



The Zwicky Transient Facility: Surveys and Scheduler

Eric C. Bellm¹, Shrinivas R. Kulkarni², Tom Barlow², Ulrich Feindt³, Matthew J. Graham², Ariel Goobar³,
Thomas Kupfer^{2,4,5}, Chow-Choong Ngeow⁶, Peter Nugent^{7,8}, Eran Ofek⁹, Thomas A. Prince², Reed Riddle¹⁰,
Richard Walters¹⁰, and Quan-Zhi Ye (葉泉志)^{2,11}

¹ DIRAC Institute, Department of Astronomy, University of Washington, 3910 15th Avenue NE, Seattle, WA 98195, USA; ecbellm@uw.edu

² Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

³ The Oskar Klein Centre, Department of Physics, Stockholm University, AlbaNova, SE-106 91 Stockholm, Sweden

⁴ Kavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, USA

⁵ Department of Physics, University of California, Santa Barbara, CA 93106, USA

⁶ Institute of Astronomy, National Central University, 32001, Taiwan

⁷ Computational Science Department, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, MS 50B-4206, Berkeley, CA 94720, USA

⁸ Department of Astronomy, University of California, Berkeley, CA 94720-3411, USA

⁹ Benoziyo Center for Astrophysics and the Helen Kimmel Center for Planetary Science, Weizmann Institute of Science, 76100 Rehovot, Israel

¹⁰ Caltech Optical Observatories, California Institute of Technology, Pasadena, CA 91125, USA

¹¹ IPAC, California Institute of Technology, 1200 E. California Boulevard, Pasadena, CA 91125, USA

Received 2019 January 2; accepted 2019 March 2; published 2019 April 29

Abstract

We present a novel algorithm for scheduling the observations of time-domain imaging surveys. Our integer linear programming approach optimizes an observing plan for an entire night by assigning targets to temporal blocks, enabling strict control of the number of exposures obtained per field and minimizing filter changes. A subsequent optimization step minimizes slew times between each observation. Our optimization metric self-consistently weights contributions from time-varying airmass, seeing, and sky brightness to maximize the transient discovery rate. We describe the implementation of this algorithm on the surveys of the Zwicky Transient Facility and present its on-sky performance.

Key words: methods: observational – surveys

Online material: color figures

1. Introduction

Astronomers observe with telescopes costing millions or even billions of dollars that have finite useful lifetimes. They must accordingly decide how best to sequence observations in order to maximize the scientific output of their facilities. Despite its ubiquity, however, this scheduling problem remains challenging for both theoretical and practical reasons. With hundreds to many thousands of observations to obtain in a night or a season, the potential number of observing sequences is combinatorically large: one thousand exposures have $\sim 10^{2567}$ possible orderings. The need to slew between targets couples distinct observations together, so maintaining efficiency requires scheduling many targets at once. Observing conditions on the ground change rapidly, and requests for target of opportunity (TOO) observations can upend a carefully tuned schedule in an instant. The quality of a potential observation may vary with time (e.g., with seeing, airmass, or moon phase), and many facilities must impose complex pointing or instrument constraints. Time-domain surveys may require complex observing sequences that make future observations dependent on when past observations occurred, which is further

complicated by weather losses and other downtime. And finally it is often both difficult and impolitic to be quantitatively precise about how to measure scientific output.

Accordingly, direct, manual sequencing of observations by humans remains common at both ground- and space-based facilities. Skilled operations staff can perform complex heuristic tradeoffs to obtain observations that are “good enough” while meeting the necessary constraints. Satisficing in this manner may in some cases be the most efficient use of the human resources available, given the difficulties of developing more automated approaches. However, manual scheduling presents significant drawbacks. It is labor intensive, requiring constant staffing throughout the operation of the project. Any change to the schedule requires manual intervention, removing the ability to respond dynamically to changing observing conditions, weather losses, and TOOs. Typically the scheduling process is not optimizing in any quantitative sense, limiting clarity about its effectiveness. Because schedule generation is labor intensive, it is difficult to compare different observing plans. It is difficult to reproduce manual scheduling outcomes, which can inhibit studies of survey detection rates and efficiencies. Finally, manual scheduling provides limited

transparency to the user community about how their resources are being allocated.

“Greedy” algorithms provide a convenient entry point into automated scheduling and are widely used in astronomy. Before each observation, such algorithms compute a updated metric or score for each possible target, select the target with the current highest value, observe it, and then repeat the process. Greedy optimizers are straightforward to implement, can easily handle changes to observing plans and conditions, provide traceable quantitative optimization, and can be run in an automated fashion. The `Astroplan` package (Morris et al. 2018a) implements one such greedy scheduler. It is designed for human observers and implements a range of observational constraints.

However, it is widely recognized that local optimizers such as the greedy algorithm cannot deliver global optimization. For ground-based imaging surveys this problem manifests itself in a tendency to observe higher elevation fields as they rise instead of fields transiting at lower elevation at the same time. Some authors have recommended additional weighting schemes that account for the time until a target sets in order to penalize this behavior (e.g., Denny 2004, 2006), but this is not the same as looking ahead to determine the optimal time to observe a given field. Lookahead is especially valuable for time-domain surveys, which must understand how repeated observations of a target can be scheduled within the planning interval.

Several projects have implemented more sophisticated schedulers (see Solar et al. 2016, for a review). Of particular note is the scheduling approach of the Las Cumbres Observatory (LCO).¹² LCO operates a global network of replicated 0.6, 1.0, and 2.0 m telescopes with identical imagers and spectrographs. Rather than manually directing their observations to a specific telescope, LCO users make requests to the entire network, leaving to the scheduler the task of determining which facility to use to observe a target. Lampoudi et al. (2015) describe the scheduling algorithm, which uses integer linear programming (ILP¹³) techniques to assign requested observations to telescopes subject to any observability or cadence constraints. The LCO scheduler optimizes the assignments in order to maximize the total number of observations obtained, weighted by the priority assigned to them by the Time Allocation Committee (TAC). Notably, the scheduler’s ability to rapidly re-solve the entire network within minutes allows rapid TOO observations to be integrated into the scheduling process without disruption, as each new optimization run starts *de novo*, integrating any new targets that have arrived in the meantime.

¹² <https://lco.global/>

¹³ ILP problems have variables which take only discrete integer values, linear objective functions, and linear constraints. Mixed ILP problems include some non-discrete variables.

Solar et al. (2016) presents a similar Mixed ILP solution in the scheduler for ALMA. The ALMA scheduler discretizes time into scheduling blocks and assigns observations to them in order to maximize TAC-assigned scientific priorities, program completion, and telescope utilization. However, this scheme makes scheduling some types of observations relevant for time-domain followup challenging (Alexander et al. 2017).

Finally, Naghib et al. (2019) casts the scheduling problem of the Large Synoptic Survey Telescope (LSST) as a memoryless Markov Decision Process, using hand-designed features to reduce the dimensionality of the state space and optimizing the feature weights with a simple throughput-based objective function.

In this work we consider the specific scheduling problem of a single-telescope ground-based wide-field imaging survey. We are focused on its application to the Zwicky Transient Facility (ZTF; Bellm et al. 2019; Graham et al. 2019) project, which imposes some specific requirements (Section 2), but our formalism is relevant for other time-domain surveys, such as those conducted with the LSST (Ivezić et al. 2019), the Dark Energy Camera (Flaugher et al. 2015), and Hyper Suprime-cam (Miyazaki et al. 2018). Minor modifications would enable its use by multi-telescope surveys such as the Asteroid Terrestrial-impact Last Alert System (Tonry et al. 2018), PanSTARRS (Kaiser et al. 2010), the All-sky Automated Survey for Supernovae (Shappee et al. 2014), and BlackGEM (Bloemen et al. 2016).

Simply stated, the scheduling problem to be solved is to determine which fields to observe in what order, with a goal of maximizing an objective function (Section 3; here, a proxy for the transient discovery rate) while achieving the desired temporal spacing of observations (“cadence”). Optimizing the survey schedule provides a greater quantity of high-quality data, increasing the scientific output of the survey. During the development of the ZTF survey camera and observing system (R. Dekany et al. 2019, in preparation), the engineering team devoted substantial effort to developing percent-scale improvements in throughput and efficiency. Preserving these gains requires similar attention to the operation of the survey itself. Our approach provides a self-consistent means of scheduling an entire night of ZTF observations.

In this paper we outline the scheduling approach used by ZTF and its application to the surveys undertaken during the early operations period. In Section 2 we outline the requirements we used to guide our scheduler development. Section 3 describes the scalar survey speed metric we optimize for and discuss its applicability to other surveys and optimization approaches. Section 4 presents the ILP formalism we use to optimize ZTF observations for an entire night. In Section 5 we describe the practical implementation the algorithm in the ZTF scheduler. Section 6 details the surveys executed by ZTF in its first year of on-sky operations. Section 7 assesses the performance of the scheduler. We conclude in Section 8.

2. ZTF Requirements

The requirements for the ZTF scheduler grew from the experience of the Palomar Transient Factory (PTF; Law et al. 2009) and Intermediate Palomar Transient Factory (iPTF) surveys. PTF used a greedy scheduler. Its objective function is described in Law et al. (2009); it includes ad-hoc weightings for Sun altitude, sky brightness excess due to the moon, moon phase, telescope and dome slews, airmass, and a cadence term. In practice it proved unpredictable and prone to long slews. Operations staff frequently applied manual weights to ensure fields were observed.

For iPTF, a single member of the operations staff scheduled each night manually. This procedure reduced the number of long slews and (in conjunction with other technical improvements) increased the overall number of exposures taken.

During ZTF development, we began evaluating scheduling approaches in conjunction with other efforts at maximizing survey efficiency (R. Dekany et al. 2019, in preparation). The vastly improved readout speed of ZTF (8 s) relative to PTF/iPTF (40 s) made limiting scheduling overheads a higher priority. Additionally, the large number of simultaneous survey programs (Section 6), some of them public, also necessitated the ability to simulate and test survey plans.

Specific requirements imposed on the ZTF scheduler included the following:

1. Select pointings from a fixed field grid (see Masci et al. 2019).
2. Operate in both simulation mode and on-sky using the same scheduling code.
3. Conduct several surveys (Section 6), maintaining strict independence of their observations and balancing observing time between programs.
4. Provide interfaces for conducting TOO observations and monitoring scheduler status.
5. Recover appropriately from interruptions and weather losses.
6. Maximize an observing efficiency metric and prioritize cadence control.

3. Optimization Metric

The ZTF scheduler attempts to maximize (Section 4) the total number of exposures taken per night, weighted by the spatial volume probed by each, and subject to the constraints imposed by program balance and cadence (Section 6). If the observing cadences are well chosen, maximizing this quantity will maximize the transient discovery rate. Bellm (2016) explores the relationship between the chosen observing cadences, a survey’s volumetric and areal survey rates, and the transient detection rate.

Neglecting cosmological effects, the volume V_{lim} probed by a given exposure is proportional to the cube of the limiting distance d_{lim} a transient of fiducial absolute magnitude M can

be detected given the limiting magnitude m_{lim} : $V_{\text{lim}} \propto d_{\text{lim}}^3$, where $d = 10^{0.2(m_{\text{lim}} - M + 5)}$ pc (see Bellm 2016). The volumetric weighting per exposure is thus

$$V = 10^{0.6(m_{\text{lim}} - 21)}, \quad (1)$$

where we have absorbed constant factors and normalized to a convenient limiting magnitude for ZTF.

This weighting combines in a self-consistent way many factors that are intuitively relevant for assessing whether an image is “good”: the limiting magnitude depends on the filter, seeing, airmass, and sky brightness. We use a model (Section 5) to predict the variation in limiting magnitude and hence our metric as a function of these time-varying inputs. Accordingly, our optimization will naturally select exposures near zenith and away from the moon; but by combining them in a single scalar the optimization can coherently trade these factors against one another as they change through the night.

Our metric deliberately does not contain factors that account for relative scientific priority or cadence. These concerns have no general quantitative relationship to our objective function or each other.¹⁴ Instead, we use the structure of the optimization algorithm (Section 4) to impose these constraints.

Our optimization algorithm (Section 4) maximizes the summed metric over an entire night. In cases where a greedy algorithm is more convenient, it is simple to define an instantaneous volumetric survey speed,

$$\dot{V} \propto 10^{0.6m_{\text{lim}}} / (t_{\text{exp}} + t_{\text{OH}}) \quad (2)$$

that normalizes the volume probed in an exposure by the time required to obtain it, a sum of the exposure time t_{exp} and any readout or slew overheads t_{OH} .

Other optimization metrics will be more appropriate for surveys such as LSST (Ivezić et al. 2019) that are concerned with coadded depth in addition to transient discovery; these may easily be substituted in our algorithm (Section 4). For instance, Tonry (2011) suggests a weighting factor derived from information theory with a metric proportional to $10^{0.8m_{\text{lim}}}$.

4. Algorithm

The ZTF scheduling process begins with a set of observing programs. These are defined by their footprint on the sky (a discrete set of fields, which may be larger than the set observable in one night or even one lunation); the fraction of the telescope observing time they are allocated; the number of nights between successive revisits to this field for this program¹⁵ (i.e., the “inter-night gap,” such as a 1-day or 3-day cadence); and the number of

¹⁴ One could imagine a global model for the information contributed by a potential new observation given the past history of observations at that location, but we expect that this approach would require computationally expensive lightcurve modeling within the optimization loop and likely be limited to a single class of objects such as SN Ia.

¹⁵ Observing programs are not allowed to couple their observing sequences to the observing history of other programs; each is completely independent.

visits and filter set for observations requested within a night (e.g., two nightly visits, one in g -band and one in r -band).

At the beginning of the night, each observing program provides a list of fields that are visible long enough to obtain the requested observations and have not been observed within that program's inter-night gap. The resulting input to the scheduling algorithm are "Request Sets:" each request set is a ZTF field along with the number of observations requested per filter, the exposure time per observation, and appropriate observing program metadata. For example, one Request Set might be for field 123 with three g -band and three r -band exposures tonight, all 30 s exposures: six Requests in total. The scheduler also uses the past observation history, the fraction of time allocated to each observing program, and the length of the night to determine the number of allowed requests per observing program.

The scheduling algorithm then determines which request sets are observed, at what time to schedule the constituent observations, and how to arrange the slews and filter changes to maximize efficiency. We use ILP techniques to solve this problem; our notation and approach is inspired by that of Lampoudi et al. (2015), but there are significant differences which we discuss in Section 4.7.

4.1. Parameters

We construct the observing schedule by dividing the night into a set of temporal blocks T . This discretization is necessary to make scheduling the entire night computationally tractable: rather than determining an exact sequence of hundreds or thousands of exposures, we must merely assign the observations to 15–25 blocks. The block structure also provides a useful means of applying cadence constraints and minimizing filter changes (Section 4.4).

The length of the block T_{block} is set to the minimum desired separation between exposures. For ZTF we set the time block size to 30 minutes, sufficient to identify the motion of main-belt asteroids with ZTF's moderate image quality ($\sim 2''$ FWHM).

The set of available filters in the camera is F . The set of Request Sets from all observing programs P is R . For each Request Set we use Equation (1) to calculate the volumetric weighting factor V_{rff} for an observation of that field at time block $t \in T$ for filter $f \in F$. The weight of an observation thus changes through the night: image quality, atmospheric transmission, and sky brightness change as fields rise and set, and the sky brightness also changes with the motion of the Sun and moon. We approximate the weight factor as constant within any single time block. Filter changes only occur at the block boundaries.

4.2. Decision Variables

We solve for binary decision variables:

1. $Y_{rff} = 1$ if Request Set $r \in R$ has an observation scheduled at time block $t \in T$ using filter $f \in F$, and 0 otherwise.

We also define resultant variables used to apply constraints (Section 4.4):

1. $Y_{tf} = 1$ if observations in time block $t \in T$ are conducted using filter $f \in F$, and 0 otherwise.
2. $Y_s = 1$ if the filter changes between time blocks $s \in T$ and $s + 1 \in T$, and 0 otherwise.

4.3. Objective

The optimizer maximizes an objective function which sums the volume-weighted (Equation (1)) number of exposures scheduled through the night. Because of how we constrain the number of exposures in a temporal block (Section 4.4), we also penalize for exposures lost due to filter changes. The objective function is thus

$$\max \left(\left(\sum_{r \in R} \sum_{t \in T} \sum_{f \in F} V_{rff} Y_{rff} \right) - \left(\frac{t_{\text{filt}}}{t_{\text{exp}} + t_{\text{OH}}} w \sum_{t \in T} Y_s \right) \right), \quad (3)$$

where t_{filt} is the time required to change filters and w is a weight factor ($\approx \max(V_{rff})$) accounting for the value of each lost exposure.

4.4. Constraints

Each scheduled Request Set should have exactly the requested number of observations per filter n_{rf} :

$$\sum_{t \in T} Y_{rff} = n_{rf}, \forall f \in F \forall r \in R. \quad (4)$$

Only one filter should be used within a given time block:

$$\sum_{f \in F} Y_{tf} = 1, \forall t \in T. \quad (5)$$

The time required to execute the observations assigned to a block should be less than or equal to the length of the block:

$$\sum_{r \in R} \sum_{f \in F} Y_{rff} (t_{\text{exp},r} + t_{\text{OH}}) \leq T_{\text{block}}, \forall t \in T, \quad (6)$$

where we have here allowed for variable exposure times per request set. Because we have not yet *sequenced* the observations in a block (see Section 4.5), we do not know the exact slew times required and so use a fiducial value of 9 s (corresponding to the limit imposed by CCD readout) for the overhead time. In practice this means a few more exposures may be scheduled in a block than can practically be observed.

Finally we apply a constraint to limit the number of scheduled requests per observing program to enforce the desired balance between programs:

$$\sum_{r \in R, p' = p} \sum_{t \in T} \sum_{f \in F} Y_{rff} \leq n_p, \forall p \in P \quad (7)$$

where the number of allowed exposures n_p for a given observing program p is determined each night from the

fractional observing time assigned to the program, the length of the requested exposures, and the past observing history.

4.5. Sequencing Exposures within a Block

The solution to the ILP program is a list of observations assigned to each time block in the night. We use a second optimization process to sequence observations efficiently within each block. We compute the pairwise slew times between all fields assigned to a block, and then solve the Traveling Salesman Problem (TSP) as an ILP problem in Gurobi¹⁶ using the cutting plane method (Dantzig et al. 1954). In our application the quantity to be minimized is not the total length of the salesman’s tour, but the total time spent slewing between fields within the block.

Since the P48 is an equatorial telescope, the slew time between fields using the hour angle and decl. axes of the telescope do not vary with time. However, slews of the dome are azimuthal and so must be computed for each time block individually. Because the same field may be requested by multiple observing programs, we apply a penalty factor to prevent the same field from being observed multiple times consecutively within a block, reducing the redundancy of the repeated exposures.

4.6. Re-solving within the Night

It is computationally feasible to resolve the entire optimization problem repeatedly within the night to account for time lost to weather, TOOs, or other schedule disruptions. However, once time is lost during the night there is a complex tradeoff in determining which observing sequences to complete. One option would be complete some observing sequences exactly as requested and omit others entirely. Another possibility would be accept partial completion of the remaining request sets, but this may limit the scientific usefulness of the observations. To avoid making program-dependent decisions, we implement recomputes in a more limited way: at each block boundary, the best un-observed requests from earlier in the evening are reassigned to any unused time in the current block.

4.7. Comparison with Other ILP Scheduling Algorithms

Our ILP algorithm differs in important ways from those of LCO (Lampoudi et al. 2015) and ALMA (Solar et al. 2016). Because LCO and ALMA are scheduling scientifically disparate observations, both schedulers use the TAC-assigned priority to provide an overall objective function. Beyond simple acceptability constraints, the schedulers do not weight by the relative quality of an observation at any given time. In contrast, because ZTF is simply an imaging survey and all surveys have equal priority, we are free to optimize an

objective function (Section 3) that explicitly and self-consistently accounts for the time-varying quality (and hence scientific value) of any given exposure. Additionally, because the observations scheduled by LCO and ALMA are long relative to the time to transition between them, their scheduling algorithms do not attempt to account for these transitions. For ZTF, readout and slew overheads account for about 25% of any given exposure, and long slews and filter changes create even larger losses. Accordingly our approach sequences exposures within a block to explicitly minimize the time spent slewing, and our objective function penalizes filter changes for the time lost.

4.8. Summary of Algorithm Features

Our choice of this ILP algorithm was motivated by its strengths in handling cadenced observing within a night and in balancing several simultaneous surveys (Section 6). To our knowledge ZTF must attempt to execute more independent observing programs simultaneously than any other wide-field imaging surveys (typically five, in addition to TOO observations), so rigorous cadence control is required. The complete night lookahead provided by our algorithm ensures that observations are scheduled for the best time in the night, accounting for the number of observations required, variations in airmass and sky brightness, and the competing demands of other surveys. Our ILP constraints (Equation (4)) guarantee that the scheduler will provide the requested number of observations if a field is observed. This capability is vital for the success of observing programs requiring many observations during the night. For example, the ZTF Collaboration’s Extragalactic High Cadence survey (Section 6) requires six nightly observations per field in two filters over three or more hours, which would be challenging to schedule effectively without the lookahead provided by our algorithm. The scheduler uses the past observing history to rigorously maintain night-to-night cadences and to enforce the time allocated to the various surveys. The scheduler treats each survey uniformly and interleaves the requested observations. The algorithm self-consistently trades the observing time lost to filter changes against potential improvements in the quality of the observations. Finally, while these values are only a component of our scientifically motivated optimization metric, we note that the scheduler is effective at observing near zenith and minimizing slew time (Section 7).

4.9. Limitations

Obtaining these characteristics required accepting some tradeoffs in the capabilities of the scheduler. Notably, our algorithm does not enable exact cadences or filter sequences within a night. For example, it is not possible to schedule a *g*-band observation followed 12–15 minutes later by an *i*-band observation. Rather, a total number of observations per filter is

¹⁶ See <http://examples.gurobi.com/traveling-salesman-problem/> for a TSP solver implemented with Gurobi.

guaranteed, each separated by roughly the time block size (here, 30 minutes). Even that minimum separation is not strictly guaranteed, as the sequencing of the fields within a block is independent, and observations may occur near the end of one block and near the beginning of the next. We do schedule a minority of surveys that require more controlled within-night cadences; we implement these as pre-defined queues that interrupt the operations of the primary scheduler (Section 5).

While our Traveling Salesman solution (Section 4.5) minimizes the slew time within a given block, the initial block assignment does not account for the slew time between the fields. Accordingly our algorithm cannot be said to globally minimize slew time, although in practice we find that slew overheads are small (Section 7).

It is possible that the scheduler does not assign enough observations to a specific block to fill it.¹⁷ This is because our constraint on the number of observations per block (Equation (6)) is less than or equal to the number of observations that would fill the block, not a strict equality. Strict equality creates overconstrained models that cannot be solved. In general, the scheduling algorithm is subject to the details of the input observing programs. If a large fraction of the observing time is concentrated on a narrow area of the sky, for instance, there is no way for the scheduler to manufacture unrequested observations to fill other parts of the night. In practice, we manage this issue by simulating potential observing strategies in advance when possible. Additionally, re-solves during the night (Section 4.6) can fill in previously unscheduled time with scheduled observations that were missed. Finally, we implement a “fallback queue” to ensure that useful observations can be obtained if the main queue runs empty. To date this fallback time has largely been used to improve sky coverage for reference image building. Without re-solves, typically the amount of slack in the schedule is a few percent if the input observing programs are well-balanced.

Finally, our current scheduler implementation does not yet dynamically adapt to changing observing conditions within the night due to the additional operational complexity and potential for schedule thrashing. We do not attempt cloud avoidance, for instance, or adjust to changes in seeing. Such extensions are possible. One approach would be to maintain the overall block structure but conduct more extensive re-optimization at the block boundaries, and use a more dynamic selection of the next target field within a block to handle short-timescale variations.

5. Implementation

We have implemented the scheduling algorithm as a Python library, which is publicly available¹⁸ under an open source license. The scheduler code takes advantage of a range of open-

source Python libraries, including *Astropy* (The Astropy Collaboration et al. 2018), *Astroplan* (Morris et al. 2018b), and *pandas* (McKinney 2010). We use the commercial optimization package *Gurobi*¹⁹ (Gurobi Optimization 2018) under an academic license to perform the core ILP optimization. While some attempt has been made to make the scheduler interfaces telescope agnostic, the library does encode assumptions specific to the ZTF use case.

Our objective function (Section 3) requires a detailed sky brightness model. We trained a gradient boosted tree model as implemented in *xgboost* (Chen & Guestrin 2016) on historical data from ZTF (and initially PTF). Our model predicts the sky brightness in each filter as a function of telescope pointing altitude and azimuth, Sun altitude, and moon altitude, moon distance, and moon illumination fraction.

The scheduler library can be run both in simulation mode (using historical weather data from PTF) as well as in operations. For on-sky scheduling, we use an *aiohttp*²⁰ webserver on the primary host computer of the ZTF Robotic Observing System (ROS; R. Dekany et al. 2019, in preparation). The webserver calls the scheduling library and provides a RESTful interface for command and status information.

We run the optimizer for five minutes before the start of the night’s observations using two cores of the host machine, which yields satisfactory results without interfering with other robotic operations. Moving the scheduler to a dedicated host would enable us to obtain equivalent performance in a shorter time by parallelizing over a larger number of CPUs. The Gurobi solver library offers native parallelization by initializing multiple candidate solutions on different threads and concurrently optimizing each, terminating when one thread obtains a solution. The memory footprint during nightly optimization can be as high as 700 MB, dropping to about 500 MB in sustained operations, although we have not attempted to optimize these values.

The ROS system obtains the required evening and morning calibration observations; the scheduler is responsible for selecting on-sky science observations. The ROS software also updates focus through the night using telemetry from the $2k \times 2k$ focus CCDs on the perimeter of the mosaic. Focus observations and updates occur concurrently with science observations and do not create additional overheads.

While most ($\gtrsim 90\%$) of ZTF’s observing programs are scheduled using our ILP algorithm, a subset require precise sequencing over continuous time blocks. These programs we implement as simple “list queues” that prescribe an expected start and stop time and a defined sequence of exposures to take. We use this mechanism both for pre-planned observations as well as TOO triggers. A monitoring thread checks for the presence of such timed queues every ten seconds and switches to the appropriate queue if its validity window has started. If list queues are planned before the

¹⁷ This is known as “slack” in the optimization literature, and is also a feature of the LCO scheduler (Lampoudi et al. 2015).

¹⁸ https://github.com/ZwickyTransientFacility/ztf_sim

¹⁹ <http://www.gurobi.com/>

²⁰ <https://aiohttp.readthedocs.io/>

Table 1
Major ZTF Surveys, Year One

Survey	Total Survey Footprint	Inter-night Cadence	Nightly Cadence	Average Nightly Area	Time Allocated
Public Surveys					
Northern Sky Survey	23675 deg ²	3 days	1 <i>g</i> , 1 <i>r</i>	4325 deg ²	40% × 85%
Galactic Plane Survey	2800 deg ²	1 day	1 <i>g</i> , 1 <i>r</i>	1475 deg ²	40% × 15%
ZTF Collaboration Surveys (Year One)					
Extragalactic High Cadence	3000 deg ²	1 day	3 <i>g</i> , 3 <i>r</i>	1725 deg ²	40% × 67.5%, Mar–Nov
<i>i</i> -band	10725 deg ²	4 day	1 <i>i</i>	1975 deg ²	40% × 22.5%, Mar–Nov
Target of Opportunity	Varies	varies	Varies	varies	40% × 10%
High-cadence Plane Survey	~2100 deg ²	N/A	≥2.5 hr continuous, <i>r</i>	95 deg ²	40% × 80%, Aug, Nov–Jan
Twilight Survey	N/A	N/A	4 <i>r</i>	425 deg ²	12°–18° twilight, Nov–Feb
Asteroid Rotation Period	N/A	N/A	>25 <i>r</i>	600 deg ²	40% × 80%, Jan–Feb

Note. Partnership surveys transition on the fifteenth of the month. The high-cadence plane survey substituted for the extragalactic high cadence survey for two weeks in 2018 August.

start of the night’s observing and will take at least one complete block, the ILP optimizer omits those blocks from scheduling the primary ZTF programs.

6. ZTF Surveys

ZTF observing time is divided between three major programs: public surveys facilitated by an award from the NSF Mid-Scale Innovations Program (MSIP; 40% of the telescope time); surveys designed by the members of the ZTF Collaboration (40%); and surveys selected each semester by the Caltech TAC (20%). The ZTF scheduler attempts to achieve this balance each calendar month, roughly the interval in which collaboration and Caltech sub-programs change. The scientific goals of the surveys are discussed in Graham et al. (2019).

All surveys select fields from a discrete field grid.²¹ The current surveys all use the “primary” grid, which covers the entire sky with an average overlap between fields of about 0°29 in R.A. and 0°26 in decl. The average spacing between fields in the primary grid is 7°2 north–south and 7°0 east–west. Upgrades to the P48 drive motors enables slews between adjacent fields within the 8.3 s CCD readout time. The primary grid is arranged to align with $b = 0^\circ$ of the Galactic Plane to improve the efficiency of the MSIP surveys. It also ensures good coverage of nearby galaxies (M31, M33, M51, M101, etc.). A secondary grid, offset from the first by roughly half a field in R.A. and decl., fills in missing sky coverage due to the gaps between CCDs and provides additional depth for sky areas covered by vignettted corners of the focal plane in the primary grid. The

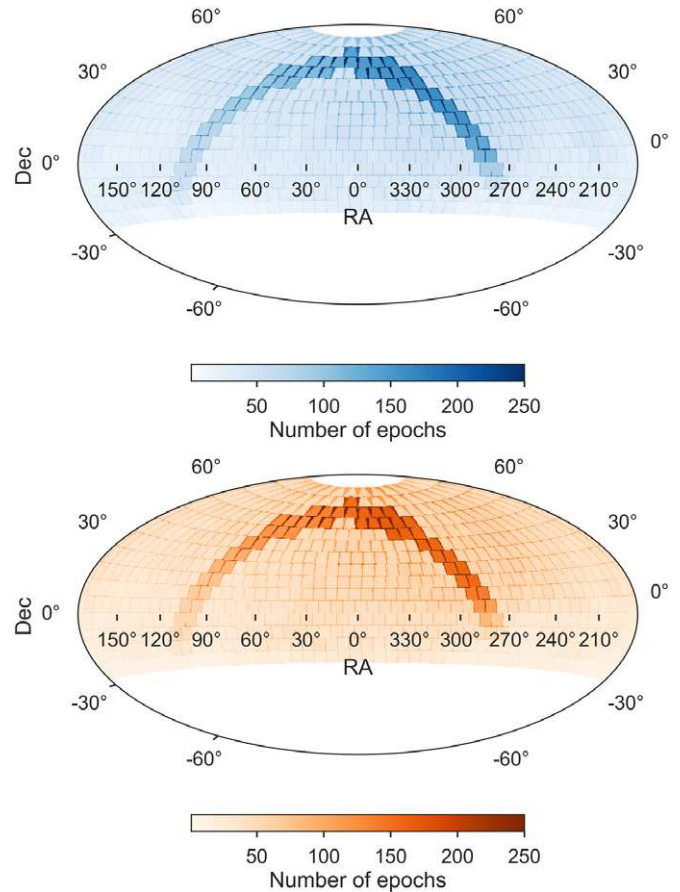


Figure 1. Number of epochs obtained by the MSIP surveys across the sky in *g*-band (blue, top) and *r*-band (orange, bottom) to date. (A color version of this figure is available in the online journal.)

²¹ See https://github.com/ZwickyTransientFacility/ztf_information.

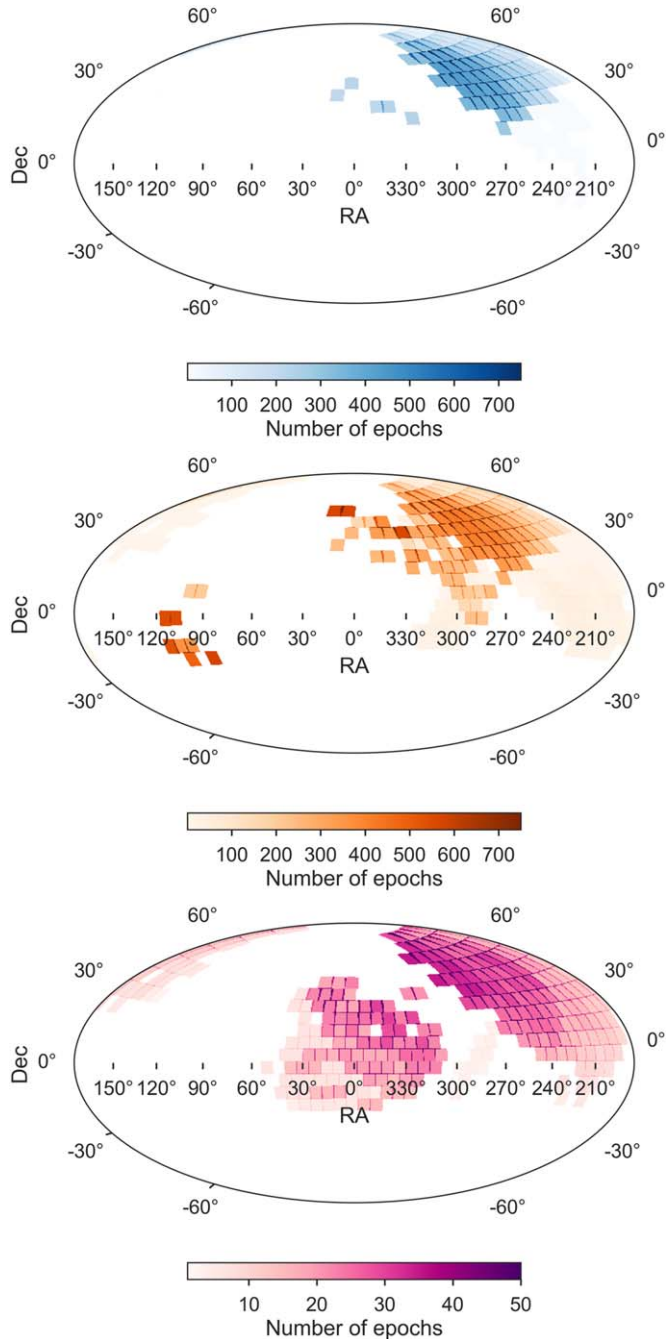


Figure 2. Number of epochs obtained by the ZTF Collaboration surveys across the sky in *g*-band (blue, top), *r*-band (orange, middle), and *i*-band (pink, bottom) to date. *g*-band observations are almost exclusively in the Extragalactic High-cadence fields, and *i*-band observations in the *i*-band survey fields. Several surveys contribute to the *r*-band coverage, including the Extragalactic High Cadence survey, the High-cadence Plane Survey, and the Twilight Survey.

(A color version of this figure is available in the online journal.)




Figure 3. Histogram of time elapsed between the end of one observation and the start of the next. The vertical dashed line indicates the shortest possible time between exposures (~ 9.1 s) due to readout time and the shutter opening and closing.

(A color version of this figure is available in the online journal.)




Figure 4. Histogram of total distance slewed between observations. The vertical dashed line at 7° indicates the average East–West distance between two adjacent fields in the same grid, although the exact grid spacing varies slightly with declination.

(A color version of this figure is available in the online journal.)

primary grid alone covers 87.5% of the sky; with the addition of the secondary grid, spatial coverage increases to 99.2%.

Table 1 provides a high-level overview of the major public and collaboration surveys.

6.1. Public Surveys

The ZTF public surveys were defined in the ZTF proposal to the NSF MSIP program. A “Northern Sky Survey” covers all fields with centers $\delta \geq -31^\circ$ and $|b| > 7^\circ$. When a field is up, on every third night it is observed once in *g*-band and once in

r-band, with a spacing of at least 30 minutes between observations to discriminate between transients and moving objects (see Miller et al. 2017). The Northern Sky Survey is allocated 85% of the public time (34% of telescope time).

ZTF also conducts a Galactic Plane Survey using the remaining 15% of the public time (6% of telescope time). Fields with $\delta \geq -31^\circ$ and $|b| \leq 7^\circ$ are visited twice each night they are visible, with one observation in *g*-band and one in *r*-band, again separated by at least 30 minutes.

Figure 1 shows the *g* and *r* band sky coverage of the MSIP surveys to date.

Public alerts are issued in near real-time for all sources identified in image differencing from the public surveys (Masci et al. 2019; Patterson et al. 2019). Additionally, images, catalogs, and direct imaging lightcurves (Masci et al. 2019) will be released in data releases beginning in 2019.

We plan to continue these surveys in their present form through the first half of the three-year ZTF survey. At that point we will assess the scientific returns from ZTF and the broader time-domain landscape and evolve the public surveys accordingly.²²

6.2. ZTF Collaboration Surveys

The ZTF Collaboration defined an initial slate of surveys for the first year of ZTF operations, although an extended commissioning period meant that the total time the surveys were executed is about 11 months. Five major surveys were approved, with approximately two collaboration surveys plus TOO observations active at any one time.

The bulk of the time (mid-March to mid-November) was dedicated to two extragalactic surveys: a high-cadence survey of 6 visits nightly (3 in *g*-band and 3 in *r*-band; 67.5% of the collaboration time, or 27% of the total time) optimized for the discovery of young supernovae and other fast transients, and a slow, wide *i*-band survey (one visit per field every four nights; 22.5% of the collaboration time, or 9% of the total time) designed to improve the cosmological constraining power of ZTF SNe Ia. In the future, co-adding multiple images taken by the high-cadence survey within a night can provide additional sensitivity to faint transients as well as strongly lensed supernovae (Goldstein et al. 2018).

Additionally, 10% of the collaboration time (4% of the total time) was reserved for TOO observations of gamma-ray bursts, neutrino counterparts, gravitational wave triggers from LIGO and VIRGO, and Near-Earth Objects.

Two weeks in August and two months from mid-November to mid-January were allocated to very high cadence observations of Galactic Plane fields. A typical observation pattern was to alternate between two adjacent fields continuously for 2.5 hr on two consecutive nights in *r*-band. These “continuous




Figure 5. Histogram of the number of filter exchanges per hour, computed on a nightly basis. Nights shortened by weather may have no filter exchanges and hence appear as zero exchanges per hour.

(A color version of this figure is available in the online journal.)




Figure 6. Histogram of airmass values for ZTF (filled blue), PTF (black), and iPTF (orange). The structured peaks in the ZTF histogram are due to the wider spacing of the fields compared to PTF. During the late spring to early fall ZTF observed at lower airmass (light blue histogram) due to the distribution of collaboration fields and the shorter nights.

(A color version of this figure is available in the online journal.)

cadence” observations enabled more sensitive searches for short-period binaries and stellar outbursts.

For three months from mid-November to mid-February, the period from 12° to 18° evening and morning twilight was devoted to the search for Near-Earth Objects at small solar elongation, with four visits over a 30 minute period in *r*-band separated by 5–10 minutes.

Finally, during the period from mid-January to mid-February, a high-cadence survey near opposition will obtain tens of nightly observations per field in order to identify fast-rotating asteroids.

Figure 2 shows the *g* and *r* band sky coverage of the collaboration surveys to date.

New ZTF Collaboration surveys will be selected for observations in 2019.

²² The ZTF MSIP PI will select the public surveys in consultation with the ZTF Community Science Advisory Committee.




Figure 7. Histogram of hour angle values for ZTF.
(A color version of this figure is available in the online journal.)

Images, catalogs, and lightcurves for data obtained during collaboration surveys will be released publicly during scheduled data releases after an 18-month proprietary period.

6.3. Caltech Surveys

Surveys selected by the Caltech TAC have included programs optimized for the discovery of transient, variable, and moving objects, with particular priority given to cadences and sky areas not being surveyed by the collaboration. As these surveys are proposed and led by individuals we do not detail them further in this manuscript. Data releases for these surveys are the responsibility of the proposer.

7. Performance

7.1. Simulated Performance

To compare the performance of our ILP algorithm to a simple greedy optimizer, we simulated the 2018 May observing programs using both optimizers with realistic weather losses. Both algorithms attempted to maximize our survey speed metric (Equation (1)): the ILP algorithm optimized the form of the objective function in Equation (3), while the greedy algorithm optimized the instantaneous volumetric survey speed \dot{V} (Equation (2)). Both approaches yielded comparable numbers of exposures per hour. However, the ILP approach provided an 9% increase in the metric (Equation (1)) summed over all exposures. It also scheduled observations closer to zenith, with a median airmass of 1.11 compared to 1.20 for the greedy approach. Perhaps surprisingly, the greedy scheduler yielded fewer filter changes, averaging 2.0 per night compared to 3.6 per night. Additionally, the ILP solution produced 4% slack before within-night re-optimization.

The importance of the lookahead provided by the ILP algorithm is most clearly demonstrated by the sequence completion fraction—the fraction of observed fields for which




Figure 8. Histogram of the metric values per image. Colors indicate tertiles of moon phase, with dark blue corresponding to dark time (0%–33% moon phase), medium blue indicating gray time (33%–66%), and light blue bright time (66%–100%). Smaller scale structure is due to the discrete spacing of the ZTF field grid: some fields transit at higher airmasses depending on their declination.

(A color version of this figure is available in the online journal.)

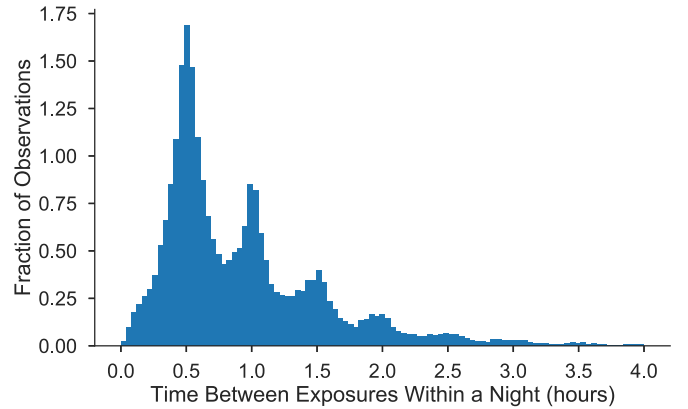


Figure 9. Histogram of times between successive observations of a field by a given program within a night.

(A color version of this figure is available in the online journal.)

the scheduler obtains all of the desired nightly observations. Including the effects of weather losses, the greedy algorithm completed an average of 64% of the MSIP Northern Sky Survey observations, 79% of the MSIP Galactic Plane Survey observations, and 72% of the collaboration Extragalactic High-cadence observations. In contrast, the ILP scheduler completed 81% of the requested observations for each of the same surveys.

7.2. On-sky Performance

The scheduler has performed effectively during on-sky operations. It has scheduled more than 120,000 observations since the start of formal survey operations. Overall balance




Figure 10. Histogram of times between successive observations of a field by a given program from night to night. Left: MSIP Northern Sky Survey (3-day cadence) and Galactic Plane Survey (1-day cadence). Right: partnership *i*-band Survey (4-night cadence) and Extragalactic High Cadence Survey (1-day cadence). Revisit times longer than the target cadence are due to weather and scheduling effects. (A color version of this figure is available in the online journal.)




Figure 11. Histogram of fractional sequence completion for the major ZTF surveys. The MSIP Northern Sky Survey (All Sky) and Galactic Plane Survey (Nightly Plane) each request two observations per field, so have fractional completion of 0.5 or 1.0 on nights the field is observed. The collaboration Extragalactic High Cadence Survey (High Cadence) has six observations nightly, so the fractional completion can range from 1/6 to 6/6. (A color version of this figure is available in the online journal.)

between the MSIP, ZTF Collaboration, and Caltech observing programs was maintained, with 42% of scheduled observations conducted in the MSIP surveys, 40.2% in the collaboration surveys, and 18.8% in the Caltech surveys. The slight shortfall in the Caltech programs can be attributed in part to short intervals when no Caltech programs were available or they did not fill the entire time allocation.

The scheduler uses the telescope efficiently, with the median time between observations of 9.9 s (Figure 3) and most slews of one field offset (Figure 4). The tenth–90th percentile overhead times and slew distances are 9.4 s–14.9 s and 6°0–16°1

respectively. Repeated exposures of the same field without slews have a median time between exposures of 9.4 s. Filter exchanges occur less than once per hour during the vast majority of nights (Figure 5). While the ZTF filter exchanger hardware is designed to support a higher rate of filter changes, our penalty factor (Equation (3)) self-consistently trades the need for filter changes against the time lost during the exchange and prevents filter changes from occurring on every block boundary. The choice of optimization metric leads observations to be preferentially scheduled around zenith (Figures 6 and 7). Figure 8 shows the resulting metric values, which vary sharply with moon phase.

The scheduler delivers the desired cadences. Eighty percent of all observations are spaced by at least 30 minutes as desired for asteroid discrimination (Figure 9). The intra-night cadences for the major surveys are delivered as expected (Figure 10), with minimal tailing to longer-than-desired revisit times.

Finally, the scheduler delivers a high fraction of completed observation sequences, averaging 84.6% completion for the MSIP surveys and the collaboration high-cadence surveys (Figure 11).

8. Conclusions

We have implemented a scheduling algorithm for wide-field imaging time-domain surveys that cleanly delineates three core concerns:

1. The intrinsic quality of a specific image, as specified by signal-to-noise ratio or spatial volume probed (Section 3); this encapsulates image quality, sky background, airmass, and related terms.

2. The scientific value of obtaining an image of a given field at a specific time; these desired cadences are specified a priori.
3. The means of sequencing observations to maximize efficiency and throughput.

A survey focused on a single class of astrophysical object could combine the first two goals, trading off the information gained from a high SNR observation now versus a low SNR observation later, using knowledge of lightcurve shape, periodicity, etc. However, this combination is not possible for a general-purpose, wide-field survey. Similarly, long-term planning could account for the uncertain availability of future observations (due to weather, instrument failures, etc.).²³

We suggest that this formalism would provide useful clarity to the problem of scheduling observations for LSST. In particular, we argue that an appropriate scheduler for LSST would attempt to maximize the contribution of a night's observing to the total coadded depth of the survey, subject to the desired cadence constraints. This is simply the approach developed here with a slightly modified objective function (Section 3). It directly optimizes the metric of interest without requiring intermediary features which intermix concerns of image quality, cadence, and efficiency (see Naghib et al. 2019). Since the number of exposures scheduled nightly for ZTF and LSST are comparable, our on-sky implementation demonstrates directly that this algorithm could be feasibly applied to LSST. Further work would be needed to adapt our algorithm to meet all LSST requirements and rigorously compare its performance to other scheduling approaches, however.

The coming decade will see new surveys of unprecedented scale—imaging and spectroscopy, on the ground and in space. To fully reap the scientific value of these large investments, astronomers must give sustained attention to the scheduling problems unique to each survey. Cross-fertilization with research in the field of operations research may be of particular value. Different surveys will necessarily require different algorithms and metrics, but thanks to increasingly powerful computing resources, new optimization approaches are now feasible. Careful attention to scheduling can provide some of the most cost-effective improvements in science throughput available.

E.C.B. thanks Eric Saunders, Ben Shappee, Lynne Jones, and Peter Yoachim for useful conversations.

Based on observations obtained with the Samuel Oschin Telescope 48 inch and the 60 inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility project. Z.T.F. is supported by the National Science Foundation under grant No. AST-1440341 and a collaboration including Caltech, IPAC, the Weizmann Institute for Science, the Oskar Klein

Center at Stockholm University, the University of Maryland, the University of Washington, Deutsches Elektronen-Synchrotron and Humboldt University, Los Alamos National Laboratories, the TANGO Program of the University System of Taiwan, the University of Wisconsin at Milwaukee, and Lawrence Berkeley National Laboratories. Operations are conducted by COO, IPAC, and UW.

E.C.B. is supported in part by the NSF AAG grant 1812779 and grant #2018-0908 from the Heising-Simons Foundation. E.C.B. also acknowledges support from the Large Synoptic Survey Telescope, which is supported in part by the National Science Foundation through Cooperative Agreement 1258333, managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory. Additional LSST funding comes from private donations, grants to universities, and in-kind support from LSSTC Institutional Members. E.C.B. is grateful for further support from the University of Washington College of Arts and Sciences, Department of Astronomy, and the DIRAC Institute. University of Washington's DIRAC Institute is supported through generous gifts from the Charles and Lisa Simonyi Fund for Arts and Sciences, and the Washington Research Foundation.

E.O.O. is grateful for support by grants from the Willner Family Leadership Institute, Ilan Gluzman (Secaucus NJ), the Israel Science Foundation, Minerva, BSF, BSF-transformative, Weizmann-UK, and the I-Core program by the Israeli Committee for Planning and Budgeting and the Israel Science Foundation (ISF).

Facility: PO:1.2 m.

Software: Astropy (The Astropy Collaboration et al. 2018), Astropplan (Morris et al. 2018b), Numpy (Van Der Walt et al. 2011), Scipy (Jones et al. 2001), pandas (McKinney 2010), Matplotlib (Hunter 2007), Seaborn (Waskom et al. 2018), Scikit-Learn (Pedregosa et al. 2011), xgboost (Chen & Guestrin 2016), Gurobi (Gurobi Optimization 2018), makecrite (Price-Whelan et al. 2018).


ORCID iDs

Eric C. Bellm  <https://orcid.org/0000-0001-8018-5348>

Shrinivas R. Kulkarni  <https://orcid.org/0000-0001-5390-8563>

Ulrich Feindt  <https://orcid.org/0000-0002-9435-2167>


Matthew J. Graham  <https://orcid.org/0000-0002-3168-0139>

Ariel Goobar  <https://orcid.org/0000-0002-4163-4996>

Thomas Kupfer  <https://orcid.org/0000-0002-6540-1484>

Chow-Choong Ngeow  <https://orcid.org/0000-0001-8771-7554>

Peter Nugent  <https://orcid.org/0000-0002-3389-0586>

Quan-Zhi Ye (葉泉志)  <https://orcid.org/0000-0002-4838-7676>

²³ I.e., “Expected Future-discounted Information Gain” (Hogg 2018).

References

- Alexander, K. D., Berger, E., Bower, G., et al. 2017, arXiv:1703.04692
- Bellm, E. C. 2016, *PASP*, **128**, 084501
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, *PASP*, **131**, 018002
- Bloemen, S., Groot, P., Woudt, P., et al. 2016, *Proc. SPIE*, **9906**, 990664
- Chen, T., & Guestrin, C. 2016, in Proc. 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining, KDD '16, ed. B. Krishnapuram et al. (New York: ACM), 785, doi:10.1145/2939672.2939785
- Dantzig, G., Fulkerson, R., & Johnson, S. 1954, *J. Oper. Res. Soc. Am.*, **2**, 393
- Denny, R. B. 2004, *SASS*, **23**, 35
- Denny, R. B. 2006, Dispatch Scheduling of Automated Telescopes, <http://scheduler.dc3.com/Scheduler2006.pdf>
- Flaugher, B., Diehl, H. T., Honscheid, K., et al. 2015, *AJ*, **150**, 150
- Goldstein, D. A., Nugent, P. E., & Goobar, A. 2018, arXiv:1809.10147
- Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, in press (arXiv:1902.01945)
- Gurobi Optimization, L. 2018, Gurobi Optimizer Reference Manual, <http://www.gurobi.com>
- Hogg, D. 2018, Information, Interpretability, Emulation, Isocurvature, <http://hoggresearch.blogspot.com/2018/11/information-interpretability-emulation.html>
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Ivezić, Z., Kahn, S. M., Tyson, J., et al. 2019, *ApJ*, **873**, 111
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python, <http://www.scipy.org/>
- Kaiser, N., Burgett, W., Chambers, K., et al. 2010, *Proc. SPIE*, **7733**, 77330E
- Lampoudi, S., Saunders, E., & Eastman, J. 2015, arXiv:1503.07170
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, *PASP*, **121**, 1395
- Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *PASP*, **131**, 018003
- McKinney, W. 2010, in Proc. 9th Python in Science Conf., ed. S. van der Walt & J. Millman (Washington, D.C.: IEEE), 51
- Miller, A. A., Kasliwal, M. M., Cao, Y., et al. 2017, *ApJ*, **848**, 59
- Miyazaki, S., Komiyama, Y., Kawanomoto, S., et al. 2018, *PASJ*, **70**, S1
- Morris, B. M., Tollerud, E., Sipőcz, B., et al. 2018a, *AJ*, **155**, 128
- Morris, B. M., Tollerud, E., Sipőcz, B., et al. 2018b, *AJ*, **155**, 128
- Naghieb, E., Yoachim, P., Vanderbei, R. J., Connolly, A. J., & Jones, R. L. 2019, *AJ*, **157**, 151
- Patterson, M. T., Bellm, E. C., Rusholme, B., et al. 2019, *PASP*, **131**, 018001
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825
- Price-Whelan, A., Mechev, A., & Jumerog 2018, adm/makecite: v0.1, Zenodo, doi:10.5281/zenodo.1343295
- Shappee, B. J., Prieto, J. L., Grupe, D., et al. 2014, *ApJ*, **788**, 48
- Solar, M., Michelon, P., Avarias, J., & Garcés, M. 2016, *A&C*, **15**, 90
- The Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
- Tonry, J. L. 2011, *PASP*, **123**, 58
- Tonry, J. L., Denneau, L., Heinze, A. N., et al. 2018, *PASP*, **130**, 064505
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Waskom, M., Botvinnik, O., O’Kane, D., et al. 2018, Seaborn v0.9.0, Zenodo, doi:10.5281/zenodo.1313201