

Mary, a Pipeline to Aid Discovery of Optical Transients

I. Andreoni^{1,2,3,4,5}, C. Jacobs¹, S. Hegarty¹, T. Pritchard^{1,3}, J. Cooke^{1,2,3} and S. Ryder⁴

¹Swinburne University of Technology, PO Box 218, Hawthorn, VIC 3122, Australia

²ARC Centre of Excellence for All Sky Astrophysics (CAASTRO), Australia

³ARC Centre of Excellence for Gravitational Wave Discovery (OzGrav), Australia

⁴Australian Astronomical Observatory, 105 Delhi Rd, North Ryde, NSW 2113, Australia

⁵Email: igor.andreoni@gmail.com

(RECEIVED December 24, 2016; ACCEPTED August 7, 2017)

Abstract

The ability to quickly detect transient sources in optical images and trigger multi-wavelength follow up is key for the discovery of fast transients. These include events rare and difficult to detect such as kilonovae, supernova shock breakout, and ‘orphan’ Gamma-ray Burst afterglows. We present the *Mary* pipeline, a (mostly) automated tool to discover transients during high-cadenced observations with the Dark Energy Camera at Cerro Tololo Inter-American Observatory (CTIO). The observations are part of the ‘Deeper Wider Faster’ programme, a multi-facility, multi-wavelength programme designed to discover fast transients, including counterparts to Fast Radio Bursts and gravitational waves. Our tests of the *Mary* pipeline on Dark Energy Camera images return a false positive rate of $\sim 2.2\%$ and a missed fraction of $\sim 3.4\%$ obtained in less than 2 min, which proves the pipeline to be suitable for rapid and high-quality transient searches. The pipeline can be adapted to search for transients in data obtained with imagers other than Dark Energy Camera.

Keywords: Gravitational waves – Methods: data analysis – novae, cataclysmic variables – supernovae: general – techniques: image processing

1 INTRODUCTION

The direct detection of gravitational waves and the design of new facilities to identify fast optical transients open fascinating perspectives for the future of time-domain astronomy. Both large observing programmes and small ‘boutique’ experiments are taking part in the search for electromagnetic counterparts to gravitational wave events discovered by the LIGO/Virgo collaboration (e.g., Abbott et al. 2016). These programmes can involve large aperture, narrow-field facilities such as the Keck telescopes, as well as smaller telescopes with wide field of view such as the Australian National University’s SkyMapper telescope (Keller et al. 2007).

Fast and faint transients in the nearby Universe are among the most intriguing but elusive events to detect in time-domain surveys. For example, they are considered to be the most promising counterparts to gravitational waves at optical wavelengths. These sources include kilonovae (Tanvir et al. 2013; Metzger 2016 and references therein), gamma-ray bursts—detected as prompt emission, afterglow, or orphans (Vestrand et al. 2014; Cenko et al. 2015; Ghirlanda et al. 2015), and supernova shock breakouts (Garnavich et al. 2016) which are transients difficult to detect because of their

expected low-luminosity and/or short timescale. A confident identification and rapid follow-up of such transients becomes the key to successful multi-messenger studies (Chu et al. 2016). The most outstanding discoveries in the near future are likely to depend on the specifics and performance of the telescopes involved, as well as the speed of the data analysis. In fact, the prompt identification of interesting fast transients allows rapid follow-up that can provide much more valuable information than the detection of the transient alone, even if the transient is detected with a large telescope.

Currently, the Dark Energy Camera (DECam; Flaugher et al. 2015) is one of the best imagers in the southern hemisphere to catch fast and faint optical transients. DECam combines a ~ 3 deg² field of view per pointing with the depth reachable with the 4-m Blanco telescope at CTIO, allowing to probe extremely large volumes, being sensitive to extragalactic fast transients out to redshift ~ 2 . As such, DECam is an excellent instrument to perform observations aimed at detecting optical counterparts to Fast Radio Bursts (Petroff et al. 2017), or counterparts to gravitational wave events in response to LIGO triggers (Cowperthwaite et al. 2016). Its potential to detect kilonovae (Doctor et al. 2016), supernova shock breakout (Förster et al. 2016), and even new classes of fast transients makes DECam suitable also to initiate

‘reverse’ searches in LIGO data for gravitational wave signals at lower significance (Waş et al. 2012).

Several pipelines have been designed to detect optical transients, usually focusing on objects evolving on timescales of a few days to months. Such pipelines are usually designed for surveys such as, among others, the Supernova Legacy Survey (Astier et al. 2006), the Dark Energy Survey (Kessler et al. 2015), SkyMapper (Keller et al. 2007), the Palomar Transient Factory (Law et al. 2009), Pan-STARRS (Rest et al. 2005), and the Catalina Sky Survey (Djorgovski et al. 2010). These programmes unveiled many new classes of Galactic and extragalactic events by regularly surveying large swaths of the sky. In some cases, pipelines were optimised to allow near real-time analysis and detection of transients within minutes from the data acquisition (e.g., Perrett et al. 2010; Rest et al. 2014; Cao, Nugent, & Kasliwal 2016; Förster et al. 2016), leading to the detection of fast transients such as orphan gamma-ray bursts (Cenko et al. 2015). The development of new pipelines able to process huge amounts of data quickly becomes even more important for programmes using the upcoming Zwicky Transient Facility (ZTF, Smith et al. 2014; Bellm 2014) and the future Large Synoptic Survey Telescope (LSST, Ivezic et al. 2008).

In this paper, we present *Mary*¹, an example of a custom-made pipeline that quickly and effectively analyses optical images to detect transient events. We coded *Mary* to search for transients in near real time with DECam, in the framework of the ‘Deeper Wider Faster’ (DWF) programme. We give a brief overview of the ‘Deeper Wider Faster’ programme in Section 2, we describe the *Mary* code in Section 3 and we present the tests performed on DECam images in Section 4. A discussion of the parameters choices, the tests results, the possible application of *Mary* to data from other optical telescopes, and concluding remarks (Sections 5 and 6) complete the paper.

2 THE ‘DEEPER WIDER FASTER’ PROGRAMME

The DWF programme (Cooke et al., in preparation; Vohl et al. 2017; Andreoni et al. 2017a, 2017b) organises and conducts several components to detect and study fast transients.

The first component coordinates *simultaneous* observations of the same targeted field, among a list of target fields, with multiple facilities operating at different wavelengths. In 2017 February, DWF included simultaneous observations with the DECam optical imager, the Parkes, Molonglo, and Australia Telescope Compact Array radio telescopes, the Rapid Eye Mount optical/infrared telescope in Chile, and the optical/UV/X-Ray/Gamma ray telescopes mounted on the NASA *Swift*² satellite. The observing strategy with DECam is based on fast (20 s exposure time), continuous sampling of

successive target fields for 1–3 h each per night. The fields are chosen depending on their visibility from both Chile and Australia and are observed for multiple consecutive nights, usually six per semester.

The second component consists of the near real-time data analysis and candidate assessment that takes place in the collaborative workspace environment organised for DWF at Melbourne University or at the Swinburne University of Technology (Meade et al. 2017).

The third component of DWF consists of rapid response and conventional Target of Opportunity (ToO, imaging, and spectroscopy) follow-up. While some of these facilities can follow up transients, only a few hours after their identification due to the geographical location of the observatories, the Gemini South 8-m telescope can perform rapid spectroscopy minutes after significant discoveries due to its proximity to the Blanco telescope, as well as its rapid response ToO programme. Classically scheduled follow-up observations (imaging and spectroscopy) constitute the fourth component of the DWF programme.

Fast and effective data analysis is crucial to catch fast transients and successfully obtain rapid response, ‘flash’ spectroscopy, which can add valuable information to understand the mechanisms at the basis of supernova explosions (Gal-Yam 2014), including optical shock breakout. It also offers our only opportunity to study and characterise new exotic fast transients. DWF organised a network of small to large telescopes to follow up the transients discovered with the *Mary* code, including the Anglo-Australian Telescope (AAT), Gemini South, the Southern African Large Telescope (SALT), and the SkyMapper and Zadko telescopes (Coward et al. 2016) in Australia.

During DWF campaigns, *Mary* automatically generates lists of highly significant candidates, providing information that populates the database described in Section 3.6. Experts validate and further prioritise the candidates using light curves, catalogue information, and products for the visualisation generated with *Mary*. This process leads to manual triggers of large, narrow-field facilities to follow up the most interesting events.

3 THE MARY PIPELINE

We wrote the *Mary* pipeline to serve as a fast, simple, and flexible tool to analyse the 60 CCD mosaic of DECam in near real time (seconds to minutes after the data are acquired) during DWF observations. The pipeline is also suitable to perform accurate late-time analysis on hundreds of images, with the goal of discovering transients of any type, evolving at nearly any timescale and with nearly any trend in their light curve. However, the main focus of the DWF programme is to unveil fast transients, especially the least understood ones. The pipeline is adaptable to telescopes/imagers other than Blanco/DECam, as we discuss in Section 5.2.

¹ The name *Mary* is a reference to Mary Shelley, author of the ‘Frankenstein’ novel; the name was chosen to express the idea of several pieces of code brought together to form the final ‘Creature’.

² swift.gsfc.nasa.gov

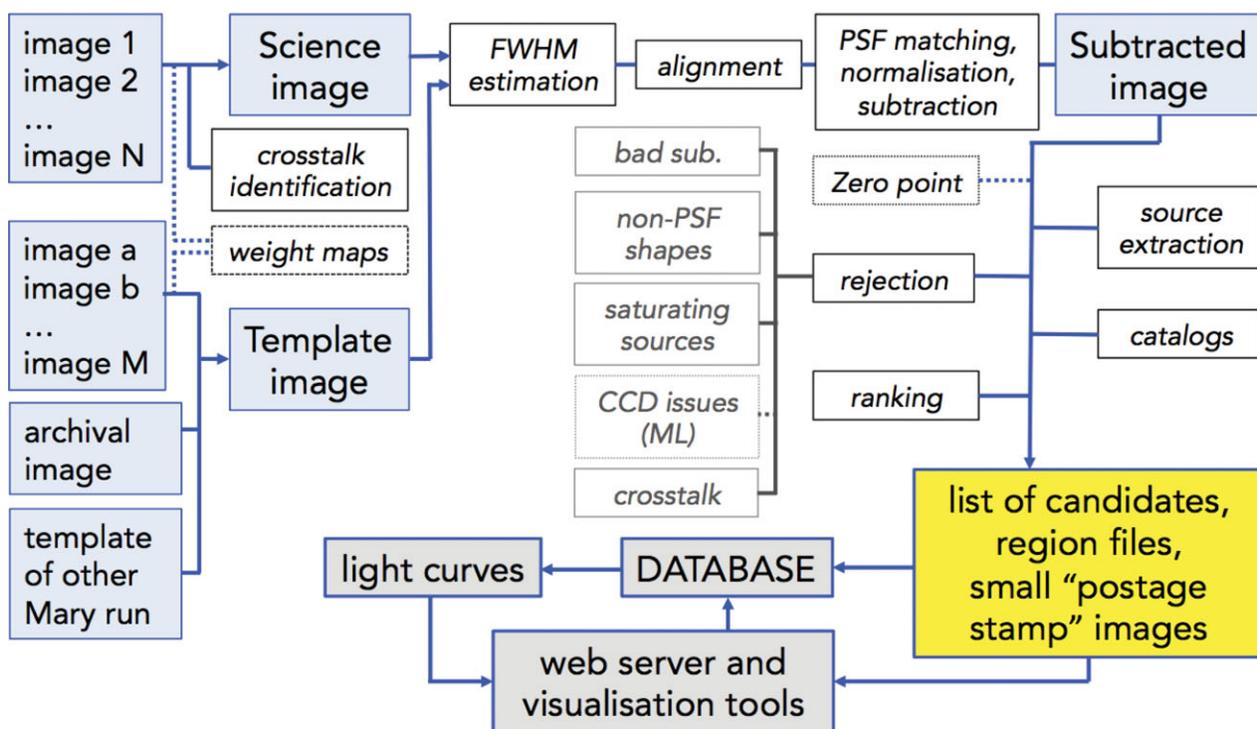


Figure 1. The *Mary* workflow, highlighting the main steps leading to the identification of transient events. Dashed lines indicate optional steps that the user can activate. All the indicated operations happen in parallel for each of the 60 functioning CCDs of DECam. The main steps for the selection of good candidates include the rejection of those sources which are badly subtracted, non-PSF shaped, saturating, possibly associated with crosstalk effects, or flagged as possible CCD artefacts by the machine learning (ML) classifier. The storage of the information in a database, the generation of the light curves, and the display of the products for the visualisation take place outside the parallel process, thus we marked them with a grey background.

The lack of a well-defined ‘training set’ at the beginning of the novel DWF project forced us to think of a new way to distinguish real transients from cosmic rays and bogus detections without adopting machine-learning techniques (Section 3.3). A machine-learning step was introduced in 2017 February to reject a specific type of CCD artefact, as described in Section 3.4. We present how *Mary* prioritises the candidates in Section 3.5 and the results of timing and efficiency tests in Section 4.1–2. Finally, we compare the results with the performance achieved by some of the best pipelines in Section 5.1. In the Appendix, we report a table of the parameters—set automatically or by the user—that fully regulate the analysis that *Mary* performs.

Mary is coded in the IDL programming language and runs on CPU machines on the Green II supercomputer at Swinburne University of Technology. It ties together popular packages such as the AstrOmatic³ packages Swarp and SExtractor (Bertin & Arnouts 1996; Bertin et al. 2002), HOTPANTS (Becker 2015), and new custom programmes that facilitate the workflow and the specific selection of candidate transient sources. In this paper, we assume that the images have already been astrometrically calibrated and pre-processed with bias and flat field corrections. We will discuss individually the main stages driving the workflow (shown in its most

basic structure in Figure 1) of the pipeline: the preparation of a ‘template’ image to use as a reference and a ‘science’ image, the image subtraction, the selection of the candidates, and the creation of products to facilitate the visual inspection.

3.1. ‘Template’ and ‘science’ images

Mary takes advantage of image subtraction techniques to identify sources changing their flux between different epochs. The science image is obtained coadding a number ‘N’ of calibrated images (Figure 1), or using an individual image when only one is available or during the search for fast transients. The template image can be obtained in the following three ways:

- Using a deep archival template image, well separated in time with respect to the epoch of the observation, depending on the timescale at which the target transients evolve. This method usually allows the discovery of the largest number of transients.
- Coadding a number ‘M’ of individual calibrated images. This becomes particularly useful when an ‘old’ template image is not available, for example, during the follow-up of a gravitational wave trigger in a region of the sky not previously covered.

³ <http://www.astromatic.net/>

- Copying the template from a previous *Mary* run. This method becomes handy when the analysis is performed in near real time with continuous imaging of the same field, as it sensibly reduces the *Mary* running time. The time saved depends on the number of images to be coadded when the user chooses to avoid copying the template.

Mary coadds the input images using the Swarp package. The user can switch on/off the optional automatic creation and usage of weight maps. After the creation of the science and template images, *Mary* automatically estimates the seeing FWHM of the images in the following way:

- Running SExtractor on the images.
- Selecting those sources with a Star/Galaxy classification ranking >0.98 . If less than three sources fulfil this requirement, sources with a ranking >0.96 are selected.
- Calculating the average FWHM_IMAGE that SExtractor outputs for the selected sources.

The automatic estimation of the seeing FWHM allows the definition of other parameters that, otherwise, would have to be set manually by the user for each set of images to search with *Mary*: those parameters are among those marked with the letter ‘A’ in Table A1. Also, this module prepares a list of bright and saturating sources that can be removed from the final list of candidates at later stages (see Section 3.3).

When science and template images are ready to use, they are aligned in image coordinates using the resampling functions of Swarp. Usually, it is convenient to ‘frame’ the common area between science and the template images (using the *xdim*, *ydim*, *fixoldchoice* parameters; see Appendix A) to allow (i) a proper alignment in image coordinates; and (ii) a more polished image subtraction, thanks to a maximisation of the common sky area over the non-overlapping area. This framing operation is not needed when the template image is the coaddition of many largely dithered images, covering the whole area of the science image. During DWF observations, we rarely have such images available, because even the best templates available leave ‘stripes’ (a few pixels wide) of the new science images without coverage. Therefore, knowing the shifts occurring during the observations and the precision of the astrometric calibration, the pipeline can be set up to different sizes for the final aligned images. The optimal size usually adopted in near real time and during the tests presented in this paper is $3\,800 \times 1\,800$ pixel, while for other works, we adopt $4\,000 \times 2\,000$ pixel or the exact dimensions of the CCD. The aligned images are kept available for the user to be displayed with region files overplotted (Section 3.6).

A bad alignment ($\geq 1.0 \times$ seeing FWHM) would result in the failure of the image subtraction, or in a non-perfect subtraction that results in an error message, usually triggered by the unexpectedly large number of candidates exceeding the *maxsources* number set in the parameters file (with a default value of 60 per CCD during the near real-time analysis). The images are stored rather than deleted as temporary files, in

order to be checked at any time for quality assessment and to facilitate the generation of light curves with programmes that we coded separately to the *Mary* pipeline.

3.2. Image subtraction

When science and template images are aligned in physical coordinates, *Mary* uses the HOTPANTS package to perform the main steps of image subtraction, such as sky subtraction, normalisation of the flux and PSF matching. *Mary* computes input parameters such as the average sky level and its standard deviation. The user can indicate the number of standard deviations of the sky to set the threshold below which pixels are considered ‘bad pixels’ to be masked. Similarly, the user can set the maximum pixel value; otherwise, HOTPANTS extracts this information from the header of the image (‘SATURATION’ keyword). Finally, the user can switch on/off the creation of data quality masks with HOTPANTS depending on the minimum and maximum pixel values to limit the effect of saturating sources, bad columns, and bad pixels. An optional step (activated with the *photchoice* parameter) determines the zero point of the template image by measuring the magnitude of a number (usually 30–50 per CCD) of pre-selected stars present in the target field and comparing them with the values reported in existing catalogues. Specifically, we use the USNO-B1 or Gaia catalogues (determined with the *catalogcalib* parameter), as the Sloan Digital Sky Survey is largely incomplete in the southern hemisphere.

It may be possible to implement the optimal subtraction algorithm (Zackay, Ofek, & Gal-Yam 2016) in future versions of the *Mary* code, with the intent to compare the results against those presented in this paper.

3.3. Selection of the candidates

Sources varying their luminosity between the template and the science images leave residuals on the images generated by the subtraction with HOTPANTS. We run the SExtractor package on the image resulting from the subtraction to detect and measure such residuals. The user controls the main extraction and photometry parameters (DETECT_MINAREA, THRESH_TYPE, DETECT_THRESH, ANALYSIS_THRESH) from the *Mary* parameter file. The result consists of a catalogue of a (typically) large number of sources, dominated by non-astrophysical detections.

We first apply a number of criteria to pre-select candidates, for the specific instrumental parameters of DECam and the seeing FWHM at CTIO, such as a shape elongation upper limit, a minimum isophotal area, and minimum FWHM value for the source (regulated by the *elmax*, *isoareafmin*, *fwhm* parameters, see Appendix). The default values we set during the DWF observations are *elmax*=1.8, *isoareafmin*=20, and a minimum FWHM value for the residual equal to half the *fwhm* parameter associated with the subtracted image. This initial step allows the rejection of most of the remaining

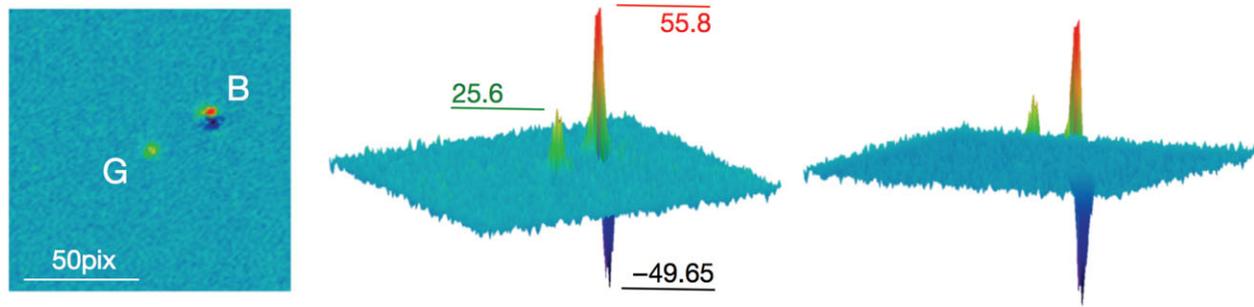


Figure 2. Example of image subtraction leaving ‘good’ or ‘bad’ residuals, indicated with the letters G and B in the left figure, respectively. The figures at the centre and at the right present a 3D rendering (generated with the AAOGLIMPSE software, Shortridge 2012) of the result of the image subtraction, from two different viewing angles. The G residual lacks the negative peak associated with the B residual. ADU values of the positive and negative peaks are indicated in the central figure.

cosmic rays (especially when the science image is created without median stacking ≥ 3 images) and highly asymmetric artefacts. In fact, bright and saturated point sources rarely subtract properly and usually leave a ‘negative’ flux signature adjacent to a ‘positive’ flux excess (Figure 2). Data quality masks generated with HOTPANTS (see Section 3.2) help reject such bogus sources. In addition, *Mary* creates an additional catalogue on the inverse (i.e., taking $-1.0 \times \text{VALUE}$ per pixel) of the subtraction products and, after matching the two catalogues, rejects all candidates associated with detection in the ‘negative’ image. After these selection steps, more than 99% of bogus sources are rejected without the need of machine learning, when the images are aligned within \sim the FWHM of the image with poorer seeing. Completeness and efficiency tests of these selection methods are reported in Section 4.2.

In the specific case of the double-amplifier CCDs of DECam, electronic crosstalk⁴ associated with bright point sources can cause ‘ghost’ sources to appear at a location on the CCD symmetrical with respect to the contact line between the two amplifiers. Crosstalk features are often similar in appearance to real-point sources. Before the coaddition of the images, *Mary* identifies bright sources (saturating or brighter than a user-defined threshold) and flags the regions of the CCD where crosstalk effects are expected.

3.4. Machine-learning classifier to reject CCD artefacts

In order to further purify our candidate sample, we employ machine learning to reject bad pixels and CCD artefacts, such as those shown in the left panel of Figure 3. Machine-learning techniques take a data-driven approach to problems such as the classification and clustering of data in various parameter spaces, instead of relying on human intuition regarding the significant features in the data (see Jordan & Mitchell 2015, for an overview). Here we consider a *supervised learning*

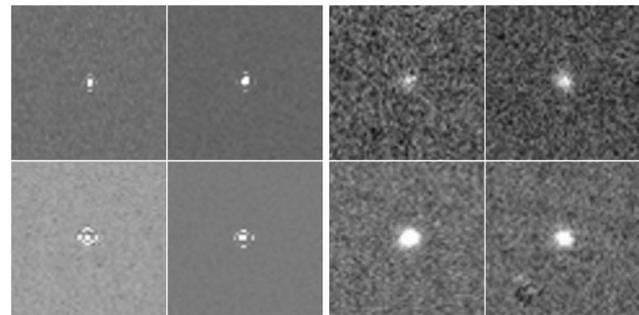


Figure 3. Examples of artefacts (left) and possible good candidates (right) present in the training set used for machine learning. Our training set consisted of 559 false positives and 524 potential transients selected by visual inspection. Only the central 16×16 pixel area was considered to train the machine-learning algorithm.

problem, where we used a pre-prepared, labelled data set to train our algorithm to classify the candidates into one of two categories (candidate or artefact).

To prepare the training set, we visually inspected $\sim 1\,000$ candidates from the selection described in Section 3.3 and classified them as ‘good’ (possible transients) and ‘artefacts’ (unlikely to be genuine sources). This resulted in a set of 524 good candidates and 559 false positives. We augmented this training set by rotating each image through three 90° rotations, for a total training set of size 4 332. Given that we are searching for point sources (with typical FWHM = 1–2 arcsec), and we want to keep good candidates even if they occur close to one of these artefacts, we restricted our training images to a 16×16 pixel stamp (equivalent to 4.2 arcsec on a side) centred on the candidate object. We reserved 20% of the training images as a test set, in order to evaluate the performance of trained algorithms on data not used for training purposes. Some sample images from our training set are depicted in Figure 3.

We used this training set to train two machine-learning methods. As the inputs are 16×16 pixel areas of a single-band FITS file—256 floating point numbers—they are of a dimensionality tractable by most traditional classification

⁴ See the DECam known problems web page: <http://www.ctio.noao.edu/noao/node/2630>

Table 1. Training and test set accuracy for the machine-learning methods, trained on a training set of artefacts and transient candidates. SVM = Support Vector Machine; CNN = Convolutional Neural Network. We incorporated the CNN-based classifier into the *Mary* pipeline.

	Training accuracy	Test set accuracy
SVM	98.11%	97.21%
CNN	97.96%	98.79%

techniques. We first employed a support vector machine (SVM, Cortes & Vapnik 1995), an efficient classifier able to distinguish between classes with highly non-linear decision boundaries. We also applied a convolutional neural network (CNN, Fukushima 1980; Lecun et al. 1998), a specialisation of an artificial neural network (Rosenblatt 1957; LeCun, Bengio, & Hinton 2015, for a more recent review) designed for complex image data and computer vision applications. We trained the SVM using the scikit-learn (Pedregosa et al. 2011) PYTHON library. We constructed the CNN with three convolutional layers and two fully connected layers of 128 neurons each. We trained our network using the Keras package on an NVidia GTX1070 GPU.

The accuracy, defined as the number of correct classifications divided by the total number of classifications, is presented in Table 1. Due to the limited size of the test set, the true misclassification rate could be as high as 4.5% within a 95% confidence interval. Despite the small size of the input images, we found that the convolutional neural network achieved slightly better performance on our test set of 866 images, divided evenly between the two classes. Although understanding the decision criteria inside the network is non-trivial, the network appears to have a developed robust representation of the gaussianity of the genuine sources and the less continuous artefacts in the data. The CNN-based classifier was incorporated into the *Mary* pipeline. The user can choose to take advantage of this module by setting the *ml-choice* parameter and can define the threshold of the classification with the *clasthresh* parameter.

For example, during the 2017 February 2–7 DWF run (see Cooke et al. in preparation, for more details), the CNN-based classifier classified 24 153/89 491 (27%) of candidates shown as artefacts, leaving 73% to undergo further automatic selection with *Mary* and finally be inspected by project astronomers and volunteers. At this time, we have not performed a systematic inspection of all rejected candidates, but given the conservative settings employed and current analysis we expect false positives to be very low. Performance of the machine-learning step can be improved with a larger and more diverse training set, allowing better optimisation of the network training parameters. Future development could include an online learning component, where false positives are fed back into the algorithm in real time to allow ‘on the fly’ retraining.

PASA, 34, e037 (2017)
doi:10.1017/pasa.2017.33

3.5. Catalogues query and ranking of the candidates

When large numbers of candidates pass the selection criteria, the user may find it useful to see a priority value attached to each candidate. *Mary* assigns the priority values by combining two types of information: those retrieved by querying online or downloaded catalogues, and information about the number of times that a transient was detected. In fact, transient searches usually imply target fields to be observed several times with a defined cadence; thus knowing if, when, and how many times a transient was discovered in the past greatly helps to understand the nature of the detected sources. During DWF observing runs aimed at discovering extragalactic transients, we set the Guide Star Catalog 2.3 (GSC, Lasker et al. 2008) as a preferred catalogue for *Mary* to query, as it provides a basic star/non-star classifier for objects in the southern hemisphere down to relatively faint (>20) magnitudes. The search radius (expressed in arcsec) is defined as $\text{radius}_{\text{GSC}} = \sqrt{\text{seeing}^2 + u^2}$, where the seeing is defined as $\text{fwhm} \times \text{scale}$ and u is the angular resolution of the GSC catalogue, $u = 1$ arcsec.

Mary prioritises the detected sources by assigning a score from 1 to 5, where higher value indicates a higher priority for the source to be followed up. The ranking system is optimised for the typical ~6 consecutive observing nights of DWF, and it works as follows:

5. Source detected only in the current observing night and not reported in the GSC catalogue.
4. Source detected only in the current observing night and reported as possible non-star in the GSC catalogue.
3. Source detected in the current observing night and in the previous observing night, not reported in the GSC catalogue or reported as possible non-star.
2. Source detected in the current observing night and in two or more previous observing nights, not reported in the GSC catalogue or reported as possible non-star.
1. Source classified as ‘star’ in the GSC catalogue.

The query of the GSC catalogue is less effective than the query of catalogues generated on DECam images using the SExtractor star/galaxy classifier, or more sophisticated classifiers such the algorithm described in Miller et al. (2017). However, the query of online catalogues becomes extremely valuable when searching for transient sources with template images acquired only a few minutes or hours before the science images, for example, during the first nights of DWF observations or during the follow-up of GW events. The programme can be slightly modified to allow the query of different or additional catalogues in the future.

3.6. Products for visualisation

The list of candidate transients comes with a set of products to facilitate the visual inspection of the candidates. Along with the priority assigned to each candidate, the products

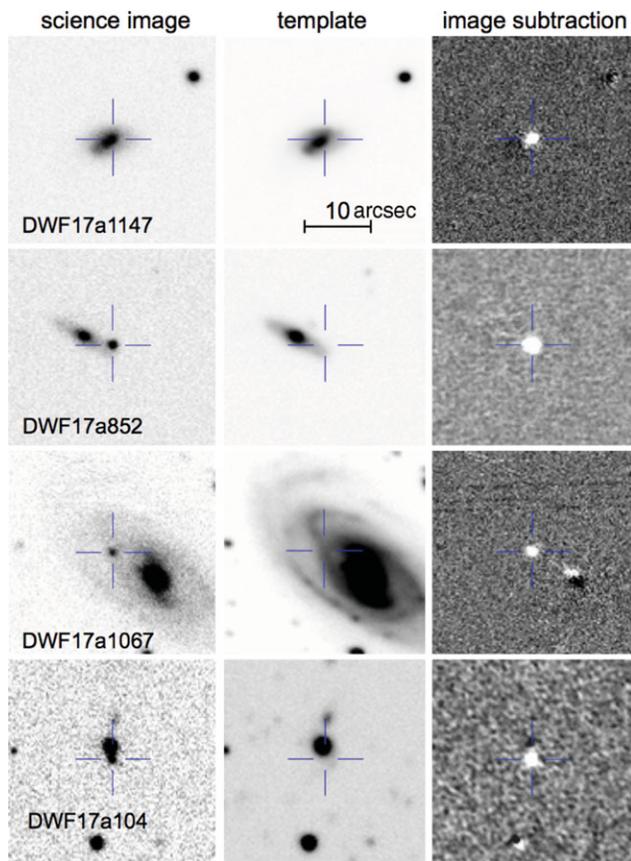


Figure 4. Examples of transient sources detected with the *Mary* pipeline in near real time. In particular, these images show the small ‘postage stamp’ images that allowed the identification of the possible supernovae DWF17a1147, DWF17a852, DWF17a1067, DWF17a104 presented in Andreoni et al. (2017b). The science images, where the transient is present, and the template images populate the first two columns (in inverse-colour grayscale). The image subtraction is presented in the last column for each candidate.

to help identify and classify interesting sources include the following:

- Small ‘postage stamp’ images with user-defined size cut from the template, the science, and the subtracted images, centred on the coordinates of the candidates; some examples are presented in Figure 4.
- Region identification files suitable for ‘ds9’ image visualisation tool with both circles and projection shapes. These files mark the locations of objects on the images, with different colours depending on the object class as identified in the catalogues such as the GSC.
- Light curves built by collecting magnitude values extracted from the subtracted images (with the same template) created during a series of *Mary* runs. These light curves are calibrated if the *photochoice* parameter is set, otherwise a zero point = 25 is assumed.
- Calibrated aperture and PSF photometry light curves. To optimise the fast completion of the transient identi-

cation, we consider this part of code detached from the *Mary* pipeline (but can be automatically started for each detected source), because the calibration can take minutes of computational time, depending on the number of images available—usually a few hundreds in the DWF programme context.

These *Mary* products become particularly useful when observations are taken with only a single (or no) filter, as it is not possible to classify the discoveries based on colour information, or during searches in near real time. For instance, the DWF programme uses images in the *g* filter for fast transient detection, complemented by colour information taken at a much slower cadence.

A custom programme consolidates *Mary*’s results into an SQL database on the Green II supercomputer. Once each *Mary* run is complete, all data pertaining to each identified candidate (including information about its associated image products) is stored to the database. This facilitates simple, powerful data analysis, including straightforward generation of light curves for each observed candidate. The database can be queried from the command line, programmatically or via a custom-built web interface, which allows the user to request and view data, candidate ‘postage stamps’, or compressed CCD images, and to add manual candidate classifications.

We coded additional custom programmes to cluster the detections into groups of interest or to better characterise them, for example, building well-calibrated light curves using all the images available, by ‘unfolding’ the images coadded to generate the science images during one or multiple *Mary* runs. These additional programmes and tools are not discussed in this paper, because their application strongly depends on the type of observations performed and lies outside the essential structure of *Mary*.

4 TESTS ON DECAM IMAGES

The *Mary* pipeline is designed to work both in near real time during the observations, as well as at later times. We developed *Mary* independently from other pipelines built to analyse DECam data, such as PhotPipe (Rest et al. 2005, 2014), DiffImg (Kessler et al. 2015), or pipelines designed for other transients projects (e.g., Förster et al. 2016). *Mary* is designed to detect transients in images already calibrated, thus we established a profitable collaboration to calibrate the DECam images in near real-time analysis with part of the PhotPipe reduction code. All the subsequent steps (from the creation of science and template images to the end of the pipeline) do not depend on PhotPipe or any other pipeline. For a more refined and complete late-time analysis, we take advantage of images calibrated with the NOAO High-Performance Pipeline System (Valdes & Swaters 2007; Swaters & Valdes 2007). The main differences between the near real-time and late-time transient detection with *Mary* lies in the quality of the pre-processing of the images available and the percentage

of CCDs successfully accomplishing the calibration without further intervention.

Fast and efficient analysis of the images was the goal at the outset of the *Mary* design, which becomes particularly challenging when dealing with the DECam data products. These consist of a collection of 60 functioning $2k \times 4k$ pixel CCDs (32 MB each), one of which (S7) we ignore due to calibration problems. For DWF, we analyse each CCD in parallel, using one CPU core of the Green II supercomputer per CCD. The individual processes undergo no further parallelisation. We coded a script that manages the submission of the jobs to the Portable Batch System queue of the Swinburne supercomputer: dedicated nodes make the time between the submission and the running of each job negligible (<1 s) during the DWF observations. The user can switch between the two main architectures by setting the *pipsetup* parameter to ‘RT’ to process in near real time the images calibrated with PhotPipe, or to ‘NOAO’ to analyse multi-extension images processed with the Community Pipeline. The *path_original* parameter allows a choice of path to find the NOAO processed images, while the products are always organised with the same structure to facilitate their retrieval.

The time needed to complete the data processing depends on the number of images to be coadded, the generation and usage of weight maps, the detection threshold, and the set of parameters chosen for the desired configuration of the pipeline. We tested the pipeline to present an example of timing and completeness estimates in the specific configuration described below.

4.1. Timing test

In this section, we tested the timing performance of the *Mary* pipeline, without accounting for the data transfer time. The data compression and transfer from CTIO (in Chile) to the Swinburne supercomputer (in Australia) affects the near real-time observations in a way extensively discussed in two separate papers (Cooke et al. in preparation; Vohl et al., 2017) specifically dedicated to the DWF programme. In particular, Vohl et al. present the custom compression code based on JPEG2000 standard for the compression of DECam data that significantly reduces the time needed to transfer the images from CTIO to the Swinburne supercomputer in Australia. Improved data transfer speed allows a successful near real-time analysis during fast-cadence time-domain surveys such as DWF, without the need of a supercomputing facility at the telescope.

We ran our timing test on images taken on 2015 December 21 UT, during a DWF observing campaign, and calibrated with the NOAO pipeline. The conditions were stable, with seeing around ~ 1.5 arcsec, and we acquired 84 continuous exposures of 20 s in *g* band. We chose to use self-generated weight maps for the coaddition, and downloaded catalogues to calibrate the zero point and to match the sources with the GSC catalogue. We set the following

PASA, 34, e037 (2017)
doi:10.1017/pasa.2017.33

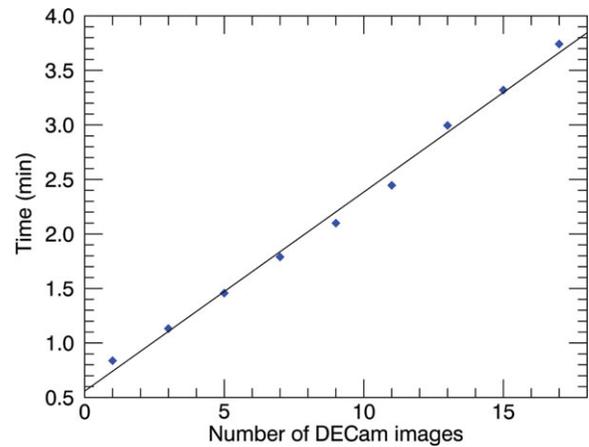


Figure 5. Median time needed to complete all the steps of the pipeline in parallel for 59 functioning CCDs of DECam. We timed the pipeline using different numbers of images to be coadded to form the science image, always using the same 11-month-old template image.

parameters for the source extraction: DETECT_MINAREA = 8, DETECT_THRESHOLD = 1.8 (for ‘positive’ images, which corresponds to an effective $S/N \sim 7\sigma$ for point sources), DETECT_THRESHOLD = 1.3 (for ‘negative’ images). The results of the timing test are shown in Figure 5, which shows the average time taken to complete all the operations, from the coaddition of science images (whose number constitutes the horizontal axis) to the building of the products for the visualisation. In order to get a larger sample, each point represents the median of the time values obtained during three runs of the pipeline, keeping exactly the same configuration. A line with slope $m = 0.18$ and intercept $c = 0.56$ fits the data quite well, as the χ -squared of the fit returns a p -value = 1.00 for eight degrees of freedom. This means that the probability of the scatter of the points from the line lies within the statistically expected scatter. It is important to note that the crosstalk rejection needed for our data set strongly affected the timing results, adding about 2–3 s per individual image. This could be avoided by either (a) performing dithered pointings during the observations; (b) further improving the crosstalk correction during the calibration; (c) running the pipeline on images taken with other cameras not affected by CCD operating system crosstalk effects.

4.2. Completeness and efficiency tests

We tested the performance of the *Mary* pipeline in different conditions of seeing and depth, using 59 CPU cores and 1 GPU of the Swinburne Green II supercomputer. The data we used were taken between 2015 December 19–22 during high-cadenced DWF observations. Seeing conditions, limiting magnitudes, and the results of the tests are summarised in Table 2. We estimated the significance above the background by performing PSF photometry of the sources in the science image, then evaluating an approximate signal-to-noise ratio

Table 2. Date (UTC, YYMMDD), seeing (arcsec), number of coadded images, and magnitude limit (g -band, 5σ) of the science images used during the completeness (μ) and efficiency (ε) tests. We compute μ as the average number of mock sources with $S/N > 10\sigma$ recovered by the pipeline. We base the estimate of ε as the ratio between the number of candidates that pass the visual inspection and the total number of candidates.

Date	Seeing	N	σ_{det}	mag_{det}	$\text{mag}_{5\sigma}$	$\text{mag}_{10\sigma}$	$\mu_{>10\sigma}$	ε
15-12-19	1.7 arcsec	9	7.0	22.6	23.0	22.0	95.7%	98.2%
15-12-20	1.3 arcsec	9	7.5	23.2	23.7	22.7	97.0%	99.0%
15-12-21 ^a	1.5 arcsec	3	7.6	22.5	23.1	22.2	96.7%	97.1%
15-12-21		9	6.7	22.4	23.4	22.4	97.0%	96.2%
15-12-21 ^a		35	6.3	23.5	23.8	22.8	96.0%	98.3%
15-12-22	1.1 arcsec	9	9.5	22.3	23.0	22.2	93.5%	97.9%

^aUsed to compare the performance on shallow and deep images.

(S/N) using the formula⁵:

$$S/N = (10^{-\frac{\text{emag}}{2.5}} - 1)^{-1},$$

where ‘emag’ is the magnitude error of the measurement. The PSF extraction and photometry were obtained using the PythonPhot package (Jones, Scolnic, & Rodney 2015). To perform completeness tests, we injected a flat distribution (in magnitude) of 100 synthetic mock point sources per CCD for the 59 best functioning CCDs. The shape of the PSF of the injected sources is defined by a double-Gaussian function whose parameters are measured separately on each CCD, thus the shape of the PSF varies across the field of view. We injected the same amount of flux at the same location for all the synthetic sources in different cases to compare the behaviour of the pipeline in different situations. Each individual result was obtained running the *Mary* pipeline twice on the same field on two sub-sets of images taken on the same night. The injection of fake sources at random locations leads to a probability of blending with other bright or saturating sources that increases with the depth of the images, making the results of the tests conservative. Nevertheless, the tests do not account for possible failure of the processing of CCDs during the analysis, which mainly depends on the goodness of the calibration of the images.

We used ‘real’ (i.e., not-synthetic) sources to estimate the detection efficiency of the pipeline. Three people independently assessed the transient candidates that *Mary* output on the images in Table 2 using a template image taken on 2015 January 17, about 11 months before the science images were acquired. The detection efficiency ε is the ratio of the average number of the ‘good’ candidates that passed the visual inspection, over the total number of candidates. We summarise the results of our completeness and efficiency tests in Table 2. In addition, we plot the completeness test results in Figures 6 and 7, where we binned the fraction of detected sources in bins of 0.3 mag.

We compared the response of *Mary* in different seeing conditions by running the pipeline on a set of nine images during each run of the pipeline (Figure 6), keeping the same parameters adopted for the timing test (Section 4.1) which set the

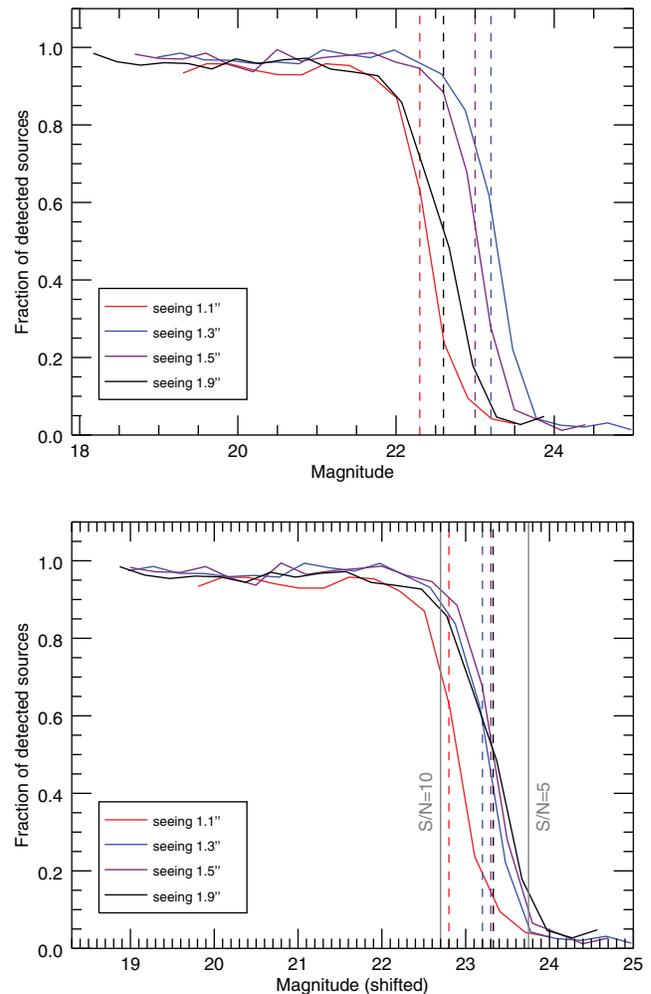


Figure 6. Completeness test runs at different observing conditions. The vertical dashed lines intercept the solid-lined curves of the same colour where the detection fraction is equal to 0.5, thus where half of the injected sources are recovered by the pipeline. *Bottom panel:* The curves are rigidly shifted to make their 10σ limits match the 10σ limit of the blue curve (seeing = 1.3 arcsec) to facilitate the comparison of the shape of the curves. The $S/N = 5$ vertical line indicates the average magnitude of the 5σ limits of the curves, after being shifted.

⁵ see, for example, <http://www.eso.org/~ohainaut/ccd/sn.html>

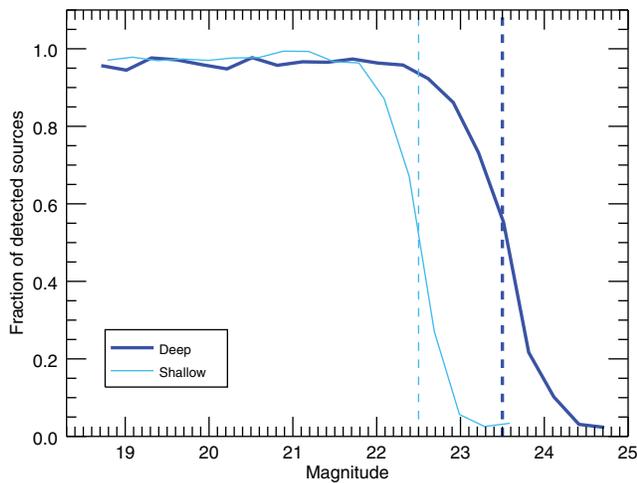


Figure 7. Completeness test performed on stacks of 3 (shallow) and 35 (deep) images taken on 2015 December 21. The vertical dashed lines intercept the solid-lined curve of the same colour where the detection fraction is equal to 0.5, thus where half of the injected sources are recovered by the pipeline.

detection limit at $\sigma_{\text{det}} = 7.1_{-0.4}^{+0.5}$. We shifted all the curves to make their 10σ mag limit match the 10σ limit of the deepest image (data taken on 2015 December 20, seeing = 1.3 arcsec), obtaining the bottom plot of Figure 6. The shape of the decaying curve is consistent for three different seeing conditions, which demonstrates the adaptability of the pipeline to different observing situations even without the user acting directly on the control parameters, placing the detection limit at $\sigma_{\text{det}} = 7.1_{-0.4}^{+0.5}$. The outlier red curve (data taken on 2015 December 22, seeing = 1.1 arcsec) shows a steeper decay and lower, more scattered values in the part of the curve preceding the decay. In particular, the mean fraction of detected sources with $S/N > 10$ is $\mu_{>10\sigma} = 93.5\%$, against a mean value of 96.6% for the other cases. This can be the result of thin clouds present during the observations, and a consequence of the seeing of the science image being better than the seeing of the template image (1.1 arcsec vs. 1.3 arcsec). This type of problem could be overcome with the implementation of the optimal photometry algorithm (Zackay et al. 2016) for the image subtraction.

We tested the response of the *Mary* pipeline to a variable number of images to be coadded and searched for transients. Figure 7 highlights the large gap of 1.0 magnitude in the detection limit between the two completeness curves obtained by coadding 3 and 35 images, while maintaining the same parameters for the pipeline.

5 DISCUSSION

The *Mary* pipeline was designed to identify optical transient and variable objects in seconds to minutes during the multi-facility observations part of the DWF programme. In this framework, the source identification priority can quickly change, for example, in case a LIGO trigger is received and

a follow-up is performed. Also, observing and data transfer conditions may dramatically change in little time. During DWF observations, a team of astronomers is always ready to modify the parameters of the pipeline according to the priorities of the moment, even if little or no action is required in standard observing conditions. Arbitrary changes may follow decisions to modify the balance between processing speed, completeness down to faint magnitudes, or efficiency in allowing the visual inspection team to confidently identify bright (mag ~ 21) sources of particular interest.

For example, a functionality such as the creation of weight maps for the image coaddition can be switched off to save data processing time, especially if images are acquired in a continuous and uniform sequence, improving the performance presented in Section 4.1. If a multi-object spectrograph such as the AAT/AAOmega can provide spectroscopy of hundreds of candidates less than 2 h after their detection with DECam, the astronomers can modify the selection criteria to allow a larger completeness. In other cases, multi-object spectrographs and 8-m-class telescopes may not be available to follow-up DECam candidates due to bad weather, while only smaller telescopes at different observing sites would. In this case, the detection threshold should be moved up in order to facilitate the visual inspection team to focus only on the brightest sources. In general, the users of the pipeline can decide to be more or less conservative in the completeness and candidate selection requirements mainly by modifying the parameters that regulate the source extraction (DETECT_THRESH_POS/NEG and DETECT_MINAREA_POS/NEG), the shape selection (especially *elmax*), and the bogus rejection parameters (*isoreafmin* and *side*). During late-time data analysis on archival images (*pipsetup*=NOAO) better completeness and lower efficiency are usually preferred, in order to carefully explore the whole data set, reducing the risk to miss faint or interesting transients.

In Appendix A, we present the default values used for the tests whose results are shown in this paper. These are the values commonly used during DWF observations in standard conditions and are rarely to be modified before or during an observing night. The most sensitive parameters, depending on the FWHM measured for each CCD are automatically computed by the pipeline.

5.1. Comparison with other pipelines

The results of the tests of the *Mary* pipeline with DECam data to facilitate the comparison with other existing top-level pipelines can be briefly summarised as follows:

- Timing: $\bar{t} \sim 1.1$ min, corresponding to the coaddition of three images, optimal for near real-time analysis of the continuous 20 s exposure images and 20 s readout time acquired during the DWF observations.
- Completeness: We find an average completeness rate for sources with $S/N > 10$ and coadding nine images, where the seeing is comparable to or worse than the template,

of $\bar{\mu}_{>10\sigma} \sim 96.6$ that corresponds to a missed fraction of 3.4%.

- Detection efficiency: the average value $\bar{\epsilon} = 97.8\%$, obtained with the coaddition of nine images at different seeing conditions, corresponds to a false positive rate of 2.2%.

The \bar{t} we compute for *Mary* is lower than, or comparable with, the timing to run the post-calibration steps of the iPTF pipeline reported by Cao et al. (2016), specifically 154.3 s (or 2.57 min) per image. Our \bar{t} is slightly larger than the 60 s required by the HiTS pipeline to cover the steps between the astrometric calibration and the visual inspection (Förster et al. 2016). The *Mary* pipeline closely approaches the best performing, fully machine-learning-based pipelines in terms of completeness and real/bogus classification ability. In fact, the tests we performed returned a false positive rate of 2.2% and missed fraction of 3.4% against, for example, the 1 and 5% of iPTF (Cao et al. 2016). These results demonstrate that a pipeline with the design presented in this work is suitable to perform high-quality searches for transients during observing campaigns with DECam such as the DWF programme. The *Mary* pipeline should produce similar performance for other telescopes and CCDs.

5.2. Application to other programmes and facilities

The *Mary* pipeline was originally designed to work on DECam images and was optimised for the success of DWF observing campaigns. However, the pipeline can be adapted to search for transients in data acquired with different facilities. In this paper, we presented *Mary* crafted specifically for DWF DECam data, instead of presenting it as a more general pipeline, in part because (1) small or major necessary changes cannot be predicted for each individual telescope; (2) the pipeline is optimised for DECam images, and includes steps (for example, the crosstalk rejection) that may not apply to other instruments; (3) the pipeline was designed to run on the Swinburne high-performance supercomputer, with large availability of computational resources and parallel processing.

Nevertheless, the *Mary* pipeline was used in contexts other than DWF, in order to discover slowly evolving transients. In particular, *Mary* was used during the follow-up of Fast Radio Bursts with DECam (Bhandari et al., in preparation; Petroff et al. 2017). We also adapted the *Mary* pipeline to analyse pre-processed images acquired with the Magellan Fourstar infrared camera (Petroff et al. 2017), discovering tens of variable objects between two consecutive observing nights. Finally, we successfully used *Mary* to search for transients in data collected with the 1-m Zadko telescope, which is among the facilities that make up the DWF follow-up network. This type of analysis was not performed in near real time and dedicated nodes of the Swinburne supercomputer were not required. The *Mary* pipeline is currently available upon request, but will be uploaded on GitHub in the future.

PASA, 34, e037 (2017)
doi:10.1017/pasa.2017.33

6 CONCLUSION

In this paper, we presented the *Mary* pipeline, designed to discover optical transients in DECam CCD images within seconds to minutes for the DWF programme. The pipeline relies on popular astronomical packages to perform key operations and does not involve machine-learning techniques. The performance of such a pipeline (with a false positive rate of $\sim 2.2\%$ and a missed fraction of $\sim 3.4\%$) is similar to the results achieved by the most popular pipelines used by large programmes continuously surveying the sky. The *Mary* pipeline matches the needs of programmes, like DWF, that aim to search for transients over large areas of sky in near real time, with a flexible strategy and variable conditions.

Such a pipeline is crucial to trigger rapid response spectroscopic and imaging observations to study and understand the fast transient Universe. In addition, the *Mary* pipeline was successfully used to search for an optical afterglow of a Fast Radio Burst (Petroff et al. 2017) with CTIO/DECam and with the FourStar infrared camera on Magellan.

We are planning to make the code available to astronomers interested in analysing public DWF data. Finally, pipelines designed on the model of *Mary* as presented in this paper will allow more groups of astronomers to search for counterparts to gravitational wave events, Fast Radio Bursts, and new classes of fast transients using small to large facilities.

ACKNOWLEDGEMENTS

The authors acknowledge Uros Mestric (SUT) for contributing to the visual inspection of the candidates for this paper, and the group of volunteers who constituted the ‘visual inspection team’ during the DWF observations in near real time. IA acknowledges Armin Rest (STScI), Phil Cowperthwhite (Harvard), and Frank Masci (IPAC/Caltech) for sharing their knowledge and expertise.

This research was conducted by the Australian Research Council Centre of Excellence for All-sky Astrophysics (CAASTRO), through project number CE110001020. Research support to IA is provided by the Australian Astronomical Observatory. This work was performed on the gSTAR national facility at Swinburne University of Technology. gSTAR is funded by Swinburne and the Australian Governments Education Investment Fund. This research was funded, in part, by the Australian Research Council Future Fellowship grant FT130101219 and by the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav) through project number CE170100004.

References

- Abbott, B. P., et al. 2016, ApJ, 826, L13
- Andreoni, I., et al. 2017a, The Astronomer’s Telegram, No. 10078, 78
- Andreoni, I., et al. 2017b, The Astronomer’s Telegram, No. 10072, 72
- Astier, P., et al. 2006, A&A, 447, 31
- Becker, A. 2015, HOTPANTS: High Order Transform of PSF ANd Template Subtraction, Astrophysics Source Code Library, ascl:1504.004

- Bellm, E. 2014, in *The Third Hot-wiring the Transient Universe Workshop*, eds. P. R. Woźniak, M. J. Graham, A. A. Mahabal, & R. Seaman (held in Sante Fe, 13–15 November), 27
- Bertin, E., & Arnouts, S. 1996, *A&AS*, 117, 393
- Bertin, E., et al. 2002, in *ASP Conf. Ser.*, Vol. 281, *Astronomical Data Analysis Software and Systems XI*, eds. D. A. Bohlender, D. Durand, & T. H. Handley (San Francisco: ASP), 228
- Cao, Y., Nugent, P. E., & Kasliwal, M. M. 2016, *PASP*, 128, 114502
- Cenko, S. B., et al. 2015, *ApJL*, 803, L24
- Chu, Q., et al. 2016, *MNRAS*, 459, 121
- Cortes, C., & Vapnik, V. 1995, *Mach Learn*, 20, 273
- Coward, D. M., et al. 2016, *ASA*, 34, e005
- Cowperthwaite, P. S., et al. 2016, *ApJL*, 826, L29
- Djorgovski, S. G., et al. 2010, in *The First Year of MAXI: Monitoring Variable X-ray Sources*, eds. T. Mihara & N. Kawai (Tokyo: JAXA Special Publ.) 32
- Doctor, Z., et al. 2016, *ApJ*, 837, 57
- Flaugher, B., et al. 2015, *AJ*, 150, 150
- Förster, F., et al. 2016, *ApJ*, 832, 155
- Fukushima, K. 1980, *Biol. Cybernetics*, 36, 193
- Gal-Yam, A. 2014, in *American Astronomical Society Meeting Abstracts*, Vol. 223, *American Astronomical Society Meeting Abstracts #223*, 235.02
- Garnavich, P. M., et al. 2016, *ApJ*, 820, 23
- Ghirlanda, G., et al. 2015, *A&A*, 578, A71
- Ivezic, Z., et al. 2008, *SerAJ*, 176, 1
- Jones, D. O., Scolnic, D. M., & Rodney, S. A. 2015, *PythonPhot: Simple DAOPHOT-type photometry in Python*, *Astrophysics Source Code Library*, ascl:1501.010
- Jordan, M., & Mitchell, T. 2015, *Science*, 349, 255
- Keller, S. C., et al. 2007, *PASA*, 24, 1
- Kessler, R., et al. 2015, *AJ*, 150, 172
- Lasker, B. M., et al. 2008, *AJ*, 136, 735
- Law, N. M., et al. 2009, *PASP*, 121, 1395
- LeCun, Y., Bengio, Y., & Hinton, G. 2015, *Nature*, 521, 436
- Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, *Proc. IEEE*, 86, 2278
- Meade, B., et al. 2017, *ASA*, 34, e023
- Metzger, B. D. 2016, arXiv:1610.09381
- Miller, A. A., et al. 2017, *AJ*, 153, 73
- Pedregosa, F., et al. 2011, *J. Mach. Learn. Res.*, 12, 2825
- Perrett, K., et al. 2010, *AJ*, 140, 518
- Petroff, E., et al. 2017, *MNRAS*, 469, 4465
- Rest, A., et al. 2005, *ApJ*, 634, 1103
- Rest, A., et al. 2014, *ApJ*, 795, 44
- Rosenblatt, F. 1957, Technical Report No. 85-460-1 (Buffalo: Cornell Aeronautical Laboratory Inc.)
- Shortridge, K. 2012, in *ASP Conf. Ser.*, Vol. 461, *Astronomical Data Analysis Software and Systems XXI*, eds. P. Ballester, D. Egret, & N. P. F. Lorente (San Francisco: ASP), 49
- Smith, R. M., et al. 2014, in *Proc. SPIE*, Vol. 9147, *Ground-based and Airborne Instrumentation for Astronomy V*, eds. S. K. Ramsay, I. S. McLean, & H. Takami (Bellingham: SPIE) 914779
- Swaters, R. A., & Valdes, F. G. 2007, in *ASP Conf. Ser.*, Vol. 376, *Astronomical Data Analysis Software and Systems XVI*, eds. R. A. Shaw, F. Hill, & D. J. Bell (Tucson: ASP), 269
- Tanvir, N. R., et al. 2013, *Nature*, 500, 547
- Valdes, F. G., & Swaters, R. A. 2007, in *ASP Conf. Ser.*, Vol. 376, *Astronomical Data Analysis Software and Systems XVI*, eds. R. A. Shaw, F. Hill, & D. J. Bell (Tucson: ASP), 273
- Vestrand, W. T., et al. 2014, *Science*, 343, 38
- Vohl, D., Pritchard, T., Andreoni, I., Cooke, J., & Meade, B. 2017, *PASA*, 34, 38
- Waş, M., Sutton, P. J., Jones, G., & Leonor, I. 2012, *PhRvD*, 86, 022003
- Zackay, B., Ofek, E. O., & Gal-Yam, A. 2016, *ApJ*, 830, 27

APPENDIX A THE MARY PARAMETERS

Table A1. The majority of parameters that regulate the *Mary* pipeline must be set before starting the analysis of the images from a given observing night (S), or are automatically computed (A). In particular, most of the automatically computed parameters rely on the estimation of the average FWHM of each CCD described in Section 3.1. No manual, ‘on the fly’ intervention is required for the user, while *Mary* allows the user to tweak any of the parameters at any time during an observing campaign. In particular, the *sequencenumber* parameter, that can uniquely identify the current data processing, was manually (M) modified for each set of images during the past DWF observations in near real time. Default values are indicated for sensitive parameters, as used for the tests reported in this paper, which represent a standard setup during DWF observations.

Parameter	S/A/M	Description
PIPELINE SETUP		
<i>pipsetup</i>	S	Choose to use science images acquired in near real time (RT) or processed with the NOAO pipeline (NOAO).
<i>path_original</i>	S	Path where the images processed with the NOAO pipeline are stored.
<i>wmapchoice</i>	S	Choose to create and use weight maps (YES/NO, default = YES).
<i>checkwcs</i>	S	Choose to proceed only in case of successful WCS calibration with PhotPipe (YES/NO, used only with <i>pipsetup</i> = RT, default = NO).
<i>extnum</i>	S	Number of extensions of the FITS files (default = 60).
<i>walltime</i>	S	Maximum ‘wall’ time for the processing on the Swinburne Green II supercomputer.
<i>list_images_sci</i>	S	Path to the list of science images to process.
<i>list_images_temp</i>	S	Path to the list of template images to process (only if <i>useoldtemplate</i> = <i>usemarytemplate</i> = NO).
<i>sequencenumber</i>	M/A	Code associated to the <i>Mary</i> run.
<i>field</i>	S	Name of the observed field.
<i>date</i>	S	Date of the observations.
<i>filter</i>	S	Filter used during the observations.
TEMPLATE IMAGES		
<i>datetemp</i>	S	Date of the template image(s) (only if <i>useoldtemplate</i> = <i>usemarytemplate</i> = NO).
<i>useoldtemplate</i>	S	Choose to use a full-field, well-processed image as template (YES/NO).
<i>oldtemp</i>	S	Specify the path to the template image (only if <i>useoldtemplate</i> = YES).
<i>usemarytemplate</i>	S	Choose to use the same template used for another <i>Mary</i> run (YES/NO).
<i>otherdate</i>	S	Date of the <i>Mary</i> run to copy the template from (default = <i>datetemp</i>).
<i>othersequencenumber</i>	S	Code of the <i>Mary</i> run to copy the template from.
<i>otherfield</i>	S/A	Target field of the <i>Mary</i> run to copy the template from (default = <i>field</i>).
<i>otherfilter</i>	S/A	Filter used in the <i>Mary</i> run to copy the template from (default = <i>filter</i>).
SETRACTOR AND HOTPANTS		
<i>DETECT_MINAREA_POS</i>	S	SExtractor DETECT_MINAREA parameter for the subtracted image (default = 8).
<i>THRESH_TYPE_POS</i>	S	SExtractor THRESH_TYPE parameter for the subtracted image (default = RELATIVE).
<i>DETECT_THRESH_POS</i>	S	SExtractor DETECT_THRESH parameter for the subtracted image (default = 1.8).
<i>ANALYSIS_THRESH_POS</i>	S	SExtractor ANALYSIS_THRESH parameter for the subtracted image (default = 1.8).
<i>DETECT_MINAREA_NEG</i>	S	SExtractor DETECT_MINAREA parameter for the inverse of the subtracted image (default = 8).
<i>THRESH_TYPE_NEG</i>	S	SExtractor THRESH_TYPE parameter for the inverse of the subtracted image (default = RELATIVE).
<i>DETECT_THRESH_NEG</i>	S	SExtractor DETECT_THRESH parameter for the inverse of the subtracted image (default = 1.3).
<i>ANALYSIS_THRESH_NEG</i>	S	SExtractor ANALYSIS_THRESH parameter for the inverse of the subtracted image (default = 1.3).
<i>ilsigma</i>	S	Number of standard deviations from the background value at which the good pixel lower limit is set in HOTPANTS (default = 5).
<i>tlsigma</i>	S	Number of standard deviations from the background value at which the good pixel upper limit is set in HOTPANTS (default = 5).
<i>ker</i>	A	Convolution kernel half width (default = $2.5 \times fwhmsub$).

Table A1. Continued

Parameter	S/A/M	Description
IMAGE		
<i>xdim</i>	S/A	Number of pixels in the <i>X</i> axis of the aligned images (default = 3 800).
<i>ydim</i>	S/A	Number of pixels in the <i>Y</i> axis of the aligned images (default = 1 800).
<i>fixoldchoice</i>	S	Choose to manually define the size of the aligned images (with <i>xmin</i> , <i>ymin</i>) even when <i>useoldtemplate</i> = YES.
SELECTION OF THE CANDIDATES		
<i>elmax</i>	S	Maximum elongation allowed, computed with SExtractor (ELONGATION output key, default = 1.8).
<i>fwhmtemp</i>	A	FWHM of the template image.
<i>fwhmsci</i>	A	FWHM of the science image.
<i>fwhmsub</i>	A	FWHM of the subtracted image.
<i>scale</i>	S	Pixel scale (arcsec/pix, default = 0.263).
<i>seeing</i>	A	Seeing of the science image (in arcsec, $fwhmsci \times scale$).
<i>isoareafmin</i>	S	Minimum value allowed for the filtered isophotal area of the sources, computed with SExtractor (ISOAREAF_IMAGE, default = 20).
<i>side</i>	A	Side of the box for the rejection of sources with a counterpart in the inverse of the subtracted image (default = $2.0 \times fwhmsub$).
MACHINE LEARNING CLASSIFIER		
<i>mlchoice</i>	S	Choose whether to use the machine learning classifier to remove CCD artefacts (YES/NO).
<i>classthresh</i>	S	Classification threshold for the machine learning output (between 0 and 1, default = 0.5).
CROSTALK AND CATALOGUES		
<i>crosschoice</i>	S	Choose whether to keep or eliminate those candidates flagged as possible crosstalk (ELIMINATE/KEEP).
<i>satlevelchoice</i>	S	Choose whether to use a fixed value (in ADU) above which sources generate crosstalk effects, instead of using the value of the SATURATE header keyword (YES/NO).
<i>fixsatlevel</i>	S	Fixed value (in ADU) above which sources are considered to generate crosstalk effects (used only if <i>satlevelchoice</i> = YES, default = 10 000).
<i>saturationchoice</i>	S	Choose whether to keep or eliminate those candidates flagged as saturating (ELIMINATE/KEEP).
<i>maxsources</i>	S	Maximum number of sources allowed per CCD; an error message is generated if the number of sources is greater than <i>maxsources</i> .
<i>querychoice</i>	S	Choose whether to query the catalogues online or their downloaded version (O/D).
<i>pathdownloadedGSC</i>	S/A	Path to the downloaded GSC catalogue (only if <i>querychoice</i> = D).
<i>radiusGSC</i>	A	Search radius for the GSC catalogue, automatically defined as $\sqrt{seeing^2 + u^2}$, where <i>u</i> is the angular resolution of the GSC catalogue (<i>u</i> = 1 arcsec).
RANKING		
<i>nightsarray</i>	S	Array of dates to be accounted for the ranking of the candidates.
<i>matchradius</i>	S	Radius (in arcsec) to match detections at different epochs.
PHOTOMETRY		
<i>photochoice</i>	S	Choose to compute the zero point of the images.
<i>catalogcalib</i>	S	Catalog to use to calibrate the zero point of the images (GAIA/USNO-B1).
<i>path_star</i>	S	Path to the stars pre-selected to perform the zero-point calibration.
<i>zpt</i>	S	Nominal zeropoint to be used if <i>photochoice</i> = NO (default = 25).
VISUALISATION		
<i>radcir_i</i>	S	radius (pixels) of the circular regions centred on the candidates (default = 12 pixels).
<i>radcir_w</i>	S	radius (arcsec) of the circular regions centred on the candidates (default = 8 arcsec).
<i>sidestamp</i>	S	Measure of the side of the postage stamp centred on the candidate (in pixel, default = 120).