

1 Evaluation and comparison of methods for

2 neuronal parameter optimization using the

3 Neuroptimus software framework

4 Máté Mohácsi^{1,2}, Márk Patrik Török^{1,2}, Sára Sáray^{1,2}, Luca Tar^{1,2}, Szabolcs Káli^{1,2,*}

5 ¹ Institute of Experimental Medicine, Budapest, Hungary

6 ² Faculty of Information Technology and Bionics, Pázmány Péter Catholic University,
7 Budapest, Hungary

8

9 * Corresponding author

10 Email: kali.szabolcs@koki.hun-ren.hu (SK)

11

12 Abstract

13 Finding optimal parameters for detailed neuronal models is a ubiquitous challenge in
14 neuroscientific research. Recently, manual model tuning has been replaced by automated
15 parameter search using a variety of different tools and methods. However, using most of these
16 software tools and choosing the most appropriate algorithm for a given optimization task
17 require substantial technical expertise, which prevents the majority of researchers from using
18 these methods effectively. To address these issues, we developed a generic platform (called
19 Neuroptimus) that allows users to set up neural parameter optimization tasks via a graphical
20 interface, and to solve these tasks using a wide selection of state-of-the-art parameter search
21 methods implemented by five different Python packages. Neuroptimus also offers several
22 features to support more advanced usage, including the ability to run most algorithms in
23 parallel, which allows it to take advantage of high-performance computing architectures. We
24 used the common interface provided by Neuroptimus to conduct a detailed comparison of more
25 than twenty different algorithms (and implementations) on six distinct benchmarks that
26 represent typical scenarios in neuronal parameter search. We quantified the performance of the
27 algorithms in terms of the best solutions found and in terms of convergence speed. We identified
28 several algorithms, including covariance matrix adaptation evolution strategy and particle
29 swarm optimization, that consistently found good solutions in all of our use cases. By contrast,
30 some other algorithms including all local search methods provided good solutions only for the
31 simplest use cases, and failed completely on more complex problems. Finally, we created an
32 online database that allows uploading, querying and analyzing the results of optimization runs
33 performed by Neuroptimus, which enables all researchers to update and extend the current

34 benchmarking study. The tools and analysis we provide should aid members of the neuroscience
35 community to apply parameter search methods more effectively in their research.

36 **Author summary**

37 Model fitting is a widely used method in scientific research. It involves tuning the free
38 parameters of a model until its output best matches the corresponding experimental data.
39 Finding the optimal parameter combination can be a difficult task for more complex models
40 with many unknown parameters, and a large variety of different approaches have been proposed
41 to solve this problem. However, setting up a parameter search task and employing an efficient
42 algorithm for its solution requires considerable technical expertise. We have developed a
43 software framework that helps users solve this task, focusing on the domain of detailed models
44 of single neurons. Our open-source software, called Neuroptimus, has a graphical interface that
45 guides users through the steps of setting up a parameter optimization task, and allows them to
46 select from more than twenty different algorithms to solve the problem. We have also compared
47 the performance of these algorithms on a set of six parameter search tasks that are typical in
48 neuroscience, and identified several algorithms that delivered consistently good performance.
49 Finally, we designed and implemented a website that allows users to view and analyze our
50 results and to add their own results to the database.

51

52 Introduction

53 The construction and simulation of data-driven models has become a standard tool in
54 neuroscience [1, 2, 3]. Such models can be employed, among other things, to consolidate the
55 knowledge obtained from various experimental approaches into a common framework, to test
56 the consistency of the data, and to make novel predictions by examining the response of the
57 model to arbitrary inputs and by applying clean manipulations. Models at a given level of
58 description (e.g., individual neurons) can also be combined to form models of entities at higher
59 levels (such as networks) and thus aid the mechanistic understanding of emergent phenomena.

60 Nevertheless, these data-driven models often contain parameters that are not directly
61 constrained (or are only weakly constrained) by the available experimental data. Traditionally,
62 such unknown parameters were often tuned manually to adjust the behavior of the model
63 towards some desired target. However, this approach is typically inefficient, not quantitative,
64 and may be heavily biased to reproduce a few selected experimental results at the expense of
65 other relevant data. Consequently, in recent years, automated parameter search has emerged as
66 the preferred method for the estimation of unknown parameters of neural models [4, 5, 6, 7, 8,
67 9, 10, 11, 12, 13, 14, 15, 16]. This approach requires the definition of an error function (or cost
68 function) that measures the quality of the model with a given set of parameters, often in terms
69 of how well it approximates data obtained using a particular experimental protocol. The goal
70 of parameter optimization is then to find the set of parameters that minimizes the selected cost
71 function. The difficulty of this task can vary widely depending on the nature and complexity of
72 the model, the definition of the error function (or multiple error functions representing different
73 goals, or objectives), and the number of unknown parameters. Simple optimization problems
74 can be solved effectively by traditional gradient-based, local methods or by random search, but

75 these approaches tend to fail when there are many unknown parameters and the cost function
76 has multiple local minima [16, 17]. In fact, no algorithm is guaranteed to find the globally
77 optimal parameter combination in a short time for all problems [18], and various clever search
78 methods (called metaheuristics) have been proposed that often find good solutions in an
79 acceptable amount of time by taking advantage of various types of regularities in the cost
80 function [19].

81 Previous studies in neuroscience have used a variety of different software tools and
82 algorithms to perform parameter optimization. The general-purpose neural simulators
83 NEURON [20] and GENESIS [21] both include implementations of a few selected methods
84 that are adequate for certain parameter search tasks. In addition, several tools have been
85 developed specifically for neural parameter optimization, including Neurofitter [22],
86 BluePyOpt [23], pypet [24], and NeuroTune [25], and some more general computational
87 neuroscience tools such as NetPyNE [26] also have some support for parameter optimization.
88 However, most of these tools rely on a very limited set of parameter search methods, which
89 typically does not include many optimization algorithms that represent the state of the art in
90 global optimization and are popular in other fields of science and engineering. These new
91 methods were not included in any previous surveys of neural optimization. Systematic
92 comparisons of the existing neural optimization software tools have also been quite limited
93 [15]. Therefore, it is currently unknown which parameter search methods can be expected to
94 perform well in the parameter optimization tasks that are typical in neuroscience.

95 Furthermore, most of the existing tools for neural optimization lack any intuitive user
96 interface, and require substantial programming experience. One exception is our earlier
97 optimization software called Optimizer [15], which included a graphical user interface (GUI)
98 that was designed to guide users through the process of setting up, running, and evaluating the
99 results of a neuronal parameter optimization task. Optimizer also provided four different

100 optimization algorithms in two different Python packages, and was designed in a modular way
101 to facilitate the integration of new components including additional optimization algorithms.

102 The goal of the current study was twofold. First, we aimed to provide a general software
103 framework that allows the straightforward application of a large variety of state-of-the-art
104 parameter optimization methods to typical problems in data-driven neural modeling. This was
105 accomplished by significantly updating and extending our software tool (which is now called
106 Neuroptimus). Second, we aimed to perform a systematic comparison of parameter search
107 methods (including both previously used and novel algorithms) in the context of modeling
108 single neurons, which is probably the most common subtype of parameter optimization tasks in
109 neuroscience. To this end, we designed and implemented a test suite of neuronal parameter
110 optimization problems, and used Neuroptimus to systematically test the performance of a large
111 set of optimization algorithms on each of these benchmarks. The results of the different
112 algorithms on the test suite were systematically analyzed and compared. Finally, we designed
113 and deployed a web-accessible database that contains all the results of this study and also allows
114 users to upload, retrieve, and analyze the results of parameter optimization.

115

116 **Results**

117 The systematic evaluation of parameter optimization methods in the context of neuronal
118 modeling required the development of several interrelated methods and tools, which are
119 described in detail in the Methods section and whose main features are also summarized below.
120 The first necessary ingredient was a software tool that allows users to set up, execute, and
121 evaluate the results of a wide variety of neural parameter optimization problems in a single
122 standardized framework. The second required component was a diverse set of benchmark

123 problems that differ in the type of the model, the number of unknown parameters, and the
124 complexity of the error function, and that collectively cover many types of parameter fitting
125 problems that are often encountered in neuronal modeling. The third necessary component was
126 a set of methods that allows the consistent evaluation and comparison of optimization results
127 across the different benchmarks and algorithms. Finally, the last ingredient was a web-
128 accessible database of the optimization results that allows us to share all of our results publicly
129 and also enables us as well as other researchers to extend the study with additional optimization
130 runs and even new benchmarks.

131 **The Neural Optimization User Interface (Neuroptimus)**

132 We began our study by updating, improving and extending our previously developed
133 optimization software (Optimizer), which was already shown to be a useful tool for neuronal
134 optimization [15]. The new version (named Neuroptimus) inherited many useful features from
135 its predecessor, and added several important new capabilities. Both Optimizer and Neuroptimus
136 support the definition and solution of neural optimization problems through a graphical user
137 interface (GUI) that guides the users throughout the process. The main steps (represented by
138 different tabs in the GUI) involve selecting the target data, selecting the model and the
139 parameters to be optimized, setting up the simulations (including stimulation and recording
140 parameters), defining the cost function, selecting the optimization algorithm, running the
141 parameter search, and reviewing the results. A detailed guide to the GUI is available in the
142 online documentation of Neuroptimus (<https://neuroptimus.readthedocs.io/en/latest/>). All the
143 functionality is also accessible through a command line interface that uses configuration files
144 to set up the optimization, which enables batch processing (e.g., multiple runs with different
145 settings or random seeds). Simulations of the model can be performed either by the NEURON
146 simulator [20] (which is handled internally) or by arbitrary external code (which may include

147 running other simulators) handled as a “black box”. The modular, object-oriented structure of
148 the program makes it possible to extend its capabilities by adding new error functions and
149 optimization algorithms.

150 Neuroptimus includes several new and enhanced features compared to Optimizer. In
151 addition to specific time series (such as voltage traces), it is now also possible to use as target
152 data the statistics of features extracted (e.g., using the feature extraction module eFEL, [27])
153 from a set of experimental recordings. In this case, Neuroptimus uses eFEL to extract the same
154 features from each simulated model, computes feature errors as the difference between the
155 feature value of the model and the mean value of the experimental feature, normalized by the
156 experimental standard deviation, and uses the sum of these feature errors as the cost function
157 during parameter optimization. Weights can also be provided individually for each error
158 component.

159 While Optimizer provided four different search algorithms (two local and two global
160 algorithms implemented by the Inspyred and Scipy packages), Neuroptimus currently supports
161 more than twenty different optimization algorithms from five external Python packages (see
162 Table 2 for a complete list), plus an internally implemented random sampling algorithm, which
163 can be considered as a simple baseline method.

164 Neuroptimus also contains many enhancements “under the hood”. The new version was
165 entirely developed in Python 3 to support recent open-source Python modules, such as search
166 algorithms, graphical and parallelization interfaces. The graphical user interface was
167 completely re-implemented using the PyQt5 package, which provides a Python binding to the
168 popular cross-platform GUI toolkit Qt. In addition to the parameter search methods offered by
169 Scipy and Inspyred, Neuroptimus now also provides an interface to the algorithms implemented
170 by the widely used optimization packages Pygmo and BluePyOpt, as well as an additional

171 parallelized Python implementation of the Covariance Matrix Adaptation Evolution Strategy
172 (CMA-ES) algorithm. For many of these search algorithms, parallel evaluation of models is
173 also supported and easily configurable, which can lead to a manifold reduction in execution
174 time, especially on highly parallel architectures such as compute clusters and supercomputers.

175

176 **Neural optimization benchmarks**

177 We defined and implemented a test suite of different neuronal optimization problems to
178 demonstrate the utility of our Neuroptimus software and to quantitatively evaluate and compare
179 the effectiveness of different parameter optimization algorithms. Our aim was to identify which
180 parameter search methods (and which implementations) are able to find good solutions to each
181 of our benchmark problems, and which methods (if any) can provide consistently good
182 performance across all of these tasks. Our benchmarking use-cases differ in the complexity of
183 the models, the simulation protocol, the source and nature of the target data, the features and
184 error functions used to evaluate the model, and the number of unknown parameters. A subset
185 of our use-cases is analogous to those that were described by Friedrich et al. [15], although
186 some of these have been updated to improve their robustness. Each of the six benchmark
187 problems is described briefly below, and in more detail in the Methods section.

188 Four of the use cases involve finding the biophysical parameters of compartmental
189 models of neurons based on somatic voltage responses; however, these models differ greatly in
190 terms of the level of morphological and biophysical detail, and also in the number of unknown
191 parameters (between 3 and 12). One simple use case involves the classic single-compartment
192 Hodgkin-Huxley model with two voltage-gated conductances and a leak conductance; one uses
193 a morphologically detailed but passive model neuron; another benchmark optimizes the somatic

194 conductances of several voltage-gated ion channels in a simplified (6-compartment) model,
195 while our most complex use case involves fitting spatially varying conductance densities for a
196 large set of ion channels in a fully detailed compartmental model of a hippocampal pyramidal
197 cell. A different type of benchmark involves optimizing the parameters of a phenomenological
198 point neuron (an adaptive exponential integrate-and-fire model), and the final one simulates a
199 voltage-clamp experiment to estimate synaptic parameters.

200 Some of our benchmark problems (the Hodgkin-Huxley and the Voltage Clamp use-
201 cases) use surrogate data as the target. In this case, target data are generated by the same
202 neuronal model with known parameters; some of these parameters are then considered to be
203 unknown, and the task is to reconstruct the correct values. Therefore, in these test cases, a
204 perfect solution with zero error is known to exist, and the corresponding parameters can be
205 compared to those found by the search algorithms. However, for most of our benchmark
206 problems, the target data were recorded in electrophysiological experiments, or (in one case)
207 generated by a more complex model than the one we were fitting. In these instances, the best-
208 fitting parameters and the minimal possible error score are unknown.

209 In most of our use cases we compared the output of the model to the target data by
210 extracting several different electrophysiological features from the raw voltage traces. The
211 difference of each model feature from the corresponding (mean) experimental feature can be
212 considered as a separate error component (or objective). This allowed the direct application of
213 multi-objective optimization methods. When using single-objective algorithms, feature errors
214 were combined into a single cost function using an average with pre-defined (in most cases,
215 uniform) weights. The final best solution for multi-objective algorithms was also chosen using
216 the same weighted average of the objectives. Two of our use cases had simpler voltage or
217 current traces as their target. In these cases, the mean squared difference between the model

218 trace and the experimental trace was used as the only error function. This precluded the use of
219 multi-objective optimization methods, so only single-objective algorithms were included in the
220 comparison in these cases.

221 We used several criteria to select optimization algorithms for inclusion in our
222 benchmark study. First, we implemented a simple random search algorithm based on
223 independently repeated uniform sampling of the entire available search space defined by the
224 parameter boundaries. This algorithm can be considered as a natural baseline against which we
225 can measure the performance of more sophisticated methods. Second, we included some
226 popular local optimization algorithms (Nelder-Mead and L-BFGS-B) that are expected to be
227 efficient when the error function has a single optimum, but not for more complex problems with
228 multiple local optima. The rest of the search algorithms that we included are so-called global
229 optimization methods or meta-heuristics, which aim to take advantage of certain types of
230 regularities in the error function to find the global optimum (or another similarly good solution)
231 more efficiently than a random search or local optimization methods do. A very large selection
232 of such meta-heuristic algorithms has been developed, and many of these are included in one
233 (or several) of the Python packages that are accessible in Neuroptimus. Due to time and resource
234 constraints, not all of these algorithms were included in the current study, but we aimed to
235 include many of the algorithms that were previously used in neuronal optimization and those
236 that have proved particularly successful in other settings. More specifically, we included several
237 different types of evolutionary algorithms, several implementations of the particle swarm
238 algorithm, and also some other types of bioinspired algorithms and methods based on statistical
239 physics.

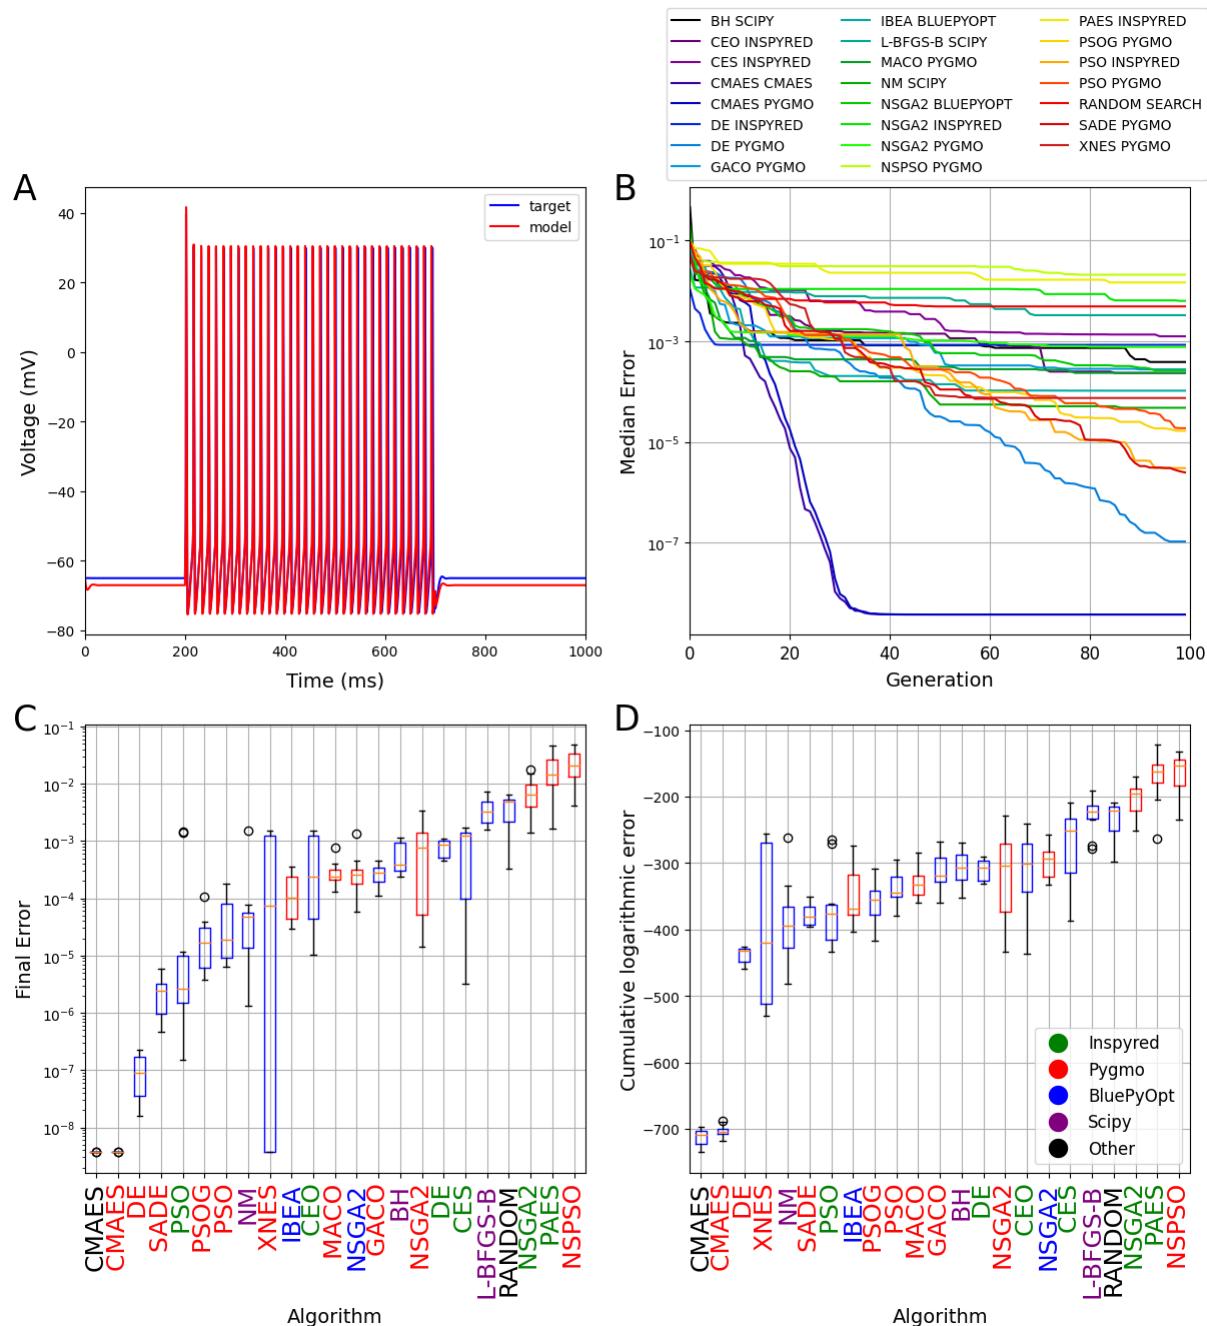
240 To ensure a fair comparison of different search methods, we allowed a maximum of
241 10,000 model evaluations in a single run of every optimization algorithm. For all the algorithms

242 that define populations of models that are evaluated as a batch in every iteration (this includes
243 both evolutionary and swarm algorithms), we set the population size to 100, and ran the
244 algorithms for 100 iterations (generations). We recorded the lowest error value achieved during
245 each run, and also looked at how the best error score evolved during the course of the
246 optimization. This allowed us to quantify the speed of convergence by calculating the area
247 under the curve showing the cumulative minimum error as a function of completed model
248 evaluations. We performed 10 repeated runs of each algorithm on every benchmark problem to
249 allow proper statistical evaluation of the results.

250 **The performance of different optimization algorithms on**
251 **individual benchmarks**

252 **Hodgkin-Huxley model**

253



254

255 **Figure 1. The results of fitting conductance densities in the Hodgkin-Huxley model.**

256 (A) Example of a comparison plot showing the voltage trace generated by the model with its original parameters
 257 (blue) and the trace given by the model using the best parameter set found by the Random Search algorithm (red).
 258 (B) Plot showing the evolution of the cumulative minimum error during the optimization. The curves show the
 259 median of 10 independent runs for each relevant algorithm. Each generation corresponds to 100 model evaluations.
 260 The colors corresponding to the different algorithms (and packages) are shown in the legend. (C) Box plot
 261 representing the distribution of the final error scores over 10 independent runs of each algorithm. (D) Box plot
 262 representing the convergence speed of the algorithms tested, measured as the area under the logarithmic cumulative

263 minimum error curve (as shown in panel B). In (C) and (D), horizontal red lines indicate the median, the boxes
264 represent the interquartile range, whiskers show the full range (excluding outliers), and circles represent outliers.
265 Boxes representing single-objective algorithms are colored blue and those of multi-objective ones are red. Results
266 are sorted by the median score, from the best to the worst. The names of the packages on the horizontal axis are
267 colored to indicate the implementing package according to the legend in (D).

268

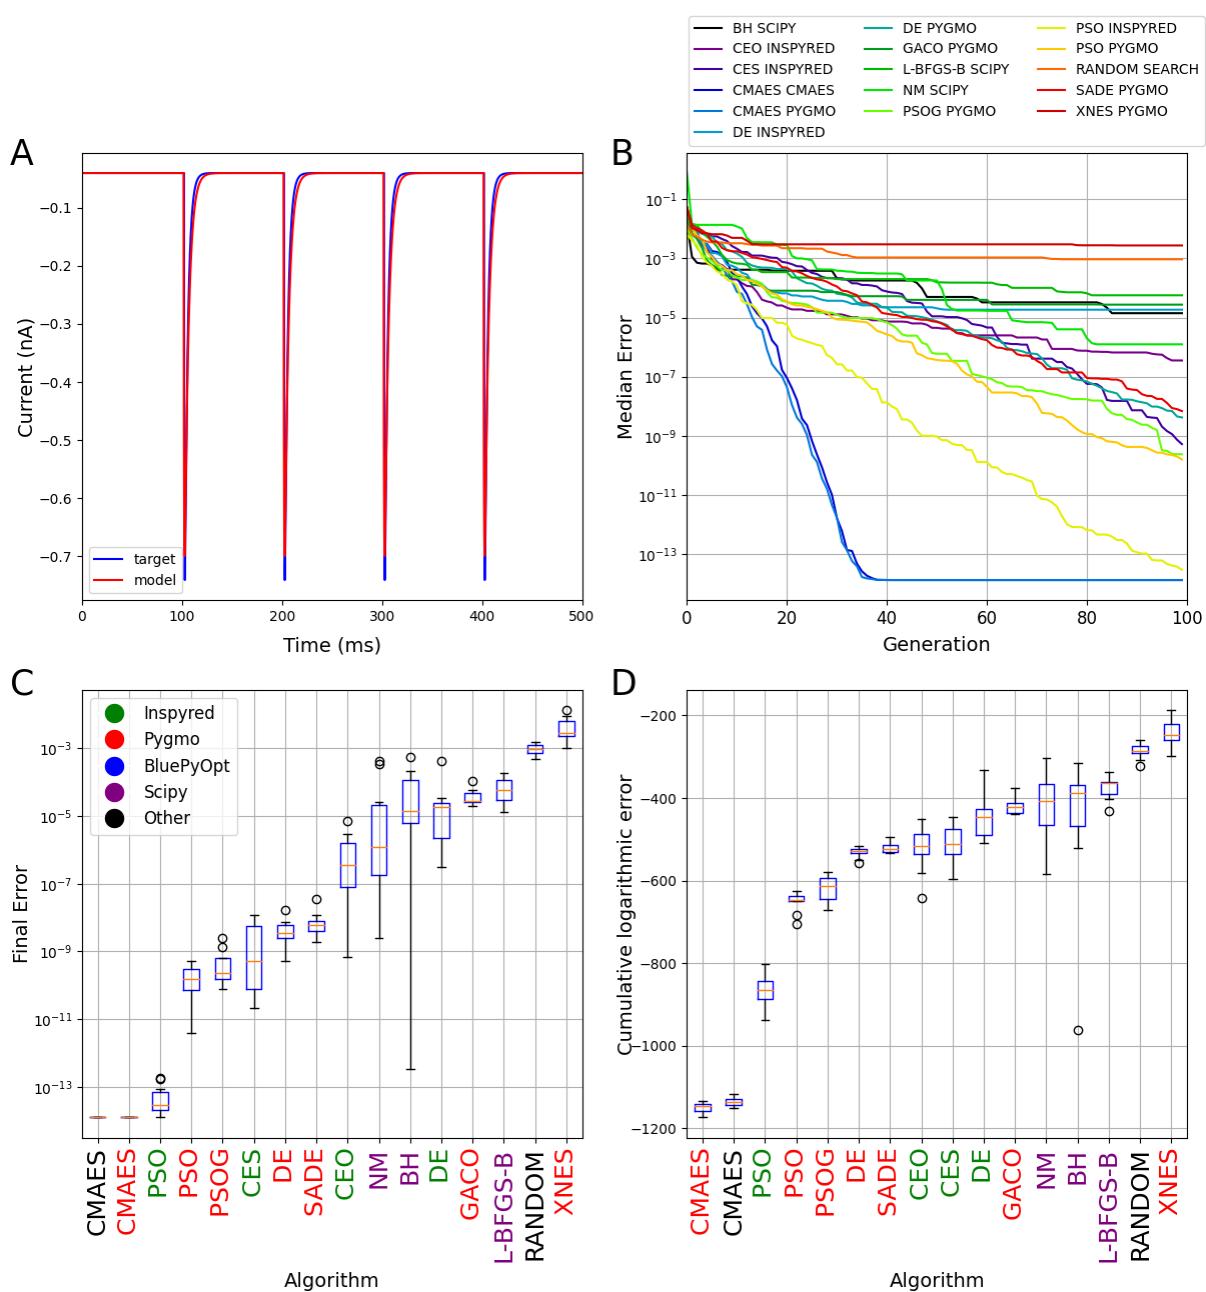
269 Our first benchmark problem involved finding the correct densities of two voltage-gated
270 conductances and a leak conductance (3 parameters overall) in the classic single-compartment
271 Hodgkin-Huxley model [28] based on the voltage response to a single current step stimulus
272 (Figure 1). We compared the response of each candidate model to that of the original model by
273 evaluating four features (spike count, spike amplitude, spike width, and mean squared error of
274 the voltage excluding spikes, evaluated using built-in error functions of Neuroptimus), which
275 also enabled the application of multi-objective optimization methods. We expected this to be a
276 relatively simple optimization problem based on the low number of parameters to fit, although
277 it is also clearly non-trivial due to the nonlinear nature of the neuronal dynamics and,
278 particularly, the complicated dependence of the extracted physiological features on the
279 conductance parameters.

280 Many of the search algorithms tested found relatively good solutions most of the time,
281 but most of them failed to converge completely in 10,000 model evaluations. The exception
282 was the CMAES algorithm, whose implementations both consistently converged to the optimal
283 solution after approximately 3,500 evaluations (the lowest possible error score was not exactly
284 zero due to rounding errors). Interestingly, multi-objective algorithms generally performed
285 worse on this use-case than single-objective ones, with Inspyred's NSGA2, PAES and Pygmo's
286 NSPSO algorithms giving worse results than Random Search. Different implementations of the
287 same algorithms (two versions for CMAES, three for PSO, and three for NSGA2) usually
288 showed similar convergence behavior, except for the implementation of NSGA2 by the

289 Insipyred package that performed significantly worse than the Pygmo and BluePyOpt versions
 290 of the same method. Overall, even this simple benchmark revealed surprisingly large
 291 differences in the performance of the different search methods that we included in our
 292 comparison.

293 **Voltage Clamp**

294



295

296 Figure 2. **The results of fitting the parameters of a synaptic connection based on simulated voltage-clamp
297 recordings.**

298 The plots in all four panels are analogous to those in Figure 1. Only single-objective methods were tested in this
299 use-case because only a single error function (mean squared difference) was used to compare model outputs to the
300 target data. Panel A shows the results of a best-fitting model found by the Random Search algorithm.

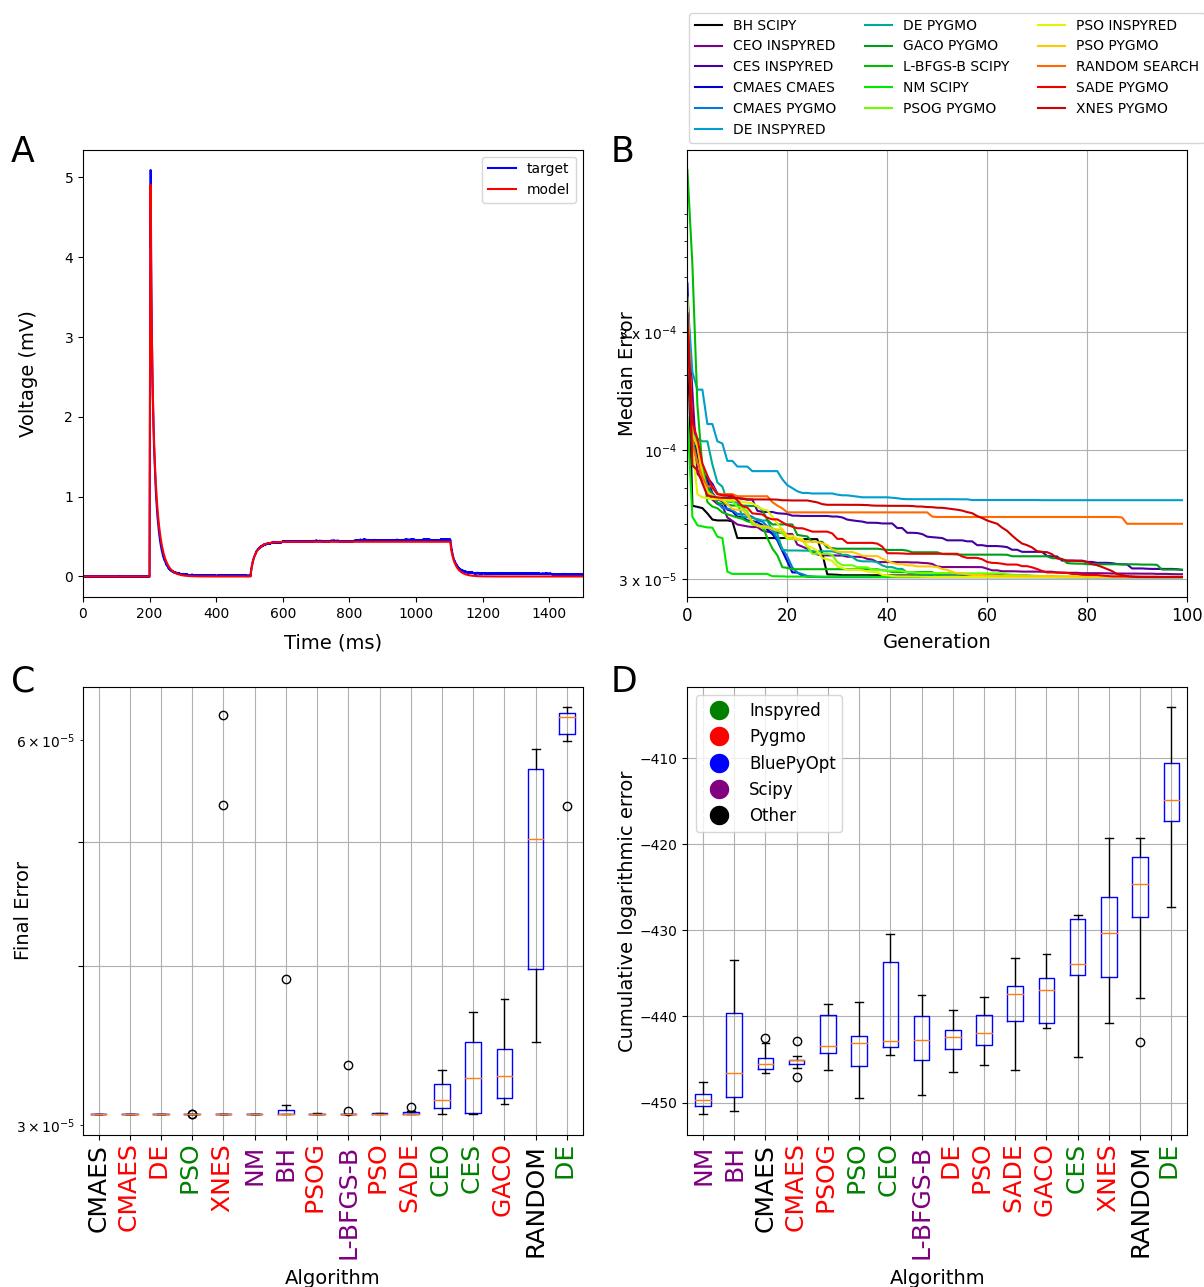
301

302 The second benchmark problem involved finding four parameters of a simulated
303 synaptic connection to a single-compartment model neuron using voltage-clamp recordings
304 (Figure 2). This use case also used surrogate data as the target, but in this case the recorded
305 variable was the current injected by the electrode during a simulated voltage-clamp experiment.
306 The parameters to be reconstructed were the maximal value (weight), delay, and rise and decay
307 times of the synaptic conductance change following each repeated activation of the synapse.
308 Due to the stereotyped nature of the data, mean squared difference was used as the only error
309 function, and thus only single-objective algorithms were tested.

310 Although this is still a relatively simple and low-dimensional problem, and the intrinsic
311 dynamics is much less complex than that of the Hodgkin-Huxley model in current clamp mode
312 in the first use-case above, we observed highly divergent performance for the set of algorithms
313 that we tested. Both implementations of CMAES reached the best possible score (again defined
314 by round-off error) in fewer than 40 generations (4000 model evaluations). The Inspyred
315 implementation of PSO also approached this limit by the end of the optimization (10,000 model
316 evaluations), but it converged substantially slower than CMAES. The Pygmo implementations
317 of PSO, two versions of the DE algorithm, and the CES algorithm of Inspyred also achieved
318 good results, but converged even more slowly. At the other end of the spectrum, local search
319 algorithms were typically not effective at solving this problem, and the XNES algorithm from
320 the Pygmo package actually performed worse than the baseline random search method.

321 **Passive, anatomically detailed neuron**

322



323

324 **Figure 3. The results of fitting the passive biophysical parameters of a morphologically detailed multi-**
 325 **compartmental model to experimental recordings from a hippocampal pyramidal neuron.**

326 The plots in all four panels are analogous to those in Figure 1. Only single-objective methods were tested in this
 327 use-case because only a single error function (mean squared difference) was used to compare model outputs to the
 328 target data. Panel A shows the results of a best-fitting model found by the CMAES algorithm.

329

330 This use-case represents an important practical problem that has been investigated in
331 several previous studies [15, 29, 30, 31]. It involves the estimation of three basic biophysical
332 parameters that determine the propagation and integration of voltage signals within neurons in
333 the subthreshold voltage range: the (specific) membrane capacitance, membrane resistance, and
334 axial resistance. The task is to estimate these three parameters based on the voltage response of
335 a neuron to a current stimulus (which, in our case, consisted of a larger short and a smaller long
336 current step) recorded from a hippocampal pyramidal cell in vitro (Figure 3). The response of
337 the model is linear in terms of the injected current, but still depends on the combination of the
338 three biophysical parameters (which are assumed to be spatially uniform within the cell) in a
339 non-trivial way due to the complex morphology of the neuron. In the absence of spikes, we
340 used the mean squared difference between the simulated and the experimentally recorded
341 voltage traces as the only error function, and restricted our attention to single-objective
342 algorithms.

343 This benchmark proved to be the easiest in our entire test suite. Many algorithms found
344 the best possible fit to the data in (almost) all the runs, and most of them also converged
345 relatively rapidly. In this case, local search methods such as the Nelder-Mead and the L-BFGS-
346 B algorithms also found the optimal solution efficiently in most runs. One curious exception
347 was the DE algorithm implemented by the Inspyred package, which achieved a worse result
348 than Random Search, even though the other implementation of the same algorithm by the
349 Pygmo package was among the high-performing methods.

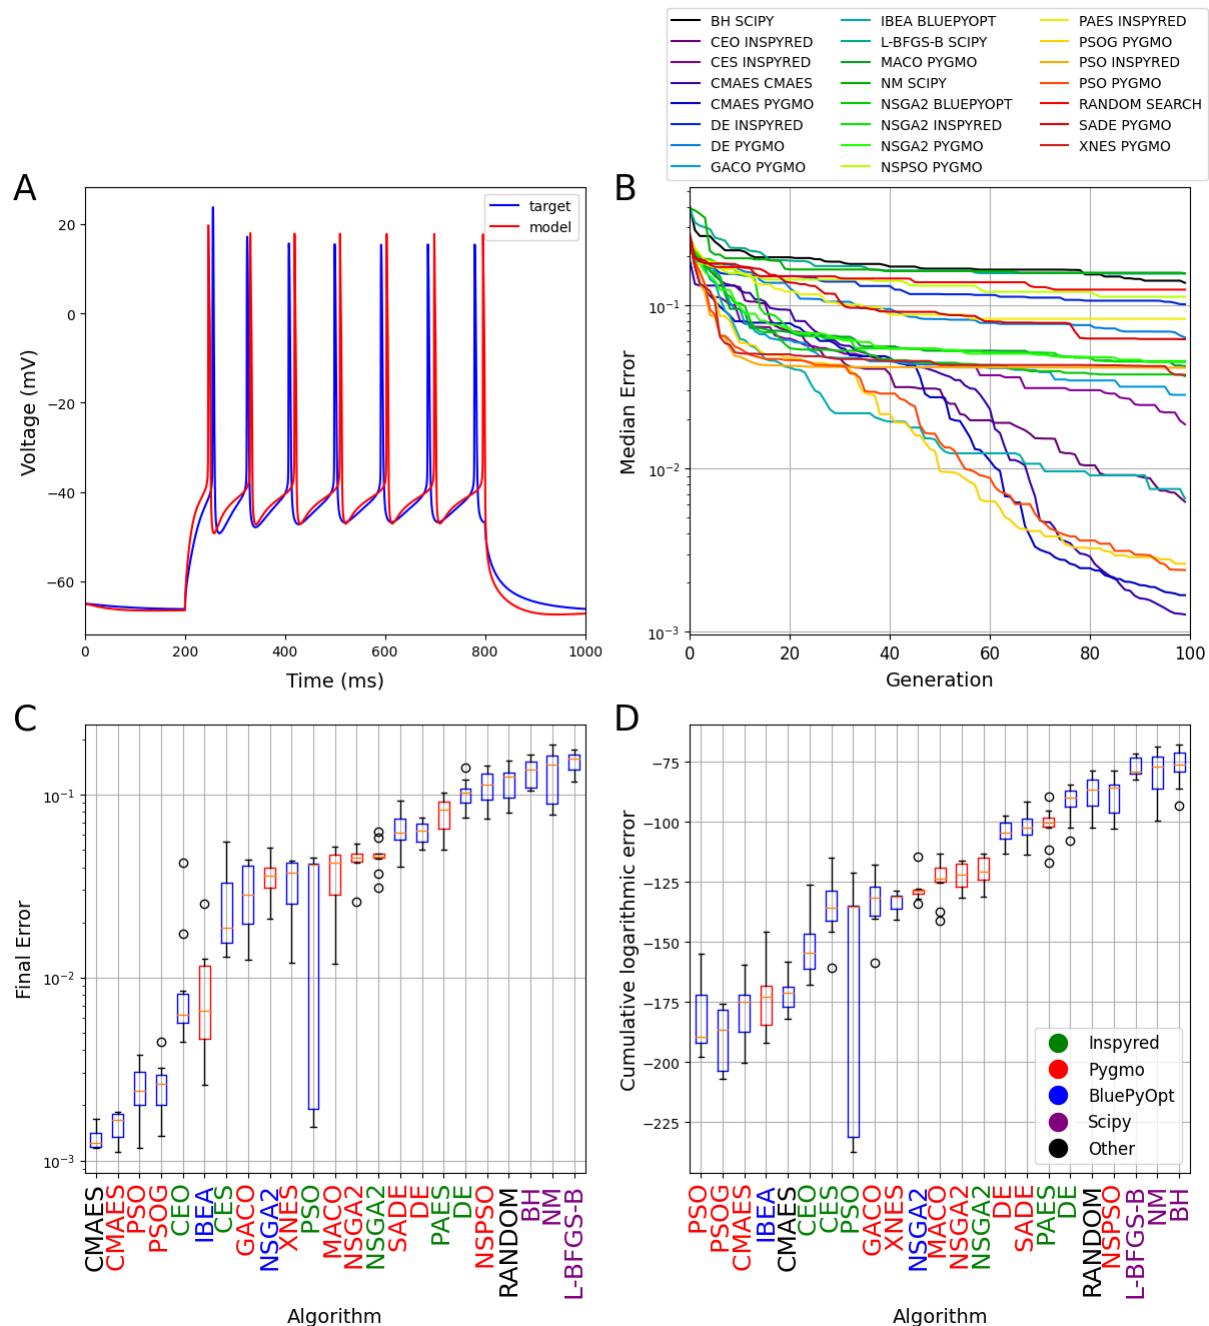
350

351

352 **Simplified active model**

353

354



355
356 **Figure 4. The results of fitting the densities of somatic voltage-gated conductances in a morphologically
357 simplified six-compartment model using a simulated voltage trace from a detailed compartmental
358 model as the target.**

359 The plots in all four panels are analogous to those in Figure 1. Panel A shows the results of a best-fitting model
360 found by the CMAES algorithm.

361
362 This benchmark problem is more complex than the previous ones in several respects.
363 The task in this use case is to determine the somatic densities of nine voltage-gated

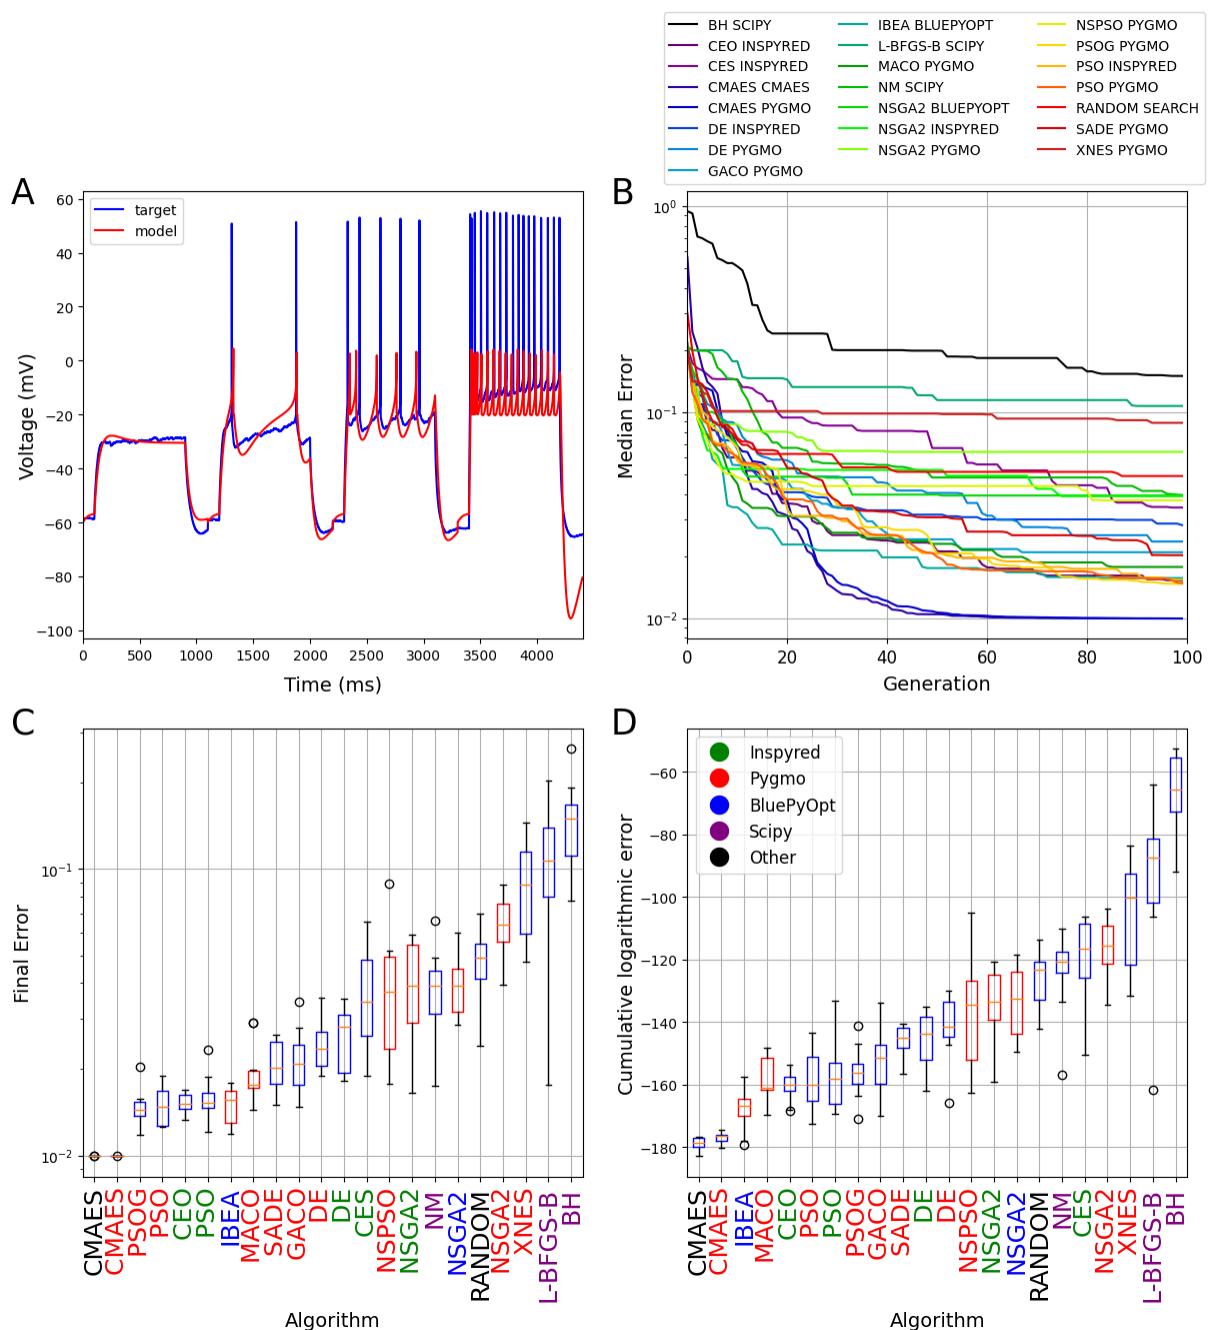
364 conductances in a model of a hippocampal CA1 pyramidal neuron with simplified morphology
365 (consisting of only six compartments) so that the somatic voltage response of the model best
366 approximates the response of a fully detailed CA1 pyramidal cell model under the same
367 conditions (Figure 4). We used five of the error functions implemented by Neuroptimus (mean
368 squared error excluding spikes, spike count, latency to first spike, action potential amplitude,
369 action potential width, and after-hyperpolarization depth) to compare the two voltage traces.
370 This also enabled us to test multi-objective algorithms besides the single-objective ones.

371 In this more complex use case, there were large differences in performance among the
372 algorithms, with two orders of magnitude difference between the final errors of the best- and
373 the worst-performing methods. Once again, implementations of the CMAES algorithm
374 achieved the best final scores, but the Pygmo implementations of PSO also delivered good final
375 scores along with the best convergence speed. Among multi-objective algorithms, IBEA
376 achieved the best final scores, and also performed quite well in terms of convergence speed. At
377 the other extreme, all local search algorithms typically performed worse than Random Search,
378 and are clearly inadequate for this type of problem. It is worth noting that all three
379 implementations of the NSGA2 algorithm gave similar results, as did the different flavors of
380 DE, although neither these algorithms nor several other bio-inspired algorithms (such as other
381 evolutionary algorithms or ant colony optimization) were capable of providing as good
382 solutions as CMAES and PSO on this benchmark.

383

384 **Extended integrate-and-fire model**

385



386

387 Figure 5. The results of fitting a phenomenological spiking neuronal model (the adaptive exponential
388 integrate-and-fire model) to capture experimental recordings with multiple traces.

389 The plots in all four panels are analogous to those in Figure 1. Panel A shows the results of a best-fitting model
390 found by the CMAES algorithm. Note that the height of action potentials is irrelevant in the integrate-and-fire
391 model, and the spikes generated by the model are not explicitly represented in the figure.

392

393 This use case involves fitting the parameters of an adaptive exponential integrate-and-
394 fire model neuron so that it captures the spiking responses of a hippocampal CA3 pyramidal

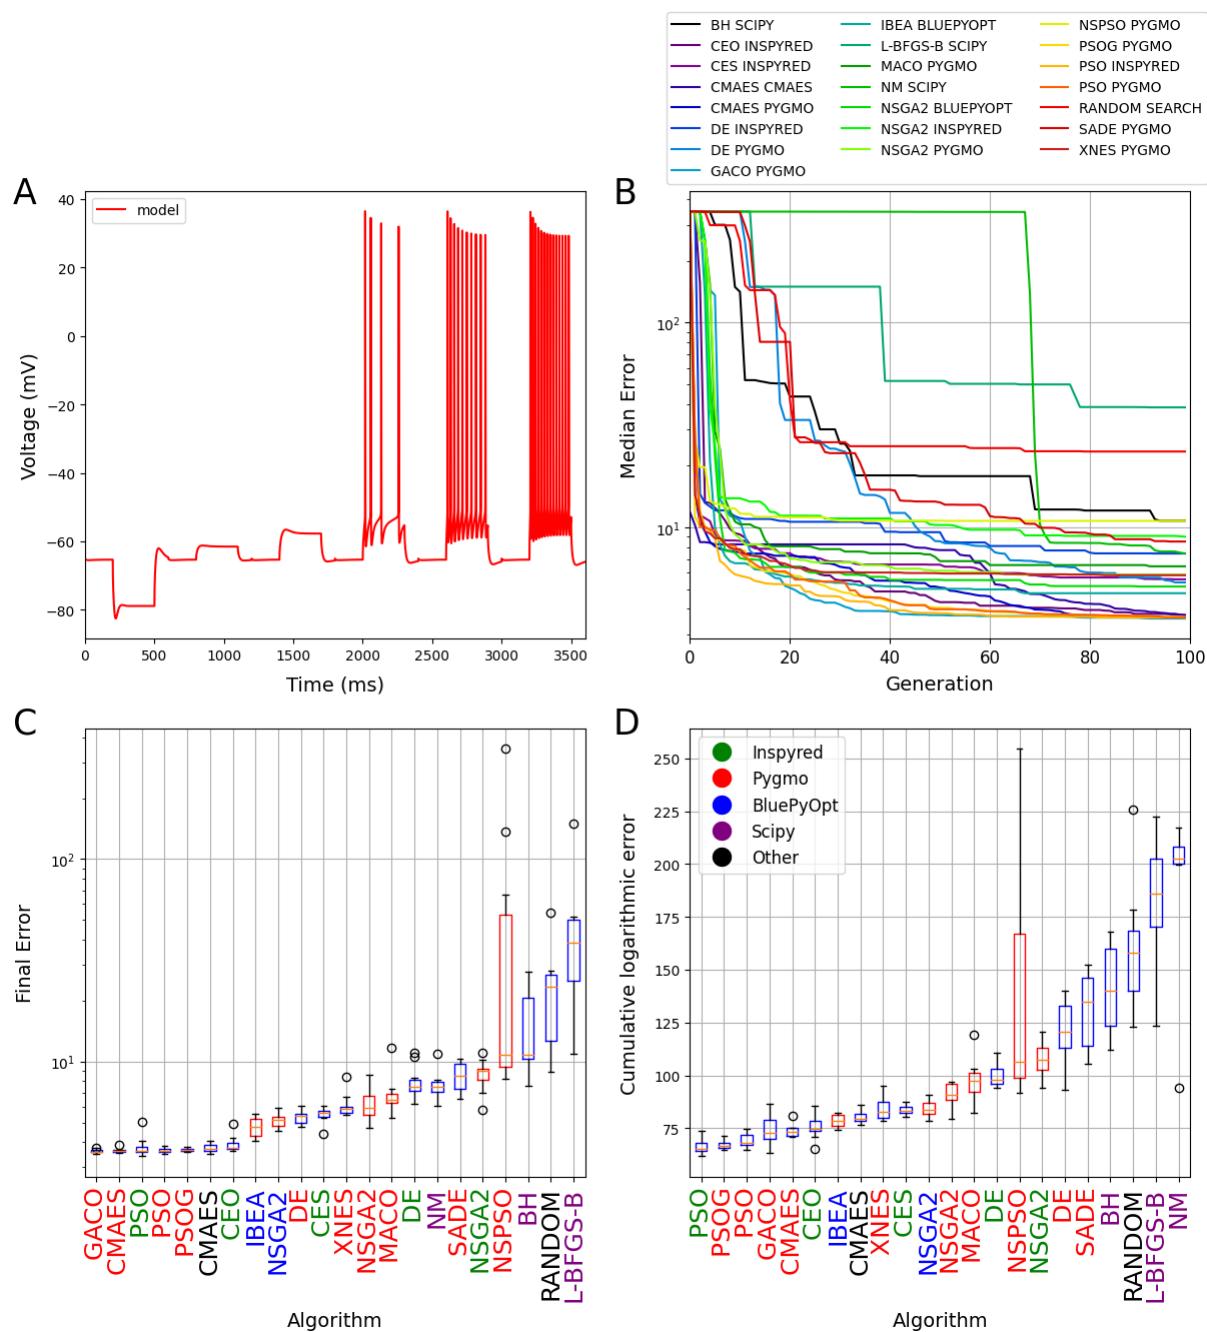
395 neuron recorded in vitro (Figure 5). This is a single-compartment model that does not include
396 detailed models of neuronal biophysics; instead, it aims to capture neuronal spiking
397 phenomenologically, using an extended integrate-and-fire formalism with an exponential term
398 in the current-voltage relationship and an adaptation variable that is also linked to spiking [32,
399 33]. This model has a total of 10 parameters that had to be fitted by the optimization algorithms.
400 Unlike the models in the other use cases (which were implemented in NEURON), this model
401 was implemented in the NEST simulator [34], and was treated as a black box by Neuroptimus.
402 The parameters generated by the optimization algorithms were passed to an external Python
403 script, which constructed the model, ran the simulations using NEST, and passed the results
404 (spike times and subthreshold voltage traces in two separate files) back to Neuroptimus for
405 evaluation and comparison with the experimental data. The data included the voltage responses
406 of a real CA3 pyramidal cell to current steps of four different amplitudes (these responses are
407 shown concatenated in blue in Figure 5A), and the model had to capture all of these responses
408 simultaneously. As integrate-and-fire models cannot (and are not expected to) reproduce spike
409 shape, we used spike count, latency to first spike, and the mean squared difference of the voltage
410 excluding spikes as three error components during the optimization.

411 On this benchmark, the two implementations of CMAES found the solutions with the
412 lowest error. In fact, they obtained the same lowest error score several times, and this was lower
413 than the scores achieved by any other algorithm, so this error score likely corresponds to the
414 best possible solution of this optimization problem. Although clearly inferior to CMAES on
415 this problem, the various implementations of the particle swarm algorithm, the multi-objective
416 algorithm IBEA, and the classical evolutionary algorithm found relatively good solutions, while
417 several methods performed substantially worse than Random Search. We note that the PAES
418 algorithm generated parameter combinations that led to errors during the NEST simulation, and
419 was therefore excluded from the current comparison.

420

421 **Morphologically and biophysically detailed CA1 pyramidal cell model**

422



423

424 **Figure 6. The results of fitting conductance densities and kinetic parameters in a CA1 pyramidal cell model.**

425 The plots in all four panels are analogous to those in Figure 1. Panel A shows the results of a best-fitting model
 426 found by the CMAES algorithm. No target trace is shown because, in this use case, the actual target is defined by
 427 the statistics of electrophysiological features that are extracted from a set of experimental recordings.

428

429 Our final use case represents a typical scenario in the construction of morphologically
430 and biophysically detailed compartmental models [1, 4, 7, 35, 36, 37, 38]. The model is based
431 on the reconstructed morphology of a CA1 pyramidal neuron [39], and contains a large set of
432 voltage-gated conductances, several of which are distributed non-uniformly within the cell (see
433 Methods for further details of the model). The goal is to find the values of 12 parameters that
434 determine the densities and biophysical properties of voltage-gated and leak conductances in
435 the model such that the features extracted from the voltage responses of the model to multiple
436 step current injections best approximate the average of the same features extracted from
437 experimental recordings under matching conditions (Figure 6). One hyperpolarizing and five
438 depolarizing current steps were used, and these yielded a total of 66 features of 20 different
439 types (Table 3) that were extracted and evaluated for each model instance during the parameter
440 search.

441 Although this is certainly the most complex model in our benchmarking suite with the
442 largest number of free parameters, finding solutions with errors close to the smallest possible
443 value was apparently easier than in the previous two use-cases (although, strictly speaking, we
444 cannot rule out the possibility that none of the algorithms tested ever came close to the unknown
445 globally optimum error score). More specifically, all three versions of PSO, both
446 implementations of CMAES, and also the GACO and CEO algorithms consistently yielded
447 similar low error scores, but several other algorithms, including the multi-objective IBEA and
448 NSGA2 methods, also gave acceptable solutions. We note that running the PAES algorithm
449 resulted in memory errors, and it was therefore omitted from the evaluation of this use case.

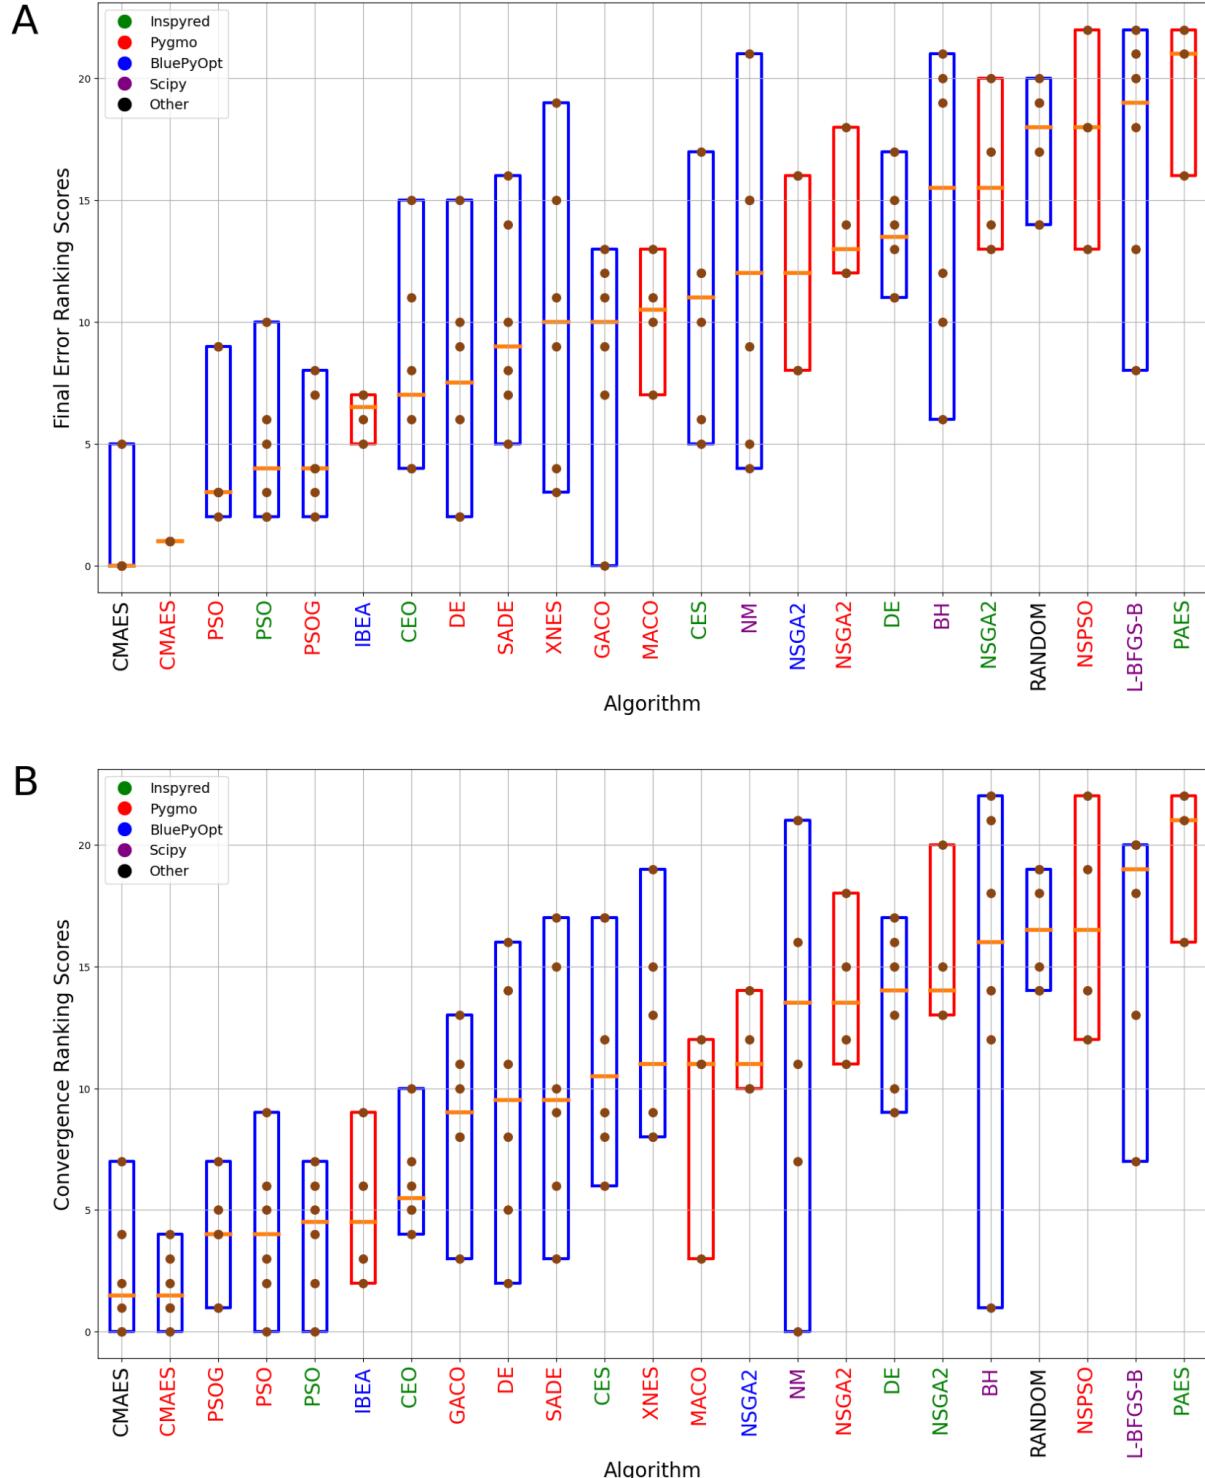
450

451 **Overall performance of the algorithms**

452 In general, no single algorithm is expected to perform well in all types of global
453 optimization problems. Popular methods can take advantage of different types of regularities in
454 the error function to speed up the search for the global optimum even in high-dimensional
455 spaces with multiple local optima. Therefore, problems with different structures may require
456 different algorithms for their efficient solution, and we can identify some signs of this
457 heterogeneity when comparing the results of the individual benchmarks described above.
458 Nevertheless, some clear patterns are evident, and we can quantify this by constructing and
459 examining summary statistics for the algorithms across all the use cases.

460 Figure 7 summarizes the rankings of the various algorithms in our study according to
461 final score and convergence speed. Individual ranks were based on the medians of the respective
462 performance measure across all runs of the algorithm in a particular benchmark (this was also
463 the basis of the placement of the algorithms along the horizontal axes in panels C and D of
464 Figures 1-6), and Figure 7 shows the statistics of these ranks for each algorithm across the
465 different benchmarks.

466



467

468 Figure 7. Overall rankings of optimization algorithms.

469 Statistics of the ranks achieved by individual optimization algorithms on the different benchmarks (Figures 1-6)
470 according to the final error (A) and convergence speed (B). Brown dots represent the ranks achieved by the
471 algorithms in each use-case; boxes indicate the full range and the orange line represents the median of these ranks.
472 The single-objective algorithms are shown in blue and the multi-objective ones in red boxes. The color of the name
473 of the algorithm indicates the implementing package, with the color code included in the legend. Algorithms are
474 sorted according to the median of their ranks.

475

476 In terms of the generally best-performing algorithms on our neuronal optimization test
477 suite, the results are quite clear. In almost all cases, CMAES delivered the best results after
478 10,000 model evaluations, and its two implementations by different packages performed quite
479 similarly. The three implementations of the particle swarm algorithm that we tested also showed
480 similar performance, and were typically better than all the other methods except for CMAES.
481 IBEA was close behind the PSO variants in the rankings, and was clearly the best among the
482 multi-objective methods that we tested. It is interesting to note that some of the algorithms,
483 including local search methods (and especially the Nelder-Mead algorithm) but also some other
484 methods such as GACO and XNES showed widely varying performance across the different
485 benchmarks, so these may be suitable for some problems but completely inadequate for others.
486 Finally, the rankings based on the convergence score are generally quite similar to those based
487 on just the final score, although there are some minor differences - for instance, PSO appears
488 to be more competitive with CMAES according to this measure.

489

490 **Online database of optimization results**

491 The results presented so far summarize the performance of a selected subset of the
492 algorithms implemented by five Python packages, using their default settings, on a pre-defined

493 suite of six neuronal optimization problems. To increase the utility and reproducibility of our
494 results, we also wanted to share the details of all the optimization runs, including the settings
495 that enable their replication as well as their detailed results. In addition, we wanted to find a
496 simple way of updating and extending the study with more optimization runs, potentially with
497 different settings or algorithms not included in the present comparison, or even involving
498 additional use cases, not just by us but also other interested researchers. We therefore designed,
499 implemented and deployed an online database with an associated, publicly accessible web
500 server (<https://neuroptimus.koki.hu>) that allows users to upload, query, and analyze
501 optimization runs performed by the Neuroptimus software tool.

502 The website allows users to browse the optimization results stored in the online
503 database, and filtering options are available to create lists of relevant results. The results of
504 optimization runs can be viewed in a detailed text-based format, and selected subsets of
505 optimizations can be analyzed and displayed graphically, similarly to the plots in Figures 1-6
506 above. Registered users can also add to the database their own optimization results by uploading
507 the JSON file (metadata.json) generated by Neuroptimus after each optimization run. Users can
508 optionally also upload the other files that belong to the optimization (including the model and
509 the target data) in the form of an archive, which creates an online record of the optimization
510 that allows its full replication.

511 The database currently contains the results of all the optimization runs from the current
512 study. This enables users of the website to replicate most of the figures in this paper, and to
513 download individual optimization runs (including their settings and results). Users can also
514 carry out custom analyses of the results, and (after registration) they can add their own
515 optimization results (created using Neuroptimus) on both existing and novel use cases and
516 compare these with other results on the same use case. This way, the website offers an

517 interactive, continuously updated, and publicly accessible “live” version of this paper, which
518 will provide a valuable online resource for researchers to explore and share methods and results
519 on neuronal optimization.

520

521 **Discussion**

522 The results of our study have the potential to advance the state of the art in neural
523 parameter optimization in several different ways. First, we have created and shared
524 Neuroptimus, a software tool that was designed to help both computational and experimental
525 neuroscientists in the complete formulation and solution of neuronal parameter search
526 problems. All the functions of Neuroptimus are accessible through a graphical user interface,
527 although there is also a command line interface to support more advanced usage. Users of
528 Neuroptimus gain uniform access to a large number of optimization algorithms implemented
529 by several widely used Python packages, including several algorithms that were used
530 successfully in previous neuronal modeling studies, and also several other state-of-the-art
531 optimization methods that are popular in other domains but have not been applied to neuronal
532 parameter optimization. This feature of Neuroptimus allowed us to systematically test the
533 performance of a wide variety of parameter search methods on six distinct neuronal
534 optimization problems, which makes it possible to offer some recommendations for future
535 neuroscientific studies that rely on parameter optimization (see below). Finally, we created an
536 online database of optimization results obtained by using Neuroptimus. This database currently
537 contains the results of the present study, but the online user interface also allows us as well as
538 others to add new results and compare them with already existing ones.

539 **Comparison of Neuroptimus with other neural**

540 **optimization tools**

541 A variety of software tools have been developed and used for the purpose of optimizing
542 the parameters of neural models. This includes the built-in optimization modules of the general-
543 purpose neural simulators NEURON [20] and GENESIS [21], the optimization-oriented
544 features of the NetPyNE neural modeling framework [26], the Neurofitter program [22], as well
545 as the Python packages BluePyOpt [23], NeuroTune [25], and pypet [24]. However, each of
546 these tools (except for Neurofitter, which is no longer actively maintained, and focuses on a
547 single specific cost function) relies on just one or a few algorithms, or a single external
548 optimization package, to perform parameter search, while Neuroptimus provides access to a
549 large variety of different algorithms from five distinct Python packages. This gives users of
550 Neuroptimus a lot of flexibility to choose the best method for any particular fitting problem. As
551 an example, Neuroptimus can take advantage of the large number of local and global
552 optimization algorithms offered by the Pygmo package [40], which is a mature and actively
553 maintained tool used, among others, by the European Space Agency. Another distinguishing
554 feature of Neuroptimus is its graphical user interface. Among the other tools, only NEURON
555 offers GUI-based access to parameter optimization, but the utility of this feature is severely
556 limited by its reliance on the local search method PRAXIS as its only available algorithm.

557 **Algorithm recommendations based on our benchmarking**

558 **results**

559 The performance of optimization algorithms in general depends quite heavily on the
560 nature of the problem, and no particular algorithm is expected to provide good solutions

561 universally. However, within the task domain that we considered here, i.e., finding the
562 biophysical parameters of models of single neurons, we can make some clear recommendations.
563 Our benchmarking results were dominated by two different metaheuristics, covariance matrix
564 adaptation evolution strategy (CMAES) and particle swarm optimization (PSO), followed by
565 the multi-objective indicator-based evolutionary algorithm (IBEA) - so we would suggest trying
566 these methods first to attack a novel neuronal parameter optimization problem. We also
567 confirmed that local optimization algorithms are generally not suitable for more complex
568 parameter search tasks, although they can be adequate and even efficient in the simplest cases.

569 Although, in principle, implementation details (particularly the default settings of
570 algorithm parameters) could influence the performance of the algorithms, we found essentially
571 no difference in the quality of solutions found by implementations of the same algorithms by
572 different packages. However, the algorithms and even implementations of the same algorithm
573 differ in the extent to which their execution can be parallelized, and this can have a large impact
574 on the runtime of the algorithms, especially on highly parallel architectures. All algorithms of
575 Inspyred and BluePyOpt, and some algorithms of Pygmo support the parallel evaluation of
576 multiple candidate models (typically those within a particular generation or iteration), and
577 Neuroptimus allows users to take advantage of these capabilities. It is worth noting in this
578 context that Pygmo contains two variants of particle swarm optimization: PSO, which is closer
579 to the original formulation of this algorithm, updates the velocities and positions of particles in
580 a serial manner, and is thus not suitable for parallelization; and the generational variant PSOG,
581 where the velocities and positions of all particles are updated in parallel. PSOG shares this
582 feature with the implementation of PSO by the Inspyred package, and both of these can be run
583 in a parallelized fashion from Neuroptimus. As a result, although all three variants of PSO
584 produced similar final results in our tests, the runtimes of Pygmo's PSOG and Inspyred's PSO
585 were significantly lower than those of Pygmo's PSO when multiple CPU cores were utilized.

586 The situation is similar for the CMAES algorithm, where the current implementation in Pygmo
587 does not support parallel evaluations. This was the reason for including the Cmaes package in
588 Neuroptimus: this module implements CMAES in a way that allows straightforward
589 parallelization, and Neuroptimus uses this implementation to support parallel execution of this
590 popular and efficient method. Finally, to demonstrate the importance of parallelized
591 implementations, we note that running a single optimization of our most complex use case (the
592 detailed CA1 pyramidal neuron) with 10,000 model evaluations took approximately 10 days on
593 our compute server for algorithms without parallelization; by contrast, a single run of this use
594 case using the same number of model evaluations (e.g., 100 generations with populations of
595 100 individuals) with algorithms that support parallel evaluations took only a few hours on a
596 single node of a supercomputer (accessed via the Neuroscience Gateway) that allowed an entire
597 generation of models to be evaluated in parallel.

598 Our findings regarding the relative performance of various optimization methods are
599 mostly in line with results of earlier studies that included such comparisons. Vanier and Bower
600 [17] compared four different algorithms on a set of use cases similar to ours. They examined
601 the performance of random search, conjugate gradient descent (a local search method), an
602 evolutionary (or genetic) algorithm, and simulated annealing. They found that their
603 evolutionary algorithm (which was similar to the CEO algorithm in our benchmark) delivered
604 good performance even for more complex use cases with a larger number of parameters. This
605 is consistent with the generally good results of evolutionary-type algorithms in our study,
606 although we found several more recent variants that outperformed the classic version. They also
607 found simulated annealing to be very effective, and this was later confirmed by Friedrich et al.
608 [15] using the implementation that is built into the GENESIS simulator. None of the packages
609 currently supported by Neuroptimus contain the traditional simulated annealing algorithm. In
610 fact, older versions of the SciPy module used to include simulated annealing, but it was later

611 deprecated and replaced by the basinhopping algorithm, which is considered to be a
612 generalization of simulated annealing. In this light, the generally poor performance of the
613 basinhopping algorithm in our tests is slightly surprising, although it may be caused by
614 implementational problems or improper default settings of the parameters.

615 Our finding that CMAES performs well in a variety of different tasks is supported by
616 several other studies. In particular, CMAES and IBEA have been compared on data-driven
617 neuronal models, and CMAES generally delivered better final scores [41]. CMAES was also
618 found to be efficient and robust in a study that involved fitting the biophysical parameters of
619 models of striatal neurons [42]. Outside the neuronal modeling domain, a recent study
620 compared two sophisticated evolution strategy variants, CMAES and xNES on different
621 problems, and the results clearly showed that CMAES consistently outperformed xNES [43].
622 Our findings also support this conclusion, and add some evidence regarding two additional
623 evolution strategy types: classic evolution strategy (CES, from the Inspyred package), which
624 performed similarly to xNES (from Pygmo), and Pareto-archived evolution strategy (PAES,
625 from Inspyred), which was one of the weakest performers (worse than random search) in our
626 comparison.

627 **Limitations of the benchmarking study**

628 The specific results that we obtained in our benchmarking study depend, to some extent,
629 on some arbitrary choices that we had to take when designing our tests. For instance, we
630 arbitrarily set the number of model evaluations to 10,000 for every algorithm to ensure a fair
631 comparison of the final results. However, not all algorithms converged completely after 10,000
632 model evaluations in some of our use cases, and thus allowing more (or fewer) evaluations
633 would likely affect the rankings based on the final score. The other performance measure that

634 we used, the convergence score, is expected to be less sensitive to the exact number of model
635 evaluations allowed, and also provides an indication of the speed of convergence.

636 We made another arbitrary choice for every algorithm with a hierarchical design, which
637 includes all population-based methods (such as evolutionary and swarm intelligence
638 algorithms) but also other nested algorithms such as basinhopping. We set the number of model
639 evaluations in the innermost loop (i.e., the size of the population, or the number of steps in the
640 local optimization) to 100, and the number of iterations in the outer loop (e.g., the number of
641 generations) also to 100. This 100x100 partitioning of the total of 10,000 evaluations is a
642 reasonable choice for most algorithms, and led to good results in most cases; however, it is
643 entirely possible that a different choice (such as 200 generations with a population size of 50,
644 or the other way around) would have resulted in improved performance for some of the
645 algorithms.

646 Similarly, almost all the algorithms that we used include some tunable parameters that
647 change the course of the optimization and may heavily influence the quality of the results. We
648 decided to use the default settings specified in the optimization packages for all algorithms
649 (with the exception of the CEO algorithm from the Inspyred package - see the description of
650 this algorithm in Methods for details). In many cases, these settings were compatible with
651 parameters recommended in the relevant literature; in some other cases, different sources
652 suggested different settings; and in some cases, no such recommendations could be found. It is
653 also known that the best settings for such algorithmic parameters can depend on the details of
654 the problem, so it is possible that distinct settings would be optimal for the various use cases.

655 Overall, we may conclude that the rankings that we provide are not just about the
656 algorithms themselves (or even about their implementations by particular packages), but are
657 likely also influenced by the settings of the parameters (including the population size) for each

658 method. In fact, we expect that the results of some algorithms could be improved substantially
659 by using different settings, and there are also many additional optimization algorithms that were
660 not included in the current comparison. Therefore, we hope to update and extend our study as
661 more optimization results become available (see below for further discussion).

662 **Possible extensions**

663 In addition to testing our use cases with more algorithms and settings as described
664 above, the results of our study could be extended in other important ways. One possible
665 direction would be to extend the set of use cases to other types of problems. All of the use cases
666 included in the current study involved single cell models; however, in principle, Neuroptimus
667 can also handle neural models at the subcellular and network levels, and supports the
668 optimization of their parameters. For example, in one project, we used Neuroptimus to tune the
669 concentrations of molecules in biochemical pathways involved in synaptic plasticity, relying
670 on the reaction-diffusion (rxn) module of NEURON to run the simulations. In another project,
671 we used Neuroptimus to tune synaptic weights in a network model that was constructed and
672 simulated using Brian as an external simulator. The main reason why we did not include these
673 examples in the current comparison is that, in both of these cases, every single simulation of
674 the model requires more substantial resources and time, so that repeatedly running these
675 problems with all the algorithms would have been beyond our current computational resources,
676 and would have been completely unfeasible with algorithms that do not support parallel model
677 evaluations.

678 There are also many useful features that we could potentially add to our optimization
679 tool, Neuroptimus. For example, Neuroptimus currently returns only a single parameter
680 combination corresponding to the lowest error score at the end of the optimization (although

681 the parameters and errors of all the models tested are also saved into a file and may be analyzed
682 outside Neuroptimus). This is the case even when multi-objective methods are used; the winner
683 in this case is selected from the final population by minimizing the weighted sum of the
684 objectives, using weights provided by the user before the optimization run. However, the final
685 population of multi-objective optimization carries much more information, as it approximates
686 the Pareto front (the parameter combinations representing the best possible tradeoffs between
687 the objectives for different choices of the weights). Therefore, it would be useful to add to
688 Neuroptimus the capability of properly representing and analyzing the results of multi-objective
689 optimization. Other useful extensions could include the ability to chain optimization algorithms
690 (e.g., by automatically running local optimization using the output of a global algorithm as the
691 starting point), and the ability to visualize the progress of optimization (in error space and in
692 the parameter space) while it is still running.

693 **Community and cooperation through the Neuroptimus**
694 **website**

695 We do not see the benchmarking results presented in this paper as the final word in
696 evaluating parameter search methods for neuroscientific problems. As we argued above, it will
697 be important to extend our study with more use cases and further evaluation of different
698 algorithms and settings. Global parameter optimization is also a fast-moving field where new
699 methods emerge regularly; the fact that Neuroptimus supports several actively developed
700 packages, and is also flexible enough to accommodate new packages guarantees that new
701 developments can be integrated with minimal effort.

702 We have developed and deployed the Neuroptimus web server to provide a platform for
703 sharing and analyzing optimization results. By allowing all users to upload results obtained by

704 using Neuroptimus, and to compare them with already uploaded results (including all the results
705 of the current paper), the web site will become a continuously updated “live” version of this
706 paper. This should facilitate meaningful, quantitative comparisons of parameter optimization
707 methods, and aid the collaboration of different research groups that are interested in this topic.
708 We encourage all interested professionals (and especially those who are experts in using
709 particular algorithms) to run the use cases with improved settings, try other algorithms, add new
710 use cases, and share their results on the Neuroptimus website. This way, we can collectively
711 track new developments, and offer reliable solutions for an increasing variety of neural
712 optimization problems.

713

714 **Methods**

715 **Software tools and services**

716 **Neural Optimization User Interface (Neuroptimus)**

717 At the core of our methodology is a software tool that we developed, called Neural
718 Optimization User Interface (or Neuroptimus). Neuroptimus implements a software framework
719 that allows users to set up and solve parameter optimization problems and analyze the results.
720 Neuroptimus performs parameter optimization mainly by providing a common interface to a
721 large number of popular parameter search algorithms implemented by various open source
722 packages. In principle, Neuroptimus can be used to optimize the parameters of all kinds of
723 systems; however, its main purpose is to aid parameter fitting in neural systems, and especially
724 in detailed models of neurons. Accordingly, it includes many features that were developed

725 specifically for this scenario, which support simulating biophysical models of neurons using
726 the NEURON simulator, and comparing their behavior to experimental data obtained with
727 common electrophysiological protocols.

728 Neuroptimus is essentially an updated and extended version of our previous tool
729 Optimizer (<https://github.com/KaliLab/optimizer>) [15]. The basic design of these two pieces of
730 software is quite similar, and they also share many details of their implementation. Therefore,
731 we will focus on the new features and other differences here, and summarize the features that
732 are used by the current benchmarking study, but we refer the reader to Friedrich et al. (2014)
733 and the Neuroptimus documentation (<https://neuroptimus.readthedocs.io/>) for further details.

734 Neuroptimus is open source software, implemented in Python3, and can be accessed at
735 the GitHub repository <https://github.com/KaliLab/neuroptimus>. Its functions are available both
736 via a graphical user interface (GUI) that guides users through the steps of setting up, running,
737 and evaluating the results of parameter optimization tasks, and via a command line interface
738 that performs these tasks based on the settings stored in a configuration file. The GUI was built
739 using the PyQt5 package that provides a Python binding to the cross-platform GUI toolkit Qt
740 (version 5).

741 The complete definition of a neural parameter optimization problem requires the
742 specification of multiple components. First, we need to provide the model whose parameters
743 we wish to optimize. Neuroptimus can load, manipulate and execute models implemented in
744 the HOC language of the NEURON simulator. The parameters to be optimized can be selected
745 from the parameters of this model, or the user can provide a function (implemented in Python)
746 that defines abstract parameters and how these should be mapped onto the concrete parameters
747 of the NEURON model. As an alternative, models can be implemented by any external program
748 that is capable of reading the variable parameters of a model candidate from a text file, setting

749 up the model accordingly, running the simulation(s), and saving the results to files that can be
750 interpreted by Neuroptimus.

751 Second, the cost function for neural parameter optimization is typically defined in terms
752 of some target data (from experiments or prior simulations) and a function (or set of functions)
753 that quantifies the difference between the output of the model and the target data. Neuroptimus
754 can handle different types of target data, including time series (such as voltage and current
755 traces), explicit spike times, and feature statistics.

756 Neuroptimus implements several error functions that can be used individually or in
757 combination to evaluate during the optimization process the discrepancy between the voltage
758 traces (or other time series) generated by the optimized model and the target data [15]. These
759 cost functions range from general ones such as the mean squared error to more specific ones
760 that are useful mainly in the context of fitting neuronal voltage responses and characterize the
761 pattern and shape of action potentials (Table 1).

762

Feature name	Definition
Mean squared error	Mean squared difference between the model trace and the target trace point by point, normalized by the squared range of the target data
Mean squared error (excluding spikes)	Same as above but excludes the parts of both traces in the vicinity of action potentials (in either trace)
Derivative difference	Normalized mean squared difference of the temporal derivatives of the two traces

Spike count	Absolute difference of the number of spikes in the entire traces, normalized by the sum of the two spike counts
Spike count (during stimulus)	Identical to spike count, except it only takes into account the action potentials during the stimulus
ISI difference	Sum of the absolute differences of the inter-spike intervals of the two traces, normalized by the length of the traces
Latency to 1st spike	Squared difference between the time to the first spike from the start of the stimulus in the two traces, normalized by the squared length of the traces
AP overshoot	First calculates the amplitudes of the action potentials in both traces as the difference between the AP peak voltage and the AP threshold, then takes the mean squared difference of the AP amplitudes normalized by the squared maximal amplitude of the target trace
AP width	Mean squared difference between the width of the action potentials at their base (at the threshold voltage level), normalized by the squared mean width of the APs in the target trace
AHP depth	The squared mean of the difference in the corresponding after-hyperpolarization depths, normalized by the squared range of subthreshold potential in the target trace

763 Table 1: **Cost functions implemented in Neuroptimus.**

764

765 The error functions above (which were already present in Optimizer; [15]) compare
766 each voltage (or current) trace generated by a model with a specific voltage (or current) trace
767 in the target data. However, a common task in single cell modeling involves finding model
768 parameters such that the behavior of the model becomes similar to the typical behavior within
769 a set of experimentally recorded neurons [5, 7]. In this case, it is more natural to define the
770 target of the optimization as the mean values of a set of pre-selected features extracted from the
771 experimental voltage traces (which may come from several experiments involving the same or
772 different neurons). Then the natural way of defining error functions is by evaluating the
773 difference between the value of a particular feature extracted from the voltage response of the
774 model and the mean value of the same feature in the experiments, divided by the standard
775 deviation of the feature in the experimental data. One additional advantage of this definition is
776 that it provides standardized, dimensionless error scores that may be combined in a
777 straightforward manner.

778 This approach based on feature statistics is now supported by Neuroptimus. To provide
779 access to a diverse array of electrophysiological features, and ensure compatibility with some
780 common workflows [5, 7, 44, 45], Neuroptimus utilizes the Electrophys Feature Extraction
781 Library (eFEL; <https://github.com/BlueBrain/eFEL>) [27] to characterize the voltage responses
782 of the models. The target data in this case contain the experimental mean and standard deviation
783 values of a predefined set of eFEL features extracted from voltage responses to specific current
784 step inputs, stored in a JSON file created from the recordings using the BluePyEfe tool
785 (<https://github.com/BlueBrain/BluePyEfe>) and a custom script that converts the output of
786 BluePyEfe to the format expected by Neuroptimus. This JSON file also contains the full
787 specification of the stimulation protocols. When the optimization is run using the GUI, the

788 settings of the stimuli and the features are automatically loaded into the GUI from this input
789 file. During the optimization process, in every model evaluation step, the features included in
790 the input file (and selected in the GUI) are extracted from the model's voltage traces, and errors
791 are computed for every feature using the feature statistics-based error function described above
792 [46].

793 Recent studies classify optimization problems according to the cardinality of objectives
794 as single, multi- (2-3 objectives) and many-objective tasks (more than 3 dimensions), which
795 affects the nature and the complexity of the problem [47]. However, we characterized our
796 problems simply as single- or multi-objective problems because these require different internal
797 representations and are solved by different algorithms. Multi-objective problems involve
798 several objective functions that are to be minimized simultaneously and require finding a set of
799 solutions that give the best tradeoffs between the objectives.

800 Neuroptimus makes it possible to use arbitrary weighted sums of error functions as the
801 ultimate objective function of the parameter search. When single-objective algorithms are used,
802 the weighted sum is calculated for every model during the optimization process, and is used as
803 the objective function. In the case of multi-objective algorithms, all the error functions are
804 treated as separate objectives during the optimization, but the weighted sum is still used after
805 running the search to select a single preferred solution from those returned by the algorithm [5,
806 23].

807 Neuroptimus supports parameter optimization algorithms implemented by five external
808 Python packages (Pygmo, Inspyred, BluePyOpt, Scipy, and Cmaes), and also contains an
809 internal implementation of a simple random search algorithm that takes independent, uniformly
810 distributed samples from the entire search space. Pygmo is a general-purpose scientific Python
811 library for optimization, based on the C++ library pagmo, which implements many different
812 optimization algorithms in a common framework [48]. Inspyred is a Python library specifically

813 developed for bio-inspired (mainly evolutionary) computation, and was already supported by
814 Optimizer [15]. The Blue Brain Python Optimization Library (BluePyOpt) is a software
815 framework developed at the Swiss Blue Brain Project [23], which implements multi-objective
816 optimization algorithms including the Indicator Based Evolutionary Algorithm (IBEA), and has
817 been applied successfully in several computational neuroscience projects [5, 7, 44, 49, 50, 51,
818 52, 53]. SciPy [54] provides implementations of various methods for scientific computation,
819 and includes several basic optimization algorithms, some of which were already supported by
820 Optimizer [15]. Finally, we also included the Cmaes package because it provides a simple,
821 robust, and easily parallelizable implementation of the Covariance Matrix Adaptation Evolution
822 Strategy (CMAES) algorithm, a popular and powerful search method that is also included in
823 Pygmo but in an implementation that does not support the parallel evaluation of models within
824 a population.

825 Some of the algorithms are local (essentially gradient-based) search methods, but most
826 of them are based on metaheuristics that attempt to find the global minimum of the cost
827 function(s). Many of the most popular single- and multi-objective optimization algorithms are
828 included. Most of the algorithms also have parameters that are configurable through the GUI
829 or the configuration file.

830 Solving nontrivial parameter optimization problems typically requires the evaluation of
831 many parameter combinations. In our case, this corresponds to running a large number of
832 simulations, which may take a prohibitively long time if simulations are performed sequentially,
833 especially for complex models such as morphologically detailed neurons, circuits, or multi-
834 scale models that include biochemical or molecular processes. Fortunately, many global
835 optimization methods (including evolutionary and swarm intelligence algorithms) can be
836 implemented in a way that populations of models can be evaluated in parallel, and several such
837 parallel (or easily parallelizable) implementations are included in the Python libraries supported

838 by Neuroptimus. However, Python provides several different methods for parallel execution of
839 code, and the optimization packages we use differ in terms of which parallelization approaches
840 they support. As a consequence, Neuroptimus uses the *multiprocessing* module for the parallel
841 execution of algorithms in Pygmo, Inspyred and Cmaes, while it relies on the IPython Parallel
842 (*ipyparallel*) package to run the algorithms of BluePyOpt in parallel. We note that some
843 optimization algorithms cannot be efficiently parallelized, while for some others (including
844 several in the Pygmo package) parallel execution is not currently supported by the optimization
845 library.

846 Batch evaluation of the models is a requisite to use one of the various parallelization
847 strategies in the Neuroptimus. Therefore both internal and external evaluations have to generate
848 results simultaneously. If we used a single model instance in every process, the results could be
849 mixed or swapped. Therefore, when simulations are carried out within Neuroptimus (using
850 NEURON), a new model instance is created for every parameter set generated by the selected
851 algorithm, and every evaluation running in parallel is performed with a separate model. In case
852 of using the external simulator, every individual is evaluated in a separate subprocess, and files
853 with unique names are used for communication between Neuroptimus and the external
854 simulation script.

855 We list all of the available algorithms along with their basic properties in Table 2. Many
856 of these algorithms were tested in our benchmarking study, and these will be described in more
857 detail below.

858

Algorithm	Objectives	Packages	Parallelization
-----------	------------	----------	-----------------

Custom Evolutionary Optimization (CEO)	Single	Inspyred	multiprocessing
Classic Evolution Strategy (CES)	Single	Inspyred	multiprocessing
Particle Swarm Optimization (PSO)	Single	Inspyred	multiprocessing
		Pygmo	None
Non-dominated Sorting Genetic Algorithm (NSGAII)	Multi	Inspyred	multiprocessing
		Pygmo	
		Bluepyopt	ipyparallel
Differential Evolution (DE)	Single	Inspyred	multiprocessing
		Pygmo	None
Pareto Archived Evolution Strategy (PAES)	Multi	Inspyred	multiprocessing
Basin-Hopping (BH)	Single	Scipy	None
Nelder-Mead (NM)	Single	Scipy	None
limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm with bound constraints (L-BFGS-B)	Single	Scipy	None

Self-Adaptive Differential Evolution (SADE)	Single	Pygmo	None
Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES)	Single	Pygmo	None
		Cmaes	multiprocessing
Exponential Natural Evolution Strategies (XNES)	Single	Pygmo	None
Extended Ant Colony Optimization (GACO)	Single	Pygmo	multiprocessing
Multi-objective Hypervolume-based Ant Colony Optimization (MACO)	Multi	Pygmo	multiprocessing
Particle Swarm Optimization Generational (PSOG)	Single	Pygmo	multiprocessing
Non-dominated Sorting Particle Swarm Optimization (NSPSO)	Multi	Pygmo	multiprocessing
Indicator Based Evolutionary Algorithm (IBEA)	Multi	Bluepyopt	ipyparallel
Simulated Annealing (SA)	Single	Inspyred	multiprocessing

Praxis	Single	Pygmo	None
Simple Genetic Algorithm (SGA)	Single	Inspyred	multiprocessing
Estimation of distribution algorithm (EDA)	Single	Inspyred	multiprocessing
Artificial Bee Colony (ABC)	Single	Pygmo	None
Differential Evolution 1220	Single	Pygmo	None

859 Table 2: **Algorithms included in Neuroptimus.** The properties listed include the full name of the algorithm, the
860 abbreviation used in this article, the type according to the number of objectives (single/multi-objective), the
861 implementing package(s), and the method of parallelization used in Neuroptimus (None if only serial execution is
862 supported).

863
864 The easiest way to perform parameter optimization using Neuroptimus is by using the
865 GUI, whose seven tabs guide the user through the steps of setting up, running, and evaluating
866 the results of the parameter search. The GUI allows the user to load the target data, select the
867 model and the parameters to be optimized, set up the stimulation and recording conditions,
868 configure the error function(s), run the parameter search, and then visualize and analyze the
869 results. The final as well as intermediate results of the optimization are also saved to files, and
870 can be analyzed outside Neuroptimus. This includes the parameters and errors of each simulated
871 model as well as the statistics of generations saved into text files, the voltage trace of the best
872 model saved into text files and in several image formats, and a final summary of the
873 optimization process and the results saved into an HTML file for visual inspection through a
874 web browser and a JSON file for a machine-readable non-SQL data representation. This final

875 metadata file created after the optimization contains automatically generated names for the
876 optimization and the model, details of the parameters of the model (name, boundaries and
877 optimal values), details of the error functions used to calculate the final error (name, value,
878 weight, weighted value), settings of the target data, the algorithm and package used for the
879 optimization, parameters used by the algorithm, and finally the statistics of each generation.

880 The program also saves the full configuration of the optimization task, and the resulting
881 configuration file can be used (directly, or after suitable modifications) by the command-line
882 interface of Neuroptimus to re-run the optimization (with the same or modified settings). This
883 method was used in our benchmarking study to run batches of the same optimization with
884 different random seeds, using a simple Python script to edit the configuration file and create
885 multiple versions of the optimization task.

886

887 **Neuroptimus server**

888 To share our results in a way that allows easy replication and further analysis, and to
889 enable the straightforward extension and updating of the current study, we created an online
890 database of optimization results that is accessible via a web interface. We designed, created and
891 deployed the Neuroptimus web-server, which can be publicly accessed at
892 <https://neuroptimus.koki.hu> and enables all users to browse, view and analyze the optimization
893 results stored in the database. Furthermore, authenticated users can also upload their
894 optimizations and compare their results with previously uploaded ones.

895 The Neuroptimus server structure consists of an Nginx web server that handles the
896 requests and responses, the frontend implemented using the JavaScript library ReactJS, the
897 backend created in the Python web framework Django, backed up by a PostgreSQL database
898 connection. The site handles the authentication of registered users, uploading of optimization

899 results via a web form, visualization of the data in a table structure, and the creation of plots for
900 comparison. The database stores information about the optimization itself, the model used for
901 the optimization and its parameters, the algorithm and its configuration, details about the target
902 data, the statistics of each generation produced by the algorithm, creation time of results and
903 upload time. The metadata JSON file created by Neuroptimus can be uploaded to the server and
904 all of its information content is transferred to the database automatically. Optionally the
905 compressed optimization files can also be uploaded and subsequently downloaded. Analysis of
906 the optimization can be created semi-automatically by selecting the desired algorithms for
907 comparison and visualizing them on the charts. Thus far generation plots, final and convergence
908 score box plots are available for online observation.

909 During the registration process users need to provide their name, affiliation, and email
910 address, choose a username, and create a password. Verifying email addresses grants
911 permission for users to upload their optimizations. Forgotten passwords can be reset on the
912 website via email verification.

913

914 Optimization algorithms

915 In the current study, we evaluated a large set of parameter search algorithms, including
916 several of the most widely used single-objective and multi-objective methods. Our optimization
917 tool supports optimization algorithms implemented by five separate Python packages: Inspyred
918 [55], Pygmo [48], BluePyOpt [23], Cmaes [56], and Scipy [54]. Table 2 shows which packages
919 implement each of the supported algorithms. The majority of these algorithms can be
920 categorized as evolutionary or nature-inspired metaheuristics.

921 Due to constraints on time and computational resources, we could not include every
922 single algorithm supported by Neuroptimus in the detailed comparison that we performed using

923 our neural benchmarking suite (see below). However, we aimed to provide good coverage of
924 algorithms that were used previously in neuronal optimization [9, 15, 17, 23], and also included
925 several additional algorithms that consistently provided good performance in other settings [57,
926 58, 59, 60].

927 Finally, we added some basic search methods such as uniform random sampling and
928 two widely used local optimization algorithms to provide a baseline against which we can
929 measure the performance of more sophisticated methods. The following algorithms were tested
930 in our neural optimization benchmark.

931 **Baseline algorithm**

932 The *Random Search (RAND)* algorithm is the simplest heuristic to discover solutions by trial
933 and error. This is our baseline method, which samples parameters from the search space
934 repeatedly based on the uniform probability distribution. Neuroptimus uses our own
935 implementation of this method [61].

936 **Local optimization algorithms**

937 The *Nelder-Mead (NM)* algorithm is a classic simplex-based direct search method to find a
938 local minimum of the cost function [62].

939 The *limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm with bound constraints (L-BFGS-B)* is considered to be a modern and efficient algorithm that aims to find a local minimum
940 of the objective function using a limited amount of computer memory [63].

942 **Single-objective global optimization algorithms**

943 The *Custom Evolutionary Optimization (CEO)* algorithm is a relatively simple member of the
944 large class of evolutionary optimization algorithms. Evolutionary algorithms are metaheuristics

945 for global optimization inspired by biological evolution. Each candidate solution, represented
946 by a particular combination of the unknown parameters, is considered to be an individual within
947 a population, and the value of the cost function for that parameter combination is treated as the
948 “fitness” of that individual (with lower costs normally associated with higher fitness). The
949 initial population typically consists of random samples from the search space. The population
950 is then updated through the application of various operators. New individuals are generated via
951 the application of genetic operators such as mutation, which introduces random variations into
952 the parameters of an individual, and crossover, which randomly combines the parameters of
953 two individuals. The size of the population is maintained by selecting individuals with higher
954 fitness. These steps are repeated iteratively for a certain number of generations. Many different
955 variants of evolutionary algorithms exist that differ in the details of the operators, and may also
956 apply additional heuristics. The CEO algorithm is based on the EvolutionaryComputation class
957 of the Inspyred package, and uses Gaussian mutation and blend crossover variators.

958 The *Classic Evolution Strategy (CES)* algorithm belongs to a subclass of evolutionary
959 optimization algorithms called evolution strategies. In these algorithms, there are distinct
960 mutation rates associated with each parameter, and these mutation rates are changed adaptively
961 during the optimization [64].

962 The *Covariance Matrix Adaptation Evolution Strategy (CMA-ES)* algorithm is an evolutionary
963 algorithm which samples candidate solutions from multivariate normal distributions with
964 adapting mean and covariance matrix [65].

965 The *Exponential Natural Evolution Strategy (XNES)* algorithm is an evolution strategy (ES)
966 that uses the natural gradient to update the search distribution [66].

967 The *Differential Evolution (DE)* algorithm is an evolutionary algorithm that generates new
968 candidate solutions from existing individuals based on some simple mathematical rules [57].

969 The *Self-Adaptive Differential Evolution (SADE)* algorithm is a version of the Differential
970 Evolution algorithm, which adjusts the mutation rate and the crossover rate adaptively [67].

971 The *Particle Swarm Optimization (PSO)* algorithm represents candidate solutions as particles
972 moving around in the search space. Each particle has a velocity and moves by adding this to
973 its current position in every iteration. Initially the velocity is random, and it is modified after
974 each iteration, influenced by the currently known best positions for the individual particles
975 and that of the entire group. In this basic implementation, velocity and position updates are
976 carried out sequentially for each particle [68].

977 The *Particle Swarm Optimization Generational (PSOG)* algorithm is similar to the PSO
978 algorithm above but, in every iteration, it first updates the velocity for all particles, then updates
979 the positions. This allows efficient parallel execution of the algorithm.

980 The *Extended Ant Colony Optimization (GACO)* algorithm is a bio-inspired algorithm based on
981 the analogy of ants finding paths from colony to food. In this algorithm, artificial agents move
982 through the parameter space, and lay down “pheromones” depending on the quality of the
983 solutions they find. These pheromones attract the other agents, making it more likely that they
984 move to locations with high amounts of pheromone. This extended version of the algorithm
985 calculates the locations of future generations of ants by sampling from a multi-kernel Gaussian
986 distribution that depends on the quality of previously found solutions [69].

987 The *Basin-Hopping (BH)* algorithm is a generalization of the Simulated Annealing algorithm
988 that was used in several earlier studies of neural parameter optimization [15, 17]. Basin-hopping
989 is a two-level algorithm: its outer loop performs stochastic jumps in the search space, while the
990 inner loop performs local optimization. The resulting new local minimum is always accepted if
991 it is better than the previous one, but it may also be accepted if it is worse with a probability
992 that depends on the increase in the cost function [70].

993 **Multi-objective global optimization algorithms**

994 The *Non-dominated Sorting Genetic Algorithm II (NSGA2)* is an evolutionary multi-objective
995 algorithm. Multi-objective optimization algorithms aim to optimize several cost functions
996 simultaneously, trying to find non-dominated (or Pareto-optimal) solutions where none of the
997 cost functions can be improved without degrading the performance on some other cost
998 functions. The algorithms also aim to create a diverse set of solutions that collectively provide
999 good coverage of the Pareto front. In NSGA2, a child population is created from the parent
1000 population using the usual genetic operators, mutation and crossover. Individuals in the next
1001 generation are then selected from the joint population based on Pareto dominance and the so-
1002 called crowding distance that penalizes closely related individuals and helps maintain diversity
1003 within the population [58].

1004 The *Pareto Archived Evolution Strategy (PAES)* algorithm is a simple multi-objective
1005 algorithm that uses local search (mutation) from the current individual(s) and maintains a
1006 reference archive of previously found non-dominated solutions to approximate the dominance
1007 ranking of candidate solutions [71].

1008 The *Indicator Based Evolutionary Algorithm (IBEA)* is a multi-objective evolutionary
1009 algorithm that computes the fitness value based on predefined binary indicators. It performs
1010 environmental selection by removing the worst individuals, chooses parents by comparing the
1011 fitness values of randomly selected pairs of individuals, and applies mutation and crossover to
1012 create offspring, repeating the process iteratively until reaching the maximum number of
1013 generations [72].

1014 The *Multi-objective Hypervolume-based Ant Colony Optimizer (MACO)* is a multi-objective
1015 optimization algorithm that extends the GACO algorithm described above, combining
1016 hypervolume computation and non-dominated fronts for ranking individuals [73].

1017 The *Non-dominated Sorting PSO (NSPSO)* algorithm extends PSO by making a better use of
1018 personal bests and offspring for non-dominated comparison [74].

1019

1020 **Use cases**

1021 To compare the efficiency of various parameter search methods in solving neuronal
1022 parameter optimization tasks, we designed and implemented a suite of six different problems
1023 that may be considered typical use cases in this domain. All of these use cases can be handled
1024 by Neuroptimus, which allowed us to run all benchmarks using every selected algorithm within
1025 the same framework, and made the subsequent evaluation of their performance quite
1026 straightforward (see below). Five of the use cases were similar (or identical) to those presented
1027 in [15], although some of them were modified to increase the robustness of the simulations
1028 (avoiding errors due to invalid parameter combinations, in the case of the AdExpIF example)
1029 or to move the target behavior of the model away from a critical boundary (the transition to
1030 repetitive firing, in the case of the Hodgkin-Huxley model). We provide a description of each
1031 use case below; all the files required to run these examples, along with detailed guides to setting
1032 up the optimizations in the Neuroptimus GUI, can be found in the corresponding subfolders of
1033 the `neuroptimus/new_test_files` directory of the Neuroptimus Github repository
1034 (<https://github.com/KaliLab/neuroptimus>).

1035 **Hodgkin-Huxley**

1036 This use case is based on a single-compartment model, which contains conductances
1037 from the original Hodgkin-Huxley model (Na⁺, K⁺, leak)[28], and is implemented in
1038 NEURON. To generate the target voltage trace, a suprathreshold step current was injected into
1039 the soma of the neuron model (amplitude = 300 pA, delay = 200 ms, duration = 500 ms, and

1040 the voltage trace duration is 1000 ms). The test case involves recovering the correct
1041 conductance densities (3 parameters) that were used to generate the target trace, while keeping
1042 the properties of the currents and the other parameters of the model constant (at their original
1043 value). A combination of four features (spike count, spike amplitude, spike width, mean squared
1044 error of voltage excluding spikes) was used to compare each simulated trace to the original
1045 (target) trace.

1046 **Voltage Clamp**

1047 In the Voltage Clamp benchmark problem the same single-compartment model with the
1048 same conductances is used as in the Hodgkin-Huxley problem. In addition, this model contains
1049 a conductance-based synapse. The goal here is to recover the synaptic parameters (weight, rise
1050 and decay time constants, delay – 4 parameters) from simulated voltage clamp recordings
1051 during synaptic stimulation (four presynaptic spikes at 10 Hz), using the mean squared error
1052 cost function to compare the current traces.

1053 **Passive, anatomically detailed neuron**

1054 This benchmark uses a morphologically detailed passive model of a hippocampal CA1
1055 pyramidal cell implemented in NEURON. During the experiment, a short (3 ms, 500 pA) and
1056 a long (600 ms, 10 pA) current pulse (separated by 300 ms) were injected into the soma, and
1057 the membrane potential was also recorded there. The neuron was filled with a dye during the
1058 recording, and was reconstructed using Neurolucida. This reconstruction defines the
1059 morphology of the model, and the task involves fitting 3 passive parameters (specific
1060 capacitance, leak conductance density, specific axial resistance, all of which are assumed to be
1061 uniform within the cell) to reproduce the experimental data recorded using the same complex
1062 current clamp stimulus. Traces are compared via the mean squared error cost function. All the

1063 experimental data for this use case, including the morphological reconstruction and the
1064 electrophysiological recordings, were provided by Miklós Szoboszlay and Zoltán Nusser.

1065 **Simplified model**

1066 This use case attempts to fit the behavior of a six-compartmental simplification of a
1067 biophysically accurate and morphologically detailed hippocampal CA1 pyramidal cell model
1068 [35] to the somatic voltage responses of the original model with full morphology. Both models
1069 contained the same set of voltage-gated conductances in their somatic and dendritic
1070 compartments: transient Na channels (separate somatic and dendritic subtypes), delayed
1071 rectifier, A-type, and M-type voltage-gated K channels, C-type and AHP-associated Ca-
1072 dependent K channels, L-type and N-type Ca channels, and the hyperpolarization-activated
1073 HCN channels. Dendrites of the full model were clustered based on their passive voltage
1074 responses, and each of these clusters defined a dendritic compartment in the simplified model.
1075 The densities of ion channels in the dendritic compartments of the simplified model were set to
1076 the average values in the corresponding clusters of the full model, while the densities of the
1077 nine somatic conductances were subject to parameter optimization. The original full model was
1078 implemented in GENESIS, while the simplified model was implemented in the NEURON
1079 simulator. The target data was the voltage response of the full model to the injection of a 200
1080 pA step current stimulus into the soma (the stimulus started at 200 ms and lasted for 600 ms,
1081 with a total recording duration of 1000 ms). The fit was evaluated via a combination of features
1082 including mean squared error (excluding spikes) weighted by 0.2, spike count (weight 0.4),
1083 latency to first spike (weight 0.1), action potential amplitude (weight 0.1), action potential width
1084 (weight 0.1), and after-hyperpolarization depth (weight 0.1).

1085 **Extended integrate-and-fire model**

1086 In this benchmark problem, the parameters of a phenomenological (adaptive
1087 exponential integrate-and-fire) spiking model [32, 33], implemented in the NEST simulator
1088 [34] were fitted to capture the somatic responses of a real neuron (hippocampal CA3 pyramidal
1089 cell) to four different inputs. Voltage traces were recorded experimentally in response to current
1090 steps of 900 ms duration, and 0.30, 0.35, 0.40, and 0.45 nA amplitudes (the step was delayed
1091 by 100 ms, and the recordings lasted for 1100 ms). Sampling frequency was 5 kHz. The
1092 unknown parameters to be optimized were the capacitance, the leak conductance, the reversal
1093 potential of the leak current, the threshold voltage, the reset voltage, the refractory period, the
1094 steepness of the exponential part of the current-voltage relation, the subthreshold adaptation
1095 conductance, the spike adaptation current, and the adaptation time constant (10 parameters).
1096 During the optimization the mean squared error (excluding spikes), the spike count (during
1097 stimulus), and the latency to first spike error functions were used with equal weights.

1098 **CA1 pyramidal cell**

1099 This is our most complex benchmark problem both regarding the number of parameters
1100 to be optimized and the complexity of the model. The test case is based on an anatomically and
1101 biophysically detailed rat hippocampal CA1 pyramidal cell model built for the NEURON
1102 simulator in our research group. The morphology of the model was from [39]. The model
1103 contained several different voltage-gated ion channels in its somatic, dendritic, and axonal
1104 compartments: a transient Na conductance, delayed rectifier, A-type, M-type, and D-type
1105 voltage-gated K conductances, and the hyperpolarization-activated current Ih. Many attributes
1106 of the model were well-constrained by experimental observations available in the literature,
1107 including the distributions and kinetic properties of the ion channels. The target data (provided

1108 by Judit Makara) consisted of the means (and associated standard deviations) of 20 different
1109 types of features extracted by eFEL from the voltage responses of five CA1 pyramidal neurons
1110 to somatic current step injections of six different amplitudes (-0.25, 0.05, 0.1, 0.15, 0.2 0.25
1111 nA), with each stimulus repeated three times for every cell. This resulted in a total of 66 feature
1112 values to be matched by the model. The eFEL features and the associated current step
1113 amplitudes are listed in Table 3.

1114

Feature name from eFEL	Feature description	Fitted current steps (nA)
Spikecount	Number of spikes during stimulus	0.05, 0.1, 0.15, 0.2, 0.25
inv_first_ISI	Reciprocal of first interspike interval	0.05, 0.1, 0.15, 0.2, 0.25
inv_last_ISI	Reciprocal of last interspike interval	0.05, 0.1, 0.15, 0.2, 0.25
inv_time_to_first_spike	Reciprocal of time to first spike	0.05, 0.1, 0.15, 0.2, 0.25
steady_state_voltage	Average voltage after stimulus	-0.25, 0.05, 0.1, 0.15, 0.2, 0.25
voltage_base	Average voltage during the last 10% of time before the stimulus	-0.25, 0.05, 0.1, 0.15, 0.2, 0.25

voltage_deflection	Difference between the voltage base and the steady-state voltage at the end of the stimulus	-0.25
voltage_deflection_begin	Difference between the voltage base and the mean voltage in the early phase (5 to 15% of duration) of the stimulus	-0.25
AHP_depth_abs	Absolute value of the depth of afterhyperpolarization	0.15, 0.2, 0.25
AHP_time_from_peak	Time from the peak of the AP to the minimum of the afterhyperpolarization	0.15, 0.2, 0.25
AP2_amp	Amplitude of second AP	0.15, 0.2, 0.25
AP_amplitude	Average height of APs	0.15, 0.2, 0.25
AP_begin_voltage	Threshold of AP initiation	0.15, 0.2, 0.25
AP_duration_half_width	Half-width of the AP	0.15, 0.2, 0.25

AP_fall_time	Time from the peak of the AP to the end of the AP	0.15, 0.2, 0.25
AP_rise_time	Time from the AP threshold to the peak of the AP	0.15, 0.2, 0.25
AP_last_amp	Amplitude of the last AP	0.15, 0.2, 0.25
time_to_last_spike	Time to the last spike	0.15, 0.2, 0.25
sag_amplitude	Difference between the minimal voltage and the steady state voltage at the end of the stimulus	-0.25
sag_ratio1	Ratio between sag amplitude and maximal hyperpolarization from voltage base	-0.25

1115 Table 3: List of eFEL features (with brief explanations) and associated current amplitudes used as the target data

1116 in the detailed CA1 pyramidal neuron use case.

1117

1118 We optimized 12 abstract parameters of the model that were mapped onto the actual

1119 parameters of the NEURON implementation by an appropriate user function. Ten parameters

1120 determined the densities of the voltage-gated and leak conductances in the different
1121 compartments (soma, dendrites, axon) of the neuron; one parameter represented the reversal
1122 potential of the leak current; and the final parameter determined the difference between the half-
1123 activation and half-inactivation potential values of the Na conductance.

1124 **Evaluation Methods**

1125 We tested the different optimization algorithms on each of the six model optimization
1126 tasks described above. To ensure a fair comparison of model performance, we allowed 10,000
1127 model evaluations for every algorithm on each task. For all population-based methods
1128 (including evolutionary algorithms and swarm intelligence-based approaches) we set the
1129 population size to 100, and the number of generations to 100 as well. We similarly set 100
1130 global and 100 local minimization steps for two-stage algorithms. Otherwise, we ran every
1131 algorithm with its default settings in Neuroptimus. These default options are typically the
1132 package default settings, with one significant exception: we observed that the default settings
1133 of the EvolutionaryComputation class of the Inspyred package that underlies our CEO
1134 algorithm led to essentially no optimization, so we adjusted the default number of elites from 0
1135 to half of the population size, changed the mutation rate from 0.1 to 0.25 and standard deviation
1136 of Gaussian mutation from 1 to 0.5.

1137 Optimization runs were parallelized for all algorithms where this is supported by
1138 Neuroptimus and the underlying packages (see Table 2). For the most resource-intensive use
1139 case (the detailed CA1 pyramidal neuron model) these parallelized runs were performed on
1140 supercomputers via the Neuroscience Gateway [75]; simpler use cases and algorithms that do
1141 not support parallelization were run on a Dell PowerEdge R730 compute server or personal
1142 computers. To allow meaningful statistical comparisons between the algorithms, we performed
1143 10 independent runs (using distinct random seeds) of each algorithm in every use case.

1144 We visualized and compared the performance of the algorithms in each use case using
1145 several different methods. All the comparisons were based on the change in the total error
1146 during the optimization. First, we visualized the convergence of the algorithms by plotting the
1147 cumulative minimum of the error function after every generation (i.e., after every 100 model
1148 evaluations). We plotted the median value across the 10 runs to see which algorithms typically
1149 find the best solutions after a given number of model evaluations. The lowest and the highest
1150 errors achieved by the 10 runs were also calculated in every iteration to observe how well the
1151 algorithm performs in the best case, and whether it gets stuck in some cases.

1152 We defined two basic scores to characterize and compare the performance of the
1153 algorithms in a concise manner. The first of these scores was defined as the lowest error
1154 achieved during the entire optimization run (these are usually, but not always, associated with
1155 members of the final population). We visualized the distribution of this measure across the 10
1156 independent runs using box plots that show the median, interquartile range, minimum, and
1157 maximum values, and also indicate apparent outliers.

1158 In the case of more complex, detailed models, each model evaluation (simulation) can
1159 be time-consuming, and thus we are also interested in which algorithms can find a reasonably
1160 good solution in a relatively short time. To characterize the convergence speed of an algorithm,
1161 we used the sum of the logarithms of the error scores achieved by the best individuals in each
1162 generation. This is essentially the area under the logarithmic convergence curve - the smaller
1163 this sum is, the faster the algorithm found a relatively good solution.

1164

1165

1166 **References**

1167

1168 1. Herz AVM, Gollisch T, Machens CK, Jaeger D. Modeling Single-Neuron Dynamics and Computations:
1169 A Balance of Detail and Abstraction. *Science*. 2006 Oct 6;314(5796):80–5.

1170 2. Einevoll GT, Destexhe A, Diesmann M, Grün S, Jirsa V, de Kamps M, et al. The Scientific Case for
1171 Brain Simulations. *Neuron*. 2019 May;102(4):735–44.

1172 3. Ramaswamy S. Data-driven multiscale computational models of cortical and subcortical regions.
1173 *Current Opinion in Neurobiology*. 2024 Apr;85:102842.

1174 4. Hay E, Hill S, Schürmann F, Markram H, Segev I. Models of Neocortical Layer 5b Pyramidal Cells
1175 Capturing a Wide Range of Dendritic and Perisomatic Active Properties. *PLoS Computational Biology*. 2011 Jul
1176 28;7(7):e1002107.

1177 5. Markram H, Muller E, Ramaswamy S, Reimann MW, Abdellah M, Sanchez CA, et al. Reconstruction
1178 and Simulation of Neocortical Microcircuitry. *Cell*. 2015 Oct;163(2):456–92.

1179 6. Gouwens NW, Berg J, Feng D, Sorensen SA, Zeng H, Hawrylycz MJ, et al. Systematic generation of
1180 biophysically detailed models for diverse cortical neuron types. *Nature Communications*. 2018 Feb 19;9(1).

1181 7. Migliore R, Lupascu CA, Bologna LL, Romani A, Courcol JD, Antonel S, et al. The physiological
1182 variability of channel density in hippocampal CA1 pyramidal cells and interneurons explored using a unified
1183 data-driven modeling workflow. *PLOS Computational Biology*. 2018 Sep 17;14(9):e1006423.

1184 8. Van Geit W, De Schutter E, Achard P. Automated neuron model optimization techniques: a review.
1185 *Biological Cybernetics*. 2008 Nov;99(4–5):241–51.

1186 9. Druckmann S. A novel multiple objective optimization framework for constraining conductance-based
1187 neuron models by experimental data. *Frontiers in Neuroscience*. 2007 Nov 1;1(1):7–18.

1188 10. Keren N, Peled N, Korngreen A. Constraining Compartmental Models Using Multiple Voltage
1189 Recordings and Genetic Algorithms. *Journal of Neurophysiology*. 2005 Dec;94(6):3730–42.

1190 11. Weaver CM, Wearne SL. The role of action potential shape and parameter constraints in optimization
1191 of compartment models. *Neurocomputing*. 2006 Jun;69(10–12):1053–7.

1192 12. Rossant C. Automatic fitting of spiking neuron models to electrophysiological recordings. *Frontiers in
1193 Neuroinformatics*. 2010;4.

1194 13. Hendrickson EB, Edgerton JR, Jaeger D. The use of automated parameter searches to improve ion
1195 channel kinetics for neural modeling. *Journal of Computational Neuroscience*. 2011 Jan 18;31(2):329–46.

1196 14. Svensson CM, Coombes S, Peirce JW. Using Evolutionary Algorithms for Fitting High-Dimensional
1197 Models to Neuronal Data. *Neuroinformatics*. 2012 Jan 20;10(2):199–218.

1198 15. Friedrich P, Vella M, Gulyás AI, Freund TF, Káli S. A flexible, interactive software tool for fitting the
1199 parameters of neuronal models. *Frontiers in Neuroinformatics*. 2014 Jul 10;8.

1200 16. Nogaret A, Meliza CD, Margoliash D, Abarbanel HDI. Automatic Construction of Predictive Neuron
1201 Models through Large Scale Assimilation of Electrophysiological Data. *Scientific Reports*. 2016 Sep 8;6(1).

1202 17. Vanier MC, Bower JM. *Journal of Computational Neuroscience*. 1999;7(2):149–71.

1203 18. Wolpert DH, Macready WG. No free lunch theorems for optimization. *IEEE Transactions on
1204 Evolutionary Computation*. 1997 Apr;1(1):67–82.

1205 19. Yang XS. *Nature-inspired Metaheuristic Algorithms*. Luniver Press; 2010.

1206 20. Carnevale NT, Hines ML. *The NEURON Book*. Cambridge University Press; 2006.

1207 21. Bower JM, Beeman D. *The Book of GENESIS: Exploring Realistic Neural Models with the GEneral
1208 NEural SImulation System*. Springer; 1998.

1209 22. Van Geit W. Neurofitter: A parameter tuning package for a wide range of electrophysiological neuron
1210 models. *Frontiers in Neuroinformatics*. 2007;1.

1211 23. Van Geit W, Gevaert M, Chindemi G, Rössert C, Courcol JD, Muller EB, et al. BluePyOpt: Leveraging
1212 Open Source Software and Cloud Infrastructure to Optimise Model Parameters in Neuroscience. *Frontiers in
1213 Neuroinformatics*. 2016 Jun 7;10.

1214 24. Meyer R, Obermayer K. pypet: A Python Toolkit for Data Management of Parameter Explorations.
1215 *Frontiers in Neuroinformatics*. 2016 Aug 25;10.

1216 25. Welcome to Neurotune’s documentation! — neurotune 0.2.1 documentation [Internet]. Available from:
1217 <https://neurotune.readthedocs.io/en/latest/>

1218 26. Dura-Bernal S, Suter BA, Gleeson P, Cantarelli M, Quintana A, Rodriguez F, et al. NetPyNE, a tool for
1219 data-driven multiscale modeling of brain circuits. *eLife*. 2019 Apr 26;8.

1220 27. Ranjan R, Geit V, Moor R, Rössert C, Riquelme JL, Damart T, et al. eFEL [Internet]. Zenodo. 2024.
1221 Available from: <https://doi.org/10.5281/zenodo.593869>

1222 28. Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to
1223 conduction and excitation in nerve. *The Journal of Physiology*. 1952 Aug 28;117(4):500–44.

1224 29. Major G, Larkman A, Jonas P, Sakmann B, Jack J. Detailed passive cable models of whole-cell
1225 recorded CA3 pyramidal neurons in rat hippocampal slices. *The Journal of Neuroscience*. 1994 Aug
1226 1;14(8):4613–38.

1227 30. Stuart G, Spruston N. Determinants of Voltage Attenuation in Neocortical Pyramidal Neuron Dendrites.
1228 *The Journal of Neuroscience*. 1998 May 15;18(10):3501–10.

1229 31. Golding NL, Mickus TJ, Katz Y, Kath WL, Spruston N. Factors mediating powerful voltage attenuation
1230 along CA1 pyramidal neuron dendrites. *The Journal of Physiology*. 2005 Sep 28;568(1):69–82.

1231 32. Naud R, Marcille N, Clopath C, Gerstner W. Firing patterns in the adaptive exponential integrate-and-
1232 fire model. *Biological Cybernetics*. 2008 Nov;99(4–5):335–47.

1233 33. Brette R, Gerstner W. Adaptive Exponential Integrate-and-Fire Model as an Effective Description of
1234 Neuronal Activity. *Journal of Neurophysiology*. 2005 Nov;94(5):3637–42.

1235 34. Fardet T, Vennemo SB, Mitchell J, Mørk H, Gruber S, Hahne J, et al. NEST 2.20.1 [Internet]. Zenodo.
1236 2020 [cited 2024 Feb 27]. Available from: <https://zenodo.org/record/4018718>

1237 35. Káli S, Freund TF. Distinct properties of two major excitatory inputs to hippocampal pyramidal cells: a
1238 computational study. *European Journal of Neuroscience*. 2005 Oct;22(8):2027–48.

1239 36. De Schutter E, Bower JM. An active membrane model of the cerebellar Purkinje cell. I. Simulation of
1240 current clamps in slice. *Journal of Neurophysiology*. 1994 Jan 1;71(1):375–400.

1241 37. Poirazi P, Brannon T, Mel BW. Arithmetic of Subthreshold Synaptic Summation in a Model CA1
1242 Pyramidal Cell. *Neuron*. 2003 Mar;37(6):977–87.

1243 38. Almog M, Korngreen A. Is realistic neuronal modeling realistic? *Journal of Neurophysiology*. 2016
1244 Nov 1;116(5):2180–209.

1245 39. Megías M, Emri Z, Freund TF, Gulyás AI. Total number and distribution of inhibitory and excitatory
1246 synapses on hippocampal CA1 pyramidal cells. *Neuroscience*. 2001 Feb;102(3):527–40.

1247 40. Biscani F, Izzo D. A parallel global multiobjective framework for optimization: pagmo. *Journal of
1248 Open Source Software*. 2020 Sep 13;5(53):2338.

1249 41. Damart T, Van Geit W, Markram H. Data driven building of realistic neuron model using IBEA and
1250 CMA evolution strategies. In: Proceedings of the 2020 Genetic and Evolutionary Computation Conference
1251 Companion [Internet]. New York, NY, USA: ACM; 2020 [cited 2024 Feb 28]. Available from:
1252 <http://dx.doi.org/10.1145/3377929.3398161>

1253 42. Jędrzejewski-Szmeik Z, Abrahao KP, Jędrzejewska-Szmeik J, Loveridge DM, Blackwell KT. Parameter
1254 Optimization Using Covariance Matrix Adaptation—Evolutionary Strategy (CMA-ES), an Approach to
1255 Investigate Differences in Channel Properties Between Neuron Subtypes. *Frontiers in Neuroinformatics*. 2018
1256 Jul 31;12.

1257 43. Comparing algorithms: the case of xNES and CMA-ES — pygmo 2.19.6 documentation [Internet].
1258 [cited 2024 Feb 28]. Available from: https://esa.github.io/pygmo2/tutorials/cmaes_vs_xnes.html#comparing-algorithms-the-case-of-xnes-and-cma-es

1260 44. Ecker A, Romani A, Sáray S, Káli S, Migliore M, Falck J, et al. Data-driven integration of hippocampal
1261 CA1 synaptic physiology in silico. *Hippocampus*. 2020 Jun 10;30(11):1129–45.

1262 45. Sáray S, Rössert CA, Appukuttan S, Migliore R, Vitale P, Lupascu CA, et al. HippoUnit: A software

1263 tool for the automated testing and systematic comparison of detailed models of hippocampal neurons based on
1264 electrophysiological data. *PLOS Computational Biology*. 2021 Jan 29;17(1):e1008114.

1265 46. Druckmann S, Berger TK, Hill S, Schürmann F, Markram H, Segev I. Evaluating automated parameter
1266 constraining procedures of neuron models by experimental and surrogate data. *Biological Cybernetics*. 2008
1267 Nov;99(4–5):371–9.

1268 47. Farzane K, Alireza B D. A review and evaluation of multi and many-objective optimization: Methods
1269 and algorithms. *Global Journal of Ecology*. 2022 Nov 9;7(2):104–19.

1270 48. Biscani F, Izzo D. *esa/pagmo2: pagmo 2.11.1* [Internet]. Zenodo. 2019. Available from:
1271 <https://zenodo.org/record/3364433>

1272 49. Masoli S, Tognolina M, Laforenza U, Moccia F, D’Angelo E. Parameter tuning differentiates granule
1273 cell subtypes enriching transmission properties at the cerebellum input stage. *Communications Biology*. 2020
1274 May 8;3(1).

1275 50. Masoli S, Ottaviani A, Casali S, D’Angelo E. Cerebellar Golgi cell models predict dendritic processing
1276 and mechanisms of synaptic plasticity. *PLOS Computational Biology*. 2020 Dec 30;16(12):e1007937.

1277 51. Ecker A, Bagi B, Vértes E, Steinbach-Németh O, Karlócai MR, Papp OI, et al. Hippocampal sharp
1278 wave-ripples and the associated sequence replay emerge from structured synaptic interactions in a network
1279 model of area CA3. *eLife*. 2022 Jan 18;11.

1280 52. Schneider-Mizell CM, Bodor AL, Collman F, Brittain D, Bleckert A, Dorkenwald S, et al. Structure
1281 and function of axo-axonic inhibition. *eLife*. 2021 Dec 1;10.

1282 53. Chindemi G, Abdellah M, Amsalem O, Benavides-Piccione R, Delattre V, Doron M, et al. A calcium-
1283 based plasticity model for predicting long-term potentiation and depression in the neocortex. *Nature
1284 Communications*. 2022 Jun 1;13(1).

1285 54. Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. *SciPy 1.0: fundamental algorithms for scientific computing in Python*. *Nature Methods*. 2020 Feb 3;17(3):261–72.

1286 55. Tonda A. *Inspyred: Bio-inspired algorithms in Python*. *Genetic Programming and Evolvable Machines*.
2019 Nov 2;21(1–2):269–72.

1287 56. Nomura M, Shibata M. *cmaes : A Simple yet Practical Python Library for CMA-ES* [Internet].
1288 arXiv.org. 2024. Available from: <https://arxiv.org/abs/2402.01373>

1289 57. Price K, Storn RM, Lampinen JA. *Differential Evolution: A Practical Approach to Global Optimization*.
1290 Springer Science & Business Media; 2006.

1291 58. Deb K, Pratap A, Agarwal S, Meyarivan T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*. 2002 Apr;6(2):182–97.

1292 59. Wales DJ, Scheraga HA. Global Optimization of Clusters, Crystals, and Biomolecules. *Science*. 1999
1293 Aug 27;285(5432):1368–72.

1297 60. Gad AG. Particle Swarm Optimization Algorithm and Its Applications: A Systematic Review. *Archives*
1298 of Computational Methods in Engineering. 2022 Apr 19;29(5):2531–61.

1299 61. Rastrigin LA. Random Search in Problems of Optimization, Identification and Training of Control
1300 Systems. *Journal of Cybernetics*. 1973 Jan;3(3):93–103.

1301 62. Nelder JA, Mead R. A Simplex Method for Function Minimization. *The Computer Journal*. 1965 Jan
1302 1;7(4):308–13.

1303 63. Byrd RH, Lu P, Nocedal J, Zhu C. A Limited Memory Algorithm for Bound Constrained Optimization.
1304 *SIAM Journal on Scientific Computing*. 1995 Sep;16(5):1190–208.

1305 64. Bäck T. *Evolutionary Algorithms in Theory and Practice* [Internet]. Oxford University Press; 1996
1306 [cited 2024 Feb 27]. Available from: <http://dx.doi.org/10.1093/oso/9780195099713.001.0001>

1307 65. Hansen N, Ostermeier A. Adapting arbitrary normal mutation distributions in evolution strategies: the
1308 covariance matrix adaptation. In: *Proceedings of IEEE International Conference on Evolutionary Computation*
1309 [Internet]. IEEE; [cited 2024 Feb 27]. Available from: <http://dx.doi.org/10.1109/icec.1996.542381>

1310 66. Glasmachers T, Schaul T, Yi S, Wierstra D, Schmidhuber J. Exponential natural evolution strategies.
1311 In: *Proceedings of the 12th annual conference on Genetic and evolutionary computation* [Internet]. New York,
1312 NY, USA: ACM; 2010 [cited 2024 Feb 27]. Available from: <http://dx.doi.org/10.1145/1830483.1830557>

1313 67. Qin AK, Suganthan PN. Self-adaptive Differential Evolution Algorithm for Numerical Optimization.
1314 In: *2005 IEEE Congress on Evolutionary Computation* [Internet]. IEEE; [cited 2024 Feb 27]. Available from:
1315 <http://dx.doi.org/10.1109/cec.2005.1554904>

1316 68. Kennedy J, Eberhart R. Particle swarm optimization. In: *Proceedings of ICNN'95 - International*
1317 *Conference on Neural Networks* [Internet]. IEEE; [cited 2024 Feb 27]. Available from:
1318 <http://dx.doi.org/10.1109/icnn.1995.488968>

1319 69. Schlüter M, Egea JA, Banga JR. Extended ant colony optimization for non-convex mixed integer
1320 nonlinear programming. *Computers & Operations Research*. 2009 Jul;36(7):2217–29.

1321 70. Wales DJ, Doye JPK. Global Optimization by Basin-Hopping and the Lowest Energy Structures of
1322 Lennard-Jones Clusters Containing up to 110 Atoms. *The Journal of Physical Chemistry A*. 1997 Jul
1323 1;101(28):5111–6.

1324 71. Knowles J, Corne D. The Pareto archived evolution strategy: a new baseline algorithm for Pareto
1325 multiobjective optimisation. In: *Proceedings of the 1999 Congress on Evolutionary Computation-CEC99* (Cat
1326 No 99TH8406) [Internet]. IEEE; [cited 2024 Feb 27]. Available from:
1327 <http://dx.doi.org/10.1109/cec.1999.781913>

1328 72. Zitzler E, Künzli S. Indicator-Based Selection in Multiobjective Search. In: *Lecture Notes in Computer*
1329 *Science* [Internet]. Berlin, Heidelberg: Springer Berlin Heidelberg; 2004 [cited 2024 Feb 27]. p. 832–42.
1330 Available from: http://dx.doi.org/10.1007/978-3-540-30217-9_84

1331 73. Acciarini G, Izzo D, Mooij E. MHACO: a Multi-Objective Hypervolume-Based Ant Colony Optimizer
1332 for Space Trajectory Optimization. In: 2020 IEEE Congress on Evolutionary Computation (CEC) [Internet].
1333 IEEE; 2020 [cited 2024 Feb 27]. Available from: <http://dx.doi.org/10.1109/cec48606.2020.9185694>

1334 74. Li X. A Non-dominated Sorting Particle Swarm Optimizer for Multiobjective Optimization. In: Genetic
1335 and Evolutionary Computation — GECCO 2003 [Internet]. Berlin, Heidelberg: Springer Berlin Heidelberg;
1336 2003 [cited 2024 Feb 27]. p. 37–48. Available from: http://dx.doi.org/10.1007/3-540-45105-6_4

1337 75. Sivagnanam S, Majumdar A, Yoshimoto K, Astakhov V, Bandrowski A, Martone ME, et al.
1338 Introducing the Neuroscience Gateway. Iwsg. 2013;993 of CEUR Workshop Proceedings(CEUR-WS.org).