ENERGY-EFFICIENT RECURSIVE ESTIMATION BY VARIABLE-THRESHOLD NEURONS [1]

Toby Berger
Department of Electrical and Computer Engineering
University of Virginia, Charlottesville, VA 22903

William B. Levy
Department of Neurosurgery
University of Virginia Medical School
Charlottesville, VA 22903

Galen Reeves
Electrical Engineering and Computer Science Department
University of California, Berkeley, CA 94720

Abstract. In time slot k the input X_k to an energy-efficient neuron we shall call $\mathcal{N}$ consists of a sum of contributions to $\mathcal{N}$'s post-synaptic potential (PSP) from action potentials that are afferent to $\mathcal{N}$'s synapses. The amplitude of X_k is assumed to depend upon a paramenter of interest to the neurons that comprise $\mathcal{N}$'s efferent cohort. We model this parameter as a random variable, Θ, and we assume that Θ remains effectively constant over any two successive slots; i.e., Θ varies slowly compared to the slot duration.

We adopt a simplistic model in which $\mathcal{N}$'s overall PSP in slot k is of the form $Y_k = X_k + N_k$, where the N_k are IID zero mean noise values that are statistically independent of $\{X_k\}$. Among other things this model neglects the effect of partially decayed contributions to $\mathcal{N}$'s PSP from arrivals in past slots since the last slot in which $\mathcal{N}$ generated an action potential.

$\mathcal{N}$ produces a preliminary encoding of the value of Θ by generating an action potential in slot k if and only if Y_k exceeds a threshold, T_k, that is judiciously preset within $\mathcal{N}$. Since it requires considerably more energy to generate an action potential than not to, our energy-conscious neuron $\mathcal{N}$ does not set T_k at the median of Y_k even though that would maximze the information I gathered about Θ in the current slot. Rather, $\mathcal{N}$ sets its threshold considerably higher in order to reduce the expected energy expenditure E down to the point where I/E is the quantity that gets maximized.

---

[1] This work was supported in part by NIH Planning Grant RR15205

After studying the threshold value that maximizes I/E when, as in the preceding paragraph, attention is confined to a single slot, we concentrate primarily on the case in which $\mathcal{N}$ is used in each of two successive slots, say slot 1 and slot 2. $\mathcal{N}$ is able to change its threshold from T_1 in slot 1 to a value $T_2(Q_1)$ in slot 2, where $Q_1 = 1$ if $Y_1 > T_1$ and equals 0 otherwise. We consider both the case in which $\mathcal{N}$ is capable of spiking twice in succession and the case in which refractoriness prevents it from spiking in slot 2 if it has just spiked in slot 1; i.e., in the refractory case, $T_2(Q_1 = 1) = \infty$, so $\mathcal{N}$ is able to optimize only over the value of $T_2(Q_1 = 0)$. In both cases we seek the thresholding strategy that maximizes $I(\Theta; Q_1, Q_2)/(E_1 + E_2)$, the ratio of the expected number bits of information about Θ conveyed by $\mathcal{N}$'s outputs (spikes or non-spikes) in the combination of both slots to the expected number of joules expended by $\mathcal{N}$. Quantitative results are presented for the special case in which the N_k are uniformly distributed over an interval the length of which fixes the signal-te-noise ratio.