

The SDSS Data Archive Server

Eric H. Neilsen, Jr.

Fermi National Accelerator Laboratory

The Sloan Digital Sky Survey (SDSS) Data Archive Server (DAS) provides public access to data files produced by the SDSS data reduction pipeline. This article discusses challenges in public distribution of data of this volume and complexity, and how the project addressed them.

The Sloan Digital Sky Survey (SDSS)¹ is an astronomical survey of covering roughly one quarter of the night sky. It contains images of this area, a catalog of almost 300 million objects detected in those images, and spectra of more than a million of these objects. The catalog of objects includes a variety of data on each object. These data include not only basic information but also fit parameters for a variety of models, classifications by sophisticated object classification algorithms, statistical parameters, and more. If the survey contains the spectrum of an object, the catalog includes a variety of other parameters derived from its spectrum.

Data processing and catalog generation, described more completely in the SDSS Early Data Release² paper, consists of several stages: collection of imaging data, processing of imaging data, selection of spectroscopic targets from catalogs generated from the imaging data, collection of spectroscopic data, processing of spectroscopic data, and loading of processed data into a database. Each of these stages is itself a complex process. For example, the software that processes the imaging data determines and removes some instrumental signatures in the raw images to create “corrected frames,” models the point spread function, models and removes the sky background, detects objects, measures object positions, measures the radial profile and other morphological parameters for each object, measures the brightness of each object using a variety of methods, classifies the objects, calibrates the brightness measurements against survey standards, and produces a variety of quality assurance plots and diagnostic tables. The complexity of the spectroscopic data reduction pipeline is similar. Each pipeline deposits the results in a collection of files on disk.

The Catalog Archive Server (CAS) provides an interface to a database of objects detected through the SDSS along with their properties and observational metadata. This serves the needs of most users, but some users require access to files produced by the pipelines. Some data, including the corrected frames (the pixel data itself corrected for instrumental signatures), the models for the point spread function, and an assortment of quality assurance plots, are not included in the database at all. Sometimes it is simply more convenient for a user to read data from existing files than to retrieve it using database queries. This is often the case, for example, when a user wants to download data a significant fraction of objects in the database. Users might need to perform analysis that requires more computing power than the CAS database servers can reasonably provide, and so need to download the data so that it can be analyzed with local resources. Users can derive observational parameters not measured by the standard SDSS pipeline from the corrected frames, metadata, and other data products, or simply use the output of tools with which they’re familiar.

The challenge in distributing these data is lies not in the distribution method itself, but in providing tools and support that allow users to find the data they need and interpret it properly. After introducing the data itself, this article describes how the DAS uses ubiquitous and well understood technologies to manage and distribute the data. It then discusses how it addresses the more difficult problem of helping the public

find and use the data it contains, despite its complexity of its content and organization.

The Pipelines

The SDSS uses two instruments to collect data: an imaging camera and a pair of multi-object spectrographs. The spectrographs collect light from the focal plane on the telescope through 640 fiber optic cables mounted in a metal plate. The holes in the plate in which the fiber optic cables are mounted must be drilled based on the known positions of the objects whose spectra will be collected; creating the plate requires an existing catalog of objects. We therefore begin observation of a portion of the sky by collecting images from the imaging camera. From the images collected, we then create catalogs of objects, select targets for spectroscopy, and design plates for observing these objects with the spectrographs. Finally, we collect the spectra themselves after the plates are drilled and the area of sky is well positioned for observation.

The SDSS imaging camera collects data through an array of 30 charge-coupled devices (CCDs) arranged on an array of five rows and six columns. A different filter covers each row of CCDs, so each column collects light in five different wavebands. The camera collects data in time delay and integrate (TDI) mode, so during an observation the telescope moves such that its field of view slowly scans across the sky. The image of a position on the sky first appears at the top of a CCD in the first row, and moves down the CCD with time. As the image of that position moves down the CCD, the camera controller moves charge down the CCD, continuously reading charge at the bottom of the CCD. When the image of a position reaches the bottom of the CCD, the camera controller reads the charge generated by it; the charge corresponds to the value of a pixel covering that point on the sky. A data acquisition system packages these pixels into sequences of frames, such that each frame is effectively an image of a different field in a contiguous row of fields on the circle. Because the SDSS camera collects data with 30 CCDs, a data set from an imaging run consists of 30 sequences of frames. Each row sees the sky through a different filter, so we can identify each CCD by its filter and camera column. The camera therefore produces a collection of images that we can uniquely identify by run number, filter, camera column, and field number.

Initially, the observatory in Sun Spot, New Mexico recorded these data on tape and shipped them to Fermilab for processing, but more recently an automated process transfers the data to Fermilab over the Internet. The imaging pipeline software then uses these data to produce several files with reduced data in each field. From an end user's viewpoint, the most interesting are the

1. *corrected frames*, which show the images after we remove basic instrumental effects and contaminated pixels from bad CCD columns;
2. *mask frames*, which identify the pixels in the corrected frames that we interpolated; and
3. *object catalogs*, which are the tables of objects identified in the field and a wide variety of properties measured from the images.

In addition to producing these files for each field, the pipeline produces a variety of other files, including tables of calibration parameters for each field in the run.

The data processing team can run imaging data through the pipeline multiple times—for example, to apply a new version of the software to previously reduced data. Unique identification of an individual output file therefore requires not only the information needed to identify the raw data from which it was derived (run, camera column, filter, and field) but also a rerun number, which identifies the pipeline application that produced the file.

An analogous set of parameters specifies specific files produced by the spectroscopic

pipeline. These include the plate number and the Modified Julian Date (MJD, the number of days since November 17, 1858) of the night on which the observatory collected the raw data. The spectroscopic pipeline also produces several files for each object, including plots of the spectra and tables containing the individual object's spectrum and the parameters derived from it. These files can be identified by the plate, MJD, and fiber number.

As in the case of the imaging pipeline, the data processing team can run the spectroscopic pipeline multiple times, so a rerun number is required to identify a specific file instance.

The data-processing pipeline stores files in a directory structure reflecting the input data and file types. The data set is too large to fit on a single server's disks, so we use a directory of symbolic links to data on Network File System (NFS) mounted disks to make the data set appear in a unified structure.

Public Data Access

The CAS fills most users' needs, and the many high-level interfaces help them avoid navigating the data collection details behind the file organization. However, sometimes users need access to one or more of the files the pipeline produces, rather than just the subset of this data loaded into the CAS database.

For example, many users desire access to the corrected frame from the imaging camera, which isn't included in the CAS. Although the CAS supplies a variety of properties for each object, many scientific applications require analysis of the images themselves. (The CAS Web services do produce JPEG images of the sky constructed from SDSS corrected frames. The lossy compression and limited precision used in the construction of JPEG images mean that they contain only a small fraction of the information in the corrected frames, and are not useful for scientific analysis.) The SDSS pipeline stores its data in the Flexible Image Transport System³ (FITS) format. Although relatively obscure outside the astronomical community, the FITS format is ubiquitous within it, and almost all tools written by or for astronomers support it. The FITS format derives its name from its original use, the transport of images. The standard now includes a variety of extensions supporting the storage of tables and other data; the full name of the format is now deceptive. The DAS stores not only images in FITS format, but also object catalogs, metadata, and most other information is well.

Even for simple visual inspection, studying the corrected frames using any of several commonly available FITS viewers has some advantages. For example, users of these tools can zoom the data up to the instrument resolution's limit and remap the dynamic range to show specific features of interest without reaching the limits imposed by a compressed image format; users can even compare images from all five bands to each other and to external data sets. Perhaps even more important, users can perform their own analysis on the frames using their own tools, independent of the object measurements derived by the pipeline. This not only lets them customize the algorithms used and parameters measured but also verify the pipeline's results against those they calculate themselves. Questions sent to the help desk reveal that a number of users do attempt this, but that the complexity of the pipeline is often an obstacle. With assistance from the helpdesk, these users generally meet with success. (The number of users who achieve success without help, or who give up without ever contacting the helpdesk, is unknown.)

The Data Archive Server

The DAS's basic functionality is very simple: Apache and rsync daemons on the server

let users download files using either HTTP or rsync. They can download individual files using their favorite Web browser or supply lists of files to `wget` or similar HTTP clients to perform mass downloads. We can't, however, simply point these servers at the root of the data-processing factory's data directory structure. Each data release corresponds to a specific subset of our processed data that has been completely processed through our pipelines, meets survey data quality standards, and has been handled by specific versions of the data-processing software. In some cases, imaging runs have some frames that are part of the data release but others that aren't due to data quality variability over the run.

Each data release, therefore, has its own directory tree. This tree resembles the one used by the data factory's disks, but the directories only contain links to the specific files that are part of that data release. Serving roughly 20 million files with a directory structure of symbolic links introduces several complications, so DAS developers created a collection of scripts to maintain the link trees. The first of these builds the directory structure based on a set of files that specify the data release's content. Separate scripts periodically verify link validity and check the linked file's content by using recorded checksums. Because the creation of 20 million symbolic links is much more time-consuming than creating and restoring compressed dumps of disk partitions, we store each data release's link tree in its own partition and use these partition images to replicate the directory structure on backup servers. We create archival backups even though we could regenerate the tree from scratch if necessary, again because restoration of the partition is much faster than regeneration if the link tree. (We currently use ext3 as our file system. In the past we have experimented with XFS, but in neither case was the time it took to regenerate the directory tree comparable to a simple restoration of the partition.)

The other complication lies in the data's organization. Although most CAS users don't need to know the particulars behind specific observations, these parameters are essential for navigating the DAS directory structure. To obtain imaging data, a user must determine the run, field, filter, and camera column of the data desired; if the user wants spectroscopic data, he or she must determine the plate, observation date, and fiber number used to observe the object. The user can then construct the URLs of the files of interest from these parameters.

A user with a list of objects of interest will not in general know these parameters for the objects; it is much more likely that he or she will know the coordinates of the objects instead. The DAS provides a tool, the Footprint server (shown in Figure 1), that allows a user to determine these parameters for a list of coordinates. A second form (see Figure 2) takes user-supplied observation parameters and uses them to construct lists of files in a format suitable for use with `wget` or `rsync`. The CAS itself provides a more versatile method for obtaining URLs: an object's "explore" view supplies details about the observation in which it resides along with links to the most commonly desired files in the DAS. For users who want large numbers of data files, SQL queries can return lists of URLs for those files.

Location Edit View Go Bookmarks Tools Settings Window Help

Location: <http://das.sdss.org/DR6-cgi-bin/FOOT>

SDSS Footprint Server: DR6

Use the footprint server to determine if a given (list of) RA/DEC coordinates are in the current data release. If present, the run,rerun,camcol,field of the image of the requested input coordinates are given

List of Coordinates

List of ra, dec

ra, dec
256.443154, 58.0255
29.94136, 0.08930

or select file to upload.

Output Content

Imaging (target version) run/rerun/camcol/field/rowc/colc

Imaging (best version) run/rerun/camcol/field/rowc/colc

Figure 1. The footprint server form. This form lets a users upload a list of coordinates, and returns a table listing which SDSS images cover these coordinates, if any.

Location Edit View Go Bookmarks Tools Settings Window Help

Location: <http://das.sdss.org/DR6-cgi-bin/DAS>

SDSS Data Archive Server: DR6

Available Image Products Available Spectro Products DAS Users' Guide

or, if you want to start again,

Step 1. Select one or more data products.

Data Product	File Description	run	rerun	camCol	filter	field	plate	MJD	fiberID	spRerun(*)
<input type="checkbox"/> fpC	Corrected imaging frame	X	X	X	X	X				
<input type="checkbox"/> fpBIN	Binned imaging frame	X	X	X	X	X				
<input type="checkbox"/> fpM	Mask frame	X	X	X	X	X				
<input type="checkbox"/> fpObjc	Catalog of Imaging Objects	X	X	X		X				
<input type="checkbox"/> fpAtlas	Atlas Images for a field	X	X	X		X				
<input type="checkbox"/> targetTsField	Target Summary for a field, target version	X	X	X		X				
<input type="checkbox"/> targetTsObj	Calibrated catalog for a field, target version	X	X	X		X				
<input type="checkbox"/> bestTsField	Target Summary for a field, best version	X	X	X		X				
<input type="checkbox"/> bestTsObj	Calibrated catalog for a field, best version	X	X	X		X				
<input type="checkbox"/> tsObjFromMap	Calibrated catalog for targetted objects							X	X	
<input type="checkbox"/> spPlate	extracted, calibrated spectra for a plate							X	X	x
<input type="checkbox"/> spZbest	best fit spectroscopic classifications and redshifts							X	X	x
<input type="checkbox"/> spZall	all fit spectroscopic classifications and redshifts							X	X	x
<input type="checkbox"/> spZline	line fits							X	X	x
<input type="checkbox"/> ssppPostParam	stellar parameters pipeline output							X	X	x
<input type="checkbox"/> ssppParam CSV	stellar parameters pipeline output CSV with additional photo data							X	X	x
<input type="checkbox"/> ssppLines CSV	spectral line equivalent widths from stellar parameters pipeline							X	X	x
<input type="checkbox"/> spSpec	One extracted, calibrated spectrum, with 1d parameters							X	X	X
<input type="checkbox"/> spPlotGif	Plot of one spectrum (.gif)							X	X	X
<input type="checkbox"/> spPlotPs	Plot of one spectrum (.ps.gz)							X	X	X
<input type="checkbox"/> asTrans	Astrometric calibration	X	X							
<input type="checkbox"/> nfCalib	Photometric calibration	X	X	X						
Data Product	File Description	run	rerun	camCol	filter	field	plate	MJD	fiberID	spRerun(*)

(*) spRerun note: This parameter defaults to the best available spectro rerun. It does not usually need to be given.

Step 2. Create the "upload file" which describes the data you want to be downloaded by listing the paramters required, indicated in the table of Step 1. The filename extension must be one of the following: .csv, .par, .tbl, .fit or .xml. Please refer to the [file specification section of the DAS users' guide](#) for details on the file formats.

Step 3. Upload the object list, specifying the parameters required for the data products selected above.

Either select file to upload (created in Step 2)

or enter the CSV information below.

```
run,rerun,camcol,filter,field
1339,40,5,gri,75
```

Figure 2. The Data Archive Server (DAS) form This form accepts a list of parameters from a user specifying a data set of interest, and returns a list of files in a format usable by wget or rsync. One of these programs can then be used to download the entire dataset.

Usage

Once users obtain the files, they must be able to read them and properly interpret their contents. We distribute all files either in ASCII or FITS format; most of the tools astronomers use to analyze data can read these formats. Locating the correct data in each file and using it to derive scientific results requires additional information. The SDSS data release web sites (see <http://www.sdss.org/science>) include descriptions of the algorithms used during data processing along with documentation of the contents of each file type produced. Finally, users can submit questions to the SDSS help desk at

<http://www.lib.uchicago.edu/e/crerar/using/sloan.html>.

Figure 3 shows the volume of data downloaded from the DAS per week over the past year. This value varies significantly week to week (even day to day) as a result of small-number statistics. Although many users download only a few files, some sites download large volumes of data, so the total volume transferred depends on how many of these large transfers occurred.




Figure 3. Data Archive Server bytes transferred per week. The total volume of data transferred varies dramatically week to week, because small numbers of very large transfers make up a significant fraction of the total volume.

Institutions inside and outside the SDSS collaboration have used the DAS to download local copies of SDSS data. In addition to institutions that download the data for local reference, others retrieve the data for redistribution in a different form. For example, the NASA Extragalactic Database (<http://nedwww.ipac.caltech.edu>) used the DAS interface to include the SDSS catalog in their aggregation of astronomical databases. In other cases, users download a much smaller quantity of data, sometimes only the image or spectrum of an individual object, or an image of a single field.

Conclusions

A wide variety of people use SDSS data, ranging from curious children to students to professional astronomers, and within each level of experience the SDSS data can serve a variety of purposes. A single interface to the data cannot easily fill the needs of all users, for all purposes. Therefore, the SDSS collaboration provides a variety of interfaces to the data: the visual exploration tools and forms provided by the CAS support basic exploration and common queries, the SQL interface provided by the CAS and casjobs allows more complex queries, and the DAS provides access data required to do analysis beyond what can be supported by the contents of the DAS database.

Of the various data access methods, the DAS interface is the most traditional from the point of view of a professional astronomer, and bears the greatest similarity to other archives of astronomy data; users determine what observations they need using web forms, and download FITS files with data based on those observations. We support this data distribution method using simple, well understood technologies such as http, and a modest collection of administrative scripts and web forms. Users, however, require greater resources and technical sophistication to use the data; at the very least, they will require software capable of reading FITS files and generating the results from it that they desire. This is, however, not only acceptable for the DAS, but a strength; the DAS is primarily intended for users for which the CAS interface and catalog are inadequate, and so wish to apply their own analysis and tools to the data. Because the FITS standard is well established in the astronomical community, this software is much more likely to support FITS files than any other format.

The complexity of the reduction performed by the SDSS pipeline and the presence of the online relational database introduce some features that distinguish the DAS from other data archives that distribute FITS files. First, the CAS interface can support the use of the files by helping the user locate files of interest. The DAS distributes not only the basic data and calibration information, but a wide variety of additional files produced by the data reduction pipeline. Although users of astronomy software that supports FITS should be able to read these files, properly interpreting them requires greater dependence of documentation and sometimes support from a staffed help desk. By supplying several mutually supporting data access methods, the SDSS collaboration supports highly technical uses of its data without requiring a corresponding advanced level of understanding from the general user. A large data archive should not be reluctant to distribute highly complex data to the public merely because it cannot be made accessible to most users, or because it is impractical to present it using a user friendly interface. By taking advantage of existing standards, the database that serves the higher level data, ubiquitous and easy to maintain distribution technologies, and a modest collection of administrative scripts and web forms, the SDSS is able to distribute such data to the community that desires it.

Acknowledgments

Work supported by the U.S. Department of Energy under contract No. DE-AC02-07CH11359.

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web Site is <http://www.sdss.org/>.

The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington.

Eric H. Neilsen, Jr. is a computing professional at Fermi National Accelerator Laboratory. His technical interests include astronomical surveys, image processing, and data management. Neilsen has a Ph. D. in astrophysics from Johns Hopkins University. Contact him at neilsen@fnal.gov.

¹ D. York et al., "The Sloan Digital Sky Survey: Technical Summary," The Astronomical Journal, vol. 120, 2000, pp.1579-1587.

² C. Stoughton et al., "Sloan Digital Sky Survey: Early Data Release," The Astronomical Journal, vol. 123, no. 1, 2002, pp. 485-548.

³ R. Hanisch et al., "Definition of the Flexible Image Transport System (FITS)," Astronomy and Astrophysics, vol. 376, 2001, pp. 359-380.