Personalized Digital Customer Services for Consumer Banking Call Centre using Neural Networks

Xuejie Zhang*
Analytics Center of Excellence
DBS Bank
Singapore
xuejiezhang@dbs.com

Samarth Agarwal
Analytics Center of Excellence
DBS Bank
Singapore
samarthagarwal@dbs.com

Ruth Choy
Analytics Center of Excellence
DBS Bank
Singapore
ruthchoy@dbs.com

Kay Jan Wong
Analytics Center of Excellence
DBS Bank
Singapore
kayjanwong@dbs.com

Lecia Lim
Analytics Center of Excellence
DBS Bank
Singapore
lecialim@dbs.com

Ying Yang Lee
Analytics Center of Excellence
DBS Bank
Singapore
yingyang@1bank.dbs.com

John Jianan Lu
Analytics Center of Excellence
DBS Bank
Singapore
johnlu@dbs.com

***Abstract*—Many banks and financial service companies have been transforming the way to run their businesses and serve customers. In this paper, we present a use case for digitizing customer journeys in the area of consumer banking call centre. The main objective is to provide personalized customer service experience through an integrated solution for call centre, including the Interactive Voice Response (IVR) system, SMS system, Internet Banking platform and chatbot. Topic modeling was performed on the dialogue transcript between the customers and Customer Service Officers (CSOs) to identify the customers' reason for calling. Using the customer-level profile, transaction and servicing log data, a multi-task neural network was trained to predict if a customer is going to call the bank for any customer service request in the next 10 days. In the IVR system, a personalized voice prompt will recommend relevant digital services based on the model prediction and redirect the customer to digital services through a SMS with a URL to chatbot. Through this, the number of calls reaching the CSOs has reduced and the bank can achieve considerable operational cost savings and provide a more efficient customer service experience.**

Keywords—call centre, customer service, topic modeling, neural network, digital transformation, banking

I. INTRODUCTION

Machine learning or artificial intelligence has been popular in business transformation across various industries and companies. Banks and financial service companies are also highly interested in applying machine learning techniques to improve efficiency, increase revenue and reduce operation cost.

In a highly competitive market where companies are ceaselessly competing for customer's loyalty, the role of a customer support call centre gains increasing importance, and the solutions to effective customer service comes to the fore. In the banking industry, call centres are used to meet the customers demand for various service needs such as credit card fee waiver and transaction enquiry. However, such enquiries and services can be repetitive and trivial. When there is a surge in number of calls, customers may have to wait for a long time before obtaining a response to a trivial enquiry. This results in a decrease in efficiency and increases the company's operational cost. As such, banks and e-commerce companies are turning to the use of chatbots to mitigate these issues. Cui et al. introduced a customer service chatbot that helps to improve user's online shopping experience and is more practical and cost-effective when answering repetitive questions, freeing up human support staff to answer much higher value questions [1]. In addition, Kulkarni et al. also proposed a chatbot system for banks using natural language processing (NLP) and machine learning to solve the traditional customer service problems [2].

In general, it is more costly to acquire new customers than it is to retain existing customers. Retaining customers requires banks to provide the right services quickly and conveniently. As more companies start to adopt artificial intelligence solutions within their organization to improve customer experience, one common approach is using NLP for sentiment analysis and opinion mining as demonstrated by Balu [3]. Apart from NLP, neural network application in the banking and finance industry have also been gaining traction in recent years. For example, artificial neural networks have been used for stock price prediction and credit risk scoring. [4, 5, 6, 7].

As intelligent and efficient customer service is increasingly becoming a basic expectation from customers, relevant research work can be broadly categorized into two threads:

- Intelligent customer service in call centre
- Customer behaviour prediction

Matching customer call to a Customer Service Officer (CSO) is a supply and demand problem and having optimal trade-off between available resources and customers' satisfaction is non-trivial [8]. Instead of modelling this as a resource allocation problem, it would be better to reduce demand by predicting what the customer would call for and perform pre-emptive intervention. However, even with the most accurate prediction, the execution needs to be done well for customer to adopt the solution. Identifying and predicting customer behaviour using artificial neural networks generally yields better result than traditional discriminant analysis [9].

This paper introduces a digital customer service use case to provide personalized digital services through the integration of

* Corresponding author

IVR system, a machine learning engine, the Internet Banking platform and a chatbot for customer servicing. The objective is to migrate customers calling into call centre to digital servicing channels. In order to meet the objective, a three-step approach is required:

- Detect customer needs and promote solutions before the customer calls the bank
- Promote solutions in IVR system, SMS system, Internet Banking platform and chatbot
- Deliver high-demand and sticky digital solutions, with the focus on the first theme

II. System Integration and Architecture

This section introduces the methodology used, the systems integrated in the use case and the architecture design.

A. *Methodology*

A topic model using historical call transcript data was first built to understand the customer's intentions of calling. The topics were then used as prediction labels in a neural network model to predict the customer's top 3 intentions in their next call (within 10 days). When a customer calls in, the prediction result will be played as a voice prompt, asking if he/she would like to try the recommended digital services. If the customer presses '1', a SMS with a URL to trigger the bank's chatbot on the Internet Banking website will be sent to the customer's registered mobile number. The customer can click on any of the top 3 predictions displayed in the chatbot to start a digital service in the form of a Guided Conversation (GC). The GC transforms traditional banking services that requires form-filling to guided digital service completion by allowing the customers to complete most services with a few buttons clicks. Fig. 1 shows an example of a digital service for blocking card.

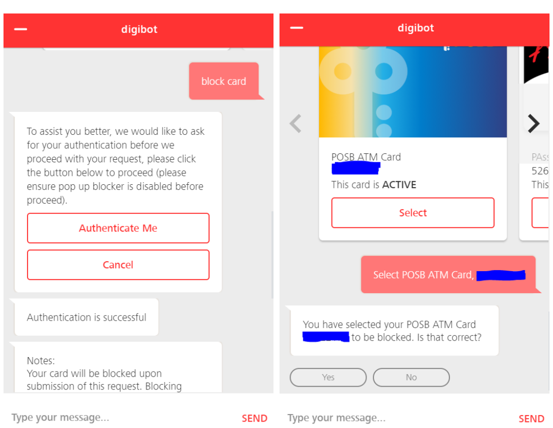

Fig. 1. Guided Conversation (GC) in the chatbot for blocking card service

B. *System Architecture*

This use case involves the integration of multiple systems in the bank to support the personalized customer journey. Specifically, digital services interventions are mainly executed through service SMS to customers that are prompted during their IVR journey, as well as chatbot pop-up in the bank's web pages and mobile applications of the customers' personal devices. Within the chatbot, embedded digital services such as GCs are personalized to the customer based on model prediction results and their recent online activity. Fig. 2 shows the main systems integrated in this use case and the corresponding flows of API calls between systems.

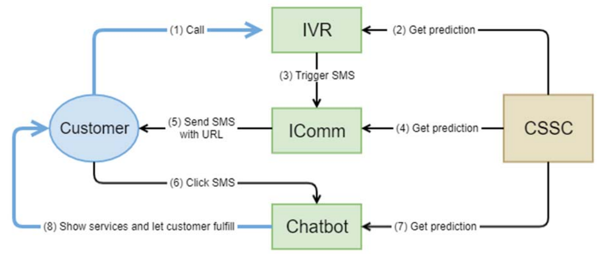

Fig. 2. Overview of architecture of systems integrated for the use case. IVR: Interactive Voice Response system. CSSC: machine learning engine serving predictions. IComm: SMS system. Chatbot: chatbot serving digital customer services in the form of GC.

III. Topic Modeling on Call Transcripts

Document classification tasks are commonly separated into three approaches: supervised, unsupervised and semi-supervised. Ji et al. developed a customer intention aware system for document analysis using a semi-supervised learning approach [10]. In our paper, we applied topic modelling, a form of unsupervised learning, to analyse the intentions of customer calls. Recorded voice call dialogues between customers and CSOs were transcribed and converted to text through a vendor service. With these transcripts, topic modelling [12] was applied to identify the different categories of customer call intents. These categories are subsequently used as prediction labels in the digital service needs prediction model introduced in the next section.

A. *Topic modeling on voice transcripts*

The call centre has its own set of call categorization labels classified by CSOs when they handle calls from customers. Although most of these labels were relatively accurate, it consists of over hundreds different categories that sometimes overlapped. Furthermore, these labels often do not reflect the specific customer's request or CSO's resolving action, which could have helped map the calls to possible digital customer service solutions. Hence, a new set of call categorization using topic modelling was generated to accurately reflect the service solutions required by customers. This new classification also provides insights towards improving the design of current and new services. An unsupervised topic modelling approach was adopted as it is not feasible and sustainable operationally to perform regular manual analysis on millions of transcripts produced annually along with constantly shifting customer service needs.

B. *Stop words removal*

Based on the nature of the conversations between CSOs and customers, a custom set of stop words and phrases for removal was first consolidated, along with custom regular expressions to mask or remove sensitive information such as account number, phone number and passport numbers. Next, informative but often wrongly transcribed terms were also identified and corrected to improve consistency. Lastly, lemmatization using

WordNet Lemmatizer from Natural Language Took Kit (NLTK) was then applied as a final pre-processing step.

C. *Term frequency*

Following the pre-processing steps, transcripts are transformed into term frequency vectors. Due to the relatively narrow vocabulary of banking terminology and spoken conversation, many informative terms have high document frequencies especially for common service needs. Usage of inverse document frequency (IDF) term weightings may not be appropriate as the importance of these informative terms will be reduced and was not applied.

D. *LDA topic modeling*

The vectors are used as input for topic modelling using Latent Dirichlet Allocation (LDA) [12]. The call centre receives a few million calls per year. An initial sampled set of over 800,000 transcripts that spanned over 6 months was used as training data. Multiple iterations of hyperparameter tuning was performed using grid search, along with analyses of the results using pyLDAvis, a python port of LDAvis [13] which introduces a definition of relevance of terms towards topics in addition to the visualization system that allows exploration and understanding of a fitted LDA model. Each set of hyperparameters were evaluated qualitatively using pyLDAvis for consistency of the topics' most relevant terms and level of granularity topics provided. More promising sets of hyperparameters were further evaluated by reading randomly sampled transcripts for each topic to determine how well the highest assigned topics aligned with transcript content.

The final number of topics after tuning is 50, which was deemed to provide a good balance between granularity of topics for caller service needs, good term consistency within topics, and topic alignment with transcript. Fig. 3 shows the visualization of topics in terms of cluster/topic size and list of important words.

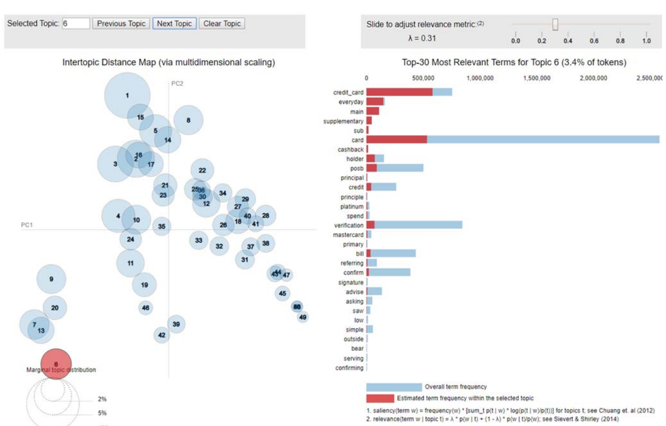

Fig. 3. Visualization of the 50 topics generated by LDA.

Each topic is represented by the top key words from the topic model. For example, the topic with top key words 'card', 'block', 'replace', 'account', 'application' et al. is related to card blocking or replacement. Examples of the topics are in Table I.

TABLE I. SAMPLE TOPICS GENERATED

Topic	Top 10 keywords	Description
1	payment make pay bill credit_card paid date due outstanding made	Enquiry on details on expected bank payment/status of settled bank payment
2	transaction merchant online cent purchase pending made claim history recent	Dispute or enquiry of details of known pending or posted transaction made
3	back money department refund contact credited return relevant colleague report	Enquiry for status of ad-hoc fund transfer/change in recurrent fund transfer
4	credit_card card everyday main posb holder verification supplementary credit bill confirm	Enquiry for general details of credit card (transactions, balance, points accumulated, cashback, limits, fees)
5	card lost passion replacement block debit_card atm_card immediate hotline report	Block or replace card
…	…	…

To verify the quality of the topics generated, 100 transcripts were randomly sampled for each of the 50 topic and were read by the customer servicing team to confirm that the transcript matches the topics assigned. The overall matching accuracy was about 85%.

From the 50 topics generated, 14 topics were selected based on the topics' call volume and mapped to the existing digital banking services that can be provided in the chatbot. These 14 topics covers about 55% of the total call volume. As the customer centre team continues to build new digital services, the number of services is expected to increase to 25 next year, thus covering more than 80% of the calls.

The transcripts data were processed monthly using Apache Spark, which is approximately 5 times faster than a Python version. The topic model was trained in Python and stored as artefact, which is used to generate data labels for the subsequent neural network model for customer service needs prediction.

IV. Call Intent Prediction Model

This section introduces the call intent prediction model. Due to the complexity of the data from various banking systems, the data processing pipeline and feature engineering pipeline processes a large volume of data and were implemented mainly using PySpark. The neural network model was implemented using Keras and Tensorflow. Apache Airflow was used for Spark and Tensorflow job orchestration and MLflow was used for model performance logging.

A. *Data sourcing*

The data used includes consumer banking data such as customer profiles, transactions, account balances, customer service and banking platform data. These data can be broadly categorised into two types: transaction data and non-transaction data. Transaction data includes banking transactions for credit cards, debit card, loans, call records, mobile/internet banking transactions, email records, SMS records et al. Non-transaction data includes customer demographics, profiles, product or account status updates et al. These data are pulled from over 30 data sources and are incrementally loaded daily from the respective source systems and takes 1-2 days to be fully ingested into the data warehouse. The incremental load status verification process was included in the beginning of the data pipeline to ensure that the most recent data are loaded. Processed data are stored in on-premise S3 buckets.

B. *Data processing and aggregation*

In the data processing step, raw tables are read as parquet files using PySpark. The transaction data are filtered to a common time period before extracting the daily aggregated data through relevant filters, joins and groupings. The type of aggregation used includes count, sum, minimum, maximum, mean and flag (Boolean value of flag). An end-of-month snapshot of customers demographics, profiles et al. are also captured as non-transaction data. Both aggregated and snapshot data are then saved in a sparse output format as it is more space and memory efficient, as well as less computationally expensive to perform further processing on. The sparse output has three columns: the customer identifier, feature name and feature value. The feature name column captures information such as data source, feature generated and type of data aggregation. For example, a feature generated from the SMS data source would be the number of card activation SMS messages sent to a specific customer on a specific day. Hence, the feature name in the sparse output would be '*sms__num_card_activation_msg__count*'.

Once data processing completes for each data source, a success status file will be created in the data warehouse which then triggers the data validation pipeline. The data validation pipeline checks for null entries in each columns and missing data from various source systems, which may occur due to upstream data loading or processing.

C. *Feature Engineering*

The processed data from each data source are combined to generate customer-level dense feature set through a time-windowed aggregation approach.

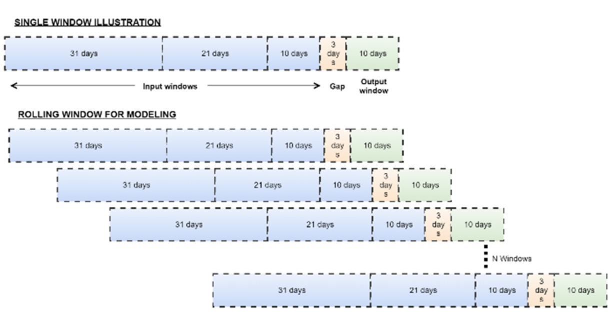

Fig. 4. The time-windowed feature engineering approach. Blue color: input data period. Orange: gap period. Green: output period.

Features were created using a backward-looking approach whereby the output window of predicting the customer's top 3 call intentions was fixed at 10 days. A 3-day gap between the input feature window and output prediction window was also taken into consideration to account for the data ingestion lag time and modelling run time during production. The model input features were then generated by aggregating the processed data with 3 window sizes (10, 21, 31 days) to capture the short-, medium- and long-term features. As shown in Fig. 4, given a gap time of 3 days and output prediction window of 10 days, the input features will be generated from the past 4^{th} day to 66^{th} day.

There are three main categories of input features: transaction features, recency features, and profile features. Transaction features are simple aggregations of previously processed data at the various input window levels. The aggregation function for each feature depends on the type of aggregation used in the data processing and aggregation phase. Recency features such as the number of days since the customers' last credit card transaction are features that describes the customer behaviour and are also generated from transaction data. Profile features are non-transactional data taken from the latest end-of-month snapshot processed data. In general, there are over 4000 features generated from these three broad categories.

The output feature was set to be a binary vector that indicates whether the customer has called for any of the top 14 topics within the output window period. Currently, a call is tagged to one intent/topic and a customer might call multiple times with different intents within the 10 days prediction period.

During model training, 12 non-overlapping output windows was created using the sliding window approach to obtain 200,000 training samples. Of which, it was observed that there was a high number of highly sparse features that may lead to computation difficulties during the modelling stage. Hence, features that contains missing or constant value for more than 99.5% of the samples was removed to prevent the model from overfitting to a rare occurring behaviour.

During the model scoring phase, this feature engineering process was applied to the entire customer base which is approximately 5 million.

D. *Call intention prediction model*

In order to predict the customer's call intention in the next 10 days, a prediction model was built using a multi-label classification model with a ranking component.

1) *Multi-label classificaiton*

A customer might call for multiple issues within a call, or multiple calls within the output window. Hence, predicting the services that a customer may require when calling in was formulated as a multi-label classification problem. Depending on the features available in the training period, the number of features can range between 4000 to 5000. Hence, each time a new model is trained, the available features in the training data will be stored in a metadata file that will be used during scoring.

2) *Hyper parameter tuning and network architecture*

A multi-task neural network was used to perform the multi-label classification, and hyperparameter tuning and searching was conducted on the following parameters to achieve the final neural network architecture shown in Fig. 5.

- Number of hidden layers after input layer: [512, 256, 128], [512, 256], [512]

- Number of layers in multi-task components: [10, 10], [10, 20], [20, 20], [20, 30], [30, 30]
- With or without residual connection component
- Residual component size: 140, 280, 700 (multiples of 14)
- Dropout: 0.2, 0.3, 0.5

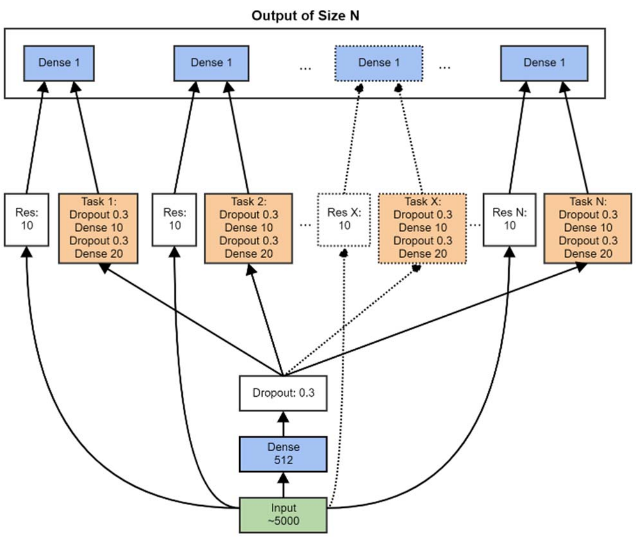

Fig. 5. Network architecture of the multi-task neural network model

3) *Ranking*

After training the multi-label classification model, the predicted confidence output from the model was used to rank and select the top 3 digital services for each customer. Several ranking criteria were evaluated using out-of-time period data during the training phase, and the best performing ranking metric was the original confidence score from the neural network.

4) *Early stopping*

Learning rate decay was included, and early stopping was conducted if performance doesn't increase for 5 epochs.

5) *Evaluation Metric*

A customized evaluation metric, HitRate@K was used to measure the performance of the model. This metric is finalized after discussing with the business team to reflect the impact of the whole use case.

$$HitRate@K = \frac{1}{n}\sum_{c=1}^{n} H_K$$

$$H_K = \begin{cases} 1, & if\ any\ 1\ of\ 3\ predictions\ is\ correct \\ 0, & otherwise \end{cases}$$

where n is the number of customers who called for any of our 14 intents/topics selected, H_K is the hit rate of if any of the 3 predictions matches with the customer's actual calling intent.

V. Experimental Results

This section presents the experimental results, including model performance evaluation, model analysis results and impact evaluation after deployment.

A. *HitRate@3*

The performance of our model was evaluated using an out-of-time period test data, and compared against a baseline model, logistic regression model and random forest model. The baseline model is simply recommending the most popular 3 digital services to all customers. To solve the multi-label classification issue, one binary classifier was trained per label for the regression and random forest model, and the 14 binary classifier scores were then combined to rank and select the top 3 services for each customer.

Depending on the infrastructure capacity and business decisions, the proportion of customers targeted varies, which in turn influences the HitRate@3 evaluation metric. The whole customer base is ranked by the maximum of the predicted score of 14 outputs from the model and then the top X% of customers are selected and evaluated based on the HitRate@3 metric, as shown in Table II.

TABLE II. HitRate@3

Model	HitRate@3 (Out-of-Time Test)		
	Top 10%	***Top 30%***	***Top 100%***
Baseline	26.91%	26.91%	26.91%
Logistic Regression	49.09%	43.64%	40.00%
Random Forest	55.45%	49.82%	43.64%
Our model	**65.95%**	**56.33%**	**43.89%**

Table II shows the HitRate@3 for various models at different targeted population threshold. The neural-network model outperforms the other models across all population thresholds and achieved a HitRate@3 of 66% when the top 10% of the customer population was targeted. This translates to a 2.5 times lift compared to the baseline model. Our business users eventually decided to target the top 30% of customer population as it covers a larger population base while still maintaining a relatively high HitRate@3. By providing personalized call centre digital service experience to the top 30% of customers, the model can correctly target ('hit') 56% of these customers.

The evaluation above was based on a training dataset constructed in the first phase of the project. In subsequent phases, using datasets with longer training period and model optimization increases the HitRate@3 of our model to 65% for the top 30% of customers.

B. *Different Gaps*

In the initial phases of this project, a 3-day gap window was used to account for the data ingestion lag time and modelling run time. This implies that the customers' predictions were made with input features that were at least 3 days ago. Since hit rate is expected to improve if predictions were made with more recent input features, it would be ideal to reduce this gap. However, the gap can only be reduced if there were time improvements made to the data pipeline. Therefore, an experiment was conducted to determine how hit rate varies with different gap windows. For each experiment, the output window size was set to 3 days, and gaps between the input and output windows were set to [0, 3, 6, 9, 12] days.

From the experiment results shown in Table III, evaluation metric of HitRate@3 decreases slightly as gap increases. After several iterations of the experiments, a minimum gap of 3 days was decided to be a good compromise and achievable for the data pipeline, and a maximum gap of 9 days was allowed. This means that we have to complete our data ingestion, data processing and modelling in 9 days before the results deteriorate.

TABLE III. EXPERIMENT RESULTS FOR DIFFERENT GAPS

Gap (days)	Baseline@3		Hit-Rate@3	
	Training Period	***Testing Period***	***Training Period***	***Testing Period***
0	23.0%	23.8%	30.5%	30.3%
3	22.6%	23.8%	29.4%	30.0%
6	22.9%	23.8%	29.8%	29.8%
9	22.8%	23.8%	28.7%	29.5%
12	22.8%	23.8%	29.1%	29.4%

C. *Within 10 days, the prediction performance drop*

An output window size of 10 days was used, implying that predictions made were only valid for a 10-day period. By observing the decay of HitRate@3 across individual dates for the 10-day period, the ideal output duration can be determined, although it is still bounded by the model refresh rate.

In this experiment, predictions were only considered for the top 30% of callers with the highest prediction score. In addition, HitRate@3 for individual dates was also normalized by the number of calls on that day. To observe the decay across 10 days, the output window size was set to 10 days, with 12 windows for training and 3 windows for testing, and the gap window was set to 3 days. Hit rate for individual dates will be observed for the first 10-day window of testing output period.

As shown in Fig. 6, hit rate generally decreases across the 10 days, but decays after the 7^{th} day. Hence, it is ideal to refresh the model with a new 10-day period prediction within 7 days.

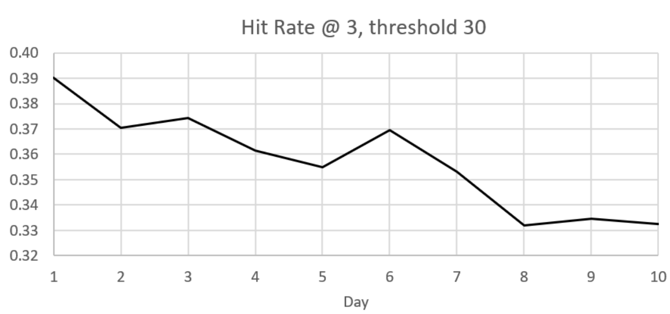

Fig. 6. HitRate@3 across each day for the 10-day output period.

D. *Feedback Data Analysis and Impact Analysis*

Digital solutions from services recommended to customers is available on the banks' chatbot. Customers can access these solutions through different channels such as SMS redirection through IVR system to chatbot, Digital Service Menu (DSM) redirection to chatbot, internet banking chatbot or mobile banking chatbot. Instead of going by the approach of measuring the amount of call reduced or redirected, it is more sustainable to measure the impact metric by how the digital solutions were used and fulfilled by the customers. Channels leading to the digital solutions are analysed separately to determine the final impact metric to call centre.

1) Callback or pass through to customer service officer

A customer might choose not to use the digital solution provided and would either be redirected to a CSO in the same call or call back again. To analyse the response rate from this channel, drop-offs are analysed at each step of the customers' journey from their press-1 click and SMS-click. The usage and fulfilment of digital services of these customers are tracked and churn calls are also accounted for to determine the success of customers' adoption of the digital solution.

2) Digital Service Menu

Customers might drop off from chatbot and proceed to the bank's central digital customer service webpage for digital services. The customers' usage, fulfilment and drop-offs of the digital services is tracked to determine the success of customers' adoption of the digital solution.

3) Chatbot transcript

Digital solutions are offered through chatbot GCs, which give customers' automated prompts and can help customers perform banking operations through the chatbot. GC can be triggered through 3 methods: (1) Chatbot Verbatim, whereby chatbot shows the digital services based on predictions for the customer, and clicking on the prompt will either lead to a help and support page or GC depending on the customer journey, (2) Utterance, whereby GC is triggered through free text entered in the chatbot, (3) Digital Service Menu (DSM), whereby the services clicked will lead to a chatbot popup enquiring if customers want to proceed with the banking operation

From this flow analysis, it can be determined if the customer fulfilled his/her service request, or attempted to do so, and at which stage did each customer drop off. Combined with the churn call analysis, it is also possible to check if the customer called back after attempting to fulfil the digital services.

By combining the analysis of the 3 channels involved, the model performance was found to be approximately 65% accurate in predicting customers intentions. This was measured by checking the feedback data of all three channels to determine if a customer performed any digital services related to the 3 predictions provided by our model.

VI. TECHNOLOGY STACK

Apache Spark was chosen as the framework for big data processing in order to handle approximately 3 TB of data across 30 data sources, and Alluxio with on-premise S3 was used as the object storage layer. The neural network model was implemented using Keras and Tensorflow. Apache Airflow was used for Spark and Tensorflow job orchestration. An example Airflow pipeline is shown in Fig. 7.

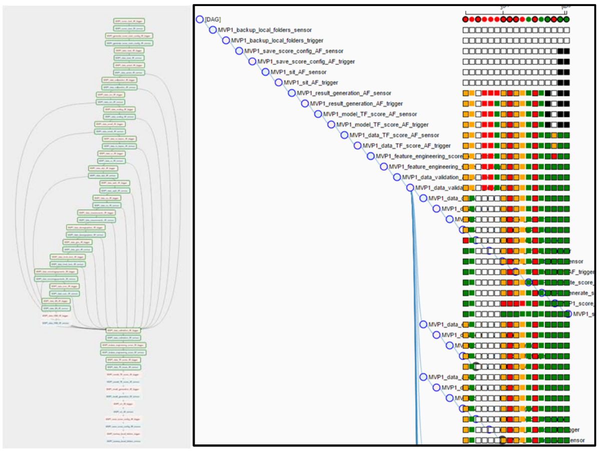

Fig. 7. Apache Airflow DAGs of scoring pipeline

VII. CONCLUSIONS

In this paper, we present a use case for digitizing customer journeys in the consumer banking call centre by providing a personalized customer service experience through an integrated solution and a machine learning platform. A multi-task neural network was trained to predict customer's call intent and subsequently migrate the customer to digital channels. Through this machine learning solution, the bank saved a significant amount of operation cost in the form of reduced number of calls to the call centre. In 2019, the broader call reduction initiative of the bank (including other parallel workstreams such as campaign-based targeting and IVR-based targeting) reduced about 400,000 calls which translates to about 2 million SGD cost savings.

ACKNOWLEDGMENT

We would like to thank the collaborating departments/teams and sponsors within DBS bank: Kailash Ramalingam, Sameer Gupta, Mike Power, Ali Jinah Mohanmed Yusuf, Jodon Santos Aguiling, Zhiyang Lim, Hui Li New, Xu Ren Sng, Shaoyao Cai, Yi Hock Yap, Chen Guo, Angie Hong Pei Phua, Gautam Gorki, Vinod Bhat, Joanne Hui Ping Chan, Brandon Ho, Henry Leqi Soon and other team members.

REFERENCES

[1] L. Cui, S. Huang, F. Wei, C. Tan, C. Duan and M. Zhou, "SuperAgent: A Customer Service Chatbot for E-commerce Websites", *Proceedings of ACL 2017, System Demonstrations.*, Vancouver, Canada, 2017, pp. 97-102.

[2] Chaitrali S. Kulkarni, Amruta U. Bhavsar, Savita R. Pingale, Prof. Satish S. Kumbhar, "BANK CHAT BOT - An Intelligent Assistant System Using NLP and Machine Learning", International Research Journal of Engineering and Technology, Volume 4, Issue 5, May 2017.

[3] E. Angelini, G. D. Tollo and A. Roli, "A Neural Network Approach for Credit Risk Evaluation", The Quarterly Review of Economics and Finance Volumn 48, Issue 4, 2008, pp. 733-755.

[4] Youngohc Yoon and G. Swales, "Predicting stock price performance: a neural network approach," Proceedings of the Twenty-Fourth Annual Hawaii International Conference on System Sciences, Kauai, HI, USA, 1991, pp. 156-162 vol.4.

[5] Y. Li and W. Ma, "Applications of Artificial Neural Networks in Financial Economics: A Survey," 2010 International Symposium on Computational Intelligence and Design, Hangzhou, 2010, pp. 211-214.

[6] A. Natasha, D. D. Prastyo and Suhartono, "Credit Scoring to Classify Consumer Loan using Machine Learning", AIP Conference Proceedings 2194, 020070 (2019).

[7] R. Balu, "Using Natural Language Processing to Improve Customer Experience", International Journay of Computer Science Engineering (IJCSE) Volume 8, Issue 1.

[8] C. Long, L. Zining, L. Xiaolu et al. "CRSRL: Cusomer Routing System Using Reinforcement Learning", IJCAI, Macau, China, August 2019.

[9] S. Laura. "Predicting Consumer Behavior with Artificial Neural Networks", Emerging Market Queries in Finance and Business, Romania, 2013.

[10] J. Ji, D. Kunita, and Q. Zhao, "A customer intention aware system for document analysis", IJCNN, Bacelona, Spain, July 2010.

[11] S. B. Garavaglia, Health Care Customer Satisfaction Survey Analysis Using Self-Organizing Maps and "Exponentially Smeared" Data Vectors. IJCNN, 2000.

[12] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," J. Mach. Learn. Res., vol. 3, pp. 993--1022, 2003.

[13] Sievert, C. and Shirley, K. "LDAvis: a method for visualizing and interpreting topics", In Proc. ACL Workshop on Interactive Language Learning, Visualization, and Interfaces, 2014, pp. 63-70.