Success Story

Developing an AI-based Recommendation Engine for India’s Largest Automobile Manufacturer

Customer	Country	Industry
India’s largest automobile manufacturer	India	Automobile

About the client

The client is India's largest automobile manufacturer, clocking-in a revenue of more than US$17 billion. It is credited with having ushered in the automobile revolution in the country and has established dominance in the sector with over 40% market share and more than 28 million customers till date.

Tech Stack

python™ | Streamlit | jupyter | amazon EC2 | databricks

Business Situation

Research shows that 71% of customers today expect brands to personalize interactions and as many as 76% move away, if they don't get any. While 91% of customers prefer to shop from brands who regularly send them relevant offers.

Personalization is the key to cutting through the noise and capturing customer's attention. Numerous research studies highlight the pivotal role of personalization & how it is essential in retaining customers in the long run. Recognizing this, the client, a leading automobile manufacturer, sought to enhance user journeys and drive loyalty through a personalized messaging initiative for its vast user base.

The client's vision was to send customized recommendations to its customers on specific occasions. Customers who've owned the brand's car for more than 5 years, were to receive a new car recommendation while recent owners would get service-related offers and discounts.

However, there were many hurdles on the way.

Firstly, being the largest automobile manufacturer in India, the client has a huge customer base of millions of users, covering over 40% of the total market share; and delivering tailored suggestions to every customer was a daunting task for the client's team.

Secondly, the client took a manual approach to messaging. This required the team to sift through the vast datasets like name, car model, year of purchase etc., filter out the relevant users, identify their specific needs and manually curate & send recommendations and offers to each one of them.

Scaling the initiative was another challenge for the client. As the process was manual, this task was not just time and resource intensive, but generated little-to-no value and had a negligible impact on the client's growth.

Major challenges that the client faced were:

- *The lengthy data-search process and labor-intensive manual messaging was a major bottleneck.*
- *Achieving true personalization was a distant dream for the client - required more than just changing names and details in static messages*
- *Maintaining recommendation accuracy for the client's growing customer base, their changing needs and evolving times was yet another challenge for the team*

The Solution

Team Daffodil collaborated with the client's team to have an in-depth understanding of their goals and requirements and identify bottlenecks in the process. Our team kickstarted the development of a comprehensive web application, post a comprehensive Discover and Frame workshop - where we defined the entire scope, tech stacks to be used and UI/UX of the solution.

In order to deliver tailored, compelling and conversion-worthy messages efficiently and at speed, Team Daffodil harnessed GenAI-based LLM models, Python & Streamlit for application development. Further, OpenAI's GPT-3.5 was utilized to create personalized messages for the end users.

The application was deployed on Amazon's EC2 Instances, capable of handling high user traffic and data volume. DataBricks provided a central workspace for model training while Jupyter Notebook helped in training models and visualizing data.

Team Daffodil encountered several challenges during the development stages. A few of them were:

Insufficient Data:
The data shared by the client was insufficient for model training and posed a significant challenge in accurately recommending suggestions. To overcome this, Team Daffodil employed synthetic data utilization techniques & expanded the dataset leveraging machine learning (ML), to improve model performance.

Time Constraints & Need for Accuracy:
The team had to deliver the application with accurate predictions in a short span of time. Our team did careful prioritization of features & adopted efficient development practices to overcome this challenge.

Evolving Times:
During the development, there were many changes on the client's production side, such as discontinuation of certain models and car types. It was necessary to accommodate these changes for the recommendations to be accurate and relevant. To tackle the same & stop the system from suggesting outdated vehicle choices, the irrelevant data was eliminated from the system from time to time, thereby enhancing the overall reliability of the solution.

How the solution works?

On the basis of the targeted filters such as birthday, anniversary, service due date, etc., the system showcases the list of highly-probable users/customers whom the client could send the messages to. The client could then view the customer's profile with details such as anniversary, region, current car, budget etc. and select the most suitable messaging from the list of available, AI-generated options.

The recommendations could either be for a new car for users who've owned the vehicle for over 5 years; or of a service for recent buyers.

Key features of the solution are:

Car Recommendation Engine:
Leveraging collaborative filtering, the car recommendation engine offers personalized vehicle suggestions to customers, based on their past purchase history and preferences.

In order to predict if a customer would like a product recommendation or not, the system employs matrix-factorization and forecasts how the customers would rate the cars even when they haven't interacted with them physically. It analyzes data types such as user data (age, gender, location etc.), car-related information such as pricing, fuel-efficiency etc., past purchases and models owned, along with feedback on test-drives, if they've taken any.

Further, the recommendation engine is trained with user-profile to offer relevant suggestions.

True Value Prediction:
The solution also offers the feature to predict the true value of a customer's current vehicle in case they're looking out for an exchange/ upgrade or a trade-in opportunity. Leveraging the same, the system can offer relevant messaging to customers and help to increase customer retention and loyalty towards the brand.

Best Time Prediction:
To enhance the efficiency of responses, Daffodil leveraged the Generative AI capabilities of Large Language Models (LLMs) along with automation to generate human-like messages. Through a rule based system, it became possible to predict the best time to reach out to customers.

Here's how the rule-based system works:

- **Special Occasions:** The rule-based system identifies special days such as car purchase anniversary, birthdays etc. to send messages around that time for maximum user engagement.
- **Festive Days:** The system schedules messages and greetings around major cultural holidays to connect with users & resonate with their religious practices.
- **Service Reminders:** By analyzing a customer's vehicle purchase date and service history, the system predicts the right time to send service and maintenance reminders to the target customers.
- **New Offers:** For customers who are due for an upgrade or have shown interest in newer car models, the system directs regular offers, launch info and promotional messages to those users.

Message Generation:
The LLM-based message generation process collates the data and insights and transforms them into personalized messages which reflect the client's tone & values. For the same, Team Daffodil had set clear guidelines while training the solution to showcase results based on the brand's style of language and emotional undertones such as professional, friendly, exciting etc. This is done to generate brand-like tonality into the messages.

Next, leveraging prompt engineering, our team enhanced the solution by utilizing content-based filtering & eliminating certain keywords that did not adhere to client's brand policies. These phrases and keywords were then integrating into the LLM's operational framework.

The recommendation engine was then trained with previously successful messages which were further personalized for the customers unique needs.

For the solution to work seamlessly, the system is made to intelligently integrate specific data such as car recommendations, service due date etc. at the suitable time in a natural, human-like tonality that feels personal to each user.

Impact

The end result was an efficient, streamlined and customer-centric marketing process that empowered the client to reach the right customers at the right time. The automated solution reduced the time consumed to sort users and send recommendations. The AI-driven recommendation engine allowed the client to tailor suggestions & match them to their customers' preference.

The developed solution not only improved customer satisfaction but also drove higher engagement & conversion for the client, establishing personalization as one of the key pillars of the client's marketing strategy.

Product Screenshots

Feature-rich user interface

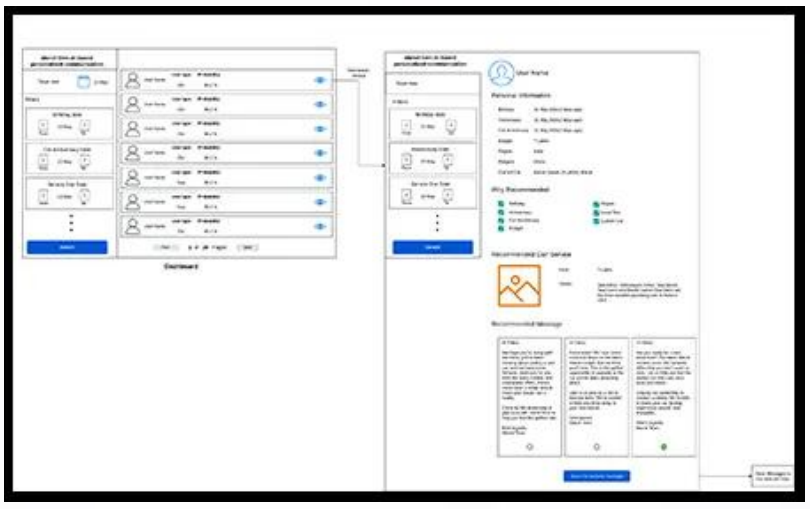

Recommendation engine workflow

User profile shortlisting

True value prediction

Best time prediction

LLM-based message generation

Services Used

- Generative AI Services & Solutions
- AI Development Services
- Software Development Services

Have a software product vision in mind?

Setup a personalized consultation with our technology expert.

Let’s Talk