

# NDI: A platform-independent data interface and database for neuroscience physiology and imaging experiments

<https://doi.org/10.1523/ENEURO.0073-21.2022>

**Cite as:** eNeuro 2022; 10.1523/ENEURO.0073-21.2022

Received: 19 February 2021

Revised: 5 January 2022

Accepted: 10 January 2022

---

*This Early Release article has been peer-reviewed and accepted, but has not been through the composition and copyediting processes. The final version may differ slightly in style or formatting and will contain links to any extended data.*

**Alerts:** Sign up at [www.eneuro.org/alerts](http://www.eneuro.org/alerts) to receive customized email alerts when the fully formatted version of this article is published.

Copyright © 2022 Murillo et al.

This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International license, which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

**Title:** NDI: A platform-independent data interface and database for neuroscience physiology and imaging experiments

**Abbreviated title:** Data interface for neuroscience experiments

**Authors:** Daniel García Murillo (0000-0002-6802-4784)<sup>1,2,5</sup>, Yixin Zhao (0000-0002-8931-350X)<sup>5</sup>, Ora S. Rogovin (0000-0001-8423-2332)<sup>2</sup>, Kelly Zhang (0000-0001-5610-9601)<sup>1,5</sup>, Andrew W. Hu (0000-0001-8758-0316)<sup>1</sup>, Mo Re Kim (0000-0002-5911-8201)<sup>1</sup>, Shufei Chen (0000-0003-2104-1461)<sup>2,5</sup>, Ziqi Wang<sup>1,5</sup>, Zoey C. Keeley (0000-0003-1468-6763)<sup>2</sup>, Daniel I. Shin (0000-0001-5021-4101)<sup>2</sup>, Victor M. Suárez Casanova (0000-0002-9484-4620)<sup>2</sup>, Yannan Zhu (0000-0003-4563-0068)<sup>2</sup>, Lisandro Martin (0000-0001-9597-1843)<sup>2</sup>, Olga Papaemmanouil (0000-0003-4526-3595)<sup>3,5</sup>, Stephen D. Van Hooser (0000-0002-1112-5832)<sup>1-4</sup>

1: Department of Biology

2: Program in Neuroscience

3: Volen Center for Complex Systems

4: Sloan-Swartz Center for Theoretical Neurobiology

5: Michtom School of Computer Science

Brandeis University, Waltham, MA 02454 USA

**Corresponding author:**

Stephen D. Van Hooser, vanhooser@brandeis.edu

Brandeis University, 415 South St. MS008, Waltham, MA 02454 USA

**Number of Pages:** 33

**Number of Figures and Tables:** 12 Figures, 1 Table, 0 Multimedia

**Words:** Abstract: 150 words, Intro: 645, Discussion: 2427, Total: 6728

**Conflicts of interest:** The authors declare no competing financial interests.

**Acknowledgements:** This work was funded by an NIH BRAIN Grant (MH114678). We thank members of the Van Hooser lab and the Brandeis systems neuroscience community for comments. We thank Eve Marder's lab, Alessandra Angelucci's lab, and Don Katz's lab for sharing data for demonstration purposes.

**Contributions:** SDV, ZW, DGM, and OP designed the system, with input from all authors. OP and SDV obtained funding. SDV, ZW, DGM, KZ, YZ, SC, LM wrote the code. ZCK, DS, VMSC, YZ, and SDV analyzed data. SDV, MRK, AH, KZ wrote and edited the online tutorials. SDV wrote the paper with input from all authors.

38 **Keywords:** data acquisition, queries, data archive, BRAIN Initiative, brain science

## 39 **Abstract**

40

41

42 Collaboration in neuroscience is impeded by the difficulty of sharing primary data, results,  
43 and software across labs. Here we introduce Neuroscience Data Interface (NDI), a platform-  
44 independent standard that allows an analyst to use and create software that functions  
45 independently from the format of the raw data or the manner in which the data is organized  
46 into files. The interface is rooted in a simple vocabulary that describes common apparatus  
47 and storage devices used in neuroscience experiments. Results of analyses – and analyses of  
48 analyses – are stored as documents in a scalable, queryable database that stores the  
49 relationships and history among the experiment elements and documents. The interface  
50 allows the development of an application ecosystem where applications can focus on  
51 calculation rather than data format or organization. This tool can be used by individual labs  
52 to exchange and analyze data, and it can serve to curate neuroscience data for searchable  
53 archives.

54

## 55 **Significance Statement**

56

57 Neuroscience experiments generate heterogeneous data, and each lab typically stores its data  
58 and analyses in their own idiosyncratic formats and organizations. We introduce an interface  
59 standard - the Neuroscience Data Interface - that allows the user to specify these formats and

60 organizations so that data and analyses can easily be shared among labs or posted to journals  
61 and archives.

## 62 Introduction

63

64 Despite its importance, collaboration and sharing of data and primary results is very difficult  
65 in the neurosciences, particularly for physiology experiments. At present, physiology  
66 experiments are usually performed on custom experimental rigs that acquire data in unique,  
67 creative, and idiosyncratic ways. Neurophysiology or neuroimaging rigs often employ several  
68 pieces of equipment from different eras of time and with vastly different degrees of  
69 engineering refinement. Each data acquisition (DAQ) device on a rig usually has its own  
70 sampling rate, clock, and means of storing data to disk. On top of this physical heterogeneity  
71 are at least 2 types of digital heterogeneity: the digital format of the data, that typically varies  
72 from device to device, and the organization of data and metadata into files or folders, that  
73 differs greatly from device to device and from lab to lab.

74

75 While the current state of affairs allows for significant creativity on the measurement side of  
76 experiments, it presents substantial challenges for data analysis and its reproducibility. Most  
77 laboratories cannot analyze the data of other laboratories without perhaps a month or more  
78 of effort writing conversion software (Teeters et al., 2008; Garcia et al., 2014; Wiener et al.,  
79 2016; Rübél et al., 2019; Sprenger et al., 2019). This barrier has meant that most labs or  
80 investigators write their own analysis software that they test themselves in only a limited  
81 manner. Further, this barrier impedes the development and utility of common, best-of-breed  
82 analysis packages that are dedicated to analyzing certain classes of data (Wiener et al., 2016).  
83 There are some important efforts to develop file format standards (Teeters et al., 2015; Rübél  
84 et al., 2019) that, if followed, would allow for the development of these packages. However,  
85 these standards typically require users to first convert their data into the common format,  
86 which is itself a barrier to adoption. Heretofore, these packages have been used by relatively  
87 few labs, although this situation is improving. It would be ideal to have a tool that allows an

88 analyst to quickly read and analyze data regardless of whether it is organized  
89 idiosyncratically or stored in standardized container formats.

90

91 Here, we introduce a new approach that allows the development of common analysis tools  
92 without requiring a common file format: a Neuroscience Data Interface (**NDI**). The interface  
93 provides a standard means of specifying and addressing the data that are collected in  
94 neuroscience experiments. At the highest level, the interface provides a vocabulary and  
95 conceptual framework for specifying recordings and analyses. At the implementation level,  
96 the interface contains an extendable set of open source code and interface standards for  
97 reading from a variety of data formats and for specifying the manner in which the  
98 experimental data is organized on disk. The interface is platform- and computing language-  
99 independent. The interface includes a scalable database for storing results of calculations on  
100 the raw data, and user-designed or commercial applications can read and write from the  
101 database in order to build complex, layered analyses. These database entries are specified  
102 using platform-independent metadata that is human- and machine-readable, and database  
103 entries can exist on a user's computer or in the cloud. NDI is designed to serve analysts who  
104 want to be able to quickly read data from a variety of collaborators; if it were widely adopted  
105 by the community, it also has the capability to act as a data curation and archive system for  
106 neuroscience data.

107

108 In this article, we demonstrate the interface in a Matlab prototype. Our purpose here is not  
109 to showcase a completed system that works at scale, but is instead to propose a solution to  
110 the scientific problem about the level of abstraction that is most useful for wide scale  
111 curation and sharing of neuroscience data that allows for the development of common tools.  
112 We view this as an important scientific problem at the boundaries of computer science,  
113 library science, and neuroscience.

114

## 115 **Materials and Methods**

116

### 117 *Design of the interface*

118

119 The neural data interface in its current form was designed and revised over the course of 5  
 120 years. The conceptual framework of the system was developed through discussions with  
 121 Brandeis neuroscience and computer science graduate and undergraduate students. The  
 122 system began from a Lab Information Management System (LIMS) in the Van Hooser lab,  
 123 and was rebuilt twice from scratch to incorporate necessary features and simplify the  
 124 interface and external concepts.

125

126 The interface was prototyped in Matlab (The MathWorks) and is available at  
 127 <https://neurodatainterface.org>. The website provides installation instructions and several  
 128 tutorials that demonstrate how to use NDI. NDI was used extensively to analyze the data of  
 129 Roy et al. (2020), and NDI was revised and debugged as necessary to allow a full pipeline  
 130 analysis. In addition, the process of developing tutorials for user feedback also identified  
 131 unnecessary complexity and bugs that were revised or simplified. Third party libraries such  
 132 as sigTOOL (Lidierth, 2009) (<https://sourceforge.net/projects/sigtool/>) are extensively used to  
 133 read a variety of data formats. Functions in NDI also depend on the VH Lab toolbox  
 134 <http://github.com/VH-Lab/vhlab-toolbox-matlab> and a set of third-party tools:  
 135 <http://github.com/VH-Lab/vhlab-thirdparty-matlab>.

136

137

138 The code for reading data from the Marder, Angelluci, and Katz labs is included in the  
 139 distribution in the `ndi.setups` package.

140

141



142 **Table 1: Key resources table**

143

Reagent type	Designation	Source or reference	Identifiers	Additional information
Software	Matlab	The Math Works, Natick, MA	RRID: <a href="#">SCR_001622</a>	Software language
Software	GitHub	GitHub	RRID: <a href="#">SCR_002630</a>	Software repository
Software	Python3	<a href="http://www.python.org">www.python.org</a>	RRID: <a href="#">SCR_008394</a>	Software language
Software	sigTool	<a href="https://sourceforge.net/projects/sigtool/">https://sourceforge.net/projects/sigtool/</a>	Lidierth, 2009	Open source software product
Software	Neo	<a href="http://neuralensemble.org/neo/">http://neuralensemble.org/neo/</a>	RRID:SCR_000634	Open source software product

144

145

## Results

### *Concepts and vocabulary – probes, subjects, elements, DAQ systems, and epochs*

Before designing a software interface to experiments, we first sought to codify the elements of an experiment using easy concepts and defined terms, in an effort to take inspiration from the graphical user interfaces developed by Xerox PARC and Apple. We define a **probe** to be any instrument that makes a measurement of or produces a stimulus for a **subject**. *Probes* are part of a broader class of experiment items that we term **elements**, which include concrete physical objects like *probes* but also inferred objects that are not observed directly, such as neurons in an extracellular recording experiment, or abstract quantities, such as simulated data, or a model of the information that an animal has about a stimulus at a given time. Each element must have a *subject*, which can be an experimental subject or an inanimate object like a test resistor. We define a **DAQ system** as an instrument or a set of instruments that digitally records the measurements or the stimulus history of a *probe*. These *DAQ systems* record data from *probes* each time the *DAQ systems* are switched into record mode, and we use the term **epoch** to signify each of these recording periods.

The conceptual framework of the interface is applied to a simple experimental situation in **Figure 1**. Here, a *probe* (an extracellular electrode) is used to record activity in the cerebral cortex of a *subject*, a ferret. The *probe* is wired to a *DAQ system* (data acquisition system, DAQ), that is turned on and off 3 times, resulting in 3 *epochs* of sampled probe data that is saved to disk. The probe has been given the name `cortex` and a reference number of 1 in metadata, in this case provided by the user.

In this framework, a large variety of experimental apparatus are considered *probes*. Examples of *probes* that make measurements include a whole cell pipette, a sharp electrode, a single

173 channel extracellular electrode, multichannel electrodes with either known or unknown  
 174 geometries, cameras, 2-photon microscopes, fMRI machines, nose-poke detectors, EMG  
 175 electrodes, and EEG electrodes. Examples of *probes* that provide stimulation are odor ports,  
 176 valve-driven interaural cannulae, food reward dispensers, visual stimulus monitors, audio  
 177 speakers, and stimulating electrodes.

178

179 In an experiment, we also deal with items that we do not observe directly, or abstract items,  
 180 or simulated data. We term all of these items as experiment *elements* (avoiding the term  
 181 “object” to minimize confusion with the software objects in the implementation). An  
 182 example of an inferred *element* is the activity of a neuron derived from an extracellular  
 183 recording. We do not observe the neuron directly, so while we have some certainty that it  
 184 corresponds to a physical entity, this is really an inference, and different analysts may  
 185 disagree as to whether it exists. Another type of quantity that we may wish to use in our  
 186 analysis is a model, such as a calculation of the information that the animal has about a  
 187 stimulus at a given time. Moreover, we may wish to generate artificial data or simulated data  
 188 that will go through the same pipelines as experimental data. Thus, experiment *elements*  
 189 encompass a broad class of items, including *probes*.

190

191 To read the data generated by a *probe*, **NDI** must access data from the data acquisition device  
 192 or devices that recorded the probe, which we term a *DAQ system*. A *DAQ system* can either  
 193 be a single data acquisition system, such as a data acquisition device made by a major  
 194 company, or it can describe the collective recordings of a set of these systems, such as a  
 195 home-brew system that might use a few data acquisition devices at a time. In our own lab,  
 196 our visual stimulation system relies on data from 2 data acquisition systems (our stimulus  
 197 computer and a multifunction data acquisition system that records digital triggers), but  
 198 logically these are treated together as a single *DAQ system* in NDI (**Figure 1**).

199

Each time a *DAQ system* is switched on and off, an *epoch* of data is logged. The *epochs* are numbered (1, 2, etc) and assigned a unique identifier that never changes, so that the *epoch* can be unambiguously referenced even if other *epochs* are added or deleted later. It is also necessary to specify, for each *epoch*, the mapping between any *probes* that are present and the channels of the *DAQ system* that correspond to the *probes*. Commonly, this information must be specified manually using a data type that we have created, but some multifunction data acquisition systems (such as SpikeGadgets MFDAQs) and file formats include this *epoch* metadata in their native file formats, and this metadata can be processed from the files directly.

With a vocabulary to describe the real-world items in an experimental session, we can describe the necessary computational features of the interface (**Figure 2**). While the specification of the *probes*, *subjects*, *elements*, *DAQ systems*, and *epochs* is sufficient to allow the interface to read the data from the *probes* in the experiment, it would be useful to the analyst and his/her collaborators to have a space to store the results of analyses of this data. This space is provided by the **database** (**Figure 2**), which allows the user to store any type of text or binary data related to the experiment in entries called **documents**. For example, one may have a *document* that stores the responses of a neuron to a family of stimuli, and another *document* that stores the results of a model fit of that neuron's responses to the stimulus family. Still another *document* might store the aggregate statistics of the responses to all the neurons in a given study. *Documents* in NDI have a human-readable portion and the option of a binary blob, so that they can be understood easily by humans and programs.

The *interface* with the *database* allows the creation of an **application ecosystem** (**Figure 2**) that can read the raw data and read and write to the database. For example, one common set of early analyses that must be performed by physiologists examining extracellular data is to identify spike waveforms from the raw data and to make an inference as to which spike

227 waveforms arise from the same neuron(s). The NDI *document schema* specifies a *document*  
 228 type that includes common spike detection parameters, including threshold algorithm, filter  
 229 frequencies, the amount of time around each spike to extract, refractory period, etc. These  
 230 parameters can be used by a variety of spike extraction applications, including the example  
 231 “*spikeExtractor*” app shown in **Figure 2** but also other related applications that may be  
 232 developed. There is also a *document schema* for storing extracted spike waveforms and the  
 233 spike times, and another *schema* for spike shape features. These *documents* can be used by  
 234 spike sorting applications, such as the example “*spikeCluster*”, to produce assignments of  
 235 spikes to clusters. One can imagine another application that automatically performs neuron  
 236 assignment from these clusters (“*autoSpikeSort*”), and so on. The *document schemas* are  
 237 flexible and expandable, but must contain certain fields that applications can count on being  
 238 present. In this way, developers and scientists can write applications that perform a  
 239 particular job well, and mix and match their desired applications. The *database* and  
 240 *document schema* allows for powerful collaboration across applications, and allows for a  
 241 healthy competition and interchangeability among applications that perform similar jobs.

242  
 243 The *database* is also designed to allow for the curation and examination of neuroscience data  
 244 and computations at scale. Because each *database document* contains the identifier of the  
 245 experimental session, the *documents* can be combined and searched across the cloud so that  
 246 data and analyses from multiple experiments can be queried, allowing third parties to easily  
 247 perform analyses or meta analyses of a wide variety of experimental data.

248  
 249 The interface is also meant to be used in a similar manner during on-line evaluation of data  
 250 and off-line evaluation of data. The data is addressed in the same manner regardless of  
 251 whether it has been acquired in the last few seconds or a long time ago. This design choice  
 252 has the advantage that all applications can be used on-line or off-line, and removes the

253 necessity of any second “curation” step before making data available to the world on a data  
 254 archive. The data can be curated live, during the experiment.

255

### 256 *Implementation - high level*

257

258 The Neuroscience Data Interface is both an idea, as described above, and an evolving open-  
 259 source software product that implements the concepts. The current software implementation  
 260 of NDI has two layers: a high-level layer of core objects that are described here, and a low-  
 261 level of objects that implement the details of the high-level objects. The separation between  
 262 the high-level and low-level objects has been made so that the external interface of NDI can  
 263 be stable, while the open-source products that implement file reading or the database can be  
 264 switched in and out over time without greatly impacting the user/analyst’s use of the  
 265 interface. The high-level interface is intended as a sort of “neural data operating system” on  
 266 which GUIs and other programs can build, but the core of NDI does not define any particular  
 267 graphical user interface or stipulate the use of any particular underlying database product.

268

269 The goal of this paper is to describe the high-level objects in brief so that the ideas of the  
 270 interface can be discussed or criticized. This paper is not meant to serve as a software  
 271 tutorial. For tutorials on using the software with neuroscience data, please see the repository  
 272 of our current software at <http://github.com/VH-Lab/NDI-matlab>.

273

### 274 *Reading from data acquisition systems: `ndi.daq.system`*

275

276 An `ndi.daq.system` object is a means of addressing and reading the files that are stored  
 277 by the DAQ devices that comprise a *DAQ system*. Different high-level subclasses of  
 278 `ndi.daq.system` allow the user to read from multifunction data acquisition systems (with  
 279 analog and/or digital channels and sampling rates: `ndi.daq.system_mfdaq`), from

280 imaging systems (with image channels and frames: `ndi.daq.system.image`), or from  
 281 stimulus systems (with events and parameters: `ndi.daq.system.stimulus`).

282

283 All `ndi.daq.system` objects rely on 2 key software objects that determine the  
 284 `ndi.daq.system` object's input and output. The first of these is an  
 285 `ndi.file.navigator` object, which allows the user to specify, with a few parameters,  
 286 how the system should search for the files that correspond to each recording *epoch*. **Figure 3**  
 287 shows how different parameters and subclasses of the `ndi.file.navigator` class can be  
 288 used to navigate the different file organization schemas of different labs. Once the files are  
 289 found, another software object, the `ndi.daq.reader`, provides the services for reading  
 290 data from the particular file formats that comprise the epochs.

291

292 *Reading from probes: `ndi.element` and `ndi.probe`*

293

294 When an analyst thinks of a *probe* such as an electrode, he or she might think of the *probe* as  
 295 having the properties of the data acquisition system that records it. For example, we may  
 296 want to talk about the channels of the electrode, and even casually speak of the “sampling  
 297 rate” of an electrode despite the fact that it is the *DAQ system* that directly has a sampling  
 298 rate, not the electrode. The `ndi.element` class, of which `ndi.probe` is a member, allows  
 299 one to address the *probe* or *element* directly, without regard to the *DAQ system* that  
 300 acquired it, which is handled behind the scenes by NDI. In order to define a *probe*, it is  
 301 necessary to functionally define, for each recording *epoch*, a map between the channels of  
 302 the `ndi.daq.system` and the `ndi.probe` object. This can be done manually with the  
 303 class `ndi.epoch.epochprobemap`, or can be specified in the data files directly if the  
 304 *DAQ system* allows it. As shown in **Figure 4**, *probes* can be read by analysis programs  
 305 without any direct concern about the underlying *DAQ systems* that were employed.

306

307 The `ndi.element` class allows many types of data to be treated similarly by software  
 308 programs. For example, all time series in NDI are members of a subclass called  
 309 `ndi.element.timeseries`, which can include artificial (test) data, modeled data,  
 310 filtered data, and so on. In **Figure 5**, the user has created 2 `ndi.element.timeseries`  
 311 objects from a recording from a sharp electrode: 1 of these *elements* represents the  
 312 membrane voltage where the spikes have been removed by a median filter, and the other  
 313 represents the the spiking activity of the cell that is recorded by the sharp electrode. These  
 314 `ndi.element.timeseries` objects can be passed along to an analysis application (here,  
 315 our built-in applications `ndi.app.tuning_response` and `ndi.app.oridirtuning`).  
 316 The epochs of both of these *element* objects are linked back to *epochs* in the *probe*, which  
 317 are in turn linked to the *epochs* of the *DAQ system*, so that time relationships between other  
 318 systems, such as the visual stimulus system, are automatically understood for all of the  
 319 *element* objects derived from *probes*.

320  
 321 *Clocks and time:* `ndi.time.clocktype`, `ndi.time.timereference`,  
 322 `ndi.time.syncgraph`, `ndi.time.syncrule`

323  
 324 One of the biggest challenges in experiments that involve multiple *DAQ systems* is  
 325 synchronizing time across devices that have different clocks. In general, data acquisition  
 326 devices do not share the same clocks: the current time reported by each device will differ  
 327 from others at any given time, and the drift rate of these clocks differs very slightly in a  
 328 matter that may alter the timing of samples in long recordings. Many current data  
 329 standardization schemas sidestep this issue and simply insist that the user must convert all  
 330 times into a standard clock, and NDI is rare in building clocks and synchronization into the  
 331 interface.

332



333 NDI defines several types of clocks (`ndi.time.clocktype`). The most common type of  
 334 clock is “device local time” (`dev_local_time`), which means that a *DAQ system* has a  
 335 local clock that, for each *epoch*, starts a time  $t_0$  and ends at a time  $t_1$ . In most cases,  $t_0$  is 0, and  
 336  $t_1$  is the duration of the recording. Some devices may further keep a “device global” time, so  
 337 that the device has a sub-millisecond record of the relationship between the  $t_0$  of a given  
 338 recording *epoch* and the  $t_0$  of a second recording *epoch* on the same device, but this is  
 339 unusual. We also define the possibility that a device has a record of some “global  
 340 experimental time” or that it keeps “universal controlled time” (UTC).

341

342 As analysts, we’d like to be able to refer unambiguously to a time  $t$  on the clock of a given  
 343 *DAQ system*, and effortlessly know the corresponding time  $t'$  on the clock of another *DAQ*  
 344 system. Therefore, built into every call to the function `readtimeseries`, which reads data  
 345 from a time  $t_i$  to a time  $t_f$  from an `ndi.element`, `ndi.probe`, or `ndi.daqsystm`, is an  
 346 input that specifies the time reference (`ndi.time.timereference`) being used.  
 347 `ndi.time.timereference` objects include the referent (the `ndi.element`,  
 348 `ndi.probe`, or `ndi.daqsystm` being referred to), the clock type, an *epoch* id (if the  
 349 `ndi.clocktype` is `dev_local_time`, which is most common), and an offset time.

350

351 The system is illustrated in **Figure 4**. Here, the user reads samples from a sharp electrode  
 352 *probe* using `readtimeseries`, which returns the *time reference* that was used. Next, the  
 353 user wants to extract stimulus times from the visual stimulus *probe*, which has a different  
 354 clock. The user simply passes the *time reference* object that was returned from the sharp  
 355 electrode *probe* to the `readtimeseries` call to the visual stimulus *probe*, and NDI  
 356 converts the input and output times appropriately so that the output returned is relative to  
 357 the sharp electrode *probe*’s clock.

358

359 The interface solves these conversions from a given clock to another clock by computing  
 360 paths through a directed graph that contains all recorded *epochs* as nodes and the mappings  
 361 between *epochs* as edges. The object that performs this computation is called  
 362 `ndi.time.syncgraph`. The mappings across *epochs* recorded on different *DAQ systems*  
 363 are typically calculated by examining recordings of the same signal (such as a set of digital  
 364 triggers) on both *DAQ systems*. One can also specify rules of synchronization  
 365 (`ndi.time.syncrule`) among devices, and `ndi.time.syncgraph` will automatically  
 366 calculate possible mappings from its set of `ndi.time.syncrule` objects and solve the  
 367 paths through the graph. An `ndi.time.syncrule` might specify the channels of 2 *DAQ*  
 368 *systems* that record digital triggers in common, or might specify that 2 *DAQ systems* have  
 369 the same clock if one of their data files is shared between the 2 systems (such that the same  
 370 *DAQ* hardware is being used in service of 2 *DAQ systems*). Sometimes, if *DAQ* systems were  
 371 not used simultaneously, or if there is no `ndi.time.syncrule`, there is no known  
 372 mapping between different *epochs*. For example, if a *DAQ* system only has a local clock,  
 373 then we usually do not understand the time relationship between subsequent epochs of that  
 374 system (and usually there is no need to understand this relationship). Example cases of  
 375 synchronization relationships are shown in **Figure 6** and **Figure 7**, and a demo of using  
 376 `ndi.time.syncgraph` is shown online in Tutorials 2.1-2.5.

377  
 378 *Database, documents:* `ndi.database` and `ndi.document`

379  
 380 All of the interface that we have described so far is necessary for reading raw  
 381 electrophysiology or imaging files, but does not allow the user to store the results of analysis  
 382 in a convenient and well-documented manner. For this purpose, each experiment is linked to  
 383 a *database* that can, in principle, be running on the local computer or in the cloud. The  
 384 database class `ndi.database` provides standardized methods for adding *documents* to the  
 385 database that conform to a validated, open schema, searching the *database*, and removing

386 *documents* from the *database*. As of this writing, the online version of NDI-matlab offers a  
 387 database using a file system on the local computer, and subclass implementations of  
 388 `ndi.database` that allow cloud access using Postgres and MonogDB are in early testing.

389  
 390 The fundamental unit of the *database* is the *document*, which is implemented by the  
 391 software class `ndi.document`. All *documents* include a core structure of fields that  
 392 describe the unique identifier of the experiment session, the unique identifier of the  
 393 *document*, the time of creation, the schema of the *document*, and a history of how the  
 394 *document* was created so that the calculation can be traced back to the raw data or  
 395 antecedent computations in other *documents*. *Document schemas* are specified in a platform  
 396 independent, human-readable format so they can be read and interpreted on any platform  
 397 and be read and understood by human readers easily. *Document* classes can be composed so  
 398 that one can build *documents* that refer to common elements (such as *epoch* ids or app  
 399 properties) in a consistent manner across *documents* (**Figure 8**). Dependencies among  
 400 *documents* can also be expressed so that relationships among *documents* in a pipeline are  
 401 clear. Finally, each *document* has its own binary stream that can be used to store large binary  
 402 data.

403  
 404 Note that the idea for an extendable, local- or cloud-based database of this type is not new.  
 405 For example, the open-source program DataJoint (Yatsenko et al., 2015) uses a similar design,  
 406 although the underlying data are organized into smaller units called tables rather than  
 407 *documents*. The tables in DataJoint are similar to the substructures of NDI *documents*.

408  
 409 *Analysis pipelines: `ndi.app` and `ndi.query`*

410  
 411 To understand the power of the interface and the potential app ecosystem, it is useful to  
 412 examine a simple analysis pipeline. In this pipeline, we will use a simple spike detection app

413 that is included in the base distribution of NDI called `ndi.app.spikeextractor` to  
 414 detect spikes in sharp electrode data, and then user code to plot the spike shapes.

415

416 The steps of the code that produces the pipeline are illustrated in **Figure 9**, along with the  
 417 *database documents* that are produced at each step. First, the user opens an experiment  
 418 session and identifies the sharp electrode data for each epoch. The data here has been  
 419 normalized by subtraction so that the voltage activity during the preceding interstimulus  
 420 interval (blank screen) is 0. Then, the user creates an instance of the application  
 421 `ndi.app.spikeextractor` (Step 1), builds a *document* that has a set of parameters that  
 422 the app will use in identifying spikes, and adds this *document* to the *database* (Step 2). Next,  
 423 the user calls the app's `extract` method to find and extract the spike data from the *element*;  
 424 the results of the extraction, including spike times and spike shapes for each epoch, are added  
 425 to the *database* as a *document* (Step 3).

426

427 To see what results have been computed, it is necessary to search the *database* for the  
 428 analysis *documents* that currently exist. The *database documents* can be queried with a  
 429 search object called `ndi.query`, which allows the user to perform many types of searches.  
 430 For example, the user can search any text field for several types of matches (exact, partial,  
 431 regular expression match) or search any number field for several types of matches (equal to,  
 432 greater than, less than, etc). The user can also search for *documents* of specific types,  
 433 membership in a particular *session*, and search for *documents* that "depend on" specific other  
 434 *documents*. **Figure 10** shows a short example of the user using `ndi.query` to check for the  
 435 existence of a spike extraction *document* for a particular `ndi.element` object, and then, if  
 436 one is found, plotting the spike waveforms.

437

438 Developing pipelines in NDI becomes a task of writing small programs that read raw data  
 439 and/or existing *database documents*, perform computation, and write results back to the

440 *database* in the form of new *documents*. The *documents* exhibit a beautiful structure when  
 441 plotted as a graph with nodes corresponding to *documents* and edges corresponding to  
 442 dependencies among *documents*. A representative graph from an experimental session in the  
 443 study by Roy et al. (2020) is shown in **Figure 11**. Online tutorials at  
 444 <https://neurodatainterface.org> showcase 4 applications and how to use them with NDI.

#### 447 *Implementation - lower level*

448  
 449 The software product implementation of the interface is currently released in Matlab (see  
 450 **Materials and Methods**). The low-level database implementation is only a slow prototype,  
 451 and is currently being modified to use external SQL databases to allow the system to be used  
 452 at a larger scale. Database documents in the prototype are JSON-based (with a binary blob)  
 453 but will have stricter typing as the external database options come online. The system has  
 454 been used to analyze data for a paper (Roy et al., 2020) and will be tested with data from  
 455 other labs in 2021. The software product is continuously updated on GitHub (see Materials  
 456 and Methods).

#### 458 *Case studies – reading data from many labs*

459  
 460 How easy or difficult is it to read data from other labs in NDI? We present in **Figure 12** an  
 461 example of reading data from 3 laboratories: the Marder Lab at Brandeis (Hamood et al.,  
 462 2015), the Angelucci Lab at the University of Utah (unpublished data), and the Katz Lab at  
 463 Brandeis (Mukherjee et al., 2019).

464  
 465 The Marder lab recorded signals from the stomatogastric ganglion of the crab *Cancer*  
 466 *borealis*. The lab used a common data acquisition system (Spike2 software from Cambridge  
 467 Electronic Design), and the data can be specified by creating an `ndi.daq.system` with the

468 `ndi.daq.reader.mfdaq.cedspike2` reader and describing where the files for  
 469 different epochs are found on disk using an `ndi.file.navigator` object. It requires only  
 470 3 instructions (**Figure 12a**) to create the `ndi.daq.system` once, and this  
 471 `ndi.daq.system` can be used over and over again to access all the data from the  
 472 experiments in the Hamood et al. study (2015) and many current and past experimental  
 473 sessions in the Marder lab.

474  
 475 The Angelucci lab recorded 96-channel data from a Utah array in the marmoset  
 476 (unpublished data courtesy Alessandra Angelucci and Andrew M. Clark). The Angelucci lab  
 477 used a commercial data acquisition system (from Blackrock Microsystems) and, like many  
 478 visual labs, use their own visual stimulus system. The Angelucci stimulus system stores its  
 479 files in Matlab with a time clock that matches the Blackrock Microsystems time clock. For  
 480 this data, we had to follow a template to make a customized stimulus metadata reader (15  
 481 lines of code from a template), and it took 6 instructions to specify the 2 `ndi.daq.system`  
 482 objects needed to access the Utah array data and visual stimulus parameters and timing data  
 483 (**Figure 12B**).

484 The Mukherjee data (2019) included several probes in rat, including dual 32-channel  
 485 electrode arrays that recorded gustatory cortex bilaterally, dual optical fibers that  
 486 ontogenetically manipulated activity in gustatory cortex bilaterally, dual EMG electrodes for  
 487 observing licks and gapes, and intraoral cannulae for delivering tastants directly to the  
 488 tongue. The Katz lab used a commercial Intan Technologies multifunction data acquisition  
 489 system, and the code that specifies the `ndi.daq.system` takes just 6 instructions. Again,  
 490 this `ndi.daq.system` is made once and can be re-used by other members of the Katz lab  
 491 (**Figure 12C**).

492 Thus, an analyst who receives data from another lab, regardless of whether that data  
 493 is packaged in a standard format such as NWB or in custom formats, can gain easy access to  
 494 the data of other researchers and begin analyses the same day using software that follows the

495 NDI conventions, including apps and custom code. Data that is passed on as an  
496 `ndi.session` can be immediately read by other researchers.

497

## 498 Discussion

499

500 We have designed a neuroscience data interface (NDI) that greatly reduces the burden of  
501 analyzing datasets from other labs. The interface allows an analyst to quickly address data  
502 that is acquired in a variety of formats and stored with a variety of organization schemes on  
503 disk. It provides tools for time synchronization across data acquisition systems, and allows  
504 experimental *probes* to be addressed directly by the analyst, while the interface performs the  
505 necessary reading from underlying *DAQ systems*. The interface contains a *database* that  
506 allows experiment objects, analyses, and analyses of analyses to be stored as *documents*,  
507 enabling the development of an application ecosystem that performs analysis independently  
508 of the format or organization of the underlying data. The results of the dataset can be  
509 accessed widely by anyone using the interface, such that the dataset and its analyses are  
510 curated for wide distribution.

511

### 512 *An interface with low barriers for curation and exchange*

513

514 This neuroscience data interface offers several advantages relative to the current  
515 neurophysiological data standardization approaches of which we are aware. 1) NDI is  
516 grounded in concepts and a vocabulary that is easy for non-coders and coders to grasp. 2)  
517 NDI reads data in its native formats, so there are no restrictions for experimental data  
518 collection other than a requirement for using a logically consistent scheme and, once,  
519 locating or writing an open-source reader for each data type. 3) Reading native formats also  
520 offers the significant advantage that the interface can be used regardless of whether the lab  
521 performing the data collection wishes or has the expertise to explicitly convert and curate

522 their own data for analysis by others: an experienced data analyst will be able to quickly  
 523 analyze data using the tools provided by NDI. 4) Reading native formats does not preclude  
 524 the development of excellent file formats, and implementations of NDI can take partial  
 525 advantage of fast code created for existing or future formats. 5) There is a *database document*  
 526 framework so that users and applications can create and abide by *document* templates for  
 527 saved analyses, so that other users and applications can read and interpret the results of  
 528 classes of data analyses in a consistent manner. 6) The *database* is scalable and can exist on a  
 529 user's computer or in the cloud, and data from multiple experiments can easily be combined  
 530 in the cloud to form large, searchable *databases* of neuroscience data and analyses. 7) The  
 531 *database* offers methods for auditing computations and analyses, such that the code and raw  
 532 data that underlie computations and analyses can be fully tracked and reconstructed. Finally,  
 533 like many standardization efforts, we aim for the development of an ecosystem of  
 534 neuroscience analysis apps that will improve reliability, reproducibility, and ease of  
 535 discovery through re-analysis of data by scientists or amateurs.

536  
 537 *Why not simply a file format?*  
 538

539 Why not simply require users to convert their data into a common, standard file format? A  
 540 standard file format provides several advantages. It provides a common target for  
 541 development for device manufacturers and for companies and scientists writing analysis  
 542 software. As the number of channels on some devices become larger, it may be prudent to  
 543 include hardware in analysis, and a common format facilitates this process. Converting to a  
 544 common file format also puts the burden of solving the synchronization of different devices  
 545 outside the scope of the file format, as common file formats such as Neurodata Without  
 546 Borders (Teeters et al., 2015; Rübel et al., 2019) require the user to import data from various  
 547 devices into the format, and the scientist performing data analysis is freed from considering  
 548 these problems.



549

550 However, there are many reasons why, in our opinion, a common file format should not be  
551 the only tool in our toolbox. The first set of arguments against a common file format is  
552 technical in nature. We take it as a given that the most appropriate way to store raw data  
553 from an acquisition device (or simulation) will vary according to the particular  
554 computational and hardware needs of the device, and these needs may evolve in ways that  
555 we cannot imagine at present. For example, the optimal way to compress and store full 3-d  
556 voxel images from a calcium imaging experiment involving a major portion of the macaque  
557 brain (which may be possible in the future) may be very different from those required to  
558 store 3-d voxel images from a  $500\text{ }\mu\text{m} \times 500\text{ }\mu\text{m} \times 10\text{ }\mu\text{m}$  cube. By specifying a common  
559 interface standard but leaving the implementation to vary from *DAQ system* to *DAQ system*,  
560 we gain most of the benefits of a common file format without the liabilities of imposing a  
561 particular storage structure. One may suggest that one could always export the data from a  
562 device's native format to a common file format, but one must remember that a) this is an  
563 extra step for the experimenter, and b) this step could be prohibitively expensive (in time)  
564 for experiments that require somewhat "online" access to neural responses. Having direct  
565 read access via a common reader interface allows the data to be examined "in place" in any  
566 file format. Our own experience waiting an hour to convert a few minutes of 1000-channel  
567 recordings from a prototype acquisition system in order to perform "online" analysis makes  
568 us very enthusiastic about "in place" analysis.

569

570 A second set of arguments against a common file format relates to the ease of workflow for  
571 the scientists. Our goal was to create a system that can be used at the time of data acquisition.  
572 There should be no forced separation between on-line and off-line analysis, so that one can  
573 develop best-of-breed tools for either application that do not depend strongly upon the  
574 platform or devices being used.

575

576 Finally, data curation is clearly a major burden, as there exist file formats that could be used  
 577 for exchange but very few people use them, although this is improving. The requirement of  
 578 an extra step at the conclusion of analysis to “export” the data is a barrier to adoption. In  
 579 NDI, there is no curation step, it is an inherent part of using the data interface.

580

581 An interface can bring on board some of the best benefits of an excellent file format, because  
 582 an interface can read from any file format. As excellent file formats (such as Neurodata  
 583 Without Borders) are developed, interfaces such as NDI can still read them, and these  
 584 formats can be used as a target for future development of hardware and software. The NDI  
 585 approach allows data from these sources to be integrated easily with data from older devices,  
 586 or newer devices that use a different format for whatever reason (technical, creative, or  
 587 historical/idiosyncratic). NDI also allows arbitrary time relationships among epochs to be  
 588 specified and navigated by the interface (local or global), so there are no limits on the data  
 589 that can be easily included and referenced.

590

591 *Stress points: the first DAQ system, `ndi.daq.reader`, `ndi.file.navigator`*

592

593 NDI was designed so that an experienced analyst can specify only a few parameters about the  
 594 file format (`ndi.daq.reader`) and data organization (`ndi.file.navigator`) in order  
 595 to get started (**Figure 3**). For most labs, this will entail a small time investment by a user with  
 596 coding experience to set up the initial *DAQ system* for a lab, or less if the lab uses file formats  
 597 for which `ndi.daq.reader` objects are already available. After this initial setup, a *DAQ*  
 598 *system* definition can be re-used as often as necessary, so a majority of lab users will not need  
 599 this initial expertise.

600

601 *Comparisons and synergies with other efforts*

602

603 This work builds on the experience and expertise of past and current efforts to ease the  
 604 sharing of data in the neurosciences. A scholarly list of efforts to organize and share  
 605 neuroscience data is presented in Table 1 of Teeters et al. (2015), and we won't attempt to  
 606 enumerate a list of all such projects here. Instead, we will draw comparisons with a few  
 607 ongoing efforts.

608  
 609 The idea of an open-source system that can read a variety of file formats is not new. The  
 610 Matlab project sigTOOL (Lidierth, 2009) and the Python-based projects Neo (Garcia et al.,  
 611 2014) and SpikeInterface (Buccino et al., 2020) are already capable of reading a wide variety  
 612 of data formats, and we are using the open source libraries of sigTOOL, Neo, and  
 613 SpikeInterface extensively in our construction of the Matlab- and Python-based versions of  
 614 NDI. On top of reading different file formats, NDI adds the ability to deal with different file  
 615 organizations and explicit management of different timebases on top of managing different  
 616 file formats or collections. That is, in NDI, you specify a rule that describes the arrangements  
 617 of the files without explicitly instructing the software where each file is located. Neo and  
 618 SpikeInterface manage their raw data output in terms of quantities that are similar to NDI's  
 619 epochs.

620  
 621 Neurodata Without Borders (NWB) is an ongoing effort to devise a file format for  
 622 neuroscience data and analyses (Teeters et al., 2015; Rübel et al., 2019). At present, it  
 623 requires users to use or write conversion software to save data into a single file that is  
 624 organized in HDF5 format and that employs a consistent data schema. In NWB, there is no  
 625 equivalent of the NDI *daq system*; instead, users save what NDI calls *probe* and *element* data  
 626 directly to the file. The system also offers spaces to save results of "processing" and "analysis".  
 627 NWB does not allow for multiple time bases, which simplifies the format greatly for the  
 628 analyst, but it means that it is difficult to specify situations where probes or other elements  
 629 have time bases that can be only partially mapped to each other (such as multiple

630 synchronized devices that have only local clocks and no way of mapping to a global time).

631 The format is at present very tied to a file system (1 file per session), although it can be used

632 in conjunction with databases like DataJoint. NWB continues to evolve to broaden its

633 functions and extension capability.

634

635 NWB and many other efforts use an HDF5 file format, which offers some advantages but the

636 notable disadvantages that controlling versions is relatively difficult as is accessing partial

637 datasets in the cloud. Some of these disadvantages can be overcome with approaches like

638 Exdir (Dragly et al., 2018), which offers all of the advantages of HDF5 but without using a

639 single file to store all information.

640

641 Expipe (Lepperod et al., 2020) is another data model that uses the easy object concepts of

642 Projects, Actions, and Entities to organize experimental data. It is a lightweight approach

643 that is highly customizable.

644

645 The *document* space of the NDI *database* has commonalities with the tables in the database

646 DataJoint (Yatsenko et al., 2015). For example, the *document* in **Figure 8** can be built by 5

647 related tables in DataJoint (document classes `ndi_document`, `ndi_epochid`, `ndi_app`,

648 `spikewaves`, `document_class`). Different users may prefer the *table* arrangement of

649 DataJoint or the *documents* of NDI. We designed our *documents* independently of DataJoint

650 and noticed the similarities later. We think that the *document* structure of NDI might be

651 easier for non-programmers to grasp and no more difficult for programmers to query, but the

652 *database* forms share similar forms, including the ability to have dependencies across *table*

653 entries or *documents*. Both DataJoint and NDI lend themselves to the creation of exploration

654 tools that allow users to examine the analyses that have been run and the creation of

655 pipelines – compositions of analyses – that can speed analyses and improve reliability and

656 reproducibility.

657

658 At the other extreme of these approach is a curation-free (or non-curated) database, such as  
 659 that proposed in an article by Cannon and colleagues (Cannon et al., 2002). In such an  
 660 implementation, there is minimal standardization and the data is downloaded from the  
 661 original investigators. While this approach has the advantage of nearly eliminating the  
 662 “curation” step, it does not easily allow an app ecosystem. NDI allows the user to flexibly  
 663 specify the organization and format of their raw data, but it is accessed through a fixed API.

664

665 *Big challenge: A culture of digital annotation*

666

667 Although NDI was designed to tackle the heterogeneity of the digital organization of data,  
 668 our own experience and several colleagues have commented that another barrier to  
 669 analyzing the data of others is the lack of any consistent digital annotation of data (Teeters et  
 670 al., 2008; Grewe et al., 2011; Wiener et al., 2016; Sprenger et al., 2019). Often, the only copy  
 671 of important metadata is written in a physical notebook and is not expressed digitally.  
 672 Hopefully, as investigators see the utility of common analysis tools, the need to have  
 673 consistent digital annotations of data and metadata will become clearer and more ingrained  
 674 in experimental culture.

675

676 *Big challenge: Common database schemas for analyses, analyses of analyses*

677

678 As data interfaces allow more streamlined access to data formats, a new problem arises: how  
 679 do we read analyses or analyses of analyses from other labs? The database’s flexibility in  
 680 creating new schemas and *document* types is a double-edged sword. Imagine that one lab  
 681 develops a set of *database documents* that describes several responses indexes that  
 682 characterize the response of a neuron to a class of stimuli. Now, imagine that another lab  
 683 develops its own set of *database documents* for the same purpose, but gives the fields

684 different names and organizes these indexes into a different *document* set. Someone doing a  
685 meta-analysis of data from the different labs would either have to recompute the index  
686 values from the raw activity of the neurons, or write analysis code that would search the  
687 *database* for the *document* schemas of both labs. For example, users are free to design their  
688 own schemas in DataJoint, NWB, NDI, odML, or NeuroSys (Pittendrigh and Jacobs, 2003;  
689 Grewe et al., 2011; Sobolev et al., 2014; Sprenger et al., 2019), but there is no requirement  
690 that these schemas be similar or be able to exchange with one another.

691

692 Efforts to standardize schemas for certain sub disciplines (such as visual physiologists, or  
693 cellular physiologists) could be quite useful, but will take time (Wiener et al., 2016). In our  
694 opinions, the development of these schemas have the best chance for broad adoption if they  
695 are created independently of software implementation and are not tied to any specific  
696 software product. Each software tool may have its own particular advantages for certain  
697 applications, and it would be very powerful if users could form queries that make sense  
698 across multiple tools. If there were a standard list of metadata for common data types, an  
699 interface or file format or database could say it was “ACME 12345”-compliant (where ACME  
700 is the name of the organization making the standard, and 12345 was the version of the  
701 standard), and users could make common searches across these systems.

702

703 The field of fMRI is several years ahead of the physiology and imaging communities in the  
704 development of these systems (Cox, 1996; Saad et al., 2006; Gorgolewski et al., 2016; Farber,  
705 2017; Gorgolewski et al., 2017; Nichols et al., 2017; Poldrack and Gorgolewski, 2017;  
706 Markiewicz et al., 2021). Some of these approaches have been extended to support human  
707 EEG data in a similar manner (Holdgraf et al., 2019; Pernet et al., 2019).

708

709 *Summary:*

710

711 As experimentalists and theorists in neuroscience enter the era of big data, it is necessary to  
 712 lower barriers of data exchange and to increase access and the ability to search and aggregate  
 713 data across labs and studies. Some labs have already developed pipelines and tools for  
 714 exchange of neurophysiology and imaging data (Teeters et al., 2008; Teeters et al., 2015;  
 715 Yatsenko et al., 2015; R  bel et al., 2019), while the great majority of labs and investigators  
 716 still use custom or idiosyncratic schemas. Data interfaces allow analysts to quickly work with  
 717 both types of data, greatly speeding collaborations that might otherwise be too cumbersome.  
 718 Data interfaces also allow the development of best-of-breed tools that focus on analysis  
 719 rather than being burdened with the format or organization of the underlying digital data.  
 720 As more neuroscientists gravitate towards sharing data, utility and ease of use will be  
 721 important determining factors in adoption and the degree to which users with different  
 722 levels of computer expertise (users, novice programmers, advanced programmers) can do  
 723 science with each system. NDI was designed to address all these considerations through  
 724 conceptual design first, and implementation second, using an interface framework that can  
 725 reach back into the data of the past and into the data of the future.

726

727

728

## 729 **Figure Captions**

730

731 **Figure 1. A vocabulary for neuroscience experiments that forms the basis of the Neuroscience Data Interface (NDI). Top-**  
 732 **left)** An example experiment. A **probe** is any instrument that can make a measurement from or provide stimulation to a  
 733 **subject**. In this case, an electrode with an amplifier is monitoring signals in cerebral cortex of a ferret and the electrode is a  
 734 *probe* and the ferret is a *subject*. A **DAQ system** is an instrument that digitally logs the measurements or stimulus history of  
 735 a *probe*. In this case, a data acquisition system (DAQ) is logging the voltage values produced by the electrode's amplifier and  
 736 storing the results in a file on a computer. An **epoch** is an interval of time during which a *DAQ system* is switched on and  
 737 then off to make a recording. In this case, 3 *epochs* have been sampled. The experiment has additional experiment **elements**.  
 738 One of these *elements* is a filtered version of the electrode data. A second *element* is a neuron, whose existence and spike  
 739 times have been inferred by a spike analysis application and recorded in the experiment. **Bottom)** In NDI, a wide variety of  
 740 experiment items are called *elements*, of which *probes* are a subset. Examples of *probes* include multi-channel extracellular  
 741 electrodes, reward wells, 2-photon microscopes, intrinsic signal imaging systems, intracellular or extracellular single

electrodes, and visual stimulus monitors. Other *elements* include items that are directly linked to *probes*, such as filtered versions of signals, or inferred objects like neurons whose activity are inferred from extracellular recordings or images. Still other *elements* have no physical derivation, such as artificial data or purely simulated data; nevertheless, we want to be able to treat these items identically in analysis pipelines. Finally, *elements* might be the result of complex modeling that depends on many other experiment *elements*, such as an inferred phenomenological model of the amount of information that an animal has about whether a stimulus is a grating. **Top-Right** *DAQ systems* digitally record *probe* measurements or histories of stimulator activity. In NDI, *DAQ systems* are logical entities, which could correspond physically to a single DAQ device made by a particular company (**top**), or a collection of home-brewed devices that operate together to have the behavior of a single DAQ device (**bottom**). In the bottom example, information from an electrode *probe* and digital triggers from a visual stimulation *probe* are acquired on a single DAQ device, but digital information from both systems (in separate files) is needed to fully describe the activity in each *epoch*.

**Figure 2. An overview of the Neuroscience Data Interface (NDI).** **Top-left** The physical experiment from **Figure 1**. A *probe* (electrode) is used to sample data from the visual cortex of a *subject* ferret. A *DAQ system* digitally logs the measurements. 3 *epochs* of data have been recorded by the *DAQ system*. **Top-right** An experiment *session* is contained in a software object that has a link to the **raw data** (red), an internal set of NDI objects that have information about *DAQ systems* and synchronization methods (green), and link to a **database** (dark blue). Upon creation, each **ndi.daq.system** object is provided with an **ndi.file.navigator** object, which is a parameterized set of instructions for locating the raw files or links that contain the data for a given *epoch*. Therefore, the same *ndi.daq.system* can manage data that is organized into *epochs* on disk according to different schemas. Metadata associated with each *epoch*, in a type called **ndi.epoch.epochprobemap**, specifies the *probes* that are present in each recorded *epoch* and indicates the *probe*'s name, a unique reference, and the channel mapping between the *ndi.daq.system* and the *probe*. This data can be added manually by the user or analyst, or can be read from the *epoch* data files if the *ndi.daq.system*'s data format or a Laboratory Information Management System (LIMS) encodes this information. The database stores documents, which are platform-independent representations of analyses, analyses of analyses, and NDI internal objects. **Bottom-right** Applications can use NDI to read raw data and read the results of previous analyses from the *database* and write the results of new analyses back to the *database* as *documents*. The *database* and *documents* therefore support the construction of pipelines that may be linear or integrated. Applications are free to focus on single analysis problems instead of the raw data format or organization of their input.

**Figure 3. *DAQ systems* allow an analyst to read data in a variety of formats and with a variety of file organizations on disk or in the cloud.** All labs begin by initializing the main data management object, an **ndi.session**. **A)** In lab 1, data from an ACME DAQ device (*.acme* files) is organized in a single, flat directory. With a search parameter (the regular expression *\*\acme\>*), an **ndi.file.navigator** object is instructed how to find the data for each epoch. The file for epoch 2 is requested and shown. **B)** In lab 2, data from a home-brewed configuration using an ACME DAQ device that writes *.acme* files and a custom stimulation system that writes *.stim* files are organized in a single DAQ system. In this lab, data from individual epochs are contained in subdirectories. A subclass **ndi.file.navigator.epochdir** is used to restrict



epochs to the contents of subdirectories, and the search parameters indicate that an epoch must have both a `.acme` file and  
`.stim` file to be valid. **C)** Lab 3 uses an integrated file format, such as that from SpikeGadgets. **D)** After setting up the *DAQ*  
*systems*, data for all the labs is read using the same code, which is independent of the file format or the organization on the  
disk or server.

**Figure 4. Probes.** **A)** When *probes* are defined by providing **B)** a mapping between the channels of the *probe* and the  
channels of the *DAQ system*, the data can be read through direct calls, and NDI manages the necessary calls to the *DAQ*  
*systems*. **C)** Code snippet that loads *probe* objects for a visual stimulus system and a sharp electrode, and reads time series  
data from the sharp electrode *probe*. The code returns a time reference for the sharp *probe*'s epoch, and that reference is  
used to request a time series with the corresponding time intervals from the visual stimulus system (even though the  
systems likely do not have the same clocks). **D)** The raw data and stimulus information are plotted together.

**Figure 5. `ndi.element` objects allow different types of data to go through identical analysis pipelines.** **A)** Code that reads  
and **B)** plots time series data from 2 `ndi.element` objects derived from a single sharp electrode probe: voltage membrane  
data where spikes have been “chopped” out with a median filter (top) and thresholded spike data (bottom). **C)** The objects  
can be sent through analysis applications identically and the same type of summary data generated and plotted. **D)**  
Orientation and direction tuning curves for the subthreshold membrane voltage and spiking activity of the same cell. Note  
that filtered data, modeled data, or artificial test data can be sent through the same analysis pipelines with `ndi.element`.

**Figure 6. Epochs and `ndi.time.syncgraph`.** Illustration of an example experiment with 2  
`ndi.daq.system` objects (`elec_mfdaq` and `vis_stim_daq`) that are each connected to a probe  
(`elec_probe` and `vis_stim_probe`, respectively). The *DAQ systems* have their own clocks that are not  
linked to any global time system. 3 epochs have been recorded by each *DAQ system*. The electrode probe has  
been analyzed and an `ndi.element` object (a neuron, `elec_neuron`) has been created from it. The clock  
and time of each of the epochs for the neuron is inherited from its underlying *probe*, which is in turn inherited  
from the underlying *DAQ system*. The 2 *DAQ systems* each record the same set of digital triggers, and  
`ndi.time.syncgraph` has used its list of `ndi.time.syncrule` objects to compute a mapping  
(`ndi.time.timemapping`) between epochs of those *DAQ systems*. Time can be converted between epochs  
that are recorded simultaneously on the 2 *DAQ systems*, but we do not know how the other epochs are related  
to each other, or how any epoch is related to a global time system like universal controlled time (UTC), shown  
below.

811 **Figure 7. Epochs and `ndi.time.syncgraph`.** Illustration of an example experiment similar to that in **Figure**  
812 **6**, except that the `vis_stim_daq DAQ system` also keeps UTC time in addition to its own local clock. Here,  
813 time can be converted among any epoch because there is a mapping between the epochs of `vis_stim_daq`  
814 and UTC, and there are `ndi.time.timemapping` mappings between the DAQ system. The time in any  
815 epoch can be computed according to the clock of any other epoch, by solving the transformations in the  
816 `syncgraph`. The mappings shown are `ndi.time.timemapping` objects built by a) an  
817 `ndi.time.syncrule`, b) inheritance (e.g., a *probe* inherits the *epoch* information of the *DAQ system* that  
818 acquired it); and c) same units (UTC is a global time system).

819  
820 **Figure 8. Illustration of `ndi_documents` and the creation of new classes of `ndi_documents` by composition. Top panel)**  
821 *Document* definitions, with fields. Several *document* classes are created by composition: for example, the *spikewaves*  
822 type has its own fields plus those of document classes `ndi_document`, `ndi_epochid`, and `ndi_app`. **Bottom panel)** A  
823 specific *spikewaves document* from a *database*. The *document* includes a description of the *document* definition, a unique  
824 ID and timestamp, the app that created it, the parameters that were used, a link to the `ndi.element` that was analyzed  
825 and other parameters.

826  
827 **Figure 9. Analysis pipelines build *database documents*. A)** Code snippet that creates an instance of the NDI spike extractor  
828 app (Step 1), creates a *document* that contains the parameters to be used for spike waveform extraction (Step 2), and extracts  
829 the spikes (Step 3). **B)** The *database documents* that are present at each Step. Initially, the experiment has an  
830 `ndi.daq.system`, 2 *probes* (a visual stimulus system and a sharp electrode), and an `ndi.element` that is a normalized  
831 version of the spiking activity. At Step 2, a *document* describing the parameters to be used for spike waveform extraction is  
832 added. At Step 3, a *document* describing the extracted spikes is added.

833  
834 **Figure 10. Accessing analysis results involves querying the database with `ndi_query`. A)** Code that uses a composition of  
835 `ndi.query` objects to look for a document that meets the following criteria: 1) it is of `ndi.document` type  
836 'spike\_extraction'; AND 2) it depends on the `ndi.element` variable named `element_vmcorrected`; and 3) it is from  
837 the session `S`. If it finds such a document, then it calls the spike extractor's method to return the spike waveforms `w` and the  
838 parameters `wp`, and spike times `t`. All spikes that have an inter-spike-interval of 100 milliseconds or greater are plotted, as  
839 shown in panel **B**.

840  
841  
842 **Figure 11. Graph structure of the *database documents* of an example experiment in NDI. A)** Full graph of documents from  
843 an experimental session from Roy et al. (2020). *Documents* are denoted by nodes (blue or green circles), and arrows point  
844 from dependent *documents* to the *documents* that they depend upon. In this graph, **a** is a visual stimulus monitor *probe*, and

845 **b** and **c** are stimulus presentation *documents* that describe the presentation of sinusoidal gratings in different directions. **d**  
 846 and **e** are sharp electrode *probes* corresponding to 2 recordings of different impaled cells. **f** and **g** are *documents* describing  
 847 the `ndi.element` objects of probe **e** where spikes are removed (**f**) and where spike times are extracted (**g**). **h** is a *document*  
 848 containing the stimulus responses of the spikes in **g** to the stimulus presentation in **c**. In **i**, these stimulus responses have  
 849 been collated into a tuning curve. Finally, these responses have been examined to extract orientation and direction index  
 850 values and to perform a double Gaussian fit, which are all stored in *document j*. **B**) Zoomed in view of the *document*  
 851 pipeline **a-j**.

852  
 853  
 854 **Figure 12. With NDI daq readers and a few parameters, one can read many different types of experiments quickly and**  
 855 **directly, without file conversion.** *Subjects* (green boxes), *probes* (blue boxes), and *daq systems* (red boxes) are shown. Wires  
 856 and terminals indicate connections of *probes* to *subjects* and *daq systems*. **A**) Activity of a central pattern generator  
 857 measured in Eve Marder's lab (stomatogastric ganglion (STG) of the crab *Cancer borealis*) (Hamood et al., 2015). Electrodes  
 858 on different nerves indicate the pyloric rhythm that controls the movement of food into the crab's stomach. The 3  
 859 instructions of code needed to specify the *daq system*, modified on a template, are shown at right. Acquisition system was by  
 860 Cambridge Electronic Design. **B**) Unpublished data snippet from Alessandra Angelluci's lab showing responses to visual  
 861 stimulation that were recorded on a 96-channel Utah array implanted in a marmoset. Traces show spikes and numbers, and  
 862 tick marks are visual stimulus identifier numbers. The 6 instructions needed to set up the 2 *daq systems* are shown; another  
 863 15 lines were needed to build a custom stimulus reader (modified from a similar reader). Acquisition system was by  
 864 Blackrock Microsystems. **C**) An experiment by Don Katz's lab (Mukherjee et al., 2019) that explored the relationship  
 865 between activity in gustatory cortex and whether a rat would choose to consume or expel a taste stimulus delivered through  
 866 interoral cannulae. The experiment also included optical fibers to optogenetically inhibit neurons projecting to the gustatory  
 867 cortex from the amygdala. Graph shows EMG recordings (green) indicating licking following sucrose delivery and gaping  
 868 following quinine delivery. Some inputs to gustatory cortex were inhibited just after quinine was delivered. The 6  
 869 instructions needed to express the *daq system* are at right. Acquisition system was by Intan Technologies. This figure shows  
 870 how diverse experiments, with different formats and different file organizations, can be read through NDI by specifying  
 871 only a few parameters. Additional experiments of these types can be read with no new code.

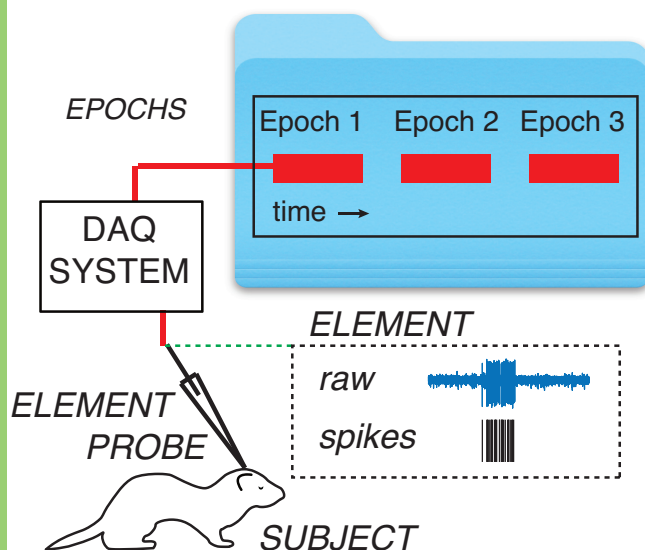
872

## Bibliography


- Buccino AP, Hurwitz CL, Garcia S, Magland J, Siegle JH, Hurwitz R, Hennig MH (2020) SpikeInterface, a unified framework for spike sorting. *eLife* 9.
- Cannon RC, Howell FW, Goddard NH, De Schutter E (2002) Non-curated distributed databases for experimental data and models in neuroscience. *Network* 13:415-428.
- Cox RW (1996) AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res* 29:162-173.
- Dragly SA, Hobbi Mobarhan M, Lepperod ME, Tennoe S, Fyhn M, Hafting T, Malthe-Sorensen A (2018) Experimental Directory Structure (Exdir): An Alternative to HDF5 Without Introducing a New File Format. *Frontiers in neuroinformatics* 12:16.
- Farber GK (2017) Can data repositories help find effective treatments for complex diseases? *Prog Neurobiol* 152:200-212.
- Garcia S, Guarino D, Jailliet F, Jennings T, Propper R, Rautenberg PL, Rodgers CC, Sobolev A, Wachtler T, Yger P, Davison AP (2014) Neo: an object model for handling electrophysiology data in multiple formats. *Frontiers in neuroinformatics* 8:10.
- Gorgolewski KJ et al. (2016) The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Sci Data* 3:160044.
- Gorgolewski KJ et al. (2017) BIDS apps: Improving ease of use, accessibility, and reproducibility of neuroimaging data analysis methods. *PLoS Comput Biol* 13:e1005209.
- Grewe J, Wachtler T, Benda J (2011) A Bottom-up Approach to Data Annotation in Neurophysiology. *Frontiers in neuroinformatics* 5:16.
- Hamood AW, Haddad SA, Otopalik AG, Rosenbaum P, Marder E (2015) Quantitative Reevaluation of the Effects of Short- and Long-Term Removal of Descending Modulatory Inputs on the Pyloric Rhythm of the Crab, *Cancer Borealis*. *eNeuro* 2.
- Holdgraf C et al. (2019) iEEG-BIDS, extending the Brain Imaging Data Structure specification to human intracranial electrophysiology. *Sci Data* 6:102.
- Lepperod ME, Dragly SA, Buccino AP, Mobarhan MH, Malthe-Sorensen A, Hafting T, Fyhn M (2020) Experimental Pipeline (Expipeline): A Lightweight Data Management Platform to Simplify the Steps From Experiment to Data Analysis. *Frontiers in neuroinformatics* 14:30.
- Lidierth M (2009) sigTOOL: A MATLAB-based environment for sharing laboratory-developed software to analyze biological signals. *J Neurosci Methods* 178:188-196.
- Markiewicz CJ, Gorgolewski KJ, Feingold F, Blair R, Halchenko YO, Miller E, Hardcastle N, Wexler J, Esteban O, Goncalves M, Jwa A, Poldrack R (2021) The OpenNeuro resource for sharing of neuroscience data. *eLife* 10.
- Mukherjee N, Wachutka J, Katz DB (2019) Impact of precisely-timed inhibition of gustatory cortex on taste behavior depends on single-trial ensemble dynamics. *eLife* 8.
- Nichols TE, Das S, Eickhoff SB, Evans AC, Glatard T, Hanke M, Kriegeskorte N, Milham MP, Poldrack RA, Poline JB, Proal E, Thirion B, Van Essen DC, White T, Yeo BT (2017) Best practices in data analysis and sharing in neuroimaging using MRI. *Nat Neurosci* 20:299-303.
- Pernet CR, Appelhoff S, Gorgolewski KJ, Flandin G, Phillips C, Delorme A, Oostenveld R (2019) EEG-BIDS, an extension to the brain imaging data structure for electroencephalography. *Sci Data* 6:103.

- 917 Pittendrigh S, Jacobs G (2003) NeuroSys: a semistructured laboratory database. *Neuroinformatics*  
 918 1:167-176.
- 919 Poldrack RA, Gorgolewski KJ (2017) OpenfMRI: Open sharing of task fMRI data. *Neuroimage*  
 920 144:259-261.
- 921 Roy A, Osik JJ, Meschede-Krasa B, Alford W, Leman DP, Van Hooser SD (2020a) Synaptic and  
 922 intrinsic mechanisms underlying development of cortical direction selectivity. *bioRxiv*:776534.
- 923 Roy A, Osik JJ, Meschede-Krasa B, Alford WT, Leman DP, Van Hooser SD (2020b) Synaptic and  
 924 intrinsic mechanisms underlying development of cortical direction selectivity. *eLife* 9.
- 925 Rübel O et al. (2019) NWB:N 2.0: An Accessible Data Standard for Neurophysiology. *bioRxiv*:523035.
- 926 Saad ZS, Chen G, Reynolds RC, Christidis PP, Hammett KR, Bellgowan PS, Cox RW (2006)  
 927 Functional imaging analysis contest (FIAC) analysis according to AFNI and SUMA. *Hum Brain Mapp*  
 928 27:417-424.
- 929 Sobolev A, Stoewer A, Leonhardt A, Rautenberg PL, Kellner CJ, Garbers C, Wachtler T (2014)  
 930 Integrated platform and API for electrophysiological data. *Frontiers in neuroinformatics* 8:32.
- 931 Sprenger J, Zehl L, Pick J, Sonntag M, Grewe J, Wachtler T, Grun S, Denker M (2019) odMLtables: A  
 932 User-Friendly Approach for Managing Metadata of Neurophysiological Experiments. *Frontiers in*  
 933 *neuroinformatics* 13:62.
- 934 Teeters JL, Harris KD, Millman KJ, Olshausen BA, Sommer FT (2008) Data sharing for computational  
 935 neuroscience. *Neuroinformatics* 6:47-55.
- 936 Teeters JL et al. (2015) Neurodata Without Borders: Creating a Common Data Format for  
 937 Neurophysiology. *Neuron* 88:629-634.
- 938 Wiener M, Sommer FT, Ives ZG, Poldrack RA, Litt B (2016) Enabling an Open Data Ecosystem for  
 939 the Neurosciences. *Neuron* 92:617-621.
- 940 Yatsenko D, Reimer J, Ecker AS, Walker EY, Sinz F, Berens P, Hoenselaar A, James Cotton R, Siapas  
 941 AS, Tolia AS (2015) DataJoint: managing big scientific data using MATLAB or Python.  
 942 *bioRxiv*:031658.


## EXPERIMENT SESSION



## DAQ SYSTEM EXAMPLES



## ELEMENT EXAMPLES



filtered data



inferred: neuron spike times



inferred model:  $P(\text{grating})$



simulated: neuron spike times

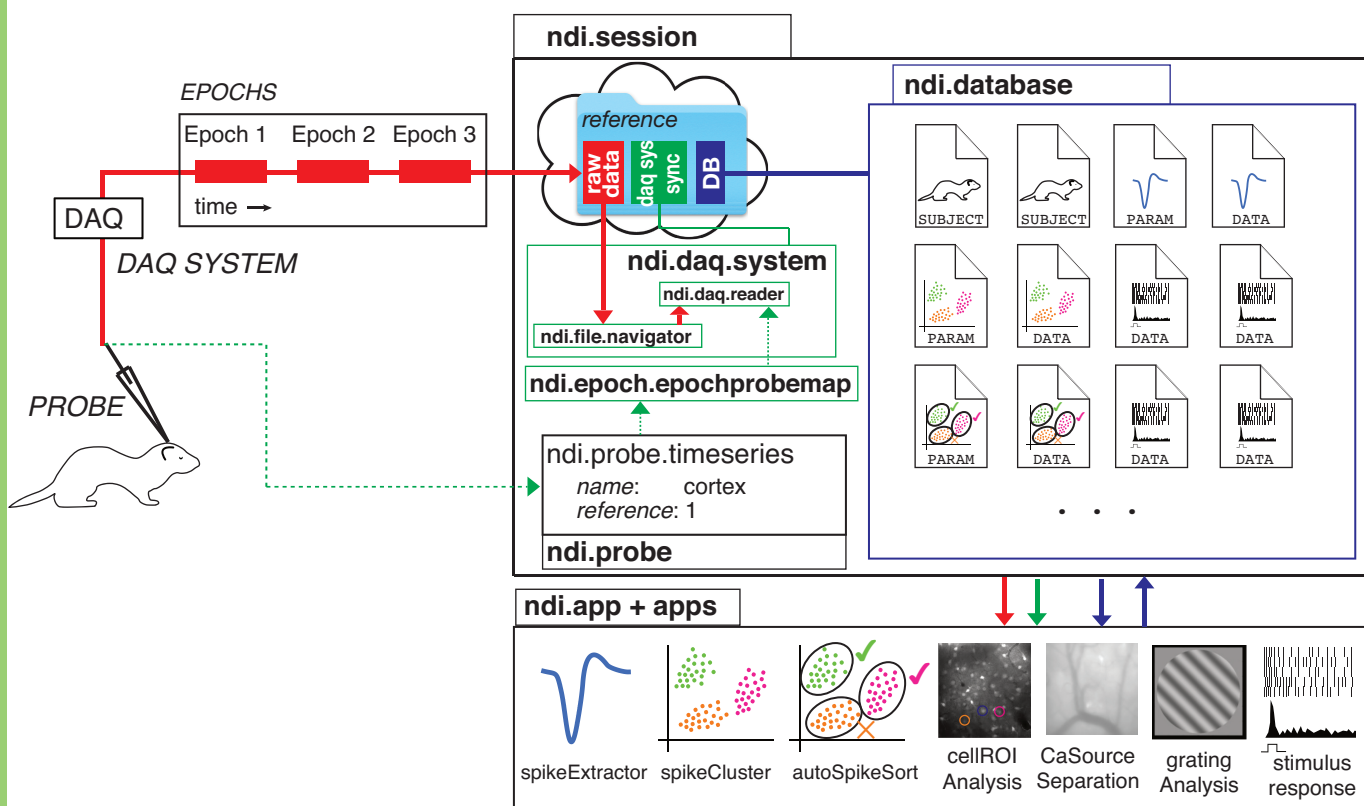


simulated: current injection

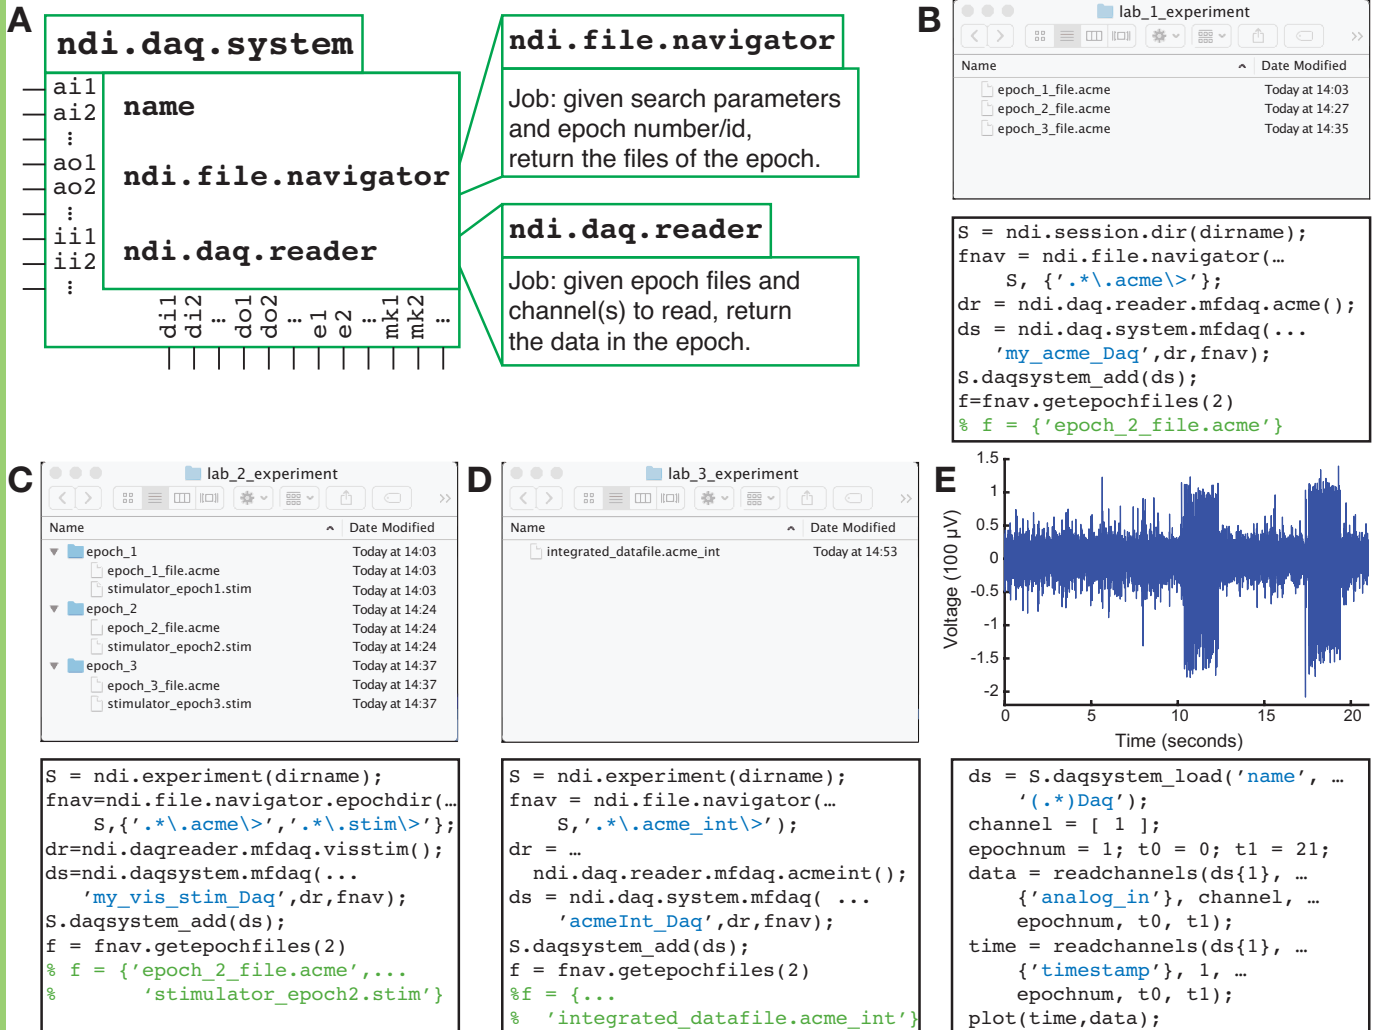


simulated: neuron LIF model

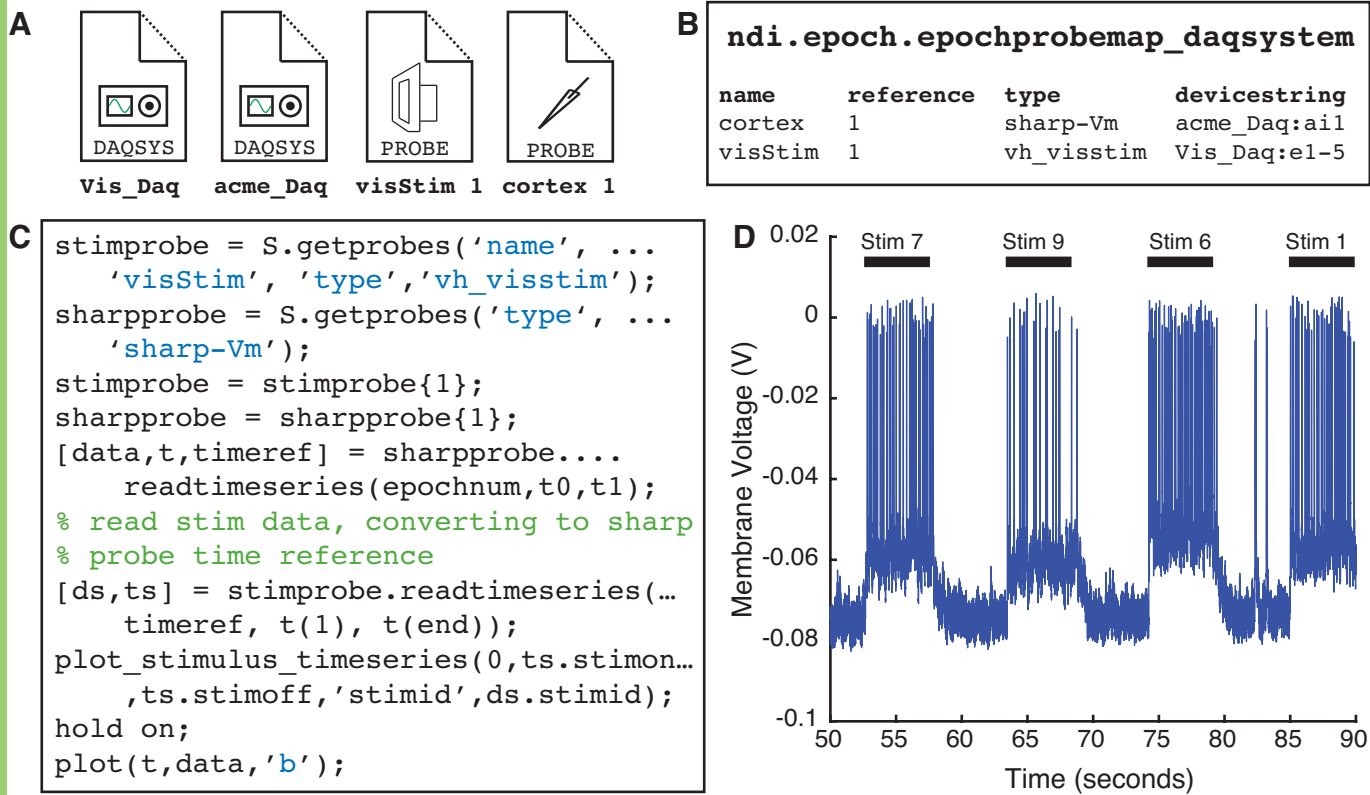












**A**

```


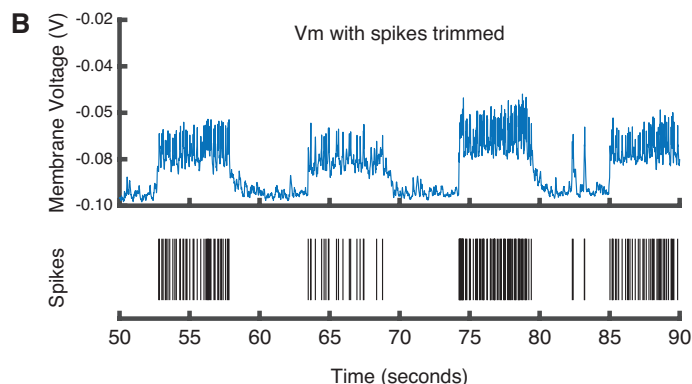
myelement_vm = S.getelements(...
    'element.type', 'Vm_without_spikes');
myelement_spike = S.getelements(...
    'element.type', 'spikes');
[data_vm, t_vm] = myelement_vm{1}...
    .readtimeseries(1, t(1), t(end));
[data_sp, t_sp]=myelement_spike{1}...
    .readtimeseries(1, t(1), t(end));
plot(t_vm,data_vm,'b');
hold on;
spiketimes_plot(t_sp);
    
```

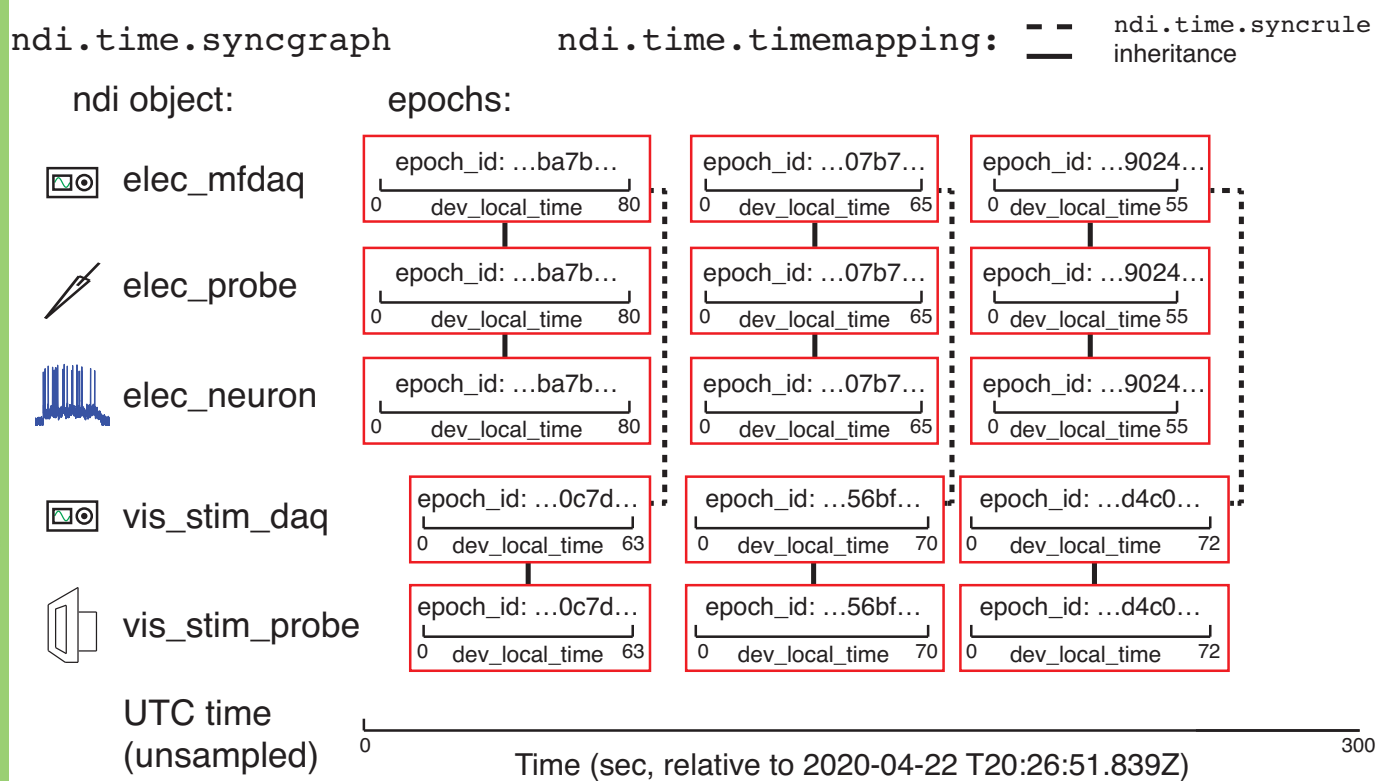
**C**

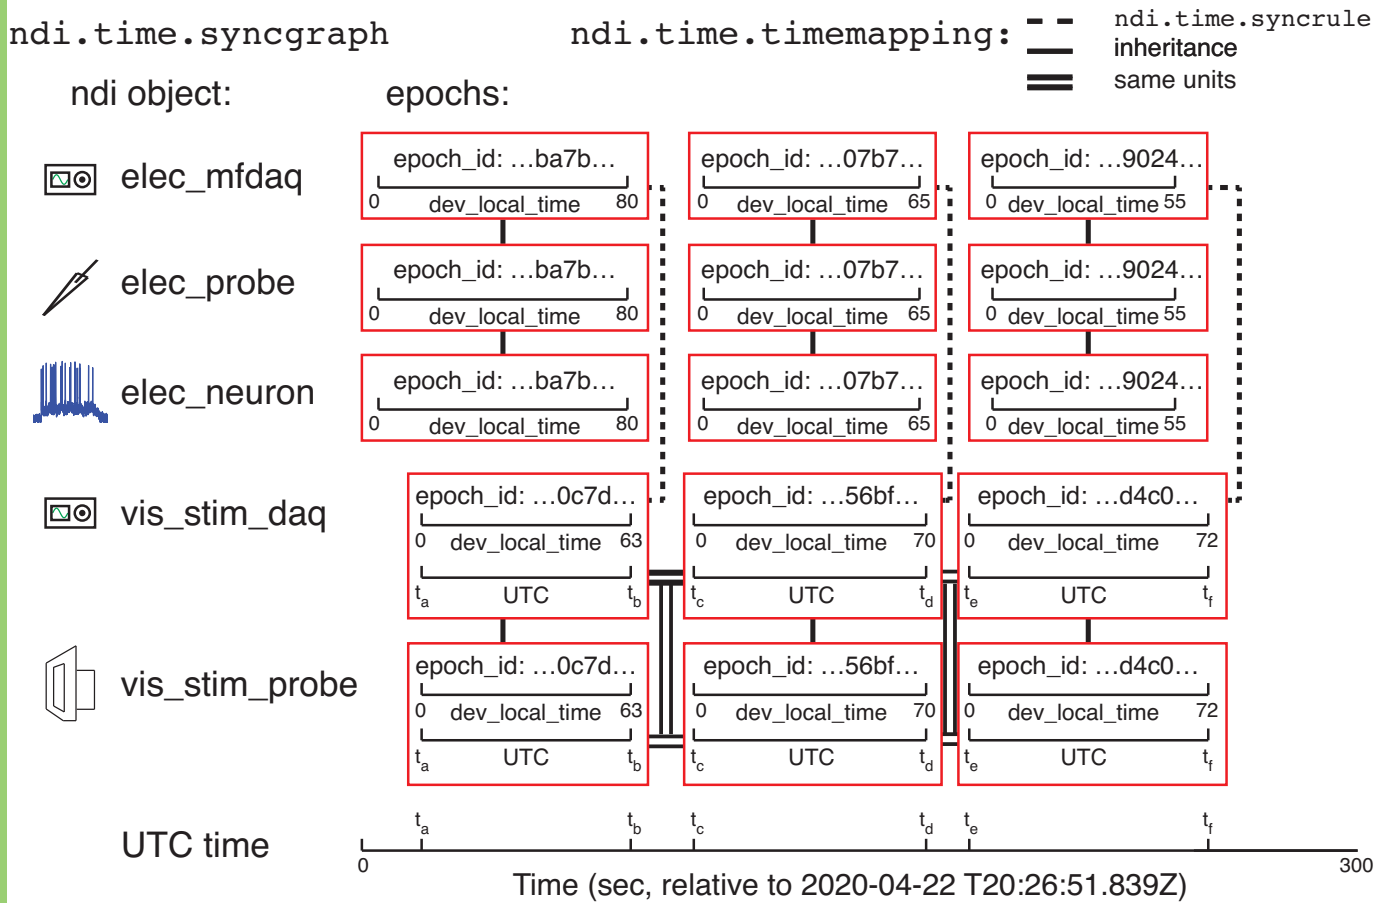
```

% use tuning_response app with things
tapp = ndi.app.tuning_response(S);
oapp = ndi.app.oridirtuning(S);

sp_resp = tapp.find_tuningcurve_document(...
    myelement_spike,1,'mean');
vm_resp = tapp.find_tuningcurve_document(...
    myelement_vm,1,'mean');
oriprop_sp = oapp.calculate_oridir(sp_resp);
oriprop_vm = oapp.calculate_oridir(vm_resp);
oapp.plot_oridir_response(sp_resp);
oapp.plot_oridir_response(vm_resp);
    
```







**A**

 **ndi\_document**, with fields:

```
"id": % unique identifier
"experiment_id": % unique identifier
"name": % name field
"type": % type field
"datestamp": %utc date stamp
"database_version": % version
```

 **spikewaves** = **ndi\_document** + **ndi\_app** + **ndi\_epochid** + fields:

```
"depends_on": [
    "extraction_parameters_id": %paramdoc
    "element_id": % element id number
]
"sample_rate": % sample rate of epoch
"s0": % time of first sample (peak:=0)
"s1": % time of last sample (peak:=0)
+ binary data
```

 **ndi\_epochid** = **ndi\_document** + fields:

```
"epochid": % unique epoch id
```

 **ndi\_app** = **ndi\_document** + fields:

```
"name": % name of app
"version": % version of app
"OS": % operating system
"OS_version": % OS version
"interpreter": % Matlab, Python3, etc
"interpreter_version": % version
```

---

**B**

 **mydoc** = S.database\_search(ndi.query('','isa','spikewaves',''));

```
mydoc{1}.document_properties.document_class:
    definition: '$NDIDOCUMENTPATH/apps/spikeextractor/spikewaves.json'
    validation: '$NDISCHEMAPATH/apps/spikeextractor/spikewaves_schema.json'
    class_name: 'ndi_document_apps_spikeextractor_spikeextractor_spikewaves'
    class_version: 1
    superclasses(1).definition: '$NDIDOCUMENTPATH/ndi_document.json'
    superclasses(2).definition: '$NDIDOCUMENTPATH/ndi_document_app.json'
    superclasses(3).definition: '$NDIDOCUMENTPATH/ndi_document_epochid.json'
mydoc{1}.document_properties.ndi_document:
    id: '41268449b95781fc_3fe0bf23a68a90a2'
    experiment_id: '2014-05-09_412684472cf40177_3feddc959c9bd904'
    name: 'manually_selected_412684472cf75018_3fe5dc9aac1a7ef0.t00012'
    type: ''
    datestamp: '2020-02-08T01:41:16.434Z'
    database_version: 1
mydoc{1}.document_properties.depends_on(1):
    name: 'extraction_parameters_id' % parameters document
    value: '41268449b92c5644_3fe16609b1bfa8f8'
mydoc{1}.document_properties.depends_on(2):
    name: 'element_id' % ndi_element that is being extracted
    value: '4126844732658ffe_3fe647147e14e1ff'
mydoc{1}.document_properties.ndi_app:
    name: 'ndi_app_spikeextractor' % our included simple spike extractor
    version: '768849c6e5a4e4b8bdfa2aef065d135222e4a93f' % git commit
    OS: 'MacOS'
    OS_version: '10.14.6 Build: 18G4032'
    Interpreter: 'Matlab'
    Interpreter_version: '9.6.0.1174912 (R2019a) Update 5'
mydoc{1}.document_properties.epochid:
    epochid: 't00012'
mydoc{1}.document_properties.spikewaves:
    sample_rate: 11111 % sample rate, Hz
    s0: -0.004 % 4ms before peak
    s1: 0.004 % 4ms after peak
```

**A Code**

```

% Job: Given a probe sharpprobe, epochid eid, and
%      threshold T, extract spikes

% Step 1: set up and load objects
% make an instance of our spike extractor
sapp = ndi.app.spikeextractor(S);
% load our normalized Vm trace
element_vmcorrected = S.getelements(...
    'element.type','Vm_corrected',...
    'element.reference',sharpprobe.reference);

% Step 2: make a spike extractor parameter
% document
extract_doc = ndi.document( ...
    'spike_extraction_parameters');
se_parameters = extract_doc.document_properties...
    .spike_extraction_parameters;

se_parameters.dofilter = 0;
se_parameters.threshold_method = 'absolute';
se_parameters.threshold_parameter = T;
se_parameters.threshold_sign = 1;
se_parameters.spike_start_time = -0.004;
se_parameters.spike_end_time = 0.004;
se_parameters.center_range_time = 0.0015;
se_parameters.read_time = 1000; % long time is faster


extract_p_name = ['manually_selected ' ...
    sharpprobe.id() '.' eid];
sapp.add_extraction_doc(extract_p_name,se_parameters);

% Step 3: do the extraction
sapp.extract(element_vmcorrected,eid,extract_p_name,1);


```

**B Database:**


At Step 1:



After Step 2:



After Step 3:



**A**

```
% prepare search queries
q_e = ndi.query(S.searchquery());
q_t = ndi.query('', 'depends_on', ...
    'element_id', ...
    element_vmcorrected .id());
q_sw = ndi.query('', 'isa', ...
    'spike_extraction', '');
% is there a document that matches
% all of these criteria?
doc = S.database_search(q_e & ...
    q_t & q_sw);

% if so, load and plot ISIs > 100ms
if ~isempty(doc),
    [w,wp]=sapp.load_appdoc('spikewaves',...
        element_vmcorrected,1,'manual');
    t = sapp.load_spiketimes_epoch(...
        element_vmcorrected,1,'manual');
    z = squeeze(w);
    indexes = 1+find(diff(t)>0.100);
    plot([wp.S0:wp.S1]/wp.sample_rate, ...
        z(:,indexes));
end;
```

