

# Discovering the Unexpected in Astronomical Survey Data

Ray P. Norris<sup>1,2</sup>

<sup>1</sup>Western Sydney University, Locked Bag 1797, Penrith South, NSW 1797, Australia

<sup>2</sup>CSIRO Astronomy & Space Science, PO Box 76, Epping, NSW 1710, Australia

Email: [ray.norris@csiro.au](mailto:ray.norris@csiro.au)

(RECEIVED November 17, 2016; ACCEPTED December 19, 2016)

## Abstract

Most major discoveries in astronomy are unplanned, and result from surveying the Universe in a new way, rather than by testing a hypothesis or conducting an investigation with planned outcomes. For example, of the ten greatest discoveries made by the *Hubble Space Telescope*, only one was listed in its key science goals. So a telescope that merely achieves its stated science goals is not achieving its potential scientific productivity.

Several next-generation astronomical survey telescopes are currently being designed and constructed that will significantly expand the volume of observational parameter space, and should in principle discover unexpected new phenomena and new types of object. However, the complexity of the telescopes and the large data volumes mean that these discoveries are unlikely to be found by chance. Therefore, it is necessary to plan explicitly for unexpected discoveries in the design and construction. Two types of discovery are recognised: unexpected objects and unexpected phenomena.

This paper argues that next-generation astronomical surveys require an explicit process for detecting the unexpected, and proposes an implementation of this process. This implementation addresses both types of discovery, and relies heavily on machine-learning techniques, and also on theory-based simulations that encapsulate our current understanding of the Universe.

Keywords: surveys, telescopes

## 1 INTRODUCTION

Popper (1959) described the scientific method as a process in which theory is used to make a prediction which is then tested by experiment. That model, and its principle of ‘falsifiability’ remains the gold standard of the scientific method, and probably drives the majority of scientific progress. Notable recent successes include the discovery of the Higgs boson (ATLAS 2012) and the detection of gravitational waves (Abbott et al. 2016). Conversely, models such as string theory are sometimes criticised (e.g. Woit 2011) for being unfalsifiable, and thus failing to adhere to this Popperian scientific method.

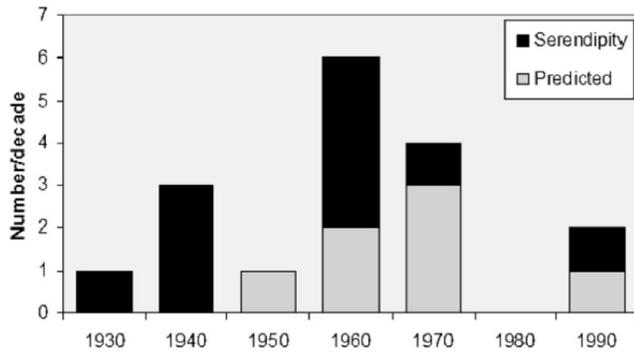
However, the Popperian scientific method is not the only one, and a number of other modes of scientific discovery have been proposed, notably by Kuhn (1962). For example, science may also proceed through a process of ‘exploration’ (e.g. Harwit 1981), in which experiments or observations are carried out in the absence of a compelling theory, in order to guide the development of theory.

Astronomy has largely developed through a process of exploration. For example, the Hertzsprung–Russell diagram (Hertzsprung 1908) was an observationally driven idea of

representing data, that led to the development of models of stellar evolution and ultimately nuclear fusion. In another example, the expanding Universe was discovered when Hubble plotted redshifts of galaxies against their brightness (Hubble 1929). More recently, the Hubble Deep Fields (Williams et al. 1996, 2000) were primarily motivated by a desire to explore the early Universe, rather than testing specific models or hypotheses.

### 1.1. The history of astronomical discovery

Astronomical discovery has often occurred as a result of technical innovation, resulting in the Universe being observed in a way that was not previously possible. Examples include the development of larger telescopes, or the opening up of a new window of the electromagnetic spectrum. More generally, we may define an  $n$ -dimensional parameter space whose  $n$  orthogonal axes correspond to observable quantities (e.g. frequency, sensitivity, polarisation, colour, spatial scale, temporal scale). Some parts of this parameter space have been well-observed and have already yielded their discoveries, whereas some parts of this space have not yet been observed. New



**Figure 1.** A plot of recent major astronomical discoveries, taken from Ekers (2009), of which seven were ‘known–unknowns’ (i.e. discoveries made by testing a prediction) and ten were ‘unknown–unknowns’ (i.e. a serendipitous result found by chance while performing an experiment with different goals). The data in this plot are taken from Wilkinson et al. (2004).

discoveries may lie in those unsampled parts of the parameter space, presumably available to new instruments able to sample that region of the parameter space. Most ‘accidental’ or ‘serendipitous’ discoveries result from observing a new part of this parameter space (Harwit 2003).

We may therefore broadly divide astronomical discoveries into (a) those which were made according to the Popperian model, in which a model or hypothesis is being tested (the known–unknowns), and (b) those which have resulted from observing the Universe in a new part of the parameter space, resulting in unexpected discoveries (the unknown–unknowns). Of course, an experiment may often be planned to test a hypothesis, but in doing so stumbles across an unexpected discovery. A classic example of this is the discovery of pulsars (Hewish et al. 1968) discussed in Section 2.2. Alternatively, data taken for an unrelated purpose may be mined for unexpected discoveries, such as the outlier detection algorithm described by Baron & Poznanski (2016) that finds ‘weird’ galaxies by searching for unusual spectra from the Sloan Digital Sky Survey.

Several studies (Harwit 1981; Wilkinson et al. 2004; Wilkinson 2007; Fabian 2010; Kellermann 2009; Ekers 2009; Wilkinson 2015) have shown that at least half the major discoveries in astronomy are unexpected, and are typically made by surveying the Universe in a new way, rather than by testing a hypothesis or conducting an investigation with planned outcomes. For example, Figure 1 shows the result of an examination (Ekers 2009) of 17 major astronomical discoveries in the last 60 yrs. Ekers concluded that only seven resulted from systematic observations designed to test a hypothesis or probe the nature of a type of object. The remaining ten were unexpected discoveries resulting either from new technology, or from observing the sky in an innovative way, exploring uncharted parameter space. In particular, experience has shown that unexpected discoveries often result when the sky is observed to a significantly greater sensitivity, or a significantly new volume of observational parameter space is explored.

PASA, 34, e007 (2017)  
doi:10.1017/pasa.2016.63

## 1.2. This paper

In Section 2 of this paper, I discuss the opportunities and challenges to making unexpected discoveries in the high data volumes and high complexity of next-generation astronomical surveys, and argue that surveys need to plan explicitly for these discoveries if they are to be successful. Section 3 proposes a process for discovering unexpected objects in astronomical surveys, and Section 4 proposes a process for discovering unexpected phenomena in astronomical surveys. Section 5 describes some preliminary attempts to implement and test some of these approaches and suggests some future directions.

To focus the discussion, this paper uses the ‘Evolutionary Map of the Universe’ survey (EMU: Norris et al. 2011) as an exemplar of next-generation surveys, but the broad conclusions and process will be relevant to all next-generation astronomical surveys.

## 2 THE PROCESS OF ASTRONOMICAL DISCOVERY

Astronomy is currently enjoying a boom in new surveys, with several next-generation astronomical survey telescopes planned, which will undoubtedly open up large new swathes of observational parameter space, potentially resulting in a large number of unexpected discoveries.

There are two quite different types of unexpected discovery:

- Type 1: Discoveries of new types of object (e.g. pulsars, quasars), identified as anomalies or unexpected objects in images or catalogues;
- Type 2: Discoveries of new phenomena (e.g. HR diagram, the expanding Universe, dark energy), identified as anomalies in the distributions of properties of objects. These are identified when the results of experiment are compared to theory (or perhaps to other observations) in some suitable parameter space.

### 2.1. Case study 1: The *Hubble space telescope*

The science goals that drove the funding, construction, and launch of the *Hubble Space Telescope* (*HST*) are listed in the *HST* funding proposal (Lallo 2012). A further four projects were planned in advance by individual scientists but not listed as key projects in the *HST* proposal. Conveniently, the National Geographic magazine selected the ten major discoveries of the *HST* (Handwerk 2005), resulting in an admittedly subjective ‘top ten’ list of *HST* discoveries (shown in Table 1). So we may compare the actual achievements of the *HST* against its planned achievements. Of these ten greatest discoveries by *HST*, only one was listed in its key science goals. In particular, the unplanned discoveries include two of the three most cited discoveries, and the only *HST* discovery (Dark Energy) to win a Nobel prize.

**Table 1.** Major discoveries made by the *Hubble Space Telescope* (*HST*). Of the *HST*'s 'top ten' discoveries (as ranked by National Geographic magazine), only one was a key project used in the *HST* funding proposal (Lallo 2012). A further four projects were planned in advance by individual scientists but not listed as key projects in the *HST* proposal. Half the 'top ten' *HST* discoveries were unplanned, including two of the three most cited discoveries, and including the only *HST* discovery (Dark Energy) to win a Nobel prize. This Table was previously published by Norris et al. (2015).

Project	Key Project?	Planned?	Nat Geo top ten?	Highly cited?	Nobel Prize?
Use cepheids to improve value of $H_0$	✓	✓	✓	✓	
UV spectroscopy of ig medium	✓	✓			
Medium-deep survey	✓	✓			
Image quasar host galaxies		✓	✓		
Measure SMBH masses		✓	✓		
Exoplanet atmospheres		✓	✓		
Planetary Nebulae		✓	✓		
Discover Dark Energy			✓	✓	✓
Comet Shoemaker-Levy			✓		
Deep fields (HDF, HDFS, GOODS, FF, etc)			✓	✓	
Proplyds in Orion			✓		
GRB Hosts			✓		

This example suggests that science goals are poor predictors of the discoveries to be made with a new telescope, and if a major new telescope merely achieves its stated science goals, it is probably performing well below its potential scientific productivity. Wilkinson et al. (2004) express this idea succinctly as *What a radio telescope was built for is almost never what it is known for*.

### 2.2. Case study 2: The discovery of pulsars

The Nobel-prize-winning discovery of pulsars by Jocelyn Bell occurred when a talented and persistent PhD student observed the radio sky for the first time with high time resolution, to study interstellar scintillation. By observing at high time resolution, she expanded the observational parameter space. She also knew her instrument intimately, enabling her to recognise that 'bits of scruff' on the chart recorder could not be due to terrestrial interference, but represented a new type of astronomical object. As a result, she discovered pulsars. She describes the process in detail in Bell-Burnell (2009).

The following critical elements were essential for this discovery:

- She explored a new area of observational parameter space.
- She knew the instrument well enough to distinguish interference from signal.
- She examined all the data by eye.
- She was observant enough to recognise something unexpected.
- She was open minded, and prepared for discovery.
- She was within a supportive environment (i.e. one that was accustomed to making new discoveries).
- She was persistent.

The value of the last three items should not be underestimated. When a PhD student obtains an observational result

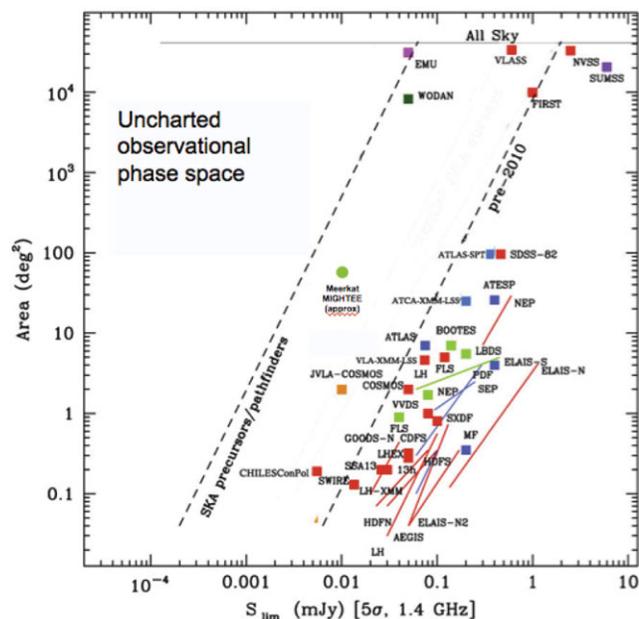
that differs from previous results or from conventional wisdom, there is a strong temptation to ascribe the difference to an error in the data.

### 2.3. Case Study 3: The evolutionary map of the Universe

Figure 2 shows the main radio surveys, both existing and planned, at frequencies close to 1.4 GHz. The largest existing radio survey, shown in the top right, is the wide but shallow NRAO VLA Sky Survey (NVSS: Condon et al. 1998). The most sensitive existing radio survey is the deep but narrow JVLA-SWIRE (Lockman hole) observation in the lower left (Condon et al. 2012). Existing surveys are bounded by a diagonal line that roughly marks the limit of available time on current-generation radio telescopes.

Many discoveries have been triggered by those surveys shown in Figure 1, ranging from the rare but paradigm-shifting discoveries (e.g. the radio-far-infrared correlation: van der Kruit 1971) to the numerous minor but still significant discoveries (e.g. the Infrared-Faint Radio Sources: Norris et al. 2006), which are now known to be very-high-redshift radio galaxies (Garn & Alexander 2008; Herzog et al. 2014; Collier et al. 2014). In the absence of any evidence to the contrary, Occam's razor would suggest that this diagram is uniformly populated with significant discoveries. Therefore, the unexplored region of observational parameter space to the left of the line presumably contains as many potential new discoveries per unit parameter-space as the region to the right. Radio surveys of that region should therefore yield many important discoveries, provided they are equipped to do so.

Within that unexplored region of parameter space are several planned next-generation radio surveys, the largest of which, in terms of numbers of sources detected, is EMU (Evolutionary Map of the Universe; Norris et al. 2011) which will



**Figure 2.** Comparison of existing and planned deep 20-cm radio continuum surveys, adapted from a diagram in Norris et al. (2013) originally drawn by Isabella Prandoni. The horizontal axis shows the  $5\text{-}\sigma$  sensitivity, and the vertical axis shows the sky coverage. The right-hand diagonal dashed line shows the approximate envelope of existing surveys, which is largely determined by the availability of telescope time. Surveys not at 20 cm are represented at the equivalent 20 cm flux density, assuming a spectral index of  $-0.8$ . The squares in the top-left represent the new radio surveys discussed in this paper. The Square Kilometre Array (Dewdney et al. 2009) will hopefully conduct even larger surveys in the next decade, extending well to the left of EMU, but such plans are not yet concrete.

use the Australian SKA Pathfinder (Johnston et al. 2008), to survey 75% of the sky to a sensitivity of  $10 \mu\text{Jy}/\text{beam rms}$ . Only a total of about  $10 \text{ deg}^2$  of the sky has been surveyed at 1.4 GHz to this sensitivity, in fields such as the *Hubble*, ATLAS, and COSMOS fields. EMU is the largest radio continuum survey so far, and will detect about 70 million galaxies, compared to the 2.5 million detected over the entire history of radio-astronomy. Not only will EMU have greater sensitivity than previous large-area surveys, but it will also have better resolution, better sensitivity to extended emission, and will measure spectral index and, courtesy of the POSSUM project (Gaensler et al. 2010), polarisation for the strongest sources.

EMU will therefore significantly expand the volume of observational parameter space, so in principle should discover unexpected new phenomena and new types of object.

However, the complexity of ASKAP and the large data volumes mean that it may be non-trivial to identify them. For example, in the list above of critical elements which led to the discovery of pulsars, EMU can satisfy all those elements except (a) knowing the instrument well enough to distinguish interference or artefacts from signal, (b) being able to examine all the data by eye, and (c) being able to recognise something unexpected.

For (a), it is likely that no human will be sufficiently familiar with ASKAP to distinguish subtle astrophysical effects from subtle instrumental artefacts. Any process to detect unexpected astrophysical effects is likely to detect unexpected artefacts. Rather than expecting to identify these *a priori*, it is likely that we will have to learn to identify them in the data, and then trace their source *a posteriori*. This process is likely to be an important component of the process of discovering the unexpected.

For (b), the petabyte data volumes from ASKAP mean that it will be impossible for an astronomer to sift through the data, looking for something unusual. Instead, the only way of extracting science from large volumes of data is to interrogate the data with a well-posed question, such as ‘plot the specific cosmic star formation rate of star-forming galaxies as a function of redshift’. So there is a danger that projects like EMU will produce good science in response to such well-posed questions (the ‘known-unknowns’), and thus achieve their science goals, but will miss the 90% of discoveries that are unexpected (the ‘unknown-unknowns’).

The final element (c), of being able to recognise something unexpected, is perhaps the hardest element. While the human brain has been exquisitely tuned by millions of years of evolution to notice anything unexpected and potentially dangerous, if we can’t sift through the data by eye, then we must rely on tools to detect the unexpected, and such tools do not currently exist.

On the other hand, if we don’t make the unexpected discoveries, then we will probably miss out on the most important science results from these telescopes. We have therefore started a project within EMU (named Widefield outlier Finder, or WTF) to develop techniques for mining large volumes of astronomical data for the unexpected, using machine-learning techniques and algorithms.

#### 2.4. The value of science goals

New telescopes or surveys are usually justified by their science goals. For example, the EMU project (Norris et al. 2011) is justified by 16 key science projects with goals such as measuring the star formation rate density over cosmic time, studying AGN evolution and the role of AGN feedback, and making independent measurements of fundamental cosmological parameters. However, as demonstrated above in the case of the *HST*, the major discoveries made with a new telescope or survey are not usually represented by such science goals.

However, science goals are still important for two reasons. First, they represent use cases. If a telescope is built that is able to address challenging science goals, then it is likely to be a high-performing telescope. Second, much of astronomy advances not by spectacular major discoveries, but by the incremental science that is usually encapsulated in science goals. Such incremental advance is also very important, and, unlike serendipitous discoveries, represents a predictable outcome from a new telescope.

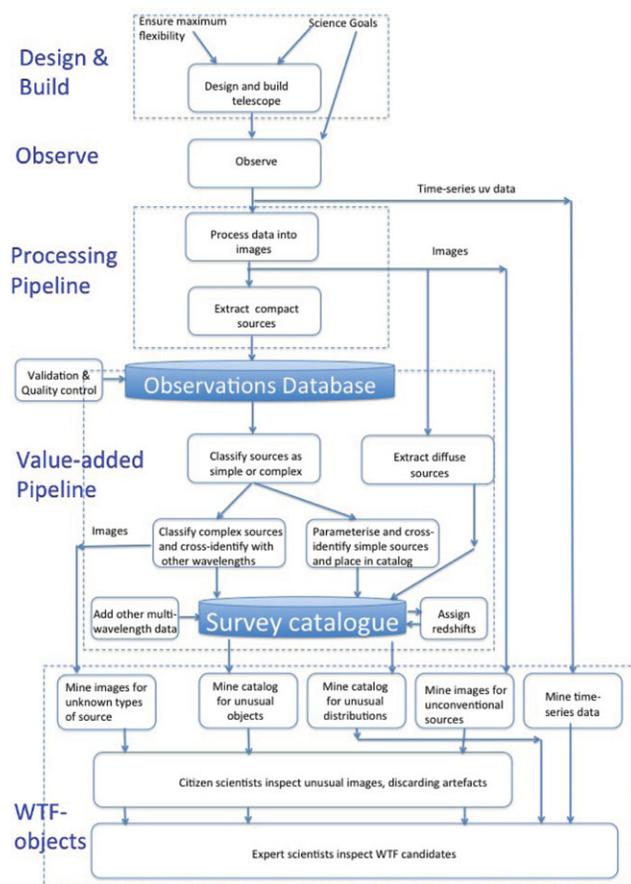


Figure 3. The flowchart for discovering unexpected objects in EMU.

For example, EMU will hopefully advance the knowledge of galaxy evolution by measuring the evolution of the cosmic star-formation rate, the evolution of active galactic nuclei, and the feedback processes that link them, and this will no doubt result in many worthwhile and highly cited papers. However, these may be dwarfed in impact by the unexpected discoveries.

### 3 TYPE 1 DISCOVERIES: UNEXPECTED OBJECTS

EMU is expected to detect about 70 million objects, compared to the current total of ~ 2.5 million known radio sources. Since the 70 million objects will probably include new unexpected classes of radio source, it is important for EMU to plan to identify new classes or phenomena, rather than hoping to stumble across them. EMU will do so through its WTF project, which has the explicit goal of discovering the unexpected.

This section describes how the WTF project will make Type 1 discoveries (unexpected objects), An overview is shown in Figure 3 and the following subsections address each of the steps in that flowchart. Although this is designed for EMU, the broad approach is applicable to any survey.

PASA, 34, e007 (2017)  
doi:10.1017/pasa.2016.63

### 3.1. Design and construction

As discussed in Section 2.4, the construction of any new telescope must necessarily be designed to optimise its performance for specific science goals. However, it is important not to design and build it so it can *only* achieve those goals, because that would limit its ability to discover the unexpected. Instead, it is important to maximise flexibility. The design of the telescope therefore needs to maximise the ultimate scientific productivity, in addition to achieving the specific science goals.

Similarly, it is sometimes necessary to process the data to reduce the volume of data to that which is necessary to achieve the science goals, discarding the excess. For example, ASKAP will generate about 70 PB of calibrated correlated time-series data each year, which is then processed into images occupying only about 4 PB per year. It is not economically possible to store all the time-series spectral-line data, and so that data is discarded.

Discarding excess data is sensible if all the information is present in the images. However, processing the time-series data to produce the images is a lossy process, and the discarded information may well be the key to an unexpected discovery. So reducing the data volume by keeping only processed data should be avoided as much as possible.

Even when time-series data must be discarded, it can still be searched in real time for time-varying phenomena such as fast radio bursts (Lorimer et al. 2007). In the case of EMU, this search is undertaken by partner projects CRAFT (Macquart et al. 2010) and VAST (Murphy et al. 2013).

### 3.2. Observations

Discoveries are thinly distributed through the observational parameter space. We cannot predict where they lie, and it is difficult to quantify the volume of parameter space being explored, but the probability of making an unexpected discovery is presumably proportional to the volume of new parameter space being explored. The observations should therefore be optimised, not only for the specific science goals, but also to maximise the volume of new observational parameter space being explored, which means maximising the sensitivity to poorly explored parameters such as circular polarisation, time variability, diffuse emission, etc.

### 3.3. Data processing and compact source extraction

The first stage of ASKAP data processing, performed by the ASKAPSOFT suite of software, is to calibrate the time-series data, Fourier transform it into image data, and then deconvolve it. The resulting images are then placed in the observations database (called CASDA) for storage and retrieval by users.

It is important that this process makes as few assumptions as possible about the nature of the objects being detected. For example, we know that the vast majority of objects detected

by EMU will be less than one arcmin in extent, and so it is tempting to discard the shortest baselines corresponding to spatial scales larger than this. However, to do so will be to guarantee that EMU will not detect any objects larger than this scale, thereby limiting the volume of observation parameter space being explored.

The ASKAPSOFT real-time processing pipeline includes source extraction software to identify and measure the parameters of compact sources in the radio images. The algorithm for doing so is still being refined and tested against other source finders (Hopkins et al. 2015), but is optimised for sources that are unresolved or less than a few beamwidths in extent. The software will measure the extent of each component (an ‘island’) and fit gaussians to the peaks within the island. The measured parameters from this process are stored in a table in CASDA for storage and retrieval by users.

Diffuse sources will not normally be discovered by this process, but will be extracted in offline processing (see Section 3.5).

### 3.4. Data validation

The first stage of EMU data validation takes place in near-real-time to flag data which are affected by radio-frequency interference or hardware malfunctions. A second stage of validation is conducted on each set of observations by the EMU science survey team, checking for image artefacts, calibration errors, etc. It is important to ensure that this process does not also reject data containing unexpected discoveries. For example, a strong radio burst might be misinterpreted as interference. However, an astrophysical radio burst will take place in the far field of ASKAP, while interference generally takes place in the near field. Interference can therefore be distinguished from radio bursts by testing whether the parameters on different baselines are consistent with an astrophysical source. It is therefore important that data validation techniques use such sophisticated tests rather than simple amplitude threshold tests.

### 3.5. Diffuse source extraction

The source extraction algorithm in ASKAPSOFT is not expected to detect diffuse emission, such as cluster haloes and supernova remnants, which are notoriously difficult to detect automatically. A number of algorithms (e.g. Dabbech et al. 2015; Butler-Yeoman et al. 2016; Riggi et al. 2016) are under development for automatically detecting diffuse sources in radio-astronomical images.

### 3.6. Classification of sources as simple or complex

About 90% of EMU sources will consist of a single radio component with no nearby radio component with which it might be associated. I term these ‘simple’ sources. Physically, these are likely to be star-forming galaxies, low-luminosity AGN, or young radio-loud galaxies typically classified as

Gigahertz-peaked spectrum (GPS) or compact steep spectrum (CSS). The first stage of classification and identification is to identify such sources from their radio morphology alone. This separation into simple and complex sources will be achieved in EMU using a machine-learning algorithm, currently under development (Park, Norris & Crawford, in preparation). It is likely that the final algorithm will use one of Logistic Regression, a Support Vector Machine, or a Neural Network binary classification.

The resulting simple sources will then be matched to optical/infrared catalogues using a likelihood ratio (LR) technique (Sutherland & Saunders 1992; Weston, in preparation).

The remaining sources, which we term ‘complex’, must be classified and cross-identified in a more sophisticated process.

### 3.7. Source classification and cross-identification of complex sources

Classifying the morphology of radio sources, and cross-identifying them with their counterparts at optical/infrared wavelengths, might be regarded as being two separate processes. However, two nearby unresolved radio components might either be the two lobes of an FR II radio source, or the radio emission from two unassociated star-forming galaxies. Only by cross-identifying with multiwavelength data, particularly optical/infrared data, can these two cases be distinguished, since the pair of star-forming galaxies will have an infrared host galaxy coincident with each of the radio components, whereas the host of the FR II is likely to lie between them.

Whilst this process is easy for the expert human, the 7 million complex sources expected to be detected by EMU pose a significant challenge. Several techniques are being evaluated, using the ~5 000 sources in the ATLAS data set (Norris et al. 2006; Middelberg et al. 2008; Hales et al. 2014; Franzen et al. 2015) as a testbed, as follows:

- All sources are cross-identified and classified by eye, to provide a training and validation set.
- The sources are being cross-matched by citizen scientists in the Radio Galaxy Zoo project (Banfield et al. 2016).
- A Bayesian approach is being developed (Fan et al. 2015).
- A variety of machine-learning approaches are being explored, both supervised and unsupervised.

### 3.8. The survey catalogue

After cross-matching and classification, all sources detected in the survey are placed in the survey catalogue, which for EMU is called the EMU Value-Added Catalogue (EVACAT). To each source are added other available data such as redshifts and other multiwavelength data. Many of the redshifts are not spectroscopic, but are photometric redshifts or ‘statistical redshifts’ (Norris et al. 2011) which are best expressed as a probability distribution function rather than as a single value.

### 3.9. Mining images for unexpected objects

The source extraction algorithm in ASKAPSOFT is not expected to detect unconventional sources. An example of an unconventional source might be a ring of emission several arcmin in diameter but with an amplitude of only half the rms noise level in any one pixel. Such a structure would be invisible in the image to the human eye, or to a conventional source extraction code, but would be easily detectable at a high level of significance using a suitable matched filter, such as a Hough transform (Hollitt & Johnston-Hollitt 2012). Many other examples of potential diffuse and unconventional sources may be imagined.

To detect such sources, the WTF pipeline will retrieve images from CASDA and apply a number of different algorithms in parallel. Detecting sources with unconventional morphology is much harder and is the subject of continuing research, and several algorithms such as self-organised maps (Geach 2012) are currently being explored.

### 3.10. Mining the catalogue for unexpected objects

The catalogue will be searched for properties of objects in an  $n$ -dimensional plot with axes such as flux density, spectral index, and IR-to-radio ratio. Known types of object (e.g. stars, galaxies, quasars) will appear as clusters in this parameter space. Algorithms are being explored that will search the parameter space for clusters of objects that do not correspond to known types of objects. Although targeted specifically at EMU, such approaches are expected to have broad applicability to astronomical survey data.

## 4 TYPE 2 DISCOVERIES: UNEXPECTED PHENOMENA

Some unexpected discoveries are made when the properties of a sample of objects differ from those predicted by theory in some unexpected way. For example, dark energy was discovered (Riess et al. 1998; Perlmutter et al. 1999) when the relationship between the brightness and redshift of type 1A supernovae failed to follow the distribution predicted by theory. Here, I describe an approach in which the data is tested against theory. Although it resembles the standard Popperian technique, it differs in that what is being tested is the sum of our understanding of the Universe, rather than any particular theory.

A common way of testing theories is to derive some physically meaningful quantity, such as a luminosity function, and then compare that with the theoretical luminosity function predicted by theory. Such an approach has the advantage of yielding results which are easily compatible with other observations and other theories. It has the disadvantage that observational data has to be corrected for incompleteness, and this is often difficult to do accurately. For example, to calculate the radio luminosity function of radio sources, and compare it with other derived radio luminosity functions, Mao et al.

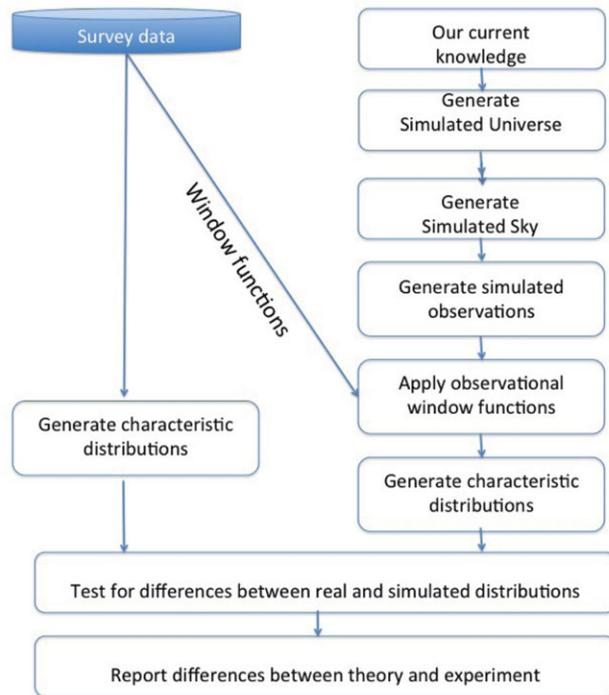
(2012) needed to correct the data not only for a variable radio sensitivity across the field, but for the incompleteness of the optical spectroscopy survey that produced the necessary redshifts. It is very difficult to account for all the selection effects accurately.

These various sources of incompleteness, which I label the ‘window function’, are generally well-understood and well-determined. For example, Mao et al. (2012) were able to use a map of the sensitivity across the radio image, and a plot of the sensitivity of the redshift survey as a function of magnitude. Thus, for a hypothetical source of a given optical magnitude and position, it is trivial to calculate the probability of it appearing in the catalogues with a measured redshift. The converse process is much harder—correcting the catalogue for these effects requires a number of approximations. It is likely that the differences between different measurements of this radio luminosity function (e.g. Mao et al. 2012; Mauch & Sadler 2007; Padovani et al. 2011) is primarily caused by these approximations.

An alternative to correcting the data to compare it with physically realistic models, is to use the theory to simulate the observations, and then apply the window function to result in simulated data that can be compared with the original data. Of course, a particular simulated galaxy will not coincide with a particular real galaxy, and so it is necessary to compare the statistical properties of the simulated data to those of the real data. But this comparison can be done in a parameter space which is close to that of the real data (e.g. source counts as a function of flux density in the survey volume), rather than transforming it to a physically meaningful parameter space (e.g. source counts as a function of luminosity in an idealised volume). This may be regarded as a Bayesian process, in that the theory is being used to predict the data, rather than the theory being inferred from the data.

In the case of searching for the unexpected, the simulations are being used to encapsulate our current understanding of astrophysics so that they can be compared with the data, to see if the data is consistent with our current understanding. Any significant difference between the two either represents an error in the data or simulation, or an unexpected discovery.

This process is shown in Figure 4, and includes the following steps. The starting point is a simulation, such as the Millennium Simulation (Springel et al. 2005) which encapsulates our knowledge about cosmology and galaxy formation. From this is generated a simulated sky, using our knowledge of the observed properties of galaxies. Tools such as the Theoretical Astrophysical Observatory (TAO: Bernyk et al. 2016) are designed to do this. However, TAO does not yet generate a radio sky, and so a simulated radio sky must be generated from the TAO sky using a semi-empirical model of radio sources. The model sky is then converted to a simulated observed sky using observational constraints such as sensitivity and resolution. The window function is then applied including factors such as area of sky observed, and any varying sensitivity across the observations.

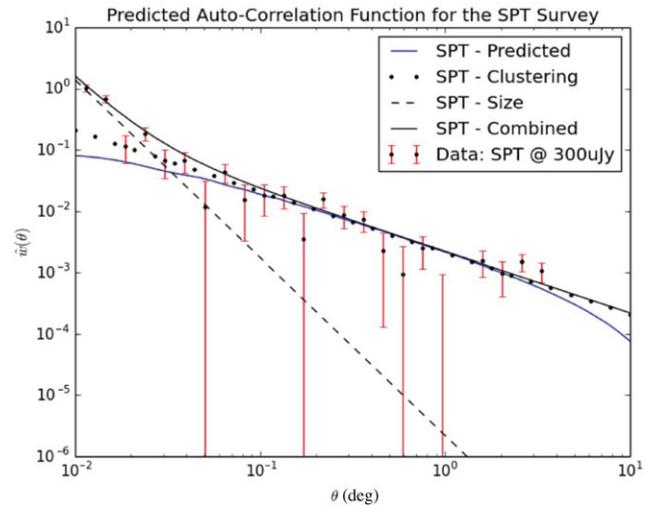


**Figure 4.** The flowchart for discovering unexpected phenomena in the EMU WTF project.

A characteristic distribution is a representation of the observational or simulated data which represents the data in some particular parameter space. Well-known examples include source count plots and angular power spectra, but in principle almost any observational quantity can be plotted against any other, and there is no need for these plots to be confined to two dimensions. To systematically search for unexpected deviations of theory from data, all combinations of observational quantities need to be searched by algorithms which will report significant anomalies to the user.

A simple example of this process, taken from Rees et al. (in preparation) is shown in Figure 5. Here, the characteristic distribution is the angular power spectrum for radio sources in the SPT (South Pole Telescope) field, using the radio observations described by O’Brien et al. (2016). The simulated data were based on the Millennium Simulation, from which a simulated sky of galaxies was generated using the TAO tool. From this, a radio sky was generated as described by Rees et al. (in preparation) using semi-empirical assumptions about the properties of radio sources based on the zFOURGE survey (Rees et al. 2016). In this case, the observational data were corrected for the window function, but the correction could equally well be applied to the simulation data. In this case, the data are found to be consistent with the simulation.

It is important to note that this process is not intended to detect outliers, or ‘Type 1’ discoveries, in the data, which are better handled using the process described in Section 3. Instead, this process is intended to detect unexpected trends or correlations in the data: the ‘Type 2’ discoveries.



**Figure 5.** The angular power spectrum for radio sources in the SPT field, taken from Rees et al. (in preparation). Points with error bars are the measured angular power spectrum of the data obtained by O’Brien et al. (2016), and the blue line shows the distribution predicted by the semi-empirical model described in the text. The dotted line shows the cosmological signal predicted by  $\Lambda$ CDM, and the dashed line show the effect of radio source size and double radio sources. The solid black line is the sum of these latter two predictions.

## 5 PRELIMINARY ATTEMPTS, AND FUTURE DIRECTIONS

To test the ideas driving this paper, a data challenge was constructed on the Amazon Web Services (AWS) cloud platform (Crawford, Norris, & Polsterer 2016). Initially, we wanted to see which algorithms and techniques are best at finding unexpected results, and so we constructed a number of data challenges in which data sets (both real and simulated, and both images and tabular data) are constructed with simulated unexpected discoveries (known as ‘eggs’) buried in them. We then invited machine-learning groups to try out their algorithms to see if they could find the simulate eggs.

This approach was less successful than expected, for the following reasons:

- We had underestimated the difficulty of non-astronomers engaging in this project. Specific difficulties included file formats, and the need to present the problem in a way accessible to non-astronomers.
- Lack of personpower: such a project requires dedicated resources.
- The most important factor was that discovering the unexpected is harder than expected.

As a result of that experiment, it was clear that a more systematic approach was needed, resulting in the process described in this paper. By breaking the problem down into building blocks, it also makes it a more tractable problem for a team-based approach. Furthermore, many of the building blocks are important tools in their own right that are

necessary to extract even the known–unknowns from EMU (e.g. classification and cross-identification of radio sources).

Other avenues of research are also likely. For example, it is likely that in the Search for Extra-terrestrial Intelligence (SETI), any detected civilisation is likely to be so much more advanced than ours (Norris 1999) than we might not recognise an intelligent signal. A better strategy may be simply to look for signals that are different from those that we expect from known astrophysical processes. In that case, a search for SETI reduces to searching for the unexpected, and can use the process proposed here.

## 6 CONCLUSION

- Most major discoveries in astronomy are unexpected.
- In the past, unexpected discoveries were made serendipitously by users pursuing other goals or exploring the parameter space. However, the complexity of next-generation instruments, and the large volumes of data generated, make it unlikely that they will make such unexpected discoveries. Instead, telescopes must be designed explicitly to maximise their ability to discover the (potentially more important) **unknown** science goals.
- The use of science goals when planning a new telescope are valuable as ‘use cases’ for helping design a good project, and are also likely to provide much of the incremental science that results from a successful project, but they are unlikely to represent the most significant science output from the telescope.
- With the exception of telescopes designed specifically to answer a particular science question, telescopes that merely achieve their stated science goals have probably failed to capture the most important scientific discoveries available to them.
- Because of the complexity and large data volumes of next-generation scientific projects, unexpected discoveries are less likely to happen by chance, but will require software designed to mine the data for unexpected discoveries.
- Unexpected discoveries may be either Type 1 (unexpected objects) or Type 2 (unexpected phenomena), and it is necessary to design processes to deal with both types.
- A process has been proposed for finding each of these types in radio survey data, and it is expected that this process may be broadly applicable to other types of astronomical survey.

## ACKNOWLEDGEMENTS

I thank Laurence Park, Evan Crawford, and Kai Polsterer for valuable discussions. I thank Amazon Web Services for grant EDU\_R\_FY2015\_Q3\_SKA\_Norris that enabled an early prototype to be constructed on the AWS cloud platform. I thank the University of Cape Town for hosting me for a period in which part of this paper was written. I acknowledge the Wajarri Yamatji people as the traditional owners of the ASKAP Observatory site.

PASA, 34, e007 (2017)  
doi:10.1017/pasa.2016.63

## REFERENCES

- Abbott, B. P., et al. 2016, *PhRvL*, 116, 061102  
ATLAS Collaboration 2012, *PhLB*, 716, 1  
Banfield, J. K., et al. 2016, *MNRAS*, 460, 2376  
Baron, D., & Poznanski, D. 2017, *MNRAS*, 465, 4530  
Bell-Burnell, J. 2009, in *Proc. Sci., Accelerating the Rate of Astronomical Discovery* (held in Rio de Janeiro, 11–14 August), 014  
Bernyk, M., et al. 2016, *ApJS*, 223, 9  
Butler-Yeoman, T., Frean, M., Hollitt, C. P., Hogg, D. W., & Johnston-Hollitt, M. 2017, in *ASP. Conf. Ser., Proc. of ADASS XXV*, eds. N. P. F. Lorente & K. Shortridge (San Francisco: ASP), in press  
Collier, J. D., et al. 2014, *MNRAS*, 439, 545  
Condon, J. J., et al. 1998, *AJ*, 115, 1693  
Condon, J. J., et al. 2012, *ApJ*, 758, 23  
Crawford, E., Norris, R. P., & Polsterer, K. 2016, arXiv:1611.02829  
Dabbech, A., et al. 2015, *Astr. Ap.*, 576, A7  
Dewdney, P. E., Hall, P. J., Schilizzi, R. T., & Lazio, T. J. L. W. 2009, *IEEEP*, 97, 1482  
Ekers, R. D. 2009, in *Proc. Sci., Accelerating the Rate of Astronomical Discovery* (held in Rio de Janeiro, 11–14 August), 007  
Fabian, A. C. 2010, in *Serendipity in Astronomy*, eds. M. de Rond & I. Morley (Cambridge: Cambridge University Press), 22–73  
Fan, D., Budavári, T., Norris, R. P., & Hopkins, A. M. 2015, *MNRAS*, 451, 1299  
Franzen, T. M. O., et al. 2015, *MNRAS*, 453, 4020  
Gaensler, B. M., Landecker, T. L., Taylor, A. R., & POSSUM Collaboration 2010, *BAAS*, 42, 470.13  
Garn, T., & Alexander, P. 2008, *MNRAS*, 391, 1000  
Geach, J. E. 2012, *MNRAS*, 419, 2633  
Hales, C. A., et al. 2014, *MNRAS*, 440, 3113  
Handwerk, B., 2005, *Hubble Space Telescope Turns 15* (Washington, DC: National Geographic Magazine)  
Harwit, M. 1981, *Cosmic Discovery* (Cambridge, MA: MIT Press)  
Harwit, M. 2003, *PhT*, 56, 38  
Hertzprung, E. 1908, *AN*, 179, 373–380.  
Herzog, A., et al. 2014, *Astr. Ap.*, 567, A104  
Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Nature*, 217, 709  
Hollitt, C., & Johnston-Hollitt, M. 2012, *PASA*, 29, 309  
Hopkins, A. M., et al. 2015, *PASA*, 32, e037  
Hubble, E. 1929, *PNAS*, 15, 168  
Johnston, S., et al. 2008, *ExA*, 22, 151  
Kellermann, K. I., et al. 2009, in *Proc. Sci., Accelerating the Rate of Astronomical*, <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=99>  
Kuhn, T. S. 1962, *The Structure of Scientific Revolutions* (Chicago: The University of Chicago Press)  
Lallo, M. D. 2012, *OptEn*, 51, 011011  
Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, *Science*, 318, 777  
Macquart, J.-P., et al., *PASA*, 27, 272.  
Mao, M. Y., et al. 2012, *MNRAS*, 426, 3334  
Mauch, T., & Sadler, E. M. 2007, *MNRAS*, 375, 931  
Middelberg, E., et al. 2008, *AJ*, 135, 1276  
Murphy, T., et al. 2013, *PASA*, 30, e006  
Norris, R. P. 1999, *AcAau*, 47, 731  
Norris, R. P., et al. 2006, *AJ*, 132, 2409  
Norris, R. P., et al. 2011, *PASA*, 28, 215

- Norris, R. P., et al. 2013, *PASA*, 30, 20
- Norris, R., et al. 2015, in *Proc. Sci., Advancing Astrophysics with the Square Kilometre Array* (held in Giardini Naxos, 9–13 June), 86
- O’Brien, A. N., Tothill, N. F. H., Norris, R. P., & Filipović, M. D. 2015, *Proc. Sci., EXTRA-RADSUR 2015* (held in Bologna, 20–23 October), 045
- Padovani, P., Miller, N., Kellermann, K. I., Mainieri, V., Rosati, P., & Tozzi, P. 2011, *ApJ*, 740, 20
- Perlmutter, S., et al. 1999, *ApJ*, 517, 565
- Popper, K. 1959, *The Logic of Scientific Discovery* (New York: Basic Books)
- Rees, G. A., et al. 2016, *MNRAS*, 455, 2731
- Riess, A. G., et al., *AJ*, 116, 1009
- Riggs, S., et al. 2016, *MNRAS*, 460, 1486
- Springel, V., et al. 2005, *Nature*, 435, 629
- Sutherland, W., & Saunders, W. 1992, *MNRAS*, 259, 413
- van der Kruit, P. C. 1971, *Astr. Ap.*, 15, 110
- Wilkinson, P. 2007, in *Proc. Sci., From Planets to Dark Energy: The Modern Radio Universe* (held in Manchester, 1–5 October), 144
- Wilkinson, P. 2015, in *Proc. Sci., Advancing Astrophysics with the Square Kilometre Array* (held in Giardini Naxos, 9–13 June), 65
- Wilkinson, P. N., et al. 2004, *NewAR*, 48, 1551
- Williams, R. E., et al. 1996, *AJ*, 112, 1335
- Williams, R. E., et al. 2000, *AJ*, 120, 2735
- Woit, P. 2011, *Not Even Wrong: The Failure of String Theory and the Continuing Challenge to Unify the Laws of Physics* (New York: Random House)