

Evaluating Predictive Models in the Orienteering Problem with Stochastic Profits: A Simulation Study

Fahrettin ÇAKIR  ^a

^a Istanbul Sabahattin Zaim University, Department of Business Administration, Istanbul, Turkey. fahrettin.cakir@izu.edu.tr

ARTICLE INFO	ABSTRACT
Keywords: Akaike information criterion Cross validation, model misspecification Prescriptive analytics Stochastic orienteering Variable neighborhood search Received 18 September 2020 Revised 21 April 2021 Accepted 1 June 2021 Article Classification: Research Article	Purpose The aim of this paper is to test the effectiveness of statistical model selection measures in terms of decision quality for the orienteering problem with stochastic profits using simulation. Design/methodology/approach – This paper is based on a quantitative numerical approach where various model selection measures are evaluated using computational experiments including model-based computer-generated random data. Findings – The findings of this paper include experimental results showing a deficiency of about 6.5 units of classical selection measures relative to a decision-based selection measure for the Tsiligrides orienteering benchmark instances. Discussion – While classical model selection measures are suitable for accuracy reasons, misspecified models sometimes do lead to better decision outcomes. From a practical perspective, in order to carry out prescriptive analytics for orienteering problems, having access to a reasonable decision algorithm at the prediction stage of data-analysis can be beneficial for downstream realized profit.

1. Introduction

An emerging trend in business analytics focuses on data-driven decision making. This trend is sometimes coined as prescriptive analytics, typically associated with downstream decision making activities that use upstream business analytics output based on data analysis. Researchers have begun to study traditional operations research problems in this context in order to improve and validate optimization solutions with data. In this paper, we focus specifically on orienteering problems where an agent traverses a network collecting rewards as a result of business activity.

From a data-driven perspective, a prediction model may be embedded in a decision procedure, typically used to evaluate a candidate solution and compare it against an incumbent solution. Most of the time, researchers use simulation when the objective is not analytically computable. Importantly, the theoretical data generating process for the simulation comes from either an assumption or some output of data analysis. As a consequence, model selection carried out at this pre-decision stage carries over ramifications for the decision making process.

A natural question that may come up is how to tailor model selection at the predictive analytics stage for downstream decision making. Classical statistical approaches are geared towards model accuracy and parsimony such as the Akaike information criterion, cross-validation error, etc. In the orienteering context, choosing the correct (statistical) model for predicting reward collection and its influence on decision quality need to be studied. The author is not aware of such an analysis in the orienteering literature.

In this study, our research objective is to test the effectiveness of various predictive model selection measures for the orienteering problem with stochastic profits using simulation and using benchmark orienteering instances from the literature. Model evaluation and selection in practice is a complex process and many different approaches exist. To address this complexity, evaluation of various measures by simulation is a reasonable approach to compare their decision quality when models are selected according to different criteria.

Suggested Citation

Çakır, F. (2021). Evaluating Predictive Models in the Orienteering Problem with Stochastic Profits: A Simulation Study, *Journal of Business Research-Turk*, 13 (2), 1067-1081.

Main contribution of this work is that we confirm the effectiveness of tailoring model selection towards decision quality for benchmark orienteering instances. We present numerical evidence showing that decision quality and predictive accuracy are not necessarily aligned in the special case of orienteering with stochastic profits.

In the following, we formulate the orienteering problem and discuss the corresponding solution approach in Section 3. We present the mathematical formulation of the prescriptive workflow in Section 4. Then, we present experimental results that evaluate model selection measures in Section 5. We make our final remarks in Section 6.

2. Literature Review

There are a wide spectrum of model evaluation techniques in traditional predictive analytics. Here, we use the term predictive analytics to encompass statistical and data mining/machine learning methods and procedures. Besides goodness-of-fit measures such as Akaike Information Criteria (AIC) (Akaike, 1998), Bayesian Information Criteria (Schwarz et al., 1978), other evaluation measures/procedures exist such as Minimum Description Length (Rissanen, 1978) and cross-validation (Geisser, 1975). We refer the reader to previous reviews in this area (Arlot et al., 2010; Claeskens et al., 2008). However, as these are traditional model evaluation methods they are garnered towards minimizing error where error is measured in various ways. See aforementioned citations for details.

Unlike predictive analytics with a well-established process and a set of objectives for data analysis (CRISPDM, accuracy, ROC, precision, recall, out-of-sample error, etc.), model selection for prescriptive analytics is not a straightforward task. There are numerous approaches proposed used in different applications. Besbes et al. (2010) design an hypothesis test for revenue management to pick models based on decision quality. Ramamurthy et al. (2012) and Liyanage and Shanthikumar (2005) take a data-driven approach to inventory management by using past-data to integrate parameter estimation and profit maximization. Hence, it does model selection (viewed as parameter estimation) jointly with the decision problem. Bastani and Bayati (2015) consider online decision making with big data where individual-specific covariates are learned using a LASSO estimator to decrease regret from optimization. We are not aware of any study that considers stochastic orienteering for prescriptive analytics.

A general approach has recently been proposed by den Boer and Sierag (2020). Under the assumption of the continuity of the decision space, they show that their method is consistent and delivers the true optimum decision, i.e. best model, with a large enough dataset. However, their method (decision-based model selection, DBMS) cannot clearly outperform standard model selection measures such as AIC and cross-validation (CV) in numerical experiments for a (discrete) product assortment problem and the newsvendor problem. We provide numerical evidence for the effectiveness of DBMS in the case of orienteering when misspecification is in terms of omitting relevant predictors.

A critical component of evaluating models for downstream decisions is to consider how decisions are generated. Golden et al. (1987) prove that the orienteering problem (OP) is NP-hard; thus, for large enough problems only heuristic approaches are viable. Adding to the complexity are stochastic variants to the OP. The heuristic solution approaches to date have been tabu search, genetic algorithm, variable neighborhood search (Campbell et al., 2011a). There are several algorithms that have been proposed for the stochastic orienteering problem. The closest problem to ours is the work by Ilhan et al. (2008). Their work considers an orienteering problem with stochastic profits subject to duration limits. They consider a probabilistic objective by maximizing the probability of achieving a minimum profit level. They assume normally distributed profits on each node. They present a multiobjective genetic algorithm to solve large instances of their model. There are also other studies that consider orienteering problems with stochastic features. Campbell et al. (2011b) introduce stochasticity by considering uncertainty in travel and service times. The reward function in their work depends whether or not a customer can be visited by a deadline. As a result, uncertainty in travel and service times lead to uncertainty in the total reward earned. They present computational results for their model and derive optimal solutions analytically under certain assumptions. Zhang et al. (2014b) present a problem that applies to the pharmaceutical industry. They use an a priori solution approach coupled with recourse actions and compute the expected total expected reward analytically. They use a variable neighborhood search heuristic to solve their problem relying on this analytical evaluation. The same authors focus on dynamic

orienteering on a network of queues (Zhang et al., 2014a) more recently. Evers et al. (2014) study a close problem with stochastic weights on the edges between nodes on a graph using a two-stage stochastic model with recourse. In general, as long as an algorithm has the means to evaluate a candidate solution, different metaheuristic techniques may be used for our problem. We compare various predictive model selection methods by using a variable neighborhood search technique as the algorithm that makes decisions. This algorithm has proven to be a strong metaheuristic for difficult problems (Campbell et al., 2011a).

3. Decision Problem and Heuristic

We use a modified version of the orienteering problem as presented in Ilhan et al. (2008) and assume we are interested in maximizing expected profits instead. We discuss the formulation below.

Let $G = (V, E)$, where V is the set of nodes and E is the set of edges. Assume Node 1 designates the start node. Assume Node N designates the end node. The challenge is to design a route over some or all of the nodes in G collecting a reward as they are visited. There is a random variable associated with every node i representing the stochastic reward (or score) $R_i(\omega)$ that is earned if node i is visited and outcome ω is realized. Each edge that connects node i to node j has a nonrandom travel time given by t_{ij} . We assume there is a deadline T_{max} before which the agent needs to arrive at the end node, N .

$$\max_{x_{ij}, y_i} \quad \mathbb{E} \left[\sum_{i=2}^{N-1} \sum_{j=2}^N R_i(\omega) x_{ij} \right] \quad (1a)$$

$$\sum_{i \in V: (i,j) \in E} x_{ij} = y_j, \forall j \in V \quad (1b)$$

$$\sum_{j \in V: (i,j) \in E} x_{ij} = y_i, \forall i \in V \quad (1c)$$

$$\sum_{(i,j) \in E} t_{ij} x_{ij} \leq T_{max} \quad (1d)$$

$$y_1 = 1, y_N = 1 \quad (1e)$$

$$x_{ij}, y_i \in \{0,1\}, \forall (i,j) \in E, i \in V \quad (1f)$$

The objective function (1a) maximizes expected profit given a particular selection of nodes that meet the criteria expressed in the constraints. Constraints (1b) and (1c) are the assignment and linkage constraints. Constraint (1d) represents the deadline for the agent. Constraint (1e) forces start and end nodes to be on a tour. Constraints (1f) are integrality constraints.

Variable neighborhood search is an effective heuristic that can address the above formulation. As the formulation includes a time constraint, we apply a penalty-based VNS shown to work well for different variants of the stochastic orienteering problem. Our purpose in this paper is to focus on model selection and therefore we refer the reader to Campbell et al. (2011a) for further detail on the specifics of the algorithm. We interpret that the decision maker's decision process consists of such an algorithm; we treat it as a given.

4. Prescriptive Analytics Workflow

In practice, the decision maker does not know the true probability distribution over the stochastic outcome, ω . The decision maker arrives at a decision as a result of a workflow. First, a predictive model is built, then it is estimated, and then the estimated models are embedded in the decision algorithm. For the following setup, we mainly follow den Boer and Sierag (2020) and specialize it for our problem.

We assume the decision maker (or agents in the business workflow responsible for prediction) develop prediction models for reward at each node in the graph, $G = (V, E)$. Let the number of nodes in G be $|V| = N$. Assume there is a reward function, $r : A \times \Omega \rightarrow \mathbb{R}$, that takes as input a feasible action $a \in A$, and a stochastic

outcome, $\omega \in \Omega$; and it returns a scalar, $r \in \mathbb{R}$. Associated with every node $i \in V$, there is a random variable R_i whose range is a suitable subset of the real numbers; thus, considering all the nodes in G , there is a collection of random variables represented in column-vector form as $\mathbf{R} = (R_1(\omega), R_2(\omega), \dots, R_N(\omega))$ (thus, $\mathbf{R}^T$ is a row-vector representation).

A decision maker tries to maximize the expected reward, $\mathbb{E}_{p_\Omega}[\mathbf{r}(a, \omega)]$ from a set of feasible actions, A ,

$$\mathbf{r}(a, \omega) = \mathbf{R}^T \cdot \mathbf{y}_a \quad (2a)$$

$$= (R_1(\omega), R_2(\omega), \dots, R_n(\omega))^T \cdot \mathbf{y}_a \quad (2b)$$

$$= (R_1(\omega), R_2(\omega), \dots, R_n(\omega))^T \cdot (y_1, y_2, \dots, y_n) \quad (2c)$$

where a has a one-to-one correspondence with feasible solutions satisfying (1b)-(1f). Moreover, Equation (2a) corresponds to the random variable in the objective in (1a).

$$a = (x_a, y_a) \quad (3a)$$

$$\mathbf{y}_a = (y_i)_{i \in N} \quad (3b)$$

$$\mathbf{x}_a = (x_{ij})_{i,j \in N} \quad (3c)$$

There are K hypothesized (competing) models, $M^{(0)}, M^{(1)}, \dots, M^{(K-1)}$ used to predict the reward over G , namely $\mathbf{R}$. Each model consists of a family of distributions over $\mathbb{R}^N$ uniquely identified by a list of parameters, θ .

Table 1 Summary of Orienteering Instances

Problem Class	Deadline (Tmax)	Number of Nodes
Tsiligirides 1	20	32
Tsiligirides 1	40	32
Tsiligirides 1	60	32
Tsiligirides 1	80	32

$$M^{(k)} = \{F_\theta^{(k)} \in \mathbf{F} \mid \theta \in \Theta^{(k)}\}, \quad k = 0, 1, \dots, K-1 \quad (4)$$

where $\mathbf{F}$ is the set of cumulative distribution functions on $\mathbb{R}^N$. Essentially, model $M^{(k)}$ is a particular subset of $\mathbf{F}$. Each model has corresponding estimator ϕ^k and optimization algorithm (possibly inexact, i.e. heuristic) H^k . The estimator, $\phi^k: \mathbb{R}^n \rightarrow \Theta^k$, maps data to parameter values. The optimization algorithm, $H^k: \Theta \rightarrow A$, maps parameter values to actions. As a result, a hypothetical dataset D_0 , induces the estimate $\phi^k(D_0)$, which in turn induces the action $\mathcal{H}^k(\phi^k(D_0))$. Furthermore, assume there is access to a (longitudinal) dataset $D_0 = (\mathbf{p}_1, \mathbf{r}_1, \dots, \mathbf{p}_n, \mathbf{r}_n), n \in \mathbb{N}$ where bold variables, $\mathbf{p}_k$, represent fixed predictors from $1 \times p$ dimensional covariates and $\mathbf{r}_i$ represent observations from $1 \times N$ dimensional response variables (i.e. rewards) over the nodes of the graph G .

5. Numerical Experiments

In this section, we use the orienteering test instances from Tsiligrides (1984) to assess the performance of predictive models for the decision maker, see Table 1. Furthermore, we assess the performance of different model selection measures. To do so, we simulate hypothetical datasets and fit a predictive model and subsequently use the fitted model (on this dataset) to produce a decision. Note that we focus on Tsiligrides (1984) instances with 32 nodes. As these instances are for the deterministic OP, reward associated with every node is fixed and known. However, in our problem the reward at each node is unknown to the decision maker; thus, each dataset we generate includes a realization of a random reward for each node. These realizations are simulated according to a general linear model, which we discuss in Section 5.1.

We assume the decision maker may follow two strategies to pick a predictive model when given a dataset. First, the decision maker uses a fixed predictive model among the competing models in Table 2 with an apriori belief in that model. As a second strategy, the decision maker applies a model selection procedure in which the model with the best measure is chosen. We analyze the numerical results of the first strategy in Section 5.2. We analyze the results of the second strategy in Section 5.3.

5.1 Dataset Generation and Prediction

The decision maker has access to a dataset of a given size to infer model parameters. The realizations are assumed to come from a normal model. While Ilhan et al. (2008) assume an ANOVA model with reward as the dependent variable, unlike their study, we use the (x,y) coordinates of a node as covariates in a normal linear model with reward as the dependent variable. More specifically, x (continuous) represents the x-coordinate of a node in a Tsiligrides instance; y (continuous) represents the y-coordinate of a node in a Tsiligrides instance. See Table 2 and equations 5 for further details.

$$\mathbf{R} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}_N(\mathbf{0}, \mathcal{R}), \quad \text{cov}(\boldsymbol{\varepsilon}) = \mathcal{R} \quad (5)$$

Table 2 Summary of Simulation Models

Model Family Name	Model ID	Covariates/Predictors	$\mathcal{R}$
Linear Model (Homoscedastic)	1	x	$\sigma^2 I_N$
Linear Model (Homoscedastic)	2	y	$\sigma^2 I_N$
Linear Model (Homoscedastic)	3	x,y	$\sigma^2 I_N$
Linear Model (Heteroscedastic)	4	x	$\sigma^2 \Lambda$ $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$
Linear Model (Heteroscedastic)	5	y	$\sigma^2 \Lambda$ $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$
Linear Model (Heteroscedastic)	6	x,y	$\sigma^2 \Lambda$ $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$

There are many ways in which a model can be misspecified. In our experiments, we focus on the effect of omitting a relevant predictor and the misspecification through incorrect variance structure (homogeneity vs heterogeneity). For this reason, we present our numerical experiments under two scenarios. In Scenario 1, the true reward generating model is Model 3 with homogeneous variance; thus, while the other models in Table 2 are misspecified in this scenario, they may still be used for prediction by the decision maker. In scenario 2, the true reward generating model is Model 6 making the other models in the same table misspecified. In

summary, each scenario is associated with a data generating model (unknown to the decision maker) but the decision maker may use any competing hypothesized model from Table 2 to predict reward.

We set the parameters of the models in Table 2 using values given in Table 3:

Table 3 Parameters for the Models in Table 2

β_0	β_x	β_y	σ^2	$\sigma_i^2 = \sigma^2 \lambda_i$
0.1	0.2	0.2	25	Nodes 1 to 16 uniformly from [16,25]
				Nodes 17 to 32 uniformly from [225,400]

where the variances are chosen identically as in Ilhan et al. (2008).

Having chosen a simulation model, we generate 1,000 (balanced) datasets of various sizes $\{n = 64, n = 128, n = 192, n = 256, n = 512\}$. We choose these sample sizes because a sample size of at least 170 is required for Model 3 to turn out to be significant. In other words, the F-test for Model 3 turns out to be overall significant at least at this sample size. We choose Model 3 to compute this threshold sample size. For every generated dataset, we apply the sequence of estimation and optimization, $\mathcal{H}^k(\varphi^k(D_0))$, for every competing model $M^{(k)}$ based on Table 2, where D_0^n is a dataset of size n from the true model $M^{(0)}$. We assume $M^{(0)}$ is the family that contains the true data generating model.

For every dataset (D_0^n) , we determine the K decisions induced by every competing model k , $\mathcal{H}^k(\varphi^k(D_0^n))$, $k = 0, 1, \dots, K-1$. Here, H^k is the same penalty-based VNS heuristic that delivers a decision, for every k . The heuristic starts from the same initial solution for all k ; however, it may traverse the solution space differently because the predictions coming from model k and model k^0 are not necessarily equal. Consequently, the heuristic is more likely to deliver different final solutions across models $k = 0, 1, \dots, K-1$.

For our performance analysis, we evaluate the true objective from decision k using Monte Carlo approximation by using the corresponding true model. We use 10,000 replicates of data generated according to the true model to exactly evaluate a decision proposed by model k .

5.2 Performance based on Fixed Predictive Model

In this section, we assume that a single model from Table 2 is chosen apriori by the decision maker. We present the performance of each alternative under the two scenarios we discussed earlier.

In Table 4, we report average reward that is *actually* earned under each hypothesized model for scenario 1. The average reward is based on 1000 data replications. For instance, for Tsiligrirides 1 ($Tmax = 20$) instances, when the given dataset size is 64, the model in C3 induces 1000 decisions; i.e. one decision for each of the 1000 replicate datasets. As a result, the average true reward earned from these 1000 decisions is 41.80 units. The decision heuristic (penalty-based VNS) does consistently better, in terms of average reward, when it uses predictions coming from Model 3 across all four Tsiligrirides instances and across all dataset sizes. Notice that Model 3 is the true model from which the datasets are generated from. Moreover, misspecification in terms of variance reduces average reward minimally which can be seen by comparing column 7 and column 8. However, the reduction in average reward due to misspecification in the mean is more nuanced. When the omitted covariate is the x coordinate, the average reward is reduced by 0.7 units to 3.1 units across all problem instances and all dataset sizes (compare column 8 with columns 4 and 6). When the omitted variable is the y coordinate, the average reward is reduced by 1.6 units to 11.1 units across the different Tsiligrirides instances. In addition to the comparison of average reward, the empirical distribution of earned reward shows near stochastic dominance across different models. For instance, Figure 1 shows the empirical cumulative distribution of the reward earned for each Tsiligrirides instance under each sample size colored by each hypothesized model. We see that Model 3 and Model 6 shows *nearly* first-order stochastic dominance to other models.

In Table 5, we report average reward that is *actually* earned under each hypothesized model for scenario 2. The average reward is based on 1000 data replications. The decision heuristic (penalty-based VNS) does consistently better, in terms of average reward, when it uses predictions coming from Model 6 across all four Tsiligrirides instances and across all dataset sizes. Notice that Model 6 is the true model. Moreover, misspecification in terms of variance does not change the average significantly (compare columns 8 and 5). The reduction in average reward due to misspecification in the mean is nearly identical to the results in Table 4.

Table 4 Average reward (True Model 3)

C 1	C 2	C 3	C 4	C 5	C 6	C 7	C 8
		Hypothesized Models					
		Misspecified					True Model
Problem	n	$\beta_0 + \beta_1 x$ (homo var)	$\beta_0 + \beta_1 y$ (homo var)	$\beta_0 + \beta_1 x$ (hetero var)	$\beta_0 + \beta_1 y$ (hetero var)	$\beta_0 + \beta_1 x + \beta_2 y$ (hetero var)	$\beta_0 + \beta_1 x + \beta_2 y$ (homo var)
Tsiligrirides 1 (Tmax = 20)	64	41.80	42.17	41.92	42.07	42.75	43.42
	128	42.34	42.42	42.40	42.53	43.58	43.95
	192	42.43	42.59	42.36	42.40	43.58	44.07
	256	42.44	42.67	42.52	42.54	43.61	44.08
	512	42.69	42.62	42.71	42.37	43.93	44.29
Tsiligrirides 1 (Tmax = 40)	64	95.84	99.31	95.52	98.57	97.95	100.01
	128	95.92	99.35	95.81	99.48	99.90	101.27
	192	96.78	100.01	96.58	100.00	100.87	101.43
	256	96.84	100.03	96.72	100.26	101.21	102.14
	512	97.18	99.95	96.66	99.63	101.57	101.89
Tsiligrirides 1 (Tmax = 60)	64	123.97	132.09	123.59	132.06	131.45	133.24
	128	124.83	133.11	124.98	132.42	133.24	135.45
	192	124.76	133.58	125.74	133.25	134.31	135.85
	256	125.55	133.56	125.69	133.75	134.87	135.69
	512	126.77	133.50	125.82	133.69	135.73	136.62
Tsiligrirides 1 (Tmax = 80)	64	146.74	152.36	146.48	152.36	151.17	152.49
	128	147.65	153.85	147.83	153.01	152.99	154.08
	192	146.98	153.41	147.89	153.34	153.76	154.35
	256	148.18	153.49	148.56	153.50	153.85	154.59
	512	148.48	153.52	148.87	154.00	154.55	154.39

Table 5 Average Reward (True Model 6)

C 1	C 2	C 3	C 4	C 5	C 6	C 7	C 8
		Hypothesized Models					
		Misspecified					True Model
Problem	n	$\beta_0 + \beta_1x$ (homo var)	$\beta_0 + \beta_1y$ (homo var)	$\beta_0 + \beta_1x + \beta_2y$ (homo var)	$\beta_0 + \beta_1x$ (hetero var)	$\beta_0 + \beta_1y$ (hetero var)	$\beta_0 + \beta_1x + \beta_2y$ (hetero var)
Tsiligirides 1 ($Tmax = 20$)	64	42.72	42.50	44.16	42.80	42.55	44.36
	128	42.71	42.32	44.36	42.78	42.33	44.46
	192	42.83	42.29	44.37	42.75	42.50	44.66
	256	43.00	42.38	44.43	42.43	42.31	44.46
	512	42.46	42.41	44.31	42.48	42.29	44.38
Tsiligirides 1 ($Tmax = 40$)	64	97.04	100.14	102.01	97.52	100.09	102.33
	128	97.48	99.43	101.85	96.64	99.70	101.94
	192	97.22	99.48	102.00	96.79	99.67	102.17
	256	97.51	99.79	101.81	96.77	99.60	102.09
	512	97.32	99.98	101.96	96.98	99.72	102.05
Tsiligirides 1 ($Tmax = 60$)	64	125.48	133.59	136.15	126.26	133.11	136.25
	128	126.45	133.37	136.46	126.17	133.54	136.62
	192	126.26	133.49	136.72	126.43	133.97	136.70
	256	126.59	133.46	136.78	126.46	133.79	136.85
	512	126.47	133.24	136.68	126.09	133.65	136.59
Tsiligirides 1 ($Tmax = 80$)	64	148.22	153.11	154.77	148.71	153.79	154.60
	128	149.21	153.40	154.81	148.42	153.32	154.78
	192	148.78	153.29	154.77	148.44	153.90	154.69
	256	149.14	153.38	154.60	149.12	153.50	154.73
	512	148.84	153.71	154.58	148.51	153.41	154.62

5.3 Performance based on Model Selection Measures

When given a dataset, (D_0^n), the decision maker may use several measures to pick a predictive model out of the competing ones. We present the prescriptive performance of score-based model selection measures such as AIC and cross-validation. Moreover, we consider the decision-based model selection (DBMS) measure proposed by den Boer and Sierag (2020). Additionally, we compute the pointwise ranking loss based on the mean rank and sum rank, PWRM and PWRS respectively (Chen et al., 2009).

In Table 6, we report average reward earned from each measure categorized by dataset size and Tsiligirides instance under scenario 1. For instance, for the Tsiligirides 1 ($Tmax = 20$) instances and dataset size 64, there are 1000 replicate datasets. For each dataset, there are a total of 6 decisions induced by each respective model from Table 2. If a model is chosen according to the AIC criterion with every dataset, the average true reward earned as a consequence of using this model selection procedure is equal to 42.85. While it is better to use predictions from the true model from an apriori perspective, the DBMS measure performs better than the other selection measures. If you consider the Tsiligirides instances with $Tmax = 20$, the DBMS selection measure leads to higher average reward by at least 2.35 units ($n = 512$, in comparison to CV). At the same time, when $Tmax = 80$, DBMS improves decisions on average by at least 6.52 units ($n = 512$, in comparison to CV). The DBMS measure occasionally chooses misspecified models when they induce a better decision. Although PWRM and PWRS behave similarly, in that they too tend to choose misspecified models relatively more frequently, the decision quality is much worse than DBMS. In fact, they tend to do poorly against measures AIC and CV, which favor predictive accuracy in terms of squared loss. See Figure 4 for the frequency of chosen models across the measures. If we observe the empirical distribution of reward earned across the 1000 datasets for each instance and dataset size, we see that DBMS nearly stochastically dominates the other measures (see Figure 2 and Figure 3). As for scenario 2, when the true model is Model 6, we observe similar results uniformly across different instances and dataset sizes (Table 7).

Table 6: Average True Reward per Model Selection Measure (True Model 3)

Problem	n	AIC	CV	DBMS	PWRM	PWRS	True Model
Tsiligirides 1 (Tmax = 20)	64	42.85	43.38	45.26	42.82	42.82	43.42
	128	43.50	43.97	45.86	43.16	43.16	43.95
	192	43.78	44.05	46.12	43.22	43.22	44.07
	256	43.99	44.07	46.27	43.45	43.45	44.08
	512	44.26	44.29	46.64	43.35	43.35	44.29
Tsiligirides 1 (Tmax = 40)	64	99.93	100.15	104.45	99.05	99.05	100.01
	128	100.30	101.24	105.07	99.34	99.34	101.27
	192	101.06	101.43	105.59	99.80	99.80	101.43
	256	101.71	102.16	105.70	99.87	99.87	102.14
	512	101.76	101.91	105.97	99.89	99.89	101.89
Tsiligirides 1 (Tmax = 60)	64	131.36	133.29	139.19	130.36	130.36	133.24
	128	134.38	135.40	140.49	132.01	132.01	135.45
	192	135.04	135.86	140.84	131.38	131.38	135.85
	256	135.35	135.71	141.16	131.68	131.68	135.69
	512	136.51	136.62	141.46	132.53	132.53	136.62
Tsiligirides 1 (Tmax = 80)	64	151.80	152.74	159.27	150.90	150.90	152.49
	128	153.95	154.07	160.53	152.48	152.48	154.08
	192	153.95	154.28	160.59	151.72	151.72	154.35
	256	154.05	154.61	160.79	152.16	152.16	154.59
	512	154.36	154.40	160.92	152.12	152.12	154.39

Table 7: Average True Reward per Model Selection Measure (True Model 6)

Problem	n	AIC	CV	DBMS	PWRM	PWRS	True Model
Tsiligirides 1 (Tmax = 20)	64	44.11	44.15	46.56	44.10	44.10	44.36
	128	44.44	44.37	46.79	44.23	44.23	44.46
	192	44.66	44.35	46.78	44.57	44.57	44.66
	256	44.46	44.40	46.67	44.46	44.46	44.46
	512	44.38	44.31	46.69	44.23	44.23	44.38
Tsiligirides 1 (Tmax = 40)	64	102.15	101.98	106.30	102.15	102.15	102.33
	128	101.93	101.78	106.17	101.64	101.64	101.94
	192	102.17	101.94	106.29	102.28	102.28	102.17
	256	102.08	101.78	106.22	101.52	101.52	102.09
	512	102.05	101.95	106.38	101.95	101.95	102.05
Tsiligirides 1 (Tmax = 60)	64	135.70	136.20	141.53	136.06	136.06	136.25
	128	136.60	136.43	141.82	136.28	136.28	136.62
	192	136.69	136.71	141.92	136.72	136.72	136.70
	256	136.86	136.77	141.90	137.04	137.04	136.85
	512	136.58	136.70	141.73	136.70	136.70	136.59
Tsiligirides 1 (Tmax = 80)	64	154.52	154.67	160.96	154.49	154.49	154.60
	128	154.76	154.80	161.00	154.60	154.60	154.78
	192	154.69	154.73	161.03	154.50	154.50	154.69
	256	154.72	154.60	161.07	154.69	154.69	154.73
	512	154.64	154.59	160.77	154.53	154.53	154.62

6. Discussion

Our numerical results are based on models whose predictors x and y have an equal impact on the mean reward (both population coefficients are 0.2). However, the reduction in decision quality created by omitting either predictor from the model is not the same. For the Tsiligirides 1 instances, omitting the x predictor leads to higher loss in decision quality. This can be explained by the fact that at the decision stage other problem parameters such as node locations, pairwise distances between nodes, and the time constraint (T_{max}) is not

taken into account while building predictive models. There are mixed results obtained in den Boer and Sierag (2020) in which the DBMS measure is shown to only occasionally perform better than AIC and CV depending on dataset size in the product assortment and newsvendor problems. However, DBMS performs uniformly better in the orienteering setting of this paper. This effect is seen across dataset sizes and across different values to the time constraint budget, T_{max} . One possible reason why DBMS may lead to higher degree of improvement is because as T_{max} increases, there is more room to search the decision space (i.e. feasible routes), and thus; the heuristic is more likely to find higher quality routes when taken into account, as the DBMS does.

From a practical perspective, evaluating models through the application of a heuristic can be computationally prohibitive if the problem dimensions are large enough; for instance, when there are a large number of customers in a network. To address such a scenario, a decision-based model selection measure that does not require access to a heuristic can be more useful. The challenge in this case, is to design a measure that solely uses raw data. The VNS heuristic explores the decision space by comparing an incumbent solution to candidate solutions. This pairwise comparison is guided by the predictions of the fitted models. While the point estimates of the models are relevant, the VNS heuristic really behaves uniquely according to the pairwise comparison of the routing solutions. This is more of a ranking problem. According to our results in Figure 4, ranking measures choose models in a more diversified manner similar to DBMS, which seems like a necessary characteristic for higher decision quality, but not sufficient.

Our numerical experiments have assumed normality which may be limiting in real life cases when rewards may in fact follow some other distribution. The normality assumption may lead to some negative realized profits in some situations. "In the context of auditing, this means the actual inventory levels are higher than the levels reported by the supplier" (Ilhan et al., 2008). Other reward distributions to consider may be of lognormal (e.g. fishing) and multinomial (e.g. product assortment) type. In such a case, it is not clear if the DBMS will be uniformly better against other measures for the orienteering problem.

7. Conclusion

Although model accuracy seems natural for decision purposes, misspecified models are sometimes better. For this reason, in a business context, historical data should not be evaluated only according to accuracy and similar measures typical of upstream business intelligence objectives. Focusing on downstream objectives while analyzing past data has a high potential value for business. This potential can be utilized by integrating the entire business intelligence workflow (descriptive, predictive, prescriptive stages) by educating agents with different roles about the implications of their own work outcome on downstream workflows. A standardized communication on the boundary between different stages of business intelligence can make experimentation and guidance easier in order to increase performance on decision quality. This can be accomplished with use of information systems that improve and quicken the information flow between these different stages.

In this paper, we compared the effect of misspecifying linear models on routing quality. We have shown that selecting predictive models based on decision quality can lead to higher quality decisions in orienteering with random reward. There is further research that lies ahead. One possible extension to the experiments may be to investigate the relationship between when a predictor is omitted and the different locations of start and end nodes in the network. How might this affect decision quality? In addition, how might this relationship be when start and end nodes are the same (e.g. single depot)? It may be valuable to study how a better diversification, akin to DBMS, can be achieved with ranking-like measures without access to a downstream heuristic. Such a study would be informative and useful for business contexts that require decisions be made quickly. For example, taxi fleet management, autopilot vehicles analyze a continuous stream of data and utilize this data for prediction in real-time. We plan to expand on such considerations in future research.

A Appendix A






Figure 1: Empirical cdf by predictive model.



Cumulative Reward Dist. by Model Selection Measure

Figure 2: Empirical cdf of Reward by Model Selection Measure (True Model 3).



Cumulative Reward Dist. by Model Selection Measure

Figure 3: Empirical cdf of Reward by Model Selection Measure (True Model 6).




Figure 4: Chosen Models by Each Measure (True Model 3).

References

- Akaike, H. (1998). Information theory and an extension of the maximum likelihood principle. In *Selected papers of hirotugu akaike*, pages 199–213. Springer.
- Arlot, S., Celisse, A. (2010). A survey of cross-validation procedures for model selection. *Statistics surveys*, 4:40–79.
- Bastani, H. and Bayati, M. (2015). Online decision-making with high-dimensional covariates. *Forthcoming in Operations Research*.
- Besbes, O., Phillips, R., and Zeevi, A. (2010). Testing the validity of a demand model: An operations perspective. *Manufacturing & Service Operations Management*, 12(1):162–183.
- Campbell, A. M., Gendreau, M., and Thomas, B. W. (2011a). The orienteering problem with stochastic travel and service times. *Annals of Operations Research*, 186(1):61–81.
- Campbell, A. M., Gendreau, M., and Thomas, B. W. (2011b). The orienteering problem with stochastic travel and service times. *Annals of Operations Research*, 186(1):61–81.
- Chen, W., Liu, T.-Y., Lan, Y., Ma, Z.-M., and Li, H. (2009). Ranking measures and loss functions in learning to rank. *Advances in Neural Information Processing Systems*, 22:315–323.
- Claeskens, G., Hjort, N. L. (2008). Model selection and model averaging. *Cambridge Books*.
- den Boer, A. V. and Sierag, D. D. (2020). Decision-based model selection. *European Journal of Operational Research*. 00006.
- Evers, L., Glorie, K., van der Ster, S., Barros, A. I., and Monsuur, H. (2014). A two-stage approach to the orienteering problem with stochastic weights. *Computers & Operations Research*, 43:248–260.
- Geisser, S. (1975). The predictive sample reuse method with applications. *Journal of the American statistical Association*, 70(350):320–328.
- Golden, B. L., Levy, L., and Vohra, R. (1987). The orienteering problem. *Naval Research Logistics (NRL)*, 34(3):307–318.
- Ilhan, T., Iravani, S. M. R., and Daskin, M. S. (2008). The orienteering problem with stochastic profits. *IIE Transactions*, 40(4):406–421. 00018.
- Liyanage, L. H. and Shanthikumar, J. G. (2005). A practical inventory control policy using operational statistics. *Operations Research Letters*, 33(4):341–348.
- Ramamurthy, V., George Shanthikumar, J., and Shen, Z.-J. M. (2012). Inventory policy with parametric demand: Operational statistics, linear correction, and regression. *Production and Operations Management*, 21(2):291–308.
- Rissanen, J. (1978). Modeling by shortest data description. *Automatica*, 14(5):465–471.
- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464.
- Tsiligrirides, T. (1984). Heuristic methods applied to orienteering. *Journal of the Operational Research Society*, 35(9):797–809.
- Zhang, S., Ohlmann, J., and Thomas, B. (2014a). Dynamic Orienteering on a Network of Queues. *Tippie College of Business Publications*. 00000.
- Zhang, S., Ohlmann, J. W., and Thomas, B. W. (2014b). A priori orienteering with time windows and stochastic wait times at customers. *European Journal of Operational Research*, 239(1):70–79.