Chapter 1
Can We Predict Every Spike?

Book chapter to appear in *Spike Timing: Mechanisms and Function*, Edited by Patricia M. Dilorenzo and Jonathan D. Victor, CRC Press, November 2012

Richard Naud and Wulfram Gerstner

Is it possible to predict the spike times of a neuron with millisecond precision? In the classical picture of rate coding (Adrian, 1928), single spikes do not play a role, and the question would have to be answered negatively. For rate coding in a single-neuron, the relevant quantity to encode a stimulus such as pressure onto a touch sensor in the skin (Adrian, 1928) or presence of a light bar in the receptive field of a visual neuron (Hubel and Wiesel, 1959) is the number of spikes a neuron emits in a short time window of, e.g., 100ms. The timing of the spikes is considered as irrelevant. However, over the last 20 years many researchers have shown that it is not only the temporally averaged firing rate that carries information about the stimulus, but also the exact timing of spikes. For example, spike timing has shown to be relevant to encode force amplitude and direction in touch sensors of the skin (Johansson and Birznieks, 2004) as well as the whole-field visual movements (Bialek et al., 1991) or object movement (Gollisch and Meister, 2008) in visual neurons.

If spike timing is important, a whole series of questions arises: What is the precision of spike timing if the same stimulus is repeated several times? Do spikes always appear at the same time? What would be a sensible measure of spike timing precision and reliability? Can a neuron model match the spike timing precision of a real neuron? Does it matter which neuron or what stimulus we take? If so, what would be a useful stimulus?

To answer these related questions, let us think of the following experimental protocol. An experimentalist injects a time-dependent input of, say, 20 second duration into a single neuron. The neuron responds with spikes. The experimentalist now repeats the same stimulus sequence several times. At each repetition, the neuron responds with a spike train that may or may not look similar to the previous one: some spikes appear at exactly the same time during the stimulus sequence, some are missing, some are shifted by a few millisecond or appear at a completely different time. The information derived from this type of experiment which dates back to Bryant and Segundo (1976) and has been popularized by Mainen and Sejnowski (1995) should be sufficient to answer questions regarding precision and reliability of spike timing.

But here comes the challenge (Fig. 1.1). Let us suppose that I give you the time course of the input as well as the neuron's response in each trial, but only for the first 10 seconds of the data. For the second half of the stimulus sequence, I give you only the time course of the input. Your task is to predict the timing of the spikes of neuron.

Will you be able to predict the timing of the spikes using an appropriate neuron model? Is your model as reliable and as precise as the real neuron? What would be the best model to choose so as to solve the task?

The above challenge has been turned into a single-neuron modeling competition that was first run by Brain Mind Institute at the Ecole Polytechnique Fédérale de Lausanne (EPFL) in Switzerland (Jolivet, Schürmann, Berger, Naud, Gerstner and Roth, 2008; Jolivet, Kobayashi, Rauch, Naud, Shinomoto and Gerstner, 2008) and was officially handed over to the International Neuroinformatic Coordinating Facility (INCF) in Sweden in 2009. In this Chapter, we recapitulate the main questions and findings related to predicting spike times, with a special focus on the spike-time prediction competition of 2009 (Gerstner and Naud, 2009). If not specified otherwise, the term spike timing competition refers in the following to part A of the 2009 competition, if we specify a different year we imply part A of the competitions in 2008 or 2007.

1.1 What is a good stimulus to probe neurons?

In classical electrophysiological experiment, an artificially generated input is injected in a neuron *in vitro* (Fig. 1.1). In principle, the time course of the input can be chosen arbitrarily and could consist of short or long steps of different amplitudes, sequences of steps, ramps, white noise, filtered noise or whatever comes to mind. But what is a 'good' stimulus?

Since the work of Hodgkin and Huxley (Hodgkin and Huxley, 1952) electrophysiologists have been using steps and ramps to characterize single-neuron responses. These types of stimulations are helpful to systematically probe the gating dynamics of ion channels under pharmacological manipulation. They can also give a qualitative classification of neuronal responses in terms of intrinsic firing patterns such as regular, fast-spiking, bursting (Connors and Gutnick, 1990; Markram et al., 2004) but they have very little resemblance with the type of stimulus a neuron would receive in its natural environment.

Inspired by signal processing theory, the pioneering studies of Bryant and Segundo (1976) and of Marmarelis and Marmarelis (1978) used white-noise stimulation instead of step currents. However, if the aim is to drive a neuron with a stimulus that resembles as much as possible the input it would receive an *in vivo* situation, a white-noise stimulus is not sufficient. Rather, a stimulus at the soma should replace the total current flowing from the synapses to the soma while the neuron receives presynaptic input. Following a line of earlier research (Stein, 1967; Poliakov et al., 1996; Destexhe et al., 2003; Jolivet et al., 2006), the first spike timing competition in

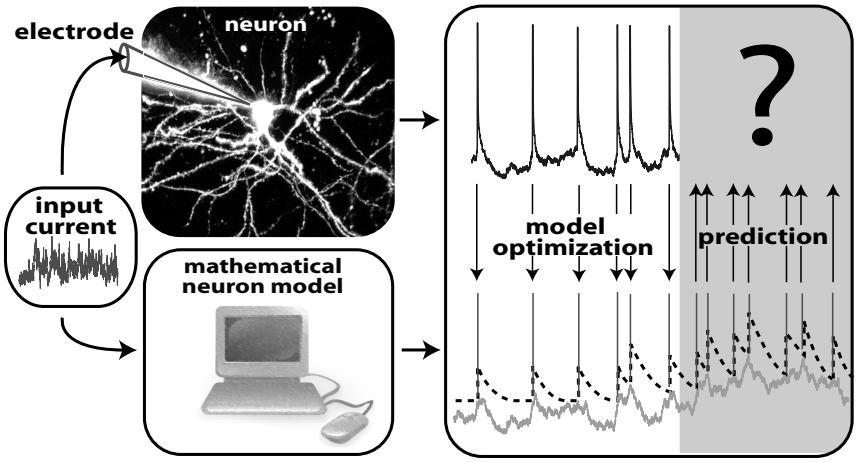

Fig. 1.1 Schematic representation of the spike-timing prediction challenge. The same time-dependent input stimulus (left) is given to a mathematical neuron model and to a real neuron in an electrophysiological experiment. Part of the response of the real neuron is used to optimize the model parameters. The remaining part of the stimulus is injected into the model so as to predict the spike-times of the real neuron. The mathematical neuron model illustrated here is made of a linear filter of the input current (bottom full trace) and a dynamic threshold (dashed black line). Fig. adapted from (Gerstner and Naud, 2009)

2007 used an Ornstein-Uhlenbeck current injection with various means and variance to mimic the combined effect of a large number of synapses (Stein, 1967; Jolivet, Kobayashi, Rauch, Naud, Shinomoto and Gerstner, 2008). In 2008, the competition was modified(Jolivet, Kobayashi, Rauch, Naud, Shinomoto and Gerstner, 2008) to replace the dynamic current by dynamic inhibitory and excitatory conductances using dynamic clampDestexhe et al. (2003). Then in 2009 the injected current was changed to a current produced by the simulation of six populations of presynaptic neurons changing their firing rate every 200-500 ms.

1.2 How can we measure spike timing precision and reliability?

Suppose that a single neuron is driven with multiple repetitions of the same time-dependent stimulus. The response of the neuron is recorded in each trial, so that the stimulation protocol builds up a database containing one spike train for each repetition.

If we compare the spike trains across several repetitions, different types of variability are seen depending on the system and the variance of the input (Bryant and Segundo, 1976; Mainen et al., 1995; Jolivet et al., 2006). Some spikes are seen at the

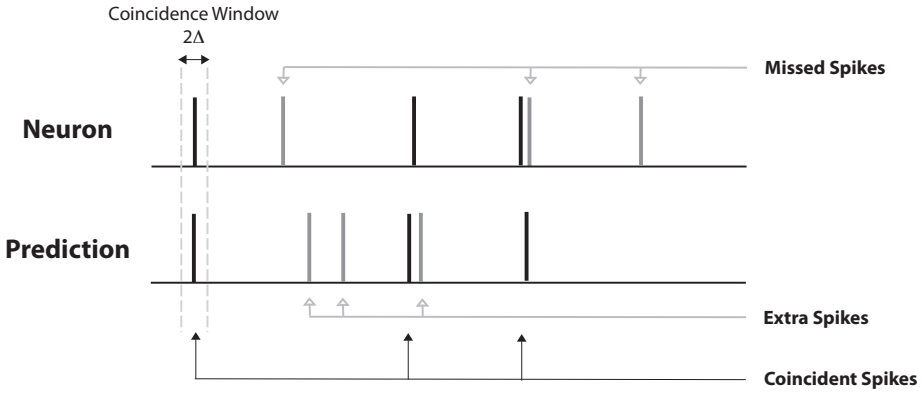

Fig. 1.2 Counting coincident spikes for the computation of the coincidence rate Γ_{nm}. The predicted spike train (bottom) is compared to the recorded spike train (top). A predicted spike is said to be coincident (black) if it falls between $\pm\Delta$ of a recorded spike and if that recorded spike was not counted as coincident with any other predicted spike. The prediction can miss recorded spikes (gray, top) or countain generate extra spikes (bottom, gray). Here the total number of coincident spikes is $N_{\text{coinc}} = 3$ while there were $N_m = 6$ spikes in the predicted and $N_n = 6$ spikes in the experimental spike train. (Adapted from Figure 1.1 in Jolivet (2005)).

same time for all repetitions, others appear at a specific time on half the repetitions and yet others do not seem to be related to a specific time. Several questions arise: First, how can we *quantitatively* compare one spike train of the neuron with another one recorded during a later repetition of the same stimulus sequence? Second, how can we quantify the reliability across the set of all spike trains recorded with the same stimulus? Finally, how can we compare the set of spike trains generated from neuronal recordings with a similar set of spike trains generated by a mathematical neuron model?

These are crucial questions which can be answered in different ways. One can compare one spike train with another one based on global features such as the intrinsic firing patterns in response to step stimuli (Connors and Gutnick, 1990; Markram et al., 2004) and ask whether a neuron model is able to reproduce the same intrinsic firing patterns (Izhikevich, 2007; Naud et al., 2008). One can focus the quantitative comparison on the shape of spikes and adaptation patterns (Druckmann et al., 2007), on the inter-spike interval distribution (Chacron et al., 2005) or the spike-count variability (Softky and Koch, 1993; Schaette et al., 2005).

If one focuses on spike timing, one may want to apply methods that compare spike trains in terms of a spike-train metrics (Victor and Purpura, 1996) or the co-

incidence rate (Kistler et al., 1997). Both measures can be used to compare a spike train from a recorded neuron in repetition n with another spike train recorded in repetition m. Both measures can also be used to compare a spike train derived from a neuron model with a spike train recorded in one of the sessions with the real neuron. Obviously, a model which achieves an optimal match in terms of spike-train metrics will automatically account for global features of the spike trains, such as inter-spike interval distributions.

In the INCF competition the average coincident rate was used to quantify spike-time prediction performance. The average coincidence rate can be seen as a similarity measure between pairs of spike trains that is finally averaged across all available pairs. To compute the pairwise coincidence rate, one first finds the number of spikes from the model that fall within an interval of plus or minus 4 ms around a spike from the real neuron. This is called the number of coincident events N_{nm} (at resolution $\Delta = 4ms$). The coincidence rate is the ratio of the number of coincident events over the averaged number of events $0.5(N_n + N_m)$, where N_n is the number of spikes in the neuron spike train and N_m is the number of spike in the model spike train (Fig. 1.2). This ratio is then scaled by the number of coincident events, $N_{\text{poisson}} = 2\Delta \cdot N_m N_n / T$, that are expected from a Poisson model that fires stochastically at a fixed rate N_m/T. The scaled coincidence rate is

$$\Gamma_{nm} = \frac{N_{\text{nm}} - N_{\text{poisson}}}{\frac{1}{2}(1 - N_{\text{poisson}}/N_n)(N_n + N_m)}. \tag{1.1}$$

Finally, the pairwise coincidence rate Γ_{nm} is then averaged across all the possible pairings of spike trains from the model with those of the neuron and this gives the averaged coincindence rate $\overline{\Gamma_{nm}}$.

The coincidence rate Γ_{nm} enables us to compare the spike timings of the model (subscript m) with that of the neuron (subscript n). If we want to know how reliable the neuron itself is, we need to measure how similar the spike trains are between two trials. To do the comparison between one neuronal spike train (index n) to another neuronal spike train (index n), we can use again the same coincidence measure that we now label Γ_{nn} so as to indicate the coincidence rate between two neuronal spike trains. The neuron-to-neuron coincidence rate, averaged over all pairs of available spike trains for the same stimulus, is a measure of the intrinsic reliability of the neuron and denoted as $\overline{\Gamma_{nn}}$. It provides an upper bound for modelling: on average, a neuron model cannot predict spikes better than the neuron itself. The averaged coincidence rate $\overline{\Gamma_{nm}}$ needs therefore to be compared to the upper bound provided by the intrinsic reliability $\overline{\Gamma_{nn}}$ of the neuron. Scaling $\overline{\Gamma_{nm}}$ by $\overline{\Gamma_{nn}}$ gives a number that can be interpreted as the fraction of the predictable spikes that are predicted by the model.

The coincidence rate, like the majority of spike-train metrics, has a time-scale parameter. In the above equation, the time-scale parameter Δ regulates the size of the coincidence window, and thus the level of precision of the prediction. For a very small coincidence window, the number of coincidences N_{coinc} goes to zero due to a jitter in spike timing and the finite number of spike trains. For a very large coinci-

dence window, the coincidence rate goes to zero because it becomes insensitive to specific times of the spikes so that the difference between the prediction of a precisely tuned neuron model and that of a Poisson model with constant rate vanishes. In a large range between these two extrema, however, the model-to-neuron coincidence rate is significantly positive. Moreover, over the range roughly from 2 to 15 ms depending on the neuron and on experimental conditions(Jolivet et al., 2006; Jolivet, Schürmann, Berger, Naud, Gerstner and Roth, 2008), the results measured in terms of $\overline{\Gamma_{nm}}$ do not depend on the choice of the time window Δ. In the INCF competition a window of $\Delta = 4ms$ was chosen.

It is useful to distinguish measures such as the coincidence rate (Kistler et al., 1997) or the spike train metrics (Victor and Purpura, 1996) which are both based on a comparison of a single spike train A with a second spike train B from measures that first average across all repetitions of an experiment to calculate the Peri-Stimulus Time-Histogram (PSTH) (Eggermont et al., 1983) before a comparison of the PSTH of neuron A with that of a neuron B (or of a model neuron) is performed. In the INCF challenge, rankings were based on the pairwise comparison of a model spike train with a real spike train, averaged *a posteriori* across all repetitions of the experiment so as to determine the average coincidence rate.

Does the average coincidence rate correspond to a comparison of the PSTH? The PSTH is calculated by averaging all the independent responses to the same input. A smoothing filter is then applied to the averaged spike trains. In many neuronal systems, the PSTH is made of a series of peaks and plateaus. The peaks correspond to spikes always coming at a precise time and the plateaus corresponds to times where spikes are emitted with no specific timing. A model reproducing such a PSTH can be said to predict the spike times because such a model will emit a spike precisely at times where the neuron emits precisely timed spikes. Indeed, normalized spike-train similarity measures such as $\overline{\Gamma_{nm}}/\overline{\Gamma_{nn}}$ calculate a quantity very similar to the variance of the experimental PSTH that is explained by the model PSTH (Naud et al., 2012). The time-scale parameter of the similarity measure is equivalent to the filter time-scale applied for smoothing the PSTH. The main difference between $\overline{\Gamma_{nm}}/\overline{\Gamma_{nn}}$ and the comparison of PSTHs is that an optimization of neuron models based on the comparison of PSTH attempts to match the spike-timing variability of the model to the variability of the data. In contrast, an optimization of neuron models based on the normalized coincidence rate gives a slight advantage to deterministic neuron models, i.e., those that do not correctly reproduce the intrinsic variability of the data (Naud et al., 2012).

Being aware of the similarity between PSTH comparison and coincidence rate scaled with intrinsic reliability is important to relate different studies to each other even though they use different evaluations criteria. It should be kept in mind, that there is an over-estimation of the prediction performance for deterministic models using the scaled coincidence rate $\overline{\Gamma_{nm}}/\overline{\Gamma_{nn}}$ with respect to variance-explained of the PSTH (Naud et al., 2012). However, there is also a small-sample bias that can overestimate the variance explained when comparing PSTHs (David and Gallant, 2005; Petersen et al., 2008; Naud et al., 2012). Values of above 100 % are possible because

of the bias, but also because a model can be fitted independently on each repetition and thus allowing to take into account the experimental drifts.

1.3 What are good neuron models?

Across several years and editions of the single-neuron modelling competition, various models participated in the challenge. The models ranged from the very simple integrate-and-fire models to complete biophysical models using the Hodgkin-and-Huxley formalism and a 3D reconstruction of the morphology from another neuron of the same class. The number of state-variables goes from one for the simplests models to a few hundred depending on the number of ion-channel types modelled and the number of compartments used in the discretization of space. Similarly, the number of parameters scaled with the number of state-variables with five parameters for the simplest models and close to a hundred parameters for the biophysical models. Between the two extremes, various other models were used. For instance, adding non-linear threshold, spike-triggered adaptation and sub-threshold adaptation, the Izhikevich model (Izhikevich, 2004) and the AdEx (Brette and Gerstner, 2005) models were used by some of the participants. Simpler versions of the complete biophysical models were also devised by reducing the number of ion-channels and the number of compartments (in the same line as Pospischil et al. (2008)).

A good prediction stems from the union of a good model and an efficient fitting method. The method for finding the optimal parameters should be efficient in the sense of providing a single set of optimal parameters with small computing resources. Many different methods were used by the various participants, some using the action potential shape, the subthreshold voltage dynamics and the spike times as the observables to fit, others using only the set of spike times as observables to fit.

The participants that used biophysical models approached the problem by first constraining a significant number of parameters with published measurements of ion channel dynamics. Most of the biophysical-model participations thus reduced the parameter space to the somatic ion-channel densities only. This leaves a number of free parameters equal to the number of ion-channel species (Druckmann et al., 2007). These remaining parameters are then fit either by hand-tuning, stochastic optimization algorithms such as the genetic algorithm or by exhaustive search when the number of ion-channels species is low.

For many of the simpler models, the optimization methods available are more efficient. Some of the most effective participations performed an exhaustive search on a small number of crucial parameters(Kobayashi et al., 2009). Convex optimization algorithms can be used to maximize the likelihood of observing the spike times (Paninski et al., 2004) this method lead to some of the top-ranking participations. Another noteworthy method for fitting involves a convex, two-step procedure where in the first step the optimal passive parameters and the spike shape are determined from the voltage trace, and on the second step the parameters regulating a dynamic

threshold are determined by maximizing the likelihood of the observed spike trains (Mensi et al., Under Review).

The model that achieved the highest performance in Challenge A 2009 was an integrate-and-fire model provided with a dynamic threshold that jumps every times there is a spike and decays back to a baseline with three different time-constants(Kobayashi et al., 2009). The ratio $\overline{\Gamma}_{nm}/\overline{\Gamma}_{nn}$ for this participation was of 76.2%. The participant winning this spike-time prediction competition had extracted the membrane time constant from the voltage trace, fixed the decay time-constants of the dynamic threshold to 10, 50 and 200 ms leaving three free parameters; one parameter regulating the amount by which the threshold jumps for each of the three time scales. The optimal set of the three parameters was found by conducting an exhaustive search for the set of parameters maximizing the average coincidence. This winning model happened to be the model with the smallest number of free parameters used in the competition. The small number of free parameters is not sufficient to explain the high performance, because for instance a simple leaky and integrate and fire cannot predict more than 38% on the same task. The winning participant used judicious insights in deciding which model to use, which parameters could be fixed *a priori* and which parameters required to be fitted to the specific neuron recorded.

A very small number of biophysical models participated in the competition, perhaps due to the difficulty of finding the optimal parameters for such complex models. There was a noteworthy participation using state-of-the-art optimization methods which was outside of the official competition because submitted after the deadline for the money prize. This submission would have ranked third if it had been submitted before the deadline. Thus, within the framework of the competition, the prediction performance of simple models is as good if not slightly better than that of the biophysical models.

The relatively high prediction performance of models is not explained by the fact that the challenge is too easy. Off-the-shelf models such as a hand-tuned models of pyramidal neurons or a leaky-integrate-and-fire achieved a performance around $\overline{\Gamma}_{nm}/\overline{\Gamma}_{nn} = 40$ %. Perhaps the most important aspect for providing a good prediction in the competition was to take into account spike-frequency adaptation. The dynamic threshold of the winning submission is just one example, there are many other ways to implemented adaptation in single neurons and models that implement adaptation were systematically better than models that did not (Jolivet, Kobayashi, Rauch, Naud, Shinomoto and Gerstner, 2008).

So what is the best model? Accurate modelling of the refractory period is essential to predict the spike times (Kistler et al., 1997; Keat et al., 2001). More generally, if the same neuron model has to predict spike timings for stimuli with different mean firing rates, the importance of spike-frequency adaptation was recognized explicitly (Pillow et al., 2005; Jolivet et al., 2006; Jolivet, Kobayashi, Rauch, Naud, Shinomoto and Gerstner, 2008). State-of-the-art models now consist of models akin to the stochastic integrate-and-fire model but upgraded with an adaptation process. The adaptation makes the firing probability dependent on the timing of all recently emitted spikes. Such models are capable of predicting 75-100% of the predictable spikes (see next Section). This leaves little room for improving the accuracy of encoding

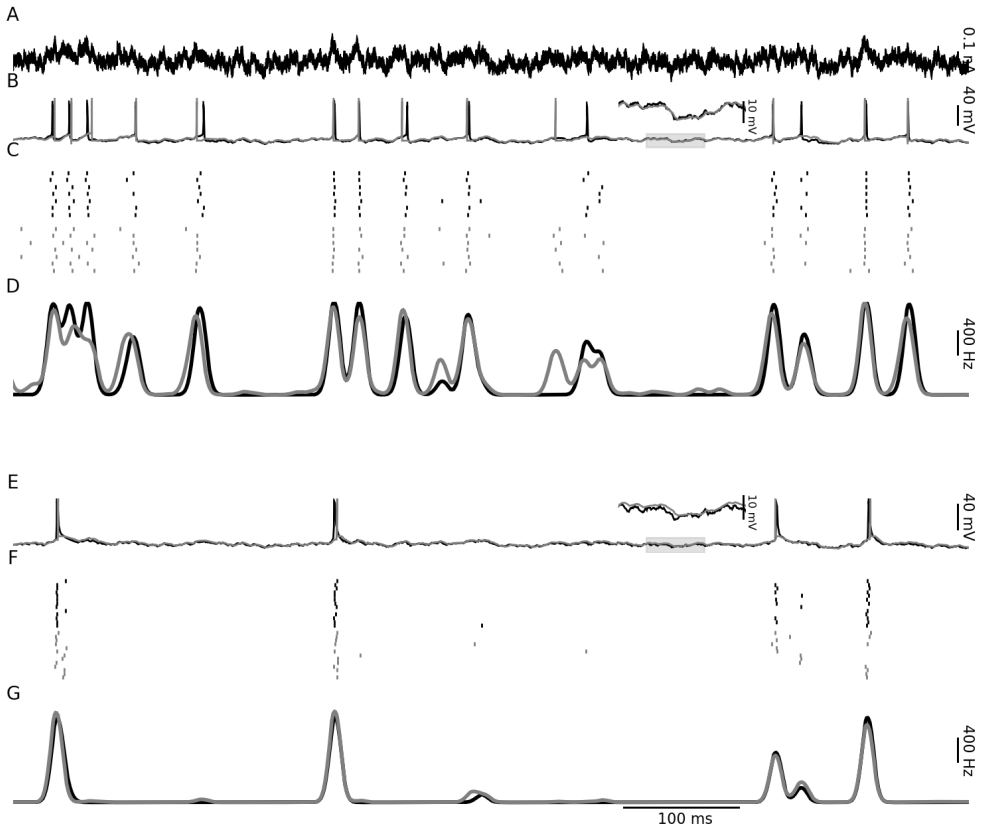

Fig. 1.3 Predicting fast-spiking GABAergic neurons and pyramidal neurons from in 5 from the 2009 competition. A: The time-dependent current input models six populations of presynaptic neurons changing their firing rate every 200-500 ms. B-D Prediction of a GABAergic fast-spiking neuron. B: modelled (gray) and recorded (black) voltage traces. A zoom of 50 ms is shown in inset. C spike trains of each 7 repetitions of the recorded (black) compared modelled (gray) spike train. D PSTH of the model (gray) overlaid on the PSTH of the data (black) calculated from the average of the spike trains that are then filtered with a gaussian of 5 ms standard deviation. For the model we computed the PSTH from one thousand independent realizations while for the data we were restricted to the number of repetitions that could be recorded in the experiment. E-G same as for B-D but predicting the activity of a pyramidal neuron from the layer 5. (Figure a courtesy of Skander Mensi).

models. Indeed, increasing the level of detail with conductance-based adaptation or Hodgkin-Huxley ion channels does not yield substantial increases in prediction performance (Druckmann et al., 2007; Mensi et al., Under Review).

1.4 Are all neurons predictable?

Can we predict the spike times in other systems than a Layer 5 pyramidal neuron in a cortical slice? The different editions of the competition showed that good spike-time prediction can be achieved for current injections of at least two types of dynamics ($\overline{\Gamma}_{nm}/\overline{\Gamma}_{nn} = 82.0$ % in 2007 for Ornstein-Uhlenbeck dynamics, $\overline{\Gamma}_{nm}/\overline{\Gamma}_{nn} =$ 71.6 % in 2009 for multiple time-scale dynamics). The predictions are slightly better when the stimulation is given as fluctuating inhibitory and excitatory conductance ($\overline{\Gamma}_{nm}/\overline{\Gamma}_{nn} = 91.4$ % for challenge A in 2008). The activity of L2/3 pyramidal neurons and non-Fast spiking GABA-ergic neurons can be predicted with similar performances (Mensi et al., Under Review). Prediction of the spike times of fast-spiking GABA-ergic neurons is systematically higher with $\overline{\Gamma}_{nm}/\overline{\Gamma}_{nn} =$ in the range of 100 percent in challenge B of 2009 (Fig. 1.3).

Real neurons receive their inputs from synapses distributed throughout their dendritic tree. The single-neuron model should model the dendritic integration of inputs. This dendritic integration is known to be highly non-linear especially in the thick tufted L5 pyramidal cells (Larkum et al., 1999). To explore the dendritic dimension, dual electrode recordings were made with two independent injection sites: one in the soma and a second high in the dendritic tree. Spike-time prediction of $\overline{\Gamma}_{nm}/\overline{\Gamma}_{nn} = 83.8$ % was achieved in Part C of the 2009 competition with a model similar to the one used by Larkum et al. (2004) .

In the retina, spiking models of the Retinal Ganglion Cells (RGC) can predict 91 % of the variance of the PSTH (Pillow et al., 2005). Similar performances have been observed in vivo where 41-92% of the variance of the PSTH Lateral Geniculate Nucleus (LGN) cells can be predicted from the activity of a single impinging RGC for spatially restricted visual stimulation(Carandini et al., 2007) as confirmed by submissions to Challenge D in the 2009 edition of the competition.

1.5 Conclusion

In summary, the prediction of precise spike timing on the millisecond time scale is similar to predicting the time-dependent firing rate on the mili-second time scale. High prediction performance is possible in many neuronal systems and depends strongly on the choice of neuron model and fitting method. One important model feature for high prediction performance is the presence of spike-frequency adaptation. The choice of the model formalism can also influence the fitting method that

can be used. High quality prediction is most of the time associated with an efficient and convex fitting method.

Can we use these results to determine the best single-neuron model? The single-neuron model of choice should be able to generalize across all the different experimental protocols and also across the possible systems and neuron types with a mere change of the model's parameters. The original competition was rewarding only the participations that could generalize across more than one of the experimental protocols. The data of the challenge will remain available in the future for bench-marking purposes, leaving the possibility for such a deed to be accomplished[1].

References

Adrian E D 1928 *The basis of sensation* W.W. Norton, New York.

Bialek W, Rieke F, de Ruyter Van Steveninck R and Warland D 1991 *Science* **252**(5014), 1854–1857.

Brette R and Gerstner W 2005 *Journal of Neurophysiology* **94**, 3637–3642.

Bryant H L and Segundo J P 1976 *Journal of Physiology* **260**(2), 279–314.

Carandini M, Horton J C and Sincich L C 2007 *Journal of Vision* **7**(14), 20.1–11.

Chacron M, Lindner B, Maler L, Longtin A and Bastian J 2005 *Proceedings of the SPIE* **5841**, 150–163.

Connors B W and Gutnick M J 1990 *Trends in Neuroscience* **13**, 99–104.

David S V and Gallant J L 2005 *Network* **16**(2-3), 239–60.

Destexhe A, Rudolph M and Pare D 2003 *Nature Reviews Neuroscience* **4**, 739–751.

Druckmann S, Banitt Y, Gidon A and Schürmann F 2007 *Frontiers in Neuroscience* **1**(1), 7–18.

Eggermont J J, Aertsen A M and Johannesma P I 1983 *Hearing Research* **10**(2), 167–90.

Gerstner W and Naud R 2009 *Science* **326**, 379–380.

Gollisch T and Meister M 2008 *Science* **319**, 1108–1111.

Hodgkin A L and Huxley A F 1952 *Journal of Physiology* **117**(4), 500–544.

Hubel D H and Wiesel T N 1959 *Journal of Phsyiology* **148**, 574–591.

Izhikevich E 2004 *IEEE transactions on neural networks* **15**(5), 1063–1070.

Izhikevich E M 2007 *Dynamical systems in neuroscience : the geometry of excitability and bursting* MIT Press Cambridge, Mass.

Johansson R S and Birznieks I 2004 *Nature Neuroscience* **7**(2), 170–177.

Jolivet R 2005 Effective Minimal Threshold Models of Neuronal Activity PhD thesis Ecole Polytechnique Federale de Lausanne.

Jolivet R, Kobayashi R, Rauch A, Naud R, Shinomoto S and Gerstner W 2008 *Journal of Neuroscience Methods* **169**, 417–424.

Jolivet R, Rauch A, Lüscher H and Gerstner W 2006 *Journal of Computational Neuroscience* **21**, 35–49.

[1] http://www.incf.org/

Jolivet R, Schürmann F, Berger T, Naud R, Gerstner W and Roth A 2008 *Biological Cybernetics* **99**(4), 417–426.

Keat J, Reinagel P, Reid R C and Meister M 2001 *Neuron* **30**, 803–817.

Kistler W, Gerstner W and Hemmen J 1997 *Neural Computation* **9**, 1015–1045.

Kobayashi R, Tsubo Y and Shinomoto S 2009 *Frontiers in Computational Neuroscience* **3**, 9.

Larkum M E, Senn W and Luscher H R 2004 *Cerebral Cortex* **14**(10), 1059–1070.

Larkum M, Zhu J and Sakmann B 1999 *Nature* **398**, 338–341.

Mainen Z F, Joerges J, Huguenard J R and Sejnowski T J 1995 *Neuron* **15**(6), 1427–1439.

Mainen Z F and Sejnowski T J 1995 *Science* **268**(5216), 1503–1506.

Markram H, Toledo-Rodriguez M, Wang Y, Gupta A, Silberberg G and Wu C 2004 *Nature Reviews Neuroscience* **5**(10), 793–807.

Marmarelis P Z and Marmarelis V Z 1978 *Analysis of Physiological Systems: The White-Noise Approach* New York: Plenum Press.

Mensi S, Naud R, Avermann M, Petersen C and Gerstner W Under Review *Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms* .

Naud R, Gerhard F, Mensi S and Gerstner W 2012 *Neural Computation* .

Naud R, Marcille N, Clopath C and Gerstner W 2008 *Biological Cybernetics* **99**, 335–347.

Paninski L, Pillow J and Simoncelli E 2004 *Neural computation* **16**, 2533–2561.

Petersen R S, Brambilla M, Bale M R, Alenda A, Panzeri S, Montemurro M A and Maravall M 2008 *Neuron* **60**(5), 890–903.

Pillow J, Paninski L, Uzzell V, Simoncelli E and Chichilnisky E 2005 *Journal of Neuroscience* **25**(47), 11003.

Poliakov A V, Powers R K, Sawczuk A and Binder M C 1996 *Journal of Physiology* **495**, 143–157.

Pospischil M, Toledo-Rodriguez M, Monier C, Bal T, Frégnac Y, Markram H and Destexhe A 2008 *Biological Cybernetics* .

Schaette R, Gollisch T and Herz A V M 2005 *Journal of Neurophysiology* **93**(6), 3270–81.

Softky W and Koch C 1993 *Journal of Neuroscience* **13**, 334–350.

Stein R B 1967 *Proceedings of the Royal Society of London. Series B, Biological Sciences* **167**(1006), 64–86.

Victor J D and Purpura K 1996 *Journal of Neurophysiology* **76**(2), 1310–1326.