ARTICLE
Neurosim: Some Thoughts on Using Computer Simulation in Teaching Electrophysiology

Bill Heitler
School of Psychology and Neuroscience, University of St Andrews, Fife KY16 9JP, United Kingdom.

Neurosim is an interactive simulation program designed for teaching electrophysiology. It was first published in 1989, but has been updated several times over the years, and v5.3.3 was released in January 2022.

Much effort has been put into making Neurosim as easy to use as possible, while at the same time offering a wide range of facilities. It contains 7 modules that simulate at biological levels ranging from single channel membrane properties, through spike and synaptic properties, small network properties, up to whole-population firing dynamics. It is highly configurable and can be useful for teaching from the beginning undergraduate level dealing with basic neuron physiology, through to the post-graduate level suitable for use as an introduction to computational neuroscience.

The article describes how Neurosim has been useful in my own teaching over the years and gives several examples of student activities that have proved effective in aiding understanding. There is a comprehensive set of tutorial exercises available on the support website.

Key words: electrophysiology, neurophysiology; simulation; action potential; voltage clamp; synapse; network; teaching neuroscience

---

This article is based on a presentation given at the Summer Virtual Meeting (2020) of the Faculty for Undergraduate Neuroscience.

For most of my career my duties included team-teaching core undergraduate neuroscience courses, within which I usually had particular responsibility for cellular electrophysiology. One early problem that I encountered was that this involved teaching some *physics* (in the form of electricity) as well as chemistry and biology. Most of my students were biology majors, and while many of them could vaguely remember hearing about voltage and current and Ohm's Law in secondary school, only a few could remember what the terms actually meant, and even fewer had internalized any understanding of the concepts involved. This meant that they initially struggled with topics such as the Nernst equation and equilibrium potentials, driving force, the voltage clamp technique, and many other core aspects of electrophysiology. (This is not a criticism of my students, they were simply victims of the universal law of knowledge – use it or lose it, and up to that point most of them had had no reason to use it.)

It also became quite rapidly apparent that the standard lecture/lab teaching methodology was not a very good way of developing deep understanding of these topics. Students could learn the facts, and even some of the concepts, and repeat them under exam conditions, but were often at a loss if presented with a problem that required actually *using* the knowledge in a novel context.

By the mid-1980s, fairly early in my career, personal computers were becoming affordable and quite powerful, and computer-based models of the nervous system were moving out of the restricted purview of institution-level mainframe computers, and into the laboratories of individual researchers (Macgregor, 1987). Early versions of powerful modelling tools (e.g., Genesis, 2019; Neuron, 2021) were becoming available, but these were very much research orientated, and not really suitable for casual use by students, or, indeed, staff. I was becoming increasingly convinced of the pedagogic benefits of simulation as a path to understanding mechanisms in complex interacting systems (which certainly include the nervous system), and in 1988 I wrote a series of computer simulations for the IBM PC which I called Neurosim. These were specifically designed as an aid for teaching electrophysiology and were intended for use by people with little knowledge of, or interest in, computing itself. In those days the internet was in its infancy, and to make these available to a wider audience (and also to gain publication credit from my university), I published them through a commercial publisher. An early version of the program was validated as a useful teaching aid in an independent study (McAteer, et al., 1996), but it has undergone many modifications since then. The latest version (5.3.3) was published in January of this year. I have now formally retired from my university appointment, although I retain an honorary position, and am publishing Neurosim myself rather than through the external publisher.

A full description of Neurosim is available on its support website (https://www.st-andrews.ac.uk/~wjh/neurosim), but in this article I will briefly describe some of its teaching-orientated features, and then illustrate some of the ways I used it in my own teaching. I will also mention some general thoughts regarding the pedagogic role of simulation which I have developed over the years of using Neurosim.

NEUROSIM
Neurosim follows the standard protocol of most simulations: the user first sets up the experimental conditions, and then runs an experiment to see what happens. It contains seven different modules, each simulating events at a different biological level:
1. Goldman: The Nernst and Goldman equations are simulated, allowing exploration of factors underlying the

membrane potential.

2. Membrane patch: This simulates the kinetics of single ion channels.

3. Passive conduction: This implements the cable equation (Jack *et al*., 1975) to simulate the properties of a uniform length of passive dendrite.

4. Hodgkin-Huxley: A direct implementation of the HH model of action potential generation in the squid giant axon (Hodgkin and Huxley, 1952), allowing both current clamp and voltage clamp experiments.

5. Advanced HH: This is like the standard HH simulation, but it allows inclusion of voltage-dependent channels with user-defined properties. This enables replication of many published models of neurons containing "non-standard" channels such as hyperpolarization-activated sodium channels, calcium-dependent potassium channels, and many more. The resulting neurons can demonstrate properties such as endogenous bursting,

plateau potentials etc.

6. Network: This allows construction of neural circuits containing an arbitrary number of identified neurons connected by chemical or electrical synapses. The neurons can have integrate-and-fire or HH-like spike mechanisms (the latter can be imported from the Advanced HH module). The synapses can have a wide range of properties, including voltage-dependency or Hebbian mechanisms. The Network module can also implement simple compartmental models by connecting a chain of neurons with powerful electrical synapses, thus making each one the equivalent of a compartment within a single neuron.

7. Wilson-Cowan. This implements a firing rate model which can contain an arbitrary number of connected populations of excitatory and inhibitory neurons, each using the W-C formalism (Wilson and Cowan, 1972).

All simulations have a standard layout involving two main
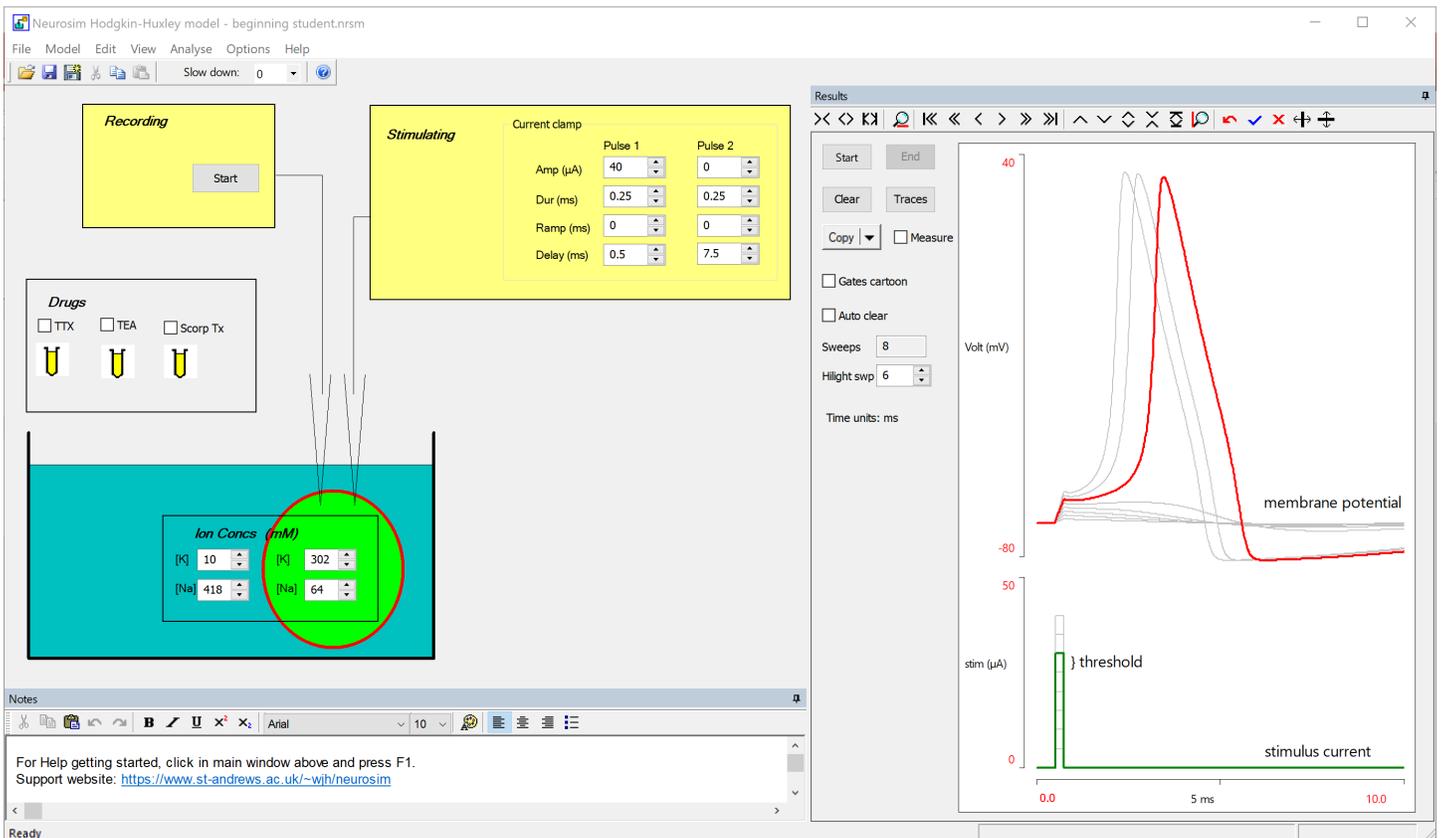
Figure 1. The HH module with a simplified configuration suitable for a beginning student. The Setup view (left) shows a nerve cell (the green circle) in a preparation bath (the blue rectangle) with various drugs poised for application. Electrodes in the nerve cell allow stimulation and recording. The Results view (right) illustrates a key learning objective for a first-year student: the all-or-none nature of the action potential. Successive experiments were run with increasing stimulus strength, with the results superimposed. The sweep showing the first above-threshold response has been highlighted. Free-form trace annotations were added within Neurosim.

This configuration would also allow the student to explore several other suitable first-year topics, including the stimulus strength-duration relationship (using the amplitude and duration settings), the refractory period (using the second stimulus option), the sodium-dependence of the action potential (by changing the external sodium concentration in the bath), the potentially lethal consequences of hyperkalemia (by increasing the external potassium concentration) and the effects of the various drugs (by clicking the test tubes to apply them). Amongst the latter, TTX is always popular since it allows discussion of fugu poisoning, zombies and numerous murder mysteries!

views (Figure 1).   The Setup View shows a visual representation of the experimental conditions in diagram. form, with key parameters exposed for easy adjustment The dockable Results View shows the outcome of the experiment in a format that resembles a real experiment. Of course, one of the key advantages of simulation is that it allows the user to look "under the hood" at the values of underlying variables such as channel conductance that cannot be directly measured in a real experiment, but both the Setup and Results views can be configured by a tutor to hide or expose features as appropriate for a particular pedagogic aim. So, the visible results can be restricted to those that could be obtained in a real experiment (as in Figure 1), or elaborated to show hidden variables (e.g., Figure 5).   The configuration settings can be stored in a parameter file, so a student just has to load the appropriate file to immediately obtain an experimental setup with only the desired features visible, and with all the parameters pre-set to suitable starting conditions. The tutor can password-protect the hidden parameters, so that the student can be set the task of discovering their values by experiment.

The overall design aims of Neurosim are thus twofold: usability and flexibility. The ability to configure and simplify the user interface is key to usability – a student is only exposed to the information that is necessary for the task in hand. There are also many GUI "tweaks" that have been included as a result of observing students actually using the program.   For instance, most parameter values can be adjusted by a spin button paired with the parameter edit box, and the delta value of the spin button (the size of the step change caused by a click) can be adjusted as needed. This is useful because a parameter in one simulation might need changing in steps of, for instance, 0.1 units, while the exact same parameter might need changing in steps of 50 units in another simulation. A fixed delta value would make the spin button useless in one or the other situation. Another tweak is inclusion of a "run-on-change" option, which runs a simulation immediately when the user changes a parameter value (normally the user has to click a *Start* button to actually run the simulation). This means that an entire experimental protocol such as that in Figure 1, can be run simply by repeatedly clicking one spin button. Such a protocol results in a series of overlaid sweeps on the screen (unless *Auto clear* is selected), and once it is completed each sweep can be highlighted in turn, which helps students keep track of which sweep goes with which condition.

The flexibility aim has two aspects. Firstly, simulation can be useful in teaching electrophysiology at a number of *biological* levels, from single channels within the plasma membrane, through to the firing rate dynamics of populations of neurons within the whole brain. This aim is met by the inclusion of the 7 different modules described above. In theory, one could probably have a single "mega simulation" engine that covered all levels, but the resulting complexity would be very likely to conflict with the usability aim. So an early design decision was to have separate simulation options within the single program, but with as much commonality in look-and-feel as possible. Secondly, I wanted Neurosim to be able to simulate experiments

suitable for different *pedagogic* levels of instruction, from simple experiments relevant to a beginning student in the first year of study, through to sophisticated models suitable for a post-graduate student wishing to improve their understanding of computational neuroscience. Again, there is a risk of this conflicting with the usability aim – "feature bloat" is notorious for making a user unable to see the wood for the trees and spoiling a once useful program. Hopefully, the Neurosim configuration options obviate this problem – the user need only be exposed to the level of detail required for the activity.

TUTORIAL LESSONS

The open-access Neurosim support website contains an extensive set of free, ready-made tutorial exercises that can be used "off the shelf" or modified to meet a particular learning objective. The tutorials are organized into five main sections - passive properties, action potentials, synapses, networks and single-channel kinetics, but each has numerous sub-sections. There is far more tutorial material than would be needed or useful for most undergraduate courses (there is a limit to the amount of time that even the
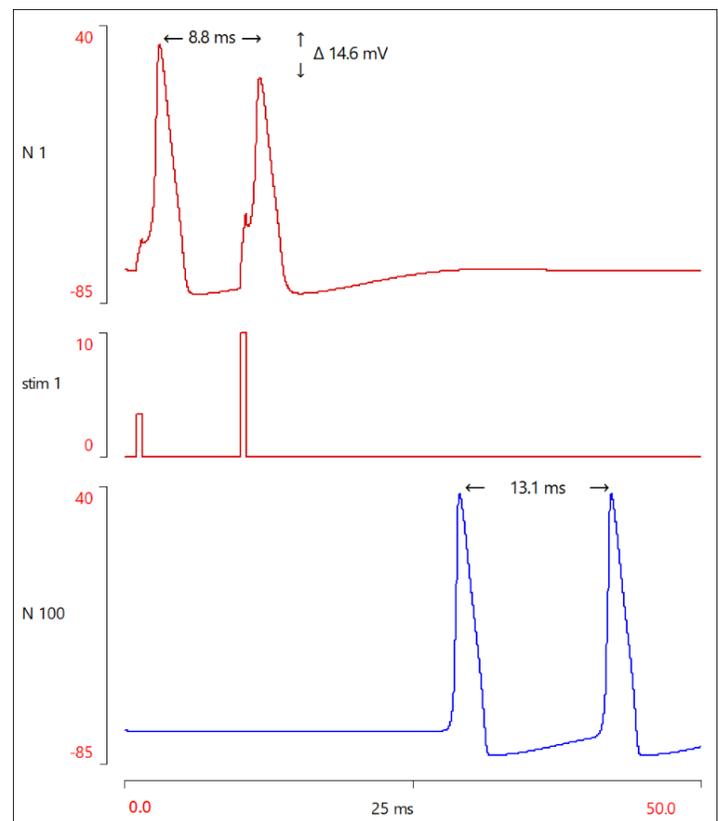
Figure 2. Refractory conduction in a 100-unit compartmental model of the squid giant axon. A spike elicited during the relative refractory period of a preceding spike has a higher stimulus threshold and a reduced peak amplitude (red traces, spike initiation site). However, by the time the spikes have propagated to the far end of the axon, the second spike, which initially conducts at a lower velocity, has dropped out of the refractory period and recovered in amplitude (blue trace, output terminal). The results are displayed in a standard voltage-vs-time format.

most dedicated student wishes to spend running computer simulations!), but hopefully most teachers will find something useful within the exercises.   The interested reader is invited to visit the website Documents tab to see a full contents listing, and/or to actually view the tutorials directly.

EXAMPLES OF USE
Lecture Illustration
Neurosim can be used as a visual aid during a standard lecture to illustrate dynamical events that are not easily conveyed with static images. An example of this concerns the propagation of an action potential elicited during the relative refractory period of a preceding spike. The standard HH model of an isopotential patch of membrane (or a space-clamped axon) clearly shows that, as expected, the second spike has a higher stimulus threshold than the preceding spike. However, it also has a reduced amplitude. This can be easily demonstrated by simulation in the HH module, and the finding can prompt useful class discussion regarding its underlying cause. However, the job of a real axon is to conduct spikes along its length, so what happens to the reduced-amplitude spike during propagation?

Spike conduction can be simulated in Neurosim by building a compartmental model using the Network module (Figure 2). This shows that the second spike, which initially has reduced amplitude, also has a reduced conduction velocity. This is because the smaller spike produces a smaller local-circuit current, and this smaller current is propagating into membrane that is partially refractory from the preceding spike and so has a higher threshold. (I must confess I did not anticipate this result when I first built the simulation – I genuinely learned something new using the program.) Consequently, as the pair of spikes propagate along the axon, the second spike drops further and further behind the first. Eventually, when it has dropped back sufficiently far, it escapes from the refractory membrane. At this point it conducts at the same velocity as the first spike, and with the same amplitude. Such a change in spike timing and amplitude between the initiation site and the output terminal may have important consequences for things like synaptic facilitation, or post-synaptic summation.

Neurosim also allows visualization of the time-varying *spatial* distribution of the membrane potential, by simultaneously displaying the potential of each compartment within the model and lining the values up like an animated bar chart (Figure 3). This gives a new perspective to axonal conduction, and this in itself can be very instructive. It is also a perspective that is hard to illustrate with static images – running the simulation and observing the evolution of the membrane potential spatial distribution over time is much more informative than static images such as the figure in this article.

It is worth pointing out that although this phenomenon is not mentioned in any standard textbook that I am aware of, it can be demonstrated quite easily in a real experiment and in fact was shown a very long time ago in the well-known earthworm giant fibre preparation (Bullock, 1951). This is a favorite preparation for undergraduate neuroscience
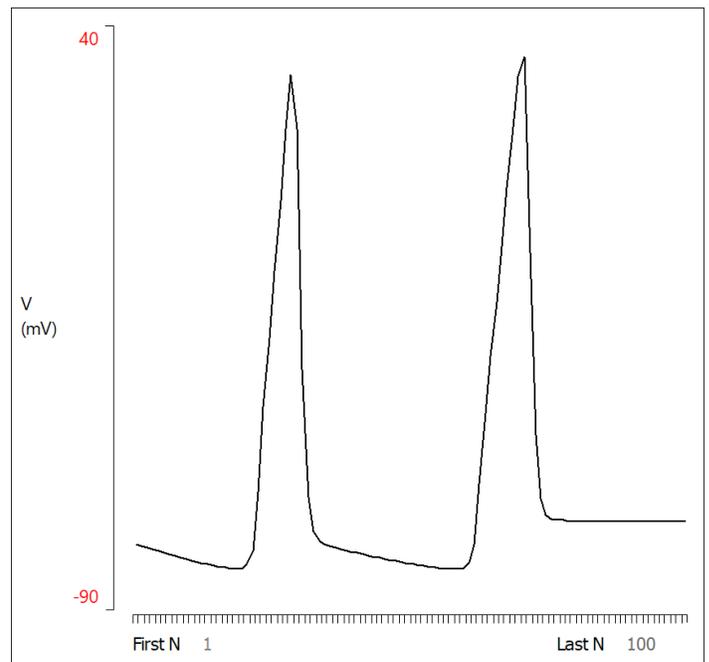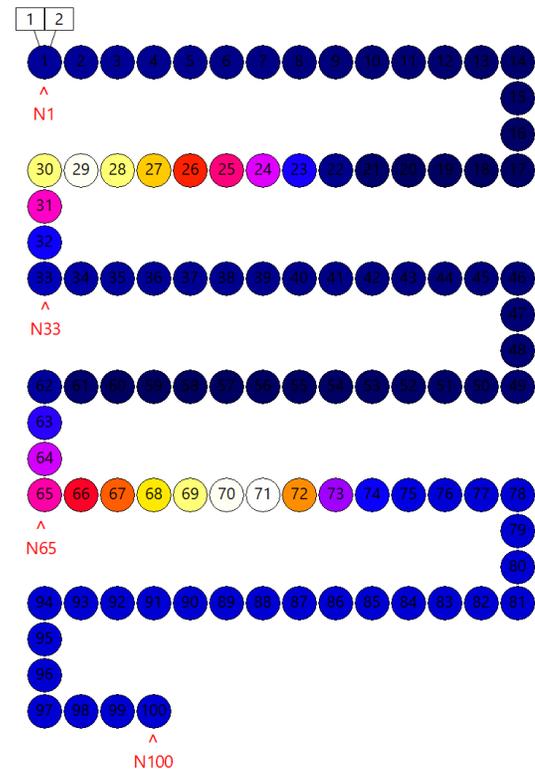
Figure 3. The spatial distribution of membrane potential in an axon propagating two action potentials. Top: In the Network Setup view the 100-compartment model is colour-coded in real time to show the changing membrane potential of each compartment. Spikes were initiated in N1 (top-left compartment) and are propagating towards N100 (bottom-right compartment). Bottom: In the Network Results view the membrane potential (vertical axis) of each compartment (horizontal axis) is shown as the spikes propagate from left (N1) to right (N100) in the view. The results are thus displayed in a voltage-vs-distance format. The display is frozen at a time when the spikes are part-way along the axon.
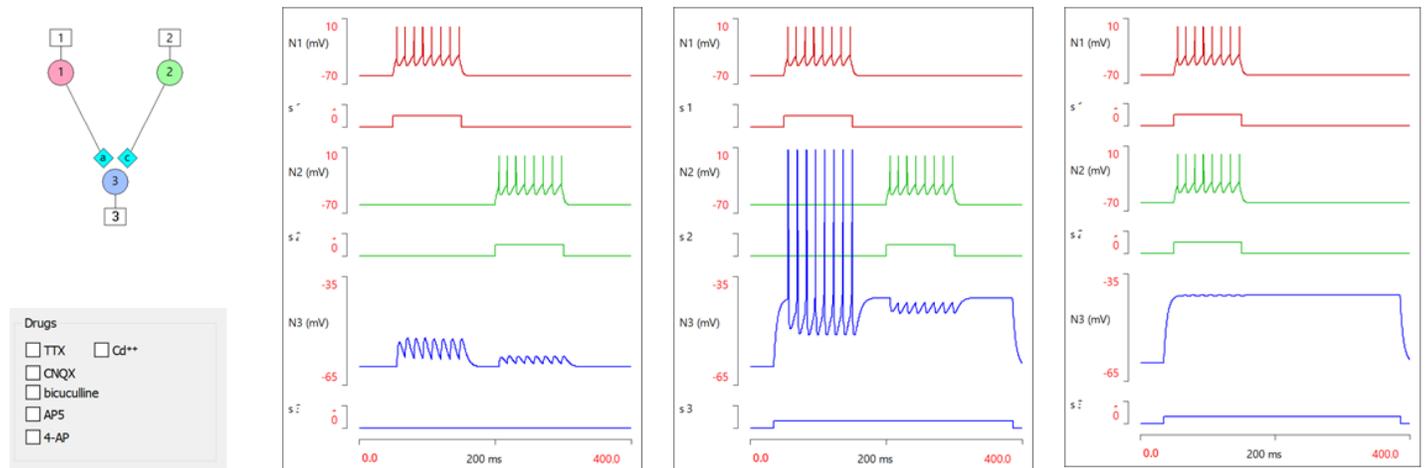
Figure 4. Synapse properties. The Network Setup view shows a simple circuit (left top), and gives the students access to a set of drugs (left bottom), whose names and properties are pre-specified by the tutor in the configuration file. Spikes in the presynaptic neurons induce what look like EPSPs in the postsynaptic neuron (left centre). However, a depolarizing current pulse applied to the postsynaptic neuron reverses the second set of PSPs, suggesting that they are in fact depolarizing IPSPs (right centre). This is confirmed by synchronizing the stimuli to the two pre-synaptic neurons, which shows that the IPSPs cancel the effect of the EPSPs (right).

laboratory classes (Kladt et al., 2010), and so the simulation could form a useful partner to a real laboratory experiment (as could simulations such as that of the threshold strength-duration relationship in the same preparation).

Laboratory Exercises
The main laboratory simulation exercises in our core neuroscience module take place over two 2-hour sessions. The first concentrates on the mechanism of the action potential and its conduction, and the second on the mechanisms of synapses and their integration. The sessions are timetabled so that each occurs shortly after the students have had the lectures that give them the necessary theoretical background. In each session the students are guided through a series of activities, each of which ends with a set of self-test questions. Each student has access to their own workstation, but the activities are not formally assessed, and the students are encouraged to discuss the answers amongst themselves before approaching an instructor to check. Consequently, the students with adjacent computers tend to self-assemble into groups of two or three as the session progresses.

One activity from the synapse session is illustrated in Figure 4. A simple circuit has been constructed using the Network module in which two pre-synaptic neurons make convergent input onto a single post-synaptic neuron. The presynaptic neurons are induced to spike with depolarizing current pulses applied to each in turn, and they each produce depolarizing potentials that look like EPSPs in the post-synaptic neuron. A suite of drugs is available, and the students are instructed to apply each in turn, and to note their effects, if any (they can google a drug name from their workstation if they have forgotten its effect). They then apply a depolarizing current pulse to the post-synaptic neuron, which is timed to bracket the synaptic potentials. They are asked to interpret the results of both the drug and stimulation

experiments, and then to check with an instructor. The hope is that the students will identify one set of PSPs as likely to be glutamate-mediated AMPA-type EPSPs, and the other as chloride-mediated depolarizing IPSPs. The counterintuitive inhibitory effect of the latter can be emphasized by synchronizing the stimuli to the pre-synaptic neurons – the EPSPs no longer generate post-synaptic spikes when summed with the IPSPs, despite the post-synaptic depolarization.

Coursework and Projects
In some of our modules, simulation is used as the basis for assessed coursework. We tend to set tasks that largely require a discursive answer, in part to reduce the risk of inappropriate collaboration that can occur with strictly quantitative questions set as coursework. One task that we have used for several years with senior neuroscience majors is to investigate the mechanism of post-inhibitory rebound (PIR), which can be readily elicited in a standard HH model simply by injecting a pulse of negative current. The students are given full access to the facilities available in the simulation (including hidden variables such as the m, h and n gate probability values), and asked to write a report in the form of a research paper (but without methods) which suggests the mechanism. As a starting point, they are told to focus on the system state at the time when the membrane potential recovers to its resting value on its way towards the rebound spike (Figure 5). At this point the potential is identical to that of the resting neuron, but clearly the system as a whole is in an unstable state because it continues to depolarize and generate a spike. The answer to the mechanism must lie in what is different at this point compared to the initial resting state.

It is important for the students to be aware (or informed) that the HH model is of the squid giant axon, and other neurons may show the same phenomenon but use
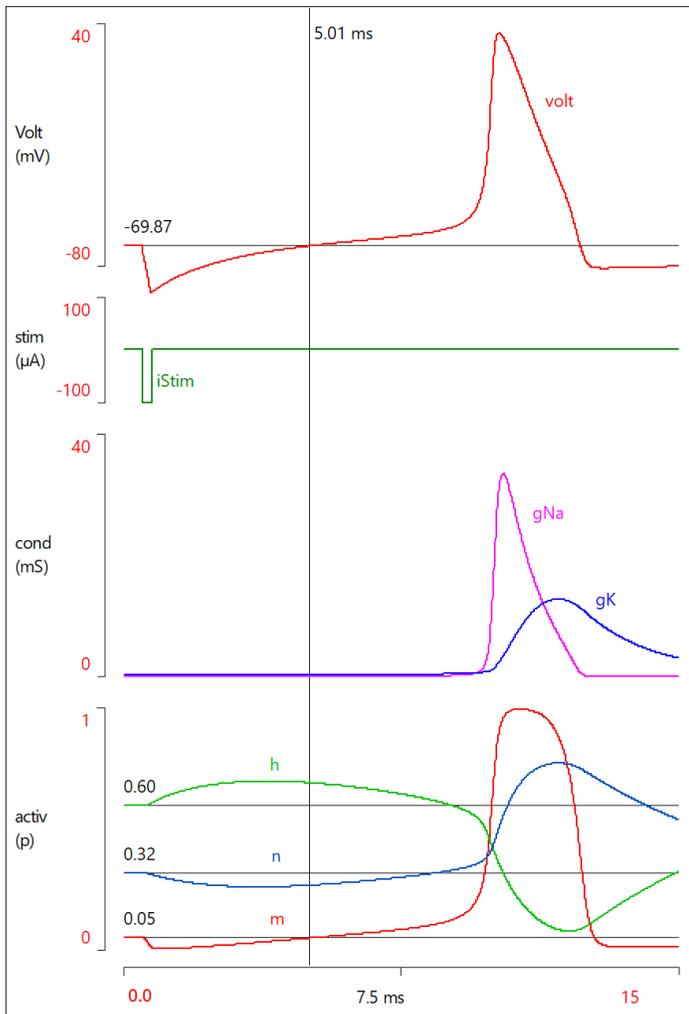
Figure 5. Rebound excitation in the HH model. A vertical cursor indicates the time at which the membrane potential crosses the resting potential on its way to the rebound spike. Students are encouraged to focus on what is different at this time compared to the initial rest condition, which is marked with horizontal cursors in several of the traces.

There are clear differences in the *h* and *n* gating variables. There are consequently small but crucial differences in sodium and potassium conductance, but to see these the student needs to increase the gain of the conductance axis. (This is not shown in the figure because it causes the traces to overlap, which is confusing if the low-gain image has not been seen previously.)

additional mechanisms. For instance, in many neurons a hyperpolarization-activated cation current (I_h) is an important component of PIR, but it is not included in the HH model. However, this can be turned to advantage in the coursework task, since it provides the students with opportunities for literature search and discussion. And of course, if desired, the tutor could always use the Advanced HH module in Neurosim to construct a neuron that *did* include I_h in the PIR mechanism

The PIR investigation is a fairly closed question, since there is only one main experiment. The challenge lies in deciding which variables to measure, and then explaining how these measurements account for the phenomenon. It is possible to set much more open questions using the
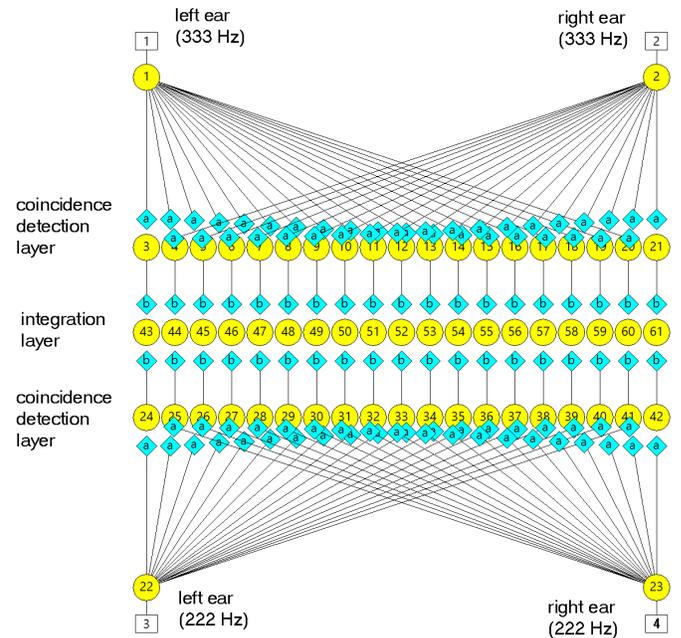
Figure 6. A Network circuit that partially implements the Jeffress mechanism for auditory localization in the azimuth plane, as used in birds such as the owl. It includes two left-right pairs of auditory receptors tuned to different frequencies, which each input through differential delay lines to a coincidence detection layer, which in turn inputs to an integration layer that produces a spatial map of the sound origin.

parameter-hiding facilities in Neurosim. For instance, a student could be presented with a neuron showing an interesting property such as endogenous bursting and asked to characterize the neuron as fully as possible just using current- and voltage-clamp protocols, and a set of tutor-determined drugs. However, in my experience such an open task would require careful preparatory groundwork for it to be successfully accomplished.

Simulation can also be used for extended coursework in the form of a project, although I would only recommend such an assignment for a student who had expressed an interest in computational neuroscience. As an example, one of our students undertook a semester-length project to generate a model of the Jeffress mechanism for auditory localization in owls (Jeffress, 1948; Takahashi, 2010), and a modified version of this is included in the tutorial lessons mentioned previously (Figure 6). The final version produced by the student included more than 300 neurons and 600 synapses, and replicated many experimental results described in the literature with good fidelity.

Data Generation for Tests and Seminar Discussion.
Simulation is an excellent method for generating realistic-looking data that can be used off-line for quantitative analysis in problem-based learning, or simply qualitative discussion and interpretation. Figure 7 shows the potassium current in a simulated voltage-clamp experiment on the squid giant axon that was used in a class discussion.

The students were divided into small groups and each group was given a few minutes to prepare and then deliver a brief explanation of what was happening at the stage of the experiment identified by a particular letter in the figure (some groups had an easier job than others!).

We have frequently used Neurosim to generate "results" such as Figure 7 (or a variety of other types of simulation output) for use in problem-solving questions in written exams. For such usage I would recommend normally including some quantitative element in the question.

SIMULATION: PROS AND CONS

There is quite an extensive literature comparing the relative educational advantages of real laboratory activity and simulation (for reviews see e.g., Ma and Nickerson, 2006; Lewis, 2014), but the general conclusion of most such studies is that both are valuable. What follows is my personal view on why this is correct.

The most obvious benefits of simulation are the nominal cost, zero animal usage, increased reliability, and the ability to do virtual experiments that are beyond the technical capability of students. However, the really important benefit, to my mind, is the increase in the depth of understanding that simulation can offer. By giving students access to underlying hidden variables, and by allowing them to carry out speculative "what if" experiments with zero cost and little time penalty, simulation can greatly enhance understanding of both core principles, and the experimental methodologies used to investigate them. For instance, even advanced students often have difficulty explaining the cause of the current profiles labelled D and E in Figure 7. However, once grasped, it can lead to a step improvement in their understanding of voltage-dependent channels, driving force, and the voltage-clamp technique itself.

So, not surprisingly, I am definitely an enthusiast for the use of simulation in teaching neuroscience. There are, however, some important caveats that both teachers and students need to be aware of.

Perhaps the most important issue is that when carrying out a simulation experiment, the student is studying the *model*, not a real biological system. Hodgkin and Huxley showed they were well aware of this when they wrote in their classic 1952 paper "...*the success of the equations is no evidence in favour of the mechanism of permeability change that we tentatively had in mind when formulating them*". Just because a model "works" (reflects reality) within the limits of the tests applied, it does not mean that it reflects reality in its mechanism.

Another more subtle problem with simulation is that the experiments always work, and always yield interpretable results (at least, barring bugs in the program). This is great for aiding understanding, but it can give a completely false impression of how science actually works. In most real experiments there are unexpected results and uninterpretable deviations from the expected outcome. Part of the job of a scientist is to decide whether an apparent glitch in a real recording is the result of a fridge switching on in the lab next door, or a potentially Nobel prize-winning observation that will open up an entirely new perspective on
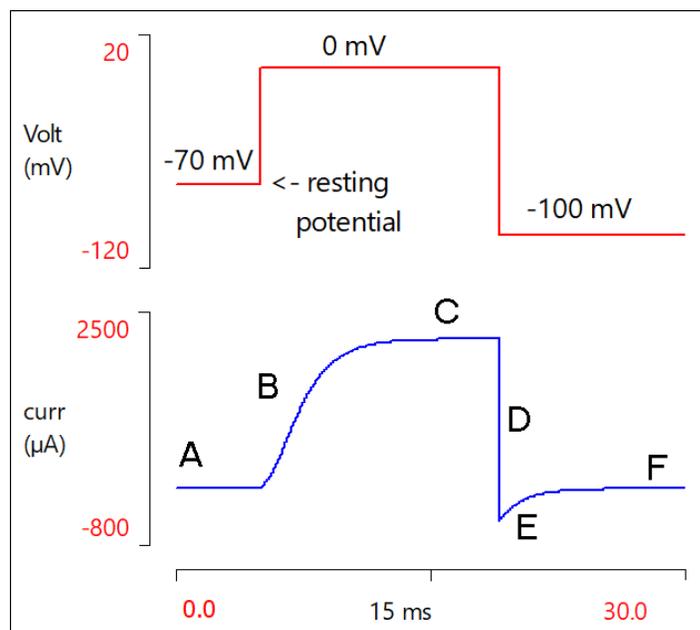

Figure 7. A simulated voltage clamp experiment in which the current flowing through voltage-dependent potassium channels has been isolated by application of TTX, with leakage and capacitive current subtraction (or, in truth, by simply choosing the K current as the trace content in the Neurosim configuration file). Upper trace: clamp potential; lower trace: clamp current.

the subject!

Related to this is the issue of noise. Real data contain noise, and, unless it is intentionally added, simulated data do not. And even if noise is added (which may actually be a key part of a simulation, as in the tutorial lesson on stochastic resonance), it is likely to be "pure" noise, rather than, for instance 50/60 Hz mains interference. Any instructor with experience of running real laboratory exercises will know the shared frustration of discovering that a student has spent significant time analyzing interference signals, rather than real neural signals. Simulation avoids the problem, which can cut down on wasted time, but it does not prepare a student for this challenge in real research.

At a very practical level, a complex simulation exercise involves a student spending an extended period of time, often on their own, sitting still and staring at a computer screen and, hopefully, thinking very hard. It is therefore not something that should be timetabled for long periods! In contrast, a real lab often requires moving around, it gives multiple things to pay attention to (bits of apparatus, chart print-outs, maybe a real animal), it usually involves some intellectual down-time while waiting for results or for drugs to take effect etc., and it is often a social activity requiring interaction with lab partners.

Finally, real experiments reify concepts that have been taught at a theoretical level and can generate a "wow" factor that is rarely present in simulation. I have, literally, heard that expression from many students when they first see action potentials generated by, for instance, the classic cockroach sensory spine experiment (Linder and Palka, 1992).

CONCLUSION

Overall, I think that simulation is probably a better tool for developing deep theoretical understanding of the current state of knowledge than the sorts of real laboratory experiments that are feasible at the student level. However, I think that real laboratory experiments are essential for motivating students, and for developing a critical awareness of the scientific process. They are also of course essential for developing the technical skills which will enable future researchers to advance the current state of knowledge, thus enabling neuroscience teachers to generate more accurate simulations! So, in my opinion, a judicious combination of simulated and real laboratory experiments provides the optimal pedagogic toolset for teaching students, from beginning to advanced levels.

REFERENCES

Bullock TH (1951) Facilitation of conduction rate in nerve fibres. J Physiol 114:89-97.

GENESIS (2019). GENESIS: The GEneral NEural SImulation System. Available at http://genesis-sim.org.

Hodgkin AL, Huxley AF (1952). A quantitative description of membrane current and its application to conduction and excitation in nerves. J Physiol 117:500-544.

Jack JJB, Noble D, Tsien, RW. (1975). Electric current flow in excitable cells. Oxford, England: Clarendon Press

Jeffress LA (1948) A place theory of sound localization. J Comp Physiol Psychol 41:35-39.Kladt N, Hanslik U, Heinzel H-G (2010) Teaching basic neurophysiology using intact earthworms. J Undergrad Neurosci Educ 9(1):A20-A35.

Lewis DI (2014) The pedagogical benefits and pitfalls of virtual tools for teaching and learning laboratory practices in the Biological Sciences. York, United Kingdom: The Higher Education Academy. Available at https://www.heacademy.ac.uk/sites/default/files/resources/the_pedagogical_benefits_and_pitfalls_of_virtual_tools_for_teaching_and_learning_laboratory_practices_in_the_biological_sciences.pdf.

Linder TM, Palka J (1992) A student apparatus for recording action potentials in cockroach legs. Am J Physiol 262:S18–22.

Ma J, Nickerson JV (2006) Hands-on, simulated, and remote laboratories: a comparative literature review. ACM Computing Surveys 38(3):7 Available at http://doi.acm.org/10.1145/1132960.1132961.

Macgregor RJ (1987) Neural and Brain Modelling. London, United Kingdom: Academic Press.

McAteer E, Neil D, Barr N, Brown M, Draper S, Henderson F (1996) Simulation software in a life sciences practical laboratory. Computers Educ 26:101-112.

NEURON (2021). NEURON: Empirically-based simulations of neurons and networks of neurons. New Haven, CT: Yale University. Available at https://neuron.yale.edu/neuron/.

Takahashi TT (2010) How the owl tracks its prey – II. J Exp Biol 213:3399-3408.

Wilson HR, Cowan JD (1972) Excitatory and inhibitory interactions in localized populations of model neurons. Biophys J 12:1-24.

Address correspondence to: Dr. W. J Heitler, School of Psychology and Neuroscience, University of St Andrews, St Mary's Quad, South Street, St Andrews, Fife KY16 9JP, U.K. Email: wjh@st-andrews.ac.uk