Electronics for electrophysiologists

Boris Barbour*†

September 19, 2018

Abstract

This article aims to present as simply as possible the fundamental principles of operation of electrophysiology amplifiers and to explain how to obtain high-quality recordings.

1 Basics

Excitable cells, including neurones, express themselves electrically. We therefore begin by recalling the electrical nature of neurones and the properties of their circuit equivalents.

A cell is delimited by its lipid membrane. A pure lipid bilayer is an excellent insulator. Both inside and outside the cell is a dilute saline solution, which is a conductor. The sequence conductor-insulator-conductor defines a capacitor. The key characteristic of a capacitor is that it can store charge (Q, measured in Coulombs) proportionally to the applied potential or voltage, giving the equation $Q = CV$, where the constant of proportionality is the capacitance C, measured in Farads (F). Typical values encountered in electrophysiology are very small fractions of a Farad—picoFarads (10^{-12} F) and nanoFarads (10^{-9} F). The *specific capacitance* of biological membranes is considered to be $1\,\mu\mathrm{F\,cm}^{-2}$.

---

*Why another text on electrophysiology? The explanations in this article have been tested and honed over many years at the Microelectrodes Techniques course run by David Ogden in Plymouth, UK, each September, and more recently the ENP course in Paris. In addition, I have been able to illustrate many of the issues using dual recordings from neurones. So I have some hope that this article will provide an accessible introduction to what is often a poorly-understood subject. I intend to extend and refine this article in the future, so interested readers are invited to check my work web site for updated versions. Substantive changes to the text are listed in §6 at the end of the document. Any corrections or suggestions will be welcomed; my email address is boris.barbour@ens.fr.
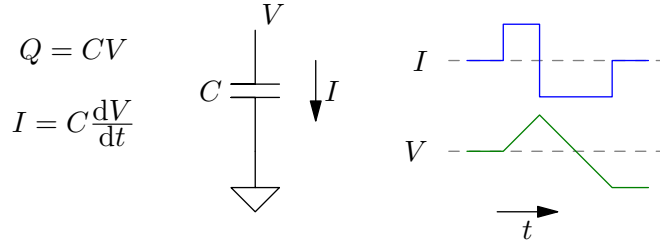
Figure 1: The capacitor and equations governing its behaviour. The example traces on the right show the voltage (green) resulting from the injection of the current above (blue). Note in particular that the capacitor prevents instantaneous changes of voltage. The imbalance of charge injected results in a different final voltage.

Differentiation of $Q = CV$ yields $I = C\frac{\mathrm{d}V}{\mathrm{d}t}$. Inspection of this equation reveals the key property of a capacitor for neuroscience: in order to change the voltage across the capacitor, a current must flow; if no current flows ($I = 0$) the voltage does not change ($\frac{\mathrm{d}V}{\mathrm{d}t} = 0$). An instantaneous or step change of voltage would imply an infinite rate of change of voltage, which would require infinite current—an impossibility. In other words, *a capacitor prevents instantaneous changes of voltage across its terminals*. The voltage across a capacitor's terminals is *continuous* in time. Fig. 1 summarises these properties. Remember that capacitances in parallel sum, while capacitances in series combine 'reciprocally' ($C_{tot} = \frac{1}{\frac{1}{C_1} + \frac{1}{C_2}}$).

The plasma membrane is not a pure capacitor, however. It contains numerous transmembrane proteins—channels—that conduct charged ions, constituting a variable conductance (or resistance) in parallel with the cell capacitance (Fig. 2). Current flow through these conductances charges or discharges the membrane capacitance and is thus responsible for variation of the membrane potential. For the purposes of the analysis we shall carry out, it suffices to consider the lumped resistance of the cell without regard to the various reversal potentials. Current flow in a resistor is of course governed by Ohm's law: $V = IR$. Resistors in series sum, while resistors in parallel combine 'reciprocally' ($R_{tot} = \frac{1}{\frac{1}{R_1} + \frac{1}{R_2}}$).

Although most explanations of recording amplifiers are based upon a simple one-compartment cell, most neurones unfortunately do not conform to this model at all well. Dendrites, but also axon(s), can only be accurately represented as a multi-compartment model. A few cells can be effectively approximated by two or three compartments (Fig. 3), though a fully accurate model can require very many compartments.
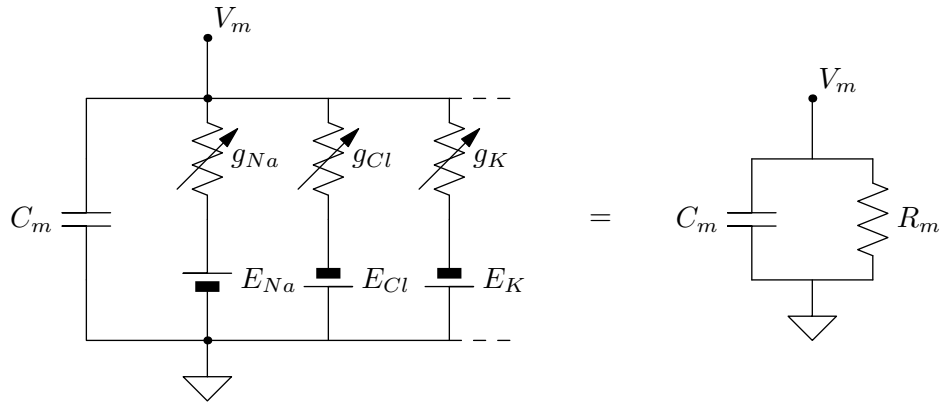
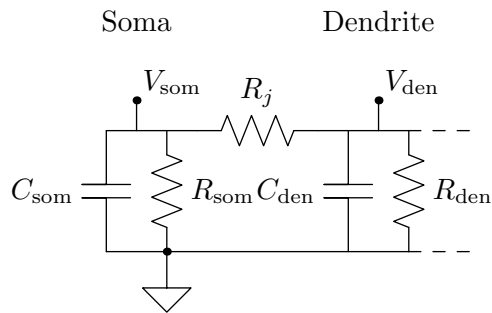Figure 2: Equivalent circuits of a one-compartment cell.

Figure 3: More realistic equivalent circuit for neurones. Usually at least somatic and dendritic compartments are required.

2 Measuring voltage (current clamp)

Before explaining the measurement of the membrane potential, we need to understand the operational amplifier. A brief explanation is available in Appendix 1.

2.1 Voltage follower

Having understood op-amps, we are now in a position to make our first measurement of the membrane potential. We connect an electrode (micro-electrode or patch electrode) to our preamplifier/headstage and impale or seal onto a cell, thereby obtaining electrical access to the inside of the cell. A simple equivalent circuit is shown in Fig. 4.

It should be noted that older patch-clamp amplifiers implement a different circuit that is less suitable for voltage recording.

The terminology used for describing the various resistances involved in electrophysiological recordings is confusing. *Series resistance* and *access resistance* are synonymous with the *electrode resistance*. Quite distinct, however, is the *input resistance*, which describes the total resistance observed by the amplifier; it is therefore equal to the sum of the *membrane resistance* and the *electrode resistance*, with the former generally dominating the sum. You may not hear of the input resistance in your work.

Injection of a step current into a cell will generate a charging curve typically composed of one or more exponential components. The slowest is the *membrane time constant*, τ_m. Faster components may represent redistribution of charge within multicompartment cells.

2.2 Bridge balance

As the circuit of Fig. 4 makes clear, the membrane potential is recorded through an electrode, which introduces several artefacts into the recording. The easiest artefact to understand arises from the electrode resistance. Any current injected into the cell from the amplifier will cause a voltage drop across the electrode resistance. By Ohm's law, $V_{\text{err}} = I_p R_e$. Clearly this error will be zero if no current is injected and can be minimised by using large electrodes containing highly-conductive solution (which will give a low resistance).

This voltage error is moreover relatively easy to correct for, since we (and the amplifier) know the injected current. Most amplifiers provide a control allowing the user to estimate the series resistance. This value is used by the amplifier to generate the expected voltage drop across the resistance and to correct the recorded voltage for this error. In practice, this is done by injecting a current step and adjusting the electrode resistance control until no instantaneous voltage change is observed at the beginning of the
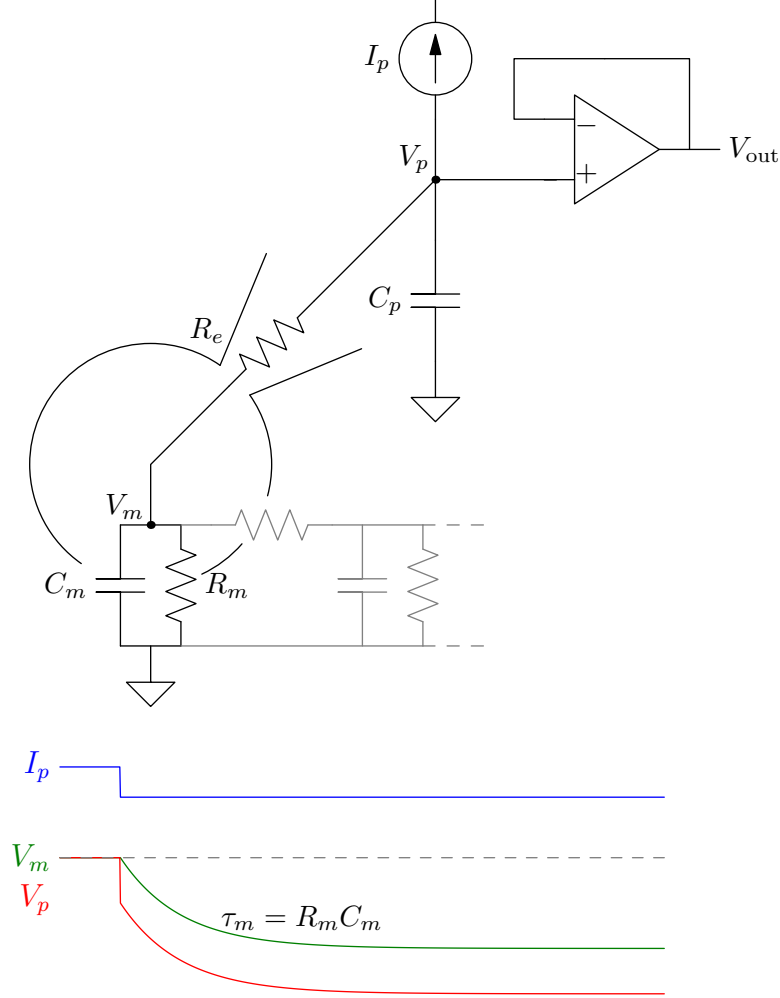
Figure 4: Voltage recording. An op-amp in the *voltage follower* configuration is connected to the cell via an electrode that has resistance R_e and contributes to the parasitic capacitance C_p. A current source controls the injected current I_p. The negative feedback connection ensures that $V_{\text{out}} = V_p$. The traces below show the expected cell and pipette voltages during a step current injection, assuming that C_p is negligible. The membrane potential charges exponentially to a steady value determined by the membrane resistance. The time constant of the exponential is the *membrane time constant* τ_m. The current flow across R_e generates a voltage drop and therefore an error in the recorded voltage.

step. With reference to Fig. 4, adjusting the control would shift the red voltage trace to the green voltage trace. This is called *bridge balance*, because the circuit was initially implemented using a Wheatstone bridge (a resistor network). All modern amplifiers use a different implementation using op-amps.

It is important to realise that balancing the bridge is an adjustment of the output signal; *it has* **no** *effect on the cell*. Throughout this article, we shall be making a distinction between adjustments that affect the output signal but not the cell—*cosmetic* operations—and those that also affect the cell, which often involve an element of *positive feedback* and therefore danger. Note that there is little consensus regarding the names for many of the adjustments available on amplifiers, with terms often being contradictory between different manufacturers ('compensation' is particularly abused in this regard). You must therefore always sort out for yourself when your amplifier is making a cosmetic or 'effective' change—the use of positive feedback is often the distinguishing feature of the latter.

It is time to show some specimen traces from a real cell. The examples shown below are from recordings of an adult cerebellar Purkinje cell. An image of this cell type is shown in Fig. 5. The Purkinje cell is among the larger neuronal types in the brain, having an extensive dendritic tree. It is therefore definitely *not* a simple one-compartment cell. It can however be usefully approximated by a two-compartment model (see Fig. 3).

In order to demonstrate the effects of various amplifier adjustments, cells were recorded with two electrodes (Fig. 6), one of which was optimally adjusted to give a true measure of the somatic membrane potential while the second electrode and amplifier underwent the full sequence of adjustments we shall demonstrate. The electrode reporting V_m was always held with zero current flow (preventing any artefacts that would result from a voltage drop across the electrode resistance). As no alterations were made to the V_m electrode during the recordings, any changes observed through it must reflect changes of cell behaviour. Thus, this approach provides clear insight into electrode adjustments that change the behaviour of the cell, as opposed to simply modifying amplifier output (i.e. the cosmetic operations).

The effect of balancing the bridge in a Purkinje cell is shown in Fig. 7, which clearly demonstrates the 'cosmetic' nature of this operation.

2.3 Capacitance neutralisation

The parasitic capacitance C_p in the circuit of Fig. 4 has a more pernicious effect. The capacitor arises from the capacitance between the electrode filling solution and the bath saline combined with some unavoidable capacitances associated with the headstage and its electronics.

Remember that a capacitor requires current flow to change its voltage. Thus, in order to record a change of membrane potential at the amplifier,
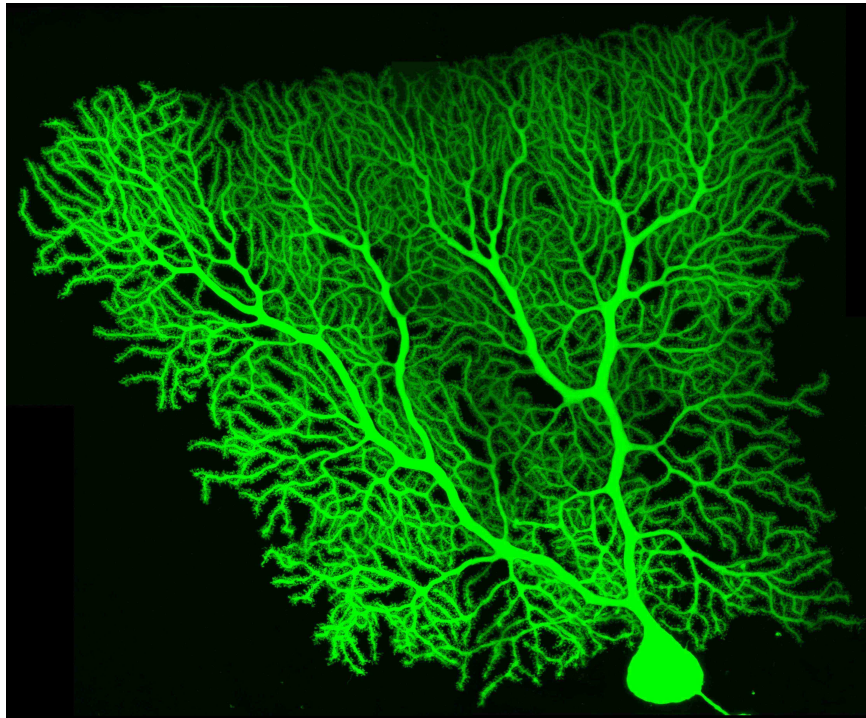
Figure 5: Cerebellar Purkinje cell. Specimen recordings shown in this article were obtained from such a cell, which is slightly more than $200\,\mu$m square.
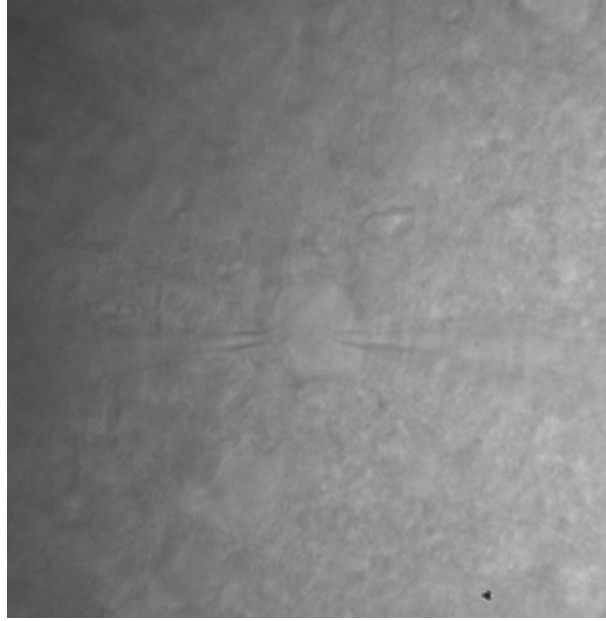
Figure 6: Purkinje cell soma with two electrodes. One was adjusted optimally to give a 'true' reading of the V_m, to show the effects of adjusting the second electrode.

the parasitic capacitance must be charged. As the circuit is drawn, the required current must be supplied by the cell through the electrode resistance R_e. This has two undesirable effects: firstly, rapid voltage changes in the cell (think action potentials) are slowed and distorted when recorded at the amplifier (the electrode resistance and parasitic capacitance combine to form a low-pass RC filter); secondly, cell behaviour is altered by the current flowing to the parasitic capacitor. I call these two distinct actions the *filtering* and *loading* effects of the parasitic capacitance. The former changes what you record, without altering cell behaviour. The latter, however, alters cell behaviour and, as such, must be corrected for as much as possible before recording.

A partial solution to this problem is to use a circuit called *capacitance neutralisation*, in which the output is used to charge the parasitic capacitance on the input (Fig. 8). This is a positive feedback circuit that can cause oscillation and cell death if incorrectly adjusted. The following argument gives some insight into how the circuit operates. Imagine that $C_n = C_p$ and that V_{out} is amplified two-fold ($G = 2$). Assume initially that $V_m = V_p = V_{out} = 0$ and then a step change of the membrane potential to V occurs. In the steady-state, $V_m = V_p = V_{out} = V$. A charge $Q_p = C_p V$ has been deposited on C_p. The voltage across C_n has changed from 0 to
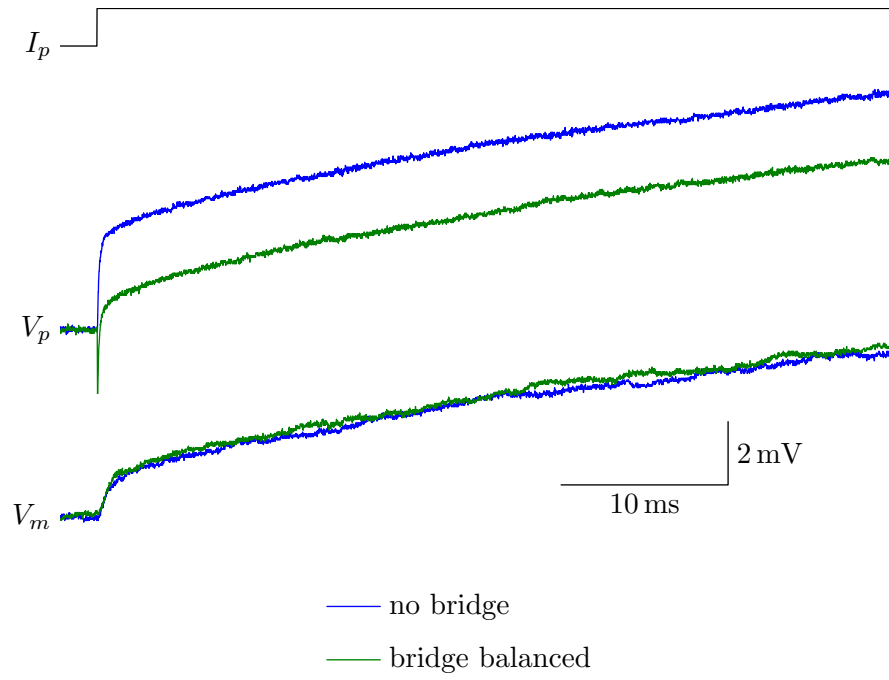
Figure 7: Balancing the bridge in a Purkinje cell. Note the double-exponential charging curve reflecting firstly rapid charge distribution within the cell, then charging of the whole cell according to τ_m. Balancing the bridge has no effect on the cell (V_m), only on the recorded voltage (V_p).

$2V - V = V$, so a charge $Q_n = C_n V$ must have flowed across that capacitor. But since $C_n = C_p$, that charge is the same as that required to charge C_p. In other words, the amplifier now charges C_p instead of the cell supplying the necessary current. This has the effect of very significantly reducing C_p, diminishing both the filtering and loading effects. Of course, perfect neutralisation is not possible, but in practice the circuit works quite well, because the positive feedback around the amplifier is much quicker (MHz) than any cellular signal (kHz).

It might erroneously be supposed that the loading effect will only influence small cells. Obviously, conductances that usually charge a cell capacitance of 2 pF (such as for the cerebellar granule cell) will be strongly affected by an additional load of 5–10 pF (typical parasitic capacitance values). Conversely, for cells such as the Purkinje cell which may total 1500–2000 pF, the addition of 10 pF would at first appear negligible. But this is not always the case, as shown in the specimen traces from the Purkinje cell (Fig. 9, in which the V_m electrode sees slightly faster and taller action potentials when the other electrode has its capacitance neutralised). The reason for this is that action potentials are initiated in the somatic region, which has a much smaller capacitance, of the same order of magnitude as the parasitic capacitance.

3   Voltage clamp

The study of voltage-dependent conductances requires the ability to control the membrane potential—to clamp the voltage. There are classical techniques for doing this using two electrodes, one to measure voltage (as described above) and one to pass the current required to ensure that the voltage recorded remains constant whatever currents flow in the cell. However, it is often impractical or impossible to use two electrodes. Here we analyse a method for approximate voltage-clamp using a single patch electrode. This makes use of the current follower op-amp configuration (see Appendix 1). The basic circuit is shown in Fig. 10. Negative feedback ensures that the inverting input of the amplifier is always at the command voltage V_c. In order for this to occur, any current entering the inverting node must be removed via the feedback resistor, and, by Ohm's law, the output voltage is therefore $V_{\text{out}} = V_c + I_p R_f$. In exotic amplifiers, the feedback current can be supplied via a capacitor ('integrating headstage') or even light-sensitive diodes, instead of via a resistance.

3.1   Capacity transient

As for voltage recording, the properties of the electrode introduce unwanted artefacts. In this voltage-clamp mode, the current follower clamps the end of the electrode, not the membrane potential. So, if any current flows across
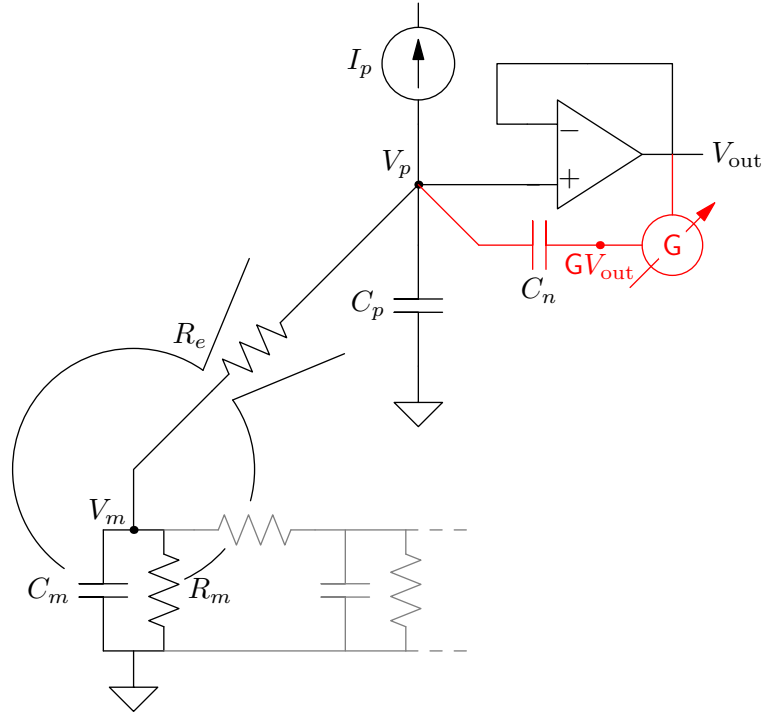
Figure 8: Voltage recording with capacitance neutralisation. A positive feedback circuit is added (red) which injects an amplified version of the output back into the input via a capacitor. When the gain of the feedback is correctly adjusted, a voltage change of the input very quickly causes an injection of charge at the input such that the amplifier and not the cell charges C_p. This reduces both the filtering and the loading effects of the parasitic capacitance.
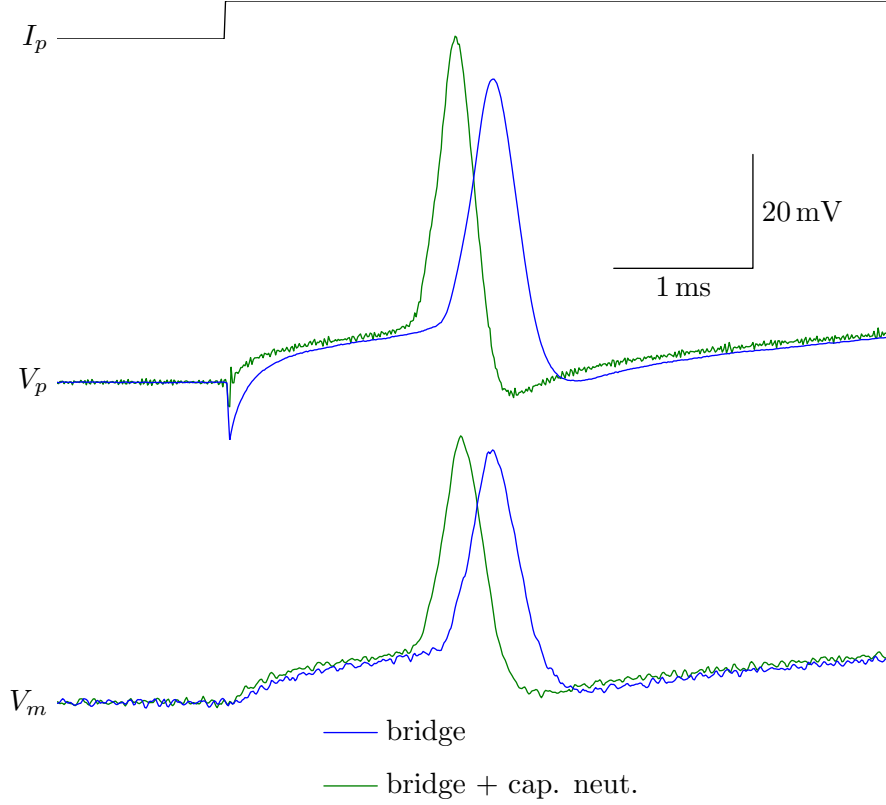
Figure 9: Effect of capacitance neutralisation in the Purkinje cell. The V_p traces with and without capacitance neutralisation show the combined filtering and loading effects of the parasitic capacitance, while the V_m traces show only the loading effect. In the adjusted electrode, capacitance neutralisation sharpens the electrode artefact (at the beginning of the current step) and reveals briefer action potentials of greater amplitude. Threshold would be expected to be attained earlier (but the jitter in action potential timing means that the change shown in the figure is not necessarily significant). Despite the large size of the Purkinje cell, its relatively small soma *is* sensitive to the loading effect when the voltage changes quickly. (The difference of action potential size between the two electrodes after compensation is unexplained.)
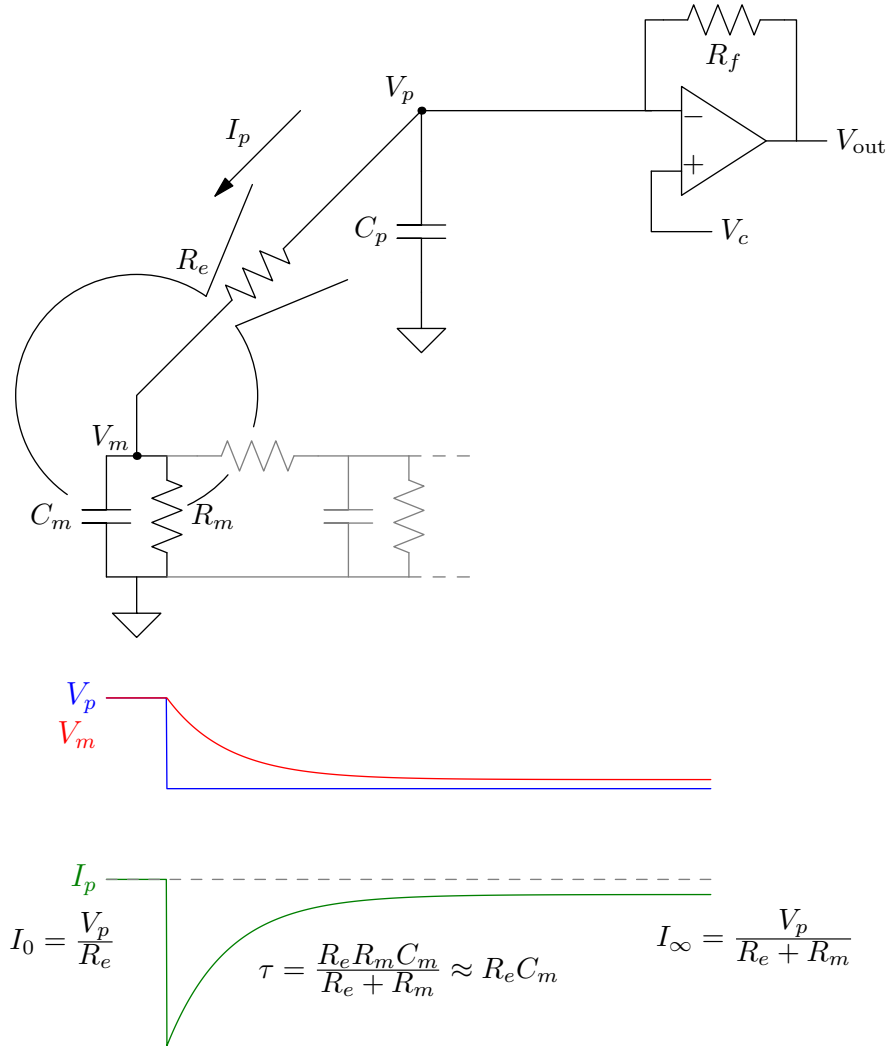
Figure 10: Voltage clamp. Single-electrode voltage clamp implemented using an op-amp in the current follower configuration. The inverting input is clamped to the command voltage V_c by the negative feedback. Because all pipette current must pass through the feedback resistor to ensure this clamp, the output voltage is proportional to the current flowing. The traces below show the current and voltage transients in response to a command voltage step.

the electrode resistance, there will be a voltage drop. More subtly, the interaction of the electrode resistance and the cell capacitance cause a strong low-pass filtering effect. This can be seen by examining the current and voltage transients in response to a command voltage step. The current and membrane potential relax to their final values according to an exponential decay with a time constant $\tau \approx R_e C_m$. This time constant governs all interactions between cell and amplifier. A rapid conductance change in the cell will be filtered by this time constant. Equally, a voltage step at the amplifier will not be instantly transmitted to the cell.

3.2 RC cancellation

Users of patch-clamp amplifiers are often bewildered by the numerous adjustments to be made at the beginning of a recording. The initial adjustments are 'cosmetic', according to our classification, but are useful, even necessary, preliminaries to those interventions that really change cell behaviour.

The first adjustments are to cancel the pipette and cell capacity transients. The circuits for both of these operations make use of the same principle, in that the amplified command voltage is coupled directly to the input via a variable capacitor and resistor in series (or equivalent circuit). In this way a capacity current exactly equal to the transient required to charge the capacitance of interest (electrode or cell) can bypass the recording amplifier, hence the term cancellation. See Fig. 11. Note that the electrode and cell remain attached to a node whose voltage time course remains exactly equal to the command voltage; therefore the cell's behaviour is not altered. All that does change is that the required current can be transferred between the recording circuit and the resistor/capacitor networks connecting the command to the headstage input.

The cancellation of the electrode capacitance is important for subsequent compensation of the electrode resistance—otherwise the capacitance present at the inverting input of the amplifier will cause it to oscillate. Cancellation of the cell transients or part of them can also prevent amplifier saturation during large voltage steps (this is also true for cancellation of the electrode capacitance). Saturation occurs when the finite power supply of the operational amplifier prevents the required current from being driven across the large feedback resistor. When this occurs, the amplifier neither clamps voltage nor measures current correctly, and recovery from this state may be delayed, though this delay is less of a problem with modern equipment. Adjustment of the cancellation yields a good estimate of the electrode resistance, which the amplifier requires for other adjustments. It also helps the user to judge the quality of the recording and to assess any changes of electrode resistance.
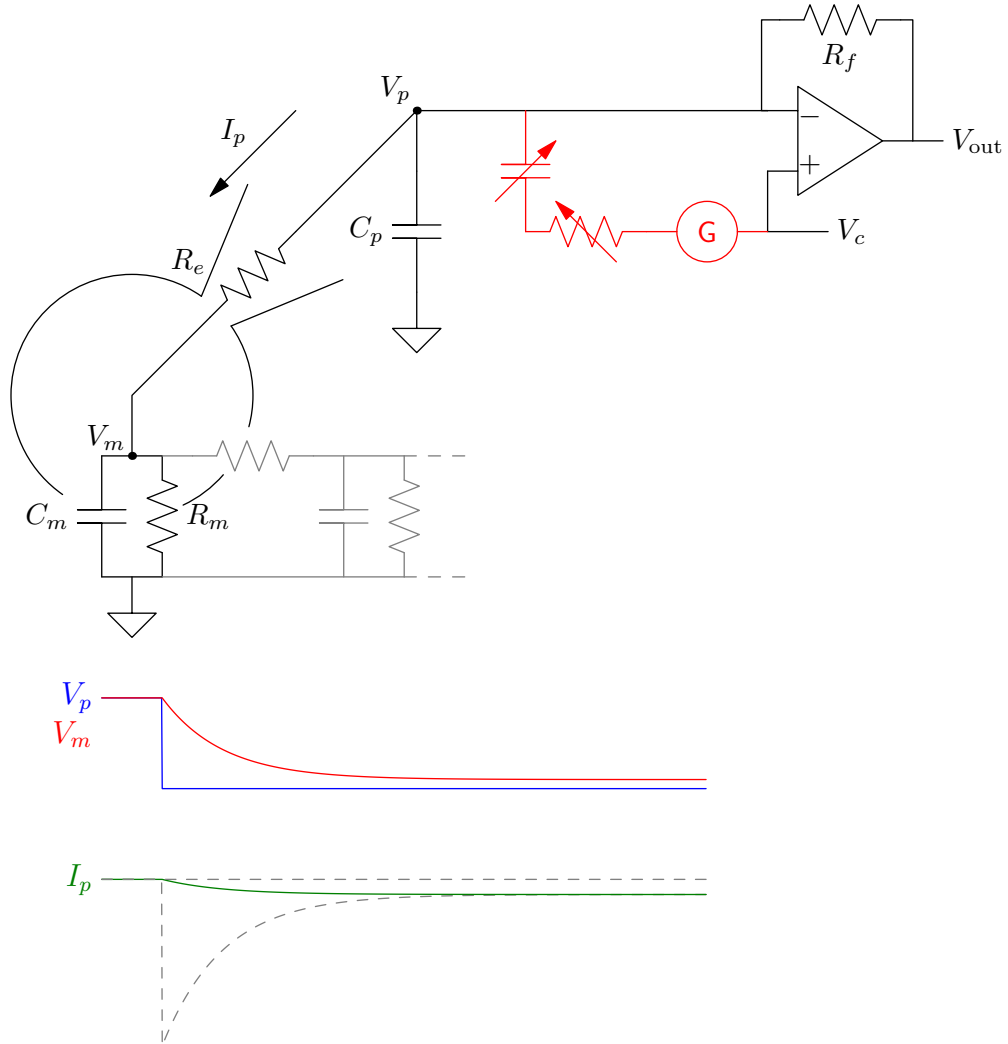
Figure 11: Cancellation of capacity transients. The equivalent of a variable resistor and capacitor (red; G represents an amplification) provide a parallel pathway for capacity currents required to charge the pipette or cell (there are often two or three such circuits connected in parallel, with different parameter ranges appropriate for cancelling the electrode or cellular capacity transients). These currents bypass the amplifier and are therefore not reflected in the output (the dashed grey transient is without cancellation and the solid green line is with cancellation). However, as the cell remains connected to the inverting input, which is still always clamped at the command voltage, the behaviour of the membrane potential is unchanged.
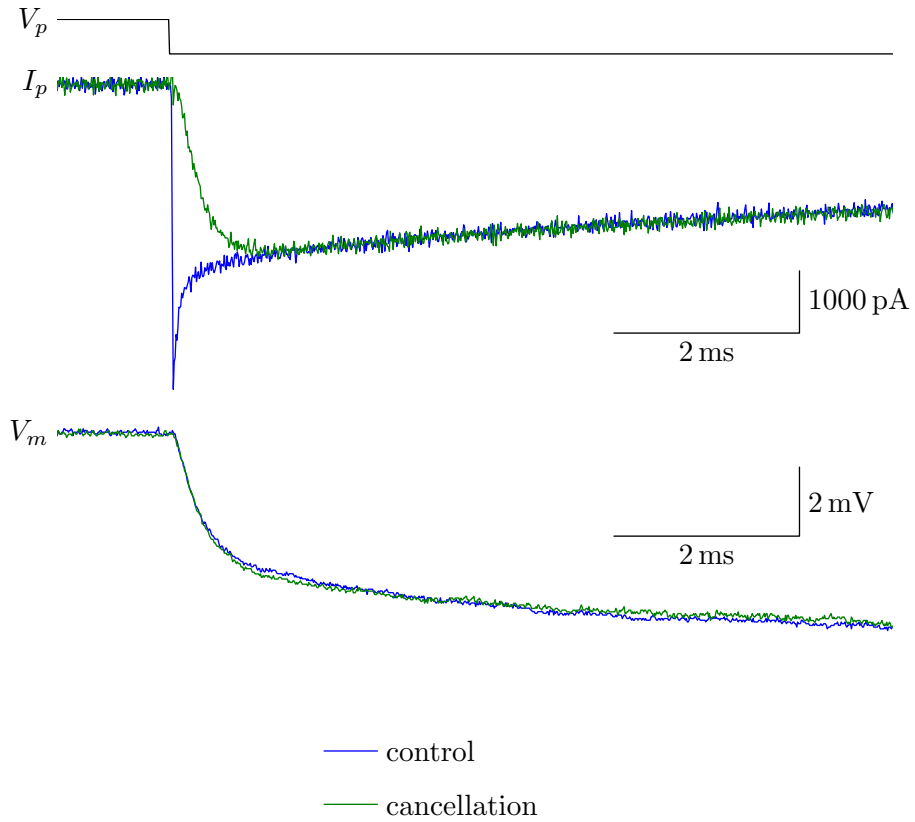
Figure 12: Effect of cancellation of capacity transients in the Purkinje cell. The Purkinje cell has a complex capacity transient, reflecting its ramified morphology. It is therefore only possible to cancel an initial component of the transient. Note that the cancellation has no effect on the membrane potential.

3.3 Series resistance compensation

The interaction of the electrode resistance with the cell capacitance causes voltage errors and unwanted filtering. Clearly, the first step to reducing this problem is to arrange for electrodes to have as low a physical resistance as possible, but in patch-clamping this can only be done by using larger electrodes, so the reduction of resistance that can be attained is limited. There is, however, an electronic trick that can effectively reduce the electrode resistance. This is variously called compensation or correction of the series resistance (a synonym for the electrode resistance). This mechanism involves positive feedback and does change cell behaviour. The degree of positive feedback is adjusted as a percentage of the electrode resistance that was previously estimated by setting the cancellation controls. As with all positive feedback, overcompensation will lead to oscillation and cell death.

The positive feedback adds a fraction of the output to the command potential (strictly, it is a signal proportional to the current which is added to the command; see Fig. 13). This tends to increase the pipette voltage when current flow and the voltage error is greatest, driving extra current through the electrode, which is exactly what would happen if the electrode had a lower resistance. This argument can be made exact. Without compensation, the simple application of Ohm's law governs current flow through the electrode: $V_p = V_c = I_p R_e$ (assuming $V_m = 0$ for simplicity). After compensation, $V_p = V_c + kI_p'$ (where the prime indicates a new value) and $V_p = I_p' R_e$. Eliminating V_p and rearranging, we obtain $V_c = I_p'(R_e - k)$, which has exactly the form of Ohm's law, but with an apparently lesser electrode resistance of $R_e - k$. So, idealised compensation of the electrode resistance causes the system to behave exactly as if the electrode resistance had been reduced. Solution of the model cell for the current and voltage time courses after compensation show the transient current is larger and faster and V_m approaches the command value more quickly and more closely (Fig. 13).

Because the compensation circuit reacts to the current flowing in the recording circuit (as opposed to the cancellation circuit), it should be tested with the cancellation switched off, so that the circuit can react to the capacity current. This point is rarely explained correctly in manuals for patch-clamp amplifiers.

Fig. 14 shows the effect of compensating the series resistance in a Purkinje cell. Without compensation, the membrane potential only 'quickly' reaches about 50 % of the command potential and then approaches its final value very slowly. Application of optimal series resistance compensation clearly increases the peak capacity current (indicating an apparent reduction of the electrode resistance) and the recording of the membrane potential shows that clamp of the soma is greatly improved, rapidly approaching some 80 % of its target value.

In practice, series resistance compensation only works well (70–80 % com-

pensation) for large cells, though some improvement can be obtained with small cells (20–40 %). In most cases it is necessary to cancel carefully the electrode capacitance and to make use of a filter (often called 'lag') that filters the positive feedback signal; this improves stability at the cost of a filtered signal, so the compromise needs to be optimised with care. A reasonable choice is to increase compensation until some ringing is observed and then to reduce the compensation slightly. In the specimen traces from the Purkinje cell, this ringing is just apparent at the break between the two exponential components of the current transient.

Series resistance compensation will increase the high-frequency noise in the recording, and this is often the reason given for not using compensation. However, at least part of the increased noise is to be expected, since the recording will be less filtered. It is preferable to use compensation, obtaining a better voltage clamp (see below) during the recording, then to filter out any bothersome high frequencies afterwards.
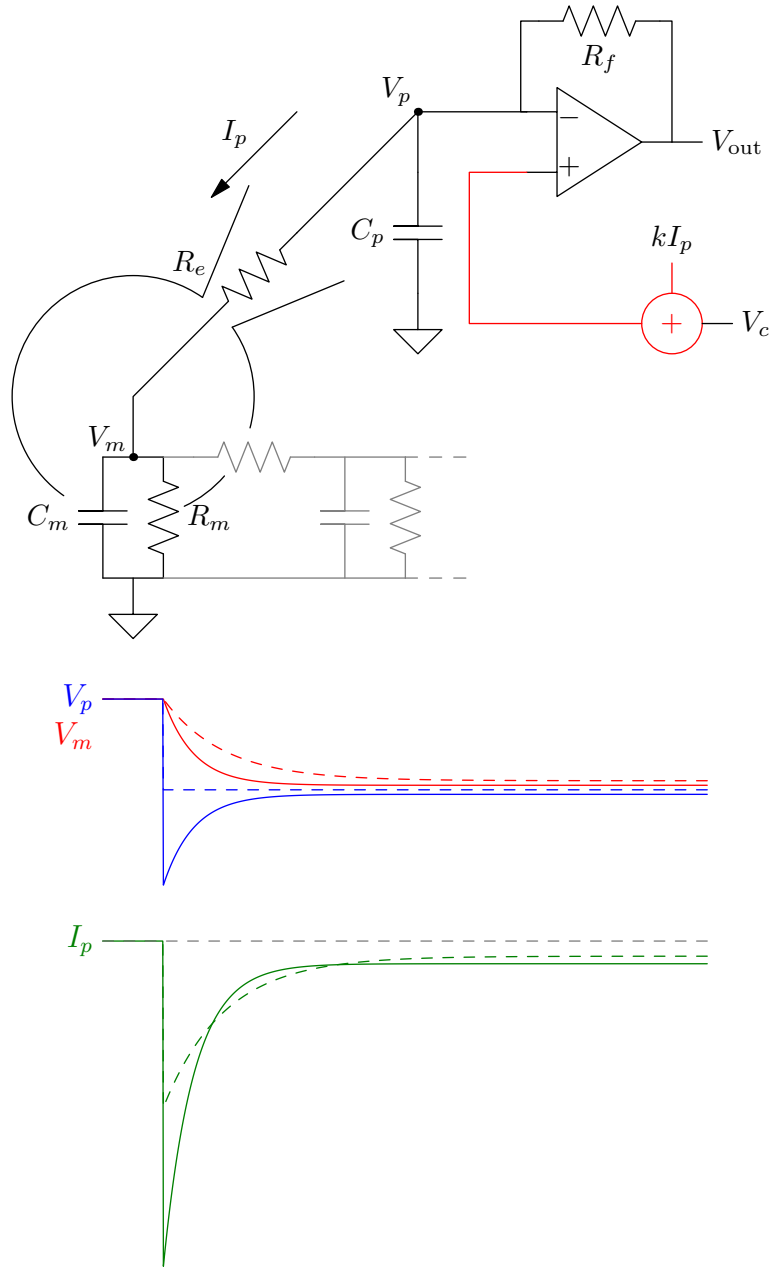
Figure 13: Compensation of the electrode resistance. A positive feedback (red) adds a signal proportional to the current (extracted from V_{out}) to the command potential. This tends to increase the pipette voltage when current flow and therefore the voltage error is greatest. The resulting current and voltage transients reflect the 'apparent' decrease of the electrode resistance. The dashed traces are those predicted without compensation, and the solid lines those for about $50\,\%$ compensation of R_e.
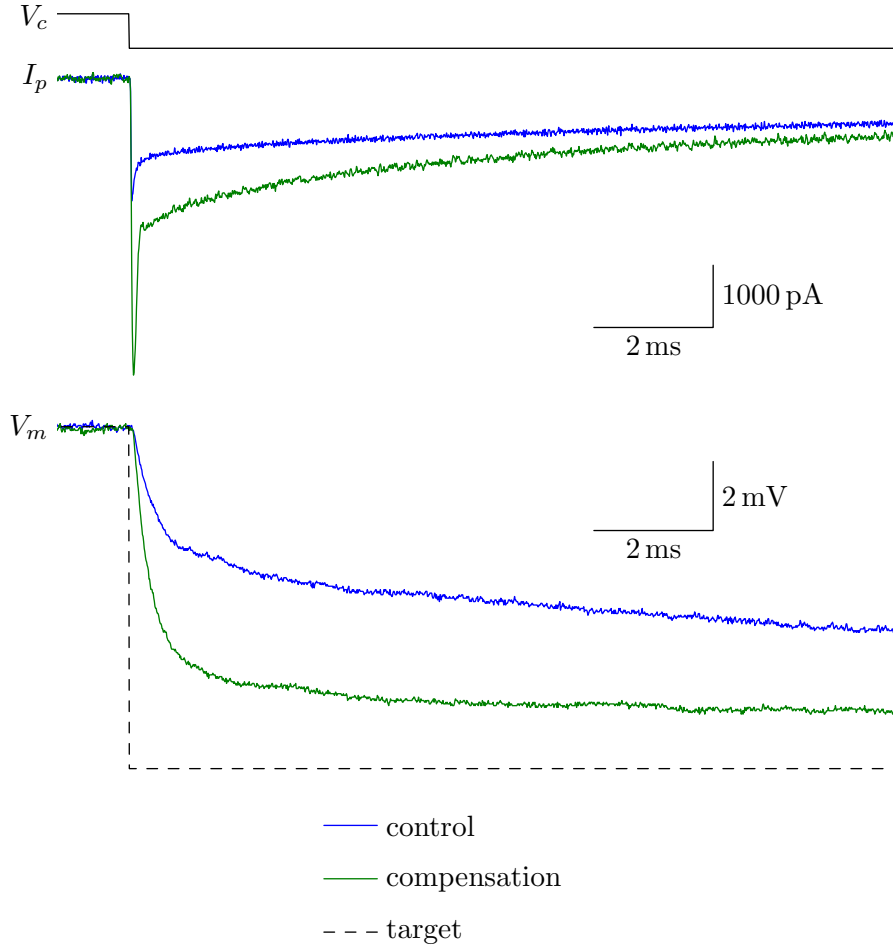
Figure 14: Effect of electrode resistance compensation on the capacity transients in the Purkinje cell. The nominal compensation was $90\,\%$ of R_e, which is unlikely to be exact. The clamp of the soma remains approximate at best (compare to the desired voltage represented by the dashed line). Nevertheless, the clamp of the somatic compartment is clearly improved by the compensation.

4   Not discussed . . .

Study of the following techniques and issues is left as an exercise for the reader:

- Supercharging

- Two-electrode voltage-clamp

- Switch-clamp

- Extracellular recording

- Extracellular stimulation

- Loose-cell-attached stimulation and recording

- Electrode non-linearity

- Filters
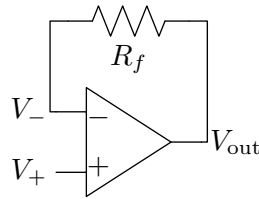
- $V_{\text{out}} = G(V_+ - V_-)$, where $G \approx \infty$

- Inputs (V_+ and V_-) draw no current

- Negative feedback ensures $V_- = V_+$

Figure 15: The ideal op-amp and the 'golden rules' for analysing op-amp circuits. Note that the negative feedback does not necessarily occur via a resistor.

5 Appendix 1: operational amplifiers

We first consider the 'ideal op-amp' (Fig. 15). In order to understand op-amp operation and when first analysing a circuit, one should start with the 'golden rules' listed in the figure. The output voltage of the op-amp is the difference between the voltages present at the non-inverting and inverting inputs multiplied by the gain, assumed here to be infinite, but in reality of the order of 10^5 or more. This gain is not particularly useful until the output is connected back to the inverting input in some way (a simple case is via a feedback resistor). In the presence of the negative feedback, the output tends to displace V_- towards V_+. The larger the gain, the closer the two become; at infinite gain we can assume that the the two are equal. Thus, negative feedback ensures that the voltage at the negative input equals that at the non-inverting input. The great usefulness of op-amps derives from the way this equality conditions the output voltage.

As an example of a real op-amp circuit, consider Fig. 16, which depicts an op-amp configuration called a *current follower*. It provides a voltage output proportional to the current flowing into the node at the inverting input (see analysis in the legend).

Op-amps are of course not ideal, so here is a short list of some of their deviations from the ideal:

- Saturation. Op-amps require power supplies and the output voltage therefore cannot exceed the supply voltage; usually the useful range is significantly less. Thus, for $\pm 15\,\text{V}$ supplies, the op-amp will work correctly between, say, $\pm 13\,\text{V}$. If the output is driven beyond this *linear operating range*, the ap-amp is said to saturate; it no longer

22
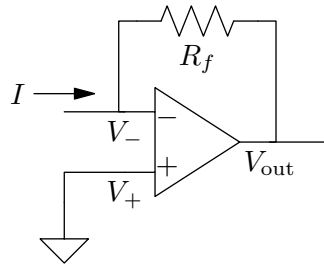
$$V_{\text{out}} = -IR_f$$

Figure 16: Example op-amp circuit—a current follower. By negative feedback, the voltage at the inverting input must be equal to that at the noninverting input. In order to prevent the voltage at the inverting input changing, the amplifier must immediately remove through the feedback resistor any current arriving at the inverting input (i.e. from the cell). As the inputs draw no current, I can only be removed from the input via the feedback resistor. Given that $V_- = V_+ = 0$, we can use Ohm's law to deduce that $0 - V_{\text{out}} = IR_f$. So the voltage output of the op-amp is proportional to the current flowing into the node at V_-.

works correctly. Some amplifiers work with high-voltage supplies to offer greater current passing ability.

- Oscillation. The negative feedback is not instantaneous. This delay and the high open-loop gain can induce a run-away *positive* feedback that causes the op-amp output to oscillate wildly. A good way to induce this behaviour is to attach a capacitor to ground at the inverting input, as this will further delay the arrival of the negative feedback.

- Bias. Real op-amps generate small positive or negative currents at their inputs and also have finite impedance. Nevertheless, op-amps designed with field-effect transistor (FET) inputs can have bias currents of a picoAmpère or less. Similarly, the voltage measurement is not perfect, but the errors, which are systematic, are also small ($\leq \text{mV}$) compared to other sources of error in electrophysiological recordings.

- Noise. Op-amps generate small amounts of noise that can nevertheless limit sensitive recordings (single channels). Both voltage and current noise are generated.

23

6   Changes to the text

2014/09/16 — The circuit for voltage-clamp capacity transient cancellation in §3.2 and Fig. 11 was made more realistic by the addition of a gain stage applied to V_c. Note that practical implementations may be designed with different circuits, but that shown will at least produce the desired cancellation.

2018/09/19 — Minor corrections thanks to Louise Thiaville and Miren Tamayo Elizalde.