CAPCO

APPLIED DATA SCIENCE IN FINANCIAL SERVICES

ARTICLE 1:

FORECASTING CAPACITY REQUIREMENTS IN OPERATIONS

a wipro company

Capco has combined intelligent heuristics and machine learning to accurately forecast resource utilization in manually intensive functions.

The financial services industry has witnessed considerable hype around data science and machine learning in recent years. However, a quick Google search will confirm that there are very few concrete examples of it being put in practice in large institutions and delivering tangible results.

We strongly believe that data science can add significant value in financial services across multiple functions with high returns on investment. Our content series, 'Applied Data Science in Financial Services', aims to highlight the common yet painful problems which Capco has solved using advanced analytics techniques.

Our first article focuses on an age-old problem: capacity management. Banks with global operational teams still find it difficult to forecast how much time and resources are required to perform required operational tasks, from back-office trade functions to KYC remediation. In the end, the result is often the same: mad scrambles and massive overtime fees to vendors as deadlines loom and patchy temporary solutions to appease regulators.

We have helped multiple clients solve this problem successfully by applying modern data analytics techniques, reaping significant benefits in short spaces of time.

Here's our story.

THE PROBLEM

Reliably and effectively forecasting capacity is critical to successfully running an operational team. Operational leads attest that it is nearly impossible to run a large-scale global service without allocating too few resources to it, or alternatively wasting too many resources on it. Getting it just right has been an elusive goal.

Banks have thus far resorted solely to capacity planning tools that are often static, localized and involve no measure of the expected result. They resort to the law of averages versus utilizing processing times broken down by task, user, location, system and other parameters. Using such solutions restricts businesses' ability to estimate how much capacity is required to complete a task, let alone modulating capacity changes in real time.

Sub-optimal resourcing has other detriments too. Take the example of back-office operations in a capital markets organization, where incorrectly estimating resources can lead to:

- Missed service level agreements (SLAs), resulting in financial loss as well as reputational damage
- Higher cost per trade due to inefficient resourcing
- Higher rate of human error, or financial losses caused by under-resourced teams
- Regulatory fines for missing essential deadlines.

Similar scenarios are faced within any large operational team; from call centers to KYC remediation teams.

The problem is the number and permutations of the parameters which go into forecasting required capacity. Batching, up and down times of every team and every resource, varying complexity of every task, differing global systems and processes to do the same task – all of these make an impact in predicting capacity requirements and it is difficult to factor them all into one comprehensive model.

Take seasonality as a factor that is not included in current capacity planning. An approach which doesn't take into account seasonality and relies instead on an average resource requirement throughout the year will result in under- and over- resourcing at most points in time.

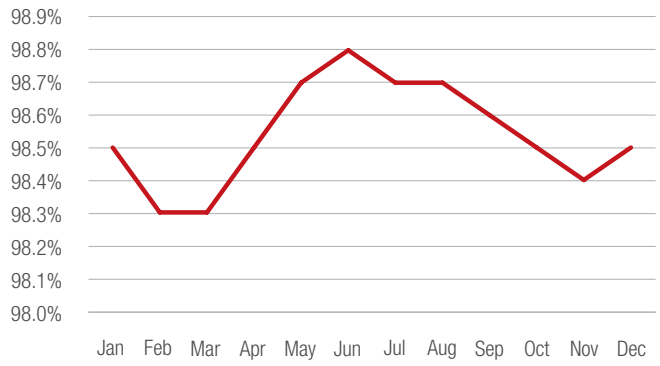

Figure 1: an example of STP rate seasonal variation

Like seasonality, multiple other buried and invisible factors are involved in forecasting capacity. We use data science methods to extract and calculate the impact of these parameters in our capacity models.

Given the focus on cost optimization within the industry in recent years, making the shift towards more intelligent capacity planning is not just a fruitful exercise, but an imperative one.

CAPACITY PLANNING DONE SUCCESSFULLY OUTSIDE OF FINANCIAL SERVICES

Capacity planning solutions have been tried and tested extensively in industries outside of FS.

Amazon, for example, uses capacity prediction in number of different ways[1]: they predict overall retail demand (see Figure 2) to plan their investments in fulfilment centers and hiring new staff, they predict demand for individual product lines to optimize their inventory and they predict demand for their cloud computing capacity.

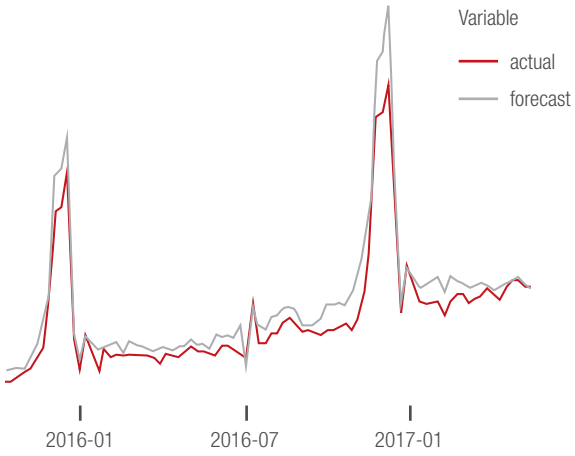

Figure 2: Actual vs forecast detail at Amazon

In a similar example, Uber makes extensive use of capacity planning algorithms[2]. With daily, weekly and monthly seasonality as well as a general upward trend in the number of bookings, this is not an easy task. What's more, Uber not only needs to be able to predict overall demand but be able to predict this demand geographically. Their marketplace forecasting algorithm enables them to direct their drivers towards areas with high predicted future demand (see Figure 3).

The opportunity for companies outside of the technology sector is that many of these technology companies have 'open-sourced' their capacity planning algorithms. Companies in other industries can therefore benefit from this forecasting experience at minimal cost.

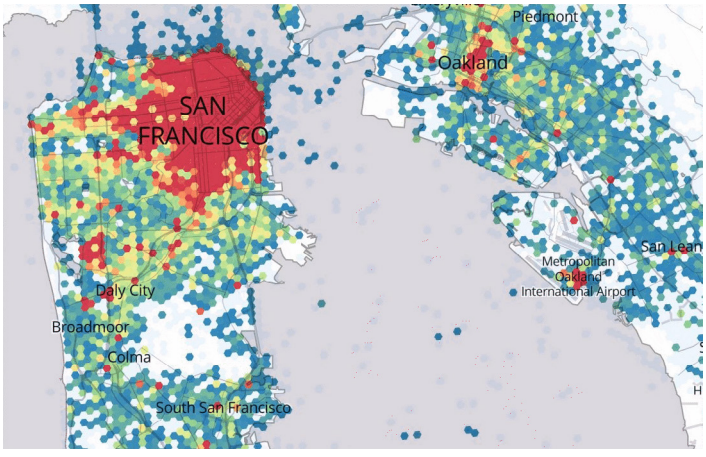

Figure 3: Result of marketplace forecasting algorithm at Uber

1. Forecasting at Amazon – Tim Januschowski: https://forecasters.org/wp-content/uploads/gravity_forms/7-c6dd08fee7f0065037affb5b74fec20a/2017/07/Januschowski_Tim_ISF2017.pdf
2. Forecasting at Uber: https://eng.uber.com/forecasting-introduction/

CAPCO'S APPROACH

Overlaying the tools and methodologies used outside the industry with our deep expertise in financial services, Capco looked to bring the benefits of capacity modelling to our clients.

A shift away from using static spreadsheets to plan capacity, Capco's capacity models were implemented at multiple clients as a platform, leveraging front-end tools such as ReactJS in addition to Spark clusters to effectively process large amounts of data. Operational leads were able to more effective plan and optimize allocation as well as assess opportunity costs of reallocation of resources.

Features of Capco's capacity models

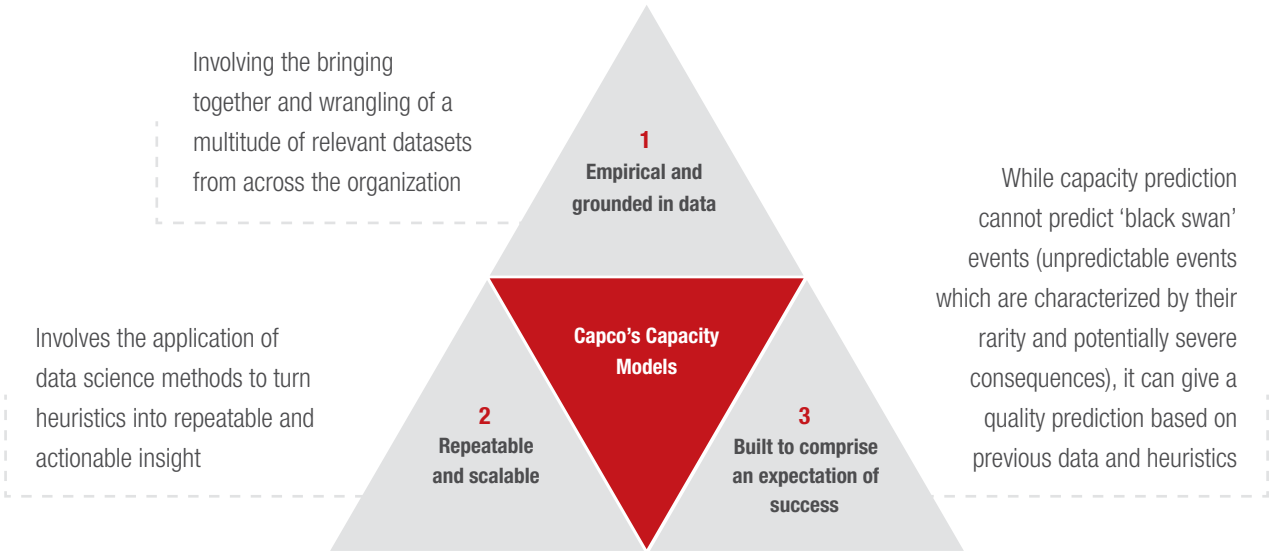

Figure 4 – Capacity modelling: the key differentiating features.

CASE STUDY: MAXIMIZING EMPLOYEE UTILIZATION WITHOUT IMPACTING OPERATIONAL RISK

Capco was engaged by a tier 1 global bank who sought to make intelligent use of their data to improve resource planning for their operations team in the investment bank. Previously, the client had been tracking manual operations and times spent on various processes in order to manage and plan process capacity but struggled to consolidate all the data and draw meaningful insights.

The team followed a four-step agile iterative approach to stand up a functional capacity model combining intelligent heuristics, data science techniques and a forecasting model. The model forecasted the time required to process trades which required manual intervention, breaking this down by product, the team processing the trade and the region.

Step 1: Business understanding

The team began to build a well-rounded understanding of how current capacity planning worked for the bank. This included knowledge of what resources are used and when, understanding of the 'primary drivers' of the demand for the operations team, and familiarity with the key metrics that impact capacity supply. No detail was too small to overlook; from the number of hours worked per team to the holidays booked by team members – every level of granularity which impact capacity was explored as potential to be incorporated into the model. In addition, the team gained an understanding of the specific actions that were required for each type of intervention, be it automated or manual, as well as which applications were used to conduct these actions and how long each of those took.

As the team got a detailed appreciation for current processes, an idea for target state processes was emerging.

Step 2: Exploration of data sources & data lake

Next, a deep dive into the relevant data sources was performed.

Capco typically works with data of different sizes and stored using different technologies, either on premise or on the cloud. Capco's data scientists can analyze data in CSVs or SQL databases but are also adept at using 'big data' technologies such as Hive, Pig and Spark.

In this particular case, relevant data was pulled from several diverse sources. The datasets were then joined together to create a single enriched dataset, stored in a centralized cloud location, which enabled a thorough investigation to realize patterns and interrelations within the data.

The team then proceeded to interrogate the data and ask key questions of it, such as:

- Which factors affected trading volumes for different asset classes?
- What are the key factors that impact STP rates?
- To what extent did volumes or rates exhibit seasonality?
- What were the variations between different regions?

Value was already unlocked at this point, as the team could now provide deep insights to the client as to how ops resources across regions use their time to execute various operational activities.

CASE STUDY: MAXIMIZING EMPLOYEE UTILIZATION WITHOUT IMPACTING OPERATIONAL RISK - CONTINUED

Step 3: Building the capacity model

The rich understanding of current processes was instrumental in building a volume forecasting algorithm, foundational to the bespoke capacity model.

A number of different time-series forecasting algorithms were tested in order to choose the algorithm that best suited this particular use case. Depending on the problem, Capco's data scientists will often begin with applying simpler models such as AR (auto regression) or by making use of Prophet (Facebook's open-sourced time-series forecasting package). In certain situations, and for more complex use cases, deep learning approaches can be used, such as an LSTM (long short term memory) model.

In this particular model, a blended approach was used to forecast future trading volumes and STP rates. Both of these forecast values enabled Capco to estimate the number of trades expected, and the subset of those which would require manual processing. Capco was then able to compute the different times needed to complete the different actions required to process these trades (including automated channels and manual tasks).

All these pieces above were combined with intelligent heuristics and other insights gained to result in the development of a capacity 'calculator', able to forecast future trade volumes across various products and regions.

Step 4: Development of a dashboard to access the model

Capco's goal was for the capacity model, now completely functional, to be deployed onto a platform where it could be accessed and leveraged by relevant users.

While many of our clients have off-the-shelf business intelligence tools to produce visualizations, these usually lack the versatility and flexibility to interact with the dynamic nature of complex models. For this reason, Capco's recommendation was to build a web application, which could be optimized specifically for visualizing the outputs of the model.

Capco's blended team of front-end engineers and data scientists came together to build a user interface which would enable an operations manager to easily and swiftly obtain insights from the calculator.

This dashboard was built using ReactJS which enabled users to interact with the calculator, and crucially, tweak key parameters in order to evaluate different scenarios. All the analysis was presented in a series of easy-to-consume interactive visualizations.

The mock-up (in Figure 5) shows an example of an operational dashboard.

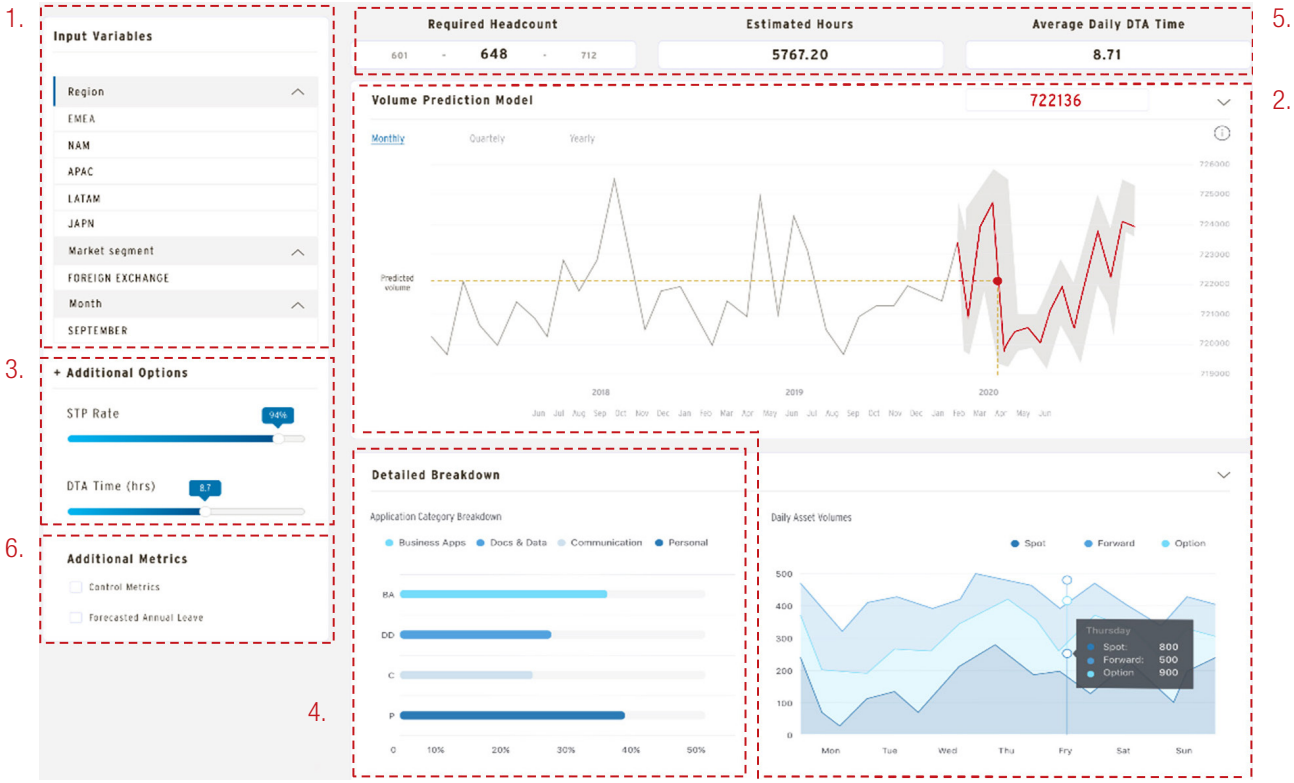

Figure 5 – Example dashboard to enable end-users to easily interact with the model

The user interface above is used in the following way:

1. A user chooses some input variables to forecast against, for example:
    a. Location of interest
    b. Market segment they are interested in
    c. The month of year
2. According to the selected input parameters, the model calculates the predicted volume. This is displayed alongside previous volumes for that year, as well as a detailed breakdown of predicted volumes per asset class (seen at the bottom right of dashboard)
3. The capacity model enables scenario modelling, such that users can tweak the prediction and inspect different configurations by altering other parameters such as:
    a. Current or expected STP Rate (if known)
    b. Average time staff spent on applications
4. Detailed breakdown of time spent on various applications categories
5. The number of work hours required to process the predicted volumes in the given month is derived, and subsequently the number of required FTEs
6. On demand, users can choose to incorporate data from other sources in order to enrich the findings further or explore further correlations; corresponding visualizations then appear below the current graphs.

CLIENT REACTIONS TO THE BESPOKE CAPACITY MODEL

The capacity model allowed the client to:

- Accurately forecast the resource strength required in different regions and teams
- Plan their resourcing against the peaks and troughs expected over the year
- Effectively reallocate resources in real time.

The application that was most lauded was the integrated scenario model. This feature empowered each manager to personalize the capacity calculator to suit the needs for their particular team, down to the dimensions of the region and function they operated in.

This improved their capacity planning going forward, bringing about both budgeting and scaling benefits. As a result, managers were able to develop financial projections and effectively allocate resources and costs, and also cut costs where teams were over-resourced.

CONCLUSION

Whilst building a strategic capacity model might seem like an intimidating process, the cost of running operational teams as is currently being done is simply inefficient. Allocating more resources is not an effective solution to the problem, as years of experience has shown.

In our experience, we have seen that financial services organizations can benefit by saving hundreds of thousands, even millions, through accurate planning and cost avoidances associated of underutilized resources.

CONTACT

Intrigued by our solution? Get in touch with our Data Science capability lead,
Nadir Basma, nadir.basma@capco.com or our Data Partner, **Chris Probert** on chris.probert@capco.com

AUTHORS

Nadir Basma, Data Science Capability Lead
nadir.basma@capco.com

Edward Pease, Data Science Consultant
edward.pease@capco.com

ABOUT CAPCO

Capco, a Wipro company, is a global technology and management consultancy specializing in driving digital transformation in the financial services industry. With a growing client portfolio comprising of over 100 global organizations, Capco operates at the intersection of business and technology by combining innovative thinking with unrivalled industry knowledge to deliver end-to-end data-driven solutions and fast-track digital initiatives for banking and payments, capital markets, wealth and asset management, insurance, and the energy sector. Capco's cutting-edge ingenuity is brought to life through its Innovation Labs and award-winning Be Yourself At Work culture and diverse talent.

To learn more, visit www.capco.com or follow us on Twitter, Facebook, YouTube, LinkedIn Instagram, and Xing.

WORLDWIDE OFFICES

APAC
Bangalore
Bangkok
Gurgaon
Hong Kong
Kuala Lumpur
Mumbai
Pune
Singapore

EUROPE
Berlin
Bratislava
Brussels
Dusseldorf
Edinburgh
Frankfurt
Geneva
London
Munich
Paris
Vienna
Warsaw
Zurich

NORTH AMERICA
Charlotte
Chicago
Dallas
Hartford
Houston
New York
Orlando
Toronto
Tysons Corner
Washington, DC

SOUTH AMERICA
São Paulo

WWW.CAPCO.COM

CAPCO
a wipro company