

GrepNova: A tool for amateur supernova hunting

Dominic Ford

This paper presents *GrepNova*, a software package which assists amateur supernova hunters by allowing new observations of galaxies to be compared against historical library images in a highly automated fashion. As each new observation is imported, *GrepNova* automatically identifies a suitable comparison image and rotates it into a common orientation with the new image. The pair can then be blinked on the computer's display to allow a rapid visual search to be made for stars in outburst. *GrepNova* has been in use by Tom Boles at his observatory in Coddenham, Suffolk since 2005 August, where it has assisted in the discovery of 50 supernovae up to 2011 October.



Introduction

Over the past fifteen years, three UK-based amateur astronomers have excelled in the discovery of extra-galactic supernovae. Mark Armstrong made the UK's first amateur supernova co-discovery (1996bo)^{1,2} in 1996 October and with Ron Arbour and Tom Boles, has since gone on to amass a total of 220 discoveries and 16 co-discoveries as of 2011 October 17 (see Table 1 and Figure 1). Boles,

the discoverer of 135 of these, has heavily optimised his Suffolk observatory (above) for his three-telescope patrol programme and has earned the distinction of having discovered more supernovae than any other individual worldwide.

These represent a significant contribution to the total worldwide accumulation of supernova discoveries (see Figure 2). Though only around one in twenty of the supernovae recorded to date have been discovered by amateur astronomers – a fact made particularly apparent by the sharp increase in worldwide discoveries since the advent of the 0.75m Katzman Automatic Imaging Telescope (KAIT) of the highly automated Lick Observatory Supernova Search (LOSS) programme, which has discovered over 850 supernovae since 1997 – much interest surrounds the question of whether some amateur discoveries represent events which would otherwise have been missed by professional surveys. In addition to the discoveries of Boles, Armstrong, Arbour and their colleagues elsewhere in the world, of particular interest in this regard are the 42 discoveries and 5 co-discoveries made by the exclusively visual search of Australian amateur astronomer Revd Robert Evans (see Figure 3).

This paper describes *GrepNova*,³ a software package which I wrote in 2005 to assist amateur supernova hunters in the staple task of comparing new observations of galaxies against historical library images. At the end of an observing run, newly acquired observations can be imported into *GrepNova* as a batch for efficient sequential processing. *GrepNova* automatically searches a database of library images to identify a suitable historical comparison image for each observation and then uses an image registration algorithm to rotate each of these comparison im-

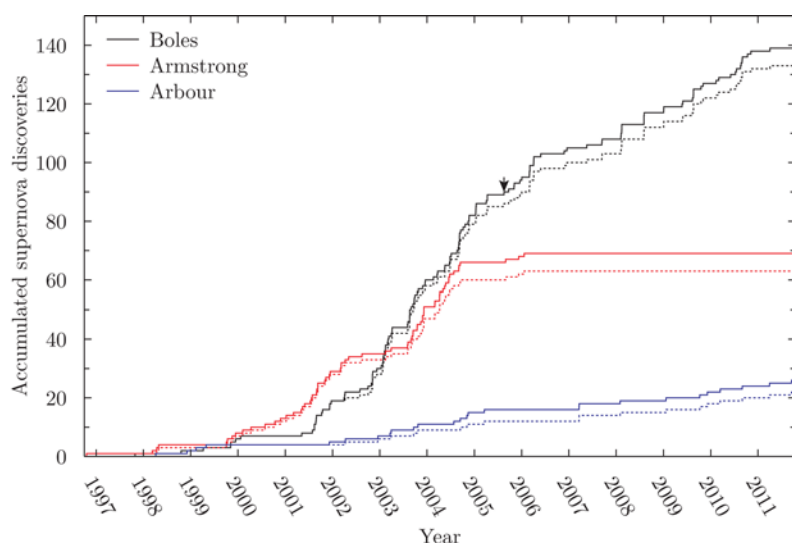


Figure 1. The accumulation of supernova discoveries and co-discoveries by UK-based amateur astronomers up to 2011 October. The dotted lines include only discoveries; the solid lines also include co-discoveries. An arrow indicates the point at which Boles adopted *GrepNova* in 2005 August. Data courtesy of the IAU Central Bureau for Astronomical Telegrams.¹⁷

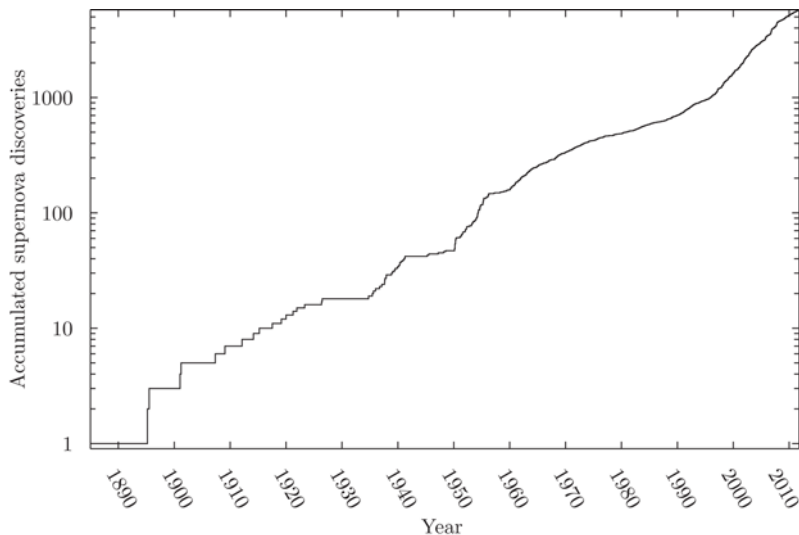


Figure 2. The total worldwide accumulation of supernova discoveries, including both amateur and professional discoveries, up to 2011 October 17. Data courtesy of the IAU Central Bureau for Astronomical Telegrams.¹⁷

ages into a common orientation with the corresponding new observation. The pair are then displayed in *GrepNova*'s image viewer, which includes an option to blink between the two images to ease the process of making a quick and efficient visual search for stars in outburst.

GrepNova has been in use by Boles since 2005 August and was used in the discovery of its first supernova, 2005dj, on 2005 August 18 (see the arrow on Figure 1). Boles has since used it in the discovery of 50 supernovae and in the co-discovery of two further events. *GrepNova* is freely available to all amateur astronomers; it can be downloaded from my website⁵ under the GNU General Public License⁶ (GPL). A binary executable is provided for use under Microsoft Windows and source code is provided for compilation on other platforms.

This paper is structured as follows. The following section describes the scientific motivation for supernova search programmes. Descriptions then follow of the observing technique used to identify them, how *GrepNova* works, and how it can be used.

The astrophysics of supernovae

Supernovae are explosions triggered at the ends of the lives of high-mass stars, which typically release so much energy that they are able for a brief time to match or exceed the total brightnesses of the host galaxies in which they occur. They were first distinguished from common novae by Walter Baade & Fritz Zwicky in 1934,⁷ and spectroscopic examination in the mid-twentieth century led to the realisation that they can be further sub-divided into at least three distinguishable classes.⁸ Supernovae of type II, which show

emission lines due to hydrogen, are understood to occur at the ends of the lives of stars with masses of between around 9 and 16 times that of the Sun. As such stars finally run out of nuclear fuel, their cores undergo catastrophic gravitational collapse to form neutron stars, and their outer parts are ejected with phenomenal force, producing the observed supernovae. Supernovae of types Ib and Ic, which do not show emission lines due to hydrogen, are formed by a similar mechanism, but from progenitor stars with masses greater than 16 solar masses.

The progenitors for supernovae of type Ia are white dwarfs – remnants left over at the ends of the lives of lower mass stars. These remnants have reached a deadlock state in which they are composed primarily of carbon and oxygen which cannot undergo any further nuclear fusion, and yet their internal pressures are great enough that they are also unable to

collapse under gravity to any greater density. A white dwarf can escape from this deadlock state, however, if its mass drifts above around 1.4 solar masses – perhaps because material from a companion star accretes gravitationally onto its surface – in which case the result is a catastrophic nuclear ignition, ending in a thermonuclear explosion which thrusts the entire star apart within a few seconds to leave no central core at all.

Understanding supernovae is important to many areas of astrophysics. For example, they are believed to play a central rôle in the evolution of galaxies, enriching the interstellar gas – and any future stellar systems which form out of it – with elements such as carbon which are vital for life, and at the same time producing shock waves which can either inhibit or aid the formation of stars in the interstellar gas they pass through under different circumstances. In recent years, supernovae of type Ia have attracted especial interest because of their use in estimating the distances of very distant galax-

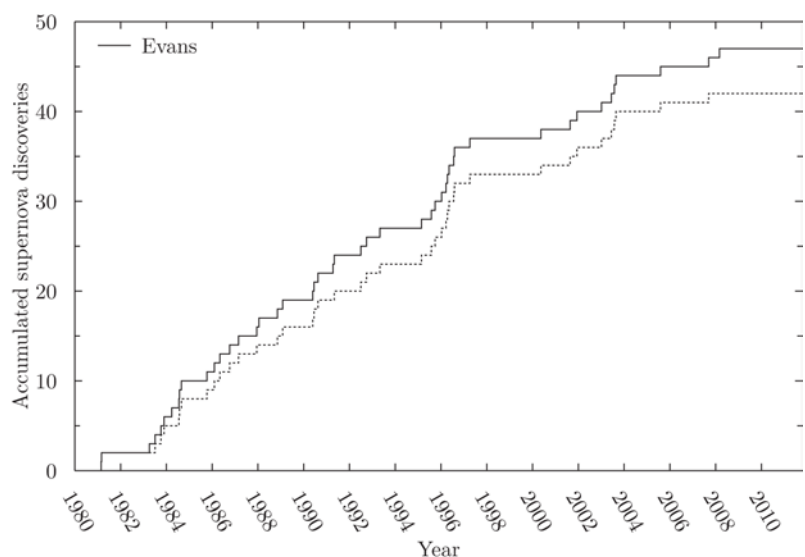


Figure 3. The accumulation of supernova discoveries and co-discoveries by the visual search programme of Robert Evans, up to 2011 October. The dotted line includes only discoveries; the solid line also includes co-discoveries. Data courtesy of the IAU Central Bureau for Astronomical Telegrams.¹⁷

Table 1. Numbers of supernova discoveries and co-discoveries by UK-based amateur astronomers up to 2011 October

Nos. of discoveries made using GrepNova are shown in brackets. Data courtesy the IAU Central Bureau for Astronomical Telegrams.¹⁷

| Observer | Discoveries | Co-discoveries |
|-----------|-------------|----------------|
| Arbour | 22 | 4 |
| Armstrong | 63 | 6 |
| Boles | 135 (50) | 6 (2) |

ies, which have in turn been central to the evidence which has been put forward for the existence of dark energy.

Though their extreme luminosities might suggest them to be rather conspicuous events, the rate at which supernovae occur in any given galaxy is difficult to estimate. It is almost certain that many events go completely unobserved. This is illustrated by the fact that no supernovae have been observed in the Milky Way since those seen in 1572 by Tycho Brahe, and in 1604 by Johannes Kepler, despite other evidence which leaves little doubt that many supernovae have occurred in our Galaxy since then. The sheer number of galactic supernova remnants visible,⁹ and especially individual examples such as the radio source Cas A which can be confidently dated at around 300 years old, indicates that a significant population of more recent supernovae has gone entirely unobserved. Given the active state of telescopic astronomy in the late 17th century, it is simply implausible that the appearance of a bright new star in Cassiopeia could have gone unnoticed or unrecorded, so it seems likely that it – like many other supernovae – was masked from optical visibility by the obscuring dust grains which pervade the Galaxy. A similar conclusion must also be reached for the remnant G1.9+0.3, which has recently been dated to around 150 years old.¹⁰

Whilst our position within the Milky Way, in the midst of a disc of obscuring dust, is especially unpropitious for observing supernovae within it, search programs for supernovae in other galaxies are almost certainly plagued by similar incompleteness. Observed supernova rates in spiral galaxies are highly dependent upon the inclinations of the galaxies to our line of sight, with the highest rates recorded among face-on galaxies.^{11,12} In the extreme case of the edge-on dwarf M82, though radio observations suggest a supernova rate of around one every ten years, none has ever been detected optically.¹³ Further evidence for incompleteness stems from the uncomfortable fact that the closest supernova to have been observed in modern times – 1987A in the Large Magellanic Cloud (LMC) – is also the least luminous to have been observed. It remains a matter for debate whether SN 1987A was highly atypical, or whether similar events routinely go unnoticed in more distant galaxies because of their faintness.¹⁴

It is probable that, once the incompleteness of supernova catalogues has been taken into account, around one or two events per century occur in spiral galaxies such as the Milky Way. Rates are certainly somewhat higher in prodigiously star-forming systems such as M82, and much less in elliptical galaxies. A more thorough discussion of the evidence leading to such estimates can be found in, for example, van den Bergh & Tammann (1991).⁸

Search techniques

The infrequency with which supernovae occur in any single galaxy means that to observe an appreciable number, it is necessary to monitor a large number of galaxies. Search programmes typically entail the regular inspection of a large number of nearby galaxies, each of which is compared against a library image to check for the appearance of new stars. The degree of automation in this process varies. At one extreme, Revd Evans searches visually, memorising the morphologies of galaxies and taking advantage of the tremendous dynamic range of the human eye to scan for faint supernovae. At the opposite extreme, professional searches such as LOSS are automated to the point where promising candidates can be identified and automatically re-observed before being brought to the attention of any human operator.

Most amateur searches lie between these extremes. For example, Boles uses three Celestron C14 (355mm) Schmidt–Cassegrain telescopes, each mounted on robotically-controlled Paramount German equatorial mounts and equipped with Apogee AP7 CCD cameras. A combination of 30-second and 60-second exposures is taken, yielding around 210 or 135 images per hour respectively once slewing time has been accounted for.¹⁵ Over the course of a year, these telescopes monitor a sample of 12,000 galaxies as regularly as the weather and the changing seasons permit.

Though computerised search programmes such as LOSS may appear more efficient, amateur searches have the arguable advantage that the human eyeball and brain is much closer to the raw data and in a better position to detect events of unusual or marginal appearance. This is scientifically important when supernova surveys are so incomplete and potentially biased towards particular kinds of event. For example, the catalogue of discoveries from Evans' visual search for supernovae in 855 Shapley–Ames galaxies between 1980 and 1988¹⁶ has been argued by some to differ in its average properties as compared with samples found by photographic searches.⁸

Image comparison using *GrepNova*

To use *GrepNova*, it is necessary first to compile a library of reference images of the galaxies which are to be patrolled, against which future observations can be compared. These should be stored in FITS format in a common directory in the computer's filing system, whose path is given as a setting to *GrepNova*. Each time *GrepNova* is started, this directory is scanned and a database of the available library images compiled, using the OBJECT header field in each of the FITS files – normally set automatically by telescope control software – to determine a galaxy name to associate with each image.

At the end of an observing run, batches of new observations can be imported into *GrepNova* in a similar fashion, by placing them in a common directory whose file path is passed

to *GrepNova*. The new observations are then processed in sequence. For each image, the FITS headers are interrogated to determine the name of the galaxy observed, which is used to identify a suitable library image against which the image might be compared. To ease the visual comparison of the two images, the library image is shifted into a common orientation with the new image using the pattern recognition algorithm described below, which works using the positions of field stars in the two images.

Once this automated procedure is completed, the user is able to visually inspect each of the new observations in turn in *GrepNova*'s image viewer, which can be set to blink slowly between the new and comparison images, or toggled between them manually. The brightnesses of the two images can be changed individually to maximise the contrast in specific structures of interest and to match the brightnesses of the two images as closely as possible.

Figure 4 shows a screenshot of *GrepNova*'s user interface displaying the discovery image of SN 2009B, the 28th supernova discovered using *GrepNova*.

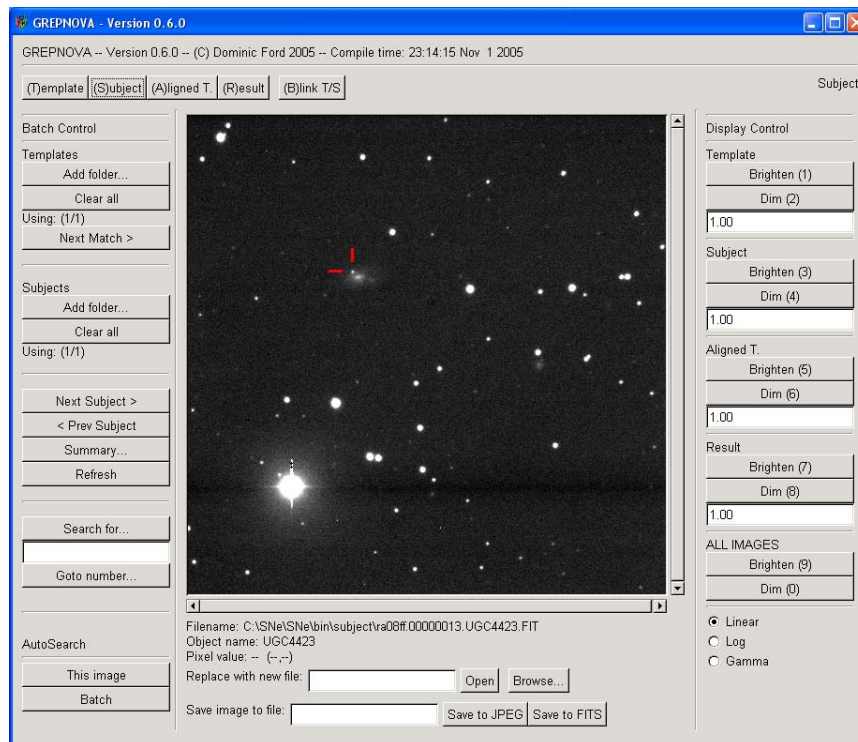


Figure 4. A screenshot of *GrepNova*'s user interface, showing the discovery image of SN 2009B (marked) in UGC 4423. This was the 28th supernova to have been discovered using *GrepNova*.

Registration algorithm

Rotating two images into a common orientation requires the identification of common features between the two images which should be brought into alignment. In the images processed by *GrepNova*, the most readily available such features are field stars, of which there are always at least several tens visible, even in fields well removed from the galactic plane. *GrepNova* applies a simple pattern recognition algorithm to these stars, beginning by identifying the twenty brightest stars in each image. For the purposes of this task, the brightness of each star is crudely defined by adding the pixel brightnesses within a circular aperture whose radius is set to equal an estimate of the size of the point-spread function of the image.

Cycling through all of the possible groupings of three stars among these twenty, a total of $20 \times 19 \times 18 / 3! = 1,140$ distinct triangles can be constructed with one star at each vertex. Each of these may be characterised by its interior angles, which can typically be measured to within an accuracy of arcminutes when the sides measure several hundred pixels. Taking the two largest of these angles to be, respectively, α and β , the triangle can be assigned a two-dimensional fingerprint in the region of parameter space $60^\circ < \alpha < 180^\circ$; $(180^\circ - \alpha) / 2 < \beta < 180^\circ - \alpha$, which is defined by the fact that the interior angles of the triangle must sum to 180° .

In two images which are offset from one another, slightly differing sets of bright stars will have been found, but provided there is a reasonable overlap between the images, a sizeable number of common triangles will have been identified. Alignment of the two images proceeds by searching for triangles in the two images whose (α, β) -positions match to within a distance of ϕ of one another, where ϕ is chosen on a triangle-by-triangle basis to reflect the uncertainties in α and β assuming a one-pixel uncertainty in the position of each star. Each identified match is suggestive that an orientation should be sought in which the three vertices of the triangle map onto one another, but has a small probability of being spurious.

Once a complete list of such match candidates has been compiled, *GrepNova* begins searching for a way to re-orient the library image such that these matching triangles fall into alignment. In general, re-orienting photographs requires correction for the projection of the spherical sky onto a flat plane which they represent. However, this is not a serious issue for images measuring less than around a degree across, which measure the celestial sphere on sufficiently small scales that it is approximately flat. Supernova patrol images are typically smaller than this, and so *GrepNova* assumes that the problem can be simplified to one of moving and rotating a flat image around on a flat surface, with an additional parameter allowing it to be enlarged or reduced. This means that any given re-orientation can be represented by four parameters: a horizontal translation Δx , a vertical translation Δy , a rotation Θ , and a change of magnification m .

These are found using a standard numerical optimisation routine from the GNU Scientific Library (GSL) – the `nmsimplex` algorithm implemented by the `gsl_multimin_*` function set – which searches for the values of $(m, \Delta x, \Delta y, \Theta)$ which mini-

mise the sum of the squares of the distances separating the match candidates. Some of the match candidates may be spurious – because the images happen by chance to contain multiple triangles of stars with very similar shapes – and these are identified by the fact that they align significantly less well than the other triangles. These are removed from the fitting process, which is then iterated until all the match candidates are in close alignment.

Summary

GrepNova has automated much of the potentially time-consuming work which needs to be done to prepare library images for comparison with newly observed images as part of supernova patrol programmes. It allows a visual comparison to be made almost instantaneously when new images are received from telescope control software.

The search procedure could be further automated if, having aligned each pair of images, *GrepNova* automatically subtracted them, searched the resulting difference images for new stars, and flagged these up as potential supernova candidates. This degree of automation is achieved in professional search programmes such as LOSS, but is far from straightforward to implement. In order to maximise the number of galaxies which can be observed during occasional spells of exceptionally clear weather, it is highly advantageous to use an unfiltered CCD to maximise the amount of light collected across the visible spectrum, as this minimises the length of exposure required on each galaxy. However, since both atmospheric seeing and transparency are wavelength-dependent, objects of differing colours may appear to have different apparent brightnesses in white-light images on different nights. Thus, automated image comparison software must account in detail for atmospheric effects when subtracting such images.

In practice, the number of discoveries made by Boles' search programme is determined principally by the number of observations which can be made, rather than the time required to inspect the images, hence while *GrepNova* has

Ford: *GrepNova: a tool for amateur supernova hunting*

proven a tremendously labour-saving device, there would be little benefit from further automation.

Acknowledgments

The development of *GrepNova* would not have been possible without a great deal of guidance from Tom Boles, who also provided the discovery image shown in Figure 4, and made a number of useful comments during the preparation of this paper. The author is also grateful to Martin Mobberley for useful comments made during the development of the software.

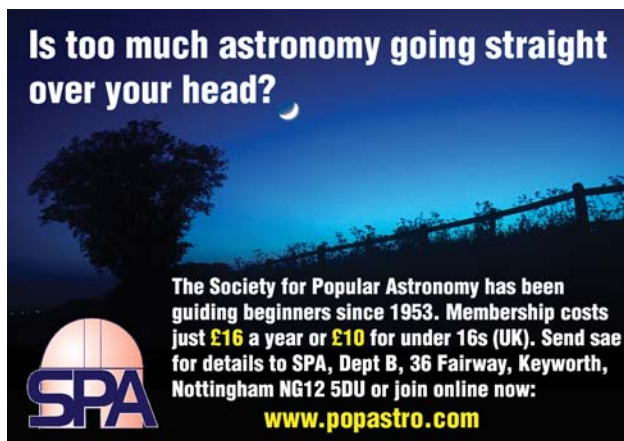
Address: 37 Coleridge Road, Cambridge, CBI 3PH. [dcf21@dcford.org.uk]

Notes and references

- 1 Throughout this paper, the *discoverer* of a supernova is taken to be the first named discoverer as listed by the International Astronomical Union (IAU). Observers whose names are listed by the IAU after a semi-colon are described as *co-discoverers*. Elsewhere in the literature, co-discoveries are commonly included in personal discovery tallies.
- 2 IAU Circular 6487 (1996)
- 3 The name *GrepNova* is derived from that of the Unix search command *grep*.
- 4 IAU observatory no. 234, located at 52°08'43"N, 01°07'43"E.
- 5 <http://dcford.org.uk/~dcf21/software.shtml>
- 6 <http://www.gnu.org/licenses/gpl.html>
- 7 Baade W. & Zwicky F, *Proc. Natl. Acad. Sci. USA*, **20**, 254 (1934)
- 8 See, e.g., van den Bergh S. & Tamman G. A., *Ann. Rev. Astron. Astrophys.*, **29**, 363 (1991)
- 9 Green D. A., *Bull. Astr. Soc. India*, **32**, 335 (2004)
- 10 Green D. A., *MNRAS*, **387**, L54 (2008)
- 11 van den Bergh S. & McClure R. D., *ApJ*, **359**, 277 (1990)
- 12 van den Bergh S., *Astron. Astrophys. Lett.*, **231**, L27 (1990)
- 13 Kronberg P. P. *et al.*, *ApJ*, **246**, 751 (1981)
- 14 Catchpole R. M. *et al.*, *MNRAS*, **229**, 15 (1987)
- 15 Boles T., *J. Brit. Astron. Assoc.*, **116**(5), 229 (2006)
- 16 Evans R. *et al.*, *ApJ*, **345**, 752 (1989)
- 17 <http://www.cfa.harvard.edu/iau/lists/Supernovae.html>

Received 2010 October 11; accepted 2010 November 20

[Editor's note: When preparing this paper for press, the supernova discovery figures were updated to reflect the situation as at 2011 October 17.]



Is too much astronomy going straight over your head?

The Society for Popular Astronomy has been guiding beginners since 1953. Membership costs just £16 a year or £10 for under 16s (UK). Send sae for details to SPA, Dept B, 36 Fairway, Keyworth, Nottingham NG12 5DU or join online now: www.popastro.com

THE ASTRONOMER

A monthly magazine providing rapid reports by amateurs in the UK and worldwide of all types of observable objects, together with discovery news from IAU reports.

CIRCULARS, available as an extra option, bring news of newly discovered objects reported to us via our IAU e-mail link.

Can you afford to be without the most comprehensive and rapid service available to observers today?

Write for details and application form, or send £1.00 for sample copy, to:

Bob Dryden, 21 Cross Road, Cholsey, Oxon. OX10 9PE.
Tel. (01491) 201620. E-mail bobdryden@ntlworld.com