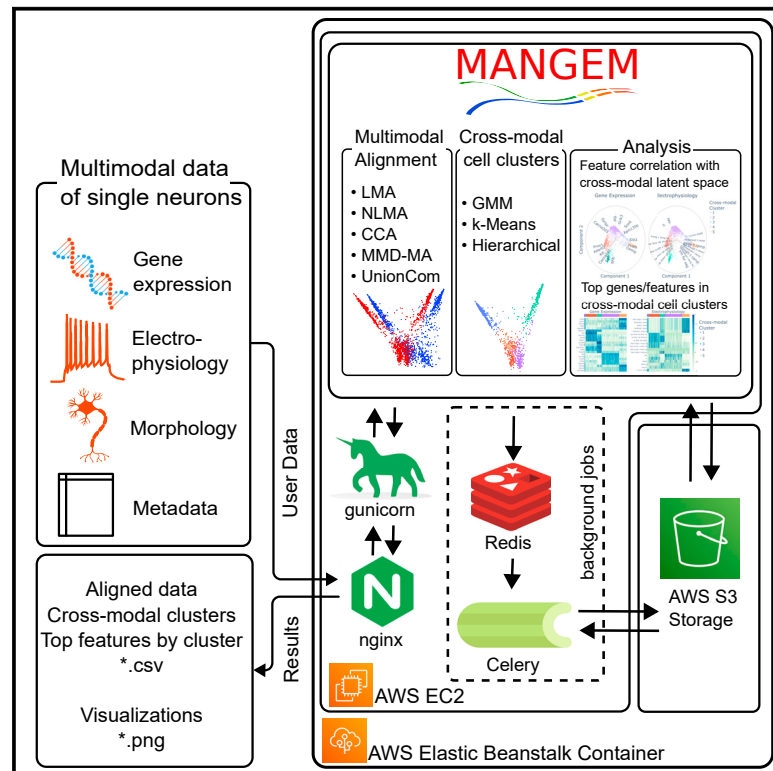


Patterns

MANGEM: A web app for multimodal analysis of neuronal gene expression, electrophysiology, and morphology

Graphical abstract



Authors

Robert Hermod Olson,
Noah Cohen Kalafut, Daifeng Wang

Correspondence

daifeng.wang@wisc.edu

In brief

Multimodal single-cell datasets, such as those from Patch-seq, can be used to identify neuronal cell types with similar characteristics, offering insight into cellular functions. Integration of multiple modalities of single-cell data is aided by machine learning methods, but they are often difficult to use. The authors present MANGEM, a web app including visualization tools that provides an easy-to-use interface for researchers to upload their multimodal data of neuronal cells, select machine learning methods for multimodal alignment, identify multimodal cell clusters, and reveal cross-modal relationships.

Highlights

- MANGEM enables single-cell multimodal learning and visualization in a cloud-based app
- Application to Patch-seq data identifies multimodal functions of neuronal cells
- Visualizations reveal cross-modal relationships of neurons
- Supports asynchronous learning and background job running for large-scale data analyses



Descriptor

MANGEM: A web app for multimodal analysis of neuronal gene expression, electrophysiology, and morphology

Robert Hermod Olson,¹ Noah Cohen Kalafut,^{1,2} and Daifeng Wang^{1,2,3,4,*}¹Waisman Center, University of Wisconsin-Madison, Madison, WI 53705, USA²Department of Computer Sciences, University of Wisconsin-Madison, Madison, WI 53706, USA³Department of Biostatistics and Medical Informatics, University of Wisconsin-Madison, Madison, WI 53706, USA⁴Lead contact*Correspondence: daifeng.wang@wisc.edu<https://doi.org/10.1016/j.patter.2023.100847>

THE BIGGER PICTURE Recently, it has become possible to obtain multiple types of data (modalities) from individual neurons, like how genes are used (gene expression), how a neuron responds to electrical signals (electrophysiology), and what it looks like (morphology). These datasets can be used to group similar neurons together and learn their functions, but the complexity of the data can make this process difficult for researchers without sufficient computational skills. Various methods have been developed specifically for combining these modalities, and open-source software tools can alleviate the computational burden on biologists performing analyses of new data. Open-source tools performing modality combination (integration), clustering, and visualization have the potential to streamline the research process. It is our hope that intuitive and freely available software will advance neuroscience research by making advanced computational methods and visualizations more accessible.



Production: Data science output is validated, understood, and regularly used for multiple domains/platforms

SUMMARY

Single-cell techniques like Patch-seq have enabled the acquisition of multimodal data from individual neuronal cells, offering systematic insights into neuronal functions. However, these data can be heterogeneous and noisy. To address this, machine learning methods have been used to align cells from different modalities onto a low-dimensional latent space, revealing multimodal cell clusters. The use of those methods can be challenging without computational expertise or suitable computing infrastructure for computationally expensive methods. To address this, we developed a cloud-based web application, MANGEM (multimodal analysis of neuronal gene expression, electrophysiology, and morphology). MANGEM provides a step-by-step accessible and user-friendly interface to machine learning alignment methods of neuronal multimodal data. It can run asynchronously for large-scale data alignment, provide users with various downstream analyses of aligned cells, and visualize the analytic results. We demonstrated the usage of MANGEM by aligning multimodal data of neuronal cells in the mouse visual cortex.

INTRODUCTION

The human brain has approximately 86 billion neurons encompassing a vast range of different functions. Understanding the roles of individual neurons is a daunting challenge that is beginning to become possible with new techniques and technologies. The development of single-cell technologies such as Patch-seq

has resulted in the ability to characterize neurons with new specificity and detail. Patch-seq enables a researcher to simultaneously obtain measures of gene expression, electrophysiology, and morphology of individual neurons. Gene expression is a measure of the extent to which different genes in a cell's DNA are transcribed to RNA and then translated to produce proteins. Electrophysiology describes the electrical behavior of a cell.



A microscopic pipette containing an electrolyte contacts the cell membrane to establish an electrical connection. Then the cell's electrical response to an applied voltage or current is measured. Morphology refers to the physical structure of a neuron, including the size and shape of the cell's axon and dendrites. Patch-seq has a wide variety of applications across different cell and tissue types¹; additionally, the modalities included are quite diverse and mesh with MANGEM's (multimodal analysis of neuronal gene expression, electrophysiology, and morphology) goal of generality. By combining microscopy, RNA sequencing, and electrophysiological recording for individual neurons, multimodal datasets can be developed with the potential to reveal relationships between neuronal function, structure, and gene expression.² Multimodal single-cell datasets are increasingly available to researchers, in part due to efforts by the Brain Research through Advancing Innovative Neurotechnologies (BRAIN) Initiative to support the development and storage of such datasets in freely accessible repositories such as the Neuroscience Multi-Omic Archive (<https://nemoarchive.org/>) for genomic data and Distributed Archives for Neurophysiology Data Integration³ (<https://dandiarchive.org/>) for neurophysiology data, including electrophysiology.

While multimodal single-cell data offers great potential for improving understanding of brain organization and function, new methods are required for integration and analysis of the data.⁴ Because cells with similar characteristics in one modality are not necessarily similar when measured by another, identification of cell clusters must incorporate disparate data types simultaneously. Machine learning methods such as manifold learning are highly applicable to the problems posed by heterogeneity of multimodal single-cell data,^{5–7} but these methods are commonly difficult to use, especially for biologists and neurologists who may not have computational expertise. Documentation and tutorials, if present, are limited in scope. The methods are often supplied as source code only, requiring coding expertise to use, which further limits their accessibility. Installation and configuration of the software adds another layer of difficulty to overcome before these methods can be applied. As an example, consider the software for UnionCom.⁸ While the UnionCom software is available in the Python package index and easily installable, its dependencies are not automatically installed. The prospective user will quickly discover that the versions of those dependencies suggested in the limited documentation are not easily installable in recent versions of Python. Given time and effort, a motivated researcher will manage to find the right combination of package versions and Python version that will be compatible, but this level of difficulty is both a significant barrier to use and common in open-source scientific software generally.⁹

An increasingly common way to address the challenges of running open-source scientific software is by implementing the methods of the software in a web application.^{10,11} Here, we present a web application named MANGEM, developed to address the challenges researchers may experience in using existing methods of aligning and analyzing multimodal single-cell data. In particular, MANGEM (1) provides an easy-to-use interface to a variety of machine learning alignment methods, (2) requires no coding to use, and (3) does not require installation of software or management of computing infrastructure.

Preloaded datasets and an interface that walks the user through each operational step provide for an accessible introduction to the use of machine learning methods to align multimodal datasets. As a cloud-based web application, MANGEM enables users to begin exploring multimodal single-cell datasets without first undertaking the challenges of software installation or management of the underlying infrastructure. While the application was designed for real-time data processing and exploration, it also supports running certain long-running methods asynchronously, providing a customized URL for users to retrieve results after computation is complete. Interactive graphical display of output facilitates exploration of the data at each step of the analysis process: raw data as uploaded, preprocessed data (e.g., standardized), aligned datasets, and cross-modal clusters. Integrated downstream analysis methods support identification of important cellular features within cross-modal cell clusters and aid interpretation of the revealed relationships within cell clusters.

RESULTS

The MANGEM web application offers a range of methods for aligning multimodal data of neuronal cells, identifying cross-modal cell clusters using the aligned data, and generating visualizations to facilitate the characterization of these cross-modal clusters, including their differentially expressed genes and correlated multimodal features (Figure 1).

The application is implemented using Plotly Dash Open Source, a Python-based framework for developing data science applications.¹² Dash is based on Plotly.js (<https://plot.ly>), React (<https://react.dev/>), and Flask,¹³ and it functions by tying user interface elements to stateless callback functions. In the case of MANGEM, some callback functions are quasi-stateless, in that uploaded and aligned datasets are stored in a file system cache to avoid repeating lengthy calculations.

Our public deployment of MANGEM is on Amazon Web Services infrastructure (Figure 2). The Elastic Beanstalk service is used to deploy the application to an Elastic Compute Cloud (EC2) instance with associated storage in Amazon Simple Storage Service (S3). In order to be accessed by a user with a web browser, MANGEM requires additional software. A reverse proxy server directs the requests from the web browser to an application server that can translate the requests to the Web Server Gateway Interface (WSGI) protocol used for communication with MANGEM. By default, the Elastic Beanstalk Python platform provides nginx¹⁴ as the reverse proxy server and Gunicorn (<https://docs.gunicorn.org/en/stable/>) as the WSGI application server; however, MANGEM does not depend on those specific programs. For example, in our development environment, we use the Apache HTTP server with mod_proxy (https://httpd.apache.org/docs/2.4/en/mod/mod_proxy.html) as the reverse proxy server and uWSGI (<https://uwsgi-docs.readthedocs.io/en/latest/>) as the WSGI application server. Most data processing occurs within the main MANGEM process, but additional software is required to enable long-running alignment jobs to run asynchronously. In this case, Celery (<https://docs.celeryq.dev/en/stable>) is used to run those background jobs, and Redis (<https://redis.io/>) is used as a message broker to communicate between MANGEM and Celery. Whether

Multimodal Analysis of Neuronal Gene Expression, Electrophysiology and Morphology

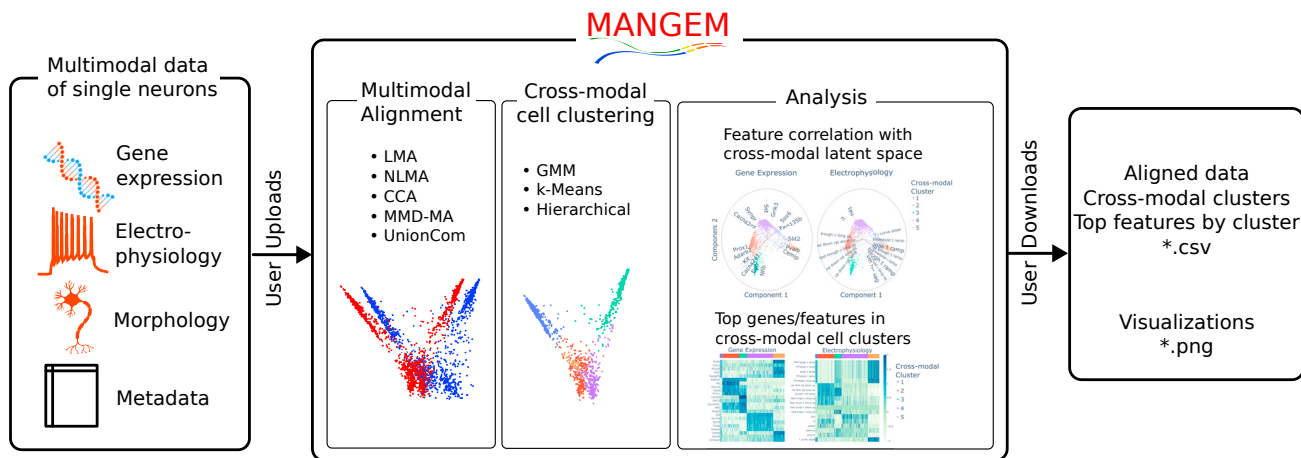


Figure 1. Overview of MANGEM (multimodal analysis of neuronal gene expression, electrophysiology, and morphology)

User input to MANGEM includes multimodal single-cell data together with cell metadata. Within MANGEM, the multimodal data are aligned using machine learning methods, projecting disparate modalities into a low-dimensional common latent space. Clustering algorithms are applied within the latent space to identify cell clusters, and then analysis methods are provided in MANGEM to characterize the clusters by differential feature expression and correlation of features with the latent space. In addition to interactive plots generated at each step of the workflow, downloadable output includes tabular data files (cell coordinates in latent space, cluster annotations, top features for each cluster) and images depicting alignment, cross-modal cell clusters, and cluster analyses.

aligned synchronously or asynchronously, aligned multimodal datasets are stored in a file system cache on AWS S3.

A major contribution of this work is the development of an easily usable and accessible interface to a fully integrated workflow running on third-party infrastructure. An alternative approach would be to develop and distribute a containerized workflow, but we have prioritized a web-based approach that relieves the user of infrastructure management concerns, maximizing accessibility to all users.

MANGEM’s layout is organized as a set of tabs on the left that contain user interface controls, while the right side contains plots or other information related to the active tab. The tabs correspond to the sequence of steps users will typically take when running the application: upload data, align data, identify cross-modal cell clusters, and perform downstream analysis of cross-modal cell clusters. Each tab contains controls that allow the user to adjust parameters relevant to the current step of the workflow and that influence the downstream results (Table 1).

At each step of the workflow (Figure 3), interactive figures are automatically generated to support understanding, and computation products are available for download as tabular data files. A video demonstration of the workflow is provided in Video S1, and the user interface is depicted in Figure S1.

In this study, we showcased the usage of MANGEM through two case studies that utilized emerging Patch-seq multimodal data of inhibitory neuronal cells in the mouse visual cortex (such as gene expression, electrophysiology, and morphology). It is worth noting that MANGEM is a general-purpose tool that can be used for any user multimodal data of neurons.

Case study 1: Neuronal gene expression and electrophysiology

We first tested MANGEM to align these neuronal cells based on gene expression and electrophysiological features. We up-

loaded two datasets, one containing the 1,302 most variable expressed genes and one containing 41 electrophysiological features for 3,654 neuronal cells, on the Upload Data tab of MANGEM. We then preprocessed the data using log transformation for gene expression and standardization for electrophysiology features.

On the Alignment tab, we set the alignment method to nonlinear manifold alignment (NLMA), the number of latent space dimensions to 5, and the number of nearest neighbors (used in construction of the similarity matrix for NLMA) to 2. Clicking the “Align Datasets” button generated two measures of alignment along with a 3D plot of the aligned cells (Figure 4A). The aligned multimodal cells were represented in the common latent space, \tilde{X} and \tilde{Y} , are 3,654 cells (rows) by 5 latent dimensions (columns). We also tested other alignment methods and found that NLMA, in addition to being one of the fastest methods to run, resulted in the smallest alignment error (Figure S2; Table S1).

Afterward, we chose to use the Gaussian Mixture Model clustering algorithm on the Clustering tab, specifying 5 clusters. Upon clicking the “Identify Cross-Modal Cell Clusters” button, the algorithm identified cross-modal cell clusters and generated side-by-side plots of the aligned cells for each modality in the latent space. These plots showed cells colored according to their respective cross-modal clusters (Figure 4B).

The Analysis tab of MANGEM offers various visualization methods for exploring cross-modal relationships between gene expression and electrophysiological features. We set the number of top features to 5 and selected the “Features of Cross-modal Clusters (Heatmap)” (Figure 4C). The resulting heatmap showed that tau and ri were the top two electrophysiological features in Cluster 4, while the top differentially expressed genes in the cluster were *Sst*, *Grin3a*, *Grik3*, *Trhde*, and *Stxbp6*.

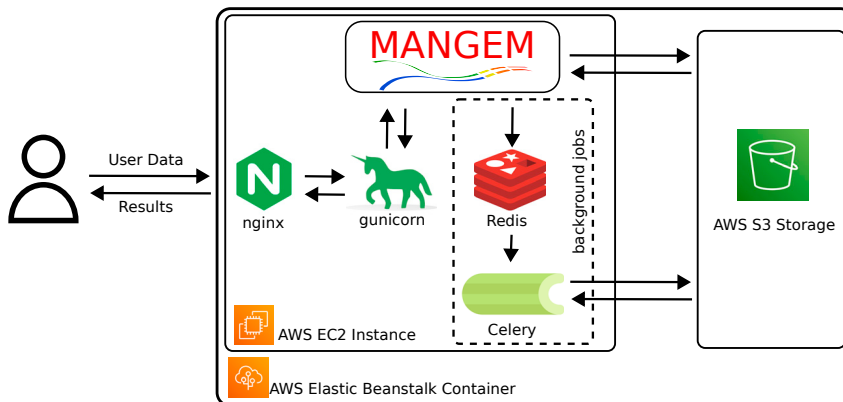


Figure 2. Cloud implementation of MANGEM using AWS infrastructure

The application runs on Amazon Cloud Services using Elastic Beanstalk to provision an EC2 instance. The web server nginx serves as a reverse proxy to the Gunicorn wsgi server. MANGEM is written in Python using the Plotly Dash framework. Long-running tasks are run in the background by Celery workers, with Redis acting as the message broker between MANGEM and Celery. Uploaded and processed data files are stored in a file system cache in AWS S3.

These shared multimodal features suggest potential functional linkages among the cells in the cluster.

To further investigate these linkages, we switched the plot type to “Top Feature Correlation with Latent Space (Biplot)” (Figure 4D) and set the number of top correlated features to 15. The biplots graphically represented the most highly correlated features from cross-modal cell clusters and allowed for interactive zooming into the Cluster 4 area on the latent space. The highly correlated features included tau and, to a lesser extent, ri among the electrophysiological features, while *Sst* and *Grik3* were among the genes associated with Cluster 4.

Case study 2: Neuronal gene expression and morphology

MANGEM was used to process gene expression of the top 1,000 variable genes and morphological features in the mouse motor

cortex.¹⁵ The data consists of 646 single cells with 42,466 genes and 63 morphological features. Each modality is formatted into a separate csv, with an additional file indicating metadata such as age, gender, etc. The data were then uploaded onto the webapp using the upload tab.

MANGEM can be used to easily test multiple integration methods. For this application, we chose NLMA. After alignment, pairwise accuracy statistics are reported (Figure 5A).

MANGEM is then used to separate the data into 5 clusters using a Gaussian mixture model. The clusters closely align with true cell types (Figure 5B). Then, differentially expressed features for each cluster may be downloaded and used for downstream analysis. The expressed genes can then be analyzed for importance in brain function.

MANGEM identifies *Pvalb*, *Vip*, *Lamp5*, and *Sst* among the top 2 most differentially expressed genes over the 5 cell clusters (Figure 5C). These genes are commonly used to identify cell type.¹⁶ So, MANGEM can be used to automatically perform cell-type clustering on multimodal datasets. In addition, MANGEM identifies *Adarb2* as a differentially expressed gene. *Adarb2* has been found to distinguish between two major branches of inhibitory neurons.¹⁷

MANGEM also allows users to create Biplots to visualize features important to the latent space (Figure 5D). These features that are highly correlated with the latent space (e.g., *SOX6* and *sp_width*) may then be the focus of future data exploration.

DISCUSSION

MANGEM is a user-friendly web application designed primarily for biologists and neuroscientists. The app comes with pre-selected general-purpose hyperparameters that can be fine-tuned by users to suit their needs. With the rapid advancements in multimodal machine learning,¹⁸ MANGEM is constantly evolving to offer more advanced alignment options.

Although the name and pre-packaged data of MANGEM are specific to gene expression, electrophysiology, and morphology of brain cells, the methods currently implemented impose no restrictions on the modalities used nor the source (e.g., brain, kidney) of the data. Generality is a key component of the usefulness of MANGEM. In the future, we plan to implement a wide variety of integration methods including joint variational autoencoders for multimodal imputation and embedding¹⁹ and cross-modal

Table 1. Key parameters of data processing and analysis in MANGEM

Tab	Parameter
Step 1. Upload Data	<ul style="list-style-type: none"> Preprocessing method for each modality (log/standardize)
Step 2. Alignment	<ul style="list-style-type: none"> Alignment algorithm Dimension of machine learning latent space (3–10) Number of nearest neighbors (1–10) used in constructing similarity matrices (LMA, NLMA). Number of iterations (MMD-MA)
Step 3. Clustering	<ul style="list-style-type: none"> Clustering algorithm Number of cell clusters (1–10)
Step 4. Analysis	<ul style="list-style-type: none"> Component selection Number of top features per cell cluster in “Features of cross-modal clusters” Number of top correlated features in “Top feature correlation with latent space”

The listed parameters all influence downstream output of MANGEM. For example, selecting a preprocessing method on the Upload Data tab will result in that method being applied to the uploaded dataset before the selected multimodal alignment method is applied.

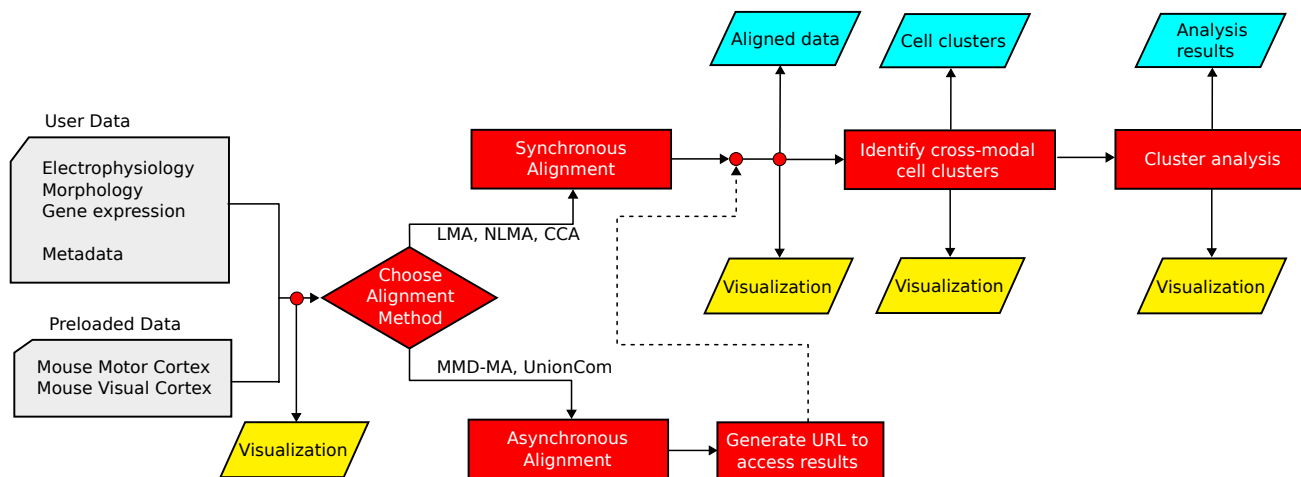


Figure 3. Data flow through MANGEM web application

Input data passes into an alignment process, which will either run in the main process or in the background, depending on the method. In the case of background (asynchronous) alignment, a URL will be supplied to the user, which will allow them to check on the job's status and access the results upon completion. Aligned data feed into a clustering algorithm, and then data analysis methods can be applied to the cell clusters. Data visualization output can be produced at each stage of the process, and tabular data files of aligned data, cell clusters, and analysis results can be downloaded.

optimal transport for multimodal inference.²⁰ Additionally, future versions may incorporate methods like deep neural networks to work with raw data (e.g., electrophysiological time-series data) or other types of data, such as genomics, epigenomics, or images.

Clustering capability in MANGEM will be enhanced by the addition of consensus clustering.²¹ This approach combines any number of clustering methods using iterative voting consensus. Such an approach even allows for merging multiple cluster configurations with differing numbers of clusters.

MANGEM is currently designed to work solely in an interactive mode, but in the future, it could be extended to support API access. One approach would enable data upload and selection of preprocessing and alignment methods through the API, and MANGEM then would return a URL where results could be viewed, just as in MANGEM's existing asynchronous mode.

MANGEM uses cloud-based computing, which in the future will enable distributed training, making computation faster and providing a smoother experience for users. To further improve the efficiency of the app, MANGEM can be optimized for parallel processing, allowing it to take advantage of multiple processors and GPUs for faster computation. In addition to its alignment capabilities, MANGEM also enables collaborative work and data sharing. The app provides a centralized repository for storing and sharing aligned data, with built-in privacy and security measures to protect sensitive data.

Instructions for deployment on AWS are included with the source code. MANGEM was developed by Waisman Center core staff, who will remain available for ongoing support.

EXPERIMENTAL PROCEDURES

Resource availability

Lead contact

Further information may be requested from Daifeng Wang (daifeng.wang@wisc.edu).

Materials availability

MANGEM is freely available for use at <https://ctc.waisman.wisc.edu/mangem>.

Data and code availability

- The data included in this study are included with the source code, which is deposited at Zenodo²² and is publicly available as of the date of publication.
- MANGEM is a web app accessible at <https://ctc.waisman.wisc.edu/mangem> and has open source code available at <https://github.com/daifengwanglab/mangem>.
- Any additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

Multimodal single-cell data analysis using MANGEM

MANGEM is designed to guide users through a multimodal analysis pipeline, step by step. These steps, described in detail in the paragraphs that follow, are

- **Step 1** Upload data. Users have the option of selecting preloaded sets of single-cell multimodal data or uploading their own data files, formatted as comma-delimited text.
- **Step 2** Multimodal alignment. One of several methods is selected that will project the data from each modality to a common latent space. After alignment, metrics of alignment quality are displayed alongside a 3D plot of the aligned data.
- **Step 3** Cross-modal cell clustering. A user-selected clustering method identifies clusters of cells.
- **Step 4** Analysis of cross-modal cell clusters. Integrated analysis methods help users to characterize the cell clusters identified in cross-modal cell clustering.

Preloaded datasets

Two main data sources are utilized for MANGEM's pre-packaged datasets. First, Patch-seq data of 3,654 single-cells from the mouse visual cortex²³ was used to generate electrophysiological and morphological features.⁵ MANGEM provides a preset combining gene expression (1,302 genes) and electrophysiological features (41 features) on this dataset. Second, 1,208 single-cells from the mouse motor cortex¹⁵ were processed similarly (1,286 genes, 29 features). MANGEM provides two presets from this dataset, each containing gene expression and one of electrophysiological or morphological features. The latter preset includes a subset of 646 cells with 1,000 genes and 61 morphological features.

Step 1 Upload data

The first data processing step in MANGEM is selecting or uploading neuronal data, accomplished on the Upload Data tab. The expected input to MANGEM

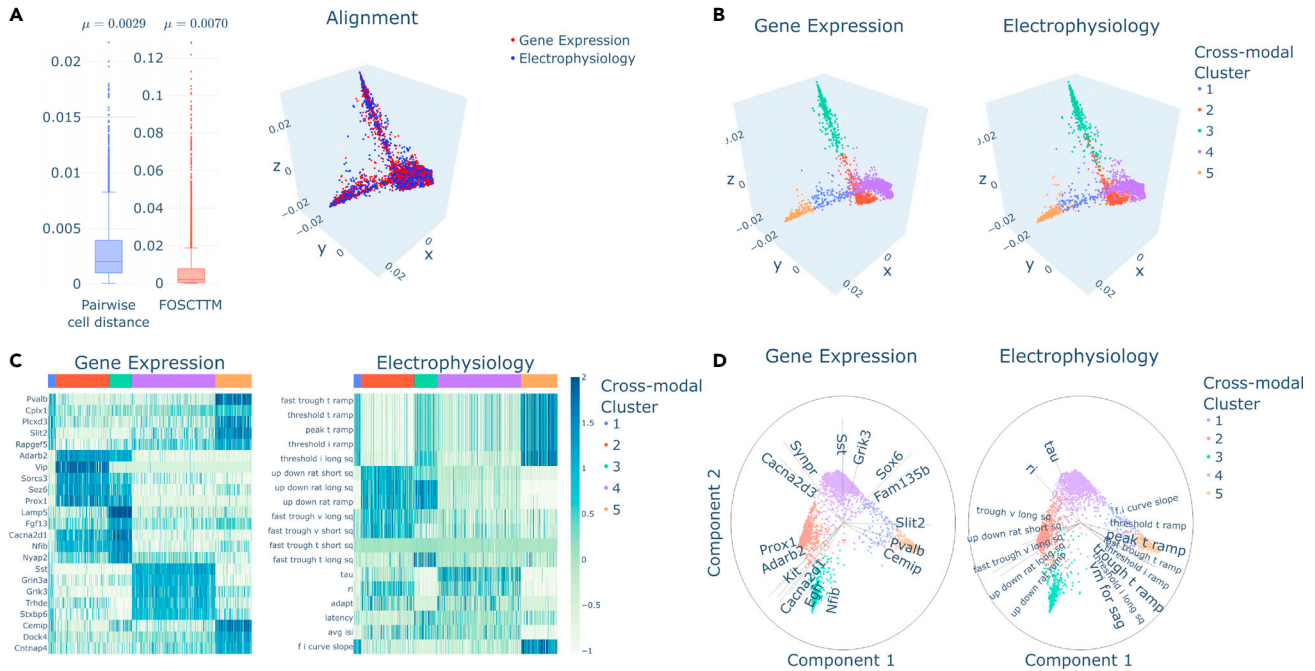


Figure 4. MANGEM analysis and visualization of neuronal gene expression and electrophysiological features in mouse visual cortex

(A) Measures of alignment error and 3D plot of superimposed aligned data in latent space are shown for the preloaded mouse visual cortex dataset after nonlinear manifold alignment. Central boxes range from the first to third quartiles, containing a tick mark for the median. The whiskers range to the farthest datapoint that falls within 1.5 times the interquartile range.

(B) Cross-modal clusters, obtained by Gaussian mixture model, are indicated by color in plots of aligned data for each modality.

(C) Feature levels across all cells for the top 5 features for each cross-modal cluster. Normalized feature magnitude was ranked using the Wilcoxon rank-sum test. Cross-modal clusters are identified by the colored bar at the top of each plot.

(D) Biplots for Gene Expression and Electrophysiological features using dimensions 1 and 2 of the latent space. The top 15 features by correlation with the latent space are shown plotted as radial lines where the length is the value of correlation (max value 1).

consists of three data files in .csv format: one file for each of two modalities and a third file of cellular metadata. Three sample datasets are preloaded in MANGEM, and links are provided within the application to download two of these. The first column of each file should contain a cell identifier, and the files are expected to have a consistent cell order.

Denote data for the first modality as X , data for the second modality, Y , and metadata, M . Each of these has n rows corresponding to n neuronal cells. X and Y have d_1 and d_2 features, respectively. Metadata matrix M has d_m cell characteristics.

$$X = [x_1, x_2, \dots, x_n]^T \in \mathbb{R}^{n \times d_1}, Y = [y_1, y_2, \dots, y_n]^T \in \mathbb{R}^{n \times d_2},$$

$$M = [m_1, m_2, \dots, m_n]^T \in \mathbb{R}^{n \times d_m}$$

When user data are uploaded, a label may be supplied for each modality; otherwise, the modalities will be identified using the default labels of “Modality 1” and “Modality 2” in plot legends.

A preprocessing operation may optionally be selected for each modality. Choices include “Log transform” and “Standardize,” which for the first modality would be

$$\text{“Log transform”}: f(X) \rightarrow \hat{X} \text{ where } \hat{x}_i^j = \log_2 x_i^j, \forall i \in [1, n], j \in [1, d_1]$$

$$\text{“Standardize”}: f(X) \rightarrow \hat{X} \text{ where } \hat{x}_i^j = \frac{x_i^j - \mu^j}{\sigma^j}, \mu^j = \frac{1}{n} \sum_{k=1}^n x_k^j, \sigma^j = \sqrt{\left(\frac{1}{n} \sum_{k=1}^n x_k^j - \mu^j \right)^2} / n$$

If a preprocessing operation is selected, that operation will be applied to the appropriate dataset prior to alignment. The default values of “Log transform” for modality 1 and “Standardize” for modality 2 are suitable for the preloaded datasets, where modality 1 is Gene Expression and modality 2 is Electrophysiology.

Data exploration

The “Explore Data” section of the Upload Data tab can generate plots to gain insight into cell features in the uploaded or selected datasets. A series of box-plots is generated for each value of a categorical metadata variable when a single cell feature is selected (Figure S3A). A particular value of that metadata variable may be selected to filter the data, in which case a violin plot is generated (Figure S3B). It is also possible to select two features to compare in a scatterplot (Figure S3C). These features could be from the same or different modalities. Similar to the single-feature case, selecting a specific value of a metadata variable filters the data so that only the points corresponding to cells having that metadata value are displayed in the scatterplot (Figure S3D). As with all plots in MANGEM, a toolbar will pop up when the cursor is placed over the plot. The toolbar has buttons to change the plot appearance (zoom or pan, for example) and also has a button with a camera icon that causes an image of the plot to be downloaded.

Step 2 Multimodal alignment

The approach used by MANGEM to find clusters of related cells is to first transform the measured cellular features into a latent space where cells having similar features are closer together. This transformation process is called multimodal alignment, and several alignment methods are implemented in MANGEM. Currently supported methods include linear manifold alignment (LMA), NLMA,²⁴ canonical correlation analysis, manifold alignment with maximum mean discrepancy (MMD-MA),²⁵ and unsupervised topological alignment for single-cell multi-omics (UnionCom).⁶ LMA and NLMA utilize similarity matrices to formulate a common latent space. MMD-MA minimizes an objective function that measures distortion and preserved representation. UnionCom infers cross-modal correspondence information before using t-SNE²⁶ to provide the final latent spaces.

of features identified for each cluster is specified using the “Number of Top Features per cluster” control on the Analysis tab. A list of the most important features can be downloaded using the “Download Top Features” button.

Top feature correlation with latent space

Top feature correlation with latent space creates a biplot (i.e., collection of biplots)¹⁵ with one biplot for each modality. Each biplot displays a 2-dimensional projection of the aligned data for the modality in the latent space while overlaying lines corresponding to the features that are most highly correlated with the latent space representation. For each modality, the correlation is computed between the original cellular data for each feature and the projection of the cellular data into the latent space dimensions selected as components X and Y on the Analysis tab. For a given feature, the correlations between that feature and its X and Y latent space representation determine the coordinates of the endpoint of that feature’s line.

The latent space dimensions in which aligned data are plotted can be selected using the “Component Selection” controls on the Analysis tab. At most, three dimensions can be plotted at one time within MANGEM, but these controls allow the user to select which dimensions are plotted to gain different perspectives on the data. Additional controls on the Analysis tab allow aligned data points to be colored either by cluster or by metadata value (for example, transcriptomic cell type).

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.patter.2023.100847>.

ACKNOWLEDGMENTS

This work was supported by National Institutes of Health grants (RF1MH128695 and R01AG067025 to D.W. and P50HD105353 to Waisman Center) and the start-up funding for D.W. from the Office of the Vice Chancellor for Research and Graduate Education at the University of Wisconsin-Madison. The funders had no role in study design, data collection and analysis, decision to publish, or manuscript preparation. The authors wish to thank all members of the Wang lab for insightful discussions on the work.

AUTHOR CONTRIBUTIONS

D.W. conceived the study. D.W., R.O., and N.K. designed the methodology and performed analysis and visualization. R.O. implemented the software. D.W., R.O., and N.K. edited and wrote the manuscript. All authors read and approved the final manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

INCLUSION AND DIVERSITY

We support inclusive, diverse, and equitable conduct of research.

Received: June 16, 2023

Revised: August 7, 2023

Accepted: September 1, 2023

Published: September 25, 2023

REFERENCES

- Lipovsek, M., Bardy, C., Cadwell, C.R., Hadley, K., Kobak, D., and Tripathy, S.J. (2021). Patch-seq: Past, Present, and Future. *J. Neurosci.* *41*, 937–946. <https://doi.org/10.1523/JNEUROSCI.1653-20.2020>.
- Marx, V. (2022). Patch-seq takes neuroscience to a multimodal place. *Nat. Methods* *19*, 1340–1344. <https://doi.org/10.1038/s41592-022-01662-5>.
- Rübel, O., Tritt, A., Ly, R., Dichter, B.K., Ghosh, S., Niu, L., Baker, P., Soltesz, I., Ng, L., Svoboda, K., et al. (2022). The Neurodata Without Borders ecosystem for neurophysiological data science. *Elife* *11*, e78362. <https://doi.org/10.7554/eLife.78362>.
- Efremova, M., and Teichmann, S.A. (2020). Computational methods for single-cell omics across modalities. *Nat. Methods* *17*, 14–17. <https://doi.org/10.1038/s41592-019-0692-4>.
- Huang, J., Sheng, J., and Wang, D. (2021). Manifold learning analysis suggests strategies to align single-cell multimodal data of neuronal electrophysiology and transcriptomics. *Commun. Biol.* *4*, 1308. <https://doi.org/10.1038/s42003-021-02807-6>.
- Gala, R., Budzillo, A., Baftizadeh, F., Miller, J., Gouwens, N., Arkipov, A., Murphy, G., Tasic, B., Zeng, H., Hawrylycz, M., and Sümbül, U. (2021). Consistent cross-modal identification of cortical neurons with coupled autoencoders. *Nat. Comput. Sci.* *1*, 120–127. <https://doi.org/10.1038/s43588-021-00030-1>.
- Gala, R., Gouwens, N., Yao, Z., Budzillo, A., Penn, O., Tasic, B., Murphy, G., Zeng, H., and Sümbül, U. (2019). A coupled autoencoder approach for multi-modal analysis of cell types. In *Advances in Neural Information Processing Systems* (Curran Associates, Inc.).
- Cao, K., Bai, X., Hong, Y., and Wan, L. (2020). Unsupervised topological alignment for single-cell multi-omics integration. *Bioinformatics* *36*, i48–i56. <https://doi.org/10.1093/bioinformatics/btaa443>.
- Swarts, J. (2019). Open-Source Software in the Sciences: The Challenge of User Support. *J. Bus. Tech. Commun.* *33*, 60–90. <https://doi.org/10.1177/1050651918780202>.
- Saia, S.M., Nelson, N.G., Young, S.N., Parham, S., and Vandegrift, M. (2022). Ten simple rules for researchers who want to develop web apps. *PLoS Comput. Biol.* *18*, e1009663. <https://doi.org/10.1371/journal.pcbi.1009663>.
- Lyons, B., Isaac, E., Choi, N.H., Do, T.P., Domingus, J., Iwasa, J., Leonard, A., Riel-Mehan, M., Rodgers, E., Schaeffbauer, L., et al. (2022). The Simularium Viewer: an interactive online tool for sharing spatiotemporal biological models. *Nat. Methods* *19*, 513–515. <https://doi.org/10.1038/s41592-022-01442-1>.
- Hossain, S. (2019). Visualization of Bioinformatics Data with Dash Bio. In *Proc. 18th Python Sci. Conf.*, pp. 126–133. <https://doi.org/10.25080/Majora-7ddc1dd1-012>.
- Grinberg, M. (2018). *Flask Web Development: Developing Web Applications with python* (O’Reilly Media, Inc.).
- Reese, W. (2008). *Nginx: the high-performance web server and reverse proxy*. *Linux J.* *2*, 2.
- Scala, F., Kobak, D., Bernabucci, M., Bernaerts, Y., Cadwell, C.R., Castro, J.R., Hartmanis, L., Jiang, X., Laturnus, S., Miranda, E., et al. (2021). Phenotypic variation of transcriptomic cell types in mouse motor cortex. *Nature* *598*, 144–150. <https://doi.org/10.1038/s41586-020-2907-3>.
- Tasic, B., Yao, Z., Grayback, L.T., Smith, K.A., Nguyen, T.N., Bertagnoli, D., Goldy, J., Garren, E., Economo, M.N., Viswanathan, S., et al. (2018). Shared and distinct transcriptomic cell types across neocortical areas. *Nature* *563*, 72–78. <https://doi.org/10.1038/s41586-018-0654-5>.
- Hodge, R.D., Bakken, T.E., Miller, J.A., Smith, K.A., Barkan, E.R., Grayback, L.T., Close, J.L., Long, B., Johansen, N., Penn, O., et al. (2019). Conserved cell types with divergent features in human versus mouse cortex. *Nature* *573*, 61–68. <https://doi.org/10.1038/s41586-019-1506-7>.
- Baltrušaitis, T., Ahuja, C., and Morency, L.-P. (2019). Multimodal Machine Learning: A Survey and Taxonomy. *IEEE Trans. Pattern Anal. Mach. Intell.* *41*, 423–443. <https://doi.org/10.1109/TPAMI.2018.2798607>.
- Cohen Kalafut, N., Huang, X., and Wang, D. (2023). Joint variational autoencoders for multimodal imputation and embedding. *Nat. Mach. Intell.* *5*, 631–642. <https://doi.org/10.1038/s42256-023-00663-z>.
- Alatkar, S., and Wang, D. (2022). CMOT: Cross Modality Optimal Transport for multimodal inference. Preprint at bioRxiv. <https://doi.org/10.1101/2022.10.15.512387>.
- Nguyen, N., and Caruana, R. (2007). Consensus Clusterings. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*, pp. 607–612. <https://doi.org/10.1109/ICDM.2007.73>.

22. Olson, R.H., Cohen Kalafut, N., and Wang, D. (2023). MANGEM: a web app for Multimodal Analysis of Neuronal Gene expression, Electrophysiology and Morphology (v1.0.0). In Zenodo. <https://doi.org/10.5281/zenodo.8048263>.
23. Gouwens, N.W., Sorensen, S.A., Baftizadeh, F., Budzillo, A., Lee, B.R., Jarsky, T., Alfiler, L., Baker, K., Barkan, E., Berry, K., et al. (2020). Integrated Morphoelectric and Transcriptomic Classification of Cortical GABAergic Cells. *Cell* 183, 935–953.e19. <https://doi.org/10.1016/j.cell.2020.09.057>.
24. Y. Ma, and Y. Fu, eds. (2011). *Manifold Learning Theory and Applications* (CRC Press). <https://doi.org/10.1201/b11431>.
25. Singh, R., Demetci, P., Bonora, G., Ramani, V., Lee, C., Fang, H., Duan, Z., Deng, X., Shendure, J., Distech, C., and Noble, W.S. (2020). Unsupervised manifold alignment for single-cell multi-omics data. *ACM Conf. Bioinforma. Comput. Biol. Biomed.* 1–10. <https://doi.org/10.1145/3388440.3412410>.
26. Maaten, L. van der, and Hinton, G. (2008). Visualizing Data using t-SNE. *J. Mach. Learn. Res.* 9, 2579–2605.
27. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830.