Solution guide

5 ways how AI can improve your demand forecasting

2021

deepsense.ai
BIG DATA SCIENCE

How can AI leverage the accuracy of your organization's demand forecasting solution?

The concept of demand forecasting is nothing new in the retail and manufacturing industries. Often performed manually by analysts with the use of ERP software, demand forecasting helps reduce out-of-stocks, increase the efficiency of production plans and improve customer satisfaction. However, with advancements in artificial intelligence, demand forecasting can be brought to the next level and will give your company an additional edge. Using our project experience and AI-expertise, here are five key ways that deepsense.ai can leverages the accuracy of your organization's demand forecasting process.

About deepsense.ai

We help companies gain competitive advantage by providing AI-based end-to-end solutions, with the main focus on computer vision, predictive analytics and natural language processing.

Our commitment and know-how have been appreciated by global clients including Nielsen, L'Oréal, Intel, Nvidia, The United Nations, BNP Paribas, Santander, and Hitachi. We also deliver machine learning and deep learning training programs to help companies build AI capabilities in-house.

For more information visit our website.

WAY 1
Feed AI algorithms with data sources including sales data, product characteristics and marketing campaign plans

Traditional demand forecasting systems typically rely on the extrapolation of historical sales information, taking into account seasonal variations or the overlay of strategic initiatives (such as that the company will push product X next month). Unlike humans, AI can manually aggregate numerous data sources and produce meaningful results. In a project we did for an FMCG retailer that carried out a lot of promotional activities, we leveraged external databases and marketing data, combining that information with historical data to produce a very precise forecast.

Combining data from various sources

The power of AI lies in its ability to process large amounts of information, which in turn enables us to use multiple sources of data to feed our models. In the case of demand forecasting, the most obvious data to use is historical sales, i.e. the amount of items sold daily and their prices. But forecasting demand is a multi-dimensional problem and a number of important factors influence demand.

Where to look for more data?

The darkest place is under the candlestick, isn't it? Look at your own internal data. It contains sales data, product characteristics, metadata, data on promotions and your marketing activities. External sources of data are also important, and include sales data from your distributors or even from points of sales or market reports provided by companies like Nielsen. These are potentially excellent data sources. A third group to look at is contextual data, such as will emerge from a calendar of

holidays, demographic or geographic data. It's all there, you just have to turn over the right stones.

Data aggregation engine

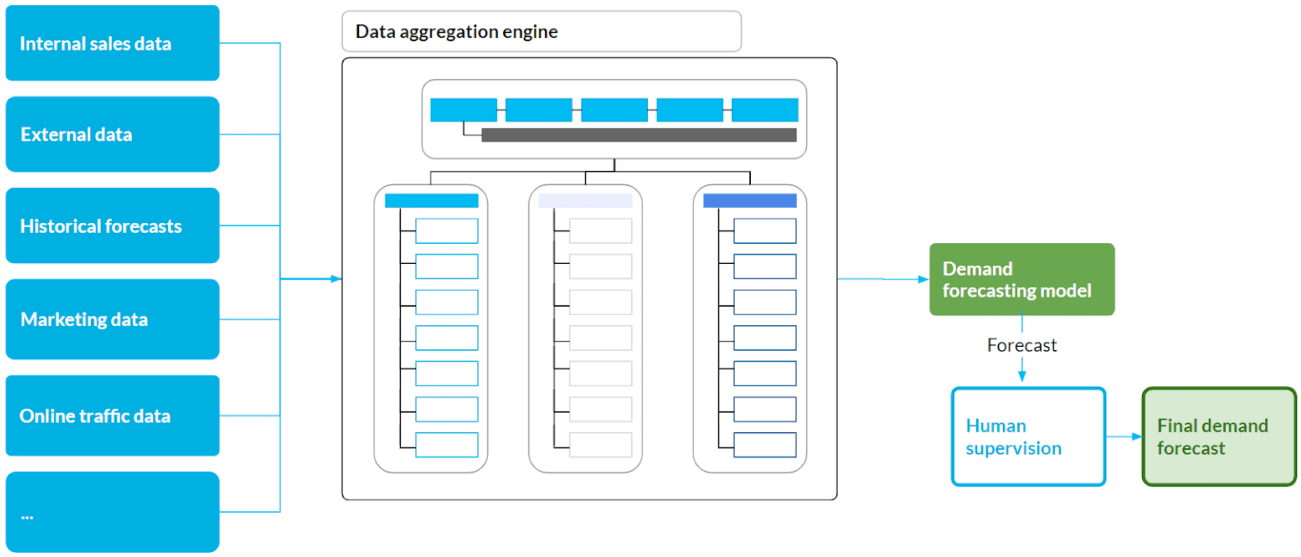

Internal data as low-hanging fruit

Probably the most valuable internal data in demand forecasting comes from promotions and marketing. You plan marketing activities and promotions, so the data is there, it's just a matter of plugging suitable data sources into a demand forecasting system. It's intuitive that low prices will increase demand, but by how much? You can figure that out by either analyzing historical data or, more effectively, feeding a demand forecasting system with promotions and marketing data to do the job for you. Is lowering the price below a specific level an unwise step, if demand and revenue aren't going to increase as a result? To find out, you can ask your demand forecasting model – just run a prediction for a hypothetical promotion level and see if your demand, or even your revenue, increases. That's right, a demand forecasting system can also act as a price elasticity tool!

Forecast features

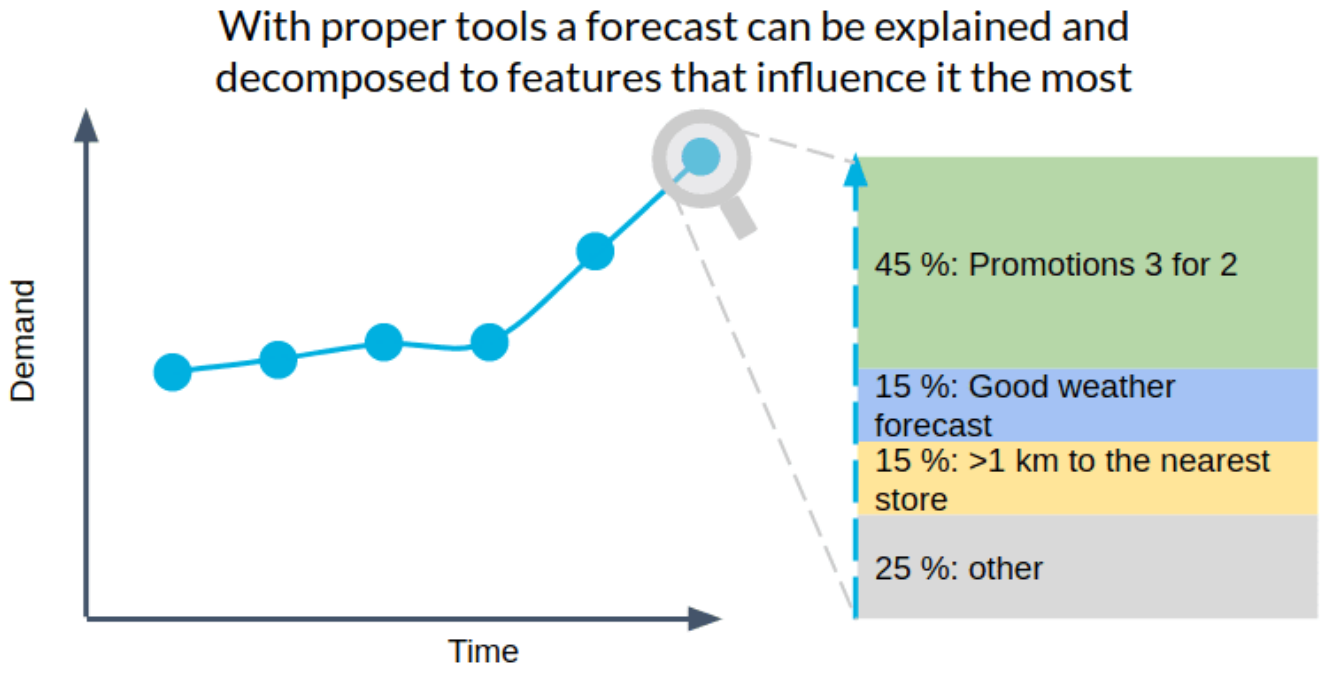

Plugging product characteristics and metadata into the system will enable it to draw on historical data on similar products while building a forecast. If you're selling online, tracking the traffic on your website is also a valuable step, as doing so may allow you to uncover trends in your demand time series.

Which external data do you really need?

If you're going to use external data, choose it wisely, as you'll have to pay for it. Some data types are virtually always helpful, like sales data from distributors or points of sales. Sometimes increased demand from a distributor can influence demand from others, and it may not be easy to spot without an automated AI system. If you sell through distributors, such a system is one way you can use the data at your fingertips.

Almost every business depends to some degree on the weather. While a daily weather forecast doesn't dictate total demand for warm shoes throughout the winter, it does influence daily demand for water in summer, especially in tourist

areas. Weather data is almost always invaluable in demand forecasting systems. But it also requires caution: because the weather changes daily, it can't actually be used to train models. You want to use a weather forecast for the day you're forecasting demand.

Another factor burdening certain industries is the level of COVID-related restrictions and the behavioural changes customer groups have exhibited. This crucial source of uncertainty can be to some extent mitigated by treating it like weather information – that is, plugging it into a forecasting model as an additional signal.

Demand is always context-specific

Demand is context-dependent for virtually every business. The same FMCG store can sell a very different amount of goods on the day before a "big" holiday and on a typical Tuesday. Although some conventional forecasting algorithms can handle changes in demand due to weekly seasonality, they can't take into account the day before a "big" holiday. AI models handle features like "is today a holiday?" or "number of days to the next big holiday?" - and improve forecasting.

Justwhat is a "big" holiday, though? That's very society-specific. Demographic data helps define more contexts in which sales take place. Using it can obviously help you decide which holidays are "big" and which are not. But demographic data can give you far more than that. If you run a brick-and-mortar store, knowing the local population's wealth can radically improve your forecasting in specific product categories – surely demand for top-shelf products will be much higher in wealthy areas. Geo data, meanwhile, allows you to use parameters like "number of similar stores in the neighborhood" or "distance to the nearest shopping mall." Such data can tell you a lot about potential local demand, and also help you uncover patterns and dependencies you wouldn't expect to be there.

Do it the right way

The more data sources being used, the more intelligently they must be connected. Data sometimes requires a lot of preprocessing (because data is missing or the time resolution presents problems. For example, hourly weather forecasts will not be suitable for a daily demand forecast unless they are properly preprocessed). Data can also be contaminated – there is a sudden change in units used or a lack of data, which manifests itself in strange patterns or constant values. Deciding which values you can actually trust is no trivial matter. deepsense.ai's experience has enabled us to produce an effective strategy for using numerous data sources and selecting only the most important.

WAY 2
Use AI to predict demand for new items not yet present on the market

While not relevant for all industries, this feature is crucial for some. Take the apparel sector, for instance. In a project for an e-commerce fashion retailer we used information about garment characteristics including cut, color, pattern, texture and material, as well as photos of garments. Finding similar items that had been sold previously helped us assess how nuances (e.g., fabric) can impact sales. Our demand forecasting models then compared these characteristics to the attributes of previously sold items. We found that enabling customers to browse products that are not yet on sale is a simple yet powerful way to predict demand even more accurately, before products hit the shelves.

Demand for new items not yet present on the market

In its simplest form, demand forecasting is about predicting future sales based on historical sales of a given item. The more often you introduce new products to the market, the more crucial demand forecasting will be for new products. But for new products, there's no historical data. Fortunately, new products can and do behave similarly to comparable products that have been sold. Our team developed a demand forecasting system for an apparel retailer for products the company had not yet made available for sale. It's fast fashion scheme was at the heart of its business, and making better predictions about what might soon be a hit was crucial.

Extracting product characteristics

To find the similarities between products, metadata is essential. With apparel, such data include an item's color, pattern and cut, while in the food industry ingredients, brand, origin, flavour, package size and weight are key data. A picture of the item is also important, particularly in the apparel industry.

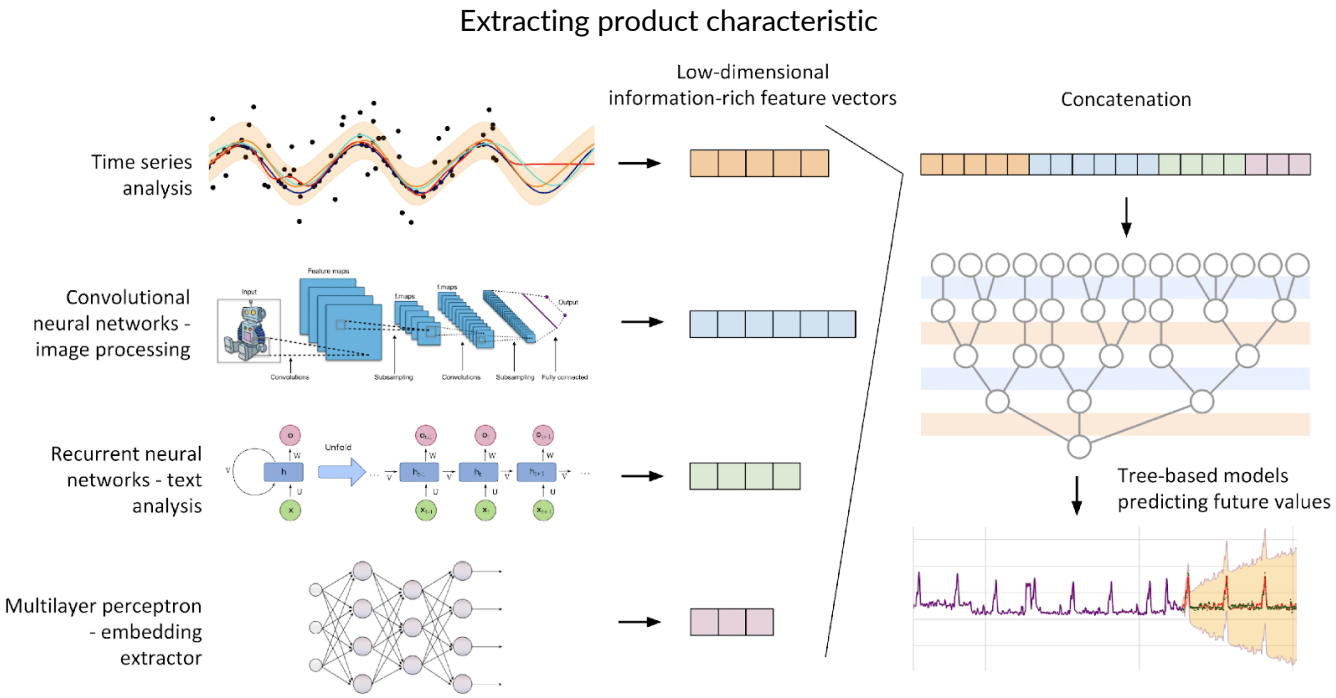

So which algorithms are used to work through all the data and images? It all depends on the model you're using. To predict a single value or to produce a single forecast for an initial sales period, the most powerful and popular algorithms are tree-based models. It's of course possible to use other models, like MLP or even simple linear regression models are also useful.

The data scientists reading this guide will know that it's pointless to pass all pixel values of an image to a tree-based model. One solution is to feed it to a convolutional neural network like ResNet and extract a feature map from the last convolutional layer. This provides a low-dimensional information-rich

representation of the image, which can be passed to the tree-based model. Metadata, which generally takes the form of categorical or numerical variables, can be easily passed to a tree-based model.

Product DNA

That's all very helpful, but we can do even better. We don't actually know whether vectors entering the model are similar for similar products, but it's good when they are. Why? Because we can anticipate that demand for similar items should also be similar, so the input vectors to the forecasting model should also be close to each other. Fortunately, there are a few ways to obtain such vectors, or embeddings, as they're known to data scientists – which will constitute specific product DNAs, and intelligently encode the most important product features.

One way to do that is by leveraging neural networks and sales information. The core of this approach is to train a neural network to predict which items are often bought together. From that network we can extract an encoded representation of the whole product. That is our product DNA. If two products are likely to be bought together, the dot product of their product DNAs will be close to 1, which means they're similar.

Product DNA

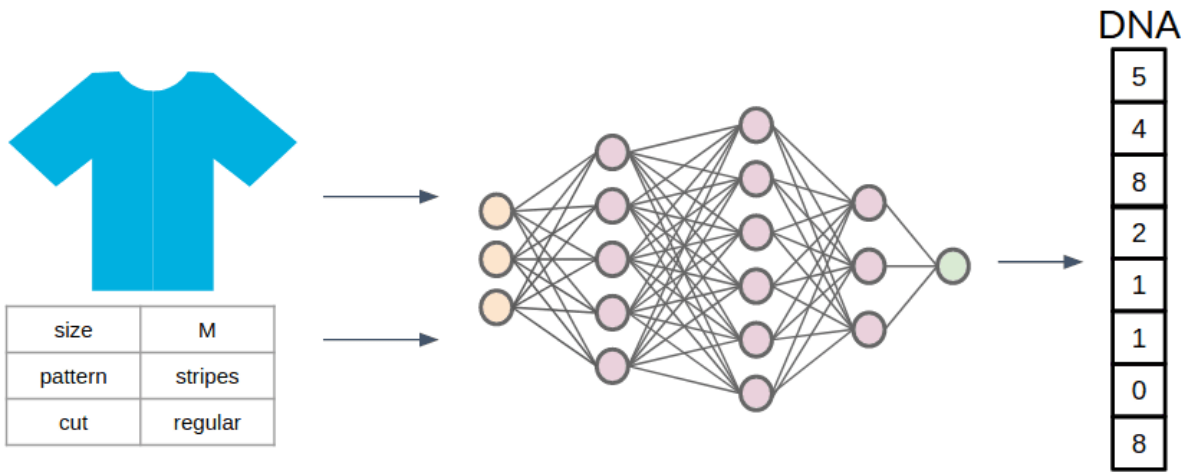

Another method to obtain product DNAs is triplet loss for neural networks, designed directly to produce an encoded representation of the input. Triplet loss directly forces the networks to produce similar vectors for similar items, and vice versa, without the use of sales data. It's based only on the product's characteristics. The first approach is more valuable when the product DNA is supposed to reflect an intrinsic customer taste, while the second one is more product-oriented.
Embedding product characteristics in a single well-behaving vector (that is, one that has similar items with similar encodings) is a fascinating yet powerful method to enhance demand forecasting models.

Use the help of your customers

There is also a very powerful feature to use when products have not yet been put on sale. Companies can place a button on their websites customers click to request a notification when a product becomes available. The number of clicks and their distribution in time yield up essential information about demand. Do exercise caution with this feature. However website layout and where the button is placed shouldn't change much. Otherwise, this could change the meaning of the feature over time.

WAY 3

Use AI demand forecasting models to predict the impact of promotions and marketing activities

An AI demand forecast model can do more than predict demand. It can also serve as a price elasticity tool and test the effects of price changes on demand. By analyzing historical data, our AI-driven demand forecasting model can break down the demand into a latent “real” component and the “promotional” effect. This knowledge is of paramount importance in the grocery and electronics industries, which rely on frequent product promotions. With our AI-driven model, retailers can test how a given promotion will affect sales and hence pick the optimal price point.

WAY 4
Track anomalies in demand levels

Demand forecasting systems can also help monitor current activities, not just forecast future ones. By integrating multiple data streams such as real-time activity on your website, we offer advanced anomaly detection. Anomalies can be caused by technical glitches (which must be fixed quickly) or a natural spike in demand. Being aware of when they may happen allows you to start fixing the problem immediately and to analyzing the root cause of the anomaly. If orders for face masks suddenly spike, and you have that information, you might want to place a substantial order with your distributor before it dawns on others to do the same, but for a higher price. Alternatively, with an advanced anomaly detection system, you should be able to avoid financial losses like those Foreo incurred when it sold ~40,000 devices for USD 9,99 instead of USD 279, before the technical glitch behind the oversight was spotted.

WAY 5
Use demand forecasting as your organization's first AI use case

Demand forecasting is all about precision: on one hand you don't want to overstock and freeze up cash in excess inventory, but neither can you afford to miss sales opportunities. According to McKinsey, AI can reduce forecasting errors by over 20%. That's a claim deepsense.ai can vouch for, as we delivered even more value in one of our projects. Working with a large CEE ecommerce retailer, we were able to reduce out-of-stocks by 30%. Depending on how your forecasting system is currently set-up, the gains could be even higher.

The benefits don't stop there. From the very first day you start using your demand forecasting model, your organization can track the improvement in key KPIs, including out-of-stock and inventory levels and the effectiveness of the promotions you run. The model can also measure "soft benefits" including fewer man-hours spent on forecasting analysis. These clear benefits stand in contrast to those offered by other recommendation engines, where attributing the real impact of an AI solution to the overall sales growth is challenging. With this in mind, demand forecasting is a perfect example for your organization's first AI use case. With its clear and substantial impact, you can convince stakeholders to invest in further AI-related projects.

deepsense.ai
BIG DATA SCIENCE

WAY 6 – BONUS

Contact deepsense.ai to tap our extensive expertise in demand forecasting

Because AI models are usually tailored to solve some of the biggest challenges in business, an increasing number of companies are recruiting data scientists and building their own AI teams. This can invite thornier problems than you might imagine. Finding a reliable partner that will share its knowledge and support development of an effective AI team may be the way to outsmart the competition and address the myriad challenges companies face today. If you would like to learn more about how deepsense.ai can help you do just that, and leverage the accuracy of your organization's demand forecasting process, get in touch.

Contact us and we'll give you a free assessment and a quotation

sales@deepsense.ai