An open, scalable, and flexible framework for automated aerial measurement of field experiments

Christophe Schnaufer, Julian L. Pistorius, and David S. LeBauer

Arizona Experiment Station, College of Agriculture and Life Sciences
University of Arizona, Tucson, USA

ABSTRACT

Unoccupied areal vehicles (UAVs or drones) are increasingly used in field research. Drones capable of routinely and consistently capturing high quality imagery of experimental fields have become relatively inexpensive. However, converting these images into scientifically useable data has become a bottleneck. A number of tools exist to support this workflow, but there is no framework for making these tools interopreable, sharable, and scalable.

Here we present an initial draft of the Drone Processing Pipeline (DPP), a framework for processing agricultural research imagery that supports best practices and interoperability. DPP emphasizes open software and data that can be shared among and used in whole or part by the research community. We are building the DPP as a distributed, scalable, and flexible pipeline for converting drone imagery into orthomosaics, point clouds, and plot level statistics. Our intent is not to replace, but to integrate components from the emerging ecosystem of utilities with a focus on end-to-end automation and scalability.

The initial focus of DPP is the measurements of experimental plots in field research. In the future we expect that standardization will enable new scientific discovery by facilitating collaboration and sharing of software and data. Our vision is to create a processing pipeline that is open, flexible, extensible, portable, and automated. With modern tools, deploying a pipeline on a laptop or HPC should only take a single command. Running a pipeline and publishing data should require only input data and a defined workflow.

Keywords: unoccupied aerial vehicles, data standards, high performance computing, open source software, high throughput phenotyping, agriculture, automation, workflows

1. INTRODUCTION

1.1 Background

Remote sensing using small Unoccupied Aerial Vehicles (UAVs or drones) is increasingly common in research across plant science domains,[1] and there is particular interest in agricultural applications.[2] It has become easy and inexpensive for researchers to use drones to collect data with a range of sensors. However, it is not yet as easy to to convert the images into data that can be analyzed and acted upon.

While drones make it easier to collect data, these data require many steps to collect and convert into useful information. It is even more difficult to process data in an open and reproducible way. Below we review a few pipelines that are being developed for this use case. We hope to leverage and support such solutions, and make it easier for the community to use and develop new ones. While there are many reasons that a researcher will want to independently develop a custom processing pipeline, the lack of shared software and protocols should not be one of them. It takes a lot of work to build a robust pipeline, and if this work can be shared we can simultaneously reduce redundancy and increase reliability of researcher-driven scientific software.

A key to developing this processing pipeline is to automate and guide contributors to best practices related to computing and standardized data and meta-data standards. Examples of this range from: automatically

Further author information: (Send correspondence to C.S.)

C.S.: E-mail: schnaufer@arizona.edu, ORCID: 0000-0002-6150-4558

J.L.P.: Email: julianp@arizona.edu, ORCID: 0000-0002-3485-0084

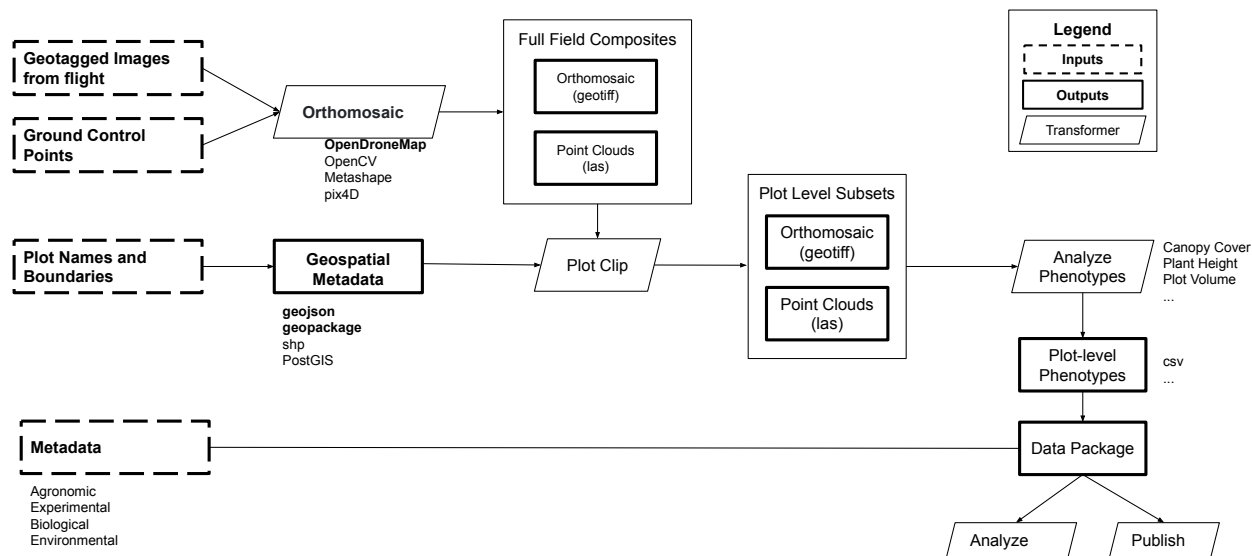D.S.L.: E-mail: dlebauer@arizona.edu, ORCID: 0000-0001-7228-053X

Figure 1. Overview of the core data processing pipeline. bold identifies default tools; a few alternatives are in normal font.

archiving code and data at each stage of processing, to capturing workflow provenance, to generating data products that are useful, by following FAIR (Findable, Accessible, Interoperable, Reusable) guidelines.[3]

Plot based field research is the cornerstone of field-based research[4] and represents a clearly defined workflow (Fig. 1) with a broad community of potential users. Field research plots on the order of 1-10 meters square are commonly used as the unit of experimental replication, and these plots are typically arranged in a grid of tens to hundreds of plots. The collection, aggregation, and analysis of replicated plots provides a common use case for collaboration and standardization. Shared software can provide a foundation for the standard formats that will make it easier to share and synthesize the results of previous research.

1.2 Existing Solutions

There are an increasing number of end-to-end pipelines that support plot based field research in agriculture and ecology. Many are commercial, though a few are open source. We want to work with and support the entire community. We focus on open source because of its role in making science more transparent, and because it is easier for a scientific community to collaborate around building open source software.

Many large research institutions have licences to powerful commercial products. Commonly used commercial software include Metashape[*] and pix4D[†] for photogrammetry (to create orthomosaics and point clouds), and GIS software such as ArcGIS Pro[‡] for geospatial analysis. A single license may cost thousands of dollars per year. A greater cost is that licences present a hurdle to reproducibility, and proprietary algorithms are not transparent. But such commercial software is widely used and work well. They can generate standard data formats, have programming interfaces, and some capacity for scaling and automating pipelines. Furthermore they have simple, well tested graphical interfaces making them easier to get started with for new users, and all three are configured to make use of GPUs to speed processing.

Many open source solutions exist; OpenDroneMap[§] for photogrammetry and QGIS[5][¶] for geospatial analysis

---

[*]Agisoft, https://www.agisoft.com/

[†]Pix4D, https://www.pix4d.com/

[‡]ESRI, https://esri.com/

[§]OpenDroneMap https://www.opendronemap.org/

[¶]QGIS https://qgis.org

are among the most popular. To illustrate the role of a scientific community in open source development: while OpenDroneMap does not currently support GPU computation, anyone is welcome to implement this feature[*].

Next we describe a few projects that have been specifically developed to support plot based experimental field research, in particular measuring plant phenotypes at the plot scale for research on agronomic management, crop improvement, and ecosystem dynamics.

Many of these tools provide an interface for creating and naming plots following a regular grid. Typically the user specifies plot dimensions and field layout, then identifies the corners of a research field, and the software generates a set of polygons that are used to clip the images. We have used UASTools[6][†] - an open source (GPL-2) standalone R package that provides a standalone solution for this particular step, but below we focus on end-to-end pipelines.

FIELDimageR[‡][7] is an R based interface with clear documentation that walks the user from image collection through the use of OpenDroneMap for photogrammetry and analyses using R based geospatial libraries. FIELD-ImageR is licensed with an open source license (GPL-2). For researchers with programming experience, and R users in particular this makes automating and modifying a pipeline accessible. The flexibility of this pipeline is demonstrated in its application to counting corn kernels and pollen as easily as it counts plants in a plot.

Another open source solution is ImageBreed.[8] ImageBreed uses ontologies and standards to facilitate interoperability. For example, it provides a BrAPI-compliant interface to its underlying relational database. This means that as long as the pipeline supports BrAPI, we can push and pull data from this interface. Individual scripts used in the pipeline are provided in a separate repository[§] with a permissive open source license (MIT) that facilitate re-use and collaborative development.

Progeny Drone is a commercial solution developed for the analysis of plot based breeding trials[9, 10][¶]. Progeny makes some field-research specific choices that simplify and improve analyses. First, it performs plot-clipping and analyses on the original images rather than the orthomosaic of images stitched together. This reduces geometric artifacts from the pixel-deforming orthomosaicing step, and it effectively provides multiple subsamples on each plot. This assumption saves time because aligning photos requires less computing power than stitching them together as an orthomosaic.

Many approaches to workflow management exist, from a scripted sequence of operations to complex webs of interactivity. Scientific workflow systems have been reviewed elsewhere including tools focused on data intensive scientific workflows[11] and the workflow functionality of desktop Geospatial Information System (GIS) software such as ArcGIS and QGIS.[5] In seeking to develop a framework that provides workflow functionality while maintaining independence of components, we are building our workflows with Makeflow.[12] Makeflow provides the features we need without the overhead of more complex systems. Like many of the workflow systems, Makeflow clearly defines the workflow - a map of inputs, processing steps, and outputs as a text file.

1.3 Our Solution and the Gaps it Fills

Given that the software described above provide so much functionality, what can DPP contribute? We are developing DPP as a framework that adds features aimed at improving reproducibility, reuseability, and collaboration[13][‖]. While we are developing a core set of algorithms that demonstrate our implementation best practices, the real value of this framework will be in leveraging the larger ecosystem of algorithms developed for drone image processing and analysis of field experiments.

Key Features of DPP include:

Portable: can be used and developed on any operating system.
Flexible: can be modified for different needs.

---

[*]OpenDroneMap Issue 92 "Image Processing using GPU" https://github.com/OpenDroneMap/ODM/issues/92

[†]UASTools https://github.com/andersst91/UAStools

[‡]FIELDImageR https://github.com/filipematias23/FIELDimageR

[§]ImageBreed Drone Image Scripts https://github.com/solgenomics/DroneImageScripts

[¶]Progeny Drone https://www.plotphenix.com/

[‖]Drone Processing Pipeline: https://osf.io/xdkcy

Automated: runs to completion with minimal intervention.
Open: supports open, transparent science and is available for reuse and extension.
Modular: components can be reused in different pipelines.
Scalable: works on laptops as well as HPC.
Reproducible: published data and software required to regenerate and build upon results.

2. IMPLEMENTATION

2.1 Supporting Best Practices

The adoption of best practices supports the needs of researchers and contributors. In this section we describe the best practices that are used to develop the pipeline and are reinforced by the use of templates. The foundation for our approach is outlined in "Best Practices for Scientific Computing."[14] Reproducibility is at the foundation of scientific discovery, and software tools are making it easier than ever for scientists to share and build on each other's computational work.[15] Using best practices in software development can provide a crucial component for reproducibility through adherence to standards, version control, and continuous integration.

Providing minimum metadata is necessary to facilitate re-use of data outputs. One approach to enhancing reproducibility is to use code templates to simplify the process of adding new algorithms to a pipeline. This approach addresses a common challenge that we have faced in enabling researchers to add new functionality to complex workflows in a way that's scalable, reliable, and consistent with the other software and data components. Code templates allow for the inclusion of best practices and standards into processes with minimal cost to the contributor.

When a scientist developing an algorithm that takes an image and calculates a metric, the scientist should be able to focus on transforming an array into a meaningful metric. Time spent writing code to automate and scale an analysis would be better spent on using their skills to test and document the algorithm itself. Details about the experimental design and other metadata required to subset and interpret the data should come from the scientist who is collecting these data.

We adopted the concept of templates following our experience developing the TERRA REF pipeline.[16] We found that getting an algorithm from proof of concept to reliable performance in an automated workflow required many steps from software engineering that may be unfamiliar to academic software authors. Much of the information required to clearly define outputs, to document software, and to write tests must come from the scientist or algorithm developer. A software engineer can use templates and tests to define these requirements and provide guidance on what to test and document, and what data formats the algorithm should accept and generate. In our initial development of templates for the TERRA REF project the use of targeted code templates made it much easier for contributors to add new algorithms to the pipeline. Future plans are to simplify our templating software by migrating to the Cookiecutter framework.[17]

2.2 Our Vision For Shared Field Phenomics Pipelines

We have experience developing pipelines that integrate heterogeneous scientific software and data into automated workflows. For example, the PEcAn Project[18,19] combines many types of data with many different crop and ecosystem models within a consistent framework for statistical analysis and prediction. The TERRA REF data processing infrastructure[16] converts data from multiple sensors into plot level phenotypes, and makes these data accessible to different research communities through a variety of portals.

We believe that creating an ecosystem of freely shared inter-operable components is in the best interest of science, and our objective is to develop a framework that allows scientists and engineers from multiple institutions to more effectively collaborate and share their work. Researchers will benefit from a portable, flexible, reproducible, open, automated, modular and scalable framework.

Our current work focuses on a number of key improvements over the TERRA REF pipeline, starting with reducing complexity and removing dependencies. One focus has been to generalize the workflow components and make them more portable. The TERRA REF pipeline was developed around one-of a kind robotics and sensing platform to generate high resolution data that is orders of magnitude larger and higher resolution than

what a typical drone user will collect. In order to make the framework more broadly useful we greatly simplified the TERRA REF pipeline. For example, we replaced the requirement for three databases with one requirement for files on a hard-drive; we also replaced the complex workflow that relied on a message queue, custom scripts, and logging databases with a single workflow engine.

The workflow engine we are using for scaling is Makeflow.[12] We are also making it possible to deploy the pipeline in one step - just as easily on a desktop as on an HPC leveraging tools including Docker.[20]

2.2.1 APIs, Standards, Formats, Conventions and all that

One key feature that has worked very well is defining standard data inputs and outputs for each stage of a workflow. This is why data and software engineers are so fond of "Application Programming Interfaces (APIs)" - the basic idea is that no matter what your data look like, as long as you can coerce it into an agreed structure and format, and develop software that reads and/or writes to the same specification, heterogeneous data and software can be combined into complex pipelines.

It is easy to agree that standards are great, but when there are so many scientific domains covered by a single pipeline, no single standard will work. Our view is that it is less important to agree on standards than it is to start with some convention and avoid generating ad-hoc data types for each pipeline or project. This reduces the number of data types we need to convert among. Here we highlight a few standards that are used by one or more communities, and that we are planning to support:

Climate Foreacst (CF) metadata conventions and NetCDF-CF file format[21] for geophysical data such as weather and land surface processes.

Open Geospatial Consortium (OGC) - provides many standards. We use GeoTIFF image formats[22] for images and other gridded data. For vector data we handle these internally as OGC compliant geojson[23] files but support many other formats, including shape (.shp) files because they are so widely used*.

Minimum Information about Plant Phenotyping Experiments (MIAPPE v1.1)[24] defines information required to support interpretation and reuse of shared phenotyping data.

Breeder's API (BrAPI): a specification for interoperability among breeder's databases.[25] BrAPI can provide all information required by MIAPPE.

AgMIP Crop Experiment / Model Output data schema (ACE / ACMO): an interchange format for crop model inputs and outputs used by the Agricultural Model Intercomparison Project;[26] these formats provide a agronomic, experimental, and environmental metadata that is not currently covered by BrAPI.

2.2.2 Pipeline Components

The two main components of the Drone Processing Pipeline are the workflow manager and Transformers. Makeflow provides the workflow engine. Docker bundles the Transformers along with Makeflow and the workflow definition in a single container image. Transformers are atomic components that analyze a standardized input and generate a standardized output. But a Transformer also encapsulates expectations for standardization and best practices.

The main driver behind any automated processing pipeline is the workflow manager - the part of the workflow that runs the processes either sequentially or asynchronously. Additional desired features for a workflow manager are: Ability to use distributed resources, automatic error detection and recovery, runtime monitoring, and ease of workflow changes. Features that provide added benefit are the ability to handle dynamic worker availability, and the categorization of workers providing for the fine tuning of processing loads. The Drone Processing Pipeline is using Makeflow.[12]

Docker makes Transformers more portable, and makes it easier to develop and deploy individual components.[20] Docker containers provide Linux environments, but can run on Windows, MacOS or Linux operating system. Docker containers can also be converted to Singularity containers that run on high performance computers.[27] Docker encapsulates computing environments into portable containers that can be shared and archived. Publishing versioned Docker containers enables provenance tracking and computational reproducibility.[20]

---

*Other OGC compliant formats: Well Known Text, sqlite geopackages, PostGIS geometry

The Makeflow file defines the network graph of processing that is to occur. This allows each Transformer to be easily organized by a higher-level process or script into a sequence required to support a particular research objective. Additional tasks such as creating directories and inserting data into a database can also be incorporated into a workflow.

We use Scientific Filesystem (SCIF)[27] recipes to specify a distinct Entry Point and Environment for each Transformer bundled in the final container image, as well as an Entry Point to run the Makeflow workflow. This means that an entire workflow, including all the algorithms can be published and archived as a functioning single container image. This is good for reproducibility and simplifies the developer experience for workflow authors.

3. CONCLUSION

The development of a freely available processing pipeline that is open, automated, portable, flexible, and is easy to adjust and enhance fills an identified need for Agricultural researchers using drones. Developing a processing pipeline using best practices from the software industry provides a robust versioned solution that allows repeated experimentation, and reproduction of found results. There are substantial engineering challenges to setting up a framework, and we expect to increase the ratio of science to software engineering. The design approaches outlined in this paper not only simplify the development of new algorithms, but also promote the distribution and use of algorithms by providing standardized interfaces.

ACKNOWLEDGMENTS

REFERENCES

[1] Cruzan, M. B., Weinstein, B. G., Grasty, M. R., Kohrn, B. F., Hendrickson, E. C., Arredondo, T. M., and Thompson, P. G., "Small unmanned aerial vehicles (micro-uavs, drones) in plant ecology," *Applications in Plant Sciences* **4**, 1600041 (sep 2016).

[2] Shi, Y., Thomasson, J. A., Murray, S. C., Pugh, N. A., Rooney, W. L., Shafian, S., Rajan, N., Rouze, G., Morgan, C. L., Neely, H. L., Rana, A., Bagavathiannan, M. V., Henrickson, J., Bowden, E., Valasek, J., Olsenholler, J., Bishop, M. P., Sheridan, R., Putman, E. B., Popescu, S., Burks, T., Cope, D., Ibrahim, A., McCutchen, B. F., Baltensperger, D. D., Avant, R. V., Vidrine, M., and Yang, C., "Unmanned aerial vehicles for high-throughput phenotyping and agronomic research," *PloS one* (2016).

[3] Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., et al., "The FAIR guiding principles for scientific data management and stewardship," *Scientific data* **3** (2016).

[4] Petersen, R. G., [*Agricultural Field Experiments: Design and Analysis*], CRC Press (1994).

[5] Dobesova, Z., "Data flow diagrams in geographic information systems: a survey," *Proceeding of 14th SGEM GeoConference on Informatics, Geoinformatics and Remote Sensing* , 705–712 (2014).

[6] Anderson, S. L., Murray, S. C., Malambo, L., Ratcliff, C., Popescu, S., Cope, D., Chang, A., Jung, J., and Thomasson, J. A., "Prediction of maize grain yield before maturity using improved temporal height estimates of unmanned aerial systems," *The Plant Phenome Journal* **2**(1) (2019).

[7] Matias, F. I., Caraza-Harter, M. V., and Endelman, J. B., "FIELDImageR: An R package to analyze orthomosaic images from agricultural field trials," *The Plant Phenome Journal* **3**(1), e20005 (2020).

[8] Morales, N., Kaczmar, N. S., Santantonio, N., Gore, M. A., Mueller, L. A., and Robbins, K. R., "ImageBreed: Open-access plant breeding web–database for image-based phenotyping," *The Plant Phenome Journal* **3**(1), e20004 (2020).

[9] Hearst, A. A., *Remote sensing of soybean canopy cover, color, and visible indicators of moisture stress using imagery from unmanned aircraft systems*, PhD thesis, Purdue University Graduate School (2019).

[10] Moreira, F. F., Hearst, A. A., Cherkauer, K. A., and Rainey, K. M., "Improving the efficiency of soybean breeding with high-throughput canopy phenotyping," *Plant Methods* **15**(1), 139 (2019).

[11] Liu, J., Pacitti, E., Valduriez, P., and Mattoso, M., "A survey of data-intensive scientific workflow management," *Journal of Grid Computing* **13**(4), 457–493 (2015).

[12] Albrecht, M., Donnelly, P., Bui, P., and Thain, D., "Makeflow: A portable abstraction for data intensive computing on clusters, clouds, and grids," in [*Proceedings of the 1st ACM SIGMOD Workshop on Scalable Workflow Execution Engines and Technologies*], 1–13 (2012).

[13] Schnaufer, C., LeBauer, D. S., and Pistorius, J., "UA Open Drone Processing Pipeline (alpha)," (May 2020).

[14] Wilson, G., Aruliah, D. A., Brown, C. T., Chue Hong, N. P., Davis, M., Guy, R. T., Haddock, S. H., Huff, K. D., Mitchell, I. M., Plumbley, M. D., Waugh, B., White, E. P., and Wilson, P., "Best practices for scientific computing," *PLoS Biology* (2014).

[15] Stodden, V., McNutt, M., Bailey, D. H., Deelman, E., Gil, Y., Hanson, B., Heroux, M. A., Ioannidis, J. P., and Taufer, M., "Enhancing reproducibility for computational methods," *Science* **354**(6317), 1240–1241 (2016).

[16] Burnette, M., Rohde, G. S., Fahlgren, N., Sagan, V., Sidike, P., Kooper, R., Terstriep, J. A., Mockler, T., Andrade-Sanchez, P., Ward, R., Maloney, J. D., Willis, C., Newcomb, M., Shakoor, N., and LeBauer, D., "TERRA-REF data processing infrastructure," in [*ACM International Conference Proceeding Series*], (2018).

[17] Pierzina, R., "Cookieutter: Better project templates." <https://github.com/cookiecutter/cookiecutter> (2020).

[18] Fer, I., Gardella, A. K., Shiklomanov, A. N., Serbin, S. P., De Kauwe, M. G., Raiho, A., Johnston, M. R., Desai, A., Viskari, T., Quaife, T., et al., "Beyond modeling: A roadmap to community cyberinfrastructure for ecological data-model integration," (2020).

[19] LeBauer, D. S., Wang, D., Richter, K. T., Davidson, C. C., and Dietze, M. C., "Facilitating feedbacks between field measurements and ecosystem models," *Ecological Monographs* **83**(2), 133–154 (2013).

[20] Boettiger, C., "An introduction to Docker for reproducible research," *ACM SIGOPS Operating Systems Review* **49**(1), 71–79 (2015).

[21] Eaton, B., Gregory, J., Drach, B., Taylor, K., Hankin, S., Caron, J., Signell, R., Bentley, P., Rappa, G., Höck, H., et al., "NetCDF Climate and Forecast (CF) metadata conventions," (2003).

[22] Ritter, N., Ruth, M., Grissom, B. B., Galang, G., Haller, J., Stephenson, G., Covington, S., Nagy, T., Moyers, J., Stickley, J., et al., "Geotiff format specification geotiff revision 1.0," *SPOT Image Corp* **1** (2000).

[23] Butler, H., Daly, M., Doyle, A., Gillies, S., Hagen, S., Schaub, T., et al., "The geojson format," *Internet Engineering Task Force (IETF)* (2016).

[24] Papoutsoglou, E. A., Faria, D., Arend, D., Arnaud, E., Athanasiadis, I. N., Chaves, I., Coppens, F., Cornut, G., Costa, B. V., Ćwiek-Kupczyńska, H., et al., "Enabling reusability of plant phenomic datasets with MIAPPE 1.1," *New Phytologist* (2020).

[25] Selby, P., Abbeloos, R., Backlund, J., Basterrechea Salido, M., Bauchet, G., Benites-Alfaro, O., Birkett, C., Calaminos, V., Carceller, P., Cornut, G., Vasques Costa, B., Edwards, J., Finkers, R., Yanxin Gao, S., Ghaffar, M., Glaser, P., Guignon, V., Hok, P., Kilian, A., König, P., Lagare, J., Lange, M., Laporte, M.-A., Larmande, P., LeBauer, D., Lyon, D., Marshall, D., Matthews, D., Milne, I., Mistry, N., Morales, N., Mueller, L., Neveu, P., Papoutsoglou, E., Pearce, B., Perez-Masias, I., Pommier, C., Ramírez-González, R., Rathore, A., Raquel, A., Raubach, S., Rife, T., Robbins, K., Rouard, M., Sarma, C., Scholz, U., Sempéré, G., Shaw, P., Simon, R., Soldevilla, N., Stephen, G., Sun, Q., Tovar, C., Uszynski, G., and Verouden, M., "BrAPI-an application programming interface for plant breeding applications," *Bioinformatics (Oxford, England)* **35**(20) (2019).

[26] Porter, C. H., Villalobos, C., Holzworth, D., Nelson, R., White, J. W., Athanasiadis, I. N., Janssen, S., Ripoche, D., Cufi, J., Raes, D., et al., "Harmonization and translation of crop modeling data to ensure interoperability," *Environmental modelling & software* **62**, 495–508 (2014).

[27] Sochat, V., "The Scientific Filesystem," *GigaScience* **7**(5), giy023 (2018).