NOVEMBER 2020

Building Personalised Shopping Experiences with Recommendations AI

Machine Learning on the Google Cloud Platform

Riley. Google Cloud

Motivations

Modern retail is being transformed by digitisation, online shopping and machine learning. Google's Recommendations AI offers state-of-the-art, personalised product recommendations for online stores designed to understand your business, your customers and your products.

Designed to lift key metrics such as click-through and conversion rates, Recommendations AI is built on Google's experience and research in predictive algorithms, overcoming many of the challenges faced by other recommendation engines. It packages everything up into a simple, scalable and results-driven product as a managed service on the Google Cloud Platform.

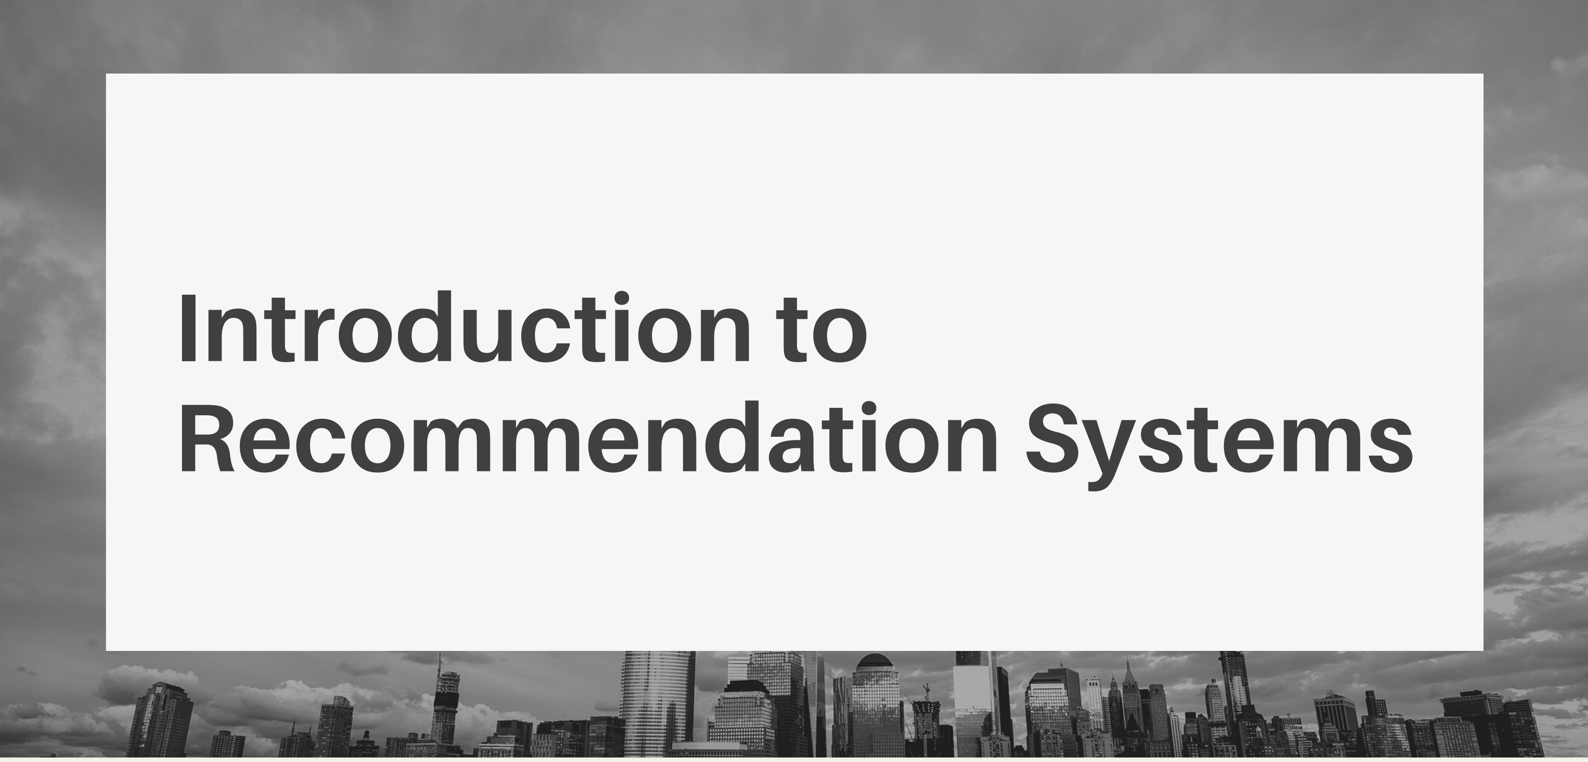

Introduction to Recommendation Systems

Personalised Recommendations

A well-placed product recommendation serves as a map for your visitors to find the next step in their customer journey. It helps drive engagement and conversions, whilst motivating customer exploration and enhancing the user experience.

Recommendation engines have proven uplift in key metrics, such as click-through-rate, revenue per session and conversion rates. A 2018 EQ report found a 70% lift in purchases rates with new customers who clicked on a recommendation.

Classes of Recommender Systems

There are a diversity of approaches to product recommendation, starting with manual recommendations decided by sales data or market-based analysis.

Automated approaches use machine learning models based on product and user event data extracted from the website. This information is passed through a model, which could be as simple as recommending items within the same product category. More complex machine learning models can form intelligent, context-based prediction by training on vast amounts of data.

In general, machine learning recommender systems use either collaborative or content-based filtering methods to recommend a product.

Collaborative Filtering

Recommender systems based on collaborative filtering leverage historical user data from all past visitors and customers. By pooling and grouping this information, clusters of similar customers are formed to guide recommendations.

Collaborative methods ensure the prediction follows real human behaviour. As the engine relies on past user data, it is vulnerable to only predicting the most popular products and avoiding new ones.

Content-based Filtering

Content-based recommender methods use the customer's own history and the metadata associated with the products they view. For example, content-based methods are able to identify products with similar keywords to the one a customer is currently viewing.

Recommendations are based on a web of connections between related products, independent of user behaviour. Models based on content-based filtering are best suited for recommending similar or alternative items rather than for shopping cart expansion.

Introduction to Recommendations AI

The Google Cloud Platform's Recommendations AI is a hybrid of collaborative and content-based methods, combining vast amounts of historical user data with intelligent product groupings. Behind each prediction is a deep neural network, trained and optimised by Google to deliver highly-personalised product recommendations.

The model is able to learn complex arbitrary relationships, such as user behaviour and how products interrelate. It considers:

- A customer's individual journey with your store, across all devices and past sessions.
- The website's user event history, including the sequence and order of critical events such as product views, add-to-carts and purchases.
- Every product and any available metadata, including the category, price and descriptions of each.

As a result, each prediction is based on a well-founded, context-based model on an individual basis.

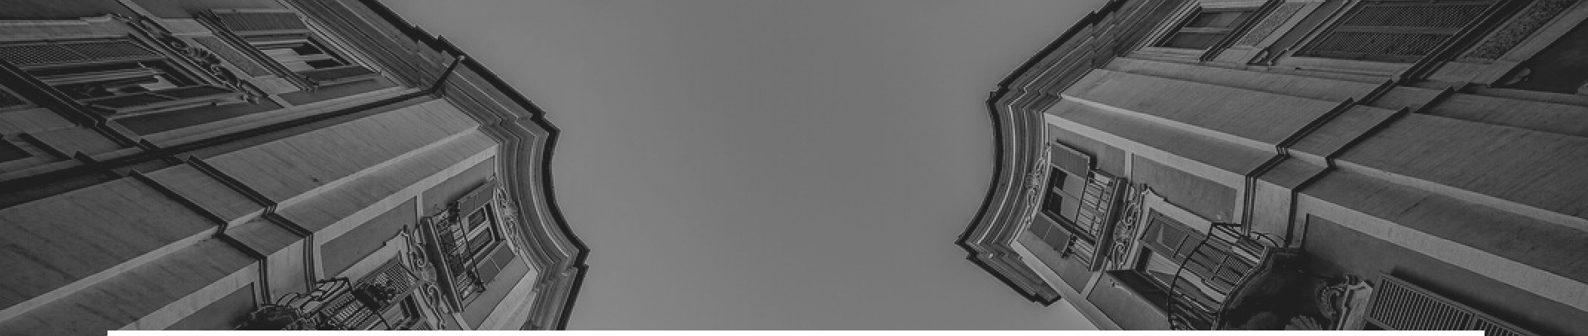

Solving Challenges of Recommender Systems

Data Ingestion

Online stores accumulate terabytes of data. Although this data is essential to training an accurate machine learning model, much of it is unstructured and difficult for humans to draw information from

For example, what effect should the release date of a product have on its recommendation rate? Recommendations AI puts the data to work, quickly learning which features are most important in leading to conversions.

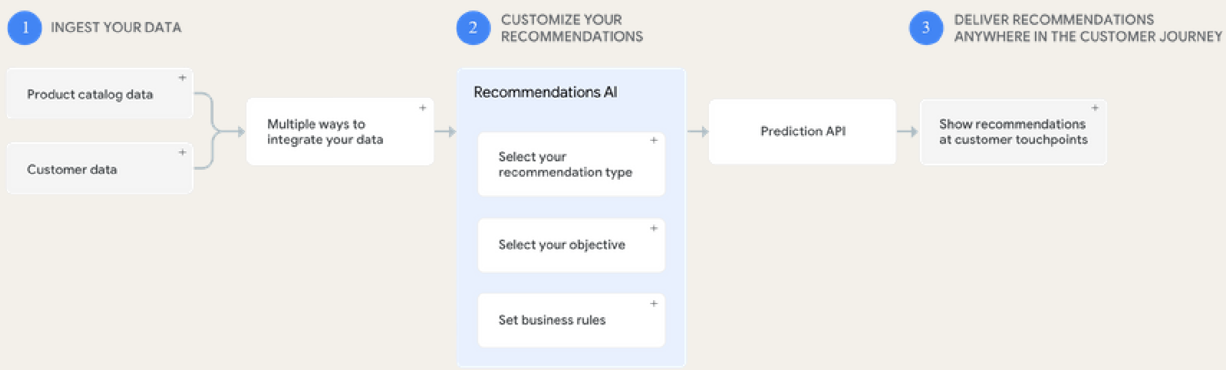

Set-up flow of the Recommendations AI

Using a machine learning model avoids time spent preprocessing and cleaning data. Instead, user events and catalog information can be ingested directly in bulk.

Long-Tail Products

Recommendation engines are easily susceptible to bias, such as favouring best-selling products. These biases lead to poorer model predictions, with niche products on the long tail being underrepresented. In machine learning, these products are said to have sparse labels and can only be properly understood by a well-designed model. This problem also extends to new products, where limited historical data reduces the engine's precision.

Recommendations AI has been specifically tuned to handle these cases, adjusting for popularity and release date. Once trained, a model will continue to re-train on a regular basis to account for new user events and updates in the product catalog.

Deployment

Once an accurate model has been trained, personalised recommendations are delivered to thousands of users worldwide. This must be achieved without increasing website loading time and causing decreased engagement. Recommendations AI is a fully managed service, automatically scaling, provisioning itself resources, and serving recommendations to customers globally. Average prediction times are often less than 100 milliseconds. For retailers, this saves having to deploy their own global, low-latency infrastructure for predictions.

Retailers also frequently face surges in traffic, up 3 to 10 times during critical periods.The Recommendations AI leverages the massive scale of the Google Cloud Platform, allocating compute power on-demand, ready to handle volatile traffic levels when it counts.

Managed by Recommendations AI	All that's left
Data Cleaning and processing	Importing products and events in bulk
Machine learning code	Calling the API for results
Globally provisioned servers	
Automatic model re-training	
Logging and alerts	
Comprehensive dashboards	

Best Practices

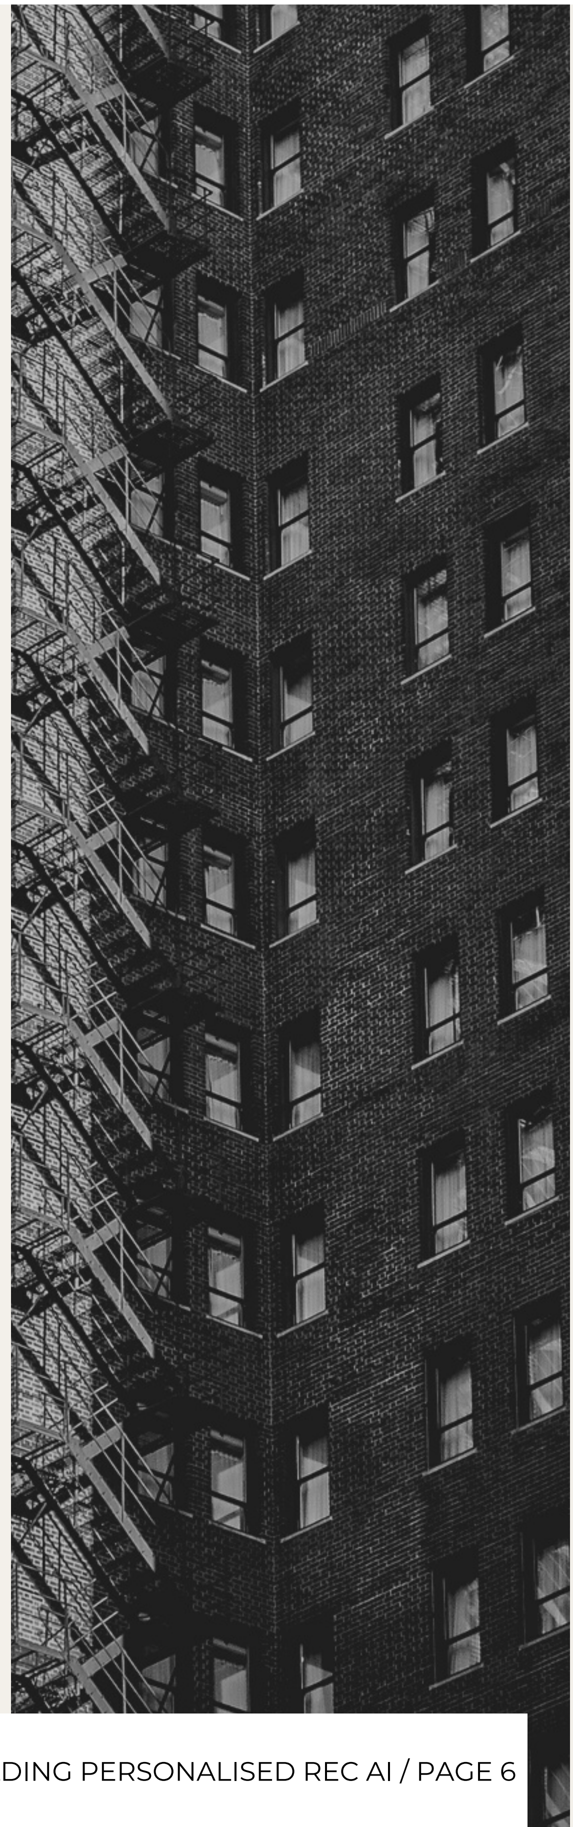

Creating the Right Model

Recommendations AI offers three distinct model types, each suitable in different placements around a website.

- “Recommended for you” recommendations are similar to content-based filtering, where the products are those the user is likely to interact with next. They are often placed on the home page.
- “Others you may like” recommendations tend to generate the highest click-through and conversion rates. Well-suited to product pages, they consider user history and related items.
- “Frequently bought together” recommendations are applicable once the user has shown intent to buy one or more products. The recommendation predicts which product the user is most likely to also buy.

Each model can be trained to maximise a particular business objective, either click-through rate, conversion rate or revenue per session. Click-through rate is the most common and reliable metric.

Once a model demonstrates success, it can be tuned further to support custom filtering rules, show products from a higher diversity of categories or re-rank on price.

Data Ingestion

To maximise accuracy, the most up-to-date user event and product catalog data needs to be provided to Recommendations AI. Historical user events can be imported in bulk before training, and real-time updates can be fed to the engine from JavaScript or Google Tag Manager.

Using a Cloud Function, product catalog items can be imported daily from a data source such as BigQuery or Google Merchant Center. An updated catalog means Recommendations AI can start gathering data and recommend the latest products

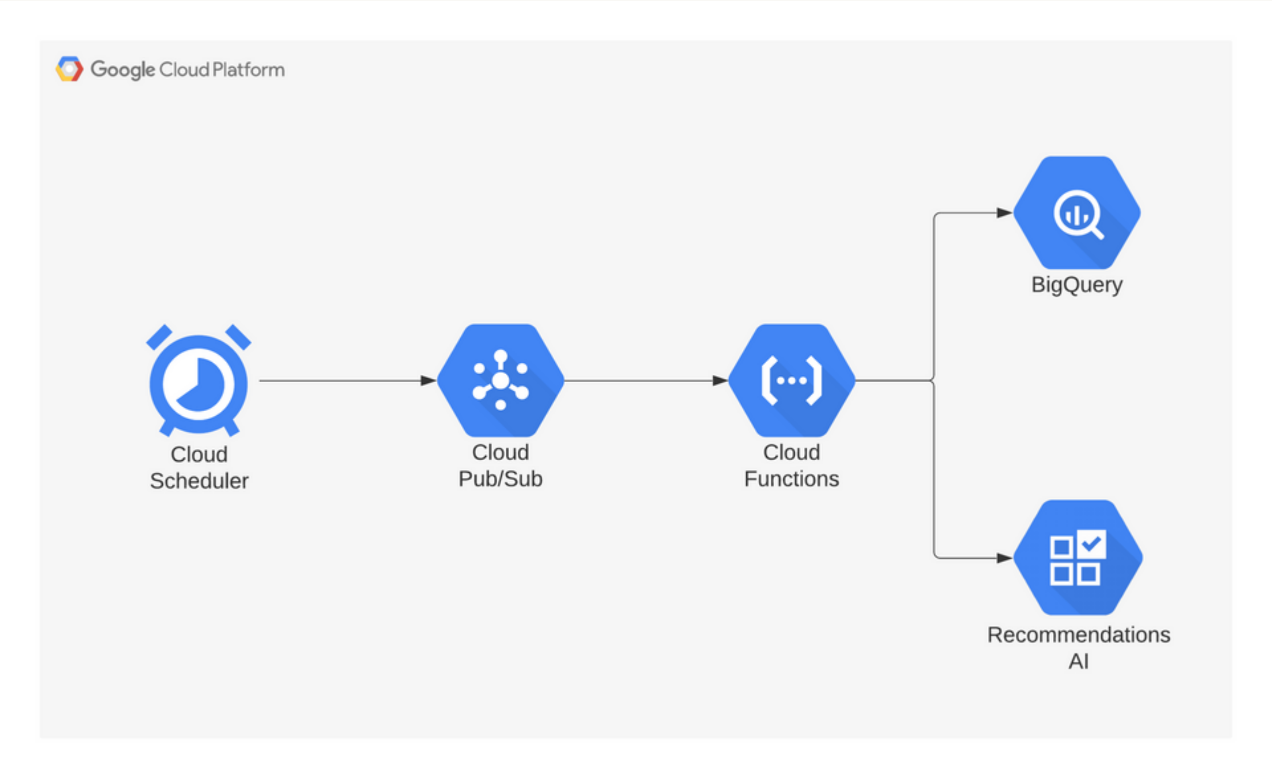

Architecture for routine product catalog imports from BigQuery.

Cost Optimisation

Recommendations AI's pricing model bills for node hours spent training models. This can be minimised by disabling automated re-training on inactive models. The only other cost is on a per-prediction basis, making the engine ideal to deploy in trial runs and tests.

Monitoring

Recommendations AI offers a comprehensive dashboard, providing metrics on the performance of each placement and their corresponding click-through-rates and revenue metrics.

The engine automatically re-trains, scales and delivers predictions, so alerts to notify of import or prediction errors should be implemented.

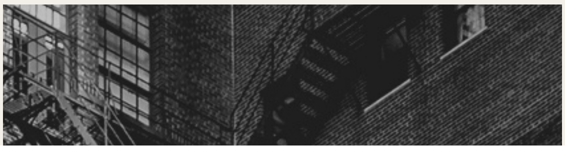

Conclusion

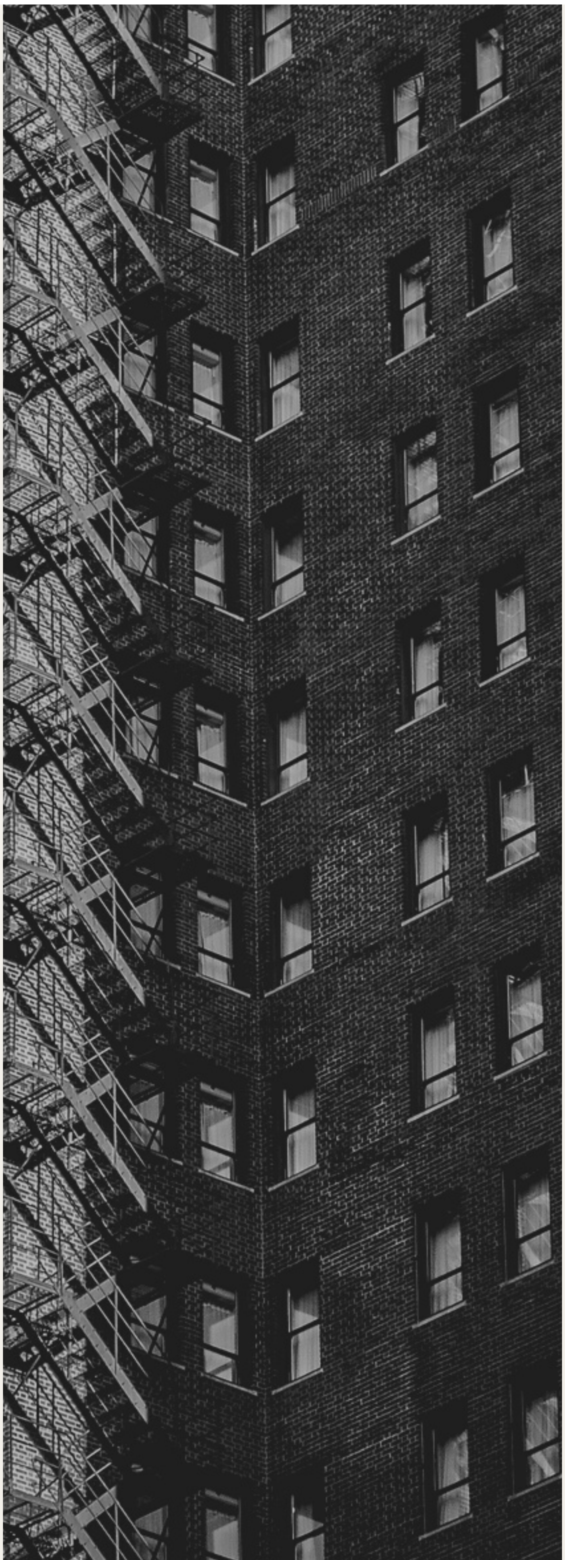

Personalised recommendations strengthen customer experiences whilst lifting key sales metrics throughout online stores. Google's Recommendations AI unifies content and collaborative filtering methods, and solves challenges faced by existing engines with data ingestion, long-tail products and deployment.

Riley can help lead digital transformation in retail and build technology needed for the modern business to thrive.

About Riley

Our purpose is to seek a positive return on creativity. We recognise the only actions worth taking are those that move toward building a better future.

Riley is an Australian technology company with a primary focus on digital transformation, data and analytics, and infrastructure modernisation. Our highly skilled team specialise in the design, deployment and management of enterprise workloads on customised AWS infrastructures.

We help customers realise the true potential of public cloud by providing tested, proven and trusted solutions

Contact

Riley was founded in June 2010 (as Data Solutions Group), as was among one of the first in Australia to migrate Enterprise workloads to the Public Cloud.
In the last 5 years we have assisted over 100 businesses across Financial Services, Retail, Manufacturing and Services industries to use cloud-native principles to transform their business.

RIley.
1300 924 076
hello@withriley.com

Melbourne
Level 4, 313 Flinders Lane
Melbourne VIC 3000
Australia

Sydney
383 George St
Sydney VIC 2000
Australia