Spyglass is a framework for reproducible and shareable neuroscience research

Kyu Hyun Lee[1,2,3†], Eric L Denovellis[1,2,3†], Ryan Ly[4], Jeremy Magland[5], Jeff Soules[5], Alison E Comrie[1,3], Daniel P Gramling[6], Jennifer A Guidera[1,3,7,8], Rhino Nevers[1,3], Philip Adenekan[1,3], Chris Brozdowski[1,3], Samuel R Bray[1,3], Emily Monroe[1], Ji Hyun Bak[1], Michael E Coulter[1,3], Xulu Sun[1,2,3], Emrey Broyles[1,3], Donghoon Shin[1,3,7], Sharon Chiang[9], Cristofer Holobetz[10], Andrew Tritt[4], Oliver Rübel[4], Thinh Nguyen[11], Dimitri Yatsenko[11], Joshua Chu[12], Caleb Kemere[12], Samuel Garcia[13], Alessio Paolo Buccino[14], Loren Frank[1,2,3]*

[1]Department of Physiology, University of California, San Francisco, San Francisco, United States; [2]Howard Hughes Medical Institute, University of California, San Francisco, San Francisco, United States; [3]Kavli Institute for Fundamental Neuroscience, University of California, San Francisco, San Francisco, United States; [4]Scientific Data Division, Lawrence Berkeley National Laboratory, Berkeley, United States; [5]Center for Computational Mathematics, Flatiron Institute, New York, United States; [6]Graduate Program in Neural and Behavioral Sciences, University of Tübingen, Tübingen, Germany; [7]UCSF-UC Berkeley Graduate Program in Bioengineering, University of California, San Francisco, San Francisco, United States; [8]Medical Scientist Training Program, University of California, San Francisco, San Francisco, United States; [9]Department of Neurology, University of California, San Francisco, San Francisco, United States; [10]Sainsbury Wellcome Centre, University College London, London, United Kingdom; [11]DataJoint, Houston, United States; [12]Department of Electrical and Computer Engineering, Rice University, Houston, United States; [13]Centre de Recherche en Neuroscience de Lyon, CNRS, Bron, France; [14]Allen Institute for Brain Science, Seattle, United States

***For correspondence:** Loren.Frank@ucsf.edu

[†]These authors contributed equally to this work

Competing interest: See page 19

Funding: See page 19

Preprint posted
16 June 2025
Sent for Review
26 June 2025
Reviewed preprint v1
26 September 2025
Reviewed preprint v2
01 May 2026
Version of Record
03 November 2026

Reviewing Editor: Adrien Peyrache, McGill University, Canada

© Copyright Lee, Denovellis *et al*. This article is distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use and redistribution provided that the original author and source are credited.

eLife Assessment

This **valuable** study presents a framework for a shareable data analysis pipeline aimed at improving reproducibility in neuroscience. The evidence for robustness and inter-laboratory operability is **convincing**. Overall, this work will be of interest to neuroscientists engaged in the analysis of large-scale neuronal recordings.

Abstract Scientific progress depends on reliable and reproducible results. Progress can be accelerated when data are shared and re-analyzed to address new questions. Current approaches to storing and analyzing neural data involve bespoke formats and software that make replication and reuse of data difficult. To address these challenges, we created Spyglass, an open-source data management and analysis framework written in Python. Spyglass provides reproducible pipelines for common neuroscience analyses and sharing of raw data, intermediate analyses, and final results within and across labs. Spyglass uses the Neurodata Without Borders (NWB) standard and includes pipelines for spectral filtering, spike sorting, pose tracking, and neural decoding. Spyglass can be extended to apply existing and newly developed pipelines to datasets from multiple sources. We demonstrate these features in the context of a cross-laboratory replication by applying advanced

state space decoding algorithms to publicly available data. New users can try out Spyglass on a Jupyter Hub hosted by HHMI and 2i2c: https://spyglass.hhmi.2i2c.cloud/.

Introduction

A central goal of neuroscience is to understand how the structure and dynamics of neural activity relate to the internal states of the organism and the external world. This understanding is derived from analyzing complex, multi-modal datasets. While the community has improved tools and algorithms for data collection and analysis (***Lopes et al., 2015***; ***Nath et al., 2019***; ***Pachitariu et al., 2016***; ***Siegle et al., 2017***; ***Yatsenko et al., 2018***; ***Buccino et al., 2019***), extracting consistent and reproducible insights from data remains a complex and time-consuming task. Often, researchers take years to collect and organize data, which is then transformed through a complicated series of analyses using custom scripts. This begins with preprocessing that isolates specific signals from the data, followed by multiple subsequent analyses that quantify properties of these signals. The outputs of these analyses are then synthesized across datasets, and when they are consistent upon limited replication, they are reported in the scientific literature—with the data and analysis scripts documented to varying degrees.

Ideally, it would be possible for another group to take the same raw datasets, apply the analyses, and rapidly and reliably reproduce the findings. In practice, this is often exceptionally challenging. Raw data are seldom shared and metadata critical for understanding the data are often not included, posing a significant challenge to replication. Essential components of the analysis pipelines, such as the manual curation of clustered spikes and artifact rejection, are often irretrievable from the written reports. Similarly, the full set of parameters used for each of the analyses is not shared or is hidden in cryptic analysis scripts. Efforts to reproduce findings are also hampered by idiosyncratic data and code organization, poor documentation, and missing vital details, including computational hardware requirements (***Abe et al., 2022***). In collaborations among multiple scientists, these problems are exacerbated due to the variability in how each participant carries out analysis. Consequently, the full validation of a result usually requires repeating the experiment and reconstructing the analysis from scratch.

These difficulties in replication of analyses incur significant costs in time and effort. A new trainee might struggle to analyze existing data because they do not understand critical details. A scientist who downloads the data from a previous study may find that the analyses they wanted to carry out are impossible because the raw data is not available. Alternatively, raw data may be available, but the scientist may need intermediate results (e.g., spike waveforms) that are not included. Similarly, shared code, including visualizations, is most often not standardized or documented, causing multiple teams to duplicate efforts and implement the same tools.

A system that addresses these challenges therefore should enable:

- compilation of raw data with sufficient metadata for analysis and reuse
- sharing of data and all intermediate analysis results in an accessible format
- reproducible analysis via well-documented, organized, and searchable pipelines
- generation of shareable visualizations to facilitate communication and collaboration
- easy use by scientists with minimal formal training in data management.

Achieving these goals would represent a major step towards meeting the FAIR guiding principles for *f*indable, *a*ccessible, *i*nteroperable, and *r*eusable (***Wilkinson et al., 2016***) data and analysis pipelines (***Goble et al., 2020***). For example, it would become possible to easily find publicly available data, analyze it with a standardized pipeline that keeps track of all the parameters, and generate a visualization to share the results over the web—a stark contrast to how science is practiced today.

In pursuit of this vision, many organizations, such as the Allen Institute for Brain Science (AIBS), Johns Hopkins Applied Physics Lab (APL), and the International Brain Laboratory (IBL), have made strides by standardizing and sharing data and analyses (***de Vries et al., 2023***; ***Hider et al., 2022***; ***Acerbi et al., 2023***). However, these efforts have not fully resolved the issues related to data sharing and reproducible analysis. For instance, a lack of raw data often precludes reproduction of early stages of the analysis, and many steps of the processing pipeline (e.g., the criteria used for manual or automatic curation of spike sorting) that can significantly affect the results (***Magland et al., 2020***) are omitted. Some of the technologies used, such as cloud computing services (e.g., CodeOcean) and

sophisticated databases and APIs (***Hider et al., 2022***; ***Matelsky et al., 2021***; ***Sanchez et al., 2022***), can be cost prohibitive or require specialized software engineering expertise that is beyond the reach of most labs. Furthermore, these existing efforts tend to be focused on the needs of specific projects, data types, and behavioral paradigms, limiting their scope. Thus, while these efforts mark important advances, there remains a need for user-friendly, integrated solutions that can be widely adopted across individual labs in the neuroscience community.

To address this need, we developed Spyglass, an open-source neuroscience data management and analysis framework written in Python. Spyglass builds on widely available community-developed tools and adopts the Neurodata Without Borders (NWB) as the standardized format (***Rübel et al., 2022***; ***Teeters et al., 2015***). It uses DataJoint (***Yatsenko et al., 2018***; ***Yatsenko et al., 2015***) to manage reproducible analysis pipelines with a relational database and incorporates novel software tools (e.g., Kachery and Figurl) for sharing data and web-based visualizations to enable collaboration within and across labs. This includes methods for exporting and uploading all raw data and intermediate results used to produce a manuscript, which, along with sharing of code, enables full replication of results. Spyglass is Python-based and thus can accommodate pipelines that use a wide array of analysis packages that have been developed by the community, including SpikeInterface (***Buccino et al., 2019***), GhostiPy (***Chu and Kemere, 2021***), DeepLabCut (***Nath et al., 2019***), and Pynapple (***Viejo et al., 2023***). Spyglass also offers ready-to-use pipelines for analyzing behavior and electrophysiological data, including spectral analysis of local field potential (LFP), spike sorting, video processing to extract position, and decoding neural data. Spyglass can be extended to support additional pipelines for behavioral, intracellular, optical physiology data, or other data types that can be stored in the NWB format. In addition to extensive documentation and tutorials, new users can try out a demo version of Spyglass hosted on the web by HHMI and 2i2c as a Jupyter Hub instance. Here, we describe the structure of Spyglass and demonstrate its potential by applying the same analysis pipelines to NWB files from different labs and comparing the results.

Results

Overview of Spyglass

Spyglass is an open-source Python-based software framework for reproducible analysis of neuroscience data and sharing of the results with collaborators and the broader community (***Figure 1***). It is designed to be used by everyone in a laboratory who works with the data, both as a general-purpose

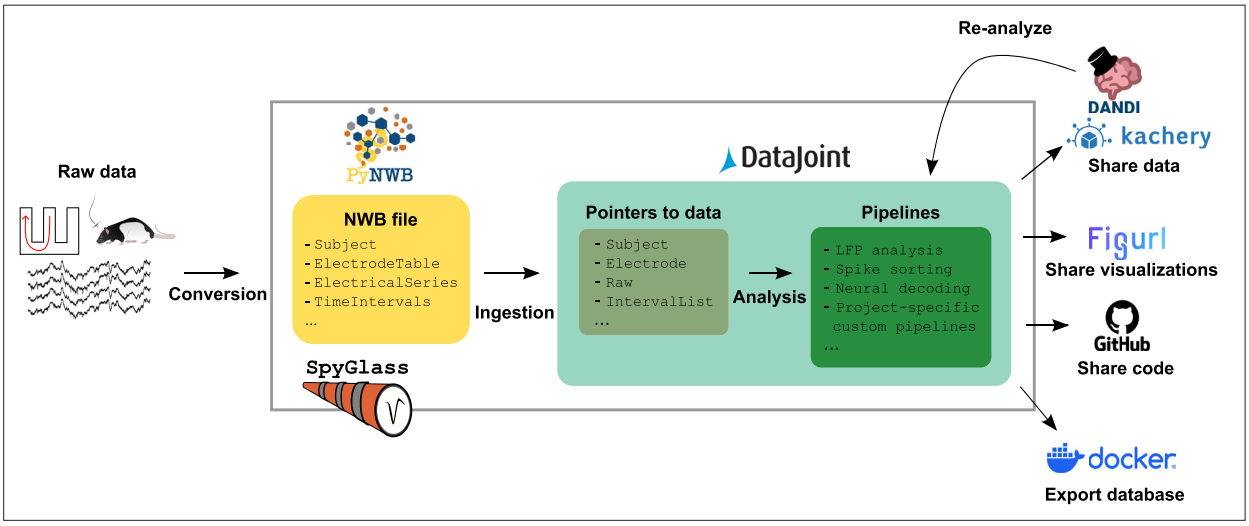


Figure 1. Overview of Spyglass. The raw data—consisting of information about the animal, the behavioral task, the neurophysiological data, etc.—is converted to the NWB format (yellow box) and ingested into the Spyglass database. The pipelines (dark green box) operate on pointers to specific data objects in the NWB file (tan box). The raw and processed data are then shared with the community by depositing them to public archives like DANDI or shared with collaborators via Kachery. Visualizations of key analysis steps can be shared over the web via Figurl. Code is shared by hosting the codebase for Spyglass and project-specific pipelines on online repositories like GitHub. Finally, the populated database may be shared by exporting it to a Docker container.

tool to enable the development of new analysis pipelines and a tool that allows those pipelines and associated results to be frozen and packaged to enable reproducibility. It can be run locally or in the cloud. Analyzing data with Spyglass begins with raw data and experimental metadata stored in the NWB format (***Rübel et al., 2022***; ***BRAIN Initiative, 2019***). These NWB files are ingested into a relational database and processed using DataJoint-enabled pipelines. Existing pipelines are built around common neuroscience tasks such as spectral filtering, spike sorting, pose tracking, and neural decoding, and each user can extend these pipelines to carry out the specific sets of analyses needed for their project. DataJoint stores parameters of each analysis and tracks the intermediate analysis results, which are also stored as NWB files to maintain a shareable standardized data format. Spyglass enables sharing results and interactive visualization of the data over the web via Kachery and Figurl. Finally, Spyglass supports exporting specific parts of the database required to reproduce the results and figures of a manuscript and the upload of the associated raw data and analysis outputs to a public repository. In the following sections, we provide detailed descriptions of these components and the design decisions behind them.

The NWB format as data specification

Why NWB?

A typical neuroscience experiment consists of multiple data streams stored in different formats. Managing such heterogeneous data in a shareable and accessible manner is challenging. A practical solution is to save the data in a community-supported format like NWB, which is emerging as a standard for neurophysiology and behavior data (***Rübel et al., 2022***; ***BRAIN Initiative, 2019***). We have chosen NWB as the data specification in Spyglass for the following reasons:

- The versatility of NWB accommodates various data types and allows metadata to be saved with the data in a single self-annotated file.
- NWB files are immediately shareable.
- Public data archives like DANDI (***Rübel et al., 2019***; ***Callaway et al., 2021***; ***Hawrylycz et al., 2023***; ***Park et al., 2024***) accept the NWB format and provide APIs to easily stream file contents for local analysis.
- Tools developed for NWB files are immediately accessible to users.

Conversion to NWB can be done using software tools developed by the community, such as the NeuroConv package (***Baker et al., 2026***) or NWB GUIDE, a desktop app for converting data to NWB without having to write code.

Importantly, Spyglass requires all raw data—including neurophysiology, behavioral task, interaction with the environment—to be in the NWB format *prior to any analysis*. This ensures reproducibility of all subsequent analyses by sharing the NWB file containing the raw data and the analysis pipelines. Furthermore, Spyglass stores virtually all intermediate results from downstream analysis pipelines in NWB. This ensures that all data associated with the analysis can be shared and read using the same software tools.

Spyglass-specific NWB requirements

NWB allows some flexibility in the specification of data to accommodate a broad range of experiments and lab-specific requirements. For example, the names of data types within the NWB file can differ from those expected by Spyglass. We have fully described the Spyglass-specific NWB conventions on the documentation website. To further accommodate NWB files from many sources, we have also developed a system that makes it possible to ingest NWB files into Spyglass even when they do not adhere to our naming conventions or best practices by including a configuration YAML file (see Materials and methods and ***Table 1***, 02_Insert_Data).

Relational database as analysis pipelines

Why a relational database?

One significant challenge with data analysis is in managing its complexity. Most results derive from an extended series of steps, including 'preprocessing' (e.g., spike sorting for electrophysiological data, region-of-interest identification for optical physiological data, video processing for behavioral data, etc.) and downstream analyses. Each step depends on a different algorithm with a specific set of

Table 1. Tutorials included in Spyglass and their descriptions.
All available from https://github.com/LorenFrankLab/spyglass, ***Broz et al., 2026***.

Pipeline	Tutorial notebook	Description
Data ingestion	00_Setup	Setting up Spyglass
	01_Concepts	Basic concepts of using Spyglass
	02_Insert_Data	How to insert data into Spyglass
	03_Data_Sync	How to share data with collaborators who have access to the database
	04_Merge_Tables	A new table type unique to Spyglass that allows the user to use different versions of pipelines on the same data
	05_Export	How to export parts of the database used to generate figures and upload corresponding NWB files to DANDI
Spike sorting	10_Spike_SortingV0	Detect spikes from electrophysiological recording and assign them to individual neurons (example of multiple versions of the same pipeline)
	10_Spike_SortingV1	
	11_Spike_Sorting_Analysis	Curate the results of spike sorting manually for V0
	12_Burst_Merge_Curation	Automatically merge units separated due to bursting
Position processing	20_Position_Trodes.ipynb	Process information about animal's position from video recording of the behavior using Trodes
	21_DLC	Detect keypoint markers with DeepLabCut
	22_DLC_Loop	Detect keypoint markers with DeepLabCut over multiple epochs
	23_Linearization	Convert 2D position to 1D position using track geometry
LFP analysis	30_LFP	Filter broadband electrophysiology data to isolate low-frequency LFP bands
	31_Theta	Filter LFP to isolate the theta band
	32_Ripple_Detection	Detect sharp-wave ripples from filtered LFP
Decoding	40_Extracting_Clusterless_Waveform_Features	Extract waveform features for clusterless decoding
	41_Decoding_Clusterless	Apply the decoding algorithm using clusterless waveform features
	42_Decoding_SortedSpikes	Apply the decoding algorithm from spikes of sorted and curated units
MUA	50_MUA_Detection	Detect times of high multi-unit firing
MoSeq	60_MoSeq	Tutorial on using the MoSeq pipeline

parameters and generates distinct intermediate data. Tracking these numerous components is difficult, and understanding how another scientist has managed them can be even more daunting. This complexity hinders collaboration, verification of results, and data reuse.

These issues motivated our use of a formal software system: the relational database, a well-established data structure that uses tables to organize data. To construct an analysis pipeline, we use DataJoint (***Yatsenko et al., 2018***; ***Yatsenko et al., 2015***) to define a series of database tables with a dependency structure. Associated with the database tables is code that carries out an analysis using a specific set of parameters and a specific part of the data. This code then stores the results as a new row in the table with a pointer to the results stored on disk as an NWB file. Creating a new table row is referred to as 'populating' the table. Thus, data analysis becomes a matter of populating and interacting with the database. This style of data analysis offers many advantages:

- It lowers effort for users seeking to apply the same analysis to multiple datasets, as they only need to specify the data and parameters for computation ('what') independent of the execution details ('how').
- It provides a structure to organize and systematize the analysis parameters, data, and outputs into different tables. This contrasts with user-generated configuration files where each user could adopt their own idiosyncratic approach to specifying parameters and data.

- It enables easy access to multiple datasets via queries (e.g., to find all datasets with recordings from a particular brain region or that used a particular behavioral paradigm)
- It is concurrently accessible to multiple users.

Because DataJoint binds the code for running the computation with the table that will store the result, populating the same table will execute the same code. DataJoint also provides additional features for reproducible data analysis, such as maintaining data integrity of the database (e.g., deleting a table entry causes cascading deletion of dependent entries in downstream tables) and the files containing the results (e.g., by checksum verification).

How does Spyglass differ from DataJoint?

While Spyglass is based on DataJoint, it offers many useful features that DataJoint lacks. These include:

- A tight integration with the NWB format: When the NWB file is ingested into Spyglass, pointers to the data types appear as rows in a set of predefined tables. These serve as the starting point for analyses and an interface for the users to access the raw data within Spyglass. We provide the mapping between NWB data types and corresponding Spyglass tables on the documentation website.
- Extending table types: Spyglass provides a mix-in class, which allows different tables to inherit shared behaviors without duplicating code, for defining table types that are not included in DataJoint or extending the function of existing table types. This is used to implement many key table types such as Merge tables, which allow multiple upstream pipelines to feed into the same downstream pipeline. This example is illustrated in the description of the spike sorting pipeline below.
- Permission-based delete: Spyglass enables the deletion of individual rows in a table based on pre-defined user permission. This is not naturally supported by MySQL, the underlying relational database management system used by DataJoint.
- Improved searching based on restrictions on non-primary keys: Spyglass allows the users to conveniently track the provenance of a particular row in a downstream table across multiple upstream tables with only partial information.
- Export system for publishing: Spyglass provides a convenient way to export only the part of the database used for generating results and figures for a publication. This is done by caching the information about tables that are accessed when generating figures.
- The inclusion of various helper functions, which are detailed on the Spyglass documentation website.

Setup and hardware requirements

Spyglass can be installed in any setting that can support Python via the Python Package Index (PyPI). We provide detailed installation instructions on the documentation website, including a complete list of software requirements. In addition to the Python package, using Spyglass requires running a relational database (currently MySQL backend is supported). In our laboratory, we run this from a Docker container provided by DataJoint on a lab-wide server and grant access to members of the lab and other collaborators. This local configuration is recommended for use cases involving ~10 users. For a larger scale deployment, one could also run the Spyglass database in the cloud using services such as AWS.

Practical use cases and extensions

Spyglass comes with many pre-defined pipelines that implement common analysis tasks for electrophysiological and behavioral data. For users interested in using these pipelines, they can do so as soon as they ingest their NWB files into the database. Spyglass can also serve as a jumping off point for exploratory data analysis. For example, the user can conveniently read specific data types from the NWB file by first ingesting it into Spyglass and accessing database tables with Spyglass functions (e.g., `fetch_nwb`) or even load those objects in a format compatible with Pynapple (***Viejo et al., 2023***; `fetch_pynapple`). If they need to pre-process the data first, they can do so by running the relevant pipelines. Once the user has decided to formalize a particular analysis that is not yet supported by Spyglass, they can extend Spyglass and create user-generated custom pipelines. These could include data types from NWB files not currently supported by Spyglass (e.g., photometry, optical physiology, etc.) or build on existing Spyglass pipelines. Because the raw data and intermediate results are in

NWB format, the custom pipelines can take advantage of analysis software packages within the NWB ecosystem.

Organization of analysis pipelines

Here, we delve deeper into the design and organization of analysis pipelines in Spyglass. As mentioned previously, the analysis pipelines are defined as a set of tables in the relational database. Specifically, Spyglass uses DataJoint syntax to define tables as Python classes (see online documentation on Custom Pipelines and this video for examples). The code for executing the analysis is associated with these tables as class methods, enabling a tight integration of the database structure with the code for populating it. We refer the reader to the DataJoint documentation for more details on specific commands to interact with the database.

When an NWB file is first ingested into Spyglass, pointers to the data types in the NWB file are stored in database tables of the `Common` module. Each `Common` table corresponds to a data object in the NWB file and serves as an interface to retrieve it with simple function calls (`fetch_nwb`). The retrieval is 'lazy' in the sense that only a specific part of the data is loaded for analysis instead of the entire NWB file.

An analysis pipeline consists of sets of tables downstream of the `Common` tables. In each step in the analysis, the user populates one of four table types (***Figure 2A***):

- `Data` tables contain pointers to data objects in either the original NWB file or ones generated by an upstream analysis.
- `Parameter` tables contain a list of the parameters needed to fully specify the desired analysis.
- `Selection` tables allow users to select and pair a data entry and a parameter entry, defining the input to the `Compute` table.
- `Compute` tables execute the computations to carry out the analysis using the `Data` and `Parameters` specified in the `Selection` table entry. These results are then stored and can serve as `Data` for downstream analysis.

This design has multiple features that we have found to be beneficial. First, `Parameter` tables store the full set of parameters needed to specify a given analysis. For example, a `Parameter` table entry for a firing rate analysis of a single neuron might specify the bin size and smoothing to be used for that analysis. Multiple such entries can be defined, allowing a user to select the most appropriate one for the question being addressed. Second, because `Selection` tables specify which `Parameter` table entry was used for a given analysis on the associated `Data` table entry, they provide the key information needed to know which parameters were used to generate the entry in the downstream `Compute` table. Third, it is simple to associate a given `Data` table entry with multiple `Parameter` table entries and then re-run the analysis on those pairs. This enables a user to understand how their choice of parameters impacts their results, something that is otherwise difficult to manage and track.

Spyglass includes pipelines for a diverse range of analysis tasks in systems neuroscience, such as the analysis of LFP, spike sorting, video and position processing, and fitting state-space models for decoding neural data. Tutorials for all pipelines are available on the Spyglass documentation website (***Table 1***). Our goal was to take advantage of other open-source packages, and we have therefore integrated support for Pynapple (***Viejo et al., 2023***), a general-purpose neural data analysis package. We also built our pipelines to take advantage of other community-developed, open-source packages, like GhostiPy (***Chu and Kemere, 2021***), SpikeInterface (***Buccino et al., 2019***), DeepLabCut (***Nath et al., 2019***), and MoSeq (***Weinreb et al., 2024***). These pipelines store a complete record of the analysis and simplify the application of these tools. Furthermore, multiple versions of the pipelines can co-exist to apply different algorithms to a single data set, making it easy to probe the robustness of the results (see *Merge motif* below). Finally, the pipelines are modular as long as they process different kinds of data stored in the NWB files.

Next, we provide a detailed description about the implementation of three common analysis tasks in Spyglass pipelines: (i) filtering broadband extracellular voltage traces to extract the lower-frequency LFP bands; (ii) detecting discrete events (e.g., sharp-wave ripples [SWRs], a hippocampal event marking the time of bursts of population activity) in the LFP signals; and (iii) spike sorting and curation.

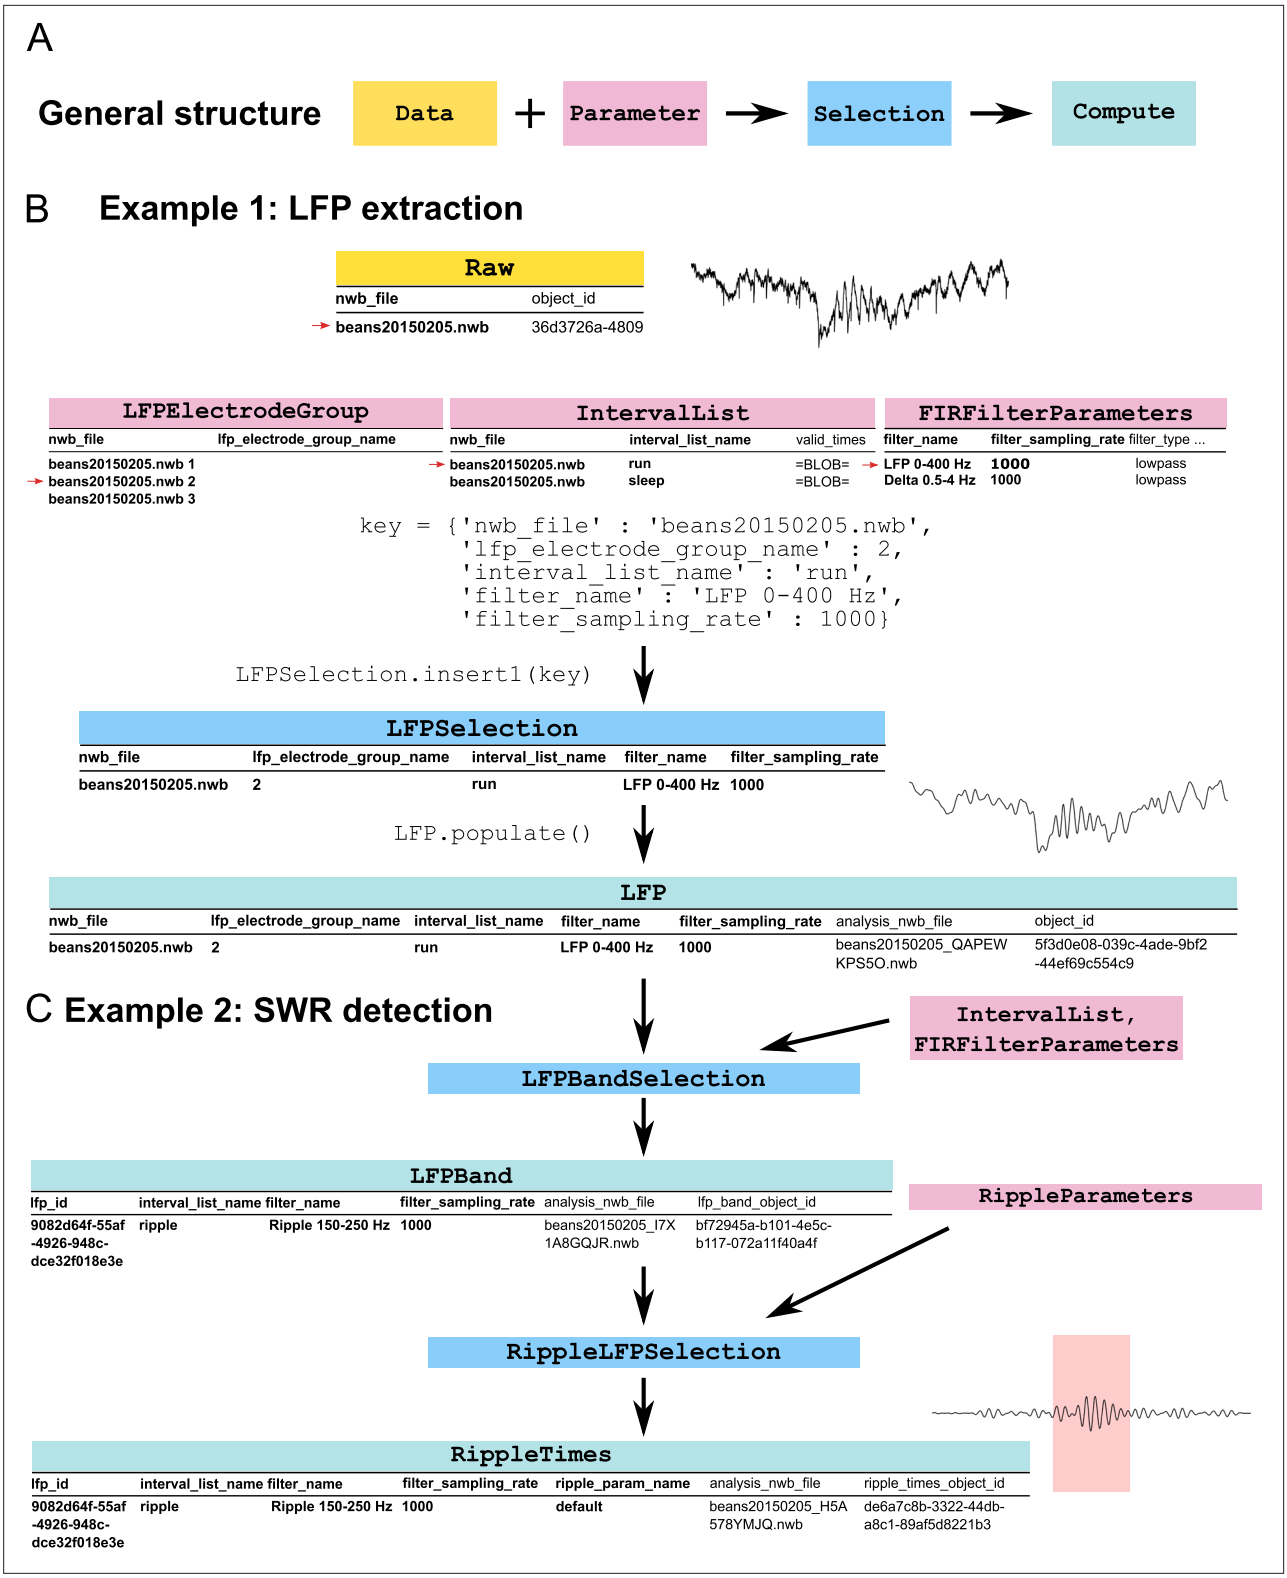


Figure 2. Analysis pipelines in Spyglass. (**A**) A general structure for a Spyglass pipeline. (**B**) Example 1: LFP extraction. Note the correspondence to the pipeline structure in (**A**) as shown by the color scheme. The trace next to the Raw table is raw data sampled at 30 kHz and is represented by a row in the Raw table. This, along with parameters from LFPElectrodeGroup, IntervalList, and FIRFilterParameters tables (red arrow), are defined in a Python dictionary and the LFPSelection.insert1(key) call is used to insert the reference to the raw data and the now associated parameters into the LFPSelection table. When the populate method is called on the LFP table, the filtering is initiated and the output is inserted into the database. The results (e.g. the

Figure 2 continued on next page

Figure 2 continued

trace above the LFP table) are stored in NWB format and its object ID within the file is also stored as a row in the LFP table, enabling easy retrieval. (**C**) Example 2: Sharp-wave ripple (SWR) detection. Note that the key specification, insert, and populate calls are omitted for simplicity. This pipeline is downstream of the LFP extraction pipeline and consists of two steps: (i) further extraction of a frequency band for SWR (LFPBand) and (ii) detection of SWR events in that band (RippleTimes). Note that the output of LFP extraction serves as the input data for the SWR detection pipeline and can thus be thought of as both Compute and Data types. As in (**B**), for each step, the results are saved in NWB files and the object IDs of the analysis result within the NWB file are stored as rows in the corresponding Compute tables. The trace above the RippleTimes table is the SWR-filtered LFP around the time of a single SWR event (pink shade). In each table, columns in bold are the primary keys. Arrows depict dependency structure within the pipeline.

Example 1: LFP extraction (***Figure 2B***)

To extract the LFP signal (below 400 Hz), we use the pipeline shown in ***Figure 2B***. First, we select a row from the `Raw` table, a `Data` table that points to an `ElectricalSeries` object in the NWB file. We then specify the parameters of the analysis in the `Parameter` tables: the list of channels for which LFP should be extracted (`LFPElectrodeGroup`), the time interval for the LFP extraction (`IntervalList`), and the coefficients for the filter that will be used on the data (`FIRFilterParameters`). These parameters are associated with the entry in the `Raw` table by defining a Python dictionary object that specifies the `Data` and `Parameter` entries and inserting it into a `Selection` table (`LFPSelection`) by calling the `LFPSelection.insert1` method (***Figure 2B***). Finally, we apply the filter to the selected data over the selected interval using the LFP table (a `Compute` table) by calling the `LFP.populate` method. The resulting filtered data is saved to disk in the NWB format, and the object ID associated with the `LFP` object within the NWB file is also stored in the `LFP` table for easy retrieval. Thus, the corresponding entry in the LFP table contains all the details about the data and the parameters, allowing a user to fully track the provenance of the output.

Example 2: sharp-wave ripple detection (***Figure 2C***)

Once the LFP extraction is completed, we can build on the results by applying another filter to isolate a specific frequency band and identifying SWRs, a prominent LFP event within hippocampal data. This pipeline is illustrated in ***Figure 2C***. It applies two additional steps to a row in the `LFP` table: another band-pass filter to isolate the 150–250 Hz band and a subsequent detection of SWR events. Each step uses the same basic scheme shown in ***Figure 2A***. These include defining a specific band-pass filter in the `Parameter` tables; selecting a time interval for the bandpass filtering; and adding an entry to `LFPBandSelection` table that binds both the filter parameters and the time interval with a row in the LFP table. A call to `LFPBand.populate` generates an NWB file containing the ripple-band data and an entry in the `LFPBand` table with information about which data and parameters were used. Next, the user selects an entry in `RippleParameters` to define the parameters for detecting the ripple events (e.g., threshold over the spectral power) and associates it with filtered data in LFPBand in the `RippleLFPSelection` table. Finally, the `RippleTimes` table is populated (by `RippleTimes.populate`), which identifies the start and end times of each ripple event and saves these to a new NWB file.

Example 3: spike sorting and curation (***Figure 3***)

The spike sorting pipeline (***Figure 3***) combines the principles of analysis pipeline design we outlined previously with additional design features. This pipeline uses SpikeInterface (***Buccino et al., 2019***) to perform the operations critical for spike sorting, but also tracks all of the parameters used and provides a system for tracking multiple sorting curations. The pipeline includes the following steps: (1) preprocess the recording (e.g., filter and whiten to remove noise); (2) apply spike sorting algorithm (e.g., MountainSort4, Kilosort3, etc.); (3) curate the results (e.g., either manually or automatically by computing quality metrics); and (4) consolidate the output with other sources of sorted units (e.g., those already present in the NWB file) for downstream analysis. Each of these steps follows the general design shown in ***Figure 2A***. We also detail additional features that have not been discussed previously.

Global `Parameter` tables (e.g., `IntervalList`)

An important object in any analysis is the time interval during which the data were collected or to which analysis procedures should be applied. To avoid having a separate table for time intervals

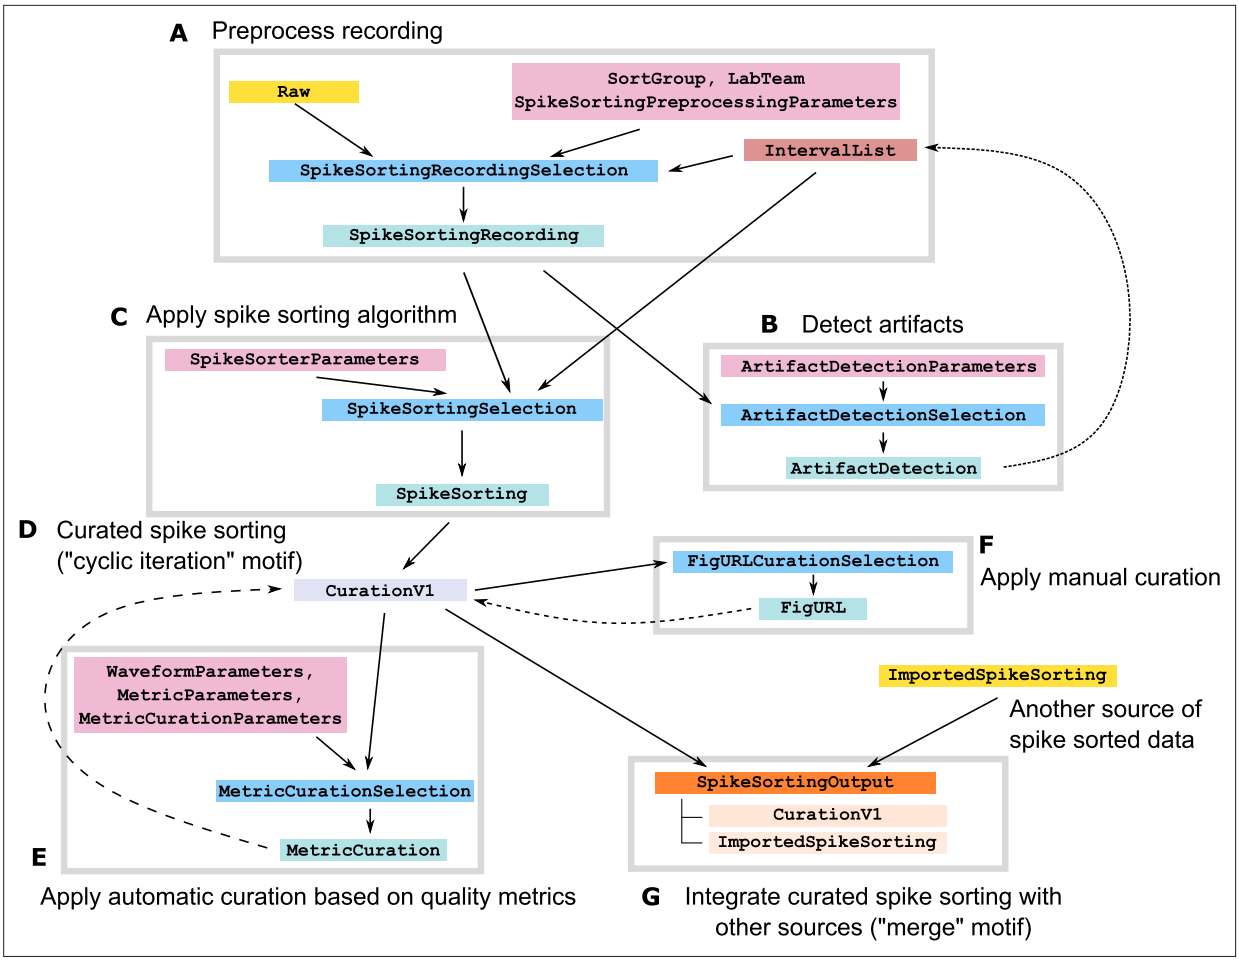


Figure 3. Spike sorting pipeline. The Spyglass spike sorting pipeline consists of seven components (large gray boxes), all of which take advantage of the SpikeInterface library: preprocess recording (**A**); detect artifacts to omit from sorting (**B**); apply spike sorting algorithm (**C**); curate spike sorting (**D**), either with quality metrics (**E**), or manually (**F**); and merge with other sources of spike sorting for downstream processing (**G**). Solid arrows describe dependency relationships and dashed arrows indicate that the data is re-inserted upstream for iterative processing. Note the two design motifs (see text): 'cyclic iteration' for curation and 'merge' for consolidating data streams. The color scheme is the same as ***Figure 2***, except for light purple (cyclic iteration table), orange (merge table), and peach (parts table of the merge table).

in every pipeline, we store them in the `IntervalList` table of the `Common` module for all pipelines. For example, in the spike sorting pipeline (***Figure 3***), `IntervalList` provides a time interval for both preprocessing the recording (`SpikeSortingRecordingSelection`) and running a spike sorting algorithm (SpikeSortingSelection). In addition, the intervals during which artifacts (i.e., high-amplitude voltage transients from behavioral events such as licking) occur are identified and fed back into `IntervalList` (dashed arrow in ***Figure 3***).

'Cyclic iteration' motif for curation

Certain pipelines, such as curating the output of spike sorting, may need to be run multiple times on the same data. For example, one might first compute quality metrics to identify noise clusters and potential candidates for merging over-clustered units (Automatic); then inspect, merge, and apply curation labels to the result with an external viewer (Manual); and finally, compute a final set of metrics to describe the quality of each unit (Automatic). This results in a sequence of curation steps: Automatic, Manual, Automatic. Depending on the data, the user may choose a different curation sequence, and the order and length of these sequences might change as new algorithms and metrics are developed. This presents a challenge in modeling the pipeline within the relational database.

We therefore developed a specific design motif to enable this iterative curation with a finite number of tables (***Figure 3***). First, a given row of the `CurationV1` table (the output of the spike

sorting step) is taken through automatic or manual curation steps downstream. Upon completion, the spike sorting object may enter this curation pipeline again as a new row in the `CurationV1` table. Importantly, the new row has information about previous curation from which it descended. This allows the user to track each round of curation while applying as many steps as desired. It can also be easily extended; if new automatic curation algorithms are developed in the future, it can simply be added downstream to the `CurationV1` table, enabling application of the latest methods to previously collected data.

'Merge' motif for consolidating data streams and versioning pipelines

A different challenge arises when the user wants to feed multiple streams of data of the same type into a single downstream pipeline. For example, once curation is completed, the spike sorting is saved in `CurationV1`. But some NWB files may already contain curated spike sorting (as a row in the table `ImportedSpikeSorting`), and one may want to apply the same downstream pipeline to both data sources to compare the results. In yet another case, the other data stream could be a different version of the spike sorting pipeline (e.g., `CurationV2`) that uses different algorithms but produces output of the same type. Adding the same downstream pipeline to each of these separately would result in code redundancy and database bloat. Simply having these converge onto a single downstream table is not desirable either, as it will require modifying an existing table to add new columns every time a new version or new data stream is added.

To solve this problem, we have designed a 'merge' table type (***Figure 3***). Here `Parts` tables (a table type within DataJoint tightly associated with a parent table) are used to implement the merging of multiple data streams into a single table. The downstream pipeline then gets data from this table without any duplication. More details for the implementation and helper functions to maintain data integrity can be found in the tutorial notebook (***Table 1***, 04_Merge_Tables).

Sharing data, analysis, and visualization

Complete sharing of data and analysis at the end of projects

A key goal of our system is to simplify sharing data and analyses when results are ready to be published. Because all raw and intermediate data are in the NWB format, they can be directly deposited to DANDI (***Rübel et al., 2022***; ***Callaway et al., 2021***; ***Hawrylycz et al., 2023***; ***Park et al., 2024***), a NIH-supported public archive for neuroscience data. Sharing the analysis code is also easy: simply share the codebase for the analysis pipelines (i.e., Spyglass plus any project-specific pipelines) and the scripts used to populate the database. Others can then download the raw data from DANDI, set up the database with Spyglass, and recreate all results locally by executing the population script. Alternatively, users may want to share the Spyglass database in its populated state so that the community can access it directly without going through the setup procedures or re-running time-consuming analysis steps. This can be done by (i) hosting the database on the cloud and granting access to users outside the lab or (ii) exporting and sharing parts of the database that were used by the project. Spyglass facilitates the second option by providing functions that automatically log the table entries and NWB files used for creating figures of a manuscript in a Python environment (***Table 1***, 05_Export). The dependencies of these entries are traced through the database to compile the complete set of raw, intermediate, and plotted NWB files and their corresponding database entries. These are stored in the 'Export' table, which also generates a bash script to create SQL dumps of the identified database entries.

To upload these files to DANDI, users must first register a new dandiset for their project and record their API and dandiset ID. With this information, they can then use the method `DandiPath.compile_dandiset()` to automatically validate, organize, and upload all project files to the DANDI archive. Additionally, this process stores the archive information for each file in the `DandiPath` table, allowing `fetch_nwb` to automatically stream data from the DANDI cloud storage when not available locally.

To create a sharable Docker image of the project, we provide a template repository spyglass-export-docker (***Broz, 2024***). Users first download a local copy of this repo and copy the SQL dump file, environment YAML, and figure-generating notebooks generated during spyglass export into the appropriate folders. Running the provided docker compose scripts then generates two linked docker containers: one running the reconstructed spyglass SQL database, and a second connected to this database and running a JupyterHub—with a Python environment matching that used when generating

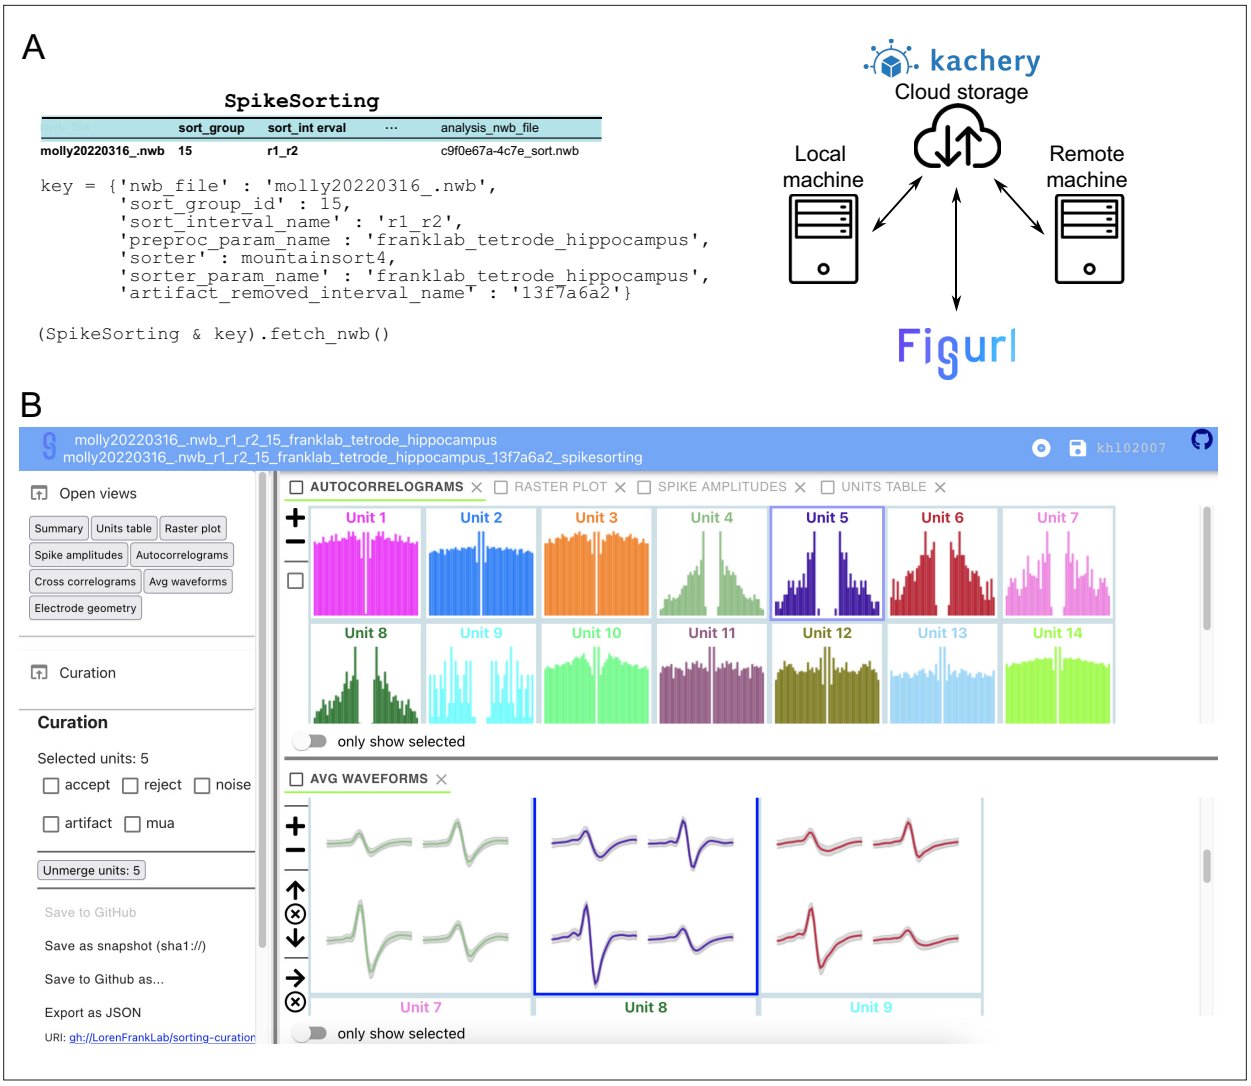


Figure 4. Sharing data and visualizations. (**A**) Kachery provides a convenient Python API to share data over a content-addressable cloud storage network. To retrieve data from a collaborator's Spyglass database, one can make a simple function call (fetch_nwb) that pulls the data from a node in the Kachery Zone to the local machine. (**B**) Example of a Figurl interactive figure for visualizing and applying curation labels to spike sorting over the web.

the figures. These can be readily shared with new users to provide them immediate access to all steps of the analysis process and the corresponding data through DANDI streaming.

Controlled sharing for ongoing projects

For ongoing projects, users may want to limit the sharing of the analyses to their collaborators. This requires controlling access to the database and the underlying NWB files that contain the raw or intermediate data. This is straightforward to manage in Spyglass. DataJoint handles access to the database natively by requiring a username and a password. Managing access to the NWB files is handled by Kachery (***Magland et al., 2024***), a content-addressed sharing tool for scientific data (***Figure 4A***). Specifically, the user selects the NWB files to be shared by inserting pointers to them into `NwbKachery` and `AnalysisNwbKachery` tables within Spyglass. When the collaborator attempts to access these files, Kachery first looks for them in their local system. If not found, the corresponding files are automatically uploaded from the user's system to a cloud storage server and then downloaded to the collaborator's computer. This feature is detailed in a tutorial (***Table 1***, 03_Data_Sync). Critically, the downloaded files are never modified locally within Spyglass, and an attempt to access a modified file would result in a DataJoint error. This ensures that each user is working on the same underlying data even if they are at different sites. More generally, Kachery offers advantages over

other file hosting services (e.g., Dropbox and Google Drive) or alternative architectures (e.g., IBL data architecture) by not requiring a central location to track available files and providing a user-friendly Python API. We point interested readers to the Kachery GitHub repo (https://github.com/magland/kachery, ***Magland, 2026***) for further descriptions.

Sharing visualizations

Spyglass enables users to create and share interactive visualizations of final and intermediate analysis results through the Figurl package(***Magland et al., 2025***). These visualizations facilitate understanding complex, multi-modal neuroscience datasets by allowing users to (i) quickly compare different stages of processed data to spot issues with their data, and (ii) align multimodal information sources to get a more holistic view of their dataset. Figurl is integrated within Spyglass as dedicated tables attached to specific pipelines such as spike sorting (***Figure 4B***) and neural decoding (***Figure 5***). Populating these tables generates a URL to web-based visualizations for exploring complex, multi-dimensional time series across multiple views whose time axes can be linked. Sharing them is also easy, as the URL can be accessed from any browser without the need for local software installation or specialized hardware. This allows collaborators anywhere in the world to easily access and explore the data.

Demonstration of generalizability: neural decoding of position in multiple datasets

A major goal of Spyglass is to facilitate the analysis of data across multiple datasets that may come from different laboratories. To illustrate this, we ingested and analyzed two NWB files containing single-neuron recordings from rat hippocampus, one from our laboratory and another from the Buzsáki laboratory at NYU (***Petersen and Buzsáki, 2020***). Specifically, we applied a switching state space model (***Denovellis et al., 2021***; ***Denovellis et al., 2019***) to decode the animal's position from spikes and infer periods of different types of non-local representations (such as replay and theta sequences), during which the decoded position deviates from the animal's true position. This is a complex analysis that involves integrating multiple data sources, including position and neural spiking activity, and applying an advanced statistical model with many user-defined parameters. The decoding pipeline in Spyglass enables the user to carry out every step of this analysis, including 'preprocessing' of the data (e.g., linearize the 2D position of the animal, perform spike sorting, or import units that have already been sorted) and fitting of the model (see ***Figure 5—figure supplement 1*** for a visualization of the steps involved). After running the decoding pipeline, we visualize the results on the browser via Figurl and generate plots to further reproduce the reported results.

The UCSF dataset contains large-scale hippocampal recordings in an animal performing a foraging task in a maze with six reward sites and dynamic reward probabilities (***Figure 5A***, top panel). Applying the decoding pipeline to these data yields a probability distribution over space in 2 ms bins that describes our estimate of the 'mental' position of the animal. This mental position tracks the animal as it traverses the maze (***Figure 5A***, second panel from top; see interactive visualization via Figurl) but also shows interesting systematic deviations from actual position. Computing the distance between the peak of the probability distribution and the actual location reveals characteristic patterns of such deviations from the actual position (***Figure 5A***, third panel from top) in which the decoded position sweeps ahead of the actual position and then back during movement bouts. This pattern recurs at ~8 Hz, reflecting the well-known 'theta sequences' seen in the hippocampus (***Skaggs et al., 1996***; ***Foster and Wilson, 2007***).

We then applied this same pipeline to the NYU dataset, where animals performed a spatial alternation task on a maze with a "figure-8" topology (***Figure 5B***, top panel). As expected, we could identify theta sequences in these data as well, highlighting the robustness of these phenomena (***Figure 5B***, second and third panels from top, see interactive visualization via Figurl). Moreover, the NYU dataset includes a specific manipulation in which the medial septum, a brain region critical for pacing the theta rhythm, was cooled, reducing the theta frequency from 8–10 Hz to 5–8 Hz. The authors originally carried out several detailed analyses to demonstrate that cooling reduced theta frequency and impaired behavior without changing the overall spatial tuning of single neurons or their tendency to fire sequentially within theta cycles. However, the authors did not apply state-space decoding methods and did not characterize the effects of cooling on the decoded representation of space in relation to the animal's actual position. We therefore applied our decoding pipeline to the cooling

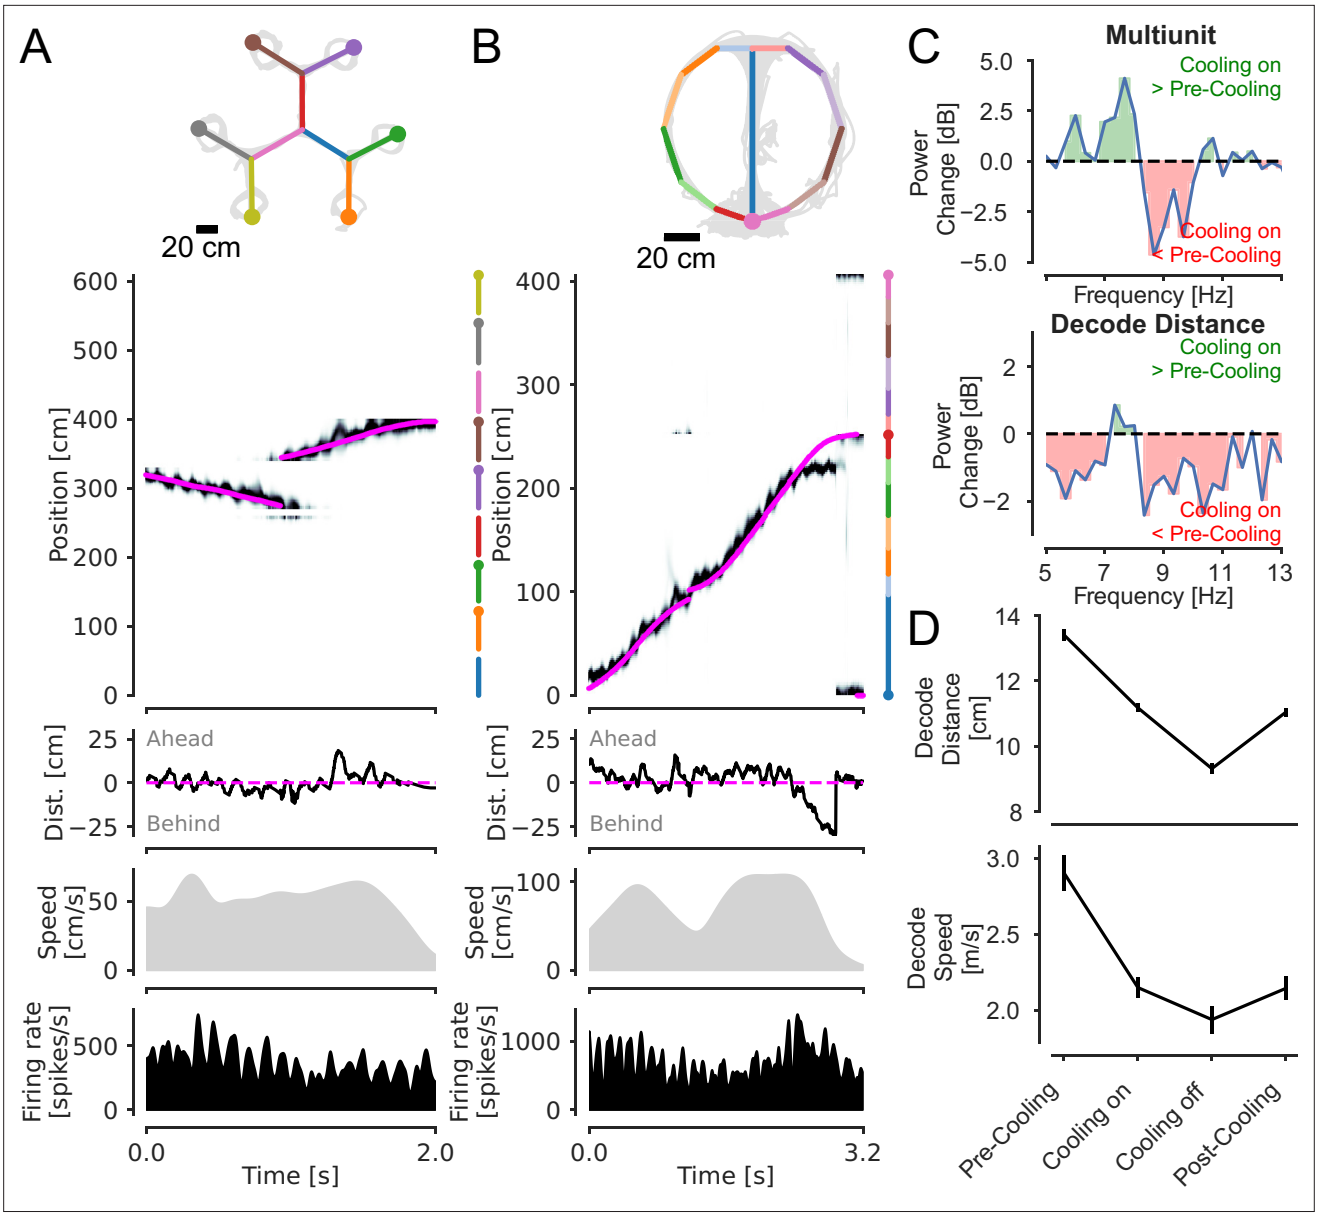


Figure 5. Applying decoding pipelines to multiple datasets from different labs. (**A**) Decoding neural position from rat hippocampal CA1 using a clusterless state space model (UCSF dataset). In the top panel, grey lines represent positions the rat has occupied in the spatial environment. Overlaid lines in color are the track segments used to linearize position for decoding. Filled circles represent reward wells. The second panel from the top shows the posterior probability of the latent neural position over time. The magenta line represents the animal's actual position. The vertical lines on the right represent the linearized track segments with the colors corresponding to the top panel. The third panel from the top shows the distance of the most likely decoded position from the animal's actual position and the sign indicates the direction relative to the animal's head position. The fourth panel from the top is the animal's speed. The final panel is the multi-unit firing rate. (**B**) Decoding from rat hippocampal CA1 using existing spike sorted units (NYU dataset). Conventions are the same as in A. The filled circle in the linearization represents the reward zone rather than the reward well. (**C**) Decoding analysis of the NYU dataset. The top panel shows the power difference of the multi-unit firing rate between the medial septal cooling period and the pre-cooling period in the 5–13 Hz range. The power at 8–10 Hz is attenuated during cooling while the power at 5–8 Hz is enhanced, showing a slowing of the theta rhythm during cooling. The bottom panel shows that the power of the distance between decoded and actual position (decode distance) is mostly reduced throughout the 5–13 Hz range. (**D**) Cooling decreases the decode distance and speed, and this effect may only recover partially after cooling. Bars represent 95% confidence intervals.

The online version of this article includes the following figure supplement(s) for figure 5:

Figure supplement 1. Spyglass pipeline workflow for ***Figure 5A and B***.

trials ('cooling on') and the control trials preceding it ('pre-cooling'), just after it ('cooling off'), and the recovery trials 10–12 min after cooling ('post-cooling').

The results of these analyses were consistent with the published findings and provided new characterizations that could serve as the foundation for additional discoveries. We first estimated the multi-unit firing rate as a proxy for the theta LFP and characterized its power spectrum before and after cooling. As expected, cooling decreased the power above ~8 Hz and increased the power below ~8 Hz, consistent with the slowing of theta LFP shown in the original manuscript (***Figure 5C***, top panel). We then applied the same analysis described above to the distance between the decoded and the actual position during movement ('decode distance'), expecting cooling to have a similar effect on its power spectrum. Interestingly, here cooling led to a decrease in power at essentially all frequencies (***Figure 5C***, bottom panel). Consistent with this result, the decode distance decreased from the pre-cooling to cooling period, with a partial recovery during the post-cooling period (***Figure 5D***, top panel). Similarly, the average speed at which the decoded position moved ahead and behind the animal was also reduced during cooling and showed a partial recovery after the cooling period (***Figure 5D***, bottom panel). These results indicate that cooling reduces both the extent and the rate at which the decoded position deviates from the actual position. This was unexpected given that cooling had no effect on the average spatial tuning of these cells (***Petersen and Buzsáki, 2020***). It also raises an interesting hypothesis: hippocampal representations of distant locations may be exquisitely tuned to the specific frequency of the rhythmic input from medial septum, such that slowing the rhythm down by just 2–3 Hz significantly limits their expression. More broadly, these findings illustrate the power of our framework that enables both replication of results across datasets and the re-analysis of previously collected data.

Discussion

Summary of results

Science is a social enterprise that relies heavily on collaboration and transparency among researchers (***Committee on Responsible Science et al., 2017***; ***OECD, 2015***). Reproducible and shareable data analysis plays a critical role in this context, as it ensures that scientific findings can be independently verified and built upon by others. To facilitate this for the neuroscience community, we built Spyglass, a software framework that combines the NWB format and the relational database structure. Building on many community-developed tools, it provides useful features to design complex analysis pipelines, share raw and processed data, generate web-based visualizations, and analyze data from multiple sources. As a result, it simplifies collaboration within and across labs, making it well suited as a community framework for neurophysiological and behavioral data analysis.

Comparison to prior work

Our work builds on many previous approaches that have been proposed for scientific data management and reproducible analysis pipelines. This includes work from individual laboratories that have illustrated how a few elements of an NWB file could be read into a DataJoint database, and publications highlighting datasets available in NWB (***Chandravadia et al., 2020***). More broadly, DataJoint is used by many labs with lab-specific pipelines (***Yatsenko et al., 2021***), but to our knowledge none of these efforts integrate cross-laboratory data and visualization tools or use NWB as the foundation to facilitate sharing. Our system also contains elements similar to those developed by large collaborative groups like The International Brain Laboratory (IBL) that are designed to organize neurophysiology data for sharing with collaborators and a module to automatically run analyses (***Acerbi et al., 2023***). But the conversion to a standardized format (outside the collaboration or group) and public data sharing are only done following substantial analysis in the IBL system, complicating replication of the full analysis.

Other approaches do away with the relational database altogether. For example, DataLad uses version control tools such as git and git-annex to manage both code and data as files (***Halchenko et al., 2021***). This enables the creation of a data analysis environment and decentralized data sharing. For building analysis pipelines, it may be combined with other tools for managing the sequential execution of scripts. For example, Snakemake (***Mölder et al., 2021***) and related projects such as Cobrawap (***Gutzen et al., 2024***) allows the users to gather and define the input, output, and the

associated scripts to execute for each analysis step, thereby tracking the dependency between steps. But because these tools do not provide any formal structure for data analysis or parameter specification, they lack the advantages of the relational database that we discussed, such as being able to easily organize or search for the records of previous analysis based on specific parameters, efficient data sharing and access management to multiple users, and built-in data integrity checks based on constraints native to the database (e.g., primary keys).

By contrast, Spyglass begins with a shared data format that includes the raw data and offers both transparent data management and reproducible analysis pipelines using a formal data structure. One distinguishing feature of Spyglass is the emphasis placed on combinatorial matching of data and method in a reproducible way. For example, Spyglass makes it easy to apply multiple spike sorting algorithms to a given dataset and to compare the results, as this involves simply matching the data with different rows in the parameter tables. Spyglass also makes it straightforward to apply complex analyses like decoding to datasets from multiple labs, facilitating replication and data re-use. The system can be extended by adding new pipelines to the existing database as better tools and algorithms become available. These features enable the re-analysis of data to examine how the results depend on the choice of algorithm. We believe it is critical to provide this kind of future compatibility to maximize the impact of the years of experimental work that go into each dataset.

Limitations

Although Spyglass provides many useful features for reproducible data analysis, it has several limitations. Because of the central role played by the NWB format in Spyglass, a potential user must first convert their data to NWB, which requires time and effort (***Pierré et al., 2024***). In addition, some types of data may not have defined standards within NWB (e.g. surgical procedure details, descriptions of conditions, detailed subject information) and if the user wishes to include those details, they would need to build an NWB extension and parallel Spyglass tables to do so. NWB also allows users to choose their own names for some data types (e.g., behavioral tasks), which means that standardization across laboratories requires agreements on naming conventions.

In addition, users are expected to set up and maintain a relational database, which may require additional training. Using Spyglass includes learning to work with the structure of DataJoint, such as the strict data integrity requirement that can make modification of existing tables difficult. Spyglass also does not yet include pipelines for processing certain types of neural data, such as optical physiology or fiber photometry, and does not currently support Windows (although it may be possible to run Spyglass using the Windows Subsystem for Linux). Finally, as for all software frameworks, the evolution or lack of maintenance of other packages presents a challenge for long-term support and reproducibility.

Fortunately, there are ongoing efforts to address these challenges. These include tools to simplify the raw data conversion into NWB, such as NeuroConv, a package to convert neurophysiology data in common formats to NWB automatically, and NWB GUIDE, a desktop app that guides users through the process of converting data to NWB without writing any code. Using Spyglass could also help with standardization efforts across labs: having a database makes it easy to create lists of names used to refer to particular items and to then move toward standardization.

We also provide many tutorials on the documentation website so that the user can efficiently set up a database and learn to use Spyglass. We continue to actively maintain Spyglass and are eager to work with the community to extend it and support data types and analyses beyond what is currently available. These efforts will increase the usability and reach of Spyglass and make its adoption more attractive, particularly to early-stage investigators. Finally, even in cases where reproducing a result would require installing older versions of software, the results themselves remain accessible within NWB files referenced in Spyglass, ensuring that previous results can be built on even as packages evolve.

Future applications

Spyglass and similar tools have the potential to transform scientific data analysis. In addition to facilitating examination or extension of published results, they enable meta-analysis across studies and easy testing of novel methods across multiple datasets. The machine-readable form of data and analysis pipelines also opens doors for machine-driven analysis and hypothesis testing. As these tools

develop and become more accessible, we believe that frameworks like Spyglass will likely become essential for neuroscience researchers.

Materials and methods

Coding environment

Spyglass was developed in Python 3.9 and is compatible with version 3.10 as well. See our dependency list for a full list of Python packages used.

NWB conversion

To facilitate conversion of raw data to NWB format, we offer trodes-to-NWB (***Bray et al., 2026***), a sister package to Spyglass for converting data acquired with the SpikeGadgets hardware to NWB. This comes with a web-based GUI for conveniently generating a YAML file containing the metadata used by `trodes-to-NWB`. For converting data not acquired with SpikeGadgets, users can use NWB conversion tools developed by the NWB team, such as NeuroConv and NWB GUIDE.

NWB file conventions

We adopted a specific set of conventions for our NWB files. Some of these conventions rely on a specific set of Frank lab-specific NWB extensions (***Ly, 2026***):

- Time:
  - Spyglass inherits from the source NWB file either the explicit or implicit timestamps. NWB files from Frank lab have explicit timestamps for each sample in Unix time (seconds since 12:00 AM January 1, 1970). This lets users know exactly when data were collected. Spyglass is also compatible with other approaches, however, including implicit timestamping consisting of the start time and sampling rate.
- `ElectrodeTable` and `ElectrodeGroup`:
  - `ElectrodeGroups` are stored in a custom NWB extension that also includes the name of the targeted brain region for each group.
  - The NWB file contains information about the relative locations of each of the electrodes within each physical device used for data collection. This ensures that the relative locations of the electrodes are available for spike sorting and registration to histology.
- Video files
  - The relative path to the video files collected along with the recordings is stored in the NWB file.
- Additional files
  - Other files important to recreate the conditions of the experiments can be saved, depending on the format. For example, the code used for implementing the behavioral paradigm or reward contingency can be stored as text objects in the NWB file.

NWB file ingestion

Although the NWB format serves as a community standard for neurophysiology data and has a list of best practices, it allows some flexibility in the specification of data within NWB files to accommodate user preferences. For example, the `ElectricalSeries` object that stores the electrophysiology data may have different names depending on the convention chosen by the investigator, which may complicate programmatic access to the data. To make Spyglass interoperable with NWB files of varying degrees of NWB-compliance, we have created an option to supply or override information that is missing in the NWB file but is nevertheless required by Spyglass via a configuration file that can accompany the NWB file. We provide an example of this approach in a tutorial (***Table 1***, 02_Insert_Data).

Permission-handling and cautious delete

Spyglass is based on a relational database that is accessible to multiple users. In some cases, the type of operations that can be applied to individual data entries (i.e., rows of a table) may need to be restricted to a specified set of users. This is particularly true for operations that are irreversible or time-consuming, such as deleting a row from a table storing analysis results. However, there is no inherent mechanism within MySQL or DataJoint that allows permission handling at the level of individual rows

of a table. To solve this problem, we have implemented a `cautious_delete` function, in which the user's permission to carry out a delete operation is checked before it is applied. The permission is granted based on team membership within the lab, reflected in the `LabTeam` table. Although this is not a formal permission-management system, it serves to prevent accidental deletions. We note that this system does incur additional overhead, and while that has not been an issue for us, it is possible that this would become problematic in use for much larger cross-laboratory collaborations.

Sharing files via Kachery

One way to share the results of Spyglass analysis pipelines is to make the database publicly available. This gives anyone the permission to access the rows of the tables that make up the pipelines and inspect the metadata and the parameters associated with each step of the analysis. But because Spyglass only saves a path to the NWB files containing analysis results within the tables, external viewers cannot download the data and examine it by default.

To enable controlled external access to the data, we have created a system to share selected analysis NWB files with a specified group of users via Kachery. We define a set of tables (`KacheryZone` and `AnalysisNwbfileKachery`) where users can associate analysis NWB files to be shared with a Kachery Zone, making it available to all remote clients who are members of the zone through cloud storage services like Cloudflare R2, or self-hosted servers. Once linked, Spyglass automatically requests, downloads, and manages analysis data for remote users attempting to access shared data through Spyglass tables. This provides a convenient way to provide access to the Spyglass pipelines and associated data files to collaborators.

Customizing pipelines

To alleviate the challenges associated with database design, we have identified design principles that have been tested extensively by multiple users in the Frank lab. These are described in the text and illustrated with examples in ***Figures 2 and 3***. We recommend users adopt these design elements for building their custom pipelines. We also describe the naming conventions for the tables defined as Python classes and important methods associated with them (e.g. for multiple versions of a pipeline) in our Developer Notes available online. Once the pipeline is sufficiently mature and potentially useful to other scientists, we encourage users to submit their pipelines as a pull request to our GitHub repository.

Decoding of position from NWB files from multiple laboratories

The Frank lab data is available on the DANDI archive (DANDI:000937). The Buzsáki lab data was also obtained from DANDI (DANDI:000059/0.230907.2101). For decoding the Frank lab data, we applied the clusterless decoding pipeline by detecting the amplitude of threshold-crossing events in the tetrode recordings. For decoding the Buzsáki lab data, we applied a sorted-spikes decoding pipeline. The code for these decoding pipelines, as well as detailed tutorials describing them, are available online (***Table 1***, 40_Extracting_Clusterless_Waveform_Features, 41_Decoding_Clusterless, 42_Decoding_SortedSpikes). Code to generate ***Figure 5*** can be found at: https://github.com/Loren-FrankLab/spyglass-paper (copy archived at ***Denovellis, 2026***). Briefly, decoding the latent neural position and extracting the distance between the most likely decoded position and the animal's position used methods described in ***Denovellis et al., 2021***. We used a timestep of 4ms and a position bin size of 2 cm with a continuous (6 cm variance Gaussian random walk) and fragmented (uniform distribution) discrete state. Place intensity receptive fields were estimated using a Gaussian kernel density estimate with a standard deviation of 6 cm for position and 24 μV for amplitude space (amplitude space was used for the clusterless analysis only). We calculated the power of the multiunit firing rate and the decoded distance from the animal by using a multitaper spectrogram during the pre-cooling and cooling periods. The time resolution was 3 s and the frequency resolution of 2/3 Hz with a single taper. We excluded immobility periods by using a threshold of 10 cm/s. Power difference was calculated by converting to the Decibel scale and taking the difference of average power under the cooling and pre-cooling conditions. The decoded speed of theta sequences was calculated by taking the absolute value of the second-order difference of the decoded distance from the animal (function numpy.gradient) multiplied by the sampling frequency (250 Hz).

Acknowledgements

We thank members of the Frank and Gillespie laboratories for bug reports and testing. We also thank Peter Petersen for consultations about analyzing his publicly available data, Daniel Liu for initial discussions about standardization of pipelines, and Abhilasha Joshi for consulting on the DeepLabCut pipeline. Finally, we thank Vanessa Bender for comments on the manuscript. This work was supported by HHMI funds and NIH grants RF1MH130623 and RF1MH133778 to LMF, and a Helen Hay Whitney Postdoctoral Fellowship and NIH grant 5K99EY036953 to KHL.

Additional information

Competing interests

Thinh Nguyen, Dimitri Yatsenko: Affiliated with DataJoint Inc; no other competing interests to declare. Caleb Kemere: Reviewing editor, eLife. The other authors declare that no competing interests exist.

Funding

Funder	Grant reference number	Author
Helen Hay Whitney Foundation		Kyu Hyun Lee
National Eye Institute	1K99EY036953	Kyu Hyun Lee
Howard Hughes Medical Institute		Loren Frank
National Institutes of Health	RF1MH130623	Loren Frank
National Institutes of Health	RF1MH133778	Loren Frank

The funders had no role in study design, data collection and interpretation, or the decision to submit the work for publication.

Author contributions

Kyu Hyun Lee, Conceptualization, Resources, Data curation, Software, Formal analysis, Supervision, Funding acquisition, Validation, Investigation, Visualization, Methodology, Writing – original draft, Project administration, Writing – review and editing; Eric L Denovellis, Conceptualization, Resources, Data curation, Software, Formal analysis, Supervision, Validation, Investigation, Visualization, Methodology, Writing – original draft, Project administration, Writing – review and editing; Ryan Ly, Software, Formal analysis, Validation, Investigation, Methodology, Writing – review and editing; Jeremy Magland, Andrew Tritt, Oliver Rübel, Thinh Nguyen, Dimitri Yatsenko, Joshua Chu, Caleb Kemere, Samuel Garcia, Alessio Paolo Buccino, Resources, Software; Jeff Soules, Resources, Software, Visualization; Alison E Comrie, Data curation, Software, Investigation; Daniel P Gramling, Jennifer A Guidera, Software, Investigation; Rhino Nevers, Ji Hyun Bak, Michael E Coulter, Xulu Sun, Emrey Broyles, Donghoon Shin, Sharon Chiang, Cristofer Holobetz, Software; Philip Adenekan, Resources, Software, Methodology; Chris Brozdowski, Software, Investigation, Methodology; Samuel R Bray, Data curation, Software, Investigation, Visualization, Methodology; Emily Monroe, Data curation, Software; Loren Frank, Conceptualization, Resources, Data curation, Software, Supervision, Funding acquisition, Investigation, Methodology, Writing – original draft, Project administration, Writing – review and editing

Author ORCIDs

Kyu Hyun Lee https://orcid.org/0000-0001-6483-9444
Eric L Denovellis https://orcid.org/0000-0003-4606-087X
Ryan Ly https://orcid.org/0000-0001-9238-0642
Jeremy Magland https://orcid.org/0000-0002-5286-4375
Samuel R Bray https://orcid.org/0000-0002-9696-6920
Xulu Sun https://orcid.org/0000-0003-3076-0994
Cristofer Holobetz https://orcid.org/0009-0009-8567-3290

Caleb Kemere https://orcid.org/0000-0003-2054-0234
Samuel Garcia https://orcid.org/0000-0001-6389-9779
Alessio Paolo Buccino https://orcid.org/0000-0003-3661-527X
Loren Frank https://orcid.org/0000-0002-1752-5677

Peer review material
Version of Record:
Reviewer #1 (Public review): https://doi.org/10.7554/eLife.108089.3.sa1
Reviewer #2 (Public review): https://doi.org/10.7554/eLife.108089.3.sa2
Author response https://doi.org/10.7554/eLife.108089.3.sa3

Reviewed preprint v2
Reviewer #1 (Public review): https://doi.org/10.7554/eLife.108089.2.sa2
Reviewer #2 (Public review): https://doi.org/10.7554/eLife.108089.2.sa1
Author response: https://doi.org/10.7554/eLife.108089.2.sa0
eLife Assessment https://doi.org/10.7554/eLife.108089.2.sa3

Reviewed preprint v1
Reviewer #1 (Public review): https://doi.org/10.7554/eLife.108089.1.sa1
Reviewer #2 (Public review): https://doi.org/10.7554/eLife.108089.1.sa0
eLife Assessment https://doi.org/10.7554/eLife.108089.1.sa2

Additional files

Supplementary files

MDAR checklist

Data availability

All data used in this work is available at https://dandiarchive.org/dandiset/000937?search=000937&pos=1 and https://dandiarchive.org/dandiset/000059?search=000059&pos=1. All analysis code is available at https://doi.org/10.5281/zenodo.23541211.

The following dataset was generated:

Author(s)	Year	Dataset title	Dataset URL	Database and Identifier
Lee KH, Denovellis EL, Ly R, Magland J, Soules J, Comrie AE, Gramling DP, Guidera JA, Nevers R, Adenekan P	2026	Spyglass Paper Code	https://doi.org/10.5281/zenodo.23541211	Zenodo, 10.5281/zenodo.23541211

The following previously published datasets were used:

Author(s)	Year	Dataset title	Dataset URL	Database and Identifier
Bray S	2024	Spyglass Demo Data	https://dandiarchive.org/dandiset/000937?search=000937&pos=1	DANDI, 000937
Petersen P	2025	Cooling of Medial Septum Reveals Theta Phase Lag Coordination of Hippocampal Cell Assemblies	https://dandiarchive.org/dandiset/000059?search=000059&pos=1	DANDI, 000059

References

Abe T, Kinsella I, Saxena S, Buchanan EK, Couto J, Briggs J, Kitt SL, Glassman R, Zhou J, Paninski L, Cunningham JP. 2022. Neuroscience cloud analysis as a service: an open-source platform for scalable, reproducible data analysis. *Neuron* **110**:2771–2789. DOI: https://doi.org/10.1016/j.neuron.2022.06.018, PMID: 35870448

Acerbi L, Aguillon-Rodriguez V, Ahmadi M, Amjad J, Angelaki D, Arlandis J, Ashwood ZC, Banga K, Barrell H, Bayer HM, Benson J, Benson B, Bhagat J, Birman D, Bougrova K, Boussard J, Bruijns SA, Carandini M, Catarino J, Cazettes F, et al. 2023. A modular architecture for organizing, processing and sharing neurophysiology data. *Nature Methods* **20**:403–407. DOI: https://doi.org/10.1038/s41592-022-01742-6

Baker C, Mayorquin H, Dichter B, Weigl S. 2026 NeuroConv. ea1f362. GitHub. https://github.com/catalystneuro/neuroconv

BRAIN Initiative. 2019. Notice of data sharing policy for the BRAIN initiative. https://grants.nih.gov/grants/guide/notice-files/NOT-MH-19-010.html [Accessed January 22, 2019].

Bray S, Denovellis B, Ly R, Lee KY. 2026. Trodes_to_nwb. 6603412. GitHub. https://github.com/LorenFrankLab/trodes_to_nwb

Broz C. 2024. Spyglass-export-docker. 87efb12. GitHub. https://github.com/LorenFrankLab/spyglass-export-docker

Broz C, Denovellis E, Bray S, Adkisson-Floro P. 2026. Spyglass. da36473. GitHub. https://github.com/LorenFrankLab/spyglass

Buccino AP, Hurwitz CL, Garcia S, Magland J, Siegle JH, Hurwitz R, Hennig MH. 2019. SpikeInterface, a unified framework for spike sorting. *bioRxiv*. DOI: https://doi.org/10.1101/796599

Callaway EM, Dong H-W, Ecker JR, Hawrylycz MJ, Huang ZJ, Lein ES, Ngai J, Osten P, Ren B, Tolias AS, White O, Zeng H, Zhuang X, Ascoli GA, Behrens MM, Chun J, Feng G, Gee JC, Ghosh SS, Halchenko YO, et al. 2021. A multimodal cell census and atlas of the mammalian primary motor cortex. *Nature* **598**:86–102. DOI: https://doi.org/10.1038/s41586-021-03950-0

Chandravadia N, Liang D, Schjetnan AGP, Carlson A, Faraut M, Chung JM, Reed CM, Dichter B, Maoz U, Kalia SK, Valiante TA, Mamelak AN, Rutishauser U. 2020. A NWB-based dataset and processing pipeline of human single-neuron activity during a declarative memory task. *Scientific Data* **7**:78. DOI: https://doi.org/10.1038/s41597-020-0415-9, PMID: 32132545

Chu JP, Kemere CT. 2021. GhostiPy: an efficient signal processing and spectral analysis toolbox for large data. *eNeuro* **8**:ENEURO.0202-21.2021. DOI: https://doi.org/10.1523/ENEURO.0202-21.2021, PMID: 34556557

Committee on Responsible Science, Committee on Science, Engineering, Medicine, and Public Policy, Policy and Global Affairs, National Academies of Sciences, Engineering, and Medicine. 2017. *Fostering Integrity in Research*. National Academies Press. DOI: https://doi.org/10.17226/21896

Denovellis EL, Frank LM, Eden UT. 2019. Characterizing hippocampal replay using hybrid point process state space models. 2019 53rd Asilomar Conference on Signals, Systems, and Computers. 245–249. DOI: https://doi.org/10.1109/IEEECONF44664.2019.9048688

Denovellis EL, Gillespie AK, Coulter ME, Sosa M, Chung JE, Eden UT, Frank LM. 2021. Hippocampal replay of experience at real-world speeds. *eLife* **10**:e64505. DOI: https://doi.org/10.7554/eLife.64505, PMID: 34570699

Denovellis E. 2026. Spyglass-paper. swh:1:rev:392afd56f4bbc93286bcc8dd78af04080b6b7276. Software Heritage. https://archive.softwareheritage.org/swh:1:dir:a20c649589ae7e24f7d86ab749b645e4783f707e;origin=https://github.com/LorenFrankLab/spyglass-paper;visit=swh:1:snp:da9988dbce51f353173e6d36725f88aeb3615d0a;anchor=swh:1:rev:392afd56f4bbc93286bcc8dd78af04080b6b7276

de Vries SEJ, Siegle JH, Koch C. 2023. Sharing neurophysiology data from the allen brain observatory. *eLife* **12**:e85550. DOI: https://doi.org/10.7554/eLife.85550, PMID: 37432073

Foster DJ, Wilson MA. 2007. Hippocampal theta sequences. *Hippocampus* **17**:1093–1099. DOI: https://doi.org/10.1002/hipo.20345, PMID: 17663452

Goble C, Cohen-Boulakia S, Soiland-Reyes S, Garijo D, Gil Y, Crusoe MR, Peters K, Schober D. 2020. FAIR computational Workflows. *Data Intelligence* **2**:108–121. DOI: https://doi.org/10.1162/dint_a_00033

Gutzen R, De Bonis G, De Luca C, Pastorelli E, Capone C, Allegra Mascaro AL, Resta F, Manasanch A, Pavone FS, Sanchez-Vives MV, Mattia M, Grün S, Paolucci PS, Denker M. 2024. A modular and adaptable analysis pipeline to compare slow cerebral rhythms across heterogeneous datasets. *Cell Reports Methods* **4**:100681. DOI: https://doi.org/10.1016/j.crmeth.2023.100681, PMID: 38183979

Halchenko YO, Meyer K, Poldrack B, Solanky DS, Wagner AS, Gors J, MacFarlane D, Pustina D, Sochat V, Ghosh SS, Mönch C, Markiewicz CJ, Waite L, Shlyakhter I, de la Vega A, Hayashi S, Häusler CO, Poline JB, Kadelka T, Skytén K, et al. 2021. DataLad: distributed system for joint management of code, data, and their relationship. *Journal of Open Source Software* **6**:3262. DOI: https://doi.org/10.21105/joss.03262, PMID: 39469147

Hawrylycz M, Martone ME, Ascoli GA, Bjaalie JG, Dong H-W, Ghosh SS, Gillis J, Hertzano R, Haynor DR, Hof PR, Kim Y, Lein E, Liu Y, Miller JA, Mitra PP, Mukamel E, Ng L, Osumi-Sutherland D, Peng H, Ray PL, et al. 2023. A guide to the BRAIN initiative cell census network data ecosystem. *PLOS Biology* **21**:e3002133. DOI: https://doi.org/10.1371/journal.pbio.3002133, PMID: 37390046

Hider R, Kleissas D, Gion T, Xenes D, Matelsky J, Pryor D, Rodriguez L, Johnson EC, Gray-Roncal W, Wester B. 2022. The Brain Observatory Storage Service and Database (BossDB): a cloud-native approach for petascale neuroscience discovery. *Frontiers in Neuroinformatics* **16**:828787. DOI: https://doi.org/10.3389/fninf.2022.828787

Lopes G, Bonacchi N, Frazão J, Neto JP, Atallah BV, Soares S, Moreira L, Matias S, Itskov PM, Correia PA, Medina RE, Calcaterra L, Dreosti E, Paton JJ, Kampff AR. 2015. Bonsai: an event-based framework for processing and controlling data streams. *Frontiers in Neuroinformatics* **9**:7. DOI: https://doi.org/10.3389/fninf.2015.00007, PMID: 25904861

Ly R. 2026. Ndx-franklab-novela. db0742e. GitHub. https://github.com/LorenFrankLab/ndx-franklab-novela

Magland J, Jun JJ, Lovero E, Morley AJ, Hurwitz CL, Buccino AP, Garcia S, Barnett AH. 2020. SpikeForest, reproducible web-facing ground-truth validation of automated neural spike sorters. *eLife* **9**:e55167. DOI: https://doi.org/10.7554/eLife.55167, PMID: 32427564

Magland J, Soules J, Buccino A, Bray S. 2024 Kachery-cloud. c70e97f. GitHub. https://github.com/flatironinstitute/kachery-cloud

Magland J, Soules J, Broz C. 2025. Figurl. bfab9cd. GitHub. https://github.com/flatironinstitute/figurl

Magland J. 2026. Kachery. cd076ed. GitHub. https://github.com/magland/kachery

Matelsky JK, Rodriguez LM, Xenes D, Gion T, Hider R, Wester BA, Gray-Roncal W. 2021. An Integrated Toolkit for Extensible and Reproducible Neuroscience. 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC). 2413–2418. DOI: https://doi.org/10.1109/EMBC46164.2021.9630199

Mölder F, Jablonski KP, Letcher B, Hall MB, van Dyken PC, Tomkins-Tinch CH, Sochat V, Forster J, Vieira FG, Meesters C, Lee S, Twardziok SO, Kanitz A, VanCampen J, Malladi V, Wilm A, Holtgrewe M, Rahmann S, Nahnsen S, Köster J. 2021. Sustainable data analysis with Snakemake. *F1000Research* **10**:33. DOI: https://doi.org/10.12688/f1000research.29032.3, PMID: 34035898

Nath T, Mathis A, Chen AC, Patel A, Bethge M, Mathis MW. 2019. Using DeepLabCut for 3D markerless pose estimation across species and behaviors. *Nature Protocols* **14**:2152–2176. DOI: https://doi.org/10.1038/s41596-019-0176-0, PMID: 31227823

OECD. 2015. *Making Open Science a Reality*. OECD Publishing. DOI: https://doi.org/10.1787/5jrs2f963zs1-en

Pachitariu M, Steinmetz N, Kadir S, Carandini M, Harris KD. 2016. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv*. DOI: https://doi.org/10.1101/061481

Park J, Wang J, Guan W, Gjesteby LA, Pollack D, Kamentsky L, Evans NB, Stirman J, Gu X, Zhao C, Marx S, Kim ME, Choi SW, Snyder M, Chavez D, Su-Arcaro C, Tian Y, Park CS, Zhang Q, Yun DH, et al. 2024. Integrated platform for multiscale molecular imaging and phenotyping of the human brain. *Science* **384**:eadh9979. DOI: https://doi.org/10.1126/science.adh9979, PMID: 38870291

Petersen PC, Buzsáki G. 2020. Cooling of medial septum reveals theta phase lag coordination of hippocampal cell assemblies. *Neuron* **107**:731–744. DOI: https://doi.org/10.1016/j.neuron.2020.05.023, PMID: 32526196

Pierré A, Pham T, Pearl J, Datta SR, Ritt JT, Fleischmann A. 2024. A perspective on neuroscience data standardization with Neurodata Without Borders. *arXiv*. https://doi.org/10.48550/arXiv.2310.04317

Rübel O, Tritt A, Dichter B, Braun T, Cain N, Clack N, Davidson TJ, Dougherty M, Fillion-Robin JC, Graddis N. 2019. NWB:N 2.0: an accessible data standard for neurophysiology. *bioRxiv*. DOI: https://doi.org/10.1101/523035

Rübel O, Tritt A, Ly R, Dichter BK, Ghosh S, Niu L, Baker P, Soltesz I, Ng L, Svoboda K, Frank L, Bouchard KE. 2022. The Neurodata Without Borders ecosystem for neurophysiological data science. *eLife* **11**:e78362. DOI: https://doi.org/10.7554/eLife.78362, PMID: 36193886

Sanchez M, Moore D, Johnson EC, Wester B, Lichtman JW, Gray-Roncal W. 2022. Connectomics annotation metadata standardization for increased accessibility and queryability. *Frontiers in Neuroinformatics* **16**:828458. DOI: https://doi.org/10.3389/fninf.2022.828458, PMID: 35651719

Siegle JH, López AC, Patel YA, Abramov K, Ohayon S, Voigts J. 2017. Open Ephys: an open-source, plugin-based platform for multichannel electrophysiology. *Journal of Neural Engineering* **14**:045003. DOI: https://doi.org/10.1088/1741-2552/aa5eea, PMID: 28169219

Skaggs WE, McNaughton BL, Wilson MA, Barnes CA. 1996. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**:149–172. DOI: https://doi.org/10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6>3.0.CO;2-K

Teeters JL, Godfrey K, Young R, Dang C, Friedsam C, Wark B, Asari H, Peron S, Li N, Peyrache A, Denisov G, Siegle JH, Olsen SR, Martin C, Chun M, Tripathy S, Blanche TJ, Harris K, Buzsáki G, Koch C, et al. 2015. Neurodata without borders: creating a common data format for neurophysiology. *Neuron* **88**:629–634. DOI: https://doi.org/10.1016/j.neuron.2015.10.025, PMID: 26590340

Viejo G, Levenstein D, Skromne Carrasco S, Mehrotra D, Mahallati S, Vite GR, Denny H, Sjulson L, Battaglia FP, Peyrache A. 2023. Pynapple, a toolbox for data analysis in neuroscience. *eLife* **12**:RP85786. DOI: https://doi.org/10.7554/eLife.85786, PMID: 37843985

Weinreb C, Pearl JE, Lin S, Osman MAM, Zhang L, Annapragada S, Conlin E, Hoffmann R, Makowska S, Gillis WF, Jay M, Ye S, Mathis A, Mathis MW, Pereira T, Linderman SW, Datta SR. 2024. Keypoint-MoSeq: parsing behavior by linking point tracking to pose dynamics. *Nature Methods* **21**:1329–1339. DOI: https://doi.org/10.1038/s41592-024-02318-2, PMID: 38997595

Wilkinson MD, Dumontier M, Aalbersberg IJJ, Appleton G, Axton M, Baak A, Blomberg N, Boiten J-W, da Silva Santos LB, Bourne PE, Bouwman J, Brookes AJ, Clark T, Crosas M, Dillo I, Dumon O, Edmunds S, Evelo CT, Finkers R, Gonzalez-Beltran A, et al. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data* **3**:160018. DOI: https://doi.org/10.1038/sdata.2016.18, PMID: 26978244

Yatsenko D, Reimer J, Ecker AS, Walker EY, Sinz F, Berens P, Hoenselaar A, Cotton RJ, Siapas AS, Tolias AS. 2015. DataJoint: managing big scientific data using MATLAB or Python. *bioRxiv*. DOI: https://doi.org/10.1101/031658

Yatsenko D, Walker EY, Tolias AS. 2018. DataJoint: a simpler relational data model. *arXiv*. https://doi.org/10.48550/arXiv.1807.11104

Yatsenko D, Nguyen T, Shen S, Gunalan K, Turner CA, Guzman R, Sasaki M, Sitonic D, Reimer J, Walker EY, Tolias A. 2021. DataJoint Elements: Data Workflows for Neurophysiology. [bioRxiv]. DOI: https://doi.org/10.1101/2021.03.30.437358