



Validating Goal-Oriented Hypotheses of Business Problems Using Machine Learning: An Exploratory Study of Customer Churn

Sam Supakkul¹(✉), Robert Ahn², Ronaldo Gonçalves Junior², Diana Villarreal¹, Liping Zhao³, Tom Hill², and Lawrence Chung²

¹ NCR Corporation, Atlanta, GA, USA

{sam.supakkul,diana.villarreal}@ncr.com

² The University of Texas at Dallas, Richardson, TX, USA

{robert.sungsoo.ahn,ronaldo.goncalves,chung}@utdallas.edu,
tom.hill.fellow@gmail.com

³ University of Manchester, Manchester, UK

liping.zhao@manchester.ac.uk

Abstract. Organizations are investing in Big Data and Machine Learning (ML) projects, but most of these projects are predicted to fail. A study shows that one of the biggest obstacles is the lack of understanding of how to use data analytics to improve business value. This paper presents *Metis*, a method for ensuring that business goals and the corresponding business problems are explicitly traceable to ML projects and where potential (i.e., hypothesized) complex problems can be properly validated before investing in costly solutions. Using this method, business goals are captured to provide context for hypothesizing business problems, which can be further refined into more detailed problems to identify features of data that are suitable for ML. A Supervised ML algorithm is then used to generate a prediction model that captures the underlying patterns and insights about the business problems in the data. An ML Explainability model is used to extract from the prediction model the individual features and their degree of contribution to each problem. The extracted weighted data feature are then fed back to the goal-oriented problem model to validate the most important business problems. Our experiment results show that *Metis* can detect the most influential problem when it was not apparent through data analysis. *Metis* is illustrated using a real-world customer churn (customer attrition) problem for a bank and a publicly available customer churn dataset.

1 Introduction

Data Analytics and Machine Learning (ML) technologies benefit from a continuous improvement cycle where large amounts of data are constantly being created. Organizations are investing in Big Data and ML projects, but most of these projects are predicted to fail [1]. A study may have suggested a possible

reason which is the lack of understanding of how to use data analytics to improve business value [2]. This finding is a clear sign that stakeholders do not see the end-to-end relationship between important business goals and the emerging Big Data and ML technologies [3, 4].

Additionally, some business problems can only be hypothesized as they are difficult to validate using the traditional data analysis techniques. For example, applying data analysis on the customer churn dataset [5] during our experiment showed that there is no evidence to suggest that the customers who left the bank had a higher degree of dissatisfaction with many of the service qualities than those loyal customers.

Building on our previous approach, GOMA [6], this paper proposes *Metis*¹ to support goal-oriented hypothesis and validation of business problems. Three technical contributions are made in this paper, including 1) an ML-based approach to extracting an actual root cause hidden in the data to validate hypothesized business problems, 2) an ontology that more explicitly and formally describes the relevant modeling concepts related to business goals, problems and ML, and 3) a set of formalized validation rules for reasoning about problem hypothesis validation in a goal-oriented problem model.

The proposed approach is illustrated using a real-world banking customer churn problem, which was adapted from the example used in [6]. In the adopted case-study, a retail bank hired a company, specializing in data mining, to help address the churning problem by using insights from detailed transaction data in a newly installed powerful data warehouse [7]. The company hypothesized potential reasons to why the customers were canceling their accounts, and validated them with descriptive insights mined using a data classification technique. Since the actual dataset used by the consulting company was not available, we use a publicly available bank customer churn dataset [5] and reversed engineer to update the business problem hypotheses so that they are consistent with the dataset used. We use this example to demonstrate that *Metis* could be used to provide traceability between business problems and an ML solution, which can also reveal the insights about a root cause to a problem that may be difficult to discover using data analysis.

The rest of this paper is structured as follows. Section 2 presents the *Metis* method. Section 3 describes the experiment conducted and their results. Section 4 discusses related work along with our observed threats to validity and limitations. Finally, Sect. 5 summarizes the paper and future work.

2 Metis: A Goal-Oriented Problem Hypothesis Validation Method Using Machine Learning

To be able to hypothesize business problems and subsequently data features needed for developing an ML model, a good understanding of concepts in the domain in question is needed. In this section, we first present an example banking

¹ A Greek goddess that has been associated with prudence, wisdom, or wise counsel.

domain-specific ontology that underlies the customer churn problem that we use as a running example to illustrate Metis; we then describe the Metis domain-independent ontology used to support the modeling and validation of problem hypotheses in Metis; finally, we describe and illustrate the Metis process.

2.1 Domain-Specific Banking Ontology

This section describes an example of domain-specific ontology for the banking example, which can vary significantly depending on different organizations, processes, and so on. Figure 1 shows a typical set of banking concepts and their relationships. It is worth mentioning that this domain-specific ontology supports the understanding of the banking domain, and it does not represent a schema or model related to database design.

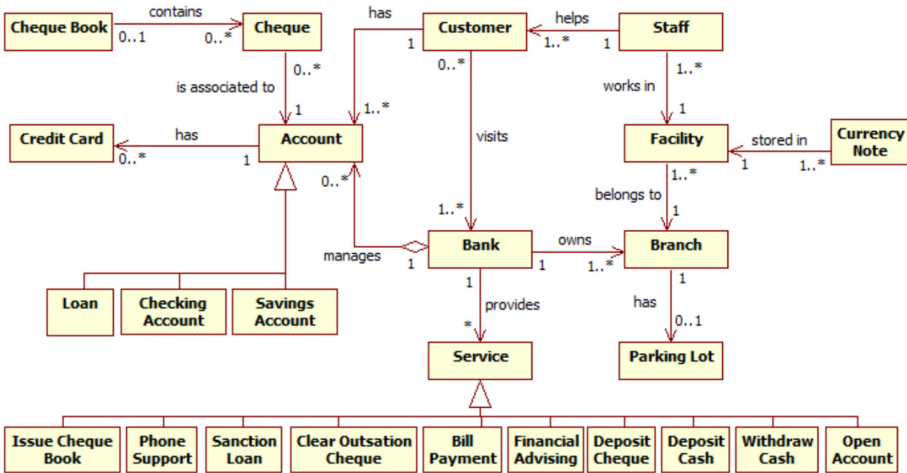


Fig. 1. Banking domain-specific ontology diagram

Some of the ontological concepts are briefly described here as examples. Banks provide numerous services, e.g., financial advising, cash withdrawal, and so on. It is important to study the qualitative aspect of these services in order to have a clear understanding of customer satisfaction. For example, customers may feel that there is not enough parking space, a lack of pleasant ambiance, no comfortable seating arrangements, lack of immediate attention, and so on. For the customer churn problem in the running example in this paper, the quality aspect of both facility and service related concepts are essential to generate hypotheses about problems.

2.2 The Metis Ontology

While modeling the mapping between a Goal-Oriented ontology and an ML-based ontology, completeness and soundness are two major concerns.

To completely and formally address these concerns for the *Metis* method, the following subsections describe the modeling concepts and the semantic reasoning formalization for the *Metis* ontology.

Modeling Concepts. A complete set of concepts and their relationships can be found in Fig. 2. To avoid omissions while mapping Goal-Oriented and ML, important concepts such as Problem, Hypothesis, and Machine Learning Model were explicitly represented. In addition, the ontology also comprehends concepts related to Big Data and Big Queries, and Features are derived from modeling concepts from a domain-dependant ontology. *Metis* is a domain-independent ontology that can be applied to a variety of domains. Section 2.1 describes a banking domain-specific ontology example.

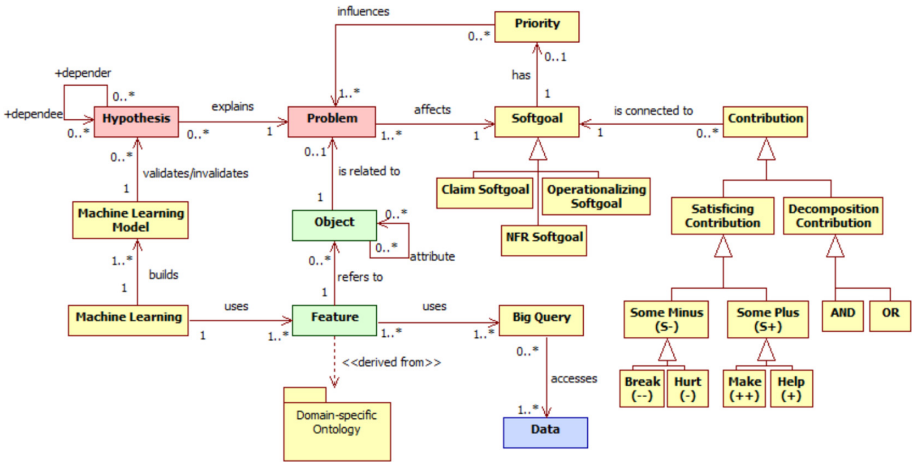


Fig. 2. Metis domain-independent ontology diagram

In a Goal-Oriented approach, problems represent undesirable situations, vulnerabilities, or threats to achieving stakeholders’ goals. For the banking customer churn example in Fig. 4, *Customer Churn* is a problem that negatively impacts the goal *Retain existing customers*. Goals and problems can be further decomposed, and we can hypothesize what might contribute to their realization based on domain knowledge, statistics, and so on.

An acceptable representation of hypotheses and problems can be generated, but ultimately, we want to determine whether we can validate these hypotheses with respect to the problems in consideration. In this context, ML is used to build models to identify the importance of features such as *Immediate Attention*. By using the relevant features, it is possible to establish how to validate or invalidate hypotheses. For instance, in Fig. 4 we hypothesize that *Lack of immediate attention* has a S+/S- contribution to the problem *Poor Service*, which in turn contributes to *Customer Churn*.

Semantic Reasoning Formalization. An important aspect of contributions in the hypothesis validation process is the propagation of validations throughout the connected hypotheses, since a hypothesis might contribute to multiple other hypotheses. This validation process starts in the lowermost level of the hypotheses and propagates until we validate or invalidate problems. A formal definition may be described as follows:

Let $validated(P_n)$ be the proposition that the problem hypothesis P_n is validated, for $n \in \mathbb{Z}^+$. Let i be an arbitrary integer, where $i > 1$. For all $j \in \mathbb{Z}^+$, let $P_{i-1,j}$ be the j th problem hypothesis directly decomposed from P_i . Assuming this decomposition is of type OR/AND, the validation propagation can be represented by the following:

$$\left(\bigvee_j validated(P_{i-1,j}) \right) \rightarrow validated(P_i) \quad (1)$$

$$\left(\bigwedge_j validated(P_{i-1,j}) \right) \rightarrow validated(P_i) \quad (2)$$

Alternatively, hypotheses and problems can be connected using a positive (S+) or negative (S-) contribution. We want to determine which hypothesis in the source set (i.e., hypotheses that originate the contributions) is more relevant to the target set (i.e., hypotheses that receive the contributions) in order to maximize the validation insights generated by the application of ML models. In this case, validating a hypothesis P_i will now depend on the validation of the selection for P_i .

$$validated(selection(P_i)) \rightarrow validated(P_i) \quad (3)$$

By following the *Metis* proposed approach, feature importance values (I) are obtained from the results of running ML algorithms and they are associated with their respective contributions, i.e., a contribution from problem hypothesis P_s to P_t has an importance value $I_{s,t}$ and its weight and type may be determined by the following:

$$weight(P_s, P_t) = I_{s,t} \quad (4)$$

$$ctr_type(I_{s,t}) = \begin{cases} \text{S+} & \text{if } I_{s,t} \geq 0 \\ \text{S-} & \text{if } I_{s,t} < 0 \end{cases} \quad (5)$$

A source hypothesis has a score based on the weight of the targeted hypotheses and their respective contributions. The function $weight(P_t)$ describes the importance weight of a target hypothesis. Hence, the overall score for a source hypothesis P_s can be given by the utility function as follows:

$$score(P_s) = \left(\sum_{t=1}^{\#targets} weight(P_t) \times weight(P_s, P_t) \right) \quad (6)$$

After computing the scores for all source hypotheses using a utility function, the selection process may be carried out in a bottom-up approach [8]. We need to select the maximum value in the lowermost source hypothesis set to propagate that validation to the target hypothesis set. In other words, the selection process for a target is represented by choosing the source with highest score:

$$selection(P_t) = \max \left(score(P_s) \right)_{s=1}^{\#sources} \tag{7}$$

After the lowermost source hypothesis set is evaluated, we proceed to the next one until the selection process covers the entire set of hypotheses.

2.3 The Metis Process

The *Metis* process consists of four steps: 1) *Model business goals and problems*, 2) *Acquire data*, 3) *Detect feature importance*, and 4) *Validate hypotheses of business problems*. As shown in Fig. 3, *Step 1: Model business goals and problems* explicitly captures stakeholders’ needs and obstacles as goals and problems using a goal-oriented modeling approach [9, 10], where potential problems are posed as problem hypotheses to be validated. The outputs from this step are problem hypotheses in the context of business goals. *Step 2: Acquire data* derives data features from the business problems hypotheses to acquire the necessary data from external and/or internal sources, for instance using a customer survey or Big Data Spark SQL if the data are already available on-line. *Step 3: Detect feature importance* uses ML to learn patterns in the data to identify how problems are associated with the data features collectively. In addition, the output from this step includes Feature Importance that identifies the degree of each feature contributing to a problem. The final step, *Step 4: Validate hypotheses of business problems* uses the Feature Importance to validate the problem hypotheses modeled during Step 1.




Fig. 3. The *Metis* process for validating goal-oriented hypotheses of business problems

Step 1: Model Business Goals and Problems. In this step, important business or stakeholders’ needs are explicitly captured as Softgoals that can be further refined using AND or OR decomposition [11]. Using Fig. 4 as an example, at the highest organizational level, *Increased profitability* is a Softgoal to be achieved, which is refined using an AND decomposition to *Increased revenue* and *Increased profit margin* sub-goals, where the former is to be operationalized by *Increase customer base* strategic level goal. *Increase customer base* is

then further AND-decomposed to more specific operationalizing goals of *Retain existing customers* and *Acquire new customers*.

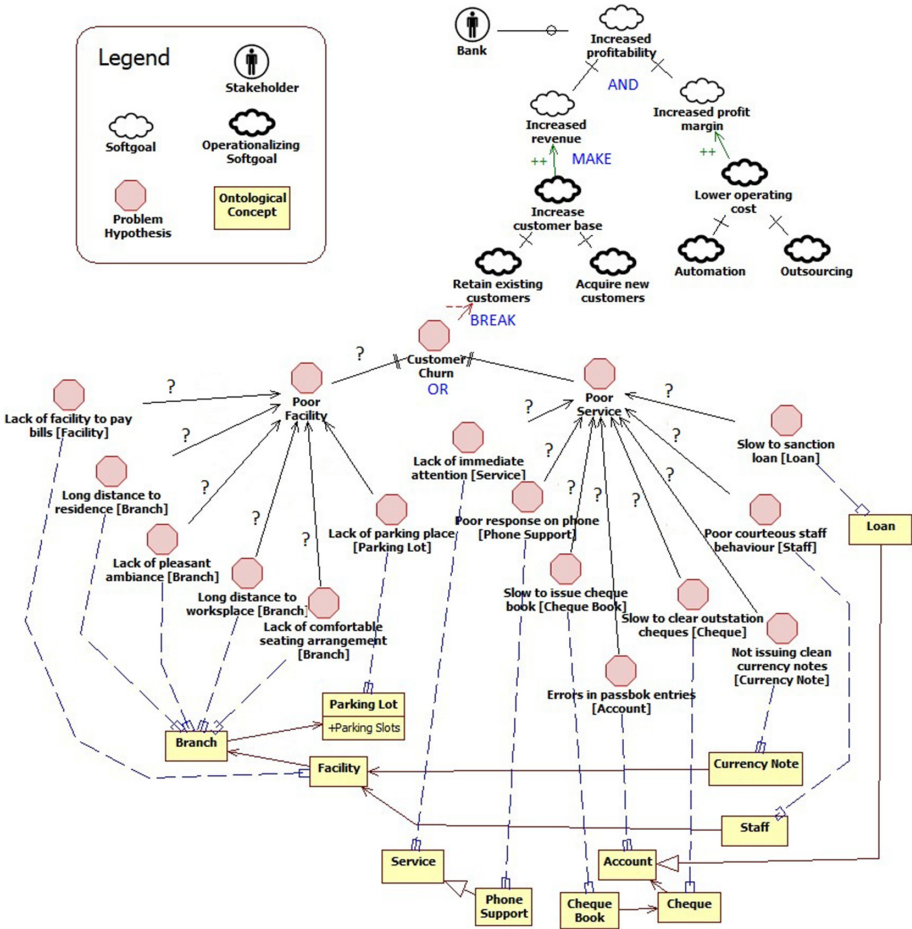


Fig. 4. Step 1: Model business goals and problems

Each lowest level goal is used as the context to identify potential problems that could hinder the goal achievement. The validity of each problem may be unknown at this point. Therefore, each problem is considered a target problem hypothesis to be validated by data. Similar to the goal refinement, each problem hypothesis may be further refined or realized by more specific problem hypotheses until they are low-level enough to identify the data features needed for data analysis or ML.

In this example, *Customer Churn* is a problem hypothesis that could *BREAK* (-) the *Retain existing customers* goal. *Customer Churn* is further refined using

an OR-decomposition to *Poor Facility* or *Poor Service* sub-problem hypotheses, which are used to identify potential causing problem hypotheses. *Poor Facility* is hypothesized to be caused by *Long distance to residence*, *Lack of pleasant ambience*, or other causes. Since each potential cause has not been validated whether it is indeed a contributing cause to the problem, the contribution link is labeled as unknown (depicted by a question mark).

Step 2: Acquire Data. This step examines the lowest level problem hypotheses to identify data features needed for data analysis or ML. Using Fig. 4 as an example, *Long distance to residence* and *Lack of pleasant ambience* may be used to identify *Distance to residence* and *Pleasant ambience* as the corresponding data features. The identified features are then used to build database queries or Big Data queries if the corresponding data are already available on-line. Otherwise, the required data features need to be acquired through other means, such as purchasing from a data provider, using a customer survey or generation from on-line sources [12].

An example of the acquired dataset is given in Fig. 5a, where F_1 – F_5 represent all features and L corresponds to the *Churner* or *Non-churner* indicator associated with the satisfaction scores for F_1 – F_5 , provided by individual customers C_1 – C_5 . For example, customer C_1 expressed dissatisfaction with feature F_1 and F_3 with scores of 2 and 3 accordingly. On the other hand, he/she expressed satisfaction with F_2 , F_4 and F_5 with scores of 8, 7, 8 accordingly. C_1 is noted by label 1 as a *Churner* customer in correlation with the given scores.

Step 3: Detect Feature Importance. The intuition for using ML is to encode the knowledge about the features hidden in the customer survey data, and then decode the knowledge representation to identify which feature is the true cause for the customer churn problem. To encode the feature knowledge, we use a Supervised ML algorithm with an assumption that an accurate prediction model represents the knowledge about features. To decode the influential features, we use a Model Explainability library [13] that was designed to explain how features contribute to the prediction outcomes.

Referring to Fig. 5b, this step splits the dataset into training and testing datasets. All features F_1 – F_5 and label L are processed by one or more Supervised ML algorithm to obtain the most desirable prediction model M_p . To determine whether M_p has been sufficiently trained to recognize the general patterns in the training dataset, it is measured on how accurate it can predict label L in the testing dataset. The accuracy is represented by an accuracy metric A_1 , which is based on the differences between predicted label L' and actual label L , where L' is generated from F_1 – F_5 in the testing dataset.

Once an accurate model M_p is obtained, it is processed by an ML Explainability algorithm to produce an Explainability model M_e , which is in turn used to detect feature importance I_1 – I_5 , where I_1 contains two pieces of information: sign and weight of the contribution F_1 makes towards the label L' as predicted M_p . The sign of the value indicates whether the corresponding feature helps or

hurts towards the predicted label, while the weight represents the amount of influence the feature has. Similarly, I_2 and I_3 represent the feature importance of F_2 and F_3 respectively. By having the highest value among all feature importance values, F_1 is considered the most influential feature, followed by F_2 and F_3 in the context of testing dataset.

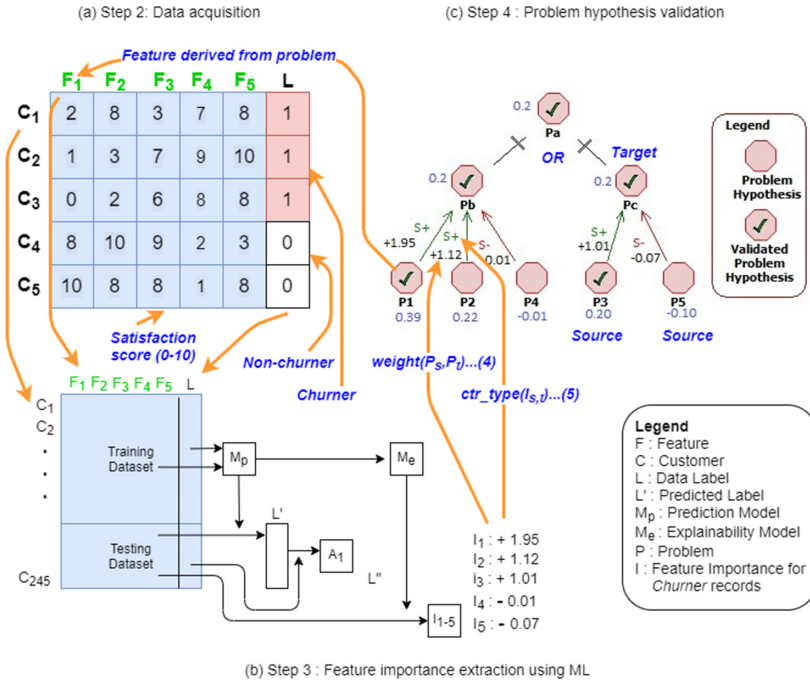


Fig. 5. Step 2, 3, 4 of the Metis process

Step 4: Validate Hypotheses of Business Problems. Referring to Fig. 5c, this step uses the feature importance values produced by the Explainability model M_e to validate problem hypotheses in the goal-problem model created in step 1, one parent-child problem set at a time in a bottom-up approach, using the quantitative and qualitative semantic reasoning formalization, as described in Sect. 2.2.

Using $P_b - (P_1, P_2, P_3)$ parent-child set as an example, the contribution link between each parent-child pair is updated by applying Formula 4 and 5 against the corresponding feature importance value where (P_1, P_2, P_3) and P_b are considered sources and target in the formulas respectively. In this example, the contribution type $ctr_type(P_1, P_b)$ is assigned with $S+$ by Formula 5 with feature importance I_1 with value +1.95 as a function parameter. I_1 is used since the corresponding F_1 was defined based on problem P_1 in step 2. To complete the contribution update, the weight of contribution is assigned with 1.95 by

Formula 4. Other contribution links with the same parent are updated in a similar fashion. Then, P_1 is selected among P_1 , P_2 and P_3 by Formula 7 to be a validated problem hypothesis since it is the most influential cause for problem P_b . After P_1 is quantitatively selected based on scores, P_b is qualitatively validated by Formula 3. Then, P_a can be qualitatively validated by Formula 1.

3 Experiment and Results

The analysis of customer feedback information may be beneficial to discerning customer satisfaction for the quality of important services. To this end, a publicly available dataset [5] acquired by Step 2 in the *Metis* process is analyzed in this section. This dataset contains typical customer information such as age and occupation. In addition, the dataset contains feedback information regarding certain banking service-related features (e.g., Immediate Attention) and facility-related features (e.g., Pleasant Ambiance). A customer can score each of these features from 0 to 10 (least to most satisfied). In this context, scores of 4 or less are used to describe some degree of dissatisfaction, an assumption that something might go wrong in a business operation, i.e., problem hypotheses. The next section describes an analysis of these problem hypotheses.

3.1 Dataset Analysis

Some examples of problem hypotheses are shown in Fig. 4, which includes *Long distance to residence*, *Lack of immediate attention*, and *Lack of pleasant ambiance*. Figure 6 shows the customer dissatisfaction with respect to these 3 features out of the 20 available features. From the total of customers that believe there is a long distance to their residence, 35% deserted the bank (churner) and 65% remained loyal (non-churner). Assessing this feature by occupation, notice that most unsatisfied customers are from the professional occupation, followed by private service and government service. Together, these three occupations represent 75% of customers unsatisfied with distance from residence.

From the customers that identified a lack of immediate attention, more than half are young customers (40 years old or younger). Analyzing pleasant ambiance by occupation, we can see that customers from the business occupation and the private service complained the most. In an overall assessment for customer dissatisfaction by loyalty, it is possible to notice that most customers remained loyal regardless of the problem hypotheses under consideration. Even though we are able to extract insights from the dataset, ultimately, there is no evidence of why customer deserted. For this purpose, Sect. 3.2 demonstrates results of using ML that can potential provide some evidence.

3.2 Prediction Models

To encode and represent knowledge about feature contributions using Supervised ML, we experimented with several ML algorithms, including Linear Regression,




Fig. 6. Analysis of customer dissatisfaction (score less than 5 in a 0 to 10 scale) for the features distance from residence, immediate attention, and pleasant ambiance

Support Vector Machine, Decision Tree, Random Forest, and XGBoost Classifier, among which XGBoost showed the highest accuracy rate in our experiment. Due to space limitation, only the results from XGBoost are discussed in this section.

The ML segments of the experiment were conducted using Python language and scikit-learn open-source ML libraries [14]. The dataset used for the experiment was a public banking customer churn dataset [5]. After data cleansing, 67% of the data (164 records) were used for model training and the 33% (81 records) for testing. Data features used included the customer responses to the survey questions, such as *Pleasant Ambiance*, *Comfortable Seating*, *Immediate Attention*, *Good Response On Phone* and others, in the scale of 0–10, excluding customer information, such as age and occupation, that were used separately for data analysis as reported in Sect. 3.1. The resulting prediction model showed an accuracy of 84% (F1 score) on the test dataset, which was better than other ML algorithms in our experiments. The modest accuracy rate was probably due to the small and highly unbalanced dataset that required a data pre-processing step that further reduced the dataset size.

3.3 Explainability Model

To extract feature contribution information from the resulting prediction model, we used SHAP (SHapley Additive exPlanations) [13], an Explainability library that uses a game theoretic approach to explain the output of many ML models. It connects optimal credit allocation with local explanations using the classic Shapley values from game theory and their related extensions.

Figure 7 is a Force Plot produced from a SHAP model (M_p in Fig. 5b) created from the most accurate XGBoost prediction model (M_e in Fig. 5b). It gives a visual representation of the influence each feature has towards the final output




Fig. 7. Features importance for one churner’s responses

value of 0.96. In this plot, the base value 0.18 is the average prediction value without any influences from the features, while output value of 0.96 is the output from the prediction model where 1 represents a churner customer. The influences of features are represented by the direction towards the output value and width of the corresponding arrow blocks. Here, *DistanceToResidence* feature has the most influence in increasing the output value away from the base value towards the final output value, which is consistent with the score of 0 (least satisfaction) given by the customer. On the other hand, *EnoughParkingSpace* has the most influence in the opposite direction, decreasing the value away from the final output value, which seems consistent with the satisfaction score of 5 (neutral satisfaction) given by the customer. It is interesting to note that *ImmediateAttn* with the value of 10 (most satisfaction) was seen as an influence towards the customer’s churner decision. SHAP does explain this counter intuitive result.

Figure 8a plots individual SHAP values for all features and all churner customers. Each dot represents a SHAP value that a feature has in support of increasing the output value towards 1 (Churner label in Fig. 5a). Visually, it is clear that *DistanceToResidence* has higher positive SHAP values than other features. For this experiment, the more positive SHAP values a feature has, the more influence it has towards the prediction outcome. This is supported by Fig. 8b where *DistanceToResidence* has the highest total(sum) SHAP value.

3.4 Validating Problem Hypotheses

By following Step 4 of the *Metis* process (Sect. 2.3), we applied Formula 4 and 5 against the sum SHAP value for the respective feature (see Fig. 8b), which led to the validation of *Long distance to residence* problem hypothesis against other features having Poor Facility as the common parent problem hypothesis. Then, Formula 3 was applied to validate Poor Service problem hypothesis. Subsequently, Formula 1 was applied to validate *Customer Churn* problem hypothesis. The resulting goal-problem model is shown in Fig. 9, with check marks to reflect the validation status.

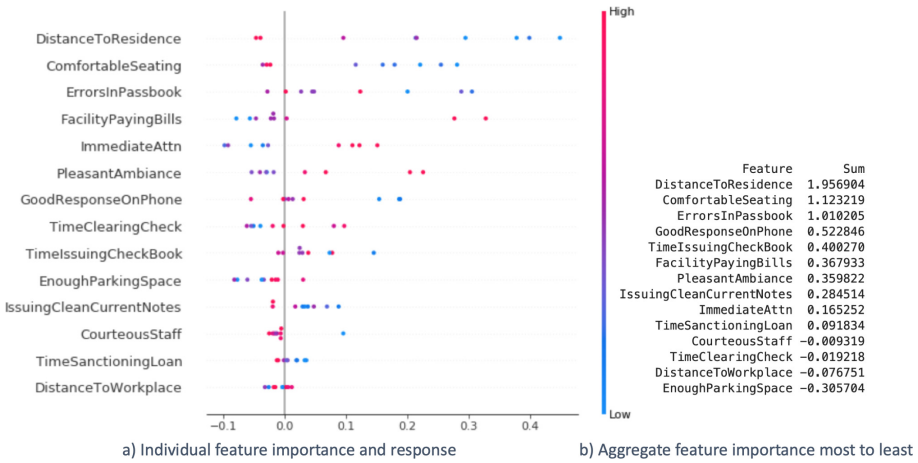


Fig. 8. Features importance for all churners' responses

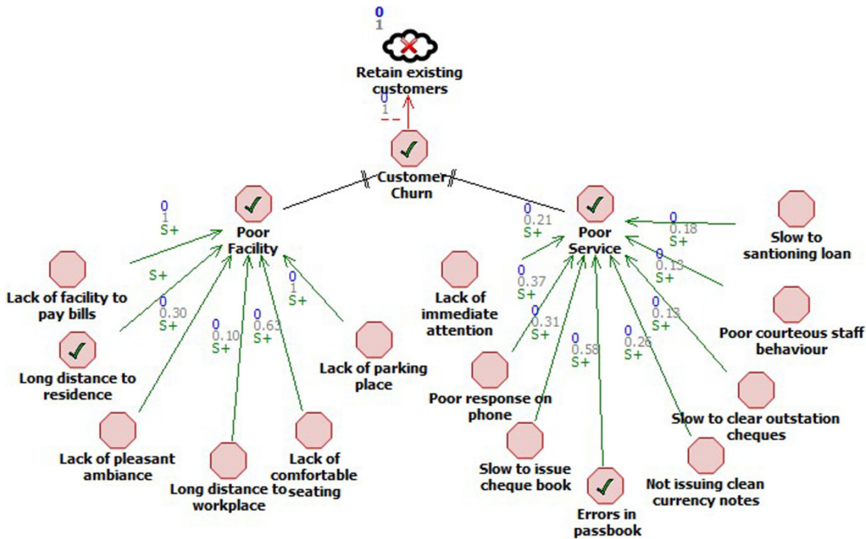


Fig. 9. Validated customer churn problems

4 Related Work and Discussion

We believe this paper is one of the first to propose an end-to-end, explicit and formal approach that provides traceability between business goals and ML. Most data mining and ML projects in practice are often based on informal identification of low-level problems [7] that may not have clear relationships with higher level goals. *Metis* allows ML solutions to be traceable to business at the highest level of business goals and corresponding problems. Using data to validate

goal-oriented models have been proposed in [15] using questionnaire and the statistical hypothesis testing to validate different model elements (e.g, actors, goals, resources) and their relationships (e.g, depends, make, hurt). The statistical method is widely accepted but has been criticized to be difficult to understand [16] and impractical to find evidence in the real-world for some hypothesis to test the null hypothesis [17]. This is especially true in the data-rich Big Data environment, where it is difficult to find evidence for both hypothesis and null hypothesis in the available business data. ML allows organizations to utilize the existing data for hypothesis validation that is grounded by the model accuracy.

Threats to Validity and Limitations. Regarding threats to internal validity, the dataset used in the experiment was highly relevant to the customer churn problem, but it was a small dataset (i.e., 245 records), which could lead to biased results. To reduce this bias, training and testing data were randomly selected and tested with stratification. We also ran several ML algorithms but got similar results. For threats to external validity, as we only applied our approach to a customer churn case, the approach may be too early to be generalized. More experimentation for different domains and datasets is needed.

This paper has presented a promising initial result with some limitations, including 1) inter-feature AND and OR relationships are not currently supported, 2) it is currently unclear whether the result would be consistent across other ML algorithms and model explainability libraries.

5 Conclusions

This paper presents, *Metis*, a novel approach that uses ML to validating hypotheses of business problems that are captured in the context of business goals. *Metis* uses Supervised ML and Model Explainability algorithms to detect feature importance information from the data. Our initial experiment results showed that *Metis* was able to detect the most influential problem root cause when it was not apparent through data analysis. The most influential root cause was then used to validate higher level problem hypotheses using the provided formalization.

Future work to address the identified threats to validity and limitations include 1) conducting additional experiments with larger datasets, 2) testing with additional ML algorithms and explainability libraries, 3) investigating solutions for encoding AND/OR relationships in the datasets for model training or exploring ML algorithms internally for an ability to extract the relationships if captured by the algorithms.

References

1. Asay, M.: 85% of big data projects fail, but your developers can help yours succeed. Techrepublic (2017)

2. LaValle, S., Lesser, E., Shockley, R., Hopkins, M.S., Kruschwitz, N.: Big data, analytics and the path from insights to value. *MIT Sloan Manag. Rev.* **52**(2), 21–32 (2011)
3. NewVantage Partners LLC: Big data and AI executive survey 2020: data-driven business transformation connecting data/AI investment to business outcomes (2020)
4. Zhou, L., Pan, S., Wang, J., Vasilakos, A.V.: Machine learning on big data: opportunities and challenges. *Neurocomputing* **237**, 350–361 (2017)
5. Tatter, P.: gpk: 100 data sets for statistics education. R package v. 1.0. (2013)
6. Supakkul, S., Zhao, L., Chung, L.: GOMA: supporting big data analytics with a goal-oriented approach. In: 2016 IEEE International Congress on Big Data (Big-Data Congress), pp. 149–156, June 2016
7. Berry, M.J., Linoff, G.S.: *Data Mining Techniques: for Marketing, Sales, and Customer Relationship Management*. Wiley, Hoboken (2004)
8. Supakkul, S.: Capturing, organizing, and reusing knowledge of non-functional requirements: an NFR pattern approach. Ph.D. thesis, University of Texas at Dallas (2010)
9. Chung, L., Nixon, B.A., Yu, E., Mylopoulos, J.: *Non-Functional Requirements in Software Engineering*, vol. 5. Springer, Cham (2012)
10. Supakkul, S., Chung, L.: Extending problem frames to deal with stakeholder problems: an agent-and goal-oriented approach. In: *Proceedings of the 2009 ACM Symposium on Applied Computing*, pp. 389–394 (2009)
11. Mylopoulos, J., Chung, L., Nixon, B.: Representing and using nonfunctional requirements: a process-oriented approach. *IEEE Trans. Softw. Eng.* **18**(6), 483–497 (1992)
12. Wang, W., Chen, L., Thirunarayan, K., Sheth, A.P.: Harnessing Twitter ‘big data’ for automatic emotion identification. In: 2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing, pp. 587–592. IEEE (2012)
13. Lundberg, S.M., et al.: From local explanations to global understanding with explainable AI for trees. *Nat. Mach. Intell.* **2**(1), 56–67 (2020)
14. Pedregosa, F., et al.: Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011)
15. Hassine, J., Amyot, D.: A questionnaire-based survey methodology for systematically validating goal-oriented models. *Requirements Eng.* **21**(2), 285–308 (2015). <https://doi.org/10.1007/s00766-015-0221-7>
16. Cohen, J.: The earth is round ($p < .05$). *Am. Psychol.* **49**, 997–1003 (1994)
17. Berkson, J.: Tests of significance considered as evidence. *J. Am. Stat. Assoc.* **37**(219), 325–335 (1942)