XAI for Product Demand Planning: Models, Experiences, and Lessons Learnt

Fenareti Lampathaki, Enrica Bosani, Evmorfia Biliri, Erifili Ichtiaroglou, Andreas Louca, Dimitris Syrrafos, Mattia Calabresi, Michele Sesana, Veronica Antonello, and Andrea Capaccioli

1 Introduction

The H2020 XMANAI project represents a unique experience in Whirlpool's Operations Excellence, within the I4.0 technology research work stream, as it really faces one of the most commonly experienced obstacles to the successful introduction of Artificial Intelligence (AI) technologies. Human engagement in the AI Loop is commonly addressed through a structured change management which acts on communication and training to achieve the readiness level that enables adoption by real users. XMANAI has developed dedicated tools that serve the customized need of the user to *understand* and *trust* the AI system.

Besides, the project's strength lies not only on the implemented solution but also in the journey to arrive there. The latter unveils the complexity of the *explainability* requirements for a user, even when one is aware of what to look for. The conversion of desired information into a format which is really usable by the business experts,

F. Lampathaki (✉) · E. Biliri · E. Ichtiaroglou · A. Louca · D. Syrrafos
Suite5 Data Intelligence Solutions, Limassol, Cyprus
e-mail: fenareti@suite5.eu; evmorfia@suite5.eu; erifili@suite5.eu; andreas@suite5.eu; dimitris.syrrafos@suite5.eu

E. Bosani
Whirlpool Management EMEA, Milan, Italy
e-mail: enrica_bosani@whirlpool.com

M. Calabresi · M. Sesana · V. Antonello
TXT e-solutions SpA, Milan, Italy
e-mail: mattia.calabresi@txtgroup.com; michele.sesana@txtgroup.com; veronica.antonello@txtgroup.com

A. Capaccioli
Deep Blue, Rome, Italy
e-mail: andrea.capaccioli@dblue.it

© The Author(s) 2024

J. Soldatos (ed.), *Artificial Intelligence in Manufacturing*,
https://doi.org/10.1007/978-3-031-46452-2_25

and sometimes also by the technical experts, has been demonstrated to be a long path. The successful navigation of this path requires the adoption of structured mapping techniques and methods, along with an agile approach to the development of the solution that ensures a smooth, step-by-step, progress.

The continuous and strict collaboration between technical experts and business experts has also been a key success factor for achieving the final goal in the manufacturing environment. This is a success factor not only for AI adoption but also for the successful deployment of most of the innovative technologies offered by I4.0.

This chapter explores and documents the experience gained by Whirlpool in the XMANAI project [1], where the explainability of AI technology is applied in a sales demand forecasting scenario. In this scenario, the level of adoption of the AI tool is highly dependent on the trust that can be generated in the users. In the Whirlpool's case, the possibility to explain AI results to users has been identified as one of the key enablers to gain users' trust and to fully engage human stakeholders within the AI loop. XAI is the door to open toward full awareness on why a specific result has been generated. The possibility to achieve a deeper understanding of the processes simulated by the AI rises the awareness and the belief in XAI: It is the spark to achieve better decisions and results in daily business management.

In this context, this chapter starts with a detailed description of the Whirlpool's use case in the H2020 XMANAI project in Sect. 2. It presents the motivation driving the project and the business context of the use cases. Then, the chapter describes the current state and the identified business requirements driving the XMANAI solution design, providing an outline of the "to-be" scenario and the key objectives to be achieved in Sect. 3.

Accordingly, Sect. 4 presents the technical implementation journey, followed by the explainability value presentation. The process used to detail the explainability requirements with the users is outlined and the meaningfulness of the explanation mode, deployed into the system, is justified in Sect. 5.

Moreover, Sect. 6 includes the description of the XMANAI platform, showing details of the users' journey in the XAI platform and in the XAI Manufacturing application of Whirlpool's use case. Finally, Sect. 7 presents the evaluation results of the demonstration sessions held with the users, along with a summary of the lessons learnt so far in Sect. 8.

2 Whirlpool as XMANAI Demonstrator

Whirlpool is the biggest player in white goods business at global level and one of the most important players in Europe, where it counts 9 industrial sites in 5 countries and more than 50 OEM (Original Equipment Manufacturers) producing and delivering in the 35 European markets and more than 100 other destinations in the world.

Overall, the white goods business is characterized by high levels of competition in the European market. Several very aggressive global competitors play a big role in the traditional B2B (Business to Business) market, where currently the key success factors are price and brand reputation. In this context, all the players started, some years ago, to enrich their product offers with services to final customers. Most of these services were partially driven by product IoT (Internet of Things) functionalities and partially leveraged on post-sales organization. Nevertheless, most of these offerings had no real and disruptive effect on business setup and market share footprint.

Today, Whirlpool's products are mainly considered a commodity. Hence, due to the maturity of the market, the overall business setup had been quite stable until 2020. However, the pandemic event of COVID19 has significantly modified such a market setup through the competitive advantage of companies, like Whirlpool, which early addressed the safety constraints for contagion avoidance and quickly restarted the production flow after the lockdowns. In addition, the pandemic lockdown has provided a unique opportunity, for white goods producers, to finally exploit the direct sales to final customers (D2C), leveraging the functionalities of the web and by-passing the B2B selling constraints. This element is opening up new opportunities for market share acquisition, and the successful competition on this new channel is expected to be crucial in the future. This new business channel is characterized by a set of specific and challenging requirements, which are mainly related to the speed of the buying experience and the hard competition on the web based on product offers and brand reputation. Specifically, key success factors include the range of the available offerings, the speed of order-to-delivery, the pricing policy with focused and personalized promotions, as well as additional services that can be offered (e.g., free installation, guarantee extension, home delivery, old product scrapping, and waste disposal).

In this context, the complexity of the overall process and the speed required in decision making can be barely addressed by humans, even if experts, with traditional analytic methods. Thus, AI becomes the key enabler for a significant improvement in decision-making process, ensuring a reliable forecasting service which may support people in "acting."

It needs to be noted that Whirlpool's experience in applying AI technology in the manufacturing environment started more than 10 years ago, mostly with applications to quality control (like vision systems or product testing) and in predictive maintenance. After every successful project implementation, the company is faced with the challenge of sustaining the implemented solutions in the mid/long term and of exploiting the successful pilot application in other areas of the factories.

Some key factors have been identified as root cause of this effect: firstly, the users' difficulty to understand AI technology fully and deeply. The awareness of AI's potential and limitations prevents a dangerous misalignment between expectations and results. This weak awareness was mainly due to the poor control and understanding of the data fed to the AI, implying a progressive and unexpected derail in the quality of AI results, which often end in systems shutdowns. In these cases, the original trust of the users in AI system was progressively consumed, sometimes even resulting in rejection or boycott. Moreover, in the case of predictive analytics applications, the user's trust was negatively affected by a weak understanding of the AI results, mainly when they were far from their usual experience or expertise. In these cases, the users' approach was to tentatively reject the recommendation, i.e., to bypass the AI system. In both cases, the capability to create and sustain the trust in AI technology has been demonstrated as a key factor to really make the most out of this approach.

In this context, eXplainable AI (XAI) may represent the key to unlock the achievement of all the AI solution objectives, as it may enable the users to deep dive into AI results with a language that makes them understandable (i.e., in terms of *how* and *why* such results are generated) and, therefore, actionable and sustainable in the long term. Thus, Whirlpool has joined the XMANAI project experience with a use case that is quite far from the traditional manufacturing domain, yet which can be used to demonstrate, in terms of business impact, the full potential of this technology.

3 White Appliances Use Case Description

Whirlpool, as other competitors in these years, has launched the D2C (Direct To Consumer) channels in the biggest European markets ensuring a wide product offer among its several important European brands (e.g., Whirlpool, KitchenAid, Hotpoint, Ariston, Ignis, Indesit) with a promised order-to-delivery time of 3 working days.

With a highly complex manufacturing footprint, composed by mono-product factories which produce one single product platform distributed in all the destination markets, the supply network imposes a transportation time, and consequent order-to-delivery time, that cannot fulfill the request of the D2C business. The traditional approach is based on inventory strategy, focused on ensuring a certain safety level of stock to be able to serve any customer, at any place, for any product, and at any moment. Due to the need to extend as much as possible the product range offer to the customer, the risk of obsolescence and the high blocked working capital must be balanced with the customer.

A reliable demand forecasting may minimize the protected inventory, maximizing the possibility for customers to find exactly what they are looking for. Thus, a reliable forecasting functionality represents the first business requirement posted to XMANAI. As explained earlier, this is not enough to ensure the full and conscious adoption of an AI solution. Here is where the explainability functionality has to enable the "actionability" of the AI results by answering users' questions about how and why an outcome is generated.

The second, perhaps more important, business requirement which XMANAI project needs to address is a deep understanding of the business process. This is based on the development of a tool capable of supporting users with a clear and user-friendly visualization of how the forecast is generated and why. This enables them to understand what they can do to change (when possible) the forecasting results toward achieving their business goals.

In summary, the Whirlpool use case in the H2020 XMANAI project is focused on the creation of a reliable, explainable, and actionable demand forecasting tool capable of providing the following:

1. Demand forecast reliability, including:
   - Getting a reliable demand forecasting
   - Minimizing inventory for D2C
   - Maximizing product availability on request
   - Maximizing customer satisfaction
2. Business dynamics understanding, including:
   - Understanding of demand evolution
   - Understanding of customers' behavior
   - Understanding of buying patterns
   - Supporting of promotional initiatives
   - Supporting simulations (i.e., execution of "what-if" scenarios)

Currently, the demand forecasting of D2C market is embedded into the full ODP process (Operational Demand Plan), which funnels all the expected sales demand of all the markets in a unique master production plan for all the factories and OEM sources. This process is coordinated by a central Demand Planning team, which enriches the unconstrained forecast generated by statistical analytics on historical data, with adjustments driven by factory capacity planning, supply base constraints, inventory and transportation strategy guidelines. They work starting from each single market demand profile to generate a total demand, which is then split to supply each single market. At this moment, the marketing and sales team of each market analyzes the forecast to take decisions about pricing strategy, promotional actions, and product range offer. The result of this second enrichment is the input for the manufacturing production plan.

In the "to-be" situation, as depicted in Fig. 1, the XMANAI platform will be used to support the decisions made by these two groups of users to achieve better results, using the reliability of AI functionality and the business dynamic understanding of XAI.

It needs to be noted that the sustainable usage of AI and XAI pipelines generates the need to include an IT role with the specific responsibility of managing the XMANAI platform daily. Data scientists and data engineers become key organizational roles to ensure that the XMANAI platform will perform according to high standard levels, granting the optimal forecast accuracy and ensuring the models' and pipelines' maintenance during the whole system lifecycle. This element is new in the organizational structure of a traditional industrial company, which, in most cases, does not have internally the competence to effectively support the business users. In this case, the answer provided by XMANAI platform highlights this organizational gap and the strategy that must be adopted to fill it.

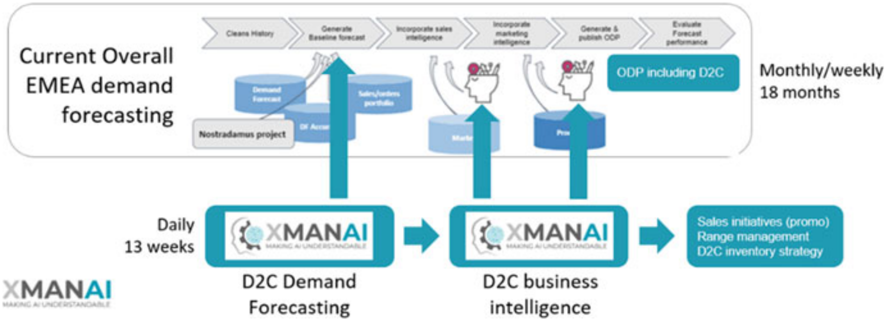

Fig. 1 XMANAI and Whirlpool ODP management process

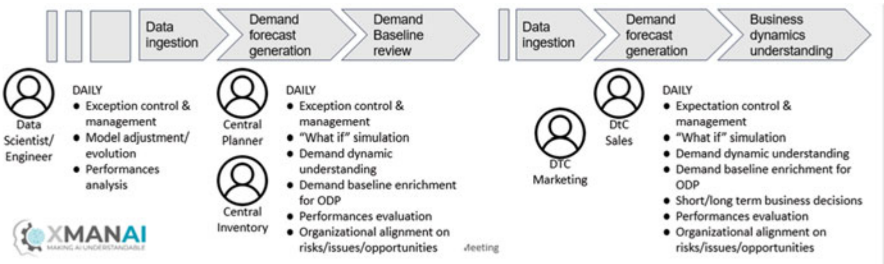

Fig. 2 Users' scenarios in Whirlpool's use case

The so-defined scenarios presented in Fig. 2 have facilitated the identification of the user's stories, which generated the specifications for the explainability requirements, as described in the following chapters, and set the foundation for demonstrator validation.

Last but not least, taking into consideration the fact that the XMANAI platform has the credentials to eventually become a very strategic business tool that can set Whirlpool apart from its competitors, it must guarantee the security and protection of the strictly confidential and sensitive data that are used. Thus, a security-by-design approach and a full GDPR (General Data Protection Regulation) compliance must be embedded into the solution development since the beginning.

4 Explainable AI Approach

The technical activities performed and the achievements reached during the preparation and execution of Whirlpool demonstrator span across different work areas, including data exploration, predictive model development and performance improvement, development of explainability methods, adaptation of the results for the needs of business users, as well as development and validation of the initial release of the corresponding XMANAI manufacturing app (i.e., Demand Forecasting Manufacturing App).

The main achievements across these axes are as follows: (a) use case detailed analysis and elaboration presented in Sect. 4.1; (b) data acquisition and exploration as explained in Sect. 4.2; (c) development and validation of hybrid XAI models outlined in Sect. 4.3; and (d) design and development of the XMANAI manufacturing app (i.e., the Demand Forecasting Manufacturing App) as elaborated in Sect. 4.4.

4.1 In-Depth Analysis

A close collaboration and discussions between the technical partners and the business users needed to occur in order to ensure that the business problem, the "as-is" situation, and the business requirements were explained in depth and that all partners acquired a common understanding in order to properly address the challenges. Various user stories were formulated based on the business needs, including the generation of an accurate demand forecast per each single product in specified time horizons and steps, in an explainable way for different target audiences. The AI needs and the data requirements were identified and discussed, along with the different technical and business aspects that needed to be aligned. The main identified AI and explainability needs were the following:

- AI requirements, including the extraction of D2C profiles and the improvement of forecasting reliability.
- General explainability requirements, including the extraction of insights about demand trends to optimize supply flow, as well as the identification of influential factors for demand and of critical situations that need to be handled to avoid stock breakage.

4.2 Data Acquisition and Exploration

The data employed for the forecasting task were acquired mainly by Google Analytics datasets, which track information from the Whirlpool's website. The data include information regarding product sales and customer-related data and are accessed through frequently updated database tables that are extracted as CSV (Comma Separated Values) files. Google Analytics generates data related to various fields that are not necessarily of use for the task at hand. Hence, there was a need to identify and isolate only a valuable subset of information that could aid the sales forecasting task and provide useful insights for the business users. Based on the Google Analytics documentation, some initial data investigation of the provided fields and various statistics, it was possible to select a subset of approximately 30 fields that were considered the most relevant. Accordingly, data from other database tables, containing Whirlpools internal product information, were also considered and added as needed to the previous extracted subset.

The next phase of data manipulation included data cleaning, identification of irregularities, and exploratory data analysis to detect interesting properties and correlations among the different data fields. The extracted results were discussed among the XMANAI partners, in order to get feedback from the business users regarding some of the findings, the assumptions, and the insights made by the technical users. Based on the findings of this procedure, the business partners agreed to provide additional data, such as product hierarchies/categories, product price, and campaigns.

Next, more features were extracted based on the fields of the selected subset including sales, visits, price, calendar information. Sales' lag features, average number of unique visits, future product price, day of the week, and month are some of the features that were created. An exploratory data analysis focused on the extracted features was performed to conclude on which of them might be the most promising for the forecasting model.

Finally, the datasets that were to be used for the next steps of the analysis were anonymized and uploaded (as CSV files) to the XMANAI platform. Note that certain extensions to the data model were required to ensure that the data uploaded are accompanied by proper semantics (i.e., explanations at the data level).

4.3 Development and Validation of Hybrid XAI Models

Once data exploration concluded and the Whirlpool data were available in the XMANAI platform, the implementation activities for an appropriate predictive model for the sales forecasting use case, along with a suitable explainability model, started.

During the initial experimentation, the focus was on the development of various forecasting models for the most sold individual products and product categories. Whirlpool's D2C channel is new and, therefore, during the first experimentation, Google Analytics contained sufficient information only for a subset of the company's products that permitted the implementation of high-performing predictive models. The examined time horizons were (a) 1 week, (b) 1 month, and (c) 3 months ahead, and the models were implemented both for weekly and daily predictions, whenever this was possible, since daily predictions for 3 months ahead were not feasible. The models were evaluated using appropriate evaluation metrics and the most promising results were obtained by boosting models and more specifically XGBoost.

As the implementation phase progressed over time, the focus moved toward improving the performance of the forecasting models. The selected horizons were fixed to 1 week and 1 month ahead with weekly steps. More data were available through Google Analytics, and the walk forward approach for training the XGBoost models was examined and found to improve the results. Additionally, different product hierarchies and categories were considered based on the input from the business users, and the hierarchical approach was employed to provide more coherent forecasts for the individual levels of the products' hierarchies.

Regarding the explainability aspect of the task, descriptive visualizations depicting correlations among the features and their influence on the target value were employed to provide initial explainability insights. After developing the predictive models, a range of explainability tools [2, 3] has been used, including SHAP, permutation importance, counterfactuals, and what-if scenarios, to generate additional explanation results. This allowed the provision of clear and meaningful insights to business users, while also offering them the flexibility to explore hypothetical scenarios as needed. The work performed included "offline" experimentation, as well as experimentation and configuration of the hybrid XAI models within XAI pipelines in the XMANAI platform.

4.4 Delivery of the XMANAI Demand Forecasting Manufacturing App

Through dedicated brainstorming sessions among the business users and the technical partners of the XMANAI project, the detailed design of mockups for the user interface of the manufacturing app has been defined. The identification of the right visualization tools has been supported by the QFD (Quality Function Deployment) methodology and led to the definition of the type of diagrams and the exact information to be displayed.

The subsequent development activities of the manufacturing app fulfilled the specified visualization requirements, through front-end and back-end development. In addition, they involved integration activities with the XMANAI Platform to retrieve necessary data, results, and explanations from pre-configured XAI pipelines (as described in Sect. 4.3) and display them in the dedicated dashboard presented in Fig. 3. Finally, they insured that only authorized partners could utilize the app through single sign-on functionalities with the XMANAI Platform.

5 Explainable AI Implications and Added Value

For the Whirlpool's use case, the purpose of the explainability activities is to address the needs of the business users who want to comprehend how product sales are influenced by multiple parameters. The concerned business users are the central demand planner, the D2C marketing and sales specialist, and the data scientist/engineer. Different explainability needs have been identified for each of the users. For example, the central demand planner is interested in understanding the causes of critical situations in forecasting process and the root causes that affect

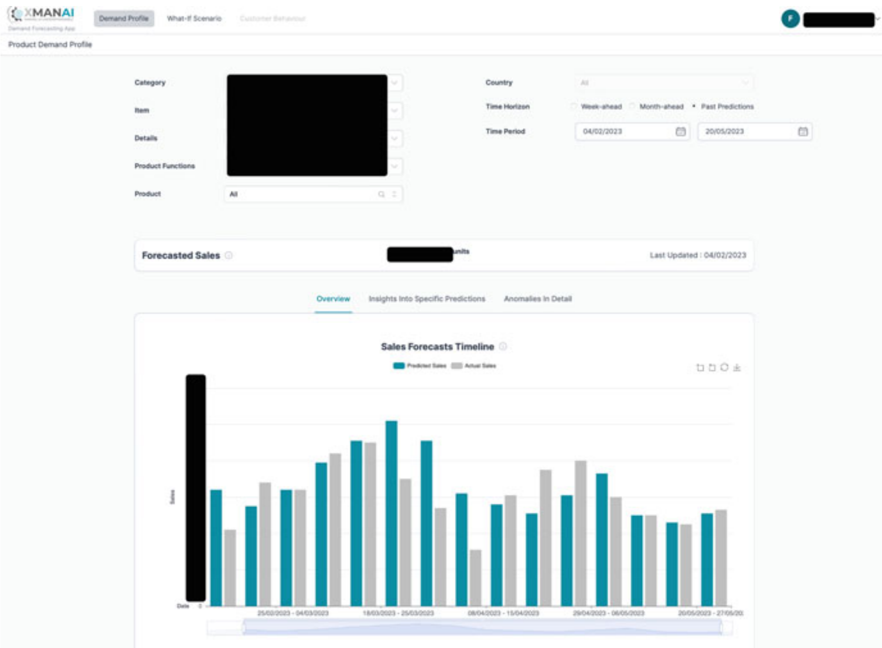

Fig. 3 XMANAI Demand Forecasting Manufacturing App

the forecast accuracy the most, visualizing forecasting plots and getting the trends and the directions based on the past, understanding customer behaviors and the impact of marketing strategies (campaigns, promotions), and finally being able to examine how the output would change based on different company's decisions. The D2C marketing and sales specialist is interested in understanding demand evolution, customers' buying patterns, and the effect of the marketing strategies.

The user's stories developed in the aforementioned activities and indicatively presented in Fig. 4 have been the key drivers to guide the finalization of the visualization tools for XAI deployment. The design methodology was extracted based on user's questionnaires, personas identification, and user's stories definition. All these were integrated into the user's journey description, which embeds the full user's experience within XMANAI platform and Manufacturing App.

The main explainability functionalities that have been identified for business users are the following:

- Demand forecasting visibility
- Demand root cause analysis
- Feature relations visualization
- Feature impact visualization
- Target scenario simulation
- Demand trends identification

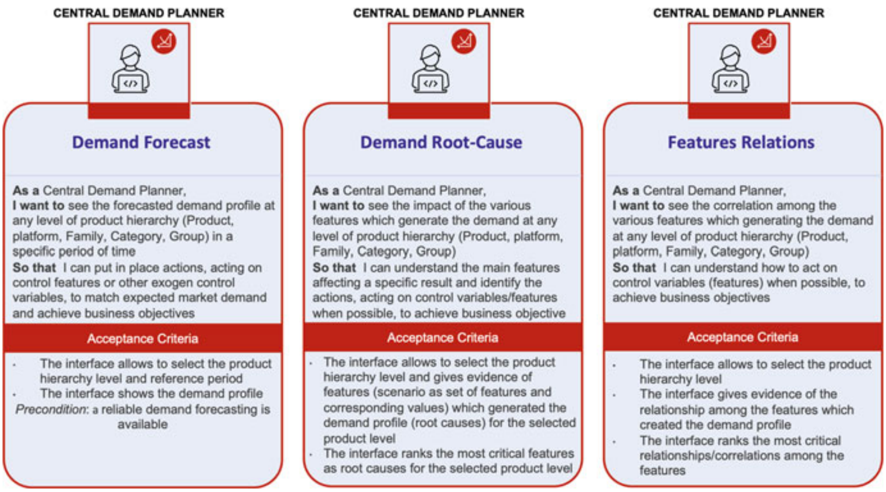

Fig. 4 User's stories example

- Demand trends root cause analysis
- Demand anomalies visualization
- Demand anomalies root cause analysis
- Forecast accuracy visualization
- Buying patterns identification
- Buying patterns root cause analysis
- Customer's behaviors visualization
- Customer's behavior root cause

In order to meet the above-mentioned business needs, the implementations of the explainability approaches were divided into two categories:

(a) Explanation at data level
(b) Explanations at instance and model levels

The primary focus was on the generation of appropriate explanations, followed by the selection and development of suitable visualizations that are more user-friendly and easily interpreted by the business users.

5.1 *Explanations at Data Level*

In the case of explanations at the data level, the aim has been to extract descriptive visualizations that provide insights in a comprehensive way for the business needs. At this stage, line plots and histograms were employed to explain demand evolution, historical and forecasting patterns, trend lines, and seasonality. Pie charts were employed to explain fluctuations in the forecasts' accuracy based on the inclusion/exclusion of individual features. Partial dependency plots were also examined to illustrate the influence of the features on the output. Finally, heat maps were used to depict the relationships among the features, the sales, and the buying patterns of the customers, as well as the sales with respect to different calendar information. Upon discussion with the business users, the technical partners included only the visualizations that were most appropriate for their needs.

5.2 Explanations at Instance and Model Levels

In the case of explanations at instance and model levels, the aim has been to carry out a post-hoc processing of the already trained models using SHAP and permutation importance explainability techniques. In addition, a tool for the creation of what-if scenarios has been provided to inspect the effect of features on hypothetical test cases. The instance-level explanations provide insights into the direction and the contribution of individual features to the model's output. The explanations at this level are significant to the business users as they are capable of understanding individual extreme situations and the factors that were the most influential. The model-level explanations provide a more holistic view of how the features affect the outputs, showing their relevance. Instance-level explanations are typically provided by SHAP force plots. The forecasted value of the model is explained by showing in a graph the feature contributions and the impact direction. Shapley values (average marginal contribution of a feature value over all possible coalitions) are calculated locally for specific instances and shown in a graph provided by the SHAP library (Fig. 5).

After discussing with the business users, a more self-explanatory graph than the default SHAP force plots was requested and designed, with a less data scientist–oriented approach: alluvial plots, as depicted in Fig. 6.

The possibility to generate and explore what-if scenarios is also provided to the business users, who can examine how the predicted value of an instance would change in response to a modification of the input features. Such scenarios (as indicatively presented in Fig. 7) allow the business users to understand how strong the correlation of the input features to the models output is, potentially even suggesting changes that could increase the company's sales.

The model-level explanations are indicatively provided by SHAP and permutation importance. The SHAP library provides and visualizes global explanations by creating a bees-warm plot and aggregating all local feature contributions, while permutation importance generates global feature contributions by computing the change in the forecasting error, as it randomly shuffles each input feature. The results are presented in a table with the most influential features and their weights. In the context of the first phase of the manufacturing app, global explanation results were decided to be depicted through bar plots, upon discussions with the business users (as indicatively shown in Fig. 8).

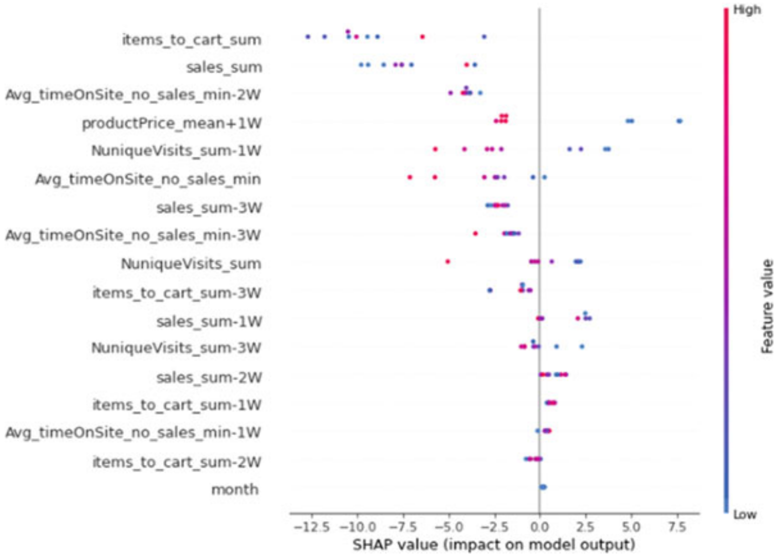

Fig. 5 Instance-level explanations through a SHAP force plots – data scientist–oriented

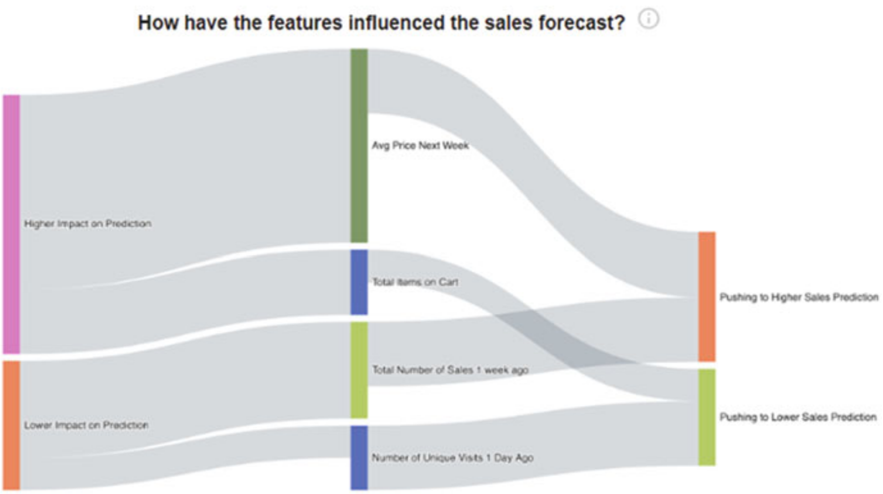

Fig. 6 Alluvial plot example presenting the feature contributions on the output at instance level – business user–oriented

Fig. 7 What-if scenario example. Impact of the product price to the products' sales

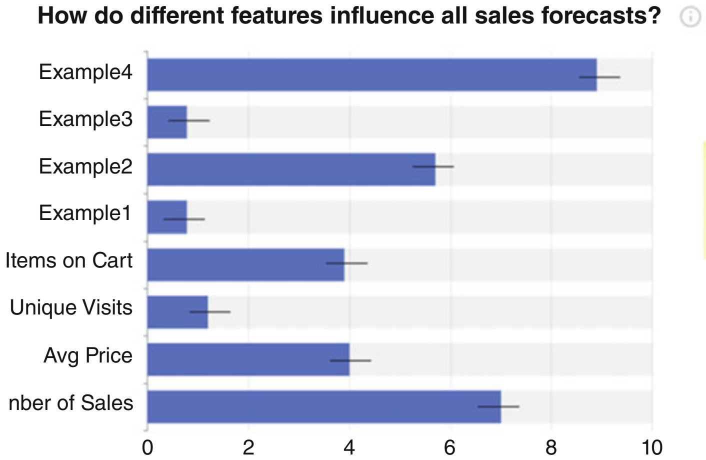

Fig. 8 Bar plot example presenting the feature contributions on the output at model level

6 Application of the XMANAI Platform and Manufacturing App

In order to deliver explainable AI results to the target end-users, Whirlpool has leveraged (a) the XMANAI Platform that provides a wide range of functionalities to address explainability from a data, model, and result perspective [4] (as presented in Sect. 6.1) and (b) a dedicated Manufacturing App that offers the user interface with appropriate visualizations specifically selected for demand analysis, what-if simulation, and customer behavior visualization (as outlined in Sect. 6.2).

6.1 XMANAI Platform

In brief, the main aim of the XMANAI XAI platform is to enable the efficient XAI pipelines lifecycle management and to ensure the capability of the user to maintain the best performance level. It's mainly dedicated to IT roles (data scientists and data

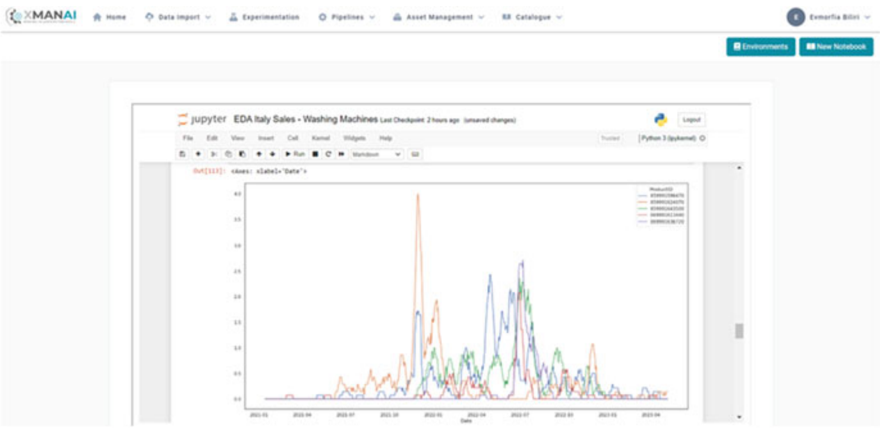

Fig. 9 XMANAI Platform – data exploration over a selected Whirlpool dataset

engineers) and business roles (demonstrator's business/technical users) and follows the specific users' journey identified during the previous design phases [5].

In particular, for the Whirlpool demonstrator, the relevant CSV files that have been extracted from the Whirlpool databases and systems (as explained in Sect. 4.2) have been uploaded in the XAI Platform and then mapped to their equivalent fields in XMANAI Data Model. Once the dataset has been ingested, the metadata have been provided by Whirlpool data scientists. As a last step, an appropriate access policy has been defined for the uploaded datasets by Whirlpool users, in order to ensure full data protection: a dedicated sharing contract signature management function has been created to ensure a secure access to allowed partners. Once the available datasets are available and access rights have been granted, the data exploration phase may start in the Interactive Exploration & Experimentation menu depicted in Fig. 9.

Based on the data exploration outcomes, different models have been trained in the XMANAI Platform and different XAI pipelines have been configured, utilizing both the data preparation functionalities and the ML/XAI functionalities. Once the pipelines were configured, the different execution logs were leveraged by the data scientists of the technical partner to check the progresses made (Fig. 10).

6.2 *XMANAI Manufacturing App*

The main aim of the XMANAI manufacturing app of Whirlpool is to offer business users the possibility to get access to the explainability dashboard, which is targeted to the "demand forecasting" problem and consists in 3 different web pages offering the functionalities of Demand profile, What-If Scenario and Customer's behavior that map to corresponding menu items.

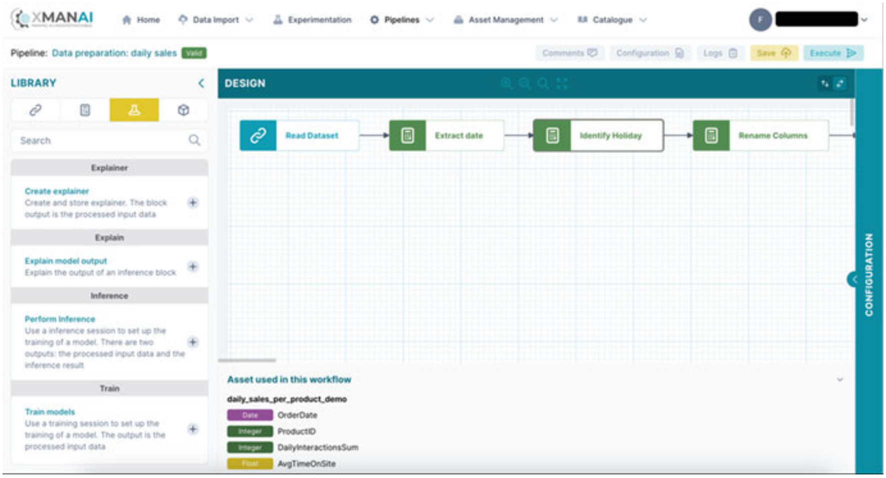

Fig. 10 XMANAI Platform – XAI Pipeline utilizing the Whirlpool data

The "Demand Profile" is dedicated to the demand forecast analysis results visualization for both central planning users and market users. It includes information regarding the sales predictions and the features' influence on the predictions. The information may be presented with a weekly scale in different relevant time horizons in the future (1 week, 1 month) and in the past. Here, the user may get important information not only on predicted sales value for specific products or product families but may also see which features have higher impact on the prediction value and prediction accuracy. This is relevant, as it allows them to understand which levers to use to drive the values. Also, they may get visibility of the correlations among features and awareness about the final result of a potential action. An extract of the "Demand Profile" page is presented in Fig. 11.

The "What-If Scenario" page provides insights into how the sales predictions change depending on the values of the input features, as well as on which features are to be changed (and how) in order to reach a predefined sales prediction. The user can also create new scenarios using a wizard, edit the configuration of the existing ones, view an existing scenario of their interest (as depicted in Fig. 12), or delete those which are no longer useful. All the created scenarios provide the possibility to see the "Demand Profile" diagrams. This function, dedicated to central planning users and market users, allows to see the effect on the output by modifying one or more features. Conversely, it also provides information about the specific values of the modifiable features required to achieve the desired forecast.

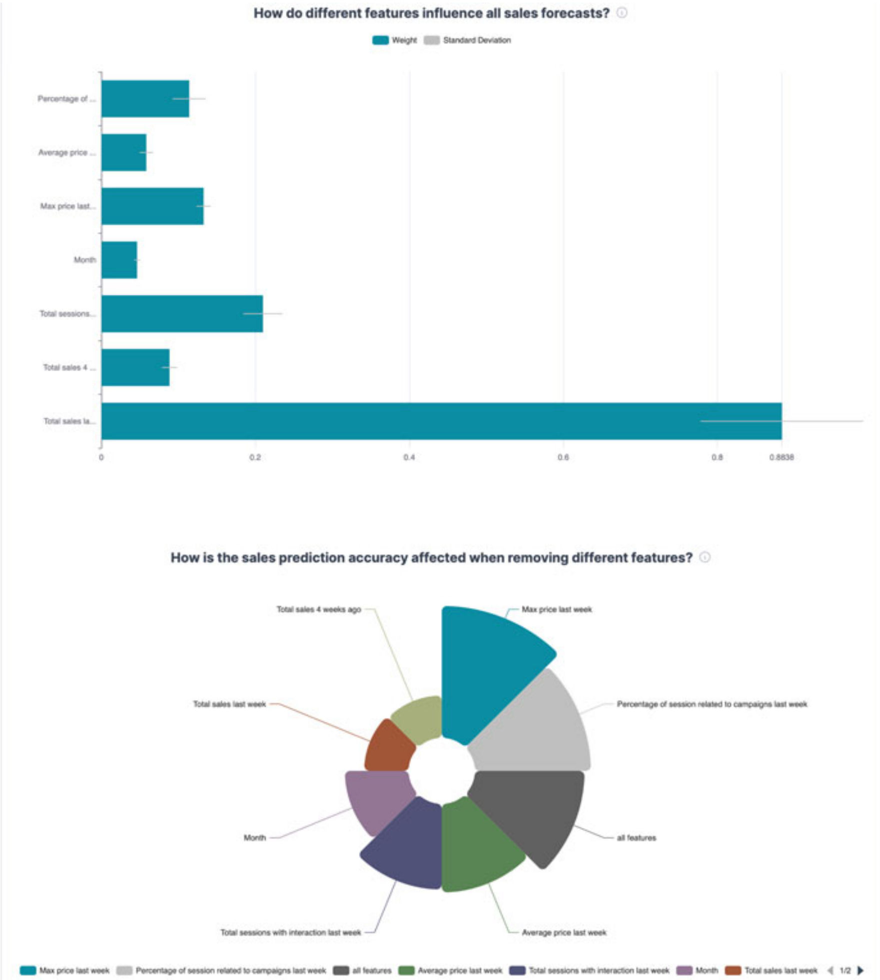

Fig. 11 XMANAI Manufacturing App view

7 Evaluation of the XMANAI Solution

The validity of the XMANAI solution and the status of achievement of the project objectives have been captured, after the completion of the demonstrator sessions held by the users, through questionnaires that identified key results from a user's perspective on different dimensions (in a scale 1–5, where 1 = not at all and 5 = completely).

As the main strategic objective of the XMANAI platform is to gain the users' trust through the explainability associated with artificial intelligence, the users'

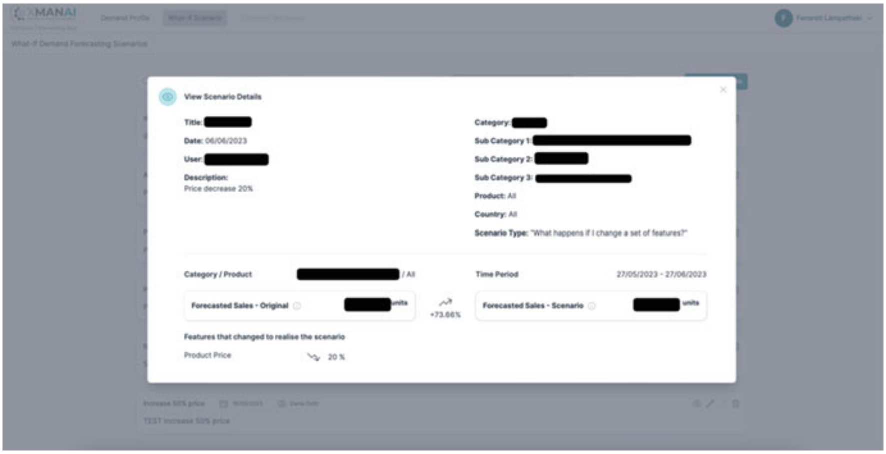

Fig. 12 View of a configured what-if scenario

Fig. 13 Users' XAI Manufacturing App evaluation questionnaire for Whirlpool's use case

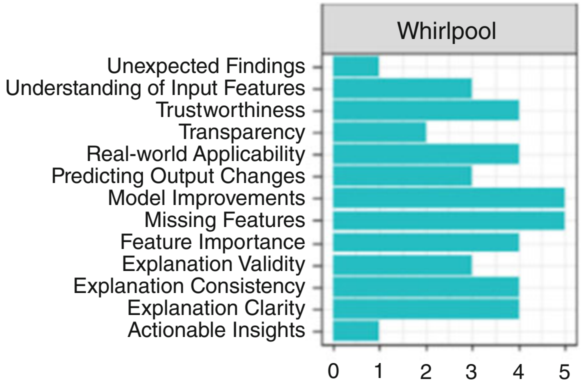

feedback has been particularly meaningful, even in a first release of the XMANAI results.

As depicted in detail in Fig. 13, the questionnaire results highlighted a strong improvement on the business impact driven by XAI solution, even if the alpha release that was assessed still needs to be significantly improved to achieve the expected excellence. The effort spent on explainability and the refinement path (which started from visualization diagrams fit for data scientists to a dashboard fit for business users) has been appreciated even if, at the moment, the expectations are not fully satisfied. The questionnaire focused on what is still missing and what can be done to fully reach the goals.

The business objectives that have been identified during the preliminary project phases and are mainly related to the reliability of the forecast and the explainability of the results have been summarized as follows:

Fig. 14 Whirlpool's use case KPIs

Demonstrator KPI	Meaning
DFE Monthly target (Lag1, Lag2)	Demand Forecast Error at SKU level 1/2 months before
DFE Weekly target (Lag 2,3,4)	Demand Forecast Error at SKU/week/market level 2/3/4 weeks before
ATP	Availability To Promise
Sales trend	Revenue's monthly variation
5STARS	Customer appreciation

– Optimization of the **order-to-delivery process**, maximizing the **customer satisfaction** (revenues, margins) with minimal **required resources** (inventory minimization, supply management just-in-time)
– Business **dynamics knowledge** acquisition for people empowerment in driving process

As a consequence, the corresponding business KPIs (Key Performance Indicators) set has been identified in Fig. 14 in order to capture the business impact of XMANAI usage.

In this initial demonstration phase, a detailed and solid measurement of the KPIs set has not been feasible, but the data gathering related to weekly DFE (Demand Forecast Error) generated by XMANAI platform vs the demand generated in the actual management process showed a significant improvement as shown in Fig. 15, notwithstanding an overall fluctuation effect.

To conclude, the users' feedback has been very positive, demonstrating a high confidence on the possibility of further improvement in reliability of prediction, completeness, and effectiveness of the visualization tools.

Finally, an evaluation tool based on 6P methodology [6] has been applied to the demonstrator session after the sessions with the end users. The gathered results have been compared with the same questionnaire submitted at the beginning of the project, in order to capture the "as-is" situation. The assessment has been focused on various aspects of AI and XAI adoption to measure the progress in the development journey during the project lifecycle. The specificity of Whirlpool's use case excluded the possibility to measure some of the dimensions, specific for manufacturing environment. However, the final result has provided clear evidence of the main gap in the dimensions of people readiness for AI and explainability technologies (as depicted in Fig. 16). The strategy to address this gap coverage

KPI	Base Level value (Without XMANAI)*	Measured value with first alpha release*	Expected Value
DFE weekly (Lag1) (reference at total IT market level)	w9: 67% w10: 61.2% w11: 68.2% w12: 81.9% w13: 82% w14: 83.4% w15: 71.1% w16: 76.3% w17: 74.8% w18: 66.3% w19: 71.6%	W9: 0% w10: -6% w11: 1.1% w12: 45.7% w13: 68.5% w14: 100% w15: -8.2% w16: -32% w17: -22.8% w18: 21.7% w19: 0%	<50%

Fig. 15 Whirlpool's use case – preliminary KPIs results for example product range

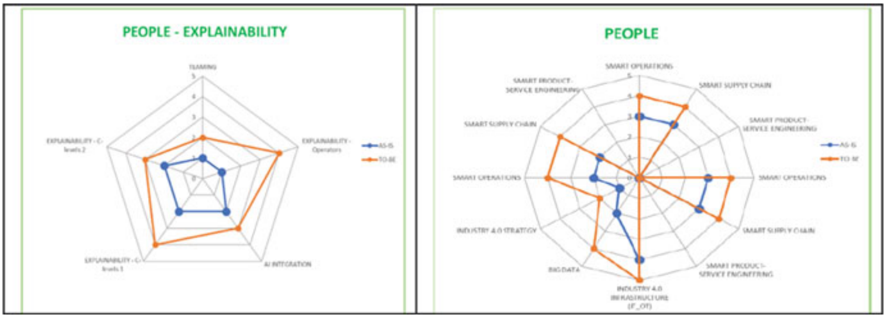

Fig. 16 Whirlpool's use case – 6P assessment result for the PEOPLE dimension after the initial demonstrator phase

represents one of the most interesting indirect results expected out of the XMANAI project experience.

8 Conclusions and Lessons Learnt

Following the completion of the first demonstration phase with the end users of the XMANAI solution, some key lessons learnt have been captured and are to be used for a more effective continuation both of the next project phases and of any XAI initiative:

- As with any AI initiative, the availability of appropriate data is instrumental. In the case of Whirlpool, the need to enrich the underlying data pipelines has been further highlighted (after the initial results) to improve accuracy and explainability of the demand forecasts.

- The end users have significant difficulty in vocalizing and describing their explainability requirements: a preliminary preparation of the target users is recommended in terms of change management to enable awareness of XAI earlier and then competence to get the most out of it.
- Predictions at multiple product hierarchy levels are not trivial and need to be reconciliated with appropriate techniques to ensure homogeneous results. Such an activity requires additional work on the XAI pipelines to ensure that all predictions across all product hierarchy levels (e.g., category, item, product function) are consistent, while always in consultation with the business experts to understand the business dynamics.
- The product demand forecasting problem has significant complexity since each product has its own peculiarities and one-size-fits-all XAI models do not work. Therefore, the selection of a single product hierarchy that brings maximum impact and development of targeted XAI models to improve accuracy is a most recommended approach (instead of trying to solve the full product demand forecasting problem at once).

Overall, from the business perspective, the support of a structured change management, which includes communication and focused training actions, seems to be the key success factor for the introduction of XMANAI platform into a business organization. Besides, from a technical point of view, the communication between data scientists and business experts represents some of the real challenges to be addressed and to which XMANAI has tried to contribute.

Acknowledgments The research leading to this work has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No: 957362.

References

1. Lampathaki, F., Agostinho, C., Glikman, Y., Sesana, M.: Moving from 'black box' to 'glass box' artificial intelligence in manufacturing with XMANAI. In: 2021 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC), Cardiff, United Kingdom, vol. 2021, pp. 1–6 (2021). https://doi.org/10.1109/ICE/ITMC52061.2021.9570236
2. Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions. In: Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17), pp. 4768–4777. Curran Associates Inc., Red Hook (2017)
3. Arrieta, A.B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R.: Explainable Artificial Intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI. Inf. Fusion. **58**, 82–115 (2020)
4. Miltiadou, D., Perakis, K., Sesana, M., Calabresi, M., Lampathaki, F., Biliri, E.: A novel explainable artificial intelligence and secure artificial intelligence asset sharing platform for the manufacturing industry. In: Proceedings of 2023 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC) (2023)

5. Branco, R., Agostinho, C., Gusmeroli, S., Lavasa, E., Dikopoulou, Z., Monzo, D., Lampathaki, F.: Explainable AI in manufacturing: an analysis of transparency and interpretability methods for the XMANAI platform. In: Proceedings of 2023 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC) (2023)
6. Spaltini, M., Acerbi, F., Pinzone, M., Gusmeroli, S., Taisch, M.: Defining the roadmap towards industry 4.0: the 6Ps maturity model for manufacturing SMEs. Procedia CIRP. **105**, 631–636 (2022)

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.