
POLY//MATICA

RECOMMENDATION ENGINES

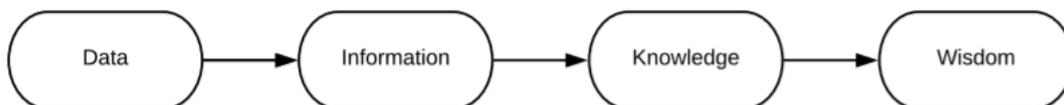
THE POLYMATICA WAY



RECOMMENDATION ENGINES AND THE POLYMATICA WAY

Our philosophy leads from data to wisdom

Every day, humans process data, make sense of it and base our decisions around it. The process happens in four stages: Data, Information, Knowledge, Wisdom.



Seeing shapes on a computer screen, our brains are processing data and turning it into information by transforming shapes into letters. By registering the letters as words and sentences, we're turning the information into knowledge: from looking at a timetable, for example, we can conclude our train departs at 7pm. By leveraging that knowledge and making it useful to us, we turn it into wisdom: we decide to leave the house at 6:30pm to make sure we catch the train.

Scaling things up, the same logic can be applied to business. In business, data needs to be transformed into wisdom every minute of the day. For a bank, billions of financial transactions are recorded as raw data and, with context, the data becomes incredibly powerful. Let's use a banking customer – Bob – as an example. Bob has 10 rows stored in binary in a bank's database. The rows relate to his transactions. That's data – so far it makes little sense to humans. To turn the data into information, we need to first extract it then view it in an understandable way: for example, as a table, perhaps showing the exact time, amount and location of Bob's financial transactions. By looking at the table and slicing and dicing the information, we can turn it into knowledge: Bob goes out to lunch every weekday at 1pm. The last step is turning the newfound knowledge into wisdom by acting on it. As the clock approaches 12:45pm, we send Bob a discount voucher for a nearby lunch place – and Bob is delighted to accept.

Of course, a single human would never have the capacity to scale the above process to tens of millions of customers. Without the assistance of technology, the feat would be impossible for even an army of Business Analysts, no matter how clever and/or hardworking they might be. The technology necessary is Business Intelligence (BI) and, according to Gartner, investment in the field will grow to \$22.8 billion by 2020.

Most BI solutions turn data into information through visuals but go no further. It's understandable. Historically, two sizeable obstacles have always limited the power of BI: scarcity of data and data processing speeds.

The former has been solved by the declining prices of data storage. In fact, many argue that data is the highest-price commodity in our changing world, and most organisations have adopted the approach of collecting as much data as possible, hoping to use it in the future.

The obstacle of data processing speed has only been solved more recently. Following the intelligent application of GPUs, as pioneered by Polymatica, it's now possible to process billions of rows of data in seconds.

Processing vast amounts of information can be extremely challenging for Business Analysts. Many organisations turn to expensive Data Scientists to help with automating the process.

At Polymatica, we're paving the way for what we call the *Citizen Data Scientist* – someone able to gain knowledge from information without being a Programmer, a Business Analyst or a trained Data Scientist. By using point-and-click advanced analytical modules such as clustering, association rules and forecasting, Citizen Data Scientists are able to quickly make sense of data, finding patterns that are usually hidden within the sheer volumes of data.

The final step is assisting Citizen Data Scientists in acting on their knowledge, and thus attaining business wisdom. Polymatica's vision is to stay with the Citizen Data Scientist every step of the way, and we achieve this through the continued development and optimisation of our Recommendation Engine.

What are recommendation engines and why do we need them?

A recommendation engine is a general term and can vary in complexity massively. From simply grouping users based on their age and gender to creating a unique persona for each customer, recommendation engines are on track to begin predicting behaviour at least as well as, and often better than, a human with access to the same information.

Considering the information overload a typical person experiences daily, personalised marketing becomes more and more essential as generic marketing emails go unopened and ads blocked. Businesses no longer need to worry solely about signing up new customers – retaining existing customers is now just as important. This is especially relevant in the financial sector in light of open data mandates such as Basel III and PSD2, which make switching financial providers simple and transparent.

While the short-term aim of recommendation engines is to sell more products, the longer-term aim is for the customers to come to trust the recommendations and ultimately to rely on them – just as they once did following years of using the same financial advisor. If customers are consistently recommended products they need, they're naturally more likely to read the recommendations and try suggestions out.

With technologies like Polymatica now able to process massive amounts of data, often in real-time, recommendation engines can finally fulfil their potential of making every marketing moment matter. It's an advance set to enable progressive brands to build deep and valued relationships with all their customers.

Delivering the perfect marketing campaign

Effective management of a recommendation engine is no simple task. It usually requires Data Scientists to build the individual components used in both the models and training algorithms.

Polymatica shifts the task from Data Scientists to Citizen Data Scientists by building the advanced analytical modules necessary for a recommendation engine beforehand. From instantaneous data exploration and analysis to automated association rules, clustering and forecasting algorithms and full access by APIs, Polymatica is paving the way for Citizen Data Scientists to leverage the power of recommendation engines.

A simple recommendation engine could look at customers based on their previous behaviour. It could run association rules to determine product dependencies and needs. If Bob has just bought a house, for example, Polymatica could recommend he buy home insurance, as the two usually go together in that particular order.

Going a step further, a recommendation engine could analyse the explicit data Bob provides: being a 32-year-old male earning £150,000 a year in West London, Polymatica's Recommendation Engine might recognise Bob as a time-poor, high-net worth individual. It might recommend an upmarket insurance product with the emphasis on a high quality service.

Bob's persona is limited in this case. While much better than "traditional" marketing, simple recommendation engines can sometimes misfire, especially when lacking business rules and impact analysis. A customer who has already purchased home insurance being bombarded with fresh quotes for the next two weeks is frustrating and common.

Due to its ability to quickly process massive amounts of data, Polymatica's Recommendation Engine will improve targeting accuracy even further by including implicit data – something that the customer has not explicitly revealed, that either appeared in their behaviour, or that's taken from another source altogether. This might include the products they browse, the amount of time they spend looking at each product and/or external data such as UK public Companies House data.

A more sophisticated Polymatica Recommendation Engine could take Bob's existing persona (a 32-year-old male making £150,000 a year in West London) and blend it with implicit data such as the fact that he goes out to lunch at 1pm almost every weekday and spends around £10 on lunch. Knowing that Bob is on Oxford Street at 12:45pm on a Tuesday, Polymatica's engine could send him an offer for a nearby restaurant.

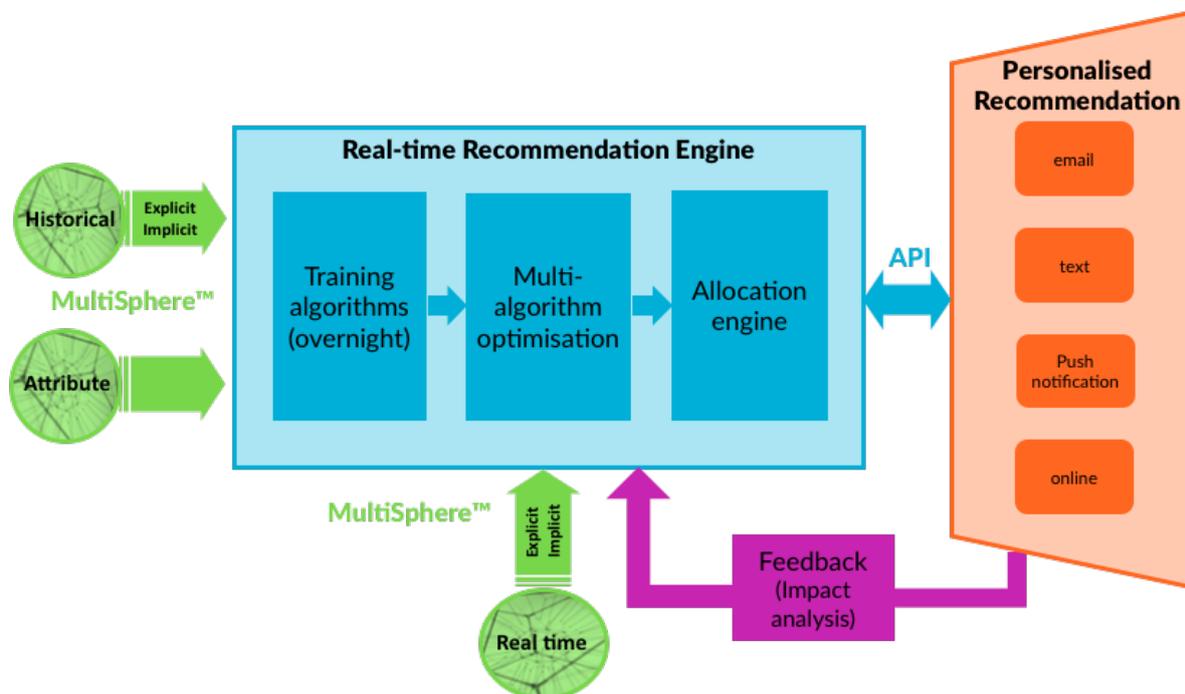
The final piece of the puzzle is including impact analysis and business rules in the equation. Polymatica's Impact Analysis engine could determine Bob rarely opens marketing emails but frequently responds to SMS notifications, and therefore suggest future discounts be delivered via SMS.

As a customer interacts with the platform more, their personas become increasingly populated and models increasingly optimised. Over time, recommendations become increasingly accurate meaning customers become more likely to respond positively, leading to ever-improving business outcomes.

Our Progress to date and platform overview

Polymatica is currently used extensively to make recommendations that significantly increase the conversion rates of targeted marketing, with much more relevant content and lower frustration levels. We're currently in the process of developing the next phase of our Recommendation Engine. The new phase will provide a fully automated solution that consumes streaming data into the platform, is able to make a recommendations in real-time, can connect to a broad variety of systems via an API or as an ODBC connector, and is self-learning.

Schematic 1: A high level overview of Polymatica's Recommendation Engine



The Data Science core of our engine will consist of three parts: the training algorithms, the multi-algorithm optimisation aspect, and the allocation engine. The training algorithm runs periodically (typically on fifteen-minute cycles), updating multiple models according to the most up-to-date data. The model houses the information necessary to make recommendations, accepting new pieces of information in real-time. The model optimisation algorithm will then evaluate available models against each other and will decide on the optimal model on a case-by-case basis. Finally, the allocation engine will contain business rules and traffic control, regulating factors like the maximum number of messages a customer sees in a day, or the means of communication that take priority.

The Impact Analysis module analyses customer behaviour before and after a targeted campaign and, based on all information, optimises the type of campaign delivery. By continuously analysing trends in customer behaviour, the module finds and recommends the most effective campaigns, as well as the optimal campaign delivery methods.

To ensure we make the best possible recommendations, we plan to implement a mixture of classical models and models that utilise deep machine learning. These models are continuously updated and validated using response feedback data. The models compete with each other on a regular basis and are being designed to readily consume new sources of data to help improve their accuracy.

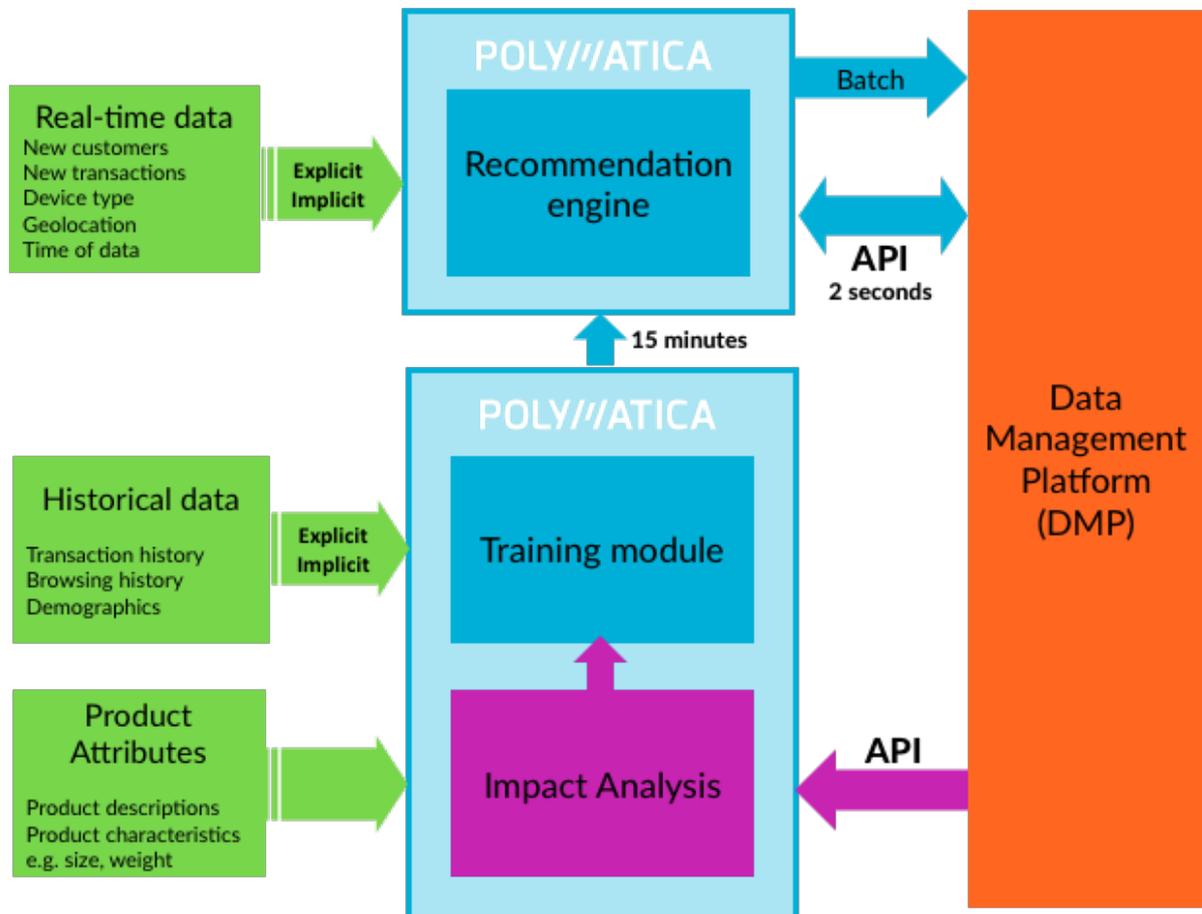
Connecting to the outside world

Polymatica is designed to connect easily to any system using Application Programming Interface (APIs), or as an ODBC. This makes the integration of an existing Data Management Platform (DMP) with Polymatica extremely flexible.

Polymatica's distributed architecture allows the training algorithms and the recommendation models to run on two separate servers. The training algorithm module can ingest data continuously, including historical data, new product data and behavioural data. These models can be scheduled to update as frequently as every 15 minutes. A DMP can communicate feedback information on the impact of campaigns, which the Impact Analysis module can digest and add to the training model.

The recommendation engine itself can run on a separate server, and can provide recommendations based on real-time data, such as geo-location or device type. Recommendations can either be retrieved from Polymatica on-demand in real-time, or pushed into a DMP in a periodic batch process.

Schematic 2: Polymatica's Dual Server Recommendation Engine Configuration



Polymatica’s modular architecture means adding new functionality is easy and doesn't affect core modules. In combination with open APIs, the modular architecture means Polymatica’s Recommendation Engine can be adjusted to a client’s specific needs, business rules, or preferred models.

Reimagining Targeted Marketing

Recommendation engines allow companies to turn their data into wisdom. As the obstacles of data scarcity and processing speed have been removed, now is the time to reimagine targeted marketing by adopting recommendation engines. The good news is we've already done the groundwork and are continuously building and improving the components necessary for advanced recommendation engines, alongside making the tuning simple for the new breed of analysts – Citizen Data Scientists.

We believe the keys to the next generation of recommendation engines will be:

1. Unlimited scale in terms of data volumes and speed, making every piece of data and customer interaction matter
2. An engines' ability to provide streaming recommendations in real-time across all channels
3. Models that are continuously updated based on the latest data
4. A design fit for Citizen Data Scientists, enabling people to easily develop and enhance their own unique models, rather than rely on generic out-of-the-box frameworks

All four are core to our Recommendation Engine roadmap.

We're passionate about building for the future and helping our clients create deep and lasting relationships with all of their customers. The evolution of our Recommendation Engine is our passion in practice.

About us

Polymatica is a Data Science and Visual Analytics solution with embedded intelligence that works at scale and speed using revolutionary GPU and CPU hybrid architecture. Engineered for everyone from Business Executives and their teams to Citizen Data Scientists and trained data professionals, Polymatica offer a comprehensive portfolio of frontend analysis, predictive and visualisation tools, as well as comprehensive training programmes and in-depth masterclasses.

For more information on Polymatica, please visit www.polymatica.com or call +44 (0) 20 3468 1974.