

JANOS: An Integrated Predictive and Prescriptive Modeling Framework

Bergman, David; Huang, Teng; Brooks, Philip; Lodi, Andrea; Raghunathan, Arvind

TR2021-025 April 04, 2021

Abstract

Business research practice is witnessing a surge in the integration of predictive modeling and prescriptive analysis. We describe a modeling framework JANOS that seamlessly integrates the two streams of analytics, allowing researchers and practitioners to embed machine learning models in an end-to-end optimization framework. JANOS allows for specifying a prescriptive model using standard optimization modeling elements such as constraints and variables. The key novelty lies in providing modeling constructs that enable the specification of commonly used predictive models within an optimization model, have the features of the predictive model as variables in the optimization model, and incorporate the output of the predictive models as part of the objective. The framework considers two sets of decision variables; regular and predicted. The relationship between the regular and the predicted variables are specified by the user as pre-trained predictive models. JANOS currently supports linear regression, logistic regression, and neural network with rectified linear activation functions. In this paper, we demonstrate the flexibility of the framework through an example on scholarship allocation in a student enrollment problem and provide a numeric performance evaluation.

INFORMS Journal on Computing

© 2021 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

JANOS : An Integrated Predictive and Prescriptive Modeling Framework

David Bergman, Teng Huang*
Philip Brooks†
Andrea Lodi‡
Arvind U. Raghunathan§

September 4, 2020

Abstract

Business research practice is witnessing a surge in the integration of predictive modeling and prescriptive analysis. We describe a modeling framework JANOS that seamlessly integrates the two streams of analytics, allowing researchers and practitioners to embed machine learning models in an end-to-end optimization framework. JANOS allows for specifying a prescriptive model using standard optimization modeling elements such as constraints and variables. The key novelty lies in providing modeling constructs that enable the specification of commonly used predictive models within an optimization model, have the features of the predictive model as variables in the optimization model, and incorporate the output of the predictive models as part of the objective. The framework considers two sets of decision variables; *regular* and *predicted*. The relationship between the regular and the predicted variables are specified by the user as pre-trained predictive models. JANOS currently supports linear regression, logistic regression, and neural network with rectified linear activation functions. In this paper, we demonstrate the flexibility of the framework through an example on scholarship allocation in a student enrollment problem and provide a numeric performance evaluation.

1 Introduction

There has been significant proliferation of research in and application of machine learning and discrete optimization. These two analytical domains have frequently been used in a single business decision-making problem but for different purposes. Machine learning techniques have typically been used to predict what is likely to happen in the future, while optimization methods have been used to strategically search through feasible solutions.

*Department of Operations and Information Management, School of Business, University of Connecticut, United States {david.bergman, teng.huang}@uconn.edu

†Optimized Operations, LLC., phil.brooks24@me.com

‡Department of Mathematical and Industrial Engineering, Polytechnique Montreal, Canada, andrea.lodi@polymtl.ca

§Mitsubishi Electric Research Laboratories, 201 Broadway, Cambridge, MA 02139. raghunathan@merl.com

However, the integration of the two analytics streams is sometimes handled sequentially without full integration (Deng et al. 2018), with the outputs of machine learning models only used as parameters in an optimization model. As an example, suppose a neural network is built to predict customer churn of a telecommunication company, with features consisting of demographic information together with service price to the customer (the lower the price, the lower the probability of customer churn). How then should the company set service price in order to maximize revenue? The features of the predictive model are variables in the decision problem and the output of the predictive model for any choice of the decision variables is a part of the objective function. This makes the optimization problem of maximizing revenue particularly challenging. There are few, if any, techniques, much less tools, available for solving such an optimization problem in the way proposed. This paper seeks to fill that void by introducing JANOS¹, a modeling framework for joint predictive-and-prescriptive analytics. Figure 1 depicts how one could use JANOS for a similar business problem of minimizing customer churn. First, the user would train a machine learning model, in this case a neural network, for predicting customer churn. Then, the trained model would be replicated any number of times to represent either each customer or each segment of customers, and an optimization model where the neural network is embedded is formulated and solved.

JANOS allows a user to specify portions of an objective function as commonly utilized predictive models, whose features are fixed or are variables in the optimization model. JANOS currently supports predictive models of the following three forms: linear regression, logistic regression, and neural networks (NNs) with rectified linear (ReLU) activation functions. These models are commonly used in practice, and the framework is easily extensible to other predictive models.

We decided to only include the output of the predictive models in the objective function because the typical setting is that a predictive model provides just an expected value for the dependent variable. Adding the outputs in the constraints is possible, but it would result in satisfying constraints only in expectation. Constrained stochastic optimization problems sometimes require chance constraints or other considerations, so in this iteration of the solver we opted for a simplistic setting.

To embed these predictive models in an optimization model, we utilized linearization techniques. For linear regression this is straightforward. For logistic regression, we employ a piece-wise linear approximation. Details are provided in Section 4.2. For NNs, we make use of recent work that formulates a NN as a *mixed integer programming* (MIP) problem (Serra et al. 2017, Bienstock et al. 2018, Fischetti and Jo 2017, Anderson et al. 2020); we do not use the MIP to train the NN, but rather utilize the network reformulation to produce outputs of the NN based on the input features. Details are provided in Section 4.1.

A key advantage of JANOS is that it automates the transcription of common predictive models so that they can be handled by mixed integer programming solvers. Thus, researchers and practitioners are relieved of this onerous task of reformulating the predictive models into tractable constraints. A different model for predictive variables can quickly be substituted out without much effort from the modeler, enabling the user to quickly compare the optimal decisions when different predictive models are used. Finally, a researcher or practitioner looking

¹A play on Janus, who according to ancient Roman mythology is the god of beginnings, gates, transitions, time, duality, doorways and is usually depicted with two faces one looking to the past (predictive) and one to the future (prescriptive). (Source: Wikipedia)

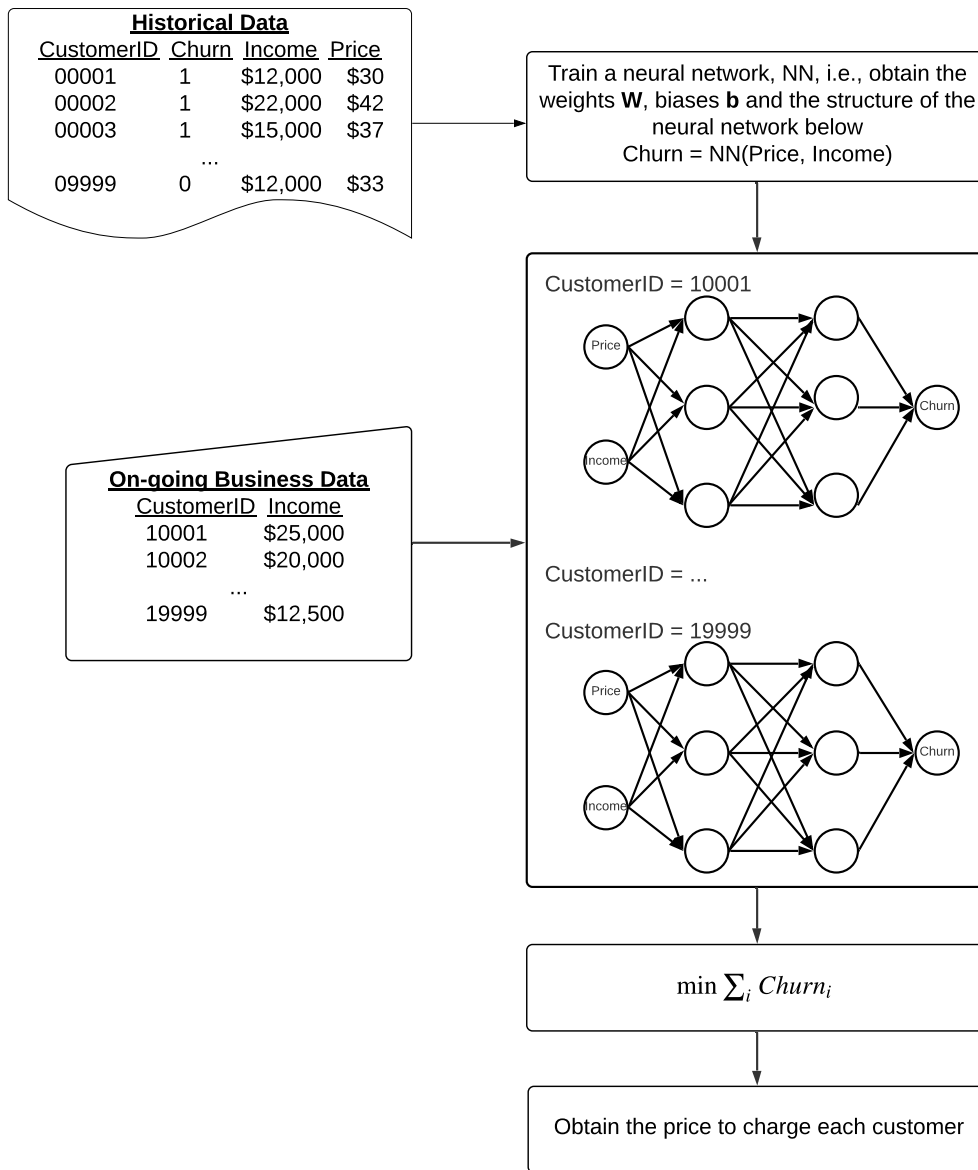


Figure 1: A schematic of how JANOS works

to apply an algorithmic approach to the type of decision-making problem discussed in this paper would not have to add all parameters of a predictive model with an optimization model; for example, a neural network with 3 layers of 10 nodes each would have hundreds of parameters, and JANOS automates that transformation. In future releases, additional predictive models will be added as well as more advanced reformulations and algorithmic implementations.

1.1 Contributions

JANOS is built in Python and calls Gurobi (Gurobi Optimization 2019) to solve MIPs. A user specifies an optimization model through standard modeling constructs that share similarities to those in other common optimization modeling languages, for example Gurobi’s Python interface, Pyomo Gurobi (Hart et al. 2011) or JuMP (Dunning et al. 2015) in Julia (Bezanson et al. 2017).

We partition the variables in the model into two sets—the *regular variables* and the *predicted variables*. The regular variables are used to model operational constraints and objective function terms as typical variables in MIP. The predicted variables are specified through predictive models wherein some of the features depend on regular variables. The predictive models are pre-trained by the user, who can load any of the three permissible predictive model forms, together with a *mapping* between the regular variables and the features. We eventually plan to integrate automated machine learning (Feurer et al. 2015) and have JANOS determine the best predictive model to associate with a given data frame. To exhibit how the framework can be used, we present as an example the allocation of scholarship offers to admitted students to optimize the enrollment.

The impact of the JANOS solver on research is twofold. First, framing a problem as optimizing over standard predictive models where the decision variables are features of the predictive models requires much attention. Suppose the predictive model is a support vector regression with a radial kernel. How does one solve the optimization problem? Each predictive model added to JANOS will require research effort to investigate linearizations or advanced optimization methodology that lead to efficient solution times.

More broadly, we envision that JANOS will be used by researchers in fields other than optimization to solve the problems they are interested in to derive the insights they desire. Currently, a researcher familiar with optimization who has a machine learning model that they would like to optimize over is unable to solve the problem of interest. Contrast that with a researcher in logistics that requires the solution of a routing problem to identify how many trucks are needed by the company. This researcher can use a standard commercial integer programming solver, solve the routing problem to identify the number of trucks that are needed by the company, and then make decisions based on that output. The logistics researcher need not know how the optimization model is solved, just that the problem of interest can be solved. We envision the same thing for the former researcher. JANOS allows the researcher to solve an optimization problem with machine learning models so that optimal decisions can be identified for real-world decision making problems that previously could not.

The rest of the paper is organized as follows. We first review the literature related to the joint reasoning in predictive and prescriptive analysis in Section 2. The general problem addressed by JANOS is provided in Section 3. The algorithmic details of how we optimize over linear regression models, logistic regression models and NNs are given in Section 4. The

student enrollment example and a collection of experiments designed to test the efficiency of JANOS are described in Section 5. We then describe how JANOS can be downloaded and installed in Section 6. We conclude in Section 7.

2 Literature Review

Existing studies on the combination of predictive and prescriptive analytics take predictions as fixed and then make choices based on fixed predictions, for example, predictions are parameters in an optimization model (Ferreira et al. 2015). Ferreira et al. (2015) first predict sales based on a chosen number of price values and then uses those fixed estimates to determine the optimal price. In their application, the predicted sales are parameters in the optimization model. This modeling approach adapts, but still lacks in full integration between the two analytics disciplines. Sales will generally depend on price, and an optimization algorithm for tackling this problem should incorporate such reasoning. Approaches using fixed-point estimates of parameters are feasible when full enumeration or partial enumeration of the collection of feasible solutions is practical. However, in instances where enumeration is not possible, aspects of the optimization algorithm need to be integrated into the predictive model. For example, Huang et al. (2019) model a real-world problem in such a way, but only propose an exact optimization approach when simple linear regression models are used for prediction.

There are additional streams of research that combine predictive modeling and optimization. First of all, machine learning algorithms are powered by optimization techniques. For example, when fitting a simple linear regression model, the *ordinary least square* method determines the unknown parameters by minimizing the sum of the squares of the difference between the observed dependent variable values and the fitted values. In machine learning, various optimization techniques are applied so that the learning process is efficient and achieves desired accuracy (Boyd et al. 2011). For example, Koh et al. (2007) propose an interior-point method for solving large-scale ℓ_1 -regularized logistic regression problems; ALAMO (Cozad et al. 2014) uses mixed integer optimization to learn algebraic models from data; and, linear programming (LP) and integer programming (IP) based methods have been proposed to assist training NNs (Bienstock et al. 2018). Additionally, there have been recent efforts on using machine learning to improve optimization algorithms. For example, Cappart et al. (2019) utilize deep reinforcement learning to improve optimization bounds, and Khalil et al. (2017) proposes a generic method for learning combinatorial optimization over graphs, with a plethora of papers arising in this area (Nazari et al. 2018, Lemos et al. 2019, Bengio et al. 2018). These papers either leverage machine learning to improve optimization, or leverage optimization to improve machine learning—in our setting, we integrate the two into a unified decision-making framework. There are multiple streams of research that consider machine learning and operations research in one decision-making pipeline. One stream of research is to estimate machine learning models while minimizing the regret of decisions (Elmachtoub and Grigas 2017, Wilder et al. 2019, Demirović et al. 2019, Mandi et al. 2019). Another is a framework for using auxiliary data to prescribe business decisions (Bertsimas and Kallus 2020). In particular, a conditional stochastic optimization problem minimizing uncertain costs is formulated, while our framework does not explicitly model the uncertainty in the predictions.

There appear many recently-published and on-going works that optimize over supervised learning models. For example, Paul et al. (2018) optimize over a non-parametric tree choice

model; Mišić (2017) and Biggs and Hariss (2018) optimize over tree ensembles; and Aouad et al. (2019) and Feldman et al. (2019) optimize over multinomial logistic models.

There exist some studies on optimizing over neural network structures. For example, in Lam et al. (2000), a dual-objective optimization algorithm is formulated to optimize two neural-network outputs. They use a heuristic algorithm for finding the optimal values of the decision variables by searching over a small enumerated collection of possible values. Schweidtmann and Mitsos (2019) present a method for a global optimization problem with a single embedded neural network. They utilize the convex and concave envelopes of the activation function they employ, the hyperbolic tangent function. In JANOS, we restrict to only *rectified linear unit* (ReLU) functions as activation functions, enabling us to formulate a network flow model to optimize over any number neural networks through a *mixed-integer linear programming* (MILP) transformation.

3 Problem Description

JANOS seeks to solve problems formulated as (**PROBLEM-ORI**):

$$\max_x \quad \sum_{j=1}^{n_1} c_j x_j + \sum_{k=1}^{n_2} d_k y_k \quad (\text{PROBLEM-ORI})$$

$$\text{s.t.} \quad \sum_{j=1}^{n_1} a_j^i x_j \leq b_i, \quad \forall i \in \{1, \dots, m\} \quad (1)$$

$$y_k = g_k(\alpha_1^k, \dots, \alpha_{p_k}^k; \theta_k), \quad \forall k \in \{1, \dots, n_2\} \quad (2)$$

$$\alpha_l^k \text{ is given}, \quad \forall l \in \{1, \dots, q_k\}, \forall k \in \{1, \dots, n_2\} \quad (3)$$

$$\alpha_l^k = e_l^k \cdot x, \quad \forall l \in \{(q_k + 1), \dots, p_k\}, \forall k \in \{1, \dots, n_2\} \quad (4)$$

$$x_j \in X_j, \quad \forall j \in \{1, \dots, n_1\}. \quad (5)$$

The variables $x = (x_1, \dots, x_{n_1})$ are *regular variables* and the variables $y = (y_1, \dots, y_{n_2})$ are *predicted variables*. Each variable x_j belongs to a finite or continuous set X_j and are constrained via linear inequalities. Each predicted variable y_k is associated with a predictive model g_k , with features $\alpha^k = (\alpha_1^k, \dots, \alpha_{p_k}^k)$. Each g_k is assumed to be a pre-trained predictive model (a linear regression, logistic regression, or neural network with a rectified linear activation function) so that the parameters θ_k are fit prior to optimization by the user. Model g_k has p_k features, $(\alpha_1^k, \dots, \alpha_{p_k}^k)$. The first $q_k, 1 \leq q_k \leq p_k$, features of predictive model g_k are fixed and given, while each of the remaining $(p_k - q_k)$ features are regular variables, linked through the function e_l^k , a n_1 -length binary unit-vector with a 1 in the coordinate of the associated regular variable. Note that a single pre-trained model can be used as multiple g_k 's.

4 Algorithmic Details

In this section, we summarize how JANOS handles linear regression, logistic regression, and NN models. If y is determined by a linear regression model, the function g is linear and the construction is straightforward. We construct (**PROBLEM-ORI**) and feed the model to

Gurobi. If y is a predicted value of a NN, y is obtained from a network flow model. The details are in Section 4.1. If y is a predicted value of a logistic regression model, we partition the range of the *log-odds*, i.e., $X\beta$ as in $y = \frac{1}{1+e^{-X\beta}}$ into a collection of intervals, and use the mean of the **sigmoid** value of $X\beta$ within the corresponding interval to approximate y in the objective function. The details are in Section 4.2.

4.1 Optimization over Neural Networks

For every predicted variable y that is determined by a NN prediction, we associate a distinct network flow model. A NN can be viewed as a multi-source single-terminal² arc-weighted acyclic layered digraph $\mathbf{N} = (\mathbf{V}, \mathbf{A})$. The node set \mathbf{V} is partitioned into a collection of layers $\mathbf{V}_1 \cup \dots \cup \mathbf{V}_l$. There is a one-to-one mapping between input features α of g and nodes $\mathbf{v} \in \mathbf{V}_1$. For any node $\mathbf{v} \in \mathbf{V}_1$, we denote the corresponding feature by $\alpha(\mathbf{v})$. Set \mathbf{V}_l consists of a single terminal node \mathbf{t} . For $j \in \{2, \dots, l\}$, every node $\mathbf{u} \in \mathbf{V}_j$ has a *bias* $\mathbf{B}(\mathbf{u})$ learned during training.

Each arc $\mathbf{a} = (\mathbf{u}, \mathbf{v}) \in \mathbf{A}$ is directed from a node \mathbf{u} in layer \mathbf{V}_j to a node \mathbf{v} in layer \mathbf{V}_{j+1} for some $j \in \{1, \dots, l-1\}$. Every arc has a *weight* $w(\mathbf{a})$ learned during training.

Given values $\alpha(\mathbf{v})$ for all nodes $\mathbf{v} \in \mathbf{V}_1$, the prediction from a NN with a ReLU activation function is calculated recursively by assigning a value $\mathbf{F}_\mathbf{v}$ to all nodes in the NN via the following iterative procedure and returning $\mathbf{F}_\mathbf{t}$:

- $\forall \mathbf{v} \in \mathbf{V}_1, \mathbf{F}_\mathbf{v} = \mathbf{G}_\mathbf{v} = \alpha(\mathbf{v});$
- For $j = 2, \dots, l-1, \forall \mathbf{v} \in \mathbf{V}_j, \mathbf{G}_\mathbf{v} = \sum_{\mathbf{u} \in \mathbf{V}_{j-1}} w((\mathbf{u}, \mathbf{v})) \cdot \mathbf{F}_\mathbf{u} + \mathbf{B}(\mathbf{v}), \mathbf{F}_\mathbf{v} = \max\{0, \mathbf{G}_\mathbf{v}\};$ and
- $\mathbf{F}_\mathbf{t} = \mathbf{G}_\mathbf{t} = \sum_{\mathbf{u} \in \mathbf{V}_{l-1}} w((\mathbf{u}, \mathbf{t})) \cdot \mathbf{F}_\mathbf{u} + \mathbf{B}(\mathbf{t}).$

Here $\mathbf{G}_\mathbf{v}$ is the input of the ReLU function, and $\mathbf{F}_\mathbf{v}$ is the output of the ReLU function.

We further define $z_\mathbf{u} \forall \mathbf{u} \in \mathbf{V}$ as a binary variable indicating if $\mathbf{G}_\mathbf{u} > 0$. With this interpretation, one can formulate the following model (**MOD-NN**) to calculate y based on the inputs $\alpha(\mathbf{v}), \forall \mathbf{v} \in \mathbf{V}_1$, that are the features, which can be fixed constants or functions of the decision variables.

$$y = \mathbf{F}_\mathbf{t} \tag{MOD-NN}$$

$$\mathbf{F}_\mathbf{v} = \mathbf{G}_\mathbf{v}, \quad \forall \mathbf{v} \in \mathbf{V}_1 \cup \mathbf{V}_l \tag{6}$$

$$\mathbf{G}_\mathbf{v} = \alpha(\mathbf{v}), \quad \forall \mathbf{v} \in \mathbf{V}_1 \tag{7}$$

$$\mathbf{G}_\mathbf{v} = \sum_{\mathbf{u} \in \mathbf{V}_{j-1}} w((\mathbf{u}, \mathbf{v})) \cdot \mathbf{F}_\mathbf{u} + \mathbf{B}(\mathbf{v}), \quad \forall j \in \{2, \dots, l\}, \forall \mathbf{v} \in \mathbf{V}_j \tag{8}$$

$$-M \cdot (1 - z_\mathbf{v}) \leq \mathbf{G}_\mathbf{v} \leq M \cdot z_\mathbf{v}, \quad \forall \mathbf{v} \in \mathbf{V}_2 \cup \dots \cup \mathbf{V}_{l-1} \tag{9}$$

$$\mathbf{G}_\mathbf{v} - M(1 - z_\mathbf{v}) \leq \mathbf{F}_\mathbf{v} \leq \mathbf{G}_\mathbf{v} + M \cdot (1 - z_\mathbf{v}), \quad \forall \mathbf{v} \in \mathbf{V}_2 \cup \dots \cup \mathbf{V}_{l-1} \tag{10}$$

$$0 \leq \mathbf{F}_\mathbf{v} \leq M \cdot z_\mathbf{v}, \quad \forall \mathbf{v} \in \mathbf{V}_2 \cup \dots \cup \mathbf{V}_{l-1} \tag{11}$$

$$z_\mathbf{v} \in \{0, 1\}, \quad \forall \mathbf{v} \in \mathbf{V} \tag{12}$$

$$\mathbf{G}_\mathbf{v} \text{ unconstrained}, \quad \forall \mathbf{v} \in \mathbf{V} \tag{13}$$

$$\mathbf{F}_\mathbf{v} \text{ unconstrained}, \quad \forall \mathbf{v} \in \mathbf{V}. \tag{14}$$

²The NN does not always have a single terminal. In our case, we only have one output, and so the output layer in our trained NN only has one node.

Constraints (6) guarantee that the ReLU activation function is not enforced in the input and output layers. Constraints (7) enforce that the values of the nodes in the input layer is the values of each input variable of the predictive model. These will either be fixed constants or the value determined by the optimization model for a single problem variable. Constraints (8) compute the input of the ReLU function of each node that is not on the first layer. Constraints (9) to (11) enforce the ReLU activation function on each hidden layer.

In **(PROBLEM-ORI)**, $\alpha(v)$ can be a regular variable or a constant, and there will be a separate network flow formulation for each of the predicted variables that are outcomes of neural networks. We test the impact of the size of the neural network and the number of such predicted variables on the performance of JANOS in Section 5.

4.2 Optimization over Logistic Regression Models

JANOS provides a parameterized discretization for handling logistic regression prediction. Specifically, it represents the nonlinear function of a logistic regression model using a piece-wise linear approximation, partitioning the range of the *log-odds*, i.e., $X\beta$ as in $y_k = \frac{1}{1+e^{-X\beta}}$ into a collection of mutually exclusive and collectively exhaustive intervals. The number of intervals is a parameter that can be specified by users. This is done by taking any number of intervals, and forcing the output of the logistic regression model to take a fixed value within a small window of what the actual predicted value is. This idea is illustrated on Figure 2.




Figure 2: A logistic function curve for a collection of points. Each point represent an element of a training set, with coordinates given by the dot product of the coefficients and their features, and the dependent variable. The curve shows a fit logistic regression, with confidence interval in estimation around the points. The horizontal lines depict the linearization currently implemented in JANOS, which associates the mean value of the function with every interval in the discretization.

We partition the range of the *log-odds* into Δ intervals, $[L_\delta, U_\delta]$, for $\delta \in \{1, \dots, \Delta\}$, where $L_{\delta+1} = U_\delta$ and we assume that the length of each interval $[\mathbf{sigmoid}(L_\delta), \mathbf{sigmoid}(U_\delta)]$, $\forall \delta \in$

$\{1, \dots, \Delta\}$ is uniform. The range of the *log-odds* is computed based on the bounds of the features. We use the mean of the **sigmoid** function value within an interval to serve as a piece-wise linear approximation of the actual predicted value of the logistic regression model. Specifically, for interval δ , let

$$\mathcal{V}_\delta = \frac{\log(1 + e^{U_\delta}) - \log(1 + e^{L_\delta})}{U_\delta - L_\delta}.$$

The value \mathcal{V}_δ is the average value of the **sigmoid** function over all values between L_δ and U_δ , where

$$\mathbf{sigmoid}(\mathbf{a}) = \frac{e^{\mathbf{a}}}{1 + e^{\mathbf{a}}}.$$

We define $z_\delta, \forall \delta \in \{1, \dots, \Delta\}$ as a binary variable indicating if we select a value for y_k in interval δ . Let F be the vector of features. Let β be the vector of estimated coefficients in the logistic regression model. With this interpretation, one can formulate (**MOD-LOGIT**) to maximize over a logistic regression model *approximately*, using the following transformation:

$$\max \quad y \quad (\mathbf{MOD-LOGIT})$$

$$\text{s.t.} \quad \sum_{\delta=1}^{\Delta} z_\delta = 1, \quad (15)$$

$$(L_\delta - L_1) \cdot z_\delta + L_1 \leq F\beta, \quad \forall \delta \in \{1, \dots, \Delta\} \quad (16)$$

$$(U_\delta - U_\Delta) \cdot z_\delta + U_\Delta \geq F\beta, \quad \forall \delta \in \{1, \dots, \Delta\} \quad (17)$$

$$\mathbf{sigmoid}(L_1) + (\mathcal{V}_\delta - \mathbf{sigmoid}(L_1)) \cdot z_\delta \leq y, \quad \forall \delta \in \{1, \dots, \Delta\} \quad (18)$$

$$\mathbf{sigmoid}(U_\Delta) + (\mathcal{V}_\delta - \mathbf{sigmoid}(U_\Delta)) \cdot z_\delta \geq y, \quad \forall \delta \in \{1, \dots, \Delta\} \quad (19)$$

$$z_\delta \in \{0, 1\}. \quad \forall \delta \in \{1, \dots, \Delta\} \quad (20)$$

The value assumed by y will be approximately equal to $\mathbf{sigmoid}(F\beta)$, used as the value predicted by a logistic regression model.

Constraint (15) ensures that only one interval is selected. Constraints (16) to (17) select the interval that contains the linear combination $F\beta$. Constraints (18) to (19) make sure that y equals the mean outcome value of the selected interval. Recall that F will be determined partially through fixed features and partially through decision variables.

5 Example Applications

In this section, we explore an example of allocating offered scholarship to admitted college students to exhibit the capability and flexibility of JANOS v0.0.9. All predictive models were built in Python3.6.9 using `scikit-learn` 0.22.2 (Pedregosa et al. 2011) and all optimization models were solved with `Gurobi Optimizer` v9.0.2 (Gurobi Optimization 2019) setting the `MIPGap` parameter to 0.1%. All experiments were run in Ubuntu 18.04.4 LTS on a Intel(R) Xeon(R) Gold 6130 CPU @ 2.10GHz processor with 16 GB RAM. Data, code, and results are available on GitHub³.

³https://github.com/iveyhuang/janos_paper_experiments

5.1 Problem Description

The Admission Office of a university wants to offer scholarship to its admitted students in order to bolster the class profile, such as academic metric, and often to simply maximize the expected class size (Maltz et al. 2007).

The Admission Office has collected from previous enrollment years the applicants’ SAT, GPA, scholarship offer, and matriculation result, i.e., whether the student accepted the offer or not. This year, suppose the school is issuing N offers of admission. Moreover, suppose the budget available for offered scholarship is $\$0.2 \cdot N \cdot 10^4$, denoted by **BUDGET** henceforth. The amount of scholarship that can be assigned to any particular applicant is between $\$0$ and $\$25,000$. The Admission Office wants to maximize the incoming class size.

To solve this optimization problem using JANOS, one can pre-train a model to predict the probability of a candidate accepting an offer given this student’s SAT, GPA and scholarship offered. The decision to make is the third feature *for each student*: the amount of scholarship to offer to each student.

We model this problem as follows (**STUDENT-ENROLL**):

$$\begin{aligned} \max \quad & \sum_{i=1}^N y_i && \text{(STUDENT-ENROLL)} \\ \text{s.t.} \quad & \sum_{i=1}^N x_i \leq \mathbf{BUDGET} && (21) \\ & y_i = g(\mathbf{s}_i, \mathbf{g}_i, x_i; \theta), && \forall i \in \{1, \dots, N\} && (22) \\ & 0 \leq x_i \leq 25,000, && \forall i \in \{1, \dots, N\} && (23) \end{aligned}$$

where,

- x_i is the decision variable, i.e., the amount of scholarship assigned to each student accepted;
- \mathbf{s}_i is the SAT score of applicant i (standardized using z -score);
- \mathbf{g}_i is the GPA score of applicant i (standardized using z -score);
- y_i is a predicted variable per admitted student, the outcome of a predicted model representing the probability of a candidate accepting the offer; and
- g is a predicted model pre-trained to predict any candidates’ probabilities of accepting an offer. The parameters \mathbf{s}_i , \mathbf{g}_i and x_i are the predictive model’s inputs. The vector θ represents the parameters of the predictive model, which we assume to be the same for each applicant. The function g can be any of the permissible predictive models with θ determined prior to optimization.

5.2 Experimental Results

We utilize randomly generated realistic student records to train predictive models and test the efficiency of the solver when solving different-sized problems with variations in parameters as

well. We build the three permissible models (linear regression, logistic regression, and neural networks) with various parameter settings, i.e., the number of intervals for logistic regression models and the hidden layer sizes for neural networks. Each of the models is trained on 20,000 randomly generated student records.

We then generate sets of random student records of different sizes to test the scalability of JANOS, 50, 100, 500 and 1,000 admitted students. These experiment instances are produced using the same data-generating scheme as was used for building the training set. In the test data sets, each record contains only a SAT score and a GPA score, where the decision is to assign scholarship based on prediction of enrollment. We document how long it takes JANOS to solve problems of different sizes with different predictive models.

We first provide an analysis of the total runtime using various parameters for each of the models. For the linear regression there are no configurable parameters, and so we have only one setting, **LinReg**. For the logistic regression model, the main parameter of interest is the number of intervals in the discretization. For a fixed number of intervals Δ , let **LogReg**(Δ) refer to solving the model with logistic regression prediction with Δ intervals. For neural networks, there are several parameters one might tune, most apparent being the configuration of the neurons. We fix 10 nodes per hidden layer and vary the number of hidden layers to 1, 2, 3. Let **NN**(h) refer to solving the model with neural networks with h hidden layers.

For each of the predictive model specifications, we run 5 instances and take the mean runtime. Figure 3 depicts the runtimes. On the x -axis we indicate the number of admitted students in the admitted pool. On the y -axis we report the runtime in seconds. **LinReg** yields the most efficient model, taking up to a second to solve. **LogReg**(Δ) takes an increasing amount of time as Δ grows, as does **NN**(h) for increasing h , but the runtimes are not prohibitively large. Note that largest instances have 1,000 logistic regression approximations or 1,000 neural network flow models.




Figure 3: The average runtimes of linear regression models, logistic regression models, and neural networks with different scales.

We also evaluate how well the approximation of the logistic regression performs at obtaining

optimal solutions. We apply the logistic regression approximation for 10 instances with 50 students. The results are reported on Figure 4. On the x -axis we indicate the number of intervals in the approximation, and on the y -axis we report the *root mean squared error* (RMSE) of the probability estimates given by the approximation and the actual learned logistic regression evaluation at the optimal solutions obtained by the approximation. As the number of intervals increases, the approximation becomes stronger, but as discussed earlier, increases runtime.




Figure 4: The quality of the linear approximation of the logistic regression function at optimal solutions.

In order to evaluate the expected improvement in solutions obtained from JANOS over what might be done in practice, we evaluate the solution obtained by JANOS and compare with the following heuristic that can be employed for any predictive model g for this application:

1. Sort the accepted students in non-decreasing order of

$$g(\mathbf{s}_i, \mathbf{g}_i, \$25,000|\theta) - g(\mathbf{s}_i, \mathbf{g}_i, \$0|\theta).$$

2. Following the above order, allocate the maximum permissible aid (i.e., \$25,000) until reaching the maximum budget.

This is a realistic heuristic because it greedily assigns scholarship to the students in the order of those that are most sensitive to scholarship. Results obtained using this heuristic are listed under “Heuristic” in Table 1.

There are other ways that a practitioner who is well-versed in both optimization and predictive modeling might address a decision problem of this sort. For example, in this application, one could discretize the domain of the decision variables to take values $D = \{0, 1, \dots, 25,000\}$ and then evaluate the predictive model $g(\cdot)$ at each value for each admitted student. However, such a transformation results in a pseudo-polynomial size model and also admits only an approximation. Note that if desired, one can model the problem in this way directly in JANOS by declaring the decision variables as discrete and setting their domain to D . So we provide results evaluated using JANOS by setting the regular variables as discrete in $\{\$0, \$5,000, \$10,000, \$15,000, \$20,000, \$25,000\}$, listed under

“JANOS_Discrete” in Table 1. Results obtained by solving **STUDENT-ENROLL** are listed under “JANOS_Continuous” in Table 1.

Table 1 reports results from the experimental evaluation. In particular, for logistic regression (**LogReg**(Δ)) and neural network prediction (**NN**(1)) models, we report for 500 and 1000 admitted students the expected number of enrolled students based on the allocation determined by JANOS and the heuristic described above. We also report for both models and for each N the percent reduction in admitted student declination of admission.

The results indicate that simply by a more careful assignment of scholarship and making no other changes, JANOS can provide a substantial improvement in expected matriculation rates. In particular, we see that the heuristic described above would yield under a 50% expected matriculation rate on average, while the solutions obtained using JANOS yield approximately 60% expected matriculation. Using a discretization of monetary increments within JANOS yields slightly lower expected matriculation rates than a continuous domain, but is a realistic setting for a university where giving scholarships in fixed increments is more common, even though more expected enrollments can be realized from allowing the university any level of scholarship for any student, i.e., allowing continuous assignment of scholarship. In summary, going from the heuristic to using a discretization can reduce the number of declined admissions by about 20%, and moving from the discretized model to the continuous model can reduced declined admissions by an additional 6%. This example exemplifies the improvement in decision-making capability that JANOS can obtain.

Table 1: Improvement by JANOS over a basic heuristic.

LogReg	Heuristic	JANOS_Discrete	JANOS_Continuous	Expected % Reduction in Declination	
				JANOS_Discrete vs. Heuristic	JANOS_Continuous vs. JANOS_Discrete
500	243.17	297.45	308.59	21.13%	5.50%
1000	467.16	576.35	599.95	20.49%	5.57%
NN	Heuristic	JANOS_Discrete	JANOS_Continuous	Expected % Reduction in Declination	
				JANOS_Discrete vs. Heuristic	JANOS_Continuous vs. JANOS_Discrete
500	243.28	299.67	311.86	21.97%	6.08%
1000	467.44	580.66	606.57	21.26%	6.18%

6 Accessing the Solver

JANOS works with Python3 and currently requires **Gurobi** for optimization and **sklearn** for predictive modeling. You also must have **numpy** and **matplotlib** installed. Please refer to JANOS’s website (<http://janos.opt-operations.com>) for more information, where a user manual, quick start guide, and examples are provided, including all data used for the experimental analysis in this paper.

7 Conclusions and Future Work

We propose a modeling framework JANOS that integrates predictive modeling and prescriptive analytics. JANOS is a useful tool both for practitioners and researchers who are seeking to inte-

grate machine learning models within a discrete optimization model. JANOS can be expanded in many directions. Adding the capability to handle additional predictive models like random forests will greatly expand the potential for use in industry. Investigating how other predictive models, like support vector regressions with an arbitrary kernel function, can be incorporated is also an interesting avenue for research. Another direction is to explore alternative MIP formulations for existing predictive models, such as the MIP formulations for trained NNs in Anderson et al. (2020). Incorporating uncertainty in the predictions and making decisions that account for their impact as suggested in Bertsimas and Kallus (2020) is a further direction to explore.

References

- Anderson R, Huchette J, Ma W, Tjandraatmadja C, Vielma JP (2020) Strong mixed-integer programming formulations for trained neural networks. *Mathematical Programming* 1–37.
- Aouad A, Feldman J, Segev D, Zhang D (2019) Click-based mnl: Algorithmic frameworks for modeling click data in assortment optimization. *Available at SSRN 3340620* .
- Bengio Y, Lodi A, Prouvost A (2018) Machine learning for combinatorial optimization: a methodological tour d’horizon. *arXiv preprint arXiv:1811.06128* .
- Bertsimas D, Kallus N (2020) From predictive to prescriptive analytics. *Management Science* 66(3):1025–1044.
- Bezanson J, Edelman A, Karpinski S, Shah VB (2017) Julia: A fresh approach to numerical computing. *SIAM review* 59(1):65–98, URL <https://doi.org/10.1137/141000671>.
- Bienstock D, Muñoz G, Pokutta S (2018) Principled deep neural network training through linear programming. *arXiv preprint arXiv:1810.03218* .
- Biggs M, Hariss R (2018) Optimizing objective functions determined from random forests. *Available at SSRN 2986630* .
- Boyd S, Parikh N, Chu E, Peleato B, Eckstein J, et al. (2011) Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 3(1):1–122.
- Cappart Q, E G, Bergman D, Rousseau L (2019) Improving Optimization Bounds using Machine Learning: Decision Diagrams Meet Deep Reinforcement Learning. *ArXiv e-prints* .
- Cozad A, Sahinidis NV, Miller DC (2014) Learning surrogate models for simulation-based optimization. *AIChE Journal* 60(6):2211–2227.
- Demirović E, Stuckey PJ, Bailey J, Chan J, Leckie C, Ramamohanarao K, Guns T (2019) An investigation into prediction+ optimisation for the knapsack problem. *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, 241–257 (Springer).
- Deng Y, Liu J, Sen S (2018) Coalescing data and decision sciences for analytics. *Recent Advances in Optimization and Modeling of Contemporary Problems*, 20–49 (INFORMS).
- Dunning I, Huchette J, Lubin M (2015) Jump: A modeling language for mathematical optimization. *SIAM Review* 59, URL <http://dx.doi.org/10.1137/15M1020575>.
- Elmachtoub AN, Grigas P (2017) Smart” predict, then optimize”. *arXiv preprint arXiv:1710.08005* .
- Feldman J, Zhang D, Liu X, Zhang N (2019) Customer choice models versus machine learning: Finding optimal product displays on alibaba. *Available at SSRN* .
- Ferreira KJ, Lee BHA, Simchi-Levi D (2015) Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* 18(1):69–88.

- Feurer M, Klein A, Eggenberger K, Springenberg J, Blum M, Hutter F (2015) Efficient and robust automated machine learning. *Advances in neural information processing systems*, 2962–2970.
- Fischetti M, Jo J (2017) Deep neural networks as 0-1 mixed integer linear programs: A feasibility study. *arXiv preprint arXiv:1712.06174* .
- Gurobi Optimization L (2019) Gurobi optimizer reference manual. URL <http://www.gurobi.com>.
- Hart WE, Watson JP, Woodruff DL (2011) Pyomo: modeling and solving mathematical programs in python. *Mathematical Programming Computation* 3(3):219.
- Huang T, Bergman D, Gopal R (2019) Predictive and prescriptive analytics for location selection of add-on retail products. *Production and Operations Management* 28(7):1858–1877.
- Khalil EB, Dilkina B, Nemhauser GL, Ahmed S, Shao Y (2017) Learning to run heuristics in tree search. *IJCAI*, 659–666.
- Koh K, Kim SJ, Boyd S (2007) An interior-point method for large-scale l1-regularized logistic regression. *Journal of Machine Learning Research* 8(Jul):1519–1555.
- Lam SS, Petri KL, Smith AE (2000) Prediction and optimization of a ceramic casting process using a hierarchical hybrid system of neural networks and fuzzy logic. *Iie Transactions* 32(1):83–91.
- Lemos H, Prates M, Avelar P, Lamb L (2019) Graph colouring meets deep learning: Effective graph neural network models for combinatorial problems. *arXiv preprint arXiv:1903.04598* .
- Maltz EN, Murphy KE, Hand ML (2007) Decision support for university enrollment management: Implementation and experience. *Decision Support Systems* 44(1):106–123.
- Mandi J, Demirović E, Stuckey P, Guns T, et al. (2019) Smart predict-and-optimize for hard combinatorial optimization problems. *arXiv preprint arXiv:1911.10092* .
- Mišić VV (2017) Optimization of tree ensembles. *arXiv preprint arXiv:1705.10883* .
- Nazari M, Oroojlooy A, Snyder L, Takác M (2018) Reinforcement learning for solving the vehicle routing problem. *Advances in Neural Information Processing Systems*, 9839–9849.
- Paul A, Feldman J, Davis JM (2018) Assortment optimization and pricing under a nonparametric tree choice model. *Manufacturing & Service Operations Management* 20(3):550–565.
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E (2011) Scikit-learn: Machine Learning in Python . *Journal of Machine Learning Research* 12:2825–2830.
- Schweidtmann AM, Mitsos A (2019) Deterministic global optimization with artificial neural networks embedded. *Journal of Optimization Theory and Applications* 180(3):925–948.
- Serra T, Tjandraatmadja C, Ramalingam S (2017) Bounding and counting linear regions of deep neural networks. *arXiv preprint arXiv:1711.02114* .
- Wilder B, Dilkina B, Tambe M (2019) Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1658–1665.