

The Segment of One

Real-Time Context Awareness and Hyper-Personalization

TATA
CONSULTANCY
SERVICES

Lipika Dey
Rekha Singhal

IN BRIEF

While digital-born retailers are the poster children of hyper-personalization, no solution is perfect. Typically, soft computing methods are suitable, as data is incomplete and situations are uncertain.

We look at two approaches: recommendation for the next best offer and information fusion for recommending the next action.

As personalization gets more ambitious, its complexities are likely to keep researchers in the area of Computational Intelligence busy for quite some time to come.

Amazon started personalizing our shopping experience twenty years ago. Now, Netflix creates our own home theatre, experimenting even to the extent of changing billboard images. For instance, if you watch *Doctor Strange* and *The Imitation Game*, it offers the next Cumberbatch movie; no surprises there. However, what is surprising is that it may change the poster of *Twelve Years a Slave* from showing Chiwetel Ejiofor (the lead) to showing Benedict Cumberbatch (a minor role as William Ford), just so you will pick it up. You may indeed be happy to watch a movie you never knew had Cumberbatch in it; and Netflix sold you a new genre of content. Digital businesses are trying to do special things for hundreds of millions of customers

and over millions of products, *over each interaction*. The segment of one has arrived.

Powering the segment of one

The power behind the segment of one lies in predictive analytics powered by big data technologies. These work with large volumes of transactional and operational data containing a mix of structured and unstructured content. Predictive analysis helps in identifying different ways to keep a customer engaged on a site. Knowledge and insights extracted from analysis of a customer's past history, including his or her browsing history, past transactions, items abandoned in their carts, buying

Fact File

TCS Research: Data and Decision Sciences

Outcomes: Automatic Email Classification for Support Centers, Early Warning Systems, Next Best Offer Recommendation Systems, Opinion Mining and Sentiment Analysis from Consumer-generated Content, Information Extraction, Process Analytics, Social Media Behavioral Analytics

Principal Investigator: Lipika Dey

Academic Partners: IIT Kharagpur

Techniques used: Machine Learning, Natural Language Processing, Deep Learning

Industries benefited: All industries

Patents: 10 filed, 2 granted (USA)

Papers: 70

behavior, etc., are combined with data tracked during a current visit to generate recommendations and targeted offers. Adding profile information obtained from customer relationship management (CRM) systems or gathered through surveys or customer-provided feedback is often found to significantly improve the results.

Despite the strides personalization has made, it is still far from perfect. As personalization becomes more widespread, the many complexities of dealing with this many-headed monster are coming to the forefront, which are likely to keep researchers in the field of Computational Intelligence busy for quite some time to come.

Micro-segmentation

With humongous computational resources at their disposal, eCommerce establishments are taking personalization to the logical extreme, where micro-segmentation culminates at the "group of one". It is also possible to do powerful long-tail analysis,

that takes into account seemingly unpopular but niche and novel products that interest only specific customer segments.

Historically, business analysts have tried to segment customers into groups based on similarity in behavioral patterns. This is still a valid grouping. Based on historical and current action gathered through real-time tracking, micro-segmentation aims at finding small groups of customers, who exhibit fine-grained behavioral similarities. Coupled with customer profile data and ample details about available inventory, predictive analytics models built for each of these segments can provide the right inputs for targeted marketing to optimize nuanced customer experience. These models can also generate context-sensitive next best actions or offers.

Combining micro-segmentation with other ancillary data about cross-sell, upsell, etc., has led to the development of rich predictive models that are suggesting the next best offers to customers with high success rate.

.....

Probabilistic reasoning techniques are key to designing personalized content delivery systems.

.....

Computational intelligence and hyper-personalization

Computational Intelligence usually refers to an array of techniques that are employed to emulate complex real-world phenomena, which are difficult to represent as pure mathematical models. Traditionally, soft computing methods included Fuzzy and Rough Logic, Evolutionary Computation, Machine Learning, and Probabilistic Reasoning, all of which can be modeled to work with uncertainties and incompleteness of knowledge. These methods were found to be highly suitable while working with unstructured data such as images, videos, or text. Soft-computing methods help in extracting data about personal preferences of users, based on product images or descriptions chosen by them, textual feedback given, content of documents read, and also from content generated in the form of status messages or communications. Since user

actions are rarely deterministic and rather largely dependent on the surrounding physical environment and mental states, probabilistic reasoning techniques play a key role in designing personalized content delivery systems.

High-performance computing and hyper-personalization

The dream of hyper-personalization can be attained only if it is coupled with a delivery platform that can track millions of users and predict their next needs or actions, based on past history. High-performance computing practices aggregate computing power to deliver application performances that are not possible to be delivered otherwise over typical workstations. Building application-agnostic hyper-personalization platforms that can adapt to multiple recommendation engines and domains, while scaling up robustly, is a research challenge.

A Next-Best Offer System

TCS has developed a high-performance prescriptive system and configurable framework for generating personalized offers for customers, while optimizing business objectives. This can be easily customized for any B2C system, such as for retail, telecom, hospitality, and insurance. The capabilities include:

1. Configurability according to domain or segment: The configurable framework captures business data semantics and its sensitivity to customer repeat probability

for a given business objective, using metamodels, so that it can be easily configured for any B2C system. For instance, you can configure it for your apparel line differently from your electronics one and, likewise, for flight tickets differently from trains; corporate data users or individual voice subscribers; and so on. Metamodels use a set of scalable library functions for building recommendation models, using deep learning and machine learning algorithms. Figure 1 shows the use of metamodel for building domain-specific models.

2. Ensemble technology: The recommendation model is built as an ensemble of gradient regression machine learning and long short term memory (LSTM) deep learning classifiers, trained on users' historical action or transaction data. To generate the right offers at the right time, the system creates temporal and non-temporal features for users, which are fed to the LSTM and the machine learning algorithm to build the recommendation model. LSTM algorithm implementation is optimized to reduce inference latency, while improving the user conversion rate.

3. Scalability: The system is co-located in a B2C environment such that it senses all actions or transactions across different channels and responds with the next best offer or action. The prescriptive system can cope easily with large workloads and high performance demands of the B2C system. To ensure low recommendation latency and a scalable architecture, the system is built using tuned big data technology stacks, which can be deployed on the cloud as well. The system architecture is shown in Figure 1.

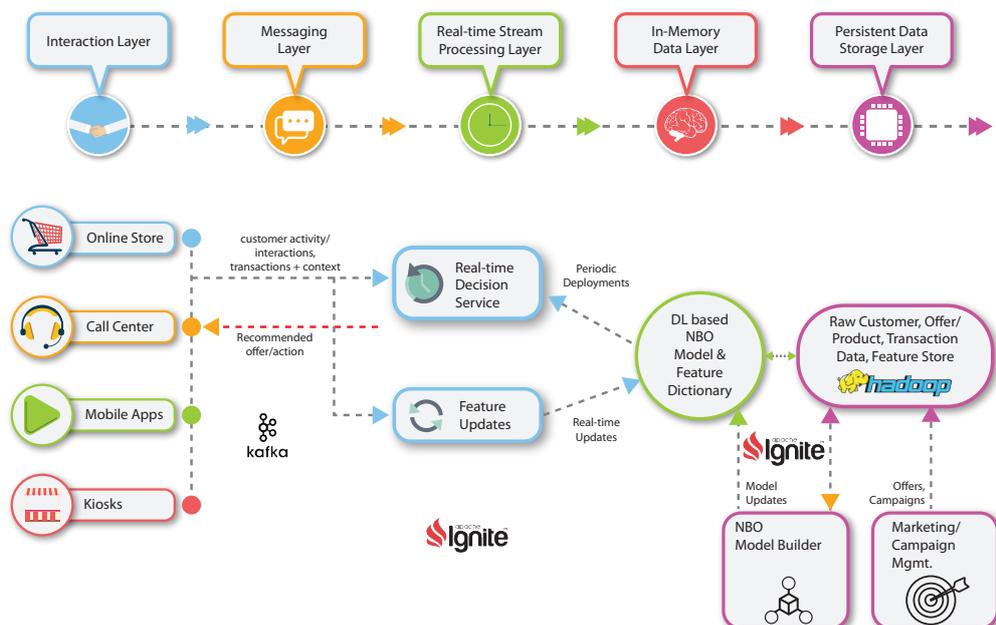


Figure 1: The TCS intelligent prescriptive system architecture

The system is planned for use in retail, telecom, and hospitality, for user's home page personalization. The system is under test for integration with TCS' telecom digital platform "Hosted OSS/BSS solution" (HOBS) for giving the right telecom package to customers, based on the history of their usage of services for call, data, and tickets raised in the past. This will improve customer experience for our telecom clients.

Projects are underway for retail and hospitality clients.

Role-based personalization in the enterprise

Digital-born retail has been the poster child for personalization. But for businesses with global footprints, the enterprise and its extensions (vendor, partners, supply chains, customers), information overload has created the need for personalized content delivery.

Contextually intelligent systems fuse multi-structured information gathered from heterogeneous sources in real time to predict the next likely action of the user, and thereby, optimize the uptake possibility of the recommended content, not just for a single application but for a multitude of applications.

Personalization—future challenges

While the recommendation models discussed here appear to be fairly generic in design, their applications across all fields have not seen the same level of success as eRetail. The primary reason for this is the lack of mature behavioral models for other domains. Though the challenges for each area appear to be different from the others, there are some common problems that

personalized recommenders are trying to address.

(i) Handling novelty:

Though repetition and frequency derived from personal preferences drive the bulk of personalized recommendations, sectors like travel and hospitality, and entertainment, which have seen a huge surge in the number of online transactions, have had limited success of personalization, to date. A major reason for this is the incapability of the current models to incorporate novelty as a user preference. Novel or unique elements, which have no history of being chosen or, more importantly, are not even similar to products chosen earlier, cannot be accommodated easily in the predictive model. However, novelty is a key factor that drives customer choices for these sectors. This is referred to as the cold-start problem in the science of recommendation. Determining the novelty of an item is comparatively easier than determining which novel items to recommend to whom, how, and when. These are some

Contextually intelligent systems fuse multi-structured information from heterogeneous sources in real time.

Personalization for information recommendation

We designed a system to deliver contextually relevant business-critical information gathered from open sources for enterprise users for augmented intelligence. It combines individual preferences of users, along with role-based personalization, to determine the criticality of open-source information that is continuously gathered and processed at the back end. Information collected from multiple sources is fused using computational intelligence. The relevance of a piece of information to a particular

user at a particular time is derived through contextual intelligence. The causal analytics framework exploits business rules and domain ontology, which are plugged into the framework during system configuration. The system is trained to generate alerts and alarms as well as to recommend the next best action to users. User reactions are monitored and system intelligence is continuously enhanced through reinforcement learning.

During the implementation of this system, personalized information recommendations are generated for a human wealth manager who manages several customer portfolios and is interested in receiving real time alerts about critical business events reported for companies with these portfolios, to generate appropriate advisories for clients. Alerts and action suggestions for the wealth manager are delivered through multiple platforms such as emails, personal messaging services, and dedicated mobile apps.

In another instance, we personalized a news alert generation system for insurance underwriters who compute premiums for corporate clients. This system keeps track of all news that can have negative impact on company financials and performs a risk assessment for each reported event, taking into account social media reactions to these events, if any. The risk score, along with its genesis, is then presented to the human underwriter for appropriate action. The system has an inbuilt feedback mechanism that provides feedback to the risk assessment module.

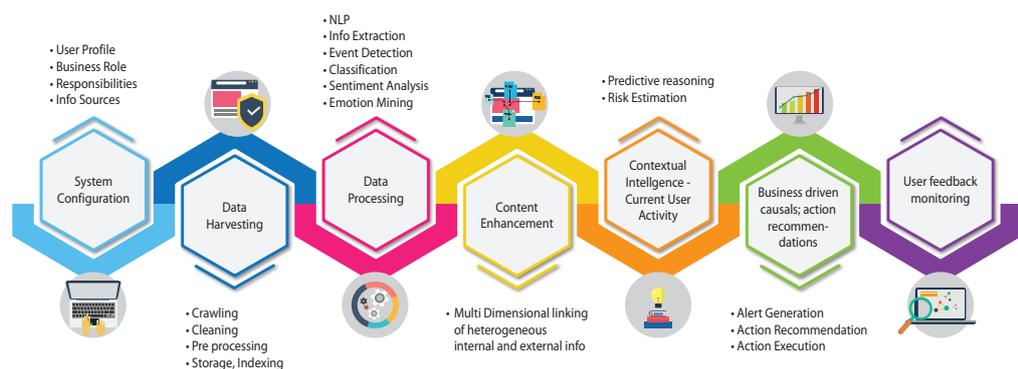


Figure 2: A Framework to deliver business news alerts and action recommendations, based on enterprise profile

of the major challenges that personalization systems are trying to deal with.

(ii) Objective-driven information recommendation:

Recommending information or textual content to users faces another type of challenge. Due to the explosive growth in the amount of available digital information, filtering systems that can detect the right information for a user are in great demand. Along with user preferences that can be learnt from their past behavior, personalization, in this case, also has to determine the intangible quantity that can measure the worth of a piece of information in the user's personal and professional life. While role-based information filtering systems are being designed, the challenge lies in formally specifying the concepts of role and responsibility and integrating them to create an objective-driven information filtering framework.

(iii) Security and privacy concerns:

The recent success of personalization in many sectors can be attributed to the data that

a user shares with content providers voluntarily and involuntarily. With rising concerns over privacy and security breaches, research in the area of personalization will have to deal with this seriously. An immediate point of concern is the enforcement of the new EU-wide privacy rules, the General Data Protection Regulation (GDPR) which is expected to create a seismic shift not only in the way personal information is defined but also in how it is collected, processed, used, and transferred. Several key data elements, like sexual orientation, organizational memberships and so on will require the explicit consent of users before they can be used. An important technological aspect that would be affected by the introduction of GDPR is the regulation on using "persistent online identifiers", such as cookies, which are small text files storing tiny pieces of data in a user's browser. What it would mean for personalization platforms, which store these identifiers, will have to be seen, since violators of GDPR provisions are likely to face huge penalties.



Lipika Dey

Lipika Dey is a Chief Scientist at TCS Research and Innovation and heads the Analytics and Insights Practices. Her research interests are in the areas of Natural Language Processing, Text and Data Mining, Machine Learning, and Semantic Search. Along with publishing prolifically in academic conferences and journals, Lipika has been invited to speak at many business and technology conferences around the world. Lipika did her Integrated M.Sc. in Mathematics, M.Tech in Computer Science and Data Processing, and Ph.D. in Computer Science and Engineering—from IIT Kharagpur.



Rekha Singhal

Rekha Singhal has worked with both corporate and academic research. Her research interests are High Performance Data Analytics systems, Big Data Performance, High Performance Computing, DB Performance Modeling, Query Optimization, Storage Area Networks and Distributed Systems. She headed the development of the disaster recovery product “Revival 1000” which was launched at the Centre for Development of Advanced Computing. The invention was also nominated for the NASSCOM 2011 award. Rekha has around 50 presentations and publications in international and national conferences, workshops, and journals. She has patents filed and granted in India and international territories. Rekha did her B.E. in Computer Science from Netaji Subhash Institute of Technology. She has received her M.Tech and Ph.D., Computer Science from IIT, Delhi.



All content / information in The Segment of One is the exclusive property of Tata Consultancy Services Limited (TCS) and/or its licensors. This publication is made available to you for your personal, non-commercial use for reference purposes only; any other use of the work is strictly prohibited. Except as permitted under the Copyright law, this publication or any part or portion thereof may not be copied, modified, adapted, translated, reproduced, republished, uploaded, transmitted, posted, created as derivative work, sold, distributed or communicated in any form or by any means without prior written permission from TCS. Unauthorized use of the content/information appearing here may violate copyright, trademark and other applicable laws, and could result in criminal or civil penalties.

TCS attempts to be as accurate as possible in providing information and insights through this publication, however, TCS and its licensors do not warrant that the content/information of this publication, including any information that can be accessed via QR codes, links, references or otherwise is accurate, adequate, complete, reliable, current, or error-free and expressly disclaim any warranty, express or implied, including but not limited to implied warranties of merchantability or fitness for a particular purpose. In no event shall TCS and/or its licensors be liable for any direct, indirect, punitive, incidental, special, consequential damages or any damages whatsoever including, without limitation, damages for loss of use, data or profits, arising out of or in any way connected with the use of this publication or any information contained herein.

©2019 Tata Consultancy Services Limited. All Rights Reserved.

Tata Consultancy Services (name and logo), TCS (name and logo), and related trade dress used in this publication are the trademarks or registered trademarks of TCS and its affiliates in India and other countries and may not be used without express written consent of TCS. All other trademarks used in this publication are property of their respective owners and are not associated with any of TCS' products or services. Rather than put a trademark or registered trademark symbol after every occurrence of a trademarked name, names are used in an editorial fashion only, and to the benefit of the trademark owner, with no intention of infringement of the trademark.