Competing for Amazon's Buy Box: A Machine-Learning Approach

Álvaro Gómez-Losada^(✉) and Néstor Duch-Brown

Joint Research Centre, European Commission, Seville, Spain
alvaro.gomez.losada@gmail.com,
nestor.duch-brown@ec.europa.eu

Abstract. A key feature of the Amazon marketplace is that multiple sellers can sell the same product. In such cases, Amazon recommends one of the sellers to customers in the so-called ‘buy-box’. In this study, the dynamics among sellers for occupying the buy-box was modelled using a classification approach. Italy’s Amazon webpage was crawled during ten months and features from products analyzed to estimate the more relevant ones Amazon could consider for a seller occupy the buy-box. Predictive models showed that the more relevant features are the ratio between consecutive prices in products and their number of assessment received by customers.

Keywords: Buy-box · Amazon · Machine learning · Classification · Data science

1 Introduction

Amazon is currently the leading online e-commerce platform in Europe [1] and attracts exceptionally large public interest. In this platform, Amazon acts as a re-seller of some products and also as a marketplace that enables independent sellers to sell new or used products. Sellers have the possibility to use the platform’s different services, including among others management of inventories, advertisement of products and the *Fulfilled by Amazon* program, by which Amazon handle logistics for independent sellers’ products. A key feature of the Amazon marketplace is that multiple sellers can offer the same product. In such cases, Amazon recommends one of the competing sellers to customers in the so-called ‘buy-box’ [2]. This box is placed on every product detail page so customers can begin the purchase process by adding items to their shopping carts directly.

One of the main challenges in Amazon marketplace operation is the lack of understanding of the empirical mechanisms adopted by sellers for market competition. To the best to the authors’ knowledge, just one study has focused on studying mechanisms that sellers adopt to gain exposure to consumers in Amazon marketplace from a machine learning perspective. Chen et al. [3] explored algorithmic pricing strategies from sellers in Amazon marketplace by obtaining products and seller characteristics. These authors examined the empirical behavior of the buy-box to understand relevant features by which Amazon selects a given seller among competitors to occupy the buy-box. These authors studied sellers’ and products’ characteristics and

estimated that the product price related characteristics are the most relevant for winning the buy-box.

In this study, the dynamics of competition among sellers for occupying the buy-box was modeled using a predictive approach based on classification. After scraping Italy's Amazon webpage, longitudinal datasets were created for selected products describing over time the different default sellers selected by Amazon to sell such products (i.e., to occupy the buy-box). The classification models built on each longitudinal dataset aimed to predict when a change of seller for a product occurs at the buy-box.

The objective of this work is twofold. First, we aim to contribute to the empirical understanding of the Amazon's buy-box algorithm, detecting the more relevant features that explain that a given seller is featured to occupy the buy-box. Second, we offer a characterization of the more representative sellers according to these features. The remainder of this study is organized as follow. In Sect. 2, the methodology applied to estimate the more relevant features from sellers for becoming buy-box eligible is described. A discussion about these estimated features is provided in Sect. 3, together with the characterization of sellers according to them. Finally, in Sect. 4, the main conclusions are presented.

2 Research Methodology

The methodology used in this study is illustrated in Fig. 1 which summarize: (i) the characteristics of the database obtained, (ii) the creation of a longitudinal dataset for every analyzed product, (iii) the creation and selection of features in each dataset,

Fig. 1. Methodology and analyses performed. Blue boxes indicate the main obtained results (RF: random forest; SVM: support vector machine; NN: neural network; Acc: Accuracy; BAcc: Balanced accuracy; k : kappa statistic; PCA: Principal component analysis, k-M: k-means). (Color figure online)

(iv) the modeling step using three different classifiers, (v) the estimation of relevant features for winning the buy-box by analyzing all the individual results from products, and (vi) the clustering of sellers considering these latter features.

2.1 Data Collection

Data for this study was obtained using a crawler that navigated through 26 categories of new products in Italy’s Amazon webpage, over the course of 10 months, beginning on April 4, 2018. These categories were selected according to the results from a previous crawling experiment to detect the more dynamic ones in terms of products’ price changing and presence of sellers offering the same products (some digital categories including apps for Android, *Kindle* or products offered primarily by Amazon like videos were not considered). For each category, just the details from most popular products were extracted. The interest focused on the most demanded ones which are likely to describe seller’s strategies for maximizing benefits, and therefore providing more information and easing the modeling process. Due to Amazon’s strategic commercial reasons, the number of most popular products displayed for each category varied along the crawling period (*top-20*, *top-50* or *top-100*) causing the frequency of crawling cycles to differ (≈ 1 h to complete the crawling of *top-20* products displayed for each category). Each time the crawler visited a product page, it recorded the time of crawling, the characteristics of the product and those from the seller featured in the buy-box. An example of product characteristics and the buy-box is illustrated in Fig. 2. To consider the time the crawler extracted information from products through navigation allowed for obtaining a longitudinal approach for each of them.

Olio di Cocco Biologico Extra (1)

di NaturaleBio

★★★★★ (3)

1,561 recensioni clienti

89 domande con risposta

Amazon's Choice (7)

"olio di cocco"

Prezzo: EUR 11,97

Spedizione GRATUITA per ordini superiori a EUR 29.

Tutti i prezzi includono l'IVA.

5% di sconto sull'acquisto di 3 Pezzi

1 promozione

EUR 11,97 (2)

Spedizione GRATUITA per ordini superiori a EUR 29. Maggiori informazioni

Vuoi riceverlo domani? Ordina entro 29 min e scegli la spedizione 1 giorno. Maggiori informazioni

(6) Disponibilità immediata. (5)

Venduto da NaturaleBio e spedito da Amazon. Con (3) come regalo disponibile.

Il tuo indirizzo di consegna: Spagna

Quantità: 1

Aggiungi al carrello

Acquista ora

Fig. 2. Characteristics from a given product (A) and buy-box in Italy’s Amazon webpage (B). Numbers in brackets are explained in Table 1.

Table 1. Features from products extracted during the crawling process.

Feature	Description
Product	Unique identifier of the product (ASIN code)
Price	Seller price for the product
Ratings	Customers rating of the product
Assessments	Number of assessments (opinions) the product received by customers
Fulfilled	Fulfilment of the product by Amazon
Stock	Availability of the product
Choice	Recommended product by Amazon (Amazon's choice)
Seller	Unique identifier of the seller selected by Amazon to sell the Product
Time	Time instant the product was crawled

2.2 Longitudinal Datasets for Products

Once semi-structured data from crawling was processed, an indexed-by-time database containing the products' and sellers' characteristics was obtained. Such database was filtered by products, obtaining an independent longitudinal dataset for each of them. The aim was to describe the change over time of characteristics of products featured in Table 2, and importantly, to monitor the different sellers that have sold each product. For modelling purposes a total of 461 products were selected to be studied. These products fulfilled a double criterion. Firstly, each product was offered by more than three sellers, and secondly, information from products by crawling was obtained a minimum of 110 occasions. This latter threshold was stablished according to the one in five ratio of examples (instances) per predictor variable [4, 5], considering that a maximum of 21 candidate predictors were considered to build the predictive models (Table 2) in each longitudinal dataset. These predictors are analyzed next.

2.3 Feature Building and Selection

This section describes the creation and selection of features used for building the predictive models. For every longitudinal dataset of products, a set of 21 candidate predictors (Table 2) were created to build a binary classification model. These predictors included all features from Table 1 (except ASIN code) as well as a set of 13 dynamic features generated from price, aiming to study the possible implications of product prices' trends and variations applied by sellers for winning the buy-box. In order to reduce the high dimensionality in these datasets, those predictors with near-zero variance, showing a linear dependency or correlated (Pearson's r higher than ± 0.8) were removed. Additionally, using a Principal Component Analysis (PCA), those predictors failing to explain a 95% of the variance in the longitudinal dataset were excluded for building the predictive models. These techniques for feature selection were used in each longitudinal dataset from products. To detect variables with variance close to zero, correlated or showing linear dependency, the *nearZeroVar*, *findLinearCombos* and *findCorrelation* functions from the *caret* package [6, 7] in *R* [8] were used, respectively. PCA for feature selection was applied with the *preProcess* function from the same package.

Labelling. For each instance in every longitudinal dataset, a positive or negative class was assigned according to the next criteria. The first time a new seller was selected to occupy the buy-box for selling a product, the value “1” was assigned to the instance (representing the positive class), and “0” otherwise. An example of the labeling process is illustrated in Fig. 3. In this study, the selection by Amazon of a new merchant to occupy the buy-box is considered the more decisive instant to understand the algorithm that performs such selection. It is assumed that once a new merchant is selected, its continuity in time by winning the boy-box does not provide conclusive information, since this could be achieved simply by either inactivity (e.g., out of stock) or unknown strategic decisions from other competing sellers.

2.4 Classification

Three well-known classifiers, namely, random forest (RF), support vector machine (SVM) and neural networks (NN), were used independently to predict if a competing seller is likely to replace the current one occupying the buy-box. It is important to remind that the purpose of this predictive approach is not to identify the seller likely to win the buy-box among others offering the same product, but to estimate if the current merchant selling a product is going to be replaced by another competing seller also offering the same product. For each longitudinal dataset from products, such classification models were built and the prediction accuracy evaluated. Once accomplished this prediction exercise, the relative importance of the involved predictors (after the feature selection process) was estimated.

Product	Seller	Label
1	A	1
1	A	0
1	A	0
1	B	1
1	B	0
1	C	1
1	D	1
1	D	0

Fig. 3. Example of positive and negative class assignments (“1”: positive class) in a longitudinal dataset.

Model Building. Longitudinal datasets from each product were split in training and test sets following a 70:30 proportion, respectively. For tuning parameters of classifiers, ten trials were carried out using the training data, repeated three times, following a 10-fold cross-validation (CV) resampling scheme. Every tuned classifier was built on nine training folds using the CV scheme, and predictions were evaluated on the remaining held-out fold using the ROC (receiver operator characteristic) metric, obtaining a performance profile based on all the tuning parameters tested. After identifying the

optimum model among the candidates based on their performance, this one was then re-built using the whole training data set (without fractions of the training data being held-out). To handle possibly unbalance between positive and negative classes, which could bias the model towards the positive class, a stratified version of the CV scheme was adopted. To that end, instances containing the majority (negative) class were randomly removed to ensure an approximate equal presence of both classes [9]. The *caret* package from the R software was used to build classifiers (*train* function), and the *down-sampling* argument in the *preProcess* function was used to implement the stratified CV.

Table 2. Candidate features used to build classification models for products.

Abbreviation	Type	Role	Value
<i>Change in seller</i>	Categorical	Response	1: The seller winning the buy-box has changed (positive class) 0: The seller winning the buy-box remains (negative class)
<i>dPrLowestPr</i>	Numeric	Predictor	Difference between the current price and lowest price
<i>dPrHighestPrP</i>	Numeric	Predictor	Difference between the highest price and current price
<i>dPrRM32pr</i>	Numeric	Predictor	Difference between current price and 32-prices rolling mean
<i>dPrRM16pr</i>	Numeric	Predictor	Difference between current price and 16-prices rolling mean
<i>dPrRM8pr</i>	Numeric	Predictor	Difference between current price and 8-prices rolling mean
<i>dPrRM4pr</i>	Numeric	Predictor	Difference between current price and 4-prices rolling mean
<i>rPrLowestPr</i>	Numeric	Predictor	Ratio between the current price and the lowest price
<i>rPrHighestPr</i>	Numeric	Predictor	Ratio between the current price and the highest price
<i>rPrRM32pr</i>	Numeric	Predictor	Ratio between the current price and 32-prices rolling mean
<i>rPrRM16pr</i>	Numeric	Predictor	Ratio between the current price and 16-prices rolling mean
<i>rPrRM8pr</i>	Numeric	Predictor	Ratio between the current price and 8-prices rolling mean
<i>rPrRM4pr</i>	Numeric	Predictor	Ratio between the current price and 4-prices rolling mean
<i>rPr</i>	Numeric	Predictor	Ratio between the current and previous price
<i>Ops</i>	Numeric	Predictor	Number of assessment the product has received by customers
<i>ProdRating</i>	Numeric	Predictor	0 to 5. Rating of the product received by customers

(continued)

Table 2. (continued)

Abbreviation	Type	Role	Value
<i>Day</i>	Numeric	Predictor	1 to 31
<i>Wday</i>	Categorical	Predictor	1: The product is displayed on a working day; 0: otherwise
<i>AmChoice</i>	Categorical	Predictor	1: The product is Amazon's choice; 0: otherwise
<i>Bestseller</i>	Categorical	Predictor	1: The product is a bestseller; 0: otherwise
<i>Fulfilled</i>	Categorical	Predictor	1: The product is fulfilled by Amazon; 0: otherwise
<i>Stock</i>	Categorical	Predictor	1: The product is in stock; 0: otherwise

Evaluation. To measure the classification accuracy, a 2×2 confusion matrix was set to summarize the number of instances which class label is predicted correctly or incorrectly. This classification accuracy was estimated for every predictive model built on each longitudinal dataset from products, and for every classifier. Considering that no one criteria is sufficient to assess the performance of a classifier [10], the accuracy, balanced accuracy and the kappa statistic (Cohen's kappa) [11] were used in this work.

The first metric, accuracy, measures the proportion of examples correctly classified. Balanced accuracy is estimated as the average value of sensitivity and specificity, and it provides an estimation on the average accuracy obtained on either class [12]. The third one, the kappa statistic takes into account the accuracy that would be generated simply by chance. The form of the statistic is $k = (O - E) / (1 - E)$, where O is the observed accuracy and E is the expected accuracy based on the marginal totals of the confusion matrix. The statistic can take on values between -1 and 1 : a value of 0 means there is no agreement between the observed and predicted classes, while a value of 1 indicates perfect concordance of the model prediction and the observed classes [7].

Relative Importance of Predictors. Once every predictive model was built on each longitudinal datasets from products, using the three different classifiers, their accuracy was evaluated on test sets and the relative importance of the predictors in the models estimated. On each predictor, a ROC curve analysis was conducted and a series of cut-offs applied to the predictor data to predict the class. ROC analysis is a useful metric since allows for estimating the optimal model independently from the class distribution. The area under to ROC curve was used as the measure of variable importance. Relative importance of the variables was estimated using the *varImp* function from the *caret* package.

The selection of the more important variables for sellers to win the buy-box after building a predictive model for each of 461 products, and for each classification model, was accomplished in a two-step process. Arbitrarily, the five more relevant variables (1-highest to 5-lowest) from classification models were estimated for every product. Then, considering all the products, the frequency distribution of the variables with

highest relevance obtained. Secondly, from this frequency distribution, a tabulation scheme aggregated the range of highest frequencies containing the four more decisive variables for winning the buy-box. Thus, the most relevant features after considering all products were obtained.

Relevant Feature Analysis and Clustering of Sellers. Database was filtered by seller and their sold products obtained. Each seller was characterized by a single vector containing the average value of the relevant variables obtained from each product in its portfolio. Less representative sellers were removed and just those selling more than four products studied, obtaining a total of 525 sellers to be analyzed. From this characterization two analyses were conducted. The first one was a *bi-plot* study of relevant variables to understand their linear relation, using the *factoextra* library [13] in *R*, and secondly, a *k-means* cluster analysis of sellers. Intuitively, the latter analysis tries to detect patterns in sellers that can be aware of these relevant variables for considering them in their selling strategic decisions. The optimal number of clusters for grouping of sellers was obtained using the *NbClust* package [14] from *R*. The function *kmeans* from base *R* was used to perform the k-mean analysis.

3 Results and Discussion

3.1 Importance of Features

Figure 4 shows the frequency of the five more important features estimated after building the 461 predictive models (x-axis: from the 1-highest to 5-lowest importance). In Table 3, from variables with highest importance, the four more frequent features from each model are indicated as well as their percentage of occurrence. Some of these

Fig. 4. Frequencies of variable importance in 461 predictive models using three different classifiers (support vector machine -SVM-, random forest -RF- and neural networks -NN-). In the x-axis, the relative importance of features is shown from left (1-highest) to right (5-lowest). By using a color code, each square indicated the number of times a variable is considered important in this ranking (1-highest to 5-lowest). (Color figure online)

variables (*Ops* and *Day*) are considered relevant in all classifiers and *rPr* in two of them (RF and NN). The predictive accuracy of each classifier was estimated after averaging the quality of the predictions accomplished in the 461 datasets (Table 4). As can be appreciated, RF provides the more accurate predictions in all the quality measures used.

According to the more accurate classifier (RF) the more decisive feature to be considered for gaining the buy-box is the ratio between the current price and the previous price of a product (*rPr*), followed by the number of assessment received by customers (*Ops*). According to [3], who used RF as only classifier using a different set of features, the most important one is the price difference to the lowest price (*dPrLowestPr*), followed by the price ratio to the lowest price (*rPrLowestPr*). Results of RF and NN are similar for the three more relevant features. Even SVM classifier provides the less accurate predictions, the variable importance estimations from this algorithm were also considered. The rest of features *Day*, *dPrLowestPr*, *rPrHighestPr* and *rPrRM32pr* seem to have a secondary, but worth to consider, role in this regard.

Table 3. Relevant features obtained from predictive models. Support vector machine (SVM), random forest (RF) and neural networks (NN). Relative frequency (%) is indicated in brackets.

Order of Importance	Classifiers		
	SVM	RF	NN
1	<i>Ops</i> (20%)	<i>rPr</i> (88%)	<i>rPr</i> (33%)
2	<i>Day</i> (9%)	<i>Ops</i> (4%)	<i>Ops</i> (28%)
3	<i>dPrLowestPr</i> (8%)	<i>Day</i> (3%)	<i>Day</i> (5%)
4	<i>Stock</i> (8%)	<i>rPrRM32pr</i> (1%)	<i>rPrHighestPr</i> (5%)

Complementarily to the estimation of the more relevant variables from predictive models, these variables were alternatively used to characterize sellers. Values of these variables in each product from seller's portfolio were averaged. Thus, it was possible to characterize each seller by a single vector defining its general behavior with respect the most relevant features to win the buy-box. Importantly, this approach allow for performing a bi-plot analysis of sellers to understand the linear relation among relevant variables. In Fig. 5, this relation is illustrated as well as the contribution of each of them to explain the variability of seller using two principal components (PCA-biplot). As expected, variables describing the relation of prices to the lowest (*dPrLowestPr*) and highest (*rPrHighestPr*) prices for products are negatively correlated, even quantifying a different measure (difference and ratio, respectively). Conversely, a high correlation is found between the variables studying the ratio between the price (*rPr*) and the 32-prices rolling mean (*rPrRM32pr*), indicating that this latter statistics describes effectively the trend of the prices for analyzed products. Remaining studied variables (*Stock*, *Day* and *Ops*) are not significantly useful to describe variability in sellers and do not show a remarkable association. All variables contributions explain around the 86% of variability for seller characterization.

Fig. 5. Biplot analysis of the more relevant features. Color in arrows indicates contribution of each variable to explain the variability in seller characterization. Positive correlated variables point to the same side of the plot. (Color figure online)

Table 4. Quality measure results for predictive modeling. Support vector machine -SVM-, random forest -RF- and neural networks -NN-.

Classifiers			
Classifiers	Accuracy	Balanced accuracy	kappa
SVM	0.91	0.56	0.15
RF	0.94	0.76	0.52
NN	0.92	0.59	0.20

3.2 Importance of Features

A *k-means* analysis was carried out using the seller characterization from Sect. 3.1, considering six the optimum number of clusters. In the interest of space, graphical results are omitted. The value of the more relevant features (Table 3) from each cluster of sellers is indicated in Table 5, together with the size of each cluster.

Table 5. Value of analyzed features in cluster of sellers, including size of clusters (values in € when applicable).

		Features							
Cluster	Size	Ops	Day	rPr	dPrLowestPr	rPrHighestPr	Stock	rPrRM32pr	
1	32	2027	15	-0.3	2.1	-4.2	0.8	-0.03	
2	24	71	15	4.0	5.9	-19.6	0.9	0.8	
3	103	220	15	0.06	0.9	-3.4	0.4	-0.02	
4	1	202	19	24.6	0.4	-16.7	1.0	9.3	
5	1	15	17	0.6	62.0	-23.1	0.9	-0.5	
6	364	160	16	0.1	1.2	-4.3	0.9	0.05	

First, third and sixth clusters are poorly delimited which centres are close to the axes origin. These three clusters hold most of analysed seller and as expected, they do not differ greatly in the value of the analyses features (Table 5). It is worth to note that Amazon seller is included in the sixth cluster, and also that 50% of sellers grouped in the second cluster are devoted to sell imitation jewellery items. Fourth and fifth clusters are composed by a single seller. Value of *Day* feature in clusters has approximately same value, indicating that most of the products are offered along the all days of the month. Short variations in these values are not conclusive enough to obtain an indication of product selling during specific days of the month and its analysis remains open for further investigation.

Seller in fourth cluster offered just six products competing by winning the buy-box, and none of its competitor was Amazon. The more relevant feature of this seller is the high value of the *rPr* feature and their products' permanent availability (stock presence). Surprisingly, this seller offers 6000 low-cost products through all product categories in Italy's Amazon webpage, enjoying a particular shop-window. After visiting many products displayed in this shop-window, it seems that those products are exclusively sold and fulfilled by this seller. Seller from fifth cluster is represented by an United State technological corporation presented in US's and all Europe's Amazon webpages, specialized in providing video and image processing chips for connected consumer camera applications. In the database, it was presenting just nine products competing for winning the buy-box. No other rival than Amazon was detected for selling products offered by this merchant.

4 Conclusions

This study aims to estimate the more relevant features of sellers to be eligible to occupy the buy-box, and to characterize sellers according to these features. Predictive models showed that the more relevant features are the ratio between consecutive prices in products (*rPr*) and the number of assessment received by customers for products (*Ops*). Cluster analysis of sellers considering these features showed a set of overlapped clusters which does not provide relevant information about these characteristics. However, this analysis revealed two clusters represented by singles sellers, together with another one, all of them with distinct characteristics.

Disclaimer. The views expressed are purely those of the authors and may not in any circumstances be regarded as stating an official position of the European Commission.

References

1. E-commerce News Europe (2018). <https://ecommercenews.eu/the-most-visited-online-marketplaces-by-country/>. Accessed 15 Feb 2019
2. Amazon: Buy box eligible status (2017). <https://sellercentral-europe.amazon.com/forums/t/when-will-i-be-buy-box-eligible/157401>. Accessed 20 Feb 2019

3. Chen et al.: An empirical analysis of algorithmic pricing on Amazon marketplace. In: WWW 2016 Proceedings of the 25th International Conference on World Wide Web, Montreal, Québec, Canada, 11–15 April 2016, pp. 1339–1349 (2016)
4. Vittinghoff, E., McCulloch, C.E.: Relaxing the rule of ten events per variable in logistic and cox regression. *Am. J. Epidemiol.* **165**(6), 710–718 (2007)
5. Austin, P.C., Steyerberg, E.W.: The number of subjects per variable required in linear regression analyses. *J. Clin. Epidemiol.* **68**(6), 627–636 (2015)
6. Kuhn, M.: Building predictive models in r using the caret package. *J. Stat. Softw.* **28**(5), 1–26 (2008)
7. Kuhn, M., Johnson, K.: *Applied Predictive Modelling*. Springer, New York (2013). <https://doi.org/10.1007/978-1-4614-6849-3>
8. R Development Core Team: *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria (2018). <http://www.R-project.org/>. ISBN 3-900051-07-0
9. Galar, M., Fernández, A., Barrenechea, E., Bustince, H., Herrera, F.: A review of ensembles for the class imbalance problem: bagging-, boosting-, and hybrid-based approaches. *IEEE Trans. Syst. Man Cybern. Part C (Appl. Rev.)* **42**(4), 463–484 (2012)
10. Diettrich, T.G.: Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Comput.* **7**(10), 1895–1924 (1998)
11. Carletta, J.: Assessing agreement on classification tasks: the kappa statistic. *Comput. Linguist.* **22**(2), 249–254 (1996)
12. Brodersen, K.H., et al.: The balanced accuracy and its posterior distribution. In: 2010 20th International Conference on Pattern Recognition. IEEE Computer Society, Istanbul, Turkey (2010)
13. Kassambra, A., Mundt, F.: *Factor extra: extract and Visualize the Results of Multivariate Data Analyses* (2017). <https://rpkgs.datanovia.com/factoextra/index.html>
14. Charrad, M., Ghazzali, N., Boiteau, V., Niknafs, A.: NbClust: an R package for determining the relevant number of clusters in a data set. *J. Stat. Softw.* **61**(6), 1–32 (2014)

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

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