Decision Support Systems 107 (2018) 64–77

Contents lists available at ScienceDirect

Decision Support Systems

journal homepage: www.elsevier.com/locate/dss

Failure pattern-based ensembles applied to bankruptcy forecasting

Philippe du Jardin

Edhec Business School, 393 Promenade des Anglais, BP 3116, Nice Cedex 3 06202, France

ARTICLE INFO

Article history:
Received 27 August 2017
Received in revised form 16 January 2018
Accepted 17 January 2018

Keywords:
Decision support systems
Ensemble-based models
Self-organizing map
Bankruptcy forecasting

ABSTRACT

Bankruptcy prediction models that rely on ensemble techniques have been studied in depth over the last 20 years. Within most studies that have been performed on this topic, it appears that any ensemble-based model often achieves better results than those estimated with a single model designed using the base classifier of the ensemble, but it is not uncommon that the results of the former model do not outperform those of a single model when estimated with any other classifier. Indeed, an ensemble of decision trees is almost always more accurate than a single tree but not necessarily more than a neural network or a support vector machine. We know that the accuracy of an ensemble used to forecast firm bankruptcy is closely related to its ability to capture the variety of bankruptcy situations. But the fact that it may not be more efficient than a single model suggests that current techniques used to handle such a variety are not completely satisfactory. This is why we have looked for a method that makes it possible to better embody this diversity than current ones do. The technique proposed in this article relies on the quantification, using Kohonen maps, of temporal patterns that characterized the financial health of a set of companies, and on the use of an ensemble of incremental size maps to make forecasts. The results show that such models lead to better predictions than those that can be achieved with traditional methods.

1. Introduction

Models that have been studied in the financial literature and that are used to forecast bankruptcy are primarily default models: a firm goes bankrupt when it lacks sufficient resources to meet its financial obligations, hence when it becomes insolvent. Most empirical studies that have focused on bankruptcy prediction have therefore attempted to find measures that characterize a risk of default. The first models developed in the 1960s, following the study by Altman [4], have sought to assess this risk by estimating the distance between the financial situation of a given firm and a standard bankruptcy situation. Virtually all data-mining techniques that have been developed for classification purposes have been used to design failure models that share almost all the same characteristics: models are dichotomous, have good forecasting abilities and are easy to estimate. However, what can be considered the main factor of their success is also their main weakness. They essentially rely on a single rule and are estimated using financial data that solely characterize a unique period of firm life. This type of modeling reflects a rather rudimentary view of bankruptcy; it is considered the result of a a-historical process [39] that does not depend on time and that is reducible to a limited number of measures. But reality is a bit different. One knows that firms that apparently share the same financial profile, from the point of view of a model, may in reality have a very different probability of failure. Over time, some of them may have gained a certain resilience that gives them the ability to withstand failure. Some others may have received from their environment a sort of carrying capacity that has changed their fate at the very moment where their situation worsened, or have managed to recover even though nothing suggested they were able to do so [11]. All these factors, which can solely be analyzed over time, cannot be properly embodied by traditional models.

The historical dimension of failure and the multiplicity of the situations that lead to bankruptcy have given rise to a large body of literature. We can find, on the one side, studies that focused on the temporal dimension of financial failure. They analyzed the way variables that measure firm activity over several years [22] may influence model accuracy, assuming that taking time into account with multi-period data would be sufficient to embody the dynamics of the phenomenon. We can also find, on the other side, studies that were interested in modeling the different financial situations that lead to bankruptcy. They especially analyzed how to embody at-risk situations using ensemble-based models, this time assuming that the multiplication of forecasting rules would make it possible to model the diversity of failure symptoms [32,46,61]. In both cases, models on the whole lead to better results than those estimated with single models, but it seems that multi-rule models sound more promising

E-mail address: philippe.dujardin@edhec.edu.

https://doi.org/10.1016/j.dss.2018.01.003
0167-9236/

than multi-period ones. Indeed, the sole use of historical data hardly changes model physiognomy; they are still singular and still stumble on the fact that they embody a unique measure of a distance to bankruptcy which is too simple to be truly effective. In fact, ensemble models make it possible to represent the different facets of the decision boundary that separate failed from non-failed firms [9,49] and, to a certain extent, are able to embody the different patterns of decline that may lead to failure and that traditional models cannot assess.

For all that, when one looks closely at the results achieved with all these models and compares those estimated with ensemble techniques to those calculated with the best single models, in many cases, discrepancy between them does exist but it is relatively low (between 2 and 3%) and is often is not statistically significant [21,23,27,43,49,63]. Of course, a model designed with an ensemble of decision trees is often more accurate than one estimated with a single tree, but not necessarily much more powerful than a single model estimated with a support vector machine. This shows that ensembles represent a source of performance that is related to the way they are able to capture the variety of bankruptcy situations. But the fact that ensembles are not systematically more accurate than single models suggests that they do not handle variety in a completely optimal way. These are the reasons why we have studied a method that makes it possible to better embody the different bankruptcy situations, using what we call “failure patterns”, and to use these “patterns” to make forecasts. It relies on the quantification, with Kohonen maps, of temporal patterns that characterize both failed and non-failed firms, and on the use of an ensemble of incremental size maps to make forecasts.

2. Literature review

The very first bankruptcy prediction models that were developed, following that of Altman [4], represent failure as if it were a critical situation that might be captured using a measure of the discrepancy that separates the financial situation of a given firm from that of a standard critical situation. These models also make the assumption that failure is a phenomenon that can be estimated using a unique measure of firm financial health. Finally, they assume that failure corresponds to a particular event that can be explained using a single classification rule, and to a certain extent that bankruptcy is the result of a unique process of decline. The limits of these models are the direct consequences of their assumptions. We know that bankruptcy has multiple causes and symptoms, and that a model with variables that are solely measured over a single period would probably not be able to embody such diversity. We also know that failure does not depend on the sole situation of a firm at a given period of its life, but is the result a protracted process that cannot be captured properly by models that do not take into account a temporal dimension. Finally, we know that different paths to firm failure exist [11,39,48], but their complexity cannot be properly assessed with simple single-rule models.

In order to overcome these limitations, some research works have sought to estimate and use the historical dimension of bankruptcy through multi-period data, but still using traditional single-rule models. They have shown that models designed with financial variables measured over several years lead to better predictions than those achieved with models designed with single period variables [22,65]. But their structure can solely solve a part of the issue: the uniqueness of the rule leads to models that still embody a single standard bankruptcy situation. This is why some other research works have attempted to model the variety of bankruptcy situations using multiple classification rules. The techniques that have been used for this purpose aim at designing a meta-model where each component has a particular expertise on a certain region of the decision space. If these components are sufficiently diverse [38], they make it possible to estimate models that are on average more accurate than single models.

The characteristics of ensemble-based techniques rely on the way classification rules are developed and combined. With certain techniques, rules can be estimated with the same modeling method. This is the case of bagging, boosting, random subspace... where all models are estimated using either a decision tree [23], logistic regression [43], a feed-forward neural network [32], a survival model [16], k-nearest neighbors [49] or a support vector machine [61]. Rules can also be computed with different methods: for example, a model estimated with a logistic regression can be used in conjunction with other types of models that are assessed either with a support vector machine [25], or with both a support vector machine and a neural network [70], or with a neural network combined with a decision tree and discriminant analysis [46]. Sometimes, combinations are more complex; a first set of models is computed with different methods and their results are then used as inputs of a final model estimated with a neural network [7]. Finally, rules can be estimated after a prior segmentation of the decision space. The technique consists of grouping observations into a few classes and then calculating models where each of them fits the characteristics of a given class, using the same methods as those presented above [15,16,63].

On the whole, ensemble models present a better ability to make accurate forecasts than single models do. However, the absolute gain brought by these techniques is relatively low compared to that of traditional models, even if it remains real. We have measured the average gain calculated over 31 studies published between 2000 and 2017, and presented in Table 1[1]. If one calculates the difference, for each study, between the correct classification rate of the best ensemble-based model and that of the best single model from among those that have been estimated while taking into account the sample size used for their validation[2], the average gain does not even reach 2.4%. And, above all, if one estimates whether the differences between these models are significant, one can notice that among the 31 differences calculated in Table 1, only 10 are statistically significant.

The literature shows that the ability of a model to capture the whole variety of bankruptcy situations is a key factor of its performance. But it also shows that usual ensemble-based models are not able to easily embody this variety because, each rule, solely represents a boundary between two groups, although we might likely consider more subtle modeling. Literature has long shown the existence of different profiles or failure “patterns” which represent prototype situations that firms may experience over their life, and where some of them may lead to bankruptcy. But it has mainly focused on very general patterns that are shared by the largest number of firms. Those that were estimated by D’Aveni [11], Laitinen [39] or by Lukason et al. [48] illustrate this finding. However one may think that these patterns are far too general and could be refined. A widespread cause of bankruptcy, such as a lack of liquidity, can be embodied in many different financial situations that are very unequally distributed within a population of firms. A firm may not be liquid because of a lack of cash, a lack of permanent capital, a problem of balance between payables and receivables... This suggests that non-liquid firms can be represented through a wide variety of profiles and that a large number of illiquidity patterns should

[1] Table 1 lists the main studies that have been published since 2000 and that have studied the accuracy of ensemble-based models when it comes to forecasting bankruptcy. For each study, this table presents the results achieved with the best ensemble-based models (maximum 3) and those achieved with the best single model from among all single models that have been estimated.

[2] When the size of the test sample was not indicated, we estimated the gain using the size of the learning sample, despite the positive bias introduced in the estimation.

Table 1
Best results of the main studies that have estimated ensemble-based models to forecast financial failure.

Studies	Ensemble-based models	Accuracy (1)(%)	Single models	Accuracy (2)(%)	Difference (1−2)(%)	Forecasting horizon*	Sample size	
							F**	NF
Alfaro et al. [3]	AdaBoost (DT)	91.10	BPNN	87.29	3.81	1 Y	590	590
Alfaro et al. [2]	AdaBoost (DT)	89.16	DT	88.56	0.60	1 Y	1365	1365
Cho et al. [7]	(DA+LR+BPNN+RI)+BPNN	78.92	DA	78.15	0.77		900	900
Chuang [9]	RS-CBR	94.10	LR	85.30	8.80		42	279
	DT-CBR	90.40			5.10			
du Jardin [15]	SOM-Random subspace (BPNN)	87.04	BPNN	83.97	3.07	1 Y	>8000	>8000
	SOM-Boosting (BPNN)	85.93			1.96			
	SOM-Bagging (BPNN)	84.98			1.01			
du Jardin [16]	SOM-Boosting (SVM)	91.13	ELM	89.11	2.02	1 Y	>95000	>1800
	SOM-Rotation forest (SVM)	90.87			1.76			
	SOM-Random subspace (Cox)	90.85			1.74			
Fedorova et al. [19]	AdaBoost (BPNN)	88.80	BPNN	87.80	1.00		444	888
Geng et al. [21]	DT+SVM+BPNN	78.40	BPNN	78.80	−0.40	3 Y	107	107
Heo and Yang [23]	AdaBoost (DT)	78.52	BPNN	77.08	1.44	1 Y	1381	1381
Hua et al. [25]	LR+SVM	94.66	SVM	91.38	3.28		60	60
	Bagging (SVM)	86.82			0.21			
Huang et al. [27]	MultiBoost (DT)	88.13	SVM	86.61	1.52		50	100
	AdaBoost (DT)	86.74			0.13			
Hung and Chen [28]	DT-BPNN-SVM	72.50	BPNN	72.37	0.13		112	128
Karthik-Chandra et al. [30]	Random forest+SVM+BPNN	93.33	DT	75.83	17.50		120	120
	Random forest	81.67			5.84			
Kim et al. [34]	AdaBoost (SVM)	95.20				1 Y	500	2500
Kim and Kang [32]	Bagging (BPNN)	75.97	BPNN	71.02	4.95	1 Y	728	729
	AdaBoost (BPNN)	75.10			4.08			
Kim and Kang [33]	Boosting (SVM)	77.53	SVM	72.45	5.08	1 Y	600	600
	Bagging (SVM)	77.23			4.78			
Kim and Upneja [35]	AdaBoost (DT)	97.70	DT	96.73	0.97	1 Y	42	779
Li et al. [43]	Random subspace (LR)	83.97	DA	82.98	0.99	1 Y	135	135
Li and Sun [44]	CBR-based ensemble	89.33	CBR	89.16	0.17	1 Y	135	135
Liao et al. [45]	SVM+BPNN	96.94	BN	94.82	2.12	1 Y	63	2680
	Random forest (DT)	94.91			0.09			
Lin and McClean [46]	DA+LR+BPNN+DT	89.60	DT	88.70	0.90	1 Y	154	979
Lopez-Iturriaga and Pastor-Sanz [47]	MLP-SOM	96.15	BPNN	93.27	2.88	1 Y	386	386
Marques et al. [49]	Rotation forest (k-NN)	79.17	k-NN	75.42	3.75	2 Y	112	128
	Bagging (BPNN)	78.75			3.33			
	Decorate (k-NN)	77.50			2.08			
Sun et al. [61]	AdaBoost (SVM)	86.29	SVM	79.77	6.52	2 Y	466	466
Sun et al. [62]	AdaBoost (SAT)	97.22	DT	96.54	0.68	1 Y	346	346
Tsai [63]	SOM-(LR+BPNN+DT)	90.14	DT	85.96	-8.43	2 Y	112	128
Tsai and Hsu [64]	Meta-classifier (BPNN)	97.33	BPNN	69.58	27.75	2 Y	112	128
	Stacking (BPNN)	70.83			1.25			
Wang et al. [66]	Boosting (DT)	81.50	DA	74.04	7.46	2 Y	112	128
	Bagging (DT)	76.21			2.17			
Wang and Wu [67]	SOM-based ensemble	94.54				2 Y	108	108
West et al. [68]	Bagging (BPNN)	87.37	BPNN	86.86	0.51		93	236
	Boosting (BPNN)	87.24			0.38			
	BPNN-based ensemble	87.08			0.22			
Xiao et al. [70]	DSE+(LR+BPNN+SVM)	87.80	SVM	85.43	2.37	2 Y	92	161
	LR+BPNN+SVM	86.25			0.82			

BN: Bayes network; BPNN: back-propagation neural network; CBR: case-based reasoning; Cox: Cox's model; DA: discriminant analysis; DSE: Dempster-Shafer evidence; DT: decision tree; GRA: gray relational analysis; k-NN: k-nearest neighbors; LR: logistic regression; RI: rule induction; RS: rough set; SAT: single-attribute test; SOM: self-organizing map; SVM: support vector machine. * Y: Years - ** F: failed; NF: non-failed.

also exist. The same reasoning applies to all causes of bankruptcy. Therefore, we assume that there exists a large variety of pre-bankruptcy situations we call failure "patterns" and that a model, which would make it possible to represent the information that characterizes each of these situations, would likely make better forecasts than those made with traditional methods. In this study, we propose to quantify these "patterns" using an ensemble of Kohonen maps and financial data that characterize firms over several years. The particularity of this ensemble lies in the fact that it is made up of incremental size maps, where each map has one more neuron than the former. Once the quantification is finished, neurons are labeled according to the class (failed vs. non-failed) for which they can be considered a prototype, and are used to make forecasts.

The way we use Kohonen maps breaks with usual practices. In the field of business failure, Kohonen's maps were first used, as a clustering technique, to better understand what characterized failed and non-failed firms [13] or to identify companies that would eventually go bankrupt [1,57]. But they were also used to perform classification tasks. Thus, Huysmans et al. [29] used a single map to make forecasts; du Jardin and Severin [17] designed a map to estimate a set of bankruptcy trajectories and used them as classification rules, du Jardin [14–16] designed maps to cluster firms that share the same bankruptcy trajectory, then built classification models that fit each trajectory; Tsai [63] estimated a map to divide a sample of firms into a few groups and then designed, using traditional methods, as many prediction models as there were groups; Lee et al. [40] and Lopez-Iturriaga and Pastor-Sanz [47] used a map in conjunction with a neural network to make forecasts. In the field of classification, Kohonen maps were also used to estimate prediction models, either with single maps or with an ensemble of maps. When ensembles of maps were considered, the generation of the different components often relied on two basic techniques: each map was either trained using a subset of variables drawn from an initial set [8] or using a re-sampling scheme applied to an original learning sample [56]. By contrast, in this study, we use ensembles of incremental size maps (each map has one more neuron than the former), assuming that the way we produce variety among classifiers would improve model accuracy.

In order to provide a robust estimation of the performance of our method, we applied it to multiple samples and we compared the different results with those calculated with usual modeling techniques. We used discriminant analysis, logistic regression, a decision tree, Cox's model, a feed-forward neural network, an extreme learning machine and a support vector machine, and all these techniques were used to design single models but also ensemble-based models in conjunction with bagging, boosting, random subspace and rotation forest. We also designed hybrid models that were estimated with different modeling methods [7,14,30,46,64].

3. Samples and variables

3.1. Samples

Samples were collected using a database managed by bureau Van Dijk, Diane (https://diane.bvdinfo.com), which provides balance sheets and income statements of firms that filed their annual accounts with the French commercial courts. Data were gathered over different years, between 2006 and 2014, to control for the economic environment effect on model performance. This period of time is interesting since over these 9 years, the French economy experienced several periods of growth (2007, 2010, 2011, 2013 and 2014) but also periods of downturn (2008) or recession (2009, 2012), and we know that when variations occur within the firm environment, models tend to become less accurate than they are when the environment remains stable [53]. For each year, we selected two samples: one to estimate model parameters, the other to test model accuracy. Learning and test samples do not share any common data and were not collected over the same period, but with a lag of one year, to estimate an out-of sample and out-of time error [60]. Table 2 presents the characteristics of each sample. Thus, for the first period studied (2006–2007), firms from the learning sample were selected in 2006 and those from the test sample were collected in 2007. Companies were chosen at time t, if they were still operating at time $t+1$, or if they were liquidated or reorganized at time $t+1$ by court decision. None of the failed firms filed for bankruptcy protection. Firms were drawn at random from among those for which at least 5 consecutive years of financial data were available in the database. We chose 5 years because it is over the last 5 years of their life that the differences between the two groups of firms truly widen.

Learning samples are made up of as many failed firms as non-failed firms, as is commonly done in the literature (Table 1). It is true that samples are (sometimes) randomly drawn to avoid a choice-based sample bias. However, the failure patterns-based models we propose in this study were designed using Kohonen maps. With this method, if the proportion of firms belonging to one class is very different from that of the other class, then the minority class cannot be quantified properly. This is why we chose to balance the proportions with each learning sample and each modeling method. However, since the proportion of failed firms is, each year, on average slightly lower than 2%, we estimated model accuracy using this proportion. Test samples were then built in two steps. First, we collected, for each period, as many failed as non-failed firms. Then, to make the estimation of the error as independent from sampling variations as possible, we bootstrapped each test sample 100 times and we averaged the error over these samples, each of which is made up of 6000 non-failed firms and 120 failed firms.

3.2. Explanatory variables

Variables were chosen within those commonly used in the literature [41] using two criteria to avoid, as much as possible, any arbitrary choice. First, variables must belong to one of the different dimensions that are used in financial analysis to explain bankruptcy. We chose 6 dimensions: activity, financial structure, liquidity, profitability, solvency and turnover. Second, they must present a good mid-term discrimination ability to make models as robust as possible against any variations that may occur within a firm's macroeconomic environment. Variables are presented in Table 3. They are all financial ratios estimated using firm balance sheets and income statements.

4. Modeling methods

4.1. Single and ensemble-based methods

The accuracy of the models that were developed using Kohonen maps, and that we call "failure pattern-based models", was benchmarked against that of models designed using common methods [58]. We chose two types of methods: some used to build single models – discriminant analysis [4], logistic regression [52], a decision tree with Cart [6], Cox's method [10], a feed-forward neural network [69], a support vector machine [18] and an extreme learning machine [26] – and also some used to design ensemble-based models [38] – bagging [5], boosting [20,55] with AdaBoost, random subspace [24], rotation forest [54]. We also estimated hybrid ensemble-based models proposed in the literature and that were built with different modeling methods: two of them rely on a combination of models [30,46], two others use a sequence of models [7,64] and the last one relies on models based on an *a priori* segmentation of data [14][3] The methods used in this study to design single models but also most of those used to estimate ensemble-based models such as bagging, boosting, random subspace and rotation forest are well-known. This is the reason why they are not presented. We just give an overview of the Kohonen algorithm to present the parameters used to design maps[4].

[3] Lin and McClean [46] designed models with four methods (logistic regression, discriminant analysis, a decision tree and a feed-forward neural network), then estimated the correct classification rate of each model, and finally calculated the accuracy of the combination of the four models by weighting each individual prediction with the correct classification rate of its corresponding model. Karthik-Chandra et al. [30] estimated five models (using random forest, logistic regression, a decision tree, a support vector machine and a feed-forward neural network), then chose the three best models (assessed with the AUC of a ROC curve) and combined their individual predictions using a majority vote. Cho et al. [7] designed four models (with logistic regression, discriminant analysis, a decision tree and a feed-forward neural network). With each model, they estimated the predicted class of each observation and they weighted this prediction using the correct classification rate of the model. The weighted predictions (one per model) were assigned to each observation. Then, a feed-forward neural network was estimated with all explanatory variables and was trained to estimate the weighted predictions. This network was used to calculate the final forecasts. Tsai and Hsu [64] calculated a set of models (with logistic regression, a decision tree and a feed-forward neural network) and selected the data that were correctly classified by at least two models. Then, these data were used to estimate a final model computed with a feed-forward neural network. Finally, du Jardin [14] estimated the evolution of the financial situation of a set of firms over 3 years, then quantified these individual evolutions into a set of prototype sequences of evolution, using a Kohonen map, and finally designed as many models as there were different sequences (using each time discriminant analysis, Cox's method, a decision tree and a feed-forward neural network).

[4] We used Weka to design and test all models. We also developed several routines with Eclipse that were then embedded into Weka. All computations related to the estimation of the Kohonen maps were developed using Visual Basic. Computations of financial ratios and their analysis were conducted with SPSS.

Table 2
Characteristics of the different samples.

Firm status	Periods															
	2006–2007		2007–2008		2008–2009		2009–2010		2010–2011		2011–2012		2012–2013		2013–2014	
	LS*	ITS	LS	ITS	LS	ITS	LS	ITS	LS	ITS	LS	ITS	LS	ITS	LS	ITS
Failed	7400	7200	7200	7500	7500	7200	7200	7450	7450	6850	6850	7050	7050	7400	7400	7350
Non-failed	7400	7200	7200	7500	7500	7200	7200	7450	7450	6850	6850	7050	7050	7400	7400	7350
		FTS**		FTS		FTS		FTS		FTS		FTS		FTS		FTS
Failed		120		120		120		120		120		120		120		120
Non-failed		6000		6000		6000		6000		6000		6000		6000		6000

* LS: learning sample; ITS: initial test sample. ** FTS: final test sample. Final test samples are bootstrap samples: firms were randomly drawn 100 times from the initial test samples using the same % of failed firms, which is on average slightly lower than 2%, as that which characterizes the different periods.

4.2. Self-organizing map

A Kohonen map [37] is made up of a set of neurons and each neuron is represented with a set of weights: each weight corresponds to a variable that characterizes a sample of observations. A map is estimated using a learning process during which each weight is adjusted gradually to finally reflect the initial structure of the data. Over this process, each observation is first compared to each neuron, through a distance calculation. Once the nearest neuron to a given observation is found, its weights are adjusted to decrease the distance between them. Then, the weights of the neurons located in the neighborhood of the latter neuron are also adjusted, but the magnitude of the modification is inversely proportional to their distance on the map. The process is repeated a certain number of times, which corresponds to a number of iterations that is to be set up *a priori*. Weights are updated with the following rule:

$$w_k(t+1) = w_k(t) + \alpha(t)h_{ck}(t)\left[x(t) - w_k(t)\right]$$

where t represents a given iteration, $\alpha(t)$ the learning step, $x(t)$ the input vector, $w_k(t)$ a neuron and $h_{ck}(t)$ the neighborhood function between the neuron $w_c(t)$, which is the closest to the input $x(t)$, and the neuron $w_k(t)$. We used the following decreasing function of the learning step:

$$\alpha(t) = \alpha_0 / \left(1 + 100 * t/tmax\right)$$

with α_0 the initial value of the learning step and *tmax* the length of the learning phase that corresponds to the number of times a learning sample is presented to the map.

To choose the length of the learning process (*tmax*), and set up the initial value of the learning step (α_0), we conducted a set of experiments. With each learning sample, we drew at random 50 sets of variables, and with each set, we designed several maps that ranged in size from 20 to 200, with a step of 20. We tested different values of *tmax* between 10 000 and 200 000, with a step of 10 000. With these values, we also tested five different values of α_0: between 0.1 and 0.5, with a step of 0.1. On average, the quantification error of each map was the lowest when the number of iterations was around 600 times the number of neurons, and the learning step was set up using 0.1. These are the settings we used to design all maps.

And we used the following neighborhood function:

$$h_{ck}(t) = max\left[0, 1 - (\sigma(t) - d_{ck})^2\right]$$

where $\sigma(t)$ is the neighborhood radius with $\sigma(t) = \sigma_0 * (\sigma_n/\sigma_0)^{t/tmax}$, σ_0 and σ_n are the initial and final value of the neighborhood radius, and d_{ck} is the distance on the map between neuron $w_c(t)$ and neuron $w_k(t)$. With each map, the initial value of the neighborhood radius was set up using the number of neurons of the map, and the final value using 1.

While observations are companies that are characterized by financial variables, each neuron can be considered a pattern that embodies the financial situation of a subset of firms.

Table 3
Initial set of variables.

Ratios		Dimensions*	Ratios		Dimensions
AP/TS	Accounts Payable/Total sales	TO	FE/TA	Financial expenses/Total assets	SO
C/CA	Cash/Current assets	LI	FE/VA	Financial expenses/Value added	SO
C/CL	Cash/Current liabilities	LI	I/TS	Inventories/Total sales	TO
C/TA	Cash/Total assets	LI	LTD/TA	Long Term debt/Total assets	FS
CA/TS	Current assets/Total sales	TO	NI/TS	Net income/Total sales	AC
CF/SF	Cash flow/Shareholder funds	PR	NI/VA	Net income/Value added	AC
CL/TA	Current liabilities/Total assets	FS	PBT/SF	Profit before Tax/Shareholder funds	PR
CL/TS	Current liabilities/Total sales	TO	QA/CL	Quick assets/Current liabilities	LI
EBIT/VA	EBIT/Value added	AC	R/TS	Receivables/Total sales	TO
EBITDA/PE	EBITDA/Permanent equity	PR	SF/PE	Shareholder funds/Permanent equity	FS
EBITDA/TA	EBITDA/Total assets	PR	SF/TA	Shareholder funds/Total assets	FS
EBITDA/TS	EBITDA/Total sales	AC	TD/TA	Total debt/Total assets	FS
FE/CF	Financial expenses/Cash flow	SO	TS/TA	Total sales/Total assets	TO
FE/EBITDA	Financial expenses/EBITDA	SO	VA/FA	Value added/Fixed assets	PR
FE/NI	Financial expenses/Net income	SO	VA/TS	Value added/Total sales	AC

EBIT: Earnings before interest and taxes; EBITDA: Earnings before interest, taxes, depreciation and amortization.
* AC: activity; FS: financial structure; LI: liquidity; PR: profitability; SO: solvency; TO: turnover.

5. Experimental settings

5.1. Variable selection

"Failure patterns" correspond to prototypic forms that characterize different financial situations and where some of them embody weak situations that may arise many years before bankruptcy occurs. These patterns were estimated with financial ratios that were measured over several consecutive years to take into account the influence of firm history on their fate. A pattern thus represents a temporal situation that is shared by a subset of companies. Patterns and all models were built with the same variables. We made this choice to control for the influence of the number of variables but also for that of variable informational content on model accuracy. We selected an initial set of 30 ratios (Table 3). Since we collected firm financial data over 5 years, each ratio was computed with data from each of these 5 years. We finally selected, from among the 150 variables we computed (30 ratios × 5 years), and with each sample, the variables that we were going to use to create failure patterns and models. For this purpose, we chose those with the best discriminatory ability and the lowest possible correlations with others [31]. We thus calculated a Mann-Whitney test for differences between failed and non-failed firms and selected, for each dimension, two variables with the best discrimination power among those for which the difference was significant at the threshold of 0.01%. We chose the same number of variables from each dimension to be sure that each financial dimension would be taken into account. To select this number, we designed models using four sets of variables with each learning sample: one with one variable per dimension, one with two, then one with three and the last one with four variables. With each set, we removed the variables that had correlations with others that were larger than 0.6. Then we estimated the forecasting ability of each set. On average, the best classification rates were achieved by far with the sets that were made up of two variables per financial dimension. We thus used this setting to design our models. Table 4 presents the p-values of a Mann-Whitney test calculated to assess the differences between failed and non-failed firms for each variable used to design failure patterns and models with data from 2006.

5.2. Model design

5.2.1. Single models

We designed as many single models as we collected samples and chose modeling methods, and with each method, we used the same variables. With discriminant analysis, logistic regression and Cox's model, we did not use any specific setup. With Cart, we did not perform any pruning because we wished for all models to be used with the same variables to avoid any informational distortion and we used the Gini index as a measure of heterogeneity to split each node. By contrast, we had to choose a few parameters with the other techniques. With the feed-forward neural network and the extreme learning machine, we used an architecture that was made up of one hidden layer, the size of which was experimentally determined before models were estimated [42], and we also used the hyperbolic tangent as an activation function of neurons. As such, we tested several sizes of the hidden layer (between 5 and 30 neurons) using different learning rates (between 0.01 and 0.5, with a step of 0.01). Model parameters were assessed using 50% of each learning sample and model accuracy was estimated with the remaining 50%. We computed the error achieved with each model and we chose the architecture that led to the lowest error. With the support vector machine, we used a radial basis function and we also conducted a set of experiments to determine the values of the function parameters: C and gamma. Here as well, each learning sample was divided into two equal parts, one to estimate parameters, one to assess model accuracy. We tested the values that ranged from 10^{-5}, 10^{-4} . . .to 10^{5}, and we chose the model that led to the lowest error.

5.2.2. Ensemble-based models

We also built as many ensemble-based models as we collected samples and chose modeling methods. Hence, with discriminant analysis and a given learning sample, we designed an ensemble with bagging, one with boosting, another with random subspace and a final one with rotation forest. With bagging and boosting, we used all variables to design models. With random subspace, variables that belonged to each model were selected by the algorithm itself. With rotation forest, variables were also chosen by the algorithm, knowing that each model was made up of 5 variables. To choose the size of each ensemble, we created a set of ensembles made up of between 2 and 200 models and then we tested their accuracy. As such, each learning sample was divided into two equal parts, one to estimate model parameters, and one to test their performance. Once the error of all sets was assessed, we chose the ensemble that led to the lowest error. With hybrid techniques, each single model was estimated as mentioned previously, but the combinations of their results were performed as presented in Section 4.1.

5.2.3. Failure pattern-based models

Pattern-based models were designed with ensembles of Kohonen maps, the sizes of which were different from each other, but the dimension of which was defined *a priori*: each map is a one-dimensional map made up of a single row of neurons. We first designed, with each learning sample, a set of maps with sizes ranging from 2 to 200 neurons. Then, each set was pruned, by removing maps one at a time. The estimation process was performed with 50% of learning data and the pruning process was performed with the remaining 50%.

Over the learning process, we ensured that the quantification of each map had been performed properly. Indeed, the quality of the quantification depends on the initial parameters which can sometimes lead to maps that are distorted due to twists, the butterfly effect… [12], even if the dimension(s) of the maps match those of the data. We therefore estimated 20 maps of the same size each time and chose the one leading to the lowest quantification error estimated with learning data. Once a map was estimated, neurons were

Table 4
P-values of a Mann-Whitney U test for differences between failed and non-failed firms calculated for each variable used to design failure patterns and models with data from 2006.

Year	Variables								
1	C/CA	CL/TS	TS/TA	EBIT/VA	FE/TA	FE/VA	TD/TA	EBITDA/TA	EBITDA/PE
	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
2	QA/CL	R/TS	NI/TS	VA/TS	FE/VA	EBITDA/TA	EBITDA/PE		
	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000		
3	EBIT/VA	FE/CF	SF/PE	EBITDA/TA	EBITDA/PE				
	0.0000	0.0000	0.0000	0.0000	0.0000				
4	AP/TS	FE/VA	FE/NI	EBITDA/TA	EBITDA/PE				
	0.0000	0.0000	0.0000	0.0000	0.0000				
5	C/CA	NI/TS	FE/VA	FE/EBITDA	CL/TA	VA/FA	EBITDA/TA		
	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000		

labeled according to the class (failed vs. non-failed) for which they can be considered a prototype. For this purpose, we calculated the number of firms of each group that were the closest to each neuron ($c1$ and $c2$) and we assigned to each of them the label of the majority class (-1 if $c_1 > c_2$, 1 if $c_2 < c_1$, 0 if $c_1 = c_2$). We also assigned to each neuron a statistic d that represented the degree of similarity between the neuron and the class it was assigned to. This statistic was calculated by dividing the number of firms that were the closest to a given neuron, and that belonged to the majority class, with the total number of firms that were the closest to the same neuron ($d = c_1/(c_1 + c_2)$ if the label $= -1$, $d = c_2/(c_1 + c_2)$ if the label $= 1$). These two parameters (the label and d) were used to compute forecasts made with each map.

Over the pruning process, several maps were removed from the initial set of maps to select the best performing subset. As such, we conducted two pruning sequences. First, maps were removed one at a time, starting with a map with the largest number of neurons, until no map remained available. At the end of this procedure, we chose the ensemble with the lowest error. Second, with the previously designed ensemble, maps were again removed one at a time, but this time in the other direction, starting with the map with the smallest number of neurons. And we selected the ensemble that led to the lowest error. The estimation process can be presented as follows.

Design sequence:

- Step 1: select a sub sample $L = (L_1, L_2, \ldots, L_p)$ made up of 50% of a learning sample;
- Step 2: set up the number n of neurons of the map to be designed using $n = 2$;
- Step 3: repeat k times, with $k = 1, 2 \ldots 20$;
- Step 3.1: estimate a one-dimensional map $M_{n,k}(l)$ with n neurons using sample L;
- Step 3.2: estimate the quantification error of the map $M_{n,k}(l)$;
- Step 4: choose the map $M_n(l)$ with the lowest quantification error;
- Step 5: compute the distance between each neuron and each firm l;
- Step 6: calculate the percentage of firms in each class that are the closest to each neuron: c_1 and c_2. The neurons λ_{l*} that are the closest to each firm satisfy the following condition:

$$\| x - \lambda_{l*} \| = \min_{l} \| x - \lambda_l \|$$

- Step 7: label each neuron with the label of the class that has the highest percentage: -1 if $c_1 > c_2$, 1 if $c_2 < c_1$, 0 if $c_1 = c_2$;
- Step 7: repeat Step 3 to Step 7 with $n = n + 1$ until $n = 200$.

Pruning sequence:

- Step 1: select the remaining 50% of a learning sample;
- Step 2: set up the maximum size of the maps to be removed to $s = 200$;
- Step 3: remove all maps with a size that is $>s$;
- Step 4: with each map, calculate the predicted class of each firm;
- Step 4.1: for firm l, compute its weighted distance D to each neuron n_w and find the closest neuron n_c;

$$D = d * ||l - n_c|| = d * min||l - n_w||$$

  with $d = c_1/(c_1 + c_2)$ if the label $= -1$, $d = c_2/(c_1 + c_2)$ if the label $= 1$;
- Step 4.2: assign to firm l the label of the class of neuron n_c: this label becomes the predicted class made by the map, the size of which is n;
- Step 5: calculate the final predicted class of all firms by combining the forecasts of all maps $M_p(l)$ of the set under consideration using a majority vote;

$$Prev(l) = \arg \max_{g \in -1;1} \sum_{p=t}^{s} \delta_{sign(M_p(l)),g}$$

- Step 6: estimate the correct classification rates achieved with all maps made up of between 2 and s neurons by comparing the predicted class of each observation to its current class;
- Step 7: repeat Step 3 to Step 6 with $n = n - 1$ until $n = 2$;
- Step 8: choose the subset of maps leading to the lowest error with 2 the size of the smallest map and s the size of the largest one;
- Step 9: set up the minimum size of the maps to be removed to $t = 2$;
- Step 10: remove all maps with a size that is $\leq t$;
- Step 11: repeat Step 4 to Step 5;
- Step 12: estimate the correct classification rates achieved with all maps made up of between t and s neurons by comparing the predicted class of each observation to its current class;
- Step 13: repeat Step 10 to Step 12 with $t = t + 1$ until $t = s$;
- Step 14: choose the subset of maps leading to the lowest error with t the size of the smallest map ($2 \leq t \leq s$), and s the size of the largest one ($t \leq s \leq 200$).

5.3. Estimation of model performance

The estimation of single model accuracy was performed by comparing the label of the true class of each firm to that which was forecasted by a model. With ensemble-based models, the same comparison was performed and we used a majority vote, with bagging, random subspace and rotation forest, or a weighted majority vote, with boosting, to perform the final forecast. With hybrid models, the predictions were estimated as mentioned in Section 4.1. With ensembles of maps, the predictions were done as follows. First, with each map, we calculated the weighted distance between each firm of a sample and each neuron of a map using the d statistics estimated previously. Then, we looked for the neuron that showed the lowest weighted distance to each firm and we assigned to the firm the label of this neuron: this label was considered the label of the predicted class made by a given map. Then, when the forecasts of all maps of an ensemble were calculated, we used a majority vote to combine them and calculate the final prediction. We used different measures to estimate model performance. First, we calculated the overall error achieved by models in order to assess their general discrimination ability. Then, since this error can be divided into type-I and type-II errors, the respective costs of which are far different from each other, we calculated the area under the receiver operating characteristic curve (AUC) that makes it possible to estimate a measure of performance which is totally independent from these costs. Finally, to get a precise overview about the main factors that may explain an error, we decomposed type-I and type-II errors using the method designed by Kohavi and Wolpert [36], who suggested decomposing the error of a model into bias, variance and noise terms. Since, in practice, the bias and the noise terms cannot be estimated separately, we solely computed the bias and the variance of each error[5].

[5] The bias represents the part of the error that is due to the distance between the performance of a given classifier and that of the best one within the same class of classifiers. The variance represents the part that is due to the different samples that were used. And the noise is the statistical uncertainty.

Table 5
Model accuracy (%) achieved with traditional models.

Periods	Single rule-based models							Bagging-based models						
	DA	LR	DT	Cox	SVM	FNN	ELM	DA	LR	DT	Cox	SVM	FNN	ELM
2007	79.74	80.28	78.49	81.66	81.57	81.67	80.70	79.28	80.64	81.06	82.99	81.92	82.18	82.75
2008	78.17	80.11	75.61	81.77	81.80	83.75	81.25	77.29	80.57	80.70	80.14	81.73	81.23	80.81
2009	76.30	77.53	76.14	79.46	76.25	77.46	79.36	75.45	78.50	78.97	78.35	79.43	78.69	79.81
2010	77.59	78.62	79.52	78.60	79.98	80.60	80.57	77.49	76.10	78.17	80.11	80.31	81.04	82.38
2011	77.33	78.36	79.56	78.88	81.14	79.33	81.48	80.54	83.47	84.24	83.38	82.43	81.92	81.11
2012	75.71	77.24	76.89	77.62	79.44	80.61	80.23	77.43	80.96	80.61	79.70	80.58	82.14	80.67
2013	74.09	76.90	76.32	74.11	76.64	75.76	77.16	77.30	77.32	75.02	76.37	80.15	77.71	79.07
2014	75.54	77.13	75.06	75.89	78.11	78.75	79.59	81.72	80.18	82.54	81.96	79.67	83.03	82.43
Mean	76.81	78.27	77.20	78.50	79.37	79.74	80.04	78.31	79.72	80.16	80.38	80.78	80.99	81.13

Periods	Boosting-based models							Random subspace-based models						
	DA	LR	DT	Cox	SVM	FNN	ELM	DA	LR	DT	Cox	SVM	FNN	ELM
2007	80.09	80.07	79.37	80.10	81.50	82.65	83.42	80.16	81.31	80.34	81.83	79.55	81.65	84.42
2008	80.17	82.49	80.02	83.49	82.62	82.17	80.67	79.94	80.58	80.74	78.10	81.66	80.49	83.96
2009	77.31	78.84	80.65	80.15	78.18	81.85	79.06	78.22	77.35	79.69	80.42	80.04	79.84	80.07
2010	76.83	82.20	81.43	82.44	80.55	78.46	82.74	78.13	78.56	78.29	79.86	81.27	78.30	81.41
2011	80.21	81.55	82.50	81.28	82.75	81.58	83.53	79.69	81.22	82.49	81.70	83.42	82.57	82.73
2012	80.19	78.82	79.24	78.25	78.06	80.78	82.85	79.58	82.62	81.58	80.07	82.07	83.93	81.09
2013	80.36	82.57	82.56	81.36	80.60	80.41	80.07	79.16	82.40	80.56	79.41	80.69	84.08	80.17
2014	82.04	79.09	81.64	82.67	81.09	83.74	82.16	79.31	81.40	80.82	80.77	78.42	80.22	78.20
Mean	79.65	80.70	80.93	81.22	80.67	81.45	81.81	79.27	80.68	80.57	80.27	80.89	81.38	81.51

Periods	Rotation forest-based models							Hybrid ensemble-based models				
	DA	LR	DT	Cox	SVM	FNN	ELM	CO1	CO2	SE1	SE2	SG
2007	83.31	82.67	83.36	84.67	83.93	84.37	84.76	81.55	80.65	79.21	82.27	82.92
2008	83.27	84.12	83.05	85.24	82.26	84.64	84.63	80.05	80.29	80.94	81.33	83.43
2009	76.43	78.55	77.68	79.35	77.69	79.08	79.35	78.56	79.59	78.38	79.11	80.51
2010	78.95	79.56	79.29	80.86	81.31	81.35	80.29	77.79	78.30	79.20	80.44	81.58
2011	82.67	80.35	80.30	81.61	81.80	82.01	81.08	80.97	80.81	80.58	81.51	82.53
2012	79.81	81.65	78.17	80.11	81.49	82.48	80.68	81.28	81.58	80.38	79.80	81.39
2013	79.94	78.26	78.12	79.38	77.77	79.60	79.24	80.69	79.22	78.85	80.78	81.07
2014	80.62	81.39	80.64	81.50	79.08	78.71	80.51	80.41	80.02	80.51	82.05	81.30
Mean	80.63	80.82	80.07	81.59	80.67	81.53	81.32	80.16	80.06	79.76	80.91	81.84

DA: discriminant analysis; DT: decision tree; Cox: Cox's model; ELM: extreme learning machine; FNN: feed-forward neural network; LR: logistic regression; SVM: support vector machine. CO: combined models (CO1 represents models designed with Lin and McClean [46]'s method and CO2 models designed with Karthik-Chandra et al. [30]'s method). SE: sequential models (SE1 represents models estimated with Cho et al. [7]'s method and SE2 models estimated with Tsai and Hsu [64]'s method). SG: segmentation-based models (represents models assessed with du Jardin [14]'s method).

6. Results and discussion

6.1. Results by period

We first calculated the correct classification rates achieved by period with each model estimated using a traditional method: single rule-, bagging-, boosting-, random subspace-, rotation forest-based models and hybrid models. The classification rates were estimated using a cut-off value that minimized the overall classification error and were averaged over the 100 samples used for their estimation. The results are presented in Table 5. We complemented Table 5 with Fig. 1 which shows the distribution of the correct classification rates estimated over 100 samples[6].

Three conclusions can be drawn from these results. First, one clearly notices that all models are sensitive to changes that occurred within the firm macroeconomic environment because the error increases over periods of trouble and reaches a maximum over the two critical years of the period studied (2009 and 2013), which correspond to recession phases. Obviously, none of the models appear to be better at achieving stable forecasts over time than others. This situation was often noticed in the literature [53]. Second, among all models, the extreme learning machine leads, on the whole, to the best results, whether it was used to design single models (mean = 80.04%, with the lowest variance among all results achieved with single models) or ensemble-based models (their means range from 81.13% to 81.81%). The extreme learning machine is followed by the feed-forward neural network and the support vector machine, whereas discriminant analysis leads to the worst results. This indicates the advantage of the extreme learning machine over traditional optimization techniques used to estimate neural network parameters. Third, one may notice that ensemble-based models are more accurate than single models but none of them seem to be radically better than others, even when one analyzes the distribution of classification rates presented in Fig. 1.

Then, we compared all results that have just been presented to those estimated with failure pattern-based models. They are presented in Table 6. We added to these results those computed with the best single models, that is to say those achieved with the extreme learning machine. Table 6 shows that failure pattern-based

[6] We do not present all results that were achieved with hybrid models, but solely the best ones when different combinations of models or different methods were used. As such, with the method designed by Karthik-Chandra et al. [30], we present the results estimated with the combination of three techniques: random forest, a support vector machine and a feed-forward neural network. With that by Tsai and Hsu [64], we present the results estimated using a combination of logistic regression, a decision tree and a feed-forward neural network. And with that by du Jardin [14], we show the results computed with a feed-forward neural network.

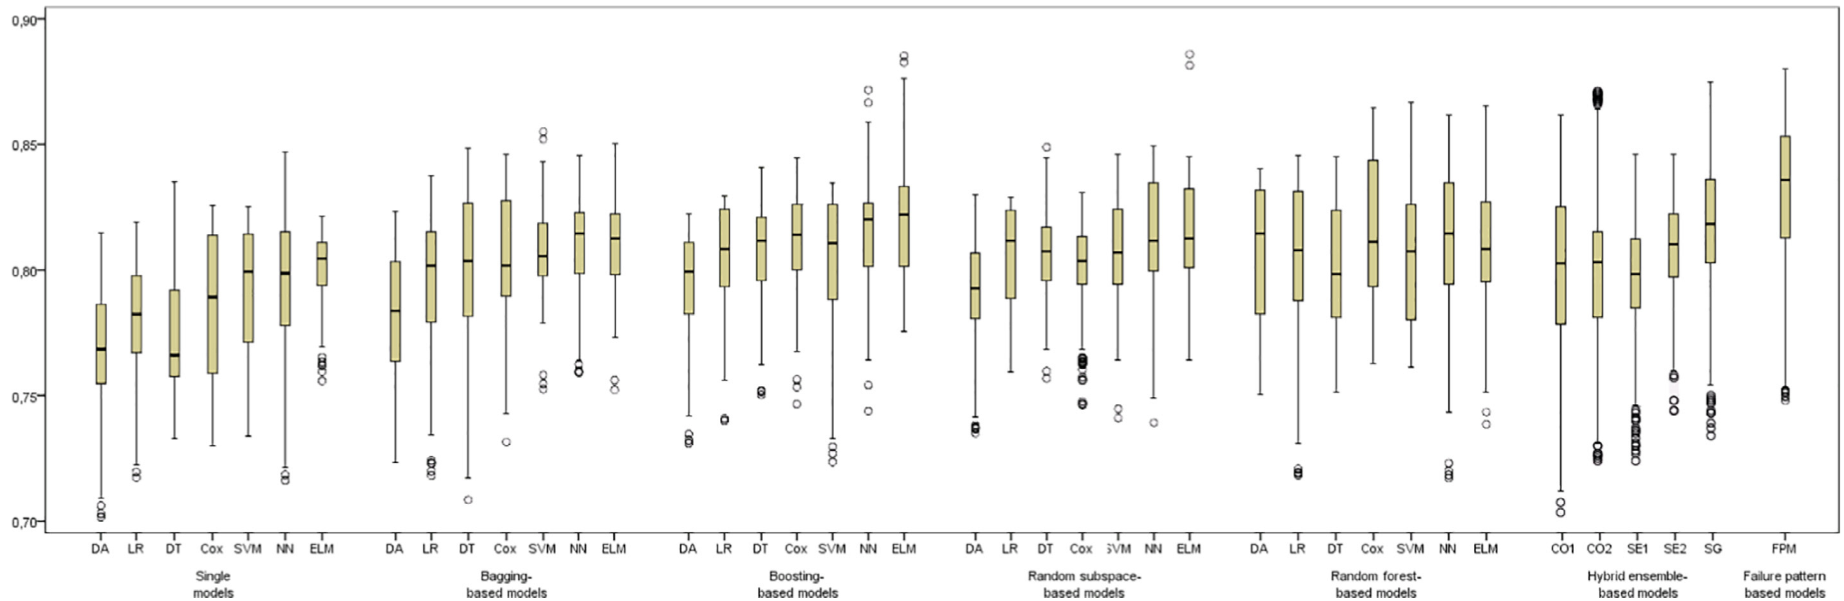

Fig. 1. Distribution of correct classification rates by model.

Table 6
Model accuracy (%) by type of model and by period.

Periods	SM	BA	BO	RS	RF	HM	bSM*	FM
2007	80.59	81.55	81.03	81.32	83.87	81.32	80.70	82.92
2008	80.35	80.35	81.66	80.78	83.89	81.21	81.25	83.74
2009	77.50	78.46	79.43	79.37	78.30	79.23	79.36	79.94
2010	79.35	79.37	80.66	79.40	80.23	79.46	80.57	84.56
2011	79.44	82.44	81.91	81.98	81.40	81.28	81.48	83.54
2012	78.25	80.30	79.74	81.56	80.62	80.88	80.23	84.91
2013	75.85	77.56	81.13	80.93	78.90	80.12	77.16	81.34
2014	77.15	81.65	81.78	79.88	80.35	80.86	79.59	84.78
Mean	78.56	80.21	80.92	80.65	80.95	80.55	80.04	83.22

BA: bagging-based models; BO: boosting-based models; bSM: best SM (corresponds to ELM); FM: failure pattern-based models; HM: hybrid ensemble-based models; RF: rotation forest-based models; RS: random subspace-based models; SM: single models.

models are also sensitive to macroeconomic fluctuations, like other models are, but in a somewhat different way. When there is a sharp decline within the economic situation between the period when a model is estimated and when it is used, its accuracy decreases more significantly than for other models. By contrast, when the economic situation improves between the estimation period of a model and its period of use, failure pattern-based models are more likely to better forecast firm fate since the increase of their accuracy is larger than that of other models. This is the case between 2009 and 2010, but also between 2013 and 2014.

To deepen the analysis, we computed the differences between the results calculated with failure pattern-based models and those estimated with all other models, by period. They are presented in Table 7, Panel A. In this table, Panel A is complemented with Panel B which presents the p-values of a test for differences between proportions to assess the statistical significance of the discrepancies mentioned in Panel A. The results indicate that, whatever the period, failure pattern-based models are more accurate than all single models but also than most (82%) of ensemble-based models. This underlines that the hypotheses we made in the introduction about failure pattern-based models and their informational added value seem to be verified. Because failure pattern-based models propose a subtle way of modeling firm financial situations, they are more likely than others to properly typify the symptoms that lead to failure.

We also calculated the average correct classification rate by type of model, then the difference between each pair of rates and finally the significance level of these differences. Results are shown in Table 8. In this table, Panel A indicates the average rates, Panel B the differences between these rates, and Panel C the p-values of a test for differences between proportions that make it possible to estimate the statistical significance of the discrepancies mentioned in Panel B.

These figures show that correct prediction rates achieved with failure pattern-based models are significantly higher than those calculated with any other model (at the threshold of 0.001), with differences that range from 2.27 percentage points to 3.01. They also show that ensembles of Kohonen maps are more accurate than traditional models but also than the best single model, with an average significant difference of 3.17 percentage points, and of course more accurate than all single models, with a difference of 4.66 percentage points which is also significant.

These findings shed light on a characteristic of failure pattern-based models that is ultimately not commonly shared by failure models. The literature shows that ensemble-based models are often more accurate than a single model when it is designed with the base classifier of the ensemble, but much less frequently when the single model is designed with another classifier. However, here, Kohonen maps manage to be more efficient than ensemble models as well as very good single models such as those designed with an extreme learning machine. In addition, these findings show how large the discrepancy is between the accuracy of failure pattern-based models and that of models that are used by banks and financial companies.

6.2. Model predictive power

We have also estimated model performance using the area under the ROC curve (AUC) which is a measure that is completely independent from any misclassification costs. Table 9 presents the AUC calculated by type of model and by period but it also shows the p-values of a test for differences between AUCs, making it possible to assess the statistical significance of the differences between the AUCs estimated with failure pattern-based models and those computed with traditional models.

This table shows that the differences are significant in more than 85% of the cases, at the threshold of 0.05. Of the remaining 15%, more than 3/4 are significant at the threshold of 0.1, since p-values range from 0.06 to 0.088, whereas 1/4 are not at all significant: this is the case of rotation forest with data from 2007 and of bagging with data from 2011. This finding is rather interesting because financial institutions may experience very different misclassification costs and, whatever these costs, our models provide a modeling framework that leads to less costly errors than those that can be estimated with any kind of model, and especially with single models. Indeed, financial companies mainly use discriminant analysis or logistic regression to build failure models, and with these two methods all differences we computed between their results and those achieved with failure pattern-based models are significant.

6.3. Results by firm status and decomposition of type-I and type-II error

Finally, we wanted to know whether the error due to failure pattern-based models was mainly made up of bias or of variance and how these two components compared with those achieved with traditional models. Indeed, failure pattern-based models require the estimation of a number of parameters far larger than that required by any other model used in this study, and one knows that the variance of an estimate is partly related to this number. In order to obtain a precise view of the influence of a given model on the components of the error, we decomposed type-I and type-II errors. For the sake

Table 7
Differences (% point) between correct classification rates achieved with failure process-based models and those achieved with other models.

Panel A: Differences

Periods	FM-SM	FM-BA	FM-BO	FM-RS	FM-RF	FM-HM	FM-bSM
2007	2.34	1.38	1.90	1.60	−0.94	1.60	2.22
2008	3.38	3.38	2.07	2.95	−0.15	2.53	2.49
2009	2.44	1.49	0.51	0.57	1.64	0.71	0.58
2010	5.20	5.19	3.90	5.16	4.33	5.10	3.99
2011	4.10	1.10	1.63	1.57	2.14	2.27	2.06
2012	6.66	4.61	5.17	3.34	4.28	4.02	4.68
2013	5.48	3.78	0.21	0.41	2.44	1.22	4.17
2014	7.62	3.13	3.00	4.90	4.43	3.92	5.19

Panel B: P-values of a test for differences

Periods	FM-SM	FM-BA	FM-BO	FM-RS	FM-RF	FM-HM	FM-bSM
2007	0.001	0.046	0.006	0.021	0.161	0.021	0.001
2008	0.000	0.000	0.002	0.000	0.820	0.000	0.000
2009	0.001	0.043	0.483	0.433	0.026	0.327	0.424
2010	0.000	0.000	0.000	0.000	0.000	0.000	0.000
2011	0.000	0.105	0.017	0.022	0.002	0.001	0.003
2012	0.000	0.000	0.000	0.000	0.000	0.088	0.000
2013	0.000	0.000	0.770	0.560	0.001	0.000	0.000
2014	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table 8
Average correct classification rates by type of model and differences between correct classification rates.

Panel A: Average correct classification rates (%)

SM	BA	BO	RS	RF	HM	bSM	FM
78.56	80.21	80.92	80.65	80.95	80.55	80.04	83.22

Panel B: Differences (percentage point)

	BA	BO	RS	RF	HM	bSM	FM
SM	1.65	2.36	2.09	2.38	1.98	1.48	4.66
BA		0.71	0.44	0.74	0.34	−0.17	3.01
BO			−0.27	0.03	−0.37	−0.88	2.30
RS				0.29	−0.11	−0.61	2.56
RF					−0.40	−0.90	2.27
HM						−0.50	2.67
bSM							3.17

Panel C: P-values of a test for differences

	BA	BO	RS	RF	HM	bSM	FM
SM	0.000	0.000	0.000	0.000	0.000	0.000	0.000
BA		0.005	0.080	0.004	0.185	0.514	0.000
BO			0.291	0.915	0.139	0.001	0.000
RS				0.245	0.671	0.992	0.000
RF					0.113	0.000	0.000
HM						0.048	0.000
bSM							0.000

of simplicity, we solely report the decomposition of these errors by type of model. The results are given in Table 10. It presents, for each error, its value, its bias and its variance as well as the ratio between the variance and the bias.

It emerges from Table 10 that failure pattern-based models barely reduce type-II error compared to single models but also compared to ensemble-based models. The variance/bias ratio is 24.14% with failure pattern-based models, 23.65% with single models, and that of all other models ranges only from 23.31% to 24.08%. By contrast, the variance/bias ratio of type-I error achieved with failure pattern-based models is lower than that calculated with all other types of model: the ratio is 41.16% with failure pattern-based models whereas it ranges with all others from 34.89% to 37.39%. This demonstrates the contribution of the modeling technique we propose in this article: it seems to better decode the information that derives from financial accounts and that embodies different symptoms of failure than other techniques do.

All measures we used to characterize the performance of failure models lead to the same conclusion. Our modeling technique leads to more accurate predictions than those performed with current advanced methods such as extreme learning machines or support vector machines, as well as more accurate predictions than those of ensemble methods. And the gain is partly due to its ability to better forecast the fate of firms that are likely to go bankrupt, and thus the fate of those that may lead to the highest misclassification costs a bank may face.

6.4. Analysis of failure patterns

To study failure patterns, we clustered all maps that were estimated with a given sample using a hierarchical ascending classification and a Ward criterion. Each time, we looked for the most homogenous partition, using the three best indices mentioned in the research done by [51], among a set of initial partitions that ranged in size between 2 and 15 clusters. We intentionally studied small partitions because it would have been extremely difficult to analyze the characteristics of large sets of clusters. Clusters were mainly analyzed based on the way they quantified the different financial dimensions that explain failure (profitability, solvency…). On the whole, in each sample, we can find three general survival patterns and four general failure patterns.

Among survival patterns, the first one characterizes companies that are very solvent and profitable, with a good financial structure, somewhat good solvency, a low debt level and a relatively low level of supplier and trade credits. The second one consists of firms that are also very profitable, but are less solvent and liquid than the former, with an average financial structure, a rather low level of long-term debt and an average level of short-term debt. The third pattern represents companies with average financial health, rather

Table 9
AUC (Area under the ROC curve) by type of model and by period.

Panel A: AUC

Periods	SM	BA	BO	RS	RF	HM	bSM	FM
2007	0.781	0.764	0.776	0.793	0.811	0.782	0.788	0.829
2008	0.746	0.755	0.767	0.761	0.766	0.769	0.762	0.808
2009	0.740	0.728	0.724	0.736	0.742	0.743	0.734	0.792
2010	0.778	0.755	0.768	0.765	0.755	0.738	0.776	0.837
2011	0.753	0.798	0.754	0.761	0.751	0.744	0.758	0.805
2012	0.749	0.763	0.761	0.768	0.764	0.761	0.757	0.812
2013	0.712	0.708	0.744	0.740	0.738	0.740	0.684	0.789
2014	0.739	0.751	0.745	0.735	0.752	0.745	0.740	0.801

Panel B: P-values of a test for differences

Periods	FM-SM	FM-BA	FM-BO	FM-RS	FM-RF	FM-HM	FM-bSM
2007	0.024	0.003	0.013	0.088	0.381	0.028	0.049
2008	0.008	0.022	0.071	0.041	0.065	0.080	0.045
2009	0.034	0.011	0.007	0.023	0.040	0.045	0.019
2010	0.005	0.000	0.001	0.001	0.000	0.000	0.004
2011	0.027	0.739	0.029	0.060	0.023	0.011	0.042
2012	0.007	0.032	0.027	0.050	0.036	0.027	0.018
2013	0.002	0.001	0.063	0.045	0.038	0.043	0.000
2014	0.009	0.034	0.018	0.006	0.036	0.018	0.011

Table 10
Decomposition of type-I and type-II error (%) by type of model.

	Type-II				Type-I			
	Error	Bias	Variance	Variance/Bias	Error	Bias	Variance	Variance/Bias
SM	21.33	17.25	4.08	23.65	26.71	19.80	6.91	34.89
BA	19.68	15.96	3.72	23.31	25.48	18.75	6.74	35.95
BO	18.94	15.31	3.63	23.67	26.12	19.05	7.07	37.14
RS	19.21	15.48	3.73	24.08	26.38	19.20	7.18	37.39
RF	18.93	15.28	3.65	23.86	25.29	18.64	6.65	35.68
HM	19.33	15.62	3.71	23.72	25.87	19.03	6.84	35.94
FM	16.66	13.42	3.24	24.14	23.08	16.35	6.73	41.16

good liquidity, but low profitability, and a relatively fragile financial situation because of a relatively high level of debt compared to their assets.

Among failure patterns, the first one represents firms that are rather profitable, have average solvency, a poor financial structure, a low level of liquidity and a high level of long-term debt. The second one consists of companies that are roughly in the same financial situation as those of the first pattern, but that have a very low level of liquidity, with a high amount of supplier and trade credits, and a lack of shareholder funds. The third one corresponds to firms with a much-deteriorated financial situation, with an average level of liquidity, but with low profitability and solvency levels, and a very fragile financial structure. The fourth and last one corresponds to very unhealthy firms, with a huge amount of short-, mid- and long-term debt, very weak solvency and profitability, and an extremely low level of liquidity. Of course, all these general patterns could be divided into a few sub-patterns, but their analysis is beyond the scope of this article. We can, however, note that failure sub-patterns are far more diverse than those of survival are. This is certainly the reason why models that are not able to embody such diversity lead to forecasts that are, on the whole, weaker than those achieved using the method we present in this article.

6.5. Contributions to the literature and limits

The contribution of this study to the literature is threefold. From a conceptual stand point, it shows that the way a model is likely to embody both the variety and the historical nature of firm financial situations can significantly improve its performance. Some rare studies have already observed that firms that are going bankrupt do not follow the same path before failing, and that the ability of a model to account for such differences could affect its accuracy. The concept of “failure patterns” on which our models rely is rather similar to the concept of “bankruptcy trajectories”, which can be found in the literature. Indeed, the former makes it possible to model the existence of groups of firms where each of them can be characterized with a particular temporal and financial profile. The latter makes it also possible to represent different groups, but this time each of them share the same dynamics that govern the way their financial situation may change over time. The research works that have been carried out to study bankruptcy trajectories but also our current study reasserts that one of the key factors that can improve predictions is to be sought in the way models are able to account for both firm history and symptoms of failure.

From a methodological stand point, this study provides a means of quantifying the diversity of firms’ financial situations and the way these situations can be assessed over time, which breaks with usual practices, and which illustrates how this variety can affect model performance. Actually, failure patterns make it possible to capture different variations of firms’ financial situations which cannot be embodied by traditional ensemble-based models. Ensemble-based models are able to account for some variations using an estimation of the magnitude of changes that may affect a set of variables or using different sets of variables. But in a way, these two calculations make it possible to solely represent a unique level of representation of these variations. By contrast, failure patterns allow a model to embody different levels: some patterns are very general and correspond to common financial situations that are shared by a lot of firms, some others are rarer and are solely shared by a few subgroups of firms. This is where the difference lies between these two types of modeling.

Finally, from an empirical point of view, it shows that the limits of multi-period models, whatever the modeling techniques, can be partially overcome using an appropriate modeling of bankruptcy symptoms.

However, our study presents some limits. The first one is conceptual. Failure patterns are not able to account for all types of failure symptoms since they solely represent those that can be detected using firm financial accounts. This suggests that they might be well complemented with other factors that are related either to the firm strategy, organization or management. The second one is methodological. The error achieved with our models is rather unstable over time, as is that of other models, despite the fact that models have been estimated every year. This raises the question of their sensitivity to economic conditions and requires a thorough study to analyze the part of the error attributable to macroeconomic changes and find a means to overcome this influence. The third and last one is practical. Failure pattern-based models, such as any ensemble-based models, are not understandable by any financial analyst. This situation can be quite problematic [50], especially if the user of a model must be able to justify a decision that relies on one of its forecasts.

7. Conclusion

In this study, we propose a new bankruptcy prediction model that relies on the estimation of failure patterns that are quantified with ensembles of Kohonen maps. The results show that this type of model is, on the whole, more efficient than single or ensemble-based models and this efficiency seems to be rather robust since it was assessed using different samples collected over different time periods. Moreover, failure pattern-based models embody a predictive function that appears to be more complex than that embodied by any type of model since the bias component of their error is markedly lower than that of all other models. One may then think that their performance is likely due to this complexity.

Our findings have two main implications. First, they reinforce the idea that emerged in the literature a few years ago and which suggests that model accuracy does not solely rely on data mining techniques but also on the way one will use some knowledge about the bankruptcy phenomenon during a modeling process; incorporating domain knowledge into classification models may indeed improve model performance [59]. Here, domain knowledge relies on the concept of "failure patterns", which has been highlighted in the literature that deals with financial or organizational issues, but also on the limits of the techniques that have been used so far to assess these patterns. Thus, a modeling framework can really benefit from conceptual developments that make it possible to enrich the knowledge that one can have about a firm failure. Our findings invite the scientific community that is interested in this issue to renew the conceptual basis used to design models. Second, they remind financial institutions of the true added value of ensemble-based models by once again showing them the limits of their own models and by providing them with a reliable modeling framework that may fit their needs.

Acknowledgment

We are very grateful to the two anonymous reviewers for their substantial contribution to the improvement of this article.

References

[1] P. Alam, D. Booth, K. Lee, T. Thordarson, The use of fuzzy clustering algorithm and self-organizing neural networks for identifying potentially failing banks: an experimental study, Expert Syst. Appl. 18 (2000) 185–199.
[2] E. Alfaro, M. Gamez, N. Garcia, A boosting approach for corporate failure prediction, Appl. Intell. 27 (2007) 29–37.
[3] E. Alfaro, N. Garcia, M. Gamez, D. Elizondo, Bankruptcy forecasting: an empirical comparison of AdaBoost and neural networks, Decis. Support. Syst. 45 (2008) 110–122.
[4] E.I. Altman, Financial ratios, discriminant analysis and the prediction of corporate bankruptcy, J. Financ. 23 (1968) 589–609.
[5] L. Breiman, Bagging predictors, Mach. Learn. 24 (1996) 123–140.
[6] L. Breiman, J.H. Friedman, R.A. Olshen, C.J. Stone, Classification and Regression Trees, Chapman and Hall-CRC, Boca Raton, Florida, 1984.
[7] S. Cho, J. Kim, J.K. Bae, An integrative model with subject weight based on neural network learning for bankruptcy prediction, Expert Syst. Appl. 36 (2009) 403–410.
[8] S.B. Cho, Ensemble of structure-adaptive self-organizing maps for high performance classification, Inf. Sci. 123 (2000) 103–114.
[9] C.L. Chuang, Application of hybrid case-based reasoning for enhanced performance in bankruptcy prediction, Inf. Sci. 236 (2013) 174–185.
[10] D.R. Cox, Regression models and lifetables, J. R. Stat. Soc. 34 (1972) 187–220.
[11] R.A. D'Aveni, The aftermath of organizational decline: a longitudinal study of the strategic and managerial characteristics of declining firms, Acad. Manag. J. 32 (1989) 577–605.
[12] E. de Bodt, M. Cottrell, M. Verleysen, Statistical tools to assess the reliability of self-organizing maps, Neural Netw. 15 (2002) 967–978.
[13] B.M. del Brio, C. Serrano-Cinca, Self-organizing neural networks for the analysis and representation of data: some financial cases, Neural Comput. & Applic. 1 (1993) 193–206.
[14] P. du Jardin, Bankruptcy prediction using terminal failure processes, Eur. J. Oper. Res. 242 (2015) 286–303.
[15] P. du Jardin, A two-stage classification technique for bankruptcy prediction, Eur. J. Oper. Res. 254 (2016) 236–252.
[16] P. du Jardin, Dynamics of firm financial evolution and bankruptcy prediction, Expert Syst. Appl. 75 (2017) 25–43.
[17] P. du Jardin, E. Severin, Predicting corporate bankruptcy using a self-organizing map: an empirical study to improve the forecasting horizon of a financial failure model, Decis. Support. Syst. 51 (2011) 701–711.
[18] A. Fan, M. Palaniswami, Selecting bankruptcy predictors using a support vector machine approach, IEEE International Joint Conference on Neural Networks, vol. 6, 2000. pp. 354–359.
[19] E. Fedorova, E. Gilenko, S. Dovzhenko, Bankruptcy prediction for Russian companies: application of combined classifiers, Expert Syst. Appl. 40 (2013) 7285–7293.
[20] Y. Freund, Boosting a weak learning algorithm by majority, COLT'90: The third annual workshop on computational learning theory, 1990. pp. 202–216.
[21] R. Geng, I. Bose, X. Chen, Prediction of financial distress: an empirical study of listed Chinese companies using data mining, Eur. J. Oper. Res. 241 (2015) 236–247.
[22] A. Gepp, K. Kumar, The role of survival analysis in financial distress prediction, Int. Res. J. Financ. Econ. 16 (2008) 13–34.
[23] J. Heo, J.Y. Yang, AdaBoost based bankruptcy forecasting of Korean construction companies, Appl. Soft Comput. 24 (2014) 494–499.
[24] T.K. Ho, The random subspace method for constructing decision forests, IEEE Trans. Pattern Anal. Mach. Intell. 20 (1998) 832–844.
[25] Z. Hua, Y. Wang, X. Xu, B. Zhang, L. Liang, Predicting corporate financial distress based on integration of support vector machine and logistic regression, Expert Syst. Appl. 33 (2007) 434–440.
[26] G.B. Huang, Q.Y. Zhu, C.K. Siew, Extreme Learning machine: theory and applications, Neurocomputing 70 (2006) 489–501.
[27] S.C. Huang, Y.C. Tang, C.W. Lee, M.J. Chang, Kernel local Fisher discriminant analysis based manifold-regularized SVM model for financial distress predictions, Expert Syst. Appl. 39 (2012) 3855–3861.
[28] C. Hung, J.H. Chen, A selective ensemble based on expected probabilities for bankruptcy prediction, Expert Syst. Appl. 36 (2009) 5297–5303.
[29] J. Huysmans, B. Baesens, J. Vanthienen, T. van Gestel, Failure prediction with self-organizing maps, Expert Systems with Applications 30 (2006) 479–487.
[30] D. Karthik-Chandra, V. Ravi, I. Bose, Failure prediction of dotcom companies using hybrid intelligent techniques, Expert Syst. Appl. 36 (2009) 4830–4837.
[31] M.Y. Kiang, A comparative assessment of classification methods, Decis. Support. Syst. 35 (2003) 441–454.
[32] M.J. Kim, D.K. Kang, Ensemble with neural networks for bankruptcy prediction, Expert Syst. Appl. 37 (2010) 3373–3379.
[33] M.J. Kim, D.K. Kang, Classifiers selection in ensembles using genetic algorithms for bankruptcy prediction, Expert Syst. Appl. 39 (2012) 9308–9314.
[34] M.J. Kim, D.K. Kang, H.B. Kim, Geometric mean based boosting algorithm with over-sampling to resolve data imbalance problem for bankruptcy prediction, Expert Syst. Appl. 42 (2015) 1074–1082.
[35] S.Y. Kim, A. Upneja, Predicting restaurant financial distress using decision tree and AdaBoosted decision tree models, Econ. Model. 36 (2014) 354–362.
[36] R. Kohavi, D. Wolpert, Bias plus variance decomposition for zero-one loss functions, Proceedings of the thirteenth international conference on machine learning, 1996. pp. 275–283.
[37] T. Kohonen, Self-organizing maps, Springer Series in Information Sciences, vol. 30, Springer. 2001,
[38] L.I. Kuncheva, C.J. Whitaker, Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy, Mach. Learn. 51 (2003) 181–207.
[39] T. Laitinen, Financial ratios and different failure processes, J. Bus. Financ. Account. 18 (1991) 649–673.
[40] K.C. Lee, I. Han, Y. Kwon, Hybrid neural network models for bankruptcy predictions, Decis. Support. Syst. 18 (1996) 63–72.
[41] T. Lensberg, A. Eilifsen, T.E. McKee, Bankruptcy theory development and classification via genetic programming, Eur. J. Oper. Res. 169 (2006) 677–697.
[42] P. Leray, P. Gallinari, Feature selection with neural networks, Behaviormetrika 26 (1998) 145–166.
[43] H. Li, Y.C. Lee, Y.C. Zhou, J. Sun, The random subspace binary logit (RSBL) model for bankruptcy prediction, Knowl.-Based Syst. 24 (2011) 1380–1388.
[44] H. Li, J. Sun, Case-based reasoning ensemble and business application: a computational approach from multiple case representations driven by randomness, Expert Syst. Appl. 39 (2012) 3298–3310.
[45] J.J. Liao, C.H. Shih, T.F. Chen, M.F. Hsu, An ensemble-based model for two-class imbalanced financial problem, Econ. Model. 37 (2014) 175–183.
[46] F.Y. Lin, S. McClean, A data mining approach to the prediction of corporate failure, Knowl.-Based Syst. 14 (2001) 189–195.
[47] F.J. Lopez-Iturriaga, I. Pastor-Sanz, Bankruptcy visualization and prediction using neural networks: a study of U.S. commercial banks, Expert Syst. Appl. 42 (2015) 2857–2869.
[48] O. Lukason, E.K. Laitinen, A. Suvas, Failure processes of young manufacturing micro firms in Europe, Manag. Decis. 54 (2016) 1966–1985.
[49] A.I. Marques, V. Garcia, J.S. Sanchez, Exploring the behaviour of base classifiers in credit scoring ensembles, Expert Syst. Appl. 39 (2012) 10244–10250.
[50] D. Martens, J. Vanthienen, W. Verbeke, B. B, Performance of classification models from a user perspective, Decis. Support. Syst. 51 (2011) 782–793.
[51] G.W. Milligan, A Monte-Carlo study of thirty internal criterion measures for cluster analysis, Psychometrika 46 (1981) 187–199.
[52] J.A. Ohlson, Financial ratios and the probabilistic prediction of bankruptcy, J. Account. Res. 18 (1980) 109–131.
[53] P.P.M. Pompe, J. Bilderbeek, Bankruptcy prediction: the influence of the year prior to failure selected for model building and the effects in a period of economic decline, Intell. Syst. Account. Finance Manag. 13 (2005) 95–112.
[54] J.J. Rodriguez, L.I. Kuncheva, Rotation forest: a new classifier ensemble method, IEEE Trans. Pattern Anal. Mach. Intell. 10 (2006) 1619–1630.
[55] R.E. Schapire, The strength of weak learnability, Mach. Learn. 5 (1990) 197–227.
[56] A. Scherbart, T.W. Nattkemper, Looking inside self-organizing map ensembles with resampling and negative correlation learning, Neural Netw. 24 (2011) 130–141.
[57] C. Serrano-Cinca, Self-organizing neural networks for financial diagnosis, Decis. Support. Syst. 17 (1996) 227–238.
[58] C. Serrano-Cinca, B. Gutierrez-Nieto, Partial least square discriminant analysis for bankruptcy prediction, Decis. Support. Syst. 54 (2013) 1245–1255.
[59] A.P. Sinha, H. Zhao, Incorporating domain knowledge into data mining classifiers: an application in indirect lending, Decis. Support. Syst. 46 (2008) 287–299.
[60] R.M. Stein, Benchmarking default prediction models: pitfalls and remedies in model validation, J. Risk Model Val. 1 (2007) 77–113.
[61] J. Sun, H. Fujita, P. Chen, H. Li, Dynamic financial distress prediction with concept drift based on time weighting combined with AdaBoost support vector machine ensemble, Knowl.-Based Syst. 120 (2017) 4–14.
[62] J. Sun, M.Y. Jia, H. Li, AdaBoost ensemble for financial distress prediction: an empirical comparison with data from Chinese listed companies, Expert Syst. Appl. 38 (2011) 9305–9312.
[63] C.F. Tsai, Combining cluster analysis with classifier ensembles to predict financial distress, Inform. Fusion 16 (2014) 46–58.

[64] C.F. Tsai, Y.F. Hsu, A meta-learning framework for bankruptcy prediction, J. Forecast. 32 (2013) 167–179.
[65] A. Volkov, D.F. Benoit, D.V. den Poel, Incorporating sequential information in bankruptcy prediction with predictors based on Markov for discrimination, Decis. Support. Syst. 98 (2017) 59–68.
[66] G. Wang, J. Ma, S. Yang, An improved boosting based on feature selection for corporate bankruptcy prediction, Expert Syst. Appl. 41 (2014) 2353–2361.
[67] L. Wang, C. Wu, Business failure prediction based on two-stage selective ensemble with manifold learning algorithm and kernel-based fuzzy self-organizing map, Knowl.-Based Syst. 121 (2017) 99–110.
[68] D. West, S. Dellana, J. Qian, Neural network ensemble strategies for financial decision application, Comput. Oper. Res. 32 (2005) 2543–2559.
[69] R.L. Wilson, R. Sharda, Bankruptcy prediction using neural networks, Decis. Support. Syst. 11 (1994) 545–557.
[70] Z. Xiao, X. Yang, Y. Pang, X. Dang, The prediction for listed companies' financial distress by using multiple prediction methods with rough set and Dempster-Shafer evidence theory, Knowl.-Based Syst. 26 (2012) 196–206.

Philippe du Jardin is currently a professor at EDHEC Business School, in charge of the information technology department, and conducts his research works within the EDHEC Financial Analysis and Accounting Research Centre. He holds three Master's degrees (business administration, research and information technology) as well as a Ph.D. in business administration from Nice Sophia-Antipolis University and a post-doctotal degree (HdR) from the Lille University. Prior to coming to EDHEC, he has held academic appointments at SKEMA Business School, Aix-Marseille University and Nice University and has served as a consultant for various companies for questions related to IT and data-mining issues. His research interests focus on credit risk and company financial failure and he is interested in neural networks and nonlinear models.