single circadian filter. This model goes beyond current entrainment models in that it correctly predicts the relationship between the feeding period and the anticipatory interval. An alternative approach using multiple filters is examined that can account for animals' ability to correctly anticipate two daily feeding times.

Although neural networks have been successfully applied to problems of a spatial nature (in the sense of spatial pattern recognition), they have been deficient in their temporal-processing capabilities. Many attempts at solving the temporal deficiencies have been profoundly non-biological and have not taken seriously either the primary role that time plays in animal cognition (Gallistel, 1990) or physiological evidence suggesting that timing mechanisms are a cellular phenomenon (Jacklet, 1989; Takahashi & Hoffman, 1995).

The research presented here takes at face value the ethological and physiological evidence that points to animals having built-in, self-sustaining oscillators that form the foundation for their temporal-processing capabilities. In this paper, we examine how by placing biologically defensible clocks into the simulated neuron, we give individual neurons a multiphasic transfer function that makes them act like temporal filters. The outputs of banks of such filters could capture the temporal structure of the animal's experience. In this paper we explore how such a model can account for animals' ability to anticipate daily feeding times. We hope that this approach may ultimately unify the analysis of a variety of temporal-processing problems that animals routinely face, including the learning of temporal intervals in classical conditioning.

An Animal's Anticipation of Daily Feeding Times

Rosenwasser, Pelchat, and Adler (1984) examined the ability of rats to anticipate a daily feeding time. Initially a rat was given constant access to food. Under these conditions, the rat showed its regular onset of activity (wheel running) when the lights went out. After about 2 weeks, the rat was given access to food only for a 2-h period during its subjective day, a time it would not normally be active. After a few days of training, the rat reliably showed anticipatory activity (wheel running) before the scheduled feeding time. On Day 43, the rat was returned to constant food access conditions, and it immediately went back to its regular activity cycle. On Day 77, the rat was exposed to complete food deprivation, and remarkably, the rat returned to the anticipatory behavior it had shown previously. This reinstatement phenomenon was demonstrated up to after 50 days of constant feeding conditions, the longest period tested.

The feeding anticipatory response demonstrated in the Rosenwasser et al. (1984) experiment shows that animals have the ability to remember the time of day that a significant event occurs and make use of this information in their decision-making process. This ability has been demonstrated even in constant light conditions and is not dependent on external cuing. Some salient aspects of the feeding anticipation that we will attempt to explain are as follows: (1) It takes 3–6 days of regular feedings before rats show clear signs of anticipatory activity. (2) The feeding period does not have to be exactly 24 h, but rats will not

Correspondence should be addressed to A. P. King, University of California, Los Angeles, Computer Science Department, 3732 Boelter Hall, Los Angeles, CA 90024-1600 (e-mail: king@cs.ucla.edu).

learn to anticipate meals that recur with periods shorter than about 22 h or longer than 27–29 h. This has been one of the principal findings thought to support the notion that anticipatory feeding is mediated by food-entrainable circadian oscillators (Mistlberger, 1994). (3) As the period of the feeding cycle gets shorter, the amount by which the onset of the rat's anticipatory activity precedes the onset of food availability (the length of the anticipatory interval) gets shorter; as the cycle gets longer, the anticipatory interval gets longer. Most importantly, the slope of the function relating the length of the anticipatory interval to the period of the feeding schedule is significantly greater than 1 (Aschoff, Goetz, & Honmäh, 1983). For each hour increment in the period of the feeding schedule, there is an increase of about 1 h and 20 min in the duration of the anticipatory interval. (4) Rats can learn to anticipate each of two different daily feedings 6 or more hours apart (Bolles & Lorge, 1962; Bolles & Moot, 1973; Edmunds & Adler, 1977).

Neural Network Models and Time

Artificial neural network researchers have proceeded with the belief that the information-processing aspects of the nervous system can be successfully abstracted away from the biological particulars. The models usually consist of (1) units or "neurons" that communicate with each other via weighted links; (2) a transfer function within each neuron, which specifies how it computes its state of activation or output as a function of the many inputs received; and (3) a learning rule, which specifies how the weights of the links change as a function of experience. Traditionally, the transfer function has performed a sum and threshold, causing the cell to output a scaled sum of the inputs if this sum exceeds some critical value.

Although this abstraction has proven quite successful for a number of models, the commitment to this view is poorly suited to explain the temporal processing seen in the above examples. Connectionist networks are usually discretely clocked in the sense that at every time step, each neuron computes an updated value. This clocking is known to be biologically indefensible. In addition, it does not provide for any form of *time keeping* within the network. To add temporal capabilities to the connectionist models without doing violence to the abstraction, some researchers have worked with converting time into spatial inputs (Lang, Waibel, & Hinton, 1990) while others have studied recurrent networks (Elman, 1991). They have been somewhat successful at capturing the ordering property that time has, but they have not been able to capture the metric aspect of time. Some success in this area has been achieved by parameterizing the conduction time in the connections between units (Beroule, 1987). This approach shows some promise in concurrency detection over short time intervals—for example, in binaural hearing—but it is difficult to see how this idea could plausibly be extended to handle times substantially longer than 1 sec.

Thus, traditional neural network models do not allow the animal to represent the temporal properties of its

experience without severe stretching of their assumptions. A notable exception to the general timelessness of neural net models is the Church and Broadbent (1990) model for temporal interval learning. It is similar to our model in that it takes endogenous oscillations of the neurons as fundamental to the timing process. These oscillators are used to create binary-valued vectors that represent experienced temporal intervals (interval vectors). Their model stores and retrieves the values of these interval vectors by creating a matrix of connection weights in an autoassociative memory. The stored vector is read by passing a retrieval vector through the autoassociative memory.

The Phase-Quadrature Filter Pair Model

Our model uses the oscillators in a very different way, a way that avoids storing and retrieving the intervals. In our model, the oscillatory capacity of the neurons (units) is exploited to make them into temporal filters. Information about the temporal structure of the past is preserved not in the form of stored interval vectors but rather in the form of the on-line output from banks of temporal filters.

Phase-quadrature filter pairs. We base our modeling effort on the well-established temporal-filtering capacity of the nervous system—the ability of neurons and neural circuits to respond selectively to stimuli of certain frequencies. We propose that oscillator mechanisms exist at the cellular level and that the weight the neuron applies to an input signal at a given moment in time is determined by the state of the neuron's endogenous oscillator at that moment. It gives positive weight to inputs received during one phase of its cycle but negative weight to inputs received during the opposite phase of its cycle. These weighted internal reactions to past values of the input decay exponentially. The output of the filter at the current moment in time (the signal it supplies to the rest of the system) is the sum of its exponentially decaying past reactions. Thus, the neuron has the transfer function of a temporal filter tuned to the frequency of the cell's endogenous oscillation. Like Church and Broadbent (1990), we follow (Gallistel, 1990) in assuming that the periods of the oscillators in the filter bank—hence the

Figure 1. Two filter functions with different decay rates. The present moment (a moment at which the filter output is to be measured or evaluated) is designated t ; any moment in the past at which some input arrived is designated τ . The interval $(t - \tau)$ is the retrospective interval.

frequencies to which they are sensitive—range from milliseconds to years.

The weight that such a neuron applies to a stimulus input is a sinusoidally varying function of time, τ . If we choose some present moment, t , to measure the output of the filter, the effect of an input at some previous moment, τ , is an exponentially decaying function of the interval $(t - \tau)$ between the present moment and the earlier moment τ . Figure 1 shows two filter weighting functions with different rates of decay. The contribution of a weighted input to the current output of the neuron is an exponentially decaying function of the interval $(t - \tau)$ that has elapsed since that input was received.

A well-known feature of endogenous oscillators is that they may be coupled so that they cycle in some fixed phase relationship. We assume that our filter neurons come in pairs, matched for frequency and decay rate, and coupled so that they run 90° (a quarter cycle) out of phase, that is, in a sine–cosine phase relationship. We call such a pair of matched and coupled filters a phase-quadrature or p-q filter pair.

Analytic characterization. Our model, unlike conventional neural networks, is a continuous (unlocked) system with real and specified dynamics. In this regard, it is much closer to neurobiological reality. The dynamics are given by the transfer functions of the filters, which are continuous functions of the past history of the input. However, in computing filter outputs, we follow the usual numerical approximation approach; that is, we break time up into small increments and evaluate the filter's output, F , at successive small increments in time. The evaluation increment ($\Delta\tau$) is a small fraction of the period of the filter (1/20th or less). The updating function is as follows:

$$F(\tau + \Delta\tau) = \alpha^{\Delta\tau\omega} F(\tau) + (1 - \alpha) f(\tau) S(\tau), \quad (1)$$

where τ is time relative to some 0 point in the distant past, $F(\tau)$ is the filter's output at time τ , $f(\tau)$ is the weight applied by the filter to an input at time τ , and ω is the frequency of the filter (cycles per day). $S(\tau)$ is the magnitude of the input stimulus at time τ , and α is the decay rate (the fraction of the output remaining after an interval equal to one period of the filter's cycle). The increment, $\Delta\tau$, is some fraction (e.g., $1/144$) of the period of the filter. In other words, the size of the temporal increment is scaled by the period of the filter.

The impulse response of the filter is an exponentially decaying cosine or sine function, depending on which member of the p-q filter pair one is considering. Both the cosine and the sine impulse responses are given by the complex exponential $f(t - \tau) = e^{-p(t - \tau)}$, where p is a complex number and $(t - \tau)$ is a retrospective temporal interval. The real part of p , $Re[p]$, is the decay constant (the neper frequency), which specifies the rate at which the influence of an input on the current output of the filter declines as the time of occurrence of that input recedes into the past. The further back in time the input occurred, the greater the value of the term $(t - \tau)$. The imaginary part of the complex coefficient, $Im[p]$, is the frequency of sinusoidal variation, because by Euler's identity, $e^{-i(t - \tau)} =$

$-(\cos(t - \tau) + i \sin(t - \tau))$, where $(t - \tau)$ is scaled in radians. In this expression for the impulse responses of the filter pair, the phases of the cosine and sine variations are fixed with respect to the arbitrarily chosen present moment t , because when $\tau = t$, $(t - \tau) = 0$, and $(t - \tau)$ is the argument of the cosine and sine functions. In other words, these functions always have 0 phase at the present moment. This is physically impossible; the phase of a cell's endogenous oscillation is fixed relative to absolute time, τ , not relative to an arbitrarily chosen present moment t . However, the complex exponential may be rewritten in terms of the product of a real exponential and an imaginary exponential:

$$\begin{aligned} e^{-px} &= e^{-Re[p]x} e^{-iIm[p]x} \\ &= e^{-Re[p]x} (-\cos Im[p]x - i \sin Im[p]x). \end{aligned} \quad (2)$$

This enables us to make absolute time τ the argument of the oscillatory component of the filter's impulse response (the cosine and sine functions) while keeping $(t - \tau)$ as the argument of the decaying component. Keeping $(t - \tau)$ as the argument of the decaying component makes inputs that occurred further back in the past have less influence on the current output of the filter, because as $(t - \tau)$ gets large, $e^{-(t - \tau)}$ gets small.

If we use Equation 2 as the expression for the impulse responses of the filters, then, as $\Delta\tau$ becomes arbitrarily small, the outputs calculated by the discrete updating Equation 1 become equivalent to the convolution of the filters' (complex) impulse response with $S(\tau)$, the signal generated by stimulus onsets:

$$F_{\omega,d}(t) = \int_0^t S(\tau) (-\cos \omega\tau - i \sin \omega\tau) e^{-\omega d(t - \tau)} d\tau, \quad (3)$$

where ω is the filtering frequency in radians per unit of time and d is the decay constant. Notice that there are three factors inside the integral. The first, $S(\tau)$, is the input stimulus as a function of time. The second, $(-\cos \omega\tau - i \sin \omega\tau)$, specifies the varying weight that the filter gives an input depending on where in the filter's cycle the input arrives. The third, $e^{-\omega d(t - \tau)}$, causes the effect of an input at time τ to decay as time t advances, that is, as the input recedes into the past. Notice that ω , the filter frequency, multiplies d , the decay constant. This means that d specifies the amount of decay that will occur in one cycle; thus, ωd plays the same role as α in the discrete updating equation.

$F_{\omega,d}(t)$ is a complex variable; its real part is the output of the cosine filter and its imaginary part is the output of the sine filter. These filter outputs are in essence Laplace coefficients, because Equation 3 is closely related to the Laplace transform:

$$F(p) = \int_0^\infty S(t) e^{-pt} dt.$$

The Laplace transform is just the Fourier transform:

$$F(\omega) = \int_{-\infty}^\infty S(t) e^{-i\omega t} dt,$$

but with a finite lower limit of integration and a non-zero value for the real part of the complex coefficient in the

exponential. This non-zero real part of the exponential coefficient makes the sinusoidal components decay exponentially, rather than extending from past infinity to future infinity, as do Fourier components. Equation 3 is a real-time or on-line version of the Laplace transform, with the upper limit of integration set at t , the real part of the exponential set to $\omega d(t - \tau)$, and the imaginary part set to $\omega\tau$.

Laplace coefficients, when normalized for the decay rate, are like Fourier coefficients; they specify the phase and magnitude of the energy in the input signal at or near the filter frequency. By comparing the signs and relative strengths of these two values, we can determine important properties of the input signal. The quadrant in which the stimulus tended to fall is reflected by the signs of the filter outputs. The relative magnitudes of these outputs indicate where within the quadrant the stimulus tends to occur. Even more simply, one can think of the output from the filters (sign included) as indicating a vector that points to the time on a clock that has the period of the p-q filter pair.

The Laplace transform reflects the temporal structure of the series seen through filters that decay exponentially. That is, it converts the stimulus sequence from the time domain to a localized frequency domain, the frequency content of the stimulus in the recent past. Each p-q filter pair yields information about the time sequence for a particular frequency. How far back a filter pair views as recent is reflected in the decay constant. The complete model consists of a two-dimensional p-q filter pair grid, in which one dimension consists of p-q filter pairs of varying frequencies and the other dimension consists of varying decay constants. The result is that for any particular stimulus, the system will be calculating in parallel a whole range of frequency information on varying time scales of varying temporal locality.

Implementation. The p-q filter pair grid model is implemented in C++ on a Macintosh. The program creates a grid with a variable number of frequencies (where frequency $1 = 1 \text{ cycle/day} = 2\pi \text{ radians/day}$) and a variable number of decay rates (specified as a percentage decay per unit cycle). Given a stimulus function that indicates the presence of an impulse stimulus as a function of time, and a total number of days to run, the program generates the Laplace coefficients for the grid using the updating function presented in Equation 1. For each filter, $\Delta\tau$ was set to be $1/144$, chosen to produce a sampling rate equivalent to 1 sample per 10 min in the circadian oscillator. A filter's position in the grid is defined by the values for its frequency, ω , and its decay rate, α . The output of such a filter at time t is $A(\omega, \alpha, t)$ (the cosine coefficient) and $B(\omega, \alpha, t)$ (the sine coefficient).

The model assumes that the filter banks remain phase locked to the circadian schedule, and that the circadian oscillator is itself entrained by a 24-h day/night cycle. Noise was added to the filters by adding a random normal variable to each oscillator's phase at the time of input. The random phase perturbation had a mean of 0 and

some standard deviation (ψ). Unless otherwise noted, simulations in this paper were run with $\psi = .007$, which reflects a standard deviation of 5 min in the accuracy of the circadian clock.

In each of the simulations below, we assume that, in addition to the outputs from the filter bank, there is also a single bank of paired oscillators. These oscillators run at the same frequencies as the filter oscillators but with a phase advance ϕ equal to the desired anticipation interval. The vectors that describe the current state of these oscillators are designated $C(\omega, t + \phi)$ and $S(\omega, t + \phi)$. A value for ϕ of 0.083 days gives the 2-h anticipation interval typical of food anticipatory activity. We use the vectors from these phase-advanced oscillators together with the output vectors from the filter grid to generate inputs to the decision processes that determine the onset of anticipatory activity.

Results

Anticipating a single daily feeding. The model was able to anticipate a single feeding time delivered at 24-h periods, with the use of a single circadian oscillator pair generating the coefficients $A(1, \alpha, t)$ and $B(1, \alpha, t)$ —the cosine and sine outputs from the 1-cycle-per-day (circadian) filter with decay rate α . Figure 2 shows a clock face representation of this single filter, receiving a single impulse stimulus every 24 h. Results are shown for 4 different days. Each clock face represents a 24-h clock with a single hand that moves clockwise around the circle, with 12:00 a.m. facing north. The stimulus was delivered at 1800 h (due west on the clock faces). The days are shown from left to right, and sets are shown for 5 differ-

Figure 2. Outputs of circadian phase-quadrature filters with different decay rates for an impulse stimulus delivered once every 24 h. The stimulus was delivered at 1800 h, which is indicated by the triangle. The filter output was sampled immediately after each feeding.

ent decay rates, $\alpha = .3, .5, .7, .9$, and $.95$. The output from the filters shown was taken immediately after the stimulus impulse was delivered. This particular simulation was run with the noise, $\psi = 0.05$. This high degree of noise (standard deviation of 72 min) was chosen only for illustrative purposes. When the model was tested for its predictions, the more reasonable level of $\psi = 0.007$ (standard deviation of 5 min) was used.

The first thing to notice is that each clock face has its hand pointing in the direction of the time of the input stimulus. By comparing the output of these filters to the time on the clock that is phase advanced 2 h into the future, the model can determine whether or not the predicted feeding time is within 2 h of the current time. This calculation is made by normalizing the filter's output vector to Length 1. Then, we use the trigonometric formulae for the cosines and sines of the difference of two angles to obtain the cosine and sine of the phase difference, $\cos_{\text{diff}}(t)$ and $\sin_{\text{diff}}(t)$. With $\cos_{\text{diff}}(t) > 0$ as an enabling condition, we can use $\sin_{\text{diff}}(t) > 0$ as our criterion for the onset of anticipatory activity.

Although the above equations will correctly show anticipatory responses 2 h prior to the predicted phase angle of feeding, whether or not to show anticipation at this time must also be a function of the strength of the output vector, which is calculated as the amplitude or length of the vector:

$$\text{strength}(t) = \sqrt{A(t)^2 + B(t)^2}.$$

The strength of the filter's response is indicated in Figure 2 by the length of the clock's hand. For each decay rate, each hand develops a stronger response over the successive days. (The responses of the filters are calculated with normalization so that the asymptotic response to that filter's ideal input is the same for every filter.) The decay rate determines the amount of time it takes a filter to achieve its asymptotic value. It is assumed that the animal will show anticipatory behavior only if the strength of the filter's output exceeds some threshold, T . For $\alpha = .3$, we were able to predict that anticipatory behavior would occur after 3 days of training using $T = 0.9$.

Noncircadian feeding schedules. To test the model's prediction on noncircadian feeding schedules—schedules whose period did not match the 24-h period of the filter—11 simulations were run with the feeding period varying from 19 to 29 h in integer steps. Only the single circadian filter pair was used, with the same parametric assumptions as those made above ($\alpha = .3$, $T = 0.9$). Figure 3 shows the results of these runs. The line labeled "slope = 1" indicates the prediction of a model that only took into account one previous day's feeding in its estimation of the following day's feeding event. The line labeled "slope = 1.3" shows the findings taken from the experimental literature (Aschoff et al., 1983). When the feeding interval was less than 22 or greater than about 27½ h, the rats showed very little or no anticipatory behavior. The predictions of the model are shown plotted as

Figure 3. The relationship between feeding period and the anticipatory interval. The boxes indicate the model's predictions and the line with slope = 1.3 represents the experimental data.

little black boxes. If the filter's output did not exceed the threshold levels, then no box was plotted.

Figure 3 shows that the model can predict correctly both the slope of the function relating the length of the anticipatory interval to the period of the feeding schedule and the limits at which feeding periods will cease to produce anticipatory activity. The model's prediction of a slope greater than 1 is significant because entrainment models of the feeding anticipatory response have difficulty accounting for this result. The slope of this function will vary with the choice of α , although it will always be greater than 1. However, the limits on the feeding period that will be anticipated seems to be fairly insensitive to the parameters. This result needs further study.

Anticipating multiple feeding times. The above results demonstrate that for a single daily feeding event, a model using only one circadian p-q filter pair can account for a great deal of the salient findings regarding the feeding anticipatory response. However, rats have a demonstrated ability to anticipate each of two daily feedings when they are at least 6 h apart. This result places great difficulty in the path of the popular entrainment model for food anticipation. It requires a phase response curve that somehow makes it possible for two different food-entrainable oscillators to entrain separately, one to each of the feeding events, which are indistinguishable except by their time of onset. It has yet to be shown that this is possible.

In our approach, it is clear that the single circadian filter pair used above will not account for the animal's ability to anticipate two feeding events, because the output vector from a single filter pair can point to only one feeding time. In fact, the circadian filter's output vector will point to a time that is intermediate between the two feeding events.

To anticipate two feeding times per day, we adopt a different approach. We use a number of p - q filter pairs that vary in their resonant frequency but have the same decay rate per cycle. This set of filters is represented as $A(\omega, \alpha, \tau)$ and $B(\omega, \alpha, \tau)$, with ω varying in discrete steps equal to the difference in frequencies between filters. This model synthesizes the food delivery schedule to create a food-anticipation function. If the synthesis is satisfactory, the peaks in the food-anticipation function will lead feeding onsets by ϕ . The decision process triggers anticipatory activity when the anticipation function exceeds a threshold. The following equation is used to calculate the anticipation function derived from filters that vary in frequency but share a common decay rate per cycle, α :

$$\text{Anticipation}(t) = \sum_{i=1}^n A(\omega_i, t) C(\omega_i, t + \phi) + B(\omega_i, t) S(\omega_i, t + \phi). \quad (4)$$

This function is derived from the fact that the Laplace transform is invertible—the stimulus function $S(\tau)$ can be recovered from the Laplace coefficients. The Laplace coefficients represent the information regarding the availability of food in the localized frequency domain. Equation 4 converts this representation back into the time domain.

Equation 4 is very similar to the equation for Fourier synthesis (the inversion of a Fourier series) except that we are using a finite upper bound on the summation (there can only be a finite number of filter pairs) and the coefficients are Laplace rather than Fourier coefficients. In the Fourier series, all functions are periodic and the frequencies used are all the harmonics of this basic period (in our case, one cycle per day, two cycles per day, three cycles per day, etc.). Following Church and Broadbent (1990), we contemplated the use of filter frequencies that progressed in logarithmic nonharmonic fashion. This, however, results in an anticipation function that does not have a sustained circadian periodicity. We also explored using only harmonics that form a logarithmic series (1, 2, 4, 8...). This resulted in spurious predictions of feeding times. We have tentatively concluded that our approach will work only if one uses frequencies that are the first several harmonics of the circadian base frequency. The results below are based on the use of the first 10 harmonics, although it appears that only about 4 or 5 are needed to accurately predict the anticipation of two daily feeding times separated by at least 6 h.

Figure 4 shows the result of this model run over 8 simulated days. On Days 1–4, the model received an impulse stimulus at both 1200 and 2200 h. Thus the model received two feeding events that both occurred every 24 h and that were separated by 10 h. On Days 5–8, the feedings were omitted. The y -axis gives the value of the anticipation function. The bars indicate feeding events, with onset occurring at the leading edge. This simulation was run with $\alpha = .8$, and $T = 1.7$. As the figure demonstrates, the model correctly predicted anticipatory activity approxi-

Figure 4. The model's predictions for anticipatory activity shown for two daily feeding times. On Days 1–4 the feedings came on 24-h periods separated by 10 h. On Days 5–8 the feedings were omitted.

mately 2 h in advance of both feedings after Day 3. When the stimulus was taken away, the model still predicted anticipatory behavior for the 1st day that the feedings were omitted. After Day 5, the model no longer predicted an anticipatory response.

Conclusions

By adding a multiphasic transfer function to model neurons, we were able to have neurons in our model behave like temporal filters that capture the temporal experience with an input stimulus. Using just one of these filters, we were able to account for a number of salient features of the feeding anticipatory response. By using the outputs of multiple p - q filter pairs to resynthesize the time domain, we were also able to account for an animal's ability to correctly anticipate two daily feeding events. Although these results are encouraging, the need to use harmonic series—which make the solution unique to a particular base period—makes us doubt that a single model can deal with phase learning, interval learning, and the temporal properties captured in classical conditioning. We do expect, however, that by using the outputs from filter banks with varying decay rates, we can begin to approach the issue of how animals deal with nonstationarity—such as step changes in the feeding time. The explanation of circadian phase learning is an interesting

challenge to those who want to model learning using what we currently understand about the nervous system as their starting point.

REFERENCES

- ASCHOFF, J., GOETZ, C. V., & HONMAH, K. I. (1983). Restricted feeding in rats: Effects of varying feeding cycle. *Zeitschrift für Tierpsychologie*, **63**, 91-111.
- BEROULE, D. (1987). Guided propagation inside a topographic memory. *International Conference on Neural Networks*, **4**, 469-476.
- BOLLES, R. C., & LORGE, J. D. (1962). The rat's adjustment to a-diurnal feeding cycles. *Journal of Comparative & Physiological Psychology*, **55**, 760-762.
- BOLLES, R. C., & MOOT, S. A. (1973). The rat's anticipation of two meals a day. *Journal of Comparative & Physiological Psychology*, **83**, 510-514.
- CHURCH, R. M., & BROADBENT, H. A. (1990). Alternative representations of time, number, and rate. In C. R. Gallistel (Ed.), *Animal cognition* (pp. 55-81). Cambridge, MA: MIT Press.
- EDMUNDS, S. C., & ADLER, N. T. (1977). The multiplicity of biological oscillators in the control of circadian running activity in the rat. *Physiological Behavior*, **18**, 921-930.
- ELMAN, J. L. (1991). Finding structures in time. *Cognitive Science*, **14**, 179-211.
- GALLISTEL, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- JACKLET, J. W. (Ed.) (1989). *Neuronal and cellular oscillators*. New York: Marcel Dekker.
- LANG, K. J., WAIBEL, A. H., & HINTON, G. E. (1990). A time-delay neural-network architecture for isolated word recognition. *Neural Networks*, **3**, 23-44.
- MISTLBERGER, R. E. (1994). Circadian food-anticipatory activity: Formal models and physiological mechanisms. *Neuroscience & Biobehavioral Reviews*, **18**, 171-195.
- ROSENWASSER, A. M., PELCHAT, R. J., & ADLER, N. T. (1984). Memory for feeding time: Possible dependence on coupled circadian oscillators. *Physiology & Behavior*, **32**, 25-30.
- TAKAHASHI, J. S., & HOFFMAN, M. (1995). Molecular biology of clocks. *American Scientist*, **83**, 158-165.

(Manuscript received November 13, 1995;
revision accepted for publication February 6, 1996.)