On combining machine learning with decision making

Theja Tulabandhula · Cynthia Rudin

Received: 2 March 2013 / Accepted: 7 June 2014 / Published online: 28 June 2014
© The Author(s) 2014

Abstract We present a new application and covering number bound for the framework of “Machine Learning with Operational Costs (MLOC),” which is an exploratory form of decision theory. The MLOC framework incorporates knowledge about how a predictive model will be used for a subsequent task, thus combining machine learning with the decision that is made afterwards. In this work, we use the MLOC framework to study a problem that has implications for power grid reliability and maintenance, called the *Machine Learning and Traveling Repairman Problem* (ML&TRP). The goal of the ML&TRP is to determine a route for a “repair crew,” which repairs nodes on a graph. The repair crew aims to minimize the cost of failures at the nodes, but as in many real situations, the failure probabilities are not known and must be estimated. The MLOC framework allows us to understand how this uncertainty influences the repair route. We also present new covering number generalization bounds for the MLOC framework.

Keywords Decision theory · Generalization bound · Constrained linear function classes · Covering numbers · Traveling repairman · Mixed-integer programming

1 Introduction

In many domains, it is essential to understand how uncertainty in predictions influences decision-making. In that sense, one would like to explore the space of possible reasonable

Editors: Hendrik Blockeel, Kristian Kersting, Siegfried Nijssen, and Filip Železný.

T. Tulabandhula (✉)
Department of Electrical Engineering and Computer Science,
Massachusetts Institute of Technology, Cambridge, MA 02139, USA
e-mail: theja@mit.edu

C. Rudin
Computer Science and Artificial Intelligence Laboratory and Sloan School of Management,
Massachusetts Institute of Technology, Cambridge, MA 02139, USA
e-mail: rudin@mit.edu

predictions and understand the range of reasonable policies and their costs. The new framework of Machine Learning with Operational Costs (MLOC) (Tulabandhula and Rudin 2013) provides a mechanism to do this, and is a type of exploratory decision theory. Where usual decision theories provide a single policy that minimizes expected costs, the MLOC framework is able to produce a range of reasonable policies that span the full set of reasonable costs. To do this, the operational cost becomes a regularization term within the machine learning model, and adjusting the regularization constant allows us to explore solutions for all reasonable costs. This gives decision makers a way to understand the uncertainty in their predictive model in terms of something they can grasp—uncertainty in the cost to solve the problem.

The MLOC framework can also be used in another way, namely to incorporate prior knowledge about the cost to produce a better predictive model. In that sense, knowledge about the cost translates into a more restricted hypothesis space, which potentially translates into better generalization. In particular, if the hypothesis space is restricted, then upper bounds on the complexity of the hypothesis space are smaller, leading to better generalization bounds.

In this work, we provide an application of the MLOC framework to power grid engineering and reliability. This problem, called the *Machine Learning and Traveling Repairman Problem* (ML&TRP), has a machine learning component and a decision-making component. The machine learning component is to predict future power grid failures before they occur, where these failures occur at equipment that is distributed throughout the city. The decision-making component is to determine in what order the equipment should be inspected. We could use the MLOC framework in either of the two ways outlined above: either to understand the range of reasonable costs for the power company, or to use prior knowledge that the costs are high or low in order to choose a more predictive and cost-effective route.

To be more precise, the ML&TRP *prediction* problem is to determine the failure probability for each node on a graph, using features of each node and past failure data. The *decision* problem is to determine a route for a “repair crew” on the graph, where there is some travel time between each pair of nodes. There are many possible applications of the ML&TRP, including the scheduling of safety inspections or repair work for the electrical grid, oil rigs, underground mining, machines in a factory, or airlines. In our experiments, we use data from an ongoing project with Con Edison, which is NYC’s power utility company.

We also provide a generalization bound for the MLOC framework based on covering numbers. These bounds are different than those of Tulabandhula and Rudin (2013) which use concentration of Rademacher complexity and Dudley’s entropy integral, and are not directly comparable. The bounds here have a much more geometric flavor looking at the hypothesis space as a volumetric object. Neither of the two bounds are tighter in all situations. We find the bounds here to be more intuitive, as the geometry is more transparent.

The ML&TRP relates to literature on both machine learning and optimization (time-dependent traveling salesman problems). In machine learning, our work bears a slight resemblance to work on graph-based regularization (Agarwal 2006; Belkin et al. 2006; Zhou et al. 2004), but their goal is to obtain probability estimates that are smoothed on a graph with suitably designed edge weights. On the other hand, our goal is to obtain, in addition to probability estimates, a low-cost route for traversing a very different graph with edge weights that are physical distances. Our regularization is vastly different from popular ones (ℓ_1 or ℓ_2 norm) because our regularization comes from beliefs on decision-making costs. We use unlabeled data as does semi-supervised learning (Chapelle et al. 2006) but differ in the motivation as well as the way we use these additional data. For example, we do not extract distributional information from the unlabeled data. Our work contributes to the literature on the TRP (Traveling Repairman Problem) and related problems by adding the new dimension

of probabilistic estimation at the nodes. We create new adaptations of modern techniques (Fischetti et al. 1993; van Eijl 1995; Lechmann 2009) within our work for solving the TRP part of the ML&TRP.

There is a body of literature regarding cost models for maintenance in the reliability modeling literature, though the emphasis in those works is usually to design a model that accurately represents the stochastic process for the failures. In that literature, for instance, a maintenance schedule would be created from the predicted condition of the equipment (but not on the cost of performing the repairs in a certain order or routing a vehicle between the equipment). Barbera et al. (1996) develop a model that assumes that equipment have exponential rates of failure and fail only once in an inspection interval, and they use this model to determine a maintenance schedule. Marseguerra et al. (2002) introduces a model for degradation leading to failure for a continuous complex system, and use Monte Carlo simulations to determine the optimal degradation level to perform an inspection. Their work uses a very different cost model from ours; the cost is the long run average maintenance cost and cost of failures. A neural-network based maintenance model was developed by Heng et al. (2009). A related work on routing for emergency maintenance on the electrical grid is the heuristic algorithm of Weintraub et al. (1999) that dispatches vehicles to areas where there are currently breakdowns and where there are likely to be breakdowns in the future. Ertekin et al. (2013) propose a model for failures of power grid equipment and use this model to simulate the cost of various inspection policies.

One can view the MLOC framework to be analogous to a Bayesian approach, in the sense that prior knowledge is being used when not enough data are available.

In Sect. 2 we review the MLOC framework. In Sect. 3 we will motivate and outline the new application of the MLOC framework to the ML&TRP, providing two ways of modeling failure cost. In Sect. 4 we provide mixed-integer nonlinear (MINLP) formulations and discuss algorithms an illustrative example. Section 5 gives experimental results on data from the NYC power grid, showing the benefit of the ML&TRP over traditional methods. Section 6 contains the theoretical generalization result for the MLOC framework with proofs. Section 8 concludes the paper. The conference paper of Tulabandhula et al. (2011) contains a summary of work on the ML&TRP, and the paper by Tulabandhula and Rudin (2013) provides a more complete explanation of the MLOC framework, with other illustrations and connections to robust optimization.

2 Review of framework for machine learning with operational costs

In the MLOC framework we have the standard supervised training set of labeled instances, $\{(x_i, y_i)\}_{i=1}^m$, where $x_i \in \mathcal{X}$, $y_i \in \mathcal{Y}$. For simplicity, $\mathcal{X} \subset \mathbb{R}^d$. To have nonlinear functions, we could simply have the j^{th} component of x replaced by a nonlinear function $h_j(x)$. Also $\mathcal{Y} \subset \mathbb{R}$. We wish to learn a function $f^* : \mathcal{X} \rightarrow \mathcal{Y}$. This is ordinarily done by solving a minimization problem:

$$f^* \in \operatorname{argmin}_{f \in \mathcal{F}^{unc}} \left(\sum_{i=1}^m l(f(x_i), y_i) + C_2 R(f) \right), \quad (1)$$

for some loss function $l : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$, regularizer $R : \mathcal{F}^{unc} \rightarrow \mathbb{R}$, constant C_2 and function class $\mathcal{F}^{unc}$. $\mathcal{F}^{unc}$ is the set of all linear functionals, where $f \in \mathcal{F}^{unc}$ is of the form $\lambda \cdot x$, $\lambda \in \mathbb{R}^d$. The superscript ‘unc’ refers to the word “unconstrained.”

Consider an organization making a policy decision regarding a new collection of unlabeled instances $\{\tilde{x}_i\}_{i=1}^M \in \mathcal{X}^M$. The cost to enact a policy is not exactly known, because the labels for the $\{\tilde{x}_i\}_i$ are not known. Instead the model's predictions are used, which are the $f^*(\tilde{x}_i)$'s. The goal of the organization is then to create a policy π^* that minimizes operational cost $\text{OpCost}(\pi, f^*, \{\tilde{x}_i\}_i)$. The operational cost $\text{OpCost}(\pi, f^*, \{\tilde{x}_i\}_i)$ is how much will be spent if policy π is chosen in response to the $\{f^*(\tilde{x}_i)\}_i$'s. When there is uncertainty in f^* , there is uncertainty in the cost to enact the optimal policy π^* . This uncertainty is what we would like to explore. A typical way that companies make decisions is using what we call the sequential process, which computes the policy according to two steps:

Step 1: Create function f^* based on $\{(x_i, y_i)\}_i$ according to (1). That is:

$$f^* \in \operatorname{argmin}_{f \in \mathcal{F}^{unc}} \left(\sum_{i=1}^m l(f(x_i), y_i) + C_2 R(f) \right).$$

Step 2: Choose policy π^* to minimize the operational cost,

$$\pi^* \in \operatorname{argmin}_{\pi \in \Pi} \text{OpCost}(\pi, f^*, \{\tilde{x}_i\}_i).$$

On the other hand, the MLOC framework is based around a simultaneous process, which combines Steps 1 and 2 of the sequential process. To do this, the operational cost becomes a regularization term, and its regularization parameter C_1 controls the amount of optimism or pessimism for the operational cost.

Step 1: Choose a model f^* obeying the following:

$$f^* \in \operatorname{argmin}_{f \in \mathcal{F}^{unc}} \left[\sum_{i=1}^m l(f(x_i), y_i) + C_2 R(f) + C_1 \min_{\pi \in \Pi} \text{OpCost}(\pi, f, \{\tilde{x}_i\}_i) \right].$$

Step 2: Compute the policy:

$$\pi^* \in \operatorname{argmin}_{\pi \in \Pi} \text{OpCost}(\pi, f^*, \{\tilde{x}_i\}_i).$$

The case $C_1 = 0$ for the simultaneous process is precisely the sequential process; thus, the sequential process is a special case of the simultaneous process. Our ability to solve the MLOC simultaneous process depends on the tractability of the optimization problem $\operatorname{argmin}_{\pi \in \Pi} \text{OpCost}(\pi, f^*, \{\tilde{x}_i\}_i)$. However, if this problem is intractable, then the sequential process is also intractable, and the organization will not be able to choose an optimized policy at all. The simultaneous process requires this subproblem to be solved several times, whereas the sequential process only requires the subproblem to be solved once. If the number of unlabeled instances is small, then Step 1 can be solved without a problem, even if the training set is large. As C_1 varies over its full range, it maps out the full range of costs for all reasonable solutions. If C_1 is set to a number that is too large (either positive or negative), the solution of the simultaneous process will have empirical error that is too high to be reasonable. In that case, we know that by varying C_1 within a smaller range will lead to the full range of costs for reasonable predictive models.

As with any regularization term, the new operational cost term can be interpreted as a prior belief about the model - in this case, a belief that the operating costs should be lower or higher on the current set of unlabeled instances $\{\tilde{x}_i\}_i$. In that sense, MLOC regularization may have a closer connection to reality than typical (e.g., ℓ_1 or ℓ_2 norm) regularizers. If one asks a manager at a company what prior belief they have about the estimation model, it is not likely

they would give an answer in terms of coefficients for a linear model. Even managers who are not mathematicians or computer scientists might have some belief - they could perhaps believe that they are expecting to spend a certain amount to enact the policy. It is possible that this type of belief, which relies on direct experience, might be more practical, and more accurate, than the more abstract prior information that we are typically used to dealing with. In the ML&TRP, the training error term is derived from data from the past, and the OpCost term is calculated on data from the present. The OpCost term is the only term that deals with routing.

3 The machine learning and traveling repairman problem

The US Department of Energy's Grid 2030 document states that "America's electric system, 'the supreme engineering achievement of the twentieth century,' is aging, inefficient, and congested, and incapable of meeting the future energy needs of the Information Economy without operational changes and substantial capital investment over the next several decades" (United States Department of Energy and Distribution 2003). Since 2004, many power utility companies are implementing new inspection and repair programs for preemptive maintenance, whereas in the past, all repair work was done reactively (Urbina 2004). New York City has the oldest power system in the world, and the largest underground electric system, with enough electrical cable to go three and a half times around the world. In New York City, there are several separate new preemptive maintenance programs, including the targeted inspection program for electrical service structures (manholes), programs that perform extensive repairs that were placed on a waiting list after the manhole was inspected, and the *vented cover replacement program*, where each manhole is replaced with a vented cover that allows gases to escape, mitigating the possibility and effects of serious events including fires and explosions. Con Edison, the power company in NYC, has the ability to use machine learning models in Manhattan, Brooklyn and the Bronx for scheduling of manhole inspection and repair work (Rudin et al. 2010, 2012, 2011, 2014). This project was the motivation for the development of the ML&TRP and we use data from the NYC power grid for our experiments. Features for the NYC model are derived from physical characteristics of the manhole (e.g., number of electrical cables entering the manhole), and features derived from its history of involvement in past events. Repeat failures (serious and non-serious events) can occur on the same manhole. We take the possibility of repeat failures into account in the ML&TRP (in Cost 1 given below). That said, failures are rare events, and it is not easy to accurately estimate the probability that a given manhole will fail within a given period of time. Because of this uncertainty, we can use the MLOC framework to assist in decision-making. The result $\pi^* \in \Pi$ from the algorithm would be a route that could be used for the repair crew to fix a pre-specified set of manholes corresponding to $\{\tilde{x}_i\}_{i=1}^M$, which are assumed to need a particular repair.

3.1 Learning

In what follows, we will use descriptions and terminology that match the power grid application. In the ML&TRP, data from the past will be used to train the model, denoted $\{(x_i, y_i)\}_{i=1}^m$, whereas the $\tilde{x}_i$ are calculated from the present, whose labels are from the future and thus not known. Let x_i^j indicate the j -th coordinate of the feature vector for manhole i calculated at a time period from the past. The x_i vector encodes the number and types of electrical cables, number and types of previous events, etc. The label for manhole i from the past is denoted y_i , where $y_i \in \{-1, 1\}$ indicating whether the manhole had a failure (fire, explo-

sion, smoking manhole) within a specific period of time in the past. More details about the features and labels can be found in Sect. 5. The other instances $\{\tilde{x}_i\}_{i=1}^M$ (with M unrelated to m), are unlabeled data that are each associated with a node on a graph G . The nodes of the graph G indexed by $i = 1, \dots, M$ represent manholes on which we want to design a route. Note that M can be substantially smaller than m , e.g., $M < 10$ and $m > 20,000$; e.g., for a repair truck that carries supplies for at most M repairs. We are also given physical distances $d_{i,j} \in \mathbb{R}_+$ between each pair of nodes i and j . A route on G is represented by a permutation π of the node indices $1, \dots, M$. Let Π be the set of all permutations of $\{1, \dots, M\}$. Failure probabilities will be estimated at each of the nodes and these estimates will be based on a function of the form $f_\lambda(x) = \lambda \cdot x$. The class of possible functions $\mathcal{F}$ is chosen to be: $\mathcal{F} := \{f_\lambda : \lambda \in \mathbb{R}^d, \|\lambda\|_2 \leq B_b\}$, where B_b is a fixed positive real number. We choose the logistic loss: $l(f_\lambda(x), y) := \ln(1 + e^{-y f_\lambda(x)})$ so that the probability of failure $P(y = 1|x)$, is estimated as in logistic regression by:

$$P(y = 1|x) \text{ or } p(x) := \frac{1}{1 + e^{-f_\lambda(x)}}. \quad (2)$$

Note that the routing problem is done in batch: once the route is determined, the repair truck is sent out and changes to the route are no longer possible.

3.2 Two options for the OpCost

The operational cost can be defined to match the application. In the first option (denoted as Cost 1), for each node there is a cost for (possibly repeated) failures prior to a visit by the repair crew. In this case, temporary repairs are made to fix each node before the repair crew comes to make permanent repairs. In the second option (denoted as Cost 2), for each node, there is a cost for the first failure prior to visiting it. In this case, permanent repairs are made when there is an event, or when the repair crew arrives, whichever is sooner. There is a natural interpretation of the failures as being generated by a continuous random process at each of the nodes. When discretized in time, this is approximated by a Bernoulli process with parameter $p(\tilde{x}_i)$. Both Cost 1 and Cost 2 are appropriate for power grid applications. Cost 2 is also appropriate for delivery truck routing applications, where perishable items can fail (once an item has spoiled, it cannot spoil again).

For convenience, we assume that after the repair crew visits all the nodes, it returns to the starting node (node 1) which is fixed beforehand. Scenarios where one is not interested in beginning from or returning to the starting node would be modeled slightly differently (the computational complexity remains the same). Let a route be represented by $\pi : \{1, \dots, M\} \mapsto \{1, \dots, M\}$, this means that $\pi(i)$ is the i th node to be visited. For example, let $M = 4$, $\pi = [2, 3, 4, 1]$. This means, $\pi(1) = 2$, node 2 is the first node to be visited, $\pi(2) = 3$, node 3 is the second node on the route, and so on. Since the final node visited is the first node, we append the following to the definition of π : $\pi(M + 1) = \pi(1)$. Let the distances be scaled appropriately so that a unit of distance is traversed in a unit of time. Given a route, the *latency* of a node $\pi(i)$ is the time (or equivalently distance) from the start at which node $\pi(i)$ is visited. It is the sum of distances traversed before position i on the route:

$$L_\pi(\pi(i)) := \begin{cases} \sum_{k=1}^M d_{\pi(k)\pi(k+1)} \mathbf{1}_{[k < i]} & i = 2, \dots, M \\ \sum_{k=1}^M d_{\pi(k)\pi(k+1)} & i = 1. \end{cases} \quad (3)$$

The starting node $\pi(1)$ thus has a latency $L_\pi(\pi(1))$ which is the total length of the route starting at node $\pi(1)$ and ending at node $\pi(1)$ after visiting all other nodes.

Cost 1: cost is proportional to expected number of failures before the visit

Up to the time that node $\pi(i)$ is visited by the repair crew, there is a probability $p(\tilde{x}_{\pi(i)})$ that a failure will occur within each unit time interval. Equivalently, within each unit time interval, failures are determined by a Bernoulli random variable with parameter $p(\tilde{x}_{\pi(i)})$. Thus, in a time interval of length $L_{\pi}(\pi(i))$ units, the number of node failures follows the binomial distribution $\text{Bin}(L_{\pi}(\pi(i)), p(\tilde{x}_{\pi(i)}))$. For each node, we will associate a cost proportional to the expected number of failures before the repair crew's visit, as follows:

$$\begin{aligned} \text{Cost of node } \pi(i) &\propto E(\text{number failures in } L_{\pi}(\pi(i)) \text{ time units}) \\ &= \text{mean of Bin}(L_{\pi}(\pi(i)), p(\tilde{x}_{\pi(i)})) = p(\tilde{x}_{\pi(i)})L_{\pi}(\pi(i)). \end{aligned} \quad (4)$$

Using this cost, if the failure probability for node $\pi(i)$ is small, we can afford to visit it later on, trading off its latency $L_{\pi}(\pi(i))$. If $p(\tilde{x}_{\pi(i)})$ is large, we should visit node $\pi(i)$ earlier to keep our overall failure cost low. The failure cost of route π is then $\text{OpCost}(\pi, f_{\lambda}, \{\tilde{x}_i\}_{i=1}^M, \{d_{i,j}\}_{i,j=1}^M) = \sum_{i=1}^M p(\tilde{x}_{\pi(i)})L_{\pi}(\pi(i))$.

Substituting the definition of $L_{\pi}(\pi(i))$ from (3):

$$\begin{aligned} \text{OpCost}(\pi, f_{\lambda}, \{\tilde{x}_i\}_{i=1}^M, \{d_{i,j}\}_{i,j=1}^M) \\ = \sum_{i=2}^M p(\tilde{x}_{\pi(i)}) \sum_{k=1}^M d_{\pi(k)\pi(k+1)} \mathbf{1}_{[k < i]} + p(\tilde{x}_{\pi(1)}) \sum_{k=1}^M d_{\pi(k)\pi(k+1)}, \end{aligned} \quad (5)$$

where $p(\tilde{x}_{\pi(i)})$ is given in (2). This will be Cost 1. There are ways to make Cost 1 more general. The individual node cost in (4) assumes that the node's failure probability $p(\tilde{x}_{\pi(i)})$ becomes zero after the repair crew's visit, so that for the remainder of the route, the cost incurred at this node is $\propto 0 \times (L_{\pi}(\pi(1)) - L_{\pi}(\pi(i)))$. We could relax this by assuming $p(\tilde{x}_{\pi(i)})$ does not vanish after the repair crew's visit and adding an additional cost for the expected failures in this period. That is, if β is a constant of proportionality for the cost after visiting node $\pi(i)$, then the cost would become:

$$\text{Cost of node } \pi(i) = \beta [L_{\pi}(\pi(1)) - L_{\pi}(\pi(i))] p(\tilde{x}_{\pi(i)}) + L_{\pi}(\pi(i)) p(\tilde{x}_{\pi(i)}).$$

If $\beta = 1$, then the repair crew does not have any effect and cost of each node is independent of its expected number of failures before the repair crew's visit. Typically, we expect that the repair crew will repair the node so that it will not fail, and the second term above is much larger than the first. Taking the constant of proportionality as $\beta = 0$, we return to the individual costs given by (4).

Note that since the cost is a sum of M terms, it is invariant to ordering or indexing (caused by π). Thus we can rewrite the cost as

$$\text{OpCost}(\pi, f_{\lambda}, \{\tilde{x}_i\}_{i=1}^M, \{d_{i,j}\}_{i,j=1}^M) = \sum_{i=1}^M p(\tilde{x}_i) L_{\pi}(i). \quad (6)$$

Cost 2: cost is proportional to probability that the first failure is before the visit

This cost reflects the penalty for not visiting a node before the first failure occurs there. This model is governed by the geometric distribution. Let the parameter of the distribution be p . Then the probability that the first failure for node $\pi(i)$ occurs at time index $t > 0$ is $p(1-p)^{t-1}$. The probability that the first failure for node $\pi(i)$ occurs before time $L_{\pi}(\pi(i))$ is

then the sum of the failure probabilities from $t = 1, \dots, L_\pi(\pi(i)) : \sum_{t=1}^{L_\pi(\pi(i))} p(1-p)^{t-1} = 1 - (1-p)^{L_\pi(\pi(i))}$. Thus, substituting the expression (2) for p , we have:

$$\begin{aligned} P(\text{first failure occurs before time } L_\pi(\pi(i))) &= 1 - (1 - p(\tilde{x}_{\pi(i)}))^{L_\pi(\pi(i))} \\ &= 1 - \left(1 - \frac{1}{1 + e^{-f_\lambda(\tilde{x}_{\pi(i)})}}\right)^{L_\pi(\pi(i))} = 1 - \left(1 + e^{f_\lambda(\tilde{x}_{\pi(i)})}\right)^{-L_\pi(\pi(i))}. \end{aligned}$$

The cost of visiting node $\pi(i)$ will be proportional to this quantity:

$$\text{Cost of node } \pi(i) \propto \left(1 - \left(1 + e^{f_\lambda(\tilde{x}_{\pi(i)})}\right)^{-L_\pi(\pi(i))}\right). \quad (7)$$

Similarly to Cost 1, $L_\pi(\pi(i))$ influences the cost at each node. If we visit a node early in the route, then the cost incurred is small because the node is less likely to fail before we reach it. Similarly, if we schedule a visit later on in the tour, the cost is higher because the node has a higher chance of failing prior to the repair crew's visit.

The total failure cost is thus:

$$\text{OpCost}(\pi, f_\lambda, \{\tilde{x}_i\}_{i=1}^M, \{d_{i,j}\}_{i,j=1}^M) = \sum_{i=1}^M \left(1 - \left(1 + e^{f_\lambda(\tilde{x}_{\pi(i)})}\right)^{-L_\pi(\pi(i))}\right). \quad (8)$$

This cost is not directly related to a weighted TRP cost in its present form. That is, when the failure probabilities of the nodes are all the same, the total cost is not linear in the latencies, as is the case for Cost 1. Building on this cost, we will derive a cost that is the same as a weighted TRP in Sect. 4.2, of the form:

$$\text{Cost of node } \pi(i) \propto L_\pi(\pi(i)) \log \left(1 + e^{f_\lambda(\tilde{x}_{\pi(i)})}\right), \quad (9)$$

as an alternative to (7).

There is a slightly more general version of this formulation (as there was for Cost 1), which is to take the cost for each node to be a function of two quantities: the probability of failure before the visit, and the probability of failure after the visit. Let us redefine β to be a constant of proportionality for the cost of visiting before the failure event. From the geometric distribution, $P(\text{failure occurs after time } L_\pi(\pi(i))) = (1 - p(\tilde{x}_{\pi(i)}))^{L_\pi(\pi(i))}$, and the cost of visiting node $\pi(i)$ becomes:

$$\text{Cost of node } \pi(i) \propto P(\text{failure before } L_\pi(\pi(i))) + \beta \times P(\text{failure after } L_\pi(\pi(i))).$$

If $\beta = 1$, then the sum above is 1 for all nodes regardless of node failures or latencies. More realistically, the cost of visiting the node after the failure is more than the cost of visiting proactively, $\beta \ll 1$ leading to (7). We could again have written the summation to hide the dependence on π :

$$\text{OpCost}(\pi, f_\lambda, \{\tilde{x}_i\}_{i=1}^M, \{d_{i,j}\}_{i,j=1}^M) = \sum_{i=1}^M \left(1 - \left(1 + e^{f_\lambda(\tilde{x}_i)}\right)^{-L_\pi(i)}\right).$$

Remark 1 The costs defined above are by no means exhaustive. We chose to define operational costs this way because they mimic the well known minimum latency objective in routing problems. For instance, we could have used a Poisson failure model at each node instead of binomial or geometric as in Costs 1 and 2. Let us assume that the Poisson rate

parameter $\mu(\tilde{x}_{\pi(i)})$ is the output of the estimation problem (say proportional to $p(\tilde{x}_{\pi(i)})$). Then

$$P(k \text{ failures occur in time } L_{\pi}(\pi(i))) = \frac{(\mu(\tilde{x}_{\pi(i)})L_{\pi}(\pi(i)))^k e^{-\mu(\tilde{x}_{\pi(i)})L_{\pi}(\pi(i))}}{k!}.$$

From this we can get the probability that at least one failure occurs in time interval $[0, L_{\pi}(\pi(i))]$ at node $\pi(i)$. Now we can define the operational cost to be the sum of these probabilities which depend on the routing and proceed in the same way as Cost 2. That is, we can minimize this cost to get the optimal routing π^* .

Remark 2 The operational cost must depend on graph properties like latency. We would not like to minimize an objective of the form $\sum_{i=1}^M \frac{1}{p(\tilde{x}_{\pi(i)})}$ (or any other function of just $p(\tilde{x}_{\pi(i)})$, the output of the estimation problem) as this does not lead to an operational cost in the true sense. This operational cost does not make use of latency information or other graph properties related to routing unless $p(\tilde{x}_{\pi(i)})$ implicitly depends on them (which is not the case here).

Now that the major steps for both formulations have been defined, we will discuss methods for optimizing the objectives.

4 Optimization

We start by formulating mixed-integer linear programs (MILP's) for the TRP subproblem.

4.1 Mixed-integer optimization for Cost 1

For either the sequential or simultaneous processes, we need the solution of the subproblem: $\pi^* \in \operatorname{argmin}_{\pi \in \Pi} \operatorname{OpCost}(\pi, f_{\lambda}^*, \{\tilde{x}_i\}_{i=1}^M, \{d_{i,j}\}_{i,j=1}^M)$, or equivalently,

$$\pi^* \in \operatorname{argmin}_{\pi \in \Pi} \sum_{i=2}^M p(\tilde{x}_{\pi(i)}) \sum_{k=1}^M d_{\pi(k)\pi(k+1)} \mathbf{1}_{[k < i]} + p(\tilde{x}_{\pi(1)}) \sum_{k=1}^M d_{\pi(k)\pi(k+1)}. \quad (10)$$

Let us compare this to the standard traveling repairman problem (TRP) problem (see [Blum et al. 1994](#)):

$$\pi^* \in \operatorname{argmin}_{\pi \in \Pi} \sum_{k=1}^M d_{\pi(k)\pi(k+1)} (M + 1 - k). \quad (11)$$

The standard TRP objective (11) is a special case of the weighted TRP (10) when $\forall i = 1, \dots, M$, $p(\tilde{x}_i) = p$:

$$\begin{aligned} & \sum_{i=2}^M p(\tilde{x}_{\pi(i)}) \sum_{k=1}^M d_{\pi(k)\pi(k+1)} \mathbf{1}_{[k < i]} + p(\tilde{x}_{\pi(1)}) \sum_{k=1}^M d_{\pi(k)\pi(k+1)} \\ &= p \sum_{i=2}^M \sum_{k=1}^M d_{\pi(k)\pi(k+1)} \mathbf{1}_{[k < i]} + p \sum_{k=1}^M d_{\pi(k)\pi(k+1)} \\ &= p \sum_{i=2}^M \sum_{k=1}^M d_{\pi(k)\pi(k+1)} \mathbf{1}_{[k < i]} + p \sum_{k=1}^M d_{\pi(k)\pi(k+1)} \mathbf{1}_{[k < M+1]} \end{aligned}$$

$$= p \sum_{k=1}^M d_{\pi(k)\pi(k+1)} \sum_{i=2}^{M+1} \mathbf{1}_{[k < i]} = p \sum_{k=1}^M d_{\pi(k)\pi(k+1)} (M+1-k).$$

The TRP is different from the traveling salesman problem (TSP); the goal of the traveling salesman problem is to minimize the total traversal time (in this case, this is the same as the distance traveled) needed to visit all nodes once, whereas the goal of the traveling repairman problem is to minimize the sum of the waiting times to visit each node. Both the TSP and the TRP are known to be NP-complete in the general case (Blum et al. 1994). Intuitively, a TRP route cost objective captures the total waiting cost of a service system from the customer's (the node's) point of view. For example, consider a truck carrying prioritized items to be delivered to customers. At each customer's stop, that customer's item is removed from the truck. The goal of the TRP is to minimize the total waiting time of these customers.

We start by extending an integer programming formulation of standard TRP (Fischetti et al. 1993) to include “unequal flow values” so that we can solve (10) (there are many other integer programming formulations in the literature as well, see for instance Méndez-Díaz et al. 2008). The weights $\{\bar{p}(\tilde{x}_i)\}_i$ within the formulation below will be defined later. For interpretation, consider the sum of the probabilities $\sum_{i=1}^M \bar{p}(\tilde{x}_i)$ as the total “flow” through a route. At the beginning of the tour, the repair crew has flow $\sum_{i=1}^M \bar{p}(\tilde{x}_i)$. Along the tour, flow of the amount $\bar{p}(\tilde{x}_i)$ is dropped when the repair crew visits node $\pi(i)$ at latency $L_{\pi}(\pi(i))$. In this way, the amount of flow during the tour is the sum of the probabilities $\bar{p}(\tilde{x}_i)$ for nodes that the repair crew has not yet visited. We introduce two sets of variables $\{z_{i,j}\}_{i,j}$ and $\{y_{i,j}\}_{i,j}$ that together represent a route (instead of the π notation). Let $z_{i,j}$ represent the flow on edge (i, j) and let a binary variable $y_{i,j}$ represent whether there exists a flow on edge (i, j) . (There will only be a flow along the route, and there will not be a flow along edges that are not in the route.) The mixed-integer program is as follows:

$$\min_{z,y} \sum_{i=1}^M \sum_{j=1}^M d_{i,j} z_{i,j} \text{ s.t.} \quad (12)$$

$$\text{No flow from node } i \text{ to itself: } z_{i,i} = 0 \quad \forall i = 1, \dots, M \quad (13)$$

$$\text{No edge from node } i \text{ to itself: } y_{i,i} = 0 \quad \forall i = 1, \dots, M \quad (14)$$

$$\text{Exactly one edge into each node: } \sum_{i=1}^M y_{i,j} = 1 \quad \forall j = 1, \dots, M \quad (15)$$

$$\text{Exactly one edge out from each node: } \sum_{j=1}^M y_{i,j} = 1 \quad \forall i = 1, \dots, M \quad (16)$$

$$\text{Flow coming back to initial point at the end of loop: } \sum_{i=1}^M z_{i,1} = \bar{p}(\tilde{x}_1) \quad (17)$$

Change of flow after crossing node k :

$$\sum_{i=1}^M z_{i,k} - \sum_{j=1}^M z_{k,j} = \begin{cases} \bar{p}(\tilde{x}_1) - \sum_{i=1}^M \bar{p}(\tilde{x}_i) & k = 1 \\ \bar{p}(\tilde{x}_k) & k = 2, \dots, M \end{cases} \quad (18)$$

$$\text{Connects flows } z \text{ to indicators of edge } y : z_{i,j} \leq r_{i,j} y_{i,j} \quad (19)$$

$$\text{where } r_{i,j} = \begin{cases} \bar{p}(\tilde{x}_1) & j = 1 \\ \sum_{i=1}^M \bar{p}(\tilde{x}_i) & i = 1 \\ \sum_{i=2}^M \bar{p}(\tilde{x}_i) & \text{otherwise.} \end{cases}$$

Constraints (13) and (14) restrict self-loops from forming. Constraints (15) and (16) ensure that every node should have exactly one edge coming in and one going out. Constraint (17) represents the flow on the last edge coming back to the starting node. Constraint (18) quantifies the flow change after traversing a node k . Constraint (19) represents an upper bound on $z_{i,j}$ relating it to the corresponding binary variable $y_{i,j}$. We can define the weights $\bar{p}(\tilde{x}_i)$, for example, for Cost 1, to be equal to the estimated failure probabilities $1/(1 + e^{-\lambda \cdot \tilde{x}_i})$.

4.2 Mixed integer optimization for Cost 2

Here we reason about the choice for changing the cost per node in (7) to resemble (9). Starting with the sum (8) over node costs (7), we apply the log function to the second term of the cost of each node (7) to get a new cost $(1 - \log(1 + e^{f_\lambda(\tilde{x}_{\pi(i)})})^{-L_\pi(\pi(i))})$, and the new minimization problem is:

$$\begin{aligned} \min_{\pi} \sum_{i=1}^M & \left(1 - \log \left(1 + e^{f_\lambda(\tilde{x}_{\pi(i)})} \right)^{-L_\pi(\pi(i))} \right) \\ &= - \max_{\pi} \left(\sum_{i=1}^M \log \left(1 + e^{f_\lambda(\tilde{x}_{\pi(i)})} \right)^{-L_\pi(\pi(i))} - \text{const} \right) \\ &= \min_{\pi} \left[\sum_{i=1}^M L_\pi(\pi(i)) \log \left(1 + e^{f_\lambda(\tilde{x}_{\pi(i)})} \right) \right] + \text{const}, \end{aligned}$$

where the first term is the sum over nodes of the expression (9). This failure cost term is now a weighted sum of latencies where the weights are of the form $\log(1 + e^{f_\lambda(\tilde{x}_{\pi(i)})})$. We can thus reuse the mixed integer program (12)–(19) where the weights are redefined as $\bar{p}(\tilde{x}_i) := \log(1 + e^{\lambda \cdot \tilde{x}_i})$.

Our choices for the cost and failure models above allow us to use a weighted version of the intuitive minimum latency or TRP problem for routing. In particular, the log transformation of individual terms in the original version of Cost 2, (8), precisely serves this purpose. In general, depending on the way we define the operational cost and the failure model, they may not necessarily map back to popular routing problems like the TRP as we have here. Nonetheless, there are many valid approaches beyond what we pursue this in this paper.

Now that the TRP subproblem has been completely defined for both Cost 1 and Cost 2, we will discuss first how to solve the subproblem alone, which is Step 2 of the sequential process. Then we will discuss the solvers for the simultaneous process.

4.3 Solving the weighted TRP subproblem

A generic MILP solver like CPLEX¹ or Gurobi² can produce an exact solution using branch-and-bound or other related exact methods. We use Gurobi. The weighted TRP problem is NP-hard (can be shown by a reduction to the Hamiltonian cycle problem) and hence most likely not solvable by polynomial-time algorithms. The standard unweighted (all weights equal)

¹ IBM ILOG CPLEX Optimization Studio v12.2.0.2 2010.

² Gurobi Optimizer v3.0, Gurobi Optimization, Inc. 2010.

TRP can be encoded by different mixed-integer programming formulations (see Fischetti et al. 1993; van Eijl 1995; Méndez-Díaz et al. 2008) each with different performance guarantees (e.g., solving 15–60 nodes), which could be adapted for our purpose. There are also techniques for producing constant factor approximate solutions to the unweighted TRP (Goemans and Kleinberg 1998; Blum et al. 1994; Arora and Karakostas 2006; Archer et al. 2008; Archer and Blasiak 2010), which could run faster than the MILP solvers for large problems. If the weights $\{w_i\}_i$ are integers, we can adapt these faster techniques for the standard problem to the weighted TRP problem by replicating each node w_i times. If the weights are rational, as is the case in (20) and (21), we can use rounding and discretization in order to apply the faster solution techniques for solving the standard TRP.

4.4 Solving mixed-integer nonlinear programs (MINLPs)

For the simultaneous process, the inputs to the program are training data $\{x_i, y_i\}_{i=1}^m$, unlabeled nodes $\{\tilde{x}_i\}_{i=1}^M$ the distances between them $\{d_{i,j}\}_{i,j=1}^M$ and constants C_1 and C_2 . The full simultaneous process formulation using Cost 1 is:

$$\min_{\lambda} \left(\sum_{i=1}^m \ln \left(1 + e^{-y_i f_{\lambda}(x_i)} \right) + C_2 \|\lambda\|_2^2 + C_1 \min_{\{z_{i,j}, y_{i,j}\}} \sum_{i=1}^M \sum_{j=1}^M d_{i,j} z_{i,j} \right) \\ \text{subject to constraints (13) to (19), where } \bar{p}(\tilde{x}_i) = \frac{1}{1 + e^{-\lambda \cdot \tilde{x}_i}}. \quad (20)$$

The full formulation using the modified version of Cost 2 is:

$$\min_{\lambda} \left(\sum_{i=1}^m \ln \left(1 + e^{-y_i f_{\lambda}(x_i)} \right) + C_2 \|\lambda\|_2^2 + C_1 \min_{\{z_{i,j}, y_{i,j}\}} \sum_{i=1}^M \sum_{j=1}^M d_{i,j} z_{i,j} \right) \\ \text{subject to constraints (13) to (19) hold, where } \bar{p}(\tilde{x}_i) = \log \left(1 + e^{\lambda \cdot \tilde{x}_i} \right). \quad (21)$$

If we have an algorithm for solving (20), then the same scheme can be used to solve (21). There are multiple ways of solving (or approximately solving) a mixed integer nonlinear optimization problem of the form (20) or (21). We consider three methods in this paper for solving (20) and (21).

- Generic mixed integer non-linear programming (MINLP) solver (Bonmin).
- Nelder–Mead (NM) which is a iterative scheme over the λ parameter space, solving a weighted TRP subproblem in each iteration.
- Alternating Minimization (AM) which alternatively minimizes over λ and π optimization variables.

Method 1: MINLP solver

For our experiments we directly use a MINLP solver called Bonmin (Bonami et al. 2008). These types of solvers typically use general MILP solving techniques like branch and bound or dynamic programming interleaved with continuous optimization. Since the general MILP solving techniques, as discussed, can take exponential time when applied directly to our formulations, the MINLP solvers which use them can in turn, be inefficient if the graph is moderate to large in size. However, when the graph is small, for instance when we want to schedule a tour over only a few nodes, the MINLP solver can generally compute a solution to the problems (20) or (21) in a manageable period of time.

Algorithm 1 AM: Alternating minimization algorithm

Inputs: $\{x_i, y_i\}_1^m, \{\tilde{x}_i\}_1^M, \{d_{ij}\}_{ij}, C_1, C_2, T$ and initial vector λ_0 .
for $t=1:T$ **do**
 Compute $\pi_t \in \operatorname{argmin}_{\pi \in \Pi} \operatorname{Obj}(\lambda_{t-1}, \pi)$.
 Compute $\lambda_t \in \operatorname{argmin}_{\lambda \in \mathbb{R}^d} \operatorname{Obj}(\lambda, \pi_t)$.
end for
Output: π_T .

Method 2: Nelder–Mead in λ -space (NM)

The Nelder–Mead minimization algorithm requires only function evaluations (Nelder and Mead 1965). The ML&TRP can be viewed as a minimization in the space of all λ vectors; since we have solvers for the weighted TRP subproblem, we are able to evaluate the ML&TRP objective for a given value of λ . In our experiments we use the MILP solver (Gurobi) for the subproblem. Note that the ML&TRP objective can have non-differentiable kinks arising from discontinuities in the failure cost term; a method that relies on the gradient or Hessian of the objective function might get stuck in narrow local minima, whereas methods that use only function evaluations may not have this problem. The generic Nelder–Mead scheme can have disadvantages with respect to performance (Rios 2009), in which case, other schemes like Multilevel Coordinated Search (MCS) (Huyer and Neumaier 1999) can be used in place of Nelder–Mead. Note that since the objective is non-convex, all solutions obtained by NM are only guaranteed to be locally optimal.

Method 3: alternating minimization in λ – π space (AM)

Our alternating minimization scheme also operates in the λ – π space as follows. Define the objective Obj as a function of λ and π :

$$\operatorname{Obj}(\lambda, \pi) = \sum_{i=1}^m \ln \left(1 + e^{-y_i f_{\lambda}(x_i)} \right) + C_2 \|\lambda\|_2^2 + C_1 \operatorname{OpCost} \left(\pi, f_{\lambda}, \{\tilde{x}_i\}_{i=1}^M, \{d_{i,j}\}_{i,j=1}^M \right).$$

Starting from an initial vector λ_0 , Obj is minimized alternately with respect to λ and then with respect to π , as shown in Algorithm 1. The second step, solving for π , is the same as solving the TRP subproblem, and we again use the MILP solver for this. Conditions for convergence and correctness for such iterative schemes are given by Csiszár and Tusnády (1984); again, it is not possible to guarantee globally optimal solutions using this method.

4.5 Illustrative experiment

We will use the ML&TRP to show the fundamental property motivating the MLOC framework: that a large change in the probability model does not necessarily lead to a large change in overall prediction accuracy, but may lead to very different solutions.

The training set was chosen uniformly at random from a distribution that is uniform over two triangles pointing end to end. We used six unlabeled points as the nodes. See Fig. 1a. In addition a level set, colored black, is also plotted. It is the estimated level set for $P(y = 1|x) = 0.5$ learned from ℓ_2 -regularized logistic regression. A second level set, colored red, also drawn at probability estimate 0.5, is learned from the simultaneous process, with failure cost modeled according to Cost 1. Now, node 6 (triangle with label “ $\tilde{x}_6$ ”) lies in a low density region of feature space, so its probability cannot be well estimated. For the

Fig. 1 Left x^1 and x^2 represent the first and second coordinates respectively of the 2D feature space. The triangles represent the unlabeled data $\{\tilde{x}_i\}_{i=1}^6$. Right the numbers in the nodes indicate their probability of failure, and the numbers on the edges indicate distances. The optimal route 1-2-3-6-4-5-1 as determined by the sequential formulation is highlighted

Fig. 2 The optimal route 1-2-3-4-5-6-1 determined by the simultaneous process is highlighted

sequential formulation, node 6 was assigned $p(\tilde{x}_6) = 0.5$ and the optimal route obtained by solving the weighted TRP problem is 1-2-3-6-4-5-1, shown in Fig. 1b. The node represented by $\tilde{x}_1$ is chosen to be the starting point. For the simultaneous process, node 6 has been assigned a new probability value $p(\tilde{x}_6) = 0.29$. This change is possible because node 6's probability estimate can vary quite a lot without changing the probability estimates of others. This changes the route to 1-2-3-4-5-6-1 as shown in Fig. 2.

In the simultaneous process, we chose C_1 large enough so that the tour route visits 4 and 5 before 6. This results in a $\sim 9\%$ decrease in the failure cost (Cost 1), with a $\sim 3\%$ change in the training error (logistic loss). In particular, for the sequential process, Cost 1 is 4.7 units and the training error is 15.7 units; for the simultaneous process, Cost 1 is 4.25 units and the learning error is 16.2 units ($C_1 = 5 \times 10^{-4}$). This is an illustration of the core of MLOC: both predictive models are good, and a range of operational costs and decisions exist between them.

5 ML&TRP on the NYC power grid

We now show how the MLOC framework might be used to assist companies like Con Edison, which is NYC's power utility company. We pursue three sets of experiments. The first experiment demonstrates the use of the simultaneous process when given a specific routing

problem. This shows how a practitioner would use the simultaneous process in practice. In the second experiment, we randomize over the training sample and routing problems. This experiment shows that the simultaneous process can find models that are equally predictive or better than the sequential method when operational costs are included. In the third experiment, we look at scaling issues.

In all these experiments, we are predicting the probability of failure over the course of a year. While using the predicted failure probabilities in the routing problem, we will assume that these are probabilities of failures in an arbitrary unit interval of time. In particular, they can be the probability of failures over an hour, a day etc. We make the approximation that the probabilities at finer time scales (required for the routing problem) are proportional to the probabilities at coarser time scales for the purpose of our experiments.

5.1 The dataset

The dataset we use is described by [Rudin et al. \(2010\)](#), which was developed in order to assist Con Edison with its maintenance and repair programs on the secondary electrical distribution network in NYC; specifically, it was designed for the purpose of predicting manhole fires and explosions. We chose to use all manholes from the Bronx (~23K manholes). Each manhole is represented by (4-dimensional) features that encode the number and type of electrical cables entering the manhole and the number and type of past events involving the manhole. The event features encode how often in the past the manhole was the source of partial outages, full outages and/or underground burnouts. The training features encode events prior to 2008, and the training labels are 1 if the manhole was the source of a serious event (fire, explosion, smoke) during 2008. The prediction task is to predict events in 2009. The test set (for evaluating the performance of the predictive model) consists of features derived from the time period before 2009, and labels from 2009. In our experiments, for both training and test we had a large sample (23,217 instances). There were 211 and 132 failure instances in the test and training data respectively.

5.2 Performance of the simultaneous process for a seven node decision problem

In this experiment, the operational task is to design a route for a repair crew that is equipped to fix seven relatively more vulnerable manholes in 2009. The distances between the nodes were obtained from Google Maps, by querying the driving distance between each pair of nodes. Note that we do not want ‘flying’ distance between two coordinates as this can be very different from the actual driving distance, especially in New York City.

The limited resources for inspection and repair of manholes should generally be designated to the most vulnerable manholes. With uncertainty in many of the probability estimates, if we are not careful, it is possible that most of these resources will be spent in dealing with outliers whose probabilities are overestimated. The simultaneous process will generally prevent this from happening if we choose C_1 to have a sufficiently large positive value.

Manhole failures are rare events. This means there are many more negative labels than positive labels. Using a logistic model gives probability estimates which are low overall, so the misclassification error is almost always the size of the whole positive class. Because of this, we evaluate the quality of the predictions from f_{λ^*} using the area under the ROC curve (AUC), for both training and test. AUC is a measure of ranking quality; it is sensitive to the rank-ordering of the nodes in terms of their probability to fail, and it is not as sensitive to changes in the values of these probabilities. This means that as the parameter C_1 increases, the estimated probability values will tend to decrease, and thus the failure cost will decrease.

Fig. 3 *Left* the AUC values corresponding to models (parameterized by C_1) obtained from the simultaneous process using Cost 1 by NM and AM and MINLP techniques. The AUC values on the training data decrease slightly and the same values for test data increase marginally. The two horizontal lines represent the training and test AUC values obtained by ℓ_2 -penalized logistic regression are constant with respect to C_1 . *Right* similar AUC values obtained from the simultaneous process, using Cost 2

For the experiment, a specific decision problem was sampled and fixed a priori, involving repairs on a handful of relatively more vulnerable manholes in the Bronx. We solved (20) and (21) for a range of values for the regularization parameter C_1 , for both costs and all three methods, with the goal of seeing whether for the same level of estimation performance, we can get a range in the cost of failures. In particular, we wanted to know if we could see a substantial reduction in the cost. We varied C_1 so that the variation in the training error term across the methods was small, about 2 % away from the solution of the sequential process ($C_1 = 0$), see Fig. 4a. For that range, the test AUC values for the simultaneous process were all within 1 % of each other; this is true for both Cost 1 and Cost 2, for each of the AM, NM, and MINLP solvers, see Fig. 3a, b. So, changing C_1 did not dramatically impact the prediction quality as measured by the AUC. On the other hand, the failure costs varied widely over the different methods and settings of C_1 , as a result of the change in the probability estimates, as shown in Fig. 4b. As C_1 was increased from 0.05 to 0.5, Cost 1 went from 27.5 units to 3.2 units, which is over eight times smaller. This means that with a 1–2 % variation in the predictive model's AUC, the operational cost can decrease a lot, yielding a completely different possible route for inspection and/or repair work. The reason for an order of magnitude change in the failure cost is because the probability estimates vary by an order of magnitude due to uncertainty at the nodes. This uncertainty in costs is what the MLOC allows us to uncover.

In Fig. 5a–c we show the routes according to the different algorithms. We first provide the naïve route in Fig. 5a, which was obtained by estimating probabilities using ℓ_2 -penalized logistic regression, and then simply visiting nodes according to decreasing values of these probabilities. Figure 5b shows the route provided by the sequential process. When the failure term starts influencing the optimal solution of the objective (20) because of an increase in C_1 , we get a new route, depicted in Fig. 5c. In most applications relevant to this problem, we suspect that the solution used in practice is somewhere in between the naïve route and the sequential route, in that a human views the naïve solution and adjusts it by hand to be closer to the sequential route (without solving the TRP). For the application to electrical grid maintenance, the simultaneous process was able to find a substantially lower cost route than the naïve or sequential process, with little (if any) change in the AUC prediction quality. This demonstration on data from the Bronx indicates that it is possible to better understand uncer-

Fig. 4 *Left* the ℓ_2 -regularized logistic loss increases as a function of increasing C_1 . The *horizontal line* represents the loss value from ℓ_2 -penalized logistic regression with no regularization ($C_1 = 0$). *Right* the failure costs decrease as a function of the regularization parameter C_1 . The *horizontal lines* in the figure represent the sequential formulation solution; the lower horizontal line is Cost 1 of the solution obtained by ℓ_2 -penalized logistic regression, and the upper line is Cost 2 of that solution

Fig. 5 *Left* a naïve route: 1-5-4-3-2-6-7-1 obtained by sorting the probability estimates in decreasing order and visiting the corresponding nodes. *Center* sequential process route: 1-5-3-4-2-6-7-1. The simultaneous process also chooses this route when C_1 is small. *Right* route chosen by the simultaneous process when C_1 is larger: 1-6-7-5-3-4-2-1. Prediction performance is only slightly influenced by the route change, but the routing cost (Cost 1) decreases a lot

tainty in modeling. If engineers truly believe the costs will be lower, their belief, combined with the route we found, can be used to justify a much more cost-effective solution.

5.3 Performance of the simultaneous process across randomly generated decision problems

In this experiment, we varied the size of the training data and characterized its effect on learning for both the sequential process and the simultaneous process. We expect to see that when the sample size is small, the operational cost regularization can lead to better performance for the simultaneous process for some C_1 . That is, we are showing that some type of knowledge on the operational cost can be helpful in prediction. (When the sample size is large, the regularization term of the simultaneous process should not have much of an effect, and the sequential and simultaneous process models should perform similarly, which is unsurprisingly what we observe).

To conduct the experiment, we considered training samples ranging from 10 % of the original training set size to 100 % of the original training set size. For each training set we generated, we then generated 100 seven node decision problems (TRP problems) from a

Fig. 6 Performance of the two processes on randomly generated decision problems at various training sample sizes with Cost 1 as the routing cost. The evaluation is over a separate held out test set. The *green solid line* is the zero mark. For each size of the training sample on the x-axis (varying from 10 to 100% of the original training sample size), we solved the simultaneous process for 100 random seven node decision problems and the performances of the corresponding models relative to the sequential process models are plotted as a box-plot (Color figure online)

separate held out test set. Each decision problem was generated by randomly picking the nodes (whose labels are not known during training) and computing the distances between each pair of them. For each new training sample size and for each random decision problem, we solved the sequential process and the simultaneous process for both Cost 1 and Cost 2. In particular, this involved the following.

- For the sequential process we performed a fivefold cross validation to pick the coefficient for the ℓ_2 regularization term. Once the optimal regularization constant was chosen, we computed the predicted probabilities of failure and solved the corresponding weighted TRP subproblem.
- We solved the simultaneous process using the AM algorithm for 4 different C_1 values, and the one achieving the best test performance (on a separate held out test set) was reported. This encodes the notion that one of the C_1 values, namely the one which gives the best test performance, encodes the right prior knowledge. In total, 8,000 mixed integer nonlinear programs were solved [4 C_1 value settings per decision problem (100) per training sample size (10) per decision cost type (Cost 1 and Cost 2)].

Figure 6 shows how the simultaneous process compares with respect to the sequential process in terms of AUC on a held out test set as the size of the training sample is varied for Cost 1. The x-axis shows different training sample sizes and the y-axis shows the difference between the AUC of a simultaneous process model (one for each training size and decision problem) and the AUC of the corresponding sequential process model, where 0 means that the AUC's for the two processes were identical. From the figures, we can infer the following:

- The test performance of the simultaneous process can often be better than that of the sequential process for smaller training sets. This is because at lower sample sizes, the

Fig. 7 Performance of the models output by the two processes on randomly generated decision problems at various training sample sizes with Cost 2 as the routing cost. The evaluation is on a separate held out test set. The green solid line is the zero mark. The box-plots at each training sample size represent the distribution of performances (relative AUC) of the models obtained by the simultaneous process (Color figure online)

simultaneous process gains an advantage from the prior knowledge about operational costs.

- At larger training set sizes, the logistic models from the simultaneous process and the sequential process performed similarly. Again this is not surprising, as the regularization becomes less influential as the training set size increases.

At each training sample size, we tested two hypotheses using the (nonparametric) sign test, with significance level $\alpha = 0.05$. In the first test, the null hypothesis was that the median AUC performance of the two processes was the same versus the alternative that the median AUC performance of the simultaneous process is greater than the median AUC performance of the sequential process. For three of the larger training sample sizes (namely 0.6, 0.7 and 0.9 of the original), we could not reject the null as the corresponding p values were greater than the significance level and for the remaining 7 training sample sizes, we could reject the null that the median performance of the two methods is the same. In the second test, the null hypothesis was that the median routing cost using the two processes was the same versus the alternative that the median routing cost of the simultaneous process is smaller than the median routing cost of the sequential process. Here, we were able to reject the null hypothesis for all 10 training sample sizes.

We ran this experiment again with Cost 2 as the routing cost, and solved the same 100 decision problems for 4 different C_1 values for each of the 10 different training samples of different sizes. Figure 7 summarizes the performance of these models. The inferences one can draw from this plot are similar to the previous case.

5.4 Scalability of MLOC for routing

In this experiment, we varied the size of the training sample and decision problem and characterized their effect on time to obtain a solution. All experiments were carried out in a cluster environment (128-256GB RAM, 16-32 core machines).

Fig. 8 *Left* boxplot of times taken to solve randomly generated 7 node decision problems for various training sample sizes (from 10 to 100 % of the original), when Cost 1 is used. For each training sample size, we solved the simultaneous process for 100 random decision problems and recorded the times. As shown, the time for solving the simultaneous process depends mildly on the size of the training sample size. *Right* boxplot of times taken to solve randomly generated 7 node decision problems for various training sample sizes when Cost 2 is used

In the first case, we analyzed the effect of training sample size when the decision problem size was fixed to 7 nodes. In particular, we generated 100 seven node decision problems for each of the 10 training sample sizes (varying from 10 to 100 % of the original) and solved the corresponding MINLPs using the AM method discussed in Sect. 4.4. As discussed before, a decision problem was created by randomly picking a set of seven nodes and computing the distances between them. Additionally, the C_2 parameter was set using fivefold cross validation. A fixed value of C_1 was also chosen a-priori. Thus a total of 1000 MINLPs were solved for each Cost 1 and Cost 2. Figure 8a, b show the box plots for the time taken in seconds to solve each simultaneous process problem for Cost 1 and Cost 2 respectively. From the figures, we can infer that as the training sample size increases, the time taken to solve the MINLP increases only mildly for both cost options. This is because the AM method can efficiently scale with the number of examples.

In the second case, we analyzed the effect of decision problem size. In particular, we generated 100 decision problems for node sizes $M = 7, 8, 9, 10, 11, 12, 13$ and 10 decision problems for node size $M = 15$. We solved the MINLPs of Eqs. (20) and (21) using the AM method. Similar to the previous experiment, a decision problem of a given size was created by randomly picking a set of nodes and computing the distances between them. The C_2 parameter was set using fivefold cross validation. The MINLPs were then solved for a fixed value of C_1 chosen a-priori. Thus a total of 710 MINLPs were solved for each Cost 1 and Cost 2. Figure 9a, b show the box plots for the time taken in seconds (in log scale) to solve each simultaneous process problem for Cost 1 and Cost 2 respectively. From the figures, we can infer that as the decision problem size (M nodes) increases, the time taken to solve the MINLP increases exponentially for both cost models. As mentioned earlier, this is because TRP—and generally routing—problems are hard. One needs to solve the TRP anyway, regardless of whether the sequential or simultaneous process is used, to determine the route.

Remark 3 A note on the performance of other methods (Method 1 and Method 3): For a given C_1 , the computation times to solve a typical problem with $\sim 23K$ examples in training and 6, 7, 8, or 10 nodes for the routing problem are about 30, 130, 140, 240s respectively using Method 2 (NM). NM took $\sim 1,000$ iterations to reach a solution where each iteration

Fig. 9 *Left* boxplot of times taken to solve the randomly generated decision problems for various values of M , the number of decision problem nodes, when Cost 1 is used. For each decision problem size (varying from 7 to 15), we solved the simultaneous process for 100 random decision problems (10 problems for the 15 node setting) and recorded the times. As shown, the time for solving the decision problem grows exponentially in the size of the decision/routing problem (since the trend is linear in log scale). *Right* boxplot of times taken to solve the randomly generated decision problems for various values of M when Cost 2 is used

involved solving a weighted TRP subproblem within ~ 2 s. The computation times for solving the MINLP formulation given in (20) directly (Method 1) for a given C_1 were ~ 100 times slower. Since the computation times for Method 2 (AM) were the best among the three, we used it to benchmark scalability of MLOC for our application.

6 Generalization bound

We initially introduced the failure cost regularization term in order to find scenarios where the data would support low-cost (more actionable) repair routes. From a learning theoretic point of view, incorporating regularization reduces the size of the hypothesis space and may thus promote generalization. In our case, we can think of decision makers having prior knowledge about how much it should cost for an optimal routing solution. This information should constrain the size of the hypothesis space via the parameter C_1 . Increasing C_1 may thus assist in predicting failure probabilities. In what follows, we will provide a generalization bound for the MLOC framework, and specifically for the ML&TRP.

We seek to bound the true risk $R^{\text{true}}(f_\lambda) := E_{(x,y) \sim \mu_{\mathcal{X} \times \mathcal{Y}}} l(f_\lambda(x), y)$ with empirical risk $R^{\text{emp}}(f_\lambda, \{x_i, y_i\}_1^m) = \frac{1}{m} \sum_{i=1}^m l(f_\lambda(x_i), y_i)$ plus a complexity term capturing the size of the hypothesis space. Here $l : f_\lambda(\mathcal{X}) \times \mathcal{Y} \rightarrow \mathbb{R}$ is logistic loss, instance (x, y) is drawn from an unknown distribution $\mu_{\mathcal{X} \times \mathcal{Y}}$ and the initial hypothesis space is $\mathcal{F} := \{f_\lambda : f_\lambda(x) = \lambda \cdot x, \lambda \in \mathbb{R}^d, \|\lambda\|_2 \leq B_b\}$.

6.1 Hypothesis sets for Cost 1 and Cost 2

Consider the ML&TRP with Cost 1 in (20). The hypothesis space for the ML&TRP is smaller than $\mathcal{F}$, since we have also the constraint on the failure cost. Replacing the Lagrange multiplier C_1 with an explicit constraint on the failure cost (6), we have that for the ML&TRP, f_λ is subject to the failure cost constraint: $\min_{\pi} \sum_{i=1}^M p(\tilde{x}_{\pi(i)}) L_{\pi}(\pi(i)) \leq C_{\text{budget}}$, where C_{budget} is inversely related to C_1 , controlling a “budget” for the failure cost. This gives us the restricted hypothesis space:

$$\mathcal{F}_0 := \left\{ f_\lambda : f_\lambda \in \mathcal{F}, \min_{\pi \in \Pi} \sum_{i=1}^M L_\pi(\pi(i)) \frac{1}{1 + e^{-f_\lambda(\tilde{x}_{\pi(i)})}} \leq C_{\text{budget}} \right\}.$$

Even though $\mathcal{F}_0$ is smaller than $\mathcal{F}$, it is difficult to construct a tight bound on its covering number. So we enlarge $\mathcal{F}_0$ just enough so that a bound on its covering number can be calculated. In particular, we will enlarge the set $\mathcal{F}_0$ to the set $\mathcal{F}_2$. We define set $\mathcal{F}_2$ parametrized by a vector $a_{\text{budget}} \in \mathbb{R}^d$ as follows:

$$\mathcal{F}_2 := \{ f_\lambda : f_\lambda \in \mathcal{F}, a_{\text{budget}} \cdot \lambda \leq 1 \},$$

where vector a_{budget} is a function of C_{budget} , the graph and the unlabeled data $\{\tilde{x}_i\}_i$.

$\mathcal{F}_2$ is the intersection of the ball $\mathcal{F}$ with the halfspace defined by a_{budget} ; it is a ball that is missing a spherical cap. The vector a_{budget} will capture the effect of C_{budget} in such a way that $\mathcal{F}_0 \subset \mathcal{F}_2$, which we will show within the proof of the Theorem 1. $\mathcal{F}_2$ is the space whose complexity we will bound, again within the proof of Theorem 1.

We will now define the vector a_{budget} in terms of C_{budget} and provide a proof later. Let d_i be the shortest distance from the starting node (node 1) to node i for $i = 2, \dots, M$ and d_1 be the length of the shortest tour that visits all the nodes and returns to node 1. This means $d_i \leq L_\pi(i)$; $i = 1, \dots, M$ with equality if the physical graph can be embedded into 1-dimensional Euclidean space. The vector a_{budget} is then related to C_{budget} defined elementwise as:

$$a_{\text{budget}}^j = \frac{1}{C_{\text{budget}} - a_0} \left(\frac{e^{B_b X_b}}{(1 + e^{B_b X_b})^2} \right) \left(\sum_i d_i \tilde{x}_i^j \right) \text{ for } j = 1, \dots, d \quad (22)$$

$$\text{where } a_0 = \left(B_b X_b \frac{e^{B_b X_b}}{(1 + e^{B_b X_b})^2} + \frac{1}{1 + e^{B_b X_b}} \right) \sum_i d_i.$$

Remark 4 (Defining $\mathcal{F}_0$, $\mathcal{F}_2$ and a_{budget} for Cost 2): The definitions of $\mathcal{F}_0$ and $\mathcal{F}_2$ can be easily adapted to Cost 2 in (21) of the ML&TRP. Here too, the hypothesis space for the ML&TRP is smaller than $\mathcal{F}$ because of the constraint on the failure cost. Again replacing the Lagrange multiplier C_1 with an explicit constraint on the failure cost, we have that for the ML&TRP, f_λ is subject to the failure cost constraint: $\min_{\pi} \sum_{i=1}^M \log(1 + e^{\lambda \cdot \tilde{x}_{\pi(i)}}) L_\pi(\pi(i)) \leq C_{\text{budget}}$, where C_{budget} is inversely related to C_1 , controlling a “budget” for the failure cost. This gives us the restricted hypothesis space:

$$\mathcal{F}_0 := \left\{ f_\lambda : f_\lambda \in \mathcal{F}, \min_{\pi \in \Pi} \sum_{i=1}^M L_\pi(\pi(i)) \log(1 + e^{f_\lambda(\tilde{x}_{\pi(i)})}) \leq C_{\text{budget}} \right\}.$$

We can again enlarge this class of functions just enough so that a bound on the covering number of $\mathcal{F}_0$ can be calculated. The enlarged set $\mathcal{F}_2$ will have the same form as for Cost 1 except for a different definition of a_{budget} (we will derive this later):

$$a_{\text{budget}}^j = \frac{1}{C_{\text{budget}} - a_0} \left(\frac{e^{-B_b X_b}}{1 + e^{-B_b X_b}} \right) \left(\sum_i d_i \tilde{x}_i^j \right) \text{ for } j = 1, \dots, d \quad (23)$$

$$\text{where } a_0 = \left(B_b X_b \frac{e^{-B_b X_b}}{1 + e^{-B_b X_b}} + \log(1 + e^{-B_b X_b}) \right) \sum_i d_i.$$

Since Cost 2 can be handled in the same way as Cost 1, we will focus on Cost 1 for the rest of this section.

6.2 Main generalization result

Recall that we would like to establish that generalization can depend on C_{budget} . The following theorem shows this explicitly. C_{budget} enters the bound through the vector a_{budget} .

Theorem 1 (Main Result) *Let $\mathcal{X} = \{x \in \mathbb{R}^d : \|x\|_2 \leq X_b\}$, $\mathcal{Y} = \{-1, 1\}$. Let $\mathcal{F}_0$ be defined as above with respect to $\{\tilde{x}_i\}_{i=1}^M$, $\tilde{x}_i \in \mathcal{X}$ (not necessarily random) and a corresponding physical graph. Let $\{x_i, y_i\}_{i=1}^m$ be a sequence of m instances drawn independently according to unknown distribution $\mu_{\mathcal{X} \times \mathcal{Y}}$ and*

$M_{\text{bound}} := B_b X_b + \log 2$. For any $\epsilon > 0$,

$$\begin{aligned} &P\left(\exists f \in \mathcal{F}_0 : |R^{\text{emp}}(f_\lambda, \{x_i, y_i\}_1^m) - R^{\text{true}}(f_\lambda)| > \epsilon\right) \\ &\leq 4\alpha(d, a_{\text{budget}}(C_{\text{budget}})) \left(\frac{32B_b X_b}{\epsilon} + 1\right)^d \exp\left(\frac{-m\epsilon^2}{128M_{\text{bound}}^2}\right), \end{aligned}$$

where $\alpha(d, a_{\text{budget}}(C_{\text{budget}}))$ is equal to

$$\frac{1}{2} + \frac{\|a_{\text{budget}}\|_2^{-1} + \frac{\epsilon}{32X_b}}{B_b + \frac{\epsilon}{32X_b}} \frac{\Gamma\left[1 + \frac{d}{2}\right]}{\sqrt{\pi} \Gamma\left[\frac{d+1}{2}\right]} {}_2F_1\left(\frac{1}{2}, \frac{1-d}{2}, \frac{3}{2}; \left(\frac{\|a_{\text{budget}}\|_2^{-1} + \frac{\epsilon}{32X_b}}{B_b + \frac{\epsilon}{32X_b}}\right)^2\right) \quad (24)$$

$$\text{or equivalently, } 1 - \frac{1}{2} I_{1 - \left(\|a_{\text{budget}}\|_2^{-1} + \frac{\epsilon}{32X_b}\right)^2 / \left(B_b + \frac{\epsilon}{32X_b}\right)^2} \left(\frac{d+1}{2}, \frac{1}{2}\right) \quad (25)$$

and where ${}_2F_1(a, b; c; d)$ and $I_x(a, b)$ are the hypergeometric function and the regularized incomplete beta functions respectively.

The term $\alpha(d, a_{\text{budget}}(C_{\text{budget}}))$ comes directly from formulae for the volume of spherical caps. As C_{budget} decreases, the norm $\|a_{\text{budget}}\|_2$ increases, and thus $\|a_{\text{budget}}\|_2^{-1}$ decreases, (24) and (25) decrease, and the whole bound decreases. This is the mechanism by which decreasing C_{budget} may improve generalization ability.

Theorem 1 is specific to the ML&TRP because $\mathcal{F}_0$ was defined based on the ML&TRP and a_{budget} was defined in (22) for Cost 1 and (23) for Cost 2.

The technique of Theorem 1 applies much more broadly than the ML&TRP. In fact, we can derive a general bound that applies to any problem with a similar hypothesis space constraint. Specifically, the hypothesis space should be bounded by the intersection of a ball with a half-space.

Corollary 1 (Bound for General MLOC Framework) *Consider any operational cost constraint such that the hypothesis space lies within $\mathcal{F}_2$ defined by $\mathcal{F}_2 = \{f_\lambda \in \mathcal{F} : a_{\text{budget}} \cdot \lambda \leq 1\}$ for some $a_{\text{budget}} \in \mathbb{R}^d$. Then, for any $\epsilon > 0$,*

$$\begin{aligned} &P\left(\exists f \in \mathcal{F}_2 : |R^{\text{emp}}(f_\lambda, \{x_i, y_i\}_1^m) - R^{\text{true}}(f_\lambda)| > \epsilon\right) \\ &\leq 4\alpha(d, a_{\text{budget}}) \left(\frac{32B_b X_b}{\epsilon} + 1\right)^d \exp\left(\frac{-m\epsilon^2}{128M_{\text{bound}}^2}\right), \end{aligned}$$

where $\alpha(d, a_{\text{budget}})$ equals

$$\frac{1}{2} + \frac{\|a_{\text{budget}}\|_2^{-1} + \frac{\epsilon}{32X_b}}{B_b + \frac{\epsilon}{32X_b}} \frac{\Gamma\left[1 + \frac{d}{2}\right]}{\sqrt{\pi} \Gamma\left[\frac{d+1}{2}\right]} {}_2F_1\left(\frac{1}{2}, \frac{1-d}{2}, \frac{3}{2}; \left(\frac{\|a_{\text{budget}}\|_2^{-1} + \frac{\epsilon}{32X_b}}{B_b + \frac{\epsilon}{32X_b}}\right)^2\right)$$

$$\text{or equivalently, } 1 - \frac{1}{2} I_{1 - \left(\|a_{\text{budget}}\|_2^{-1} + \frac{\epsilon}{32\bar{X}_b} \right)^2 / \left(B_b + \frac{\epsilon}{32\bar{X}_b} \right)^2} \left(\frac{d+1}{2}, \frac{1}{2} \right)$$

and where ${}_2F_1(a, b; c; d)$ and $I_x(a, b)$ are the hypergeometric function and the regularized incomplete beta functions respectively.

The $\alpha(d, a_{\text{budget}})$ is influenced by our belief on the operational cost. Thus, by being able to specify something about the operational cost, we are able to have a better guarantee on generalization. In the case where we are not able to specify anything about the operational cost, the quantity $\alpha(d, a_{\text{budget}})$ is equal to 1 giving us the standard generalization result for norm constrained linear function classes.

6.3 Proof

The proof outline is as follows. We will construct two classes, $\mathcal{F}_1$ and $\mathcal{F}_2$ that are slightly larger than $\mathcal{F}_0$, but smaller than $\mathcal{F}$ when C_{budget} is small enough. Then we will use a volumetric argument to bound the covering number of $\mathcal{F}_2$, which uses the volumes of spherical caps; the idea is to show that the value of C_{budget} affects the volume of the hypothesis space, and thus the covering number. The covering number bound is then applied to a uniform bound of Pollard (1984) to obtain a generalization bound. The fact that the covering number of $\mathcal{F}_2$ can be below that of $\mathcal{F}$ indicates that using functions from $\mathcal{F}_2$ may provide improvements in generalization over using the full set $\mathcal{F}$.

Let us lead up to the proof of Theorem 1.

Definition 1 Let $A \subseteq X$ be an arbitrary set and (X, dist) a (pseudo) metric space. Let $|\cdot|$ denote set size.

- For any $\epsilon > 0$, an ϵ -cover for A is a finite set $U \subseteq X$ (not necessarily $\subseteq A$) s.t. $\forall x \in A, \exists u \in U$ with $\text{dist}(x, u) \leq \epsilon$.
- A is totally bounded if A has a finite ϵ -cover for all $\epsilon > 0$. The *covering number* of A is $N(\epsilon, A, \text{dist}) := \inf_{U \in \mathcal{U}} |U|$ where $\mathcal{U}$ is the set of all ϵ -covers for A .
- A set $R \subseteq X$ is ϵ -separated if $\forall x, y \in R, \text{dist}(x, y) > \epsilon$. The *packing number* $M(\epsilon, A, \text{dist}) := \sup_{R \in \mathcal{R}} |R|$, where $\mathcal{R}$ is the set of all ϵ -separated subsets of A .

Consider Cost 1. Since, for any collection of values $p(\tilde{x}_i) \geq 0, \sum_i d_i p(\tilde{x}_i) \leq \sum_i L_\pi(i) p(\tilde{x}_i) \leq C_{\text{budget}}$, the class of functions which obey the constraint $\sum_i d_i p(\tilde{x}_i) \leq C_{\text{budget}}$ is larger than the class obeying $\sum_i L_\pi(i) p(\tilde{x}_i) \leq C_{\text{budget}}$. That is, $\mathcal{F}_0 \subseteq \mathcal{F}_1$ where

$$\mathcal{F}_1 := \left\{ f_\lambda : f_\lambda \in \mathcal{F}, \sum_{i=1}^M d_i \frac{1}{1 + e^{-f_\lambda(\tilde{x}_i)}} \leq C_{\text{budget}} \right\}.$$

As long as $C_{\text{budget}} \leq \sum_{i=1}^M d_i$, the constraint in $\mathcal{F}_1$ is not vacuous. The choice of the vector a_{budget} ensures that $\mathcal{F}_1$ is a subset of $\mathcal{F}_2$ as we will prove below.

Lemma 1 ($\mathcal{F}_0$ is contained in $\mathcal{F}_2$)

$$N(\epsilon, \mathcal{F}_0, \|\cdot\|_{L_2(\mu_{\mathcal{X}}^m)}) \leq N(\epsilon, \mathcal{F}_1, \|\cdot\|_{L_2(\mu_{\mathcal{X}}^m)}) \leq N(\epsilon, \mathcal{F}_2, \|\cdot\|_{L_2(\mu_{\mathcal{X}}^m)}).$$

Proof It is sufficient to show $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \mathcal{F}_2$. The first inequality was discussed earlier; since $d_i = \inf_{\pi \in \Pi} L_\pi(i)$, this implies:

$$\sum_{i=1}^M d_i p(\tilde{x}_i) \leq \sum_{i=1}^M L_\pi(i) p(\tilde{x}_i) \leq C_{\text{budget}} \Rightarrow \mathcal{F}_0 \subseteq \mathcal{F}_1.$$

We now show $\mathcal{F}_1 \subseteq \mathcal{F}_2$. We first lower bound $p(\tilde{x}_i)$ by a line with slope $m_1 := \frac{e^{B_b X_b}}{(1+e^{B_b X_b})^2}$ and intercept $m_0 := B_b X_b \frac{e^{B_b X_b}}{(1+e^{B_b X_b})^2} + \frac{1}{1+e^{B_b X_b}}$ such that $m_1 f_\lambda(\tilde{x}_i) + m_0 \leq p(\tilde{x}_i)$ within the function range $[-B_b X_b, B_b X_b]$.

This leads to the definition of a_{budget} as we show now:

$$\sum_i d_i p(\tilde{x}_i) \geq \sum_i d_i (m_1(\lambda \cdot \tilde{x}_i) + m_0) = \tilde{a} \cdot \lambda + a_0, \quad (26)$$

$$\text{where } \tilde{a}^j := m_1 \left(\sum_i d_i \tilde{x}_i^j \right) = \frac{e^{B_b X_b}}{(1+e^{B_b X_b})^2} \left(\sum_i d_i \tilde{x}_i^j \right) \text{ for } j = 1, \dots, d \quad (27)$$

$$\text{and } a_0 = m_0 \sum_i d_i = \left(B_b X_b \frac{e^{B_b X_b}}{(1+e^{B_b X_b})^2} + \frac{1}{1+e^{B_b X_b}} \right) \sum_i d_i.$$

$$\text{Thus } \forall \lambda \in \mathcal{F}_1, \tilde{a} \cdot \lambda + a_0 \leq \sum_{i=1}^M d_i p(\tilde{x}_i) \leq C_{\text{budget}}, \quad (28)$$

which implies $\tilde{a} \cdot \lambda \leq C_{\text{budget}} - a_0$ or equivalently, $\frac{1}{C_{\text{budget}} - a_0} \tilde{a} \cdot \lambda \leq 1$.

This allows us to define a_{budget} using (27) as

$$a_{\text{budget}}^j = \frac{1}{C_{\text{budget}} - a_0} \left(\frac{e^{B_b X_b}}{(1+e^{B_b X_b})^2} \right) \left(\sum_i d_i \tilde{x}_i^j \right) \text{ for } j = 1, \dots, d,$$

which is the same as (22). This vector is such that the set $\mathcal{F}_2$ is larger than $\mathcal{F}_1$. $\square$

Remark 5 (Deriving a_{budget} for Cost 2): The above lemma can be adapted to Cost 2 to give the corresponding a_{budget} that we had defined earlier. In particular, for any collection of values $\log(1 + e^{\lambda \cdot \tilde{x}_i}) \geq 0$ for all i ,

$$\sum_i d_i \log(1 + e^{\lambda \cdot \tilde{x}_i}) \leq \sum_i L_\pi(i) \log(1 + e^{\lambda \cdot \tilde{x}_i}).$$

Thus the class of functions that obey the constraint $\sum_i d_i \log(1 + e^{\lambda \cdot \tilde{x}_i}) \leq C_{\text{budget}}$ is larger than the class obeying $\sum_i L_\pi(i) \log(1 + e^{\lambda \cdot \tilde{x}_i}) \leq C_{\text{budget}}$, which is $\mathcal{F}_0$. $\mathcal{F}_1$ will be the set corresponding to the former constraint:

$$\mathcal{F}_1 := \left\{ f_\lambda \in \mathcal{F} : \sum_{i=1}^M d_i \log(1 + e^{\lambda \cdot \tilde{x}_i}) \leq C_{\text{budget}} \right\}.$$

We now define $\mathcal{F}_2$ and a_{budget} as follows. We can also see that $\log(1 + e^{\lambda \cdot \tilde{x}_i})$ can be lower bounded by a line with slope $m_1 := \frac{e^{-B_b X_b}}{1+e^{-B_b X_b}}$ and intercept $m_0 := B_b X_b \frac{e^{-B_b X_b}}{1+e^{-B_b X_b}} + \log(1 + e^{-B_b X_b})$ in the function range $[-B_b X_b, B_b X_b]$ giving us the definition of a_{budget} for Cost 2 as follows:

$$C_{\text{budget}} \geq \sum_i d_i \log(1 + e^{\lambda \cdot \tilde{x}_i}) \geq \sum_i d_i (m_1(\lambda \cdot \tilde{x}_i) + m_0) = \tilde{a} \cdot \lambda + a_0,$$

$$\text{where } \tilde{a}^j := m_1 \left(\sum_i d_i \tilde{x}_i^j \right) = \frac{e^{-B_b X_b}}{1+e^{-B_b X_b}} \left(\sum_i d_i \tilde{x}_i^j \right) \text{ for } j = 1, \dots, d$$

$$\text{and } a_0 = m_0 \sum_i d_i = \left(B_b X_b \frac{e^{-B_b X_b}}{1 + e^{-B_b X_b}} + \log(1 + e^{-B_b X_b}) \right) \sum_i d_i.$$

Thus, $\frac{1}{C_{\text{budget}} - a_0} \tilde{a} \cdot \lambda \leq 1$, and since we wanted to have $a_{\text{budget}} \cdot \lambda \leq 1$ we define a_{budget} element-wise as:

$$a_{\text{budget}}^j = \frac{1}{C_{\text{budget}} - a_0} \left(\frac{e^{-B_b X_b}}{1 + e^{-B_b X_b}} \right) \left(\sum_i d_i \tilde{x}_i^j \right) \text{ for } j = 1, \dots, d.$$

Note that we have produced two a_{budget} vectors for each of the two costs: Cost 1 and Cost 2 above.

Let $B(0, B_b) := \{\lambda : \lambda \in \mathbb{R}^d, \|\lambda\|_2 \leq B_b\}$. Let the half space corresponding to $\mathcal{F}_2$ be $H_{\|a_{\text{budget}}\|_2^{-1}} := \{\lambda : a_{\text{budget}} \cdot \lambda \leq 1\}$. The lemma below relates covering numbers of $\mathcal{F}$ and $\mathcal{F}_2$ in function space to covering numbers of $B(0, B_b)$ and $B(0, B_b) \cap H_{\|a_{\text{budget}}\|_2^{-1}}$ in $\mathbb{R}^d$.

Lemma 2 (Relating covering numbers in $\|\cdot\|_{L_2(\mu_{\mathcal{X}}^m)}$ to $\|\cdot\|_2$)

- $\sup_{\mu_{\mathcal{X}}^m} N(\epsilon, \mathcal{F}, \|\cdot\|_{L_2(\mu_{\mathcal{X}}^m)}) \leq N(\epsilon/X_b, B(0, B_b), \|\cdot\|_2)$, and
- $\sup_{\mu_{\mathcal{X}}^m} N(\epsilon, \mathcal{F}_2, \|\cdot\|_{L_2(\mu_{\mathcal{X}}^m)}) \leq N(\epsilon/X_b, B(0, B_b) \cap H_{\|a_{\text{budget}}\|_2^{-1}}, \|\cdot\|_2)$.

Proof Each element $f \in \mathcal{F}$ corresponds to at least one element of $B(0, B_b)$ by definition of $\mathcal{F}$. Choose any distribution $\mu_{\mathcal{X}}^m$. Consider two elements $\lambda_f, \lambda_g \in B(0, B_b)$ corresponding to functions $f, g \in \mathcal{F} \subset L_2(\mu_{\mathcal{X}}^m)$. Then,

$$\begin{aligned} \|f - g\|_{L_2(\mu_{\mathcal{X}}^m)}^2 &= \frac{1}{m} \sum_{i=1}^m (f(x_i) - g(x_i))^2 \\ &= \frac{1}{m} \sum_{i=1}^m ((\lambda_f - \lambda_g) \cdot x_i)^2 \\ &\leq \frac{1}{m} \sum_{i=1}^m \|\lambda_f - \lambda_g\|_2^2 \|x_i\|_2^2 \quad (\text{Cauchy-Schwarz to each term}) \\ &\leq \|\lambda_f - \lambda_g\|_2^2 \left(\frac{1}{m} \sum_{i=1}^m X_b^2 \right) \quad (\text{since } \sup_{x \in \mathcal{X}} \|x\|_2 \leq X_b) \\ &= \|\lambda_f - \lambda_g\|_2^2 X_b^2. \end{aligned}$$

Consider a minimal ϵ/X_b -cover $\{\lambda_r\}_r$ for $B(0, B_b)$ where λ_r corresponds to a function $r \in \mathcal{F}$. Then by definition, $\forall \lambda \in B(0, B_b)$, $\exists \lambda_r : \|\lambda - \lambda_r\|_2 \leq \epsilon/X_b$. Thus, picking any two such elements λ_f, λ_g in a ball of radius ϵ/X_b around λ_r , we see that, the corresponding functions f, g belong to a ball of radius ϵ measured using distance in $L_2(\mu_{\mathcal{X}}^m)$ by the inequality above. The centers of these ϵ -balls in $L_2(\mu_{\mathcal{X}}^m)$ form an ϵ -cover for $\mathcal{F}$. The size of this set is equal to $N(\epsilon/X_b, B(0, B_b), \|\cdot\|_2)$ (which is the size of ϵ/X_b -cover for $B(0, B_b)$). The size of the minimal ϵ -cover of $\mathcal{F}$ will be less than or equal to this size. Hence, $N(\epsilon, \mathcal{F}, \|\cdot\|_{L_2(\mu_{\mathcal{X}}^m)}) \leq N(\epsilon/X_b, B(0, B_b), \|\cdot\|_2)$. Taking a supremum over all $\mu_{\mathcal{X}}^m$, we obtain the first inequality of the lemma. The same argument also works for the second inequality. $\square$

Because of rotational symmetry of $B(0, B_b)$, the volume cut off by a hyperplane $a_{\text{budget}} \cdot \lambda = 1$ from $B(0, B_b)$ is determined only by its distance from the origin, which is $1/\|a_{\text{budget}}\|_2$.

Such a portion (or its complement, if smaller) of a ball obtained from slicing the ball with a hyperplane is called a spherical cap. It can be parameterized by the distance of its (hyper)plane base from the center of the ball as shown below. For notation, let the volume of a set $A \subset \mathbb{R}^d$ be represented as $\text{Vol}(A)$. For example, $\text{Vol}(B_1) = \frac{\pi^{d/2}}{\Gamma[d/2+1]}$.

Lemma 3 (Volume of spherical caps) *Let the volume of ball $B(0, B_b)$ in $\mathbb{R}^d$ be denoted as $\text{Vol}(B(0, B_b))$. Given a d -dimensional vector a , let $z = \|a\|_2^{-1}$ be a number and $H_z = \{\lambda : a \cdot \lambda \leq 1\}$ be a half space parameterized by z . Let the spherical cap be denoted by $B(0, B_b) \cap H'_z$ where the cap is at a distance z (measured from the base of the cap to the center of the ball), and H'_z represents the complement half space ($H_z \cup H'_z = \mathbb{R}^d$). Then, $\text{Vol}(B(0, B_b) \cap H'_z) / \text{Vol}(B(0, B_b))$ is equal to two expressions:*

$$\left(\frac{1}{2} - \frac{z}{B_b} \frac{\Gamma\left[\frac{1+d}{2}\right]}{\sqrt{\pi}\Gamma\left[\frac{d+1}{2}\right]} {}_2F_1\left(\frac{1}{2}, \frac{1-d}{2}; \frac{3}{2}; \left(\frac{z}{B_b}\right)^2\right) \right) = \frac{1}{2} I_{1-z^2/B_b^2}\left(\frac{d+1}{2}, \frac{1}{2}\right),$$

where ${}_2F_1(a, b; c; d)$ and $I_x(e, f)$ are the hypergeometric and regularized incomplete beta functions respectively.

Proof See Li (2011) and references therein.

Next, we need the relationship between packing numbers and covering numbers to prove Theorem 2:

Lemma 4 (Packing and covering numbers) *For every (pseudo) metric space (X, dist) , $A \subseteq X$, and $\epsilon > 0$,*

$$N(\epsilon, A, \text{dist}) \leq M(\epsilon, A, \text{dist}).$$

Proof See Theorem 4 in Kolmogorov and Tikhomirov (1959) or Theorem 12.1 in Anthony and Bartlett (1999) for a proof of this classical result.

We use the above lemma to obtain bounds for the covering numbers of subsets of $\mathbb{R}^d$ which appeared in Lemma 2.

Theorem 2 (Bound on Covering Numbers)

$$N(\epsilon/X_b, B(0, B_b), \|\cdot\|_2) \leq \left(\frac{2B_b X_b}{\epsilon} + 1 \right)^d, \text{ and}$$

$$N\left(\epsilon/X_b, B(0, B_b) \cap H_{\|a\|_2^{-1}}, \|\cdot\|_2\right) \leq \left(\frac{\text{Vol}\left(B_{B_b + \frac{\epsilon}{2X_b}} \cap H_{\|a\|_2^{-1} + \frac{\epsilon}{2X_b}}\right)}{\text{Vol}\left(B_{B_b + \frac{\epsilon}{2X_b}}\right)} \right) \left(\frac{2B_b X_b}{\epsilon} + 1 \right)^d.$$

Proof Both statements involve a volumetric argument. For a proof of the first inequality, see Sect. 3 of Kolmogorov and Tikhomirov (1959) or Lemma 4.10 in Pisier (1989) or Lorentz (1966) or Lemma 3 in Cucker and Smale (2002).

To show the second part, let the volume of the complement of the spherical cap be $\text{Vol}(B(0, B_b) \cap H_{\|a\|_2^{-1}})$; we need to find an upper bound for the minimal ϵ/X_b -cover of this set. We can do that by scaling a minimal ϵ -cover, which we find now. By extending the boundary of $B(0, B_b) \cap H_{\|a\|_2^{-1}}$ by $\epsilon/2$ we can bound the maximal packing number $M(\epsilon, B(0, B_b) \cap H_{\|a\|_2^{-1}}, \|\cdot\|_2)$ as follows:

$$M(\epsilon, B(0, B_b) \cap H_{\|a\|_2^{-1}}, \|\cdot\|_2) \times \text{Vol}(B_1)(\epsilon/2)^d \leq \text{Vol}(B_{B_b + \epsilon/2} \cap H_{\|a\|_2^{-1} + \epsilon/2}).$$

$$\begin{aligned}
\text{Or, } M(\epsilon, B(0, B_b) \cap H_{\|a\|_2^{-1}}, \|\cdot\|_2) &\leq \frac{\text{Vol}(B_{B_b+\epsilon/2} \cap H_{\|a\|_2^{-1}+\epsilon/2})}{\text{Vol}(B_1)} \frac{1}{(\epsilon/2)^d} \\
&= \frac{\text{Vol}(B_{B_b+\epsilon/2} \cap H_{\|a\|_2^{-1}+\epsilon/2})}{\text{Vol}(B_1)} \frac{1}{(\epsilon/2)^d} \frac{(B_b + \epsilon/2)^d}{(B_b + \epsilon/2)^d} \\
&= \frac{\text{Vol}(B_{B_b+\epsilon/2} \cap H_{\|a\|_2^{-1}+\epsilon/2})}{\text{Vol}(B_{B_b+\epsilon/2})} \frac{(B_b + \epsilon/2)^d}{(\epsilon/2)^d}.
\end{aligned}$$

Again, scaling ϵ to ϵ/X_b and using the relationship between $N(\epsilon, A, \text{dist})$ and $M(\epsilon, A, \text{dist})$ in Lemma 4 yields the second result. $\square$

Thus we have so far shown the relationship between covering numbers of $\mathcal{F}_0$, $\mathcal{F}_1$, and $\mathcal{F}_2$ in terms of a certain metric in Lemma 1, we have shown how those covering numbers are related to covering numbers in $\ell_2(\mathbb{R}^d)$ in Lemma 2, we have shown how the latter covering numbers relate to volumes in $\ell_2(\mathbb{R}^d)$ in Theorem 2, and we have shown how to compute one of these volumes in Lemma 3.

To complete the proof of Theorem 1, we will use a relation between the covering number of a class of loss functions of some set $\mathcal{G}$ and the covering number of the set $\mathcal{G}$ itself. We will also use a uniform convergence bound of Pollard (1984).

Theorem 3 (Pollard 1984) *Let $l_{\mathcal{G}}$ be a set of functions on $\mathcal{X} \times \mathcal{Y}$ with $0 \leq l(f_{\lambda}(x), y) \leq M_{\text{bound}}$, $\forall l \in l_{\mathcal{G}}$ and $\forall (x, y) \in \mathcal{X} \times \mathcal{Y}$. Let $\{x_i, y_i\}_1^m$ be a sequence of m instances drawn independently according to $\mu_{\mathcal{X} \times \mathcal{Y}}$. Then for any $\epsilon > 0$,*

$$\begin{aligned}
P(\exists l \in l_{\mathcal{G}} : |R^{\text{emp}}(f_{\lambda}, \{x_i, y_i\}_1^m) - R^{\text{true}}(f_{\lambda})| > \epsilon) \\
\leq 4E \left[N \left(\epsilon/16, l_{\mathcal{G}}, \|\cdot\|_{L_1(\mu_{\mathcal{X} \times \mathcal{Y}}^m)} \right) \right] \exp \left(\frac{-m\epsilon^2}{128M_{\text{bound}}^2} \right).
\end{aligned}$$

Proof See Theorem 24 in Pollard (1984) (also in Zhang 2002 Theorem 1).

We can relate the covering number for Pollard's loss functions set $l_{\mathcal{G}}$ to the covering number for set $\mathcal{G}$ as follows.

Lemma 5 (Relating $l_{\mathcal{G}}$ to $\mathcal{G}$) *If every function from function class $l_{\mathcal{G}}$ represented as $l : f(\mathcal{X}) \times \mathcal{Y} \mapsto \mathbb{R}$, $f \in \mathcal{G}$, is Lipschitz in its first argument with Lipschitz constant $\mathcal{L}$, then the covering number of $l_{\mathcal{G}}$ is related to the covering number of $\mathcal{G}$ by*

$$\sup_{\mu_{\mathcal{X} \times \mathcal{Y}}^m} N \left(\epsilon, l_{\mathcal{G}}, \|\cdot\|_{L_1(\mu_{\mathcal{X} \times \mathcal{Y}}^m)} \right) \leq N \left(\epsilon/\mathcal{L}, \mathcal{G}, \|\cdot\|_{L_1(\mu_{\mathcal{X}}^m)} \right).$$

Proof Consider two functions $f, g \in \mathcal{G}$. Let the corresponding functions in class $l_{\mathcal{G}}$ be $l_f = l(f(x), y)$ and $l_g = l(g(x), y)$.

$$\begin{aligned}
\|l_f - l_g\|_{L_1(\mu_{\mathcal{X} \times \mathcal{Y}}^m)} &= \frac{1}{m} \sum_{i=1}^m |l(f(x_i), y_i) - l(g(x_i), y_i)| \\
&\leq \frac{1}{m} \sum_{i=1}^m \mathcal{L} |f(x_i) - g(x_i)| = \mathcal{L} \|f - g\|_{L_1(\mu_{\mathcal{X}}^m)}.
\end{aligned}$$

This implies, given $\{x_i, y_i\}_{i=1}^m$, if $\hat{\mathcal{G}}$ is a minimal $\epsilon/\mathcal{L}$ -cover of $\mathcal{G}$ in $L_1(\mu_{\mathcal{X}}^m)$, we can construct an ϵ -cover of $l_{\mathcal{G}}$ in $L_1(\mu_{\mathcal{X} \times \mathcal{Y}}^m)$ as $\hat{l}_{\mathcal{G}} = \{l_{f_i} : f_i \in \hat{\mathcal{G}}\}$. The size of the minimal ϵ -cover will be

smaller than the size of such an ϵ -cover. Taking the supremum over all empirical distributions, we get the desired result. $\square$

Theorem 3 and Lemma 5 involve L_1 covering numbers, but our covering number bounds start with an L_2 metric in Lemma 2. So we need to switch from L_1 to L_2 metric. The following lemma uses the identity $\|f - g\|_{L_1(\mu_X^m)} \leq \|f - g\|_{L_2(\mu_X^m)}$ (true because of Jensen's inequality applied to norms) to relate the two.

Lemma 6 $N(\epsilon, A, \|\cdot\|_{L_1(\mu_X^m)}) \leq N(\epsilon, A, \|\cdot\|_{L_2(\mu_X^m)})$.

Proof See for a version, Lemma 10.5 in [Anthony and Bartlett \(1999\)](#).

Finally, we can prove the main result.

Proof (Of Theorem 1)

In our setting, the loss function is logistic with Lipschitz constant $\mathcal{L} = 1$ (when viewed as a function of $f(x)$). The class of loss functions is thus defined by $\mathcal{F}_0 := \{l : f_\lambda \in \mathcal{F}_0\}$. Each $l \in \mathcal{F}_0$ is also non-negative and bounded as needed in the statement of Theorem 3.

Starting from the expectation term on the right hand side of Theorem 3 using $\mathcal{F}_0$ as $\mathcal{G}$ we get,

$$\begin{aligned} & E[N(\epsilon/16, \mathcal{F}_0, \|\cdot\|_{L_1(\mu_{\mathcal{X} \times \mathcal{Y}}^m)})] \\ & \leq \sup_{\mu_{\mathcal{X} \times \mathcal{Y}}^m} N(\epsilon/16, \mathcal{F}_0, \|\cdot\|_{L_1(\mu_{\mathcal{X} \times \mathcal{Y}}^m)}) \text{ bounding expectation by supremum} \\ & \leq \sup_{\mu_{\mathcal{X}}^m} N\left(\frac{\epsilon}{16\mathcal{L}}, \mathcal{F}_2, \|\cdot\|_{L_2(\mu_{\mathcal{X}}^m)}\right) \text{ from Lemma 5,6 and 1 respectively} \\ & \leq N\left(\frac{\epsilon}{16 \cdot 1 \cdot X_b}, B(0, B_b) \cap H_{\|a_{\text{budget}}\|_2^{-1}}, \|\cdot\|_2\right) \text{ from Lemma 2, substituting } \mathcal{L} = 1 \\ & \leq \left(\frac{\text{Vol}\left(B_{B_b + \frac{\epsilon}{32X_b}} \cap H_{\|a_{\text{budget}}\|_2^{-1} + \frac{\epsilon}{32X_b}}\right)}{\text{Vol}(B_{B_b + \frac{\epsilon}{32X_b}})}\right) \left(\frac{32B_b X_b}{\epsilon} + 1\right)^d \text{ from Theorem 2} \\ & = \alpha(d, a_{\text{budget}}(C_{\text{budget}})) \left(\frac{32B_b X_b}{\epsilon} + 1\right)^d \text{ from Lemma 3.} \end{aligned}$$

The above step uses the relation between spherical cap and its complement along with Lemma

3, $\text{Vol}\left(B(0, B_b) \cap H'_{\|a_{\text{budget}}\|_2^{-1}}\right) = \text{Vol}(B(0, B_b)) - \text{Vol}\left(B(0, B_b) \cap H_{\|a_{\text{budget}}\|_2^{-1}}\right)$.

Using the derived inequality within Theorem 3 completes the proof. $\square$

7 Future work

We provide several avenues for future work.

- *Other graph applications:* The MLOC framework is a general tool that can help decision makers translate uncertainty in prediction to uncertainty in operational costs. The ML&TRP itself is a specific application of the MLOC framework that can be applied to the power grid (as we did), but also to delivery truck routing and other physical routing problems, and can be used for more abstract routing problems such as network routing problems, where distances on the graph do not necessarily correspond to a physical distance. In the future it would be interesting to explore some of these applications.

- *Relaxing the cost constraints in the MLOC*: Our generalization bound for the ML&TRP applied to a hypothesis space that was an intersection of an l_2 ball with a halfspace. It would be interesting to consider more general operational cost constraints, such as quadratic constraints and other convex functions. As it turns out, there are many applications where such constraints naturally arise. In current work, we are constructing bounds for these types of constraints, which lead to exotic hypothesis spaces, such as an intersection of an l_2 ball with an ellipsoid (for quadratic constraints) or a general convex body (for convex constraints).
- *Sequential MLOC*: Currently the MLOC framework applies to one-shot decision problems. It would be interesting to extend it to sequential decision problems, perhaps where multiple decisions are made in a sequence of decision epochs, and training data arrive incrementally. In this case, the baseline technique analogous to the “sequential process” would be a Markov decision process (MDP). The MLOC framework would then assist in understanding the reasonable range of costs for various sequential decision policies. Note that in the current setting, there is no opportunity for exploration to improve our failure estimates. On the other hand, in a sequential MLOC setting, there can be an opportunity to get better failure estimates by collecting information. In such a case, one can take into account the “value of new information” in decision making. Since we do not have a mechanism to collect more information (and update $\{\tilde{x}_i\}_{i=1}^M$ and hence, the corresponding failure estimate), we consider only the optimistic and pessimistic decision making approaches in this paper.

8 Conclusion

In this work, we evaluated the MLOC framework in the context of a real application and demonstrated improvements over current standards. In particular, we presented an application in the domain of transportation routing called the ML&TRP. Our framework takes advantage of uncertainty in statistical modeling to explore the decision space and find potentially more practical solutions. We provide experiments quantifying the improvements and the scalability of the framework with respect to routing problem size. We provided a generalization bound for the ML&TRP (and for the general MLOC framework) indicating that a prior belief in the operational cost can potentially be beneficial to prediction ability in general.

Acknowledgments Funding for Theja Tulabandhula was provided by a Fulbright Fellowship and Xerox Fellowship. Cynthia Rudin’s work on this project was funded in part by Con Edison, by the MIT Energy Initiative Seed Fund, and NSF Grant IIS-1053407.

References

- Agarwal, S. (2006). Ranking on graph data. In *Proceedings of the 23rd International Conference on Machine Learning*.
- Anthony, M., & Bartlett, P. L. (1999). *Neural network learning: Theoretical foundations*. Cambridge, MA: Cambridge University Press.
- Archer, A., & Blasiak, A. (2010). Improved approximation algorithms for the minimum latency problem via prize-collecting strolls. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms* (pp. 429–447).
- Archer, A., Levin, A., & Williamson, D. P. (2008). A faster, better approximation algorithm for the minimum latency problem. *SIAM Journal of Computing*, 37(5), 1472–1498.

- Arora, S., & Karakostas, G. (2006). A $2 + \epsilon$ approximation algorithm for the k -MST problem. *Mathematical Programming*, 107(3), 491–504.
- Barbera, F., Schneider, H., & Kelle, P. (1996). A condition based maintenance model with exponential failures and fixed inspection intervals. *The Journal of the Operational Research Society*, 47(8), 1037–1045.
- Belkin, M., Niyogi, P., & Sindhwani, V. (2006). Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7, 2399–2434.
- Blum, A., Chalasani, P., Coppersmith, D., Pulleyblank, B., Raghavan, P., & Sudan, M. (1994). On the minimum latency problem. ArXiv Mathematics e-prints.
- Bonami, P., Biegler, L. T., Conn, A. R., Cornuéjols, G., Grossmann, I. E., Laird, C. D., et al. (2008). An algorithmic framework for convex mixed integer nonlinear programs. *Discrete Optimization*, 5(2), 186–204.
- Chapelle, O., Schölkopf, B., & Zien, A. (Eds.). (2006). *Semi-supervised learning*. Cambridge, MA: MIT Press.
- Csiszár, I., & Tusnády, G. (1984). Information geometry and alternating minimization procedures. *Statistics and Decisions*, 1(Suppl.), 205–237.
- Cucker, F., & Smale, S. (2002). On the mathematical foundations of learning. *Bulletin-American Mathematical Society*, 39(1), 1–50.
- Ertekin, Ş., Rudin, C., & McCormick, T. (2013). Predicting power failures with reactive point processes. In *Proceedings of AAAI Late Breaking Track*.
- Fischetti, M., Laporte, G., & Martello, S. (1993). The delivery man problem and cumulative matroids. *Operations Research*, 41, 1055–1064.
- Goemans, M., & Kleinberg, J. (1998). An improved approximation ratio for the minimum latency problem. *Mathematical Programming*, 82, 111–124.
- Heng, A., Tan, A. C. C., Mathew, J., Montgomery, N., Banjevic, D., & Jardine, A. K. S. (2009). Intelligent condition-based prediction of machinery reliability. *Mechanical Systems and Signal Processing*, 23(5), 1600–1614.
- Huyer, W., & Neumaier, A. (1999). Global optimization by multilevel coordinate search. *Journal of Global Optimization*, 14, 331–355.
- Kolmogorov, A., & Tikhomirov, V. (1959). ε -Entropy and ε -capacity of sets in function spaces. *Uspekhi Matematicheskikh Nauk*, 14(2), 3–86.
- Lechmann, M. (2009). *The traveling repairman problem—An overview* (pp. 1–79). Diplomarbeit: Universität Wein.
- Li, S. (2011). Concise Formulas for the Area and Volume of a Hyperspherical Cap. *Asian Journal of Mathematics & Statistics*, 4(1), 66–70.
- Lorentz, G. G. (1966). Metric entropy and approximation. *Bulletin-American Mathematical Society*, 72, 903–937.
- Marseguerra, M., Zio, E., & Podofilini, L. (2002). Condition-based maintenance optimization by means of genetic algorithms and Monte Carlo simulation. *Reliability Engineering & System Safety*, 77(2), 151–165.
- Méndez-Díaz, I., Zabala, P., & Lucena, A. (2008). A new formulation for the traveling deliveryman problem. *Discrete Applied Mathematics*, 156(17), 3223–3237.
- Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal*, 7(4), 308–313.
- Office of Electric Transmission United States Department of Energy and Distribution. (2003). Grid 2030: A national vision for electricity's second 100 years. Technical report, United States.
- Pisier, G. (1989). *The volume of convex bodies and Banach space geometry* (Vol. 94). Cambridge, MA: Cambridge University Press.
- Pollard, D. (1984). *Convergence of stochastic processes*. Berlin: Springer.
- Rios, L. M. (2009). Algorithms for derivative-free optimization. PhD thesis, University of Illinois at Urbana-Champaign, pp. 1–133.
- Rudin, C., Anderson, R. N., Boulanger, A., Salieb-Aouissi, A., Chow, M., Dutta, H., et al. (2012). Machine learning for the New York City power grid. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(2), 328–345.
- Rudin, C., Ertekin, Ş., Passonneau, R., Radeva, A., Tomar, A., Xie, B., et al. (2014). Analytics for power grid distribution reliability in New York City. *Interfaces*. doi:10.1287/inte.2014.0748.
- Rudin, C., Passonneau, R., Radeva, A., Dutta, H., Jerome, S., & Isaac, D. (2010). A process for predicting manhole events in Manhattan. *Machine Learning*, 80, 1–31.
- Rudin, C., Passonneau, R., Radeva, A., Lerome, S., & Isaac, D. (2011). 21st-century data miners meet 19th-century electrical cables. *IEEE Computer*, 44(6), 103–105.
- Tulabandhula, T., & Rudin, C. (2013). Machine learning with operational costs. *Journal of Machine Learning Research*, 14, 1989–2028.

- Tulabandhula, T., Rudin, C., & Jaillet, P. (2011). The machine learning and traveling repairman problem. In *Proceedings of the Second International Conference on Algorithmic Decision Theory*.
- Urbina, I. (2004). Mandatory safety rules are proposed for electric utilities. *New York Times*, Aug 21, Late Edition, Sec B, Col 3, Metropolitan Desk, p. 2.
- van Eijl, C. A. (1995). A polyhedral approach to the delivery man problem. Technical report, Memorandum COSOR 95–19, Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands.
- Weintraub, A., Aboud, J., Fernandez, C., Laporte, G., & Ramirez, E. (1999). An emergency vehicle dispatching system for an electric utility in Chile. *Journal of the Operational Research Society*, 50(7), 690–696.
- Zhang, T. (2002). Covering number bounds of certain regularized linear function classes. *Journal of Machine Learning Research*, 2, 527–550.
- Zhou, D., Weston, J., Gretton, A., Bousquet, O., & Schölkopf, B. (2004). Ranking on data manifolds. In S. Thrun, L. K. Saul, & B. Schölkopf (Eds.), *Advances in neural information processing systems* (Vol. 16, pp. 169–176). Cambridge, MA: MIT Press.