

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/352747944>

# Pairing conceptual modeling with machine learning

Article in *Data & Knowledge Engineering* · June 2021

DOI: 10.1016/j.datak.2021.101909

---

CITATIONS

2

READS

73

2 authors, including:



Veda C. Storey

235 PUBLICATIONS 9,627 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Ubiquitous Information Systems [View project](#)



INKASS [View project](#)

# Pairing Conceptual Modeling with Machine Learning

Wolfgang Maass<sup>1,2</sup>, Veda C. Storey<sup>3</sup>

<sup>1</sup>German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany

<sup>2</sup>Saarland University, Saarland Informatics Campus, 66123 Saarbrücken, Germany,  
wolfgang.maass@dfki.de

<sup>3</sup>J. Mack Robinson College of Business, Georgia State University,  
Atlanta, GA 30302-4015 U.S.A.  
vstorey@gsu.edu

## Abstract

Both conceptual modeling and machine learning have long been recognized as important areas of research. With the increasing emphasis on digitizing and processing large amounts of data for business and other applications, it would be helpful to consider how these areas of research can complement each other. To understand how they can be paired, we provide an overview of machine learning foundations and development cycle. We then examine how conceptual modeling can be applied to machine learning and propose a framework for incorporating conceptual modeling into data science projects. The framework is illustrated by applying it to a healthcare application. For the inverse pairing, machine learning can impact conceptual modeling through text and rule mining, as well as knowledge graphs. The pairing of conceptual modeling and machine learning in this way should help lay the foundations for future research.

**Keywords:** Conceptual modeling, machine learning, methodologies and tools, models, database management, Framework for incorporating conceptual modeling into data science projects, artificial intelligence

## 1 Introduction

Machine learning (ML) emerged several decades ago as part of research in Artificial Intelligence (AI) and has recently received a surge in interest due to the increased digitization of data and processes. Machine learning uses data and algorithms to build models that carry out certain tasks without being explicitly programmed [78, 119]. While machine learning focuses on technologies, its application is part of data science. More precisely, data science uses principles, processes, and techniques for understanding phenomena via analysis of data [168]. Although data science could actually be performed manually (pen and pencil, or by calculators), the real power of data science becomes apparent by leveraging machine learning on big data. Data science supports data-driven decision making [29] and is a key driver for digital transformation [208].

Managing data, whether big or traditional, cannot be accomplished solely by humans with their limited cognitive capabilities. Rather, machine learning is important to address many business and societal problems that involve the processing of data. Machine learning has impacted many research fields, including natural sciences, medicine, management, and economics, and even humanities. In contrast to traditional software system development, machine learning does not require programming based on a given design, but rather requires fitting of parameters of generic models on data until the output (predictions, estimates, or results) minimizes or maximizes an objective function [78]. Many types of models have been developed and continue to be applied. Machine learning requires an in-depth understanding of the domains to which machine learning models and algorithms can be applied because data determines the functionality of an information system. Therefore, an assessment is required of whether the training data is representative of the domain. Otherwise, problems might arise that could contribute to biases and mistakes in machine learning models, of which there are well-documented examples, such as automatic parole decisions [175] or accidents with self-driving cars [194].

Although machine learning continues to be an important part of business and society, there are many challenges associated with the progression machine learning so that it is increasingly accessible and useful. At the same time, the development of information systems, of any kind, first requires understanding and representing the real world, which, traditionally, has been the role of conceptual modeling. Emphasis on both big data generation and traditional applications highlights the need to understand, model, and manage data. Over the past decade, research has included the role of conceptual modeling on big data, business, healthcare, and many other applications [59]. Conceptual modeling adds a perspective that starts with strategic business goals and finds translations and abstractions that finally guide software development [160, 7, 134].

The purpose of this paper is to examine how conceptual modeling and machine learning can, and should, be combined to mutually support each other and, in doing so, improve the use of and access to machine learning. We make several contributions. First, the paper provides an overview of the foundations and development cycle of machine learning. Second, we derive a framework for incorporating conceptual modeling into data science projects and demonstrate its use through an application to a specific healthcare project. Third, we examine how machine learning can contribute to conceptual modeling activities. Suggested areas of research are also proposed. This paper shall help researchers and practitioners of conceptual modeling integrate machine learning into their research and operations while also helping data scientists and machine learning experts to use conceptual modeling in their work.

The paper proceeds as follows. Section 2 provides a brief summary of conceptual modeling and its potential use in machine learning. Section 3 presents the foundations of machine learning and its development cycle. Section 4 proposes a framework for incorporating machine learning into data science projects, which it applied to a health care problem. Section 5 highlights the potential impacts of machine learning for conceptual modeling. Section 6 outlines additional research directions for pairing these two fields. Section 7 concludes the paper.

## 2 Conceptual Modeling and Machine Learning Pairing

Conceptual modeling is described as “the activity for *formally* describing some aspect of the physical and social world around us *for the purposes of understanding and communication*” (p. 51) [147]. Conceptual models attempt to capture requirements with the purpose of creating a shared understanding among various people during the design of a project within the boundaries of the application domain or an organization [132]. They help to structure reality by abstracting the relevant aspects of a domain, while ignoring those that are not relevant. A conceptual model formally represents requirements and goals. It is shaped by the perspectives of the cognitive agents whose mental representations it captures. In this way, a conceptual model can serve as a *social artifact* with respect to the need to capture a shared conceptualization of a group [81]. Much research that attempts to understand and characterize research on the development and application of conceptual modeling (e.g., [137, 123, 159, 50, 126, 36]).

The field of conceptual modeling has evolved over the past four decades and has been influenced by many disciplines including programming languages, software engineering, requirements engineering, database systems, ontologies, and philosophy. Conceptual modeling activities have been broadly applied in the development of information systems over a wide range of domains for varied purposes [50]. Activities and topics related to conceptual modeling have evolved over the past four decades [123, 101]. Notably, Jaakkola and Thalheim [104] highlight the importance of modeling, especially with the current emphasis on the development of artificial intelligence (AI) and machine learning (ML) tasks. Other research has also proposed the need for conceptual modeling to support machine learning and, in general, combining conceptual modeling with artificial intelligence [81, 137, 123, 159, 50].

Conceptual models are a “lens” through which humans gain an “intuitive, easy to understand, meaningful, direct and natural mental representation of a domain” [81]. In contrast, machine learning uses data as a “lens” through which it gains internal representations on the regularities

of data taken from a domain ([86, 109, 55]). Pairing conceptual modeling with machine learning contributes to each other by: 1) improving the quality of ML models by using conceptual models during data engineering, model training and model testing; 2) enhancing the interpretability of machine learning models by using conceptual models; and 3) enriching conceptual models by applying ML technologies.

Figure 1 summarizes the relationships among mental models, conceptual models, and machine learning (ML) models. Mental models naturally evolve by acting in domains, whereas conceptual models are shared conceptualizations of mental models [153]. For information system development, conceptual models represent shared conceptualization about a domain by means of conceptual modeling grammars and methods in given contexts [210]. For information systems based on programming approaches, conceptual models are used as requirements for implementations. Database systems are designed and realized according to requirements expressed by conceptual models, such as entity-relationship models [43]. For learning-based information systems, relationships between data and conceptual models and ML models and conceptual models are less obvious [133] (Figure 1).

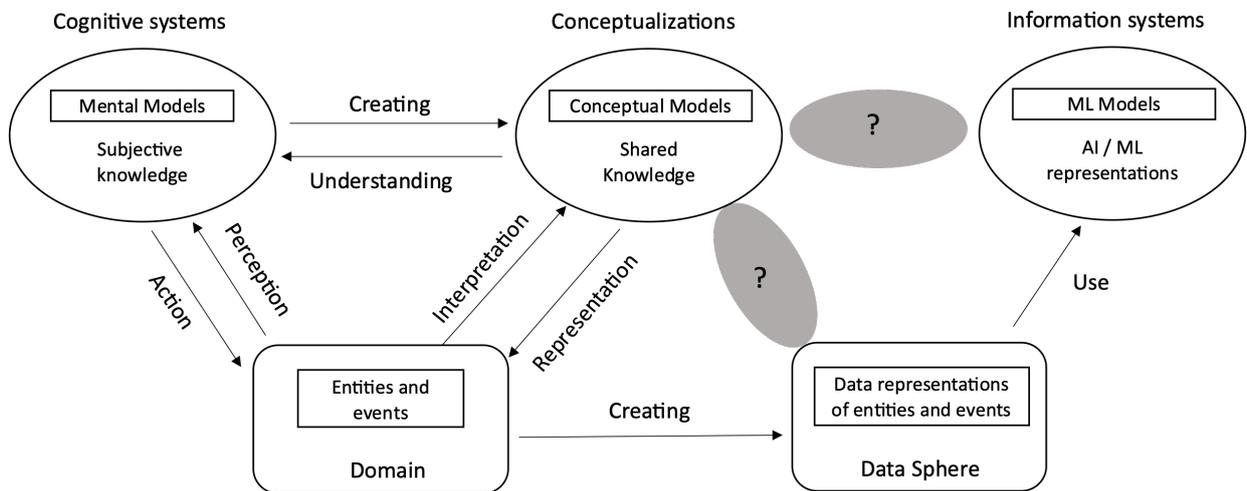


Figure 1: Progression from mental models to Artificial Intelligence / Machine Learning models .

In this paper, we examine these relationships in both directions.

1. How can conceptual modeling support the design and development of machine learning solutions?
2. How can machine learning support the development and evaluation of conceptual models?

Machine learning systems, which are applied to individual datasets, grow exponentially in both size and complexity. Conceptual modeling is instrumental in dealing with complex software development projects. Therefore, the first question to consider is whether conceptual modeling can help structure machine learning projects, create a common understanding, and thereby increase the quality of the resulting machine learning-based system. The second question reverses the direction and asks whether machine learning can provide tools that could support the development of conceptual models. Then, it would also be important to assess how accurately a conceptual model captures an application domain. This is especially challenging for automatic validation, but could take advantage of a data-driven approach to augmenting conceptual modeling.

For information systems that depend on very large datasets and increasingly complexity, machine learning systems, biases of data and uncertainty in decision-making pose threats to trust, especially when the systems are used for recommendations. Pairing machine learning and conceptual

modeling thus becomes an attempt to support *Fair Artificial Intelligence* [13]. This requires using structures and concepts during AI-based software development lifecycles that are stable and meaningful, yet have well-defined semantics, and can be interpreted by humans. We, therefore, propose that the role of conceptual modeling with respect to machine learning is (CM ML):

- *Descriptive*: informs data scientists when developing machine learning systems
- *Computational*: embedded into machine learning implementations.

Inversely, the role of machine learning with respect to conceptual modeling is (CMML):

- *Descriptive*: informs conceptual modelers
- *Computational*: constraints or creates conceptual models.

There are several ways in which conceptual models should be able to inform a data scientist’s work. Conceptual models provide conceptual semantics for concepts and relationships of a domain that govern data used for machine learning tasks. This knowledge can inform data scientists, especially during data engineering but also during model training, and model optimization. If conceptual models are expressed by some computational modeling language, they can be integrated into ML models development procedures. For example, conceptual models can be used to derive constraints on data features that are automatically evaluated during data engineering. That makes both the descriptive and computational perspectives important. Inversely, regularities found by ML models can provide insights for conceptual modelers that can be used for revision and refinement of conceptual semantics and conceptual models. Thus, conceptual models and ML models are independent means for understanding domains of interest. Conceptual models that are consistent with ML models, and vice versa, can increase trustworthiness. Inconsistencies can be indicators of flaws in conceptual models or ML models, but might also facilitate the extraction of novel insights.

### 3 Machine Learning: An Overview

This section provides the foundations necessary to understand machine learning.

#### 3.1 Model Foundations

Machine learning is part of data science projects where the results obtained from software are integrated into an information system, which is called a *ML-based information system*. A data scientist is required to understand statistics and communicate and explain the design of a machine learning system to stakeholders of a data science project [48]. Work on conceptual modeling is similar to that of a storyteller who wants to communicate about the digital and real worlds. With the increasing complexity of ML-based information systems, future data scientists should also understand: how conceptual models can govern data and ML-based information systems; how to carry out conceptual modeling activities; and how to apply knowledge representations techniques. Machine learning involves creating a **model** that is trained on a set of **training data** and is then applied to additional data to make **predictions**. Various **types of models** have been used and researched for machine learning based systems [86]. For example, a predictive model is a function of the form  $f : X^* \rightarrow Y^*$  where  $X^*$  represents a multidimensional input set and  $Y^*$  a multi-dimensional output set. Every parametric model is defined by a set of parameters (aka weights)  $w \in W$  and applied to values of an input vector  $x$ . For instance, a simple linear regression has the form:

$$f(w, x) = w^T x = w_0 + \sum_{i \in \{1, \dots, n\}} w_i x_i \tag{1}$$

Values of  $w$  are derived (or “learned”) from a set of combinations  $(x^i, y^i)$  with  $x^i = (x_1^i, \dots, x_n^i)$  the  $i$ -th input vector from the input dataset  $X$  ( $x^i \in X$ ), where  $y^i$  is  $i$ -th output from the output

dataset  $Y(y^i \in Y)$ . Each  $x_j^i \in X$  represents an attribute of an entity of interest, and is called a *feature*, e.g., age of a person or pixel of an image. For tabular data, a feature is a column. The task of *supervised learning* is to determine the weight vector  $w = (w_1, \dots, w_n)^T$  in such a way that the difference (loss)  $L(f(w, x), y)$  between an estimation of output vector  $\hat{y} = f(w, x)$  given a new input vector  $x$  and the actual value of  $y$  (often called the *ground truth*<sup>1</sup>) is minimized. Here, a loss function  $L$  is used to measure how accurate a model is with weights  $w$ . An often-used loss function is square loss:

$$L(f(w, x), y) = \frac{1}{2} \|f(w, x) - y\|_2^2 = \frac{1}{2} \sum_{i \in 0 \dots n} (f(w_i, x_i) - y_i)^2$$

A loss function is a central element for machine learning algorithms because it is used for model optimization; that is, loss minimization. A loss function is also called an error function, cost function, or objective function. The latter name emphasizes its use as a criterion for optimizing a model [78]. Some ML models are designed to make finding optima tractable, such as linear regression and support vector machines. For others, such as deep learning models, finding a global maximum or minimum is intractable [91]. Various heuristics, such as momentum or randomization, are used in combination with gradient descent algorithms to avoid being too restricted to a local minimum. Features of  $X$  with small weights relative to other weights contribute little to the outcome. Sometimes prediction accuracy is improved by setting some weights to zero (cf. [86]). Some methods start with the simple model consisting of the bias weight  $w_0$  and only add dimensions with the highest impact on the prediction until the loss stabilizes.

### Model shrinkage methods

Several reasons exist for reducing the complexity of a ML model. One reason is that it is important to identify which input variables are most important and have the strongest impact on predictions. This can be achieved by shrinking weights for variables with small to zero impact. Another reason is that the number of input variables is larger than the sample size which generally results in perfect model fit with no noise. In this case, shrinking a model contributes to model generalization.

A model that uses all parameters by setting  $w_i$  to zero has the least impact on model accuracy until the loss is stabilizing (cf. [86]). Model shrinkage methods simultaneously adjust all weights  $w_i$  by optimizing a risk function  $R$  that is defined by the loss function  $L(f(x), y)$  and an additional function  $J(f(x))$  that defines a penalty on model complexity. Therefore, a risk function  $R(f)$  defines a tradeoff between minimizing loss and model complexity. The aim of a learning algorithm is to find a function  $f : X \rightarrow Y$  among a class of functions  $F$  for which the risk  $R(f)$  is minimized:

$$f^* = \operatorname{argmin}_{f \in F} R(f) \tag{2}$$

where  $f^*$  gives the best expected performance for loss  $L(f(x, w), y)$  over the distribution of any function in  $F$ . Since the data generating distribution  $P(X, Y)$  is unknown, the risk  $R(f(w, x))$  cannot be computed. Therefore, minimizing the risk over a training dataset  $X$  drawn from  $P(X, Y)$  is formalized as:

$$R_{emp}(f(x, w)) = \frac{1}{n} \sum_{i \in 1 \dots n} L(f(x_i, w), y_i) + \lambda J(f(x_i, w)) \tag{3}$$

With a standard definition of  $J(f(x, w))$  as a norm:

$$J(f(x, w)) = \sum_{i \in 1 \dots m} |w_i|^p \text{ and } \lambda \text{ regularization parameter.}$$

L1 ( $p = 1$ ) and L2 ( $p = 2$ ) norms are often used.  $J(f(x, w))$  is called a *regularization function* because it controls model complexity. Using an L1 norm is called a *lasso regression*, whereas using a L2 norm is called a *ridge regression*. Several alternative regularization functions exist, such

---

<sup>1</sup>Ground truth  $Y$  is subject to uncertainties about its actual truthfulness. It only states that, for the modeling task, it is assumed that each  $y \in Y$  is true, with a probability  $\alpha$  that this assumption is false. Often  $\alpha$  is unknown.

as *elastic net* [230] and *least angle regression* [56]. In contrast to reducing model complexity by principal components regression, model parameters restrict direct correspondence to input features. Thus, models keep their ability to directly support model explanations.

### 3.1.1 Prediction types

Classification and regression supervised learning tasks are similar in that they both have numerical input but differ in the type of target variable being predicted. The numerical form is transformed based on a dependent variable that estimates output  $y$  based on input vector  $x$ . Different types of target variable that determine different learning tasks are explained in the following sections.

#### Classification

For classification,  $f(x)$  projects input  $x$  on categorical values given by a discrete dimension of  $y$ . The following types of classifications are used:

1. *Binary Classification*: For every input  $x$ , there are only two possible output values for  $y$ .
2. *Multi-class Classification*: For every input  $x$ , there are more than two possible output values  $y$ .
3. *Multi-label Classification*: For every input  $x$ , there are more than two target values with each input sample associated with one or more labels, i.e.  $y \in Y^K$  with  $K$  being the set of all class labels.

An example for a binary classification is a binary logistic regression used for finding a function  $f(x)$  that separates positive from negative illness cases.

$$p(x) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 X_1)}} \quad (4)$$

The logistic function provides the probability that the estimate of class  $k$  is correct given input  $x$  with  $k \in \{0, 1\}$ . Figure 2 shows that  $p(x)$  perfectly separates people younger than 25 years, and older than 26, with several misclassifications made in-between. The logistic function separates two topological spaces with associated labels  $k \in \{0, 1\}$  and selects the class with the highest probability.

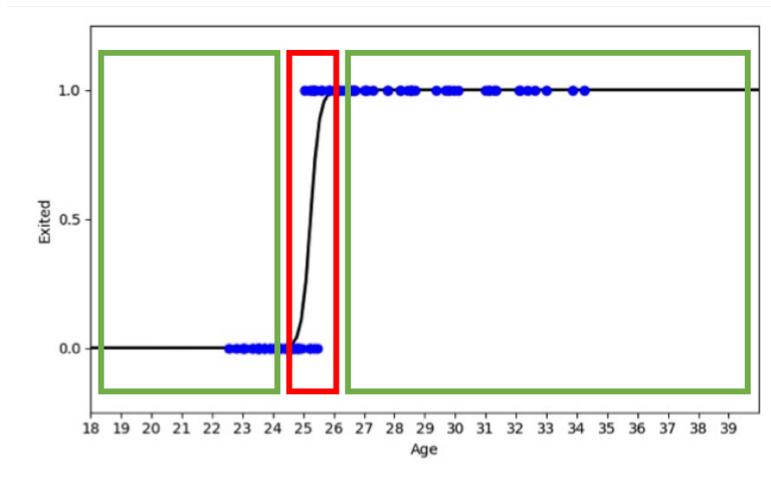


Figure 2: Classification by using a binary logistic regression function; orange line: separation line of two classes.

## Regression

For regression,  $f(x)$  projects input  $x$  on numerical values given by the numerical dimension of  $y$ . Often the dimension of  $y$  is real values ( $y \in \mathbb{R}$ ). Examples are estimates for stock prices or number of people swimming in a pool per day.

### 3.1.2 Model selection

Machine learning has a long history, dating as far back as the beginning of AI when the 1956 Dartmouth Summer Research Project contained *neuron nets* as well as other topics. Because machine learning is built on mathematical statistics, the terms *statistical learning* and *machine learning* are often used synonymously [109]. Given this long history, it is not surprising that a large variety of machine learning model types are available, which can be generally classified as follows.

- *Supervised learning*: learning a function that maps an input to an output based on example input-output pairs [178].
- *Unsupervised learning*: learning a function that maps an input to an output without prior labels available for output variables. Example models are k-nearest neighbors (KNN), k-means and clustering models.
- *Reinforcement learning*: learning a function on how to take actions in an environment in order to maximize the notion of cumulative reward [200]. Examples are Q-learning [200], Monte-Carlo [200] or DQN [144].

Recent approaches for unsupervised learning make use of reinforcement learning models so that the difference between both classes becomes less strict [188]. In the following, we focus on supervised learning, the most common form.

## Supervised Learning

Commonly used model types for supervised learning are decision trees, ensemble learning and neural networks, each of which is discussed briefly below.

### Decision Trees

Decision trees are widely used machine learning models. Particularly important are additive combinations of so-called *weak learners*, such as boosting (AdaBoost [68] and XGBoost [44]). An early decision tree model is the Classification and Regression Trees (CART) model [27]. In the example (cf. Figure 3), a patient who was infected before week 11.5, in a province labeled smaller than 2.5 (i.e. provinces Busan, Chungcheongbuk-do and Chungcheongnam-do) is classified as being released (data source: [110]).

Decision trees are based on recursive decomposition mechanisms that optimize on a separation step. For classification, in each node with a data set  $N_i$ , the dimension  $d \in D$  is selected that best separates data in cell  $N_i$  according to a loss function  $L(\cdot)$ . The most basic approach is to systematically consider one attribute after another, use any value of the dataset as threshold, and assess the loss  $L(\cdot)$ . Loss functions are used for measuring the *impurity* of resulting nodes after a split. For categorization, a node has a low impurity if it contains many samples from one category and few from others.

Typical loss functions [86] are defined by using class probabilities  $\hat{p}_{mk} = 1/N_m \sum_{x_i \in R_m} I(y_i = k)$ , with a node  $m$  holding a region  $R_m$  with  $N_m$  data samples [86]:

- *Gini index*:  $\sum_{k \neq k'}^K \hat{p}_{mk} \hat{p}_{mk'} = \sum_{i=1}^K \hat{p}_{mk} (1 - \hat{p}_{mk})$  for all classes  $k \in K$
- *Cross-entropy*:  $-\sum_{i=1}^K \hat{p}_{mk} \log(\hat{p}_{mk})$  for all classes  $k \in K$ .

Node  $N_0$  contains 2538 patients of which 56 are labeled “deceased”, 1054 as “isolated” and “1428” as “released”. After reformulation of the equation of the Gini index, impurity of node  $N_0$  calculates as follows:  $G(N_0) = 1 - \left(\frac{56}{2538}\right)^2 - \left(\frac{1054}{2538}\right)^2 - \left(\frac{1428}{2538}\right)^2 = 0.51$ .

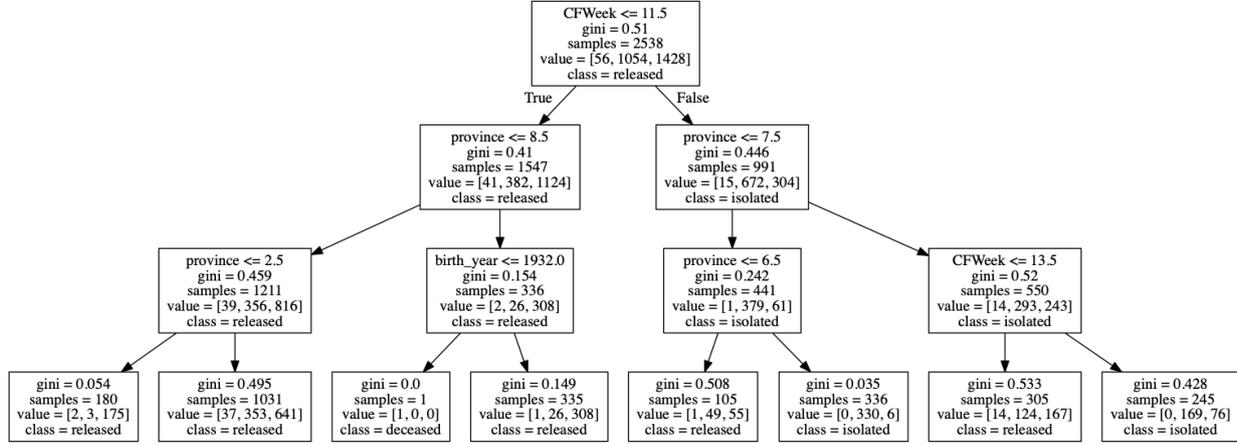


Figure 3: Korean Covid-19 disease data.

The decision criterium for a split is the impurity of resulting nodes that is calculated as weighted Gini impurity  $I_G$ . For the example,  $I_G(N_0) = \frac{1547}{2538} * 0.41 + \frac{992}{2538} * 0.446 = 0.424$ . By calculation of  $I_G$ , a feature and a threshold are selected resulting in the smallest lost, ie. smallest  $I_G$ .

Figure 3 shows a decision tree derived from data of COVID-19 infection cases in South Korea trained for three categories ("deceased", "isolated", "released") based on eleven input variables, including year of birth, sex, age and confirmed day and month of infection.

### Ensemble Learning

Ensemble learning integrates base learners, such as decision trees, into complex models [179]. The advantage of ensemble learning is in improved accuracy for the additional cost of increased latency due to testing a series of models at runtime and lack of interpretability. Instead of independent decision trees, *boosting* iteratively integrates decision trees by focusing on mis-classified samples. For instance, *AdaBoost* is a popular boosting model that only uses decision trees with one node and two leaves, called *stumps* ([67]). In each step, all samples of a dataset are associated with an equal weight  $w_i = \frac{1}{\text{number of samples}}$ . The attribute with the smallest Gini index is selected for the next stump. For each stump, the percentage of errors is computed:  $\text{error}_m$  is the sum of weights  $w_i$  misclassified divided by all weights  $w_i$ .

Additionally, each stump  $m$  has a weight  $\alpha_m = \log\left(\frac{1-\text{error}_m}{\text{error}_m}\right)$  that represents the importance of a stump. Small  $\alpha_m$  means that a stump has little impact on the final result. From a stump, weights for each sample are updated according to a correct classification result: misclassified samples:  $w_i = w_i * e^{\alpha_m}$  and correct classified samples:  $w_i = w_i * e^{-\alpha_m}$ . Then, weights for misclassified samples are increased and decreased for correctly classified samples. A new dataset is determined by drawing from samples with replacement by using normalized weights as probabilities; that is, misclassified samples have a higher chance of getting drawn multiple times. This new dataset resets all weights to equal weights again and the process restarts until the maximum number of iterations  $M$  is reached.

*Gradient boosting* extends the idea of AdaBoost by allowing larger trees with fixed size than stumps [68]. Gradient boosting starts with the mean  $\gamma$  of the dependent variable that minimizes the loss function  $L(y_i, \gamma) = \frac{1}{2}(y_i - \bar{y})^2$ . Next, the so-called *pseudo residuals* of each sample are determined which are the differences of the sample value and mean  $\bar{y}$ . Formally, the pseudo residual is the derivative of the squared loss function  $L$  which is  $-\frac{\partial L(y_i, f(x_i))}{\partial f(x_i)} = f(x_i) - y_i$  with  $f(x_i) = \gamma$  [86].

Instead of growing a tree on the dependent variable, it is grown on predicting the pseudo residuals. Thus, previous sample errors are corrected to a certain extent. For each leaf of terminal node  $j$ , a value  $\gamma_{jm}$  is determined that is the minimum of the loss function over all samples in this terminal node. The final estimate is the summation of all  $\gamma_{jk}$  of regions where the input is associated with all trees of the gradient boosting model. When computing the estimate, the contribution of each tree is scaled by a learning rate  $\eta$  for each tree result that controls for overfitting. Gradient boosting also supports classification by converting class labels into probabilities by applying the logistic function on the log of relative occurrence of classes in the dataset. The pseudo residuals are the differences between data values; the mean probability minimizes the negative log likelihood function used as loss function [68].

*XGBoost* is an extension of gradient boosting and optimization on the residuals [44]. It uses an alternative for calculating the gain of a split based on the squared sum of residuals. A subtree is pruned if the gain is smaller than a threshold value  $\gamma$ . Regularization is achieved by another hyperparameter  $\lambda$  used for decreasing gain values and, thus, the size of trees. Hence, *XGBoost* has two hyperparameters for controlling model complexity. Estimations are calculated such as gradient boosting controlled by a learning rate  $\epsilon$ . *XGBoost* is designed for parallelization that is useful for large datasets. For classification, gain is computed by sample probabilities, similar to gradient boosting.

### **Neural Networks**

Neural networks represent a general class of learning models that can be adapted to different problems. For instance, convolutional neural networks (CNN) for visual computing (e.g., ResNet [87]) and transformer models for natural language processing (e.g., BERT [52]). Neural networks make use of parallel execution of weak learners and, thus, train universal approximators of any function given sufficient data and resources [96]. Various forms of neural nets, such as recurrent neural nets (RNN) are *Turing complete* [184]. The proof of Turing completeness of more sophisticated model types, such as Transformer and Neural GPU [171], is built on foundational mechanisms; that is, residual connections for Transformer and gates for Neural GPUs. These results provide evidence for the hypothesis that all, non-trivial machine learning model types are Turing complete. Decision trees and ensemble models based on decision trees lack the concept of loops and memory, so they are not Turing complete. However, the class of all machine learning model types as a whole, is Turing complete. Therefore, defining a model architecture is not a question of the complexity classes of computable functions, but of performance. A single node, called a neuron, consists of an application of a non-parametric non-linear function  $\sigma$ , called an *activation function*, on a linear function with weights  $w_j$ :

$$\hat{y} = \sigma\left(\sum_{j=1}^p w_j x_j + b\right) \quad (5)$$

A neural network model is trained by fitting weights so that a corresponding loss function is minimized. Similar to gradient boosting, optimization of the loss function means to minimize residuals. Optimization of loss with respect to weights of the neural network is typically performed using a form of a gradient descent, such as gradient descent with momentum, or more sophisticated variants, such as the Adam optimization algorithm [111].

For classification tasks, a softmax function transforms output of the final activation layer to valid probabilities for each class, for input vector  $x$ , weight matrix  $W$ , and bias vector  $b$  into probabilities. The output vector of the softmax layer is a vector  $a$  with probabilities for each class.

Recurrent neural networks allow information produced by a neuron to be used as input together with inputs to the neural networks at the next time step [64]. Long Short-Term Memory (LSTM) neural networks are RNNs, with a complex structure combining various activation functions. LSTM include memory cells for keeping gradient information that can be fed back into the neuron's activation functions [92]. Various network topologies exist. For example, modifiable self-connections decide whether to overwrite a memory cell, retrieve it, or retain it for the next time

step [72]. LSTM are widely used for natural language processing and other tasks with time-variant data, even over long periods of time [180].

Convolutional neural networks (CNN) are variants of neural networks that specializes on multi-dimensional data<sup>2</sup> that supports processing of image datasets [119]. Convolutional layers transform input data filter matrices (kernels). Low-level kernels detect simple entities, such as edges whereas higher order kernels are sensitive to more complex visual structures [116]. Thus, CNN filters and transforms data with the help of massive application of low-level mechanisms, such as max pooling, padding, and striding. Stacking layers of large amounts of neurons enables complex visual computing operations with high quality, such as object recognition and object tracking in videos.

### Unsupervised learning

In unsupervised learning, teaching a model can be accomplished without ground truth data of dependent variables  $y$ . Hence, loss functions cannot be used for assessing model quality, but need to be replaced by heuristics or other quality metrics. The focus is only on direct inference of properties of the probability density function of dataset  $X$  [86]. From complex datasets, simpler approximation models are derived by using principal components models, multidimensional scaling, self-organizing maps, and principal curves. Other classes for extracting model abstractions are clustering and association rules.

Cluster models are centered around the concept of proximity; that is, the distance  $d(x_i, x_j)$  between two instances of a dataset  $X$ . The overall distance  $D$  with normalized weights  $w_k$  for each dimension is generally formulated as follows:  $D(x_i, x_j) = \sum_{k=1}^p w_k * d_j(x_{ik}, x_{jk})$  [86]. Distance function  $d$  can be instantiated by the Minkowski distance with different settings for parameter  $p$  ( $p = 1$ : Manhattan distance,  $p = 2$ : Euclidean distance) or self-defined functions. Clusters are abstractions of probability density function  $P(X)$  and mediate interpretation of datapoints by domain experts that goes beyond model properties. Unknown datapoints are directly associated with clusters and, thus, inherit cluster interpretations.

K-means is a simple model often used for clustering. It uses an iterative procedure that stop if a threshold  $\varepsilon$  is undercut or if a maximum number of iterations is reached. For a given  $k$ , K-means determines the distance between all datapoint  $x \in X$  and datapoint  $k_i \in K \subseteq V$  with  $V$  the  $p$ -dimensional vector space of  $X$ ,  $|K| = k$  and associates label  $k_j$  to a datapoint  $x_i$  if distance  $d(k_j, x_i)$  is the smallest for all  $k \in K$ . Then, for each label,  $k_j$  is replaced by  $k'_j \in V$  that is the center of all datapoints  $x_i$  with label  $k_j$ . This process is repeated until  $\sum_{i,j}^k d(k_i, k_j) < \varepsilon$  or a maximum number of iterations is reached. Datapoints  $x_i, x_j$  connected to a cluster  $k_s$  ( $x_i, x_j \in X(k_s)$ ) are stronger connected than with any  $x_m \in X(k_r)$  with  $s \neq r$ .

Hierarchical clustering iteratively merges clusters based on a group distance metric. For classification tasks, loss functions are defined on the percentage of correct and incorrect classifications and how these change with different  $k$  values.

### Generative models

While standard approaches of unsupervised learning models attempt to find unknown labels for instances, generative models are used to learn underlying distributions of given data that can be used for sampling from these models; that is, the generation of novel instances indistinguishable from input data. For instance, a *generative adversarial neural network* (GAN) [77] learns a generator's distribution  $p_G$  over data  $x$  by a neural network  $G(z)$ , with  $z$  taken from an input noise variable of  $p_z(z)$ , and uses a discriminator  $D(x)$  neural network for making decisions as to whether the examples are real or fake. The goal of the generator is to generate artificial examples from the learned distribution that the discriminator cannot distinguish from real examples, i.e.  $D(G(z)) \approx 1$  with  $z$  random noise. Models  $G$  and  $D$  compete with each other while increasingly becoming better; that is,  $G$  generates more realistic output and  $D$  gets better at discriminating fake from real. The

---

<sup>2</sup>CNN works on 2D matrices but can handle several of those matrices, called *channels*. Therefore, it can handle, for instance, RGB images with one channel for each color. Standard ML models operate on 1D vectors; CNN operate on 2D matrices. With data streams and tensors this becomes multi-dimensional.

situation between  $G$  and  $D$  is modeled as a game-theoretic min-max game [77]:

$$\min_{\theta_G} \max_{\theta_D} [E_{x \sim p_{data}} \log D_{\theta_D}(x) - E_{z \sim p(z)} \log (1 - D_{\theta_D}(G_{\theta_G}(z)))]$$

with samples from the distribution of the real world  $p_{data}$ . The discriminator tries to maximize correct classification of real versus generated data samples by training model weights  $\theta_D$  of the discriminator neural network. On the other hand, the generator tries to minimize the success of the discriminator by training a model weights  $\theta_G$  of the generator neural network so that  $D_{\theta_D}(G_{\theta_G}(z))$  becomes close to 1. GANs are used for generating 2D images [77], 3D models [18], music [226] and even support arithmetical operations [172].

*Autoencoders* consist of one neural network used for encoding input data  $x$  into  $z(x)$  and one that is used for decoding results of the encoder  $\hat{x} = d(z)$ . The goal is to encode only the information in  $z$  that is critical for reconstructing  $\hat{x}$  with the smallest error. The smaller  $z(\cdot)$ , the less information used. This shows the resemblance to principal component analysis (PCA) while, unlike PCA, a non-linear mapping between compressed representation and the original representation can be achieved. Autoencoder architectures are effective for reduction of dimensionalities. Variational autoencoders (VAEs) [112] are used for content generation by using distributions of latent space  $z$ , i.e.  $p(z|x)$  instead of  $z(x)$ .<sup>3</sup> Complex distributions are approximated by variational inference that defines a set of Gaussian distributions  $\mathcal{N}(g(x), h(x))$ , with  $g$  mean and  $h$  variance dependent on  $x$ . Mean and variance functions are optimized by minimizing the Kullback-Leibler divergence between the approximation and the target distribution  $p(z|x)$ .

## Reinforcement learning

Reinforcement learning combines machine learning with agent-based systems [200]. A reinforcement learning model tries to maximize a reward function, such as winning a game (Silver et al. 2017) or a robot walking up a slope [70]. The machine learning environment, including model and training mechanism, takes the role of an agent that predicts the best action according to an internal strategy without internal representation of the environment (model-free reinforcement learning). An action is performed in a specific situation and the agent receives a reward for this action (cf. Figure 4). This procedure ends when a final situation is reached. The goal of the agent is to maximize reward. The environment is an abstraction of the world in which an agent operates; that is, it is a model of the world. The world could be digital; for example, artificial, as in a chess game, or fully realistic and physical such as used for research on robotics. An agent learns by performing actions in an environment and receiving feedback in the form of rewards. By adjusting the model, the agent tries to find means for maximizing the total reward.

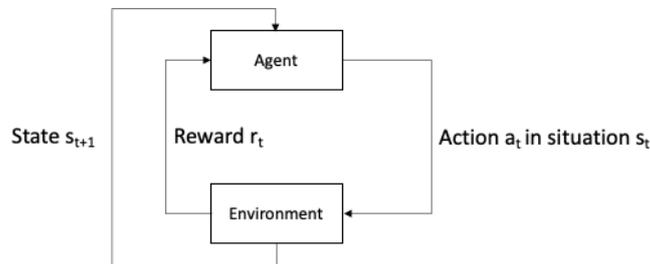


Figure 4: Reinforcement learning.

A reinforcement learning system is based on the concept of a Markov decision process. A state of a Markov decision process completely characterizes the state of the world under investigation

<sup>3</sup>Both VAEs and GANs result in generative models, although neither is a subset of the other. VAEs model the data generating distribution explicitly as an infinite mixture of Gaussians, whereas GANs do so implicitly.

(Markov property) by  $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathbb{P}, \gamma)$  with  $\mathcal{S}$  the set of possible states,  $\mathcal{A}$  set of possible actions,  $\mathcal{R}$  distribution of reward, i.e. (state, action) pairs,  $\mathbb{P}$  transition probability on actions in the environment and  $\gamma$  discount factor. A policy  $\Pi$  is a function from  $\mathcal{S}$  to  $\mathcal{A}$  that specifies which action is best to take in a given state. The objective is to find a policy  $\Pi^*$  that maximizes the cumulative discounted reward  $\sum_{t>0} \gamma^t r_t$  by taking a series of future actions.

In contrast to supervised learning with minimizing loss, reinforcement learning tries to find a sequence of actions and situations that maximizes the sum of rewards:  $\Pi^* = \arg \max_{\Pi} \mathbb{E} \left[ \sum_{t \geq 0} \gamma^t r_t | \Pi \right]$  with  $s_0 \sim p(s_0)$ ,  $a_t \sim \pi(\cdot | s_t)$ ,  $s_{t+1} \sim p(\cdot | s_t, a_t)$ . The value of a situation  $s$  is calculated by the sum of rewards the agent expects under the given policy  $\pi$ :  $V^{\Pi}(s) = \mathbb{E} \left[ \sum_{t \geq 0} \gamma^t r_t | s_0 = s, \pi \right]$  and the value of action  $a$  performed in situation  $s$  under a policy  $\pi$  given by the Q-value function by accumulating the expected discounted reward  $r$ :  $Q^{\Pi}(s, a) = \mathbb{E} \left[ \sum_{t \geq 0} \gamma^t r_t | s_0 = s, a_0 = a, \pi \right]$ . From this objective, the agent tries to find a policy that maximizes  $Q$ :  $Q^*$ . A standard approach optimizes  $Q^*$  by satisfying the *Bellman equation*:  $Q^*(s, a) = r + \gamma \max_{a'} (Q^*(s', a'))$  in any situation  $s$  and using  $Q^*$  for defining policy  $\pi$ :  $\pi(s) = \max_a (Q^*(s, a))$ . For problems with small action set  $A$  and small state set  $S$ ,  $Q^*$  is defined by a look-up table, whereas, for complex environments, such as robotics,  $Q^*$  is approximated by neural networks (deep q-learning):  $Q(s, a, \theta) \approx Q^*(s, a)$  with parameters  $\theta$ . The difference to standard neural networks in supervised learning is that labels  $y$  are unknown. The challenge is to use the Bellman equation with the current function approximation  $Q(\cdot)$  for calculating labels  $y$ ; for example, Baron Munchausen getting himself out of the mud by pulling on his hair (Münchhausen trilemma)<sup>4</sup>. By using a squared loss function  $L_i(\theta_i)$ , the difference between  $y$  and the estimate of the neural network is used to determine the loss. The gradient of the loss (with holding  $Q(\cdot)$  fixed) is used for improving  $Q(\cdot)$  so that the next label  $y$  is closer to  $Q^*(\cdot)$ , assuming convexity.

### 3.1.3 Model training

Parametric machine learning models iteratively adjust model parameters so that the error on estimates for unknown input data is minimized. Examples include: linear regression, support vector machines, decision trees and neural networks. This iterative adjustment process is called *model training*, which searches through combinations of weights and model architectures, i.e. the configuration of components that will be fitted to data, such as the number, type, sequence and size of layers in a neural network. Model training compresses datasets into model parameters. In practice, model sizes can range from few bytes to over 100GB. GPT-3 [28], for example, has approximately 325GB with 175B parameters of the language generation model represented with 16Bit precision floating point number [49]. The model fitting process creates a data generating function for generative tasks, that replicates output data similar to the unknown function underlying the data used for training. For classification, the function clearly discriminates. Fitting methods require at least as many training samples as there are parameters to be trained.<sup>5</sup>

Several methods are used for optimizing model performance that involve either growing a model (e.g., gradient boosting) or adjusting weights of a fixed model (e.g., neural networks). In all cases, model optimization is guided by its associated objective function. Linear regression uses ordinary least square method. Maximum likelihood is a well-known method for finding the optimum values for the parameters by maximizing a negative log-likelihood function derived from the training data [22]. This is also used for logistic regression models. Decision trees and ensemble learning models, such as AdaBoost, use an additive expansion approach that adds weak learners for reduction of prediction errors [86]. For neural networks, the loss function is optimized by using gradient descent where the gradient of the loss function on a dataset at hand is calculated with *backpropagation*

<sup>4</sup>[https://en.wikipedia.org/wiki/M%C3%BCnchhausen\\_trilemma](https://en.wikipedia.org/wiki/M%C3%BCnchhausen_trilemma)

<sup>5</sup>This, at least, was the case for traditional machine learning, although recent research in neural networks suggests this may be false with shot learning and zero shot learning looking promising for some models.

[86]. With gradient descent, each weight is slightly adjusted along the negative gradient according to its contribution to the result to minimize the loss function.

### Loss function

Beside mean squared loss error (MSE), several other loss functions exist for regression tasks, such as mean absolute error (MAE) and a combination of both (cf. Figure 5), which is called the Huber loss [99]. For binary and multiclass classification, cross entropy and corresponding Kullback-Leibler divergence are used that assess the difference in entropy of training and testing data and the entropy of predictions (figure 5). Other loss functions are commonly used, such as hinge loss or exponential loss [86].

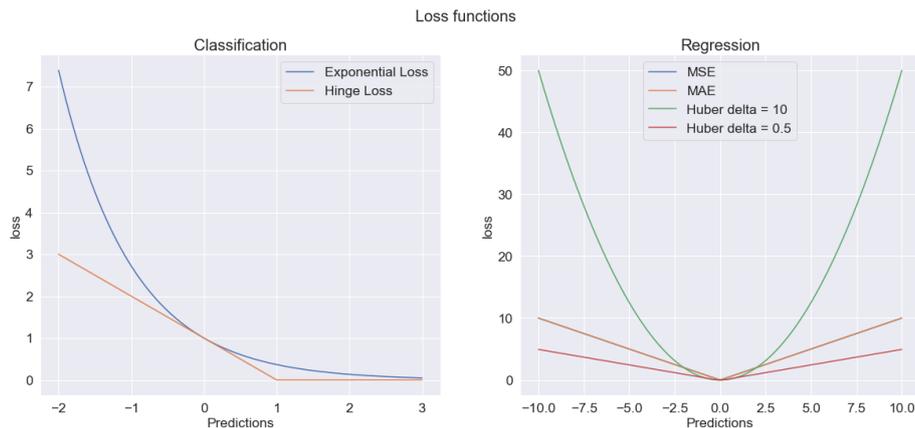


Figure 5: Example loss functions for classification and regression.

The goal of model training is to find a model parameter that minimizes the error of the selected loss function  $L(f)$ :

$$\min_{w \in W} \sum_{i \in 1 \dots n} L(f(w, x), y_i) \quad (6)$$

### Gradient descent procedure

Instead of progressing through the entire search space of weights  $W$ , gradient descent is used for finding local minima for a given weight vector  $w$  by calculating partial derivatives of  $L$  on  $w$ . Because  $L$  is to be minimized, the negative of the derivatives is used. For updating the weight vector  $W$ , a scalar step size  $\eta$  is used that determines the size of adjustment. If stepsize  $\eta$  is set too large, the risk of destabilizing the optimization procedure increases, whereas a step size that is too small assumes the risk of slowing down training and reaching the maximum number of iterations too early, before reaching the minimum (For details cf. section 4.3 in [78]). This equation shows the dependency of the training algorithm on the definition of the loss function.

$$w := w - \eta \nabla_w L(f(w, x), y)$$

Function  $f(w, x)$  contains the logic for computing an estimate  $\hat{y}$ . The difference between  $\hat{y}$  and  $y$  is the core of a loss function in supervised learning, because the loss is dependent on  $f(\cdot)$ , which, in turn, is dependent on parameters  $w$  (or  $\theta$ ). Partial derivatives of  $L(\cdot)$  on  $w$  adjust weights in the direction of a minimum of loss function  $L(\cdot)$ .

### Model overfitting, bias and variance

Not all machine learning models have the same capacity for capturing the signals embedded in a dataset. Linear models can only capture linear functions, whereas neural networks generally capture non-linear functions. However, learning algorithms that can produce models that can learn arbitrary relationships between inputs and outputs, so they might adapt to idiosyncratic data and outliers, and hence not generalize to new data. This trade-off is characterized by *bias* and *variance*. A model that is too simple for capturing the complexity of a function underlying a dataset has high bias (cf. linear function  $f_0$  in Figure 6a)); that is, on average, it shows high error. A trivial model (cf. function  $f_2$  in Figure 6a)) shows no error and has a bias of 0. Function  $f_1$  is in the middle between  $f_0$  and  $f_2$  with respect to bias. When applied to unseen data (Figure 6b)), function  $f_2$  shows a large error whereas function  $f_1$  is better. However, function  $f_0$  is not able to estimate new data well. The degree by which models perform worse on testing data than on training data is called the *variance of a model*. Function  $f_2$  shows low bias, but high variance. Function  $f_0$  has high bias and low variance (because it performs poorly on training and testing data) whereas  $f_1$  has low bias and lower variance than  $f_2$ . Function  $f_1$  generalizes better than  $f_2$  because it has a lower loss on new data not present in the training data. Alternatively,  $f_1$  is suffering less from *overfitting* to training data than  $f_2$ . Function  $f_0$  is *underfitting* the data due to high bias and, thus, is not being specific enough to capture the underlying function. In practice, the greater the complexity of a model, the greater the tendency to overfit. This is accounted for by adding a regularization function  $J(\cdot)$  to the risk function  $R(\cdot)$  that adds a penalty to more complex models. This is shown in Figure 6.

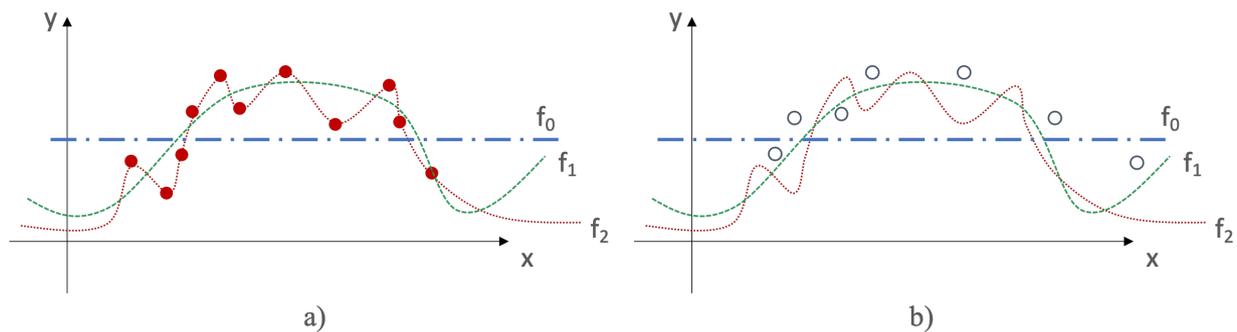


Figure 6: Bias and variance; circles in b) indicate unseen data used for testing the models.

In general, model search trees is a process to optimize the search for a model that minimizes the loss on training and unknown testing data via analyzing underfitting and overfitting behavior (cf. Figure 7).

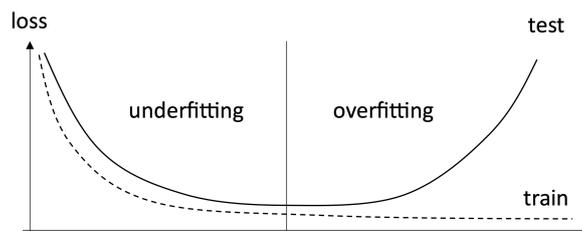


Figure 7: Overfitting and underfitting.

Generic model Kurgan and Musilek [117]	Shmueli and Koppius [183]	Chambers and Dinsmore [39]	Goodfellow et al. [78]	Data science development process
Application domain understanding	Goal definition	Defines business needs	Determination of goals	Problem understanding
Data understanding	Data collection and study design	Build analysis data set	Establish a working end-to-end pipeline	Data collection
Data preparation and identification of data mining technology	Data preparation			Data engineering
	Exploratory data analysis			
	Choice of variables			
Data mining	Choice of potential methods	Build predictive model	Instrument the system well to determine bottlenecks in performance	Model training
Evaluation	Evaluation, validation and model selection		Repeatedly make incremental changes such as gathering new data, adjusting hyperparameters, or changing algorithms	Model optimization
Knowledge consolidation and deployment	Model use and reporting	Deploy predictive model	<i>Beyond discussion</i>	<i>Model Integration</i>
<i>Beyond discussion</i>	<i>Beyond discussion</i>	<i>Beyond discussion</i>	<i>Beyond discussion</i>	Analytical decision making

Table 1: Data science development processes.

### 3.2 Data Science Development Cycle

The development of information systems based on machine learning is still progressing. Major platform providers have published their own processes, such as Google’s *Train-Evaluate-Tune-Deploy* workflow<sup>6</sup> or Amazon’s *Build-Train-Deploy* model<sup>7</sup>. By analyzing five development models, including CRISP-DM (cross-industry standard process for data mining) [217], we identify six phases in the machine learning development cycle, as shown in Table 1 [117]. Proposed models should map business requirements into data requirements. Technically, data is prepared according to data requirements and processed with appropriate data mining technologies. Deployment mainly consists of presenting discovered knowledge. Research in information systems has traditionally adopted data mining processes which focus more on the variables under investigation and less on technologies.

<sup>6</sup><https://cloud.google.com/ai-platform/docs/ml-solutions-overview>

<sup>7</sup><https://aws.amazon.com/getting-started/hands-on/build-train-deploy-machine-learning-model-sagemaker>

Research on deep learning adds development processes that focus on data processing pipelines because ML models grow excessively in size, sophistication, and training costs. Therefore, focusing on performance issues, identifying bottlenecks, and optimizing hyperparameters, are critically important for deep learning models (cf. Table 1) [78].

The goal of a data science project is to *approximate an unknown function* by a fitted statistical model that exhibits an estimated function, which generates results (predictions) by obeying domain and technical constraints, while maximizing performance goals. A learner abstracts from all possible learning models and makes hypotheses on the unknown functions that relate input data with results (in a hypotheses space) [55]. This means that data scientists choose data (aka features) and data representations. They also select model candidates that are hypothetically capable of finding a fitted function,  $f$ , showing satisfying performance. This step requires explicit representations for data, functions, constraints, and performance goals including objective functions that can be scrutinized as part of subsequent analysis and explanations of results. It is especially important that the objective function used for assessing the quality of a trained model is, not only defined based on its technical purposes, but also supports domain requirements. For example, if domain experts are interested in identifying all features with an impact on the results, this would be contradicted by using a L1/Lasso norm that tends to eliminate features with small impacts and, thus, favor a sparse function  $f$ . Domain experts and data scientists must agree on project goals, generally, and on the level of complexity of the trained model, specifically, in accordance with the problem statement.

By integrating the different views, we now propose a *data science development process model* that consists of the phases identified in Figure 8.

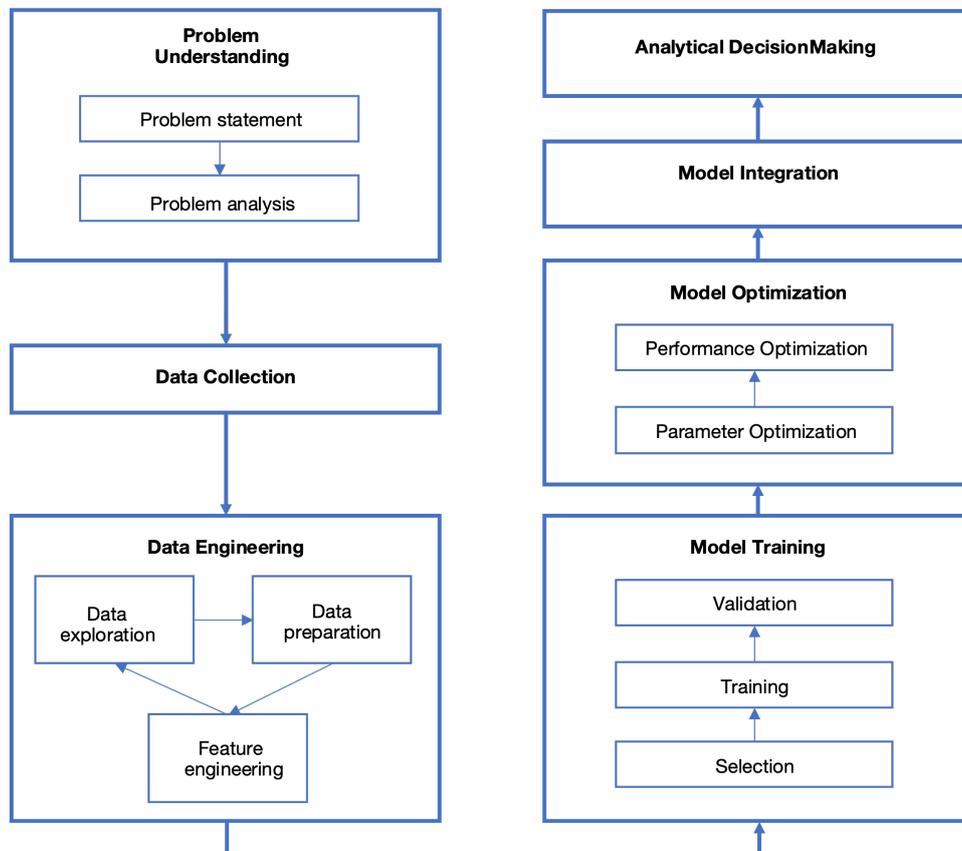


Figure 8: Data science development process.

### 3.2.1 Problem understanding

Complex environments require managers to make decisions under increasing uncertainty. By digitalizing many common processes, business environments have been able to adopt machine learning technologies, such as manufacturing [222], finance [176], marketing [98] and also social media [182]. Data science is applied to decision problems that can be addressed by statistics and machine learning technologies. Data science projects translate problem statements provided by domain experts into project definitions for data scientists. Given a problem statement, a data scientist tries to find solutions to a decision problem by posing three questions: (1) is there a mathematical formalization for this decision problem and a solution path based on linear algebra and statistics, and, if so, (2) is this solution path implementable on some software platforms and (3) is this implementation scalable for production?

*Problem understanding* starts with a problem statement and description of the data science project by domain experts. The results from both are discussed with data scientists until a shared understanding exists. Software engineering has many examples that support the importance of shared understanding [8]. Even though software engineering and conceptual modeling have investigated means for properly building shared knowledge, little has diffused into machine learning research. A strong emphasis on data alone diminishes the importance of domain knowledge and the role that domain experts play in designing ML-based information systems.

A *problem statement* is a hypothesis, created by a domain expert, which asserts that a decision problem can be solved by a computational process. A data scientist analyzes a problem statement for gaining a proper understanding of the problem. The problem understanding phase is highly iterative. Typically, neither the problem statements nor the data scientist's understanding of the domain is sufficient. Conceptual modeling provides a rich toolbox for supporting shared understanding between domains and technical experts. A problem statement describes decision making situations and parameters that influence decision making. Uncertainties and external influences might influence the decision-making process. Because decision making is embedded into a business context, performance goals, such as key performance indicators (KPIs) or response time behavior of decision processes are defined [71].

During *problem analysis* it is important to assess whether a problem statement can be translated into a data science problem that is feasible to solve, given desired performance goals. The problem analysis phase also includes project management issues, such as negotiation and definition of human, data and computational resources, milestones, and time plans. During problem analysis, data scientists start investigating whether the problem can be understood as a classification or regression task and whether this becomes accessible by supervised, unsupervised or reinforcement learning approaches. For any data science project, it is crucial to determine the accessibility of data, data size, and data quality. Care is required if data needs to be collected. In general, the resources required for data collection are grossly underestimated, but have a direct impact on data quality.

### 3.2.2 Data collection

Data is the core object in data science projects. Data is not just collected by some business processes, but also by sophisticated means, such as the Internet of Things (IoT) sensors, remote sensing, social media, financial markets, weather data, supply chains, and so forth. In this sense, data becomes an economic asset (data product [212]) exchanged via data ecosystems [157].

Data collection is constrained by data requirements derived by a problem statement and problem analysis, specification of data sources and corresponding data types, and volume, and quality requirements. Data with sufficient quality is a precondition for quality results of data science projects. Data quality is described using four main categories: (1) *intrinsic* including *accuracy*, (2) *contextual* including *relevancy* and *completeness*, (3) *representational* including *interpretability*, and (4) *accessibility* [213]. Additionally, some researchers have added *availability* as another main category [33]. Data quality strategies are distinguished as: (1) *data-driven* and (2) *process-driven*.

A data-driven strategy improves data quality by data modification. A process-driven strategy tries to modify the process by which data is collected [14]. Numerous data quality methods exist that span steps for evaluation of costs, assignments, improvement solutions, and monitoring, with numerous quality metrics [14].

### 3.2.3 Data engineering

After data has been made available, it is cleansed, explored, and curated. For univariate data these procedures overlap with data mining (cf. CRISP-DM) [217]. This process is more demanding for multi-variate data, such as image and video data with multi-dimensional features with channels (e.g., for RGB colors). Time series data, such as that provided by sensors, often contain missing data that needs to be replaced by meaningful data fitting with a temporal context [129]. More sophisticated exploration and preparation methods are used for unstructured data, such as texts and auditory data. Auditory data is usually transformed into textual data from which core structures are extracted by text mining, including keyword selection and linguistic preprocessing (e.g., part-of-speech tagging, word sense disambiguation) [97]. An example from health care is found in Palacio and Lopez (2018) [158].

Data exploration is used to understand a dataset in detail with respect to the domain. Descriptive statistical analysis provides standard metrics, such as mean, standard deviation of variance of single features, and correlation values between two features, whose values must remain within the boundaries of the application domain. In addition to statistical analysis, semantic analysis exploits domain requirements for assessing data validity. Ontological representations [82] may be associated with features to support domain experts in understanding the datasets. The domain can also restrict constraints on data values and the range of acceptable values. For instance, the concept “blood pressure” has an associated constraint that blood pressure values cannot be negative. Thus, domain requirements enable domain experts to understand data sets and assess their quality and, in this way, reflect some of the semantics of the real world. Domain requirements can be simple statements, such as feature ranges, or complex conceptual models with cascades of requirements that need to be tested carefully. For some domains, theories with formal representations exist.

Datasets are rarely collected in highly controlled laboratory environments but, instead, collected in different environments under dynamically changing conditions. Datasets are mixed, merged, and added with features that are not necessary for the data science task at hand. Feature engineering provides methods for identifying the features that are relevant and those that are not. However, this task is highly dependent on both the domain and problem statement. Technically, feature engineering is domain-specific and requires intuition, creativity, and “black art” [55]. Technical feature engineering alone can result in negative side-effects if, for example, features are dropped that are relevant or merged by incorrect means. Relevant features subsequently increase performance of the trained model [24].

Large univariate and multivariate datasets are generally difficult to analyze at an item level. However, flawed results of data science projects are often caused by missing an understanding of the structure and meaning of a dataset. Research in statistics has developed standard visualizations of probabilistic data, such as visualization of density functions plus visualization of statistical measures, such as box plots, histograms, scatter plots, normal Q-Q plots and quasi-visualization, such as correlation matrix and confusion matrix. Legendary are Hans Rosling’s data visualizations that make transparent what is hidden in raw data.<sup>8</sup> Exploration of multivariate data is much more complex. For instance, analyzing whether images in an animal data set actually show buildings requires either many people [58] or models that have been developed on other datasets.

When a domain expert and data scientist have a common understanding of the data set and single features, data is prepared for analytical processing. Data preparation includes: (1) *data exploration*, (2) *data preparation*, and (3) *feature scaling*. Similar to ETL (extract-transform-load) in data mining, the *data exploration* phase processes and transforms raw input data into a

---

<sup>8</sup>[https://www.ted.com/playlists/474/the\\_best\\_hans\\_rosling\\_talks\\_yo](https://www.ted.com/playlists/474/the_best_hans_rosling_talks_yo)

dataset of sufficient quality. *Data preparation* includes statistical procedures for handling missing data, data cleansing, and data transformation by normalization, standardization, and reduction of dimensions (e.g., Principal Component Analysis (PCA)). Data preparation is a “black art” that needs to be transformed into a “white art”. This goal of data interpretability aligns with the need for interpretability of models and explainable AI (XAI) for making “black boxes” of models transparent to users (e.g., [173]).

Data interpretability requires the data preparation steps to obey data constraints as part of domain constraints. For example, missing data deals with replacing unknown entries by computed values, such as a mean or median. Most frequently, “not a number” (NAN) is used or samples are deleted that have missing values. Data constraints for features guide data scientists in selecting appropriate procedures. Social sciences have experience with handling outliers according to different outlier categories: error outliers, interesting outliers, and influential outlier [4]. When lacking proper domain understanding, outliers are often deleted in data science projects because of negative effects on performance measures. However, interesting, and influential, outliers are anomalies relative to the dataset that potentially provide insights. For example, the fundamental purpose of the ATLAS (A Toroidal LHC ApparatuS) project on finding the Higgs boson [1] focused on finding anomalies, so deleting outliers would have rendered this project useless. Therefore, data constraints define the limits on what is theoretically possible in a data science project.

Data constraints are derived from domain theories. They describe ranges, rules, and invariants as well as functions on features. For instance, ranges constraint meaningful feature values, whereas rules describe dependencies between features of a sample. For example, if age is  $x$  years (today), then the date of birth is within a specified range of years  $[today - x, today]$ . Invariants are strong assertions that hold within features (for instance, feature on gender are required to be evenly balanced) or between features (for instance,  $birthday < dayofdeath$ ). In addition to textual descriptions, ranges, rules and invariants can be formally modeled using various formalisms, such as subsets of predicate logic [145], constraint logic programming [105], constraint satisfaction formalisms [203], and constraint formalisms for object models [174]. More challenging is ensuring that data constraints are valid when data transformation is applied.

Validation data annotation ensures that data preparation obeys domain constraints and generates datasets that are meaningful at both the feature level and the dataset level. The basis for semantic data preparation includes four categories for data quality: accuracy, relevancy, representation, and accessibility [212, 213].

Missing data is a major concern in almost any data preparation phase. Various imputation strategies [202] are applied for replacing missing data by random values, mean or median values, or most frequent value; using feature similarity in nearest neighbor models; removing features; or applying machine learning approaches, such as DataWig [20]. Current machine learning models only work with numerical values. Therefore, categorical or textual data is transformed into numerical representations. A standard technique for categorical data is *one-hot-encoding*, which adds binary features for each category. Preprocessed textual data is often categorical and is either mapped onto numerical indexes or transformed by one-hot-encoding.

Beyond standard imputation and encoding, data is transformed in various ways. Integration of features can lead to more expressive, additional features. For instance, if one feature is income and another is number of people per household, adding a feature that divides income by number of people per household can provide valuable information for predicting educational development. Several machine learning models, as well as gradient descent, use differences between features, such as KNN, k-means and SVM. Therefore, features with larger scales have more influence than smaller scales. Normalization (min-max scaling) and standardization are standard feature scaling procedures that tend to improve model training and prediction quality.

Normalization is applied if the data does not follow Gaussian distribution. Empirically, models that do not presuppose specific distributions, such as KNN, Perceptrons [143] and neural networks, can improve prediction performance by normalized data. Similar improvements can be achieved by standardizing data for use in distribution-dependent models.

Feature engineering is selects, transforms, adds, constructs, or replaces features in such a way that it improves model training and model performance, without changing feature semantics. Data scientist’s prior knowledge and skills are needed for organizing data representations so that discriminative information become accessible [16]. Various methods are used for creating additional features from input features, such as calculating differences, ratios, powers, logarithms, and square roots [88]. For text classification, correlation-based methods are used, such as information gain [227]. Semantic similarities of concepts derived by using ontologies [163] are used for feature ranking and feature selection (e.g., [163]).

Representation learning is a current research topic that attempts to automatically extract representations of data, such as posteriori distribution of some explanatory factors underlying observed input. These factors decrease the complexity of feature engineering because they can be used as guidance or even as input to supervised learning models [16].

### 3.2.4 Model training

Model training includes *selection*, *training* and *evaluation* of models. Training a model means adjusting the model parameters to the data. Based on an objective function. Supervised learning models adjust weights according to loss gradients for minimizing, for instance, the sum of squares for regression and minimizing cross-entropy for classification [86]. Unsupervised learning models use the sum of distances, and reinforcement learning use updates based on reward evaluations. In the early phases of data science projects, it is usually not clear which machine learning model will exhibit the best performance. Therefore, several model types with hyperparameter ranges are often tested against each other. This iterative exploration phase narrows down prime candidates for subsequent phases.

With the introduction of many different types of model architectures within a short period of time (due to the surge in the popularity of machine learning), guidelines and modeling patterns are increasingly important. To date, the architecture designs of machine learning models emerge from the practical needs of machine learning experts. This knowledge slowly diffuses to less experienced designers. Conceptual modeling, with its capabilities for abstracting the real world, can help make machine learning architectures more accessible and practically useful [191]. Elements of a model driven architecture (MDA), including UML, provide languages for describing machine learning model architectures. MDA could aid in the selection of implementation of algorithms, based on users’ requirements. Furthermore, object-oriented design patterns [69] provides a basis for technical design patterns for constructing machine learning model designs.

A general challenge for designing model architectures lies in the appeal of complex models. Even unexperienced machine learning architecture designers are inclined to prefer recent and more complex models over older and simpler models. Model design requirements are necessary that constraint minimum and maximum complexities of model designs according to problem statements and associated goal models and goal constraints. At an abstract level, model design requirements describe guidelines [191]. At a technical level, model design requirements provide information on required capabilities of model units on various levels. For instance, there could be requirements on the capability of neuron types (e.g., plain neuron or LSTM neuron) or pattern of connections between neurons; e.g., fully connected, or filters for convolutions are on the lowest level. Larger structures of layered neurons are called a topology of a network [141]). Intermediate requirements encompass the number of layers, building blocks of layers (e.g., LSTM layer, softmax layer) and general mechanisms, such as attention. Top-level requirements describe the model design space. For example, it might require the use of linear models only or models for which theoretical guarantees exist, such as those associated with complexity classes or optimality criteria.

Depending upon the datasets used, model designs have a major impact on model performance. Making requirements on performance ranges explicit will further restrict the model design space. Using performance requirements at design time is either based on heuristics or is probabilistic because of the unknown function underlying the dataset, making the actual model performance unknown. Means for expressing heuristics on the relationships between performance requirements,

datasets, and model types include heuristic rules, constraints, and logical expressions. Relationships can also be learned, given enough data on performance, models, and datasets. This multi-dependency between dataset, model design, and performance requirements carry knowledge that is important for any model designer and decision maker. The more experience that is accessible, the better a data scientist can select model designs that fulfill targeted performance ranges. A small, clarifying example for this argument is a neural network with two input features ( $x_1$ ,  $x_2$ ), one fully connected hidden layer with just two neurons, and an output layer with one neuron for adding activations. Given data from two classes that are embedded in one another (i.e., cannot be separated by a line), this model will probably not exhibit high performance (i.e., small loss) with respect to accuracy because model complexity is not specific enough and, thus, underfits the dataset.<sup>9</sup> A performance requirement is expressed as a loss on misclassified samples of less than 10%, probabilistic knowledge on performance ranges for this small neural network, and a binary classification task with 500 datapoints will inform model designers about a likely mismatch between the modeling task and performance requirements at design time. In this case, performance requirements are achieved by adding another neuron to the hidden layer.

Structural dependencies in model design tasks have an impact on resources spend on model training, parameter optimization and, subsequently energy and time consumption. So far, this knowledge is part of a data scientist’s “black art”. Making this crucial knowledge explicit by conceptual model representations is important for managing data science projects and businesses. Constraints languages, such as OCL (Object Constraint Language) [174], are means for describing and evaluating requirements between dataset, model design and performance requirements at model design time and, subsequently, for model evaluation when actual model performance is assessable. Model evaluation tests whether actual performance fulfills performance requirements.

Assessing the ability of a model to generalize is important for performing well on unseen data. The more complex a model, the better it can adjust to training data (low bias), although it might overfit and work less well on unseen testing data (high variance) [86]. The goal is finding a model and a model architecture with a minimum of absolute training and testing loss and a minimum distance between them.

Data sets are split into several parts used for training, validation, and testing. Splitting data is often based on heuristics, for instance, 50% training, 25% validation, and 25% testing. The training set is used to train as many models as there are different combinations of model hyperparameters. These models are then evaluated on the validation set, and the model with the best performance (e.g., the smallest loss or highest accuracy) on this validation set is selected as the final model. This model is retrained on training and validation data with the selected hyperparameters. Then, model performance is estimated using the test set. It is assumed that the model generalizes well if the validation error is similar to the testing error. Finally, the model is trained on the full data. If datasets are small, training and validation is carried out with the same dataset. Folded *cross-validation* and *bootstrapping* are used for iterative assessment of model accuracy. Cross-validation separates the training dataset into partitions (folds) of the same size. One fold is separated for assessing model performance and the others used for training. Average performance is determined by repeating this process with all folds. Bootstrapping draws samples from the training dataset with replacement and trains a model for a specified number of times. Accuracy is assessed by averaging over all iterations [86].

### 3.2.5 Model optimization

With unlimited resources, machine learning models can train and evaluate to find an optimal system configuration. Business analytics and data science, as well as related research has continued to require attention [40, 130]. Model complexity increases excessively, making a brute-force approach infeasible. Optimization tasks are ubiquitous in machine learning. A key optimization task is finding weights that minimize loss in supervised learning or finding policies that support

---

<sup>9</sup>Example, cf. <https://t1p.de/oc2w>

a goal best in reinforcement learning. Gradient descent and stochastic gradient descent are basic algorithm for weight optimization. However, more efficient algorithms are used in practice that adds a momentum vector for speeding up gradient updates (e.g., Adam optimizer [111]).

Most machine learning tasks are controlled by external parameters, called *hyperparameters* that constrain a model's search space; for instance, the value of  $k$  in KNN, maximum depth or number of trees in random forest, or number of layers and neurons per layer in neural networks. With brute-force,  $k$  would be the range of all positive integers and a threshold for performance. Experience in the domain and previous experiments might have shown that  $k=\{4,\dots,7\}$  are most likely candidates for minimizing the loss function and optimal accuracy. Therefore, experience indicates that  $k=\{1,2,3\}$  is not worth training. For real-valued hyperparameters, this problem become even more pressing. Grid search is a greedy procedure for finding the best hyperparameter settings by using all hyperparameter combinations. This only works for small datasets and small number of hyperparameter combinations. Several approaches exist for automatic hyperparameter optimization [100]. Recently, AutoML systems have been introduced, such as Auto-sklearn [65] or AutoKeras [107], that provide automatic optimization across hyperparameter settings with an emphasis on neural networks.

Configuration parameters of relational databases barely make it into scientific discussions. The difference is that hyperparameters directly affect finding at least locally optimal models. Setting ranges for hyperparameters too large might result in excessive resource requirements whereas too small ranges might threaten the search for the best model. Requirements on hyperparameter are influenced by the domain, dataset, expertise and previous modeling tasks, but, most of all, by the model type and its implementation. Similar to performance requirements, hyperparameter requirements are an open field for conceptual modeling. Hyperparameter requirements can simply set parameter ranges. Alternatively, hyperparameter requirements can describe the complex dependencies between business requirements, goals, resource models, performance requirements, services requirements, and others. With enough knowledge captured by hyperparameter requirements, companies can optimize their resources by invested in a ML-based service development that can result in shorter time-to-market of products and services.

The *performance* of a model is assessed by analyzing results of predictions for testing data. For classification, the number of items that are correctly and incorrectly classified are analyzed. For a binary case, four cases are differentiated. Two are correct (positives and negatives are correctly classified); two cases make opposite predictions (false negatives, false positives). A confusion matrix separates these four cases: true positive (TP), true negatives (TN), false positives (FP) and false negatives (FN). Sensitivity,  $\frac{TP}{TP+FN}$ , is a measure for positive cases and specificity,  $\frac{TN}{TN+FP}$ , for negative cases. If false negatives and false positives are rare, sensitivity and specificity are close to 1. In practice, it depends on the domain and the decision task as to which metric is most important. For instance, in healthcare, there is stronger emphasis on sensitivity. Alternatively, precision  $\frac{TP}{TP+FP}$ , and recall,  $\frac{TP}{TP+FN}$ , are used with recall the same as sensitivity and precision the percentage of correct true cases modified by false positives. The F1-score combines precision and recall in one metric which is useful in cases with no clear preference for precision or recall. Finally, accuracy in binary classification is the percentage of true classifications over all samples,  $\frac{TP+TN}{TP+FP+TN+FN}$ .

The loss function for classification uses cross-entropy:  $H(p, q) = -\sum_{k=1}^K p_k * \log(q_k)$  with  $p$  probability of ground truth and  $q$  probability of predicted categories. If probability  $q$  is close to probability  $p$ , cross entropy will be close to 1. The difference between cross-entropy  $H(p, q)$  and the entropy of probability  $p$ , i.e.  $H(p)$ , is called Kullback-Leibler Divergence  $D_{KL}(p||q) = \sum_{k=1}^K p_k * \log\left(\frac{q_k}{p_k}\right) = H(p, q) - H(p)$ .  $D_{KL}$  is used as a metric for the performance of a classification model.

For a regression task, proportion of declared variance  $R^2 = \frac{\sum_{i=1}^n (\hat{y}_i - \bar{y})^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$  is often used that is close to 1 if residuals between ground truth  $y_i$  and estimates  $\hat{y}_i$  are small. Because loss functions for regression tasks normally use squared residuals, weights are adjusted for minimizing residuals and, therefore,  $R^2$ .

After training a machine learning model based on a risk function, model performance is evaluated by performance metrics. The performance metric value for a model requires domain-dependent interpretation. For instance, *sensitivity* and *specificity* results for a binary classification on healthcare diagnosis typically favors sensitivity (percentage of true positives = “ill”), over specificity (percentage of true negatives = “not ill”). Relative performance values are used for model comparison whereas absolute performance values determine whether a model is good enough. Thus, performance metrics are operationalizations of quality requirements that a model needs to satisfy; that is, model performance expressed by a performance metric is required to exceed a quality threshold. Conceptual modeling can support performance evaluation in two ways: 1) selection of performance metrics; and 2) threshold for absolute performance for selected performance metrics. Performance requirements constraining the selection of performance metrics depend on the domain and the modeling task. For instance, for classification tasks healthcare domain prefers sensitivity/specificity over precision/recall.

Performance thresholds are target of extensive debates in research domains (for instance, discussion on threshold for confirmatory factor analysis, CFA [154]) and, thus, carry deep knowledge. In the simplest form, requirements on performance thresholds are single numbers but can be expanded to intervals and distributions (similar to confidence intervals). From a scientific point of view, performance requirements must be defined before model training so that performance results of model evaluation on testing data can be assessed without bias. In practical applications, performance results are input for decision makers for making decisions on project progress and future business. If performance results seem promising by getting closer to performance requirements, positive decisions on investments in subsequent development phases are more likely. However, performance requirements are not absolute but, rather, adapt to developments in a particular field. For instance, NLP (Natural Language Processing) adopts performance metrics from computer vision (e.g. Intersection-Over-Union, IOU [62]) but also define new performance metrics, such as comprehensiveness and sufficiency within the context of explainable NLP [53]. Performance requirements could be a large area of research with descriptions of performance requirements needed. Goal models are required for mediating business goals and performance results. Specification languages are needed for properly representing, communicating, and eventually automatically reasoning on performance representations.

### 3.2.6 Model Integration and Evaluation

ML-based information systems in decision making are recognized as an important topic for both research and practice. Many applications use machine learning model types that can be directly evaluated by humans. For legal and business reasons, ML-based information systems are required to explain their results by means accessible by non-technical domain experts (explainable AI - XAI). Linear regression models, logistic regression models, decision trees and support vector machines can be all scrutinized; much more effort is required for complex ensemble models, such as those based on XGBoost. Single predictions by deep learning models and reinforcement learning models are based on myriads of simple, highly interconnected calculations that make direct understanding by domain experts impossible. For instance, a decision to stop a production line due to a ML-based prediction requires strong arguments and explanations. A recent approach involves fitting simpler surrogate models close to local areas of a prediction and using *surrogate models* for explanation, such as Individual Conditional Expectation (ICE) plots [76], Local Interpretable Model-agnostic Explanations (LIME) [173] and Shapley Additive Explanations (SHAP) [128].

### 3.2.7 Analytical Decision Making

Decision makers need more than just explanations for predictions. The performance of models strongly depends on data, so decision makers must: scrutinize raw data; pre-processed smart data [195]; identify the semantic information used for merging and processing data; identify the objectives of the data scientists who developed machine learning models; and estimate economic effects,

side effects, risks, and alternatives. Decision makers also need to understand potential semantic losses (*lost in translation*). Examples of these requirements are included in the Explainability Framework in Figure 9.

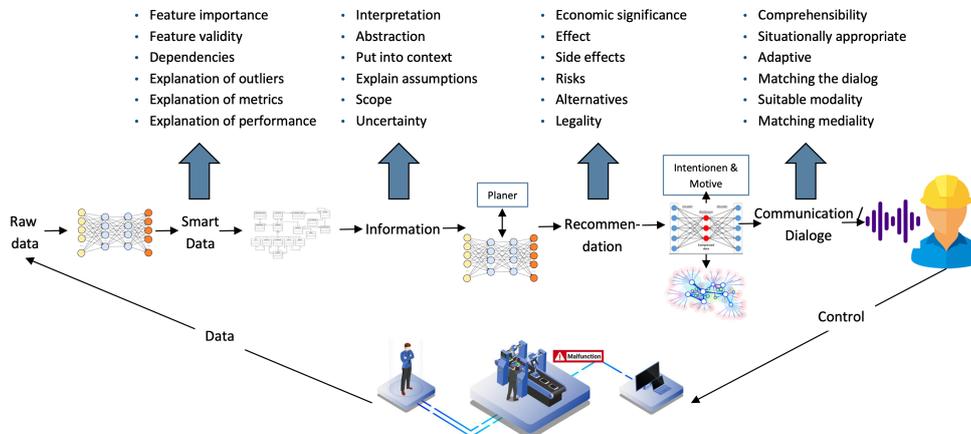


Figure 9: Explainability framework.

Figure 9 captures the important concept of *Explainable AI* (XAI), which shows how raw data is operated on to progress to information that can be used to make recommendations to a user [9]. Users need to understand the explanation of how the output is obtained. This enables the user to consider the explanation and assess whether it is necessary to rework a problem.

## 4 Conceptual Modeling for Machine Learning

Practical machine learning models are only useful within a given domain, such as games, business decisions, healthcare, politics, or education. When embedded into information systems, machine learning models must follow laws, regulations, societal values, morals, and ethics of the domain, and obey requirements derived from business objectives. This highlights the need for conceptual modeling to address the “black box” challenge of complex machine learning models. Conceptual models help transform business ideas into structured, and sometimes even formal, representations that can be used as precise guidelines for software development. Therefore, they help structure the thought processes of domain experts and software engineers, for building a shared understanding between these groups and for providing languages by which information system implementations can be understood, scrutinized, revised, and improved [132, 38]. Conceptual modeling of machine learning can also make it easier to gain skills by abstracting machine learning technologies with the help of model-driven software engineering and automatic code generation [30].

Complex machine learning models, such as deep learning models, are often considered as “black-boxes,” which are not well-scrutinized. Machine learning experts and data scientist perceive domain knowledge as a quarry from which ideas and initial guidelines can be extracted. (“We begin by training a supervised learning (SL) policy network  $p_\sigma$  directly from expert human moves.” [187]). The mechanistic nature of reinforcement learning requires exploration of any changes in environments [103] and, thus, is focused on how, rather than why, the decision-making process occurs [17]. In domains, such as video games, the basis on which a decision has been made is not important. However, in business domains, decision making requires: *trust* in recommendations; sufficient *understanding* of the reasoning processes and the underlying assumptions behind a recommendation; *legal* and *ethical* obligation adherence; and support of *stakeholder* requirements. As extreme examples, a ML-based system could recommend laying off all employees or investing in weapons

for the extermination of mankind. No serious decision maker will follow such recommendations without scrutinization. However, the decision maker will ask for explanations; look at the data used for training; ask for second opinions and recommendations of alternative models; analyze software development procedures and requirement documents; talk to software engineers and data scientists; hire external experts for unbiased views; and probably much more. This requires documentation and identification of the representations that were used for designing and building this ML-based information system, and understanding how they will help to explain a system’s behavior and recommendations.

There are differences between machine learning systems and, for instance, database systems. Database systems implement domain knowledge that has been adopted by domain experts. In contrast, ML-based information systems are not intended to implement prior knowledge, but rather, find useful patterns for making predictions given input data. These patterns may lead to theoretically interesting questions that could guide subsequent research, as is typical in biomedical research. For domains, such as gaming, designing, research and development, music and art, this freedom for finding innovative patterns and making unprecedented predictions is appreciated. For domains, such as legal decision making, production and manufacturing, healthcare, driving, operating chemical and power plants and military, highly reliable and trustworthy information systems are required that follow laws, ethics and values. Thus, conceptual modeling methods and tools should enable users to scrutinize, understand, communicate, and guide the entire lifecycle of ML-based information systems. This motivates the need for a general framework that aligns design, development, deployment, and usage of ML-based information systems for decision making.

#### 4.1 Framework for conceptual modeling in data science

The alignment between business strategy and business operations with its IT strategy and IT operations is an enabler for competitive advantages [140]. Conceptual models provide specification languages for capturing business requirements that can be translated into software requirements [201, 161]. They include primitive terms, structuring mechanisms, primitive operations, and integrity rules [149], with the entity-relationship model representative of a semantic specification language [43]. Dynamic sequences of activities are captured by process models, such as event-driven process chains, UML activity diagrams or BPBM models. For domain experts, requirement models and specifications are generally too abstract, so early requirements analysis attempts to capture stakeholders’ intentions [37] and goals [229]. Conceptual models are translated and refined by software-oriented requirements languages until they can be used as a basis for implementation. Alignment of ML-based information system development with business goals and strategies is at an early stage of research and understanding [126, 135]. Therefore, guidelines and frameworks are needed to identify the research topics that need to be studied and to support the progression of the needed research. For example, ML technologies are used to explore potential cost reduction (e.g., via predictive maintenance), but used less for business innovations. Decision makers might be reluctant to employ machine learning because of possible poor data quality and black-boxed ML algorithms [34]. Although the entity-relationship model was initially introduced to gain a “unified view of data” [43], it has, after many decades of research, been extended to business goals, intentions, processes and domain ontologies [197, 63, 60]. A domain ontology provides a set of terms and their meaning within an application domain [80], with many domain ontologies having been created and applied; however, there are many quality assessment challenges [139, 138, 32]. In response to the need to align machine learning and business, as well as the challenges in doing so, Table 2 provides a framework for incorporating conceptual models into data science projects.

##### **Problem understanding**

Using Machine learning models within a business context requires providing solutions to business problems. Research in innovation distinguishes between “technology-push” and “need-pull” [181]. From the technology-push perspective, adoption of machine learning is the driving force for competitive advantages [167]. This view is challenged by the long sequence of failure that AI has

Data Science Main Phases	Sub-Phases	Conceptual Modeling Concepts, Methods and Tools
Problem understanding	Problem statement	Business requirements, goal model
	Problem analysis	Business requirements, goal model; legal and ethical requirements; Data requirements
Data collection		Data requirements; Data quality; legal and ethical requirements, business requirements
Data engineering	Data exploration	Data requirements; Legal requirements
	Data preparation	Ontologies; Domain models
	Feature engineering	Ontologies; Domain models
Model training	Selection	Business requirements, legal requirements, performance requirements and conditions for acceptance
	Training	Performance requirements
	Validation	Performance requirements
Model optimization	Parameter optimization	Domain model, resilience requirements
	Performance Optimization	Domain model, resilience requirements
Model integration		Business requirements; Goal model; Legal requirements; Ethical requirements; Data requirements
Analytical decision making		

Table 2: Framework for incorporating conceptual models into data science projects.

suffered over many decades. For instance, Google’s Duplex dialog system that impersonates a human, raises ethical concerns and affects trust in businesses, products and services [155]. Other examples use machine learning for visual surveillance, which breaches privacy laws or uses machine learning on social media data to influencing political debates. Legal [31] and ethical requirements [26] increasingly influence design decisions on ML-based services. This resolves uncertainties for data scientists who are challenged by unclear ethical and regulatory requirements [209].

Generally, elicitation of business needs and business requirements within the context of ML-based information systems is a novel field of research. However, it is dominated more by questions and challenges, than answers, such as the lack of domain knowledge, undeclared consumers, and unclear problem and scope [15]. Proposed approaches to business intelligence have a strong overlap with ML-based information systems [94]. Because the classes of ML-based information systems are broader than business intelligence systems, a wider range of stakeholders need to provide input to problem understanding. Qualitative methods are often used to understand business needs and elicit requirements [134, 131]. In general, conceptual modeling provides a large set of modeling approaches that help heterogenous teams gain a shared an understanding of the strategic and operational business needs and goals, as well as the constraints associated with ML-based information systems. This includes shared understanding on performance and quality requirements for data, models, and predictions.

Goal modeling may become a key contribution of conceptual modeling to ML-based system development [127]. Nalchigar et al. (2021), for example, propose three views for modeling goals for ML-based system development: business view, analytics design view, data preparation view [150].

### Data collection

Any information system depends on input data. ML-based information systems even extract behavior from data, which is why data is so important. Database and information systems emphasize the importance of data schema, with web-based open data increasingly annotated with semantic markers (e.g., Gene Ontology, YAGO, dpPedia, schema.org). In specific contexts, data standards are available for different industries (e.g., eCl@ss or UNSPSC). For streaming data, as prevalent

in real-time systems and the Internet of Things applications, new standards are defined, such as OPC-UA for processing semantically annotated data in distributed environments and using machine learning based on, for instance, MLlib with Apache Spark.

Appropriation of existing data sources falls into two categories: open data and proprietary data. Many governments maintain open data repositories, such as Data.Gov in the USA and Canada, and GovData in Germany. Wikipedia extracts are provided by dpPedia (dbpedia.org). Access to proprietary data depends on contractual agreements because raw data is generally not protected by copyright laws, whereas audio and image data can claim copyright protection if deemed to be artwork. Work on digital rights management (DRM) has developed proprietary and open solutions for protecting media data, such as music and videos and enforcing license management [206]. Application of DRM on operational data, such as data by the Internet of Things systems, requires analytical run-time environments that implements DRM standards [121]. Blockchain approaches enforce immutable exchange of data and execution of contractual obligations [223]. Large Internet companies follow a business model that centralizes data via cloud infrastructures and provides access to data via market mechanisms, such as auctioning. Alternatively, federated data platforms favor decentralized data repositories that are connected via data exchange protocols (e.g., GAIA-X in Europe).

Data collection depends on data requirements [209] that provide a precise understanding of the type of data and data quality necessary for finding an application. Entity-relationship models and its derivatives are proper means to represent the requirements for data collection. These models can be used for storing, screening and interpreting of data collections in the sense of ETL-processes of data mining [196]. Data integration from multiple sources with heterogenous data schema requires ontology-based matching and mapping [61]. Data requirements for univariate data overlaps with modeling approaches. Multivariate, graph-oriented, and textual data require research on extended modeling mechanisms. Beside alignment with business requirements, data requirements also capture crucial legal and ethical requirements. Examples include constraints on the origin of the data, as well as its quality.

Data quality has a major influence on model performance and, thus, the utility of a ML-based information system. Consistency and completeness are two major indicators of data quality [118, 166] with further research on data quality needed for the adoption of external data sources. Additionally, recent developments in data ecosystems, such as GAIA-X, shows the importance of modeling legal and contractual requirements [35].

### **Data engineering**

Data requirements provide a basis from which to consider data transformations. [109]. Data requirements capture *semantical*, *structural* and *contextual* descriptions. Semantical descriptions represent information about data types, and their accepted interpretations; e.g., pressure is record in Pascal. Structural descriptions provide constraints on the form of the data; e.g., sample rate ranges, acceptable percentage of missing values, and accuracy of sensors used for collecting data. Structural descriptions are related to data quality, and also capture descriptive information, such as the time and location of data capture Contextual descriptions represent the constraints of a domain and the context within which a ML model is intended to be used, which includes requirements preventing biases or demanding coverage. Thus, data models capture semantical, structural, and contextual descriptions and provide information about obtaining data requirements.

Besides requirements, data engineering also extracts knowledge about data that has not been visible previously. The number of dimensions of a data space can be reduced (e.g., by principal component analysis) or additional dimensions added (e.g., by one-hot encoding of categorical dimensions). Combining dimensions requires theoretical understanding (e.g., of physical mechanics when combining mass and acceleration into force and velocity into energy). Dependencies between data and machine learning models require data engineering. For instance, models that use gradient descent work best if data dimensions are first standardized and normalized for multiplication.

Data requirements represent dependencies between data dimensions for constraining data trans-

formations. Results for data explorations and data transformations are fed back into enhanced data requirements for capturing additional semantics. For instance, a typical first step in data engineering is correlation analysis between input data that is visualized for data scientists, but lost afterwards. Because data exploration, data preparation, and feature engineering generate rich knowledge about data, enhanced data requirements become important. Domain experts can scrutinize this knowledge about the data before it is used for model training.

### **Model training**

Training complex machine learning model is resource-intensive and strains computation, energy and financial resources. Therefore, the declaration of functional and non-functional requirements guides model training and provides boundaries. Regulations and legal rules put requirements on data and model behavior, energy consumption, and sustainability. To date, research on legal requirements mainly focuses on the behavior of a machine learning model with respect to interpretability and explainability, especially as a consequence of European laws and the General Data Protection Regulation (GDPR) [19]. However, for commercial settings, the selection of machine learning models is tedious due to the need to avoid potential intellectual property infringements. It, thus, requires in depth technological and legal analysis, both before and after model selection. For example, decision makers might want to avoid spending extensive resources on training ML models, to later realize that they have violated license infringements. As models become more complex and stacked on top of each other, legal descriptions become even more important. After model training, various descriptions characterize functional and non-functional model behaviors, including model performance. Conceptual modeling practices have a long tradition of capturing such characterization in concise conceptual models. These models can then be used to extend or combine ML and integrate them into information systems.

### **Model optimization**

After model training, model optimization fine tunes the model parameters to ensure performance requirements are achieved. Doing so, requires updating conceptual models associated with the ML models. Resilience is a meta-requirement that describes a system's capability under disturbances, such as lower data quality or fewer parallel processes capabilities, than expected. A resilient machine learning system does not deny service under disturbances, but gracefully degrades. At the end of model optimization, all requirements and corresponding system documentations must be reviewed and updated.

### **Model integration**

Model integration resolves technical problems by addressing functional and non-functional requirements. This phase overlaps with traditional system integration that includes requirements for repair enablement, transparency, flexibility, and performance [85]. Requirements for the final analytical decision phase met business, legal and ethical requirements.

### **Analytical decision making**

Integrated ML models that fulfil model and data requirements should support business requirements as represented by conceptual models including goal requirements. Interpretability is important for any business decision making system. The entire stack of conceptual models, fully integrated with data and ML models, provide an important source for interpretability. Shallow integration only provides approximate estimates of system behavior. Full integration requires provable guarantees. Both the conceptual model stack and guarantee mechanisms require further research.

The liability for recommendations made by a ML-based information system is common in any service-oriented business. Legislators and scholars have started to demand higher levels of transparency and explainability of AI and ML technologies [20]. Technical solutions for explainable AI (XAI) (cf. section 3.2.6) are initial attempts that need to be aligned with legal and regulatory

Topic	Definition	Example specification languages
Business	description of business process that are related to the strategy and the rationale of on organization	I*, BPMB, UML, BIM, URN/GRL [7], BMM [156], DSML [74]
Legal	Goals that choices made during the ML development are compliant with the law (based on [185])	Nomos, Legal GRL [73]
Ethical [26]	Compliance with principles, such as transparency, justice and fairness, non-maleficence, responsibility and privacy [108].	textual
Data [209]	Requirements on semantics, quantity and quality of data	ER, UML, RDF, OWL, UFO, OCL
ML Model [109]	selection of architectural elements, their interactions, and the constraints on those elements and their interactions necessary to provide a framework in which to satisfy the requirements and serve as a basis for the design [164].	Finite state processes, labeled transition systems [205]
Functional	statements of services the system should provide, how the system should react to particular inputs, and how the system should behave in particular situations. [192]	BPMN, UML, EPC, KAOS, DSML [66] [74]
Non-functional	A non-functional requirement is an attribute of a constraint on a system [75]	UML, KAOS
Performance	expressed as the quantitative part of a requirement to indicate how well each product function is expected to be accomplished [46]	Rules quantified by metrics
Interpretability	Interpretable systems are explainable if their operations can be understood by humans. [2]	Qualitative rules
Resilience	ML models that gracefully degrade in performance under the influence of disturbances and resource limitations	Rules quantified by metrics

Table 3: Specification languages.

requirements.

Table 3 provides examples of specification languages known in conceptual modeling. Proven specification approaches exist for business, functional, and non-functional requirements. Nomos is used for legal requirements [186]. Data requirements resemble those for database systems and linked data. Specification approaches for ethical requirements, machine learning models, performance requirements, interpretability, and resilience also require further development and refinement.

## 4.2 Example

The framework for incorporating conceptual models into data science projects (Table 2) is illustrated by the following example.

### Problem understanding

The objective is to predict whether a female person has diabetes, recognizing that diabetes is a widespread disease that is difficult to manage. The problem is addressed based on a dataset<sup>10</sup> from the society of Pima Native Americans near Phoenix Arizona collected by the US National Institute of Diabetes and Digestive and Kidney Diseases. The Pima are a group of Native Americans living in central and southern Arizona and in Mexico in the states Sonora and Chihuahua. In the US, they live mainly on two reservations: the Gila River Indian Community (GRIC) and the Salt River

<sup>10</sup><https://www.kaggle.com/uciml/pima-indians-diabetes-database>

Pima-Maricopa Indian Community (SRPMIC). The GRIC is a sovereign tribe residing on more than 550,000 acres with six districts. They are involved in various economic development enterprises that provide entertainment and recreation: three gaming casinos, associated golf courses, a luxury resort, and a western-themed amusement park.

Two SRPMIC communities, Keli Akimel O'odham and the Onk Akimel O'odham, have various environmentally based health issues related to the decline of their traditional economy and farming. They have the highest prevalence of type 2 diabetes in the world, leading to hypotheses that diabetes is the result of: genetic predisposition [215], a sudden shift in diet during the last century from traditional agricultural crops to processed foods, and a decline in physical activity. In comparison, the genetically similar O'odham in Mexico have only a slighter higher prevalence of type 2 diabetes than non-O'odham Mexicans.

The Pima population of this study has been under continuous study since 1965 by the National Institute of Diabetes and Digestive and Kidney Diseases because of its high incidence rate of diabetes. Each community resident over 5 years of age has been asked to undergo a standardized examination every two years, which includes an oral glucose tolerance test. Diabetes was diagnosed according to World Health Organization Criteria; that is, if the 2 hour post-load plasma glucose was at least 200 mg/dl (11.1 mmol/l) on any survey examination or if the Indian Health Service Hospital serving the community found a glucose concentration of at least 200 mg/dl during the course of routine medical care [113].

In a study by Smith et al (1988) [190], eight variables were chosen to form the basis for forecasting the onset of diabetes within five years in Pima Indian women. Those variables have been found to be significant risk factors for diabetes among Pimas or other populations [113]:

1. Number of times pregnant
2. Plasma Glucose Concentration at 2 Hours in an Oral Glucose Tolerance Test (GTIT)
3. Diastolic Blood Pressure (mm Hg)
4. Triceps Skin Fold Thickness (mm)
5. 2-Hour Serum Insulin U<sub>h</sub>/ml)
6. Body Mass Index (Weight in kg / (Height in in))
7. Diabetes Pedigree Function
8. Age (years)#
9. Outcome (diabetes: binary)

The criteria applied were as follows.

- The subject was female.
- The subject was  $\geq 21$  year of age at the time of the index examination.
- Only one examination was selected per subject. That examination was one that revealed a nondiabetic Glucose Tolerance Test (GTIT) and met one of two criteria: 1) diabetes was diagnosed within five years of the examination; or 2) a GTIT performed five or more years later, failed to reveal diabetes mellitus.
- If diabetes occurred within one year of an examination, that examination was excluded from the study to remove those cases that were potentially easier to forecast from the forecasting model. In 75% of the excluded examinations, diabetes mellitus was diagnosed within six months.

The goal of the project is to develop a machine learning model that predicts diabetes with high accuracy. Business requirements, business goals or performance goals are not given. Full privacy needs to be guaranteed according to HIPAA privacy rule.<sup>11</sup> This dataset is problematic for privacy reasons because it relates to an identified tribe. Results of the analysis are associated with the tribe and can lead to discrimination. Publication of results would, most likely, require consent by the Pima people.

Development of data science solutions is generally conducted by multi-disciplinary teams consisting at least of domain experts and data scientists but usually includes software developers and functional experts, such as marketing, sales, product development and finance. Overcoming barriers given by technical complexities of machine learning and software engineering, modeling goals is an important means for shared understanding [132]. Conceptual modeling (as opposed to problem modeling) introduces various types of goals between actors: functional goals, non-functional goals [204, 228] and soft goals [148], both of which can be useful for this application. Functional goals of ML-based information systems are similar to procedural information systems while non-functional goals refer to expected system qualities and help to align understanding and work of all team members. Important goals for the development of machine learning solutions include: (1) data quality as key indicator for data engineering results; (2) accuracy and performance as key indicator for model training and model optimization; and (3) runtime behavior as key indicator for model integration and analytical decision making. Goal models can help to synchronize the work of both actors and even increase creativity [95] by clearly stating goals, events, dependencies and required resources as a means for overcoming barriers in development projects that leverage machine learning technologies. Soft goals explicate goals for the work relationship between actors. For the Pima project, medical researcher and data scientist are identified as actors. An initial goal model (cf. Figure 10) states that the data scientist assists the medical researcher in achieving the goal of finding dependencies for diabetes. The data scientist's main goal targets the collection of predictions while the medical researcher targets avoidance methods for diabetes. It mainly focusses on goals on domain level and data analytics level but abstracts from goals related to data access [150]. The principal-agent relationship between medical researcher and data scientist is modeled as a soft goal ("Be assisted"). The medical doctor is responsible for collecting unbiased data and the data scientist is responsible for data quality. Several goal dependencies between actors exist in data science projects. For instance, data engineering tries to achieve data quality requirements but this also needs to comply with the bias avoidance goal of the medical doctor. Identification of goal dependencies between actors is crucial for finding ML-based solutions for domain experts.

---

<sup>11</sup><https://www.hhs.gov/hipaa/index.html>

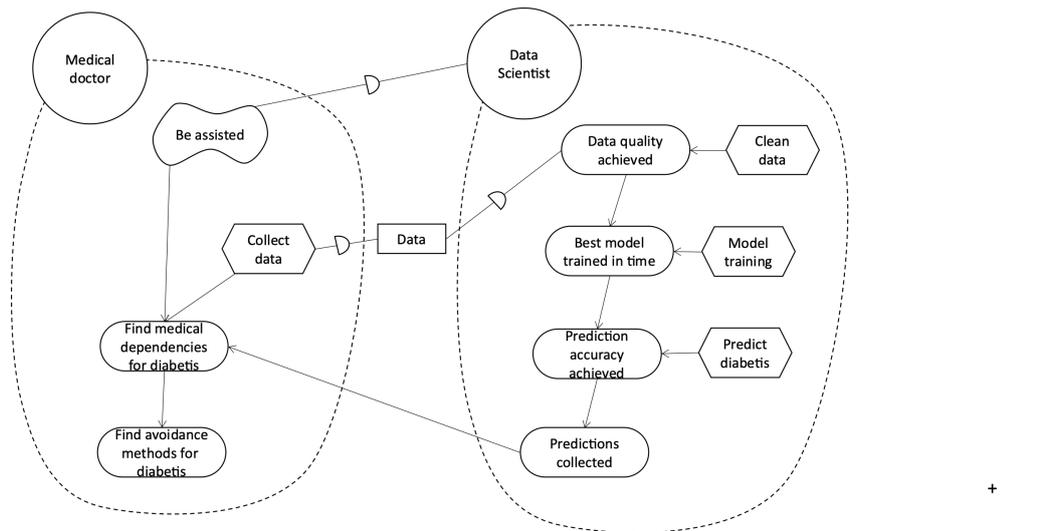


Figure 10: Initial goal model.

Several issues are identified for the initial goal model after consulting literature on medical ethics [90] and discussion with medical researchers, it becomes evident that beside pure functional goals related to avoiding diabetes, medical researchers also try to follow higher ethical principles including maintenance of integrity (cf. Figure 11). Data scientist do generally not account for goals that drive medical researchers. Therefore, linking data quality goals of data scientists with the bias avoidance goal of medical researchers is crucial for the success of the data science project. By making this explicit, both actors become aware of this relationship and can agree on measures that support goal achievement. A similar goal relationship exists between expected effect requirements on medical side and operationalization into performance requirements. Medical researchers perceive predictions as data that becomes input. This is translated into a requirement that predictions are not provided as graphics or performance measures but as tables with input and output data. Overall, the extended goal model expresses in more detail how medical researchers and data scientist intend to collaborate that reduces misunderstanding during project implementation.

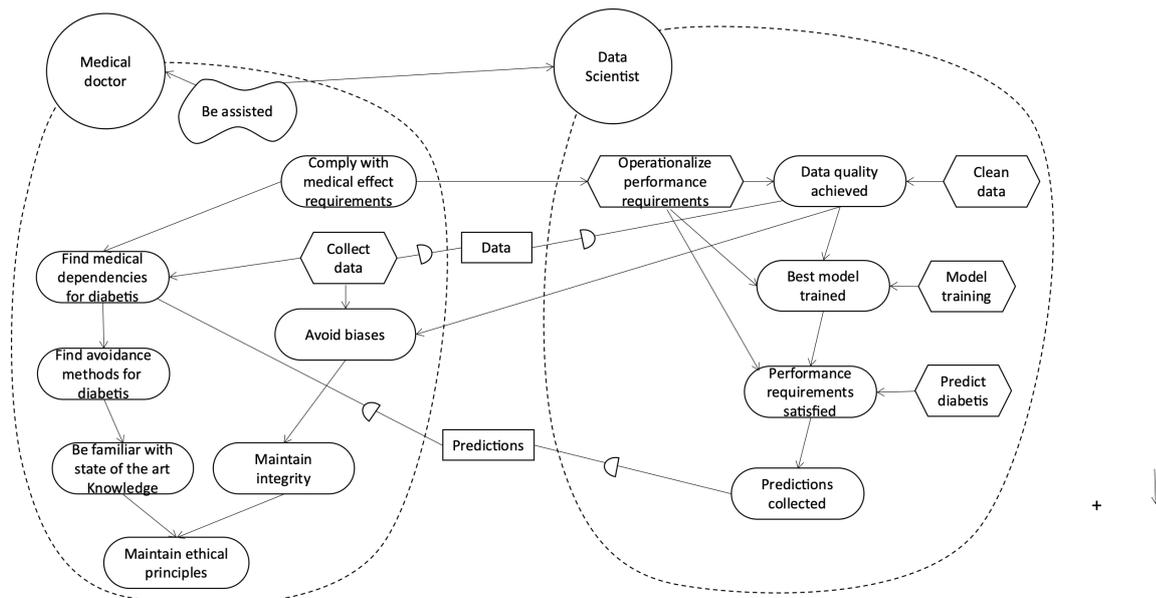


Figure 11: Extended goal model.

## Data collection

Pregnancies (number)	Glucose (Plasma glucose concentration at 2 hours in an oral glucose tolerance test)	Blood-Pressure (Diastolic blood pressure (mm Hg) )	Skin-Thickness (Triceps skin fold thickness (mm) )	Insulin (2-Hour serum insulin (mu U/ml) )	BMI (Body mass index (weight in kg/(height in m) <sup>2</sup> ) )	Diabetes-Pedigree-Function (Diabetes pedigree function )	Age (years)	Outcome (0 / 1)
6	148	72	35	0	33.6	627	50	1
1	85	66	29	0	26.6	351	31	0
8	183	64	0	0	23.3	672	32	1
1	89	66	23	94	28.1	167	21	0
0	137	40	35	168	43.1	2.288	33	1
5	116	74	0	0	25.6	201	30	0
3	78	50	32	88	31	248	26	1

Table 4: Dataset (selection).

	count	mean	std	min	25%	50%	75%	max
Pregnancies	768	3.84	3.36	0.000	1.00	3.00	6.00	17.00
Glucose	768	120.89	31.97	0.000	99.00	117.00	140.25	199.00
Blood-Pressure	768	69.10	19.35	0.000	62.00	72.00	80.00	122.00
Skin-Thickness	768	20.53	15.95	0.000	0.00	23.00	32.00	99.00
Insulin	768	79.79	115.24	0.000	0.00	30.50	127.25	846.00
BMI	768	31.99	7.88	0.000	27.30	32.00	36.60	67.10
Diabetes-Pedigree-Function	768	0.47	0.33	0.078	0.24	0.37	0.62	2.42
Age	768	33.24	11.76	21.000	24.00	29.00	41.00	81.00
Outcome	768	0.34	0.47	0.000	0.00	0.00	1.00	1.00

Table 5: Descriptive statistics.

This dataset has been provided via Kaggle.<sup>12</sup> The data set is accompanied with textual descriptions with units but provides no further semantical, structural, or contextual data requirements.

## Data engineering

This consisted of data exploration, data preparation, and feature engineering on the dataset.

*Data exploration* The dataset consists of 768 cases with 7 directly collected variables, one constructed feature and one outcome variable, as shown in Table 4 and the descriptive statistics given in Table 5.

Visualizations of the probability density functions show that some features follow a normal distribution (blood pressure, body mass index (BMI)), others are strongly skewed (DPF, age) or indicate lower quality (insulin and skin thickness) (cf. Figure 12).

<sup>12</sup><https://www.kaggle.com/uciml/pima-indians-diabetes-database>

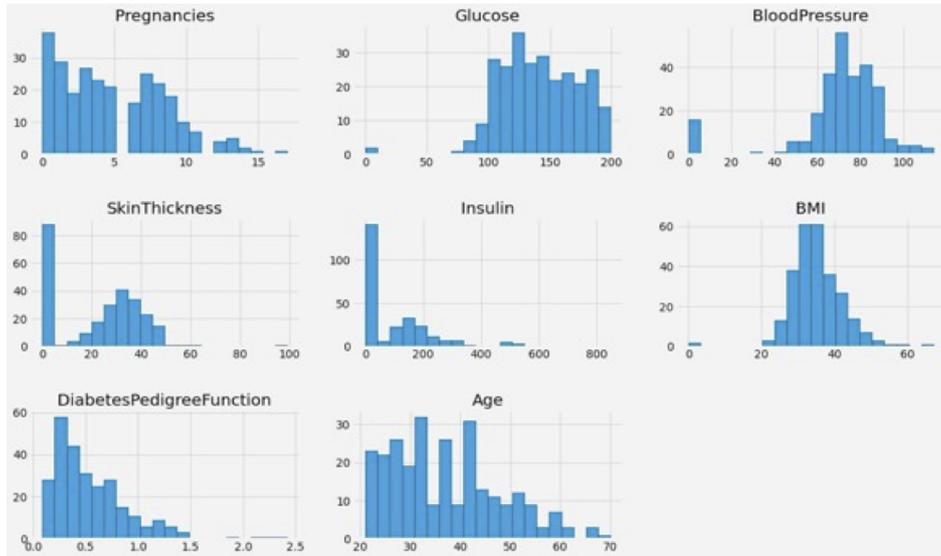


Figure 12: Probability density functions for Pima population diabetes study.

The data requirements are summarized below. Missing data has been found for BloodPressure (4.56%), SkinThickness (29.56%), and Insulin (48.7%). *All feature values must be positive*

- *Num\_of\_preg (number of pregnancies) must be recorded.*
- *Glucose (Plasma glucose concentration after 2 hours in an oral glucose tolerance test):* impaired glucose tolerance: between 7.8 mmol/L (140 mg/dL) and 11.1 mmol/L (200 mg/dL); levels  $\geq 11.1$  mmol/L at 2 hours confirm a diagnosis of diabetes.
- *BloodPressure (diastolic blood pressure (mm Hg)):* less than 120 mm Hg
- *SkinThickness (Triceps skin fold thickness (mm)):* no restrictions due to lack of knowledge
- *Insulin (2-Hour serum insulin (mu U/ml)):* categorization: 1-110, 111-150, 151-240, >241 [215]. Interpretation and derivation of data requirements requires domain expertise.
- *BMI (Body mass index (weight in kg/(height in m)<sup>2</sup>):* ranges: <18.5 underweight, 18.5 – 24.9 normal, 25.0 – 29.9 overweight, >30 obese. Ranges do not apply for athletes.
- *Age (in years):* oldest person is less than 122 years (age of oldest person ever recorded).

The correlation matrix indicates low interactions between features. This supports the assumption that features are independent and independently contribute to estimations (cf. Figure 13).

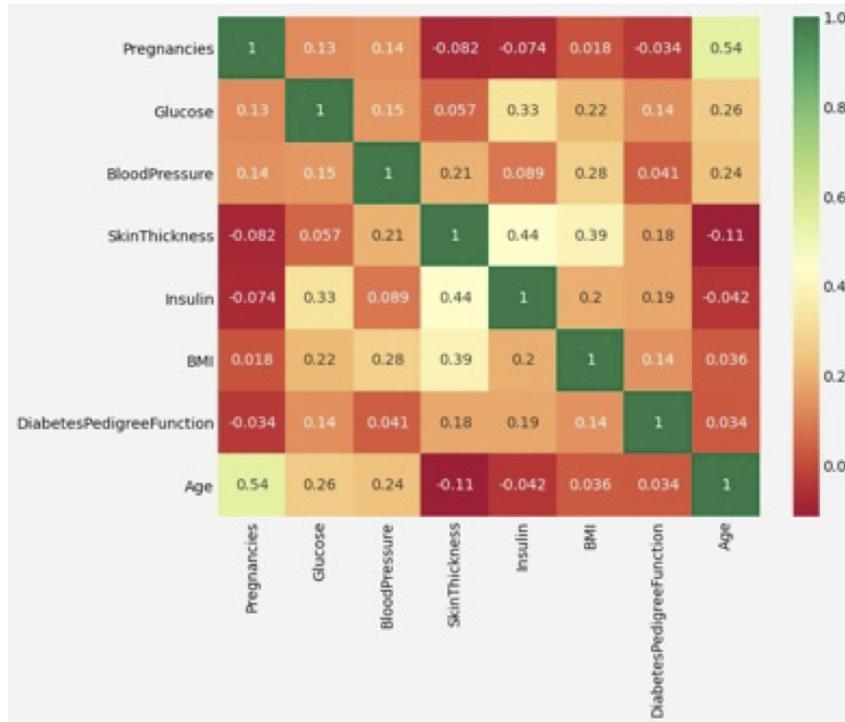


Figure 13: Correlation matrix for Pima population diabetes study.

### Data preparation

At a general level, diabetes is aligned with the diabetes diagnosis ontology (DDO) [57] that provides a rich set of concepts and relations on diabetes. DDO can be conceptually aligned [61] with the goal of the data science project; that is, the concept *diabetes diagnosis* in DDO with *diabetes* in the goal descriptions. Further variables can be extracted by analysis of the ontology. Ontology analysis of DDO shows that concept *patient* has a high centrality degree [25], with direct connection to *diabetes diagnosis* via *has diagnosis*. In order to infer independent features that can improve model performance, semantic paths can be analyzed on the basis of semantic distance in an ontology [165]. For example, *patient* is directly connected to a *diabetes symptom* with 89 associated concepts. Each concept is a candidate for enhancing the dataset.

Ontology embedding involves the following mapping to SNOMED CT<sup>13</sup>.

- *Pregnancies*: <http://purl.bioontology.org/ontology/SNOMEDCT/127362006>
- *Glucose*: <http://purl.bioontology.org/ontology/SNOMEDCT/434911002>
- *BloodPressure*: <http://purl.bioontology.org/ontology/SNOMEDCT/75367002>
- *SkinThickness*: <http://purl.bioontology.org/ontology/SNOMEDCT/247428002>
- *Insulin*: <http://purl.bioontology.org/ontology/SNOMEDCT/67866001>
- *BMI*: <http://purl.bioontology.org/ontology/SNOMEDCT/60621009>
- *Age*: <http://purl.bioontology.org/ontology/SNOMEDCT/397669002>

Additionally, DDO can be used to infer additional constraints on patients. For example, *patient* is directly related to a demographic with 9 concepts from which invariants on social status and social

<sup>13</sup><http://biportal.bioontology.org/ontologies/SNOMEDCT>

relationships can be inferred for better understanding and improving the dataset. Formal ontologies are often enriched by formal axiom specifications [89]. Invariants on datasets can be derived from formal axioms by axiom mining either directly or by propagation of axioms through ontologies; that is, axioms for relational algebra (e.g., symmetry, reflexivity, and inverse), composition of relationships, sub-relationships, and part-whole relationships [193]. This indicates that ontologies are rich sources for data exploration, improvement of data quality and data refinement. Less formal ontologies are provided by knowledge graphs that connect instance by analyzing large datasets [198]. Because knowledge graphs are often extracted from texts by text mining [83], instance-connection triples:

e.g.,  $\langle \text{DFKI}, \text{locatedAt}, \text{Saarbruecken} \rangle$

only provide weak support for ontological structures with concepts and relationships and, therefore, require knowledge graph refinement [162]. Because knowledge graphs resemble more social networks than ontologies, graph analytics uses techniques such as for finding centrality, communities, connectivity, and node similarity [102], as well as rule mining [93].

The dataset was collected before starting this data science project. Therefore, data requirements are described ex-post and data quality is assessed instead of declaring data quality requirements. UML and OCL are potential means for describing data requirements. Data features in the data set are only connected to person via sample numbers. UML representation increased understandability by declaring a relationship between a *person* and a *medical entry* (Figure 14).

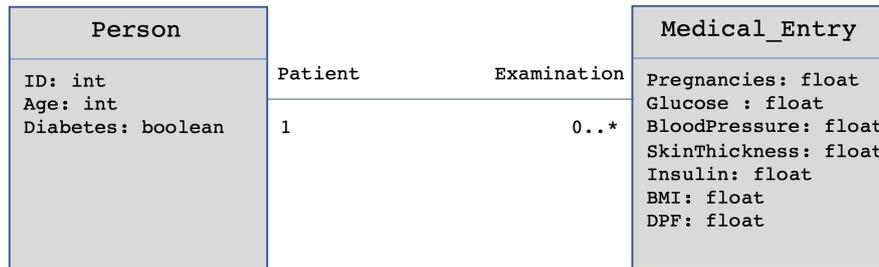


Figure 14: semantic relationships between entities.

Examples for object constraints (in OCL [174]) are as follows.

- C1: “Feature values of a person cannot be zero or negative“  
context Person  
inv: self.Examination->forAll(e | e.Glucose > 0)
- C2: “A patient’s age is between 15 and 120“  
context Person  
inv: self.Patient.age >= 15 and self.patient.age < 120
- C3: “A patient has diabetes if Glucose is above 200“  
context Person  
inv: self.Patient.Glucose >= 200 -> self.Patient.Diabetes=1

Given these data requirements, this dataset violates constraint C1 because it contains values of 0 for all features, but satisfies constraint C2. Constraint C3 raises an issue because no glucose value is above 200mg/dl, a strong indicator for diabetes.

Several data quality exist within this dataset. Glucose data values range from 0 to 199. This finding contradicts the data requirement that diabetes is diagnosed with a diabetes >200mg/dl.

Analyzing the data collection procedure [215] shows that all cases have been deleted with diabetes occurring within one year of an examination. However, the exact cut at 199mg/dl suggests that, instead, all values  $\geq 200$ mg/dl were deleted regardless of subsequent progression (cf. Figure 12). Furthermore, skin thickness and insulin are unreliable predictors due to lack of data. This is shown in Table 6.

		Number of pregnancies	Glucose	Blood Pressure	Skin Thickness	Insulin	BMI	Age
Accuracy	Believability	++	0	++	++	++	++	++
	Accuracy	++	-	++	++	++	++	++
	Objectivity	++	0	++	++	++	++	++
	Completeness	++	-	++	-	-	++	++
	Traceability	0	0	0	0	0	0	0
	Reputation	++	++	++	++	++	++	++
	Variety	0	0	0	0	0	0	0
Relevancy	Value-added	++	0	++	-	-	++	++
	Relevancy	++	++	++	++	++	++	
	Timeliness	0	0	0	0	0	0	0
	Ease of operation	++	++	++	0	0	++	++
	Appropriate amount of data	+	+	+	-	-	+	+
	Flexibility	+	+	+	+	+	+	+
Representation	Interpretability	++	++	++	++	++	++	++
	Ease of understanding	+	+	+	+	+	+	+
	Consistency	+	+	+	0	0	+	+
	Conciseness	+	+	+	+	+	+	+
Accessibility	Accessibility	++	++	++	++	++	++	++
	Cost-effectiveness	++	++	++	++	++	++	++
	Access security	++	++	++	++	++	++	++

Table 6: Data quality assessment of dataset.

Missing data is important for this dataset. For healthcare data, a popular imputation method is “multiple imputation using chained equations” (MICE) [216]. Simpler strategies are replacement by mean, median, or most frequent values. It is interesting to note that the most popular solution for this dataset on Kaggle (45,000 views out of 1067 unique solutions) uses a mix of mean and median without providing justification for doing so.

- Value 0 for *Glucose*, *BloodPressure*, *SkinThickness*, *Insulin*, *BMI* are replaced by **NAN**
- *Glucose*, *BloodPressure* NAN values replaced by **mean**
- *SkinThickness*, *BMI* NAN values replaced by **median**

#### Feature engineering

With the *Diabetes Pedigree Function* (DPF), this dataset also provides an engineered feature. In machine learning development projects this is either: provided by domain experts; or created during feature engineering in collaboration between domain experts and data scientists and then added to the dataset. Domain experts defined a Diabetes Pedigree Function (DPF) that provides a synthesis of the diabetes mellitus history in relatives and the genetic relationship of those relatives to the subject. The DPF uses information from parents, grandparents, full and half siblings, full and half aunts and uncles, and first cousins. It provides a measure of the expected genetic influence of affected and unaffected relatives on the subject’s eventual diabetes risk [215]:

$$DPF = \frac{\sum_i K_i (88 - ADM_i) + 20}{\sum_j K_j (ADM_j - 14) + 50} \quad (7)$$

i: all relatives i who had developed diabetes by the subject's examination date  
j: all relatives j who had not developed diabetes by the subject's examination date  
 $K_x$ : percent of genes shared by relative x and set at:

- 0.5 when the relative x is a parent or full sibling
- 0.25 when the relative x is a half sibling, grandparent, aunt or uncle
- 0.125 when the relative x is a half aunt, half uncle **or** cousin

$ADM_i$ : age in years of relative i when diabetes was diagnosed  
 $ACL_j$ : age in years of relative at the last non-diabetic examination  
88 / 14: maximum and minimum ages at which relatives developed DM  
20 / 50: moderating constants

This definition provides semantical and structural data requirements for the DPF variable thoroughly embedded into the domain of diabetes research (contextual data requirements). This definition could be developed into formal representations that support re-use and merging this dataset in other ML model training. Similar domain knowledge exists for glucose, blood pressure, skin thickness, and insulin.

### Model training.

Model training is governed by performance requirements as constraints for functional requirements. For building trust in medical treatment, prediction accuracy should be high with a high sensitivity (recall) value close to 100%; i.e., the percentage of missed positives should be small. Slightly less important is specificity (true negative rate); that is, not many people should receive treatment, even though they are healthy. These tradeoffs and thresholds must be based on domain knowledge. Reported sensitivity (0.78) and specificity scores (0.77) are taken as reference values.

At the beginning, a series of machine learning model types is trained and evaluated by applying default hyperparameter values. Accordingly, Gradient boosting performs best with an accuracy of more than 88%, recall (sensitivity) 87%, precision 88% and F1 value 88% [125]. Model training depends on selecting the best ML model type and leveraging domain knowledge. Selecting the best ML model type is the core of any model training phase, and requires application of different models and validation of model performance. Leveraging domain knowledge depends on the interaction between domain experts and data scientists. Domain knowledge leveraging supports: knowledge to data or data to knowledge. The first direction, *knowledge to data*, is used if a domain expert can express domain knowledge in a way that can be transformed into additional features. For instance, if a person is younger than 30 years with a plasma glucose level <120 mg/dl, then she is less likely to suffer from diabetes in the next 5 years. This heuristic rule can be expressed by an OCL rule with an added binary feature *kl1* initialized with 0 for all samples:

```
context Person
inv: self.Patient.age < 30 and
     self.medical_entry.glucose < 120
     -> forAll(e | e.kl1 = 1)
```

For *data to knowledge*, a data scientist analyzes the dataset and attempts to extract heuristic rules that will subsequently be evaluated by domain experts. For instance, if a data scientist finds support for a hypothetical rule that younger women with fewer pregnancies are less likely to suffer from diabetes in the coming years, then this is expressed as an OCL rule:

```
context Person
inv: self.Patient.age < 30 and
     self.medical_entry.pregnancies <= 6
     ->forAll(e | e.kl2 = 1)
```

Model	Accuracy	AUC	Recall / Sensitivity	Precision	F1
Light Gradient Boosting	0.89	0.94	0.84	0.86	0.85
Gradient boosting	0.89	0.95	0.81	0.85	0.83
Logistic regression	0.84	0.91	0.73	0.78	0.76
Support vector classifier	0.85	0.91	0.75	0.81	0.78
Decision tree	0.86	0.81	0.82	0.79	0.81
K nearest neighbors	0.80	0.88	0.59	0.77	0.67

Table 7: Model types.

If rule *kl2* is verified by domain experts, another binary variable is added to the dataset. From the *data to knowledge* strategy, an additional 16 binary features were found and added to the dataset.<sup>14</sup> These heuristic rules increased model performance from an accuracy of 0.73 for gradient boosting to 0.89 with a recall of 0.84 and a precision of 0.86. That is an increase of more than 20% above the initial model performance.

The following ML model types were trained as shown in Table 7.

### Model optimization.

Model optimization is also governed by performance requirements. Most machine learning model types have hyperparameters, such as the number of neighbors *k* that are used in KNN. Finding an optimal set of hyperparameters is a NP-complete search problem.<sup>15</sup> Even if an optimal set of hyperparameters could be computed, it is not possible to assess whether a viable solution has been identified because the solution might suffer from inductive fallacy. Best practices can be expressed by heuristic rules or knowledge graphs. Model optimization is a technical task similar to optimization of a database system by configuration of database management parameters. More research is needed to understand the impact on domain knowledge and on data science knowledge for optimizing ML models.

### Model integration.

After finalizing the ML model, it is integrated into the information system. Validation procedures can be used to assess compliance with business requirements, goal requirements, data requirements, legal requirements, ethical requirements, and functional and non-functional requirements. Field tests on newly collected data are used to build trust in information system performance. Empirical studies on information systems adoption, usability, cost effectiveness and other non-functional requirements are used for practical evidence of ML development results. The diabetes machine learning model was not integrated into an information system. Therefore, model integration is not relevant in this example.

### Analytical decision making.

The value of a diabetes information system lies in its potential to support medical workers. By sampling new data, medical workers receive predictions on the risk of patients suffering from diabetes in the future so that countermeasures can be recommended, even in real-time.

## 5 Machine Learning for Conceptual Modeling

Machine learning can also contribute to conceptual modeling. Many of the challenges of applying machine learning to conceptual modeling deal with knowledge generation. Here, the term *knowledge* refers generally to the constructs of a conceptual model. Knowledge challenges can be organized into three categories: incomplete knowledge, incorrect knowledge, and inconsistent

<sup>14</sup><https://www.kaggle.com/vincentlugat/pima-indians-diabetes-eda-prediction-0-906>

<sup>15</sup>Thank you to an anonymous reviewer for this suggestion.

knowledge. Incomplete knowledge includes missing or limited entities and/or relationships. Incorrect knowledge includes incorrect entity and relationship labels or incorrect facts (e.g., cardinalities). Inconsistent knowledge includes different labels for the same entity or merging entities with the same labels.

The first, and probably easiest to understand, is missing entities or relationships. Knowledge could be extracted to identify where there is incompleteness in modeling of an application domain, or potential missing relationships, which would require interaction between a domain expert and a conceptual modeling expert.

It is possible to infer what concepts and synonyms are extracted from a text. Then, the potential entity concepts can be used to create a graph that might indicate missing pieces or something that is incorrectly labeled (wrong entity label recognition). There could also be inconsistent relationships, or the potential to make incorrect inferences. Such basic kinds of research challenges are well-known. However, anchoring such inferences in knowledge graphs could support the combining of research on knowledge graphs with data analytics and conceptual modeling. We can consider analytics on text and how to extract conceptual modeling-like structures from it, as well as rule and graph mining.

- *General supervised learning*: linear regression, support vector machines, decision trees, random forests, boosting models, multi-layer perceptrons, deep neural networks [180]
- *Sequence learning*: recurrent neural networks (incl. LSTM) [92]
- *Generative learning*: generative adversarial neural networks [77]
- *Graph learning*: graph neural networks, recurrent graph neural networks, convolutional graph neural networks ([221])
- *Unsupervised learning*: KNN, k-means, PCA [86]
- *Reinforcement learning*: dynamic learning of agents through rewards gained from actions in environments [200].

Extracting knowledge structures from datasets by means of machine learning is a fast growing research field. Table 8 provides an non-exhaustive overview of using machine learning models for extracting different conceptual model structures. Associative rules is a robust technology for extracting relational knowledge. Approaches for estimating relationships between entities as link predictions are more sophisticated. Process discovery based on analysing log files is another promising area of research. These approaches, however, do not consider semantics. This might be why ontology extraction using machine learning is still restricted to ontology matching and mapping, although there are successful language translation systems that do not have explicit semantic representations [220].

ML → CM	Rules, anomalies and explanations	Semantic models (ERM, i*)	Ontologies	Process models (EPC, BPMN)
General Supervised learning	Associative rules [3]; rule extraction [12]		Ontology mapping and matching [54, 151]	Process discovery [11]; event abstraction [207]
Sequence learning	Rule extraction [146]	Named entity recognition [45]; link prediction [42]	Ontology matching [106]	
Generative learning		Link prediction (Qin et al. 2020)		
Graph learning		Link prediction [6, 51] Relational learning [152]		
Unsupervised learning	Anomaly detection [5]	Link prediction [122]	Concept learning Mi et al. 2020)	Event abstraction [207]
Reinforcement learning	Rule extraction [169]			

Table 8: Extraction of conceptual models with machine learning.

## 5.1 Text mining

Text mining, which discovers conceptual structures from unstructured sources [218], became popular with the increasing use of social media, such as Facebook and Twitter. Various natural language processing (NLP) methods are available for filtering keywords based on domain knowledge and domain lexica. Preprocessing of textual data removes stopwords and reduces words to word stems. Shallow parsing identifies phrases and recognizes named entities with ontology mappings [21]. At a semantic level, word sense disambiguation and identification of negations are processed. Negation is difficult to deal with because it might mean that entities, relationships, and larger conceptual structures are excluded, or do not exist. Entity linking by NLP methods leverage ontologies [219]. Particularly challenging is resolving references given by anaphora or by spatio-temporal prepositions. Text mining is a specific type of data mining that focuses on unstructured text. Mining association rules [3] is often used for extracting heuristic rules. Automatic entity classification by ML models, in particular, decision trees, is another rule mining technique [84].

## 5.2 Knowledge graphs

Machine learning has focused on extracting latent representations from euclidean data including images, text and videos. The need to process graph data has also important. Graph structures are natural representations in domains, such as e-commerce [225, 224], drug discovery [41] and chemistry [47], production and manufacturing, supply chain management [120] and network optimization [177]. Graphs are a natural means for explanation of opaque machine learning models. Therefore, they are used as input to machine learning and output from machine learning, providing an important perspective for how machine learning can support conceptual modeling. A knowledge graph “(i) describes real world entities and their interrelations, organized in a graph, (ii) defines possible classes and relationships of entities in a schema, (iii) allows for potentially interrelating arbitrary entities with each other; and (iv) covers various topical domains.” [162].

Adding typed links between data from different sources is an active research topic on semantic technologies [23]. Google, for example, extended semi-automatic annotation procedures by automatic extraction of knowledge graphs [189] based on existing sources, such as dbpedia [10], YAGO [199] or WordNet [142].

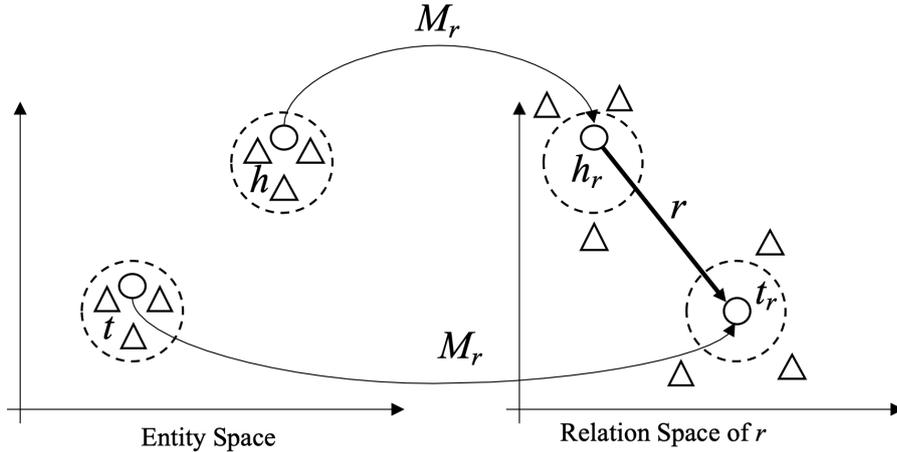


Figure 15: Link prediction with TransR. Adapted from [124].

Knowledge graphs extract named connections between instances, called *link prediction between entities* [124], which provide initial support for partial conceptual models. With large datasets, the quality of triples found by knowledge graphs is often low; that is, many links are tautological, or even meaningless. Entity resolution, collective classification, and link prediction can be used to construct consistent knowledge graphs based on probabilistic soft logic [170]. A standard approach is to transform data into vector spaces resembling principal component analysis (PCA). For instance, TransR proposed by Lin et al. [124] projects entities  $h$  (head) and  $t$  (tail) from an entity space into a  $r$ -relation space by a mapping function  $M_r$  that supports  $h_r + r \approx t_r$  and, thus, finds a set of entities  $t_r$  that fulfil a triple  $\langle h, r, t \rangle$  (cf. Figure 15).  $M_r$  is a vector embedding model that is trained on a loss function that minimizes a distance  $d(h+r, t)$  of ground truth and  $d(h+r, t)$  estimations. ML models based on embeddings are generative models (cf. section 3.1.2) that encode entities and relationships in vector spaces, make predictions into the input space, called *decoding*, and measure the reconstruction error as an indicator for model performance. Hence, graph embeddings are powerful models used for various graph analytical tasks.

*Graph analytics* use social network theory by analyzing distances and directed connections (ties) of knowledge graphs to support semantic annotations, such as similarity, centrality, community, and paths [214]. Similar to data engineering, knowledge graphs are explored and missing elements predicted (e.g., link prediction, completion and correction) [114]. When using knowledge graphs to support machine learning, graphs are embedded into multidimensional vector spaces by preserving a proximity measure defined on knowledge graph,  $G$  [79]. Graph embedding (GE) reduces the dimensionality of a graph. Currently, autoencoder models are used for embedding graph nodes and preserving non-linear dependencies [211]. Graph embeddings are used for link prediction and node classification and, thus, can be indirectly used for concept classification, identification of relationships, and ontology learning [136]. There is a growing interest in graph neural networks, which consider graphs as input, instead of tabular data. Graph neural networks (GNN) are extensions of graph embedding with emphasis on deep learning architectures, such as recurrent neural networks, convolutional neural networks and autoencoders [222].

Graph embeddings (GE) and graph neural networks (GNN) require graph input that meets quality requirements. Applications based on GE and GNN need to satisfy all requirement types, in the ML development cycle. This makes conceptual modeling methods and tools important assets for the development of GE/GNN-based information systems. Research on knowledge graphs provide important technologies for research on linked data and ontologies in general [93] and conceptual modeling in particular, including reasoning and querying over contextual data, and rule and axiom mining.

## 6 Summary and roadmap

In this paper, we align conceptual modeling and machine learning in both directions. Due to the early stages of research related to this pairing, challenges remain. Conceptual modeling was motivated by research on relational databases [43] and procedural programming whereas machine learning is a child of statistics and linear algebra. Although it is clear that conceptual modeling can support data management of machine learning, many challenges remain for supporting the development of model architectures, model training, model testing, model optimization, deployment, and maintenance in information systems. For example, deep convolutional neural network architectures consist of multiple layers with different sizes taking different roles, such as convolution and pooling layers, and application of different convolutional kernels [115].

Many technical research issues emerge including:

- Use of data ontologies and design patterns for data engineering
- Alignment of data engineering with databases and big data stores
- Models for mining data streams
- Design patterns for model architectures
- Process models for model development
- Performance models for model development.

Because ML models are central to services delivered by information systems, they also require alignment with enterprise architectures. Research issues include the following:

- Frameworks for aligning model architectures and service architectures and enterprise architectures
- Frameworks for alignment of performance metrics and key performance indicators.

Decision making, which relies on machine learning must consider:

- The quality and performance models for data-driven decision making (cf. [168])
- Conceptual modeling in real-time data-driven decision making with batch and streaming data.

Conceptual modeling is well positioned when it comes to structuring requirements for complex systems, including ML-based systems. Recent proposals structure data science development by views and goal models [127, 150]. In the direction of machine learning for conceptual modeling, research issues are in their infancies. Knowledge graphs extracted from data are promising areas with clear connections to conceptual modeling. Merging knowledge graphs with formal ontologies is a challenging, but also promising, research topic.

## 7 Conclusion

The fields of machine learning and conceptual modeling have been active areas of research for a long time, making it reasonable to expect that it might be advantageous to explore how one might complement the other. This paper has identified possible synergies between the machine learning and conceptual modeling and proposed a framework for conceptual modeling for machine learning. Conceptual modeling can be helpful in supporting the design and development phases of a machine learning-based information system. Feature sets, especially, must be consistent so that data scientists can create valid solutions. This can be best accomplished by including domain

knowledge, as represented by conceptual models. Inversely, machine learning techniques are very successful at obtaining or scraping large amounts of data, which can be used to identify concepts and patterns that could be useful for inclusion in conceptual models. There are, of course, many challenges related to incorrect, incomplete, or inconsistent knowledge. Nevertheless, it is feasible to pair conceptual modeling with machine learning. This paper has identified some of the challenges inherent in achieving this pairing in an attempt to lay the groundwork for future research on combining conceptual modeling and machine learning.

## Acknowledgements

This paper was based on a keynote presentation given by the first author at the *International Conference on Conceptual Modeling*. The authors wish to thank Peter Chen and Carson Woo, and Oscar Pastor for their support of this paper, Iaroslav Shcherbatyi for sharing his technical expertise in machine learning and Michael Schrefl for identifying this topic. We also thank the anonymous reviewers and Roman Lukyanenko for provided valuable insights and comments.

## References

- [1] G. et al. Aad, *The atlas experiment at the cern large hadron collider*, Journal of Instrumentation **3** (2008), no. S08003.
- [2] A. Adadi and M. Berrada, *Peeking inside the black-box: A survey on explainable artificial intelligence (xai)*, IEEE access **6** (2018), 52138–52160.
- [3] R. Agrawal, T. Imieliński, and A. Swami, *Mining association rules between sets of items in large databases*, Proceedings of the 1993 ACM SIGMOD international conference on Management of data, 1993, pp. 207–216.
- [4] H. Aguinis, R. K. Gottfredson, and H. Joo, *Best-practice recommendations for defining, identifying, and handling outliers*, Organizational Research Methods **16** (2013), no. 2, 270–301.
- [5] S. et al. Ahmad, *Unsupervised real-time anomaly detection for streaming data*, Neurocomputing **262** (2017), 134–147.
- [6] Mohammad Al Hasan, Vineet Chaoji, Saeed Salem, and Mohammed Zaki, *Link prediction using supervised learning*, SDM06: workshop on link analysis, counter-terrorism and security, 2006, pp. 798–805.
- [7] D. Amyot, *Evaluating goal models within the goal-oriented requirement language*, **25** (2010), no. 8, 841–877.
- [8] E. Arias, *Transcending the individual human mind—creating shared understanding through collaborative design*, ACM Transactions on Computer-Human Interaction (TOCHI) **7** (2000), no. 1, 84–113.
- [9] A. B. et al. Arrieta, *Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai*, Information Fusion **58** (2020), 82–115.
- [10] S. et al. Auer, *Dbpedia: A nucleus for a web of open data*, International semantic web conference (Berlin), Springer, 2007, pp. 722–735.
- [11] A. Augusto, *Automated discovery of process models from event logs: Review and benchmark*, IEEE transactions on knowledge and data engineering **31** (2018), no. 4, 686–705.

- [12] N. Barakat and A. P. Bradley, *Rule extraction from support vector machines: a review*, *Neurocomputing* **74** (2010), no. 1-3, 178–190.
- [13] S. Barocas and A. D. Selbst, *Big data’s disparate impact*, *Calif. L. Rev* **104** (2016), 671.
- [14] C. et al. Batini, *Methodologies for data quality assessment and improvement*, *ACM computing surveys* **41** (2009), no. 3, 1–52.
- [15] H. Belani, M. Vukovic, and Z. Car, *Requirements engineering challenges in building ai-based complex systems*, *IEEE 27th International Requirements Engineering Conference Workshops (REW).*, 2019, pp. 252–255.
- [16] Y. Bengio, A. Courville, and P. Vincent, *Representation learning: A review and new perspectives*, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35** (2013), no. 8, 1798–1828.
- [17] O. et al. Berger-Tal, *The exploration-exploitation dilemma: a multidisciplinary framework*, *PloS one* **9** (2014), no. 4, 95693.
- [18] P. Bhat, I. Shcherbatyi, and W. Maass, *Automated learning of user preferences for selection of high quality 3d designs*, *Procedia CIRP* **84**, 2019, pp. 814–819.
- [19] A. et al. Bibal, *Legal requirements on explainability in machine learning*, *Artificial Intelligence and Law* **29** (2020), 149–169.
- [20] F. et al. Biessmann, *Datawig: Missing value imputation for tables*, *Journal of Machine Learning Research* **20** (2019), no. 175, 1–6.
- [21] D. Bikel, R. Schwartz, and R. Weischedel, *An algorithm that learns what’s in a name*, *Machine Learning* **34** (1999), 211–231.
- [22] C. M. Bishop, *Pattern recognition and machine learning*, Springer, 2006.
- [23] C. Bizer, T. Heath, and T. Berners-Lee, *Linked data: The story so far*, *Semantic services, interoperability and web applications: emerging concepts*, IGI global, 2011, pp. 205–227.
- [24] A. L. Blum and P. Langley, *Selection of relevant features and examples in machine learning*, *Artificial Intelligence* **97** (1997), no. 1-2, 245–271.
- [25] S. P. Borgatti, *Centrality and network flow*, *Social networks* **27** (2005), no. 1, 55–71.
- [26] Nick Bostrom and Eliezer Yudkowsky, *The ethics of artificial intelligence*, *The Cambridge handbook of artificial intelligence* **1** (2014), 316–334.
- [27] L. et al. Breiman, *Classification and regression trees*, Wadsworth, Belmont, CA, 1984.
- [28] T. B. et al. Brown, *Language models are few-shot learners*, 2020.
- [29] E. Brynjolfsson and K. McElheran, *The rapid adoption of data-driven decision-making*, *American Economic Review* **106** (2016), 133–139.
- [30] A. et al. Bucchiarone, *Grand challenges in model-driven engineering: an analysis of the state of the research*, *Software and Systems Modeling* **19** (2020), no. 1, 5–13.
- [31] M. C. Buiten, *Towards intelligent regulation of artificial intelligence*, *European Journal of Risk Regulation* **10** (2019), no. 1, 41–59.
- [32] A. et al. Burton-Jones, *A semiotic metrics suite for assessing the quality of ontologies*, *Data & Knowledge Engineering* **55** (2005), no. 1, 84–102.

- [33] L. Cai and Y. Zhu, *The challenges of data quality and data quality assessment in the big data era*, Data science journal **14** (2015), no. 2, 1–10.
- [34] A. I. Canhoto and F. Clear, *Artificial intelligence and machine learning as business tools: A framework for diagnosing value destruction potential*, Business Horizons **63** (2020), no. 2, 183–193.
- [35] C. et al. Capiello, *Data ecosystems: sovereign data exchange among organizations (dagstuhl seminar 19391)*, in *Dagstuhl Reports Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik*, 2020.
- [36] A. Castellanos, *Improving machine learning performance based on conceptual modeling guidelines*, AAAI-Make 2021: Combining Machine Learning and Knowledge Engineering, 2021.
- [37] J. Castro, M. Kolp, and J. Mylopoulos, *Towards requirements-driven information systems engineering: the tropos project*, Information systems **27** (2002), no. 6, 365–389.
- [38] S. Chakraborty, S. Sarker, and S. Sarker, *An exploration into the process of requirements elicitation: A grounded approach*, Journal of the association for information systems **11** (2010), no. 4, 1.
- [39] M. Chambers and T. W. Dinsmore, *Advanced analytics methodologies: Driving business value with analytics*, Pearson Education, 2014.
- [40] H. Chen, R. H. Chiang, and V. C. Storey, *Business intelligence and analytics: From big data to big impact*, MIS quarterly **36** (2012), no. 4, 1165–1188.
- [41] H. et al. Chen, *The rise of deep learning in drug discovery*, Drug discovery today **23** (2018), no. 6, 1241–1250.
- [42] J. et al. Chen, *E-lstm-d: A deep learning framework for dynamic network link prediction*, IEEE Transactions on Systems, Man, and Cybernetics: Systems **51** (2019), no. 6, 3699–3712.
- [43] P. P. S. Chen, *The entity-relationship model—toward a unified view of data*, ACM transactions on database systems (TODS **1** (1976), no. 1, 9–36.
- [44] T. Chen and C. Guestrin, *Xgboost: A scalable tree boosting system*, 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 785–794.
- [45] J. P. Chiu and E. Nichols, *Named entity recognition with bidirectional lstm-cnns*, Transactions of the Association for Computational Linguistics **4** (2016), 357–370.
- [46] L. et al. Chung, *Non-functional requirements in software engineering*, Science & Business Media, Springer, 2012.
- [47] C. W. et al. Coley, *A graph-convolutional neural network model for the prediction of chemical reactivity*, Chemical science **10** (2019), no. 2, 370–377.
- [48] T. H. Davenport and D. Patil, *Data scientist*, Harvard business review **90** (2012), no. 5, 70–76.
- [49] J. Dean, *The deep learning revolution and its implications for computer architecture and chip design*, IEEE International Solid-State Circuits Conference-(ISSCC), 2020, pp. 8–14.
- [50] L. M. Delcambre, *A reference framework for conceptual modeling*, International Conference on Conceptual Modeling, Springer, 2018, pp. 27–42.
- [51] T. et al. Dettmers, *Convolutional 2d knowledge graph embeddings*, AAAI Conference on Artificial Intelligence (S. McIlraith and K. Weinberger, eds.), 2018.

- [52] J. et al. Devlin, *Bert: Pre-training of deep bidirectional transformers for language understanding*, Tech. Report 1810.04805, arXiv, 2018.
- [53] J. et al. DeYoung, *A benchmark to evaluate rationalized nlp models.*, 2019.
- [54] A Doan, Jayant Madhavan, Pedro Domingos, and Alon Halevy, *Ontology matching: A machine learning approach. staab, studer (eds) handbook on ontologies in information systems*, Springer, 2003.
- [55] P. Domingos, *A few useful things to know about machine learning*, Communications of the ACM **55** (2012), no. 10, 78–87.
- [56] B. et al. Efron, *Least angle regression*, The Annals of statistics **32** (2004), no. 2, 407–499.
- [57] S. El-Sappagh and F. Ali, *Ddo: a diabetes mellitus diagnosis ontology*, Applied Informatics **3** (2016), no. 1, 5.
- [58] J. et al. Elson, *Asirra: a captcha that exploits interest-aligned manual image categorization*, 14th ACM conference on Computer and communications security, 2007, pp. 366–374.
- [59] D. W. Embley and S. W. Liddle, *Big data—conceptual modeling to the rescue*, International Conference on Conceptual Modeling (p), Springer, 2013, pp. 1–8.
- [60] H. Estrada, A. Martínez, and O. Pastor, *Goal-based business modeling oriented towards late requirements generation*, International Conference on Conceptual Modeling, Springer, 2003, pp. 277–290.
- [61] J. Euzenat and P. Shvaiko, *Ontology matching*, Springer, Heidelberg, 2007.
- [62] M. et al. Everingham, *The pascal visual object classes (voc) challenge*, International Journal of Computer Vision **88** (2010), no. 2, 303–338.
- [63] S. et al. Fan, *A process ontology based approach to easing semantic ambiguity in business process modeling*, Data & Knowledge Engineering **102** (2016), 57–77.
- [64] L. Fausett, *Fundamentals of neural networks*, Prentice Hall, Englewood Cliffs, NJ, 1994.
- [65] M. et al. Feurer, *Auto-sklearn 2.0: The next generation*, (2020).
- [66] U. Frank, *Domain-specific modeling languages: requirements analysis and design guidelines*, Domain engineering, Springer, p, 2013, pp. 133–157.
- [67] Y. Freund and R. E. Schapire, *A decision-theoretic generalization of on-line learning and an application to boosting*, Journal of computer and system sciences **55** (1997), no. 1, 119–139.
- [68] J. H. Friedman, *Greedy function approximation: a gradient boosting machine*, Annals of statistics **29** (2001), no. 5, 1189–1232.
- [69] E. Gamma, *Elements of reusable object-oriented software*, Addison-Wesley, 1995.
- [70] Javier García and Diogo Shafie, *Teaching a humanoid robot to walk faster through safe reinforcement learning*, Engineering Applications of Artificial Intelligence **88** (2020), 103360.
- [71] A. Geron, *Hands-on machine learning with scikit-learn and tensorflow*, O’Reilly, 2017.
- [72] F. A. Gers and J. Schmidhuber, *Lstm recurrent networks learn simple context free and context sensitive languages*, IEEE Transactions on Neural Networks **12** (2001), no. 6, 1333–1340.
- [73] S. Ghanavati, D. Amyot, and A. Rifaut, *Legal goal-oriented requirement language (legal grl) for modeling regulations*, 6th international workshop on modeling in software engineering, 2014, pp. 1–6.

- [74] G. Giachetti, B. Marín, and O. Pastor, *Using uml profiles to interchange dsml and uml models*, Third International Conference on Research Challenges in Information Science, IEEE, 2009, pp. 385–394.
- [75] Martin Glinz, *On non-functional requirements*, 15th IEEE international requirements engineering conference (RE 2007), IEEE, 2007, pp. 21–26.
- [76] A. Goldstein, *Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation*, Journal of Computational and Graphical Statistics **24** (2015), no. 1, 44–65.
- [77] I. et al. Goodfellow, *Generative adversarial nets*, 2014.
- [78] ———, *Deep learning*, MIT press, Cambridge, 2016.
- [79] P. Goyal and E. Ferrara, *Graph embedding techniques, applications, and performance: A survey*, Knowledge-Based Systems **151** (2018), 78–94.
- [80] T. R. Gruber, *A translation approach to portable ontology specifications*, Knowledge acquisition **5** (1993), no. 2, 199–220.
- [81] Nicola Guarino, Giancarlo Guizzardi, and John Mylopoulos, *On the philosophical foundations of conceptual models*, Information Modelling and Knowledge Bases **31** (2020), no. 321, 1.
- [82] G. Guizzardi, *Towards ontological foundations for conceptual modeling: The unified foundational ontology (ufo) story*, Applied ontology **10** (2015), no. 3-4, 259–271.
- [83] V. Gupta and G. S. Lehal, *A survey of text mining techniques and applications*, Journal of emerging technologies in web intelligence **1** (2009), no. 1, 60–76.
- [84] J. Han and M. Kamber, *Data mining concepts and techniques*, Morgan Kaufmann, 2000.
- [85] W. Hasselbring, *Information system integration*, Communications of the ACM **43** (2000), no. 6, 32–38.
- [86] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*, Science & Business Media, Springer, 2009.
- [87] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, *Deep residual learning for image recognition*, Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
- [88] J. Heaton, *An empirical analysis of feature engineering for predictive modeling*, SoutheastCon (2016), 1–6.
- [89] H. Herre, *General formal ontology (gfo): A foundational ontology for conceptual modelling*, Theory and applications of ontology: computer applications, Dordrecht. p, Springer, 2010, pp. 297–345.
- [90] A. B. Hill, *Medical ethics and controlled trials*, British medical journal **1** (1963), no. 5337, 1043.
- [91] Christopher J Hillar and Lek-Heng Lim, *Most tensor problems are np-hard*, Journal of the ACM (JACM) **60** (2013), no. 6, 1–39.
- [92] S. Hochreiter and J. Schmidhuber, *Long short-term memory*, Neural Computation **9** (1997), no. 8, 1735–1780.
- [93] A. et al. Hogan, *Knowledge graphs*, 2020.

- [94] J. Horkoff, *Non-functional requirements for machine learning: Challenges and new directions*, IEEE 27th International Requirements Engineering Conference (RE), 2019, pp. 386–391.
- [95] J. Horkoff, N. Maiden, and J. Lockerbie, *Creativity and goal modeling for software requirements engineering*, Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition, 2015, pp. 165–168.
- [96] K. Hornik, M. Stinchcombe, and H. White, *Multilayer feedforward networks are universal approximators*, Neural networks **2** (1989), no. 5, 359–366.
- [97] A. Hotho, A. Nürnberger, and G. Paaß, *A brief survey of text mining*, Ldv Forum **20** (2005), no. 1, 19–62.
- [98] M.-H. Huang and R. T. Rust, *A strategic framework for artificial intelligence in marketing*, Journal of the Academy of Marketing Science **49** (2021), no. 1, 30–50.
- [99] P. J. Huber, *A robust version of the probability ratio test*, The Annals of Mathematical Statistics (1965), 1753–1758.
- [100] F. et al. Hutter, *Sequential model-based optimization for general algorithm configuration*, International conference on learning and intelligent optimization, 2011, pp. 507–523.
- [101] F. Härer and H.-G. Fill, *Past trends and future prospects in conceptual modeling—a bibliometric analysis*, International Conference on Conceptual Modeling, 2020.
- [102] A. et al. Iosup, *Ldbc graphalytics: A benchmark for large-scale graph analysis on parallel and distributed platforms*, Proceedings of the VLDB Endowment **9** (2016), no. 13, 1317–1328.
- [103] S. Ishii, W. Yoshida, and J. Yoshimoto, *Control of exploitation–exploration meta-parameter in reinforcement learning*, Neural networks **15** (2002), no. 4-6, 665–687.
- [104] H. Jaakkola and B. Thalheim, *Sixty years—and more—of data modelling*, Information Modelling and Knowledge Bases XXXII **333** (2021), 56.
- [105] J. Jaffar and M. J. Maher, *Constraint logic programming: A survey*, The journal of logic programming **19** (1994), 503–581.
- [106] C. Jiang and X. Xue, *Matching biomedical ontologies with long short-term memory networks*, IEEE International Conference on Bioinformatics and Biomedicine (BIBM), IEEE, 2020, pp. 2484–2489.
- [107] H. Jin, Q. Song, and X. Hu, *Auto-keras: An efficient neural architecture search system*, 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2019, pp. 1946–1956.
- [108] A. Jobin, M. Ienca, and E. Vayena, *The global landscape of ai ethics guidelines*, Nature Machine Intelligence **1** (2019), no. 9, 389–399.
- [109] M. I. Jordan and T. M. Mitchell, *Machine learning: Trends, perspectives, and prospects*, Science **349** (2015), no. 6245, 255–260.
- [110] J. et al. Kim, *Ds4c patient policy province dataset: a comprehensive covid-19 dataset for causal and epidemiological analysis*, Tech. report.
- [111] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, arXiv preprint, 2014.
- [112] D. P. Kingma and M. Welling, *Auto-encoding variational Bayes*, arXiv preprint, 2013.
- [113] W. C. Knowler, *Diabetes incidence and prevalence in pima indians: a 19-fold greater incidence than in rochester, minnesota*, Am J Epidemiol **108** (1978), 497–505.

- [114] D. et al. Koller, *Introduction to statistical relational learning*, MIT Press, 2007.
- [115] A. Krizhevsky, I. Sutskever, and G. E. Hinton, *Imagenet classification with deep convolutional neural networks*, *Advances in neural information processing systems* **25** (2012), 1097–1105.
- [116] ———, *Imagenet classification with deep convolutional neural networks*, *Communications of the ACM* **60** (2017), no. 6, 84–90.
- [117] L. A. Kurgan and P. Musilek, *A survey of knowledge discovery and data mining process models*, *The Knowledge Engineering Review* **21** (2006), no. 1, 1–24.
- [118] O. Kwon, N. Lee, and B. Shin, *Data quality management, data usage experience and acquisition intention of big data analytics*, *International journal of information management* **34** (2014), no. 3, 387–394.
- [119] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, *Deep learning*, *nature* **521** (2015), no. 7553, 436–444.
- [120] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu, *Diffusion convolutional recurrent neural network: Data-driven traffic forecasting*. *arxiv*, arXiv preprint arXiv:1707.01926 (2017).
- [121] F. et al. Liang, *A survey on big data market: Pricing, trading and protection*, *IEEE Access* **6** (2018), 15132–15154.
- [122] R. N. Lichtenwalter, J. T. Lussier, and N. V. Chawla, *New perspectives and methods in link prediction*, 16th ACM SIGKDD international conference on Knowledge discovery and data mining, 2010, pp. 243–252.
- [123] L. H. C. Lima, *An analysis of the collaboration network of the international conference on conceptual modeling at the age of 40*, *Data & Knowledge Engineering*, p. 101866, 2020.
- [124] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu, *Learning entity and relation embeddings for knowledge graph completion*, Twenty-ninth AAAI conference on artificial intelligence, 2015.
- [125] J. Lindström and J. Tuomilehto, *The diabetes risk score: a practical tool to predict type 2 diabetes risk*, *Diabetes care* **26** (2003), no. 3, 725–731.
- [126] R. Lukyanenko, J. Parsons, and V. C. Storey, *Modeling matters: Can conceptual modeling support machine learning?*, AIS SIGSAND, 2018, pp. 1–12.
- [127] R. et al. Lukyanenko, *Using conceptual modeling to support machine learning*, International Conference on Advanced Information Systems Engineering, 2019, pp. 170–181.
- [128] Scott M Lundberg and Su-In Lee, *A unified approach to interpreting model predictions*, Proceedings of the 31st international conference on neural information processing systems, 2017, pp. 4768–4777.
- [129] Y. et al. Luo, *Multivariate time series imputation with generative adversarial networks*, 32nd International Conference on Neural Information Processing Systems, 2018, pp. 1596–1607.
- [130] W. Maass, *Data-driven meets theory-driven research in the era of big data: Opportunities and challenges for information systems research*, *Journal of the Association for Information Systems* **19** (2018), no. 12, 1253–1273.
- [131] W. Maass and S. Janzen, *Pattern-based approach for designing with diagrammatic and propositional conceptual models*, International Conference on Design Science Research in Information Systems, Springer, 2011, pp. 192–206.

- [132] W. Maass, V. C. Storey, and T. Kowatsch, *Effects of external conceptual models and verbal explanations on shared understanding in small groups*, International Conference on Conceptual Modeling, Springer, 2011, pp. 92–103.
- [133] W. Maass, V. C. Storey, and R. Lukyanenko, *From mental models to machine learning via conceptual models*, Enterprise, Business-Process and Information Systems Modeling (EMM-SAD) (A. Augusto et al., eds.), Springer, 2021, pp. 293–300.
- [134] W. Maass and U. Varshney, *Design and evaluation of ubiquitous information systems and use in healthcare*, Decision Support Systems **54** (2012), no. 1, 597–609.
- [135] A. C. et al. Marcén, *Traceability link recovery between requirements and models using an evolutionary algorithm guided by a learning to rank algorithm: Train control and management case*, Journal of Systems and Software **163** (2020), 110519.
- [136] J. L. Martínez-Rodríguez, I. López-Arévalo, and A. B. Rios-Alvarado, *Openie-based approach for knowledge graph construction from text*, Expert Systems With Applications **113** (2018), 339–355.
- [137] H. C. Mayr and B. Thalheim, *The triptych of conceptual modeling*, Software and Systems Modeling **20** (2020), 7–24.
- [138] M. McDaniel and V. C. Storey, *Evaluating domain ontologies: clarification, classification, and challenges*, ACM Computing Surveys (CSUR) **52** (2019), no. 4, 1–44.
- [139] M. McDaniel, V. C. Storey, and V. Sugumaran, *Assessing the quality of domain ontologies: Metrics and an automated ranking system*, Data & Knowledge Engineering **115** (2018), 32–47.
- [140] N. et al. Melville, *Information technology and organizational performance: An integrative model of its business value*, MIS quarterly **28** (2004), no. 2, 283–322.
- [141] R. Miikkulainen, *Topology of a neural network*, Encyclopedia of Machine Learning, Boston, MA, Editors Springer, 2011.
- [142] G. A. Miller, *Wordnet: a lexical database for english*, Communications of the ACM **38** (1995), 11.
- [143] M. Minsky and S. A. Papert, *Perceptrons: An introduction to computational geometry*, MIT Press, 2017.
- [144] V. et al. Mnih, *Human-level control through deep reinforcement learning*, Nature **518** (2015), no. 7540, 529–533.
- [145] B. Motik, I. Horrocks, and U. Sattler, *Bridging the gap between owl and relational databases*, Journal of Web Semantics **7** (2009), no. 2, 74–89.
- [146] W. J. Murdoch and A. Szlam, *Automatic rule extraction from long short term memory networks*, arXiv preprint, 2017.
- [147] J. Mylopoulos, *Conceptual modeling and telos*, Conceptual Modeling, Databases, and CASE: An Integrated View of Information Systems Development (P. Loucopoulos and R. Zicari, eds.), Editors John Wiley & Sons, 1992, pp. 49–68.
- [148] J. Mylopoulos, L. Chung, and B. Nixon, *Representing and using nonfunctional requirements: A process-oriented approach*, IEEE Transactions on software engineering **18** (1992), no. 6, 483–497.

- [149] J. Mylopoulos, L. Chung, and E. Yu, *From object-oriented to goal-oriented requirements analysis*, Communications of the ACM **42** (1999), no. 1, 31–37.
- [150] S. Nalchigar, E. Yu, and K. Keshavjee, *Modeling machine learning requirements from three perspectives: a case report from the healthcare domain*, Requirements Engineering **26** (2021), no. 2, 237–254.
- [151] A. H. Nezhadi, B. Shadgar, and A. Osareh, *Ontology alignment using machine learning techniques*, International Journal of Computer Science & Information Technology **3** (2011), no. 2, 139.
- [152] M. Nickel, V. Tresp, and H. P. Kriegel, *A three-way model for collective learning on multi-relational data*, International Conference on Machine Learning, 2011.
- [153] D. A. Norman, *Some observations on mental models*, Mental models **7** (1983), no. 112, 7–14.
- [154] J. C. Nunnally, *Psychometric theory*, McGraw-Hill, New York, 1967.
- [155] D. E. O’Leary, *Google’s duplex: Pretending to be human*, Intelligent Systems in Accounting, Finance and Management **26** (2019), no. 1, 46–53.
- [156] OMG, *Business motivation model*.
- [157] B. Otto and M. Jarke, *Designing a multi-sided data platform: findings from the international data spaces case*, Electronic Markets **29** (2019), no. 4, 561–580.
- [158] A. L. Palacio and Ó.P. López, *From big data to smart data: A genomic information systems perspective*, 12th International Conference on Research Challenges in Information Science (RCIS), 2018, pp. 1–11.
- [159] O. Pastor, *Conceptual modeling of life: beyond the homo sapiens*, International Conference on Conceptual Modeling (p), Springer, 2016, pp. 18–31.
- [160] O. Pastor and J. C. Molina, *Model-driven architecture in practice: a software production environment based on conceptual modeling*, Springer, Science & Business Media, 2007.
- [161] Óscar Pastor, Marcela Ruiz, and Sergio España, *From requirements to code: A full model-driven development perspective*, International Conference on Software and Data Technologies, Springer, 2011, pp. 56–70.
- [162] H. Paulheim, *Knowledge graph refinement: A survey of approaches and evaluation methods*, Semantic web **8** (2017), no. 3, 489–508.
- [163] T. et al. Pedersen, *Measures of semantic similarity and relatedness in the biomedical domain*, Journal of biomedical informatics **40** (2007), no. 3, 288–299.
- [164] D. E. Perry and A. L. Wolf, *Foundations for the study of software architecture*, ACM SIG-SOFT Software engineering notes **17** (1992), no. 4, 40–52.
- [165] M. Pietranik and N. T. Nguyen, *Semantic distance measure between ontology concept’s attributes*, International Conference on Knowledge-Based and Intelligent Information and Engineering Systems, 2011, pp. 210–219.
- [166] L. L. Pipino, Y. W. Lee, and R. Y. Wang, *Data quality assessment*, Communications of the ACM **45** (2002), no. 4, 211–218.
- [167] M. E. Porter and V. E. Millar, *How information gives you competitive advantage*, Harvard Business Review **63** (1985), no. 4, 149–160.

- [168] F. Provost and T. Fawcett, *Data science for business: What you need to know about data mining and data-analytic thinking*, O'Reilly Media, 2013.
- [169] E. Puiutta and E. M. Veith, *Explainable reinforcement learning: A survey*, International Cross-Domain Conference for Machine Learning and Knowledge Extraction (2020), 77–95.
- [170] J. et al. Pujara, *Knowledge graph identification*, International Semantic Web Conference, 2013, pp. 542–557.
- [171] J. Pérez, J. Marinković, and P. Barceló, *On the turing completeness of modern neural network architectures*, Tech. report, arXiv preprint, 2019.
- [172] A. Radford, L. Metz, and S. Chintala, *Unsupervised representation learning with deep convolutional generative adversarial networks*, arXiv preprint, 2015.
- [173] M. T. Ribeiro, S. Singh, and C. Guestrin, *Why should i trust you?" explaining the predictions of any classifier*, 22nd ACM SIGKDD international conference on knowledge discovery and data mining, 2016, pp. 1135–1144.
- [174] M. Richters and M. Gogolla, *On formalizing the uml object constraint language ocl*, Springer, 1998.
- [175] C. Rudin, *Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead*, Nature Machine Intelligence **1** (2019), no. 5, 206–215.
- [176] F. et al. Rundo, *Machine learning for quantitative finance applications: A survey*, Applied Sciences **9** (2019), no. 24, 5574.
- [177] K. et al. Rusek, *Routenet: Leveraging graph neural networks for network modeling and optimization in sdn*, IEEE Journal on Selected Areas in Communications **38** (2020), no. 10, 2260–2270.
- [178] S. J. Russell and P. Norvig, *Artificial intelligence: A modern approach*, third ed., Prentice Hall, 2010.
- [179] O. Sagi and L. Rokach, *Ensemble learning: A survey*, Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery **8** (2018), no. 4, e1249.
- [180] J. Schmidhuber, *Deep learning in neural networks: An overview*, Neural networks **61** (2015), 85–117.
- [181] D. Schon, *Technology and social change*, Delacorte, New York, 1967.
- [182] Donghyuk Shin, Shu He, Gene Moo Lee, Andrew B Whinston, Suleyman Cetintas, and Kuang-Chih Lee, *Enhancing social media analysis with visual data analytics: A deep learning approach*, MIS Quarterly **44** (2020), no. 4, 1459–1492.
- [183] G. Shmueli and O. R. Koppius, *Predictive analytics in information systems research*, MIS quarterly **35** (2011), no. 3, 553–572.
- [184] H. T. Siegelmann and E. D. Sontag, *On the computational power of neural nets*, Journal of computer and system sciences **50** (1995), no. 1, 132–150.
- [185] A. et al. Siena, *Designing law-compliant software requirements*, International conference on conceptual modeling, 2009, pp. 472–486.
- [186] \_\_\_\_\_, *Capturing variability of law with nomos 2*, International Conference on Conceptual Modeling, Springer, 2012, pp. 383–396.

- [187] D. Silver, A. Huang, and C. Maddison, *Mastering the game of go with deep neural networks and tree search*, Nature **529** (2016), 484–489.
- [188] D. et al. Silver, *Mastering the game of go without human knowledge*, Nature **550** (2017), no. 7676, 354–359.
- [189] A. Singhal, *Introducing the knowledge graph: things, not strings*, 2012.
- [190] J. W. Smith, *Using the adap learning algorithm to forecast the onset of diabetes mellitus*, Annual Symposium on Computer Application in Medical Care, American Medical Informatics Association, 1988, p. 261–265.
- [191] L. N. Smith and N. Topin, *Deep convolutional neural network design patterns*, arXiv preprint, 2016.
- [192] I. Sommerville, *Software engineering*, Pearson, 2004.
- [193] S. Staab and A. Maedche, *Axioms are objects, too - ontology engineering beyond the modeling of concepts and relations*, 14th European Conference on Artificial Intelligence, Workshop on Applications of Ontologies and Problem-Solving Methods, 2000.
- [194] J. Stilgoe, *Machine learning, social learning and the governance of self-driving cars*, Social studies of science **48** (2018), no. 1, 25–56.
- [195] T. Stock and G. Seliger, *Opportunities of sustainable manufacturing in industry 4.0*, Procedia Cirp **40** (2016), 536–541.
- [196] V. C. Storey and I.-Y. Song, *Big data technologies and management: What conceptual modeling can do*, Data & Knowledge Engineering **108** (2017), 50–67.
- [197] V. C. Storey, J. C. Trujillo, and S. W. Liddle, *Research on conceptual modeling: Themes, topics, and introduction to the special issue*, Data and Knowledge Engineering **98** (2015), 1–7.
- [198] F. Suchanek and G. Weikum, *Knowledge harvesting in the big-data era*, ACM SIGMOD International Conference on Management of Data, 2013, pp. 933–938.
- [199] F. M. Suchanek, G. Kasneci, and G. Weikum, *Yago: a core of semantic knowledge*, 16th international conference on World Wide Web, 2007.
- [200] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*, MIT press, 2018.
- [201] A. et al. Ullah, *Modeling business goal for business/it alignment using requirements engineering*, Journal of Computer Information Systems **51** (2011), no. 3, 21–28.
- [202] S. Van Buuren, *Flexible imputation of missing data*, CRC press, 2018.
- [203] P. Van Hentenryck, *Constraint satisfaction in logic programming*, MIT Press, 1989.
- [204] A. Van Lamsweerde, *Goal-oriented requirements engineering: A guided tour*, Proceedings 5th international symposium on requirements engineering, IEEE, 2001, pp. 249–262.
- [205] ———, *From system goals to software architecture*, International School on Formal Methods for the Design of Computer, Communication and Software Systems (Berlin), Springer, 2003, pp. 25–43.
- [206] J. Van Tassel, *Digital rights management: Protecting and monetizing content*, Taylor & Francis, 2006.

- [207] S. J. et al. van Zelst, *Event abstraction in process mining: literature review and taxonomy*, Granular Computing **6** (2020), 719–736.
- [208] G. Vial, *Understanding digital transformation: A review and a research agenda*, Journal of Strategic Information Systems **28** (2019), no. 2, 118–144.
- [209] A. Vogelsang and M. Borg, *Requirements engineering for machine learning: Perspectives from data scientists*, IEEE 27th International Requirements Engineering Conference Workshops (REW), 2019, pp. 245–251.
- [210] Y. Wand and R. Weber, *Research commentary: information systems and conceptual modeling—a research agenda*, Information systems research **13** (2002), no. 4, 363–376.
- [211] D. Wang, P. Cui, and W. Zhu, *Structural deep network embedding*, 22nd ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 2016, pp. 1225–1234.
- [212] R. Y. Wang, V. C. Storey, and C. P. Firth, *A framework for analysis of data quality research*, IEEE transactions on knowledge and data engineering **7** (1995), no. 4, 623–640.
- [213] R. Y. Wang and D. M. Strong, *Beyond accuracy: What data quality means to data consumers*, Journal of management information systems **12** (1996), no. 4, 5–33.
- [214] S. Wasserman and K. Faust, *Social network analysis: Methods and applications*, vol. 8, university press, Cambridge, 1994.
- [215] K. M. Weiss, R. F. Ferrell, and C. L. Hanis, *A new world syndrome of metabolic diseases with a genetic and evolutionary basis*, American Journal of Physical Anthropology **27** (1984), no. 55, 153–178.
- [216] B. J. et al. Wells, *Strategies for handling missing data in electronic health record derived data*, Egems **1** (2013), no. 3.
- [217] R. Wirth and J. Hipp, *Crisp-dm: Towards a standard process model for data mining*, 4th international conference on the practical applications of knowledge discovery and data mining, Springer-Verlag, 2000, pp. 29–39.
- [218] I. Witten, *Text mining, practical handbook of internet computing, edited by mp singh*, CRC Press, 2004.
- [219] G.-Q. Wu, Y. He, and X. Hu, *Entity linking: An issue to extract corresponding entity with knowledge base*, IEEE Access **6** (2018), 6220–6231.
- [220] Y. et al. Wu, *Google’s neural machine translation system: Bridging the gap between human and machine translation*, 2016.
- [221] Z. et al. Wu, *A comprehensive survey on graph neural networks*, IEEE Transactions on Neural Networks and Learning Systems p (2020), 1–21.
- [222] T. et al. Wuest, *Machine learning in manufacturing: advantages, challenges, and applications*, Production & Manufacturing Research **4** (2016), no. 1, 23–45.
- [223] W. Xiong and L. Xiong, *Smart contract based data trading mode using blockchain and machine learning*, IEEE Access **7** (2019), 2331–10234.
- [224] D. et al. Xu, *Product knowledge graph embedding for e-commerce*, 13th international conference on web search and data mining, 2020, pp. 672–680.

- [225] H. Yang, *Aligraph: A comprehensive graph neural network platform*, ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2019, pp. 3165–3166.
- [226] L. C. Yang, S. Y. Chou, and Y. H. Yang, *Midinet: A convolutional generative adversarial network for symbolic-domain music generation*, arXiv preprint, 2017.
- [227] Y. Yang and J. O. Pedersen, *A comparative study on feature selection in text categorization*, Fourteenth International Conference on Machine Learning (ICML), 1997, pp. 412–420.
- [228] E. Yu, *Towards modelling and reasoning support for early-phase requirements engineering*, Proc 3rd IEEE Int Symp on Requirements Engineering, 1997, pp. 226–235.
- [229] E. S. Yu, *Towards modelling and reasoning support for early-phase requirements engineering*, Proceedings of ISRE'97: 3rd IEEE International Symposium on Requirements Engineering, IEEE, 1997, pp. 226–235.
- [230] H. Zou and T. Hastie, *Regularization and variable selection via the elastic net*, Journal of the royal statistical society: series B (statistical methodology) **67** (2005), no. 2, 301–320.