Building Models Based on Artificial Neural Networks to Predict Customers' Purchase Frequency

Galo Paiva
Departamento de Ingeniería Industrial y Sistemas
Universidad de La Frontera
Temuco, Chile
galo.paiva@uforntera.cl

Carlos Hernández
Departamento de Procesos Industriales
Universidad Católica de Temuco
Temuco, Chile
carlos.hernandez.zavala@uct.cl

Abstract

Customer loyalty depends on a combination of objective and subjective factors. For retailers, the benefit having a positive perception in the public is the credibility transmitted to their brands. This investigation is focused on the implementation of artificial neural networks (ANN) to predict customers' loyalty to supermarket chains measured in terms of their purchase frequency. The research is carried out following a classic 4-stage methodology (analysis, design, development, and validation). During the analysis, the results of a customer loyalty survey are preprocessed. During the design, the full questionnaire is divided into several domains (e.g. products' quality, buying experience, etc.). Each domain is then used to build and compare predictive models. The original dataset containing 1.050 surveys is split up in two sets: 80% of data for training and test, and the remaining 20% for validation. To predict customer loyalty twelve different ANN-based models are built. The results reveal that the proposed ANN-based models can predict correctly 60% and 72% of customers' purchase frequency. In conclusion, predictive ANN-based models help determine how often an average customer made purchases on a given supermarket chain with a reasonable degree of certainty by means of analyzing some of their habits and preferences.

Keywords
Purchase Frequency, Customer Loyalty, Private Label, Artificial Neural Network, Predictive Model.

1. Introduction
Supermarket chains permanently try to attract new customers and retain old ones by means of implementing different strategies. Some of them involve the location of their stores, the diversity in the offer of products, the ambience in the stores, marketing campaigns, and private labels. In the case of supermarkets, several authors suggest that customer loyalty depends of a combination of different factor such as: low prices, closeness, good access, product variety, discounts, and also for an emotional connection with a specific store based on the trust, environment and service (Allen and Rao, 2000; Flavián, et al., 2001; Martos-Partal and González-Benito, 2013). Other authors mention a relationship between customer loyalty and the retailer's social responsibility policies (Ailawadi et. al, 2014).
There is a correlation between customer satisfaction and brand loyalty, being the first a necessary condition to generate the latter. However, it is not enough by itself because brand loyalty depends on emotional factors (Neumeier, 2006; Allen and Rao, 2000).
Building a corporate brand implies creating a value proposition for customers based on meaningful relationships and on specific characteristics to support the group of brands of a company. The establishment of a brand can lead to customer satisfaction and satisfied customers are likely loyal to the company (Martenson, 2007).
Customer loyalty to supermarkets can be measured by the frequency of purchase (Dick and Basu, 1994), but it is not an indicator of real loyalty, because repeated purchases at a particular supermarket may be due to mere convenience if there no other closer alternative. Thus, it is a purchase based on emotional satisfaction. Even more, customers

might purchase repeatedly product of a given brand without having a positive attitude towards it. However, purchase behavior is still understood as the result of loyalty. (Allen and Rao, 2000; Flavián, et al., 2001).
The relation between private labels and loyalty to a given supermarket chain has been also studied. Private labels are brands created by retailers to attract customers and to improve their profit. These brands are owned, controlled and sold by retail chains. It is still matter of discussion if the loyalty to private label generates loyalty to retail chain or vice versa (Corstjens and Lal, 2000; Collins-Dodd and Lindley, 2003; Berkowitz et al., 2005; San Martin, 2006; Martínez and Montaner, 2008).
The model that explains customer loyalty to a supermarket chain, measured in terms of the purchase frequency (Dick, and Basu, 1994) involves the supermarket image (Grewal et al., 1998; Collins-Dodd and Lindley, 2003; Semeijn et al., 2004), purchase objectives (Baltas, 1997; Putrevu et. al, 2001), perception of the product's manufacturer brands (González et al., 2006), perception of the private labels (Burton et. al., 1998; Garretson et al., 2002; Collins-Dodd and Lindley, 2003; Dick et al., 1995), shopping experience (Martos-Partal and González-Benito, 2013), convenience of the commercial relationship (Flavián et al., 2001), and overall satisfaction (Martínez-Ruiz, 2010).

1.1 Objective
To predict the purchase frequency of an average customer on a supermarket chain by means applying machine learning techniques to build predictive models based on artificial neural networks.

2. Literature Review
Machine learning is usually referred as the branch of artificial intelligence (AI) that uses algorithms to find patterns and to learn from datasets through experience. There are supervised, unsupervised, and reinforcement algorithms. In supervised learning, the training is carried out using labelled datasets. This means that the class or the value to be predicted is included in the dataset so it can be used for training. In the case of unsupervised learning, instead, the desired class is not known.
There are several important tasks in machine learning: classification, regression, and forecasting. From these three, classification is the one that concentrate the interest of this work. It can be understood as the determination of the class, a nominal value, in an unseen dataset using a previously trained model.

2.1 Hold out and cross-validation
Holding out implies the splitting up of a dataset into a set for training and another for testing. The test dataset is employed to assess the performance of the classification model with unseen data. Usually the split up proportion is 80% for training and 20% for testing. On the other hand, cross-validation is the random split up of a dataset into k folds. During the model building, k-1 folds are employed for training while the left one is used to test model's performance. Training and testing are repeated iteratively k times until all folds have been used for testing (Figure 1). The goal is to minimize the risk of overfitting that can happen when holding out. In the case of cross-validation, each iteration produces different results because the folds for training and for testing have been interchanged. These k results are finally averaged.

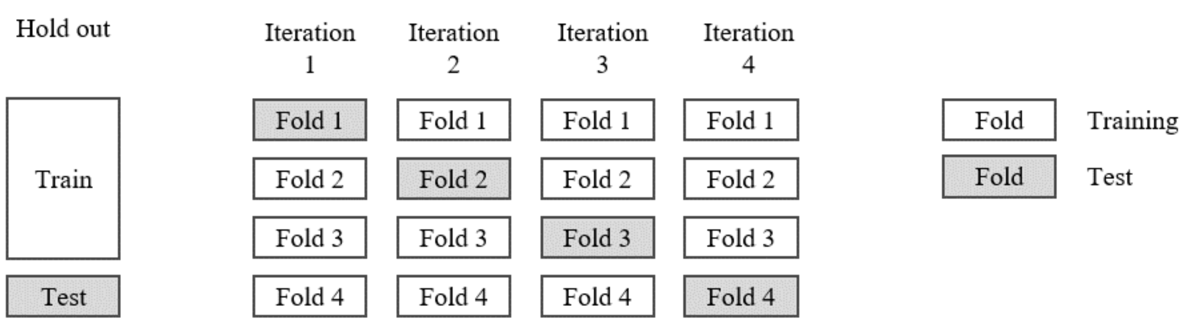

Figure 1. Hold out and cross-validation (k=4)

2.2 Overfitting and generalization

Overfitting occurs when a model learns well from the training dataset but it does not have a good performance when tested with an unseen dataset. In such situation, it is said that the model cannot generalize. This might happen due to the incorporation of many details from the training data that will not be easily found in new data (Figure 2).

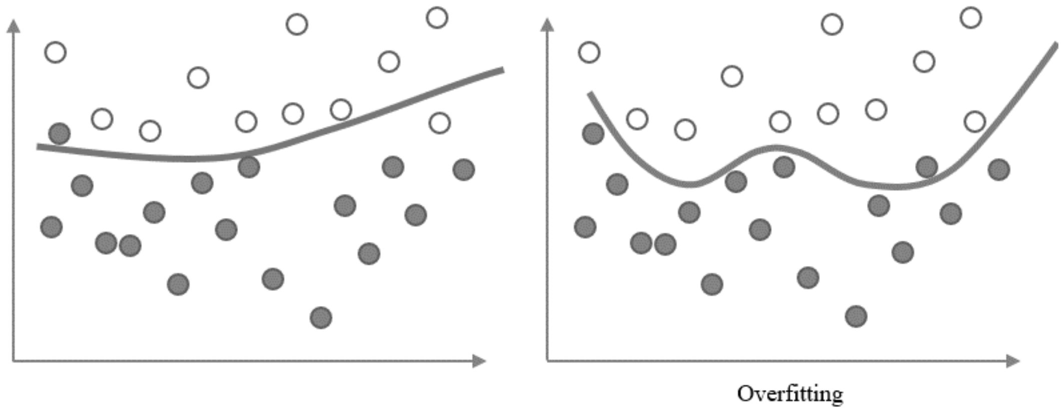

Figure 2. Overfitting

2.3 Replication

Replication is repetition of an experiment under similar conditions to estimate the variability of the results. When using cross-validation, the dataset partitioning into k folds depends on a specific seed number (Figure 3). Since different seed numbers produce different folds, the results of the training and test are different too. By means of replicating the experiments with random seeds each time, it would possible to obtain several test results from which the mean and the standard deviation can be estimated and analyzed afterwards. Thus minimizing the effect of an unfortunate partitioning.

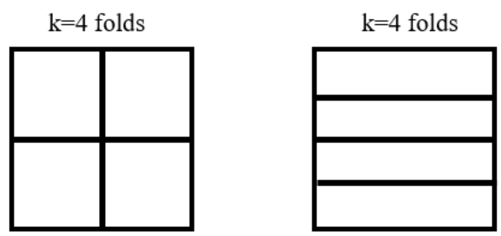

Figure 3. Different folds in cross-validation (k=4)

2.3 Artificial neural network (ANN)

An ANN is a construction made of nodes, referred as neurons, that are combined in an interconnected layered structure (Figure 4). The input layer corresponds to the nodes that receive the externa data. In the second level contains the hidden layers that transform the input data for the output layer, whose neurons are responsible for delivering the results generated by the network (Morano and Tajani 2013). The topology of an ANN is determined by the number of layers, their nodes, and a transfer function.

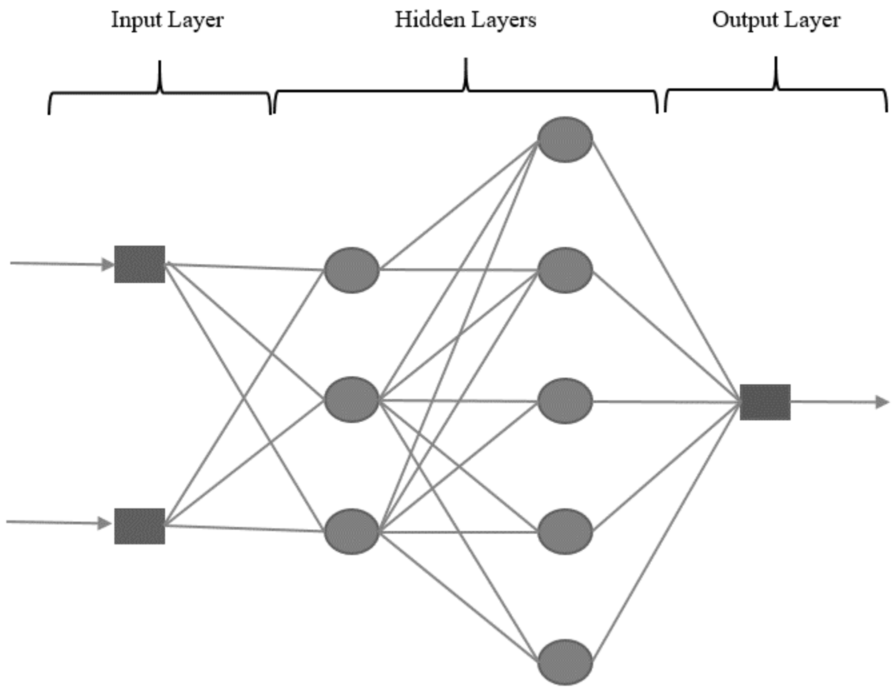

Figure 4. ANN's input, hidden, and output layers

3. Methods

This investigation is carried out following a classic 4-stage model: analysis, design, construction, and validation (Figure 5).

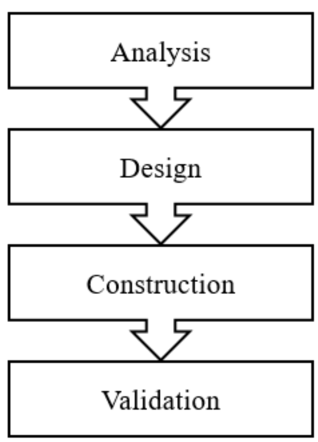

Figure 5. Four-stage model

3.1 Analysis

During the analysis, a deep review of the results of a survey about the purchasing habits and preferences of customers with respect to supermarkets and private labels is carried out. The survey was conducted in Temuco (Chile) and involved five supermarket chains with distinctive target markets. While one of them is focused on convenient stores, other offers exclusive products at higher prices. The questionnaire consists of 69 questions grouped in several domains of interest such us: product quality, product availability & variety, value/price, discount and loyalty campaign, customer service, facility organization, etc. For the purposes of the present investigation only 25 questions grouped in 5 domains are considered, all of which were answered using a scale from 1 to 7 (Table 1).

Table 1. Selected domains

	Domain	Questions
D.1	Products' quality	4
D.2	Prices and discounts	5
D.3	Facility and service	5
D.4	Products' availability and variety	5
D.5	Overall purchase experience	5

Even though five supermarket chains were originally considered in the survey, in the present work only one of them is analyzed. The decision is based in the fact that the selected chain is the middle of the price range, it counts with stores in different districts of the city and it has a well-established private label with a variety of products.
As discussed earlier, in this work customers' loyalty is measured in terms of the purchase frequency on the selected supermarket chain. The goal is to predict the frequency in which an average customer made purchases on the supermarket chain based on the answers to the questions.

3.2 Design

The original dataset, a matrix of 1073 rows (instances) by 25 columns (24 attributes or questions and the class), is split up to create 2 subsets. The first one for training and test, contains 80% of the data (860 instances). The remaining 20% of the data (213 instances) is left in a separate dataset to be used during the validation. An instance can be understood as a row containing all the answers of a single survey.

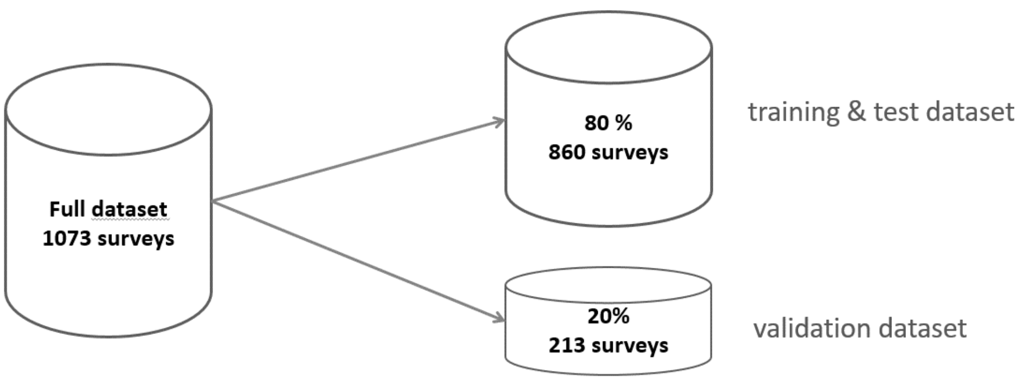

Figure 6. Dataset split up

To identity the relevance of each domain on the classification result, different combinations are defined to build and to compare 12 predictive models. While Model 0 (M0) contains all the questions, Model 6 (M6) contains question only from domains 1, 3, and 5 (Table 2). The design groups domains and questions that might be related.

Table 2. Classification models' design

		M0	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11
D.1	Product quality	✓	✓					✓	✓				✓
D.2	Prices and discounts	✓		✓						✓	✓		✓
D.3	Facility and service	✓			✓			✓		✓		✓	
D.4	Product availability and variety	✓				✓			✓		✓	✓	
D.5	Overall purchase experience	✓					✓	✓				✓	

3.3 Construction

As aforementioned the goal is to predict the frequency of purchases made on the stores of the supermarket chain under study. There are only two possible frequencies: high and low. However, there is also a third value for those customers who provided a different answer. It must be understood that it is a supervised learning task where the frequency is *a priori* known.
All models presented in this work are built using the data mining software WEKA version 3.8.5 (Witten et al., 2017). Initially, all models are trained and tested applying a cross-validation scheme of k=10 folds with a dataset of 860 instances (Table 3).

Table 3. Models' performance with training and test dataset and cross-validation k=10 (860 instances)

		Weighted Average		
Model	**%**	**Precision**	**Recall**	**ROC AUC**
M0	74.42	0.706	0.744	0.631
M1	72.56	?	0.726	0.529
M2	71.86	0.686	0.719	0.595
M3	72.67	0.656	0.727	0.535
M4	71.97	0.704	0.720	0.639
M5	71.05	0.678	0.710	0.542
M6	67.21	0.660	0.672	0.530
M7	71.28	0.699	0.713	0.637
M8	70.93	0.686	0.709	0.600
M9	71.28	0.699	0.713	0.637
M10	71.16	0.702	0.706	0.640
M11	68.37	0.668	0.684	0.544

Predictive models are compared by means of the percentage of correct predictions with unknown data from a validation dataset. Additionally, curves Precision-Recall and the area under the curve ROC (ROC AUC) are considered for comparison too (Davis and Goadrich, 2006).
When dealing with classification problems it is important to consider the class balance. In the case of heavily imbalanced datasets, some authors recommend the inclusion of additional performance metrics such as Precision-Recall curves along with ROC AUC (Saito and Rehmsmeier, 2016). In some cases, it is not possible to calculate the values of the column Precision because of the lack of instances of a given class. This is shown later in the confusion matrices (Table 8).
Although cross-validation helps reduce the risk of overfitting, the effect of the fold partitioning remains (Powers and Atyabi, 2012). The replication of the experiments might mitigate this issue by means of using different folds each iteration. For the purposes of this research 10 replications are run, which means that each model is trained and tested 100 times. In all cases, results of the replications show that the standard deviation of the percentage of correct predictions is less than 4.5% (Table 4).

Table 4. Models' performance with training and test dataset, cross-validation k=10, and 10 replications

Model	**Average Correct Predictions (%)**	**Standard Deviation (10 replications)**
M0	73.69	3.68
M1	73.06	3.32
M2	71.88	3.61
M3	72.64	3.23
M4	72.72	3.75
M5	70.33	4.26
M6	68.35	4.41
M7	72.74	3.74

M8	70.19	4.24
M9	72.74	3.74
M10	72.62	4.10
M11	68.84	4.14

3.4 Validation

The validation of the trained and tested ANN-based models is carried out with the validation dataset held out during the stage of analysis. The validation dataset contains 213 unseen instances (rows), which corresponds to 20% of the data.

4. Data Collection

The results of the validation show that approximately half of the models can generalize relatively well with unseen data (Table 5). Being close to 70% in the best cases and barely above 60% in the worst cases. In all cases the percentage of correct classification is consistent with the percentage obtained with the training dataset (Table 4).

Table 5. Models' performance with the validation dataset (213 instances)

Model	Attributes	Correct Predictions (%)	Precision	Recall	ROC AUC
M0	26	72.30	0.703	0.723	0.665
M1	5	69.48	0.553	0.695	0.593
M2	6	69.95	0.663	0.700	0.623
M3	6	63.85	0.555	0.638	0.545
M4	6	69.48	0.728	0.695	0.730
M5	6	69.48	0.668	0.695	0.642
M6	15	64.79	0.623	0.648	0.547
M7	10	60.09	0.603	0.601	0.599
M8	12	69.48	0.659	0695	0.627
M9	10	60.09	0.603	0.601	0.599
M10	16	70.42	0.706	0.704	0.686
M11	11	61.50	0.601	0.615	0.517

5. Results and Discussion

Since there are differences in the models' performance (Table 5), it might be assumed that the selection of the questions (attributes) is relevant for the prediction (Table 6). Despite of being significantly different in complexity to M1, M2, M4, and M5, model M0 does not perform exceptionally well.

Table 6. Models' size (attributes)

Model	M0	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11
Attributes	26	5	6	6	6	6	15	10	12	10	16	11

5.1 Numerical Results

Confusion matrices are useful to summarize the prediction results in tables. The diagonal of the matrix contains the number of instances correctly classified. The other cells present incorrect classifications (Table 7).

Table 7. Confusion matrix

Class 1	Class 2	Class 3
class 1 classified as class 1	class 1 classified as class 2	class 1 classified as class 3
class 2 classified as class 1	class 2 classified as class 2	class 2 classified as class 3

class 3 classified as class 1	class 3 classified as class 2	class 3 classified as class 3

Confusion matrices for both datasets, training-test (860 instances) and validation (213 instances), show that the models M0, M1, M2, M4, M5, and M10 classify misclassify between 10% and 20% of the frequencies corresponding to Low and between 40% and 60% of the frequencies corresponding to High (Table 8).
It is possible to reduce the misclassification rate by means of including a penalization for wrong classification during the training of the models.

Table 8. Resulting confusion matrices for ANN-based models

Training-test dataset (860 instances)				**Validation dataset 213 instances)**			
M0	**NA**	**Low**	**High**	M0	**NA**	**Low**	**High**
	0	12	1		0	0	0
	2	592	73		0	132	22
	0	132	48		0	37	22
M1	**NA**	**Low**	**High**	M1	**NA**	**Low**	**High**
	0	11	2		0	0	0
	0	622	65		0	147	7
	0	158	22		0	58	1
M2	**NA**	**Low**	**High**	M2	**NA**	**Low**	**High**
	0	13	0		0	0	0
	4	573	90		0	134	20
	0	135	45		0	44	15
M3	**NA**	**Low**	**High**	M3	**NA**	**Low**	**High**
	0	13	0		0	0	0
	3	605	59		1	132	21
	0	160	20		0	55	4
M4	**NA**	**Low**	**High**	M4	**NA**	**Low**	**High**
	0	11	2		0	0	0
	6	562	99		0	112	42
	2	121	57		0	23	34
M5	**NA**	**Low**	**High**	M5	**NA**	**Low**	**High**
	0	12	1		0	0	0
	2	568	97		1	131	22
	0	137	43		0	42	17
M6	**NA**	**Low**	**High**	M6	**NA**	**Low**	**High**
	1	10	2		0	0	0
	2	537	128		0	123	31
	1	139	40		0	44	15
M7	**NA**	**Low**	**High**	M7	**NA**	**Low**	**High**
	0	12	1		0	0	0
	8	55	104		0	111	43
	0	122	58		0	42	17
M8	**NA**	**Low**	**High**	M8	**NA**	**Low**	**High**
	0	11	2		0	0	0
	2	560	105		1	134	19

	0	130	50
M9	**NA**	**Low**	**High**
	0	12	1
	8	555	104
	0	122	58
M10	**NA**	**Low**	**High**
	1	10	2
	4	552	111
	0	121	59
M11	**NA**	**Low**	**High**
	0	11	2
	6	543	118
	0	135	45

	0	45	14
M9	**NA**	**Low**	**High**
	0	0	0
	0	111	43
	0	42	17
M10	**NA**	**Low**	**High**
	0	0	0
	0	122	32
	0	31	28
M11	**NA**	**Low**	**High**
	0	0	0
	0	117	37
	1	44	14

6. Conclusion

The survey about supermarket chains and private labels used this research provides enough data to build models to predict customers' loyalty in terms of their frequency of purchase. In other words, how often they visit the chain's stores. Whether the purchase frequency is the best indicator to measure the loyalty is not part this work's scope, which, instead, is limited to the building of predictive models based on ANNs.

During the building of a machine learning model it is recommendable to implement a cross-validation scheme to minimize the influence of the dataset split up instead of holding out a portion of the data. Averaging the results of k iteration will be always better than having only one result. Furthermore, replicating the experiment many times allow estimate the variability of the results, which can be expressed in terms of the standard deviation of the percentage of correctly predicted classes.

On the other hand, having a validation helps determine whether the models can generalize properly or not. Since, considering a complete rotation of k folds, during the cross-validation each instance (rows) is used for training k-1 times, data are known to the predictive. For this reason, it is crucial to have a separated validation set to evaluate the models' performance with unknown data.

The results of the experiment replications show that only models M6 and M11 predict correctly less than 70% of customers' purchase frequency with the training and test dataset. With the validation dataset, instead, only models M0, M1, M2, M4, M5 and M10 predicted correctly over around 70% of the instances. In the latter case, although prediction ratios are similar, between 69% and 72%, there are significant differences in the number of attributes used to build each model, being M0 the largest one with 26 attributes and M1 the smallest with only 5.

From the prediction rate, it seems that each domain provides enough information to build models that have a similar performance. This a particularly important result because it proves that a short and focused survey is equally useful than an extensive one when predicting purchase frequency.

Finally, prediction ratios and confusion matrices suggest that the ANN-based classification models M0, M1, M2, M4, M5 and M10 help predict customers' purchase frequency.

References

Ailawadi, K., Pauwels, K., and Steenkamp, J., Private-Label Use and Store Loyalty, *Journal of Marketing*, vol. 72, no. 6, pp.19–30, 2008.

Ailawadi, K., Neslin, S., Luan, Y., and Taylor, G., Does Retailer CSR Enhance Behavioral Loyalty? A Case for Benefit Segmentation, *International Journal Research Marketing,* vol. 31, pp.156–167, 2014.

Allen, D., and Rao, T., Analysis of Customer Satisfaction Data, Milwaukee: Quality Press, 2000.

Baltas, G., Determinants of Store Brand Choice: A Behavioral Analysis, *Journal Product Brand Management*, vol. 6, pp. 315–324, 1997.

Berkowitz, D., Bao, Y., and Allaway, A., Hispanic Consumers, Store Loyalty and Brand Preference, *Journal of Targeting, Measurement and Analysis for Marketing*, vol.14, pp. 9–24, 2005.

Burton, S., Lichtenstein, D. R., Netemeyer, R. G. & Garretson, J. A., A Scale for Measuring Attitude Toward Private Label Products and an Examination of its Psychological and Behavioral Correlates, *Journal of the Academy of Marketing Science*, vol. 26, pp. 293–306 (1998).

Collins-Dodd, C., and Lindley, T., Store brands and retail differentiation: the influence of store image and store brand attitude on store own brand perceptions, *Journal of Retailing and Consumer Services*, vol. 10, no. 6, pp. 345–352, 2003.

Corstjens, M., and Lal, R., Building Store Loyalty Through Store Brands, *Journal of Marketing Research*, vol.37, pp. 281–291, 2000.

Davis, J., Goadrich, M., The Relationship Between Precision-Recall and ROC Curves, *Proceedings of the 23rd International Conference on Machine Learning*, 2006.

Dick, A., and Basu, K., Customer Loyalty: Toward an Integrated Conceptual Framework, *Journal of the Academy of Marketing Science*, vol. 22, no. 2, pp. 99-113, 1994.

Dick, A., Jain, A., and Richardson, P., Correlates of Store Brand Proneness: Some Empirical Observations, *Journal of Product & Brand Management*, vol.4, pp. 15–22, 1995.

Flavián, C., Martínez, E., and Polo, Y., Loyalty to Grocery Stores in the Spanish Market of the 1990s, *Journal of Retailing and Consumer Services*, vol. 8, no. 2, pp. 85–93, 2001.

Garretson, J., Fisher, D., and Burton, S., Antecedents of Private Label Attitude and National Brand Promotion Attitude: Similarities and Differences, *Journal of Retailing*, vol. 78, pp. 91–99, 2002.

González M., Díaz-Martín, A., and Trespalacios, J., Antecedents of the Difference in Perceived Risk Between Store Brands and National Brands, *European Journal of Marketing*, vol. 40, pp. 61–82, 2006.

Grewal, D., Krishnan, R., Baker, J., and Borin, N., The Effect of Store Name, Brand Name and Price Discounts on Consumers' Evaluations and Purchase Intentions, *Journal of Retailing*, vol. 74, no. 3, pp. 331–352, 1998.

Martenson, R., Corporate Brand Image, Satisfaction and Store Loyalty: A Study of the Store as a Brand, Store Brands and Manufacturer Brands, *International Journal of Retail & Distribution Management*, vol. 35, no. 7, pp. 544–555, 2007.

Martínez, E., and Montaner, T., Characterisation of Spanish Store Brand Consumers, *International Journal of Retail & Distribution Management*, vol. 36, no. 6, pp. 477–493, 2008.

Martínez-Ruiz, M., Jiménez-Zarco, A., and Izquierdo-Yusta, A., Customer Satisfaction's Key Factors in Spanish Grocery Stores: Evidence from Hypermarkets and Supermarkets, *Journal of Retailing and Consumer Services*, vol. 17, no. 4, pp. 278–285, 2010.

Martos-Partal, M., and González-Benito, Ó., Studying Motivations of Store-Loyal Buyers Across Alternative Measures of Behavioral Loyalty, *European Management Journal*, vol. 31, no. 4, pp. 348–358, 2013.

Morano, P., and Tajani, F., Bare Ownership Evaluation. Hedonic Price Model v/s Artificial Neuroal Network. International Journal of Business Intelligence and Data Mining, vol. 8, no. 4, pp. 340-360, 2013.

Neumeier, M., The Brand Gap, 1st Edition, New Riders Publishing, 2006.

Powers, D., and Atyabi, A., The Problem of Cross-Validation: Averaging and Bias, Repetition and Significance. *2012 Spring World Congress on Engineering and Technology*, SCET 2012 – Proceedings, pp.1-5, 2012.

Putrevu, S., and Lord, K., Search Dimensions, Patterns and Segment Profiles of Grocery Shoppers, *Journal of Retailing and Consumer Services,* vol. 8, no. 3, pp. 127–137, 2001.

Saito, T., and Rehmsmeier, M., The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets, *PLoS ONE 10(3): e0118432*, 2015.

Semeijn, J., van Riel, A., and Ambrosini, A., Consumer Evaluations of Store Brands: Effects of Store Image and Product Attributes. *Journal of Retailing and Consumer Services,* vol. 11, pp. 247–258, 2004.

Witten, I., Frank, E., Hall, M., and Pal, C., *Data Mining: Practical Machine Learning Tools and Techniques*, 4th Edition, Morgan Kaufmann, Cambridge, 2017.

Biographies

Galo Paiva is an industrial engineer, consultant, and university professor. He earned Licentiate Degree in Engineering from Universidad de Santiago de Chile, Chile, and Doctor of Business Management from Universidad Autónoma de Madrid, Spain. He has taught lectures in Strategic Management, Operations Management, Industrial Engineering, and Project Planning & Management. His research interests include manufacturing process simulation, industrial design, business management, and entrepreneurship.

Carlos Hernández is an industrial engineer, consultant, and university professor. He earned Licentiate Degree in Engineering from Universidad de La Frontera, Temuco, Chile, Master of Sciences in Computational Engineering and Doctor of Engineering from Technische Universität Braunschweig, Brunswick, Germany. He is the author of several scientific and engineering articles. He has taught lectures in Discrete Event Simulation, Supply Chain Management, Engineering Economics, Corporate Finances, Financial Engineering, Business Analytics, Data Mining and Machine Learning for engineering students. He has developed a professional career working for large multinational companies (PricewaterhouseCoopers, BHP Billiton, and Merck Sharp & Dohme). He also worked as a scientific researcher in the Institut für Produktionsmesstechnick at TU Braunschweig, Germany. His research interests include manufacturing process simulation, transportation systems simulation, supply chain design and simulation, and machine learning for finances. He is a member of IEOM.