

1 **A deep-learning strategy to identify cell types across species from high-density**
2 **extracellular recordings**

3
4 Maxime Beau^{*1}, David J. Herzfeld^{*2}, Francisco Naveros^{*,3,4}, Marie E. Hemelt^{*2}, Federico
5 D'Agostino^{*1}, Marlies Oostland^{*1,5}, Alvaro Sánchez-López^{*,3}, Young Yoon Chung¹, Michael
6 Maibach¹, Hannah N. Stabb¹, M. Gabriela Martínez Lopera¹, Agoston Lajko¹, Marie Zedler¹,
7 Shogo Ohmae³, Nathan J. Hall², Beverley A. Clark^{&1}, Dana Cohen^{&6}, Stephen G. Lisberger^{&2},
8 Dimitar Kostadinov^{&1,7}, Court Hull^{&2}, Michael Häusser^{&1}, Javier F. Medina^{&3}

9 ¹Wolfson Institute for Biomedical Research, University College London, London, UK

10 ²Department of Neurobiology, Duke University School of Medicine, Durham, NC, USA

11 ³Department of Neuroscience, Baylor College of Medicine, Houston, TX, USA

12 ⁴Department of Computer Engineering, Automation and Robotics, Research Centre for
13 Information and Communication Technologies, University of Granada, Granada, Spain

14 ⁵Swammerdam Institute for Life Sciences, University of Amsterdam, Amsterdam, the
15 Netherlands

16 ⁶The Leslie and Susan Gonda Multidisciplinary Brain Research Center, Bar-Ilan University,
17 Ramat-Gan, Israel

18 ⁷Centre for Developmental Neurobiology, King's College London, London, UK

19
20 *Equal contribution as first authors

21 &Equal contribution as senior authors

22 **Figures:** 8, total; 8, color

23 **Word counts:** Abstract, 150; Introduction, 538; Discussion, 1,597

24 **Keywords:** cerebellum, cerebellar cortex, cell-type identification, circuit mapping, Neuropixels,
25 variational autoencoder, machine learning

26 **Acknowledgements:** We thank Bonnie Bowell, Soyoon Chun, Wenjuan Kong, Caroline Reuter,
27 and Stefanie Tokiyama for technical support. Funding provided by: NIH grants R01-NS112917
28 (SGL, JFM, CH), K99-EY030528 (DJH), R01-NS092623 (SGL), R01-MH093727 (JFM); ERC
29 AdG 695709 (MH); Wellcome Trust PRF 201225/Z/16/Z and 224668/Z/21/Z (MH) and
30 225951/Z/22/Z (DK); EMBO ALTF 914-2015 (DK); European Union's Horizon 2020 research
31 and innovation programme under the Marie Skłodowska-Curie grant agreement No 844318
32 (MO) and No 891774 (FN); SYNCH project funded by the European Commission under the
33 H2020 FET Proactive program-Grant agreement ID 824162 (DC)

34 **Correspondence to:**

35 Javier F. Medina, PhD
36 Department of Neuroscience, S709
37 Baylor College of Medicine
38 Houston, TX 77030
39 Email: jfmedina@bcm.edu
40 Phone: (713) 798-8141

41 ***Abstract***

42 High-density probes allow electrophysiological recordings from many neurons simultaneously
43 across entire brain circuits but fail to determine each recorded neuron's cell type. Here, we
44 develop a strategy to identify cell types from extracellular recordings in awake animals, opening
45 avenues to unveil the computational roles of neurons with distinct functional, molecular, and
46 anatomical properties. We combine optogenetic activation and pharmacology using the
47 cerebellum as a testbed to generate a curated ground-truth library of electrophysiological
48 properties for Purkinje cells, molecular layer interneurons, Golgi cells, and mossy fibers. We
49 train a semi-supervised deep-learning classifier that predicts cell types with greater than 95%
50 accuracy based on waveform, discharge statistics, and layer of the recorded neuron. The
51 classifier's predictions agree with expert classification on recordings using different probes, in
52 different laboratories, from functionally distinct cerebellar regions, and across animal species.
53 Our approach provides a general blueprint for cell-type identification from extracellular
54 recordings across the brain.

55

56 The nervous system comprises many molecularly, anatomically, and physiologically defined cell
57 types¹⁻⁶. Powerful modern molecular techniques now have revealed multiple sub-types even
58 within known anatomical cell classes⁷⁻¹³. Identification of cell type at multiple levels will be
59 crucial to understand how the brain works and to develop selective, targeted therapeutics for
60 brain dysfunction. Therefore, it is crucial to develop strategies to determine cell type and to
61 cross-reference different formulations of cell type across levels of analysis^{3,5,6,12,14,15}.

62 With the advent of high-density multi-contact recording probes^{16,17}, it is now possible to record
63 from hundreds of neurons simultaneously and characterize their activity during specific,
64 quantified behaviors. Simultaneous large-scale electrophysiological recordings coupled with
65 cell-type identification *in vivo* would facilitate characterization of circuit-level processing in the
66 service of behavior. Yet, identification of cell type is a particularly difficult challenge for
67 extracellular recording technologies that cannot access the transcriptional or anatomical
68 properties of neurons¹⁸. Efforts to classify neurons based on specific features of their spike
69 waveform and firing statistics have not proven robust across laboratories^{19,20}. Moreover,
70 optogenetic approaches to cell-type identification²¹⁻²⁴ currently are routine only in mice and
71 bring the technical challenges of (i) off-target expression of opsins²⁵, (ii) disambiguating direct
72 responses versus those due to recurrent connectivity within circuits²⁶, and (iii) the ability to
73 target only one or two cell types at a time in a given preparation²⁷.

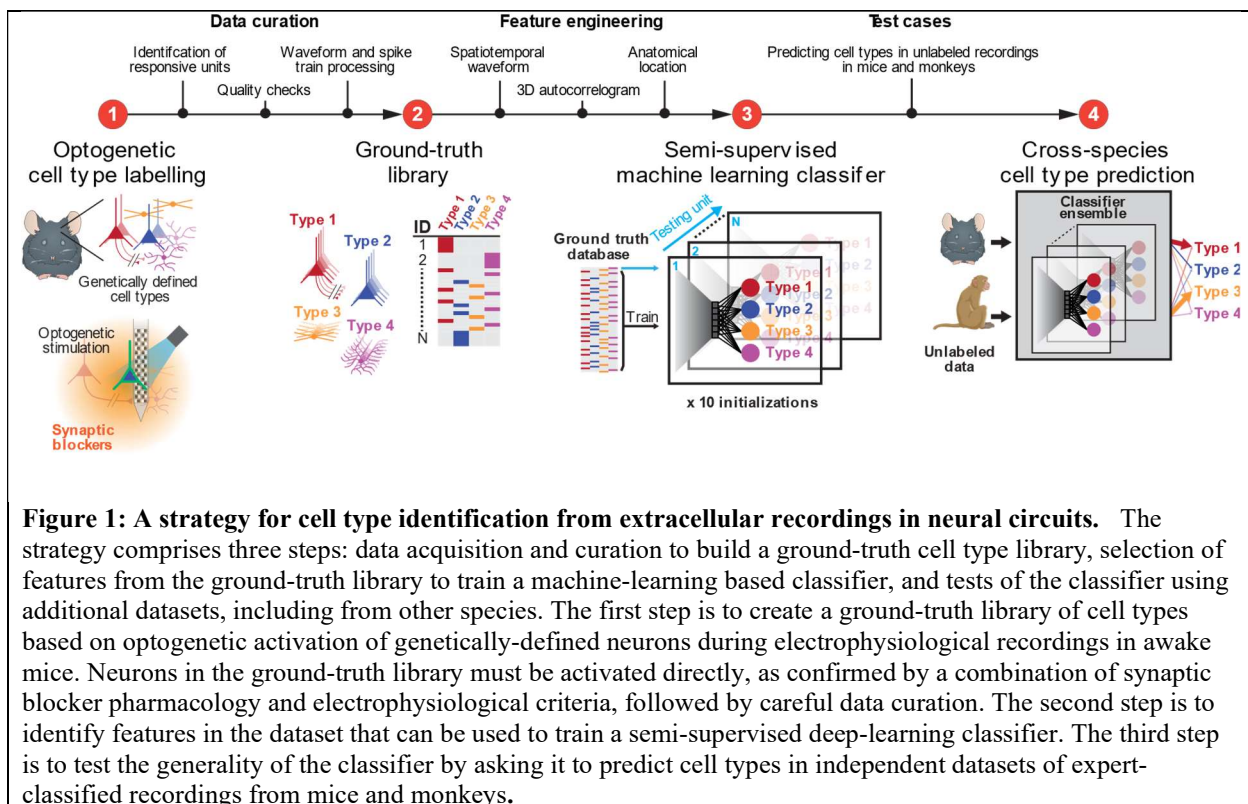
74 We assembled a collaboration of four laboratories with the single-minded goal of enabling cell-
75 type identification solely from extracellular recordings in awake animals by developing a
76 strategy that could scale across labs, probes, species, and in the future maybe across brain areas.
77 We chose to pioneer the strategy in the cerebellar cortex, which provides key advantages.
78 Specifically, the cerebellum has a crystalline architecture with well-defined neuronal
79 connectivity and a small number of anatomically-defined cell types^{1,28} that are consistent across
80 species^{29,30}, allowing direct comparison of recordings in monkeys, mice and other species. The
81 cerebellum has a range of neuron sizes from among the smallest and most densely packed
82 (granule cells) to the largest (Purkinje cells), allowing us to test the resolution of our recording
83 approaches. The cerebellum has many spontaneously firing neurons³¹⁻³³, some with high
84 spontaneous rates, allowing us to extract rigorous information about their electrophysiological
85 properties. Genetically-defined mouse Cre-lines are available for all major cell types in the
86 cerebellum³⁴⁻³⁸, allowing us to leverage optogenetic strategies for cell-type identification²¹.
87 Finally, the cerebellum has a long history of neurophysiological recording³⁹, allowing us to
88 reference our measurements and automated cell-type classifications against hard-won human
89 expertise. Strategies to solve the challenges of cell-type identification in such a testbed should
90 provide a roadmap for application to other structures, including the cerebral cortex, the
91 hippocampus, and the basal ganglia.

92 Our approach succeeded. We created a ground-truth library of identified cerebellar cell types
93 recorded in unanesthetized mice by combining rigorous spike sorting and unit curation with
94 identification through combined optogenetic activation and pharmacological synaptic blockade.
95 We demonstrate that a semi-supervised deep-learning classifier accurately predicts cell type for
96 the ground-truth library based on the waveform, discharge statistics, and anatomical layer of the
97 recording. Importantly, the classifier identifies cell type with high confidence in a high fraction
98 of expert-labeled cerebellar recordings from two different laboratories, in behaving mice and
99 macaque monkeys.

100 **Results**

101 General approach

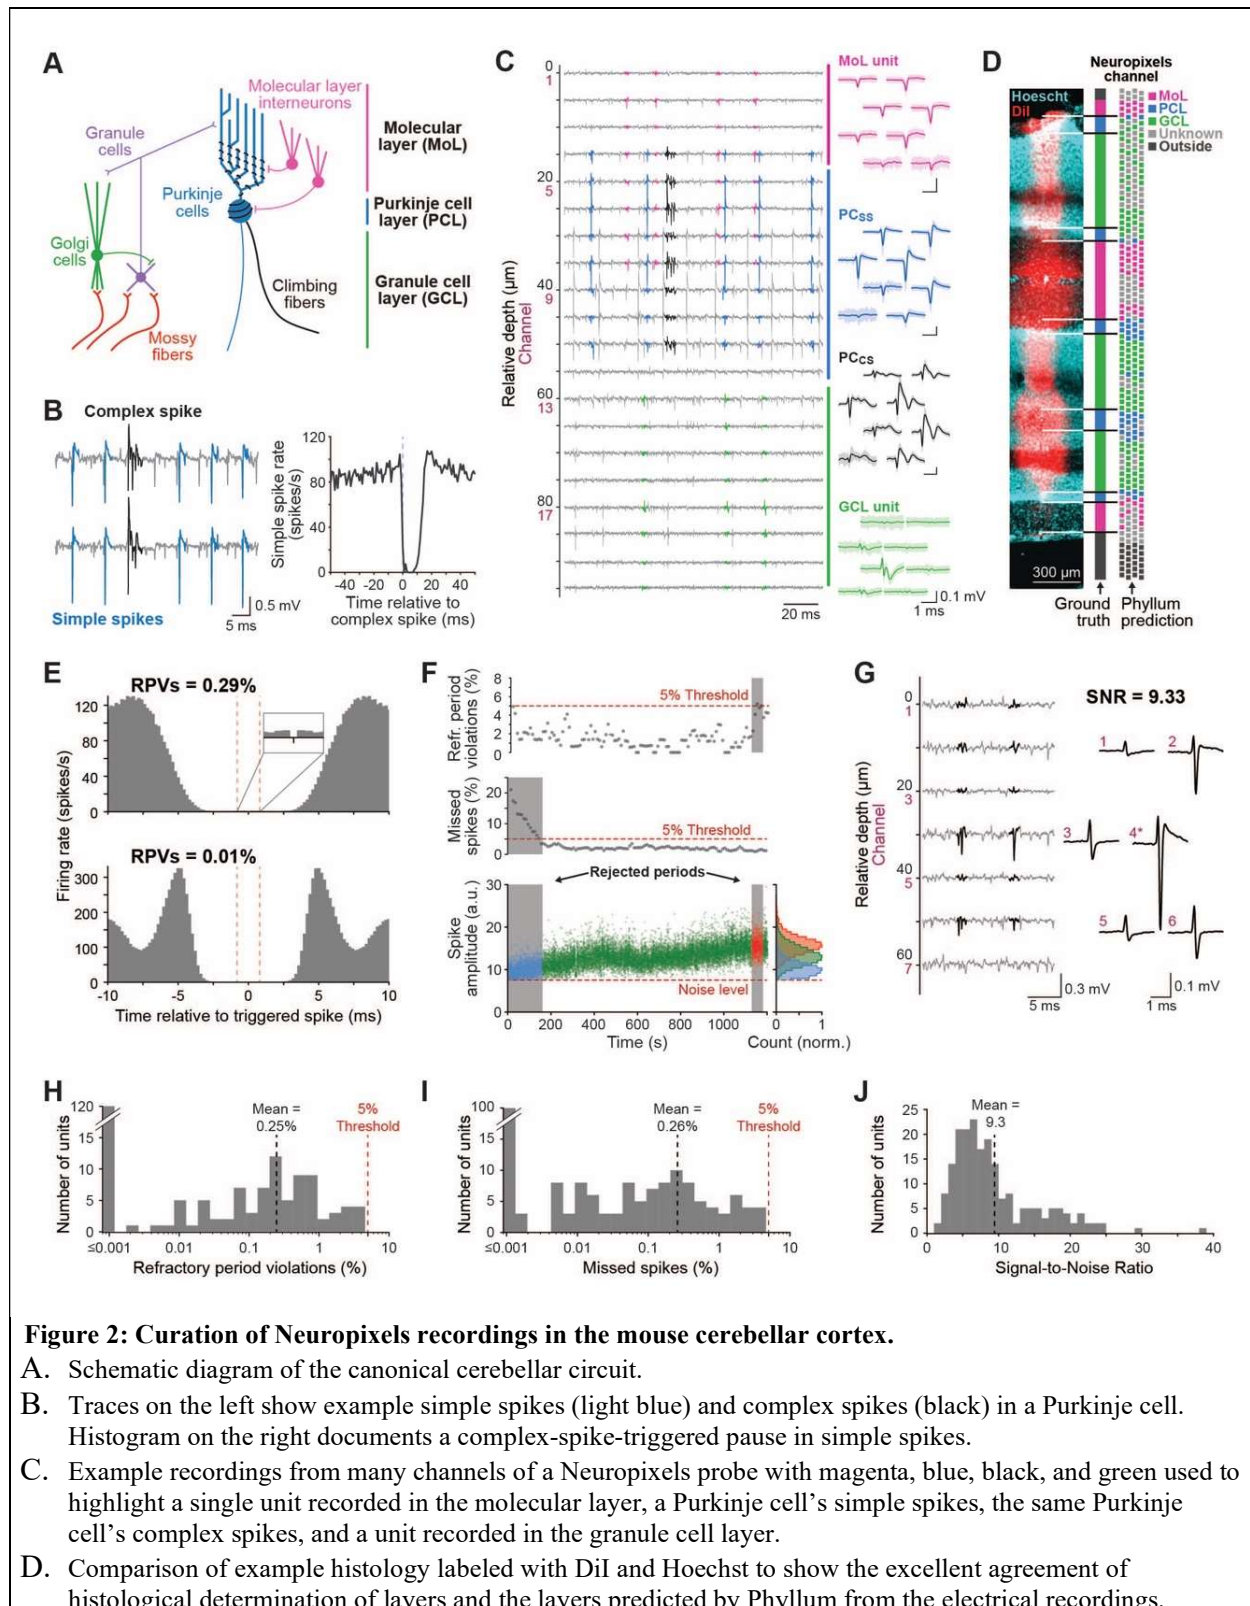
102 We start by creating a ground-truth library of extracellular recordings from neurons whose cell
103 type is established unequivocally. The analysis and classification pipeline (Figure 1) begins with
104 several data curation steps to ensure high-quality recordings and allow us to characterize with
105 high confidence each neuron's waveform, resting discharge properties, and anatomical location.
106 We use well-characterized mouse Cre-lines to identify cell types by optogenetic activation of
107 specific cell types in the presence of synaptic blockers. We then develop a semi-supervised deep-
108 learning classifier with performance evaluated with leave-one-out cross-validation. Finally, we
109 use the classifier to predict the cell types of an independent dataset of recordings made in mice
110 and macaque monkeys and we compare the performance of the classifier against cell-type
111 identification by human experts. Below, we develop the details of our strategy one step at a time.



112 Multi-contact probe recordings and data curation

113 We develop and deploy the general strategy for cell-type classification (Figure 1) in the
114 cerebellum, based on ground-truth recordings with Neuropixels probes in two laboratories
115 (Häusser and Hull labs). In the cerebellar cortex, morphologically distinct cell types reside in
116 different layers (Figure 2A). Purkinje cells comprise a monolayer and extend their planar
117 dendrites through the molecular layer. Molecular layer interneurons reside across the extent of
118 the molecular layer and include basket cells that innervate the Purkinje cell's soma and stellate
119 cells that innervate the Purkinje cell's dendrites. The granule cell layer includes mossy fiber
120 terminals, Golgi cells, and granule cells. Recent studies have identified other, less-common cell

121 types in the different layers³⁷, but we have elected to focus on the main cell types from the
 122 cerebellar circuit (Figure 2A).



Different colors on the Neuropixels schematic show: magenta, molecular layer; blue, Purkinje cell layer; green, granule cell layer; gray, unknown layer; black, outside cerebellar cortex.

- E. Autocorrelograms plotting a neuron's firing rate as a function of time from one of its own trigger spikes for two neurons with very few refractory period violations (RPVs). Note that the spike counts in the autocorrelograms have been divided by the width of the bin so that the y-axis is in spikes/s.
- F. Analysis of quality of isolation as a function of time during a recording session. From top to bottom graphs show the percentage of refractory period violations, the estimated percentage of missed spikes, and spike amplitude. Horizontal dashed lines show thresholds for acceptance. Gray regions show periods that were rejected from analysis. Blue, green, and red symbols indicate spikes that came from intervals that had too many missed spikes, acceptable isolation, and too many refractory period violations. Marginal histograms on the right show the distribution of spike amplitudes to document clipping at the noise level in the blue histogram that would be cause for rejection of a time interval.
- G. Example recording traces and spatial footprint of a representative recording with a signal-to-noise ratio (SNR) of 9.33, with the waveforms numbered according to their channel. Asterisk (*) denotes the channel with the largest peak-to-trough amplitude, used to compute the SNR.
- H. Distribution of percentage of refractory period violations across neurons accepted to the ground-truth library.
- I. Distribution of estimates of percentage of spikes that were missed across neurons accepted to the ground-truth library.
- J. Distribution of signal-to-noise ratios on the channel with the largest-amplitude waveform across neurons accepted into the ground-truth library.

123

124 Purkinje cells are the one cell type in the cerebellum that allows ground-truth identification from
125 its extracellular electrical signature. Purkinje cells show two types of action potentials (Figure
126 2B, left panel): “simple spikes” that fire at high rates and “complex spikes” driven directly by
127 climbing fiber input^{40–42}. Complex spikes occur only at ~1 Hz, and trigger a characteristic 10-50
128 ms pause in simple spikes⁴³. Thus, Purkinje cells can be identified unequivocally, and admitted
129 into the ground-truth library, if they show a pause in a complex-spike-triggered histogram of
130 simple-spike firing (Figure 2B, right panel).

131 Recordings with Neuropixels probes detect neural activity on many of the 384 channels and
132 spike sorting yields many units including non-Purkinje cells. The magenta waveforms in Figure
133 2C arise from a neuron in the molecular layer that would be a candidate to be a molecular layer
134 interneuron. The green waveforms come from a neuron recorded in the granule cell layer that
135 could be a mossy fiber, a Golgi cell, or a granule cell. The blue and black waveforms are the
136 simple spikes and complex spikes of an identified Purkinje cell.

137 Given that the soma of each cell type resides in one of the three different layers of the cerebellar
138 cortex, the first step in our analysis pipeline was an objective procedure to identify the layer of
139 each recording. The cerebellum is a foliated structure so that a single penetration with a
140 Neuropixels probe usually records from neurons in multiple repetitions of each of the 3 layers of
141 the cerebellar cortex. For example, the recording trajectory documented with DiI staining in
142 Figure 2D crossed 3 molecular layers, 5 Purkinje cell layers, and 3 granule cell layers. We
143 assigned each channel to a layer using Phyllyum, a Phy plugin that analyzes recordings across the
144 channels on a probe to infer the layer recorded by each channel (see *Methods*).

145 The layer structure inferred by Phyllyum agreed well with histological data based on simultaneous
146 DiI and cell body staining (Figure 2D). We validated Phyllyum across 21 histologically confirmed
147 penetrations and found that its conclusions agree with the histology at 99, 95, and 98% of 776,
148 367, and 1140 recording sites respectively in the molecular, Purkinje cell, and granule cell layers.

149 The layer assignments from Phylum were also consonant with the finding from single electrodes
150 of (i) clear complex-spike activity and well-isolated simple spikes in the Purkinje cell layer, (ii)
151 relative silence and abundant dendritic Purkinje cell complex spikes^{44,45} in the molecular layer,
152 and (iii) a jungle of high-intensity activity with many units in the granule cell layer.

153 We next ensured that each unit we admitted for further analysis was a well-isolated single neuron
154 with credible waveform and resting discharge properties, two of the three features we ultimately
155 would use, along with layer, to classify units. We manually curated the output from Kilosort2
156 with Phy and subsequently performed automated quality checks to ensure the quality of isolation
157 and the veracity of the waveforms and resting discharge statistics of neurons that would become
158 part of our ground-truth library. We strove to ensure that we neither missed many spikes from
159 the neuron under study nor included electrical artifacts or spikes from neighboring neurons.

160 ● We analyzed the refractory periods from each isolated neuron to assess the level of
161 contamination from other neurons or noise⁴⁶. The examples of autocorrelograms in Figure 2E
162 have vanishingly small numbers of refractory period violations and respectively represent the
163 mean (0.25%) and median (0.01%) in our dataset. We rejected from the ground-truth library
164 autocorrelograms with greater than 5% period violations (Figure 2F, red symbols and
165 histogram). Almost all accepted neurons had fewer than 1% refractory period violations with
166 a mean of 0.25% (Figure 2H).

167 ● We estimated the number of missed spikes by fitting the spike amplitude distribution with a
168 Gaussian function and quantifying the fraction of the area under the curve that was clipped at
169 noise threshold^{47,48} (Figure 2F). We estimated that few spikes were missed if the distribution
170 of spike amplitudes for the entire recording was continuous and not clipped at noise
171 threshold. In Figure 2F, we estimated that more than 5% of spikes were missed in the first
172 ~150 s of the recording (blue symbols and histogram), causing us to exclude those intervals
173 from further analysis. Among the recordings we accepted, the percentage of missed spikes
174 averaged 0.26% and almost all neurons showed fewer than 1% missed spikes (Figure 2I).

175 The requirement for few violations of the refractory period and small numbers of missed spikes
176 ensured that the units we accepted had high signal-to-noise ratios. The mean signal-to-noise ratio
177 in our accepted sample, measured as the signal-to-noise ratio on the channel with the largest unit
178 potential, was 9.3, almost identical to that of the example recording in Figure 2G. Over 90% of
179 the neurons had signal-to-noise ratios larger than 4 (Figure 2J).

180 We identified and resolved two other issues that impaired consistent and reliable estimates of
181 waveform. The *first* issue is related to an on-board hardware high-pass filter on Neuropixels
182 probes. The filter distorts the shape of waveforms and therefore hinders comparison across
183 recordings made when the filter was on versus off. We used the technical description of the
184 analog filter in the Neuropixels documentation to apply an equivalent digital filter to data
185 recorded with the filter disengaged or using other probes in monkeys (Supplementary Figure 1).
186 The *second* issue concerns temporal alignment of individual spikes, which is unreliable in
187 Kilosort's output when the signal-to-noise ratio is low to medium or when a unit drifts across
188 channels during a recording session. We resolve the alignment issue with an iterative procedure
189 we called "drift-shift-matching" to minimize waveform distortion from the averaging of
190 individual action potentials (see Supplementary Figure 1 and *Methods*).

191 *Combination of optogenetics and pharmacology for ground-truth cell-type identification*

192 For optogenetic activation, we combined genetic and viral approaches to cause expression of an
193 opsin in a specific cell type, thereby allowing these neurons to be selectively activated by light
194 and identified by photostimulation²¹. We added synaptic blockers to our experimental
195 preparations to ensure that light activates an opsin-expressing neuron directly, and not indirectly
196 by optogenetic activation of its pre-synaptic inputs. The standard criterion of short-latency
197 activation by optogenetic stimulation⁴⁹ (e.g., <10 ms) is inadequate on its own. We found ample
198 examples of short-latency responses (some <5 ms) that disappeared with synaptic blockade²¹.

199 We introduced Neuropixels probes into the lateral cerebellar cortex or the vermis of mice
200 expressing opsins (usually Channelrhodopsin-2, ChR2) in specific cell types and allowed the
201 probe to settle at a location where we recorded activity across much of its length. As illustrated
202 in Figure 3A, we performed experiments with optogenetic stimulation in multiple phases. In a
203 *baseline phase*, we recorded spontaneous activity. In a *control phase*, we applied light externally
204 to the cerebellum to activate opsins in the cell types that expressed them. We also introduced a
205 tapered optic fiber that ran alongside the recording probe in some experiments, to deliver light in
206 closer proximity to cells that expressed opsins. In an *infusion phase*, we continued to deliver
207 light to the cerebellum while we added synaptic blockers (see *Methods*) to the surface of the
208 cerebellum. In the *blockade phase*, when the synaptic blockers had permeated well into the
209 tissue, we assayed neurons for direct responses to optogenetic stimulation. The approach of
210 applying synaptic blockers on the surface of the cerebellum, instead of trying to inject them deep
211 into the tissue, has the advantage of preserving the integrity of the tissue but the disadvantage of
212 relatively slow diffusion.

213 We accepted neurons as activated directly by photostimulation only if we had strong evidence
214 that they were within the locus of successful synaptic blockade and they continued to have
215 reliable, short-latency responses to light. To determine whether synaptic blockade was effective
216 at a given recording depth, we evaluated recordings at or below that depth and looked for any of
217 the indications in the top row of Figure 3B:

- 218 1. Loss during the *blockade phase* of responses to optogenetic stimulation present in the *control*
219 *phase* (Figure 3B, top row, left). Neurons were excluded from the ground-truth library if they
220 retained their response in the *blockade phase* but were outside the region of synaptic
221 blockade (Figure 3B, bottom row, left).
- 222 2. Putative mossy fibers with loss of negative afterwaves in the *blockade phase* (Figure 3B, top
223 row, middle). Here, we rest partly on prior evidence that the negative afterwave is a post-
224 synaptic response of granule cells^{50,51} and that analogous negative afterwaves have been
225 shown to correspond to post-synaptic responses in other brain regions⁵². Also, our finding of
226 the effect of synaptic blockade on the negative afterwave in recordings from single putative
227 mossy fibers provides the strongest evidence to date that the negative afterwave represents
228 post-synaptic depolarization.
- 229 3. Substantial changes in a neuron's autocorrelogram or "coefficient of variation 2" (CV2),
230 usually due to an increase in regularity caused by a shift from synaptically- and intrinsically-
231 driven spiking to purely intrinsically-generated spiking³¹ (Figure 3B, top row, right).

232 To illustrate our strategy, we provide detailed examples from one experiment in a transgenic
 233 mouse line that expresses ChR2 in mossy fibers^{37,53} (Thy1-ChR2 line 18). Here, some neurons
 234 lost their responses to optogenetic activation with synaptic blockade (Figure 3C), while others
 235 recorded nearby retained their responses (Figure 3D). We evaluated the effect of synaptic
 236 blockade along the electrode penetration that yielded these two units (Figure 3E, region above
 237 the dashed line on the Neuropixels schematic) to confirm that the neural responses in Figure 3D
 238 were activated directly by optogenetic stimulation. At sites near and deeper than the neuron in D,
 239 neurons lost their responses with synaptic blockade, indicating that they were within the region
 240 of successful blockade. Still deeper in the penetration, we recorded two putative mossy fiber
 241 waveforms that retained their negative afterwave with synaptic blockade, indicating that they
 242 were outside the region of successful blockade. Finally, the extracellular waveforms of the
 243 activated neuron were constant across the entire experiment (Figure 3F), indicating that we had a
 244 stable recording. Therefore, we concluded that the neuron in Figure 3D was an optogenetically-
 245 activated neuron, namely a mossy fiber.

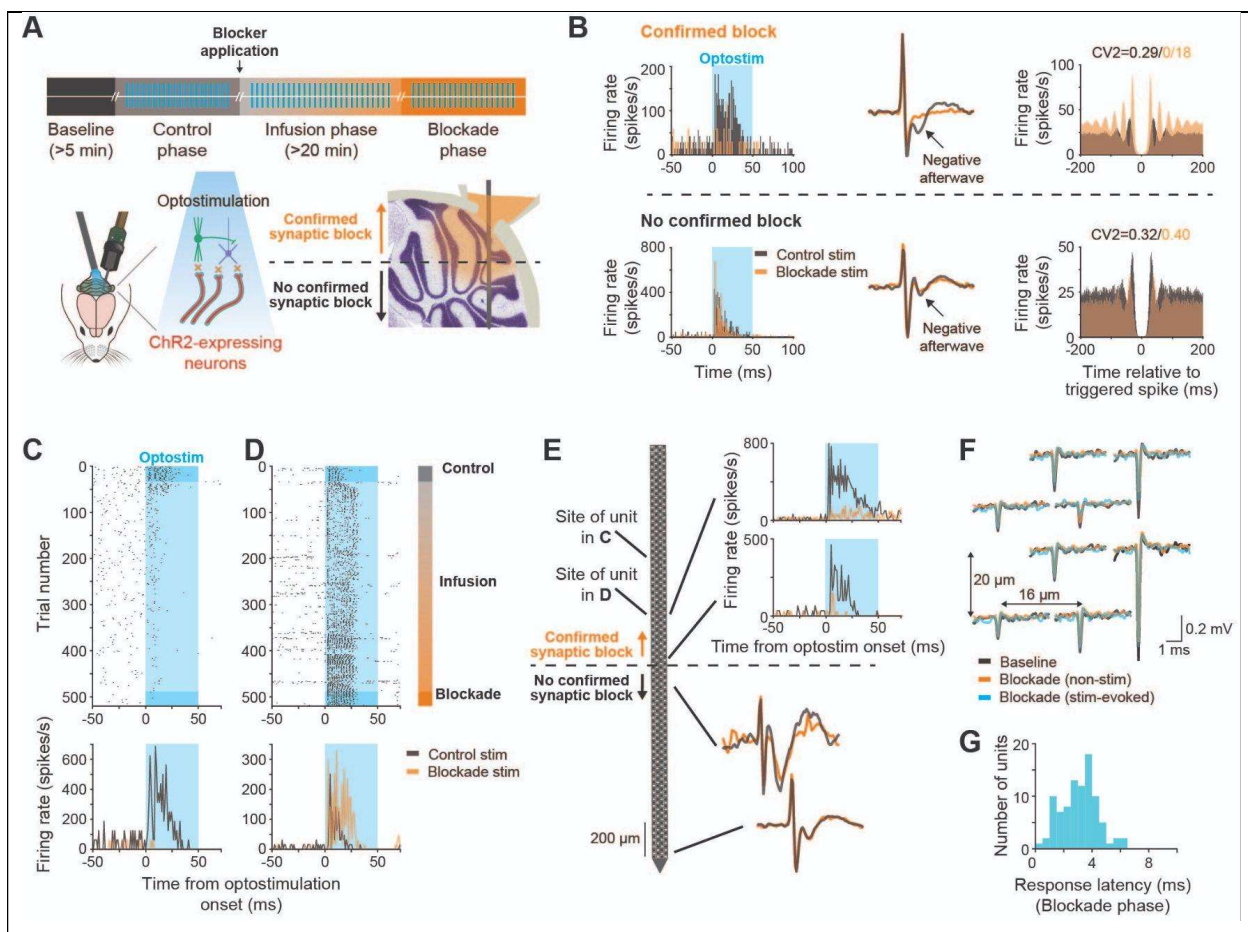


Figure 3: Strategy for ground-truth identification of cell type.

- Schematic showing the sequential phases in an experiment designed to test for optogenetic activation in the presence of synaptic blockers.
- Examples of the results used to verify the region of synaptic blockade. Examples above versus below the horizontal dashed line were taken as evidence for versus against blockade at that site. From left to right, we assayed the effect of blockade on the response to optogenetic stimulation, the negative afterwave of a

putative mossy fiber waveform, and the discharge statistics defined by autocorrelograms and the value of CV2.

- C. Raster and peri-stimulus time histogram for a neuron that lost its response to optogenetic stimulation with synaptic blockade. Trial numbers on the y-axis align with the cartoon showing the periods in the experiment to the right of D. Black versus orange histograms show responses before versus during synaptic blockade. Blue shading indicates the time of photostimulation.
- D. Same as C except for a neuron that retained its response to optogenetic stimulation during synaptic blockade.
- E. Example of how we determined whether the recordings in C and D were within the region of synaptic blockade. The cartoon schematizes a Neuropixels probe, the top histograms on the right show sites that were within the region of blockade because they lost their responses to optogenetic stimulation, and the lower waveforms show mossy fibers that were below the region of blockade because they retained their negative afterwaves.
- F. Spatial footprint of the neuron in D. Black, orange, and blue traces show the similarity of the waveforms recorded during the baseline period, during synaptic blockade without optogenetic stimulation, and during synaptic blockade with optogenetic stimulation.
- G. Distribution of neural response latencies to optogenetic stimulation of directly-activated neurons in presence of synaptic blockade.

246

247 Ground-truth neurons generally responded to optogenetic stimulation with latencies shorter than
248 5 ms (Figure 3G). However, we recommend against using latency as the sole criterion to accept
249 or reject neurons as optogenetically-activated²¹. Latency is strongly dependent on illumination
250 intensity and the density of opsin expression. We frequently observed short-latency responses to
251 optogenetic stimulation that disappeared with synaptic blockade, for example the neuron in
252 Figure 3C where the latency was less than 3 ms. Also, the neuron in the top row of Figure 3B
253 responded to optogenetic stimulation with a latency shorter than 5 ms, but lost its response with
254 synaptic blockade, indicating that it was driven by synaptic activation rather than direct
255 optogenetic stimulation.

256 Strategy to mitigate off-target expression in transgenic mouse lines

257 To varying degrees, off-target expression is a common feature of transgenic mouse lines. Often,
258 there is no ‘clean’ line available where expression is limited to a given cell-type of interest. Our
259 strategy to obtain ground-truth cell-type identification despite off-target expression was to (i)
260 characterize the anatomical specificity of opsin expression for all mouse lines under study and
261 (ii) combine identification of the recording layer based on Phyllum with optogenetic activation in
262 the confirmed presence of synaptic blockers to establish cell-type unambiguously.

263 The problem of off-target expression was most pronounced in the GlyT2-Cre line used
264 previously to image activity in Golgi cells³⁵. The GlyT2-Cre line has substantial off-target
265 expression in molecular layer interneurons and, very occasionally, Purkinje cells (Figure 4A,
266 Supplementary Figure 2): the relative density of molecular layer interneurons was higher than
267 that of Golgi cells (Figure 4B, 79 vs. 20%). Accordingly, we recorded neurons directly
268 responsive to optogenetic stimulation in both the granule cell layer and the molecular layer
269 (Figure 4C). We used Phyllum to identify the recording layers and labeled units in the granule
270 cell layer that were directly activated by optogenetic stimulation as Golgi cells (Figure 4C, green
271 PSTH). We labeled activated units in the molecular layer as molecular layer interneurons (Figure
272 4C, magenta PSTH).

273

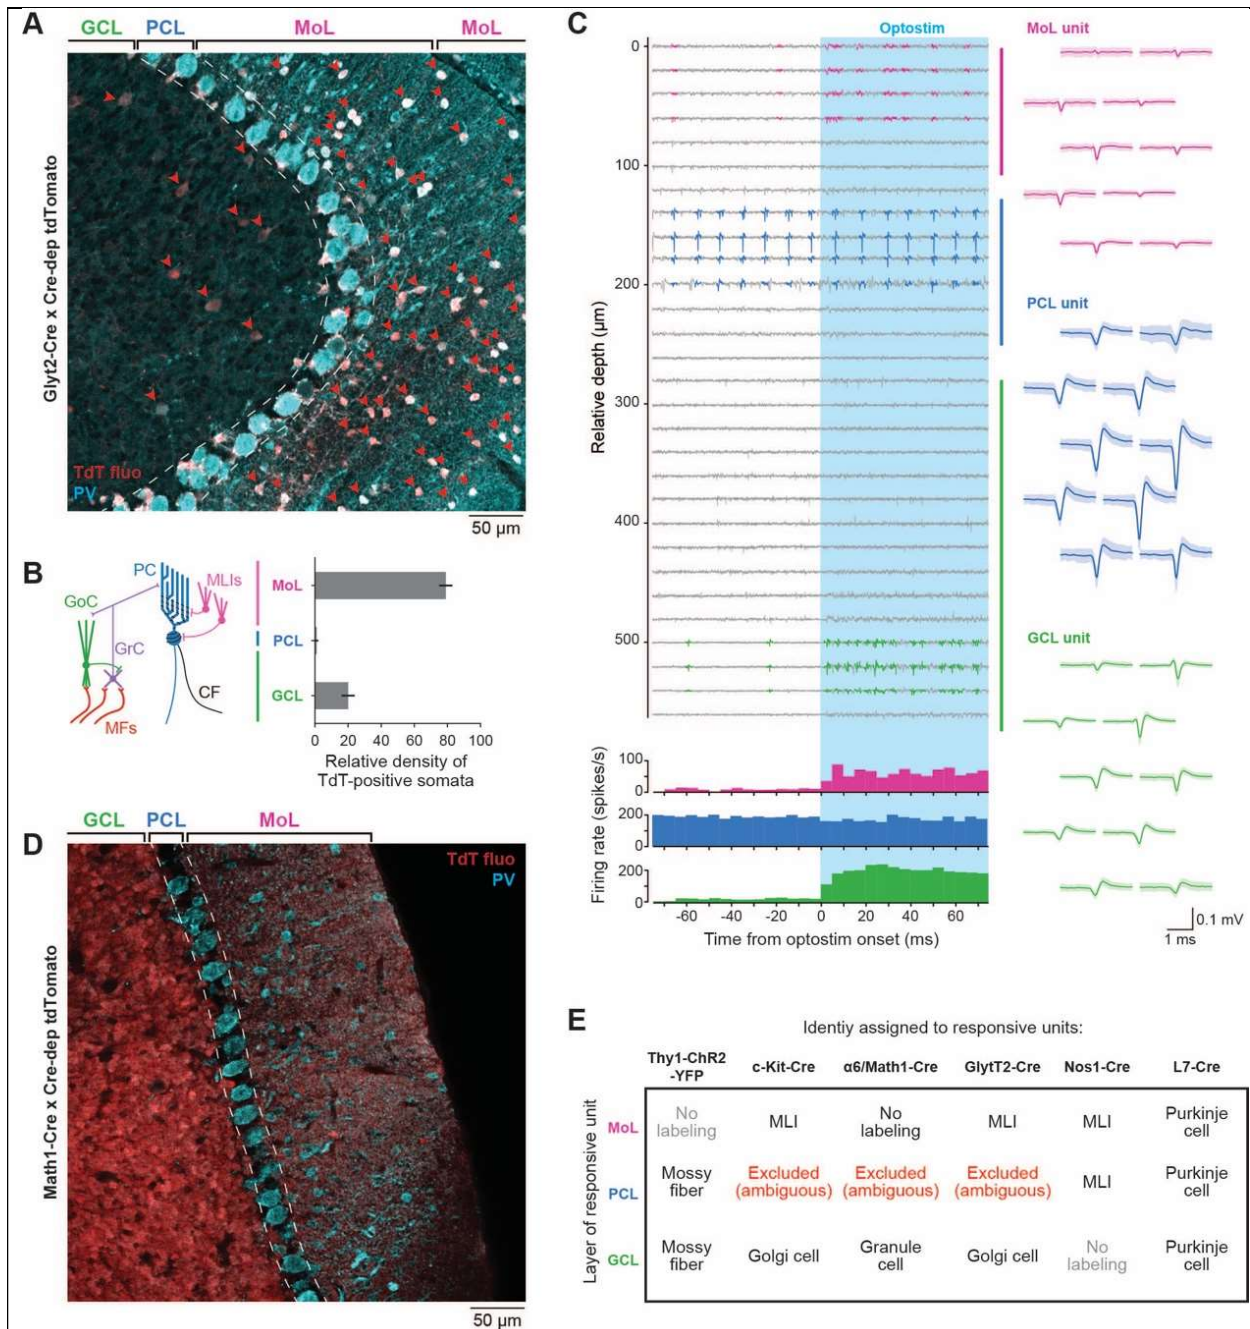


Figure 4: Analysis and mitigation of off-target expression in mouse optogenetic lines.

- A. Double stained section of cerebellum in the GlyT2-Cre line showing expression in both Golgi cells in the granule cell layer and molecular layer interneurons. Red arrows point to cells that express Td-Tomato. Blue cells express parvalbumin (PV). MoL, molecular layer; PCL, Purkinje cell layer; GCL, granule cell layer.
- B. Cartoon of cerebellar circuit and histogram showing density of TdT-positive somata in each of the three layers in a GlyT2-Cre mouse: GoC, Golgi cell; GrC, granule cell; PC, Purkinje cell; MLI, molecular layer interneuron; CF, climbing fiber; MF, mossy fiber.
- C. Representative recordings from a Neuropixels probe using optogenetics to activate neurons that express opsins in the GlyT2 line. Magenta, blue, and green waveforms on the right show the spatial footprint of neurons in the MoL, PCL, and GCL. Histograms below the voltage traces show that both the MoL and GCL layer neurons were activated by optogenetic stimulation at the time indicated by the blue shading.

D. Same as A, but for the Math1-Cre line.

E. Table outlines how we used layer information to disambiguate cell types despite some off-target expression in certain Cre-lines.

274

275 Other mouse lines also showed some off-target expression. For example, the Math1-Cre line
276 used to label granule cells was generally specific (Figure 4D) but exhibited rare labeling of
277 Purkinje cells. The c-kit-Cre line also labeled a small number of Golgi cells³⁴, and the Nos1-Cre
278 line exhibited occasional labeling of non-neuronal cells in addition to molecular layer
279 interneurons (Supplementary Figure 2). By contrast, other lines we used were cleaner, such as
280 the Thy1-ChR2-YFP line 18 and Pcp2-Cre lines used to label mossy fibers and Purkinje cells,
281 respectively (Supplementary Figure 2). Crucially, we did not observe multiple labeled cell-types
282 within a single cerebellar layer in any of our lines. Thus, the combination of an identified layer
283 with direct optogenetic activation (Figure 4E) allowed us to disambiguate cell type for all
284 experiments.

285 *The ground-truth library*

286 Across 188 Neuropixels recordings in two laboratories, we recorded a total of 3652 neurons that
287 survived the spike-sorting and curation pipeline (Figure 5A). Of these, 562 exhibited a response
288 to optogenetic stimulation but only 97 passed our rigorous criteria for direct rather than synaptic
289 activation based on reliable, short-latency responses in the presence of synaptic blockade. We
290 added the simple spikes and complex spikes of 62 Purkinje cells identified by a complex-spike
291 triggered pause in simple spikes. We removed 6 units recorded in Cre-lines with off-target
292 expression where the layer of the recording was ambiguous, and 13 units that, on final closer
293 inspection, did not have sufficiently long baseline periods due to intervals that violated our
294 missed/extra spikes criteria (Figure 2E, F). The resulting ground-truth library contained 202
295 units: 69 Purkinje cell simple spikes, 58 Purkinje cell complex spikes recorded at the same time
296 as the simple spikes, 27 molecular layer interneurons, 18 Golgi cells, and 30 mossy fibers. For
297 comparison with previous reports^{20,54,55}, Supplementary Figure 4 provides the
298 electrophysiological signatures of different cell types in our ground-truth library of cell types
299 using a range of metrics.

300 In an attempt to obtain ground-truth recordings from granule cells, we made 82 recordings with
301 Neuropixels probes in mice with the Math1-Cre or BAC α 6Cre-C lines either crossed to Cre-
302 dependent ChR2 (Ai32) or injected with AAV to confer ChR2 expression (see *Methods*). We did
303 record multiple unit activity that was responsive to photostimulation in the region of confirmed
304 synaptic blockade, but almost all putative single units found by KiloSort failed one or more of
305 our criteria for good isolation (Figure 2). After careful curation, we retained zero likely granule
306 cells from 32 recordings in the Hull lab and at most 3 from 50 recordings in the Häusser lab. The
307 yield of fewer than 0.04 granule cells per recording was significantly lower than for the other cell
308 types (Figure 5B). We conclude that it is challenging to record from granule cells using the
309 current generation of Neuropixels probes; our sample is far too small to include them in the
310 classifier we will develop next. A combination of factors may contribute to the inability to record
311 regularly from granule cells: their comparatively small size^{56,57}, the likelihood that they generate
312 a spatially-restricted closed-field extracellular potential, and the low electrode impedance⁵⁸ of
313 Neuropixels¹⁶ (150 kOhms).

314

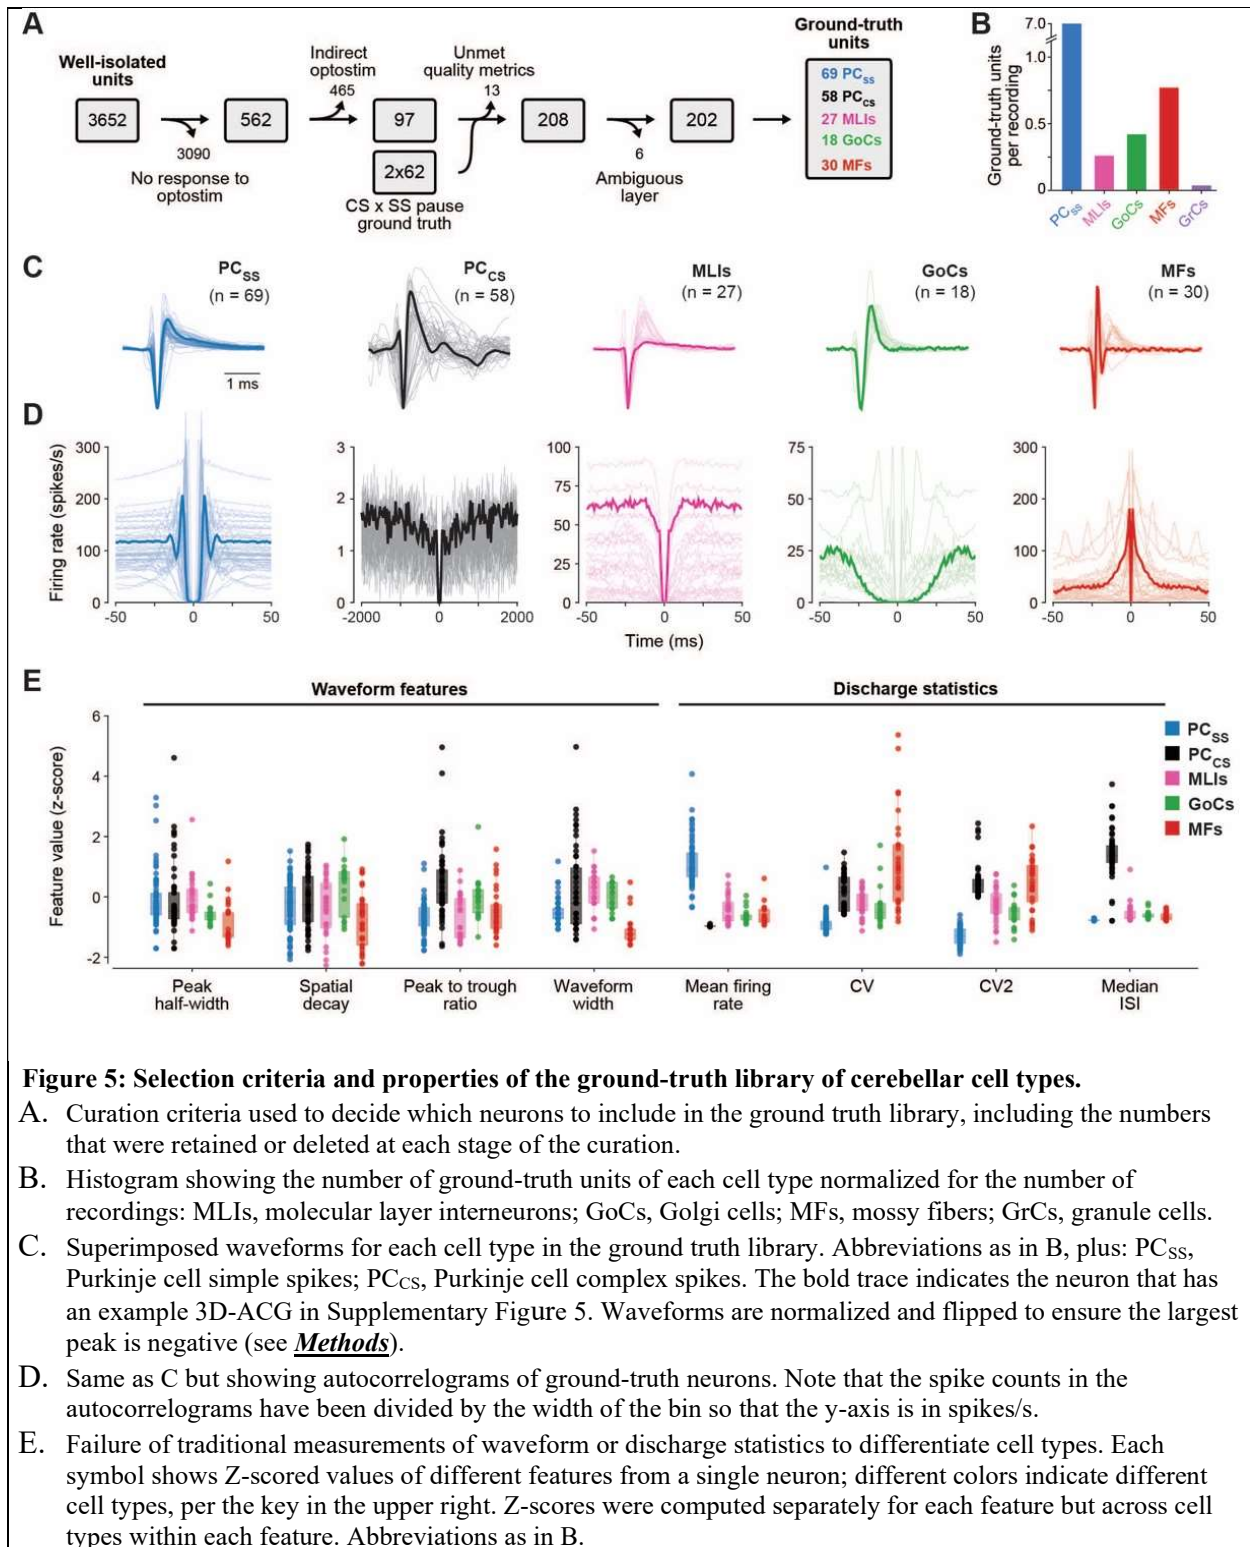


Figure 5: Selection criteria and properties of the ground-truth library of cerebellar cell types.

- Curation criteria used to decide which neurons to include in the ground truth library, including the numbers that were retained or deleted at each stage of the curation.
- Histogram showing the number of ground-truth units of each cell type normalized for the number of recordings: MLIs, molecular layer interneurons; GoCs, Golgi cells; MFs, mossy fibers; GrCs, granule cells.
- Superimposed waveforms for each cell type in the ground truth library. Abbreviations as in B, plus: PC_{SS}, Purkinje cell simple spikes; PC_{CS}, Purkinje cell complex spikes. The bold trace indicates the neuron that has an example 3D-ACG in Supplementary Figure 5. Waveforms are normalized and flipped to ensure the largest peak is negative (see [Methods](#)).
- Same as C but showing autocorrelograms of ground-truth neurons. Note that the spike counts in the autocorrelograms have been divided by the width of the bin so that the y-axis is in spikes/s.
- Failure of traditional measurements of waveform or discharge statistics to differentiate cell types. Each symbol shows Z-scored values of different features from a single neuron; different colors indicate different cell types, per the key in the upper right. Z-scores were computed separately for each feature but across cell types within each feature. Abbreviations as in B.

315 Armed with a ground truth dataset, the next challenge was to develop an accurate classification
 316 method based on consistent differences in electrophysiological features across cell types⁵⁹. To
 317 maximize the success of classification, we strove to use both waveform^{20,60,61} and discharge
 318 statistics⁶²⁻⁶⁴ as features for cell-type classification.

319 Waveform: We anticipated that the different cellular properties and morphology of different
320 cell types would lead to different waveforms^{60,61,65,66}. Patch clamp recordings *in vitro*
321 confirmed that biophysical differences across neuron classes are manifest as consistent
322 variations in the shape of the waveform (Supplementary Figure 3). Yet, as shown by the
323 comparison of Figure 5B with Supplementary Figure 3, waveforms are much more variable
324 in extracellular recordings *in vivo* than *in vitro*, and waveforms alone do not cleanly
325 distinguish cell type.

326 Discharge statistics: It is common for different cell types to have different discharge
327 statistics throughout the brain^{67,68} and the same is true in the cerebellum of anesthetized
328 animals⁶²⁻⁶⁴. In awake animals, discharge statistics are likely to vary across cerebellar regions
329 and to depend on the specific behavior or sensory input^{20,69}. Therefore, a robust
330 classification strategy should harness additional information that normalizes for the factors
331 that contribute to variation in awake animals.

332 Cell-type identification from a semi-supervised deep-learning classifier

333 Our deep-learning classifier strategy takes advantage of the rich information contained in the
334 diverse waveforms and firing statistics in the ground-truth library (Figure 5C, D), along with the
335 layer information that also provides information about cell type. Rather than using a potentially
336 biased set of investigator-chosen measurements from waveform and firing statistics, we chose to
337 use raw data because they (i) contain richer information, (ii) provide unbiased inputs for cell-
338 type identification, and (iii) are likely to generalize across regions, tasks, and species. Further,
339 Figure 5E and Supplementary Figure 4 reveal that it is difficult to guess which specific measures
340 of waveform and firing statistics would be most informative to successfully distinguish cell types
341 in awake animals.

342 We represent spike waveforms as the full time-course of the average, drift- and shift-corrected
343 waveform on the channel with the largest signal. We represent firing statistics as
344 autocorrelograms (ACGs) that assess the firing rate of a neuron as a function of time relative to
345 each spike. Because the traditional two-dimensional-ACG (2D-ACG) is subject to artifacts when
346 neural firing rate varies across a recording session or in relation to behavior, we developed
347 “three-dimensional autocorrelograms” (3D-ACGs) that normalize for firing rate (Figure 6A, see
348 **Methods**). Supplementary Figure 5 shows example 3D-ACGs for each ground-truth cell type.
349 We developed a strategy to avoid the potential issue of overfitting that is inherent in a deep-
350 learning classifier given the high dimensionality of the waveforms and 3D-ACGs and the
351 relatively small number of training examples in the ground-truth library. To address the
352 mismatch of training data relative to input dimensionality, we used an unsupervised dimension-
353 reduction technique⁷⁰ that took advantage of 3090 unlabeled single units recorded with
354 Neuropixels probes during the optogenetics experiments. We trained two variational
355 autoencoders (Figure 6B), one each for waveforms and 3D-ACGs, to reduce the input
356 dimensionality to 10 for both features (see **Methods**), thereby minimizing the number of
357 parameters in the ultimate classifier that needed to be trained *de novo*.

358 Our classifier (Figure 6C) consists of: (i) a multi-headed, normalized input layer that accepts the
359 10-dimensional representations of the waveform and 3D-ACG produced by the variational
360 autoencoders, along with a “one-hot” 3-bit binary code of the unit’s cerebellar layer; (ii) a hidden
361 layer that processes the 3 normalized inputs, and (iii) an output layer with one output unit for

362 each of the 5 cell types. The value of the output units sums to 1 so that the output of the classifier
 363 is the probability that a given set of inputs are from each of the 5 cell types. We trained the
 364 weights in the classifier on the data in the ground-truth library using gradient descent with a
 365 leave-one-out cross-validation strategy.

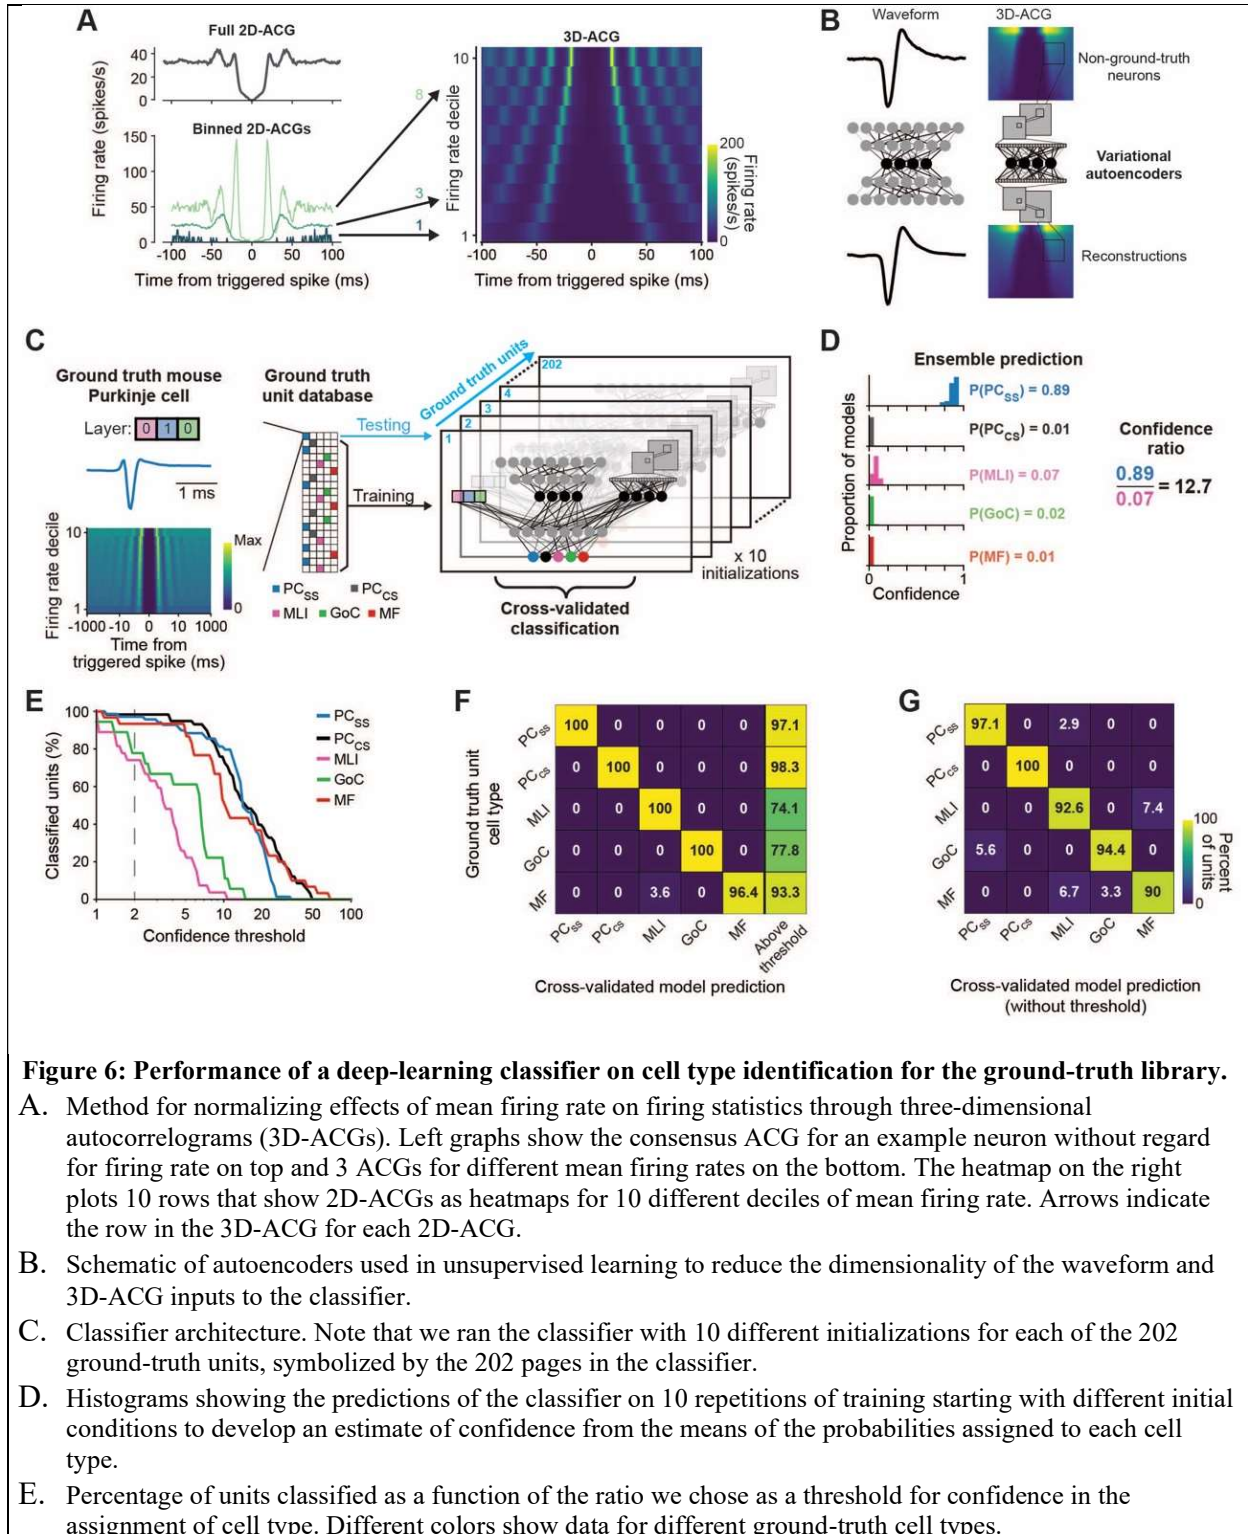


Figure 6: Performance of a deep-learning classifier on cell type identification for the ground-truth library.

- A. Method for normalizing effects of mean firing rate on firing statistics through three-dimensional autocorrelograms (3D-ACGs). Left graphs show the consensus ACG for an example neuron without regard for firing rate on top and 3 ACGs for different mean firing rates on the bottom. The heatmap on the right plots 10 rows that show 2D-ACGs as heatmaps for 10 different deciles of mean firing rate. Arrows indicate the row in the 3D-ACG for each 2D-ACG.
- B. Schematic of autoencoders used in unsupervised learning to reduce the dimensionality of the waveform and 3D-ACG inputs to the classifier.
- C. Classifier architecture. Note that we ran the classifier with 10 different initializations for each of the 202 ground-truth units, symbolized by the 202 pages in the classifier.
- D. Histograms showing the predictions of the classifier on 10 repetitions of training starting with different initial conditions to develop an estimate of confidence from the means of the probabilities assigned to each cell type.
- E. Percentage of units classified as a function of the ratio we chose as a threshold for confidence in the assignment of cell type. Different colors show data for different ground-truth cell types.

- F. Confusion matrix showing the agreement between the predictions of the classifier on a single left-out testing unit and the ground-truth cell type of that testing unit. The numbers in each cell indicate the percentage of ground-truth cell types on the y-axis for each prediction of the classifier on the x-axis, where confidence was required to be greater than 2. The rightmost column shows the percentage of ground-truth neurons that received a confidence greater than 2.
- G. Same as F, but for neurons in the ground-truth library regardless of confidence, i.e. confidence threshold = 0.

366

367 We evaluated not only the accuracy, but also the “confidence” of the output from the classifier.
368 For each leave-one-out sample ($n = 202$ ground-truth units), we trained an ensemble of 10
369 models with random initial conditions. We then averaged the classifier-predicted probability for
370 each cell type across model instantiations. For the example illustrated in Figure 6D, the
371 distributions of cell-type probability reveal repeated predictions that the held-out unit was a
372 Purkinje cell simple spike. The average probability assigned to the simple spike was 0.89 while
373 the average probability assigned to each of the other cell types was less than 0.1. However, that
374 need not have been the case: if the data for a given unit were compatible with more than one cell
375 type, then the classifier might classify the unit as highly-probable to be cell type #1 in one model
376 instance and highly-probable to be cell type #2 in another instance: the average probabilities
377 across 10 runs of the classifier might be similar and therefore closer to 0.5 for these two cell
378 types, indicative of low classifier confidence.

379 We quantified the classifier confidence for each neuron with the “confidence ratio”, computed as
380 the ratio of the mean probability of the most-likely cell type to the mean probability of the
381 second-most-likely cell type. As expected, the percentage of ground-truth units that could be
382 classified decreased as a function of the value of the confidence ratio we chose as the confidence
383 threshold (Figure 6E). Classifier confidence in general was higher for Purkinje cell simple
384 spikes, Purkinje cell complex spikes, and mossy fibers compared to Golgi cells or molecular
385 layer interneurons. Higher confidence thresholds increase the likelihood that cell-type
386 classification is correct, but also decrease the number of units that get classified. We chose a
387 confidence threshold of 2 in the remainder of our analysis because it allowed the majority of
388 neurons to be classified while providing excellent cross-validated classification performance.

389 The classifier showed impressive accuracy when applied to the units in the ground-truth library.
390 For each held-out neuron that exceeded the chosen confidence threshold, we assigned it the cell
391 type that had the highest probability, averaged across the 10 classifier runs. The classifier
392 assigned cell types to 78% of ground-truth molecular layer interneurons and 74% of ground-truth
393 Golgi cells at a confidence threshold of 2 (rightmost column of Figure 6F), almost all correctly
394 as demonstrated by the values of 100% along the diagonal of the confusion matrix (Figure 6F).
395 The classifier exceeded the confidence threshold for more than 90% of mossy fibers, Purkinje
396 cell simple spikes, and complex spikes and again it classified nearly all such units correctly. The
397 accuracy of the classifier degraded without a confidence threshold, but still performed quite well:
398 it exceeded 90% accuracy on all cell types (Figure 6G). The fact that the classifier was more
399 accurate when we required higher confidence means that 1) the classifier has a good internal
400 model of true neuron classes and 2) the choice to set a confidence threshold improves the
401 performance of the classifier. We note that there is some confusion between neurons in different
402 layers despite the use of layer as an input because their waveforms and/or autocorrelograms look
403 ambiguous to a classifier that makes a statistical prediction based on multiple inputs. The
404 classifier performed less well without layer information, mainly because of greater conflation of

405 Golgi cells and molecular layer interneurons (Supplementary Figure 6). Molecular layer
 406 interneurons exhibited two distinctly different waveforms with either a small or a large
 407 repolarization phase. The latter were nearly indistinguishable from the waveforms of Golgi cells.
 408 The two types of waveforms in molecular layer interneurons do not map onto the known
 409 subtypes of molecular layer interneurons⁷¹.

410 Classifier validation of expert-labeled datasets

411 We next evaluated how well the ground-truth classifier (Figure 7A) generalized by attempting to
 412 predict the cell type for neurons in a sample of expert-classified, non-ground-truth recordings
 413 from mice (Medina lab) and from the monkey's floccular complex (Lisberger lab).

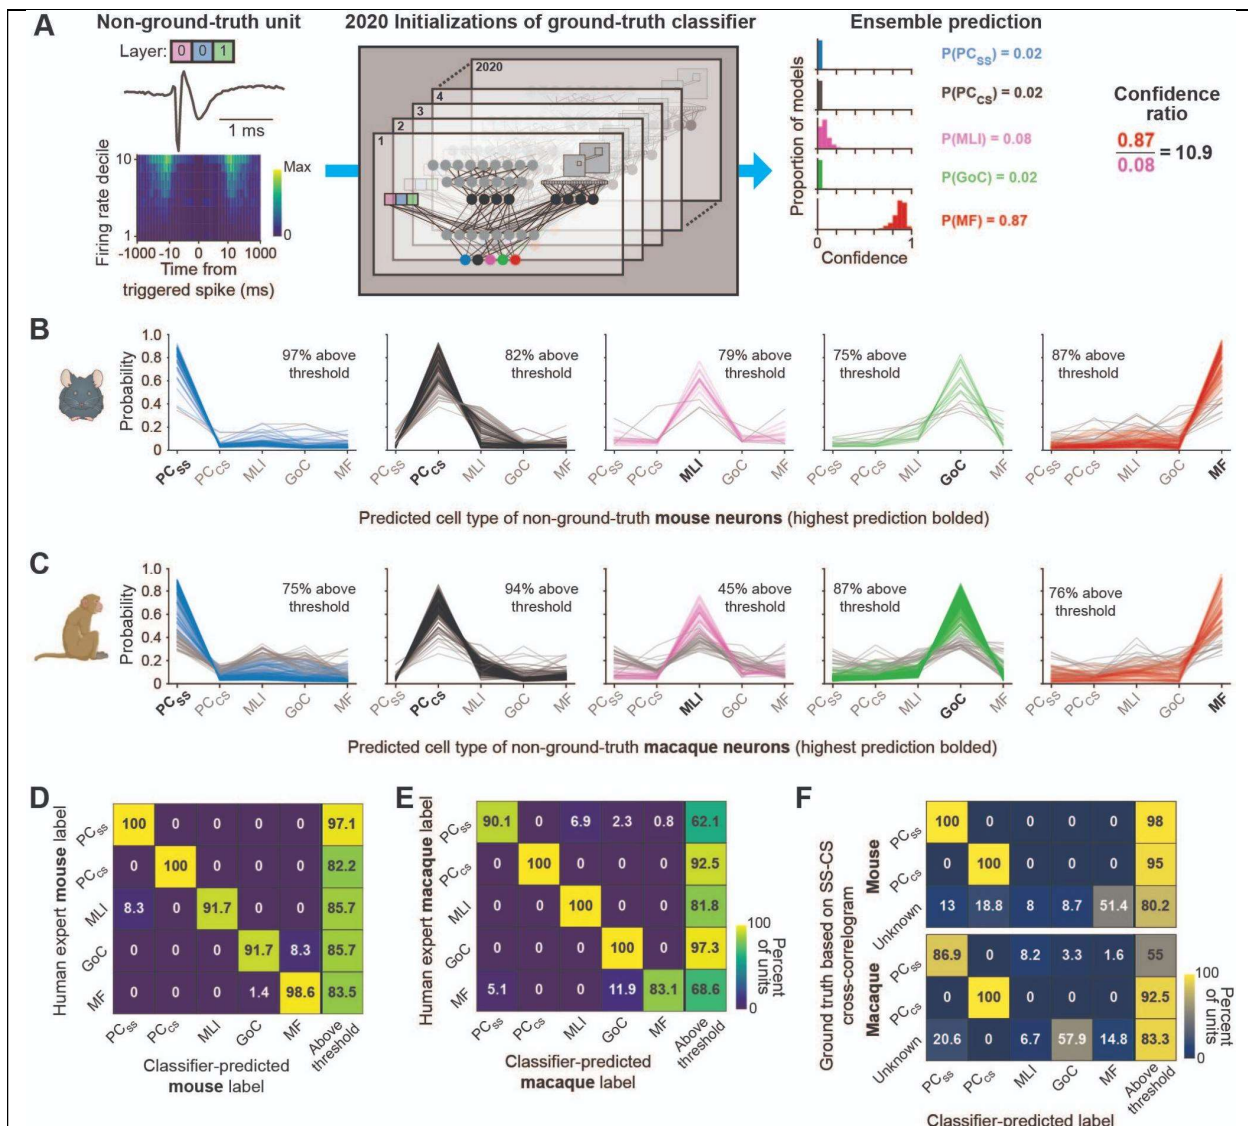


Figure 7: Ground-truth classifier performance on expert-classified datasets from mice and monkeys.

A. Schematic of the ground-truth classifier, repeated from Figure 6C, but now making predictions based on non-ground-truth data from mouse or monkey. The $n = 2020$ instantiations of the classifier arise from training the classifier 10 times with different initial conditions for each of 202 left-out ground-truth units: $10 \times 202 = 2020$.

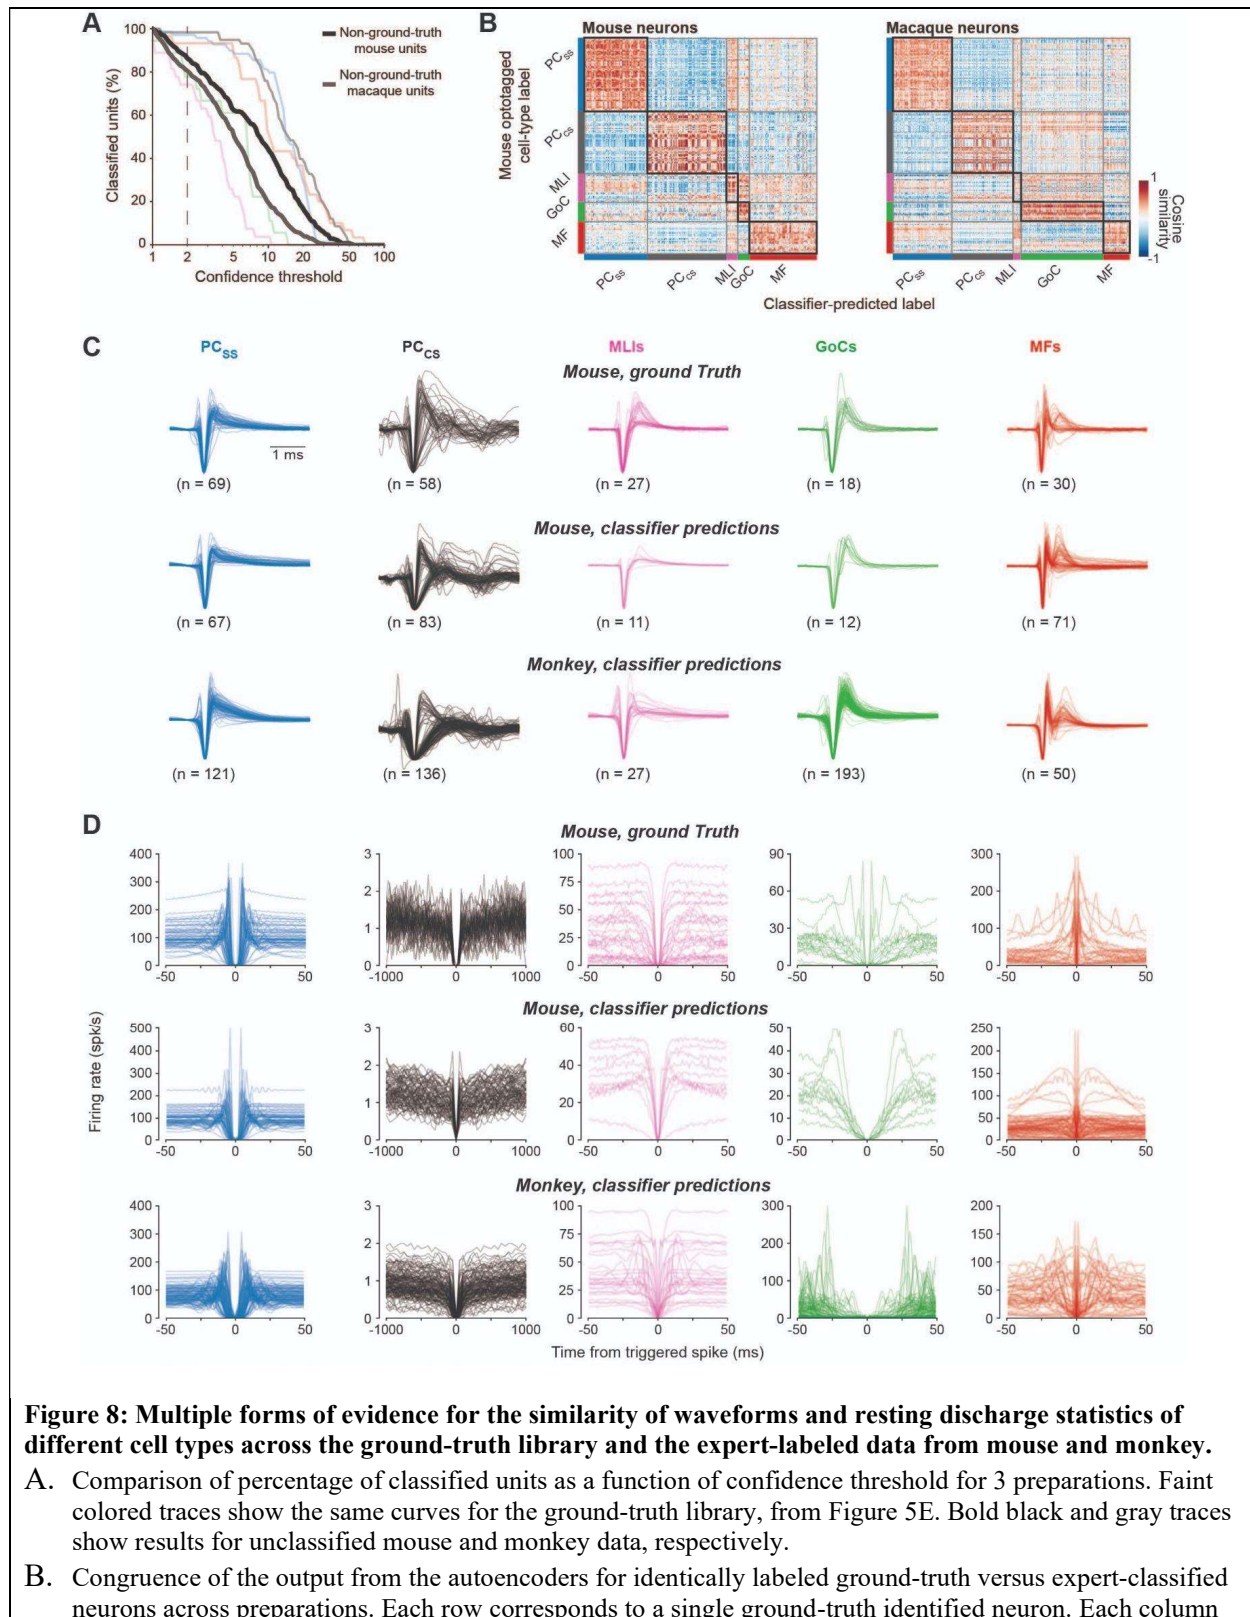
- B. Probability as a function of cell type for expert-classified neurons from mice, divided according to the cell type assigned the highest probability by the classifier. From left to right, the highest-probability cell type was a Purkinje cell simple spike (PC_{ss}), Purkinje cell complex spike (PC_{cs}), molecular layer interneuron (MLI), Golgi cell (GoC), and mossy fiber (MF). Colored versus gray traces represent neurons that exceeded versus failed the confidence threshold of 2. Probability was averaged across runs with 2020 different forms of the classifier (see Methods).
- C. Same as B, but for expert-classified neurons from monkey floccular complex of the cerebellum.
- D. Correspondence matrix showing the agreement between the predictions of the classifier on the x-axis and the expert-labeled cell type from unclassified recordings in mice. The numbers in each cell indicate the percentage of expert-classified cell types on the y-axis as a function of the predictions of the classifier on the x-axis. The rightmost column shows the percentage of expert-classified neurons that received a confidence greater than 2 from the classifier.
- E. Same as D, for expert classified neurons from monkey floccular complex.
- F. Confusion matrices showing good agreement between the output from the classifier and the ground-truth identification in mice and monkeys of Purkinje cell simple spikes and complex spikes from the presence of a complex-spike-triggered pause in simple spike firing.

414

415 Confidence is a particularly important metric for non-ground-truth data. We took advantage of
416 the large number of differentially trained and instantiated models from the ground-truth cross-
417 validation analysis to improve our computation of confidence in the expert-classified datasets.
418 The ground-truth cross-validation of the ground-truth library resulted in 2020 versions of our
419 classifier (202 ground-truth units with 10 instantiations per cross-validation). For each unit in the
420 expert-classified datasets, we averaged the cell-type probabilities predicted by the 2020
421 instantiations of the classifier and created plots of the probability assigned by the classifier as a
422 function of cell type (Figure 7B, C). Units appear in exactly one of 5 different plots, chosen
423 according to the cell type assigned by the classifier as the highest probability, not according to
424 the expert-assigned cell type. For example, the leftmost graph reports probability versus cell type
425 for all units that were classified as most probable to be simple spikes of Purkinje cells, colored
426 according to whether the confidence ratio was below or above 2 (gray versus colored lines). The
427 collection of confidence plots in Figures 7B and C underscores the points of confusion for the
428 classifier, perhaps due to subtle differences in the properties of probes, the behaviors performed
429 by the animals, or the cerebellar sites of recording. For both the mouse data and the monkey data,
430 classifier confidence was greater than 2 for the majority of units, except that only 45% of the for
431 units classified by the experts as molecular layer interneurons in the monkey data were classified
432 “correctly”. The gray curves in Figure 7C suggest that the classifier often conflated molecular
433 layer interneurons and simple spikes in the monkey data, resulting in lower confidence for these
434 neurons.

435 The ground-truth classifier agreed with the human experts about the cell types in mice and
436 monkeys of almost all units that were above confidence threshold, as demonstrated by the large
437 percentages along the diagonal in the correspondence matrices of Figure 7D and E. Further, the
438 ground-truth classifier was quite confident about the cell-types in the non-ground-truth data, as
439 illustrated by the high percentage of units above a confidence threshold of 2 in the rightmost
440 columns of Figure 7D and E. Here, it is important to explain a subtle difference in the numbers in
441 Figures 7B and C versus Figures 7D and E. The leftmost graph of Figure 7C indicates that 75%
442 of the monkey units classified as simple spikes exceeded confidence threshold. Because the units
443 included in Figure 7C are not necessarily classified by experts as simple spikes, 75% is not
444 inconsistent with the 37.9% of expert-classified simple spikes that were below confidence

445 threshold in Figure 7E. The difference in the numbers results from analysis of the space across
 446 orthogonal axes.



corresponds to a single classifier-identified neuron from mouse (left) or monkey (right). Colors at the intersections for each row and column indicate the cosine similarity of the concatenated outputs from the autoencoders for waveform and autocorrelograms, where redder colors indicate greater similarity.

- C. Waveforms of different cell-types across laboratories and species. In the first row, waveforms are divided according to ground-truth cell type in mice. In the second and third rows, cell types are divided according to classifier predictions of cell type for non-ground-truth neurons recorded in mice and monkeys.
- D. Same as C, except showing 2D-autocorrelograms. Note that the spike counts in the autocorrelograms have been normalized by the width of the bin so that the y-axis is in spikes/s.

447

448 The ground-truth classifier identified correctly, with confidence greater than 2, the mouse and
449 monkey Purkinje cell simple spikes and complex spikes from recordings with a complex-spike-
450 triggered pause in simple spikes (100/86.9% and 100/100% in mice/monkeys for simple and
451 complex spikes, respectively. Figure 7F). “Unknown” units, defined as those not identified
452 definitively as Purkinje cells, were distributed across cell-types by the classifier, as expected
453 given that they included recordings from all cell types.

454 Similar properties within cell types across species and cerebellar regions

455 Three additional analyses indicate that the success of the ground-truth classifier on the expert-
456 classified data is based on true statistical similarity of the waveforms and firing statistics of each
457 cell type across datasets. *First*, the percentage of units that we classified with confidence
458 decreased similarly as a function of the confidence threshold for the two samples of expert-
459 classified cells (Figure 8A, thick gray and black traces) and the ground-truth data set (Figure 8A,
460 colored traces). *Second*, analysis of the output of the classifier’s autoencoders revealed excellent
461 agreement between the reduced-dimension representation of expert-classified and ground-truth
462 data (Figures 8B). Here, the neurons in the ground-truth library plot along the y-axis and the
463 expert-classified neurons plot along the x-axis. Warmer colors in the heatmap indicate greater
464 alignment of the 20-dimensional vectors defined by the concatenated outputs of the two auto-
465 encoders in the classifier. *Third*, inspection of the waveforms (Figure 8C) and the 2D-ACGs
466 (Figure 8D) reveals impressive similarity across the ground-truth data, the non-ground-truth
467 mouse data, and the monkey recordings. Here, we have included only the neurons that were
468 classified with confidence greater than 2. The only real exception to the visual impression of
469 similarity is a few of the 2D-ACGs for the ground-truth Golgi cells. The similarity of resting
470 discharge properties across preparations also appears in the 3D-ACGs (Supplementary Figure 4).

471 The large fraction of non-ground-truth neurons that can be classified with confidence, and the
472 agreement with the experts, is unexpected evidence that the properties of different cerebellar cell
473 types are consistent across species and cerebellar regions. It certainly was possible, *a priori*, that
474 a different outcome might have emerged because of genuine differences in waveform or
475 discharge statistics across species and cerebellar regions, differences in data collection and
476 analysis across labs, or a failure of rigor in our procedures for curating the ground-truth and
477 expert-classified data. We anticipate that Figure 8 will serve as a useful resource for other
478 cerebellar labs to be confident of their own rigor as they assign cell types in their own data.

479

480 Discussion

481 Identification of cell type from *in vivo* extracellular recordings is a fundamental issue in systems
482 neuroscience^{5,6,20,55,60,62,63,68,72–75}. Our approach delivers a highly-reliable ground-truth library of
483 the electrophysiological properties of cerebellar cell types in awake mice based on identification
484 through optogenetic stimulation in the presence of synaptic blockers. The ground-truth library
485 consists of the waveform of the electrical recording, the statistics of the spike train, and the layer
486 of the cerebellum where we recorded each unit, information that comes from well-isolated neural
487 recordings using high-density multi-contact probes. Our semi-supervised deep-learning classifier
488 performs well in identifying the cell types in the ground-truth library, while also reporting its
489 confidence in each identification. The internal representations in the classifier reveal similar
490 statistics in the ground-truth library and in independent datasets of expert-classified recordings
491 from the mouse and monkey cerebellum. The cell types predicted by the classifier for the mouse
492 and monkey data agree with the experts' assessments. We are encouraged by the accuracy and
493 precision of our classifier and expect that it will be possible in the future to align the cell type
494 obtained from extracellular recordings with that obtained from other levels of analysis, including
495 anatomical and molecular fingerprints.

496 The strategy we developed may be more useful and important than the exact classifier. Our goal
497 at the outset of our project was to achieve cell-type identification from extracellular recordings in
498 the cerebellar cortex across laboratories and species. We think that the strategies inherent in our
499 classifier, and the classifier itself, can be used with confidence by any cerebellar recording lab
500 that is curating their electrophysiological data with sufficient rigor. However, we also point out
501 that a failure of rigorous curation will lead to noisy and unnecessarily variable inputs to the
502 classifier and will contaminate the cell-type identifications provided as its output⁶¹.

503 Others have attempted to identify the distinct signatures of discrete populations of cerebellar
504 cortical neurons^{20,62–64,76,77}. Past attempts identified either (i) neurons of interest by qualitative
505 agreement with spiking signatures found *in vitro*⁶⁴ or (ii) recordings in anesthetized preparations
506 with neurons identified anatomically via juxtacellular labeling^{20,58,62,63,76,77}. Our recordings in
507 awake and behaving mice demonstrate large variance in the discrete metrics used for
508 summarizing spiking activity both within and across ground-truth classes (Figure 5E,
509 Supplemental Figure 4). Thus, classification schemes reliant on a finite set of specific features
510 probably will not well generalize to other tasks or regions⁶⁰, or from anesthetized to behaving
511 preparations²⁰. We think our approach is more likely to generalize because it leverages the
512 informativeness of the full waveform^{60,61} and 3D-ACG to classify cell types across a wide range
513 of behaviors and stimulus-driven responses.

514 Several features of the strategy embedded in our classifier were critical to its success:

- 515 • Raw waveforms and 3D-ACGs. Raw features are an unbiased input⁶⁰ and allow the classifier
516 to take advantage of extensive information in waveform^{60,61} and discharge statistics. 3D-
517 ACGs normalize for variations in firing rate and create a statistic that can be compared across
518 cerebellar areas, experimental tasks, and species. Similarly, the choice to use single-channel
519 waveforms allows the classifier to generalize across electrode types. We think our strategies
520 are likely to generalize because they use raw features that can be measured readily in other
521 brain areas.

- 522 • *Mitigation of overfitting.* We developed a semi-supervised^{78–80} deep-learning strategy (see
523 *Methods*) to train the classifier with a relatively small number of ground-truth neurons. The
524 unsupervised training of variational autoencoders reduces the chances of overfitting⁸¹. The
525 use of a large unlabeled dataset to train the autoencoders also reduces the chances of
526 overfitting by ensuring that their architecture was designed independently from the ground-
527 truth dataset. Successful predictions of cell type that agree with two expert-classified datasets
528 supports the generalizability of the classifier on other data. Our choice of how to mitigate
529 overfitting should allow our strategy to generalize to other datasets where the dimensionality
530 of the inputs is high and the number of ground-truth neurons is comparatively small.
- 531 • *Confidence.* We were particularly cognizant of making our classifier trustworthy. To do so,
532 we established confidence by training multiple models on the same data⁸² and by using a
533 Bayesian method to calibrate confidence at the single model level^{83,84} (see *Methods*). By
534 requiring confidence above a given threshold^{85,86}, we improved the accuracy of the model on
535 the ground-truth data as well as for non-ground-truth recordings. The ability to choose a
536 confidence threshold allows the user to balance whether to include all neurons even if some
537 cell type assignments might be incorrect or to include fewer neurons with greater certainty in
538 the cell type assignment.

539 The classifier was more successful when it included layer information as an input. With layer
540 information, it classified a higher percentage of the units and, in the ground-truth data, classified
541 them with greater accuracy. However, the use of layer as an input does not make classification
542 trivial. Rather, it creates a platform that will become even more useful as we are able to achieve
543 ground-truth identification of other cell-types in the cerebellum, for example of granule cells
544 with improved recording probes. Also, because waveform and firing statistics are necessary to
545 distinguish cells that are in the same layer, the classifier makes a statistical decision about cell
546 type rather than relying solely on layer for cell identification⁸⁷. Layer is defined in a specific way
547 for the cerebellum¹, but we think of layer information more generally as a specific example of
548 “local electrical properties”. We imagine that there are other ways to quantify those properties,
549 for example LFPs and current-source-density analysis^{88,89}, that will work in brain areas without a
550 laminar structure.

551 The biggest challenge, and our bigger goal, is to see deployment of the strategy outlined here in
552 other brain areas. The use of layer information to improve classification should be relevant to
553 other structures – cerebral cortex⁹⁰, hippocampus⁹¹, superior colliculus⁹² – that have layers with
554 measurable local electrical properties. We also think that the strategies used in our classifier
555 enable generalization by showing how to reduce the dimensions of raw data used as inputs while
556 mitigating the challenges of small numbers of neurons in training sets. We hope that application
557 of our strategy in other brain areas will enable cell-type identification from extracellular
558 recordings, a key element in our collective long-term goal of understanding how neural circuits
559 work and how they generate behavior.

560 *Limitations of the study*

561 One class of limitations is related to our procedures for data collection and curation. While we
562 were scrupulous about spike sorting and criteria for data inclusion, we relied on indirect
563 measures and logic to determine whether a recording was within the region of synaptic blockade.
564 Any errors might have allowed inclusion of certain mis-identified cell types due to off-target

565 opsin expression. We assessed off-target expression thoroughly, revealing, for example,
566 extensive off-target expression in the GlyT2 line^{35,93,94}. Less obvious off-target expression could
567 have escaped our histological analysis and introduced a small fraction of incorrectly-identified
568 cells, for instance in the Math1 line⁹⁵.

569 A second class of limitations is related to the ground-truth classifier's performance for labeling
570 neurons outside of those we focused on in this study. While we trained the classifier on 5
571 identified cell types, other neuronal subtypes exist in the cerebellar cortex²⁸. We may have
572 recorded from the rare cell types in our expert-labeled datasets. Were they misclassified as one of
573 the 5 cell types in our ground-truth library? Or were they correctly recognized as "other" cell
574 types and relegated to the ~15-20% of neurons that failed to reach confidence threshold? In due
575 course, we expect to be able to augment our ground-truth library with unipolar brush cells⁹⁶,
576 candelabrum cells^{97,98}, and other Purkinje layer interneurons^{99,100} as ever more specific Cre lines
577 and viruses become available. We anticipate that advancements in recording probes with higher
578 impedance and/or more dense recording sites will enable reliable recordings from granule cells.
579 If waveforms, firing properties, and layer prove insufficient to segregate identity as other cell
580 types are incorporated, we anticipate that additional information about synaptic and electrical
581 connectivity can further improve the accuracy of our classifier.

582 A third possible limitation is related to the electrophysiological features we chose as inputs for
583 our classifier. Because our goal was to identify cell type based on extracellular recordings with
584 increasingly popular high-density probes, we used waveform, discharge statistics, and layer as
585 inputs. The use of discharge statistics might be problematic in structures with little or no
586 spontaneous discharge in some cell types, although 3D-ACGs enable assessment of discharge
587 statistics from firing related to sensory inputs or behavior. Also, our use of physiological
588 characteristics as inputs to the classifier limits our ability to align electrophysiological cell-type
589 identifications with those provided by single-cell RNAseq¹⁰¹⁻¹⁰³, juxta-cellular labeling^{62,63}, or
590 combinations of *in vivo* recording and single-cell imaging¹⁰⁴. For example, extracellular
591 electrophysiology cannot 'mark' recorded cells for post-hoc analysis of molecular identity using
592 approaches such as RNAseq. However, our approach affords several advantages over recording
593 methods that are compatible with genetic profiling but sacrifice experimental throughput, tissue
594 accessibility, temporal resolution, and/or the number of cell-types that can be simultaneously
595 recorded and identified. Thus, it is necessary to weigh the trade-offs of different strategies for
596 cell-type identification according to the particular experimental question at hand.

597 Despite its limitations, our study demonstrates a strategy that allows different cell types to be
598 identified robustly and reliably, even when analyzing independently collected data. Thus, the
599 strategy should be of great value to the growing community of cerebellar researchers using high-
600 density silicon probes. Furthermore, our strategy provides a template for principled semi-
601 automated detection of cell type, based on assembly of a ground-truth library, that can be applied
602 across other neural circuits in the brain.

603 Methods

604 We conducted experiments in four laboratories and on two species, mice and macaque monkeys.
605 All mouse procedures in the Häusser lab were approved by the local Animal Welfare and Ethical
606 Review Board at University College London and performed under license from the UK Home
607 Office in accordance with the Animals (Scientific Procedures) Act 1986 and in line with the
608 European Directive 2010/63/EU on the protection of animals used for experimental purposes.
609 Mouse procedures in the Hull and Medina labs were approved in advance by the *Institutional*
610 *Animal Care and Use Committees* at Duke University and the Baylor College of Medicine,
611 respectively, based on the guidelines of the United States' *National Institutes of Health*. Monkey
612 procedures in the Lisberger lab were approved in advance by the *Institutional Animal Care and*
613 *Use Committee* at Duke University. Every effort was made to minimize both the number of
614 animals required and any possible distress they might experience.

615 Mouse experimental procedures

616 **Mouse lines.** All transgenic mice were maintained on the C57BL/6J background. Both male and
617 female mice were used and results were pooled.

618 • *Häusser*: Mice expressing Channelrhodopsin-2 (ChR2) in various cerebellar cell types were
619 generated primarily by crossing Cre lines to a Cre-dependent ChR2-eYFP reporter line¹⁰⁵
620 (Ai32, B6.Cg-Gt(ROSA)26Sortm32(CAG-COP4*H134R/EYFP)Hze/J), or, in a subset of
621 experiments, by injecting Cre-dependent ChR2 virus (AAV1.CAGGS.Flex.ChR2-tdTomato
622 [UPenn]). Cre lines were: BAC-Pcp2-IRES-Cre (B6.Cg-Tg(Pcp2-cre)3555Jdhu/J), intended
623 to label Purkinje cells³⁶; Nos1-Cre (B6.129-Nos1tm1(cre)Mgmj/J), intended to label
624 molecular layer interneurons¹⁰⁶; Glyt2-Cre (Tg(Slc6a5-cre)1Uze), intended to label Golgi
625 cells³⁵; and Math1-Cre (B6.Cg-Tg(Atoh1-cre)1Bfri/J), intended to label granule cells⁹⁵. In
626 addition to the transgenic crosses and viral ChR2 expression, we used the Thy1-ChR2 line 18
627 (B6.Cg-Tg(Thy1-COP4/EYFP)18Gfng/J) to express ChR2 in mossy fibers³⁷. Recordings
628 using each strategy were performed as follows: L7-Cre x Ai32 – 1 recording (1 mouse), Nos1-
629 Cre x Ai32 – 40 recordings (34 mice), Nos1-Cre + AAV1.CAGGS.Flex.ChR2-tdTomato – 3
630 recordings (3 mice), GlyT2-Cre x Ai32 – 32 recording (31 mice), GlyT2-Cre +
631 AAV1.CAGGS.Flex.ChR2-tdTomato – 3 recordings (3 mice), Math1-Cre x Ai32 – 47
632 recording (38 mice), Math1-Cre + AAV1.CAGGS.Flex.ChR2-tdTomato – 3 recordings (3
633 mice), and Thy1-ChR2 line 18 – 26 recordings (22 mice). The specificity of opsin expression
634 in the cerebellum of our Cre transgenic crosses was further investigated by crossing the listed
635 Cre lines to a Cre-dependent tdTomato reporter line, Ai9 (B6.Cg-Gt(ROSA)26Sortm9(CAG-
636 tdTomato)Hze/J)¹⁰⁷, so that we could evaluate expression specificity through cytosolic, rather
637 than membrane-bound, fluorescence.

638 • *Hull*: Mice expressing ChR2 or the inhibitory opsin GtACR2 were generated by either
639 crossing the *c-kit*^{IRES-Cre}, intended to label molecular layer interneurons³⁴ or BAC α 6Cre-C,
640 intended to label granule cells¹⁰⁸, to Ai32¹⁰⁵ or a Cre-dependent ArchT-GFP reporter line,
641 Ai40 (B6.Cg-Gt(ROSA)26Sortm40.1(CAG-aop3/EGFP)Hze/J)¹⁰⁹. Alternatively, we injected
642 the same lines with Cre-dependent viruses: (AAV1.CAGGS.Flex.ChR2-tdTomato [UPenn]
643 and AAV1.Efla.Flex.GtACR2.eYFP [Duke]). In addition, we used the Thy1-ChR2 line 18 to
644 express ChR2 in mossy fibers. Recordings using each strategy were performed as follows: *c-*
645 *kit*^{IRES-Cre} + AAV1.CAGGS.Flex.ChR2-tdTomato – 8 recordings (2 mice), *c-kit*^{IRES-Cre} x Ai40

646 – 3 recordings (1 mouse), *c-kit*^{IRE5-Cre} + AAV1.Efla.Flex.GtACR2.eYFP – 11 recordings (6
647 mice), BAC α 6Cre-C x Ai32 – 22 recordings (12 mice), BAC α 6Cre-C +
648 AAV1.CAGGS.Flex.ChR2-tdTomato – 10 recordings (4 mice), and Thy1-Chr2 line 18 – 13
649 recordings (4 mice).

650 ● *Medina*: All experiments were performed in C57BL/6J mice of at least 10 weeks of age,
651 obtained from Jackson Laboratories.

652 **Experimental preparation.** To prepare mice for awake in vivo recordings, in all labs they were
653 implanted with a headplate/headpost under isoflurane anesthesia in sterile conditions. Pre-
654 operative and post-operative analgesia were administered, and mice were allowed to recover
655 from surgery for at least one week before being habituated to head-fixation and prepared for
656 recordings. Lab-specific details are as follows:

657 ● *Häusser*: We installed a custom-made aluminum headplate with a 5 mm long and 9 mm wide
658 oval inner opening over the cerebellum. Mice received a steroid anti-inflammatory drug at
659 least 1 hour before surgery (Dexamethasone, 0.5 mg/kg), followed by an analgesic NSAID
660 (Meloxicam, 5mg/kg) immediately before surgery. Anesthesia was induced and maintained
661 with 5% and 1-2% isoflurane, respectively. The headplate was positioned over the lobule
662 simplex of the left cerebellar hemisphere, angled at approximately 26° with respect to the
663 transverse plane, and attached to the skull with dental cement (Super-Bond C&B, Sun-
664 Medical). Post-operative analgesia (Carprieve, 5 mg/kg) was given for 3 days. After several
665 days of habituation on the recording apparatus, a 1 mm-diameter craniotomy and durotomy
666 were performed to allow access for Neuropixels probes into the lobule simplex (3 mm lateral
667 to the midline, anterior to the interparietal-occipital fissure). Before the craniotomy, a conical
668 nitrile rubber seal (Stock no. 749-581, RS components) was attached to the headplate with
669 dental cement to serve as a bath chamber. The exposed brain was then covered with a humid
670 gelatinous hemostatic sponge (Surgispon) and silicone sealant (Kwik-Cast, WPI) until the
671 experiment was performed (1-2 h after recovery). At the beginning of the experiment, mice
672 were head-fixed, the silicone sealant was removed, and physiological HEPES-buffered saline
673 solution was immediately applied to keep the craniotomy hydrated.

674 ● *Hull*: We installed a titanium headpost (HE Palmer, 32.6x19.4 mm) to the skull and a
675 stainless-steel ground screw (F.S. Tools) over the left cerebellum, both secured with
676 Metabond (Parkell). Mice received dexamethasone (3 mg/kg) 4-24 hours before surgery and
677 an initial dose of ketamine/xylazine (50 mg/kg and 5 mg/kg, IP) and carprofen (5 mg/kg) 20
678 min before induction with isoflurane anesthesia. Isoflurane was administered at 1-2%
679 throughout surgery to maintain appropriate breathing rates and prevent toe pinch response,
680 which were monitored throughout the duration of the surgery. Body temperature was
681 maintained with a heating pad (TC-111 CWE). Mice received buprenex and cefazolin (0.05
682 mg/kg and 50 mg/kg respectively, subq) twice daily for 48 hours after surgery and were
683 monitored daily for 4 days. After 2+ weeks of recovery, mice received dexamethasone (3
684 mg/k) 4-24 hours before recordings. Craniotomies (approx. 0.5-1.5 mm) were opened over
685 vermis or lateral cerebellum (relative to bregma: between -6.0 and -7.0mm AP, and between
686 1.0 and 2.8mm ML) on the first day of recording, under 1-2% isoflurane anesthesia, and were
687 sealed between recordings using Kwik-Cast (WPI) covered by Metabond. Craniotomies

688 could be re-opened for subsequent recordings under brief (<30 min) 1-2% isoflurane
689 anesthesia.

690 • *Medina*: Preoperative analgesia was provided (5g/kg meloxicam, subq, 0.02mL 0.5%
691 bupivacaine and 2% lidocaine, subq) and surgery was carried out under sterile conditions.
692 Mice were anesthetized with isoflurane (5% by volume in O₂ for induction and 1-2% by
693 volume for maintenance; SurgiVet) and kept on a heating pad to maintain body temperature.
694 The skull was exposed and leveled to match the stereotaxic plane before two stainless steel
695 screws were implanted (relative to bregma: AP -0.3mm, ML ±1.4mm) to anchor the whole
696 preparation. A custom-made stainless steel headplate was placed over the screws and the
697 whole preparation was secured to the skull with Metabond cement (Parkell). Additionally, a
698 craniotomy was performed (relative to bregma: AP -5.5mm) consisting of a 5x2 mm section
699 of bone removed to expose the cerebellar vermis and the right anterior and posterior lobes. A
700 chamber was then built with Metabond to cover the exposed bone around the craniotomy, the
701 dura was protected with a thin layer of biocompatible silicone (Kwik-Cast, WPI) and the
702 whole chamber sealed with silicone adhesive (Kwik-Sil, WPI). Mice were monitored until
703 fully recovered from anesthesia and analgesia was provided during the three days following
704 the surgical procedure.

705 **Recording procedures.** All labs followed the same general procedures for mouse cerebellar
706 recordings. Mice were progressively habituated to head fixation prior to Neuropixels recordings.
707 Recordings from the cerebellar cortex were made using Neuropixels 1.0 probes. Probes were
708 coated with DiI, DiO, or DiD (Cat.Nos. V22885, V22886, and V22887; Thermo Fisher
709 Scientific) by repeatedly moving a drop of dye along the probe shank using a pipette until a dye
710 residue was visible along its entire length (~20 passes). Probes were inserted into the brain at a
711 speed of 1-4 μm/s while monitoring electrophysiological signals. The recording chamber
712 surrounding the craniotomy was bathed in ACSF, with or without blockers. After each recording,
713 the probe was removed and soaked in Tergazyme, then soaked in distilled water, and finally
714 washed with isopropyl alcohol. After the last recording session, the brains of most mice were
715 fixed and processed for histology to verify recording locations.

716 In all three laboratories, Neuropixels data were acquired using SpikeGLX
717 (<https://github.com/billkarsh/SpikeGLX>). Following data acquisition, automated spike sorting
718 was performed using Kilosort 2.0^{17,110} and manual curation was performed using Phy
719 (<https://github.com/cortex-lab/phy>). Across all labs, signals were digitized at 30 kHz. Onboard
720 filtering was turned on in some but not all cases.

721 **Optogenetic stimulation and pharmacology.** The same general procedures were followed for
722 optogenetic stimulation in both the Häusser and Hull labs. This procedure consisted of four parts:
723 (1) a baseline recording period without stimulation or drug block, (2) a period of optogenetic
724 stimulation without drug block, (3) a period during which synaptic blockers were applied and
725 diffused into the brain, and (4) a period of optogenetic stimulation with blockers present. The
726 details of the procedures for ground-truth identification of cell-type varied slightly between the
727 two labs.

728 • *Häusser*: Optogenetic stimulation was performed using 1 or 2 blue LEDs (470 nm, Thorlabs
729 M470F3) and in some experiments a blue laser for surface illumination (Stradus 472,

730 Voltran). Surface illumination was performed by coupling the laser or the LED via a patch
731 cable (M95L01, Thorlabs) to a cannula (CFMXB05, Thorlabs) positioned in contact with the
732 brain surface near the probe. In some experiments a second illumination source - a tapered
733 fiber (Optogenix 0.39NA/200 μ m) glued directly to the head of the Neuropixels probe - was
734 inserted into the brain. Total power at the fiber tip (surface fiber) and coupling cannula
735 (tapered fiber) was 1-6.9 mW. Each recording session consisted of: (1) a 20 minute baseline
736 period of spontaneous activity, (2) a set of 50 optogenetic stimuli (stimulation duration: every
737 10 seconds, 1 stimulation of 250 ms or a train of 5 stimulations of 50ms at 5 Hz, depending
738 on the experiment), (3) an application of a synaptic blocker cocktail (Gabazine 0.2-0.8 mM,
739 NBQX 0.8 mM, APV 1.6 mM, MCPG 0-1.3 mM) to the surface of the cerebellum followed
740 by a 20 minute incubation, and (4) a second set of 50 optogenetic stimuli in the presence of
741 synaptic blockers. We note that we did not record any neurons in the ground-truth library
742 with the blue laser as a source of photostimulation.

743 • *Hull*: Neurons expressing ChR2 were activated and neurons expressing GtACR2 were
744 inhibited using a 450 nm laser (MDL-III, OptoEngine) using a 400 micron optic patch fiber
745 (FT400 EMT, Thorlabs) that was positioned 4-10 mm from the brain surface. Power at the
746 brain surface was approximately 2-30 mW and was calibrated for each experiment to
747 produce neuronal responses with minimal artifact. Laser stimulations lasted 50 or 100 ms and
748 were delivered at 0.1 Hz throughout the recording after the 20 minute baseline period, with
749 brief pauses to replenish ACSF or apply blockers (Gabazine 0.2-0.8 mM, NBQX 0.6-1.2
750 mM, AP-5 0.15-0.6 mM, MCPG 1-2.5 mM).

751 **Histology**

752 • *Häusser*: Mice were deeply anesthetized with ketamine/xylazine and perfused transcardially
753 with PBS followed by 4% PFA in PBS. The brains were dissected and post-fixed overnight
754 in 4% PFA, then embedded in 5% agarose (for electrode tract reconstruction) or sectioned at
755 100 μ m (for immunohistochemistry). To reconstruct electrode tracts, we imaged full 3D
756 stacks of the brains in a custom-made serial two-photon tomography microscope coupled to a
757 microtome¹¹¹, controlled using ScanImage (2017b, Vidrio Technologies) and BakingTray
758 (<https://github.com/SainsburyWellcomeCentre/BakingTray>, extension for serial sectioning).
759 The entire brain was acquired with the thickness of physical slices set at 40 μ m and that of
760 the optical sections at 20 μ m (2 optical sections/slice) using a piezo objective scanner
761 (PIFOC P-725, Physik Instrumente) in two channels (green channel: 500–550 nm, ET525/50;
762 red channel: 580–630 nm, ET605/70; Chroma). Each section was imaged in 1025 x 1025 μ m
763 tiles at 512x512-pixel identification with 7% overlap using a Nikon 16x/0.8NA objective.

764 After slicing, samples for immunohistochemistry were blocked with 2.5% normal donkey
765 serum / 2.5% normal goat serum / 0.5% Triton X-100/PBS for 4-6 hours at room
766 temperature, primary antibodies for 4-6 days at 4°C, and secondary antibodies overnight at
767 4°C. Antibodies were diluted in blocking solution. The following antibodies were used: rat
768 anti-mCherry (1:250, ThermoFisher M11217), Mouse anti-Parvalbumin (1:1000, Millipore
769 MAB1572), Donkey anti-Rat-Alexa 594 (1:1000, Invitrogen), and Goat anti-Mouse-Alexa
770 633 (1:1000, Invitrogen). Neurotrace 435/455 (1:250, ThermoFisher N21479) was added to
771 the secondary antibody solution. Sections were mounted and imaged on a Zeiss LSM 880
772 using a 20x objective in 425x425 μ m tiles at 1024x1024-pixel identification.

- 773 ● *Hull*: After the last day of recording, mice were deeply anesthetized with ketamine/xylezine
774 (350 mg/kg and 35 mg/kg, IP) and perfused with PBS followed by 4% PFA in PBS. Brains
775 were extracted and post-fixed in 4% PFA in PBS overnight, then sectioned at 100 μ m using
776 a vibratome (Pelco 102). Before sectioning, some brains were encased in a 2% agar block for
777 stability. Slices were either stained with DAPI (DAPI, Dihydrochloride, 268298, EMD
778 Millipore) and then mounted with mounting medium (Fluoromount-G, Southern Biotech) or
779 were mounted with a DAPI-containing mounting medium (DAPI Fluoromount-G, Southern
780 Biotech). Electrode tracts were visualized using a confocal microscope (Leica SP8).
- 781 ● *Medina*: After perfusion with 4% PFA in PBS, brains were extracted, post-fixed in the same
782 solution for at least 12h and then cryoprotected in 30% sucrose solution in PBS for 48h. The
783 brains were aligned so the coronal sections would match the track angle and sectioned at 50
784 μ m on a cryostat (Leica CM1950). Free floating sections were recovered in PBS and
785 incubated in Hoechst solution for 3 minutes (Hoechst 33342, 2 μ g/mL in PBS-TritonX
786 0.25%, Thermo Fisher Scientific). Sections were then washed in PBS three times using
787 fluorescence protectant medium (ProLong Diamond Antifade, Thermo Fisher Scientific).
788 Epifluorescence was acquired at 10x magnification on an Axio Imager Z1 microscope
789 (Zeiss), track reconstruction and measurements were made on specific microscopy analysis
790 software (ZEN software, Zeiss).

791 **Validation of ChR2 specificity.** To identify the classes of cerebellar neurons that expressed
792 optogenetic actuators, we determined the layer in which fluorescent neurons were present and
793 whether they expressed parvalbumin (PV), which should be present in all molecular layer
794 interneurons and Purkinje cells¹¹². The location of cerebellar layers in each image were identified
795 in the Neurotrace (fluorescent Nissl) channel. The soma locations of neurons expressing
796 tdTomato (as a proxy for Cre expression) and PV were marked manually in grayscale images
797 using Fiji (NIH). Neurons were deemed to express both tdTomato and PV if their somatic
798 locations were less than 5 μ m apart, and the layer of each neuron was determined by overlaying
799 the Neurotrace laminar mask to cell locations.

800 *Macaque experimental procedures*

801 Recordings in non-human primates were conducted in the *Lisberger* lab on three male rhesus
802 monkeys (*Macaca mulatta*) weighing 10-15 kg. A portion of the primate dataset reported here
803 have been published previously along with corresponding detailed methods¹¹³. Briefly, monkeys
804 underwent several surgical procedures under isoflurane in preparation for neurophysiological
805 recordings. In succession, we (i) affixed a head-holder to the calvarium, (ii) sutured a small coil
806 of wire to the sclera of one eye to monitor eye position and velocity using the search coil
807 technique¹¹⁴ and (iii) implanted a recording cylinder aimed at the floccular complex. Analgesics
808 were provided to the monkeys after each surgery until they had recovered.

809 Each day, we acutely inserted either tungsten single electrodes (FHC) or, for the majority of our
810 data, custom-designed Plexon s-Probes into the cerebellar floccular complex. Plexon s-Probes
811 included 16 recording contacts (tungsten, 7.5 μ m diameter) spaced in two columns separated by
812 50 μ m. Adjacent rows of contacts were also separated by 50 μ m. Once we had arrived in the
813 ventral paraflocculus, we allowed the electrode to settle for a minimum of 30 minutes. We
814 recorded continuous wideband data from all contacts at a sampling rate of 40 kHz using the
815 Plexon Omniplex system. We used a 4th order Butterworth low-pass hardware filter with a cutoff

816 frequency of 6 kHz prior to digitization to eliminate distortion of the recorded signal by the
817 electrical field produced by the eye coils. All recordings were performed while the monkey
818 tracked discrete trials of smooth motion of a single black target (0.5° diameter) on a light grey
819 background in exchange for liquid reward. All analyses of the primate neurons utilized the entire
820 recording period and were not contingent on the animal's behavior.

821 Data processing and analysis

822 **Assignment of layers with Phyllum.** For recordings in the mouse, we assigned each channel of
823 the Neuropixels probe to a layer using *Phyllum*, a custom-designed plugin for the curation
824 software *Phy*. The algorithm for layer identification in *Phyllum* starts by automatically setting
825 'anchor' channels whose recorded layer can be unambiguously identified by the presence of
826 Purkinje cell units with simple and complex spikes (Purkinje layer anchor), mossy fiber units
827 with triphasic waveforms (granule layer anchor), or low 1-2 Hz frequency units with wide
828 waveforms indicative of dendritic complex spikes (molecular layer anchor). Then, *Phyllum* fills
829 in the layer of the remaining channels via an iterative procedure based on (1) proximity to the
830 nearest Purkinje cell anchor and (2) allowed layer transitions. Every channel assigned to the
831 Purkinje cell layer must contain at least a Purkinje cell recording within 100 μm , but the channel
832 may also contain additional units located in the neighboring granule or molecular layers. If none
833 of the channels between two consecutive Purkinje cell anchors contain another anchor unit, their
834 layer is set to 'Unknown'. On average, *Phyllum* assigns 82% of all the channels on the
835 Neuropixels probe to a specific layer. Histological reconstruction of 21 recording tracks
836 confirmed that for channels that are assigned a specific layer, the assignment is highly accurate:
837 >99% for molecular layer channels, >98% for granule layer channels, and >95% for Purkinje
838 layer channels.

839 **Curation procedures**

840 ● **Mouse:** After automated sorting with Kilosort and initial manual curation with *Phy*, we
841 implemented checks to ensure that the resulting clusters selected for further analysis
842 corresponded to single units with physiological waveforms, good isolation properties, and
843 few or no refractory period violations. Rigorous curation was especially important for our
844 long recordings, which could have periods of good isolation intermixed with periods of drift
845 or poor unit isolation. We divided our recordings into overlapping segments (30 seconds
846 segments computed every 10 seconds) and computed the 'false-positive' and 'false-negative'
847 rates in each segment. False positives were defined as spikes that fell within the refractory
848 period of a unit (± 0.8 ms from a given spike) and termed refractory period violations
849 (RPVs). The proportion of false-positives was estimated as the quotient between the RPV
850 rate and the mean firing rate⁴⁶. False negatives were defined as spikes that were not detected
851 because they fell below the noise threshold of the recording. They were estimated by fitting
852 each unit's spike amplitude distribution with a Gaussian function^{47,48} and quantifying the
853 fraction of area under the curve clipped at the noise threshold. A 30-second segment was
854 deemed acceptable if it had less than 5% of false positive rate and less than 5% of false
855 negative rate. Acceptable intervals were concatenated and used for subsequent classifier
856 training. A unit was required to have 3 minutes of acceptable isolation during the baseline
857 period to be included in the sample.

858 ● **Monkey:** Following each recording session, individual action potentials were assigned to
859 putative neural units using the semi-automated "Full Binary Pursuit" sorter¹¹⁵, designed to

860 distinguish temporally and spatially overlapping spikes from different neurons. Following
861 automated sorting, we manually curated our dataset, removing neurons with significant
862 interspike interval violations or low signal-to-noise ratios. The majority of units in our
863 primate dataset significantly exceeded the metrics used for automated curation of the mouse
864 data, which potentially biases our sample of primate units towards those that are easier to
865 record.

866 **Data harmonization.** To achieve consistency of data acquired across labs and setups, we
867 implemented several procedures (Supplementary Figure 1):

- 868 1. We reprocessed the wideband voltage recordings from monkeys and mice where the
869 SpikeGLX acquisition filter was off with a causal first-order Butterworth high-pass filter
870 (300 Hz cutoff) to agree with the hardware filter used by Neuropixels probes. Following
871 filtering, we used a drift- and shift-matching algorithm to generate mean waveforms for each
872 recorded unit.
- 873 2. We sought to remove one source of waveform variability by flipping the waveform, if
874 necessary, so that the largest peak was always negative. We did so with the knowledge that
875 the polarity of the action potential waveform depends on a number of factors including the
876 proximity of the recording electrode to the dendrites, soma, and axon^{88,116} and relative
877 orientation of the recording contact and the reference.
- 878 3. We preprocessed all waveform templates by selecting the mean waveform from the highest
879 amplitude channel, resampling it to 30 kHz (if necessary), aligning it to the peak, and
880 inverting it if necessary (see #2 above) to ensure the most prominent peak in the waveform
881 was always negative. We used the harmonized waveforms to compute summary statistics
882 (Figure 5, Supp Figure 4), which have been previously used to classify cerebellar
883 neurons^{19,20,62,63}.
- 884 4. We sub-sampled the spikes of each neuron by grouping waveforms with a similar amplitude
885 on the principal channel, and therefore the same drift-state (i.e. position of probe relative to
886 the recorded neuron): “drift-matching”.
- 887 5. We re-aligned the spikes in time by maximizing the cross-correlation of each spike to a high
888 amplitude template: “shift-matching”. After alignment, the individual spikes were averaged,
889 resulting in the final mean waveform for the neuron under study. Neuropixels data processing
890 (non-manual curation, filtering, drift-shift-matching) was performed using the NeuroPyxels
891 library¹¹⁷.

892 Identification of units directly responsive to optogenetic stimulation

893 Units recorded during optogenetic activation experiments were deemed to be directly
894 responsive to photostimulation if they met the following conditions: (1) their firing rate increased
895 (ChR2) or decreased (GtACR2) more than 3.3 standard deviations from the pre-stimulus baseline
896 within 10 ms of stimulation onset in the 'post-blocker' trials (computed using 0.1 ms bins
897 smoothed with a causal Gaussian filter with a standard deviation of 0.5 ms), (2) they were
898 recorded at a depth at which pharmacological blockade was confirmed, and (3) the spike

899 waveforms evoked in the 'post-blocker' optogenetic stimulation trials matched those recorded
900 during the pre-stimulation 'baseline' period.

901 Construction of 3D autocorrelograms

902 All recordings were performed in awake animals that were either head-fixed but otherwise free
903 to move on a wheel (mice) or performing discrete trials of smooth pursuit (primates), so that
904 firing rates modulated across the experimental session. To account for the impact of changes in
905 firing rate on measures of firing statistics, we constructed “three-dimensional autocorrelograms”
906 (3D-ACGs). At each point in time, we estimated the instantaneous firing rate of the neuron as the
907 inverse interspike interval¹¹⁸. We smoothed firing rates using a boxcar filter (250 ms width) and
908 evaluated the smoothed firing rate at each spike. Finally, we determined the distribution of firing
909 rates from all interspike intervals in a recording, stratified firing rate into 10 deciles, and
910 computed separate 2D-ACGs for the spikes in each decile. We visualized the resulting 3D-ACGs
911 as a surface where the color axis corresponds to the probability of firing, the y-axis stratifies the
912 firing rate deciles so that each 3D-ACG contains 10 rows, and the x-axis represents time from
913 the trigger spike. Note that the spike counts in the autocorrelograms have been divided by the
914 width of the bin so that the y-axis or color map is calibrated in spikes/s.

915 As input to the classifier, we used log distributed bins relative to $t=0$ in contrast to the linearly
916 spaced bins shown in the Figures and Supplemental material.

917 Human expert labeling of cerebellar units

- 918 ● **Mouse.** We used Phylum to identify the layer of each recording. Most Purkinje cells were
919 identified by the presence of both simple spikes and complex spikes and complex-spike-
920 triggered histograms that showed a characteristic pause in the simple spike firing rate
921 following the complex spike. We identified a number of recordings as Purkinje cells by the
922 presence of simple spikes without a complex spike, location in a Purkinje cell layer, and
923 regular firing rate resulting in characteristic “shoulders” present in the autocorrelogram.
924 Putative molecular layer interneurons were identified by their presence in a molecular layer
925 with firing rates above 5 spikes/s, incompatible with the firing properties of the dendritic
926 complex spikes. A subset of putative molecular layer interneurons yielded a properly timed
927 spike-triggered inhibition of an identified Purkinje cell simple spike. Putative mossy fibers
928 were in a granular cell layer and some displayed a characteristic triphasic shape due to the
929 negative afterwave recorded near the glomerulus^{50,51}. Putative Golgi cells were in the
930 granular cell layer and had broad waveforms and relatively regular firing rates. In addition,
931 some pairs of putative Golgi cells showed a double peak in the millisecond range in their
932 cross-correlograms, indicative of gap-junction coupling¹¹⁹.
- 933 ● **Monkey.** We classified recordings as ground-truth Purkinje cells if they demonstrated the
934 characteristic post-complex-spike pause in simple-spike firing. Units that exhibited known
935 characteristics of Purkinje cell simple spikes but lacked a complex spike were treated as
936 “putative” Purkinje cells and used in the comparison of classifier-predicted and expert-
937 predicted labels. We included molecular layer interneurons only if they showed spike-
938 triggered inhibition of an identified Purkinje cell’s simple spikes at short latency, leaving
939 some potential molecular layer interneurons out of our sample. We included units as putative
940 mossy fibers only if the waveform showed a negative after-wave, characteristic of recording
941 near a single glomerulus^{50,51}. We note that our classification of mossy fibers is highly

942 conservative and likely leaves a large subset of mossy fiber recordings not near a glomerulus
943 as unlabeled. Putative Golgi cells were identified by their presence in the granule cell layer,
944 broad waveforms, and highly regular firing, consistent with previous recordings⁶³. Expert
945 labeling of units in the monkey were performed before collection and analysis of the ground-
946 truth units in the mouse.

947 Cross-validated cell-type classification

948 We began the design of our cell-type classifier by selecting the feature space passed to the
949 model, the model class, and model characteristics such as number of units and learning rate,
950 collectively the features that define the model’s “hyperparameters”. Our decision to select
951 hyperparameters independently from the ground truth dataset was critical to ensure
952 generalizability by minimizing overfitting¹²⁰. To construct an unbiased feature space to train the
953 model, we decided *a priori* that the model’s inputs would be anatomical location, extracellular
954 waveform, and firing statistics. We elected not to use summary statistics because they provide an
955 impoverished set of information compared to the inputs we selected. We optimized the model’s
956 architecture fully independently from our ground truth dataset by leveraging n=3,090 curated but
957 unlabeled units that were recorded in the experiments used to create the ground-truth library but
958 were not activated optogenetically. We trained variational autoencoders to reconstruct the
959 waveforms and 3D-ACGs of the unlabeled units, and optimized the architecture of the
960 autoencoders based on the quality of the reconstruction, independently from the ground truth
961 dataset. In the final classifier, we used encoder networks of the two autoencoders to reduce the
962 dimensionality of the waveforms and 3D-ACGs of the ground-truth library. The output of the
963 autoencoders, along with the layer of each neuron, served as inputs used to train the final
964 classifier on the ground truth dataset. Thus, no aspect of the model’s feature space or architecture
965 was chosen based on the model’s performance on the ground truth dataset.

966 Our classifier is a “semi-supervised” model because the variational autoencoders were tuned and
967 trained with unsupervised learning on a set of unlabeled neurons while the complete classifier
968 was trained with supervised learning on a separate set of ground-truth identified neurons. We
969 derived our strategy from the “M1” model¹²¹.

970 Variational autoencoder pre-training on unlabeled data

971 We trained two separate autoencoders to reconstruct the waveforms and log-scaled 3D-ACGs of
972 our unlabeled units. Ultimately, the encoder networks of the autoencoders, trained on our
973 unlabeled data, compressed the input data into two 10-dimensional ‘latent spaces’ for the 2 input
974 features: 3D-ACG and waveform. The training objective of the variational autoencoders was the
975 ‘Evidence Lower Bound’ loss⁷⁰ modified to include a β term to encourage disentanglement of the
976 latent space (Higgins et al. 2016). During training, we employed a Kullback–Leibler divergence
977 annealing procedure to enhance model stability and convergence¹²². Both variational
978 autoencoders were trained through gradient descent with the Adam optimizer, complemented by
979 a cosine-annealing learning rate strategy with periodic warm restarts¹²³.

980 To both facilitate model convergence and yield high-quality reconstructions, we manually
981 adjusted variational autoencoder parameters to adapt the model to our specific data
982 characteristics and improve its performance in subsequent tasks. This procedure did not rely on
983 the ground truth dataset, so we could adjust hyperparameters freely without overfitting to the
984 classification task.

985 ● The final *waveform variational autoencoder* consisted of a 2-layer multilayer perceptron
986 (MLP) encoder with Gaussian Error Linear Units (GeLU) non-linearities¹²⁴ and a 2-layer
987 MLP decoder also with GeLU non-linearities. It was trained for 60 epochs with $\eta=1e-4$, $\beta=5$
988 and a mini-batch size of 128.

989 ● The final *3D-ACG variational autoencoder* consisted of a 2-layer convolutional neural
990 network (CNN) encoder with average pooling after convolutions, batch normalization, and
991 rectified linear unit (ReLU) non-linearities, and a 2-layer MLP decoder with ReLU non-
992 linearities. It was trained for 60 epochs with $\eta=5e-4$, $\beta=5$ and a mini-batch size of 32.

993 The analysis described in Supplementary Figure 7 ensures that our trained variational
994 autoencoders accurately captured the variance in our data.

995 Semi-supervised classifier

996 The final classifier model consisted serially of: (1) the waveform and temporal feature
997 “variational autoencoders” pretrained on unlabeled data to reduce the dimensionality of the input
998 features; (2) a multi-headed input layer that accepted the latent spaces of the waveform and 3D-
999 ACG variational autoencoders, along with a “one-hot-encoded” 3-bit binary code of the unit’s
1000 cerebellar layer; (3) a single fully-connected hidden layer with 100 units that processed the 3
1001 normalized inputs; (4) an output layer with one output unit for each of the 5 cell types (Figure
1002 6C). The value of the output units sums to 1 via a softmax function so that the output of the
1003 classifier is the probability that a given set of inputs is from each of the 5 cell types. Between the
1004 input (2) and fully-connected (3) steps, we applied batch normalization⁸⁷ to equate the
1005 contributions of waveform, discharge statistics, and layer. The fully-connected hidden layer had
1006 a Gaussian prior that encouraged each network unit to have activation values across the training
1007 set with zero mean and unit variance. We trained the weights of the complete classifier on the
1008 data in the ground-truth library using gradient descent with a leave-one-out cross-validation
1009 strategy. We trained the models until convergence or for 20 epochs, whichever came first, with
1010 $\eta=1e-3$, a mini-batch size of 128 and the AdamW optimizer¹²³. We allowed the weights in the
1011 pre-trained variational autoencoders to change during optimization to allow fine-tuning that
1012 caused a small improvement in performance on the downstream classification task.

1013 Finally, we took several steps to ensure that our models are robustly trained and capable of
1014 generalizing well across datasets:

- 1015 1. To account for “class imbalance” created by the different number of neurons in each cell
1016 type, we performed random oversampling of the under-represented cell types for every
1017 model after splitting into testing and validation data¹²⁵.
- 1018 2. We assessed the performance of all models through leave-one-out cross-validation, which
1019 has a lower bias and comparable variance to other cross-validation methods^{126,127} and has
1020 been used in the past to model small datasets such as ours²⁰. Thus, leave-one-out cross-
1021 validation is better than other cross-validation methods but worse at estimating the
1022 generalization error than having an independent test set. We note that our models seemed to
1023 perform well on independent “expert-classified” test sets.

- 1024 3. We did not tune the hyperparameters of the final semi-supervised classifier or of its training
1025 procedure. We used only predefined heuristic values and trained until convergence.
- 1026 4. We adopted a strategy to prevent *confidence miscalibration*, the tendency of deep neural
1027 networks to exhibit over-confidence in their predictions¹²⁸. We corrected the overconfidence
1028 of each model instance by applying a last-layer Laplace approximation to the output
1029 layer^{83,84}. Further, for each leave-one-out sample, we created a “deep ensemble”¹²⁹ by
1030 training an ensemble of 10 models with random initial conditions. We then averaged the
1031 probability for each cell type across model instantiations. Each model generated an average
1032 prediction probability for each cell type, yielding a set of 5 values (for the 5 classes) that
1033 summed to 1. To quantify classifier confidence, we averaged the predicted probability for
1034 each cell type across the 10 instantiations of the model and computed the *confidence ratio* as
1035 the ratio of the highest- to second-highest predicted cell-type for the input features from each
1036 cell in our samples. We chose a confidence ratio of 2 as the *confidence threshold* here, but
1037 higher thresholds could be applied to increase confidence in each prediction of cell type.

1038 Generalization of prediction to unlabeled mouse and macaque cerebellar neuron cell type
1039 We predicted the cell type of mouse (*Medina*) and macaque (*Lisberger*) cerebellar neurons that
1040 were not involved in the classifier training procedures using an ensemble classifier that utilized
1041 all ground truth neurons and initial conditions (202 x 10 = 2020 models in total). Each of the
1042 2020 models was slightly different from the others due to the combination of the 10 initial
1043 conditions and the 'leave-one-out' procedure used to train them and produced slightly different
1044 results. The predicted cell-type of each neuron in the unlabeled sample was chosen as that with
1045 the maximum average prediction across the 2020 models. We applied the *confidence ratio* and
1046 *confidence threshold* as we had for the ground-truth library.

1047

1048 **Author contributions**

	Maxime Beau*	David J. Herzfeld*	Francisco Naveros Arrabal†	Marie E. Hemeit*	Federico D'Agostino*	Marijes Oosteno†	Alvaro Sánchez-López*	Young Yoon Chung	Michael Maibach	Hannah N. Stabb	M. Gabriela Martínez-Lopera	Agoston Lajko	Marie Zedler	Shogo Ohmae	Nathan J. Hall	Beverley A. Clark†	Dana Cohen†	Stephen G. Clark†	Dimitar G. Lisberger†	Court Huitt	Michael Kostadinov†	Michael Häusser†	Javier F. Medina†
Conceptualization	•	•													•	•	•	•	•	•	•	•	•
Mouse data collection	•	•	•	•	•	•	•	•			•	•								•			
Mouse data analysis	•		•	•	•	•	•		•														
Monkey data collection		•																					
Monkey data analysis		•												•									
Classifier design	•	•		•																			
Software development	•	•	•	•					•	•													
Manuscript, led writing																			•				
Manuscript, major contribution to figures	•	•	•																	•			•
Manuscript, major contribution to writing	•	•																		•	•	•	•
Manuscript, contribution to figures and/or writing	•	•	•	•	•	•	•													•	•	•	•
Supervision	•													•	•	•	•	•	•	•	•	•	•
Funding																				•	•	•	•

* co-first author
† co-last author

1049

1050

1051 **Data availability**

1052 The data that support the findings of our study will be made publicly available at the time of
1053 publication.

1054 **Code availability**

1055 The custom analysis code used in our study will be publicly available on Github at the time of
1056 publication for major packages. Other code will be available from the corresponding author upon
1057 request.

1058

1059 **References**

- 1060 1. Ramón y Cajal, S. *Histologie Du Système Nerveux de l'homme & Des Vertébrés*. (Maloine,
1061 Paris, 1909). doi:10.5962/bhl.title.48637.
- 1062 2. Fishell, G. & Heintz, N. The Neuron Identity Problem: Form Meets Function. *Neuron* **80**,
1063 602–612 (2013).
- 1064 3. Masland, R. H. Neuronal cell types. *Curr. Biol.* **14**, R497–R500 (2004).
- 1065 4. Migliore, M. & Shepherd, G. M. An integrated approach to classifying neuronal
1066 phenotypes. *Nat. Rev. Neurosci.* **6**, 810–818 (2005).
- 1067 5. Zeng, H. & Sanes, J. R. Neuronal cell-type classification: challenges, opportunities and the
1068 path forward. *Nat. Rev. Neurosci.* **18**, 530–546 (2017).
- 1069 6. Luo, L., Callaway, E. M. & Svoboda, K. Genetic Dissection of Neural Circuits. *Neuron* **57**,
1070 634–660 (2008).
- 1071 7. Siletti, K. *et al.* Transcriptomic diversity of cell types across the adult human brain. *Science*
1072 **382**, eadd7046 (2023).
- 1073 8. Hodge, R. D. *et al.* Conserved cell types with divergent features in human versus mouse
1074 cortex. *Nature* **573**, 61–68 (2019).
- 1075 9. Poulin, J.-F., Tasic, B., Hjerling-Leffler, J., Trimarchi, J. M. & Awatramani, R.
1076 Disentangling neural cell diversity using single-cell transcriptomics. *Nat. Neurosci.* **19**,
1077 1131–1141 (2016).
- 1078 10. Sugino, K. *et al.* Molecular taxonomy of major neuronal classes in the adult mouse
1079 forebrain. *Nat. Neurosci.* **9**, 99–107 (2006).
- 1080 11. Doyle, J. P. *et al.* Application of a Translational Profiling Approach for the Comparative
1081 Analysis of CNS Cell Types. *Cell* **135**, 749–762 (2008).
- 1082 12. Josh Huang, Z. & Zeng, H. Genetic Approaches to Neural Circuits in the Mouse. *Annu.*
1083 *Rev. Neurosci.* **36**, 183–215 (2013).
- 1084 13. Fujita, H., Kodama, T. & du Lac, S. Modular output circuits of the fastigial nucleus for
1085 diverse motor and nonmotor functions of the cerebellar vermis. *eLife* **9**, e58613 (2020).
- 1086 14. Ecker, J. R. *et al.* The BRAIN Initiative Cell Census Consortium: Lessons Learned toward
1087 Generating a Comprehensive Brain Cell Atlas. *Neuron* **96**, 542–557 (2017).
- 1088 15. Bota, M. & Swanson, L. W. The neuron classification problem. *Brain Res. Rev.* **56**, 79–88
1089 (2007).
- 1090 16. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity.
1091 *Nature* **551**, 232–236 (2017).
- 1092 17. Steinmetz, N. A. *et al.* Neuropixels 2.0: A miniaturized high-density probe for stable, long-
1093 term brain recordings. *Science* **372**, eabf4588 (2021).
- 1094 18. Gouwens, N. W. *et al.* Classification of electrophysiological and morphological neuron
1095 types in the mouse visual cortex. *Nat. Neurosci.* **22**, 1182–1195 (2019).
- 1096 19. Haar, S., Givon-Mayo, R., Barmack, N. H., Yakhnitsa, V. & Donchin, O. Spontaneous
1097 Activity Does Not Predict Morphological Type in Cerebellar Interneurons. *J. Neurosci.* **35**,
1098 1432–1442 (2015).
- 1099 20. Van Dijk, G. *et al.* Probabilistic Identification of Cerebellar Cortical Neurones across
1100 Species. *PLOS ONE* **8**, e57669 (2013).
- 1101 21. Lima, S., Hromádka, T., Znamenskiy, P. & Zador, A. PINP: a new method of tagging
1102 neuronal populations for identification during in vivo electrophysiological recording. *PLoS*
1103 *One* **4**, (2009).

- 1104 22. Tye, K. M. & Deisseroth, K. Optogenetic investigation of neural circuits underlying brain
1105 disease in animal models. *Nat. Rev. Neurosci.* **13**, 251–266 (2012).
- 1106 23. Petreanu, L., Huber, D., Sobczyk, A. & Svoboda, K. Channelrhodopsin-2-assisted circuit
1107 mapping of long-range callosal projections. *Nat. Neurosci.* **10**, 663–668 (2007).
- 1108 24. Petreanu, L., Mao, T., Sternson, S. M. & Svoboda, K. The subcellular organization of
1109 neocortical excitatory connections. *Nature* **457**, 1142–1145 (2009).
- 1110 25. Sjulson, L., Cassataro, D., DasGupta, S. & Miesenböck, G. Cell-Specific Targeting of
1111 Genetically Encoded Tools for Neuroscience. *Annu. Rev. Genet.* **50**, 571–594 (2016).
- 1112 26. Lee, C., Lavoie, A., Liu, J., Chen, S. X. & Liu, B. Light Up the Brain: The Application of
1113 Optogenetics in Cell-Type Specific Dissection of Mouse Brain Circuits. *Front. Neural*
1114 *Circuits* **14**, 18 (2020).
- 1115 27. Liu, X. *et al.* Near-infrared manipulation of multiple neuronal populations via trichromatic
1116 upconversion. *Nat. Commun.* **12**, 5662 (2021).
- 1117 28. Hull, C. & Regehr, W. G. The Cerebellar Cortex. *Annu. Rev. Neurosci.* **45**, 151–175 (2022).
- 1118 29. Striedter, G. F., Bullock, T. H., Preuss, T. M., Rubenstein, J. & Krubitzer, L. A. *Evolution*
1119 *of Nervous Systems*. (Academic Press, 2016).
- 1120 30. Jacobs, B. *et al.* Comparative neuronal morphology of the cerebellar cortex in afrotherians,
1121 carnivores, cetartiodactyls, and primates. *Front. Neuroanat.* **8**, (2014).
- 1122 31. Häusser, M. & Clark, B. A. Tonic Synaptic Inhibition Modulates Neuronal Output Pattern
1123 and Spatiotemporal Synaptic Integration. *Neuron* **19**, 665–678 (1997).
- 1124 32. Khaliq, Z. M., Gouwens, N. W. & Raman, I. M. The contribution of resurgent sodium
1125 current to high-frequency firing in Purkinje neurons: an experimental and modeling study.
1126 *J. Neurosci. Off. J. Soc. Neurosci.* **23**, 4899–4912 (2003).
- 1127 33. Raman, I. M. & Bean, B. P. Ionic currents underlying spontaneous action potentials in
1128 isolated cerebellar Purkinje neurons. *J. Neurosci. Off. J. Soc. Neurosci.* **19**, 1663–1674
1129 (1999).
- 1130 34. Amat, S. B. *et al.* Using c-kit to genetically target cerebellar molecular layer interneurons in
1131 adult mice. *PLOS ONE* **12**, e0179347 (2017).
- 1132 35. Gurnani, H. & Silver, R. A. Multidimensional population activity in an electrically coupled
1133 inhibitory circuit in the cerebellar cortex. *Neuron* **109**, 1739–1753.e8 (2021).
- 1134 36. Zhang, X.-M. *et al.* Highly restricted expression of Cre recombinase in cerebellar Purkinje
1135 cells. *genesis* **40**, 45–51 (2004).
- 1136 37. Hull, C. & Regehr, W. G. Identification of an Inhibitory Circuit that Regulates Cerebellar
1137 Golgi Cell Activity. *Neuron* **73**, 149–158 (2012).
- 1138 38. Chow, L. M. L. *et al.* Inducible Cre recombinase activity in mouse cerebellar granule cell
1139 precursors and inner ear hair cells. *Dev. Dyn. Off. Publ. Am. Assoc. Anat.* **235**, 2991–2998
1140 (2006).
- 1141 39. Eccles, J. C., Ito, M. & Szentágothai, J. *The Cerebellum as a Neuronal Machine*. (Springer,
1142 Berlin, Heidelberg, 1967). doi:10.1007/978-3-662-13147-3.
- 1143 40. Eccles, J. C., Llinás, R. & Sasaki, K. The excitatory synaptic action of climbing fibres on
1144 the Purkinje cells of the cerebellum. *J. Physiol.* **182**, 268–296 (1966).
- 1145 41. Thach, W. T. Discharge of cerebellar neurons related to two maintained postures and two
1146 prompt movements. II. Purkinje cell output and input. *J. Neurophysiol.* **33**, 537–547 (1970).
- 1147 42. Kostadinov, D., Beau, M., Blanco-Pozo, M. & Häusser, M. Predictive and reactive reward
1148 signals conveyed by climbing fiber inputs to cerebellar Purkinje cells. *Nat. Neurosci.* **22**,
1149 950–962 (2019).

- 1150 43. Mathy, A. *et al.* Encoding of Oscillations by Axonal Bursts in Inferior Olive Neurons.
1151 *Neuron* **62**, 388–399 (2009).
- 1152 44. Davie, J. T., Clark, B. A. & Häusser, M. The origin of the complex spike in cerebellar
1153 Purkinje cells. *J. Neurosci. Off. J. Soc. Neurosci.* **28**, 7599–7609 (2008).
- 1154 45. Llinás, R. & Sugimori, M. Electrophysiological properties of in vitro Purkinje cell dendrites
1155 in mammalian cerebellar slices. *J. Physiol.* **305**, 197–213 (1980).
- 1156 46. Hill, D. N., Mehta, S. B. & Kleinfeld, D. Quality Metrics to Accompany Spike Sorting of
1157 Extracellular Signals. *J. Neurosci.* **31**, 8699–8705 (2011).
- 1158 47. Laboy-Juárez, K. J., Ahn, S. & Feldman, D. E. A normalized template matching method for
1159 improving spike detection in extracellular voltage recordings. *Sci. Rep.* **9**, 12087 (2019).
- 1160 48. Fabre, J. M. J., Beest, E. H. van, Peters, A. J., Carandini, M. & Harris, K. D. Bombcell:
1161 automated curation and cell classification of spike-sorted electrophysiology data. Zenodo
1162 <https://doi.org/10.5281/zenodo.8172822> (2023).
- 1163 49. Cohen, J. Y., Haesler, S., Vong, L., Lowell, B. B. & Uchida, N. Neuron-type-specific
1164 signals for reward and punishment in the ventral tegmental area. *Nature* **482**, 85–88 (2012).
- 1165 50. Walsh, J. V., Houk, J. C., Atluri, R. L. & Mugnaini, E. Synaptic Transmission at Single
1166 Glomeruli in the Turtle Cerebellum. *Science* **178**, 881–883 (1972).
- 1167 51. Taylor, A., Elias, S. A. & Somjen, G. Focal synaptic potentials due to discrete mossy-fibre
1168 arrival volleys in the cerebellar cortex. *Proc. R. Soc. Lond. B Biol. Sci.* **231**, 217–230
1169 (1987).
- 1170 52. Sibille, J. *et al.* High-density electrode recordings reveal strong and specific connections
1171 between retinal ganglion cells and midbrain neurons. *Nat. Commun.* **13**, 5218 (2022).
- 1172 53. Arenkiel, B. R. *et al.* In Vivo Light-Induced Activation of Neural Circuitry in Transgenic
1173 Mice Expressing Channelrhodopsin-2. *Neuron* **54**, 205–218 (2007).
- 1174 54. Niell, C. M. & Stryker, M. P. Highly selective receptive fields in mouse visual cortex. *J.*
1175 *Neurosci. Off. J. Soc. Neurosci.* **28**, 7520–7536 (2008).
- 1176 55. Petersen, P. C., Siegle, J. H., Steinmetz, N. A., Mahallati, S. & Buzsáki, G. CellExplorer: A
1177 framework for visualizing and characterizing single neurons. *Neuron* **109**, 3594–3608.e2
1178 (2021).
- 1179 56. Badura, A. & De Zeeuw, C. I. Cerebellar Granule Cells: Dense, Rich and Evolving
1180 Representations. *Curr. Biol.* **27**, R415–R418 (2017).
- 1181 57. D'Angelo, E. *et al.* The cerebellar Golgi cell and spatiotemporal organization of granular
1182 layer activity. *Front. Neural Circuits* **7**, 93 (2013).
- 1183 58. Barmack, N. H. & Yakhnitsa, V. Functions of Interneurons in Mouse Cerebellum. *J.*
1184 *Neurosci.* **28**, 1140–1152 (2008).
- 1185 59. Ophir, O., Shefi, O. & Lindenbaum, O. Neuronal Cell Type Classification using Deep
1186 Learning. Preprint at <https://doi.org/10.48550/arXiv.2306.00528> (2023).
- 1187 60. Lee, E. K. *et al.* Non-linear dimensionality reduction on extracellular waveforms reveals
1188 cell type diversity in premotor cortex. *eLife* **10**, e67490 (2021).
- 1189 61. Lee, K., Carr, N., Perliss, A. & Chandrasekaran, C. WaveMAP for identifying putative cell
1190 types from in vivo electrophysiology. *STAR Protoc.* **4**, 102320 (2023).
- 1191 62. Ruigrok, T. J. H., Hensbroek, R. A. & Simpson, J. I. Spontaneous Activity Signatures of
1192 Morphologically Identified Interneurons in the Vestibulocerebellum. *J. Neurosci.* **31**, 712–
1193 724 (2011).
- 1194 63. Hensbroek, R. A. *et al.* Identifying Purkinje cells using only their spontaneous simple spike
1195 activity. *J. Neurosci. Methods* **232**, 173–180 (2014).

- 1196 64. Prsa, M., Dash, S., Catz, N., Dicke, P. W. & Thier, P. Characteristics of Responses of Golgi
1197 Cells and Mossy Fibers to Eye Saccades and Saccadic Adaptation Recorded from the
1198 Posterior Vermis of the Cerebellum. *J. Neurosci.* **29**, 250–262 (2009).
- 1199 65. Mountcastle, V. B. Modality and topographic properties of single neurons of cat's somatic
1200 sensory cortex. *J. Neurophysiol.* **20**, 408–434 (1957).
- 1201 66. Gold, C., Henze, D. A., Koch, C. & Buzsáki, G. On the Origin of the Extracellular Action
1202 Potential Waveform: A Modeling Study. *J. Neurophysiol.* **95**, 3113–3128 (2006).
- 1203 67. Katai, S. *et al.* Classification of extracellularly recorded neurons by their discharge patterns
1204 and their correlates with intracellularly identified neuronal types in the frontal cortex of
1205 behaving monkeys. *Eur. J. Neurosci.* **31**, 1322–1338 (2010).
- 1206 68. Trainito, C., von Nicolai, C., Miller, E. K. & Siegel, M. Extracellular Spike Waveform
1207 Dissociates Four Functionally Distinct Cell Classes in Primate Cortex. *Curr. Biol.* **29**, 2973-
1208 2982.e5 (2019).
- 1209 69. Shin, S.-L. *et al.* Regular Patterns in Cerebellar Purkinje Cell Simple Spike Trains. *PLOS*
1210 *ONE* **2**, e485 (2007).
- 1211 70. Kingma, D. P. & Welling, M. Auto-Encoding Variational Bayes. Preprint at
1212 <https://doi.org/10.48550/arXiv.1312.6114> (2022).
- 1213 71. Lackey, E. P. *et al.* Cerebellar circuits for disinhibition and synchronous inhibition. *BioRxiv*
1214 *Prepr. Serv. Biol.* 2023.09.15.557934 (2023) doi:10.1101/2023.09.15.557934.
- 1215 72. Senzai, Y. & Buzsáki, G. Physiological Properties and Behavioral Correlates of
1216 Hippocampal Granule Cells and Mossy Cells. *Neuron* **93**, 691-704.e5 (2017).
- 1217 73. Lemon, R. N., Baker, S. N. & Kraskov, A. Classification of Cortical Neurons by Spike
1218 Shape and the Identification of Pyramidal Neurons. *Cereb. Cortex* **31**, 5131–5138 (2021).
- 1219 74. Ardid, S. *et al.* Mapping of Functionally Characterized Cell Classes onto Canonical Circuit
1220 Operations in Primate Prefrontal Cortex. *J. Neurosci.* **35**, 2975–2991 (2015).
- 1221 75. Mitchell, J. F., Sundberg, K. A. & Reynolds, J. H. Differential attention-dependent response
1222 modulation across cell classes in macaque visual area V4. *Neuron* **55**, 131–141 (2007).
- 1223 76. Simpson, J. I., Hulscher, H. C., Sabel-Goedknecht, E. & Ruigrok, T. J. H. Between in and
1224 out: linking morphology and physiology of cerebellar cortical interneurons. in *Progress in*
1225 *Brain Research* vol. 148 329–340 (Elsevier, 2005).
- 1226 77. Holtzman, T., Rajapaksa, T., Mostofi, A. & Edgley, S. A. Different responses of rat
1227 cerebellar Purkinje cells and Golgi cells evoked by widespread convergent sensory inputs.
1228 *J. Physiol.* **574**, 491–507 (2006).
- 1229 78. Hady, M. F. A. & Schwenker, F. Semi-supervised Learning. in *Handbook on Neural*
1230 *Information Processing* (eds. Bianchini, M., Maggini, M. & Jain, L. C.) 215–239 (Springer,
1231 Berlin, Heidelberg, 2013). doi:10.1007/978-3-642-36657-4_7.
- 1232 79. Peikari, M., Salama, S., Nofech-Mozes, S. & Martel, A. L. A Cluster-then-label Semi-
1233 supervised Learning Approach for Pathology Image Classification. *Sci. Rep.* **8**, 7193
1234 (2018).
- 1235 80. van Engelen, J. E. & Hoos, H. H. A survey on semi-supervised learning. *Mach. Learn.* **109**,
1236 373–440 (2020).
- 1237 81. Mahmud, M. S. & Fu, X. Unsupervised classification of high-dimension and low-sample
1238 data with variational autoencoder based dimensionality reduction. in *2019 IEEE 4th*
1239 *International Conference on Advanced Robotics and Mechatronics (ICARM)* 498–503
1240 (2019). doi:10.1109/ICARM.2019.8834333.

- 1241 82. Ganaie, M. A., Hu, M., Malik, A. K., Tanveer, M. & Suganthan, P. N. Ensemble deep
1242 learning: A review. *Eng. Appl. Artif. Intell.* **115**, 105151 (2022).
- 1243 83. Kristiadi, A., Hein, M. & Hennig, P. Being Bayesian, even just a bit, fixes overconfidence
1244 in ReLU networks. in *Proceedings of the 37th International Conference on Machine*
1245 *Learning* vol. 119 5436–5446 (JMLR.org, 2020).
- 1246 84. Daxberger, E. *et al.* Laplace Redux - Effortless Bayesian Deep Learning. in *Advances in*
1247 *Neural Information Processing Systems* vol. 34 20089–20103 (Curran Associates, Inc.,
1248 2021).
- 1249 85. Hendrycks, D. & Gimpel, K. Bridging Nonlinearities and Stochastic Regularizers with
1250 Gaussian Error Linear Units. (2016).
- 1251 86. Taha, A. A., Hennig, L. & Knoth, P. Confidence estimation of classification based on the
1252 distribution of the neural network output layer. *arXiv.org*
1253 <https://arxiv.org/abs/2210.07745v2> (2022).
- 1254 87. Ioffe, S. & Szegedy, C. Batch normalization: accelerating deep network training by
1255 reducing internal covariate shift. in *Proceedings of the 32nd International Conference on*
1256 *International Conference on Machine Learning - Volume 37* 448–456 (JMLR.org, Lille,
1257 France, 2015).
- 1258 88. Buzsáki, G., Anastassiou, C. A. & Koch, C. The origin of extracellular fields and currents
1259 — EEG, ECoG, LFP and spikes. *Nat. Rev. Neurosci.* **13**, 407–420 (2012).
- 1260 89. Senzai, Y., Fernandez-Ruiz, A. & Buzsáki, G. Layer-Specific Physiological Features and
1261 Interlaminar Interactions in the Primary Visual Cortex of the Mouse. *Neuron* **101**, 500-
1262 513.e5 (2019).
- 1263 90. Strominger, N. L., Demarest, R. J. & Laemle, L. B. Cerebral Cortex. in *Noback's Human*
1264 *Nervous System, Seventh Edition: Structure and Function* (eds. Strominger, N. L.,
1265 Demarest, R. J. & Laemle, L. B.) 429–451 (Humana Press, Totowa, NJ, 2012).
1266 doi:10.1007/978-1-61779-779-8_25.
- 1267 91. Andersen, P., Morris, R., Amaral, D., Bliss, T. & O'Keefe, J. *The Hippocampus Book*. xx,
1268 832 (Oxford University Press, New York, NY, US, 2007).
- 1269 92. Basso, M. A. & May, P. J. Circuits for Action and Cognition: A View from the Superior
1270 Colliculus. *Annu. Rev. Vis. Sci.* **3**, 197–226 (2017).
- 1271 93. Tanaka, I. & Ezure, K. Overall distribution of GLYT2 mRNA-containing versus GAD67
1272 mRNA-containing neurons and colocalization of both mRNAs in midbrain, pons, and
1273 cerebellum in rats. *Neurosci. Res.* **49**, 165–178 (2004).
- 1274 94. Ishihara, N., Armsen, W., Papadopoulos, T., Betz, H. & Eulenburg, V. Generation of a
1275 mouse line expressing Cre recombinase in glycinergic interneurons. *genesis* **48**, 437–445
1276 (2010).
- 1277 95. Wagner, M. J., Kim, T. H., Savall, J., Schnitzer, M. J. & Luo, L. Cerebellar granule cells
1278 encode the expectation of reward. *Nature* **544**, 96–100 (2017).
- 1279 96. Altman, J. & Bayer, S. A. Time of origin and distribution of a new cell type in the rat
1280 cerebellar cortex. *Exp. Brain Res.* **29**, 265–274 (1977).
- 1281 97. Lainé, J. & Axelrad, H. The candelabrum cell: A new interneuron in the cerebellar cortex.
1282 *J. Comp. Neurol.* **339**, 159–173 (1994).
- 1283 98. Osorno, T. *et al.* Candelabrum cells are ubiquitous cerebellar cortex interneurons with
1284 specialized circuit properties. *Nat. Neurosci.* **25**, 702–713 (2022).
- 1285 99. Lainé, J. & Axelrad, H. Extending the cerebellar Lugaro cell class. *Neuroscience* **115**, 363–
1286 374 (2002).

- 1287 100. Schilling, K., Oberdick, J., Rossi, F. & Baader, S. L. Besides Purkinje cells and granule
1288 neurons: an appraisal of the cell biology of the interneurons of the cerebellar cortex.
1289 *Histochem. Cell Biol.* **130**, 601–615 (2008).
- 1290 101. Lipovsek, M. *et al.* Patch-seq: Past, Present, and Future. *J. Neurosci.* **41**, 937–946 (2021).
- 1291 102. Cadwell, C. R. *et al.* Electrophysiological, transcriptomic and morphologic profiling of
1292 single neurons using Patch-seq. *Nat. Biotechnol.* **34**, 199–203 (2016).
- 1293 103. Liu, J. *et al.* Integrative analysis of in vivo recording with single-cell RNA-seq data reveals
1294 molecular properties of light-sensitive neurons in mouse V1. *Protein Cell* **11**, 417–432
1295 (2020).
- 1296 104. Turner, N. L. *et al.* Reconstruction of neocortex: Organelles, compartments, cells, circuits,
1297 and activity. *Cell* **185**, 1082-1100.e24 (2022).
- 1298 105. Madisen, L. *et al.* A toolbox of Cre-dependent optogenetic transgenic mice for light-
1299 induced activation and silencing. *Nat. Neurosci.* **15**, 793–802 (2012).
- 1300 106. Jelitai, M., Puggioni, P., Ishikawa, T., Rinaldi, A. & Duguid, I. Dendritic excitation–
1301 inhibition balance shapes cerebellar output during motor behaviour. *Nat. Commun.* **7**,
1302 13722 (2016).
- 1303 107. Madisen, L. *et al.* A robust and high-throughput Cre reporting and characterization system
1304 for the whole mouse brain. *Nat. Neurosci.* **13**, 133–140 (2010).
- 1305 108. Aller, M. i. *et al.* Cerebellar granule cell Cre recombinase expression. *genesis* **36**, 97–103
1306 (2003).
- 1307 109. Daigle, T. L. *et al.* A Suite of Transgenic Driver and Reporter Mouse Lines with Enhanced
1308 Brain-Cell-Type Targeting and Functionality. *Cell* **174**, 465-480.e22 (2018).
- 1309 110. Pachitariu, M. *et al.* MouseLand/Kilosort2: 2.0 final. Zenodo
1310 <https://doi.org/10.5281/zenodo.4147288> (2020).
- 1311 111. Mayrhofer, J. M. *et al.* Distinct Contributions of Whisker Sensory Cortex and Tongue-Jaw
1312 Motor Cortex in a Goal-Directed Sensorimotor Transformation. *Neuron* **103**, 1034-1043.e5
1313 (2019).
- 1314 112. Bastianelli, E. Distribution of calcium-binding proteins in the cerebellum. *The Cerebellum*
1315 **2**, 242–262 (2003).
- 1316 113. Herzfeld, D. J., Joshua, M. & Lisberger, S. G. Rate versus synchrony codes for cerebellar
1317 control of motor behavior. *Neuron* **111**, 2448-2460.e6 (2023).
- 1318 114. Robinson, D. A. A Method of Measuring Eye Movement Using a Scleral Search Coil in a
1319 Magnetic Field. *IEEE Trans. Bio-Med. Electron.* **10**, 137–145 (1963).
- 1320 115. Hall, N. J., Herzfeld, D. J. & Lisberger, S. G. Evaluation and resolution of many challenges
1321 of neural spike sorting: a new sorter. *J. Neurophysiol.* **126**, 2065–2090 (2021).
- 1322 116. Gold, C., Girardin, C. C., Martin, K. A. C. & Koch, C. High-Amplitude Positive Spikes
1323 Recorded Extracellularly in Cat Visual Cortex. *J. Neurophysiol.* **102**, 3340–3351 (2009).
- 1324 117. Beau, M. *et al.* NeuroPyxels: loading, processing and plotting Neuropixels data in python.
1325 Zenodo <https://doi.org/10.5281/zenodo.5509776> (2021).
- 1326 118. Lisberger, S. G. & Pavelko, T. A. Vestibular signals carried by pathways subserving
1327 plasticity of the vestibulo-ocular reflex in monkeys. *J. Neurosci.* **6**, 346–354 (1986).
- 1328 119. van Welie, I., Roth, A., Ho, S. S. N., Komai, S. & Häusser, M. Conditional Spike
1329 Transmission Mediated by Electrical Coupling Ensures Millisecond Precision-Correlated
1330 Activity among Interneurons In Vivo. *Neuron* **90**, 810–823 (2016).
- 1331 120. Verstynen, T. & Kording, K. P. Overfitting to ‘predict’ suicidal ideation. *Nat. Hum. Behav.*
1332 **7**, 680–681 (2023).

- 1333 121. Kingma, D. P., Rezende, D. J., Mohamed, S. & Welling, M. Semi-supervised learning with
1334 deep generative models. in *Proceedings of the 27th International Conference on Neural*
1335 *Information Processing Systems - Volume 2* 3581–3589 (MIT Press, Cambridge, MA, USA,
1336 2014). doi:<https://doi.org/10.48550/arXiv.1406.5298>.
- 1337 122. Bowman, S. R. *et al.* Generating Sentences from a Continuous Space. in *Proceedings of the*
1338 *20th SIGNLL Conference on Computational Natural Language Learning* (eds. Riezler, S. &
1339 Goldberg, Y.) 10–21 (Association for Computational Linguistics, Berlin, Germany, 2016).
1340 doi:10.18653/v1/K16-1002.
- 1341 123. Loshchilov, I. & Hutter, F. SGDR: Stochastic Gradient Descent with Warm Restarts.
1342 Preprint at <https://doi.org/10.48550/arXiv.1608.03983> (2017).
- 1343 124. Hendrycks, D. & Gimpel, K. Gaussian Error Linear Units (GELUs). Preprint at
1344 <https://doi.org/10.48550/arXiv.1606.08415> (2023).
- 1345 125. Lemaître, G., Nogueira, F. & Aridas, C. K. Imbalanced-learn: A Python Toolbox to Tackle
1346 the Curse of Imbalanced Datasets in Machine Learning. *J. Mach. Learn. Res.* **18**, 1–5
1347 (2017).
- 1348 126. Wang, B. & Zou, H. Honest leave-one-out cross-validation for estimating post-tuning
1349 generalization error. *Stat* **10**, e413 (2021).
- 1350 127. Zhang, Y. & Yang, Y. Cross-validation for selecting a model selection procedure. *J.*
1351 *Econom.* **187**, 95–112 (2015).
- 1352 128. Abdar, M. *et al.* A review of uncertainty quantification in deep learning: Techniques,
1353 applications and challenges. *Inf. Fusion* **76**, 243–297 (2021).
- 1354 129. Lakshminarayanan, B., Pritzel, A. & Blundell, C. Simple and Scalable Predictive
1355 Uncertainty Estimation using Deep Ensembles. in *Advances in Neural Information*
1356 *Processing Systems* vol. 30 (Curran Associates, Inc., 2017).
1357
1358