



iRODS as an Object Store for the Galaxy Platform

Kaivan Kamali, Nate Coraor, Marius van den Beek, John Chilton, Anton Nekrutenko
Penn State University

What is Galaxy?

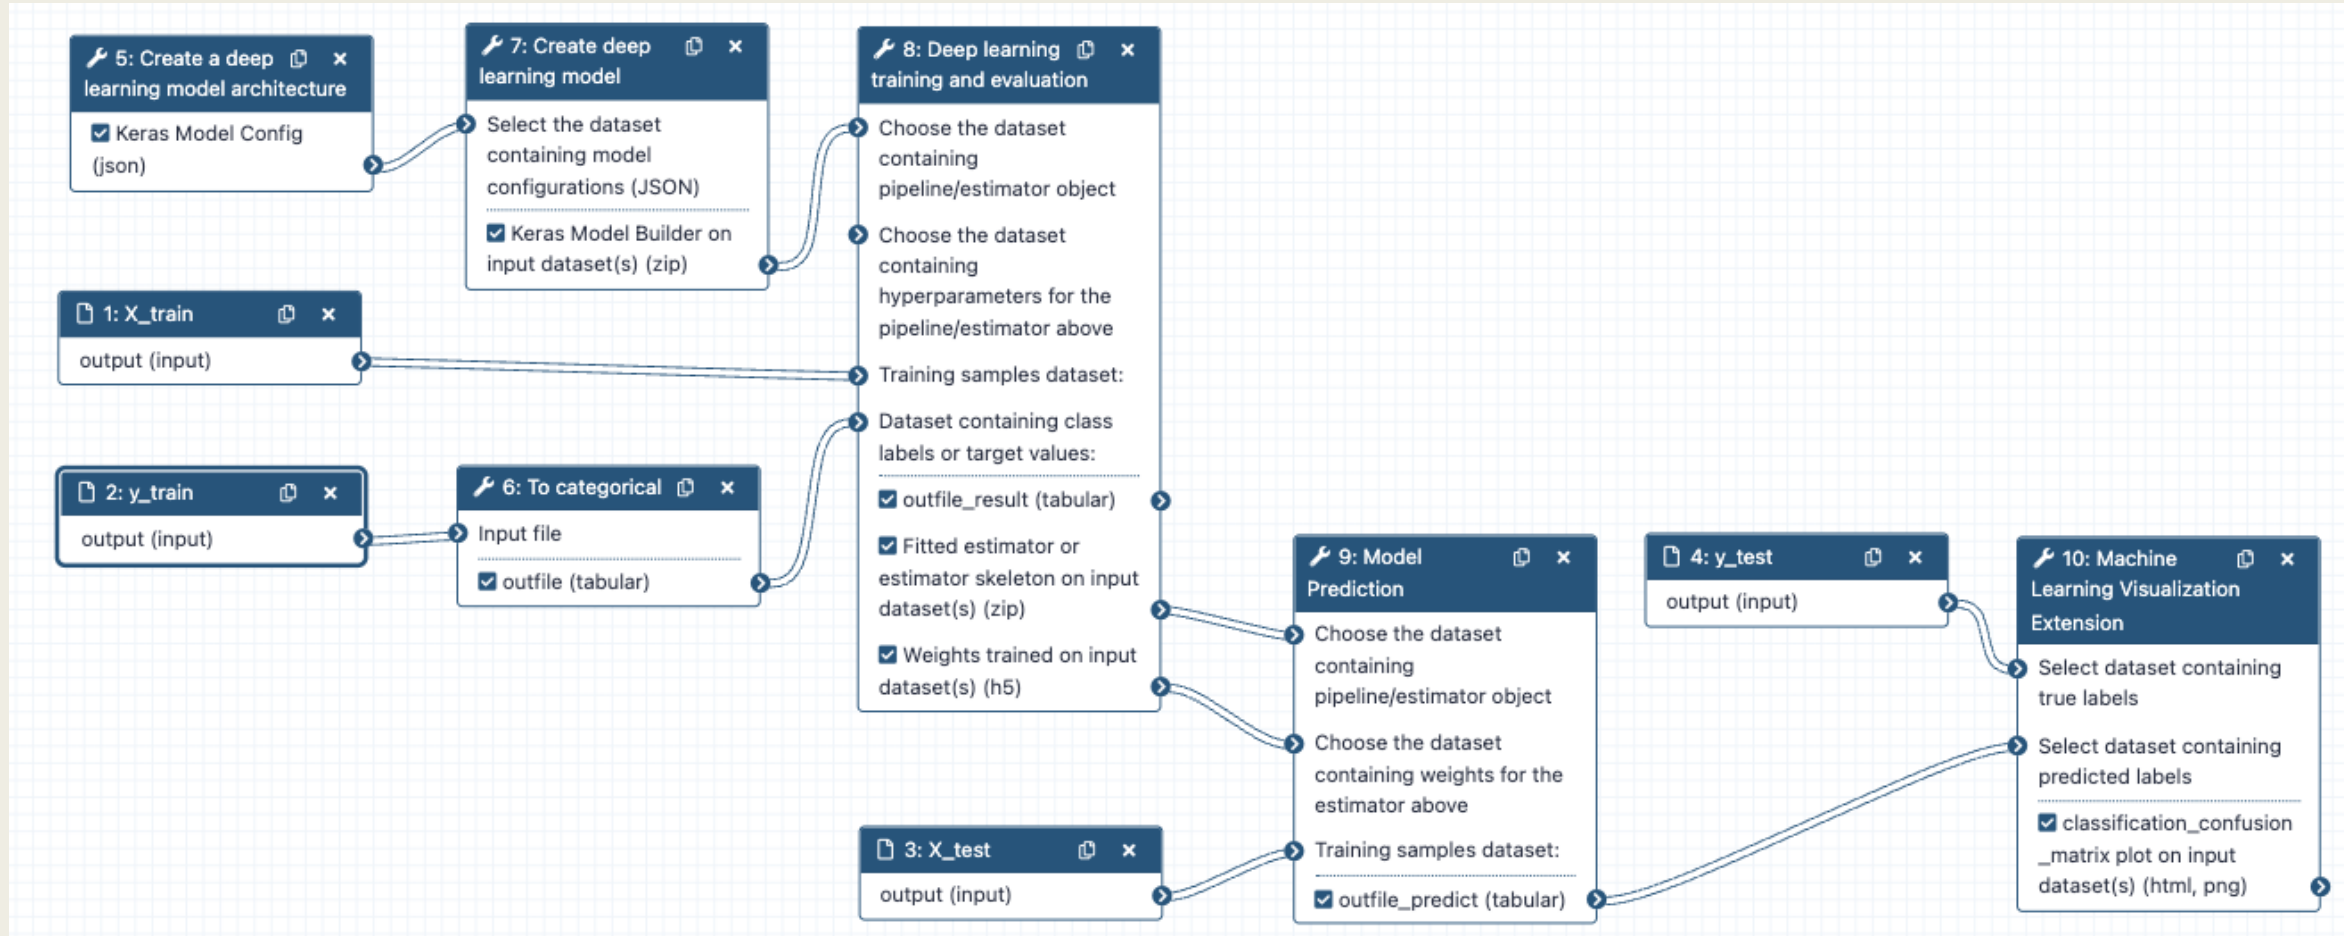
- <https://galaxyproject.org/>
- Computational workbench used by thousands of scientists across the world to analyze large heterogeneous datasets, e.g., bioinformatics, chemistry, ecology, climate science, images [1]
- Easy to use [2, 3]
 - GUI based. No programming skills required
 - Web based. No system administration skills required
- Free and Open Source
- Many tools (~8000)
- Popular (>10.000 citations)
- Extensive tutorials available via Galaxy Training Network (GTN)

Galaxy Interface

The screenshot displays the Galaxy web interface with three main sections highlighted by colored boxes:

- Tools (Left Sidebar, Blue Box):** A sidebar containing a search bar, an "Upload Data" button, and a list of tool categories including "Get Data", "Collection Operations", "GENERAL TEXT TOOLS", "Text Manipulation", "Filter and Sort", "Join, Subtract and Group", "Datamash", "GENOMIC FILE MANIPULATION", "FASTA/FASTQ", "FASTQ Quality Control", "SAM/BAM", "BED", "VCF/BCF", "Nanopore", "Convert Formats", "Lift-Over", "COMMON GENOMICS TOOLS", "Interactive tools", "Operate on Genomic Intervals", "Fetch Sequences/Alignments", "GENOMICS ANALYSIS", and "Assembly".
- Central Content (Red Box):** A main area featuring a text introduction: "Galaxy is an open source, web-based platform for data intensive biomedical research. If you are new to Galaxy start here or consult our help resources. You can install your own Galaxy by following the tutorial and choose from thousands of tools from the Tool Shed." Below this is a video player showing a woman speaking, with a blue overlay text: "James P. Taylor Foundation for Open Science. 'The most important job of senior faculty is to mentor junior faculty and students.' — @jxtx". A "Learn More" button is positioned below the video. At the bottom, a blue information box states: "Want to learn the best practices for the analysis of SARS-CoV-2 data using Galaxy? Visit the Galaxy SARS-CoV-2 portal at covid19.galaxyproject.org". Logos for PennState, Johns Hopkins University, and Oregon Health & Science University are shown at the bottom.
- History (Right Sidebar, Green Box):** A sidebar with a "search datasets" bar, a "Galaxy 101 History" section showing "2 shown" and "7.48 MB", and a list of recent jobs: "2: SNPs" and "1: Exons", each with visibility and delete icons.

Reproducible Research w/ Workflows



Transparency

- Galaxy Lets you share your histories, workflows, etc., enabling transparent research
- Sharing options
 - *Share with specific users*
 - *Share via link, with anyone who knows the link*
 - *Publish to make it visible to everybody*

Uploading Data

- Galaxy supports data imports from
 - *The user's computer,*
 - *By URL,*
 - *And, directly from many online resources, e.g., UCSC, NCBI, etc.*
- Galaxy supports a range data formats, and translation between those formats
- The Galaxy servers provide substantial CPU and disk space
 - On usegalaxy.org, the median size of the datasets created by all users per day is 8.12 TB.

Uploaded File

The screenshot displays the Galaxy web interface. At the top, the navigation bar includes the Galaxy logo, a home icon, and menu items for Workflow, Visualize, Shared Data, Admin, Help, and User. A notification bar below the navigation bar states: "Galaxy has recently been updated to a pre-release of the next Galaxy stable release (22.05), and you might encounter problems as we iron out the wrinkles. Please report any issues you encounter using the bug icon on error (red) datasets, or email galaxy-bugs." The main interface is divided into three vertical panels. The left panel is a sidebar with a "Tools" section containing a search box and an "Upload Data" button. Below this are categories for "Get Data", "Send Data", "Collection Operations", "GENERAL TEXT TOOLS", "Text Manipulation", "Filter and Sort", "Join, Subtract and Group", "Datamash", "GENOMIC FILE MANIPULATION", "FASTA/FASTQ", "FASTQ Quality Control", "SAM/BAM", "BED", "VCF/BCF", and "Nanopore". The middle panel shows a workflow with five steps, with step 5 highlighted. The right panel is titled "History" and contains a search box for datasets. Below the search box, it shows "Unnamed history" with a size of 30 B, 1 location, and 2 trash icons. A single history item, "2 : input_file", is listed with a green background and icons for viewing, editing, and deleting.

Run a Sort Tool

The screenshot displays the Galaxy web interface. At the top, the navigation bar includes the Galaxy logo, a home icon, and menu items for Workflow, Visualize, Shared Data, Admin, Help, and User. A notification banner states: "Galaxy has recently been updated to a pre-release of the next Galaxy stable release (22.05), and you might encounter problems as we iron out the wrinkles. Please report any issues you encounter using the bug icon on error (red) datasets, or email galaxy-bugs." The main content area is titled "Sort data in ascending or descending order (Galaxy Version 1.1.1)".

Tools sidebar (left):

- search tools
- Upload Data
- Select first lines from a dataset (head)
- Unique lines assuming sorted input file
- Create text file with recurring lines
- Multi-Join (combine multiple files)
- Join two files
- Text transformation with sed
- Unfold columns from a table
- Unique occurrences of each record
- Sort data in ascending or descending order
- Sort a row according to their columns
- Add line to file writes a line of text at the beginning or end of a text file.
- UniProt ID mapping and retrieval
- Search in textfiles (grep)
- SQLite to tabular for SQL query
- Query Tabular using sqlite sql
- Table Compute computes operations on table data
- annotateMyIDs annotate a generic set of identifiers
- Compute an expression on every row
- Concatenate multiple datasets tail-to-head
- Concatenate datasets tail-to-head (cat)
- Split file according to the values of a column
- Regex Find And Replace
- Column Regex Find And Replace

Sort Query configuration:

- Input: 2: input_file
- Number of header lines: 1
- Column selections: 1: Column selections
- on column: Column: 1
- in: Ascending order, Descending order
- Flavor: Fast numeric sort (-n), General numeric sort (scientific notation -g), Natural/Version sort (-V), Alphabetical sort, Human-readable numbers (-h), Random order (-R)
- + Insert Column selections
- Output unique values: No
- Ignore case: No
- Email notification: No
- Execute button

History sidebar (right):

- search datasets
- Unnamed history
- 20 B
- 2: input_file (5 lines, format tabular, database ?, uploaded txt file)
- 1: 1, 1, 5, 2, 4, 3

Output File

The screenshot displays the Galaxy web interface. At the top, a dark navigation bar contains the Galaxy logo, a home icon, and menu items for Workflow, Visualize, Shared Data, Admin, Help, and User. A notification in the top right corner indicates 'Using 8%'. Below the navigation bar, a light blue banner provides a warning about a pre-release update to version 22.05. The main interface is divided into three vertical panels. The left panel, titled 'Tools', features a search bar, an 'Upload Data' button, and a list of tool categories including 'GENERAL TEXT TOOLS' and 'GENOMIC FILE MANIPULATION'. The middle panel shows a vertical list of tool numbers 1 through 5. The right panel, titled 'History', shows a search bar and a list of datasets under 'Unnamed history', including '3 : Sort on data 2' and '2 : input_file'.

Galaxy Workflow Visualize Shared Data Admin Help User Using 8%

! Galaxy has recently been updated to a pre-release of the next Galaxy stable release (22.05), and you might encounter problems as we iron out the wrinkles. Please report any issues you encounter using the bug icon on error (red) datasets, or email galaxy-bugs.

Tools ☆

search tools x

Upload Data

Get Data

Send Data

Collection Operations

GENERAL TEXT TOOLS

Text Manipulation

Filter and Sort

Join, Subtract and Group

Datamash

GENOMIC FILE MANIPULATION

FASTA/FASTQ

FASTQ Quality Control

SAM/BAM

BED

VCF/BCF

Nanopore

Convert Formats

1

2

3

4

5

History + =

search datasets x

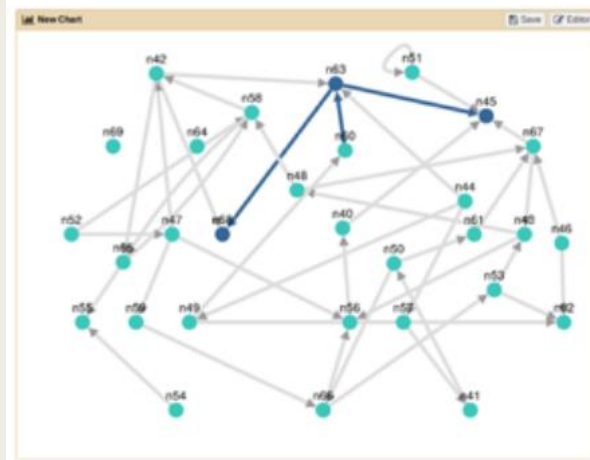
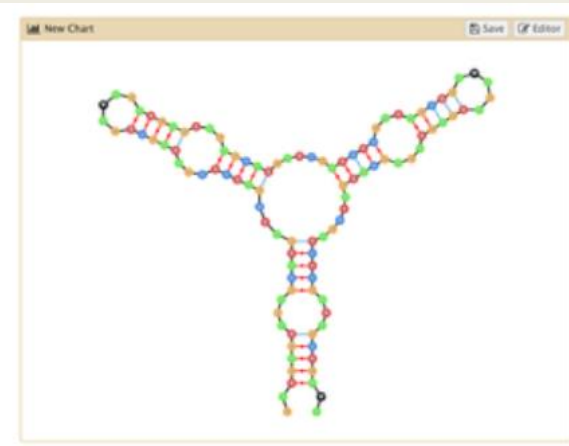
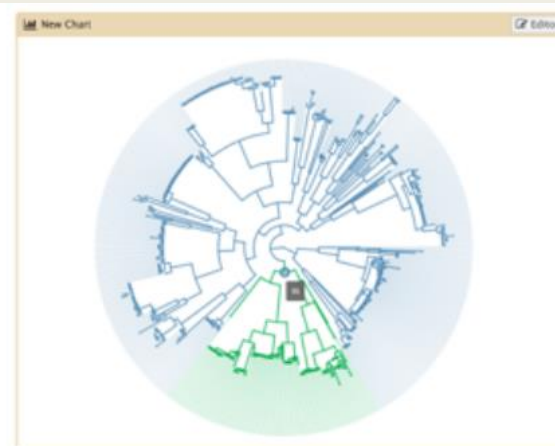
Unnamed history

30 B 2 1

3 : Sort on data 2

2 : input_file

Data Visualization in Galaxy



Galaxy Servers



Galaxy Training Network (GTN)

- Collection of tutorials developed & maintained by the worldwide Galaxy community
- Tutorials for scientists, developers, and admins
- Tutorials have slides, hands on section, datasets, workflows, and videos

28

Topics

265

Tutorials

261

Contributors

7.0

Years

Galaxy ObjectStore

- ObjectStore is Galaxy's data virtualization technology
 - Abstracts Galaxy's business logic for data persistence technology and topology
 - Makes it possible to store data on a wide-variety of persistence media, and define any data distribution policy
 - Local storage, cloud-based solutions, *etc.*
- Enables Galaxy admins to add additional persistence media to existing file system
 - Expanding a mounted filesystem (e.g., when it runs out of space) w/o moving data
 - Enables replicating data onto multiple persistence media
 - Data access fault-tolerance

Galaxy Backend

- Backend: any persistence media that ObjectStore can be configured to read/write from/to
- List of backends that ObjectStore currently supports:
 - *Local storage (e.g., disk)*
 - *Network attached storage (NAS)*
 - *Cloud (e.g. S3)*
 - *Integrated Rule-Oriented Data Store (iRODS)*

Data Distribution

- When you have multiple backends, can define nested relationship between them
 - *Hierarchical backends*
 - *Distributed backends*

	Where data is read from?	Where data is written to?
Hierarchical	first backend where data exists	always the first backend
Distributed	first backend where data exists	pseudo-randomly selected backend

Data Distribution

- Hierarchical
 - Useful when you have been using a backend for a while
 - Then decide to "extend" it by adding a new backend
 - **But without having to move data from previous backend to the new backend**
- Distributed
 - Randomly selects a backend to which it should persist data
 - The random selection is based on admin-specified weights for backends

Galaxy's iRODS ObjectStore

- iRODS parameters are specified in an ObjectStore XML configuration file

```
<?xml version="1.0"?>
<object_store type="irods">
  <auth username="rods" password="rods" />
  <resource name="demoResc" />
  <zone name="tempZone" />
  <connection host="localhost" port="1247" timeout="30" refresh_time="30"/>
  <cache path="database/object_store_cache_irods" size="1000" />
  <extra_dir type="job_work" path="database/job_working_directory_irods"/>
  <extra_dir type="temp" path="database/tmp_irods"/>
</object_store>
```

iRODS ObjectStore Instatiation

- Galaxy reads/parses object store XML configuration file
- Instantiates an iRODS ObjectStore class
- iRODS ObjectStore class provides various methods
 - create/get/delete a file
 - get_file_name, get_file_size, get_file_path
 - is_file_empty, does_file_exists, is_file_in_cache
 - *Etc.*
- iRODS ObjectStore class uses Python iRODS Client to interact with iRODS server

Galaxy w/ iRODS Support to Test Server

- We implemented iRODS ObjectStore class
- Wrote unit/integration tests
- Had the PR reviewed/merged
- Configured Galaxy to use iRODS server
 - Hosted on Texas Advanced Computing Center (TACC)
- Deployed Galaxy server w/ iRODS support to our Test server
 - *test.galaxyproject.org*

First Challenge

- Galaxy creates and maintains a Python iRODS client Session object upon startup
 - Session object maintains a pool of connections
 - Saw occasional NetworkException errors in the Galaxy log
 - Seemed older connections would get dropped
- Initial solution
 - Per discussions on iRODS-Chat (<https://groups.google.com/g/iROD-Chat>) created a Session object for each interaction with iRODS
 - The NetworkException errors disappeared
 - However, creating a Session object for each interaction with iRODS was not performant
 - Creating/destroying thousands of iRODS Session objects was costly

First Challenge -- Continued

- Subsequent solution
 - Maintain the session object
 - Do not create/destroy the session object for each iRODS interaction
 - Record connection creation time for each connection in the Pool
 - When interacting with iRODS, if the connection was created more than N seconds ago (N configurable, default value is 300 seconds), recreate the connection
 - Otherwise, use the connection as is
 - This resulted in NetworkException errors disappearing
 - And, the solution is also performant


Second Challenge

- Python iRODS client's Post and Get methods were not performant
 - Slowed down Galaxy Post/Get operations
 - Specially for larger files
- The client's Post/Get runtime was much slower than iCommands equivalent
- Proposed solution
 - Discussed supporting multi-threaded Post/Get on the client side only
 - Eventually decided to scrap client only solution
 - Decided on multi-threaded Post/Get support on both client AND server sides
 - Python iRODS client: v0.9.0+
 - iRODS server: v4.2.9+
 - Data object transfers using put()/get() will spawn a number of threads to optimize performance
 - File sizes larger than a default threshold value of 32 MB


Python iRODS Client Vs iCommands

- Wrote scripts that generate M (say, 50) files of size N (say, 100 GB) with random content
- Post files to iRODS server using both Python iRODS client and iCommands
- Averaged the runtime of each to compare
- Get files from iRODS server using both Python iRODS client and iCommands
- Averaged the runtime of each to compare
- Repeated the steps above for file sizes 100 bytes, 10 KB, 1 MB, 2.25 MB, 4 MB, 6.25 MB, 25 MB, 50 MB, 100 MB, and 1 GB

POST: icommand vs client



GET: icommand vs client



Third Challenge

- Galaxy server became unresponsive
 - Call stack showed it's hung in Python iRODS client code
- After extensive debugging, found a bug in client code (in `_recv_message_in_len()`)
 - The loop that iteratively reads from the socket, depended on the number of bytes read each time, to terminate the loop
 - If the iRODS server is down, the number of bytes read in each iteration of the loop is 0, causing it to loop indefinitely, halting the app
 - Revised the logic to account for this scenario
- Also, found and fixed a bug with Connection destructor
 - Memory leak fixed

Current Status

- Configured/deployed Galaxy with iRODS object store to Main (<https://usegalaxy.org>)
- Created a group of users that their object store access is overridden to iRODS
 - This is to gradually release iRODS to only a select few
 - Testing has been promising with no performance issues
- After this stage of our testing is complete, we plan to use a distributed object store
 - We use weights to split the writes between iRODS and disk
 - Say, start with 1 write to iRODS, for every 9 writes to disk
 - Reads go against both backends, depending on where the data resides

Current Status -- Continued

- We plan to incrementally increase iRODS weights and decrease disk weights
 - Say, 10% every month
 - Until all writes go to iRODS
 - Again, reads go against both backends, depending on where the data resides
 - *Can a real-time multi-user application run on top of iRODS?*
- Finally, we plan on migrating the data from disk to iRODS
 - Data migration happens via a separate script
 - Can then retire disk object store

Thank you!

- We would like to thank all the members of the iRODS team for their support and always being available
 - Special thanks to Daniel and Terrell!
- Questions/Comments?

References

1. Vahid Jalili, et. al. The Galaxy platform for accessible, reproducible and collaborative biomedical analyses: 2020 update, *Nucleic Acids Research*, Volume 48, Issue W1, 02 July 2020, Pages W395–W402, <https://doi.org/10.1093/nar/gkaa434>
2. A Short Introduction to Galaxy (<https://training.galaxyproject.org/training-material/topics/introduction/tutorials/galaxy-intro-short/slides.html#1>)
3. Introduction to Galaxy. <https://training.galaxyproject.org/training-material/topics/introduction/slides/introduction.html#1>
4. Galaxy ObjectStore. <https://galaxyproject.org/admin/objectstore/>
5. Dockerized iRODS Server. <https://github.com/kxk302/irods>