

1 SciApps: A Cloud-Based Platform for Analyses and Distribution of the MaizeCODE data

2
3 Liya Wang¹, Zhenyuan Lu¹, Melissa delaBastide¹, Peter Van Buren¹, Xiaofei Wang¹, Cornel
4 Ghiban¹, Michael Regulski¹, Jorg Drenkow¹, Xiaosa Xu¹, Carlos Ortiz Ramirez², Cristina F.
5 Marco¹, Jason Williams¹, Alexander Dobin¹, Kenneth D. Birnbaum², David P. Jackson¹, Robert
6 A. Martienssen¹, William R. McCombie¹, David A. Micklos¹, Michael C. Schatz^{1,3}, Doreen H.
7 Ware^{1,4}, Thomas R. Gingeras¹

8
9 ¹Cold Spring Harbor Laboratory, Cold Spring Harbor, NY, ²New York University, New York, NY,
10 ³Johns Hopkins University, Baltimore, MD. ⁴USDA-ARS Robert W. Holley Center for Agriculture
11 and Health, Ithaca, NY, United States

12
13 MaizeCODE is a project aimed at identifying and analyzing functional elements in the maize
14 genome. In its initial phase, MaizeCODE assayed up to five tissues from four maize strains
15 (B73, NC350, W22, TIL11) by RNA-Seq, Chip-Seq, RAMPAGE, and small RNA sequencing. To
16 facilitate reproducible science and provide both human and machine access to the MaizeCODE
17 data, we developed SciApps, a cloud-based portal, for analysis and distribution of both raw data
18 and analysis results. Based on the SciApps workflow platform, we generated new components
19 to support the complete cycle of the MaizeCODE data management. These include publicly
20 accessible scientific workflows for reproducible and shareable analysis of various functional
21 data, a RESTful API for batch processing and distribution of data and metadata, a searchable
22 data page that lists each MaizeCODE experiment as a reproducible workflow, and integrated
23 JBrowse genome browser tracks linked with workflows and metadata. The SciApps portal is a
24 flexible platform that allows integration of new analysis tools, workflows, and genomic data from
25 multiple projects. Through metadata and a ready-to-compute cloud-based platform, the portal
26 experience improves access to the MaizeCODE data and facilitates its analysis.

27
28 **Keywords:** bioinformatics, functional annotations, cloud computing, ENCODE, workflows

INTRODUCTION

Maize is one of the most biologically, socially, and economically important crop plants. Following the sequencing of its genome (Jiao et al., 2017; Schnable et al., 2009), the next critical step in understanding maize biology will involve identifying and deciphering functional sequence regions. Modeled on the Encyclopedia of DNA Elements (ENCODE) Project for the human genome (ENCODE Consortium, 2004), the MaizeCODE project is an integrated and multi-disciplinary project aimed at revealing the functional regions of the maize genome by identifying loci that are transcribed, methylated, or bound by specific modified histones and transcription factors in various tissues. In addition, MaizeCODE is designed to store, collate, display, and disseminate the data to the wider community of plant biologists worldwide.

To curate, process, and distribute the ENCODE data, the ENCODE Data Coordination Center (DCC) group established the ENCODE portal (Sloan et al., 2016), which relies on both rich metadata and commercial cloud resources through the DNAnexus platform (<https://www.dnanexus.com>). Within this platform, standard processing pipelines for human genome analysis are constructed to ensure consistent and reproducible processing of primary sequence data. However, both the ENCODE DCC and end users are required to cover the cost of the DNAnexus service and commercial cloud resources. In order to provide a cost-free data processing platform for academic users, the MaizeCODE DCC group decided to leverage two NSF-funded resources, XSEDE (Towns et al., 2014) at the Texas Advanced Computing Center (TACC) for computing power and the CyVerse Data Store (Goff et al., 2011) for cloud-based data storage.

To automate bioinformatics analysis over both the XSEDE/TACC cloud and CyVerse Data Store, we developed a bioinformatics workflow platform called SciApps (Wang et al., 2018). In

the work described here, we further improved SciApps by adding a RESTful API for automating batch processing of the MaizeCODE data and metadata management, a searchable MaizeCODE data page powered by a relational database, several analysis workflows, and Genome Browser tracks automatically generated from unique workflow identifiers via the RESTful API. The SciApps platform has been used to support both MaizeCODE DCC and end users to process/reprocess and manage multi-omics data through either the GUI or the API.

METHODS

Overview of the entire MaizeCODE data management cycle

To improve accessibility, reproducibility, reusability, and interoperability, data generated by the MaizeCODE Consortium members are uploaded to a cloud-based data storage system, the CyVerse Data Store (Goff et al., 2011). The Data Store, which is built on top of iRODS (a rule-oriented data system) ([Moore and Rajsekar 2010](#)), supports data virtualization, sharing, bulk uploads/downloads, and collaborations. Once uploaded, the experimental metadata are attached to raw data files to facilitate the reuse of data, as well as submission of data to the NCBI Short Read Archive (SRA). The metadata is later retrieved via the Terrain API (<https://github.com/cyverse-de/terrain>) to automate batch analyses via SciApps (Wang et al., 2018). SciApps provides a ready-to-compute cloud-based platform for automating complex analyses constructed using modular applications (or apps). Previously, SciApps was operated via a web GUI, but since that time we have developed SciApps RESTful APIs to support batch processing of MaizeCODE and other data. In addition, we have integrated over 20 new apps for ground-level analyses such as quality control (QC), alignment to the reference genome, filtering, quantification (e.g. for gene expression), and peak calling (if needed). Replicates (and controls if available) of each assay are organized as a single experiment (or workflow with a unique ID),

which represents an entity that chains together raw data, analysis results, experimental metadata, and computational provenance. SciApps extracts experimental metadata and attaches them to a specific workflow so that users can access them directly on the SciApps portal. Genome browser tracks are automatically generated and displayed within an integrated version of JBrowse (Skinner et al., 2009) by looping through the list of experiments/workflows via the SciApps RESTful API.

Together, the SciApps portal supports the complete cycle of MaizeCODE data management, and the level of automation it provides greatly decreases the chances of human errors in data organization, analysis, and distribution.

Processing and accessing the MaizeCODE data with the SciApps RESTful API

The cloud-based architecture of SciApps (Wang et al., 2018) enables highly scalable processing of MaizeCODE data on the XSEDE/TACC cloud. Both intermediate and final results are archived in the CyVerse Data Store, where the raw data are also hosted. As discussed above, each SciApps workflow captures experimental metadata and computational provenance along with the raw and processed data. Batch processing is supported through the RESTful API; the API endpoints are provided in **Table 1**.

Table 1. SciApps release 1.0 RESTful API

Endpoint	Method	Description
/job	GET	List all jobs
/job/new/{id}	POST	Run a new job
/workflow/build	POST	Build a workflow from jobs

/workflowJob/new	POST	Generate a workflow JSON
/workflow/new	POST	Save a new workflow
/job/{id}	GET	Return the job JSON
/job/{id}/delete	GET	Delete the job
/workflowJob/run/{id}	GET	Run a new workflow
/workflow/{id}/metadata	GET	Get the workflow metadata
/workflow/{id}/update	POST	Update the workflow with metadata etc
/workflow	GET	List all workflows
/apps/{id}	GET	Return the application JSON
/workflow/{id}/delete	GET	Delete the workflow
/workflow/{id}	GET	Return the workflow JSON
/apps	GET	List all integrated apps

1

2 The analysis workflow for a specific assay is typically built interactively within the SciApps GUI

3 using one data set as a template. Once the workflow is captured, it can then be easily and

4 automatically applied to analyze other genomes and tissues. Alternatively, users may also build

5 workflows entirely programmatically with a series of analysis job IDs via the API. Experimental

6 metadata are retrieved via the CyVerse Terrain API, and then attached to the workflow via the

7 SciApps API at runtime. The API supports the MaizeCODE DCC for automatically processing a

8 large amount of data and also supports retrieval of results and metadata by end users. For

9 example, genome browser tracks can be automatically generated given a workflow ID by the

10 following steps (**Figure 1**): 1. Retrieve job IDs and inputs with the workflow endpoint, given a

11 workflow ID; 2. Retrieve the output path with the job endpoint, given a job ID; 3. Construct the

12 browser-ready link with the retrieved information. To simplify the process, the MaizeCODE DCC

13 encodes the genome, tissue, and replicate information into the input raw data file path, which is

14 also accessible through the workflow metadata endpoints. SciApps also names the output

filename based on the input filename with the output ID (defined by the app) as the prefix. As shown in **Figure 1**, once the input filename, input path, and output path to cloud storage are retrieved by calling the API, the output file path can be constructed to build the browser links.




Figure 1. Flowchart of generating genome browser tracks via cURL for a MaizeCODE workflow. Each workflow has two replicates. **\$WF_ID** is the workflow ID, and **\$JOB_ID** is the job ID of a step of the workflow. As an example, the cURL for retrieving the input filename of replicate 1 is:

```
curl -X GET --header "Accept: application/json" "https://www.sciapps.org/workflow/$WF_ID" | jq ".data.steps[] | select(.id==1) | .inputs.inputSequence1".
```


'Relaunch' the analysis, 'Visualize' the graphic diagram of the workflow (with URLs for the raw sequence files from the input file node), 'Load' the results to the History panel, 'Share' the analysis with others, and display the experimental 'Metadata'. User can perform a keyword search for a specific dataset (e.g., B73 ears RNA-Seq). In the right panel, SciApps displays the history of the selected datasets; the visualization (eye) icon opens a panel where users can generate links to visualize the results in a web browser (e.g., a QC report) or genome browser (e.g., alignments or signal tracks). The left panel shows a list of modular apps that can be launched to perform a variety of downstream analyses with the loaded results.

Accessing the MaizeCODE data as Genome Browser tracks

Once the analysis is completed, genome browser tracks are automatically generated given the workflow ID by calling the SciApps API for an integrated version of JBrowse. The browser tracks can be accessed under the 'Tools' menu within the top navigation bar. As shown in **Figure 3**, tracks are organized by genome, tissue, replicate, and assay. Checking the box next to each track will load it into the browser. The SciApps workflow ID is embedded, so clicking on a track brings up the workflow 'Relaunch' interface, which can be used to reproduce the track signal if needed. In this interface, the user can also check the parameters used for the analysis, as well as additional results in the History panel. At the bottom of the interface, a diagram button visualizes the workflow diagram, and a metadata button displays the experimental metadata associated with the workflow. From the results, user can also generate additional browser track links through the visualization (eye) icon. For example, this can be used to verify the signal track with the alignment files (in the BAM format). As mentioned earlier, the results can also be used to perform a downstream analysis on the same interface. Finally, the browser tracks are available as a JSON file for integration into other platforms (e.g., the JSON file for B73 is available at https://data.sciapps.org/view2/data2/B73/v4/apollo_data/trackList.json).




Figure 3. Genome browser tracks for the MaizeCODE data. JBrowse is used to hold the MaizeCODE signal tracks, which are organized in the following order: genome, tissue, replicate, and assay. Clicking on each track brings up the workflow 'Relaunch' interface.

Accessing the raw reads on CyVerse Data Store

The raw sequence data is deposited into the CyVerse Data Store via iCommands (<https://docs.irods.org/4.2.1/icommands/user/>), with metadata attached before submission to the NCBI SRA. From there, users can access the raw data in several ways. Within SciApps, the input file node of the graphic diagram for a workflow/experiment is linked to the raw sequence file. Clicking on the input node will open the CyVerse Data Common landing page in a web browser. The metadata attached to the raw sequence file is also displayed on the same page. The user can further navigate through all released raw data from the landing page (<http://datacommons.cyverse.org/browse/iplant/home/shared/maizecode/released/>); the SciApps workflow ID is attached as metadata to the raw data files if it has been processed. The

user can use the ID to load the workflow on the SciApps portal. For batch downloading of raw sequence files through the GUI or the command line, we recommend CyberDuck (<https://cyberduck.io/>) or iCommands, respectively.

Analysis with reproducible workflows

Bioinformatics applications (or apps) are integrated into SciApps as modular components that can be chained with other apps into an automated workflow. Individual apps are built with Singularity images (Kurtzer et al., 2017) from BioConda recipes (Grüning et al., 2018) or directly from Dockerfiles to ensure reproducibility across different cloud resources. To support the analysis of MaizeCODE data, over 20 software tools are integrated. **Figure 4** shows two publicly accessible workflows for differential expression analysis and cytosine methylation analysis, building on the popular STAR (Dobin et al., 2013)/RSEM (Li and Dewey, 2011)/StringTie (Pertea et al., 2015) and Bismark (Krueger and Andrews, 2011) pipelines, respectively. These workflows can be constructed either with the SciApps GUI or through the API. The user can retrieve the inputs, metadata, results, and provenance of the software used in the analysis with a unique workflow ID. The interactive graph, along with the platform guide (<https://cyverse-sciapps-guide.readthedocs-hosted.com/en/latest/index.html>), helps users to understand how multiple apps are used together to analyze a specific assay. For MaizeCODE data, the graph is also helpful for visually inspecting the input–output relationship. Additionally, the user can check the input data (through the input file node) and relaunch each individual step of the analysis, or even the entire analysis, via the web interface or API.




Figure 4. Graphical workflow diagrams for differential expression analysis (top) and MethylationC-seq analysis (bottom). The interactive graph demonstrates the relationships among input–output files, displays provenance of the software tools, and provides real-time job status updates of

new analyses via the color of the app node (green: completed; blue: running; yellow: pending; red: failed).

RESULTS AND DISCUSSION

A large variety of software is needed to process the MaizeCODE data. For each experiment (consisting of two replicates), a workflow with a unique ID is provided via the SciApps platform. One major goal of SciApps is to empower anyone in the community to easily repeat an entire analysis, or use a workflow with alternative parameters for each step if so desired. A second major goal is to empower community members to process their own comparable data sets using protocols validated by MaizeCode. In the following sections, we describe how RNA-seq data is processed, how the results can be visualized, and how the primary analysis results can be used for differential expression analysis.

Processing the RNA-seq data

Besides the UCSC genome analysis tools (for format conversion and generating browser track signals), the major software used in MaizeCODE RNA-seq data analysis are bbdduk (<https://jgi.doe.gov/data-and-tools/bbtools/bb-tools-user-guide/bbdduk-guide/>) for trimming low quality reads and adapter sequences, FastQC and MultiQC (Ewels et al., 2016) for visually checking read quality, STAR (Dobin et al., 2013) for alignment, RSEM (Li and Dewey, 2011) for quantifying gene expression, and StringTie (Pertea et al., 2015) for transcriptome assembly. All tools are integrated into SciApps individually as separate apps, and also combined as a single app, MCrna, for rapid batch processing of RNA-seq data without requiring intermediate results to be transferred between the TACC and CyVerse cloud. **Figure 5** shows the relationship among these analysis tools within the MCrna app, which is used to process both replicates of an




Figure 5. MaizeCODE MCrna app for processing RNA-seq data from two replicates.

For each replicate, the MultiQC software outputs a quality report for the sequence data, before and after trimming, in an interactive HTML format. This report can be accessed via the visualization (eye) icon in the History panel (next to each loaded replicate, as shown in **Figure 2**). As with the HTML format, text, image, and other web browser-compatible files can be visualized by clicking the icon. For files that can be displayed on a Genome Browser (e.g., BAM, bigwig, etc), the user can also generate browser-ready links by clicking the same icon. These links address the cloud storage system from the CyVerse project, so they can be displayed on Genome Browsers hosted by different portals. If the user clicks on the output file name (from the History panel), they will be directed to the CyVerse Data Commons landing page, where they can preview or download the results. For files over a few GBs, we recommend that the user download their data using either iCommands or CyberDuck, using the CyVerse Data Store path available in the file URL.

Automated differential expression analysis

One of the key advantages of distributing the MaizeCODE data through SciApps is that it facilitates downstream analysis. In this section, we will show how differential expression analysis can be performed, on either the gene or isoform levels, using the primary analysis results discussed in the last section.

As shown in **Figure 6**, after loading the results for two experiments from the MaizeCODE data page (**Figure 2**), one for ear tissue and the other for the root tissue, we can launch the RSEM_de app from the 'Comparison' category (or through searching the app panel). For the analysis, users drag and drop output files with names starting with 'rsem_' into the input field for both replicates of each sample. The analysis job can then be submitted to the cloud for running, and the results (i.e., the differentially expressed genes) will be available within a few minutes. Note that the app is flexible in handling different numbers of replicates per sample. Additional input fields can be added using the '+ Insert' button. For the MaizeCODE project, most data sets are generated with two replicates.

Users can check the results through the History panel or the list of jobs (under the 'Workflow' tab from the top menu). Users can also select jobs from the History panel and save them as new workflows to organize the analysis and/or share it with others. Given that the XSEDE/TACC cloud is a shared resource, and jobs may be queued for several minutes to several hours, we have also established a local cluster (Wang et al., 2015) to quickly process small jobs requiring less than an hour to complete. Powered by the Agave API (Dooley et al., 2012), SciApps treats both the XSEDE/TACC cloud and the local cluster as virtual execution systems, allowing a scientific app to be configured to run on either cloud. As such, by using the CyVerse Data Store as a central storage hub, SciApps workflows can be seamlessly executed on a mixture of XSEDE/TACC cloud and local clusters for efficient yet scalable processing and consumption of MaizeCODE data.

The screenshot displays the RSEM_de app interface. The sidebar on the left lists various categories: Alignment, Annotation, Assembly, Calculation, Clustering, Comparison, Conversion, Data handling, Mapping, Methylation, Prediction, and Quantification. The 'RSEM_de-1.3.0' link is highlighted under the 'Comparison' category. The main panel is titled 'RSEM_de (version 1.3.0): Differential expression analysis with RSEM and Eseq'. It contains sections for specifying replicates for Sample 1 and Sample 2, a false discovery rate input, and a 'Submit job' button. Red arrows point to the 'RSEM_de-1.3.0' link in the sidebar, the 'Submit job' button, and the 'Browse DataStore' links. On the right, a list of files to be built is shown, including multiqc_report.html, rsem_ears_rep1_R1.txt, sig_f_ears_rep1_R1.bw, sig_r_ears_rep1_R1.bw, star_ears_rep1_R1.bam, star_ears_rep1_R1.bam.bai, and str_ears_rep1_R1.gtf.

Figure 6. Using the RSEM_de app for gene-level differential expression analysis

To perform differential expression on the isoform level, users can launch the workflow diagram shown in **Figure 4**. From the diagram and the relaunched app forms, the StringTie_Merge app is used to merge assembled transcripts from all replicates to generate a new annotation file. The annotation file is then passed again to the StringTie app to compute gene expression quantification results, which are then input to the Ballgown app (Frazee et al., 2015) to compute differentially expressed isoforms. Again, the user can drag and drop the alignment files and assembled transcripts from the MaizeCODE primary analysis results without repeating the time-consuming alignment and quantification steps. Given that all results are accessible through web URLs, users can also retrieve the data directly to their local server for further interactive analysis. For example, SciApps users can inspect the quantification results of each pair of

replicates to confirm that they are consistent, and if they are, proceed with the analysis of differentially expressed genes.

CONCLUSION AND FUTURE WORK

SciApps is a cloud-based platform that provides a data management infrastructure for the MaizeCODE data. The platform supports the management of both experimental and computational metadata; provides a collection of analysis apps covering analysis of multi-omics data sets; and provides both web browser and API access to data, results, metadata, and computational provenance of software tools through a unique workflow ID. Genome browser tracks are also provided to enable visualization of the results using an integrated version of JBrowse.

SciApps has been designed to integrate cloud resources for scaling and long-term stability. Currently, SciApps has been integrated with the NSF-funded XSEDE/TACC cloud and the CyVerse Data Store to provide academic users with cost-free data storage and computing services. Both resources are integrated as virtual systems via the Agave API, which also supports the integration with commercial cloud platforms like Amazon EC2/S3 and Microsoft Azure. Therefore, SciApps can be scaled for large-scale data management and analysis if needed. Additionally, SciApps can also seamlessly integrate local data and computing resources to complement cloud resources. The successful utilization of our local system suggests that SciApps can facilitate collaborative projects across different institutes for joint data production, analysis, and management with multiple local systems, thereby avoiding high financial costs.

To process the data sets that are continually being generated by the MaizeCODE project, several new analysis tools have been integrated into SciApps. Future goals include developing

new analysis workflows, supporting sophisticated queries against metadata, reanalyzing and distributing published sequencing data sets from raw data, and conducting training and community outreach.

IMPLEMENTATION

The major components of the SciApps analysis portal include a web browser user interface, a MySQL database, a workflow engine, and a web API. The user interface was built with the React-Bootstrap front-end framework for data and workflow interactions. The workflow engine is written in Perl and uses the MySQL database to track job status, perform submission of a subsequent job once its inputs are ready, and record the analysis into the database. The web API is written in Perl to complement the web browser user interface, especially for batch processing and metadata handling. SciApps uses the CyVerse Central Authentication Service (CAS) to grant users access to the XSEDE/TACC cloud and CyVerse Data Store. Analysis jobs are submitted to the cloud, and in addition all released raw data, processed results, metadata, and workflows, are made publicly accessible through both the GUI and API with no authentication needed.

AUTHOR CONTRIBUTIONS

LW and ZL designed, implemented, and tested the software. MD and LW managed metadata and submission of raw data to NCBI SRA. LW and PVB designed and maintained the local system consisting of a web server, a data server, and several computing servers. LW, ZL, and XW developed the scientific workflows and participated in testing through the web interface. AD and TG helped in developing the workflow. XW and LW processed the data with the automated workflows. MB, JD, XX, and COR generated the raw data. CG developed the Perl code for interacting with Agave API. CFM, CG, and JW managed the MaizeCODE website. LW, DW,

1 MS, and TG wrote the manuscript. All the authors read and approved the final manuscript.

2

3 **AVAILABILITY AND REQUIREMENTS**

4

5 Project name: SciApps

6 Project home page: <https://sciapps.org/>

7 Source code repository: <https://github.com/warelab/sciapps>

8 Operating system(s): Platform independent

9 Programming language: JavaScript, Perl

10 License: MIT

11 Any restrictions to use the data: Toronto Agreement

12

13 **FUNDING**

14 This work is supported by NSF grant IOS 1445025 for MaizeCODE. SciApps has also been
15 supported by NSF grant DBI-1265383, and USDA-ARS (1907-21000-030-00D).

16

References

- Dobin, A., Davis, C. A., Schlesinger, F., Drenkow, J., Zaleski, C., Jha, S., et al. (2013). STAR: ultrafast universal RNA-seq aligner. *Bioinformatics* 29, 15–21.
- Dooley, R., Vaughn, M., Stanzone D, T. S., and Skidmore, E. (2012). Software-as-a-service: the iPlant Foundation API. in *5th IEEE Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS)*.
- ENCODE Consortium (2004). The ENCODE (ENCyclopedia Of DNA Elements) Project. *Science* 306, 636–640.
- Ewels, P., Magnusson, M., Lundin, S., and Käller, M. (2016). MultiQC: summarize analysis results for multiple tools and samples in a single report. *Bioinformatics* 32, 3047–3048.
- Frazee, A. C., Perte, G., Jaffe, A. E., Langmead, B., Salzberg, S. L., and Leek, J. T. (2015). Ballgown bridges the gap between transcriptome assembly and expression analysis. *Nat. Biotechnol.* 33, 243–246.
- Goff, S. A., Vaughn, M., McKay, S., Lyons, E., Stapleton, A. E., Gessler, D., et al. (2011). The iPlant Collaborative: Cyberinfrastructure for Plant Biology. *Front. Plant Sci.* 2, 34.
- Grüning, B., The Bioconda Team, Dale, R., Sjödin, A., Chapman, B. A., Rowe, J., et al. (2018). Bioconda: sustainable and comprehensive software distribution for the life sciences. *Nature Methods* 15, 475–476. doi:10.1038/s41592-018-0046-7.
- Jiao, Y., Peluso, P., Shi, J., Liang, T., Stitzer, M. C., Wang, B., et al. (2017). Improved maize reference genome with single-molecule technologies. *Nature* 546, 524–527.
- Krueger, F., and Andrews, S. R. (2011). Bismark: a flexible aligner and methylation caller for Bisulfite-Seq applications. *Bioinformatics* 27, 1571–1572.

doi:10.1093/bioinformatics/btr167.

Kurtzer, G. M., Sochat, V., and Bauer, M. W. (2017). Singularity: Scientific containers for mobility of compute. *PLoS One* 12, e0177459.

Li, B., and Dewey, C. N. (2011). RSEM: accurate transcript quantification from RNA-Seq data with or without a reference genome. *BMC Bioinformatics* 12, 323.

Pertea, M., Pertea, G. M., Antonescu, C. M., Chang, T.-C., Mendell, J. T., and Salzberg, S. L. (2015). StringTie enables improved reconstruction of a transcriptome from RNA-seq reads. *Nat. Biotechnol.* 33, 290–295.

Schnable, P. S., Ware, D., Fulton, R. S., Stein, J. C., Wei, F., Pasternak, S., et al. (2009). The B73 maize genome: complexity, diversity, and dynamics. *Science* 326, 1112–1115.

Skinner, M. E., Uzilov, A. V., Stein, L. D., Mungall, C. J., and Holmes, I. H. (2009). JBrowse: A next-generation genome browser. *Genome Research* 19, 1630–1638.
doi:10.1101/gr.094607.109.

Sloan, C. A., Chan, E. T., Davidson, J. M., Malladi, V. S., Strattan, J. S., Hitz, B. C., et al. (2016). ENCODE data at the ENCODE portal. *Nucleic Acids Res.* 44, D726–32.

Towns, J., Cockerill, T., Dahan, M., Foster, I., Gaither, K., Grimshaw, A., et al. (2014). XSEDE: Accelerating Scientific Discovery. in *Computing in Science & Engineering* 5., 62–74.

Wang, L., Lu, Z., Van Buren, P., and Ware, D. (2018). SciApps: a cloud-based platform for reproducible bioinformatics workflows. *Bioinformatics* 34, 3917–3920.

Wang, L., Van Buren, P., and Ware, D. (2015). Architecting a distributed bioinformatics platform with iRODS and iPlant Agave API. in *International Conference on Computational Science and Computational Intelligence (CSCI)*, 420–423.