



A Framework for Telescope Schedulers: With Applications to the Large Synoptic Survey Telescope

Elahesadat Naghib¹ , Peter Yoachim² , Robert J. Vanderbei¹, Andrew J. Connolly² , and R. Lynne Jones²

¹ Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08540, USA; enaghib@princeton.edu, rldb@princeton.edu

² Department of Astronomy, University of Washington, Box 351580, U.W., Seattle, WA 98195, USA; yoachim@uw.edu, ajc26@uw.edu, ljones@astro.washington.edu

Received 2018 October 5; revised 2018 December 11; accepted 2019 January 7; published 2019 March 22

Abstract

How ground-based telescopes schedule their observations in response to different variables can significantly impact their efficiency. Namely, competing science priorities and constraints, variations in the weather, and the visibility of a particular part of the sky are some of the determining variables. In this paper we introduce the Feature-based telescope scheduler, which is an automated, decision-making algorithm. It offers *controllability* of the behavior, *adjustability* of the mission, and quick *recoverability* from interruptions for large ground-based telescopes. By framing the Feature-based telescope scheduler in the context of a coherent mathematical model, the functionality and performance of the algorithm are simple to interpret. Consequently, adapting this scheduler for a broad range of ground-based instruments is straightforward. This paper presents a generic version of the Feature-based scheduler, with minimal manual tailoring. We demonstrate its potential and flexibility to serve as a foundation for schedulers of the ground-based instruments. In addition, a modified version of the Feature-based scheduler for the Large Synoptic Survey Telescope (LSST) is introduced and compared to the previous LSST schedulers.

Key words: telescope – instrumentation: miscellaneous – surveys

Supporting material: animation

1. Introduction

The Large Synoptic Survey Telescope (LSST) is a large, ground-based optical survey that will image half of the sky every few nights from Cerro Pachon in Northern Chile. LSST comprises an 8.4 m primary mirror and a 3.2 gigapixel camera. With a 9.6 deg^2 field of view, it will visit each part of its $18,000 \text{ deg}^2$ primary survey area about 1000 times over the course of 10 yr. Each visit will likely comprise a 15 s pair of exposures with a single-visit depth of about 24.5 mag (AB) (in the six bands u , g , r , i , z , and y). The revolutionary role of this telescope calls for no less than optimal operation.

The algorithm that makes the sequential decisions of which filter to use and which direction to point the telescope toward is called *scheduler*. A scheduler has to maximize the scientific outcome of the telescope during its limited period of operation.

There are four primary science drivers for the LSST project: the characterization of dark energy through the multiple cosmological probes (e.g., gravitational weak lensing, luminosity distances from Type Ia supernovae, and baryon acoustic oscillations), mapping the 3D distribution of stars within our Galaxy, a census of solar system objects within the solar system, and a detailed study of the transient and variable universe. Each of these objectives has a different set of constraints and requirements on how the observations are made (e.g., the cadence of the observations, the number of filters as a function of time, the acceptable air-mass range for an observation). Generally, the mission of modern, large, ground-based telescopes, such as LSST, is constrained with various stochastic factors and contains competing objectives that are vastly different in nature, due to the different nature of the scientific expectations. In this paper, we propose a framework to formulate the problem of scheduling for the

new generation of ground-based telescopes and then introduce a scheduler based on the proposed model.

The first generation of schedulers for astronomical instruments was developed for space missions mainly to automate their operation. The *ROSAT* mission's scheduler in 1990 (Nowakowski et al. 1999), *Spike* (Johnston & Miller 1994), the *Hubble Space Telescope*'s scheduler in 1994, and *HSTS* (Muscettola et al. 1995) pioneered many of the developments in algorithmic scheduling of observations for the space missions.

Despite the similarity of the science objectives for space- and ground-based telescopes, the determining factors for the purpose of scheduling are fundamentally different. While space telescopes are required to respect kinematical and dynamical constraints, weather is the main challenge in the scheduling of ground-based telescopes. The former is predictable and efficiently computable, while the latter involves both inherent uncertainties and uncertainties due to computational limitations.

Earlier algorithmic approaches to the scheduling of ground-based telescopes are heavily based on observation *proposals*. Proposals are handcrafted sequences of scripted astronomical observations. They are generally tested only for feasibility (e.g., that a set of fields were visible, or lie within a specified air-mass range, or within a window in time), but not necessarily for optimality. For instance, the operation of Keck Telescope, 1993 (Nelson et al. 1985), is 100% based on proposals, and the Hobby–Eberly Telescope, 1997 (Shetrone et al. 2007), has a semimanual scheduling scheme.

More recently, the development of more expensive ground-based instruments with complex missions made it impossible to rely solely on handcrafted proposals. The need for more efficient use of the instrument's time led to the development of

decision-making algorithms to optimize their science output. For instance, the scheduler of the Liverpool robotic telescope was designed in 1997 to automatically allocate time slots to chunks of scripted observations. This time allocation strategy was preferred to scheduling at the single-visit level. The scheduling at the single-visit level is referred to as optimal scheduling by the authors (Steele & Carter 1997), and it is stated that the optimal scheduling requires reevaluating the future sequence of observations once it is interrupted, but the necessary extra computation is neither affordable nor fast enough. However, in this paper we show that the scheduling in the single-visit level, optimal scheduling, can be quickly recovered after an interruption, if a memoryless framework is used. Thus, the optimality does not necessarily need to be sacrificed because of the limited computational resources. Another example is the Las Cumbres Observatory, Global Telescope Network (LCOGT), with one of the most advanced telescope scheduling algorithms (Boroson et al. 2014; Saunders & Lampoudi 2014). LCOGT uses an integer linear programming (ILP) model, solved with the Gurobi solver (Gurobi Optimization 2016), to optimize the scheduling of observations over a global network of telescopes (Lampoudi et al. 2015). Due to the success of this approach, the Zwicky Transient Facility at Palomar Observatory (Bellm 2014) has also adopted an ILP scheduler.³ The ILP scheduling model performs well for observatories where slew overheads are small compared to exposure execution time. LCOGT is able to schedule observations in long contiguous blocks. In contrast, LSST plans to have observations of about 30 s with slew times up to 2 minutes (when there is a filter change); hence, it requires a scheduling algorithm that can explicitly minimize the slew times between successive observations. One could use the ILP scheduling approach to schedule large scripted blocks of observations for LSST. The blocks could be set to follow a path that only includes short slews (this is similar to the strategy taken in the scheduler developed by D. Rothchild et al.). The disadvantage of this approach and any other scripted schedule is principally the lack of recoverability from unpredictable interruptions, such as inability to dodge the clouds.

Even given the reliability of fully automatic scheduling technologies, there remain a number of modern telescopes, such as SALT, 2005 (Brink et al. 2008), and ALMA, 2013 (Wootten 2003), that are being operated based only on traditional handcrafted proposals. ALMA, in particular, requires a highly regulated structure for proposals that potentially leads to suboptimality as discussed in Alexander et al. (2017), where the authors suggest a number of corrections for the scheduling regulations to provide adaptivity to time-sensitive observations.

The LSST community, however, in addition to a proposal-based scheduler, introduced in Delgado & Reuter (2016), supported the design and implementation of proposal-free decision algorithms, such as the Feature-based scheduler, first introduced in Naghib et al. (2016), as well as a semiscripted cadence by D. Rothchild et al. (2019, in preparation).

In Section 2, first we explain the choice of Markovian framework for the Feature-based scheduler, and in Sections 2.1 and 2.2 we provide the mathematical details of the scheduler model in that framework. Section 3 presents two approaches for the optimization of the model's parameters. Sections 4

demonstrates the application of the Feature-based scheduler on LSST, which is then followed by a comparison between a modified version of the Feature-based scheduler and LSST's current official scheduler in Section 5. Finally, Section 6 presents our concluding remarks.

2. Scheduling Framework

To run a ground-based telescope with multiple science objectives, such as LSST, the scheduler has to offer *controllability*, *adjustability*, and *recoverability*.

1. *Controllability*: A telescope is controllable by a given scheduler if changes of the scheduler's design parameters are visibly manifested in the behavior and performance of the telescope. Controllability is determined by two factors: the information that is fed to the scheduler and the structure of the scheduler. Feeding the scheduler with information that is not sufficiently relevant to the high-level mission objectives leads to an output schedule that is irrelevant to the performance of the telescope. In addition, a scheduler with a flexible structure is needed to cover a large class of decision strategies, each determined by a set of its design parameters. Controllability is necessary for a scheduler to be optimizable by searching within the space of its design parameters.
2. *Adjustability*: A scheduler must be adjustable according to new conditions, environment, and scientific desirables. For a complex and multiobjective mission it is common that the scientific goals are required to be modified in the middle of the operational period. Another example is aging that causes changes in the mechanical characteristics of the telescope. Regardless of the reason, a scheduler is adjustable if the changes can be accommodated with reasonable computational cost, and preferably no or minimal human-expert intervention. Hand-tuned scheduling strategies and policies that are written in forms of instructions for human operator, or are based on observation proposals, are not fully adjustable.
3. *Recoverability*: The presence of unpredictable factors in the operation of ground-based telescopes is due to the natural stochastic processes (such as the weather) and the complexity of the mechanical facility. Unscheduled downtime and instrument failures are examples of the many unpredictable survey interruptions. In addition, there are inherently predictable interruptions, such as maintenance downtimes and cable winding, that, due to the complexity of the mechanical system, are not computationally affordable and/or valuable to keep track of. Therefore, they are considered to be stochastic variables as well. A scheduler is required to be able to make alternative decisions once a previously unpredictable event occurred. Moreover, it has to return to its optimal behavior shortly after the interrupting event is over. Such a scheduler is called recoverable, and its response to the interruptions is ideally as quick as the length of a single observation. The response time is usually limited by computational complexity of the scheduling algorithm. For instance, strategies that need to look back at the history of observations or look forward through possible sequences of observations are not fast recoverables owing to their computational

³ GitHub repository: https://github.com/ZwickyTransientFacility/ztf_sim.

complexity. Also, scripted sequences of observations, such as proposals, completely lack the attribute of recoverability.

The Feature-based scheduler is verifiably controllable. The Markovian framework provides a well-defined and tractable set of design parameters and a well-defined measure of the performance for any given choice of the design parameters; therefore, controllability is empirically verifiable. It is adjustable because the derivation of the design parameters is automatic and computationally tractable, once a new high-level mission objective is given or the telescope-environment system changes. Finally, it is recoverable owing to the inherent memorylessness of the Markovian Decision Process (MDP), which, for a decision at any time, only requires the current state of the system.

2.1. Markovian Representation

Definition 1. Let $X(\cdot)$ be a stochastic process for which X_i represents the state of the system at t_i , and let $\mathcal{S}$ be the set of all possible states that the system can take. Let $\mathbb{P}(X_i)$ be the probability distribution of X_i on $\mathcal{S}$. Then, $X(\cdot)$ is a Markovian process if and only if it satisfies the following *memorylessness* property:

$$\forall i \quad \mathbb{P}(X_{i+1}|X_i) = \mathbb{P}(X_{i+1}|X_i, X_{i-1}, \dots, X_0),$$

where $\mathbb{P}(X_{i+1}|X_i)$ is the conditional probability distribution of the system's state at t_{i+1} given its state at t_i , and $\mathbb{P}(X_{i+1}|X_i, X_{i-1}, \dots, X_0)$ is the conditional probability distribution of the system's state at t_{i+1} given all of the states that the system has been in until t_i .

The memorylessness property asserts that the system's next state depends only on its current state and is independent of its earlier history. This property is the main reason for choosing a Markovian framework for the scheduler.

Definition 2. Let $\langle \mathcal{S}, \mathcal{A}, P_a(\cdot, \cdot), R_a(\cdot, \cdot), \gamma \rangle$ be an MDP, where $\mathcal{A}$ is the set of actions and $P_a(x, y)$ is the transition probability from state x to y , which is equal to $\mathbb{P}(X_{i+1} = y|X_i = x, a)$, the conditional probability of transition from state x to state y given action $a \in \mathcal{A}$. Finally, the transition reward is denoted by $R_a(x, y)$, and $\gamma \in (0, 1]$ is the discount factor.

Definition 3. Action $a_i \in \mathcal{A}$ is *admissible* for $\langle \mathcal{S}, \mathcal{A}, P_a(\cdot, \cdot), R_a(\cdot, \cdot), \gamma \rangle$ if it is *feasible* (i.e., it is possible to be taken at t_i) and *progressively measurable* (i.e., depends only on the current state of the system, X_i , not on the future states).

To schedule a telescope is to take an admissible action (e.g., to determine the next observation) at all decision steps t_i . Notice that the decision steps t_i are not uniformly spaced. Decision steps are determined by the time that each observation takes. Also note that the scheduling is a finite time procedure from t_0 until the operation of the telescope is over, T . Our framework is general and covers all possible discretizations of the time in both finite and infinite horizon scheduling tasks.

Definition 4. A *deterministic policy* $\pi: \mathcal{S} \rightarrow \mathcal{A}_i$ is a mapping from $\mathcal{S}$ to the set of all admissible actions at t_i , denoted by $\mathcal{A}_i$.

A policy provides a time-invariant law that for all possible $x_i \in \mathcal{S}$ suggests an admissible action. The policy is the heart of the scheduler, which takes the necessary information encoded in the current state and makes a decision for the next observation. The design of an optimal scheduler is mainly to find an optimal policy.

Definition 5. A *deterministic optimal policy* π is a solution to the following optimization problem:

$$\underset{\pi}{\text{maximize}} \quad E_{\pi}[\sum_{i=0}^N \gamma^i R_{\pi(X_{i-1})}(X_{i-1}, X_i)|x_0], \quad (1)$$

where x_0 is a given initial state.

With this definition, we take the optimal policy to be a policy that maximizes the expected discounted sum of the rewards. The discount factor, $0 < \gamma \leq 1$, determines the priority of the overall gain (after an episode of observation) versus instant gains (after a single observation). Larger discount factors prioritize overall gains over instant gains. The choice of γ depends on the application and is empirically tuned.

Proposition 1. For the Markov decision process of $\langle \mathcal{S}, \mathcal{A}, P_a(\cdot, \cdot), R_a(\cdot, \cdot), \gamma \rangle$, there exists a deterministic optimal policy, and it can be written as follows:

$$\pi^* = \arg \min_{a_i \in \mathcal{A}_i} E[\Phi(X_{i+1})|a_i], \quad (2)$$

where $\Phi: \mathcal{S} \rightarrow \mathbb{R}$ is a function of the form

$$\Phi(x_i) = -R_{\pi^*(x_{i-1})}(x_{i-1}, x_i) + \gamma E_{\pi^*}[\Phi(X_{i+1})|x_i]. \quad (3)$$

Proof. See Appendix A.

For the telescope scheduler we require the policy to be deterministic, because the simulations have to be repeatable for comparison and evaluation purposes. However, it can be shown that the deterministic optimal policy is optimal not only among deterministic policies but also among stochastic policies. Therefore, the choice of deterministic policy does not harm the optimality of the scheduler.

As a result of Proposition 1, search for the optimal policy of problem (1) can be reduced from a search over the set of policies (all possible mappings from the state space to the action space) to a search over Φ functions, without loss of generality. This is a significant reduction made possible by the choice of Markovian Decision framework.

2.2. Markovian Approximation

For a decision that is inherently time dependent, such as scheduling an observation, only a maximal definition of the system's state yields a perfect Markovian system. The maximal definition of state space includes all of the possible decision sequences. In particular, LSST requires a sequence of about 1000 decisions at each night. Therefore, storing all possible scenarios requires a state space of size N_f^{1000} , where N_f is the

number of tessellation centers on the visible sky. For any N_f , N_f^{1000} number of scenarios is neither tractable nor storable in a realistic memory. In order to overcome the curse of dimensionality, we have designed a set of *features* to summarize the state of the system with only the most determining information. Thus, the system of telescope environment is only an approximated Markovian system once its state space is reduced to a *feature space*.

On the other hand, Proposition (1) shows that the search for an optimal scheduler lies within the set of functions instead of a much larger set of mappings from the state space to the action space. Despite this reduction, problem (1) is still an infinite dimensional optimization problem, because its variable is a function. To be able to numerically compute the Φ function, we assume that it can be expressed as a linear weighted summation of some basis functions. With this structure of Φ , our optimization would be further reduced to find the optimal values of the weights:

$$\tilde{\Phi}_\theta(x_i) := \sum_{j=1}^m \theta_j \Phi_j(x_i),$$

where θ is the vector of variables that fully characterizes $\tilde{\Phi}(\cdot)$, and $\Phi_j(x_i)$ are the building blocks of $\tilde{\Phi}$. We refer to $\Phi_j(x_i)$ as *basis functions*, which are handcrafted functions of the features and are designed in such a way that the domain knowledge of astronomical observation is incorporated into the decision-making strategy.

With this approximation, search space is reduced from the space of functions to a finite-dimensional vector space. This approximation replaces the original optimal policy (2) with the following approximate policy:

$$\begin{aligned} \tilde{\pi}_\theta^*(x_i) &= \arg \min_{a_i} E[\tilde{\Phi}_{\theta^*}(X_{i+1})|a_i] \\ &= \arg \min_{a_i} \sum_{j=1}^k \theta_j^* E[\Phi_j(X_{i+1})|a_i], \end{aligned} \quad (4)$$

where θ^* is a solution to the following optimization problem, in which policy π is fully determined by θ :

$$\underset{\theta}{\text{maximize}} \quad E_{\pi_\theta} \left[\sum_{i=0}^N \gamma^i R_{\pi_\theta}(X_{i-1}, X_i) | x_0 \right]. \quad (5)$$

Note that the computational time to find the optimal policy could be relatively long, but it is an offline task and can be done before the telescope starts to operate. However, evaluating the action in real-time must be at least as fast as the length of the shortest observation. A linear policy not only demands relatively small computational resources to be optimized but also is very quick to be evaluated in real time.

3. Scheduler Optimization

In this section we introduce two different approaches to solve problem (5). The solution of this problem is an optimal set of weights, θ^* , for a given set of basis functions. The first optimization approach is faster but requires the high-level mission objectives to belong to a certain class of functions. The second optimizing approach is applicable to all types of the high-level mission objectives but requires more computational resources.

3.1. Reinforcement Learning

Assume that there exists a well-defined notion of an instant reward for each state transition; then, $\Phi(x_i)$ by the definition given in Equation (3) is

$$\Phi(x_i) = -R_{\pi^*}(x_{i-1}, x_i) + \gamma E_{\pi^*}[\Phi(X_{i+1})|x_i]. \quad (6)$$

Accordingly, for the parameterized Φ function, we require the following:

$$\tilde{\Phi}_{\theta^*}(x_i) = -R_{\pi_{\theta^*}}(x_{i-1}, x_i) + \gamma E_{\pi_{\theta^*}}[\tilde{\Phi}_{\theta^*}(X_{i+1})|x_i]. \quad (7)$$

The main idea behind reinforcement learning is to optimize $\tilde{\Phi}_{\theta^*}(\cdot)$ while the decisions are being made in a simulated environment. First, the policy is initialized by an arbitrary set of variables, θ^0 . Then, the first decision is made by the policy associated with θ^0 . Then, based on the outcome reward, we update the initial set of variables to get θ^1 . Next, a decision is made with a policy associated with the new set of variables θ^1 , which yields a reward. By repeating this process, we gradually update the variables in each decision step until they converge to an optimal value. The optimal $\tilde{\Phi}$ respects Equation (7) for all $i \in \{j : t_0 \leq t^j \leq T\}$.

Note that at t_i , after the transition from x_{i-1} to x_i , we have the value of $R_{\pi_{\theta^i}}(x_{i-1}, x_i)$ already evaluated for making the decision. On the other hand, $\Phi_{\theta^i}(x_i)$ can be approximated by $E_{\pi_{\theta^i}}[\Phi_{\theta^i}(x_i)]$, which is also evaluated in the decision-making process, where θ^i is the last version of the optimization variables at t_i . Using the target value given in Equation (7), the update rule is as follows:

$$\begin{aligned} \tilde{\Phi}_{\theta^{i+1}}(x_i) &= (1 - \alpha) \tilde{\Phi}_{\theta^i}(x_i) + \alpha (\gamma E_{\pi_{\theta^i}}[\tilde{\Phi}_{\theta^i}(X_{i+1}) | \pi_{\theta^i}(x_i)] \\ &\quad - R_{\pi_{\theta^i}}(x_{i-1}, x_i)), \end{aligned} \quad (8)$$

in which $0 < \alpha < 1$ is the learning rate. The first term on the right-hand side of the equation is the most recent approximated value of $\tilde{\Phi}$ associated with θ^i . This term has $(1 - \alpha)$ amount of contribution to $\tilde{\Phi}$ of the next iteration. The second term is the value of the target $\tilde{\Phi}$ (according to Equation (7)), which has α amount of contribution to $\tilde{\Phi}$ of the next iteration. Clearly, with smaller α values this update imposes smaller adjustments from one decision to the next. The choice of α depends on the application. Higher α values are computationally preferred because they speed up the optimization; however, depending on the natural dynamic of the system and how much it changes from one decision to another, higher α values could make the process of optimization unstable and diverging. Accordingly, updates of the variables θ_j , $j = 1, \dots, k$ can be expressed as follows:

$$\begin{aligned} \theta_j^{i+1} &= \theta_j^i + (\tilde{\Phi}_{\theta^{i+1}}(x_i) - \tilde{\Phi}_{\theta^i}(x_i)) \Phi_j(x_i) \\ &= \theta_j^i + \alpha (-\Phi_{\theta^i}(x_i) + \gamma E_{\pi_{\theta^i}}[\tilde{\Phi}_{\theta^i}(X_{i+1}) | x_i] \\ &\quad - R_{\pi_{\theta^i}}(x_{i-1}, x_i)) \Phi_j(x_i). \end{aligned} \quad (9)$$

This variant of reinforcement learning is called *Temporal-Difference* (TD) learning with function approximation (Tsitsiklis & Van Roy 1997). Variants of this approach have been successfully applied to real-life problems such as training of a backgammon player (Tesauro 1995).

Recall that in order to be able to use the TD reinforcement learning method, it is necessary to have a well-behaved notion of a reward that reflects the instant gain of all possible decisions

at all of the decision steps. Moreover, the discounted sum of the instant rewards has to reflect the objective of the mission according to Equation (1). For instance, in the LSST scheduling problem, after each visit, the negative of the slew time is a well-defined instant reward that reflects how time-efficiently the telescope is being operated. This, however, does not reflect all aspects of the high-level mission's objective, such as the need to reobserve a field within a valid time window (explained in Section 4), for which there is no equivalent instant reward. For this reason, we also introduced a black-box function optimizer in the following section. We use the black-box approach when there is no well-defined notion of an instant reward with which we can build the high-level mission's objective function.

3.2. Global Optimization

In the absence of a well-defined instant reward, instead of solving problem (5), the following problem can be solved:

$$\max_{\theta} U_{\pi_{\theta}}(x_i, x_{i+1}, \dots, x_j), \quad (10)$$

where $U_{\pi_{\theta}}(x_i, x_{i+1}, \dots, x_j)$ is a utility function that measures the performance of the scheduler on a simulated episode of the operation from t_i to t_j . Unlike the previous method, the policy, π_{θ} , is fixed during an episode of simulation in this approach. The simulation episode is usually a small fraction of the total time of the telescope's operation. The idea is to test the performance of a set of candidate policies in a parallel manner, within a short episode. Then, we infer a better set of candidates and repeat until the performance cannot be improved. In general, $U(\cdot)$ cannot be explicitly expressed as a function of θ ; therefore, a global optimizer that can maximize a black-box function is required. Evolutionary optimizers have successfully been applied to numerous real-life problems involving black-box function optimization, and specifically astronomical mission planning such as the scheduling of the Exoplanet Characterization Observatory (Garcia-Piquer et al. 2015). We used the *e*DE evolutionary optimizer (Naghib & Nobakhti 2016), which is an adaptive version of the Differential Evolution (DE) algorithm (Storn & Price 1997). DE is generally one of the most efficient evolutionary algorithms, and the *e*DE variant uses a notion of entropy to automatically preserve the diversity of the candidate solutions. As a result, in contrast with DE, it does not require the user to choose any tuning parameters for the algorithm, which is the most time-consuming task in using an evolutionary optimizer. In addition, *e*DE, similar to many other evolutionary algorithms, is highly parallelizable, and the computational time can be almost linearly decreased with respect to the number of computational cores.

4. Problem of Scheduling for LSST

LSST's mission is to uniformly scan the visible sky within five different regions shown in Figure 1. Each region, also referred to as a *survey*, has certain science-driven goals and constraints, defined and precisely described in Ivezić et al. (2008). The notion of the features enables the scheduler to systematically fetch all of the various bits of information and turn them into comparable quantities for the purpose of the decision-making. The proposed feature space of LSST contains seven features, each of which can be evaluated given a field i , a

filter f , and a time t . The fields discretize the visible sky through a fixed sky tessellation, with N_f number of tessellation centers. There are six choices of filter, $[u, g, r, i, z, y]$, for each visit. Finally, the time is discretized by the natural timing of the process. In other words, the time interval between two decision steps is the time that it takes to execute the observation in between the two decisions. Given that a consecutive visit of the same field-filter is not allowed in the main survey, there is a slew time between any two decisions, and therefore $t_j - t_{j-1} > 0$. On the other hand, the operation is over a limited time horizon, T ; thus, the number of decision time steps is finite. In conclusion, a finitely discretized sky, a finite number of filters, and a finite number of time steps pose a finite feature space, denoted by $\{(f_1(i, f, t_j) \dots f_7(i, f, t_j)) : i = 1 \dots N_f, f \in \{u, g, r, i, z, y\}, t \in [t_0, T]\}$. With this feature space the implication of the policy, stated in Equation (4), is as follows:

$$\tilde{\pi}^*(x_j) = \arg \min_{(i,f) \in \mathcal{A}_j} \sum_{k=1}^5 \theta_k^* E_{\pi}[\Phi_k(X_{j+1}) | x_j], \quad (11)$$

where $x_j = [f_1, \dots, f_7(\text{id}(t_j), \text{ft}(t_j), t_j)]$ is a seven-dimensional state at t_j , and (i, f) is an admissible field-filter pair. Section 4.3 introduces the constraints under which a field-filter pair is admissible at t_j . Accordingly, $\mathcal{A}_j$ is the set of all field-filter pairs that are admissible at t_j .

In the implementation of the scheduling software we took a modular approach. The expected values of the basis functions, $E_{\pi}[\Phi_k(x_{j+1}) | x_j]$ for $k = 1, \dots, 5$, are evaluated in separate software routines. Then, they are delivered to the scheduler at the stage of the decision. The basis functions that address the environmental parameters are developed by the LSST community. For example, see Gressler et al. (2014) for the parameters that capture the status of the LSST site, Sebag et al. (2007, 2008) for cloud cover measurements that were used to develop a predictive cloud model, and Yoachim et al. (2016) for the sky brightness model.

Generally speaking, making a decision for a visit at t for the LSST scheduling problem is mainly determined by the following factors:

1. The amount of the time it takes to redirect the telescope and the dome to move from one target to the next target.
2. The short-term science-driven requirements, such as the same-night revisit of a field.
3. The long-term mission-driven requirements, such as maintaining a uniform coverage of all field-filter pairs within each region.
4. The observational quality of a field-filter at the time of decision, such as the expected depth of the resulting image.
5. The general preference for observing the fields around the meridian.

Accordingly, the basis functions of the LSST scheduler are designed to formalize the above factors. For the full definition of the basis functions of the Feature-based scheduler for LSST, refer to Section 4.2.

In what follows, first we precisely define the features and basis functions of the LSST scheduler, and then we show how the two optimization procedures, described in Section 3, are applied. The optimization results are associated with two sample objective functions. However, the LSST community, and principally any individual, can design their own mission

Table 1

Key Terms and Notations Used in the Definition of the Features and the Basis Functions

fields	Fixed-point configuration on the sky such that the visible sky could be completely covered by pointing the telescope toward all of those directions
N_f	Total number of the fields
t	Coordinated universal Time (UTC)
$\tau_s(t)$ ($\tau_e(t)$)	Beginning (end) of the night that t lies within
$\tau_{\text{rise(set)}}(i, f, t)$	Rising (setting) time of field-filter (i, f) above (below) the acceptable air-mass horizon at current night
$\tau_n(i, f, t)$	Time of the last visit of field-filter (i, f) before $\tau_s(t)$
$\text{id}(t)$	ID number of the field that is visited at t
$\text{ft}(t)$	Camera's filter at t
$n(i, f, t)$	Total number of the visits of field-filter (i, f) before t
$\text{slew}(i, j)$	Slew time from field i to field j in seconds
$\text{settling}(i, j)$	Mechanical settling time after slewing from field i to field j
Δt_f	Time needed to change filter
$t_{\text{dome}}(i)$	Time needed to move the dome to make field i visible to the telescope
$\text{ha}(i, t)$	Hour angle of the center of field i at t in hours, $-12 \leq \text{ha}(i, t) \leq 12$
$\text{am}(i, t)$	Air-mass of the center of field i at t
$\text{br}(i, t)$	Brightness of the sky at the center of field i at t
$\sigma(i, t)$	Seeing of the sky at the center of field i at t
$K(i, f, t)$	Atmospheric extinction coefficient at the center of field i at t
W_1, W_2	Given constant time window within which a revisit is valid

objective function and repeat the optimization procedure to obtain a scheduler that optimizes the probabilistic expectation of their objective function.

4.1. Features of the Telescope Scheduler

For designing the features, it is important to avoid redundancy in the information that features contain. It is also critical to hold a modular approach in the delivery of the information to the decision stage. For instance, consider the amount of time, Δt , it takes for a telescope to move from one visit to another. In the LSST problem, Δt mainly depends on the slew time, the mechanical settling time, the dome placement time, and the time it takes to change the filter. All of these timings are available through a precise simulation of the LSST model (Delgado et al. 2014). A modular design would be to bring the summation of the operational timings to the stage of the decision instead of bringing them separately as different features. This approach makes the implementation significantly simpler and more readable. Conceptually, the modular approach makes it possible to track the effect of the operational costs of a similar nature in the overall performance of the scheduling. Particularly, the operational cost is independent of the amount to which each cause contributes to the overall Δt . Hence, bringing the timing of each procedure separately in the decision-making level adds unnecessary complications to the design.

Using the notations in Table 1, we propose seven features for the LSST scheduler in Table 2. Features are designed to efficiently carry the determining information with a modular approach. Each feature is denoted by $f_k(i, f, t)$ for $k = 1 \dots 7$ and indexed by the triplet of (i, f, t), field, filter, and time. To make a decision at t , the scheduler computes all seven features for all of the (i, f) pairs. There are some features that do not change in every time step, for instance, if i is not visited at t^j ,

then $f_5(i, f, t^j) = f_5(i, f, t^{j+1})$. For such cases, the implementation has categorized updating routines to avoid redundant computations. For features such as the cloud coverage that are inherently random variables, we have separate predictive modules that evaluate the expectation of their values at the time of the next observation.

The main computational time during the real-time scheduling is spent on the evaluation of the features. Some features such as the slew times are essentially look-up tables. Some features, such as the co-added depth, are stored in the memory and indexed by a pair of field-filters, which the scheduler updates for only one field-filter after every decision step. The rest are continuous variables such as the location of the clouds, which can be interpolated from their past values, even if the exact value at a given time is not available.

4.2. Basis Functions of the Telescope Scheduler

Basis functions, Φ_k for $k = 1 \dots 5$, are fully determined by the value of the features. Hence, they are indexed in the same way that the features are, by a triplet of (i, f, t), field, filter, and time. Similar to the update of the features, all five basis functions should be evaluated for all pairs of (i, f) for a decision at time t , except that in this case admissibility of a field-filter can be evaluated (see Section 4.3 for the list of constraints that ensure admissibility). Therefore, while it is required to evaluate the features for all possible pairs of (i, f) at all decision steps, the number of basis function evaluations for admissible pairs of field-filter, is on average a factor of three times less than the number of feature computations.

Common basis functions are shared among all of the regions of the sky. They are designed to reflect the five general decision factors described in Section 4. The exact definitions of the common basis functions are presented in Table 3, and the key terms and notations can be found in Table 1.

LSST's mission poses different requirements on different regions of the sky, shown in Figure 1. First, we modify Φ_2 , for the Wide Fast Deep (WFD) and North Ecliptic Spur (NES) regions. Because they require the telescope to observe a field twice on the same night⁴ within a valid time window, $[W_1, W_2]$, the following modification prioritizes the fields that have received a first visit but not a second visit on the same night. There are two regimes to prioritize these fields, encoded in $\mathbb{1}_{(f_4 \geq W_2)}$. This function is zero if the expected remaining time for a field to become invisible, f_4 , is less than W_2 ; hence, the associated cost is minimal ($=0$) to ensure that this field will receive its second visit before it becomes invisible. If f_4 is larger than or equal to W_2 , it means that the expected remaining time for the field to become invisible is longer than the revisit deadline; therefore, $\mathbb{1}_{(f_4 \geq W_2)}$ is 1 and the cost is evaluated by a different regime defined by Φ^{pair} in the following:

$$\Phi_2^{\text{WFD}}(f_2, f_3, f_4) = \begin{cases} \Phi^{\text{pair}}(f_3) \mathbb{1}_{(f_4 \geq W_2)}, & \text{if } \sum_f f_2 = 1, \\ \Phi_2(f_2), & \text{else,} \end{cases}$$

where $\Phi^{\text{pair}}(f_3) = \exp(-\frac{\min_{\phi} f_3(i, \phi, t)}{W_2})$. Fields that have received their first visit of the night are distinguished via $\sum_f f_2 = 1$, the

⁴ Later, during the development of the scheduler software, it was decided that this constraint should be applied to all of the regions.

Table 2
Features of the Approximated Markovian Model for the Telescope-environment System

Notation	Definition/Description
$f_1(i, f, t)$	$(\text{slew}(\text{id}(t), i) + \text{settling}(\text{id}(t), i) + \Delta t_f I_{f(t) \neq f} \vee t_{\text{dome}}(i))$: either the time required to point the telescope to i and change the filter to f , or the time required to relocate the dome to make i visible, whichever is larger
$f_2(i, f, t)$	Total number of the same-night visits of field-filter (i, f) until t
$f_3(i, f, t)$	$(t - \tau_n(i, f, t)) I_{\theta(i, f, t) > \tau_n(i)}$, time since the last same-night visit of (f, i)
$f_4(i, f, t)$	Remaining time for field-filter (i, f) to become invisible, either by passing the air-mass or the moon-separation limit or by being covered by temporary objects such as clouds, as projected at t
$f_5(i, f, t)$	Co-added depth, a measure of cumulative quality of past visits of field-filter (i, f) until t
$f_6(i, f, t)$	5σ depth, a measure for quality of visiting field-filter (i, f) at t , depending on seeing, sky brightness, and air-mass; $f_6(i, f, t) = C_m + 2.5 \log(\frac{0.7}{\sigma(i, t)}) + 0.50(\text{br}(i, t) - 21) - K(i, f)\text{am}(i, t)$, where C_m is a scaling coefficient
$f_7(i, t)$	Hour angle of field i at t

Note. Features $f_{1 \dots 7}$ provide a memoryless, approximate description of the system's state.

sum of the number of visits in all filters. Finally, Φ^{pair} prioritizes the fields that received their first visit earlier.

The Deep Drilling Field (DDF) region contains a very small fraction of the visible sky's area (it is about 10 individual fields). Hence, it is unnecessary to adjust the basis functions that yield a separate generic policy for such fields. Instead, we treat the observation of DDFs as interruptions to the scheduler's regular operation, with each interruption comprising a sequence of DDF observations. This scheme is computationally more efficient and reduces the structural complexity of the scheduler. The recoverability attribute of the Feature-based scheduler enables the scheme of the interruptions to be a part of the optimal scheduling.

4.2.1. Controllability of the Scheduler

As discussed in Section 2, the scheduler is optimizable if the telescope is controllable. In this section we discuss the empirical controllability of LSST given the Feature-based scheduler. What we observe are the variations of two sample objective functions with respect to the variations of the design parameters, θ . If there are no meaningful variations, then the telescope is not controllable with this scheduler. As a result, the scheduler is not amenable to any form of optimization. On the other hand, if the objective function is extremely sensitive to the changes in the design parameters, the solution of the optimization is not fully reliable, because the objective is not a well-behaved function of the optimization variables.

In order to observe the variability of the objective functions with respect to changes in the scheduler's design parameters, first we defined a sequence of equidistant values for $\theta_i \in \{0, 0.5, \dots, 8.5, 9\}$ and kept the other θ values fixed at 2 or 5 or 8 in separate experiments. Then, we scheduled a simulated episode of observation for $t_n - t_0 = 4.8$ hr, for each θ_i separately. Finally, we evaluated the objective functions U_1 and U_2 of each schedule, where

$$U_1(x_0, x_1, \dots, x_n) = \sum_{\{i: t_0 < t^i < t_n\}} -\text{slew}(\text{id}(t_i), \text{id}(t_i)) - 10\text{am}(\text{id}(t_i)), \quad (12)$$

$$U_2(x_0, x_1, \dots, x_n) = n. \quad (13)$$

The first objective function reflects the slew time, $\text{slew}()$, and air-mass, $\text{am}()$, aversion. The second objective function, U_2 , reflects the time efficiency of the operation by counting the total number of observations. Note that the second objective function belongs to the class of black-box functions, and the first one belongs to the other class, where there exists a decomposition based on the discounted sum of the rewards.

Figure 2 contains slices of the five-dimensional U_1 and U_2 . Both of the simple objective functions reasonably respond to the changes in all five dimensions of the variable θ , which is the evidence of controllability. Moreover, the smaller variations for slices closer to the boundaries of the search space suggest that the design and normalizing the basis functions provide a desirable behavior within the proposed search space.

4.3. Survey-specific Constraints

The scheduler's decision at each time step is an admissible (feasible and measurable) pair of field-filter (i, f) . Feasibility of a candidate (i, f) is driven by the following factors:

1. Visibility: the candidate field-filter has to be visible.
2. Quality: the expected observational quality of a field-filter, such as the expected depth of the resulting image, has to be better than a given threshold.
3. Survey's timing: the science-driven revisit constraints have to be respected.

To ensure measurability, the above criteria must be evaluated based only on the information that is encoded in the feature space. Exact expressions of the proposed constraints for the LSST scheduler are presented in Table 4.

4.4. Scheduler Optimization

In this section we present two simple choices for the high-level mission objective in order to demonstrate the application of the proposed optimization approaches, discussed in Section 3. (More sophisticated mission objective functions can be defined based on LSST performance studies such as Grav et al. 2016, Graham et al. 2018, Jacklin et al. 2017, and Oluseyi et al. 2012). The choice of optimization algorithm depends on the nature of the mission objective. The first experimental mission objective function in this section can be expressed as the discounted sum of instant rewards $R(s_{i-1}, s_i)$;

Table 3
Basis Functions of the Feature-based Scheduler, the Building Blocks of the Decision Function

Notation	Definition/Description
$\Phi_1(f_1(i, f, t))$	$s_1 f_1(i, f, t)$, the cost of the required time for visiting field-filter (i, f) . Scale factor, s_1 , is chosen to be 0.43. This is based on the values of f_1 in our simulations and is chosen to ensure that 80% of the Φ_1 values are between 0 and 1. Scaling is a regulation that improves the rate of training convergence.
$\Phi_2(f_2(i, f, t))$	$\begin{cases} 0.5, & \text{if } \sum_f f_2(i, f, t) = 0 \\ 1, & \text{if } \sum_f f_2(i, f, t) = 2 \\ 0, & \text{else,} \end{cases}$ reflects the short-term visit/revisit priority of field i , conditioned on the total number of the previous same-night visits.
$\Phi_3(f_3(i, t))$	$\left(1 - \frac{f_3(i, f, t)}{\max_i \max_{\phi} f_4(i, \phi, t)}\right)$ reflects the long-term visit priority of field-filter (i, f) , based on the ratio of its co-added depth to the maximum co-added depth of all pairs of field-filters until t .
$\Phi_4(f_6(i, f, t))$	$1 - \Pr(f_6(i, \phi, t) \leq f_6(i, f, t))$, empirical complimentary cumulative density function (CDF) of 5σ -depth of all (i, f) pairs at t . Φ_4 assigns a cost to field-filter (i, f) based on its relative visiting quality compared to the other field-filter pairs at t .
$\Phi_5(f_6(i, f, t))$	$\frac{ hr(i, t) }{12}$ encourages visiting of the fields near the meridian.




Figure 1. Regions of the sky with different requirements and constraints for scheduling: (1) Galactic Plane Region (GP), (2) universal or Wide Fast Deep (WFD), (3) South Celestial Pole (SCP), (4) North Ecliptic Spur (NES), and (5) Deep Drilling Fields (DD).

thus, the reinforcement learning is applied to find the scheduler's parameters θ . The second objective function cannot be decomposed as a discounted sum of instant rewards; thus, we used the global optimizer approach. For the following experiment, the reinforcement learning is about 10 times faster than the global optimization and requires 50 times less memory. From the practical point of view, however, for some missions it is impossible to define an objective function that can be expressed as a discounted sum of a well-defined instant reward. In that case, the mission objective can be optimized only via a global optimizer.

In the following experiments, for both of the optimizations we used a simulated model of the telescope (Connolly et al. 2014) and the environment, including the brightness of the sky and coverage of the clouds that are developed based on the measurements at the LSST site.

4.4.1. Reinforcement Learning for the First Choice of Mission Objective

Let the instant reward, $R(i-1, i)$, be $-\text{slew}(\text{id}(t_{i-1})) - \text{am}(\text{id}(t_i))$. It is defined as a linear combination of the slew time to point the telescope from the $(i-1)$ th field to the i th field and the air-mass of the destination. Since both factors have a negative effect on the quality of the observation, the reward would be measured by the negative of each. Then, the mission objective function can be simply defined as $\sum_{i=0}^N \gamma^i R(i-1, i)$.

The simulation for the reinforcement learning starts at $t_0 = 2462867.5 \text{ MJD}$ (2021 January 1), with $\theta^0 = (5, 5, 5, 5, 5)$, initialized at the midrange values, and continues until θ converges. Figure 3 is the training curve for all of the

variables over a course of 3000 decisions. The discount rate $\gamma = 0.9$ and learning rate $\frac{0.01}{\log^3(i)}$ are chosen empirically.

For the above choices of reward, learning rate, discount factor, and initialization, θ converges to $\theta^* = (8.18, 1.04, 3.26, 7.59, 1.13)$. With a personal computer,⁵ each decision and its associated update take about 0.8 s;⁶ thus, the time of the convergence for the simulation presented in this section is $3000 \times 0.8 \text{ s} = 40 \text{ minutes}$. Note that the optimization time is linear with respect to the number of decision steps.

4.4.2. Global Optimization for the Second Choice of Mission Objective

One of the important simple objective functions that cannot be expressed as the discounted sum of the rewards is the total number of the observations from a given t_i to a given t_j , which can be expressed as $U_{\pi_{\theta}}(x_i, x_{i+1}, \dots, x_j) = j - i$. To find a set of parameters, θ , that optimizes this objective function, we applied the global optimization approach, explained in Section 3. To decrease the computational time, we used the following regulatory constraints:

1. $\theta \geq 0$: Positive coefficients for the basis functions are assumed in the design of the basis functions, because in the context of the telescope scheduling it is more natural to create the basis functions to reflect the cost of the operation.
2. $\theta_1 = \theta_0$: Without loss of generality, we fix the value of the first element of θ to reduce the dimension of the optimization problem by one, because homogeneity of the policy implies that if θ^* yields an optimal scheduler, then $\alpha\theta^*$ for $\alpha > 0$ yields an optimal scheduler as well.

We defined the above objective function, $U_{\pi_{\theta}}$, over a period of 10 days, from $t_i = 2462867.5 \text{ MJD}$ (2021 January 1) to $t_j = 2462877.5 \text{ MJD}$ (2021 January 11). Figure 4 shows the value of this objective function at each iteration of the eDE algorithm. The solution $\theta^* = (1.00, 0.84, 0.99, 1.34, 3.04)$ yields the best $U_{\pi_{\theta}}$ after 50 iterations.

eDE is a population-based metaheuristic algorithm. For our experiment, shown in Figure 4, the number of populations N_P is set to be 50. Each function evaluation is in fact the simulated operation of the telescope for 10 nights, with a candidate

⁵ Processor: 1.6 GHz, Memory 1600 MHz DDR3.

⁶ This reflects the efficiency of the first scheduler's prototype. The time of each decision simulation for the current efficient software is about 0.036 s.




Figure 2. One-dimensional slices of two objective functions, U_1 and U_2 , defined in Equations (12) and (13), respectively. The variation of the objective functions, especially in the midrange slices (solid line), suggests that LSST is controllable with the Feature-based Scheduler.

scheduler that takes about 8 minutes. Therefore, each iteration of the algorithm eDE takes $N_p * 8$ minutes (see footnote 5). The optimization can be manually terminated if the result is satisfactory, or it can be continued until a full convergence is achieved. In eDE (and all genetic algorithms in general), function evaluation for each individual is independent from other individuals; therefore, the parallel implementation of the same algorithm can be faster up to a factor of N_p .

5. Performance of a Modified Feature-based Scheduler for LSST

In this section, the LSST Metric Analysis Framework (Jones et al. 2014) is used, to compare⁷ the performance of a modified version of the Feature-based scheduler with opsim V4 and opsim V3, the most recent producers of the baseline schedules

of LSST. (We use two specific sequences produced by each of these schedulers, namely, astro-lsst-01-2013 and minion-1016 produced respectively by opsim V4 and opsim V3.)

The following results are based on the simulations that use a computer model of the LSST environment. The resulting schedule, however, will not be used in practice, because it is made by some realization of the random processes such as the cloud coverage. For the real-life operation of LSST, we use the same approach to optimize the scheduler; however, the decisions are made on the fly, using the real-time values of the features.

Recall that the optimization part consumes the main portion of the total computational time. Fortunately, the scheduler's optimization is an offline process and can be done before the telescope is ready to operate. On the other hand, the real-time decision-making with an already optimal scheduler is a very fast process. It only includes evaluation of the features and basis functions. The current version of the Feature-based scheduler can schedule 28 observations per second, while the length of each observation in real time is at least 30 s. Hence, the decision-making speed is roughly 900 times faster than what is required for the real-time scheduling. However, we are still working to increase the time efficiency of the software,

⁷ The sky background models and weather downtime used to benchmark the algorithms are not exactly identical because of the practical difficulties in the separation of the environment and the opsim scheduler. However, for the purpose of the comparisons in this section, the behavior of our sky and observatory model is sufficiently close to the official model. See Delgado & Reuter (2016) and Reuter et al. (2016) for the official operation simulator.

Table 4
Feasibility of Field Filter (i, f) for a Visit at t^{n+1}

Constraints	Description	Region
1 $\tau_{\text{rise}}(i, f, t^{n+1}) \leq t^{n+1} \leq \tau_{\text{set}}(i, f, t^{n+1})$	Field filter (i, f) has to be above the acceptable air-mass horizon at t^{n+1} .	All regions
2 $E[f_4(i, f, t^{n+1})] \neq 0$	Field filter (i, f) is not temporarily masked (e.g., by the moon) at t^{n+1} .	All regions
3 $\sum_f f_2(i, f, t^n) < N^{\text{survey}}$	N^{survey} poses a region-dependent upper bound on the number of the visits for each field; $N^{\text{WFD}} = N^{\text{NES}} = 3$, and $N^{\text{GPR}} = N^{\text{SCP}} = 1$	All regions
4 $E[f_6(i, f, t^{n+1})] < \sigma(\text{survey}, f)$	The expected quality of visiting field-filter (i, f) at t^{n+1} has to be better than the given threshold, $\sigma(\cdot)$, which depends on the survey and the filter	All regions
5 $f \neq \text{id}(t^n)$	Consecutive visit of the same field is not allowed	All
6 if $\sum_f f_2(i, f, t^n) = 0$ then $\max_{\phi} f_4(i, \phi, t^n) > \frac{W_1 + W_2}{2}$	The first visit of field f has to occur $\frac{W_1 + W_2}{2}$ times before it becomes invisible, so that the second visit of f can be scheduled in the valid time window	except DDF WFD and NES ^a
7 if $\max_{\phi} \theta(i, \phi, t^n) > \tau_s(t^n)$ then $W_1 \leq \min_{\phi} f_2(i, \phi, t^n) \leq W_2$	If there has been a same-night visit of field f until t^n , then the next same-night visit has to occur in the valid time window	WFD and NES
8 if $\max_{\phi} \theta(i, \phi, t^n) > \tau_s(t^n)$ then $f \notin \{y, u\}$	If there is a same-night visit of field f until t^n , then the next same-night visit cannot be with either of u or y filters	WFD
9 $f \notin \{y, u\}$	Visits with u filter and y filter are not allowed	NES

Notes. We use the probabilistic expectation, $E[\cdot]$, of the stochastic values, as evaluated at t^n .

^a Later, during the development of the scheduler software, it was decided that this constraint and the seventh constraint should be applied to all of the regions.




Figure 3. Reinforcement learning of the scheduler's parameters, ($\theta_1, \theta_2, \theta_3, \theta_4, \theta_5$). All of the parameters are initialized at the midrange value, 5. During the simulation, after each decision step there is a reward associated with the outcome of the decision that implies a small adjustment on each of the five design parameters. Then, the next decision will be taken with a slightly different set of parameters. This procedure continues until the adjustments on the design parameters are negligible.

because it directly effects the computational time of the offline optimization procedure.

The *Modified Feature-based* scheduler is under active development,⁸ and it addresses the observational details of LSST's mission through the adjustment of the constraints and the basis functions. It is designed to produce a software that can be used in practice.

Unlike the default sky tessellation for LSST, we do not require the tessellation centers to be determined based on the telescope's field of view. The default sky tessellation adopted in the baseline scheduler results in 23% of the sky being




Figure 4. Progress of the black-box objective function $U_{\pi_{\theta}}$ over the iterations of the eDE algorithm. $U_{\pi_{\theta}}$ for this simulation is the total number of the observations for 10 nights starting from 2021 January 1.

covered by more than one field, which causes serious nonuniformity in the final coverage of the sky. In the Modified Feature-based scheduler, we adopt a 50 times finer discretization of the sky using Hierarchical Equal Area isoLatitude Pixelization (HEALpix; Górski et al. 2005). The ability to adopt a finer tessellation is possible only because the real-time decision making is computationally cheap and fast. This approach allows the scheduler to handle the overlaps that cause inhomogeneity of the final sky coverage.

In addition to adopting a finer discretization, we use a spatial dithering scheme to randomize the final pointing of the telescope by a small amount around the tessellation centers to further assist the homogeneity of the coverage. Adopting the dithering scheme, the median number of observations at a typical point in the sky increases by $\sim 15\%$. Dithering is also essential for removing systematic effects for science cases such as measuring galaxy counts (see Awan et al. 2016 for more details). Moreover, the Modified Feature-based scheduler uses separate routines to decide whether an observation will need to be observed in a pair and to decide whether a Deep Drilling

⁸ GitHub repository: https://github.com/lst/sims_featureScheduler.




Figure 5. Sky coverage in each of the six filters $\{u, g, r, i, z, y\}$, measured by co-added depth. According to the mission’s objective, the scheduler has to provide a uniform coverage of the visible sky within each region and in each filter. The left panels show the opsim V4 simulation results, while the Modified Feature-based scheduler is on the right. Even without a given observation proposal, the Modified Feature-based scheduler can closely match the large-scale footprint of the official survey. (The individual yellow dots with high co-added depth are fields in the deep drilling survey, and the NES region is not visited in the u and y bands.)




Figure 6. Co-added depth coverage in a zoomed-in area of the sky map, around the border between the WFD (green) and SCP (blue) regions, in the r band. The smooth coverage of the Modified Feature-based scheduler (right) vs. the granular pattern of opsim V4 (left) further respects the uniformity of the coverage, which is one of the most fundamental objectives of the LSST mission. The middle panel is the coverage of opsim V4, with dithering of the same sequence of observations that fundamentally cannot become as smooth as the right panel, because, unlike the Feature-based scheduler, the scheme of opsim V4 does not easily allow for decision-making with arbitrarily fine tessellations of the sky.

sequence should be executed by interrupting the normal operation of the telescope.

5.1. Sky Coverage Uniformity

For a survey telescope, such as LSST, the density of the co-added depth over the visible sky should ideally be uniform in each filter and within each of the five survey regions.

Figure 5 compares the values of the co-added depth on a discretized full sky map. Figure 6 demonstrates the smoothness of the coverage by zooming in to a smaller area of the sky map. It shows the co-added depth values of opsim V4, with and without dithering, and compares them with that of the Modified Feature-based scheduler in the r band. The zoomed-in area is around the boundary of WFD and GP regions, which have different target co-added depths. The smoother coverage that




Figure 7. Each plot compares the distributions of the co-added depth coverage in one of the six filters. A dithering scheme in the Modified Feature-based scheduler in addition to a finer tessellation of the sky smoothens the density of the coverage where the fields overlap.

Table 5
The Median and Variance of the Co-added Depth Distribution on a Finely Discretized Sky

Median, Variance								
Filter	Opsim V4				WFD	Modified Feature-based		
	WFD	GP	SCP	NES		GP	SCP	NES
<i>u</i>	25.63, 0.04	25.12, 0.11	24.91, 0.12	...	25.68, 0.01	25.32, 0.05	25.10, 0.04	...
<i>g</i>	27.13, 0.04	26.41, 0.09	26.32, 0.12	26.30, 0.13	27.18, 0.01	26.69, 0.04	26.56, 0.04	26.47, 0.09
<i>r</i>	27.19, 0.04	26.01, 0.16	25.84, 0.24	26.38, 0.12	27.14, 0.01	26.21, 0.07	26.08, 0.05	26.43, 0.09
<i>i</i>	26.60, 0.04	25.44, 0.15	25.28, 0.22	25.82, 0.12	26.56, 0.01	25.68, 0.08	25.43, 0.07	25.88, 0.09
<i>z</i>	25.73, 0.04	24.62, 0.17	24.57, 0.21	24.90, 0.14	25.87, 0.01	25.03, 0.09	24.81, 0.05	25.16, 0.10
<i>y</i>	24.92, 0.04	23.81, 0.16	23.72, 0.21	...	24.92, 0.02	24.01, 0.09	23.88, 0.06	...

Note. The Modified feature-based scheduler closely matches the footprint of the official survey and also outperforms opsim V4 in terms of the uniformity of the coverage, with lower variances, especially in WFD and SCP regions.

the Modified Feature-based scheduler offers is due to the fine discretization of the sky in the decision-making stage, in addition to the dithering that takes place after the decision is made.

Figure 7 compares the distribution of the co-added depth on a (finely) discretized sky. The Modified Feature-based scheduler has paved the rightmost peak that appears in the distribution of opsim V4. This peak is the result of the field’s undesirable overlaps. Table 5 contains the median and variance of the co-added depth for both of the schedules in each of the sky regions and in each filter. The Modified Feature-based scheduler provides deeper (higher median) and more uniform (lower variance) coverage in most of the cases.

5.2. Pairs

In addition to uniformity of the coverage, the LSST mission calls for pairs of visits within a valid time window on the same

night. The main reason is to detect transient objects such as asteroids. Because the moving objects usually belong to the solar system, the pair constraint was initially imposed only on the WFD and NES regions. However, there are interesting solar system objects, such as interstellar asteroids, that can be observed in any direction of the sky. In addition, identification of the other varying objects, such as supernovae, can benefit from a follow-up visit, especially if the second visit is with a different filter. Thus, in the Modified Feature-based scheduler we made the pair constraint a universal constraint for all of the regions. The downside of this extension is the fact that it constrains the scheduler even more and the performance can be potentially less than it could be.

Note that the structure of the Feature-based scheduler allows for extension or restriction of the constraints down to the level of an individual field-filter. These extensions and restrictions do




Figure 8. Ratio of the pairs (in the g , r , and i filters) to the total number of the observations on the sky map. For the areas to which the pair constraint is applied, the ratio is desired to be 1. For the Modified Feature-based scheduler (right), the pair constraint is applied to all of the regions, whereas for opsim V4 (left), it is applied to the WFD and NES regions only.

not contradict any of the Markovian assumptions, and do not break the structure of the implementation.

Figure 8 demonstrates the distribution of the ratio of observations in pairs (in the g , r , and i filters) to the total number of the observations. For the regions to which the pair constraint is applied, this ratio can be interpreted as the success rate of the schedulers in satisfying the pair constraint. This success rate ranges from 0 to 1, and the higher values indicate the more successful pair visits.

Figure 9 compares the distribution of the pair ratio of the Modified Feature-based scheduler and opsim V4 on a (finely) discretized sky. Note that the peak of the distribution for the Modified Feature-based scheduler is closer to 1, which means that a larger area of the sky is covered by a successful pair visit; however, opsim V4 offers a sharper concentration of the values that can be interpreted as a more homogeneous pair visit, which agrees more with LSST’s mission to survey the sky in a uniform manner.




Figure 9. Distribution of the pair ratio to the total number of the observations, in the g , r , and i filters. The Feature-based scheduler covers a large area of the sky by a successful pair visit (higher median), whereas opsim V4 maintains a uniform ratio of pairs for a larger area of the sky (lower variance).

5.3. AltAz and Air-mass Distributions

Air-mass is one of the major obstacles for ground-based instruments. Zenith observations have the minimum air-mass; however, off-zenith observations cannot be avoided, for time efficiency and due to design limitations. In this case, observations around the meridian provide high-quality images and consequently result in more efficient operation of the instrument. Figure 10 compares the number density of the visits on an altitude-azimuth sky map in each of the six filters [u , g , r , i , z , y]. Clearly, in all of the filters the Modified Feature-based scheduler schedules more visits around the meridian zone. In addition, it offers a consistent concentration peak on the east wings, which is essential for a higher success rate of the pair constraint. If the first visit of the night occurs when the field is on the east side of the sky, the scheduler has a longer opportunity to schedule the second visit of the night, and hence it is more likely to perform a successful pair visit. Figure 11 demonstrates the density of visits collectively in all filters for opsim V3, opsim V4, and the Modified Feature-based scheduler. Note that the adjustability of the Feature-based scheduler allows for a significant change in the behavior of the telescope; in this case, to encourage observations around the meridian, we defined a new basis function in such a way that the Modified Feature-based scheduler prefers to observe a

contiguous set of fields that is then reobserved later in the same order.

5.4. Signal-to-noise Ratio

For a multiobjective survey telescope, such as LSST, comparing the overall performance of the different schedules is a difficult task, particularly because of the large number of competing factors that are involved in the evaluation of their performance. In some cases these criteria are not even objective or well defined, such as the importance of an area of astronomy compared to the rest. Nevertheless, we conclude this section with a general comparison of the performances using the overall signal-to-noise ratio of the schedules. Table 6 reflects the value of median throughput for three different schedulers in r and g bands. The Modified Feature-based scheduler significantly outperforms the other two.

The throughput is mainly determined by the combination of open shutter fraction (OSF) and air-mass. The OSF is the total time that the telescope camera shutter was open divided by the maximal time it could have been open. This reveals how time-efficiently the observations have been scheduled. The median air-mass reflects the overall quality of the collected data. As mentioned before, observations in lower air-mass allow for a higher data quality. Comparing the values of the OSF and air-mass for both of the baseline schedules, opsim V3 and opsim




Figure 10. Each plot is the distribution of the visits on an altitude-azimuth sky map in one of the six filters. The two left columns belong to opsim V4, and the two right columns belong to the Modified Feature-based scheduler. The higher concentration on the meridian (vertical axis) for the Modified Feature-based scheduler shows a more desirable behavior. Moreover, consistent concentration of the visits on the east wing can potentially provide a better success rate in pair observation.

V4, shows that there is a trade-off between the two values. While opsim V4 offers better median air-mass, its OSF is worse. However, its median throughput is very close to that of opsim V3. This comparison reveals that the change of meta-parameters, as well as tweaking the structures of proposal-based schedulers such as opsim V3 and opsim V4, only changes the balance of trade-off between OSF and air-mass, but not the actual performance of the scheduling.

6. Concluding Remarks

This study demonstrated that a Markovian scheduler with expert-designed features and a parameterized linear decision-making policy can be successfully applied to multimission, ground-based telescopes such as LSST. Unlike the mainstream telescope schedulers, the Feature-based scheduler does not rely on handcrafted observation proposals. Instead, bringing the decision-making process to the individual observation level improves the efficiency of the telescope’s operation.

In particular, while proposals are designed for feasibility, our approach is designed for optimality in addition to feasibility. In general, automatic schedulers such as the Feature-based scheduler are fundamentally less prone to suboptimality compared to the schedulers that heavily rely on human interaction. This is mainly due to the manual tailoring that is performed based on the inspections of the instances. Moreover, adjusting the behavior of human-dependent schedulers is inconvenient and time-consuming in practice. Furthermore,

being modeled as an MDP, the Feature-based scheduler is a systematic approach to operating the instrument under uncertainties and interruptions.

On the other hand, the building blocks of the Feature-based scheduler are designed modularly in an intuitive way for the astronomy community. This property allows for expert intervention if needed. However, the robust structure and implementation of the scheduler allow regulated interventions only. In particular, valid interventions include defining new features, new basis functions, and adjusting the design parameters, none of which break the measurability, linearity, and memorylessness of the scheduler. As a result, all of the desirable properties, such as simplicity, optimality conditions, and modularity of the design, remain valid.

In addition, due to the coherent structure, from training to online decision-making, the Feature-based scheduler is easy to understand, implement, and troubleshoot. Simplicity of the design and implementation also provides a user-friendly environment for a wide range of programming expertise to define a custom mission objective, train a scheduler, and examine the behavior of the scheduler with various mission objectives. Similarly, in a particular project, when a change in the mission’s objective is necessary, deriving a new scheduler that optimizes the new objective is principally automated. Furthermore, for the mission-planning stage of a future instrument, a scheduler with an adjustable objective can be extremely helpful, because it can answer the high-level




Figure 11. Distribution of the visits on an altitude-azimuth sky map for three schedulers. On the left, opsim V3 schedules most observations at high air-mass and very few on the meridian. In the middle, opsim V4 schedules many observations on the meridian but still executes deep drilling fields and a smattering of other observations at high air-mass. On the right, the Modified Feature-based scheduler concentrates observations around the meridian, including deep drilling observation.

Top view of the LSST site on 2015/7/2 09:02:07, GMT




Figure 12. The moving circular mark is where the telescope is pointed to at any given moment, which is shown in the title. Those parts of the sky that are covered during the same night are star-like colored marks, and the color reflects the filter with which the observation is performed. The blue colored areas are the regions that are observed a night before the current night. The gray areas are where the sky is bright and generally no observation is allowed, for instance, the halo around the moon. The white areas show where the clouds are based on a realistic model of the cloud covers at the LSST site in Cerro Pachon. Finally, the light-orange areas, which do not change as the time goes by, indicate the four main sky regions (Galactic Plane Region (GP), Universal or Wide Fast Deep (WFD), South Celestial Pole (SCP), North Ecliptic Spur (NES)).

(An animation of this figure is available.)

Table 6

Comparison of the Schedulers in a Section of Main Survey Area, WFD

Survey	Median Throughput		OSF	Median Air-mass	Dithered
	r (%)	g (%)			
Modified Feature-based	63.7	47.0	0.705	1.1	yes
opsim V3	55.3	40.0	0.736	1.2	no
opsim V4	54.4	40.8	0.715	1.1	no

trade-off questions, such as the time efficiency of the different proposed strategies.

Computationally, the optimization part of the scheduler is the most demanding task. Once the optimal values of the design parameters are known, one night of observation can be simulated

in 2–3 minutes. Equivalently, in real-time operations, decisions are made on the fly about 900 times faster than what is needed. For the optimization part, however, the necessary computational time and resources vary. If many different objective functions are being tested for planning a mission, then a quick eDE optimization, which takes a few hours, can find a sufficiently good scheduler for each mission. Even a quick manual hand tuning that reflects the intuitive importance of each basis function is possible, because they carry an astronomical meaning. On the other hand, if the objective is known and fixed and the scheduler is being trained for real-time decision-making, then one might even categorize the observation nights based on their main differences, such as the moon phase, seasonal variations, and weather patterns, and then train a scheduler specifically for each category to further increase the efficiency of the operation. A simulation of the feature-based scheduler for one week of observation is provided with this paper. Figure 12 explains the environment of this simulation.

We would like to thank the LSST team and the DIRAC Institute’s faculty and researchers for providing expertise without which this research would not have been possible, in particular Professor Željko Ivezić, whose consistent attention and insightful comments greatly assisted this work. We would also like to show our gratitude to the experts in other institutes, especially Professor Michael Strauss, Professor Christopher Stubbs, Dr. Robert Lupton, and Dr. Michael Reuter, for their insightful comments spanning from the idea stage to the end of this project. We are also grateful for the time that the anonymous referees took to provide both high-level and detailed comments that significantly improved the clarity of the paper. This work is financially supported by the National Science Foundation under 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. The DIRAC Institute is supported through generous gifts from the Charles and Lisa Simonyi Fund for Arts and Sciences and the Washington Research Foundation.

Facility: LSST.

Software: MAF (Jones et al. 2014), gorski2005healpix (Górski et al. 2005), matplotlib (Hunter 2007), astropy (The Astropy Collaboration et al. 2018), numpy/scipy (Jones et al. 2001).

Appendix A

Proof of the Main Proposition

Proof (of Proposition 1). Consider a function $C_\pi : \mathcal{S} \rightarrow \mathbb{R}$, defined as follows:

$$\begin{aligned} C_\pi(x_i) &= -E_\pi \left[\sum_{i \leq j}^N \gamma^{i-j} R_{\pi(x_j)}(X_j, X_{j+1}) | x_i \right] \\ &= -E_\pi [R_{\pi(x_i)}(x_i, X_{i+1}) | x_i] \\ &\quad - \gamma E_\pi \left[\sum_{i+1 \leq j}^N \gamma^{j-(i+1)} R_{\pi(x_j)}(X_j, X_{j+1}) | x_i \right], \end{aligned}$$

where $E_\pi[\cdot]$ is the expectation of the argument inside the brackets, if policy π is being used. Then, by applying the law of total expectation on the second term,

$$\begin{aligned} C_\pi(x_i) &= -E_\pi [R_{\pi(x_i)}(x_i, X_{i+1}) | x_i] \\ &\quad - \gamma E_\pi \left[E_\pi \left[\sum_{i+1 \leq j}^N R_{\pi(x_j)}(X_j, X_{j+1}) | X_{i+1} \right] | x_i \right] \\ &= -E_\pi [R_{\pi(x_i)}(x_i, X_{i+1}) | x_i] - \gamma E_\pi [C_\pi(X_{i+1}) | x_i]. \end{aligned}$$

By assuming a finite state space, one can expand the expectation as a finite sum,

$$\begin{aligned} C_\pi(x_i) &= -E_\pi [R_{\pi(x_i)}(x_i, X_{i+1}) | x_i] \\ &\quad - \gamma \sum_{x_{i+1} \in \mathcal{S}} \mathbb{P}(x_{i+1} | \pi(x_i), x_i, x_{i-1}, \dots, x_0) C_\pi(x_{i+1}). \end{aligned}$$

Then, by the Markov property, probabilities that are conditioned on the whole previous sequence of states can be replaced by probabilities that are conditioned only on the previous state,

$$\begin{aligned} C_\pi(x_i) &= -E_\pi [R_{\pi(x_i)}(x_i, X_{i+1}) | x_i] \\ &\quad - \gamma \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) C_\pi(x_{i+1}), \end{aligned}$$

where $P_{\pi(x_i)}(x_i, x_{i+1})$ is the transition probability from x_i to x_{i+1} , under the outcome action of policy π . Now, let $C^*(x_i) = \min_\pi C_\pi(x_i)$; then,

$$\begin{aligned} C^*(x_i) &= \min_\pi (-E_\pi [R_{\pi(x_i)}(x_i, X_{i+1}) | x_i] \\ &\quad - \gamma \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) C_\pi(x_{i+1})) \\ &= \min_{\{\pi(x_i), \pi(x_{i+1}), \dots\}} (-E_\pi [R_{\pi(x_i)}(x_i, X_{i+1}) | x_i] \\ &\quad - \gamma \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) C_\pi(x_{i+1})) \\ &= \min_{\pi(x_i)} (-E_\pi [R_{\pi(x_i)}(x_i, X_{i+1}) | x_i] \\ &\quad - \gamma \min_{\{\pi(x_{i+1}), \pi(x_{i+2}), \dots\}} \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) C_\pi(x_{i+1})). \end{aligned} \quad (14)$$

For the next time step, one can construct a function, $\hat{C}$, such that $\hat{C}(x_{i+1}) = \min_{\pi(x_{i+1})} \min_\pi C_\pi(x_{i+1})$, and then

$$\begin{aligned} &\sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) \min_{\{\pi(x_{i+1}), \pi(x_{i+2}), \dots\}} C_\pi(x_{i+1}) \\ &= \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) \hat{C}(x_{i+1}) \\ &\geq \min_\pi \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) C_\pi(x_{i+1}). \end{aligned}$$

On the other hand,

$$\begin{aligned} &\sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) \min_{\{\pi(x_{i+1}), \pi(x_{i+2}), \dots\}} C_\pi(x_{i+1}) \\ &\leq \min_\pi \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) C_\pi(x_{i+1}). \end{aligned}$$

Therefore,

$$\begin{aligned} &\min_\pi \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) C_\pi(x_{i+1}) \\ &= \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) \min_{\{\pi(x_{i+1}), \pi(x_{i+2}), \dots\}} C_\pi(x_{i+1}). \end{aligned}$$

By replacing the second term on the right-hand side of Equation (14) with the right-hand side of the above equation,

$$\begin{aligned} C^*(x_i) &= \min_{\pi(x_i)} (-E_\pi [R_{\pi(x_i)}(x_i, X_{i+1}) | x_i] \\ &\quad - \gamma \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) \min_\pi C_\pi(x_{i+1})) \\ &= \min_{\pi(x_i)} (-E_\pi [R_{\pi(x_i)}(x_i, X_{i+1}) | x_i] \\ &\quad - \gamma \sum_{x_{i+1} \in \mathcal{S}} P_{\pi(x_i)}(x_i, x_{i+1}) C^*(x_{i+1})). \end{aligned}$$

The last equality follows from the definition of C^* and is in the form of the Optimal Bellman Equation, for which a solution exists. Moreover, $C^*(x_0) = \min C^\pi(x_0)$ attains the optimal value by the construction of $C^\pi(x_0)$, which is equal to $-E[\sum_{\pi(x_i)}(x_i, X_{i+1}) | x_0]$.

Now, given a C^* , an optimal policy can be simply evaluated,

$$\begin{aligned} \pi^*(x_i) &= \arg \min_{a_i \in \mathcal{A}_i} (-E_\pi [R_{a_i}(x_i, X_{i+1}) | x_i] \\ &\quad - \gamma \sum_{x_{i+1} \in \mathcal{S}} P_{a_i}(x_i, x_{i+1}) C^*(x_{i+1})) \\ &= \arg \min_{a_i \in \mathcal{A}_i} (-E_\pi [R_{a_i}(x_i, X_{i+1}) - \gamma C^*(X_{i+1}) | x_i]). \end{aligned} \quad (15)$$

Finally, define

$$\Phi(X_{i+1}) := R_{a_i}(x_i, X_{i+1}) - \gamma C^*(X_{i+1}) \quad (16)$$

and replace $R_{a_i}(x_i, X_{i+1}) - \gamma C^*(X_{i+1})$ with $\Phi(X_{i+1})$ in Equation (15) to complete the proof. $\square$

Appendix B

Animation

With the online version of this article we provide an instance of the Feature-based scheduling simulation for seven nights of observation.

ORCID iDs

Elaheasadat Naghib  <https://orcid.org/0000-0002-4485-2673>
 Peter Yoachim  <https://orcid.org/0000-0003-2874-6464>
 Andrew J. Connolly  <https://orcid.org/0000-0001-5576-8189>
 R. Lynne Jones  <https://orcid.org/0000-0001-5916-0031>

References

- Alexander, K. D., Berger, E., Bower, G., et al. 2017, arXiv:1703.04692
- Awan, H., Gawiser, E., Kurczynski, P., et al. 2016, *ApJ*, **829**, 50
- Bellm, E. 2014, arXiv:1410.8185
- Boroson, T., Brown, T., Hjelmstrom, A., et al. 2014, *Proc. SPIE*, **9149**, 91491E
- Brink, J., Charles, A., Hettlage, C., et al. 2008, *Proc. SPIE*, **7019**, 70190N
- Connolly, A. J., Angeli, G. Z., Chandrasekharan, S., et al. 2014, *Proc. SPIE*, **9150**, 14
- Delgado, F., & Reuter, M. A. 2016, *Proc. SPIE*, **9910**, 991013
- Delgado, F., Saha, A., Chandrasekharan, S., et al. 2014, *Proc. SPIE*, **9150**, 915015
- Garcia-Piquer, A., Ribas, I., & Colomé, J. 2015, *ExA*, **40**, 671
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Graham, M. L., Connolly, A. J., Ivezić, Ž., et al. 2018, *AJ*, **155**, 1
- Grav, T., Mainzer, A. K., & Spahr, T. 2016, *AJ*, **151**, 172
- Gressler, W., DeVries, J., Hileman, E., et al. 2014, *Proc. SPIE*, **9145**, 91451A
- Gurobi Optimization, I. 2016, Gurobi Optimizer Reference Manual, <http://www.gurobi.com>
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Ivezic, Z., Tyson, J. A., Abel, B., et al. 2008, arXiv:0805.2366
- Jacklin, S. R., Lund, M. B., Pepper, J., & Stassun, K. G. 2017, *AJ*, **153**, 186
- Johnston, M. D., & Miller, G. 1994, in *Spike: Intelligent Scheduling of Hubble Space Telescope Observations* (Burlington, MA: Morgan Kaufmann), 391
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open source scientific tools for Python, <http://www.scipy.org/>
- Jones, R. L., Yoachim, P., Chandrasekharan, S., et al. 2014, *Proc. SPIE*, **9149**, 91490B
- Lampoudi, S., Saunders, E., & Eastman, J. 2015, arXiv:1503.07170
- Muscettola, N., Pell, B., Hansson, O., & Mohan, S. 1995, in *ASP Conf. Ser. 79, Robotic Telescopes. Current Capabilities, Present Developments, and Future Prospects for Automated Astronomy*, ed. G. W. Henry & J. A. Eaton (San Francisco, CA: ASP), 148
- Naghib, E., & Nobakhti, A. 2016, in *American Control Conf. (ACC)* (Piscataway, NJ: IEEE), 2440
- Naghib, E., Vanderbei, R. J., & Stubbs, C. 2016, *Proc. SPIE*, **9910**, 991011
- Nelson, J. E., Mast, T. S., & Faber, S. M. 1985, *The Design of the Keck Observatory and Telescope*, Keck Observatory Report 90
- Nowakowski, J., Schwärzler, W., & Triesch, E. 1999, *European Journal of Operational Research*, **112**, 531
- Oluseyi, H. M., Becker, A. C., Culliton, C., et al. 2012, *AJ*, **144**, 9
- Reuter, M. A., Cook, K. H., Delgado, F., Petry, C. E., & Ridgway, S. T. 2016, *Proc. SPIE*, **9911**, 991125
- Saunders, E., & Lampoudi, S. 2014, in *The Third Hot-wiring the Transient Universe Workshop*, ed. P. R. Wozniak (Los Alamos, NM: Los Alamos National Laboratory), 117
- Sebag, J., Barr, J., Saha, A., Claver, C., & Krabbendam, V. 2007, *RMxAC*, **31**, 30
- Sebag, J., Krabbendam, V. L., Claver, C. F., et al. 2008, *Proc. SPIE*, **7012**, 70123W
- Shetrone, M., Cornell, M. E., Fowler, J. R., et al. 2007, *PASP*, **119**, 556
- Steele, I. A., & Carter, D. 1997, *Proc. SPIE*, **3112**, 222
- Storn, R., & Price, K. 1997, *Journal of Global Optimization*, **11**, 341
- Tesauro, G. 1995, *Commun. ACM*, **38**, 58
- The Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
- Tsitsiklis, J. N., & Van Roy, B. 1997, *IEEE Trans. Autom. Control*, **42**, 674
- Wootten, A. 2003, *Proc. SPIE*, **4837**, 110
- Yoachim, P., Coughlin, M., Angeli, G. Z., et al. 2016, *Proc. SPIE*, **9910**, 99101A