

# Supervised Learning for the Prediction of Firm Dynamics



Falco J. Bargagli-Stoffi, Jan Niederreiter, and Massimo Riccaboni

**Abstract** Thanks to the increasing availability of granular, yet high-dimensional, firm level data, machine learning (ML) algorithms have been successfully applied to address multiple research questions related to firm dynamics. Especially supervised learning (SL), the branch of ML dealing with the prediction of labelled outcomes, has been used to better predict firms' performance. In this chapter, we will illustrate a series of SL approaches to be used for prediction tasks, relevant at different stages of the company life cycle. The stages we will focus on are (1) startup and innovation, (2) growth and performance of companies, and (3) firms' exit from the market. First, we review SL implementations to predict successful startups and R&D projects. Next, we describe how SL tools can be used to analyze company growth and performance. Finally, we review SL applications to better forecast financial distress and company failure. In the concluding section, we extend the discussion of SL methods in the light of targeted policies, result interpretability, and causality.

**Keywords** Machine learning · Firm dynamics · Innovation · Firm performance

## 1 Introduction

In recent years, the ability of machines to solve increasingly more complex tasks has grown exponentially [86]. The availability of learning algorithms that deal with tasks such as facial and voice recognition, automatic driving, and fraud detection makes the various applications of machine learning a hot topic not just in the specialized literature but also in media outlets. Since many decades, computer scientists have been using algorithms that automatically update their course of

---

F. J. Bargagli-Stoffi  
Harvard University, Boston, MA, USA  
e-mail: [fbargaglistoffi@hsph.harvard.edu](mailto:fbargaglistoffi@hsph.harvard.edu)

J. Niederreiter · M. Riccaboni (✉)  
IMT School for Advanced Studies Lucca, Lucca, Italy  
e-mail: [jan.niederreiter@alumni.imtlucca.it](mailto:jan.niederreiter@alumni.imtlucca.it); [massimo.riccaboni@imtlucca.it](mailto:massimo.riccaboni@imtlucca.it)

action to better their performance. Already in the 1950s, Arthur Samuel developed a program to play checkers that improved its performance by learning from its previous moves. The term “machine learning” (ML) is often said to have originated in that context. Since then, major technological advances in data storage, data transfer, and data processing have paved the way for learning algorithms to start playing a crucial role in our everyday life.

Nowadays, the usage of ML has become a valuable tool for enterprises’ management to predict key performance indicators and thus to support corporate decision-making across the value chain, including the appointment of directors [33], the prediction of product sales [7], and employees’ turnover [1, 85]. Using data which emerges as a by-product of economic activity has a positive impact on firms’ growth [37], and strong data analytic capabilities leverage corporate performance [75]. Simultaneously, publicly accessible data sources that cover information across firms, industries, and countries open the door for analysts and policy-makers to study firm dynamics on a broader scale such as the fate of start-ups [43], product success [79], firm growth [100], and bankruptcy [12].

Most ML methods can be divided into two main branches: (1) *unsupervised learning* (UL) and (2) *supervised learning* (SL) models. UL refers to those techniques used to draw inferences from data sets consisting of input data without labelled responses. These algorithms are used to perform tasks such as clustering and pattern mining. SL refers to the class of algorithms employed to make predictions on labelled response values (i.e., discrete and continuous outcomes). In particular, SL methods use a known data set with input data and response values, referred to as training data set, to learn how to successfully perform predictions on labelled outcomes. The learned decision rules can then be used to predict unknown outcomes of new observations. For example, an SL algorithm could be trained on a data set that contains firm-level financial accounts and information on enterprises’ solvency status in order to develop decision rules that predict the solvency of companies.

SL algorithms provide great added value in predictive tasks since they are specifically designed for such purposes [56]. Moreover, the nonparametric nature of SL algorithms makes them suited to uncover hidden relationships between the predictors and the response variable in large data sets that would be missed out by traditional econometric approaches. Indeed, the latter models, e.g., ordinary least squares and logistic regression, are built assuming a set of restrictions on the functional form of the model to guarantee statistical properties such as estimator unbiasedness and consistency. SL algorithms often relax those assumptions and the functional form is dictated by the data at hand (data-driven models). This characteristic makes SL algorithms more “adaptive” and inductive, therefore enabling more accurate predictions for future outcome realizations.

In this chapter, we focus on the traditional usage of SL for predictive tasks, excluding from our perspective the growing literature that regards the usage of SL for causal inference. As argued by Kleinberg et al. [56], researchers need to answer to both causal and predictive questions in order to inform policy-makers. An example that helps us to draw the distinction between the two is provided by

a policy-maker facing a pandemic. On the one side, if the policy-maker wants to assess whether a quarantine will prevent a pandemic to spread, he needs to answer a purely causal question (i.e., “what is the effect of quarantine on the chance that the pandemic will spread?”). On the other side, if the policy-maker wants to know if he should start a vaccination campaign, he needs to answer a purely predictive question (i.e., “Is the pandemic going to spread within the country?”). SL tools can help policy-makers navigate both these sorts of policy-relevant questions [78]. We refer to [6] and [5] for a critical review of the causal machine learning literature.

Before getting into the nuts and bolts of this chapter, we want to highlight that our goal is not to provide a comprehensive review of all the applications of SL for prediction of firm dynamics, but to describe the alternative methods used so far in this field. Namely, we selected papers based on the following inclusion criteria: (1) the usage of SL algorithm to perform a predictive task in one of our fields of interest (i.e., enterprises success, growth, or exit), (2) a clear definition of the outcome of the model and the predictors used, (3) an assessment of the quality of the prediction. The purpose of this chapter is twofold. First, we outline a general SL framework to ready the readers’ mindset to think about prediction problems from an SL-perspective (Sect. 2). Second, equipped with the general concepts of SL, we turn to real-world applications of the SL predictive power in the field of firms’ dynamics. Due to the broad range of SL applications, we organize Sect. 3 into three parts according to different stages of the firm life cycle. The prediction tasks we will focus on are about the success of new enterprises and innovation (Sect. 3.1), firm performance and growth (Sect. 3.2), and the exit of established firms (Sect. 3.3). The last section of the chapter discusses the state of the art, future trends, and relevant policy implications (Sect. 4).

## 2 Supervised Machine Learning

In a famous paper on the difference between model-based and data-driven statistical methodologies, Berkeley professor Leo Breiman, referring to the statistical community, stated that “there are two cultures in the use of statistical modeling to reach conclusions from data. One assumes that the data are generated by a given stochastic data model. The other uses algorithmic models and treats the data mechanism as unknown. [...] If our goal as a field is to use data to solve problems, then we need to move away from exclusive dependence on data models and adopt a diverse set of tools” [20, p. 199]. In this quote, Breiman catches the essence of SL algorithms: their ability to capture hidden patterns in the data by directly learning from them, without the restrictions and assumptions of model-based statistical methods.

SL algorithms employ a set of data with input data and response values, referred as training sample, to learn and make predictions (in-sample predictions), while another set of data, referred as test sample, is kept separate to validate the predictions (out-of-sample predictions). Training and testing sets are usually built by randomly sampling observations from the initial data set. In the case of panel data, the

testing sample should contain only observations that occurred later in time than the observations used to train the algorithm to avoid the so-called *look-ahead bias*. This ensures that future observations are predicted from past information, not vice versa.

When the dependent variable is categorical (e.g., yes/no or category 1–5) the task of the SL algorithm is referred as a “classification” problem, whereas in “regression” problems the dependent variable is continuous.

The common denominator of SL algorithms is that they take an information set  $\mathbf{X}_{N \times P}$ , i.e., a matrix of features (also referred to as attributes or predictors), and map it to an  $N$ -dimensional vector of outputs  $y$  (also referred to as actual values or dependent variable), where  $N$  is the number of observations  $i = 1, \dots, N$  and  $P$  is the number of features. The functional form of this relationship is very flexible and gets updated by evaluating a loss function. The functional form is usually modelled in two steps [78]:

1. pick the best in-sample loss-minimizing function  $f(\cdot)$ :

$$\operatorname{argmin} \sum_{i=1}^N L(f(x_i), y_i) \quad \text{over } f(\cdot) \in F \quad \text{s. t.} \quad R(f(\cdot)) \leq c \quad (1)$$

where  $\sum_{i=1}^N L(f(x_i), y_i)$  is the in-sample loss functional to be minimized (i.e., the mean squared error of prediction),  $f(x_i)$  are the predicted (or fitted) values,  $y_i$  are the actual values,  $f(\cdot) \in F$  is the function class of the SL algorithm, and  $R(f(\cdot))$  is the complexity functional that is constrained to be less than a certain value  $c \in \mathbb{R}$  (e.g., one can think of this parameter as a budget constraint);

2. estimate the optimal level of complexity using empirical tuning through cross-validation.

Cross-validation refers to the technique that is used to evaluate predictive models by training them on the training sample, and evaluating their performance on the test sample.<sup>1</sup> Then, on the test sample the algorithm’s performance is evaluated on how well it has learned to predict the dependent variable  $y$ . By construction, many SL algorithms tend to perform extremely well on the training data. This phenomenon is commonly referred as “overfitting the training data” because it combines very high predictive power on the training data with poor fit on the test data. This lack of generalizability of the model’s prediction from one sample to another can be addressed by penalizing the model’s complexity. The choice of a good penalization algorithm is crucial for every SL technique to avoid this class of problems.

In order to optimize the complexity of the model, the performance of the SL algorithm can be assessed by employing various performance measures on the test sample. It is important for practitioners to choose the performance measure that

---

<sup>1</sup>This technique (hold-out) can be extended from two to  $k$  folds. In  $k$ -folds cross-validation, the original data set is randomly partitioned into  $k$  different subsets. The model is constructed on  $k - 1$  folds and evaluated on onefold, repeating the procedure until all the  $k$  folds are used to evaluate the predictions.

		Observed Outcome		
		Positive	Negative	
Predicted Outcome	Positive	80 (True positive)	17 (False positive)	Positive predicted value (PPV, Precision): $80/97=82.5\%$
	Negative	2 (False negative)	1 (True negative)	Negative predicted value (NPV): $1/3=33.3\%$
		True positive rate (TPR, Recall, Sensitivity): $80/82=97.6\%$	True negative rate (TNR, Specificity): $1/18=5.6\%$	Accuracy (ACC): $81/100=81\%$ Balanced Accuracy (BACC): $(TPR+TNR)/2=51.6\%$

**Fig. 1** Exemplary confusion matrix for assessment of classification performance

best fits the prediction task at hand and the structure of the response variable. In regression tasks, different performance measures can be employed. The most common ones are the mean squared error (MSE), the mean absolute error (MAE), and the  $R^2$ . In classification tasks the most straightforward method is to compare true outcomes with predicted ones via confusion matrices from where common evaluation metrics, such as true positive rate (TPR), true negative rate (TNR), and accuracy (ACC), can be easily calculated (see Fig. 1). Another popular measure of prediction quality for binary classification tasks (i.e., positive vs. negative response), is the Area Under the receiver operating Curve (AUC) that relates how well the trade-off between the models TPR and TNR is solved. TPR refers to the proportion of positive cases that are predicted correctly by the model, while TNR refers to the proportion of negative cases that are predicted correctly. Values of AUC range between 0 and 1 (perfect prediction), where 0.5 indicates that the model has the same prediction power as a random assignment. The choice of the appropriate performance measure is key to communicate the fit of an SL model in an informative way.

Consider the example in Fig. 1 in which the testing data contains 82 positive outcomes (e.g., firm survival) and 18 negative outcomes, such as firm exit, and the algorithm predicts 80 of the positive outcomes correctly but only one of the negative ones. The simple accuracy measure would indicate 81% correct classifications, but the results suggest that the algorithm has not successfully learned how to detect negative outcomes. In such a case, a measure that considers the unbalance of outcomes in the testing set, such as balanced accuracy (BACC, defined as  $((TPR + TNR)/2) = 51.6\%$ ), or the F1-score would be more suited. Once the algorithm has been successfully trained and its out-of-sample performance has been properly tested, its decision rules can be applied to predict the outcome of new observations, for which outcome information is not (yet) known.

Choosing a specific SL algorithm is crucial since performance, complexity, computational scalability, and interpretability differ widely across available implementations. In this context, easily interpretable algorithms are those that provide

comprehensive decision rules from which a user can retrace results [62]. Usually, highly complex algorithms require the discretionary fine-tuning of some model hyperparameters, more computational resources, and their decision criteria are less straightforward. Yet, the most complex algorithms do not necessarily deliver the best predictions across applications [58]. Therefore, practitioners usually run a *horse race* on multiple algorithms and choose the one that provides the best balance between interpretability and performance on the task at hand. In some learning applications for which prediction is the sole purpose, different algorithms are combined and the contribution of each chosen so that the overall predictive performance gets maximized. Learning algorithms that are formed by multiple self-contained methods are called ensemble learners (e.g., the super-learner algorithm by Van der Laan et al. [97]).

Moreover, SL algorithms are used by scholars and practitioners to perform predictors selection in high-dimensional settings (e.g., scenarios where the number of predictors is larger than the number of observations: small  $N$  large  $P$  settings), text analytics, and natural language processing (NLP). The most widely used algorithms to perform the former task are the least absolute shrinkage and selection operator (Lasso) algorithm [93] and its related versions, such as stability selection [74] and C-Lasso [90]. The most popular supervised NLP and text analytics SL algorithms are support vector machines [89], Naive Bayes [80], and Artificial Neural Networks (ANN) [45].

Reviewing SL algorithms and their properties in detail would go beyond the scope of this chapter; however, in Table 1 we provide a basic intuition of the most widely used SL methodologies employed in the field of firm dynamics. A more detailed discussion of the selected techniques, together with a code example to implement each one of them in the statistical software R, and a toy application on real firm-level data, is provided in the following web page: <http://github.com/fbargaglistoffi/machine-learning-firm-dynamics>.

### 3 SL Prediction of Firm Dynamics

Here, we review SL applications that have leveraged inter firm data to predict various company dynamics. Due to the increasing volume of scientific contributions that employ SL for company-related prediction tasks, we split the section in three parts according to the life cycle of a firm. In Sect. 3.1 we review SL applications that deal with early-stage firm success and innovation, in Sect. 3.2 we discuss growth and firm-performance-related work, and lastly, in Sect. 3.3, we turn to firm exit prediction problems.

**Table 1** SL algorithms commonly applied in predicting firm dynamics

Method	Description	Interpretability
Decision Tree (DT)	Decision trees (DT) consist of a sequence of binary decision rules (nodes) on which the tree splits into branches (edges). At each final branch (leaf node) a decision regarding the outcome is estimated. The sequence and definition of nodes is based on minimizing a measure of node purity (e.g., Gini index, or entropy for classification tasks and MSE for regression tasks). Decision trees are easy to interpret but sensitive to changes in the features that frequently lower their predictive performance (see also [21]).	High
Random Forest (RF)	Instead of estimating just one DT, random forest (RF) re-samples the training set observations to estimate multiple trees. For each tree at each node a set of $m$ (with $m < P$ ) predictors is chosen randomly from the features space. To obtain the final prediction, the outcomes of all trees are averaged or, in the case of classification tasks, chosen by majority vote (see also [19]).	Medium
Support vector machines (SVM)	Support vector machine (SVM) algorithms estimate a hyperplane over the feature space to classify observations. The vectors that span the hyperplane are called support vectors. They are chosen such that the overall distance (referred to as margin) between the data points and the hyperplane as well as the prediction accuracy is maximized (see also [89]).	Medium
(Deep) Artificial Neural Networks (ANN)	Inspired by biological networks, every artificial neural network (ANN) consists of, at least, three layers (deep ANNs are ANNs with more than three layers): an input layer with feature information, one or more hidden layers, and an output layer returning the predicted values. Each layer consists of nodes (neurons) that are connected via edges across layers. During the learning process, edges that are more important are reinforced. Neurons may then only send a signal if the signal received is strong enough (see also [45]).	Low

### 3.1 *Entrepreneurship and Innovation*

The success of young firms (referred to as startups) plays a crucial role in our economy since these firms often act as net creators of new jobs [46] and push, through their product and process innovations, the societal frontier of technology. Success stories of Schumpeterian entrepreneurs that reshaped entire industries are very salient, yet from a probabilistic point of view it is estimated that only 10% of startups stay in business long term [42, 59].

Not only is startup success highly uncertain, but it also escapes our ability to identify the factors to predict successful ventures. Numerous contributions have

used traditional regression-based approaches to identify factors associated with the success of small businesses (e.g., [69, 68, 44]), yet do not test the predictive quality of their methods out of sample and rely on data specifically collected for the research purpose. Fortunately, open access platforms such as *Crunchbase.com* and *Kickstarter.com* provide company- and project-specific data whose high dimensionality can be exploited using predictive models [29]. SL algorithms, trained on a large amount of data, are generally suited to predict startup success, especially because success factors are commonly unknown and their interactions complex. Similarly to the prediction of success at the firm level, SL algorithms can be used to predict success for singular projects. Moreover, unstructured data, e.g., business plans, can be combined with structured data to better predict the odds of success.

Table 2 summarizes the characteristics of recent contributions in various disciplines that use SL algorithms to predict startup success (upper half of the table) and success on the project level (lower half of the table). The definition of success varies across these contributions. Some authors define successful startups as firms that receive a significant source of external funding (this can be additional financing via venture capitalists, an initial public offering, or a buyout) that would allow to scale operations [4, 15, 87, 101, 104]. Other authors define successful startups as companies that simply survive [16, 59, 72] or coin success in terms of innovative capabilities [55, 43]. As data on the project level is usually not publicly available [51, 31], research has mainly focused on two areas for which it is, namely, the project funding success of crowdfunding campaigns [34, 41, 52] and the success of pharmaceutical projects to pass clinical trials [32, 38, 67, 79].<sup>2</sup>

To successfully distinguish how to classify successes from failures, algorithms are usually fed with company-, founder-, and investor-specific inputs that can range from a handful of attributes to a couple of hundred. Most authors find the information that relate to the source of funds predictive for startup success (e.g., [15, 59, 87]), but also entrepreneurial characteristics [72] and engagement in social networks [104] seem to matter. At the project level, funding success depends on the number of investors [41] as well as on the audio/visual content provided by the owner to pitch the project [52], whereas success in R&D projects depends on an interplay between company-, market-, and product-driven factors [79].

Yet, it remains challenging to generalize early-stage success factors, as these accomplishments are often context dependent and achieved differently across heterogeneous firms. To address this heterogeneity, one approach would be to first categorize firms and then train SL algorithms for the different categories. One can manually define these categories (i.e., country, size cluster) or adopt a data-driven approach (e.g., [90]).

---

<sup>2</sup>Since 2007 the US Food and Drug Administration (FDA) requires that the outcome of clinical trials that passed “Phase I” be publicly disclosed [103]. Information on these clinical trials, and pharmaceutical companies in general, has since then been used to train SL methods to classify the outcome of R&D projects.



**Table 2** SL literature on firms' early success and innovation

References	Domain	Output	Country, time	Data set size	Primary SL-method	Attributes	GOF
Arroyo et al. [4]	CS	Startup funding	INT (2011–2018)	120,507	GTB	105	82% (ACC)
Bento [15]	BI	Startup funding	USA (1985–2014)	143,348	RF	158	93% (AUC)
Böhm et al. [16]	BI	Startup survival, -growth	USA, GER (1999–2015)	181	SVM	69	67–84% (ACC)
Guerozmi et al. [43]	ECON	Startup innovativeness	ITA (2013)	45,576	bagging, ANN	262	56% (TPR), 95% (TNR)
Kinne and Lenz [55]	ECON	Firm innovativeness	GER (2012–2016)	4481	ANN	N/A	80% (F-score)
Krishna et al. [59]	CS	Startup survival	INT (1999–2014)	13,000	RF, LR	70	73–96% (ACC)
McKenzie and Sainsone [72]	ECON	Startup survival	NIG (2014–2015)	2506	SVM	393	64% (ACC)
Sharchilev et al. [87]	CS	Startup funding	INT	21,947	GTB	49	85% (AUC)
Xiang et al. [101]	BI	Startup M&A	INT (1970–2007)	59,631	BN	27	68–89% (AUC)
Yankov et al. [102]	ECON	Startup survival	BUL	142	DT	15	67% (ACC)
Zhang et al. [104]	CS	Startup funding	INT (2015–2016)	4001	SVM	14	84% (AM)
DiMasi et al. [32]	PHARM	Project success (oncology drugs)	INT (1999–2007)	98	RF	4	92% (AUC)
Eiter et al. [34]	CS	Project funding	INT (2012–2013)	16,042	Ensemble SVM	12	> 76% (ACC)
Fejoo et al. [38]	PHARM	Project success (clinical trials)	INT (1993–2018)	6417	RF	17	80% (ACC)
Greenberg et al. [41]	CS	Project funding	INT (2012)	13,000	RF	12	67% (ACC)
Kaminski and Hopp [52]	ECON	Project funding	INT (2009–2017)	20,188	LR	200	65–71% (ACC)
Kyeambae et al. [60]	BMA	Emerging Technologies	USA (1979–2010)	11,000	SVM	7	71% (ACC)
Lo et al. [67]	CS	Project success (drugs)	INT (2003–2015)	27,800	KNN, RF	140	74–81% (AUC)
Munos et al. [79]	PHARM	Project success (drugs)	USA (2008–2018)	8,800	BART	37	91–96% (AUC)
Rouhani and Ravasan [84]	ENG	Project success (IT system)	ME (2011)	171	ANN	24	69% (ACC)

Abbreviations used—Domain: ECON: Economics, CS: Computer Science, BI: Business Informatics, ENG: Engineering, BMA: Business, Management and Accounting, PHARM: Pharmacology, Country: ITA: Italy, GER: Germany, INT: International, BUL: Bulgaria, USA: United States of America, NIG: Nigeria, ME: Middle East. Primary SL-method: ANN: (deep) neural network, SL: supervised learner, GTB: gradient tree boosting, DT: Decision Tree, SVM: support vector machine, BN: Bayesian Network, IXL: induction on eXtremely Large databases, RF: random forest, KNN: k-nearest neighbor, BART: Bayesian additive regression tree, LR: Logistic regression, TPR: true positive rate, TNR: true negative rate, ACC: Accuracy, AUC: Area under the receiver operating curve, BACC: Balanced Accuracy (average between TPR and TNR). The year was not reported when it was not possible to recover this information from the papers

The SL methods that best predict startup and project success vary vastly across reviewed applications, with random forest (RF) and support vector machine (SVM) being the most commonly used approaches. Both methods are easily implemented (see our web appendix), and despite their complexity still deliver interpretable results, including insights on the importance of singular attributes. In some applications, easily interpretable logistic regressions (LR) perform at par or better than more complex methods [36, 52, 59]. This might first seem surprising, yet it largely depends on whether complex interdependencies in the explanatory attributes are present in the data at hand. As discussed in Sect. 2 it is therefore recommendable to run a horse race to explore the prediction power of multiple algorithms that vary in terms of their interpretability.

Lastly, even if most contributions report their goodness of fit (GOF) using standard measures such as ACC and AUC, one needs to be cautious when cross-comparing results because these measures depend on the underlying data set characteristics, which may vary. Some applications use data samples, in which successes are less frequently observed than failures. Algorithms that perform well when identifying failures but have limited power when it comes to classifying successes would then be better ranked in terms of ACC and AUC than algorithms for which the opposite holds (see Sect. 2). The GOF across applications simply reflects that SL methods, on average, are useful for predicting startup and project outcomes. However, there is still considerable room for improvement that could potentially come from the quality of the used features as we do not find a meaningful correlation between data set size and GOF in the reviewed sample.

### 3.2 *Firm Performance and Growth*

Despite recent progress [22] firm growth is still an elusive problem. Table 3 schematizes the main supervised learning works in the literature on firms' growth and performance. Since the seminal contribution of Gibrat [40] firm growth is still considered, at least partially, as a random walk [28], there has been little progress in identifying the main drivers of firm growth [26], and recent empirical models have a small predictive power [98]. Moreover, firms have been found to be persistently heterogeneous, with results varying depending on their life stage and marked differences across industries and countries. Although a set of stylized facts are well established, such as the negative dependency of growth on firm age and size, it is difficult to predict the growth and performance from previous information such as balance sheet data—i.e., it remains unclear what are good predictors for what type of firm.

SL excels at using high-dimensional inputs, including nonconventional unstructured information such as textual data, and using them all as predictive inputs. Recent examples from the literature reveal a tendency in using multiple SL tools to make better predictions out of publicly available data sources, such as financial reports [82] and company web pages [57]. The main goal is to identify the key

Table 3 SL literature on firms' growth and performance

References	Domain	Output	Country, time	Data set size	Primary SL-method	Attributes	GOF
Weinblat [100]	BMA	High growth firms	INT (2004–2014)	179,970	RF	30	52%–81% (AUC)
Megaravalli and Sampagnaro [73]	BMA	High growth firms	ITA (2010–2014)	22,333	PR*	5	71% (AUC)
Coad and Srhoj [27]	BMA	High growth firms	HRV (2003–2016)	79,109	Lasso	172	76% (ACC)
Miyakawa et al. [76]	ECON	Firm exit, sales growth, profit growth	JPN (2006–2014)	1,700,000	weighted RF	50	70%, 68%, 61% (AUC)
Lam [61]	BI	ROE	USA (1985–1995)	364 firms per set	per ANN	27	Portfolio return comparison
Kolkman and van Witteloostuijn [57]	ECON	Asset growth	NL	8163 firms	RF	113	16% ( $R^2$ )
Qiu et al. [82]	CS	Groups of SAR, EPS growth	USA (1997–2003)	1276 firms	SVM	From annual reports	50% (ACC)
Bakar and Tahir [8]	BMA	ROA	MYS (2001–2006)	91	ANN	7	66.9% ( $R^2$ )
Baumann et al. [13]	CS	Customer Churn	INT	5000–93,893	Ensemble	20–359	1.5–6.8 ( $L_1$ )
Ravi et al. [83]	CS	Profit of banks	INT (1991–1993)	1000	Ensemble	54	80–93% (ACC)

Abbreviations used—Domain: ECON: Economics, CS: Computer Science, BI: Business Informatics, BMA: Business, Management and Accounting, Country: ITA: Italy, INT: International, HRV: Croatia, USA: United states of America, JPN: Japan, NL: Netherlands, MYS: Malaysia, Primary SL-method: ANN: (deep) neural network, SVM: support vector machine, RF: random forest, PR: Probit regression (simplest form of SL if out of sample performance analysis used), Lasso: Least absolute shrinkage and selection operator, Ensemble: Ensemble Learner, GOF: Accuracy, AUC: Area under the receiver operating curve,  $L_1$ : Top decile lift,  $R^2$  R-squared. The year was not reported when it was not possible to recover this information from the papers

drivers of superior firm performance in terms of profits, growth rates, and return on investments. This is particularly relevant for stakeholders, including investors and policy-makers, to devise better strategies for sustainable competitive advantage. For example, one of the objectives of the European commission is to incentivize high growth firms (HGFs) [35], which could get facilitated by classifying such companies adequately.

A prototypical example of application of SL methods to predict HGFs is Weinblat [100], who uses an RF algorithm trained on firm characteristics for different EU countries. He finds that HGFs have usually experienced prior accelerated growth and should not be confused with startups that are generally younger and smaller. Predictive performance varies substantially across country samples, suggesting that the applicability of SL approaches cannot be generalized. Similarly, Miyakawa et al. [76] show that RF can outperform traditional credit score methods to predict firm exit, growth in sales, and profits of a large sample of Japanese firms. Even if the reviewed SL literature on firms' growth and performance has introduced approaches that increment predictive performance compared to traditional forecasting methods, it should be noted that this performance stays relatively low across applications in the firms' life cycle and does not seem to correlate significantly with the size of the data sets. A firm's growth seems to depend on many interrelated factors whose quantification might still be a challenge for researchers who are interested in performing predictive analysis.

Besides identifying HGFs, other contributions attempt to maximize predictive power of future performance measures using sophisticated methods such as ANN or ensemble learners (e.g., [83, 61]). Even though these approaches achieve better results than traditional benchmarks, such as financial returns of market portfolios, a lot of variation of the performance measure is left unexplained. More importantly, the use of such "black-box" tools makes it difficult to derive useful recommendations on what options exist to better individual firm performance. The fact that data sets and algorithm implementation are usually not made publicly available adds to our impotence at using such results as a base for future investigations.

Yet, SL algorithms may help individual firms improve their performance from different perspectives. A good example in this respect is Erel et al. [33], who showed how algorithms can contribute to appoint better directors.

### ***3.3 Financial Distress and Firm Bankruptcy***

The estimation of default probabilities, financial distress, and the predictions of firms' bankruptcies based on balance sheet data and other sources of information on firms viability is a highly relevant topic for regulatory authorities, financial institutions, and banks. In fact, regulatory agencies often evaluate the ability of banks to assess enterprises viability, as this affects their capacity of best allocating financial resources and, in turn, their financial stability. Hence, the higher predictive power of SL algorithms can boost targeted financing policies that lead to safer

allocation of credit either on the extensive margin, reducing the number of borrowers by lending money just to the less risky ones, or on the intensive margin (i.e., credit granted) by setting a threshold to the amount of credit risk that banks are willing to accept.

In their seminal works in this field, Altman [3] and Ohlson [81] apply standard econometric techniques, such as multiple discriminant analysis (MDA) and logistic regression, to assess the probability of firms' default. Moreover, since the Basel II Accord in 2004, default forecasting has been based on standard reduced-form regression approaches. However, these approaches may fail, as for MDA the assumptions of linear separability and multivariate normality of the predictors may be unrealistic, and for regression models there may be pitfalls in (1) their ability to capture sudden changes in the state of the economy, (2) their limited model complexity that rules out nonlinear interactions between the predictors, and (3) their narrow capacity for the inclusion of large sets of predictors due to possible multicollinearity issues.

SL algorithms adjust for these shortcomings by providing flexible models that allow for nonlinear interactions in the predictors space and the inclusion of a large number of predictors without the need to invert the covariance matrix of predictors, thus circumventing multicollinearity [66]. Furthermore, as we saw in Sect. 2, SL models are directly optimized to perform predictive task and this leads, in many situations, to a superior predictive performance. In particular, Moscatelli et al. [77] argue that SL models outperform standard econometric models when the predictions of firms' distress is (1) based solely on financial accounts data as predictors and (2) relies on a large amount of data. In fact, as these algorithms are "model free," they need large data sets ("data-hungry algorithms") in order to extract the amount of information needed to build precise predictive models. Table 4 depicts a number of papers in the field of economics, computer science, statistics, business, and decision sciences that deal with the issue of predicting firms' bankruptcy or financial distress through SL algorithms. The former stream of literature (bankruptcy prediction)—which has its foundations in the seminal works of Udo [96], Lee et al. [63], Shin et al. [88], and Chandra et al. [23]—compares the binary predictions obtained with SL algorithms with the actual realized failure outcomes and uses this information to calibrate the predictive models. The latter stream of literature (financial distress prediction)—pioneered by Fantazzini and Figini [36]—deals with the problem of predicting default probabilities (DPs) [77, 12] or financial constraint scores [66]. Even if these streams of literature approach the issue of firms' viability from slightly different perspectives, they train their models on dependent variables that range from firms' bankruptcy (see all the "bankruptcy" papers in Table 4) to firms' insolvency [12], default [36, 14, 77], liquidation [17], dissolvency [12] and financial constraint [71, 92].

In order to perform these predictive tasks, models are built using a set of *structured* and *unstructured* predictors. With structured predictors we refer to balance sheet data and financial indicators, while unstructured predictors are, for instance, auditors' reports, management statements, and credit behavior indicators. Hansen et al. [71] show that the usage of unstructured data, in particular, auditors

**Table 4** SL literature on firms' failure and financial distress

References	Domain	Output	Country, time	Data set size	Primary SL-method	Attributes	GOF
Alaka et al. [2]	CS	Bankruptcy	UK (2001–2015)	30,000	NN	5	88% (AUC)
Barboza et al. [9]	CS	Bankruptcy	USA (1985–2014)	10,000	SVM, RF, BO, BA	11	93% (AUC)
Bargagli-Stoffi et al. [12]	ECON	Fin. distress	ITA (2008–2017)	304,000	BART	46	97% (AUC)
Behr and Weinblat [14]	ECON	Bankruptcy	INT (2010–2011)	945,062	DT, RF	20	85% (AUC)
Bonello et al. [17]	ECON	Fin. distress	USA (1996–2016)	1848	NB, DT, NN	96	78% (ACC)
Brédart [18]	BMA	Bankruptcy	BEL (2002–2012)	3728	NN	3	81% (ACC)
Chandra et al. [23]	CS	Bankruptcy	USA (2000)	240	DT	24	75% (ACC)
Cleofias-Sánchez et al. [25]	CS	Fin. distress	INT (2007)	240–8200	SVM, NN, LR	12–30	78% (ACC)
Danenas and Garsva [30]	CS	Fin. distress	USA (1999–2007)	21,487	SVM, NN, LR	51	95% (ACC)
Fantazzini and Fignini [36]	STAT	Fin. distress	DEU (1996–2004)	1003	SRF	16	93% (ACC)
Hansen et al. [71]	ECON	Fin. distress	DNK (2013–2016)	278,047	CNN, RNN	50	84% (AUC)
Heco and Yang [47]	CS	Bankruptcy	KOR (2008–2012)	30,000	ADA	12	94% (ACC)
Hosaka [48]	CS	Bankruptcy	JPN (2002–2016)	2703	CNN	14	18% (F-score)
Kim and Upneja [54]	CS	Bankruptcy	KOR (1988–2010)	10,000	ADA, DT	30	95% (ACC)
Lee et al. [63]	BMA	Bankruptcy	KOR (1979–1992)	166	NN	57	82% (ACC)
Liang et al. [65]	ECON	Bankruptcy	TWN (1999–2009)	480	SVM, KNN, DT, NB	190	82% (ACC)
Linn and Weagley [66]	ECON	Fin. distress	INT (1997–2015)	48,512	DRF	16	15% ( $R^2$ )
Moscattelli et al. [77]	ECON	Fin. distress	ITA (2011–2017)	250,000	RF	24	84% (AUC)
Shin et al. [88]	CS	Bankruptcy	KOR (1996–1999)	1160	SVM	52	77% (ACC)
Sun and Li [91]	CS	Bankruptcy	CHN	270	CBR, KNN	5	79% (ACC)
Sun et al. [92]	BMA	Fin. distress	CHN (2005–2012)	932	ADA, SVM	13	87% (ACC)
Tsai and Wu [94]	CS	Bankruptcy	INT	690–1000	NN	14–20	79–97% (ACC)
Tsai et al. [95]	CS	Bankruptcy	TWN	440	NN, SVM, BO, BA	95	86% (ACC)
Wang et al. [99]	CS	Bankruptcy	POL (1997–2001)	240	DT, NN, NB, SVM	30	82% (ACC)
Udo [96]	CS	Bankruptcy	KOR (1996–2016)	300	NN	16	91% (ACC)
Zikeba et al. [105]	CS	Bankruptcy	POL (2000–2013)	10,700	BO	64	95% (AUC)

Abbreviations used—Domain: ECON: Economics, CS: Computer Science, BMA: Business, Management, Accounting, STAT: Statistics, Country: BEL: Belgium, ITA: Italy, DEU: Germany, INT: International, KOR: Korea, USA: United states of America, TWN: Taiwan, CHN: China, UK: United Kingdom, POL: Poland, Primary SL-method: ADA: AdaBoost, ANN: Artificial neural network, CNN: Convolutional neural network, NN: Neural network, GTB: gradient tree boosting, RF: Random forest, DRF: Decision tree, SRF: Survival random forest, DT: Decision Tree, SVM: support vector machine, NB: Naive Bayes, BO: Boosting, BA: Bagging, KNN: k-nearest neighbor, BART: Bayesian additive regression tree, DT: decision tree, LR: Logistic regression, Rate: ACC: Accuracy, AUC: Area under the receiver operating curve. The year was not reported when it was not possible to recover this information from the papers

reports, can improve the performance of SL algorithms in predicting financial distress. As SL algorithms do not suffer from multicollinearity issues, researchers can keep the set of predictors as large as possible. However, when researchers wish to incorporate just a set of “meaningful” predictors, Behr and Weinblat [14] suggest to include indicators that (1) were found to be useful to predict bankruptcies in previous studies, (2) are expected to have a predictive power based on firms’ dynamics theory, and (3) were found to be important in practical applications. As, on the one side, informed choices of the predictors can boost the performance of the SL model, on the other side, economic intuition can guide researchers in the choice of the best SL algorithm to be used with the disposable data sources. Bargagli-Stoffi et al. [12] show that an SL methodology that incorporates the information on missing data into its predictive model—i.e., the BART-mia algorithm by Kapelner and Bleich [53]—can lead to staggering increases in predictive performances when the predictors are missing not at random (MNAR) and their missingness patterns are correlated with the outcome.<sup>3</sup>

As different attributes can have different predictive powers with respect to the chosen output variable, it may be the case that researchers are interested in providing to policy-makers interpretable results in terms of which are the most important variables or the marginal effects of a certain variable on the predictions. Decision-tree-based algorithms, such as random forest [19], survival random forests [50], gradient boosted trees [39], and Bayesian additive regression trees [24], provide useful tools to investigate the aforementioned dimensions (i.e., variables importance, partial dependency plots, etc.). Hence, most of the economics papers dealing with bankruptcy or financial distress predictions implement such techniques [14, 66, 77, 12] in service of policy-relevant implications. On the other side, papers in the fields of computer science and business, which are mostly interested in the quality of predictions, de-emphasizing the interpretability of the methods, are built on black box methodologies such as artificial neural networks [2, 18, 48, 91, 94, 95, 99, 63, 96]. We want to highlight that, from the analyses of selected papers, we find no evidence of a positive correlation between the number of observations and predictors included in the model and the performance of the model. Indicating that “more” is not always better in SL applications to firms’ failures and bankruptcies.

## 4 Final Discussion

SL algorithms have advanced to become effective tools for prediction tasks relevant at different stages of the company life cycle. In this chapter, we provided a general introduction into the basics of SL methodologies and highlighted how they can be

---

<sup>3</sup>Bargagli-Stoffi et al. [12] argue that oftentimes the decision not to release financial account information is driven by firms’ financial distress.

applied to improve predictions on future firm dynamics. In particular, SL methods improve over standard econometric tools in predicting firm success at an early stage, superior performance, and failure. High-dimensional, publicly available data sets have contributed in recent years to the applicability of SL methods in predicting early success on the firm level and, even more granular, success at the level of single products and projects. While the dimension and content of data sets varies across applications, SVM and RF algorithms are oftentimes found to maximize predictive accuracy. Even though the application of SL to predict superior firm performance in terms of returns and sales growth is still in its infancy, there is preliminary evidence that RF can outperform traditional regression-based models while preserving interpretability. Moreover, shrinkage methods, such as Lasso or stability selection, can help in identifying the most important drivers of firm success. Coming to SL applications in the field of bankruptcy and distress prediction, decision-tree-based algorithms and deep learning methodologies dominate the landscape, with the former widely used in economics due to their higher interpretability, and the latter more frequent in computer science where usually interpretability is de-emphasized in favor of higher predictive performance.

In general, the predictive ability of SL algorithms can play a fundamental role in boosting targeted policies at every stage of the lifespan of a firm—i.e., (1) identifying projects and companies with a high success propensity can aid the allocation of investment resources; (2) potential high growth companies can be directly targeted with supportive measures; (3) the higher ability to disentangle valuable and non-valuable firms can act as a screening device for potential lenders.

As granular data on the firm level becomes increasingly available, it will open many doors for future research directions focusing on SL applications for prediction tasks. To simplify future research in this matter, we briefly illustrated the principal SL algorithms employed in the literature of firm dynamics, namely, decision trees, random forests, support vector machines, and artificial neural networks. For a more detailed overview of these methods and their implementation in R we refer to our GitHub page (<http://github.com/fbargaglistoffi/machine-learning-firm-dynamics>), where we provide a simple tutorial to predict firms' bankruptcies.

Besides reaching a high-predictive power, it is important, especially for policy-makers, that SL methods deliver retractable and interpretable results. For instance, the US banking regulator has introduced the obligation for lenders to inform borrowers about the underlying factors that influenced their decision to not provide access to credit.<sup>4</sup> Hence, we argue that different SL techniques should be evaluated, and researchers should opt for the most interpretable method when the predictive performance of competing algorithms is not too different. This is central, as the understanding of which are the most important predictors, or which is the marginal effect of a predictor on the output (e.g., via partial dependency plots), can provide useful insights for scholars and policy-makers. Indeed, researchers and practitioners

---

<sup>4</sup>These obligations were introduced by recent modification in the Equal Credit Opportunity Act (ECOA) and the Fair Credit Reporting Act (FCRA).



can enhance models' interpretability using a set of ready-to-use models and tools that are designed to provide useful insights on the SL black box. These tools can be grouped into three different categories: tools and models for (1) complexity and dimensionality reduction (i.e., variables selection and regularization via Lasso, ridge, or elastic net regressions, see [70]); (2) model-agnostic variables' importance techniques (i.e., permutation feature importance based on how much the accuracy decreases when the variable is excluded, Shapley values, SHAP [SHapley Additive exPlanations], decrease in Gini impurity when a variable is chosen to split a node in tree-based methodologies); and (3) model-agnostic marginal effects estimation methodologies (average marginal effects, partial dependency plots, individual conditional expectations, accumulated local effects).<sup>5</sup>

In order to form a solid knowledge base derived from SL applications, scholars should put an effort in making their research as replicable as possible in the spirit of Open Science. Indeed, in the majority of papers that we analyzed, we did not find possible to replicate the reported analyses. Higher standards of replicability should be reached by releasing details about the choice of the model hyperparameters, the codes, and software used for the analyses as well as by releasing the training/testing data (to the extent that this is possible), anonymizing them in the case that the data are proprietary. Moreover, most of the datasets used for the SL analyses that we covered in this chapter were not disclosed by the authors as they are linked to proprietary data sources collected by banks, financial institutions, and business analytics firms (i.e., Bureau Van Dijk).

Here, we want to stress once more time that SL learning per se is not informative about the causal relationships between the predictors and the outcome; therefore researchers who wish to draw causal inference should carefully check the standard identification assumptions [49] and inspect whether or not they hold in the scenario at hand [6]. Besides not directly providing causal estimands, most of the reviewed SL applications focus on pointwise predictions where inference is de-emphasized. Providing a measure of uncertainty about the predictions, e.g., via confidence intervals, and assessing how sensitive predictions appear to unobserved points, are important directions to explore further [11].

In this chapter, we focus on the analysis of how SL algorithms predict various firm dynamics on "intercompany data" that cover information across firms. Yet, nowadays companies themselves apply ML algorithms for various clustering and predictive tasks [62], which will presumably become more prominent for small and medium-sized companies (SMEs) in the upcoming years. This is due to the fact that (1) SMEs start to construct proprietary data bases, (2) develop the skills to perform in-house ML analysis on this data, and (3) powerful methods are easily implemented using common statistical software.

Against this background, we want to stress that applying SL algorithms and economic intuition regarding the research question at hand should ideally com-

---

<sup>5</sup>For a more extensive discussion on interpretability, models' simplicity, and complexity, we refer the reader to [10] and [64].

plement each other. Economic intuition can aid the choice of the algorithm and the selection of relevant attributes, thus leading to better predictive performance [12]. Furthermore, it requires a deep knowledge of the studied research question to properly interpret SL results and to direct their purpose so that *intelligent machines are driven by expert human beings*.

## References

1. Ajit, P. (2016). Prediction of employee turnover in organizations using machine learning algorithms. *International Journal of Advanced Research in Artificial Intelligence*, 5(9), 22–26.
2. Alaka, H. A., Oyedele, L. O., Owolabi, H. A., Kumar, V., Ajayi, S. O., Akinade, O. O., et al. (2018). Systematic review of bankruptcy prediction models: Towards a framework for tool selection. *Expert Systems with Applications*, 94, 164–184.
3. Altman, E. I. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *The Journal of Finance*, 23(4), 589–609.
4. Arroyo, J., Corea, F., Jimenez-Diaz, G., & Recio-Garcia, J. A. (2019). Assessment of machine learning performance for decision support in venture capital investments. *IEEE Access*, 7, 124233–124243.
5. Athey, S. (2018). The impact of machine learning on economics. In *The economics of artificial intelligence: An agenda* (pp. 507–547). Chicago: University of Chicago Press.
6. Athey, S. & Imbens, G. (2019). *Machine learning methods economists should know about*, arXiv, CoRR abs/1903.10075.
7. Bajari, P., Chernozhukov, V., Hortaçsu, A., & Suzuki, J. (2019). The impact of big data on firm performance: An empirical investigation. *AEA Papers and Proceedings*, 109, 33–37.
8. Bakar, N. M. A., & Tahir, I. M. (2009). Applying multiple linear regression and neural network to predict bank performance. *International Business Research*, 2(4), 176–183.
9. Barboza, F., Kimura, H., & Altman, E. (2017). Machine learning models and bankruptcy prediction. *Expert Systems with Applications*, 83, 405–417.
10. Bargagli-Stoffi, F. J., Cevolani, G., & Gnecco, G. (2020). Should simplicity be always preferred to complexity in supervised machine learning? In *6th International Conference on machine Learning, Optimization Data Science (LOD2020), Lecture Notes in Computer Science*. (Vol. 12565, pp. 55–59). Cham: Springer.
11. Bargagli-Stoffi, F. J., De Beckker, K., De Witte, K., & Maldonado, J. E. (2021). Assessing sensitivity of predictions. A novel toolbox for machine learning with an application on financial literacy. arXiv, CoRR abs/2102.04382
12. Bargagli-Stoffi, F. J., Riccaboni, M., & Rungi, A. (2020). Machine learning for zombie hunting. firms’ failures and financial constraints. *FEB Research Report Department of Economics DPS20. 06*.
13. Baumann, A., Lessmann, S., Coussement, K., & De Bock, K. W. (2015). Maximize what matters: Predicting customer churn with decision-centric ensemble selection. In *ECIS 2015 Completed Research Papers*, Paper number 15. Available at: [https://aisel.aisnet.org/ecis2015\\_cr/15](https://aisel.aisnet.org/ecis2015_cr/15)
14. Behr, A., & Weinblat, J. (2017). Default patterns in seven EU countries: A random forest approach. *International Journal of the Economics of Business*, 24(2), 181–222.
15. Bento, F. R. d. S. R. (2018). *Predicting start-up success with machine learning*. B.S. thesis, Universidade NOVA de Lisboa. Available at: <https://run.unl.pt/bitstream/10362/33785/1/TGI0132.pdf>

16. Böhm, M., Weking, J., Fortunat, F., Müller, S., Welpel, I., & Kremer, H. (2017). The business model DNA: Towards an approach for predicting business model success. In *Int. En Tagung Wirtschafts Informatik* (pp. 1006–1020).
17. Bonello, J., Brédart, X., & Vella, V. (2018). Machine learning models for predicting financial distress. *Journal of Research in Economics*, 2(2), 174–185.
18. Brédart, X. (2014). Bankruptcy prediction model using neural networks. *Accounting and Finance Research*, 3(2), 124–128.
19. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
20. Breiman, L. (2001). Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical Science*, 16(3), 199–231.
21. Breiman, L. (2017). *Classification and regression trees*. New York: Routledge.
22. Buldyrev, S., Pammolli, F., Riccaboni, M., & Stanley, H. (2020). *The rise and fall of business firms: A stochastic framework on innovation, creative destruction and growth*. Cambridge: Cambridge University Press.
23. Chandra, D. K., Ravi, V., & Bose, I. (2009). Failure prediction of dotcom companies using hybrid intelligent techniques. *Expert Systems with Applications*, 36(3), 4830–4837.
24. Chipman, H. A., George, E. I., McCulloch, R. E. (2010). Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1), 266–298.
25. Cleofas-Sánchez, L., García, V., Marqués, A., & Sánchez, J. S. (2016). Financial distress prediction using the hybrid associative memory with translation. *Applied Soft Computing*, 44, 144–152.
26. Coad, A. (2009). *The growth of firms: A survey of theories and empirical evidence*. Northampton: Edward Elgar Publishing.
27. Coad, A., & Srhoj, S. (2020). Catching gazelles with a lasso: Big data techniques for the prediction of high-growth firms. *Small Business Economics*, 55, 541–565. <https://doi.org/10.1007/s11187-019-00203-3>
28. Coad, A., Frankish, J., Roberts, R. G., & Storey, D. J. (2013). Growth paths and survival chances: An application of gambler's ruin theory. *Journal of Business Venturing*, 28(5), 615–632.
29. Dalle, J.-M., Den Besten, M., & Menon, C. (2017). Using crunchbase for economic and managerial research. In *OECD Science, Technology and Industry Working Papers*, 2017/08. <https://doi.org/10.1787/6c418d60-en>
30. Danenas, P., & Garsva, G. (2015). Selection of support vector machines based classifiers for credit risk domain. *Expert Systems with Applications*, 42(6), 3194–3204.
31. Dellermann, D., Lipusch, N., Ebel, P., Popp, K. M., & Leimeister, J. M. (2017). Finding the unicorn: Predicting early stage startup success through a hybrid intelligence method. In *International Conference on Information Systems (ICIS)*, Seoul. Available at: <https://doi.org/10.2139/ssrn.3159123>
32. DiMasi, J., Hermann, J., Twyman, K., Kondru, R., Stergiopoulos, S., Getz, K., et al. (2015). A tool for predicting regulatory approval after phase ii testing of new oncology compounds. *Clinical Pharmacology & Therapeutics*, 98(5), 506–513.
33. Erel, I., Stern, L. H., Tan, C., & Weisbach, M. S. (2018). *Selecting directors using machine learning*. Technical report, National Bureau of Economic Research. Working paper 24435. <https://doi.org/10.3386/w24435>
34. Etter, V., Grossglauser, M., & Thiran, P. (2013). Launch hard or go home! predicting the success of kickstarter campaigns. In *Proceedings of the First ACM Conference on Online Social Networks* (pp. 177–182).
35. European Commission. (2010). *Communication from the commission: Europe 2020: A strategy for smart, sustainable and inclusive growth*. Publications Office of the European Union, 52010DC2020. Available at: <https://eur-lex.europa.eu/legal-content/en/ALL/?uri=CELEX%3A52010DC2020>
36. Fantazzini, D., & Figini, S. (2009). Random survival forests models for SME credit risk measurement. *Methodology and Computing in Applied Probability*, 11(1), 29–45.

37. Farboodi, M., Mihet, R., Philippon, T., & Veldkamp, L. (2019). Big data and firm dynamics. In *AEA Papers and Proceedings* (Vol. 109, pp. 38–42).
38. Feijoo, F., Palopoli, M., Bernstein, J., Siddiqui, S., & Albright, T. E. (2020). Key indicators of phase transition for clinical trials through machine learning. *Drug Discovery Today*, 25(2), 414–421.
39. Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, 29(5), 1189–1232.
40. Gibrat, R. (1931). *Les inégalités économiques: applications aux inégalités des richesses, à la concentration des entreprises... d'une loi nouvelle, la loi de l'effet proportionnel*. Paris: Librairie du Recueil Sirey.
41. Greenberg, M. D., Pardo, B., Hariharan, K., & Gerber, E. (2013). Crowdfunding support tools: predicting success & failure. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems* (pp. 1815–1820). New York: ACM.
42. Griffith, E. (2014). Why startups fail, according to their founders. *Fortune Magazine*, Last accessed on 12 March, 2021. Available at: <https://fortune.com/2014/09/25/why-startups-fail-according-to-their-founders/>
43. Guerzoni, M., Nava, C. R., & Nuccio, M. (2019). *The survival of start-ups in time of crisis. a machine learning approach to measure innovation*. Preprint. arXiv:1911.01073.
44. Halabi, C. E., & Lussier, R. N. (2014). A model for predicting small firm performance. *Journal of Small Business and Enterprise Development*, 21(1), 4–25.
45. Hassoun, M. H. (1995). *Fundamentals of artificial neural networks*. Cambridge: MIT Press.
46. Henrekson, M., & Johansson, D. (2010). Gazelles as job creators: a survey and interpretation of the evidence. *Small Business Economics*, 35(2), 227–244.
47. Heo, J., & Yang, J. Y. (2014). Adaboost based bankruptcy forecasting of Korean construction companies. *Applied Soft Computing*, 24, 494–499.
48. Hosaka, T. (2019). Bankruptcy prediction using imaged financial ratios and convolutional neural networks. *Expert Systems with Applications*, 117, 287–299.
49. Imbens, G. W., & Rubin, D. B. (2015). *Causal inference for statistics, social, and biomedical sciences: An introduction*. New York: Cambridge University Press.
50. Ishwaran, H., Kogalur, U. B., Blackstone, E. H., & Lauer, M. S. (2008). Random survival forests. *The Annals of Applied Statistics*, 2(3), 841–860.
51. Janssen, N. E. (2019). *A machine learning proposal for predicting the success rate of IT-projects based on project metrics before initiation*. B.Sc. thesis, University of Twente. Available at: <https://essay.utwente.nl/78526/>
52. Kaminski, J. C., & Hopp, C. (2020). Predicting outcomes in crowdfunding campaigns with textual, visual, and linguistic signals. *Small Business Economics*, 55, 627–649.
53. Kapelner, A., & Bleich, J. (2015). Prediction with missing data via Bayesian additive regression trees. *Canadian Journal of Statistics*, 43(2), 224–239.
54. Kim, S. Y., & Upneja, A. (2014). Predicting restaurant financial distress using decision tree and adaboosted decision tree models. *Economic Modelling*, 36, 354–362.
55. Kinne, J., & Lenz, D. (2019). Predicting innovative firms using web mining and deep learning. In *ZEW-Centre for European Economic Research Discussion Paper*, (19-01).
56. Kleinberg, J., Ludwig, J., Mullainathan, S., & Obermeyer, Z. (2015). Prediction policy problems. *American Economic Review*, 105(5), 491–495.
57. Kolkman, D., & van Witteloostuijn, A. (2019). Data science in strategy: Machine learning and text analysis in the study of firm growth. In *Tinbergen Institute Discussion Paper 2019-066/VI*. Available at: <https://doi.org/10.2139/ssrn.3457271>
58. Kotthoff, L. (2016). Algorithm selection for combinatorial search problems: A survey. In *Data Mining and Constraint Programming*, LNCS (Vol. 10101, pp. 149–190). Cham: Springer.
59. Krishna, A., Agrawal, A., & Choudhary, A. (2016). Predicting the outcome of startups: less failure, more success. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)* (pp. 798–805). Piscataway: IEEE.

60. Kyebambe, M. N., Cheng, G., Huang, Y., He, C., & Zhang, Z. (2017). Forecasting emerging technologies: A supervised learning approach through patent analysis. *Technological Forecasting and Social Change*, 125, 236–244.
61. Lam, M. (2004). Neural network techniques for financial performance prediction: integrating fundamental and technical analysis. *Decision support systems*, 37(4), 567–581.
62. Lee, I., & Shin, Y. J. (2020). Machine learning for enterprises: Applications, algorithm selection, and challenges. *Business Horizons*, 63(2), 157–170.
63. Lee, K. C., Han, I., & Kwon, Y. (1996). Hybrid neural network models for bankruptcy predictions. *Decision Support Systems*, 18(1), 63–72.
64. Lee, K., Bargagli-Stoffi, F. J., & Dominici, F. (2020). *Causal rule ensemble: Interpretable inference of heterogeneous treatment effects*, arXiv, CoRR abs/2009.09036
65. Liang, D., Lu, C.-C., Tsai, C.-F., & Shih, G.-A. (2016). Financial ratios and corporate governance indicators in bankruptcy prediction: A comprehensive study. *European Journal of Operational Research*, 252(2), 561–572.
66. Linn, M., & Weagley, D. (2019). Estimating financial constraints with machine learning. In *SSRN*, paper number 3375048. <https://doi.org/10.2139/ssrn.3375048>
67. Lo, A. W., Siah, K. W., & Wong, C. H. (2019). Machine learning with statistical imputation for predicting drug approvals. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.5c5f0525>
68. Lussier, R. N., & Halabi, C. E. (2010). A three-country comparison of the business success versus failure prediction model. *Journal of Small Business Management*, 48(3), 360–377.
69. Lussier, R. N., & Pfeifer, S. (2001). A cross-national prediction model for business success. *Journal of Small Business Management*, 39(3), 228–239.
70. Martínez, J. M., Escandell-Montero, P., Soria-Olivas, E., MartíN-Guerrero, J. D., Magdalena-Benedito, R., & Gómez-Sanchis, J. (2011). Regularized extreme learning machine for regression problems. *Neurocomputing*, 74(17), 3716–3721.
71. Matin, R., Hansen, C., Hansen, C., & Molgaard, P. (2019). Predicting distresses using deep learning of text segments in annual reports. *Expert Systems with Applications*, 132(15), 199–208.
72. McKenzie, D., & Sansone, D. (2017). Man vs. machine in predicting successful entrepreneurs: evidence from a business plan competition in Nigeria. In *World Bank Policy Research Working Paper No. 8271*. Available at: <https://ssrn.com/abstract=3086928>
73. Megaravalli, A. V., & Sampagnaro, G. (2019). Predicting the growth of high-growth SMEs: evidence from family business firms. *Journal of Family Business Management*, 9(1), 98–109. <https://doi.org/10.1108/JFBM-09-2017-0029>
74. Meinshausen, N., & Bühlmann, P. (2010). Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4), 417–473.
75. Mikalef, P., Boura, M., Lekakos, G., & Krogstie, J. (2019). Big data analytics and firm performance: Findings from a mixed-method approach. *Journal of Business Research*, 98, 261–276.
76. Miyakawa, D., Miyauchi, Y., & Perez, C. (2017). *Forecasting firm performance with machine learning: Evidence from Japanese firm-level data*. Technical report, Research Institute of Economy, Trade and Industry (RIETI). Discussion Paper Series 17-E-068. Available at: <https://www.rieti.go.jp/jp/publications/dp/17e068.pdf>
77. Moscatelli, M., Parlapiano, F., Narizzano, S., & Viggiano, G. (2020). Corporate default forecasting with machine learning. *Expert Systems with Applications*, 161(15), art. num. 113567
78. Mullainathan, S., & Spiess, J. (2017). Machine learning: an applied econometric approach. *Journal of Economic Perspectives*, 31(2), 87–106.
79. Munos, B., Niederreiter, J., & Riccaboni, M. (2020). Improving the prediction of clinical success using machine learning. In *EIC Working Paper Series*, number 3/2020. Available at: <http://eprints.imtlucca.it/id/eprint/4079>
80. Ng, A. Y., & Jordan, M. I. (2002). On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes. In *Advances in neural information processing systems*,

- NIPS 2001 (Vol. 14, pp. 841–848), art code 104686. Available at: <https://papers.nips.cc/paper/2001/file/7b7a53e239400a13bd6be6c91c4f6c4e-Paper.pdf>
81. Ohlson, J. A. (1980). Financial ratios and the probabilistic prediction of bankruptcy. *Journal of Accounting Research*, 18(1), 109–131.
  82. Qiu, X. Y., Srinivasan, P., & Hu, Y. (2014). Supervised learning models to predict firm performance with annual reports: An empirical study. *Journal of the Association for Information Science and Technology*, 65(2), 400–413.
  83. Ravi, V., Kurniawan, H., Thai, P. N. K., & Kumar, P. R. (2008). Soft computing system for bank performance prediction. *Applied Soft Computing*, 8(1), 305–315.
  84. Rouhani, S., & Ravasan, A. Z. (2013). ERP success prediction: An artificial neural network approach. *Scientia Iranica*, 20(3), 992–1001.
  85. Saradhi, V. V., & Palshikar, G. K. (2011). Employee churn prediction. *Expert Systems with Applications*, 38(3), 1999–2006.
  86. Sejnowski, T. J. (2018). *The deep learning revolution*. Cambridge: MIT Press.
  87. Sharchilev, B., Roizner, M., Rumyantsev, A., Ozornin, D., Serdyukov, P., & de Rijke, M. (2018). Web-based startup success prediction. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (pp. 2283–2291).
  88. Shin, K.-S., Lee, T. S., & Kim, H.-j. (2005). An application of support vector machines in bankruptcy prediction model. *Expert Systems with Applications*, 28(1), 127–135.
  89. Steinwart, I., & Christmann, A. (2008). *Support vector machines*. New York: Springer Science & Business Media.
  90. Su, L., Shi, Z., & Phillips, P. C. (2016). Identifying latent structures in panel data. *Econometrica*, 84(6), 2215–2264.
  91. Sun, J., & Li, H. (2011). Dynamic financial distress prediction using instance selection for the disposal of concept drift. *Expert Systems with Applications*, 38(3), 2566–2576.
  92. Sun, J., Fujita, H., Chen, P., & Li, H. (2017). Dynamic financial distress prediction with concept drift based on time weighting combined with Adaboost support vector machine ensemble. *Knowledge-Based Systems*, 120, 4–14.
  93. Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288.
  94. Tsai, C.-F., & Wu, J.-W. (2008). Using neural network ensembles for bankruptcy prediction and credit scoring. *Expert Systems with Applications*, 34(4), 2639–2649.
  95. Tsai, C.-F., Hsu, Y.-F., & Yen, D. C. (2014). A comparative study of classifier ensembles for bankruptcy prediction. *Applied Soft Computing*, 24, 977–984.
  96. Udo, G. (1993). Neural network performance on the bankruptcy classification problem. *Computers & Industrial Engineering*, 25(1–4), 377–380.
  97. Van der Laan, M. J., Polley, E. C., & Hubbard, A. E. (2007). Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1), Article No. 25. <https://doi.org/10.2202/1544-6115.1309>
  98. van Witteloostuijn, A., & Kolkman, D. (2019). Is firm growth random? A machine learning perspective. *Journal of Business Venturing Insights*, 11, e00107.
  99. Wang, G., Ma, J., & Yang, S. (2014). An improved boosting based on feature selection for corporate bankruptcy prediction. *Expert Systems with Applications*, 41(5), 2353–2361.
  100. Weinblat, J. (2018). Forecasting European high-growth firms—a random forest approach. *Journal of Industry, Competition and Trade*, 18(3), 253–294.
  101. Xiang, G., Zheng, Z., Wen, M., Hong, J., Rose, C., & Liu, C. (2012). A supervised approach to predict company acquisition with factual and topic features using profiles and news articles on techcrunch. In *Sixth International AAAI Conference on Weblogs and Social Media (ICWSM 2012)*. Menlo Park: The AAAI Press. Available at: <http://dblp.uni-trier.de/db/conf/icwsm/icwsm2012.html#XiangZWHRL12>
  102. Yankov, B., Ruskov, P., & Haralampiev, K. (2014). Models and tools for technology start-up companies success analysis. *Economic Alternatives*, 3, 15–24.
  103. Zarin, D. A., Tse, T., Williams, R. J., & Carr, S. (2016). Trial Reporting in ClinicalTrials.gov – The Final Rule. *New England Journal of Medicine*, 375(20), 1998–2004.

104. Zhang, Q., Ye, T., Essaidi, M., Agarwal, S., Liu, V., & Loo, B. T. (2017). Predicting startup crowdfunding success through longitudinal social engagement analysis. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management* (pp. 1937–1946).
105. Zikeba, M., Tomczak, S. K., & Tomczak, J. M. (2016). Ensemble boosted trees with synthetic features generation in application to bankruptcy prediction. *Expert Systems with Applications*, 58, 93–101.

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

