

Improvement Learning Using a Fully Integrated Neuro-Inspired Memristor Chip

Huaqiang Wu (✉ wuhq@tsinghua.edu.cn)

Tsinghua University <https://orcid.org/0000-0001-8359-7997>

Bin Gao

Tsinghua University <https://orcid.org/0000-0002-2417-983X>

Wenbin Zhang

Tsinghua University

Peng Yao

Tsinghua University

Qi Liu

Tsinghua University

Dong Wu

Tsinghua University

Qingtian Zhang

Institute of Microelectronics, Tsinghua University, Beijing, 100084, <https://orcid.org/0000-0003-2732-3419>

Minghong Xu

Tsinghua University

Ying Zhou

Tsinghua University

Zhenhua Zhu

Tsinghua University

Yi Cai

Tsinghua University

Dabin Wu

Tsinghua University

Jianshi Tang

School of Integrated Circuits, Tsinghua University

He Qian

Tsinghua University

Yu Wang

Tsinghua University

Physical Sciences - Article

Keywords:

Posted Date: May 18th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1606598/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

1 **Improvement Learning Using a Fully Integrated Neuro-Inspired** 2 **Memristor Chip**

3
4 Wenbin Zhang^{1,3}, Bin Gao^{1,3,*}, Peng Yao^{1,3}, Qi Liu¹, Dong Wu¹, Qingtian Zhang¹, Minghong
5 Xu¹, Ying Zhou¹, Zhenhua Zhu², Yi Cai², Dabin Wu¹, Jianshi Tang¹, He Qian¹, Yu Wang², and
6 Huaqiang Wu^{1,*}

7
8 ¹School of Integrated Circuits, Beijing National Research Center for Information Science and
9 Technology (BNRist), Tsinghua University, Beijing, China.

10 ²Department of Electronic Engineering, Beijing National Research Center for Information
11 Science and Technology (BNRist), Tsinghua University, Beijing, China.

12 ³These authors contributed equally: Wenbin Zhang, Bin Gao, Peng Yao.

13 *e-mail: gaobl@tsinghua.edu.cn, wuhq@tsinghua.edu.cn

14 15 **Abstract**

16 **Future neuro-inspired computing chips will feature high energy efficiency and extensive**
17 **learning capabilities. Among these features, improvement learning is of great importance**
18 **for edge intelligence devices to adapt to different application scenes and owners. However,**
19 **current technologies for training neural networks greatly rely on large-scale computers,**
20 **which are based on von Neumann architectures and complex digital circuits. The huge**
21 **energy costs of moving massive data between computing and memory units and**
22 **processing high-precision data hinder the implementation of improvement learning on**

23 **edge devices. Here, we report improvement learning with a low-power full-system-**
24 **integrated memristor-based neuro-inspired computing chip. We propose a memristor-**
25 **featured analogue in-memory computing architecture for implementing on-chip learning**
26 **efficiently and integrate ~160K memristor cells with complete peripheral circuits in a**
27 **complementary metal-oxide-semiconductor (CMOS)-compatible technology. We**
28 **demonstrate the on-chip learning of both new samples and new classes across various**
29 **tasks, including motion control, image classification, and speech recognition. The energy**
30 **consumption of the memristor chip for learning is three orders of magnitudes lower than**
31 **that of a central processing unit (CPU)-based system based on hardware system**
32 **measurements. This study is an important step toward the learning chips with much**
33 **higher energy efficiency, which will facilitate the development of future smart edge**
34 **devices that can adapt to all kinds of special scenes.**

35

36 **Main**

37 Human's ability of improvement learning, which features the fast learning of new
38 knowledge while maintaining pre-acquired knowledge, plays a vital role in the growth of
39 intelligence and the fast adaptation to dynamically changing environments. As illustrated in
40 Fig. 1a, the learning of new knowledge (e.g., new samples or new classes) is quickly realized
41 with only a few new inputs, and this is done without losing pre-acquired knowledge. Edge
42 artificial intelligence (AI) applications also require hardware to have such improvement
43 learning abilities to enable the associated devices to adapt to new scenes or user habits¹.
44 However, deep neural network (DNN) training^{2,3}, no matter improvement learning, transfer

45 learning, or learning from scratch, is typically implemented with conventional hardware based
46 on a von Neumann computing architecture and a high-precision digital computing paradigm⁴.
47 The extensive data movement between the processor chip and the off-chip main memory incurs
48 massive energy consumption and accounts for most of the latency of the whole training
49 process^{5,6}. Therefore, although cloud computing platforms can handle such energy-intensive
50 training^{4,7}, their high energy consumption hinders the implementation of improvement learning
51 on power-limited edge computing platforms¹. In contrast, memristor-based neuro-inspired
52 computing eliminates this extensive data movement via its disruptive computation-in-memory
53 architecture and analogue computing paradigm^{6,8,9}. A memristor crossbar array can store an
54 analogue synaptic weight and perform *in situ* vector-matrix multiplication (VMM) operations
55 in parallel in a single time step by exploiting Ohm's law and Kirchhoff's law. A neuro-inspired
56 computing chip that integrates multiple memristor crossbar arrays and CMOS circuits can
57 easily implement DNN inference¹⁰⁻¹³ and has great potential to handle fully on-chip learning
58 without any assistance from off-chip memory¹⁴⁻¹⁶. The significant energy efficiency
59 enhancement provided by memristor-based neuro-inspired computing makes this paradigm
60 very promising for developing future chips that can enable low-power improvement learning
61 devices.

62 Several studies¹⁷⁻²³ have demonstrated learning using memristor crossbar arrays for *in situ*
63 weight tuning while utilizing software or external digital processors to implement the
64 backpropagation (BP) algorithm². However, it is still very challenging to realize a complete
65 full-system-integrated memristor chip with strong learning ability and low energy costs. The
66 key challenge lies in the inefficiency of mapping the BP algorithm to on-chip hardware. First,

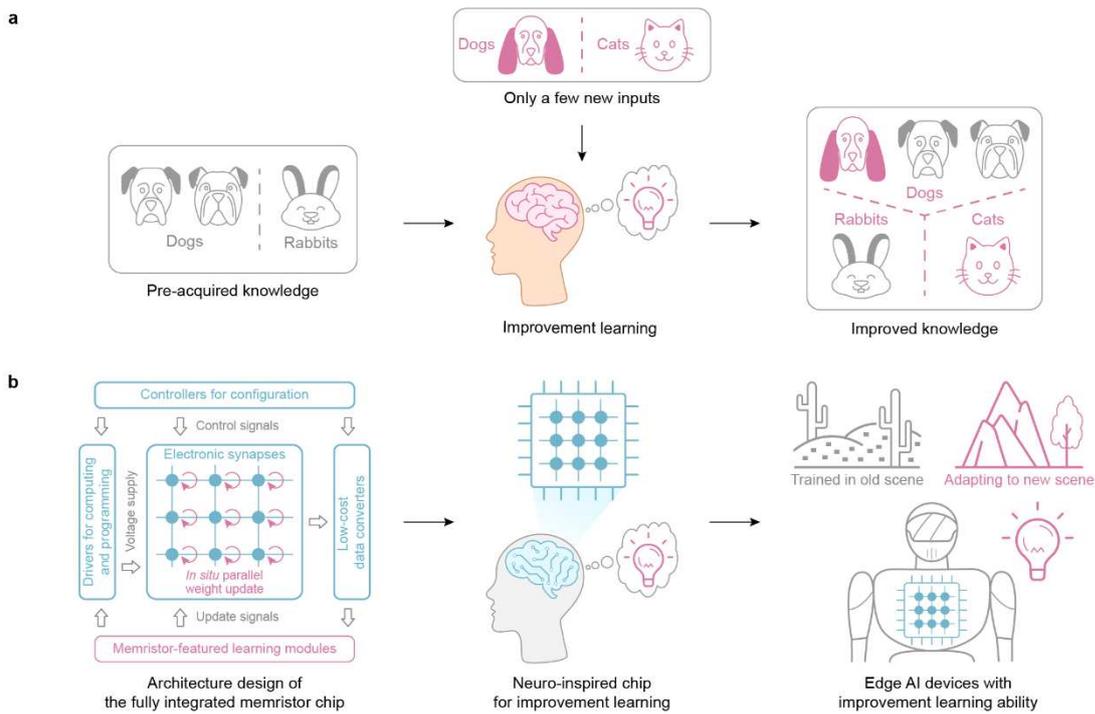
67 an in-memory implementation of the BP algorithm requires costly conductance tuning
68 operations with write verification because of device nonidealities, such as device variability
69 and nonlinear conductance modulation^{14,24–27}. Second, it is difficult to achieve efficient parallel
70 conductance tuning with write verification^{18–20,22}, which makes on-chip learning more time-
71 consuming and energy-consuming. Third, the high-precision data processing operations
72 required during weight update calculations require a large circuit area and high energy
73 consumption, leading to unacceptable overhead^{24,28}.

74 In this work, we demonstrate a memristor-based neuro-inspired computing chip that
75 enables fully on-chip improvement learning, for which a memristor-featured **sign-** and
76 **threshold-based learning** (STELLAR) architecture is proposed. In this architecture, first, the
77 on-chip updating scheme is proposed to tune the memristor without verification. This scheme
78 saves excessive write-and-read costs in the conductance tuning operations during the write
79 verification scheme. Second, the on-chip calculation module is optimized to determine the
80 weight update direction, and this process solely involves the signs of the inputs, outputs, and
81 errors instead of their high-precision formats. This design reduces the circuit design burden and
82 avoids massive overhead during on-chip learning. Third, a cycle-parallel conductance tuning
83 scheme is proposed, in which conductance tuning is performed in a row-by-row parallel fashion.
84 This scheme further reduces the induced energy consumption and latency and accommodates
85 the limited endurance of memristors.

86 The fabricated neuro-inspired computing chip integrates two memristor crossbar arrays
87 (~160K cells in total) and all the necessary circuit modules, including controllers for
88 configuration, drivers for computing and programming, low-cost data converters, and

89 memristor-featured learning modules (as illustrated in Fig. 1b). Based on the obtained
 90 hardware-measured results, the memristor chip exhibits a reduction in energy consumption of
 91 more than three orders of magnitudes relative to that of the CPU-based system (an Intel i9-
 92 10900K). The learning of both the new samples and new classes are experimentally
 93 demonstrated across various tasks, including motion control for a light-chasing car, image
 94 classification, and audio recognition. The memristor-based neuro-inspired computing chip can
 95 facilitate the development of edge AI devices that can adapt to new scenes and new users, as
 96 illustrated in Fig. 1b.

97



98

99 **Fig. 1 | Improvement learning with a neuro-inspired memristor chip.** a, Illustration of the improvement
 100 learning ability possessed by human brains. With pre-acquired knowledge about old dogs and rabbits, the
 101 learning of new samples (i.e., new dogs) and a new class (i.e., cats) is quickly realized with only a few new
 102 inputs. b, Design considerations and future applications of the memristor-based neuro-inspired computing
 103 chip. The chip, designed for fully on-chip learning, integrates all the necessary modules with memristor
 104 arrays. It equips edge AI devices with the improvement learning ability, enabling them to quickly adapt to

105 new scenes.

106

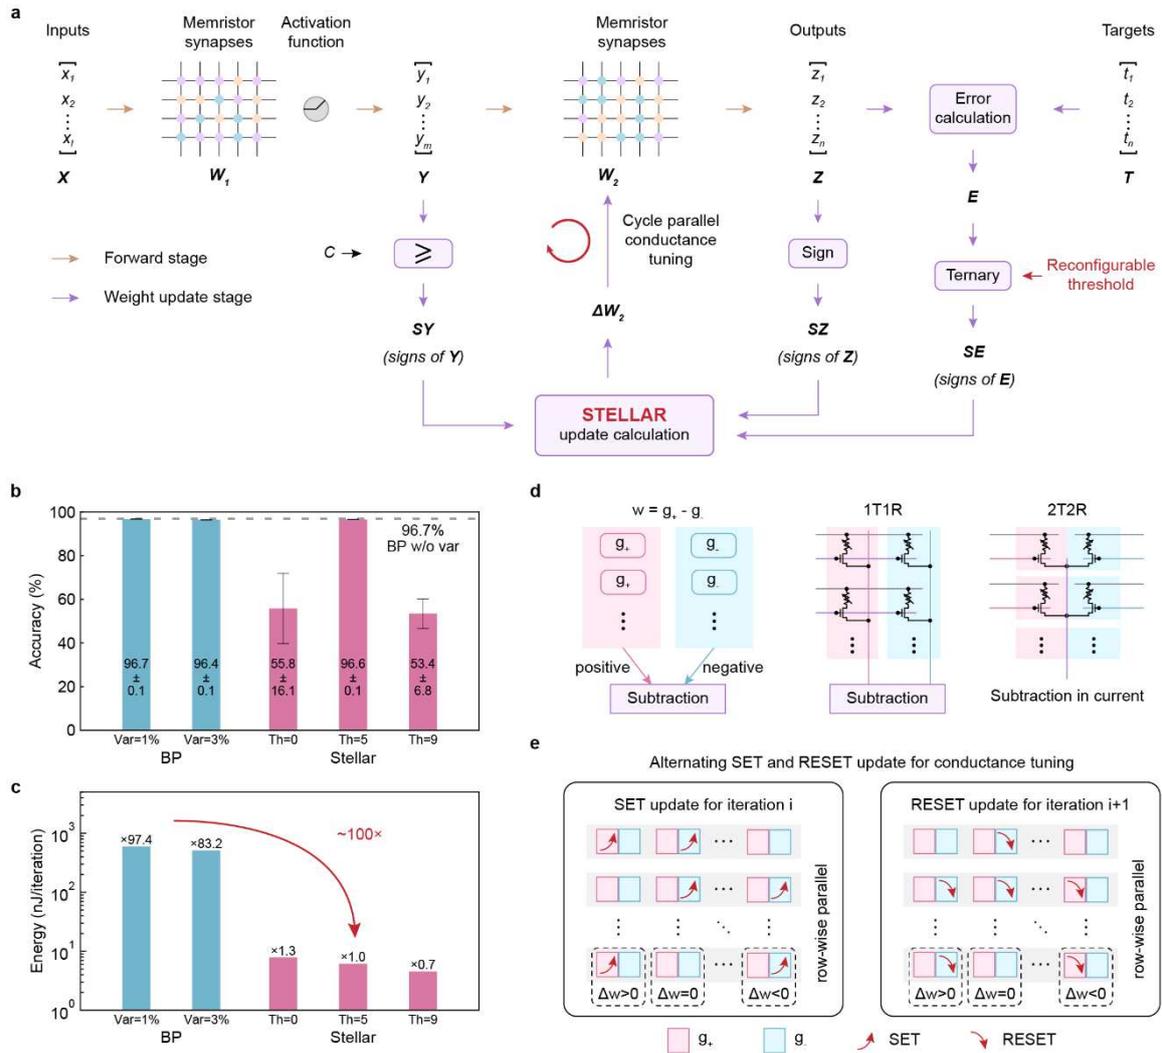
107 **Memristor-featured architecture for on-chip learning**

108 To support on-chip learning with appreciable energy efficiency, area efficiency, and
109 accuracy, we propose the STELLAR architecture, as illustrated in Fig. 2a. The STELLAR
110 architecture exploits the bidirectional analogue switching behaviour of the memristor device²⁹.
111 During the weight update stage, only the weight update direction needs to be calculated based
112 on the signs of the inputs, outputs, and errors. In addition, the architecture predefines a
113 threshold, which filters out the small errors when calculating the error signs and plays a vital
114 role in the convergence of the learning algorithm by avoiding updates that are too sensitive and
115 unnecessary. This threshold is hardware-reconfigurable to adapt to various learning tasks. The
116 algorithmic details of the STELLAR architecture can be found in the Methods section
117 (Algorithm for the STELLAR architecture). Depending on the weight update direction, a
118 corresponding identical SET or RESET pulse is applied to the memristor cells. With this
119 scheme, we can realize energy-efficient hardware by avoiding the complex precise weight
120 update calculation and write verification processes and the complex peripheral circuit design.

121 The learning performance of the STELLAR architecture is compared with that of the
122 conventional methods via simulations on the Modified National Institute of Standards and
123 Technology³⁰ (MNIST) dataset. Here, all memristors are set to random conductance states
124 before the learning process starts. Fig. 2b shows the learning accuracies of the conventional BP
125 method without and with different write variations (1% and 3%, given as the percentage of the
126 full conductance window of the memristor device) in comparison with those of the proposed

127 method under various thresholds. The simulation details can be found in the Methods section
128 (Comparison between STELLAR and the conventional BP algorithm). A properly selected
129 threshold yields improved convergence and learning accuracy, as shown in Extended Data Fig.
130 1a. A threshold that is too small leads to weight updates that are too frequent and an oscillating
131 state of the network, while a threshold that is too large leads to inadequate weight updates. The
132 comparison of the energy consumption among different methods is shown in Fig. 2c. While
133 maintaining almost the same accuracy, the energy consumption of the STELLAR architecture
134 is two orders of magnitudes lower than that of the conventional BP method due to the
135 significant reductions in the precise weight update calculations and the write verification
136 overhead.

137



138

139

140

141

142

143

144

145

146

147

148

149

150

Fig. 2 | Design of the memristor-featured architecture for on-chip learning. **a**, Diagram of the STELLAR architecture used in the memristor chip. The STELLAR algorithm features a sign-based weight update calculation and a reconfigurable threshold for the error sign calculation during the weight update stage. **b**, **c**, Comparison of the classification accuracies (**b**) and the weight update energy consumptions (**c**) using the STELLAR and the conventional BP algorithm on the MNIST dataset in simulations. **b** and **c** show the results of different write variations when utilizing the BP algorithm and the results of different thresholds in the STELLAR learning approach. The bars and error bars show the averages and standard deviations of the results of 10 repeated experiments, respectively. The dashed line in **b** shows the average accuracy obtained by BP without considering the write variations of the memristors. **d**, Representations of weights with a differential conductance pair (left) and with 1T1R (center) and 2T2R (right) configurations. **e**, Illustration of the cycle-parallel conductance tuning scheme. The red and blue rectangles represent the positive and negative memristor cells in each differential pair, respectively. The upwards and downwards arrows represent

151 the SET and RESET operations on the associated memristor cell, respectively.

152

153 The STELLAR architecture realizes positive and negative weights with differential pairs
154 of memristor cells^{19,22,31,32}, as illustrated in Fig. 2d. In conventional crossbar arrays with one-
155 transistor-one-resistor (1T1R) configurations, the two memristor cells in a differential pair are
156 connected to different source lines (SLs), and subtraction is accomplished in a digital fashion.
157 In crossbar arrays with two-transistor-two-resistor (2T2R) configurations, the two memristor
158 cells in a differential pair are connected to the same SL, and subtraction is accomplished
159 directly in the current domain¹⁰. The 2T2R design greatly reduces the SL current and thus
160 reduces IR drop issues, thus enabling a larger array size. Here, we propose a cycle-parallel
161 conductance tuning scheme for such a differential pair configuration, as illustrated in Fig. 2e.
162 In this scheme, the SET and RESET operations are performed alternatively for the learning
163 iterations of arriving input samples (e.g., images). If the SET (RESET) operation is performed
164 in the current learning iteration, then the RESET (SET) operation is performed in the next
165 learning iteration. Only one operation is performed on the whole array in each iteration; this is
166 achieved by selecting which cell of a given differential pair to operate on. Take the SET update
167 mode as an example, if the weight update direction is positive (i.e., $\Delta W > 0$), the positive cell
168 is updated with SET operation to increase the weight; if the $\Delta W < 0$, the negative cell is updated
169 with SET operation to decrease the weight; and if $\Delta W = 0$, neither the positive cell nor negative
170 is updated. The conductance tuning is performed in a row-by-row parallel fashion, and the
171 detailed schematics of the array operations are shown in Extended Data Fig. 2. The memristor
172 devices in the same row are selected through their word line (WL) signals and tuned depending

173 on the corresponding SL signals, while the bit line (BL) signals remain constant. The cycle-
174 parallel conductance tuning scheme can be applied to either 1T1R or 2T2R memristor arrays.
175 Because only one-half of the memristor devices are updated during each on-chip learning
176 iteration, the cycle-parallel conductance tuning scheme reduces the induced energy
177 consumption and alleviates the requirement regarding the memristor endurance.

178

179 **Chip design, fabrication, and measurements**

180 Fig. 3a shows the overall circuit implementation of the proposed STELLAR architecture.
181 This memristor chip consists of controllers for configuration; a 2T2R memristor array ($1568 \times$
182 100), a 1T1R memristor array (100×20); BL, WL, and SL drivers for computing and
183 programming; low-cost analogue-to-digital converters (ADCs); modules for memristor-
184 featured on-chip learning (i.e., error circulating subtractors and weight update logic); and input
185 and output buffers. The 1st-layer memristor array adopts a 2T2R configuration to reduce the IR
186 drop issues that occur in such a large array, while the 2nd-layer memristor array adopts a 1T1R
187 configuration to support more flexible *in situ* weight tuning. The controllers decode the input
188 stage selection signals and provide the output configuration signals to other circuit modules to
189 switch the chips to different working stages (see Methods: Circuit design of the memristor chip
190 and Extended Data Fig. 3). The resolutions of the ADCs are adjustable, called AR-ADC,
191 allowing for flexible threshold reconfiguration¹⁰. The error calculation is accomplished with
192 subtractors, which are realized with counters. The weight update logic determines the weight
193 update direction and the conductance tuning operations.

194 A micrograph of the fabricated chip is shown in Fig. 3b. The chip area breakdown is

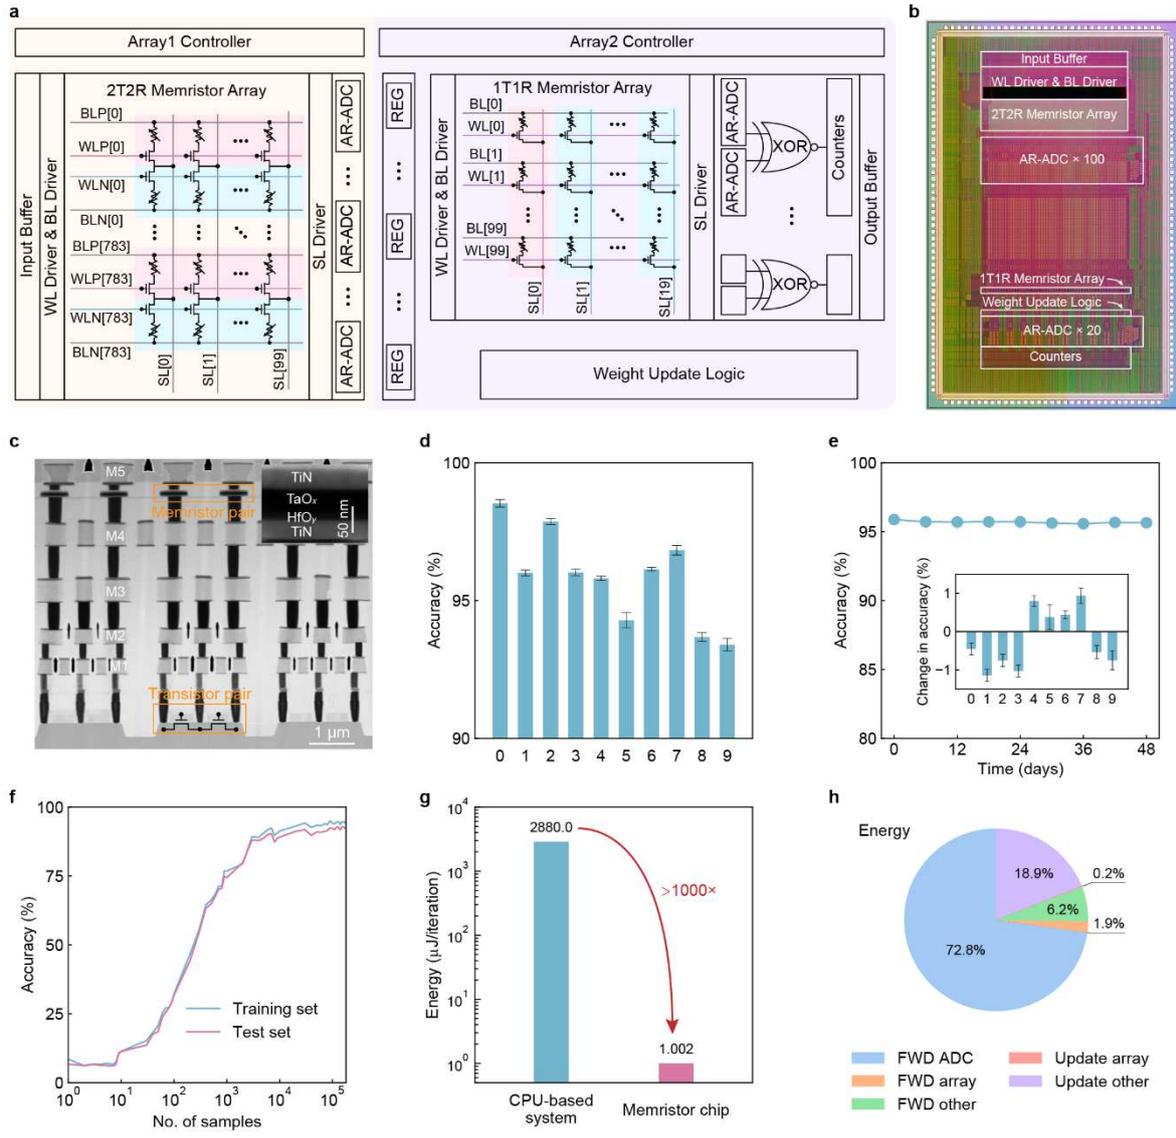
195 described in Extended Data Fig. 4b. The memristor device uses a material stack of
196 TiN/HfO_x/TaO_y/TiN, and the fabrication process is compatible with the standard CMOS
197 process (see Methods: Fabrication of the memristor chip and Extended Data Fig. 4a). As a
198 result, the memristors can be conveniently integrated with complex CMOS circuits to produce
199 an excellent yield (almost 100% of all 160K cells). The cross-section transmission electron
200 microscopy (TEM) image in Fig. 3c shows the integration of memristor cells with CMOS
201 circuitry. The fabricated memristors exhibit uniform and repeatable bidirectional analogue
202 switching with identical pulse trains, as shown in Extended Data Fig. 4d. The ~160K total on-
203 chip memristor cells can be uniformly programmed to 32 conductance states, with maximum,
204 minimum, and average success rates of 99.98%, 99.69%, and 99.90%, respectively (see
205 Methods: Measurements of the memristor devices and Extended Data Fig. 4c).

206 The on-chip inference is first demonstrated for MNIST handwritten digit classification.
207 The weights are off-chip trained and then transferred to the chip as memristor conductance (see
208 Methods: Off-chip training and on-chip inference). The measured classification accuracy of
209 each class (0~9) is presented in Fig. 3d. The average accuracy is 95.8%. The effect of cell
210 conductance fluctuations on the chip accuracy is also evaluated, as shown in Fig. 3e. We
211 monitor the accuracy for 48 days, and no obvious accuracy degradation is observed. We also
212 demonstrate real-time handwritten digit recognition with the memristor chip (see
213 Supplementary Video 3).

214 An on-chip learning task, MNIST image classification, is further demonstrated to verify
215 the on-chip learning ability based on a 784-100-10 multilayer perceptron (MLP). The weights
216 in the 1st layer are off-chip trained and then transferred to the chip as memristor conductance.

217 The memristors in the 2nd layer are firstly programmed to the high-resistance state (HRS) and
218 then updated using the STELLAR scheme. All data processing and signal control processes are
219 executed on the chip. After three epochs of on-chip learning with the training set, the
220 classification accuracies are increased from 8.6% and 8.4% to 94.9% and 92.3% on the training
221 set and test set, respectively, as shown in Fig. 3f. We then evaluate the energy consumption of
222 the on-chip learning with hardware-measured results (see Methods: Energy consumption
223 benchmark). Fig. 3g shows the comparison of energy consumption between our memristor chip
224 and the CPU-based system. The energy consumption is reduced by more than three orders of
225 magnitudes ($>2800 \mu\text{J}$ versus $1 \mu\text{J}$) for one learning iteration with the same MLP network. We
226 also evaluate the energy consumption of a graphics processing unit (GPU)-based system for
227 one training iteration with the same MLP network. The energy reduction ratio reaches 86. The
228 energy breakdown of the memristor chip during the on-chip learning process is presented in
229 Fig. 3h. The energy consumption can be further reduced by optimizing the ADC design.^{33–35}

230



231

232 **Fig. 3 | The memristor chip for on-chip learning. a**, Overview of the architecture of the memristor chip.

233 **b**, Optical microscope image of the chip, where several key components are labelled. **c**, A cross-section TEM

234 image showing 2T2R cells. The transistors in a 2T2R cell share a common source terminal. Inset: cross-

235 section TEM image of the memristor device. **d**, On-chip classification accuracy obtained for each class by

236 using the off-chip trained weights for the MNIST dataset, with bars and error bars showing the averages and

237 standard deviations of the accuracies achieved over 5 inference iterations. **e**, Changes in the classification

238 accuracy over 48 days after weight transfer, with each point showing the average accuracy over 5 inference

239 iterations. Inset: the average changes in the accuracy of each class after 48 days of weight transfer. **f**, Changes

240 in the classification accuracy over 3 learning epochs in the on-chip learning task for the MNIST dataset. **g**,

241 Comparison of the measured energy consumption between the CPU-based system and the memristor chip.

242 **h**, Energy breakdown of the memristor chip during the on-chip learning process.

243

244 **On-chip improvement learning**

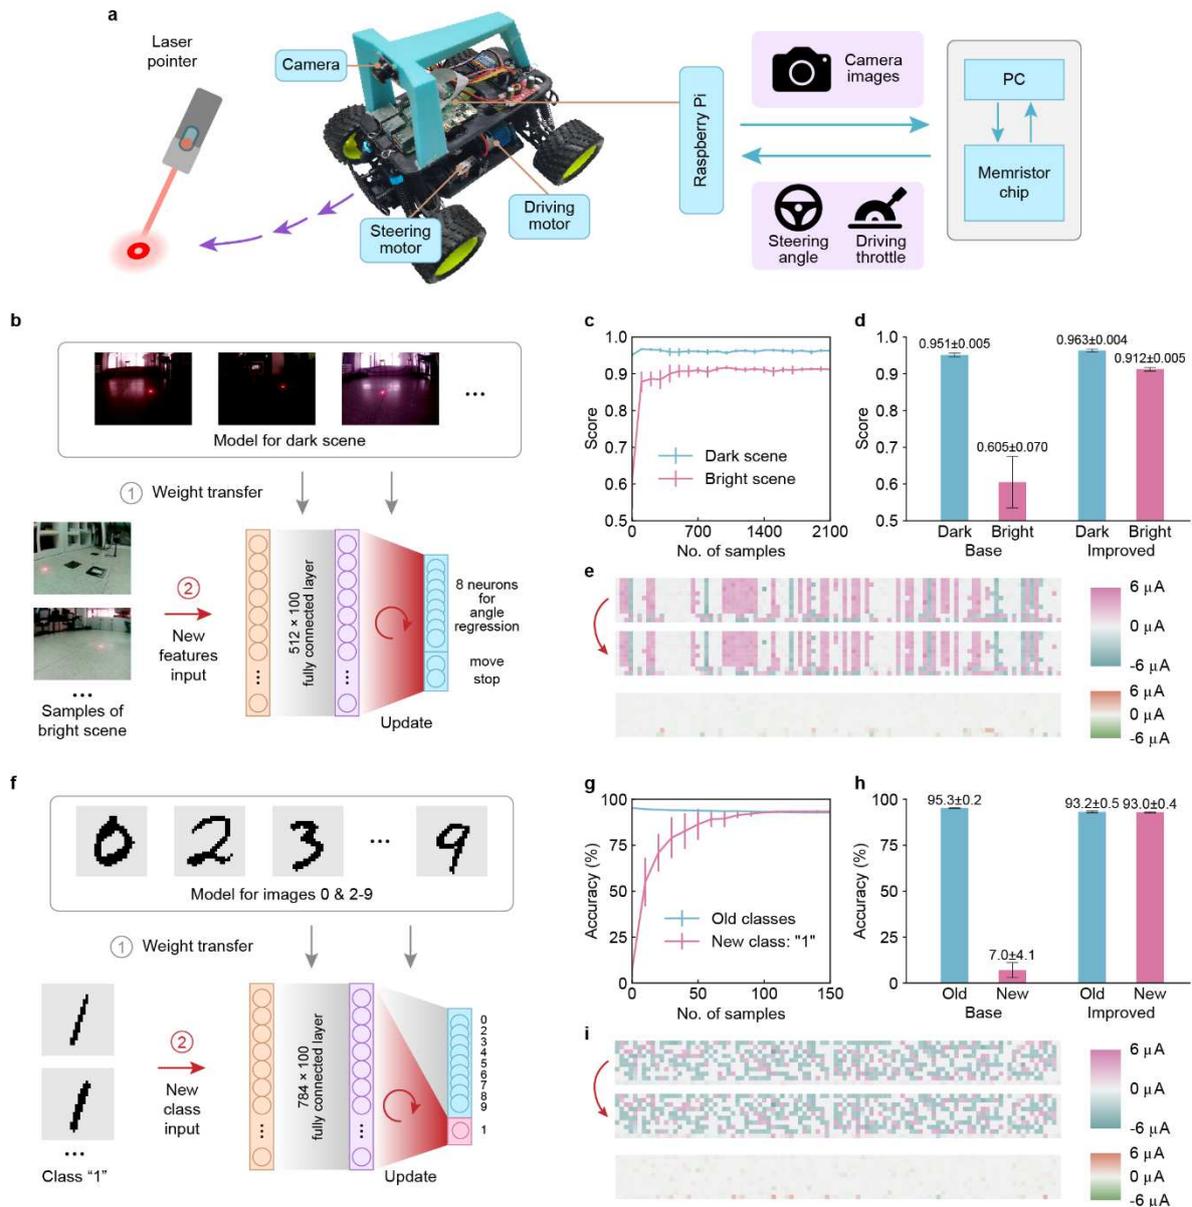
245 The memristor chip is used to further demonstrate four improvement learning tasks,
246 including the motion control task of learning new samples, the audio recognition task of
247 learning new samples, the image classification task of learning a new class, and the motion
248 control task of learning a new class. The implementation of improvement learning mainly
249 includes two stages. First, a model is off-chip trained on the base dataset, and then the model
250 is transferred to the chip. Next, on-chip improvement learning with new data is performed
251 based on the transferred model.

252 We first demonstrate the learning of new samples in a motion control task of a light-
253 chasing car, as illustrated in Fig. 4a. The car is designed to pursue the location of a laser light
254 spot; it is equipped with a camera to capture environmental images, a steering motor for
255 direction control, and a driving motor for throttle control. The memristor chip receives the input
256 features of the environmental images from the PC and provides the output control signals for
257 the steering angles and driving throttles (see Methods: Motion control task of learning new
258 samples for details). As illustrated in Fig. 4b, a convolutional neural network (CNN) including
259 two fully connected (FC) layers with dimensions of $512 \times 100 \times 10$ is first off-chip trained with
260 old scene data (i.e., dark scene data), and then the weights of these two FC layers are transferred
261 to the two corresponding layers of the memristor chip. Next, improvement learning of the new
262 scene (i.e., a bright scene) is performed on the chip by tuning the weights of the last FC layer.
263 Before improvement learning, the car sometimes loses track of the target (i.e., light spot) in the
264 bright scene: it deviates from the target or moves forward even if no target is present. After

265 improvement learning, the car adapts well to the bright scene and still performs well in the dark
266 scene (see Supplementary Video 1). Fig. 4c shows the evolution of the scores during
267 improvement learning, where a score of 1 denotes the best performance (see Methods: Motion
268 control task of learning new samples). The scores become stable after improvement learning
269 with 500 training samples from the new scene. After improvement learning with all the training
270 samples, the average score in the new scene significantly increases from 0.605 to 0.912, while
271 that of the old scene increases from 0.951 to 0.963, showing that no degradation occurs, as
272 presented in Fig. 4d. Fig. 4e shows the transitions and changes of conductance weights before
273 and after the improvement learning.

274 We next demonstrate the learning of a new class in an image classification task involving
275 the MNIST dataset, as illustrated in Fig. 4f. The base model is trained to recognize images of
276 the digits 0 and 2-9 (old classes) and then transferred to the memristor chip. Next, the
277 improvement learning of the new class (i.e., the digit 1) is performed on the chip. As shown in
278 Fig. 4g, the accuracy achieved for the new class increases dramatically during the improvement
279 learning with only a few training samples, while the accuracy achieved for the nine old classes
280 is not significantly reduced. The accuracies yielded for the new class and old classes stabilize
281 after 100 training samples. After improvement learning with 100 training samples, the average
282 accuracy for the new class increases from 7.02% to 93.0%, while that achieved for the old
283 classes decreases from 95.3% to 93.2%, as shown in Fig. 4h. The conductance weight
284 distributions before and after improvement learning and their difference are shown in Fig. 4i.

285



287

288 **Fig. 4 | Improvement learning demonstrations on the memristor chip.** **a**, Illustration of the motion control

289 task (left) and diagram of its control system (right). **b**, Illustration of the learning of new samples in the light-

290 chasing car task. **c**, Changes in the scores that the car produced in the dark and bright scenes over the course

291 of on-chip improvement learning of bright scene samples. The lines and the error bars show the averages

292 and the standard deviations of the scores of 10 repeated experiments with shuffled training sets, respectively.

293 **d**, The initial and final scores in **c**. **e**, The evolution of the weight distribution (100×10) before (top panel)

294 and after (middle panel) the on-chip improvement learning of bright scene samples. Bottom: distribution of

295 the weight changes during improvement learning. **f**, Illustration of the learning of a new class in the image

296 classification task. **g**, Changes in the classification accuracies achieved for the old classes and the new class

297 over the course of on-chip improvement learning of the new class. The lines and the error bars show the
298 averages and the standard deviations of the accuracies of 10 repeated experiments with 150 randomly
299 selected images from training sets, respectively. **h**, The initial accuracies and the accuracies achieved after
300 on-chip improvement learning with 100 samples of the digit 1 in **g**. **i**, The evolution of the weight distribution
301 (100×10) before (top panel) and after (middle panel) the on-chip improvement learning of the digit 1.
302 Bottom: distribution of the weight changes during improvement learning.

303

304 In addition, the memristor chip is also used to implement an audio recognition task of
305 learning new samples and a motion control task of learning a new class. In the audio recognition
306 task, the on-chip improvement learning of female audio samples relies on the model pretrained
307 with male audio samples, as illustrated in Extended Data Fig. 6a. In the motion control task of
308 learning a new class, the on-chip improvement learning of the “moving backwards” action
309 relies on the model knowing when to move forward or stop, as shown in Extended Data Fig.
310 7a. After improvement learning, the accuracy achieved for the new class (i.e., moving
311 backwards) increases from 0% to 95.2%, while that for the old classes (i.e., moving forward or
312 stopping) decreases from 89.5% to 89.4%. The details of the implementation of these two tasks
313 can be found in the Methods Section (Audio recognition task of learning new samples; Motion
314 control task of learning a new class).

315 In summary, we develop a full-system-integrated memristor chip with the improvement
316 learning ability and a low energy cost. The schemes in the STELLAR architecture, including
317 its learning algorithm, hardware realization, and parallel conductance tuning scheme, are
318 general approaches that facilitate on-chip learning by using a memristor crossbar array,
319 regardless of the type of memristor device. We demonstrate the improvement learning of both
320 new samples and a new class across various tasks, including motion control, image

321 classification, and speech recognition, which shows that the STELLAR architecture
322 accommodates the device nonidealities and equips the memristor chip with improvement
323 learning ability to adapt to new scenarios. Furthermore, the benchmark involving our
324 memristor chip shows that it requires three orders of magnitudes less energy consumption for
325 on-chip learning than the CPU-based system. This study is an important step toward future
326 chips with high energy efficiency and extensive learning capabilities. We hope our findings
327 will accelerate the development of future smart edge devices that can adapt to different
328 application scenarios and owners.

329

330 **Methods**

331 **Algorithm for the STELLAR architecture**

332 The STELLAR algorithm is developed to implement on-chip learning with memristors.
333 The algorithm features sign-based weight update calculations and a reconfigurable threshold
334 when calculating the signs of errors. The STELLAR algorithm is a generic approach that
335 accelerates the implementation of on-chip learning with memristors for various neural
336 networks. As the fabricated memristor chip realizes a two-layer neural network, the following
337 description of the algorithm is based on this network. The details of the algorithm for other
338 neural networks can simply be generalized from the following description.

339 A two-layer neural network with dimensions of $784 \times 100 \times 10$ is constructed; it employs
340 a rectified linear unit (ReLU) as an activation function for low-cost hardware implementation.
341 Each weight value is represented by the differential conductance of a pair of memristors. The
342 on-chip learning is composed of two stages: the forward stage and the weight update stage (as

343 shown in Extended Data Fig. 1c). The forward inference is performed with the memristor
 344 crossbar arrays (as illustrated in Fig. 2a):

$$345 \quad \mathbf{Y} = \text{ReLU}(\mathbf{W}_1^T \cdot \mathbf{X})$$

$$346 \quad \mathbf{Z} = \mathbf{W}_2^T \cdot \mathbf{Y}$$

347 where \mathbf{X} is the input vector, \mathbf{Y} is the 1st-layer output vector, \mathbf{Z} is the 2nd-layer output vector, and
 348 \mathbf{W}_1 and \mathbf{W}_2 are the conductance weight matrices of the 1st layer and 2nd layer, respectively. The
 349 weight update is performed for each input vector. The learning algorithm aims to update the
 350 conductance properly to minimize the loss function of the network output and the target. The
 351 square loss function is used in this work, so the error vector is calculated by

$$352 \quad \mathbf{E} = \mathbf{T} - \mathbf{Z}$$

353 where \mathbf{T} is the target vector and \mathbf{E} is the error vector. The signs of the errors are extracted to
 354 determine the weight update direction instead of the accurate values:

$$355 \quad \Delta \mathbf{W}_2 = \mathbf{S}\mathbf{Y} \cdot (\mathbf{S}\mathbf{Z}^T \odot \mathbf{S}\mathbf{E}^T)$$

356 where $\mathbf{S}\mathbf{Y}$ is the sign vector of the 1st-layer output, $\mathbf{S}\mathbf{Z}$ is the sign vector of the 2nd-layer output,
 357 and $\mathbf{S}\mathbf{E}$ is the ternary sign vector of the error. The elements sy , sz , and se in the sign vectors $\mathbf{S}\mathbf{Y}$,
 358 $\mathbf{S}\mathbf{Z}$, and $\mathbf{S}\mathbf{E}$ are defined as follows:

$$359 \quad sy = \begin{cases} 1, & y \geq C \\ 0, & y < C \end{cases}$$

$$360 \quad sz = \text{sign}(z)$$

$$361 \quad se = \begin{cases} +1, & e \geq Th \\ -1, & e \leq -Th \\ 0, & \text{else} \end{cases}$$

362 where y , z , and e , are the elements of vectors \mathbf{Y} , \mathbf{Z} , and \mathbf{E} , respectively. The C in the calculation
 363 of $\mathbf{S}\mathbf{Y}$ can be flexibly reconfigured, and its typical value is the maximum among the 1st-layer
 364 outputs multiplied by 0.4. The threshold Th filters out the small error values, i.e., $|se| < Th$,

365 which prevents weight updates that are too sensitive and improves the convergence of the
366 algorithm.

367

368 **Comparison between STELLAR and the conventional BP algorithm**

369 We evaluate the performance of STELLAR and the conventional BP method in an on-chip
370 learning simulation with memristors. In the STELLAR architecture, the weight update
371 calculation is accomplished by custom-designed circuits, and the conductance tuning
372 operations are accomplished without verification. In the BP method, the weight update
373 calculation is accomplished by an Intel Xeon E5-2699 processor, and the conductance tuning
374 operations are accomplished by using a write-verify scheme. We estimate the accuracies and
375 energy consumptions yielded by using the same training task for MNIST image classification.
376 An MLP with dimensions of $784 \times 100 \times 10$ is utilized when comparing these methods. We use
377 a behavioural memristor device model in the simulation of conductance tuning operation, as it
378 provides a final conductance distribution after executing a SET/RESET pulse on a given initial
379 conductance²⁸. In the BP method, the memristor is programmed to a certain margin with the
380 target conductance as the margin middle value and a given variation as the margin upper/lower
381 limit. The value of the variation is constant for all the conductance states and is defined as a
382 certain percentage of the full conductance window.

383 We focus on the energy consumption of the weight update stage. The energy
384 consumption required by the STELLAR architecture involves a weight update calculation
385 utilizing custom-designed circuits and a one-time conductance tuning operation without
386 verification. The energy consumption required by the BP method involves a weight update

387 calculation utilizing the CPU and multiple conductance tuning operations with verification,
388 requiring the ADC to verify the device conductance during the tuning process. We combine the
389 experimental data and simulation data to estimate the energy consumption. The energy
390 consumption of the custom-designed circuits is obtained from the corresponding circuits used
391 in our memristor chip with the Cadence simulator. The energy consumption of the Intel Xeon
392 E5-2699 processor is estimated as the number of calculation operations divided by its energy
393 efficiency³⁶. The energy consumption of the conductance tuning operations is estimated as the
394 number of tuning operations multiplied by the average energy of each operation. The number
395 of tuning operations is obtained from the on-chip learning simulation. The average energy of a
396 write operation is estimated from the measured results yielded by our memristor chip, while
397 the energy of a read operation is obtained from a 130-nm ADC with an 8-bit resolution³⁷.

398

399 **Circuit design of the memristor chip**

400 The STELLAR algorithm is modified for its chip implementation, and the schematic is
401 shown in Extended Data Fig. 1b. The selection of threshold is realized by configuring the
402 output scales and target values. It is predefined to take the first 8 bits of 17-bit signed errors E
403 when calculating the ternary signs SE . As a result, the selection of the threshold is equivalent
404 to adjusting the error scale, which is accomplished by adjusting the output scales and target
405 values. The output scales of Y and Z can be modulated with the resolution-adjustable ADC,
406 whose quantization scale is determined by its resolution.

407 The on-chip learning of the memristor chip involves three working stages: Forward
408 (FWD), SET, and RESET (as shown in Extended Data Fig. 1c). The controllers decode the

409 input stage selection signals to the output signals for voltage selection. The voltage selection
410 signals are taken as the control signals of the BL/WL and SL drivers; the voltage to be applied
411 to the memristor array is selected through the multiplexers in the drivers.

412 In the FWD stage, VMM operations are performed with the memristor arrays and on-chip
413 ADCs. The 1st-layer array utilizes the 2T2R configuration to reduce the IR drop. The 2nd-layer
414 array utilizes the 1T1R configuration with the two adjacent columns representing positive and
415 negative weights, respectively. The resolutions of the ADCs are adjustable, enabling the
416 flexible adjustment of the output scale. A timing-dependent scheme based on the number of
417 pulses is used in this memristor chip. If the resolutions of the 1st-layer ADCs are configured as
418 N bits, they generate 2^N pulses. For each pulse cycle of the 1st-layer output, the 2nd-layer
419 performs one VMM. The outputs of the 1st-layer ADC are stored in 1-bit registers and sampled
420 simultaneously by the next memristor array as inputs. The subtraction of the differential column
421 pair is accomplished with the logics connected to their associated ADCs, and the subtraction
422 results are sampled simultaneously by counters and stored in output buffers. The error
423 calculation between the output and target is accomplished simultaneously with another set of
424 counters, which use the preloaded targets as their initial values.

425 The signs of the 1st-layer outputs (SY) and the results of comparing the 1st-layer outputs
426 (Y) with the reconfigurable value C are also obtained in the FWD stage. Since the memristor
427 chip utilizes a pulse-based coding scheme, this function can be realized with an AND logic
428 gate, which takes the 1st-layer outputs and the reconfigurable value as inputs. The
429 reconfigurable value signal is high during the C^{th} cycle of the 1st-layer output. In this cycle, the
430 AND logic generates the sign signals, which are stored in the registers.

431 The SET and RESET stages involve determining the weight update direction and
432 performing conductance tuning operations. The conductance tuning is performed in a row-by-
433 row fashion, with a set of shift registers selecting which row to update. The calculation of the
434 weight update direction involves the signs of the 1st-layer outputs, 2nd-layer outputs, and errors,
435 where the first two groups of signs are obtained in the FWD stage and stored in registers. The
436 first 8 bits with the highest significance of error are considered when calculating the weight
437 update direction via combination logic.

438 A 2T2R memristor array is proposed to reduce the IR drop by decreasing the SL output
439 current. Voltages with opposite polarities are applied to the two memristors in the differential
440 pair during inference, where the conductance of these two memristors represent positive and
441 negative weights, respectively. The current through the positive weights and negative weights
442 connected to the same SL can be locally cancelled out. This direct subtract in current reduces
443 the IR drop, which decreases the power consumption and boosts the computing parallelism.

444 The resolution-adjustable ADC consists of three submodules: an integrator, a comparator,
445 and a shared digital-to-analogue converters (DAC) for each layer. First, the integrator converts
446 the SL current to a voltage. Then, the DAC generates a voltage that increases with the clock.
447 Simultaneously, the comparator compares the output voltages of the integrator and DAC and
448 generates a timing-dependent quantized output based on the number of pulses. The resolution
449 of the ADC is configured by setting the clock frequency of the DAC and the sampling frequency.
450 The resolution-adjustable ADC enables a flexible trade-off between system accuracy and
451 power consumption.

452

453 **Fabrication of the memristor chip**

454 The memristor arrays are monolithically integrated with CMOS peripheral circuits on a
455 chip. Each memristor device is connected in series with a selector transistor, whose gate is
456 taken as the effective control terminal during programming. As shown in Fig. 3c, the memristor
457 device is sandwiched between metal-4 and metal-5 interconnects with a material stack of
458 TiN/HfO_x/TaO_y/TiN.

459 The selector transistor, peripheral circuits, and bottom four metal interconnections are
460 fabricated in a standard CMOS foundry by using a 130-nm process. The memristors and top
461 metal interconnections are fabricated in a laboratory with a back-end-of-line process. First, the
462 TiN bottom electrode layer, HfO_x switching layer, TaO_y capping layer, and TiN top electrode
463 layer are sequentially deposited on the wafers from the foundry. The capping layer serves as a
464 thermally enhanced layer to modulate the distributions of the electric field and temperature in
465 the switching layer and contributes to the improved analogue switching behaviour. The
466 deposited layers are patterned with lithography and etching to form isolated 0.7 μm × 0.7 μm
467 memristor cells. Afterwards, a SiO₂ dielectric is deposited, polished, patterned, and etched,
468 forming vias for interconnects. Finally, the top vias are formed by sputtering tungsten and
469 conducting chemical mechanical polishing, and the top metal interconnections are formed by
470 sputtering aluminium and shaping it, which completes the fabrication process.

471

472 **Measurements of the memristor devices**

473 The test system for the memristor chip mainly integrates a field-programmable gate array
474 (FPGA) and relevant voltage generators. The FPGA generates control commands for, sends

475 inputs to, and reads results from the memristor chip. The voltage generators offer varying
476 voltages for the VMM operation and memristor programming. The FPGA communicates with
477 the host PC that realizes the user interface through an Ethernet connection. To facilitate the
478 implementation of various learning tasks on the memristor chip, we map the high-level
479 functions on PC, provided as application programming interfaces (APIs), to the fundamental
480 operations of the memristor chip. The APIs include various programming schemes, inference,
481 or learning on a given dataset.

482 The memristor device has a high yield for 5-bit (32-level) programming, with a minimum
483 success rate of 99.69% across all the programmed conductance states (see Extended Data Fig.
484 4c). During the test, ~160K memristor devices are programmed to 32 target conductance states,
485 with ~5,000 devices assigned to each state. The conductance targets range within a switching
486 window from 2 μ S to 20 μ S with a uniform interval of 0.58 μ S. The margin is set as ± 0.24 μ S
487 for each target conductance state. The memristor conductance is sensed with on-chip ADCs at
488 a 0.2 V read voltage after each programming pulse. To accelerate the high-precision multibit
489 programming of memristor conductance, we develop a row-by-row program verification
490 scheme in a parallel fashion³⁸. This parallel programming scheme utilizes the incremental gate
491 voltage programming method, which solely varies the gate voltage of the selector transistor
492 during the conductance turning process. The maximum number of programming pulses is set
493 to 130 when 10 cells are programmed in parallel; that is, 13 pulses are required per cell on
494 average for 5-bit programming in the worst case.

495 Another 1K 1T1R memristor array on the same wafer with the neuro-inspired computing
496 chip is used to characterize the analogue switching behaviour of the memristor device. Sixty-

497 four memristor cells are measured for 5 cycles consisting of alternated potentiation and
498 depression. Potentiation refers to the evolution of conductance under an identical SET pulse
499 train, which contains 128 SET pulses with 1.5 V and 50-ns pulse widths. Depression refers to
500 the evolution of conductance under an identical RESET pulse train, which contains 128 RESET
501 pulses with 1.4 V and 50-ns pulse widths.

502

503 **Off-chip training and on-chip inference**

504 The neural networks used in the following tasks are quantized to ternary weights (+1, 0,
505 and -1) during off-chip training³⁹. We inject noise into the weights during off-chip training to
506 improve the resilience of the network against nonidealities⁴⁰. The noise level here is defined as
507 the ratio of the standard deviation of the programming error to the target conductance value.
508 Based on-chip characterization involving single-side verification, the noise level is
509 approximately 30%.

510 We first train the models with different levels of noise ranging from 0% to 50% and then
511 select the model with the best inference accuracy at the 30% noise level. Extended Data Fig.
512 1a shows that the highest accuracy is achieved by the model with comparable noise injection
513 during off-chip training. The weights of the off-chip-trained model are transferred to the chip
514 as the memristor conductance during the weight transfer stage. Binarized images with
515 resolutions of 28×28 pixels obtained from the MNIST dataset are used as the input images.
516 The inference is performed on 10,000 images derived from the entire test set. We investigate
517 the change in accuracy with time by using the same model. The inference is performed 5 times
518 each day on the entire test set within 48 days after weight transfer.

519

520 **Implementation of the on-chip learning task**

521 We also use the MNIST dataset to demonstrate the on-chip learning task. The 1st-layer
522 weights of the off-chip-trained model were transferred to the 1st-layer memristor array as
523 conductance. The 2nd-layer memristors are all programmed to the HRS via single-side
524 verification, and the programming noise works as the random initialization of the 2nd-layer
525 weights. The on-chip learning is performed on 180,000 images (3 epochs with the entire
526 training set containing 60,000 images). The inference is performed on 10,000 randomly
527 selected images from the training set and the 10,000 images forming the entire test set.

528

529 **Energy consumption benchmark**

530 We measure the latency and the power consumption of the memristor chip in the FWD,
531 SET, and RESET stages. The total energy consumption for one learning iteration is the sum of
532 the energy consumption of the forward (FWD) stage and weight update stage, where the energy
533 consumption of the weight update stage is the average energy consumption of the SET and
534 RESET stages. It should be mentioned that the energy consumed by the ADCs in the SET and
535 RESET stages is not considered because the ADCs can be powered off during these stages with
536 a good chip design. From the above measurements and calculations, we obtain that the energy
537 consumption of our memristor chip for each on-chip learning iteration is 1.002 μJ . The detailed
538 metrics, including the latency, power, and energy consumption of each stage in the learning
539 mode, are listed in Extended Data Table 1. The SET pulse width is 200 ns, while the RESET
540 pulse width is 50 ns. As a result, the delays of the SET stage and RESET stage are different.

541 The energy consumption of the CPU-based system (Intel i9-10900K) for one training
542 iteration is calculated as follows: thermal design power (125 W) \times CPU usage during training
543 (50.0%) \times run time for one training iteration (46.08 μ s) = 2880.0 μ J.

544 We also estimate the energy consumption of the GPU (NVIDIA K80)-based system for
545 one training iteration. The operations required for conducting one training iteration with the
546 GPU-based system include 1) loading all the weights from dynamic random-access memory
547 (DRAM) to the GPU, 2) performing forward and weight update-related operations with the
548 GPU, and 3) writing the 2nd-layer weights back to the DRAM. We assume floating-point (32-
549 bit) weights, which are usually required for neural network training. The energy consumption
550 of the GPU is estimated as the number of calculation operations divided by its energy
551 efficiency³⁶. The energy consumption of the DRAM is estimated as the number of
552 writing/reading bits multiplied by the energy consumed by each bit, which includes the DRAM
553 writing/reading energy and the on-board interconnection energy. The DRAM writing/reading
554 energy is assumed to be 20 pJ per bit based on the reported data for Graphics Double Data Rate,
555 version 5⁴¹ (GDDR5), which is adopted in the NVIDIA K80 GPU. The on-board
556 interconnection energy is assumed to be 10 pJ per bit⁴². The estimated energy consumed by the
557 GPU, DRAM, and GPU-based system is 8.67 μ J, 77.18 μ J, and 85.85 μ J, respectively.

558

559 **Motion control task of learning new samples**

560 We demonstrate the improvement learning of new samples in the motion control task of a
561 light-chasing car, whose control algorithm involves mapping a camera image to a control
562 decision (i.e., a steering angle and a driving throttle) directly in an end-to-end learning

563 approach⁴³. As illustrated in Extended Data Fig. 5a, a traditional approach is developed to
564 generate control decisions for teaching neural networks and evaluating their performance. The
565 traditional approach divides the control task into two parts, including laser spot detection and
566 control logic generation. The camera images and corresponding control decisions are collected
567 as datasets when the traditional approach drives the car. An end-to-end neural network is then
568 trained with the collected samples to directly turn the images into control decisions. The
569 obtained network is deployed on the car to drive it autonomously in real environments. In this
570 task, the steering angles are analogue signals that range from -1 to 1, corresponding to different
571 turning angles of the steering motor from the leftmost to the rightmost. The driving throttle is
572 a binary signal corresponding to the driving motor moving forward or stopping.

573 A multitask CNN including convolutional layers and two FC layers is constructed to
574 realize the control of the light-chasing car. The dimensions of the last two FC layers are $512 \times$
575 100×10 to fit the array size of the fabricated memristor chip. Eight of the ten outputs predict
576 a regression value for the continuous steering angle, while the other two outputs classify the
577 driving throttle, indicating whether to move or stop. The FC layers are implemented on the
578 chip, while the convolutional layers are implemented with software. The on-chip learning is
579 performed on 2,100 samples collected from the bright scene. The inference is performed on
580 4,502 samples acquired from the dark scene and 900 samples collected from the bright scene.

581 A customized scoring mechanism is developed to evaluate the performance of the neural
582 network. The rule used to score each action (with the steering angle and driving throttle
583 combined) is shown in Extended Data Fig. 5b. The score of the neural network is the
584 summation of the scores over the entire test dataset of the network normalized by that of the

585 traditional approach.

586

587 **Image classification task of learning a new class**

588 We demonstrate the improvement learning of a new class in the MNIST image
589 classification task, which involves images of the digits 0 and 2-9 for off-chip training and
590 images of the digit 1 for on-chip learning. A multilayer neural network with dimensions of 784
591 $\times 100 \times 9$ (Fig. 4f) is constructed during off-chip training, and this network is then transferred
592 to the memristor chip during the weight transfer stage. The neural network is expanded to
593 dimensions of $784 \times 100 \times 10$ with a newly added output neuron for the new class. The
594 expanded parts of the conductance weights are programmed to the HRS at the differential rows
595 for the new class by using single-side verification. Only the expanded conductance weights are
596 updated, while the old conductance weights are frozen during on-chip learning. The freezing
597 of the old conductance weights is realized by using the obtained inference results as targets for
598 the nine old output neurons. The on-chip learning is performed on 150 images randomly
599 selected from the training set of the new class (6,742 images). The inference is performed on
600 the test sets of the old classes and the new class at intervals of 10 training images; the inference
601 is conducted on 1,000 randomly selected images acquired from the test set of the old classes
602 and 1,135 images (the entire test set) of the new class.

603

604 **Audio recognition task of learning new samples**

605 We demonstrate the learning of new samples in a speech recognition task conducted on
606 the AudioMNIST dataset⁴⁴, which consists of male and female audio samples of spoken digits

607 (0-9). The dataset consists of 30,000 audio samples of spoken digits (0-9) provided by 48 men
608 and 12 women, with each speaker saying each digit 50 times. Mel-frequency cepstral
609 coefficient (MFCCs) are used as the features for the audio samples, while a two-layer
610 perceptron with dimensions of $392 \times 100 \times 10$ is used to classify the MFCC features. As
611 illustrated in Extended Data Fig. 6a, the improvement learning of female audio samples is
612 based on the model trained with male audio samples. First, the base model is off-chip trained
613 with the 6,000 samples of the first 12 men and then deployed on the memristor chip. Next, the
614 improvement learning of new samples is performed on the chip with 5,400 samples (the first
615 45 audio samples of each woman).

616 Extended Data Fig. 6b shows the evolution trends of the model accuracies for male and
617 female audio during improvement learning. The on-chip inference is performed on the 18,000
618 samples of the last 36 men and 600 samples of the last 5 audio samples of each woman. After
619 improvement learning with the female audio samples, the accuracy of the model for female
620 audio increases from 85.7% to 89.5%, while the accuracy of the model for male audio slightly
621 decreases from 91.4% to 91%.

622

623 **Motion control task of learning a new class**

624 We also demonstrate the learning of a new class in another motion control task involving
625 the abovementioned light-chasing car platform. In this task, improvement learning aims to
626 make the car learn to move backwards, and both the steering angle and driving throttle are
627 taken as the classification results of the neural network. As illustrated in Extended Data Fig.
628 7a, the improvement learning of the unreachable situation is based on the model developed for

629 the reachable situation. The unreachable situation refers to the data whose camera images have
630 a light spot at their two bottom corners. In this case, the car will not reach the target (i.e., light
631 spot) even though it turns the steering motor to the rightmost (or leftmost) setting and moves
632 forward until it stops. After improvement learning of the unreachable situation, whose driving
633 throttle is labelled “move backwards”, the car learns to move backwards to align itself to the
634 target in this case. The behaviours of the car before and after improvement learning are
635 illustrated in Extended Data Fig. 7b as the trajectories of the weight transfer and improvement
636 learning, respectively, and are shown in Supplementary Video 2. A sequence of camera images
637 corresponding to the trajectory after improvement learning is shown in Extended Data Fig. 7d.
638 The driving throttle classification accuracies achieved for the old classes and the new class
639 over the course of the improvement learning are shown in Extended Data Fig. 7c.

640

641 **Data availability**

642 The MNIST dataset³⁰ and AudioMNIST dataset⁴⁴ are publicly available. The dataset for
643 the light-chasing car tasks is available from the corresponding author upon reasonable request.
644 The light-chasing car used in this work is developed based on the open-source donkeycar
645 platform. The data that support the findings of this study are available from the corresponding
646 author upon reasonable request.

647

648 **Code availability**

649 The codes that support the findings of this study are available from the corresponding
650 author upon reasonable request.

651

652 **Acknowledgements**

653 This work is supported in part by the Ministry of Science and Technology (MOST) of
654 China (2021ZD0201200), the National Natural Science Foundation of China (92064001,
655 62025111), the XPLOERER Prize, the IoT Intelligent Microsystem Center of Tsinghua
656 University-China Mobile Joint Research Institute, and the Beijing Advanced Innovation Center
657 for Integrated Circuits.

658

659 **Author Contributions**

660 B.G., Y.W., and H.W. supervised this project and proposed the overall architecture. W.Z.,
661 B.G., and P.Y. contributed to the whole experiment design. P.Y., Q.L., D.W., and D.W. designed
662 the circuits. B.G., J.T., H.Q., and H.W. contributed to the memristor device fabrication and
663 integration with CMOS. W.Z., P.Y., and Q.L. contributed to the implementation of the test
664 system. Q.Z. worked on the development of AI models. W.Z. implemented the AI models on
665 the chip and conducted the chip and system measurements. W.Z., X.M., and Y.Z. performed
666 the simulation and benchmark. W.Z., B.G., and Y.P. contributed to the writing and editing of
667 the manuscript. All authors discussed the results and reviewed the manuscript.

668

669 **Competing Interests**

670 The authors declare no competing interests.

671

672 **Materials & Correspondence**

673 Correspondence to Bin Gao (gaob1@tsinghua.edu.cn) and Huaqiang Wu
674 (wuhq@tsinghua.edu.cn).

675

676 **References**

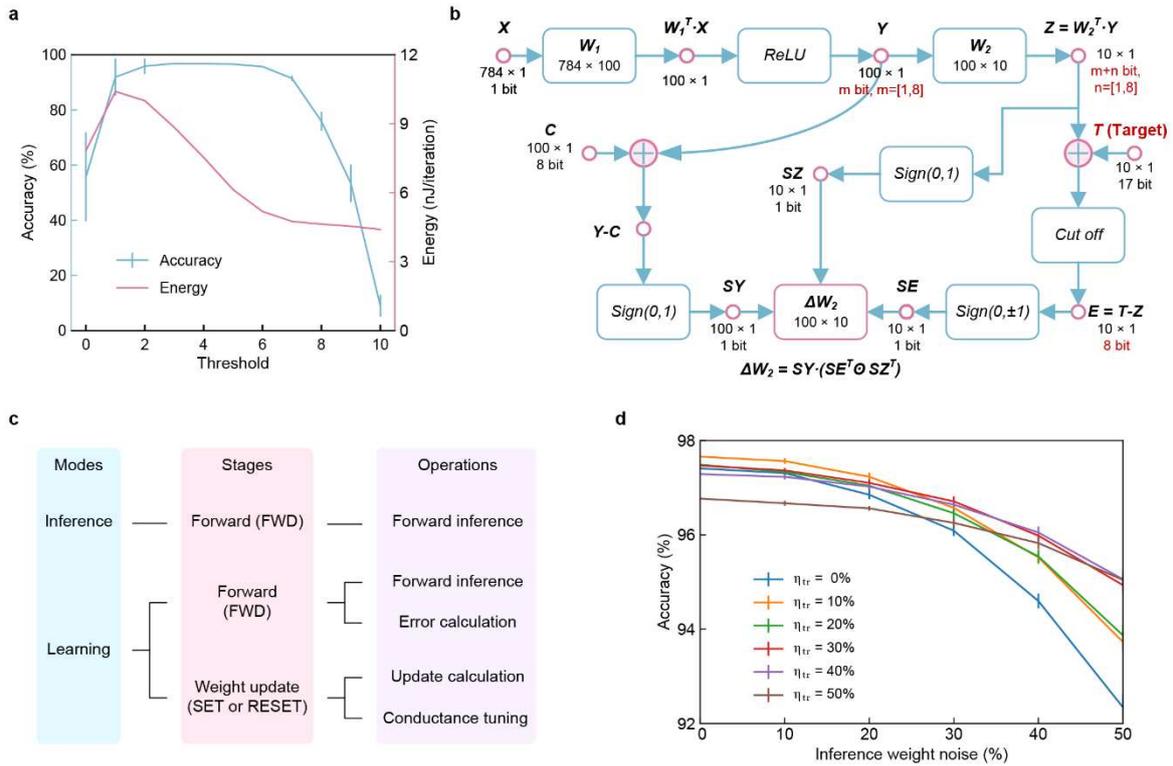
- 677 1. Shi, W., Cao, J., Member, S., Zhang, Q. & Member, S. Edge Computing : Vision and
678 Challenges. **3**, 637–646 (2016).
- 679 2. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-
680 propagating errors. *Nature* **323**, 533–536 (1986).
- 681 3. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
- 682 4. Coates, A. *et al.* Deep learning with COTS HPC systems. *30th Int. Conf. Mach. Learn.*
683 *ICML 2013* 2374–2382 (2013).
- 684 5. Horowitz, M. 1.1 Computing’s energy problem (and what we can do about it). in *2014*
685 *IEEE International Solid-State Circuits Conference Digest of Technical Papers*
686 *(ISSCC)* vol. 57 10–14 (IEEE, 2014).
- 687 6. Wong, H. S. P. & Salahuddin, S. Memory leads the way to better computing. *Nat.*
688 *Nanotechnol.* **10**, 191–194 (2015).
- 689 7. Strubell, E., Ganesh, A. & McCallum, A. Energy and policy considerations for modern
690 deep learning research. *AAAI 2020 - 34th AAAI Conf. Artif. Intell.* 1393–13696 (2020).
- 691 8. Ielmini, D. & Wong, H. S. P. In-memory computing with resistive switching devices.
692 *Nat. Electron.* **1**, 333–343 (2018).
- 693 9. Zidan, M. A., Strachan, J. P. & Lu, W. D. The future of electronics based on
694 memristive systems. *Nat. Electron.* **1**, 22–29 (2018).

- 695 10. Liu, Q. *et al.* 33.2 A Fully Integrated Analog ReRAM Based 78.4TOPS/W Compute-
696 In-Memory Chip with Fully Parallel MAC Computing. in *2020 IEEE International*
697 *Solid- State Circuits Conference - (ISSCC)* 500–502 (IEEE, 2020).
- 698 11. Wan, W. *et al.* 33.1 A 74 TMACS/W CMOS-RRAM Neurosynaptic Core with
699 Dynamically Reconfigurable Dataflow and In-situ Transposable Weights for
700 Probabilistic Graphical Models. in *2020 IEEE International Solid- State Circuits*
701 *Conference - (ISSCC)* 498–500 (IEEE, 2020).
- 702 12. Xue, C. *et al.* A CMOS-integrated compute-in-memory macro based on resistive
703 random-access memory for AI edge devices. *Nat. Electron.* **4**, 81–90 (2021).
- 704 13. Chen, W. H. *et al.* CMOS-integrated memristive non-volatile computing-in-memory
705 for AI edge processors. *Nat. Electron.* **2**, 420–428 (2019).
- 706 14. Ambrogio, S. *et al.* Equivalent-accuracy accelerated neural-network training using
707 analogue memory. *Nature* **558**, 60–67 (2018).
- 708 15. Sebastian, A., Le Gallo, M., Khaddam-Aljameh, R. & Eleftheriou, E. Memory devices
709 and applications for in-memory computing. *Nat. Nanotechnol.* **15**, 529–544 (2020).
- 710 16. Zhang, W. *et al.* Neuro-inspired computing chips. *Nat. Electron.* **3**, 371–382 (2020).
- 711 17. Prezioso, M. *et al.* Training and operation of an integrated neuromorphic network
712 based on metal-oxide memristors. *Nature* **521**, 61–64 (2015).
- 713 18. Yao, P. *et al.* Face classification using electronic synapses. *Nat. Commun.* **8**, 1–8
714 (2017).
- 715 19. Li, C. *et al.* Efficient and self-adaptive in-situ learning in multilayer memristor neural
716 networks. *Nat. Commun.* **9**, 7–14 (2018).

- 717 20. Wang, Z. *et al.* Reinforcement learning with analogue memristor arrays. *Nat. Electron.*
718 **2**, 115–124 (2019).
- 719 21. Li, C. *et al.* Long short-term memory networks in memristor crossbar arrays. *Nat.*
720 *Mach. Intell.* **1**, 49–57 (2019).
- 721 22. Yao, P. *et al.* Fully hardware-implemented memristor convolutional neural network.
722 *Nature* **577**, 641–646 (2020).
- 723 23. Cai, F. *et al.* A fully integrated reprogrammable memristor–CMOS system for efficient
724 multiply–accumulate operations. *Nat. Electron.* **2**, 290–299 (2019).
- 725 24. Burr, G. W. *et al.* Experimental Demonstration and Tolerancing of a Large-Scale
726 Neural Network (165 000 Synapses) Using Phase-Change Memory as the Synaptic
727 Weight Element. *IEEE Trans. Electron Devices* **62**, 3498–3507 (2015).
- 728 25. Wu, H. *et al.* Device and circuit optimization of RRAM for neuromorphic computing.
729 in *2017 IEEE International Electron Devices Meeting (IEDM)* 11.5.1-11.5.4 (IEEE,
730 2017).
- 731 26. Chen, P. Y., Peng, X. & Yu, S. NeuroSim: A circuit-level macro model for
732 benchmarking neuro-inspired architectures in online learning. *IEEE Trans. Comput.*
733 *Des. Integr. Circuits Syst.* **37**, 3067–3080 (2018).
- 734 27. Xi, Y. *et al.* In-memory Learning with Analog Resistive Switching Memory: A
735 Review and Perspective. *Proc. IEEE* **109**, 14–42 (2021).
- 736 28. Zhang, Q. *et al.* Sign backpropagation: An on-chip learning algorithm for analog
737 RRAM neuromorphic computing systems. *Neural Networks* **108**, 217–223 (2018).
- 738 29. Wu, W. *et al.* A Methodology to Improve Linearity of Analog RRAM for

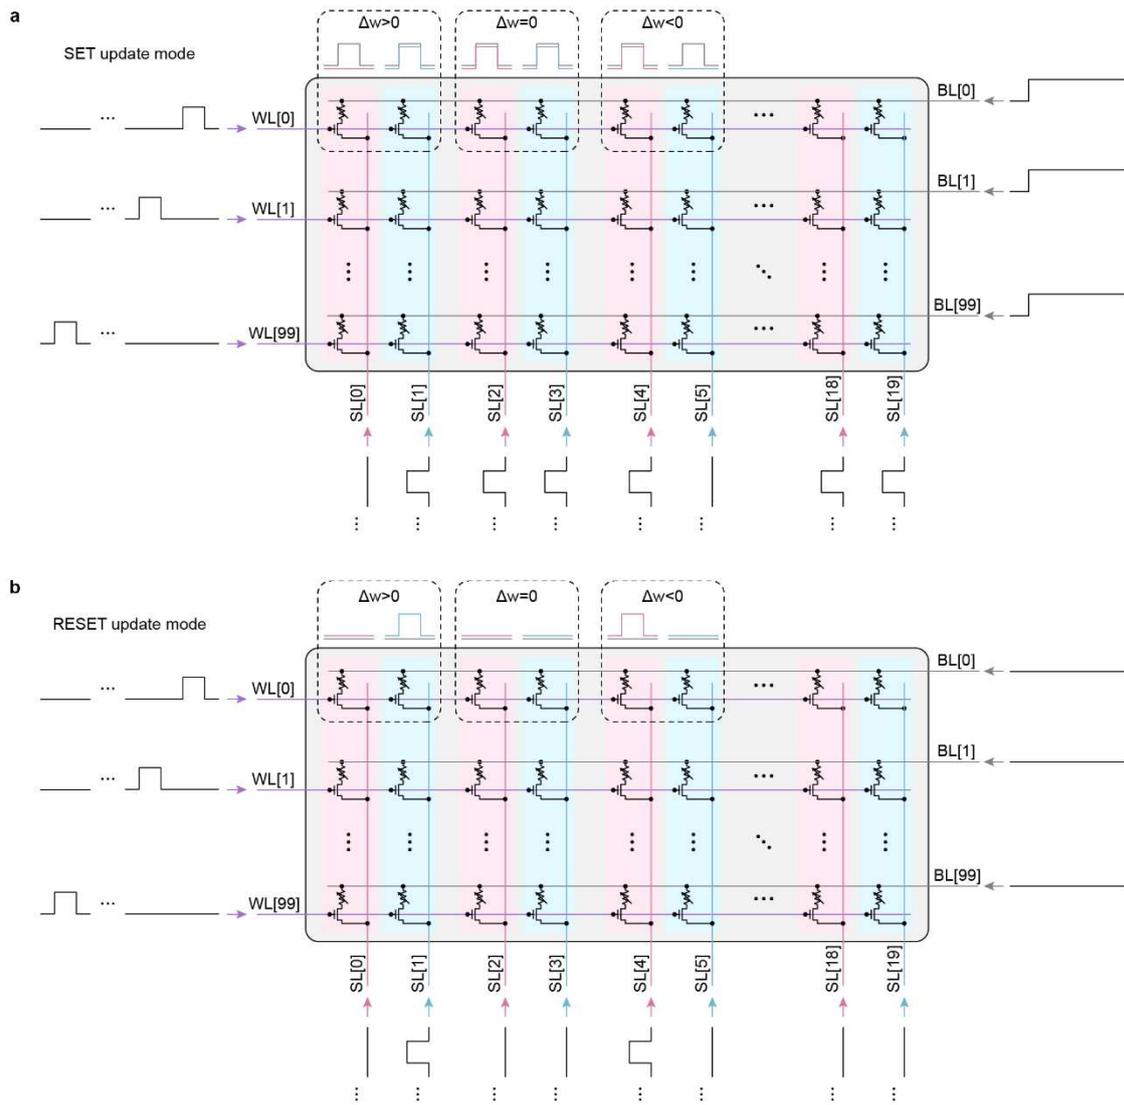
- 739 Neuromorphic Computing. in *2018 IEEE Symposium on VLSI Technology* vols 2018-
740 June 103–104 (IEEE, 2018).
- 741 30. Bengio, Y. & Haffner, P. Gradient-Based Learning Applied to Document Recognition.
742 **86**, (1998).
- 743 31. Chi, P. *et al.* PRIME: A Novel Processing-in-Memory Architecture for Neural
744 Network Computation in ReRAM-Based Main Memory. *Proc. - 2016 43rd Int. Symp.*
745 *Comput. Archit. ISCA 2016* 27–39 (2016).
- 746 32. Song, L., Qian, X., Li, H. & Chen, Y. PipeLayer: A Pipelined ReRAM-Based
747 Accelerator for Deep Learning. in *2017 IEEE International Symposium on High*
748 *Performance Computer Architecture (HPCA)* 541–552 (IEEE, 2017).
- 749 33. Khaddam-Aljameh, R. *et al.* HERMES Core – A 14nm CMOS and PCM-based In-
750 Memory Compute Core using an array of 300ps/LSB Linearized CCO-based ADCs
751 and local digital processing. in *2021 Symposium on VLSI Circuits* 1–2 (IEEE, 2021).
- 752 34. Wan, W. *et al.* A Voltage-Mode Sensing Scheme with Differential-Row Weight
753 Mapping for Energy-Efficient RRAM-Based In-Memory Computing. in *2020 IEEE*
754 *Symposium on VLSI Technology* vols 2020-June 1–2 (IEEE, 2020).
- 755 35. Hung, J.-M. *et al.* A four-megabit compute-in-memory macro with eight-bit precision
756 based on CMOS and resistive random-access memory for AI edge devices. *Nat.*
757 *Electron.* **4**, 921–930 (2021).
- 758 36. Jouppi, N. P. *et al.* In-Datacenter Performance Analysis of a Tensor Processing Unit.
759 in *Proceedings of the 44th Annual International Symposium on Computer Architecture*
760 vol. Part F1286 1–12 (ACM, New York, NY, USA, 2017).

- 761 37. Oh, T., Maghari, N. & Moon, U. K. A 5MHz BW 70.7dB SNDR noise-shaped two-
762 step quantizer based $\Delta\Sigma$ ADC. *IEEE Symp. VLSI Circuits, Dig. Tech. Pap.* 162–163
763 (2012).
- 764 38. Chen, J. *et al.* A parallel multibit programming scheme with high precision for RRAM-
765 based neuromorphic systems. *IEEE Trans. Electron Devices* **67**, 2213–2217 (2020).
- 766 39. Li, F., Zhang, B. & Liu, B. Ternary Weight Networks. (2016).
- 767 40. Joshi, V. *et al.* Accurate deep neural network inference using computational phase-
768 change memory. *Nat. Commun.* **11**, 1–13 (2020).
- 769 41. O’Connor, M. Highlights of the High- Bandwidth Memory (HBM) Standard. *Mem.*
770 *Forum* (2014).
- 771 42. Arunkumar, A. *et al.* MCM-GPU: Multi-Chip-Module GPUs for Continued
772 Performance Scalability. in *Proceedings of the 44th Annual International Symposium*
773 *on Computer Architecture* vol. 45 320–332 (ACM, New York, NY, USA, 2017).
- 774 43. Bojarski, M. *et al.* End to End Learning for Self-Driving Cars. 1–9 (2016).
- 775 44. Becker, S., Ackermann, M., Lapuschkin, S., Müller, K.-R. & Samek, W. Interpreting
776 and Explaining Deep Neural Networks for Classification of Audio Signals. 2–6 (2018).
- 777
- 778



779
 780
 781
 782
 783
 784
 785
 786
 787
 788
 789

Extended Data Fig. 1 | Details of the algorithm for the STELLAR architecture and the results of noise-aware training. **a**, The role of the threshold in the STELLAR algorithm, showing the effects of different thresholds on classification accuracy and energy consumption of on-chip learning. **b**, Schematic of the modified algorithm for chip implementation. m and n , ranging from 1 to 8, represent the resolutions of the ADCs of the 1st and 2nd layers, respectively. **c**, Dendrogram of the chip's working modes and the corresponding stages and operations. The forward stage involves a forward inference operation and an error calculation operation. The weight update stage involves a weight update calculation and a conductance tuning operation. **d**, Classification accuracies as a function of the inference weight noise for the off-chip-trained models with different levels of weight noise injected during training.

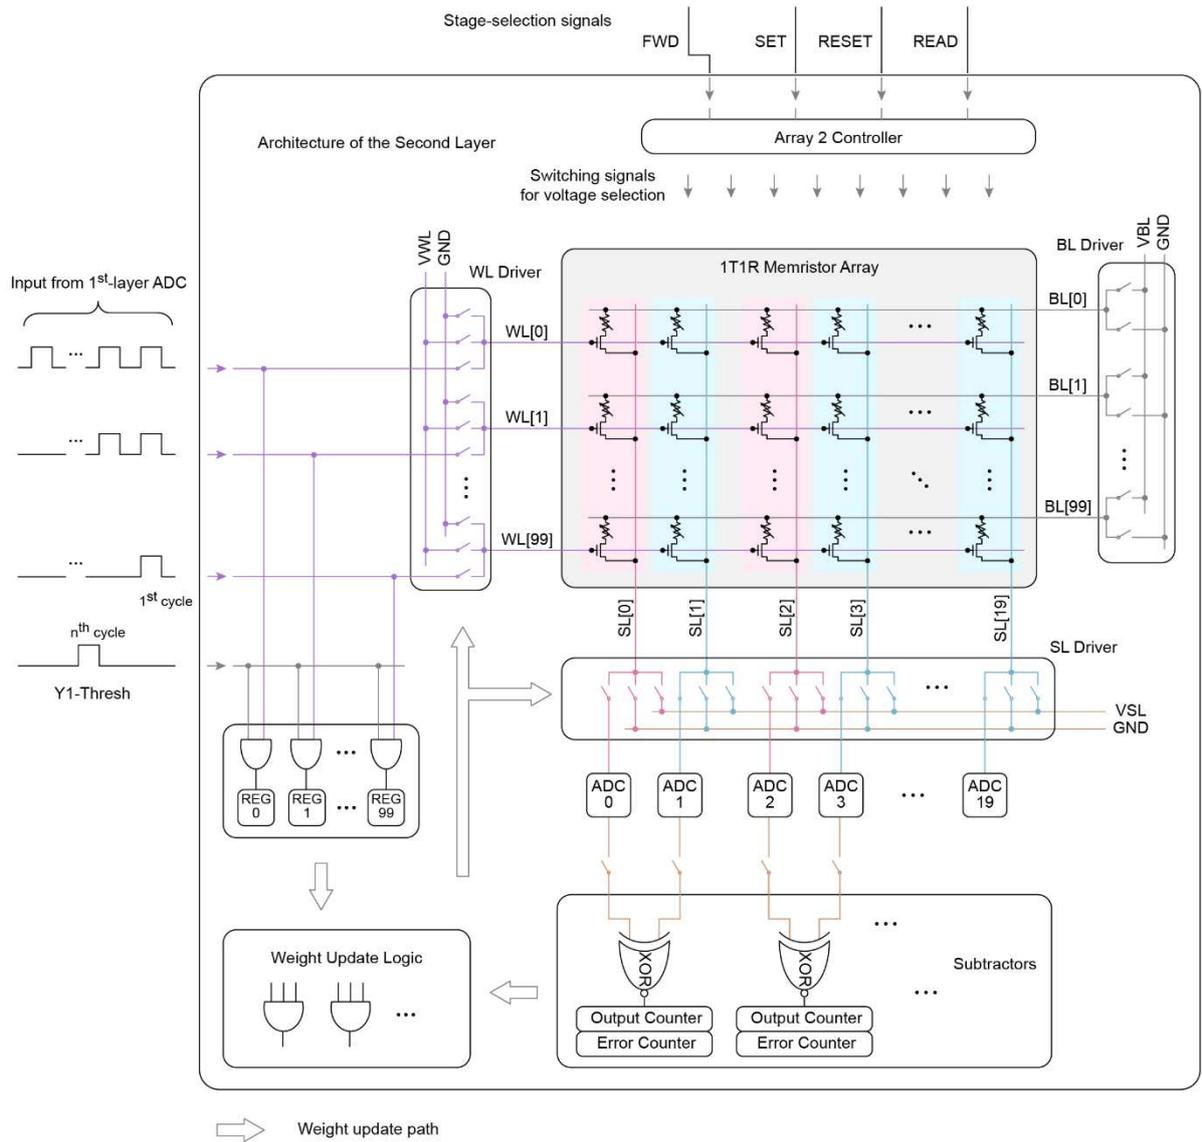


790

791 **Extended Data Fig. 2 | Schematics of array operations during update modes.** Schematics of array

792 operations in SET update modes (a) and RESET update modes (b). The WL signals select which row to

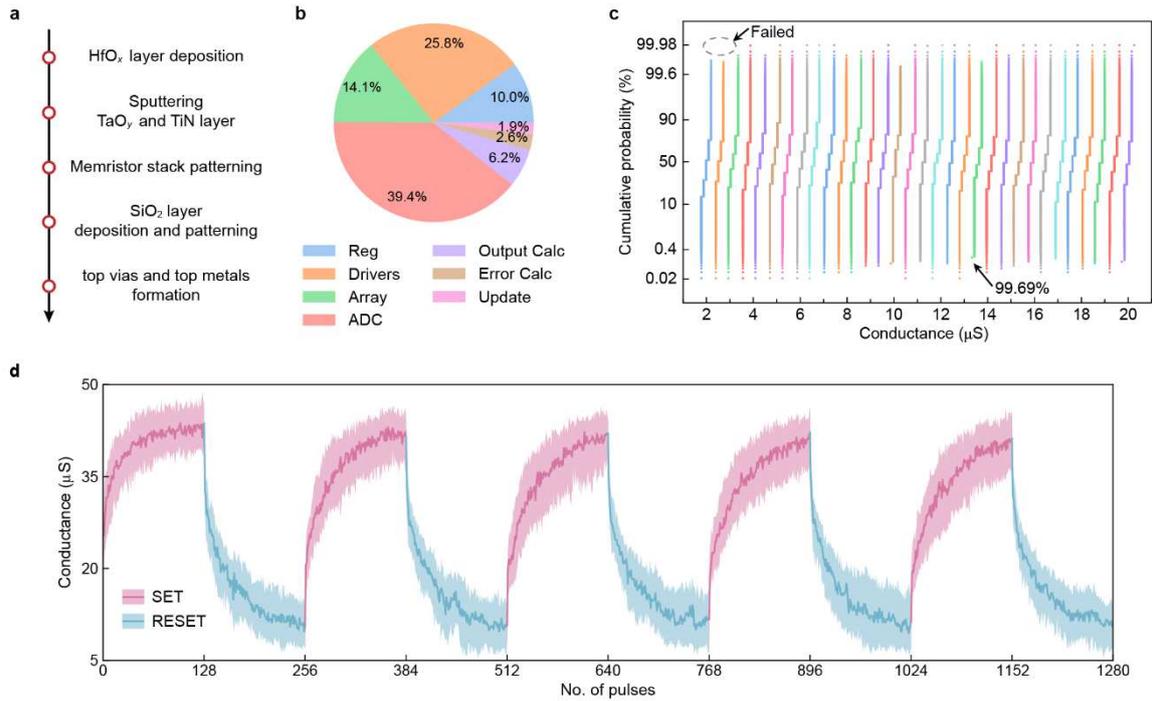
793 operate in turns, while the SL signals determine which memristors to operate on the selected row.



794

795 **Extended Data Fig. 3 | Architecture of the 2nd layer of the memristor chip.**

796



797

798 **Extended Data Fig. 4 | Fabrication and characterization of the memristor devices.** **a**, Process flow of

799 the fabrication of the memristor devices. **b**, Chip area breakdown. **c**, Cumulative probability distribution of

800 ~160K memristors after programming them to 32 conductance states. The memristors outside of the target

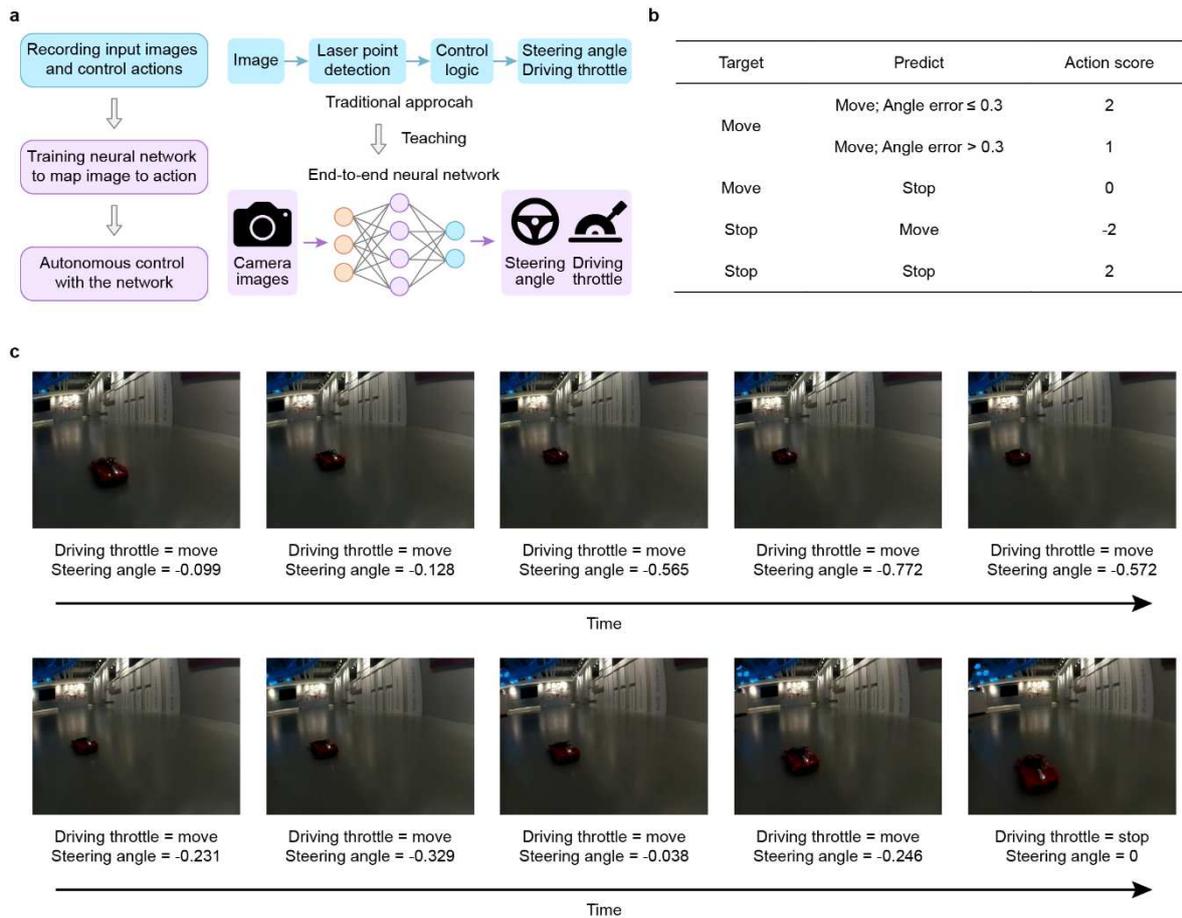
801 range are not shown, as marked by the grey circle. The minimum success rate among the 32 conductance

802 states is 99.69%. **d**, Analogue switching behaviours of 64 memristors under identical pulse trains. The dark

803 lines represent the median values of the conductance, while the light colours fill the regions between the

804 lower quartiles and the upper quartiles.

805



806

807 **Extended Data Fig. 5 | Motion control task of learning new samples. a**, Diagram of the realization of

808 end-to-end autonomous control. **b**, Table of the rule scores for different control decisions under the

809 consideration of both the steering angle and the driving throttle. Taking the first row as an example, for a

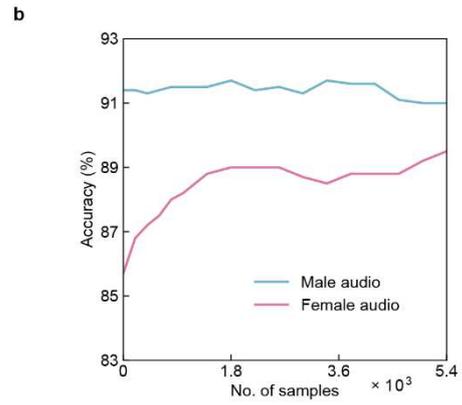
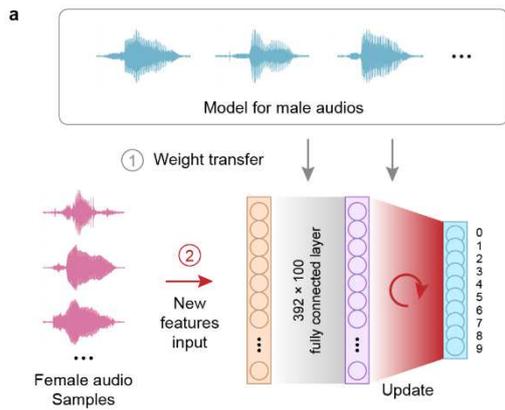
810 certain network prediction, if the predicted throttle and the corresponding labelled throttle are both “move”

811 and the difference between the predicted angle and the labelled angle is no larger than 0.3, the score of this

812 predicted control decision is 2. **c**, A sequence of input camera images and their corresponding control

813 decisions produced during real-time driving by the network deployed on the memristor chip.

814



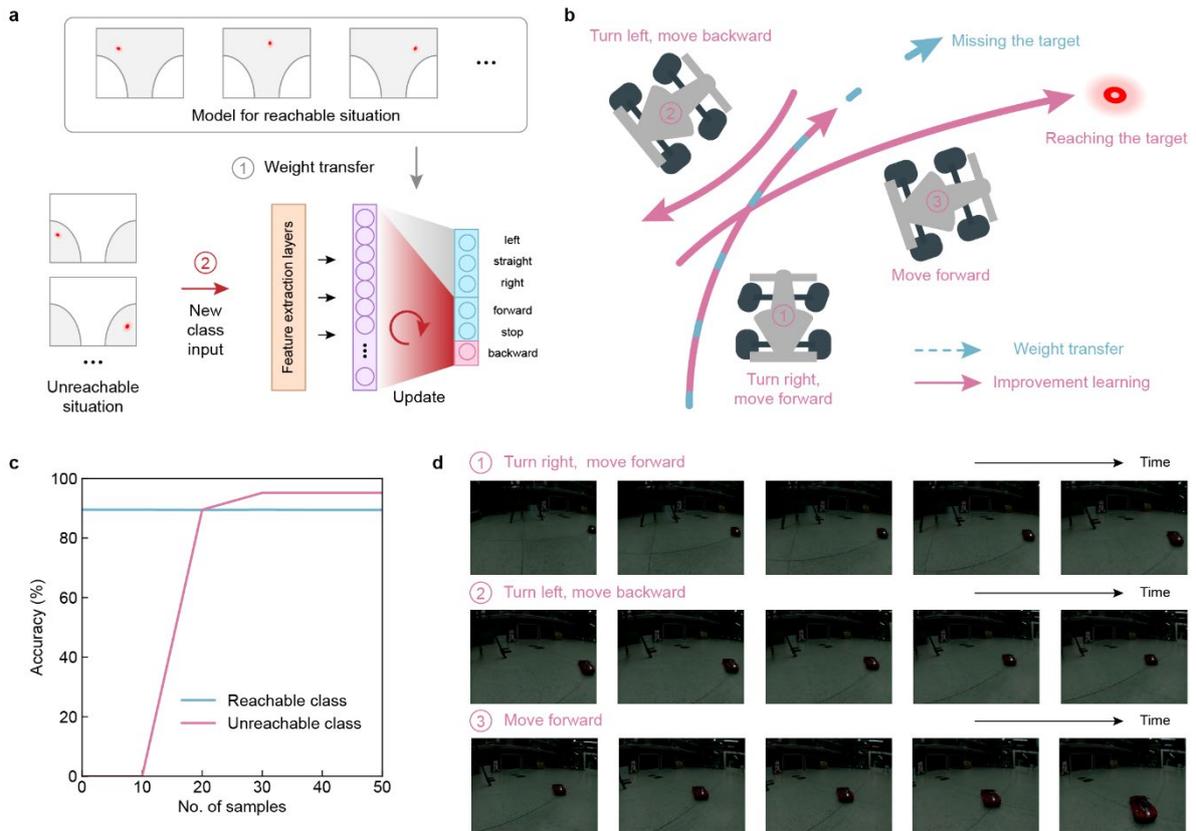
815

816 **Extended Data Fig. 6 | Audio recognition task of learning new samples. a,** Illustration of the improvement

817 learning of new samples in the audio recognition task. **b,** Changes in the recognition accuracies achieved for

818 male audio and female audio over the course of the on-chip improvement learning female audio samples.

819



820

821 **Extended Data Fig. 7 | Motion control task of learning a new class. a**, Illustration of the improvement

822 learning of new class in the light-chasing car task. **b**, Illustration of the car's movements and trajectories

823 before and after improvement learning of the new class. Before improvement learning, the car moves forward

824 even though it will miss its target. After improvement learning, the car moves backwards to align itself with

825 the target and then moves forward to reach the target. **c**, Changes in the throttle classification accuracies for

826 the old classes and the new class over the course of the on-chip improvement learning the new class. **d**. A

827 sequence of camera images corresponding to the trajectory produced after improvement learning and shown

828 in **(b)**. The car first moves forward until it will miss the target, then moves backwards to align itself with the

829 target, and finally moves forward again to reach the target.

830

Extended Data Table 1 | Detailed metrics produced during each stage of the learning mode.

	Delay/μs	Power/mW	Energy/mJ
FWD stage	14.85	54.64	811.3
SET stage	85.95	2.49	213.7
RESET stage	55.95	3.01	168.1

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SupplementaryVideo1.mp4](#)
- [SupplementaryVideo2.mp4](#)
- [SupplementaryVideo3.mp4](#)