

Utilizing Astroinformatics to Maximize the Science Return of the Next Generation Virgo Cluster Survey

Nicholas M. Ball

Abstract The Next Generation Virgo Cluster Survey is a 104 square degree survey of the Virgo Cluster, carried out using the MegaPrime camera of the Canada-France-Hawaii telescope, from semesters 2009A–2012A. The survey will provide coverage of this nearby dense environment in the universe to unprecedented depth, providing profound insights into galaxy formation and evolution, including definitive measurements of the properties of galaxies in a dense environment in the local universe, such as the luminosity function. The limiting magnitude of the survey is $g_{AB} = 25.7$ (10σ point source), and the 2σ surface brightness limit is $g_{AB} \approx 29$ mag arcsec $^{-2}$. The data volume of the survey (approximately 2.6 terabytes of images), while large by contemporary astronomical standards, is not intractable. This renders the survey amenable to the methods of astroinformatics. The enormous dynamic range of objects, from the giant elliptical galaxy M87 at $M(B) = -21.6$, to the faintest dwarf ellipticals at $M(B) \approx -6$, combined with photometry in 5 broad bands ($u^* g' r' i' z'$), and unprecedented depth revealing many previously unseen structures, creates new challenges in object detection and classification. We present results from ongoing work on the survey, including photometric redshifts, Virgo cluster membership, and the implementation of fast data mining algorithms on the infrastructure of the Canadian Astronomy Data Centre, as part of the Canadian Advanced Network for Astronomical Research (CANFAR).

1 Introduction

The Virgo Cluster is the nearest large cluster of galaxies to us, and as such provides a unique laboratory to study the properties of a dense environment in the local universe, and hence gain insights into galaxy formation and evolution. However, the

Nicholas M. Ball
National Research Council Herzberg Institute of Astrophysics, 5071 West Saanich Road, Victoria,
BC V9E 2E7, Canada. e-mail: nick.ball@nrc-cnrc.gc.ca

current state-of-the-art optical surveys of this region are the Virgo Cluster Catalogue [3], and the Sloan Digital Sky Survey (SDSS) [10], both of which are far inferior to the potential science return of applying modern survey instrumentation to a survey focusing on this particular patch of sky.

The Next Generation Virgo Cluster Survey (NGVS)¹, is a survey of 104 deg² that completely supersedes all previous optical surveys of this region. Utilizing the capabilities of the MegaCam camera on the 3.6m Canada-France-Hawaii Telescope, the area is being surveyed to a limiting magnitude of $g_{AB} = 25.7$ (10σ point source), and is detecting low surface brightness structures to an unprecedented depth of $g_{AB} \approx 29$ mag arcsec⁻² (2σ).

The survey will provide revolutionary improvements in measurements for a host of galaxy and other properties out to the cluster virial radius. The main science goals include: (1) A definitive measurement of the faint-end shape of the cluster galaxy luminosity function; (2) the characterization of galaxy scaling relations over a factor 10^7 in mass; (3) the study of stellar nuclei in galaxies and their connection to supermassive black holes; (4) the connection between the cluster, galaxies and the intracluster medium; and (5) the fossil record of star formation and chemical enrichment in dense environments. Similar to other large modern wide-field surveys, a great deal of further science will be possible, both in the optical, and in combination with the large number of surveys of the region at other wavelengths, from X-ray to radio. Much science is also possible for objects in the foreground of Virgo (e.g., the Kuiper Belt), and in the background (e.g., 2000+ galaxy clusters).

1.1 The Role of Astrominformatics Within the NGVS

The survey provides many opportunities to successfully deliver the promised science by utilizing the methods of astrominformatics² [1, 4]. At 2.6T, the data size is such that the survey is not in the petascale regime, in which handling the data volume itself is very much the subject of ongoing research, but it is also substantially larger than the data volume for which many of the traditional astronomical analysis tools are designed.

Therefore, *astrominformatics can be usefully employed on the survey to produce improved science results, without becoming the subject of the research itself.*

Examples of survey challenges that can be addressed utilizing astrominformatics as a tool include:

- Data distribution: The 2.6T dataset size is easily handled by the local infrastructure, but proper data access will still require mature database technology: the SDSS is of similar size, and justified extensive database development.
- Data processing: The large area and high resolution of the survey, combined with the large angles subtended by many galaxies, and 5 passbands, provides a

¹ https://www.astrosci.ca/NGVS/The_Next_Generation_Virgo_Cluster_Survey

² <http://www.ivoa.net/cgi-bin/twiki/bin/view/IVOA/IvoaKDDguide>

rich dataset of photometric and morphological information. Complex modelling, for example, non-axisymmetric Sérsic plus nuclear profiles, or simulations of objects, is justified, and will require large amounts of processing time.

- Object detection and classification: The objects in Virgo span an unprecedented dynamic range, from the giant elliptical galaxy M87 ($M(B) = -21.6$) to faint dwarf ellipticals at $M(B) \approx -6$. Objects subtend angles from fractions of an arc-second for point sources, to degree-scale. Many new low surface brightness features resulting from galaxy interactions are being revealed, which have an irregular, extended structure. This creates problems, e.g., SExtractor does not detect the dwarf ellipticals, and large galaxies prevent local background subtraction.
- Virgo Cluster membership: In the absence of spectroscopic data, one must deduce the probability of cluster membership from photometric data. These populate a high-dimensional space of 5 passbands (4 colours), and numerous morphological parameters. Although traditional criteria such as magnitude vs. surface brightness are quite effective, in general, algorithms that naturally deal with a higher dimensional space will be able to take advantage of the extra information provided by the survey (e.g., morphologically similar objects may have different colours).
- Visualization of results: Many objects within the survey subtend large angles, and as such ideally require the visualization of gigabyte-sized FITS images, possibly overlaid. Standard tools have difficulty coping with the 1.6G data files.

2 NGVS Survey Data

The final survey will consist of 2.6T of images supplied as FITS files. The raw survey data consist of images from the 36 CCD mosaic of the MegaCam camera. These files are pre-processed with the CFHT Elixir pipeline, to provide bias-subtraction, flat-fielding and de-darkening. The Elixir files are then processed by two independent survey pipelines, MegaPipe [6, 7], and TERAPIX [2]. MegaPipe is used to produce stacked images of each survey field, by calibrating each CCD exposure to a high photometric and astrometric precision, and combining the images. This produces, for each 1 square degree survey field, a 1.6G FITS file for each of the 5 bands, $u^* g' r' i' z'$. Catalogues are then produced from these files.

3 The CANFAR Computing Infrastructure

The NGVS survey benefits from being sited at the National Research Council Herzberg Institute of Astrophysics (HIA) in Victoria, British Columbia, which is the home of the Canadian Astronomy Data Centre (CADC). One of the largest astronomy data centres in the world, the CADC has long been at the forefront of providing services for the storage, processing, and distribution of large survey datasets.

The mixed environment of astronomers and computer specialists provides an ideal setting for applying astroinformatics to produce science results.

The Canadian Advanced Network for Astronomical Research (CANFAR) [5], a collaboration between HIA, the Universities of Victoria and British Columbia, and funded by CANARIE, is a project to provide an infrastructure for data-intensive astronomy projects. This saves a project (e.g., a survey collaboration) time, money, and effort compared to developing a processing pipeline from scratch. The aim is to provide those portions of a pipeline that can be usefully supplied in a generic manner, such as access, processing, storage, and distribution of data, without restricting the analysis that can be performed. Standard astronomical software is provided, but the user may install and run code of their own, as desired.

The system works by combining the job scheduling abilities of a batch system with cloud computing resources. Each user creates and operates one or more Virtual Machines (VMs), in the same way as one would manage a desktop machine. The user submits batch processing jobs which identify the VMs on which they could run. Well-tested, available grid technologies are employed, building on existing CADC infrastructure. Services seen by the user are implemented, where possible, using protocols compatible with the International Virtual Observatory Alliance (IVOA) standards. The CANFAR project was developed throughout with six astronomical science projects, including the NGVS, as an integral part of the process.

Components of the CANFAR infrastructure include a Virtual Image manager, which allows users to create and configure VMs, the Condor job scheduler, which gives the user powerful tools for controlling their jobs, the cloud scheduler, which automatically creates and maintains a cluster consisting of multiple instances of the user VMs, cloud functionality via Nimbus, and virtual distributed storage. The latter is provided by VOSpace, the IVOA standard. This is accessible via the command line (interactively, or from a batch job), or an intuitive graphical interface accessed via a web browser. Data within the system may be either public or proprietary.

3.1 Use of CANFAR for NGVS Science

The combination of CANFAR, VOSpace, and astroinformatics tools means that full-scale scientific analysis of real astronomical data is now feasible, by astronomers who are not experts in data-intensive computing. It may be achieved within a realistic timescale, by an international collaboration, for terascale datasets. Currently, as well as the original six projects, several others are beginning to use CANFAR for science data processing that is expected to lead to published results. A CANFAR workshop was held at HIA in May 2011³.

Within the NGVS, we have recently created a catalogue of 13,884,023 objects from the first 70 deg² of the survey (Figure 1)⁴, then used CANFAR to run the code

³ <http://www.astro.uvic.ca/~canfar/canfarw>

⁴ Although the present catalogue was not created utilizing the current version of CANFAR, such a catalogue is easily recreated there.

of Schlegel et al. [8] to provide a Galactic extinction correction for each object. Such extinctions are vital, for example, for obtaining accurate photometric redshifts when using template-based codes. We have used CANFAR to generate simulated galaxies to aid in assessing the completeness of the Virgo Cluster luminosity function, and it will similarly provide the ability to fit detailed galaxy profiles to large numbers of real galaxies within the survey. Many further uses for computationally intensive analysis of the survey that might not otherwise be feasible are planned.

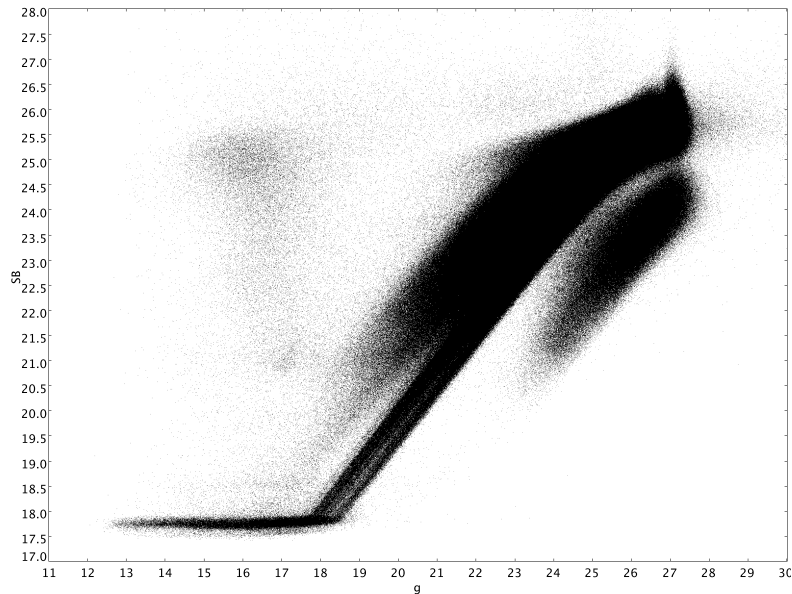


Fig. 1 Surface brightness versus magnitude for the Next Generation Virgo Cluster Survey, from a SExtractor catalogue of 13,884,023 objects. The plot shows the raw data. Such a catalogue is easily generated using CANFAR compute resources, and the 6.5G FITS file is easily stored and served to survey members via VOspace. Conversion of the FITS file to column-oriented FITS using STILTS [9] enables the full survey dataset to be plotted in seconds using TOPCAT on a desktop machine. Plotting the full data in this way shows aspects of the raw data that may not be apparent when plotting subsamples, such as the significant spread in the stellar locus, a kink in the distribution at $g \sim 23.5$, and an overdensity of (spurious) objects at $16.75 < g < 17.25$ and $20.75 < g < 21.25$.

3.2 Fast Data Mining Algorithms

By extension of the same arguments for providing a hardware infrastructure and standard software tools within CANFAR, we aim to provide a robust set of generic

tools which can be used for data analysis. Many modern data mining and machine learning algorithms are ideal for discovering useful patterns in astronomy data, but require time and effort to set up and learn, and, naively implemented, scale in a polynomial way with the number of objects (e.g., N^2 or N^3). To render the analysis tractable, they must scale no worse than $N \log N$. Such implementations are available, for example, we have confirmed that the proprietary software of the Skytree Corp. scales in this way for typical astronomy data, providing fast runtimes, a wide range of algorithms, and user support that enables its practical use.

Licensing issues may ultimately prevent the deployment of proprietary software on the distributed computing system of CANFAR, so we are also investigating open-source alternatives, such as the Toolkit for Multivariate Data Analysis (TMVA), part of the ROOT system used by the high energy physics community.

4 Virgo Cluster Membership and Photometric Redshifts

Whether or not a galaxy is a member of the cluster, or is in the background, is of fundamental importance to many NGVS science goals. For example, a luminosity function measurement for the cluster contaminated by background galaxies will yield a faint-end slope that is too steep. The faint end slope for Virgo remains unconstrained between, for the Schechter function parametrization, $\alpha \approx -1.3$ and $\alpha \approx -1.9$ because of issues like this. This is the difference between the expected value of $\alpha \approx -2$ for a Λ CDM cosmological model, and heavy suppression of galaxy formation within dark matter haloes, which implies very different physics.

Traditional methods of assigning cluster membership include statistical background subtraction using galaxy number counts, and, for Virgo, the apparent magnitude, surface brightness, luminosity class, and resolution of structure [3]. However, the NGVS is (a) much deeper than the depths for which spectra can be obtained, and (b) provides 5-band photometry and high resolution morphology. This means that a significantly larger and higher-dimensional sample of information is available, but that it may be non-trivial to extract. For example, although many populations of objects such as low surface brightness dwarf spheroidals must be in Virgo, others, such as compact ellipticals, can be ambiguous, because a small cluster galaxy can mimic a large background galaxy. In general, regions of parameter space will exist where cluster membership assignment must be probabilistic.

Results from the NGVS are still in their early stages. However, it is clear that machine learning methods can successfully separate most objects within the cluster from the background. For supervised methods, the usual limitation to the training set regime applies. Here, that is $g < 21$ mag for galaxies, from MMT/Hectospec spectra. Nevertheless, these spectra only cover the central four square degrees of the cluster, so extending this to 104 square degrees represents a considerable advance, although it is likely that environmental variations, e.g., infalling as opposed to *in situ* dwarf ellipticals, will create considerable differences in the galaxy populations

as a function of cluster environment. One must also, e.g., account for field-to-field variation, although the whole survey is calibrated with reference to the SDSS.

Figure 2 shows the results of one approach to cluster membership: empirical photometric redshifts. At first glance, the use of photometric redshift to assign cluster membership may seem odd, because traditionally, background photometric redshifts have an intrinsic spread of approximately $\Delta z \sim 0.02$, whereas the cutoff for cluster membership is $z < 0.01$. However, the use of the empirical approach with k nearest neighbours alleviates this. A support vector machine to classify objects as Virgo or the background gives similar results.

Due to poor weather, it is also now unlikely that the full 104 square degrees of the survey will be completed in all 5 bands by the nominal survey completion in semester 2012A. Hence we will obtain a more complex patchwork of results in which some bands are missing for some areas of sky. Methods that can account for missing data, such as a boosted decision tree, may thus provide a suitable approach. This may be of particular importance for regions in which the u band is not available, because the most powerful discriminant of distance, the Balmer break, drops out of g at $z < 0.2$. Hence the ability to include morphological criteria may be critical.

Clearly, there is much potential for refinement of these results. For example, one can assign full probability distributions in redshift space to each galaxy, or one can use unsupervised methods that are not subject to restriction to the spectroscopic regime. These can then be extended from the forced separation into distinct classes of k -means to fuzzy- C means or other methods to provide probabilistic membership. Regions of parameter space that remain ambiguous may be used as motivation for followup spectroscopy. The strength of the astroinformatics approach is that many such methods may be tested, with relative ease, on the entire survey.

5 Conclusions

The NGVS is an excellent example of a modern survey for which the methods of astroinformatics are ideally suited. While the data volume (2.6T) is substantially larger than that for which many of the traditional astronomical analysis tools were designed, it is not so large that it is intractable. This means that astroinformatics can be used to perform science-driven analyses of the survey, maximizing its science potential, without its use being the subject of new research in its own right.

References

1. Ball, N.M., Brunner, R.J.: Data Mining and Machine Learning in Astronomy. *Int. J. Mod. Phys. D* **19**, 1049–1106 (2010)
2. Bertin, E., Mellier, Y., Radovich, M., et al. The TERAPIX Pipeline. In: Bohlender, D.A., Durand, D., Handley, T.H. (eds.) *Astronomical Data Analysis Software and Systems XI*, ASP Conf. Proc. **281**, pp. 228–237. Astronomical Society of the Pacific, San Francisco (2002)

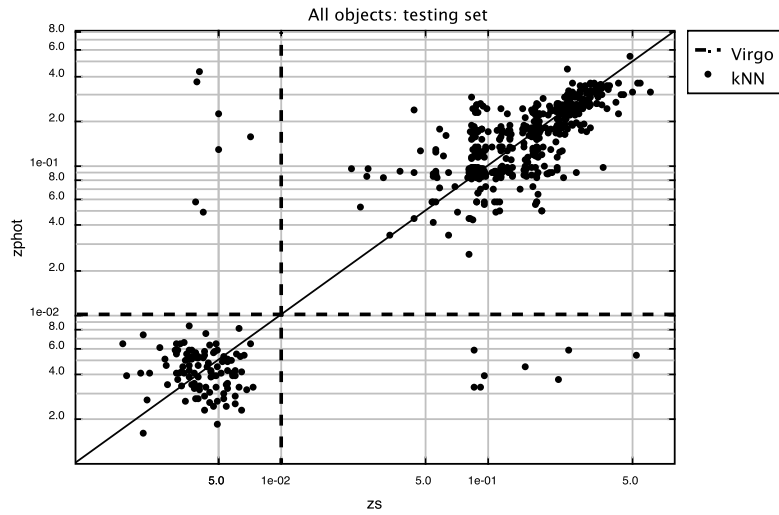


Fig. 2 Virgo cluster membership: empirical photometric redshift via k nearest neighbours versus spectroscopic redshifts to $g < 21$ mag. The lack of galaxies directly behind Virgo provides a separation such that $z = 0.01$ may be taken as a clean cutoff for cluster membership. Objects in the bottom left quadrant are thus correctly assigned Virgo members, and those in the top right are correctly assigned background galaxies. The completeness and efficiency of cluster membership assignment is approximately 90%. Note that the axes are plotted on a logarithmic scale.

3. Binggeli, B., Sandage, A., Tammann, G.A.: Studies of the Virgo Cluster. II. A Catalog of 2096 Galaxies in the Virgo Cluster Area. *Astron. J.* **90** 1681–1758 (1985)
4. Borne, K.: Scientific Data Mining in Astronomy. *Data Mining and Knowledge Discovery Series*, Taylor & Francis: CRC Press, Boca Raton, FL, Ch. 5, pp. 91–114 (2009)
5. Gaudet, S., Hill, N., Armstrong, P., et al. CANFAR: the Canadian Advanced Network for Astronomical Research. In: Radziwill, N.M., Bridger, A. (eds.) *Proc. SPIE, Software and Cyberinfrastructure in Astronomy*, **7740**, 1L
6. Gwyn, S.D.J.: MegaPipe: The MegaCam Image Stacking Pipeline at the Canadian Astronomical Data Centre. *Publ. Astron. Soc. Pac.* **120** 212–223 (2008)
7. Gwyn, S.D.J.: The CFHT Legacy Survey: Stacked Images and Catalogs. *arXiv/1101.1084* (2011)
8. Schlegel, D.J., Finkbeiner, D.P., Davis, M.: Maps of Dust Infrared Emission for Use in Estimation of Reddening and Cosmic Microwave Background Radiation Foregrounds. *Astrophys. J.* **500**, 525–553 (1998)
9. Taylor, M.B.: STILTS — A Package for Command-Line Processing of Tabular Data. In: Gabriel, C., et al. (eds.) *Astronomical Data Analysis Software and Systems XV*, ASP Conf. Proc. **351**, pp. 666–669. Astronomical Society of the Pacific, San Francisco (2006)
10. York, D., Adelman, J., Anderson, J.E., et al. The Sloan Digital Sky Survey: Technical Summary. *Astron. J.* **120**, 1579–1587 (2000)