A&A 652, A76 (2021)
https://doi.org/10.1051/0004-6361/202140735

Astronomy & Astrophysics

Special issue

Gaia Early Data Release 3

Gaia Early Data Release 3

Gaia photometric science alerts⋆

S. T. Hodgkin[1], D. L. Harrison[1,2], E. Breedt[1], T. Wevers[1,3], G. Rixon[1], A. Delgado[1,4], A. Yoldas[1], Z. Kostrzewa-Rutkowska[5,6], Ł. Wyrzykowski[7], M. van Leeuwen[1], N. Blagorodnova[8], H. Campbell[9], D. Eappachen[6,8], M. Fraser[10], N. Ihanec[7,11], S. E. Koposov[12,1], K. Kruszyńska[7], G. Marton[13], K. A. Rybicki[7], A. G. A. Brown[5], P. W. Burgess[1], G. Busso[1], S. Cowell[1], F. De Angeli[1], C. Diener[1], D. W. Evans[1], G. Gilmore[1], G. Holland[1], P. G. Jonker[8,6], F. van Leeuwen[1], F. Mignard[14], P. J. Osborne[1], J. Portell[15], T. Prusti[16], P. J. Richards[17], M. Riello[1], G. M. Seabroke[18], N. A. Walton[1], P. Ábrahám[13,19], G. Altavilla[20,21], S. G. Baker[18], U. Bastian[22], P. O'Brien[23], J. de Bruijne[16], T. Butterley[24], J. M. Carrasco[15], J. Castañeda[25], J. S. Clark[26], G. Clementini[27], C. M. Copperwheat[28], M. Cropper[18], G. Damljanovic[29], M. Davidson[12], C. J. Davis[30], M. Dennefeld[31], V. S. Dhillon[32,33], C. Dolding[18], M. Dominik[34], P. Esquej[4], L. Eyer[35], C. Fabricius[15], M. Fridman[36,37], D. Froebrich[38], N. Garralda[15], A. Gomboc[36], J. J. González-Vidal[25], R. Guerra[39], N. C. Hambly[12], L. K. Hardy[32], B. Holl[35], A. Hourihane[1], J. Japelj[40], D. A. Kann[41], C. Kiss[13], C. Knigge[42], U. Kolb[26], S. Komossa[43], Á. Kóspál[13,44,19], G. Kovács[45], M. Kun[13], G. Leto[46], F. Lewis[28,47], S. P. Littlefair[32], A. A. Mahabal[48,49], C. G. Mundell[50], Z. Nagy[13], D. Padeletti[51], L. Palaversa[35,52], A. Pigulski[53], M. L. Pretorius[54], W. van Reeven[39], V. A. R. M. Ribeiro[55,56,8], M. Roelens[35], N. Rowell[12], N. Schartel[39], A. Scholz[34], A. Schwope[57], B. M. Sipőcz[45], S. J. Smartt[58], M. D. Smith[38], I. Serraller[15], D. Steeghs[59], M. Sullivan[42], L. Szabados[13], E. Szegedi-Elek[13], P. Tisserand[31], L. Tomasella[60], S. van Velzen[5], P. A Whitelock[54,61], R. W. Wilson[24], and D. R. Young[58]

(Affiliations can be found after the references)

Received 5 March 2021 / Accepted 20 May 2021

ABSTRACT

Context. Since July 2014, the *Gaia* mission has been engaged in a high-spatial-resolution, time-resolved, precise, accurate astrometric, and photometric survey of the entire sky.
Aims. We present the *Gaia* Science Alerts project, which has been in operation since 1 June 2016. We describe the system which has been developed to enable the discovery and publication of transient photometric events as seen by *Gaia*.
Methods. We outline the data handling, timings, and performances, and we describe the transient detection algorithms and filtering procedures needed to manage the high false alarm rate. We identify two classes of events: (1) sources which are new to *Gaia* and (2) *Gaia* sources which have undergone a significant brightening or fading. Validation of the *Gaia* transit astrometry and photometry was performed, followed by testing of the source environment to minimise contamination from Solar System objects, bright stars, and fainter near-neighbours.
Results. We show that the *Gaia* Science Alerts project suffers from very low contamination, that is there are very few false-positives. We find that the external completeness for supernovae, $C_E = 0.46$, is dominated by the *Gaia* scanning law and the requirement of detections from both fields-of-view. Where we have two or more scans the internal completeness is $C_I = 0.79$ at 3 arcsec or larger from the centres of galaxies, but it drops closer in, especially within 1 arcsec.
Conclusions. The per-transit photometry for *Gaia* transients is precise to 1% at $G = 13$, and 3% at $G = 19$. The per-transit astrometry is accurate to 55 mas when compared to *Gaia* DR2. The *Gaia* Science Alerts project is one of the most homogeneous and productive transient surveys in operation, and it is the only survey which covers the whole sky at high spatial resolution (subarcsecond), including the Galactic plane and bulge.

Key words. surveys – supernovae: general – quasars: general – stars: variables: general

1. Introduction

On 19 December 2013, the European Space Agency (ESA) launched its *Gaia* satellite, which was the start of an ambitious project to measure the parallaxes of a billion stars in the Milky Way. *Gaia* started scientific operations in July 2014 and completed the 5-yr nominal mission on 16 July 2019, but the spacecraft is in good health and the data collection and processing is still ongoing as an extended mission phase. Although the final data release of the nominal mission is still to come (DR4, the extended mission will be released as DR5), the survey has already had a transformational impact on a broad range of fields, including white dwarfs (Gentile Fusillo et al. 2019), hypervelocity stars (Boubert et al. 2018), cosmological

⋆ Classification tables are only available at the CDS via anonymous ftp to cdsarc.u-strasbg.fr (130.79.128.5) or via http://cdsarc.u-strasbg.fr/viz-bin/cat/J/A+A/652/A76.

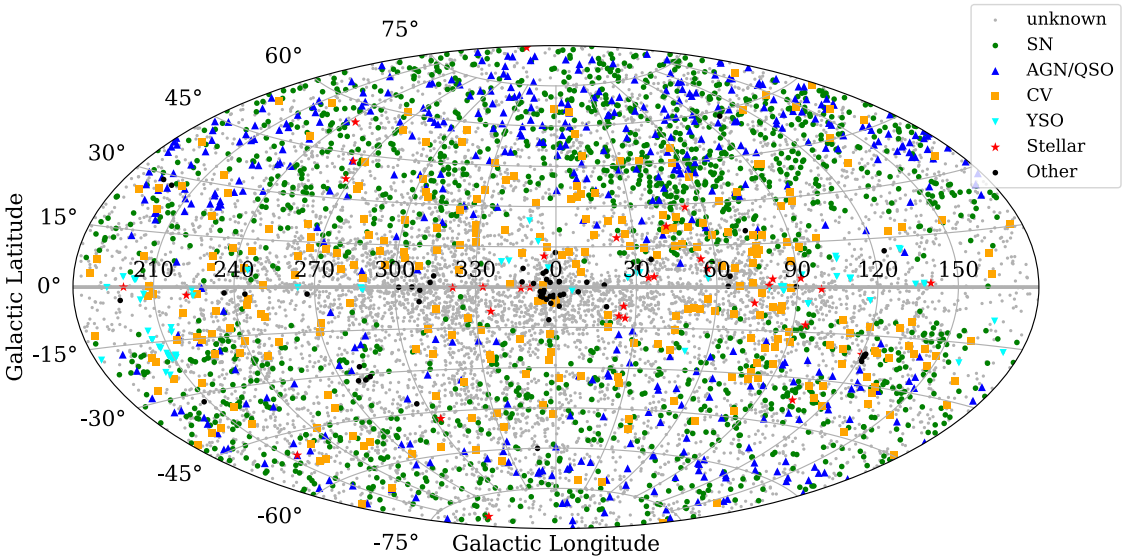

Fig. 1. Alerts detected by *Gaia* up to the end of 2019, plotted in Galactic coordinates. Alerts with unknown identifications are shown in grey, and spectroscopically confirmed alerts are highlighted in colour. The category 'other' includes microlensing events, galactic novae and X-ray binaries.

gravitational lensing (Lemon et al. 2019), and the merger history of the Galaxy (Belokurov et al. 2018).

In order to make the astrometric measurements, *Gaia* scans the full sky repeatedly. The exact number of observations and observing cadence of a given source depends on its location on sky, but each source will be observed ∼140 times over the lifetime of the survey (see e.g. Boubert et al. 2020). Typically a pair of observations, separated by 106.5 min, is followed by another pair of observations 2–4 weeks later. Each observation consists of a 50-s long white-light (*G*-band) lightcurve, sampled every 5 s, which can also be used for variability detection on very short timescales (Wevers et al. 2018; Roelens et al. 2017, 2018). Hence *Gaia* samples the sky on a range of timescales, allowing us to search these time series observations for transient variables. The detected transients are published as a public alerts stream, known as *Gaia* science alerts (GSA). Throughout the lifetime of the survey so far, GSA has undergone several changes, in particular as more data became available, making it possible to introduce more reliable and efficient detection algorithms. This paper focuses on the current operational state, but it includes the details of important changes throughout the development.

GSA has been designed to produce notifications for transient phenomena, that is to say any event which would benefit from a timely reaction, and thus to avoid a potential science loss. GSA is an added-value science product to the main astrometric goals of the satellite mission; the survey is not optimised for transient detection or completeness of transient populations. Nevertheless, *Gaia* has numerous advantages compared to ground-based transient surveys. Its space-based location means that biases due to weather or variable seeing are eliminated. It also benefits from a high dynamic range, high spatial resolution (∼0.1″), high photometric precision (1% at $G = 13$; 3% at $G = 19$), and all-sky coverage, including the Galactic plane which most ground-based surveys avoid because of crowding. Each observation also includes a low resolution ($R \sim 100$) 'blue photometer/red photometer' (BP/RP) slitless spectrum, which provides colour information at every epoch. A comparison between GSA and a number of other existing and planned transient surveys is presented in Table 1.

Up to 31 December 2019, 10765 alerts have been published, covering the full sky (Fig. 1). The alert detection is ongoing and currently alerts are published at a rate of 12 per day (see Sect. 3). We note that pulsating stars, regular variables, and eclipsing binary stars are excluded from this alert stream, as far as possible, as such variables are processed and published separately by the *Gaia* collaboration (e.g. Gaia Collaboration 2019).

Approximately 25% of the *Gaia* alerts have been classified (Sect. 3), including previously known objects. The majority of classifications are from ground-based spectroscopic observations, with a small sample classified mainly through photometry (e.g. microlensing event classification includes a model fit to the lightcurves). The alert stream is currently available in its entirety to the public[1], so that alerts can be followed up on by anyone interested. Currently, the majority of spectroscopically identified alerts are supernovae due to large-scale supernova follow-up by, for example, PESSTO (Public ESO Spectroscopic Survey of Transient Objects; Smartt et al. 2015), NUTS (Nordic Optical Telescope Unbiased Transient Survey), and the Zwicky Transient

Table 1. Comparison between *Gaia* and other existing or planned transient surveys (Bellm 2016).

Survey	$\Omega_{\rm fov}$ (deg^2)	Platescale (arcsec)	$m_{\rm lim}$	$\dot{\Omega}$ (deg^2 h^{-1})
Gaia (2 FOVs)	0.9	0.06×0.18	20.7	81.4
ASAS-SN	73.0	7.8	17	1294
ATLAS	60.0	1.9	20.0	5684
CRTS-2	19.0	1.5	19.5	1628
PS1	7.0	0.3	21.8	630
ZTF	47	1.0	20.4	3760
LSST	9.6	0.2	24.7	842

Notes. For each survey we list the instantaneous field-of-view ($\Omega_{\rm fov}$; *Gaia* has two fields-of-view), the size of the pixels, the limiting magnitude, and the areal survey rate ($\dot{\Omega}$). We note that *Gaia* and ASAS-SN are the only surveys which cover the whole sky. We further note that all the other transient surveys employ difference-imaging techniques to identify transients, while GSA is a purely catalogue driven survey (see discussion in Sect. 2.7).

[1] http://gsaweb.ast.cam.ac.uk/alerts

Facility (ZTF; Bellm et al. 2019). Among the large number of Type-I and Type-II supernovae observed so far, *Gaia* also discovered a number of unusual supernovae, such as the extremely UV-bright super-luminous supernova (SLSN) Gaia16apd (Kangas et al. 2016; Nicholl et al. 2017), Gaia17biu, which is a hydrogen-poor SLSN and by a factor of almost 3, the nearest SLSN known to date (Xiang et al. 2017; Dong et al. 2017; Bose et al. 2018), and Gaia16bvd, the first example of a pair-instability supernova (Gomez et al. 2019). GSA is currently the second-largest contributor of transients to the IAU Transient Name Server[2].

Other highlights so far include the discovery of the first fully-eclipsing AM CVn binary Gaia14aae (Rixon et al. 2014; Campbell et al. 2015), the fifth alert that was published by GSA. The outburst that led to this discovery is the only outburst of this object that has been observed so far. Subsequent follow-up observations have resulted in high-precision measurements of the binary parameters that had not been possible for this class of object before (Green et al. 2018).

The high photometric and astrometric precision (~50 milli-arcsec per transit) also makes *Gaia* sensitive to gravitational microlensing events, and several microlensing candidates have already been alerted on. Microlensing events occur when a star crosses our line of sight towards a distant background star and is observed as a temporary magnification of the background starlight. In 2016 *Gaia* detected the first binary microlensing event in the Galactic disc, Gaia16aye. The *Gaia* data, along with subsequent time series follow-up observations, afforded a full solution of the binary parameters, showing that this is a K giant doubly lensed by a main sequence binary (Wyrzykowski et al. 2020). The observations illustrate the potential for measuring the mass function of dark objects through microlensing.

A unique feature of GSA is that it is also able to alert on sources that fade significantly. In this way, many new young stellar objects (YSOs) and other 'dipping' sources, such as VY Scl stars, have been discovered or alerted on. Gaia17aeq is shown as an example in Fig. 2. This is an EXor variable – a YSO with a large proto-stellar accretion disc, characterised by large amplitude eruptive variability. It was originally discovered in outburst by the ASAS-SN survey as ASASSN-13db. A second, long-lasting outburst was underway when *Gaia*'s nominal observations started (see Sicilia-Aguilar et al. 2017) and GSA detected the accretion state change when it started to fade again towards quiescence (star symbol in Fig. 2). The time-series BP/RP spectra clearly illustrate the dramatic colour and spectral changes that accompany the flux variation in accretion events like these. ASASSN-13db/Gaia17aeq is the lowest mass star known to show outbursts like these (Holoien et al. 2014). Kashi et al. (2019) suggested that ASASSN-13db/Gaia17aeq may also be a luminous red nova, with the long-lasting outburst resulting from the disruption of the inner accretion disc or the accretion of a planet, but Cieza et al. (2018) confirmed its nature as an EXor variable, using ALMA observations of its dust disc. Several other YSO outbursts have been discovered as a result of the flaring activity observed by *Gaia* (e.g. Gaia18dvy – Szegedi-Elek et al. 2020, Gaia18dvz – Hodapp et al. 2019 and Gaia19ajj – Hillenbrand et al. 2019). A detailed study of Gaia17bpi (Hillenbrand et al. 2018) showed that this FU Ori-type outburst started in the mid-infrared, appearing at optical wavelengths approximately 1.5 yr later. This is the first of these outbursts to be detected at both

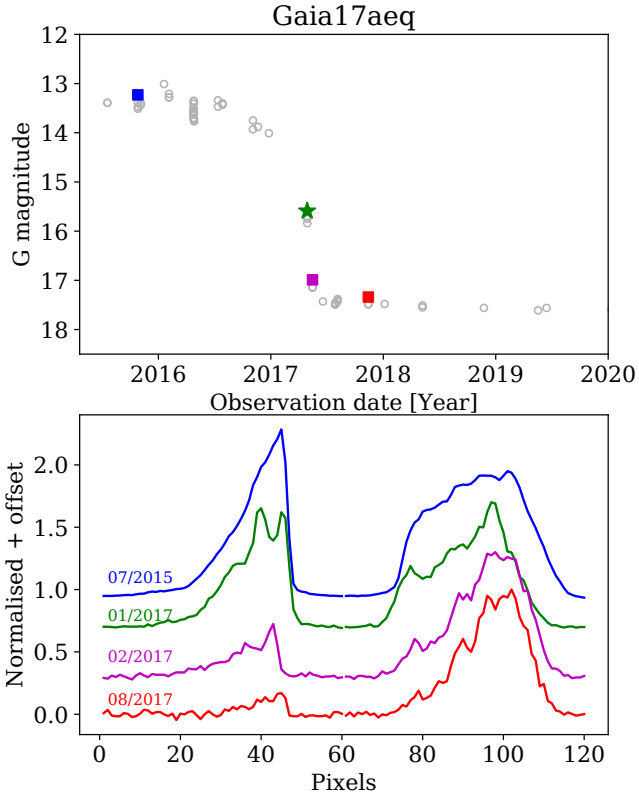

Fig. 2. Lightcurve (*upper panel*) and spectral variation (*lower panel*) of ASASSN-13db/Gaia17aeq. The points in the lightcurve for which the BP (lower left) and RP (lower right) spectra are shown are indicated with filled symbols in the same colour. The *Gaia* alert was issued when the target faded, at the point indicated by the star symbol.

parts of the spectrum and it serves as direct tests of accretion disc models in these large discs.

Finally, *Gaia* is making contributions to the growing field of transients occurring in the very centres of galaxies (in spite of incompleteness in these regions, see Kostrzewa-Rutkowska et al. 2018 and Sect. 4.2). One such event – Gaia16aax – has been detected in a galaxy hosting a known QSO where the centre brightened by about 1 mag over 1 yr, before fading back to its pre-outburst state over more than 2 yr. Both the photometric and spectroscopic variability show a dramatic change. The outburst of Gaia16aax can be explained by a change in the accretion flow onto the central black hole or could have been caused by a tidal disruption event (Cannizzaro et al. 2020).

In this paper, we describe the operational state of the *Gaia* Science Alerts survey. Section 2 gives a full technical description of the data flow including the ingestion of the main data, the alert detectors, filtering methods, eyeballing and publication. The main results are described in Sect. 3, which includes a summary of the GSA event rate, the photometric and astrometric precision of the candidates, and their main properties. The purity and completeness of the survey is discussed in Sect. 4 and we summarise in Sect. 5. We also include appendices with additional information on the cyclic processing of *Gaia* data and subsequent catalogue changes (Appendix A), the photometric calibration of GSA (Appendix B), details of the computing cluster (Appendix C), and a complete list of abbreviations used in this paper (Appendix D).

Throughout the paper, where we have performed analysis of the GSA detection rates (and contamination rates), or considered the performance of the photometry or astrometry, we have set a

[2] TNS; the official IAU mechanism for reporting new astronomical transients, https://www.wis-tns.org/

fixed range of observational dates, encompassed in a fixed set of Initial Data Treatment (IDT, Fabricius et al. 2016) runs. These runs and dates are: run 1046 (earliest data point: 2016-07-11 04:45:53) to run 4724 (latest data point: 2019-12-30 09:35:49). The starting point was set as the point in time when the largest part of our system had stabilised.

2. Data flow: from observation to alert

Gaia is at heart a time-domain experiment, measuring exquisitely precise astrometry and photometry with a well-defined observational depth and cadence. However, the daily processing of GSA cannot accumulate, and iteratively calibrate, data in the same way that is used for the main *Gaia* data releases. In this section we discuss how GSA proceeds from the on-board measurements taken by the *Gaia* spacecraft to the eventual publication of transient astronomical phenomena. We pay attention to how we curate the large data flow, apply simple calibrations, and filter out spurious detections, resulting in a viable and scientifically useful stream of transient events. An overview of the principal steps is described here (see also Fig. 3). Firstly, sources are detected and observed by *Gaia* as the spacecraft rotation and precession brings them through the fields of view (FOV, Sect. 2.1). Next, observations are downlinked and forwarded via the Mission Operations Centre (MOC) to ESA's Science Operations Centre (SOC) for processing (Sect. 2.2). SOC collates the telemetry from *Gaia* and performs the Initial Data Treatment (IDT), extracting positions and fluxes of the sources from the pixel data. The results are copied to the various data-processing centres of the Data Processing and Analysis Consortium (DPAC, Mignard et al. 2008), including the one at Institute of Astronomy Cambridge (known as DPCI) where alerts processing takes place (Sect. 2.3). GSA processes the data of the current IDT run[3], filtering the observations by quality, applying an on-the-fly photometric calibration, detecting transient features in the lightcurves, and flagging events suspected to arise from specific instrumental effects, as well as transients of astrophysical sources that are not worthy of alerts (e.g. known periodic variable stars and Solar System objects). This stage produces a list of candidate alerts (see Sect. 2.4 and after for details). All data are handled by the GSA PostgreSQL database which make use of the Quad Tree Cube (Q3C) software (Koposov & Bartunov 2006). Further filtering removes the alert candidates that are probably due to interference effects from neighbouring sources (Sect. 2.7). Human inspection (i.e. eyeballing) identifies those candidates suitable for publication (Sect. 2.8). Finally, the chosen alerts are published immediately to the World Wide Web via the Alerts Website, TNS entries and VOEvents (see Sect. 2.9).

Each alert is published with a timestamp corresponding to the observation time by *Gaia* (in barycentric coordinate time, TCB) as well as the time of publication of the alert (in Coordinated Universal Time, UTC). The latency between the two timestamps is the sum of: (1) the time from observation until downlink of the data to MOC (commonly less than 12 h, but significantly more in exceptional cases), (2) processing time at MOC and SOC, mainly in IDT (typically around 10 h), (3) time for automatic processing at DPCI (typically from 3 to 6 h, but rising to ~24 hours for scans that run tangentially along the Galactic plane), and (4) time for human evaluation at DPCI (see Sect. 2.8).

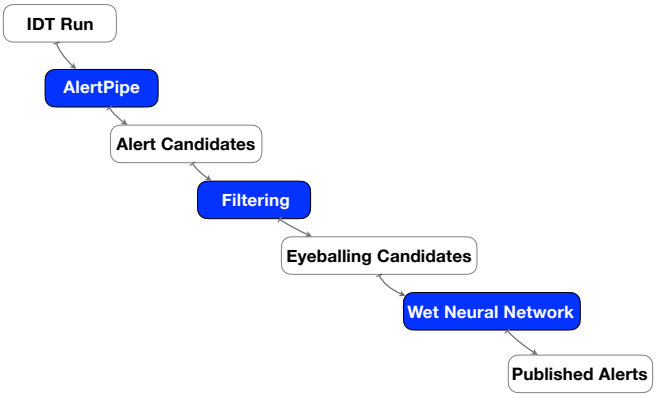

Fig. 3. Schematic of the data flow and processing performed by the GSA project. Unfilled boxes indicate data, blue boxes (dark grey) show processes. Wet Neural Network refers to the eyeballing, voting and commenting process performed by humans.

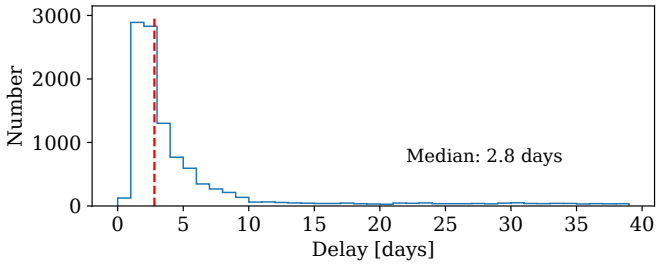

Fig. 4. Histogram of the delay (in days) between the *Gaia* observation and the publication of an alert. The long tail is the result of allowing the two FOVs that we require the alert to be seen in, to be separated by up to 40 days. Twelve per cent of the alerts have a publication delay longer than 10 days. The median delay is 2.8 days and is indicated by a dashed vertical line.

Alerts, therefore, typically appear between 24 and 96 h after the triggering observations (median delay is 2.8 days, see Fig. 4). There is also a long tail, which corresponds to the delay between detection in two different FOVs (up to 40 days), discussed in Sect. 2.6.1.

2.1. Observations and data types used in alerts processing

Gaia is a drift-scan survey with two telescopes whose FOVs are separated by 106.5°. The closely controlled rotation of the spacecraft scans the two FOVs, which are both pointed perpendicular to the spin axis, across the sky once every 6 h. Precession of the spin axis, and the satellite's orbit around the Sun, varies the part of the sky observed on each rotation.

Each pass of a source across a FOV is termed a 'transit', and this is the fundamental unit of observation. In a transit, a source crosses first the sky mapper (SM) CCD, then nine CCDs (except row 4 which has eight CCDs) of the astrometric field (AF), then the CCDs of the blue and red photometers (BP and RP) where the light is dispersed by prisms to obtain low-resolution spectra, then finally the grid of the Radial Velocity Spectrometer (RVS). The SM and AF measurements are in white light (covering 330–1050 nm, Evans et al. 2018).

The on-board algorithms responsible for the detection, selection and confirmation of sources are described in de Bruijne et al. (2015) and Gaia Collaboration (2016). The magnitude limit for retaining an observation is $G = 20.7$.

[3] Processing requires having all of the IDT output relating to a given source, but IDT's output is not organised cleanly by source, because to do so would be expensive and inefficient. In practice, the alerts pipeline runs once per IDT run (typically one day of observation) and cannot begin until the last output for that run arrives at DPCI.

In alerts processing, use is made of the following IDT data: (1) fluxes measured on each AF CCD; (2) positions of the source on each AF CCD along and across the scan direction, extracted by means of a PSF/LSF fitting (Point/Line Spread Function, Fabricius et al. 2016); (3) the calculated RA, Dec; (4) integrated fluxes for the transit in BP and RP, plus the colour derived from their combination; (5) individual pixels of the BP and RP measurements; (6) matching of transits to sources in the working catalogue; (7) status flags describing the reliability of the IDT results.

Alerts processing does not use the SM data, the raw pixel values from the AF CCDs, or the RVS data (although RVS data were reported for a small number of alerts, for a limited time, see Sect. 2.9).

2.2. Downlinking of data

Typically, data from the *Gaia* spacecraft can be transmitted to three ground-stations (operated by ESA) at Malargüe (Argentina), Cebreros (Spain), and New Norcia (Australia). More recently, NASA Deep Space Network stations have also been used during some of the recent Galactic plane scans. The actual contact time is adjusted to match the predicted downlink data volume for the day, typically ~8–10 h, covered by one of the three antennae (two are used if the data rate is very high).

It is worth noting that the typical amount of (compressed) science data downlinked to the ground is some 40GB per day. Small onboard data losses (photometry and astrometry) can be caused by shortages of ground-station contact periods (e.g. in times when *Gaia* scans along the Galactic plane), amounting to zero for bright objects ($G < 16$ mag), a few per cent for $G = 16$–20 mag, around 10% for $G = 20$–20.5 mag, and ~25% for fainter objects (see Gaia Collaboration 2016 for details).

The MOC, located at the European Space Operations Centre (ESOC) in Darmstadt buffers the data packets and forwards them to the SOC near Madrid. SOC marshals the data into the standard formats of DPAC, and runs IDT.

2.3. Initial data treatment

The main role of IDT is to generate self-contained raw data records, extract the fluxes and centroids for SM, AF and BP/RP CCDs, and to match transits to catalogue sources (Fabricius et al. 2016). These processes are done in a time-constrained computer system where fully consistent processing is foregone in favour of prompt delivery to other data processing centres; both are subject to data artefacts that can cause false alerts. IDT breaks its operations into runs, where a typical run contains roughly one mission day of data.

IDT reconstructs the spacecraft attitude to enable generation of the first on-ground attitude (OGA1), and thus the computation of source positions in sky coordinates (RA, Dec), to a required accuracy of ≤100 milliarcsec (Fabricius et al. 2016, but see Sect. 3.4 for a discussion of the GSA astrometric precision which we find to be ~55 milliarcsec). These reasonably accurate coordinates are used in the cross-match between the transits in the current IDT run and the *Gaia* working catalogue of the current data reduction cycle[4].

[4] *Gaia* Data Release 3 (DR3) is based on Cycle 03 processing, while the *Gaia* alerts included in DR3 are based on the Cycle 01 and Cycle 02 IDT working catalogues.

2.3.1. IDT new sources

A transit which can be associated with an IDT working catalogue source is assigned the appropriate sourceId (defined in Bastian 2013), while one that cannot, triggers the generation of a new sourceId (which is added into this catalogue).

The magnitude limit for detection of a source by *Gaia* is $G = 20.7$. Some 15–20% of all *Gaia* detections are spurious detections on board (Fabricius et al. 2016), and ~80% of these cases are flagged in IDT. The most common causes of spurious detection include: diffraction spikes, bright sources from the other FOV, major planets (especially Venus), diffuse objects, duplicated detections, cosmic rays and hot CCD columns (see Fabricius et al. 2016 for a detailed description of the causes and mitigation strategies). Occasionally, large numbers of new sources can be generated when the OGA1 attitude solution for the spacecraft suffers an excursion. This can arise when the spacecraft suffers disturbances from external micro-meteoroid hits. Later processing, and in particular the Astrometric Global Iterative Solution (AGIS, Lindegren et al. 2016) do a much better job of modelling these excursions, but these are beyond the timescale constraints of the IDT and GSA systems. Rarely, *Gaia* also detects very large numbers of prompt particle events, associated with Solar coronal mass ejections. The largest such event occurred on 10 September 2017[5], which resulted in our system flagging spacecraft revolutions 5651.4 to 5659 as bad.

At the beginning of operations, the working catalogue was the Input *Gaia* Source List (IGSL, Smart & Nicastro 2014), and the maximum radius for cross-match was set to 2.0 arcsec. During this first phase, large numbers (up to millions) of new sources were generated in each IDT run, due to incompletenesses and inaccuracies in the IGSL, as well as from spurious detections. Over time, as the working catalogue has been improved, this cross-match radius has been reduced to 1.0 arcsec. The number of new sources arising via this channel has been significantly reduced.

2.3.2. Cyclical processing and catalogue update

IDT is a daily process using a working catalogue which is updated as new detections arise. The primary data-releases of *Gaia* are derived from cyclical reprocessing of the whole data set, in which a new catalogue is formed from consideration of all observations. When IDT replaces its working catalogue with the new, cyclical catalogue, GSA experiences disruption. This can lead to gaps in the alert lightcurves, or to lightcurves being the union of observations of several sources in the new catalogue. See Appendix A for details.

2.4. GSA lightcurve processing

The data in an IDT run represent new observations of sources (identified by their sourceId) in the *Gaia* working catalogue, and observations of new sources. GSA processing starts with the building of lightcurves for all these sources. A full lightcurve is the union of all observations assigned to that sourceId by IDT over all runs, with photometric calibration applied on-the-fly, and precise to 3% at $G = 19$ (see Appendix B). A typical run contains on average 60 million observations (transits) arising from some 37 million sources, but the amount of data in each

[5] https://blogs.esa.int/rocketscience/2017/11/03/unexpected-view-from-gaia-the-galactic-surveyor/

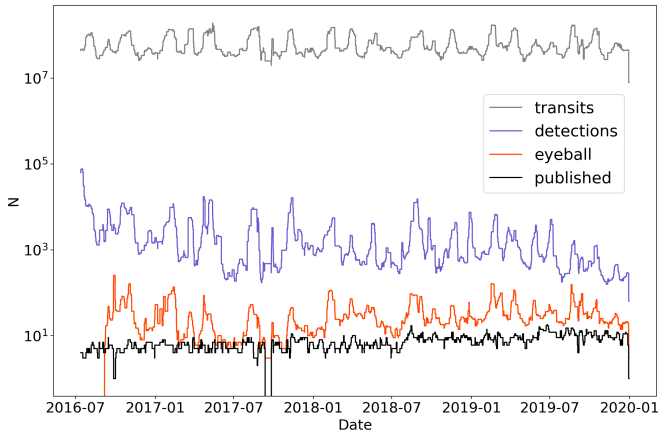

Fig. 5. Histograms showing the time evolution of four quantities (all are totals per-IDT run): (1) numbers of transits processed by AlertPipe (in grey), (2) numbers of automated alert detections from AlertPipe (in blue), (3) numbers of alerts presented to eyeballers after additional (mostly environmental) filtering (in red), (4) numbers of alerts published from each run (in black). A 7-day running median filter has been applied to all totals. Note that no records of eyeballing statistics were preserved for the first months of 2016.

run varies widely according to the current direction of *Gaia*'s scanning across the Galactic disc (see Fig. 5).

With the lightcurves formed, GSA processing proceeds to the evaluation of each source that has received new data in the current run. The principal steps for each lightcurve are: (1) filtering to remove untrusted transits from the lightcurve (Sect. 2.5); (2) detection of transients, using four different algorithms (see Sect. 2.6); (3) automated classification of transients to flag possible artefacts from instrumental effects and astrophysical transients not suitable for an alert. These can include excursions in the spacecraft attitude from nominal pointing, proximity to minor planets, or already classified long period variable stars (from DR2 Gaia Collaboration 2019, see also Sect. 2.7).

This gives a list of alert candidates for the run. A typical run produces a few thousand candidates (see Fig. 5). Over the duration of the mission, the number of candidates detected by GSA processing has decreased, while the number of published candidates has increased, demonstrating a trend of increased efficiency.

2.5. Filtering of bad transits

An observed transit may be eliminated from a lightcurve for a number of reasons. For example, the details (flux, position) of the transit may be flagged by IDT as improperly extracted from the pixel data. Alternatively, the transit may have been observed when *Gaia* was not in a stable state, as when the mirrors were being heated to remove condensates. Sometimes, the readout parameters may be inconsistent with the magnitude of the source[6]. Another example occurs when the scatter in the distribution of fluxes obtained from the individual CCDs is significantly higher than expected from photon statistics. This is evidence of interference from sources in the opposite FOV.

[6] *Gaia* may observe multiple sources simultaneously on the same TDI (time-delayed integration) line (de Bruijne et al. 2015). The readout parameters are set to suit the brightest of such sources, and this may compromise observations of the fainter sources.

Where a transit is filtered, its flux is not used in detecting transients. If an alert is published for that source, the transit appears in the published lightcurve with no stated magnitude.

2.6. Alert detection algorithms

Transients are detected in the lightcurves formed from the white light G-band fluxes measured by the AF CCDs. Four detection algorithms are applied to detect different kinds of events.

2.6.1. New source detector

This reacts to sources that brighten from below *Gaia*'s detection threshold. A source not previously seen and rising to $G < 19$ is considered as an alert candidate.

To defeat the many sources of systematic noise, some other criteria must be met for the detector to report a candidate. The source must be seen in both FOVs; many effects result in a spurious detection in only one FOV (see e.g. Wevers et al. 2018 and Kostrzewa-Rutkowska et al. 2020 for more detailed discussions). The location of the source must have passed through *Gaia*'s FOV at least 10 times previously without detection (calculated using HEALpix with a resolution of ∼40 arcsec). Due to instrumental and resource limits, not all transits of all sources are recorded, with fainter sources (in crowded regions in particular) more likely to be lost before transmission to Earth.

Observations of a newly-visible source may be split between IDT runs, and this would cause the detector to miss them if no single run contains detections in both FOVs. To avoid this, the detector aggregates all observations of the source in the current IDT-run and in all previous runs. Any transits in the current run and in the preceding 40 days are potentially from a new source. Any older transits are taken to indicate a previously known source and the NewSource detector is not triggered.

Requiring detection in both FOVs improves the cleanliness of the alert stream at the expense of completeness. An alternative approach (Kostrzewa-Rutkowska et al. 2020) would be to alert on each detection of a previously unknown source: one alert per FOV transit. This would be suited to finding brief, faint transients such as possible optical counterparts of gravitational wave events, but at the expense of increased contamination. Work on implementing this detector is ongoing.

2.6.2. Old source delta-magnitude detector

This detects gross changes in the brightness of sources already in the IDT working catalogue. It reacts to the more extreme events (e.g. cataclysmic variables) but can also detect supernovae that are not resolved spatially from the nuclei of their host galaxies (where the galaxy is in the *Gaia* catalogue, otherwise this would be a NewSource alert).

Measurements in the lightcurve obtained within 40 days of the most recent measurement are analysed for transient behaviour, while the mean and standard deviation of older measurements are taken as a historic baseline for comparison. To become an alert candidate, the lightcurve must have at least two transits that differ from the historic mean by at least one magnitude and by three times the standard deviation of the baseline.

The scatter of measured positions on the sky is used to rule out cases where transits of two separate (barely-resolved) sources have been mixed. To survive as an alert candidate, the source must have a standard deviation in position of less than

0.1 arcsec. This may have a negative impact on transients arising in marginally resolved sources such as galaxies.

2.6.3. Old source mean-rms detector

This detector is similar to the old source delta-magnitude detector, above, but detects smaller changes in the quieter lightcurves. The minimum change in brightness is reduced to 0.15 magnitudes, but the deviating transits must change by at least six times the standard deviation of the baseline flux.

2.6.4. Skewness/Von Neumann detector

This detector, hereafter called OldSourceSkewVonN, exploits the available source history to search for slower photometric variability. It was designed to cover a parameter space that is complementary to the other detectors. It is based on slicing a parameter space consisting of the third moment of the distribution of magnitudes (the skewness) and the von Neumann statistic η. The latter is defined as the ratio of the mean square successive difference to the variance (von Neumann 1941):

$$\eta = \frac{\delta^2}{s^2} = \frac{\frac{1}{n-1}\sum_{j=1}^{n-1}(m_{j+1} - m_j)^2}{s^2}, \quad (1)$$

where n is the number of datapoints in each lightcurve, s is the standard deviation of the lightcurve, and m_j are measured magnitudes in the G-band. A strong positive serial correlation between datapoints leads to a low von Neumann statistic, which signifies smooth variability, as opposed to single outliers or non-variable lightcurves which result in large η values (see e.g. Wevers et al. 2018, Kostrzewa-Rutkowska et al. 2018 for an application to *Gaia* data). The skewness metric can be used to remove stochastic/periodic variability.

One advantage of the OldSourceSkewVonN detector is that it is well suited to finding relatively low amplitude events with high fidelity, such as microlensing events, variable AGN, and YSOs. The need for a sustained upward/downward trend in the lightcurve makes this detector robust against artefacts and outliers. The downsides are that (i) it requires sufficient history – it was only brought into operation in May 2019 – and (ii) several outlying data points are required before detection can be triggered, thus there is a delay between the start of the event and its detection.

2.7. Spurious alerts

To all intents and purposes, GSA is a catalogue-driven transient survey, because two-dimensional AF pixel data are not available for the vast majority of sources. The strength of many of the extant ground-based transient surveys, including ZTF (Bellm et al. 2019), ASAS-SN (Shappee et al. 2014) and PanSTARRS (Chambers et al. 2016), is that they employ difference-imaging techniques, thus the operators and users can ultimately inspect the images, and decide on the veracity of each event. For GSA this lack of an image, and constraints on the release of *Gaia* data ahead of formal data releases, pushes us to deliver a high-purity alert stream, whereby a high degree of candidate vetting and rejection is performed in house.

Some statistics for the processing of GSA are shown in Fig. 5. Each run of AlertPipe handles on average 60 million transits for 37 million sources (maximum values can reach in excess of 300 million transits for 200 million sources). The vast majority of these measurements are not unusual, or are easily identifiable as spurious (e.g. big dippers, attitude excursions etc, see below), leading to a median raw detection rate of ∼1000 alerts per run, thus about 30 per million sources show anomalous flux behaviour. More detailed filtering, particularly exploring the environment of the candidates (see Sect. 2.7.1) leads to a reduction in the median number of candidate alerts by a further factor of ∼50. Thus, about 20 candidates per run survive to the phase of human eyeballing, and about half of these are published.

Not everything which is found by the detectors is something we wish to alert on and publish. There are many types of false positives, some of which are the real behaviour of real sources (such as periodic variables and asteroids), some of which are spurious behaviour of real sources (such as an increase in flux due to a bright star or planet lying nearby in the along-scan (AL) or across-scan (AC) direction from the source), and some are completely spurious sources (such as apparent new sources reported during attitude excursions, which are in effect the misplaced detections of old sources). Here we describe the mitigations we have put in place for some of the leading causes of false positive alerts. It is worth noting that there has been significant evolution in the rates of the differing types of false positives throughout the *Gaia* mission. These have arisen from (i) changes to the on-board *Gaia* detection parameters, (ii) improved mitigations in IDT, combined with updates to the working catalogue, (iii) evolution in our own understanding of the data and identification of spurious events. As an example, in the first half of operations during 2016, we employed a source-density map of the sky to reject all transients found in the most crowded regions (the density map was constructed from the GSA database). Once more thorough environmental filters were developed and tested, the use of the density map was discontinued (i.e. by June 2016).

In the following sections we detail the most common types of false positive, which are either trapped and rejected in AlertPipe, or in two cases, flagged for inspection by the eyeball team. The most common types of automatically rejected alerts are summarised in Fig. 6. The two classes of candidate false alert that require a human decision are (i) Solar System objects (SSO), and (ii) variable star in *Gaia* DR2 (Sect. 2.7.6). The first case is very rare, and almost always is unrelated to the alert (e.g. a faint SSO is reasonably close to a bright CV candidate). For the second case, the human almost always follows the flag, however due to occasional misclassifications in the variable catalogue, we do not automatically reject candidate alerts that are cross-matched to classified variable stars in *Gaia* DR2.

2.7.1. Environment: alignments in AL/AC directions, planets and bright stars

While planets and the very brightest of stars can induce spurious alert candidates over a large area (∼2° radius, Fabricius et al. 2016), less bright stars still have an impact, albeit over a smaller area around their locations. A bright star can cause a spurious alert candidate as a result of flux from a diffraction spike entering into the window of the alert candidate, and producing an apparent increase in magnitude. The amount of additional flux depends on the magnitude of the bright star, the separation, and the alignment with respect to the orientation of the scanning direction (the spikes are asymmetric and aligned in AL and AC).

Alignments in the AC direction between the bright star and the candidate are particularly difficult to deal with as there is no noticeable impact on the goodness-of-fit statistics of the candidate. A common arrangement which produces a large population of spurious OldSourceMeanRms candidates is a neighbour

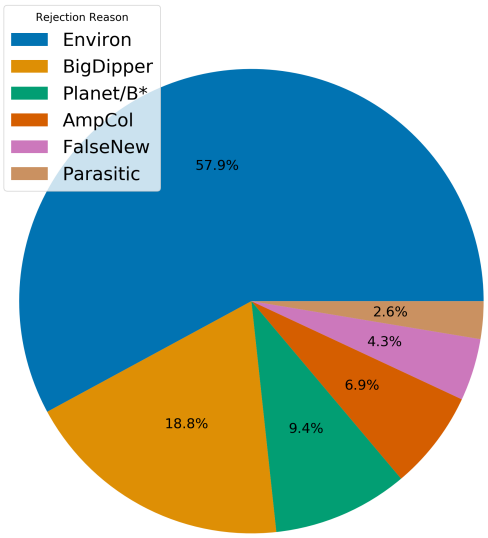

Fig. 6. Most common classes of spurious alert that are rejected automatically by the GSA system (Sect. 2.7). Environ: alert rejected after assessment of near neighbour(s) within 10 arcsec, BigDipper: alerting transit is in wing of bright star, which leads to a fainter measurement being associated with the star (i.e. the window is effectively misplaced, Sect. 2.7.2), Planet/B*: alert likely caused by influence of bright star or planet in vicinity, AmpCol: likely Mira-like variable on the basis of historic flux-scatter and extreme red colour, FalseNew: source was new to the GSA database, but not new to IDT. This could be due to missing or late arriving data for an IDT run which was not ingested into the database, Parasitic: second FOV source affecting flux within transit for alerting source.

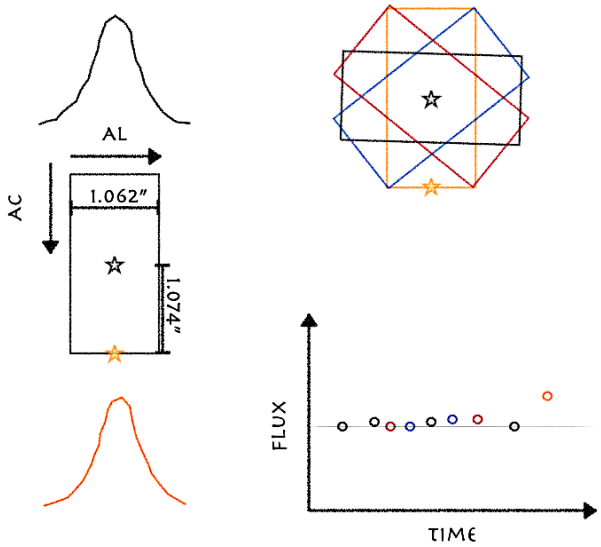

Fig. 7. Transits of sources fainter than $G = 13$ are one-dimensional with the assigned window divided into different samples in the AL direction with each sample spanning the full length of the window in AC. We illustrate, above and below the acquisition window (*left panel*), the approximate shape of the LSF provided through these one-dimensional samples. If a near neighbour aligns with the source in the AC direction this can cause an enhancement in the flux recorded for the source without disturbing the goodness-of-fit statistics of the transit in question. This sketch shows the impact of this arrangement on the lightcurve of the source; how the rectangularly shaped windows can capture flux from neighbouring sources in preferential directions, and hence how this may generate a spurious alert.

between 1 and 2 arcsec away in the AC direction. This alignment can result in a significant amount of flux from the neighbouring source entering into the window of the alerting candidate source and producing artificial brightness variations. This arrangement and its effects is illustrated in Fig. 7.

The environment of every alert candidate, therefore, must be examined to reject such artificial flux variations. This assessment is performed in sky coordinates, rather than *Gaia* detector coordinates. Although sources from both FOVs can be adjacent in pixel coordinates, their differential motion will vary their separation, and thus lead to a variation in flux across the AFs (i.e. within a transit). These alerts are weeded out. The amount of additional flux required to induce an alert candidate depends on the historical magnitude of the candidate and the detector type (faint OldSourceMeanRms alert candidates are the most vulnerable to this effect). Mitigation of this effect may then be expected to depend on the detector, the historical magnitude of the alert candidate and the magnitudes and angular separations to neighbouring stars which are as bright or brighter than the alert candidate.

The exact implementation of the environmental assessment is based on empirically derived magnitudes and angular separation distances as well as computational considerations. Extending the environmental search out beyond 10 arcsec for every alert candidate becomes infeasible in terms of CPU time. Instead for the brightest of sources (planets and the top 30 brightest stars) the environmental search is done in reverse by finding all the alert candidates near them. There is a subset of alert candidates caused by environmental effects, therefore, which could survive to the eyeballing stage (see Sect. 2.8) should there be a bright enough star beyond 10 arcsec. However, these are sufficiently few in number to be dealt with at that stage.

While not an every day occurrence, when the location of Jupiter, Saturn or Venus is near the scanning path of *Gaia* they can cause many spurious alert candidates by increasing the apparent fluxes of sources. The same is true for some of the brightest stars in the sky. Hence all alert candidates within 2 degrees of a planet or one of the top 30 brightest stars in the sky are assessed.

The local (within 10 arcsec) environmental assessment is performed as follows:

– All alert candidates are rejected if they are not the brightest source by at least 1 magnitude in a 1.5 arcsec radius about their median position. For OldSourceDeltaMag and OldSourceMeanRms alert candidates this radius is extended to 2 arcsec;

– NewSource, OldSourceDeltaMag and OldSourceMeanRms alert candidates are rejected if there is a neighbouring source within 10 arcsec which is brighter than $G = 12$ mag;

– NewSource alert candidates with a source in the AL or AC direction within 10 arcsec and $G < 17$ may be rejected depending on the relative magnitudes of the two sources, as spurious non-blacklisted detections may still occur due to the AL/AC PSF spikes of these sources;

– The rejection criteria for OldSourceDeltaMag and OldSourceMeanRms alert candidates are stricter, as sources fainter than $G = 17$ in the AL/AC directions may still cause a brightening in an existing source even if they cannot cause an entirely spurious detection. For these alert candidates, if the alert is due to the brightening of the candidate, any source within 10 arcsec in the AL/AC direction may lead to a rejection, again depending on the relative magnitudes of the two sources;

– Note that OldSourceSkewVonN candidates undergo less filtering as the detector is sensitive to long-term changes rather than short term ones produced by unfavourable alignments.

All OldSourceDeltaMag and OldSourceMeanRms alert candidates with an historic magnitude fainter than 19, and within 2° of a planet or one of the top 30 brightest stars in the sky, are immediately discarded. For the remaining alert candidates, their positions in AL and AC with respect to the planet or bright star are evaluated, and any candidates in a predetermined box around the planet are rejected. The size of the box is determined by the area in which there is a clear excess of alerts. The size of the box is larger in area for OldSourceMeanRms candidates, than for OldSourceDeltaMag candidates, whereas OldSourceSkewVonN alert candidates use the same exclusion region as the OldSourceDeltaMag candidates if the change in magnitude from the historic magnitude is less than 0.5 magnitudes. If the change in magnitude is greater than this, they are not automatically discarded. The box size is always at least a degree wide in AL and always more than 1.5° in AC.

2.7.2. Big dippers

Early on it was noticed that there were large numbers of alert candidates which had alerted due to their associated source having dimmed by several magnitudes. These sources were predominately in the magnitude range $13 \leq G \leq 17$. Further investigations revealed the position of the alerting transit to be offset from the median source position. It is thought that these observations are due to bright-star artefacts, where the onboard algorithm detects the spikes of the point-spread function, resulting in a fainter measurement for the same sourceId.

For brighter stars ($G < 13$) these spurious offset transits are successfully blacklisted by IDT, at least in the vicinity of the star itself, and hence removed from the data-stream for AlertPipe. However, this proved not to be the case for fainter stars and additional processing was required to remove the resultant alerts from the list of candidates. This was done by evaluating the median position of the source and rejecting any alert caused by a drop in magnitude and a transit located more than 0.3 arcsec from this position. Note a transit which brightens and is offset is not rejected, to allow the discovery of supernovae whose host galaxies are detected by *Gaia*. This source of spurious alerts was significantly reduced once IDT updated their algorithm to include the region around fainter sources when blacklisting transits which are due to this effect.

2.7.3. Attitude excursion: hits, clanks

Large scale attitude excursions are rare events, but when they occur they can render the data unusable. In GSA, an indicative measure of the reconstructed attitude is achieved per IDT run by accumulating the offsets in AL and AC of each transit of a NewSource alert candidate to the median position of the source in question (we recall that a NewSource alert must have at least two transits in order to be an alert candidate). The width of the distribution of the AL and AC offset may then be compared against that expected. These offsets may also be displayed as a function of time, highlighting periods of excess error. Diagnostic plots are created for each IDT run, and form part of the final verification process described in Sect. 2.8. Additionally, as large scale attitude excursions generate many spurious NewSource alert candidates, any NewSource alert candidate which does not have at least two transits located within 0.3 arcsec from the median position of the source, is rejected automatically. Smaller shorter-term attitude excursions rely on the final inspection step prior to publication for their rejection, where the location of the alerting transits are compared to those of the other transits belonging to the source.

2.7.4. Prompt particle events and parasitics

Prompt particle events (PPEs) are high-energy particles, such as cosmic rays or trapped protons from the Solar wind, which may cause noise in the signal read out from *Gaia*'s CCDs. Parasitics are instances where a source from the other FOV happens to be projected onto the same location on the AF CCDs (Wevers et al. 2018). As the AC rate is different for the other FOV, and thus the star-path is not parallel, this projection only contributes to a few of the AF CCDs along the transit rather than all of them. It is for this reason that we require eight reliable (as defined using IDT's flags) AF flux measurements per transit and take the median value (and its error computed by median absolute deviation statistics) for the value of the transit's flux and its error.

In addition, for the OldSourceMeanRms alert candidates, the goodness-of-fit (GoF) measures of the PSF/LSF to the transits are used as an additional means to reject suspicious candidates. The GoFs belonging to the alerting transits are compared against the expected GoF from the historical transits, and if there are too many significant outliers in the alerting transits the candidate is rejected. Note that the GoF has a magnitude dependence so this method is not applied to the OldSourceDeltaMag alert candidates.

2.7.5. Solar system objects

As part of the DPAC processing system, the predicted *Gaia* transits of SSOs are calculated roughly every year and shared with GSA (Mignard 2016). The transit times are accurate to < 0.02 s and account for planned changes to the *Gaia* Nominal Scanning Law.

If an alert candidate is found to be within 2 arcminutes of the expected location of a known SSO as seen by *Gaia* the candidate is flagged as a tentative match, if it is within 2 arcsec then it is flagged as a probable match. An associated match probability is calculated, which depends not only on angular separation, but also on the magnitude difference between SSO and alert candidate. The flagging does not remove the candidate automatically, but this information is retained for the final verification step prior to publications, see Sect. 2.8, where the likelihood that the alert candidate is due to the observation of the SSO may be assessed.

2.7.6. High amplitude variables: known and unknown

Gaia DR2 included classifications for more than 550 000 variable stars, many of which are periodic (Holl et al. 2018). From 2018, we began to compare the GSA candidates to these DR2 tables, and flag them if already classified (see Fig. 6). But the DR2 candidates, drawn from 22 months of data, are not a complete sample, so additional strategies were devised to automatically identify the large numbers of high amplitude variables (such as Miras) which were still being seen in the OldSourceDeltaMag alert candidates.

Searching for periodicity proved problematic given the poor and non-uniform sampling of the lightcurves, however cuts in the colours and in statistics which are indicative of a high scatter in the lightcurve have proved useful in removing many Miras prior to the final verification procedure (see Sect. 2.8). These

cuts were empirically derived using the data itself, selecting cuts on parameters which would remove as many candidates previously rejected by the final verification step as possible without resulting in the loss of any published alerts. If the median colour (BP–RP) of the source is >4.0 and the median absolute deviation (MAD) of the magnitude is >0.3, the alert candidate is rejected. Additionally, if the median colour is >4.4 and the kurtosis of the magnitude >0.4, the alert candidate is rejected. This results in a reduction of the order of 40% in the number of high amplitude long period variables surviving to the verification (eyeballing) stage.

2.7.7. Salvaged alerts

Our filtering approach errs on the side of caution, to avoid placing excessive burden on the eyeballing process (see Sect. 2.8). This would suggest that we generate a pure sample of events, but with reduced completeness (see Sects. 4.1 and 4.2 for more discussion). We know (Kostrzewa-Rutkowska et al. 2018) that an independent search for transients in galactic nuclei can find bona-fide events missed by GSA, however the extra eyeballing required prohibits daily operation.

To mitigate against some of these lost events, we introduced (27 June 2017) a method for salvaging alerts discovered by our detectors, but rejected by the filters. There are four scenarios which we include, and which are passed to the eyeballers:

1. Transients which are near a known galaxy in the LEDA catalogue (Makarov et al. 2014).

2. Transients which are spatially-coincident with externally reported events. We maintain a comprehensive list of events discovered in other surveys within the GSA database (see Sect. 2.8 for more details).

3. An additional and independent filter (LWfilter) was brought into operation in July 2018. LWfilter uses auxiliary data from other surveys to classify the source (star, galaxy, AGN). Additionally, each alerting lightcurve is fit with a microlensing model (Paczynski 1996) in order to identify potential microlensing events. BP–RP colour is also used to identify blue flares (e.g. CVs and Be stars) and very red variables (e.g. long-period variables, such as Miras). Alerts are then inspected visually (by the Warsaw team) and added into the list for eyeballing. Until the end of 2019 (IDT run 4724, i.e. spanning 18 months) this filter added 323 alerts for Eyeballing, primarily (85%) from the OldSourceSkewVonN detector.

4. We also salvage candidates that are spatially coincident with a bespoke set of catalogues of YSOs (compiled by some of the authors). These include: (1) a catalogue of optically selected YSOs, (2) a catalogue of YSOs based on Spitzer observations (compiled from published articles), (3) a catalogue of confirmed YSOs from the Spitzer c2d survey (Young et al. 2015), (4) a catalogue of candidate YSOs from the Spitzer c2d survey, and (5) a list of candidate YSOs published by Marton et al. (2019).

Salvaging does not make a large difference to the numbers of alerts we publish. In a 2-yr period from 1-Jul-2018 to 30-Jun-2020 (IDT runs 4026–4956), we published a total of 7568 alerts, of which 945 (around 12%) came through the salvaging route. The breakdown for the 4 channels listed above are: (1) 187 events, (2) 365 events, (3) 340 events, (4) 53 events. Over half of the candidates were rejected by the filters because there were nearby neighbour sources, or because the IDT cross-match split the event across multiple sourceIds (sometimes in error).

2.8. Eyeballing

After detection and filtering, surviving alert candidates are subject to human evaluation using a web application, the Eyeballing App. This presents team-members with a series of figures and charts displaying *Gaia* and ancillary data. These data are used by the eyeballer to rank the candidate with a score between −1 and +1. A comment box is provided for the eyeballer to describe the event for the community[7], and a dialogue box enables internal discussion between eyeballers. Votes from a minimum of two eyeballers, with a net score of +2, are required for an alert to be deemed publishable. A total of 15 people have contributed to the eyeballing of *Gaia* Alerts over the duration.

The *Gaia* data made available to the eyeballer include: The calibrated lightcurve, including the photometric scatter within a transit; The line spread function goodness-of-fit vs. time of the alerting source, derived from the image parameter determination in IDT (Fabricius et al. 2016); All near-neighbour *Gaia* transits within 10 arcsec of the alert, projected in RA-Dec and AL-AC directions; Radial distribution of all neighbour transits out to 10 arcsec (magnitude versus separation); Uncalibrated BP/RP spectra showing the evolution of the source before and after alert (if available); The probability of a known Solar System object crossing the FOV; a flag if the source is already classified as a long-period variable star in *Gaia* DR2 (Holl et al. 2018); *Gaia* DR2 parameters (including parallax, proper motion, BP/RP colour); HR diagram with the candidate superimposed (when possible).

Between 2014 and August 2018, we also applied a classifier (GS-TEC, Blagorodnova et al. 2014) to the raw BP/RP spectra, and shared the results with the eyeballers. GS-TEC takes a Bayesian approach to model observed spectra, using a constructed reference spectral library and literature-driven priors. GS-TEC can classify SN, AGN and stars down to $G = 19$, however the classifier was disabled due to its significant execution time.

Auxiliary data are parsed from a variety of sources, and presented to the eyeballers, to help understand the context of a *Gaia* transient detection:

– To allow a visual inspection of the alert's location, the Eyeballing App shows the Aladin Lite (Boch & Fernique 2014) and SDSS finding charts;

– Results of positional queries to the Simbad (Wenger et al. 2000), NED and VSX databases[8], to determine whether it is an already-known transient or variable object;

– The list of YSOs described in Sect. 2.7.7;

– To aid the rejection of spurious transients arising from contamination by Solar System objects, we also display data on nearby planets, their satellites, and minor planets. In the early phases of GSA we used SkyBot (Berthier et al. 2006), but we now exploit ephemerides shared within DPAC (see also Sect. 2.7.5);

– Results of positional cross-match against our own tables of transient events, assembled from the hourly parsing of a significant collection of other publicly available transient surveys.

[7] The comment is limited to 100 characters, and draws on the eyeballer's experience to try to describe the event as succinctly as possible. The eyeballer may sometimes make an estimate of a possible classification.

[8] The NASA/IPAC Extragalactic Database (NED) is funded by the National Aeronautics and Space Administration and operated by the California Institute of Technology. VSX is the International Variable Star Index database, operated at AAVSO, Cambridge, Massachusetts, USA.

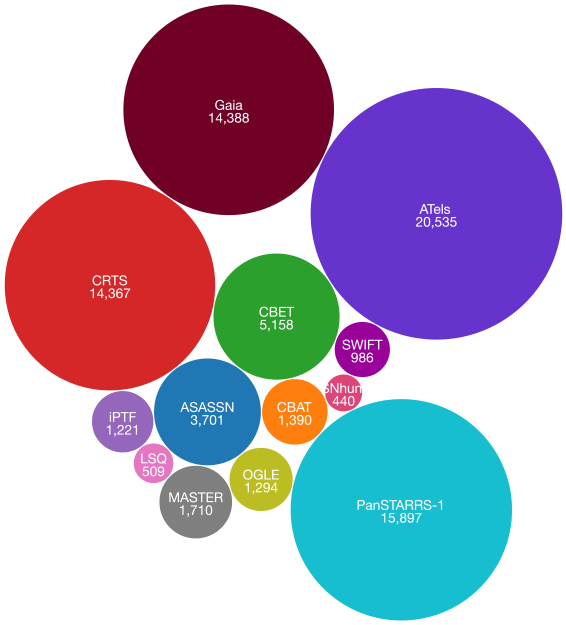

Fig. 8. Circles proportional in area to the unique numbers of objects/events (also shown in text), compiled for the GSA database, and colour-coded by the data source. The circle for *Gaia* is shown for comparison. References for the surveys are given in the text. The data are taken from a snapshot of our archive on 24 November 2020.

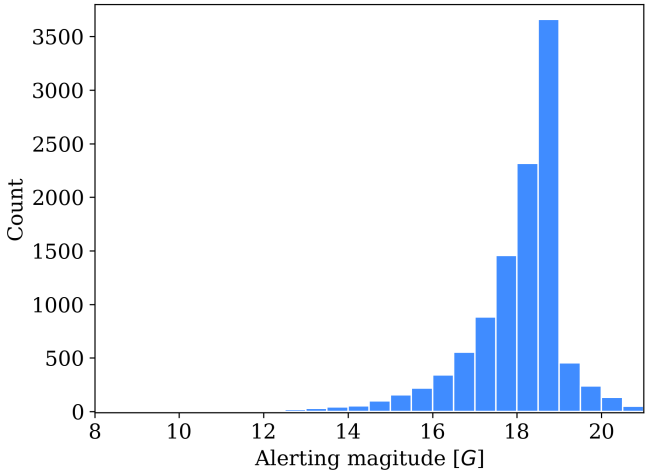

Fig. 9. Histogram showing the number of published alerts as a function of the alerting magnitude, covering the start of operations (September 2014) until the end of 2019.

An ETL (Extract Transform Load) system gathers discoveries reported by the major transient survey websites: Transient Name Server (TNS), Catalina Real-Time Transients (Drake et al. 2009), ASAS-SN (Shappee et al. 2014, Pan-STARRS1 (Kaiser et al. 2010), OGLE IV (Kozłowski et al. 2013, Wyrzykowski et al. 2014), MASTER (Lipunov et al. 2010), iPTF (Law et al. 2009), La Silla Quest (Baltay et al. 2013) and IAU Central Bureau for Astronomical Telegrams (CBAT[9]). Every hour, a total of 27 websites are scraped for data that are transformed, cleaned, homogenised and stored in the GSA database. In a similar manner, Astronomer's Telegrams are automatically parsed and stored in the database, accounting for the very diverse formats in the content of these HTML pages. The data stored in the GSA database for the external transient surveys is shown in Fig. 8;

– These data also contain classification information for large numbers of transient events which are shared with the eyeballer, and used at the point of publication. Classifications often arrive to the database after publication of an alert. As part of the publisher app, these can be viewed and the alert record updated (at the discretion of the operator). The bulk of classifications are reported via TNS (supernovae for the most part), but we also receive classifications on microlensing events from the Warsaw group through the publisher app (see Wyrzykowski et al. 2020).

2.9. Publication

Once eyeballing is complete, successful alerts are made publicly available to the astronomical community in several formats: via a dedicated website in CSV, HTML and RSS formats with permanent URLs for every published alert[10]; via the IAU-Transient Name Server[11]; as VOEvents using the 4 Pi Sky broker[12]. The GSA web application (Delgado et al. 2019b,a) has a public facing side where a set of the information is published, and a restricted area for administration and bookkeeping. The *Gaia* alerts catalogue can also be visualised on an All-Sky interface developed using Aladin Lite (Boch & Fernique 2014) enabling the display of alerts by time or individually.

For a limited time, the set of information published also included a small number of RVS transit spectra (27 spectra for 12 alerts: see Seabroke et al. 2020). This number is small for several reasons: (1) most detected alerts are much fainter than the limiting magnitude of RVS ($G_{\rm RVS}$ = 16.2 mag, while $G \sim 17$ mag for the alerts) see Fig. 9; (2) RVS covers only four of the seven *Gaia* CCD rows; (3) RVS spectra have much lower signal-to-noise than the other *Gaia* measurements at the same magnitude; (4) the pipeline used to produce the RVS spectra of alerts did not process blended windows or take into account flux from sources without windows (an issue because the majority of alerts with RVS spectra are close to the Galactic plane). The RVS pipeline now treats these issues and all RVS transit spectra will be published in *Gaia*'s fourth data release. This should provide additional useful diagnostic information for the brightest alerts.

Once an alert is published for a source, the alert page is permanent. New data from *Gaia* concerning that source are added to its lightcurve as they become available. Hence, the published description of the source is mutable and represents the most recent information available. The state of the source at the time of the alert is preserved in the VOEvent document released to the 4 PI SKY event-broker (Staley & Fender 2016) at the time that the alert is first raised. Once an alert is raised on a source, no second alert can be raised on the same source, even when subsequent events occur, for example in the case of repeated outbursts. There are a handful of exceptions (e.g. Gaia16acr≡Gaia16adx and Gaia16ade≡Gaia16aey) where new events in the same source are attached to a new sourceId arising from the IDT cross-match algorithm (see Sect. 2.3). Note that these duplicated alerts will also be included in the DR3 data release.

If an alert candidate does not pass the aforementioned filtering/eyeballing steps, future observations can raise another alert for the same source, which will then be re-evaluated, possibly leading to publication. Between the IDT runs 1046 and 4724 inclusive there were 556 published alerts which had previously

[9] http://www.cbat.eps.harvard.edu/index.html
[10] http://gsaweb.ast.cam.ac.uk/alerts
[11] https://www.wis-tns.org/
[12] https://4pisky.org/voevents/