

RoboTAP - target priorities for robotic microlensing observations

M. Hundertmark⁶, R. A. Street¹, Y. Tsapras⁶, E. Bachelet¹, M. Dominik³, K. Horne³, V. Bozza^{7,8}, D. M. Bramich, A. Cassan¹⁰, G. D'Ago¹⁶, R. Figuera Jaimes^{3,2}, N. Kains^{12,2}, C. Ranc¹⁵, R. W. Schmidt⁶, C. Snodgrass⁹, J. Wambsganss⁶, I. A. Steele⁴, S. Mao¹¹, K. Ment^{17,6}, J. Menzies⁵, Z. Li¹, S. Cross¹, D. Maoz¹³, Y. Shvartzvald^{14, and 13}

(Affiliations can be found after the references)

ABSTRACT

Context. The ability to automatically select scientifically-important transient events from an alert stream of many such events, and to conduct follow-up observations in response, will become increasingly important in astronomy. With wide-angle time domain surveys pushing to fainter limiting magnitudes, the capability to follow-up on transient alerts far exceeds our follow-up telescope resources, and effective target prioritization becomes essential. The RoboNet-II microlensing program is a pathfinder project which has developed an automated target selection process (RoboTAP) for gravitational microlensing events which are observed in real-time using the Las Cumbres Observatory telescope network.

Aims. Follow-up telescopes typically have a much smaller field-of-view compared to surveys, therefore the most promising microlensing events must be automatically selected at any given time from an annual sample exceeding 2000 events. The main challenge is to select between events with a high planet detection sensitivity, aiming at the detection of many planets and characterizing planetary anomalies.

Methods. Our target selection algorithm is a hybrid system based on estimates of the planet detection zones around a microlens. It follows automatic anomaly alerts and respects the expected survey coverage of specific events.

Results. We introduce the RoboTAP algorithm, whose purpose is to select and prioritize microlensing events with high sensitivity to planetary companions. In this work, we determine the planet sensitivity of the RoboNet follow-up program and provide a working example of how a broker can be designed for a real-life transient science program conducting follow-up observations in response to alerts, exploring the issues that will confront similar programs being developed for the Large Synoptic Survey Telescope (LSST) and other time domain surveys.

Key words. Gravitational lensing: micro – Methods: observational – Methods: statistical

1. Introduction

1.1. RoboTAP in context

In an era of increasing sky coverage and telescope étendue, surveys monitor billions of stars and alert whenever an astronomical object changes its brightness. Most science cases will require follow-up observations responding to these alerts. This is especially true for transient events that, by definition, occur only once. The sheer number of alerts issued by the Large Synoptic Survey Telescope (LSST), for example, can be on the order of one million per night (Ridgway et al., 2014) and therefore can not be handled by humans. Brokers for processing and filtering the alert stream for rapid response telescopes already exist, for instance the Arizona-NOAO Temporal Analysis and Response to Events System (ANTARES) (Saha et al., 2016), but more work is needed to address the variety of science cases. The demand for such systems will increase due to emerging projects such as the Evryscope (Law et al., 2015), the Zwicky Transient Facility (Smith et al., 2014; Bellm, 2014), the BlackGEM survey and last but not least the Large Synoptic Survey Telescope (LSST Science Collaboration et al., 2009). The RoboTAP system is such an automated broker system and a prototypical example for a range of transient science cases.

The philosophy behind the RoboTAP implementation is to provide a robust and computationally lightweight algorithm that only takes seconds for an event feed with ~ 1000 events. In order to achieve that, we followed a small set of basic working principles. We start by carrying out computationally inexpensive operations first and apply increasingly complex criteria later; we replace estimates that are only required to be roughly known by suitable interpolants. Incoming event parameters are kept and

updated in memory all the time. The same technical considerations apply to selecting a suitable and algebraically simple priority function.

Fig. 1 depicts the simplified workflow of RoboTAP acting as a broker between the event stream and the observation manager. The main source for the alert stream are existing microlensing surveys such as the Optical Gravitational Lensing Experiment (OGLE) survey (Udalski et al., 1994) or the Microlensing Observations in Astrophysics (MOA) survey (Bond et al., 2004). The workflow illustrates how an incoming stream of events is subjected to filters and triggers observations which in turn can refine the event parameters in a closed-loop system. It also highlights that further building blocks are necessary to efficiently process the event stream. Only events that have been classified as microlensing are processed. Moreover, we introduce two sub-classes - *regular* and *anomalous* microlensing events. For that purpose, we rely on the ARTEMiS system (Dominik et al., 2008) providing us with event parameters and anomaly triggers for events with irregular light curve shapes. In this two-tier approach, we generate preliminary target lists by reducing the number of transient (microlensing) events based on observability, survey coverage and expected duration. Finally, we rank all events based on our priority function. Anomalous and non-anomalous are allocated a pre-defined fraction of the observing time on the network. The corresponding target lists are then formatted as separate observing requests and submitted to the telescope network.

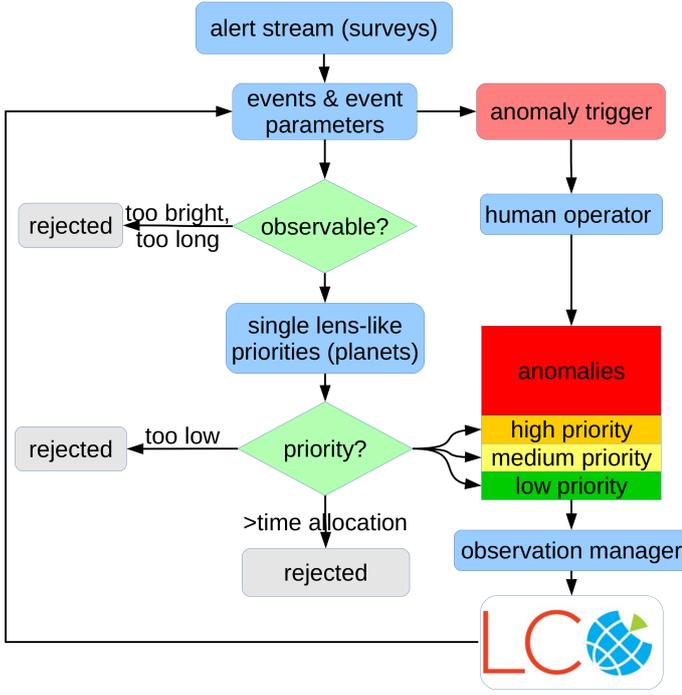


Fig. 1: Workflow of the RoboTAP system processing the incoming alert stream of regular and anomalous microlensing events from ARTEMiS.

1.2. Related observatories

While our approach is of general interest for all astronomical observatories, the actual implementation was customized to the Las Cumbres Observatory. The Las Cumbres Observatory (LCO) is an organization carrying out astronomical research and operating a global network consisting of homogeneous telescopes of different aperture classes (Brown et al., 2013). In addition to the two 2m Faulkes telescopes originally used by the RoboNet collaboration (Burgdorf et al., 2007; Snodgrass et al., 2008; Bramich, 2008; Tsapras et al., 2009; Street et al., 2011), $9 \times 1\text{m}$ telescopes became available in 2013 to the microlensing program of which 8 are on sites in the Southern Hemisphere, where most of the microlensing targets lie. The geographical location of telescopes as shown in Fig. 2, ensures that targets can be observed nearly seamlessly if weather permits. The lower density of observations above Australia is a consequence of having only $2 \times 1\text{m}$ telescopes available as well as the characteristic weather pattern.

With an emphasis on robotic time-domain astronomy, gravitational microlensing perfectly fits into the observing program of LCO. Moreover, with its telescope clusters in Chile, South Africa and Australia, it achieves unprecedented coverage of events in the Galactic center. Consequently, microlensing was selected to be part of the 2013 pilot program of the 1m network and was awarded Key Project¹ status on the LCO network between 2014–2016. At the beginning of the 2013 pilot phase the multi-purpose scheduler was not available and the telescopes were initially equipped with SBIG STZ-16803 cameras². The Key Project was interrupted by the 2015 Spitzer campaign as described by Yee et al. (2015) for obtaining simultaneous ground

¹ PI: R.A. Street

² <http://LCO.net/network/instrumentation/1m-sbig-camera>

and space-based observations of microlensing events. This paper will focus on the pilot phase and the first year of operations with its lower number of competing projects, direct scheduling and corresponding camera systems.

1.3. Gravitational Microlensing

Gravitational microlensing is one of the well established methods of exoplanet detection, probing the Galactic planet population down to Earth-mass range with ground-based observations (Bennett & Rhie, 1996; Wambsganss, 1997). Light from a distant source star is gravitationally deflected by a foreground lens star, producing an increase of brightness of the source from the point of view of the observer. Planets orbiting faint (even unseen) lens stars can be discovered in this manner, because the light curve shape is related to the mass of the lens-system and depends on the mass ratio between planet and host star as well as their angular separation. Moreover, the planet population beyond the snow-line can be probed due to the increased sensitivity of this detection technique (Park et al., 2006). This is the main science driver of the aforementioned Key Project.

The sample of microlensing events and their absolute sensitivity to planets provides a statistical and independent measure of the planet abundance in our Galaxy (Snodgrass et al., 2004; Gaudi et al., 2002; Gould et al., 2010; Sumi, 2010; Sumi et al., 2011; Cassan et al., 2012). The detection of microlensing events itself is subject to statistics, as only ~ 1 in a million stars in the Galactic bulge is sufficiently aligned with a deflecting foreground star as to ensure that more light will be redirected to the observer as predicted by the seminal paper of Paczynski (1986). Both planet detections and non-detections provide important constraints on the frequency and distribution of planets in the Galaxy.

Most microlensing events follow a characteristic symmetric light curve which is a direct consequence of Einstein’s deflection angle. This deflection can be described by a simple lens equation, relating source position, image position and deflection angle. The solutions to the lens equation provide the image positions. The radius of the ring-like image in a co-linear lens-source-observer configuration defines the Einstein radius

$$\theta_E = \sqrt{\frac{4GM_L}{c^2} (D_L^{-1} - D_S^{-1})} \quad (1)$$

constraining the typical angular scale of the effect, where D_L denotes the distance from observer to deflecting lens of mass M_L and D_S denotes the distance to the source star. The total observable brightness increase is time dependent due to the relative source-lens proper motion. While multiple images can be detected and resolved for massive, extra-galactic source-lens configurations, *microlensing* on Galactic scales means that the image separation is of the order of ~ 1 mas and remains unresolved in the optical regime with current imaging technology. The total flux $F(t)$ of a microlensing event with a contribution of source flux F_S and blend flux F_B can be written as

$$F(t) = F_S A(u(t)) + F_B, \quad (2)$$

where the magnification A of a single lens is obtained by adding the differential magnifications of each image A_{\pm} which depends on the actual image positions in units of the angular Einstein radius

$$u_{\pm} = \frac{u \pm (u^2 + 4)^{1/2}}{2}. \quad (3)$$

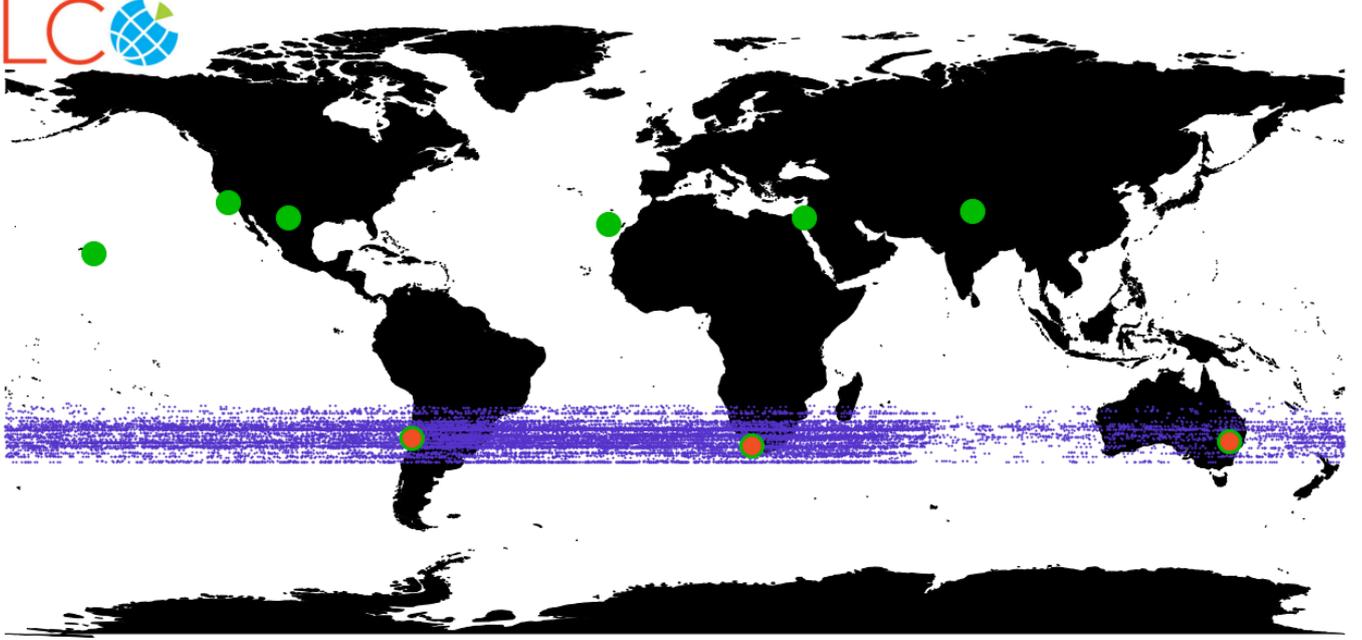


Fig. 2: Sites hosting LCO telescopes (green) and LCO sites used for the Key project (red) are shown along with observed microlensing target positions at zenith over Earth for the 2015 season (blue dots). The Australian site hosts only two telescopes which leads to less coverage, otherwise the event distribution would be homogeneous along all longitudes.

The image magnification for the minor image A_- and the major image A_+ is

$$A_{\pm} = \frac{u_{\pm}^2}{u_+^2 - u^2}. \quad (4)$$

For practical purposes, we use the blend ratio

$$g = \frac{F_B}{F_S}, \quad (5)$$

and express the lens position in units of the angular Einstein radius. The time-dependent source-lens distance denoted as $u = u(t)$ defines the shape of the point-source point-lens (PSPL) Paczyński light curve

$$A(u) = A_+ + A_- = \frac{u^2 + 2}{u\sqrt{u^2 + 4}} \quad (6)$$

in units of the angular Einstein radius. The latter expression can be derived analytically, because the lens equation of a single object can be solved for the image positions. The relative lens-source proper motion plays a role for the light curve shape but for events that do not exhibit parallax effect (Gould, 1992) a uniform motion along a straight line serves as reasonable approximation on time scales of several days or weeks with the source-lens distance

$$u = \sqrt{u_0^2 + \left(\frac{t - t_0}{t_E}\right)^2}. \quad (7)$$

The time of maximum magnification is denoted as t_0 , the Einstein-radius crossing time is t_E and the minimum separation of lens and source is u_0 . As an analytic model for all further considerations, Eq. 6 plays a crucial role for prioritizing events and detecting anomalies. In this context, we consider each event as

anomalous if it deviates from the standard point-source point-lens (PSPL) model as defined in Eq. 6 and is triggered by the anomaly detector algorithm³ described by Dominik et al. (2007).

The science driver of the LCO Key project “Exploring Cool Planets Beyond the Snowline” is to find more planets with orbits between 0.5 and 10 AU which corresponds to the orbital region where the temperature drops below the freezing point of water. Gravitational microlensing provides us with mass measurements and the projected orbital radii, thus populating the mass vs. orbital radius diagram. Microlensing is complementary to the transit and the radial-velocity methods, since it is sensitive to Earth mass planets beyond the water snow line.

2. Prioritizing microlensing events

2.1. Priority and planet detection zones

Microlensing follow-up programs can be adapted to different scientific objectives such as studying the brown-dwarf population (Street et al., 2013), providing mass measurements of distant stellar remnants (Mao et al., 2002; Wyrzykowski et al., 2016) or searching for planets beyond the habitable zone (Tsapras et al., 2014). Searching for planets beyond the snowline is the main motivation behind the RoboNet-II program for following-up microlensing events with telescopes from the Las Cumbres Observatory Global Telescope Network (LCO) and the Liverpool Telescope (Brown et al., 2013; Steele et al., 2004).

For the pilot phase of the 1m network in 2013 we combined the approach of prioritizing events based on planet probability estimates with network specific observational constraints (Horne et al., 2009; Dominik et al., 2010). This approach addresses the challenge of selecting the right targets out of hundreds of

³ in a typical scenario more than 5 clearly deviating points from at least 2 consecutive nights are required to confirm an anomaly.

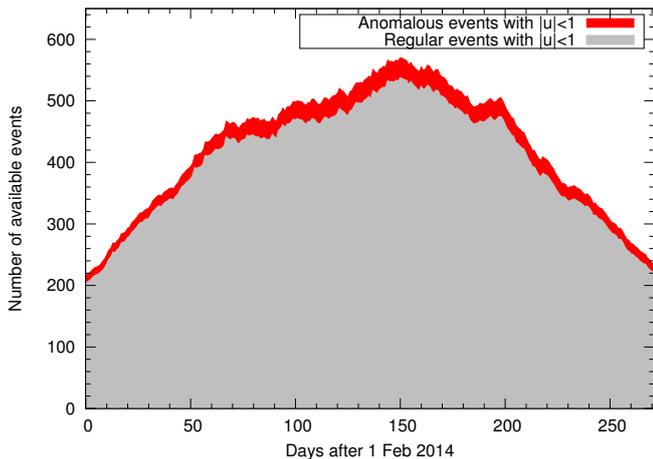


Fig. 3: Number of ongoing regular and anomalous events in the season 2014 before and after t_0 .

ongoing events. For reasons of simplicity, one often considers an event to be ongoing when the separation between lens and source is smaller than the Einstein radius and the source thus experiences a brightening by at least 0.3 mag with respect to baseline. In practice, we permit observations for source-lens distances $|u| < 1.5$ when an unblended source appears to be 0.1 mag brighter. Fig. 3 illustrates the number of available events at any time during the microlensing season (April to October). These 200-600 events include 5-10% anomalous events deviating from a simple static point lens and are reported by anomaly detectors. Examples of different anomaly types can be found in Tsapras et al. (2016). Not all anomalous events can be monitored continuously, and even with two-to-three follow-up telescopes per site we can only cover a handful of events with high enough cadence. A suitable prioritization is therefore required.

The idea of defining the planet detection probability of a given microlensing event via its detection zone goes back to the initial suggestions how likely it is that planets orbiting microlenses can be found (Gould & Loeb, 1992). A single-lens microlensing event creates a pair of images propagating through different parts of the lens plane where they can be perturbed by the presence of a planet. Estimating when such a perturbation becomes photometrically detectable gives the chance of finding such a planet. Fig. 4 illustrates the locations, where the extra planetary deflection is powerful enough to generate detectable deviations. The detection requires comparison with an underlying single-lens model, but the detection zone can shrink or grow depending on photometric accuracy. Therefore, one more commonly uses a $\Delta\chi^2$ threshold. For reporting a discovery one would require that observations lead to a change in $\Delta\chi^2 > 25$, corresponding to 5σ detections, if the underlying uncertainties are appropriately assigned and follow a Gaussian distribution. The complicated way how uncertainties are affected by crowded field photometry (Bachelet et al., 2015) frequently leads to a distribution of residuals which differs from that assumption and needs rescaling in the form of

$$\sigma_{\text{rescaled}}^2 = k(\sigma_0^2 + \sigma_{\text{reported}}^2), \quad (8)$$

where σ_{reported} refers to the original uncertainties and the parameters k, σ_0 are chosen so that the reduced χ^2 becomes one for each independent light curve. Since some light curves have uncertainties that can be off by a factor of 2, one commonly requires a threshold of $\Delta\chi^2 > 100$ for a detection.

In order to quickly calculate the planet detection probability Ψ , the exposure time t_{exp} is adjusted according to the expected brightness. We thereby ensure that a photometric accuracy of 1.5 to 5% is reached and a planetary perturbation is unlikely to be missed. One of the main changes between the empirical relations by Horne et al. (2009) and Dominik et al. (2010), which are both based on the planet detection zone area, is the underlying logarithmic distribution of planets along the orbital axis which is also consistent with observations (Cassan et al., 2012). The planet probability Ψ for major and minor image positions u_{\pm} and their corresponding magnification A_{\pm} is given by:

$$\Psi(u) = \frac{2A_+ - 1}{u_+^\gamma} + \frac{2A_-}{u_-^\gamma}, \quad (9)$$

where the parameter γ is chosen to be 2 for the assumed population of planets, which comes from the aforementioned empirical relations combined with the assumption of a logarithmically distributed orbital axis. For $\gamma = 2$ the detection probability can be simplified:

$$\Psi(u) = \frac{4}{u\sqrt{u^2 + 4}} - \frac{2}{u^2 + 2 + u\sqrt{u^2 + 4}}. \quad (10)$$

We consider the gain of observing an event to be given by the ratio of the planet detection probability and the invested observing time for a photon-noise dominated light curve. The priority function then reads

$$\Omega_S \propto \frac{A^{3/2}\Psi}{(A + F_B/F_S)(1 + F_B/F_S)}. \quad (11)$$

The latter expression makes implicitly the assumption that the suggested exposure time can be reached, the noise model is correct and that the underlying magnification A is sufficiently constrained by fitting a point-source point-lens model to the data. In the next section we will discuss to what extent these assumptions affect our results. One of the strategic differences between RoboNet and the strategy implemented in ARTEMiS is the absence of spending observing time for checking events when they slightly deviate from a point-source point-lens model. Based on estimates of earlier seasons we expect that the priority function alone will deliver and select high priority events that eventually will turn out to be anomalous - including those of potential planetary nature. Because of the two-step multi-tier strategy, it is more difficult to evaluate the sensitivity of our follow-up program. In order to ensure that anomalous events are not completely dominating our Key program, we have put a cap on the fraction of allocated anomalies.

To get a better understanding of where microlensing can find planets, we illustrate the detection zone in the right panel of Fig. 4. For each data-point a single point lens creates two images sweeping along the Einstein ring. Depending on the proximity of these images to a potential planet the light curve can be perturbed. For planets at various positions the corresponding deviation can be recorded and if it exceeds a certain threshold a planetary anomaly can be detected. The threshold depends on the achievable photometric accuracy as well as the desired number of deviating points. We assume that planets deviating by more than 1.5% (0.015 mag) can be detected if the exposure time can be adjusted accordingly.

For typical host stars ranging from 0.1 to 0.5 M_{\odot} we simulate the detection zone area for a mass ratio of $q \approx 10^{-4}$. This corresponds to planets in the range of 3 to 17 Earth masses covering super-Earths and mini-Neptunes. For typical point lens parameters ($u_0 = 0.15, t_E = 25 \text{ d}, \rho = 10^{-3} \theta_E$) the 1.5% detection zone

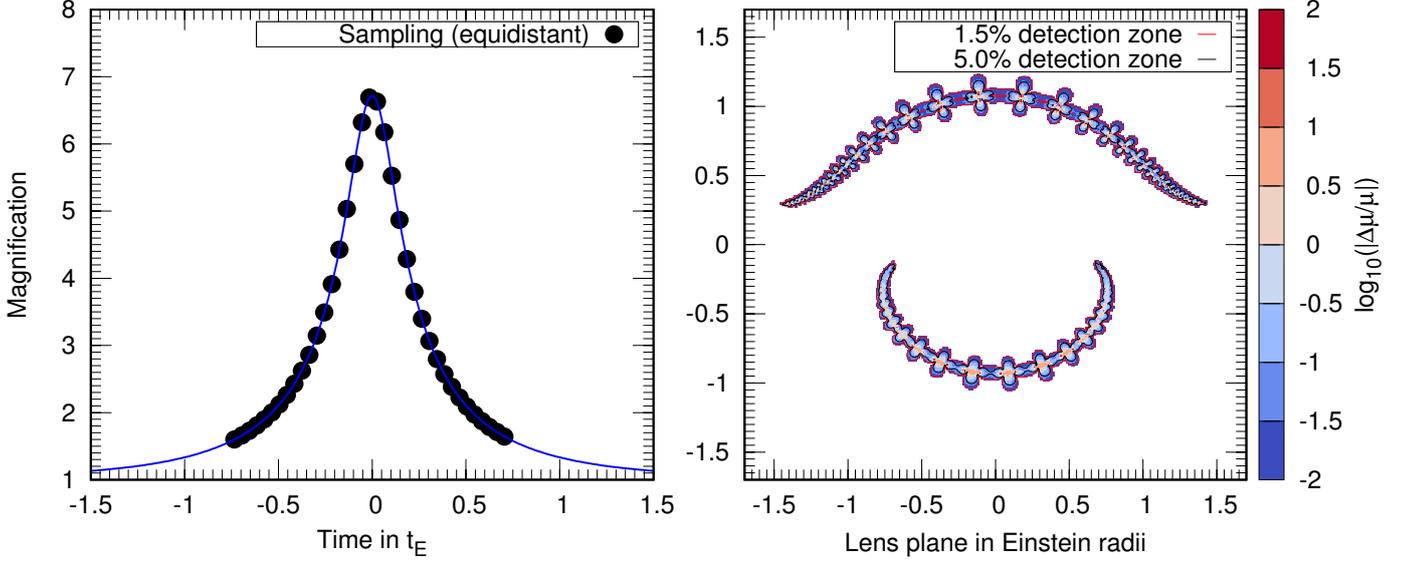


Fig. 4: Left: the light curve of a standard point lens light curve with a given equidistant sampling rate of 1 visit per night is shown. Right: the corresponding planet detection sensitivity of planets perturbing the light curve is shown for the lens plane assuming a low-mass planet with mass ratio $q = 10^{-4}$.

area shown in Fig. 4 also accounts for overlapping detection zones covering the same region. Finite source effects have been considered in the form of a uniform disk using the binary lens model by Bozza (2010)⁴. Most light curves can be simulated under the assumption that the source star is an infinitesimally small box and the distortion of that box is given by the inverse determinant of the lens equation. Whenever an image approaches a so-called critical line, the change of magnification is so dramatic that the surface of the source star is differentially magnified, which also ensures that the predicted brightness change remains finite.

Dominik (2009) hints at the unpredictability of events or in other words the difficulty of predicting the light curve days before the event reaches its maximum magnification. To successfully prioritize events, we expect that order relations

$$\Omega_{S,i} > \Omega_{S,j} \quad (12)$$

hold for each pair of consecutive events i, j . A heuristic estimate can be achieved assuming that the uncertainty of priority is limited by the uncertainty of u as Eq. 11 monotonically increases with u . If the peak has not been reached and the blend flux is low, it is safe to assume that $u_0 \in [0, F_{\text{now}}/F_S]$. The corresponding uniform distribution and its variance $\Delta u/12$ can be used to estimate the uncertainty in u and for the last reported flux F_{now} that provides an upper bound for u that gives the uncertainty

$$\sigma_u^2 \approx \frac{F_{\text{now}}^2}{12F_S^2}. \quad (13)$$

If multiple events are considered before reaching their peak, those with the highest magnification estimate have automatically the lowest uncertainty and as long as the available events are separated by a factor of $1/\sqrt{12}$ in u_0 the priority order is appropriate and that corresponds to the aforementioned variance of a uniform distribution. For events with $t > t_0$, u is much better

constrained. After observing the peak on a single occasion with a relative accuracy of $\delta_F = 1.5\%$ and by estimating that $A \approx 1/u$ and $A_{\text{true}} > A_{\text{peak}}$

$$\sigma_u^2 < \delta_F^2 u^2. \quad (14)$$

This can be interpreted by comparison with Eq. 13 for $F_{\text{now}} \approx F_{\text{peak}}$ where the uncertainty σ_u roughly follows $1/u$ before the peak and u after the peak.

Before the peak is characterized, it is difficult to optimize the exposure time regardless of our knowledge of g . After the peak, a higher underlying signal-to-noise ratio in flux units is required where events with $g \gg 10$ may end up with more than 1% noise in A .

2.2. Simulating a fiducial season

Before analyzing the outcome of real observations we shall provide an estimate of the expected number of events and achievable number of planets based on a simple model of basic microlensing parameters. For that purpose, we simulate ensembles of events with event parameters $t_E, u_0, l_{\text{baseline}}, g$ where samples are drawn using a Gaussian kernel density estimate of PSPL parameters obtained from post-season ARTEMiS fits of the 2012 season. We will show at a later point that the real-time ranking is sufficiently close to the post-season estimate. For each event we simulate a random time in the season and let the system determine its current priority state. Based on that simulation we determine the priority thresholds for the season. The cumulative distribution of the priority function as shown in Fig. 5 indicates that 1%, 2% and 1% of events would fall in our *low*, *medium* and *high* priority categories. Consequently, we would expect to request observations of 20, 40 and 20 non-anomalous events assuming that 2000 events are provided by survey teams.

In practice, the assumed number of underlying events turned out to be too low, because the revised simulation based on 2013 season parameters hints at an excess of low-priority targets. One explanation for the corresponding offset is a new discovery channel of the Optical Gravitational Lensing Experiment (OGLE-

⁴ <http://www.fisica.unisa.it/gravitationAstrophysics/VBBinaryLensing.htm>

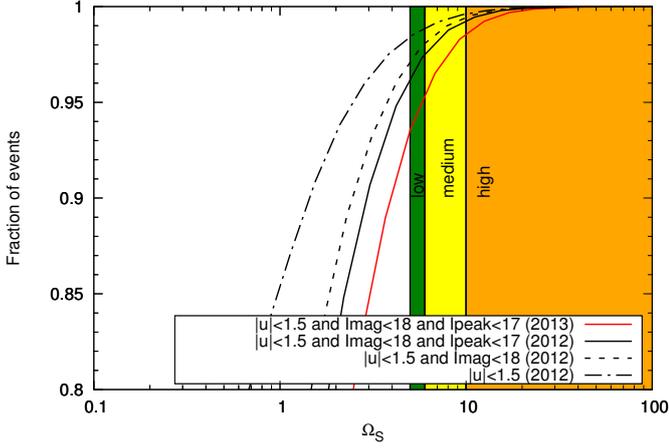


Fig. 5: Cumulative priorities based on simulated events matching the season 2012 are shown for the underlying observing constraints and after the season 2013 has finished.

IV) in its fourth phase (Udalski et al., 2015) reporting fainter microlensing events which has led to an increase of 15% in terms of discovered events, so effectively we should expect to see 40,70 and 30 events within our thresholds. Fainter source stars that are magnified beyond the limiting magnitude are by definition of higher magnification and thus more likely to contribute to events with higher priorities.

Based on the confirmed planets on the exoplanet archive⁵, we have seen 6 planetary events with 7 planets in the OGLE fields for 2012 and about 1700 stellar microlensing events in total. One event contains two planets (Han et al., 2013). With no assumption on the optimal choice for selecting events, we have a 1/5 chance to see a single planet when covering 50 events. Our choice of parameters and detection probability ψ can be calculated at the peak of each light curve which is an approximation of the planet detection sensitivity of each event. That gives us a relative number of how many more planets we expect to see.

Our strategy automatically incorporates that approach by assigning a priority level Ω_S . The planets shown are based on the parameter records and not on published values to give us an indication whether the system would have picked the target. While 99.9% of the peak planet detection probability is concentrated in 100 targets, two thirds of published planets instead belong to the 200 highest priority targets. We have limited that estimate to data from 2013 onwards, because older data would misrepresent the input from the anomaly detection and available survey data. Planets in the wings can still be detected although the planet detection probability is several orders of magnitude lower, if it is sampled in a way that prevents overlapping detection zone areas for many events.

2.3. Characterizing events

The most important physical parameter that is accessible through gravitational microlensing is the planet mass:

$$M_p = \frac{\mu t_E}{\kappa \pi_E} \frac{q}{1+q}; \kappa \approx 8.14 \frac{\text{mas}}{M_\odot}, \quad (15)$$

where μ denotes the lens-source relative proper motion, π_E is the magnitude of the microlensing parallax vector expressed in

units of the Einstein radius, $q = M_p/M_{\text{host}}$ is the mass ratio and κ a constant (Gould, 2000). Propagating the uncertainties and writing all uncertainty contributions relative to their quantities yields:

$$\sigma_{M_p} \approx \frac{\mu q t_E}{\kappa \pi_E (1+q)} \left(\frac{\sigma_{t_E}^2}{t_E^2} + \frac{\sigma_{\pi_E}^2}{\pi_E^2} + \frac{\sigma_\mu^2}{\mu^2} + \frac{\sigma_q^2}{q^2 (1+q)} \right)^{1/2}. \quad (16)$$

For small mass ratios $1/(1+q)$ is sufficiently close to one and thus the planet mass uncertainty is dominated by the term with the largest relative uncertainty. In the following, we will focus on the uncertainties of mass-ratio and timescale since these can be constrained better for most events. In order to quantify how the uncertainties in the fit-parameters can be reduced, we use the total information content expressed in the Fisher matrix:

$$I_{i,j}(t) = \left\langle \left(\frac{\partial \log \mathcal{L}}{\partial p_i} \right) \left(\frac{\partial \log \mathcal{L}}{\partial p_j} \right) \right\rangle, \quad (17)$$

which is the inverse of the covariance matrix depending on likelihood \mathcal{L} and estimated parameters p_i . The expected value of this expression can be used to calculate how each data point contributes to the event parameters extracted from a fiducial model fitted to the light curve. The Fisher matrix simplifies in the χ^2 -case for a given time t to:

$$I_{i,j}(t) = \frac{1}{\sigma_f^2} \left(\frac{\partial f_{\text{model}}}{\partial p_i} \right) \left(\frac{\partial f_{\text{model}}}{\partial p_j} \right) \quad (18)$$

Evaluating this expression requires only the best-fit model f_{model} and the derivatives with respect to the parameters p_i . The information content of the full light curve is the sum of all matrices $I_{i,j}$. The microlensing model is usually given in the form shown in Eq. 2 and thus depends on $F(A(u(p)))$. Horne et al. (2009) and Dominik et al. (2010) focus on finding planets rather than characterizing parameters. Fig. 6 shows areas associated with high probabilities of planet detection, for a typical detection threshold of 1.5% in photometric accuracy and mass ratios $q > 10^{-4}$.

The sensitivity levels shown in Fig. 6 can be understood as inverse parameter uncertainties. In case of the mass ratio q , it follows directly the detection zone area, while for the Einstein time one needs a better coverage in the wings. The exact location of information maxima is not known before the peak and based on the distribution of these maxima a sampling interval of at least $t_E/10$ is advisable. That holds for achieving about 10% accuracy in q , but there is an obvious caveat. The corresponding measurements need to be affected by the planet. For that purpose, our highest sampling interval is 15 min and follows $A^{1/2} t_E$ (Horne et al., 2009; Dominik et al., 2010). We would also like to highlight, that the azimuthal lobes in Fig. 4 overlap for the equidistant sampling. The radial detection zone lobes do not overlap to the same extent, which explains why the increased sampling rate at the peak is beneficial.

For low-mass planets we expect only a small perturbation in the light curve. Therefore, we estimate the contribution to the timescale based on a PSPL model. If the sampling rate would aim to maximize our understanding of t_E the sampling strategy needs to be adjusted. The sensitivity with respect to a parameter p , such as the Einstein time t_E , can be obtained from

$$\frac{\partial F_{\text{PSPL}}}{\partial t_E} = F_S \frac{\partial A}{\partial u} \frac{\partial u}{\partial t_E}, \quad (19)$$

⁵ <http://exoplanetarchive.ipac.caltech.edu>

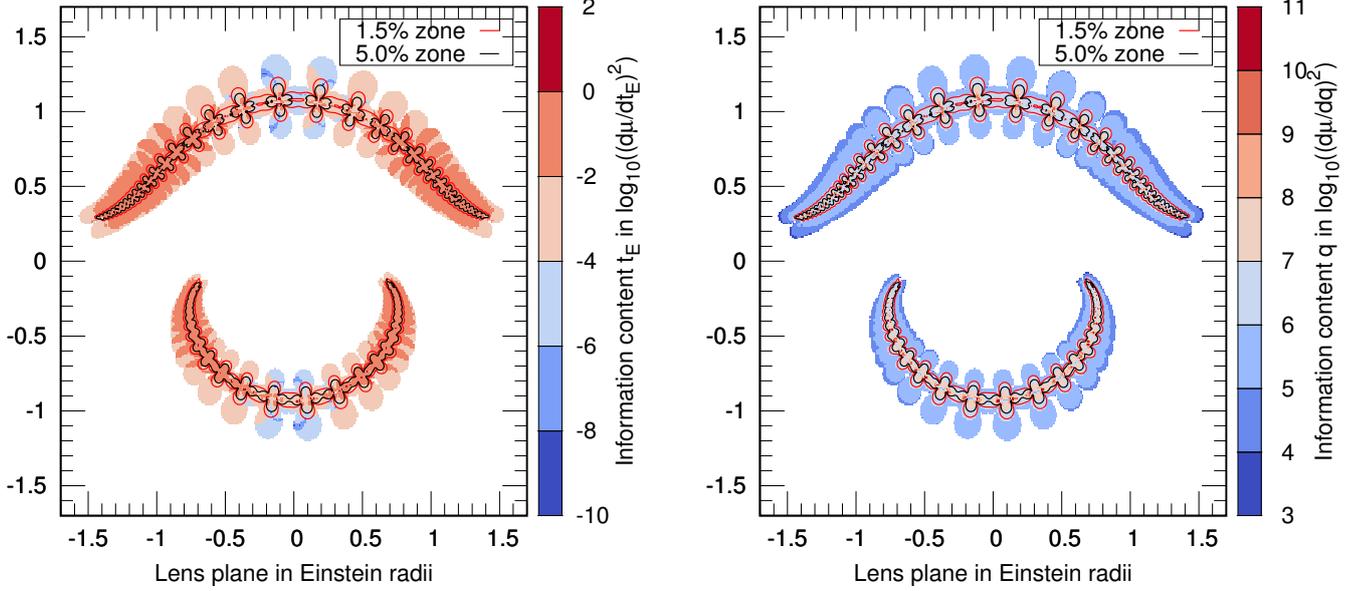


Fig. 6: The information content for characterizing t_E (left) and q (right) is shown for an unblended $q = 10^{-4}$ light curve with known parameters. Each of the color-coded points corresponds to a potential planet and its maximum contribution for the given sampling (cf. Fig. 4). The mass ratio is optimally characterized by exactly following the detection zone area, while the Einstein time t_E could benefit from more observations in the wings. For 1% photometric accuracy, the color-coded logarithmic sensitivity just needs to exceed 1 in order to ensure 10% accuracy in each parameter. Since information is additive, one can estimate the number of required points.

where

$$\frac{\partial A}{\partial u} = \frac{-8}{\sqrt{u^2 + 4(u^4 + 4u^2)}}, \quad (20)$$

and the inner derivative is

$$\frac{\partial u}{\partial t_E} = \frac{-(t - t_0)^2}{t_E^3 u}. \quad (21)$$

This implies that the sensitivity vanishes at $t = t_0$. The magnification A can be approximated as $A \approx 1 + (3u^2)^{-1}$ (Horne et al., 2009) which leads to a simple estimate of sensitivity:

$$\frac{\partial F_{\text{PSPL}}}{\partial t_E} = 2F_S \frac{(t - t_0)^2}{3t_E^3 u^4} \quad (22)$$

The latter expression has two implications for the sampling rate: to characterize t_E it should follow roughly $(t_E^3 u^4)^{-1}$ and the sampling interval T should be scaled with t_E . For the 2013 observing season, the requested sampling rate was kept independent of the Einstein time, but in 2014 t_E was used to intensify the sampling of short events. By using the weighted average of the a priori variance w_0 from the t_E distribution of all events and the inverse variance from the PSPL-fit $w_{t_E} = 1/\sigma_{t_E}^2$ we ensure that the sampling rate is not adjusted before the peak when the parameter uncertainty in t_E is high.

$$T \propto \sqrt{A} \frac{w_{t_E} t_E / \langle t_E \rangle + w_0}{w_{t_E} + w_0}, \quad (23)$$

3. Implemented strategy in 2013 and 2014

3.1. Observing constraints - pre-selection

In order to achieve a sensible pre-selection of events, we use empirical thresholds based on the achieved signal-to-noise ratio

during the start of the pilot phase:

$$t_{\text{exp}} \approx \exp(0.995 I_{\text{mag}} - 11.042) \text{ s}, \quad (24)$$

where the maximal exposure time was required to stay between 30 and 200 s. The resulting expected accuracy has lead us to request events with a baseline $I_{\text{mag}} < 18$ and a predicted peak $I_{\text{mag}} < 17$. In addition, very long $t_E > 400$ d events are excluded because they do not need to be observed at high cadence from follow-up observing teams. Most of the time we expect the telescope to be fully occupied with high priority targets and anomalies. If the event is well-past the peak, non-anomalous events with $u > 1.5$ are dropped from further consideration.

Independent of the constraints imposed by the system for automatic events, users can interact with the system through our online portal. As each site housed two or more telescopes, the target priority generator preferentially assigned events to telescopes that had already observed them. Additional observing resources on the 2m Faulkes telescopes could be activated by the person responsible for daily monitoring of operations. The exposure time and sampling rate for anomalous events was obtained from ARTEMiS (Dominik et al., 2008) and the SIGNALMEN anomaly detection algorithm, but exposure time estimates and sampling rates were adjusted for the 1 m network.

In order to assess how useful the suggested strategy is, we calculate the detection zone area as shown by Horne et al. (2009) and define the detection zone area as that area on the lens plane where a potential planet causes a deviation of at least 1.5% of magnification A . That implicitly assumes that we are able to adjust our exposure times accordingly. For that purpose, we use typical event parameters from 2012 ($u_0 = 0.15$, $t_E = 25$ d) and introduce finite source effects as $\rho = 10^{-3}$ Einstein radii for the binary model but not for the underlying point lens model.

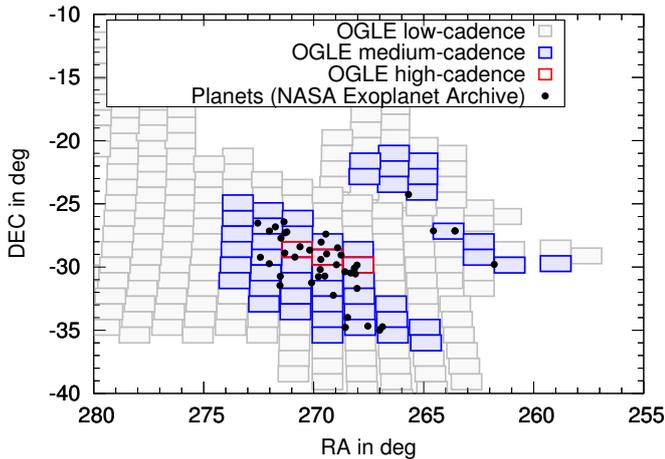


Fig. 7: The number and location of planets confirmed by 2017 (NASA Exoplanet Archive) is shown along with the survey cadence.

3.2. Precedence for low-cadence fields

Despite the increasing footprint of survey teams, there are patches on the sky that cannot be visited more than once or twice per night. Therefore, we observe events more or less intensively, i.e. events that are visited at least 3 times per night from OGLE are observed only once a night from the 1 m network to ensure early baseline coverage and reference frame. This rule only applies to regular events before reaching the peak and is our contribution to the microlensing community.

The priority flag requested by the network⁶ is set to *low* for events with $\Omega_s < 6$, to *medium* for $6 < \Omega_s < 10$ and to *high* for $\Omega_s > 10$. The latter limit supersedes the nominal cadence by OGLE and ensures that all events can be observed on a regular basis. Events with $\Omega_s < 5$ are not queued at all. These limits do not refer to the current priority but to the maximum priority before uploading the target list. In the pilot phase of 2013, the estimates were made for a response time of 12 and later 6 hours. Since 2014, target lists are updated hourly. In addition to the 1 m telescopes, RoboNet was awarded observing time on a robotic 2 m telescope network consisting of the Faulkes Telescope North and South and the Liverpool Telescope. RoboTAP automatically requested 2 m data in case of $\Omega_s > 30$ for 2013 and for a manual selection of the events in 2014.

The target list for each telescope or for the telescope network is populated after estimating the total available observing time on the network. Anomalous events are added until 30% of the maximum available observing time is used up or one anomaly is requested per telescope. The latter can have a huge impact at the beginning and at the end of the season, where only a single anomalous event could be requested. The sampling interval is independent of anomalous events and is set on an event by event basis, partly automatically relying on the sampling rate recommended by ARTEMiS. At later stages team members can adjust the sampling interval depending on the events exact shape of the light curve and the expected duration of the event. There are automatic and public fitting systems such as RTModel⁷ helping to assess the nature of the event but at early stages the true nature remains uncertain.

⁶ see <http://robonet.lco.global/temp/tap1mlist.html>

⁷ <http://www.fisica.unisa.it/GravitationAstrophysics/RTModel.htm>

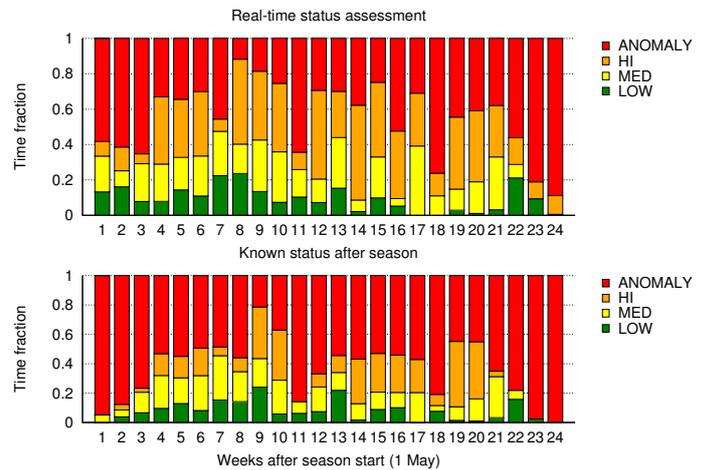


Fig. 8: Weekly fraction of events of each category based on the real-time assessment of the system (top panel) and the actual priority levels based on the parameters and the event status obtained at the end of the season.

The rest of the available time is populated with all other active events sorted by the aforementioned priority. In the pilot phase 2013 and at the beginning of the Key Project in 2014, there was not an immediate targets of opportunity (ToO) system available. For extremely magnified events or likely planets, continuous observations of a single target were requested on a single telescope, redirecting all other events to the remaining telescopes on a given site. Since 2014 we can request urgent observations as target of opportunity but due to the fast response of the scheduler (usually less than one hour) this is hardly ever necessary.

4. RoboTAP in action

RoboTAP generates target lists that are submitted to a target and observation manager code (ObsControl) which interfaces with the telescope network to request observations. In 2013 ObsControl read the telescope schedules directly per telescope and submitted microlensing observations into the gaps; once the network manager came online it “just” had to submit observation requests. The resulting requested observations were sent on a daily basis to the telescope network in 2013 and on an hourly basis in 2014. As soon as observations have been taken, the system informs the team about its current state. For the Key project the exact behavior of the scheduler changed slightly in the course of the project and no full-fledged simulator of the scheduler was available. Instead, we have monitored how the scheduling worked in practice and if our desired target success rate was matched.

The pilot phase season in 2013 lasted from 1st May until the end of September. During the season we logged the requested observing time by the target priority generator. Fig. 8 displays a weekly chart of how the real-time priority knowledge compares to the known final state of the events at the end of the season. As expected, some events that have been identified as regular events turn out to be anomalous, and particularly those of high priority. The actual fraction of anomalous events, including manually alerted ones differs a bit, because single anomalous events at the beginning and end of the season can override the fraction of anomaly-time we have envisaged. The agreement between the predicted priority level and the priority level for the final PSPL fit matches 82% of the requested observing time for regular events.

For anomalous events, a successful coverage with either anomalous or high-priority status was achieved in 89% of the time of anomalous events. These high success rates testify to the efficiency of the algorithm.

The field of view of follow-up telescopes constrains the number of targets that can be followed-up. In the microlensing pilot program up to 12 targets can be observed at the same time and on average 7 events were on our target lists. Thresholds for rejecting events have been chosen in order to match the expected available observing time. The decision-making was relying on PSPL fits using the SIGNALMEN anomaly detector (Dominik et al., 2007). For consistency, we have decided to focus our regular monitoring campaign on OGLE-IV⁸ events. The network approach enables us to follow more than one strategy and events from other survey teams are separately followed as targets of opportunity.

Fig. 9 shows the time per week spent on microlensing targets. The change in overall shape between the two seasons can be explained by the deployment of the network scheduler distributing observing time between different projects. Nevertheless, the same overall pattern is apparent in both years, with the largest fraction of observing time allocated to anomalous and high priority events.

5. Expected number of planet detections

To assess the number of planets that can be detected, we apply the following approach. Starting from the known PSPL model fit parameters for microlensing events in a given observing season, we employ the simplified Galactic model presented in Tsapras et al. (2016) to derive fiducial physical masses and circular orbital separations. The sample of planets drawn is calibrated to the orbital and mass distribution given in Cassan et al. (2012). Our results are still consistent when simulating orbits $a \in [0.1, 10]$ AU and planet masses $M_p \in [5, 3178] M_\oplus$ distributed uniformly in $\log a, \log M_p$.

Depending on the expected host star mass, we randomly assume that a multiple-star is present. The multiplicity fraction of stars and brown dwarfs is roughly approximated by a linear decrease in logarithmic mass from 80 % for a $10 M_\odot$ star to 20 % for a $0.1 M_\odot$ which is loosely motivated by (Lada, 2006; Raghavan et al., 2010). For a uniform hourly sampling of events and based on the LCO exposure time calculator we obtain for the longest useful exposure time of 300 s with the SBIG cameras 1.5 % for average conditions⁹. The actual uncertainty can be larger than that and we can still detect a signal but for uncertainties beyond 10 % effectively no planet can be detected any longer with the aforementioned criterion. As a side-remark, we are exclusively focusing on planets orbiting single-star hosts and exclude circumbinary planets (Bennett et al., 2016).

Instead of simulating complete binary-lens light curves, we inject planets into a point-source point-lens model and test those 7 points that have been observed closest to the anomaly to test if their cumulative $\Delta\chi^2 > 100$. In order to prevent single-outlier events, we request each observation to contribute by $\Delta\chi^2 > 100/7$. Requesting more points would exclude planets inducing only a short anomaly (Beaulieu et al., 2006) while discarding the threshold for each single observation would permit detecting events that are only weakly varying and difficult to distinguish from systematic or long-term effects such as annual parallax.

⁸ <http://ogle.astrouw.edu.pl/ogle4/ews/ews.html>

⁹ https://lco.global/files/etc/exposure_time_calculator.html

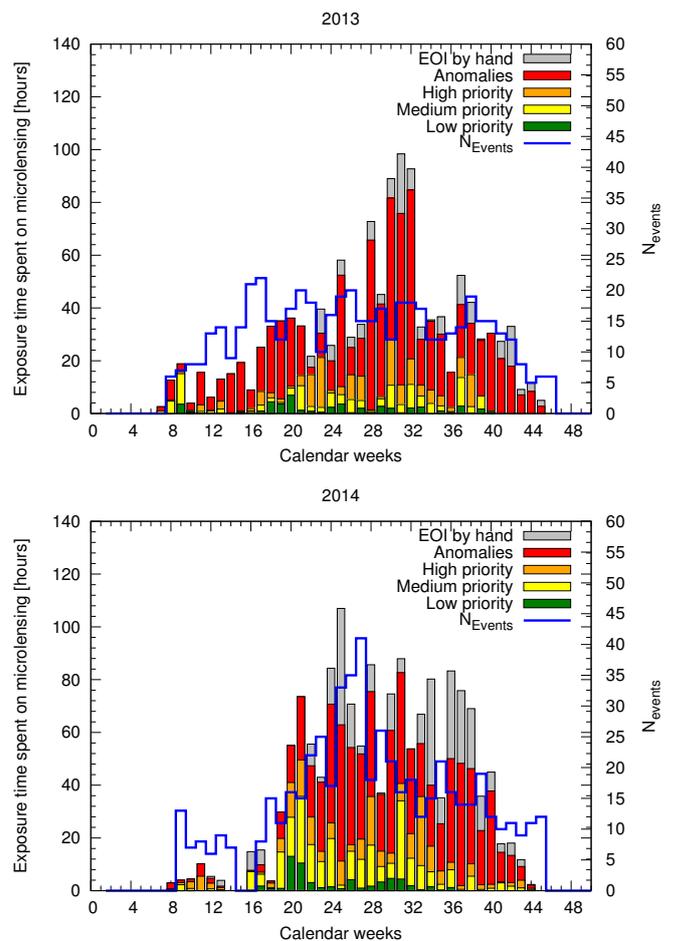


Fig. 9: Weekly observations with the LCO 1 m network of microlensing events in each category including manually requested observations, shortly denoted as events of interest (EOI), shown for the 2013 and the 2014 season. The continuous blue curve provides the number of followed events. The LCO scheduler came online since calendar week 18 in 2014.

If the expected annual sample of 2000 events could be used for our detections, we expect that 14 to 46 planets would be detectable orbiting single stars. The range of detectable planets corresponds to the 1σ uncertainty of Cassan et al. (2012) and is used for all subsequent estimates. Only single-star hosts were considered and for that purpose 36 % of all planets were discarded because they are orbiting potential multiple stars and further 8 % were discarded because their host mass was below $70 M_{\text{Jupiter}}$. As a side-remark, the observed fraction of identified binary star lenses is 2 % (Tsapras et al., 2016) and that means that one third of our perceived single-lens events are actually multiple stars. Our most important idealization is an uninterrupted hourly sampling of all 2000 events with constant exposure time so that the SBIG cameras installed on the LCO 1m-network reach 1.5 % photometric accuracy for an $I = 18$ mag star. Our initial priority threshold for Ω_S reduces the number of detected planets by 15 % while reducing the requested time by a factor of 10. The invested time still exceeds the scope of a Key project by a factor of 50.

All requirements together exceed the capabilities of existing observing teams. Estimating the yield for the anomaly-triggered part is more difficult. Our former estimates give a hint on the

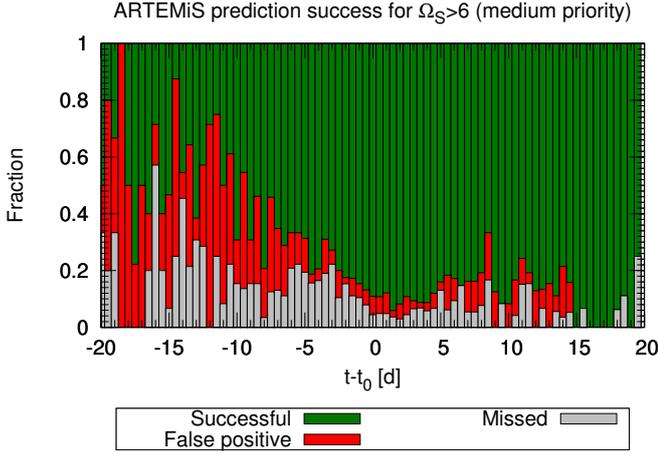


Fig. 10: Successful predictions, false positives and missed predictions for PSPL parameters are shown as a stacked histogram, i.e. the color-shaded areas correspond to the relative fractions.

fraction of constituents. Based on 10% non-planetary anomalies and 2% of anomalous events, we can expect to find a planet whenever 50 events have been sufficiently covered.

We are still missing the contribution of the actual sampling strategy. For that purpose, we simulate light curves with the full range of sampling intervals and assign 60% of our allocated time for monitoring. So far we have assumed that the parameters of our light curves are appropriately prioritizing events at any given time. In practice, we are affected by uncertain event parameters. For a typical season, we determine the number of missed and incorrectly assigned priority categories as shown in Fig. 10. The estimates allow on the event status at any given time compared to the final season parameter estimates. This leads us to a conservative estimate of 50% loss for events selected before the peak and roughly 20% after the peak. The expected planet yield decreases accordingly. In practice, the detection probability is higher in our selection of events and in the next section we will see how that supports our case.

6. Achieved planet sensitivity

From the observations taken in the seasons 2013 and 2014, we can estimate the expected number of planets from the non-anomalous part of our sample, using the end of season point-source point-lens parameter estimates. For the pilot phase 2013, we would have been sensitive to 2 to 8 planets (for the reported light curve uncertainties), while the sensitivity in the regular observing season 2014 was increased to 3 to 11 planets neglecting the aforementioned fraction of missed and overrated events. Both estimates assume the Cassan et al. (2012) PMF and our 7-point detection criterion.

We start by simulating planets uniformly in $\log q$, $\log s$ where the separation $s \in [0.1, 10]$ Einstein radii and $q \in [0.03, 10^{-5}]$ (Bond et al., 2004). A planet is detected whenever the resulting $\Delta\chi^2$ exceeds 100. Our reported photometric uncertainties are rescaled with one parameter so that $\chi_{\text{red}}^2 = 1$ because our light curves can come from up to 8 telescopes. In addition, we keep the finite source radius constant for the binary model ($\rho = 10^{-3}$) and compare it with a point-source point-lens model only. In a second step, we fit a point lens model to the simulated planetary light curve and check if the detection can be smoothed out and the $\Delta\chi^2$ can fall below 100. As a side remark, we have started the

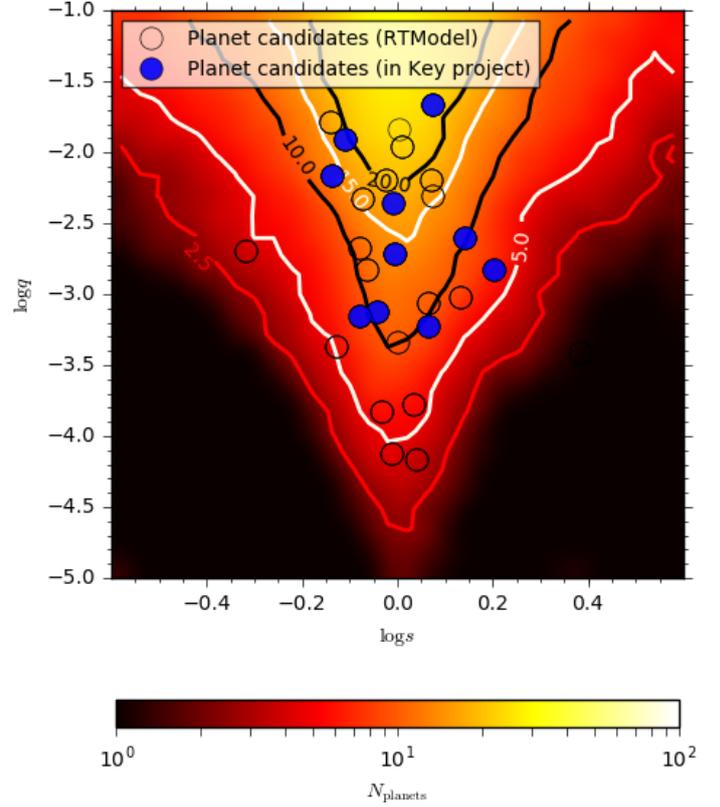


Fig. 11: Key project sensitivity illustrated as expected number of detections assuming each event has a planet in the range $\log s \in [-1, 1]$ Einstein radii and $\log q \in [\log 0.03, -5]$. To compare that with the number of detected planet candidates, we plot all planet candidates from real-time modeling efforts (RTModel) that have been reported 2013-2014 as part of the key project (filled) and all other planet candidates that were reported 2013-2016 (empty circles).

simulation using the ARTEMiS fit parameters which includes survey data. In the second step, we fit the model only to follow-up data. The detection is only accepted if it can be sustained by our follow-up data, but we do not claim that we can characterize the event without survey data. All accepted events are then divided by the number of samples per event and by that we obtain the expectation for the 118 regular events observed in the 2013-2014 season. Assuming each event contains one planet in the given range, we expect to find 4 in the season 2013 and 5 planets in the season 2014. Our earlier considerations are consistent as far as the expected planet yield is concerned. Defining an even more conservative threshold of $\Delta\chi^2 > 500$ as suggested by Yee et al. (2012) would reduce that by a factor of 1.4 and thus to 6 planets for both seasons combined.

In order to obtain an immediate impression how many planet candidates we have actually covered, we plot all events classified as planet candidates found on the RTModel web-page¹⁰ in Fig. 11. The number of covered events exceeds one planet per year, but half of the events were observed only after receiving an anomaly trigger which also roughly matches the time allocation. The achieved planet detections nearly matches our expectation for the underlying regular event distribution and is consistent with Cassan et al. (2012).

¹⁰ <http://www.fisica.unisa.it/GravitationAstrophysics/RTModel.htm>

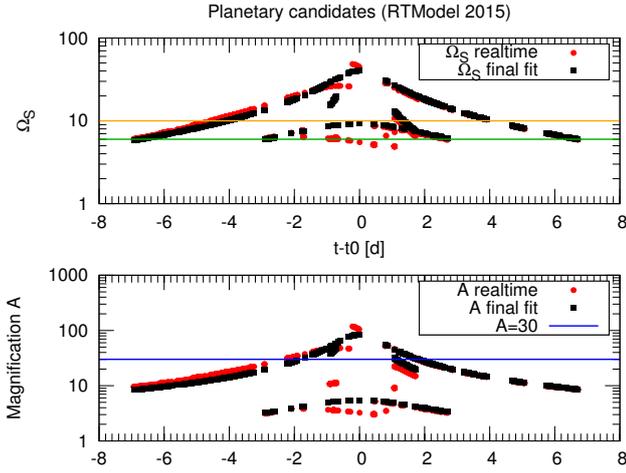


Fig. 12: The priority function and magnification is shown for those planet candidates that have been identified by RTModel (2015). High priority would correspond to a priority value above 10 and means that the event is observed.

In the season 2015, a non-negligible fraction of observing time was spent on following Spitzer targets. This gives us the opportunity to see how a season with a partly different sampling strategy behaves. In order to observe all Spitzer targets, the observing cadence was gradually reduced in order to secure points on all ongoing events. In principle, that should double our planet sensitivity, but in practice we see a similar number of covered planet candidates and one event with crucial data for characterizing the event (Street et al., 2016). The unusually good coverage on that event also enables us to reach our target mass ratio of 10^{-4} .

Independent of the exact sampling strategy required, we show in Fig. 12 how the actual assessment of the priority has changed over time and what the actual priority should have been based on the final parameter estimates. All targets shown in Fig. 12 have been reported online by the real-time modeling platform RTModel¹¹ and observed by RoboNet. The priority thresholds discussed and chosen in this work are evidently allocating more observing time to events with lower magnification (< 50).

7. Conclusions

We have reported how follow-up observations of gravitational microlensing events can be automatically requested by our RoboTAP system. In that context, RoboTAP acts as an example of a broker between an incoming event stream and a robotic telescope network such as the LCO 1 m network. The 2013 pilot phase provided some valuable and well-tested insights into the way targets can be selected. One of the insights shown in Fig. 13 is our large contribution of observing time to help cover regular (PSPL-like) events in low-cadence fields, which are fields observed less than once per night by survey teams. Similar considerations will be required for all-sky surveys and their varying field cadence.

In the broader context of observing transient events triggered by future survey telescopes, like the ZTF and the LSST, we have shown that respecting the survey cadence is very useful as a selector: it helps to avoid regions where surveys achieve sufficient

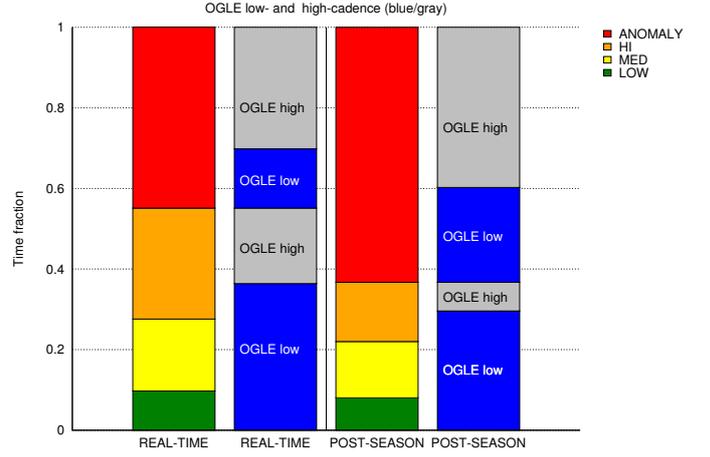


Fig. 13: Time fraction of different priority levels for regular and anomalous events based on the observing logs (left) and recalculated for parameters at the end-of-season (right). The figure highlights how many anomalous events have been identified as regular, high priority targets. It also suggests that events from high-cadence survey fields are more likely to be assessed as an anomalous event.

coverage for the respective science goal. At the same time targets should not be selected when there is only little chance to trigger an anomaly in due time. In retrospect, 80% of our coverage went to regular targets. It is also noteworthy that regular events in the high-cadence zone in the real-time assessment more often turned out to be anomalous events and thus half of all observed events exhibited some form of anomaly.

Restricting follow-up observations to microlensing fields that survey teams can only visit once a night is not a viable option because event parameters and anomaly triggers remain elusive. For the fields defined in this work, the number of potential planetary events is reduced by a factor of 2. Two factors contribute to that: larger event uncertainties, and thus inappropriately low or high priority levels, and the differing blend ratio. That former is also an indicator for the number of available source stars.

The anomaly detector requires multiple data points from two consecutive nights to deviate. Whenever surveys achieve only 1-3 visits per night, it is beneficial to follow anomalies from all-fields. The overhead of low-cadence follow-up for the initial phase of the Key project is shown in Fig. 14. Starting in 2015, we have changed the interpretation of what belongs to each cadence category. Only 3 fields with the very highest coverage by OGLE are subject to a reduced cadence during regular follow-up because we have concluded that OGLE and other survey teams will be able to characterize these targets by themselves.

Selecting microlensing events with high planet detection probability automatically ensures that the potentially limiting parameters for the planet mass (q, t_E) are sufficiently characterized. This is another way of expressing that some change of brightness that permits us to find a planet also contributes to the characterization of its parameters. However, a better coverage in the wings is the preferable choice for the Einstein time. For future observing programs we suggest to study in more detail how to routinely include estimates of the expected achievable parallax signal in the selection process in a way that avoids a selection bias tending towards long duration events. Nevertheless,

¹¹ <http://www.fisica.unisa.it/GravitationAstrophysics/RTModel/2015/RTModel.htm>

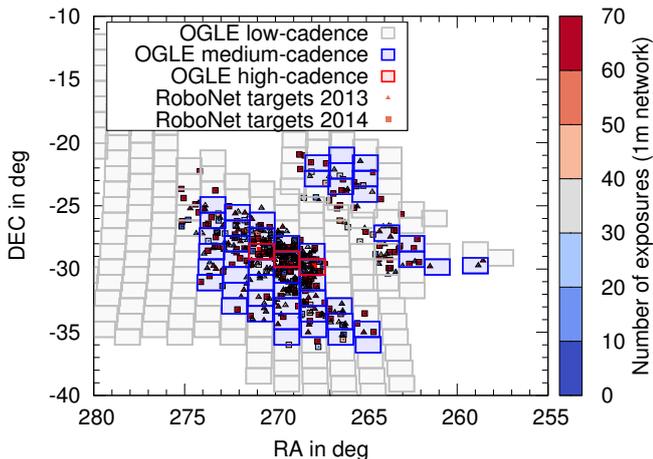


Fig. 14: The distribution of observed events and their number of visits during the RoboNet campaign 2013-2014 is shown along with high-, medium- and low-coverage zones of the OGLE-IV survey.

with current technology only half of the planets will benefit from that. In all other cases, we expect our approach to be appropriate.

When comparing the requested and observed fractions of all observing requests submitted to the telescope scheduler, we find that in 64% of all cases the requested next observation was either immediately after the requested sampling interval or at most 15 min later. That would be sufficient for detecting low-mass planets, but is not representative of present and future systems due to modifications in the target scheduling and a larger number of parallel projects. Our Key project sensitivity suggests that low-mass planets with $q < 10^{-4}$ should have been detected, but actual candidates are above a mass-ratio of $q > 10^{-3.5}$. Achieving a higher cadence, a higher photometric accuracy and covering more events at the same time is usually hard to accomplish. The Spitzer campaign in 2015 was scientifically fruitful, not least because of the planet discovery of OGLE-2015-BLG-0966 (Street et al., 2016). At the same time our nominal chance to detect planets was larger while the actual number of planet detections was reduced.

As far as finding and characterizing cool planets is concerned, we expected to detect ~ 5 planets per year. This estimate was consistent with the planet candidates identified in real-time modeling. Half of these candidates were identified in events with anomaly-triggered observations, while the other half were found in regularly monitored events. Further refinement of the results reported here will come from analyzing the data obtained by the next microlensing Key Project on the LCO network (2017-2020).

Acknowledgements. This publication was made possible by NPRP grant # X-019-1-006 from the Qatar National Research Fund (a member of Qatar Foundation). Work by SM has been supported by the Strategic Priority Research Program “The Emergence of Cosmological Structures” of the Chinese Academy of Sciences Grant No. XDB09000000, and by the National Natural Science Foundation of China (NSFC) under grant numbers 11333003 and 11390372. Clément Ranc’s research was supported by an appointment to the NASA Postdoctoral Program at the NASA Goddard Space Flight Center, administered by Universities Space Research Association under contract with NASA. KH acknowledges support from STFC grant ST/M001296/1. This work makes use of observations from the LCO network, which includes three SUPAScopes owned by the University of St Andrews. The RoboNet programme is an LCO Key Project using time allocations from the University of St Andrews, LCO and the University of Heidelberg together with time on the Liverpool Telescope through the Science and Technology Facilities Council (STFC), UK. This research has

made use of the NASA Exoplanet Archive, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program.

We would like to thank the anonymous referee for his/her helpful comments and suggestions. Finally, we would like to thank all teams - and in particular the survey teams OGLE and MOA - for making their real-time data public.

References

- Bachelet, E., Bramich, D. M., Han, C., et al. 2015, *ApJ*, 812, 136
 Bellm, E. 2014, *The Third Hot-wiring the Transient Universe Workshop*, 27
 Bennett, D. P., & Rhie, S. H. 1996, *ApJ*, 472, 660
 Bond, I. A., Udalski, A., Jaroszyński, M., et al. 2004, *ApJ*, 606, L155
 Bozza V., 2010, *MNRAS*, 408, 2188
 Bramich, D. M. 2008, *MNRAS*, 386, L77
 Beaulieu, J.-P., Bennett, D. P., Fouqué, P., et al. 2006, *Nature*, 439, 437
 Bennett, D. P., Rhie, S. H., Udalski, A., et al. 2016, *AJ*, 152, 125
 Brown, T. M., Baliber, N., Bianco, F. B., et al. 2013, arXiv:1305.2437
 Burgdorf, M. J., Bramich, D. M., Dominik, M., et al. 2007, *Planet. Space Sci.*, 55, 582
 Cassan, A., Kubas, D., Beaulieu, J.-P., et al. 2012, *Nature*, 481, 167
 Dominik, M., Rattenbury, N. J., Allan, A., et al. 2007, *MNRAS*, 380, 792
 Dominik, M., Horne, K., Allan, A., et al. 2008, *Astronomische Nachrichten*, 329, 248
 Dominik, M. 2009, *MNRAS*, 393, 816
 Dominik, M., Jørgensen, U. G., Rattenbury, N. J., et al. 2010, *Astronomische Nachrichten*, 331, 671
 Gaudi, B. S., Albrow, M. D., An, J., et al. 2002, *ApJ*, 566, 463
 Gould, A. 1992, *ApJ*, 392, 442
 Gould, A., & Loeb, A. 1992, *ApJ*, 396, 104
 Gould, A. 2000, *ApJ*, 542, 785
 Gould, A., Dong, S., Gaudi, B. S., et al. 2010, *ApJ*, 720, 1073
 Han, C., Udalski, A., Choi, J.-Y., et al. 2013, *ApJ*, 762, L28
 Henderson, C. B., Gaudi, B. S., Han, C., et al. 2014, *ApJ*, 794, 52
 Horne, K., Snodgrass, C., & Tsapras, Y. 2009, *MNRAS*, 396, 2087
 Kerins, E., Robin, A. C., & Marshall, D. J. 2009, *MNRAS*, 396, 1202
 Lada, C. J. 2006, *ApJ*, 640, L63
 Law, N. M., Fors, O., Ratzloff, J., et al. 2015, *PASP*, 127, 234
 LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201
 Mao, S., Smith, M. C., Woźniak, P., et al. 2002, *MNRAS*, 329, 349
 Paczyński, B. 1986, *ApJ*, 304, 1
 Park B.-G., Jeon Y.-B., Lee C.-U., Han C., 2006, *ApJ*, 643, 1233
 Raghavan, D., McAlister, H. A., Henry, T. J., et al. 2010, *ApJS*, 190, 1
 Ridgway, S. T., Matheson, T., Mighell, K. J., Olsen, K. A., & Howell, S. B. 2014, *ApJ*, 796, 53
 Smith, R. M., Dekany, R. G., Bebek, C., et al. 2014, *Proc. SPIE*, 9147, 914779
 Saha, A., Wang, Z., Matheson, T., et al. 2016, *Proc. SPIE*, 9910, 99100F
 Snodgrass, C., Horne, K., & Tsapras, Y. 2004, *MNRAS*, 351, 967
 Snodgrass, C., Tsapras, Y., Street, R., et al. 2008, *Manchester Microlensing Conference*, 56
 Shvartzvald, Y., & Maoz, D. 2012, *MNRAS*, 419, 3631
 Soszyński, I., Udalski, A., Poleski, R., et al. 2012, *Acta Astron.*, 62, 219
 Steele, I. A., Smith, R. J., Rees, P. C., et al. 2004, *Proc. SPIE*, 5489, 679
 Street, R. A., Tsapras, Y., Horne, K., et al. 2011, *IAU Symposium*, 276, 459
 Street, R. A., Choi, J.-Y., Tsapras, Y., et al. 2013, *ApJ*, 763, 67
 Street, R. A., Udalski, A., Calchi Novati, S., et al. 2016, *ApJ*, 819, 93
 Sumi, T., Woźniak, P. R., Udalski, A., et al. 2006, *ApJ*, 636, 240
 Sumi, T. 2010, *Pathways Towards Habitable Planets*, 430, 225
 Sumi, T., Kamiya, K., Bennett, D. P., et al. 2011, *Nature*, 473, 349
 Tsapras, Y., Street, R., Horne, K., et al. 2009, *Astronomische Nachrichten*, 330, 4
 Tsapras, Y., Choi, J.-Y., Street, R. A., et al. 2014, *ApJ*, 782, 48
 Tsapras, Y., Hundertmark, M., Wyrzykowski, Ł., et al. 2016, *MNRAS*, 457, 1320
 Udalski, A., Szymański, M., Kaluzny, J., et al. 1994, *Acta Astron.*, 44, 227
 Udalski, A., Szymański, M. K., & Szymański, G. 2015, *Acta Astron.*, 65, 1
 Wambsganss, J. 1997, *MNRAS*, 284, 172
 Wyrzykowski, Ł., Kostrzewa-Rutkowska, Z., Skowron, J., et al. 2016, *MNRAS*, 458, 3012
 Yee, J. C., Shvartzvald, Y., Gal-Yam, A., et al. 2012, *ApJ*, 755, 102
 Yee, J. C., Gould, A., Beichman, C., et al. 2015, *ApJ*, 810, 155

-
- ¹ Las Cumbres Observatory Global Telescope Network, 6740 Cortona Drive, suite 102, Goleta, CA 93117, USA
 - ² European Southern Observatory, Karl-Schwarzschild-Str. 2, 85748 Garching bei München, Germany
 - ³ SUPA, School of Physics & Astronomy, University of St Andrews, North Haugh, St Andrews KY16 9SS, UK
 - ⁴ Astrophysics Research Institute, Liverpool John Moores University, Liverpool CH41 1LD, UK
 - ⁵ South African Astronomical Observatory, PO Box 9, Observatory 7935, South Africa
 - ⁶ Astronomisches Rechen-Institut, Zentrum für Astronomie der Universität Heidelberg (ZAH), 69120 Heidelberg, Germany
 - ⁷ Dipartimento di Fisica "E.R. Caianiello", Università di Salerno, Via Giovanni Paolo II 132, 84084-Fisciano (SA), Italy
 - ⁸ Istituto Nazionale di Fisica Nucleare, Sezione di Napoli, Via Cintia, 80126 Napoli, Italy
 - ⁹ Planetary and Space Sciences, School of Physical Sciences, The Open University, Milton Keynes, MK7 6AA, UK
 - ¹⁰ Sorbonne Universités, UPMC Univ Paris 6 et CNRS, UMR 7095, Institut d'Astrophysique de Paris, 98 bis bd Arago, 75014 Paris, France
 - ¹¹ National Astronomical Observatories, Chinese Academy of Sciences, 100012 Beijing, China
 - ¹² Space Telescope Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA
 - ¹³ School of Physics and Astronomy, Tel Aviv University, Tel Aviv 69978, Israel
 - ¹⁴ Jet Propulsion Laboratory, M/S 169-506, 4800 Oak Grove Drive, Pasadena, CA 91109
 - ¹⁵ NASA Postdoctoral Program Fellow, NASA Goddard Space Flight Center, Mail Code 661, Greenbelt, MD 20771, USA
 - ¹⁶ INAF - Observatory of Capodimonte, Salita Moiariello, 16, 80131, Naples, Italy
 - ¹⁷ Department of Astronomy, Yale University, 52 Hillhouse Avenue, New Haven, CT 06511, USA