

Persona-Based App Content Personalisation

Technical Field

[0001] This disclosure generally relates to methods and systems for personalizing content in an app. More particularly, the disclosure relates to clustering consumers into groups and developing personalized content for each cluster of consumers.

Background

[0002] This background description is provided for the purpose of generally presenting the context of the disclosure. Contents of this background section are neither expressly nor impliedly admitted as prior art against the present disclosure.

[0003] With the growth in super apps – apps that offer multiple services in a single interface – when new services are activated there is a need to alert consumers of those services and attract their attention. As the number of services offered by an app increases, smartphone screen real estate becomes difficult to manage – not all services can be concurrently displayed, and a user may overlook some information since it does not appear as relevant to them as other information.

[0004] Push notifications, CVPs (Customer Value Propositions), tiles, nudges, are often too generic and unable to encourage users to undertake specific actions – e.g. activating a new product, cross-selling a service, or changing to low cost-of-funds payment methods. Current notification tools are only able to provide solutions that fit consumers as a whole, or very large groups of consumers, and use rule-based intelligence. Such tools are traditionally safe, and generic, to cater for a wide audience due to the time-consuming nature of servicing subgroups of that audience.

[0005] Some recommender engines endeavor to recommend services. However, they often overlook business goals and face challenges in promoting new products/services due to the absence of historical data. This leads to recommendation bias against new offerings since the engine lacks relevant purchase or activation history to leverage for recommendations.

[0006] It is desired to address or ameliorate one or more disadvantages or limitations associated with the conventional systems and methods for personalizing content in an app, or to at least provide a useful alternative.

Summary

[0007] Disclosed is a system for personalising content in an app, the system comprising:

one or more processor (processor(s));

a memory comprising instructions that when executed by the processor(s) cause the processor(s) to:

retrieve historical consumer data for a plurality of consumers, the historical consumer data comprising, for each consumer, a behavioural attribute and a transactional attribute;

segment the consumers, using a segmentation engine, into a plurality of personas based on similarity between the respective consumers, each consumer associated with a said persona;

identify a business use case accessible via the app;

generate, for each persona and using a copywriter engine, a notification for the business use case; and

display to each consumer the notification corresponding to the persona associated with the respective consumer.

[0008] Also disclosed is a method for personalising content in an app, the method comprising:

retrieving historical consumer data for a plurality of consumers, the historical consumer data comprising, for each consumer, a behavioural attribute and a transactional attribute;

segmenting the consumers, using a segmentation engine, into a plurality of personas based on similarity between the respective consumers, each consumer associated with a said persona;

identifying a business use case accessible via the app;

generating, for each persona and using a copywriter engine, a notification for the business use case; and

displaying to each consumer the notification corresponding to the persona associated with the respective consumer.

Brief Description of the Drawings

[0009] Some embodiments of systems and methods for personalizing content in a super app, in accordance with present disclosure, will now be described, by way of non-limiting example only, with reference to the accompanying drawings in which:

[0010] Figure 1 illustrates a block diagram of a network including a system for personalizing content in an app and its associated components;

[0011] Figure 2 is a system for personalizing content in an app and its associated components;

[0012] Figure 3 illustrates a flowchart for a method for personalizing content in an app; and

[0013] Figure 4 shows various example screens of notifications, tiles and others, with content personalized based on personas identified by the method of Figure 2.

Detailed Description

[0014] Embodiments relate to systems and methods for personalizing content in an app such as push notifications, tiles and the like. Embodiments may include a machine learning model that generates content corresponding to a business case and personalized for a particular consumer, based on a consumer segment or cluster to which that consumer belongs. The machine learning model is trained using historical data relating to consumer behaviours such as rides, transactions, and cancelled rides, payment methods and others. The embodiments advantageously provide notifications that have a higher likelihood of garnering interaction from consumers, to activate and/or use more services through the app. This improves the experience and engagement of users with the app. Embodiments also provide a mechanism for prioritizing business use cases (i.e. products and services it is desirable for consumers to engage with) depending on the persona of each consumer.

[0015] Figure 1 illustrates a block diagram of a system for personalizing content in an app and its associated components. The system 100 comprises at least one processor 102, memory 104 accessible to the processor 102 and a network interface 108 to facilitate communication with a plurality of consumer's computing devices 160. Program code 106 provided in memory 104 comprises instructions executable by the processor 102 to perform at least a part of the method of the embodiments described herein. Notably, while individual computer systems are described in Figure 1, any such computer system may be distributed across multiple servers or multiple devices, or some functionality may be consolidated into a single server or device, without departing from the purposive intent of the present disclosure.

[0016] The consumer's computing device 160 comprises one or more processors 162, a memory 164, a GPS device 167 and a network interface 169. The device 160 hosts the app for which the content is to be personalized using the methods described herein, and may also host a digital wallet for transactions to be made on the app. The memory 164 comprises program code 166 comprising instructions executable by the processor 162 to facilitate interactions with the system 100. The consumer's computing device may include a personal or handheld computing device such as a smartphone or a tablet. Network 130 facilitates communication between the various devices and may include one or more communication networks including the internet, cell phone networks etc. The system 100 communicates with the consumer's computing device 160 over network 130, to deliver notifications to the consumer's computing device.

[0017] Each notification may comprise a push notification, a nudge (in-app nudges are subtle, non-intrusive directions or assistance for a user, relevant to the user's context at all times. These nudges highlight important app features, make them discoverable, and spur meaningful engagement on the app.), a tile or CVP.

[0018] One or more database 120 are also accessible to the system 100. The database 120 comprises historical consumer data comprising consumer attributes. A consumer attribute is a data point or data points describing the consumer, from which a determination can be made of a consumer's behavior, demographic and other characteristics. Consumer attributes include:

transactional attributes – transaction categories (food, transport, deliveries types, express delivery usage, premium service usage, online shopping, investment or

insurance product purchase), payment method or pay-later, sensitivity to promotions (i.e. whether purchases are frequently made during promotional periods when compared with full cost periods) and transaction RFM (Recency, Frequency, Monetary attributes) level;

behavioural attributes – ridesharing and ride taking behaviour, transactions, ride and service cancellation behavior – e.g. proportion of service cancellations compared with service uses, and reason for cancellation – delivery types and frequency, promotion interaction and time factors such as time taken to cancel a ride or purchase;

demographic attributes – age, gender, user acquisition channels, mobility, new or existing consumer, consumer system and device preferences, language; and

use cases – awareness (e.g. increasing publicity or availability to market for the product), activation (creating more opportunities for users to use the product), repayment (repaying any "Paylater" (i.e. deferred payment) bills) and retention (creating opportunities to return and use the product); and

others – location of consumer's computing device 160 as determined by GPS 167.

[0019] Historical consumer attribute records serve as a foundation for generating insights on consumer behavior patterns and groups. These records are used to train machine learning models, which in turn personalize app content to enhance user engagement and interaction with the app's services and products. The machine learning model not only considers consumer attributes but also incorporates business use cases, such as acquisition, retention, and cross-selling. During acquisition, consumers receive information about the app's offerings. For retention, app content is tailored to ensure consumer satisfaction with the products and services they engage with. Cross-selling leverages understanding of a consumer's preferences to offer related products and services

[0020] Figure 2 shows modules in system 200, corresponding to system 100 of Figure 1. The modules may have hardware components, implemented in software, separate modules or may be combined in any desired manner, without loss of generality and without departing from the functions described herein. The system retrieves, e.g. from database 202 historical consumer data. Database 202 may be referred to as an Online

Analytical Processing (OLAP) database. The OLAP consolidates historical data of consumers on the app. The historical data comprise the attributes mentioned above, being a non-exhaustive list of attributes.

[0021] In other embodiments, the system 200 comprises memory (e.g. memory 104 of system 100) that stores the historical consumer data, the historical consumer data is retrieved from individual consumer devices in real time, or the historical consumer data is retrieved from any other suitable source.

[0022] The system 200 includes segmentation engine 204. The segmentation engine 204 determines personas 208 for a group of consumers (also referred to as users) based on the historical consumer data. The segmentation engine 204 engine uses a machine learning model 210 to provide each consumer with a persona. Each persona forms or is associated with a user profile that informs the copywriter engine 206 on how to personalized notification (i.e. communications to each consumer). The segmentation engine 204 segments consumers into clusters. This can be achieved using any desired method, such as through unsupervised segmentation. For example, historical consumer data (transactional attributes 212, behavioural attributes 214 and demographic attributes 216) may form class labels or feature vectors. Segmentation can then be performed based on a distance between consumers in a feature space – e.g. Euclidean distance in a feature space determined from the class labels or feature vectors. Clustering may be used to cluster consumers into personas, where each cluster corresponds to a respective persona. Clustering may involve *k*-means clustering, hierarchical clustering, density-based spatial clustering or a combination thereof. As such, consumers with common attributes will be grouped closely together and potentially form clusters, whereas users with dissimilar attributes will be spread apart and potentially fall into different clusters. Segmentation may also take into account a specific number of personas, or the personas may be predetermined and each consumer is allocated to the persona to which their attributes most closely relate.

[0023] Segmentation may take into account a business use case. The business use case may result in weights being applied to particular attributes to emphasise or de-emphasise each attribute. Thus, segmentation may involve receiving a business use case and segmenting the consumers based on consumer attributes and the business use case. For example, specific attributes such as payment method and paylater (i.e. lay-away) may be more influential on a consumer's uptake of financial services. Therefore, if a business

use case relates to the promotion of financial services, a weight may be applied to attributes that align with the business use case such that users are clustered based on relevant, financial attributes rather than, for example, food preferences. The system 200 may use a machine learning model to analyse a business use case and determine which attribute or attributes are relevant to that business use case, and/or which attributes are less relevant, or irrelevant, to that business use case.

[0024] Examples of personas include Foodie (users who are highly interested in foods), Flash (users who are time-savers), Motorhead (users who are automotive enthusiasts), Travellers (inter and intra Southeast Asia) and many others. The personas may also include an indication of whether the services to which those user are attracted are offered in the app. Table I provides examples of personas.

Persona	Definition	Currently in app/Not in app
Foodie	A person who loves to explore new food	In-app food delivery services, non-app food & beverage
Flash	A person who often uses the app for ride hailing services and likes things to be fast	In-app transport services, non-app transport services
Grocer	A homemaker who cooks	In-app food delivery services, non-app food delivery services
Traveller	A person who is using the app not in his/her native country	In-app services
Shopper	A person who is into fashion and accessories	Non-app services
Gamer	A person who loves games, toys and movies	Non-app services
Sportie	A person who loves playing sports	Non-app services
Motorhead	A person who is into automobiles and cars	Non-app services
Others	All Others	

Table I – example personas and definitions

[0025] The segmentation engine 204 may be periodically updated to either reclassify existing consumers into personas based on recent behavior of each consumer – consumer behavior changes over time – and/or to identify new clusters (personas). Update may occur one a month or at any other frequency as desired. Notably, updating can include leveraging previous knowledge of clusters, or may involve wholesale replacement of existing clusters with newly identified clusters.

[0026] The personas are fed into the copywriter engine 206. Copywriter engine 206 provides a text output for all personas 208 and business use cases 218 – there may be only a single business use case, or a plurality of business use cases depending on the desired outcome. This will create the notification or content personalization needed for different groups of users and use cases. The copywriter engine 206 engineers prompts that are designed to attract and engage consumers of each respective persona. Prompts may be engineered using a machine learning language model (220) to analyse a business use case (e.g. a description of a service or product being offered in the app) and produce a generic prompt to engage or attract consumers to the business use case (i.e. to the product or service being offered under that business use case). The generic prompt may be generated by model 220 using a restricted vocabulary common to the app in question, or may be based on a language-specific, but otherwise unlimited, vocabulary. The copywriter engine 206 then uses a large-language model (222) to generate, for each persona, a notification for each business use case. Since the notification generated by model 222 is based on the persona of a consumer, the notification is personalized based on that persona and is therefore more relevant to the consumer. Notification generation may involve model 222 injecting descriptions of personas (e.g. attributes that define a cluster corresponding to a particular persona), business use cases, and types of content, into generic prompts. The generic prompts are thereby converted into personalized notifications, one personalized notification for each persona, for each business case. These personas are then displayed to consumers based on the persona of the consumer.

[0027] In some embodiments, machine learning model 222 may also receive an input temperature. An input temperature defines the degree with which the wording of the notification must align with the business use case and/or generic prompt. By having a temperature other than "1" (absolute or tightest alignment with the business use case or generic prompt), the temperature range being from 0 to 1, multiple notifications can be generated for each business use case, for each persona. For any particular consumer, the system 200 may then randomly select a notification, from a plurality of notifications generated for the persona of the use for a particular business use case, and display the selected notification to the consumer. On subsequent occasions when a notification for the same business use case is presented to the consumer, another random selection may be made.

[0028]Data may be collected on the frequency of interactions of particular personas with each notification and the wording of notifications, to improve the ability of machine learning model 222 to generate future notifications that will be interacted with, with higher frequency. To this end, the model 222 may learn, for each persona, wording and/or phraseology associated with the persona – i.e. that engages or attracts the attention of consumers of a particular persona. The model 222 may then generate personalized notification using that wording and/or phraseology or similar wording and/or phraseology.

[0029]All models may be implemented in any desired language. For example, models may be implemented using Python script and dictionaries.

[0030]An optional step is to validate notifications using validation module 224, prior to displaying the notifications to consumers. The validation module 224 may employ a machine learning model that accepts notifications and manual validation flags – e.g. "good" or "bad" – to validate notifications. In this context, 'validating' a notification refers to approving its wording or content (which may in some instances include an image or audio prompt) for display. Once the model in module 224 is trained, validation may occur automatically. The model in module 224 may be updated based on data mentioned above, relating to consumer interaction with notifications. In some embodiments, the model in module 224 receives no manual inputs and instead learns from consumer interactions with notifications, or employs a language model to compare each notification with the business use case to ensure a meaning of the notification aligns with the business use case.

[0031]The copywriter engine 206 then sends the notifications over network 226 to each consumer device 228 for display as a push notification 230, CVP 232, tile 234 or other form of notification.

[0032]System 200 also includes a ranking engine 236. The ranking engine provides an automated ranking of business use cases (cross-selling, up-selling, retention, acquisition, etc.) to be prioritized. This helps balance the competition between products and services in a super app, for fixed screen real estate. Tiles, nudges, tooltips are usually competing for the same real estate in the app. The ranking engine 236 highlights the most appropriate business use case to be shown in a given on-screen space. For a single business use case, no ranking needs to take place. However, when multiple business use cases are being promoted, they should be prioritized.

[0033] The ranking engine may identify a plurality of business use cases accessible via the app. This can be done by retrieving business use cases from database 202, or through any other method. The business use cases are then ranked according to at least one business metric, such as new products or services being marketed before existing products or services, projected profitability of a product or service, seasonal variation – e.g. in-app purchases may be more prevalent around the end of a calendar year, Valentine's Day – and others. Ranking may involve using retrieving the same attributes 212, 214, 216 used by the segmentation engine 204, and treating them as metrics 238 to a ranking (optimization) algorithm 240. Ranking algorithm 240 optimizes a business metric (objective) function that will solve a system of linear equations from a set of relevant metrics 238, to maximize the objective – e.g. profitability, or likelihood of consumer interaction with a business use case. For example, the objective function may maximize consumer interactions with a business use case like Grab Unlimited, and thus for Grab Unlimited business use case, transactional attributes will be weighted higher than other attributes when providing metrics as inputs to the ranking algorithm 240. To that end, a prioritizing engine may learn a relationship between business use cases and personas, and prioritise business use cases based on the relationship and the ranking. For example, flash personas may prefer express delivery products, one-touch online shopping and the like. Thus, business use cases for express delivery and online shopping will rank higher than other business use cases. After ranking, the business use cases may then be prioritized at 242 based on a business priority, such as maximizing frequency of consumer interaction or consumer time in app. Ranking and prioritizing may, in some instances, refer to the same function – namely, ranking the business use cases for each consumer based on the consumer's persona (i.e. consumer attributes of the persona that align with consumer attributes relevant to particular business use cases), to help prioritise which business use case should be maximised for a given consumer. In other instances, the ranking process and prioritizing process each apply separate weights such that business use case ranking influences, but does not determine, the order in which business use cases are displayed. For example, ranking business use cases may involve applying weights to business use cases such that they appear in the order in which they are desired to be promoted, so that particular business use cases are more heavily promoted than others – ranking therefore applies to all consumers equally. Prioritising may then involve applying consumer specific weights to business use cases depending on the degree with which consumer attributes relevant to those business use cases align with consumer attributes of the persona of the consumer. Thus, the ranking and

prioritizing steps apply separate weights to the business use cases, and the combined weights determine the order of appearance or promotion of those business use cases. As such, a balance is made between the order in which a business would like business use cases promoted, and the types of business use cases each individual consumer will find attractive.

[0034] Figure 3 illustrates a flowchart of a method 300 for personalizing content in an app, implemented by the system 100, 200. Particular embodiments may repeat one or more steps of the method of Figure 3, where appropriate. Although this disclosure describes and illustrates particular steps of the method of Figure 3 as occurring in a particular order, this disclosure contemplates any suitable steps of the method of Figure 3 occurring in any suitable order.

[0035] At step 302, the system 100 retrieves historical consumer data from the database 120. The data includes, for each consumer, at least a behavioural attribute and a transactional attribute.

[0036] At step 304, the consumers are segmented into a plurality of personas 306 based on similarity between the respective consumers. As a result, each consumer associated with a said persona. Typically, a consumer will be associated with only one persona, though in some instances multiple personas may be applied if the personas are generated with respect to different business use cases. Segmentation can involve clustering consumers based on a distance between consumers in feature space. The feature space is defined by the attributes for each consumer – e.g. the attributes may be used as a feature vector or a latent feature space may be learned based on the consumer attributes extracted from the historical consumer data. Clustering can employ any appropriate clustering algorithm.

[0037] Step 308 involves identifying a business use case accessible via the app. This business use case can be retrieved from a database or server, or can be generated and inputted manually. This step can involve identifying a plurality of business use cases 314, so that a notification can be generated for each persona, for each business use case. The business use cases may be ranked according to a business metric or metrics, and prioritized based on the ranking, to ensure the notifications are displayed in a way that best meets business goals.

[0038] At step 310, the copywriter engine 206 generates, for each persona, a notification for the business use case or for each business use case. The copywriter engine 206 first generates a generic notification for each business use case and then personalizes that notification based on the personas. In other embodiments, the copywriter engine 206 generates the personalized notifications without first generating a generic notification. Notification generation may involve learning wording and/or phraseology that appeals to a particular persona – appeal can be determined based on wording and/or phraseology that attracted greater consumer interaction and engagement (determined through time in app or other metrics) – and generating the notifications based on that wording and/or phraseology.

[0039] Notifications may be generated in the language of the consumer, regardless of their geographical location. Also, historical consumer data may be extracted and grouped, prior to segmentation, based on demographics. For example, users above a certain age threshold may be grouped and segmentation performed by group. In that way, products and services that are specific to a demographic will be marketed to personas within that demographic – e.g. products marketed to older flash personas will differ from products marketed to younger flash personas.

[0040] At step 312 each consumer is displayed the notification that has been personalized for their persona. Display may involve using the priorities, discussed below, to determine where to place notifications on the screen of the consumer's computer device 160. Higher priority business use cases will be displayed higher or otherwise more prominently than those of lower priority.

[0041] Step 320, is an optional validation step that involves comparing a meaning of the notification against the business use case to determine whether or not the notification actually markets or promotes the product or service as intended by the business use case.

[0042] Figure 4 depicts examples of consumer computing device (smartphone) displays with content personalized using the methods described herein. In particular, image (a) shows a promotion nudge for payment services for a Flash persona when ordering food, and image (b) shows a similar promotion nudge for payment services, but for a Foodie persona when ordering food. The phrasing differs and the service offering may vary slightly. Similarly, for image (c) a food order basket shows three alternatives for

notifications (only one would be displayed at any one time, in practice) for a generic notification 400, a notification for a Foodie persona 402 and a notification for a traveler 404.

[0043] The present methods automatically gather information about consumers from a database and prioritizes the business use cases accordingly. Combined with the use of machine learning, consumers are segmented into different personas. Finally, we use LLM to create personalized content based on the user persona and use case.

[0044] Proposed algorithms use state of art machine learning to segment consumers, linear optimization to prioritize business use cases, and Large Language Models (LLM) to personalize the content a consumer engages with on an app. By leveraging consumer data (transactional, experiential, etc) and behavioral attributes, consumers can be segmented into different groups based on their interests, and preferences. A LLM can then be used to generate personalized content and recommendations that are tailored to each consumer's specific needs or interests. This can be refined by location: for example, travelers can receive recommendations and in-app experiences that can be curated in real time.

[0045] For example, we can send targeted push notifications to guide users through the onboarding process based on their individual preferences and history, and suggest alternative payment methods through contextual nudges based on their previous behavior.

[0046] With personalized notification, push notifications can be customized and sent to a consumer that enjoys exploring new food options, to attract them to click through and activate a now product that may enhance their food purchasing experience – e.g. a product for buying now and paying later.

[0047] Similarly, when users with different personas, categorized by the segmentation engine, open the app, the experience from tile names, feed card, onboarding content, etc. is personalized and curated specifically for that user.

[0048] The personalized notifications are intended to lead to greater consumer engagement, loyalty, retention, interaction and satisfaction, leading to higher conversion for products and likelihood to transact. This can assist with increasing revenue through

the app as well as cost optimization – e.g. moving consumers to low cost payment channels.

[0049] In summary, the present methods help personalize all content on an app, focusing particularly in this instance on notifications. This solution will be a long-term strategic moat for Grab as it continuously provides a refreshing change in content and engagement while scaling to various old and new products, thereby providing a superior experience to all its users in a personalized and automated way.

[0050] High degree of personalization would leads users with better engagement leading to higher conversion rates leading to higher revenues and cost optimisation by moving to lower cost channels. This can also help automate lot of manual processes, leading to efficiency and high output.

[0051] Any content in the app, particularly content with which a user is to interact, can be personalized using the system 100, 200 and method 300. Notifications, such as push notifications, nudges, CVPs, tiles and others may be curated for each persona, to achieve greater consumer engagement.

[0052] The reference in this specification to any prior publication (or information derived from it), or to any matter which is known, is not, and should not be taken as an acknowledgment or admission or any form of suggestion that that prior publication (or information derived from it) or known matter forms part of the common general knowledge in the field of endeavor to which this specification relates.

[0053] Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

[0054] The scope of this disclosure encompasses all changes, substitutions, variations, alterations, and modifications to the example embodiments described or illustrated herein that a person having ordinary skill in the art would comprehend. The scope of this disclosure is not limited to the example embodiments described or illustrated herein. Moreover, although this disclosure describes and illustrates respective embodiments herein as including particular components, elements, feature, functions, operations, or

steps, any of these embodiments may include any combination or permutation of any of the components, elements, features, functions, operations, or steps described or illustrated anywhere herein that a person having ordinary skill in the art would comprehend. Although this disclosure describes or illustrates particular embodiments as providing particular advantages, particular embodiments may provide none, some, or all of these advantages.

Claims

1. A system for personalising content in an app, the system comprising:
 - one or more processor (processor(s));
 - a memory comprising instructions that when executed by the processor(s) cause the processor(s) to:
 - retrieve historical consumer data for a plurality of consumers, the historical consumer data comprising, for each consumer, a behavioural attribute and a transactional attribute;
 - segment the consumers, using a segmentation engine, into a plurality of personas based on similarity between the respective consumers, each consumer associated with a said persona;
 - identify a business use case accessible via the app;
 - generate, for each persona and using a copywriter engine, a notification for the business use case; and
 - display to each consumer the notification corresponding to the persona associated with the respective consumer.
2. The system of claim 1, wherein identification of the business use case comprises identifying a plurality of business use cases accessible via the app, and wherein generation of the notification comprises generating, for each persona and using the copywriter engine, a notification for each business use case.
3. The system of claim 1, wherein segmentation of the consumers comprises segmenting the consumers based on a distance between respective consumers in a feature space.
4. The system of claim 3, wherein segmentation comprises clustering the consumers into clusters, each cluster corresponding to a persona, using one or more of:
 - k*-means clustering;

hierarchical clustering; and

density-based spatial clustering.

5. The system of any one of claims 1 to 4, wherein generation of a notification comprises:
 - learning, for each respective persona, at least one of wording and phraseology associated with the persona; and
 - generating the notification using the wording and/or phraseology.
6. The system of any one of claims 1 to 5, wherein identifying a business use case comprises:
 - retrieving a plurality of business use cases;
 - ranking the business use cases according to at least one business metric; and
 - prioritising the business use cases based on the ranking.
7. The system of claim 6, wherein prioritising the business use cases comprises:
 - learning a relationship between the business use cases and each respective persona; and
 - prioritising the business use cases based on the relationship and the ranking.
8. The system of any one of claims 1 to 7, being configured to validate each notification, by comparing a meaning of the notification against the respective business use case.
9. A method for personalising content in an app, the method comprising:
 - retrieving historical consumer data for a plurality of consumers, the historical consumer data comprising, for each consumer, a behavioural attribute and a transactional attribute;
 - segmenting the consumers, using a segmentation engine, into a plurality of personas based on similarity between the respective consumers, each consumer associated with a said persona;
 - identifying a business use case accessible via the app;

generating, for each persona and using a copywriter engine, a notification for the business use case; and

displaying to each consumer the notification corresponding to the persona associated with the respective consumer.

10. The method of claim 9, wherein identifying the business use case comprises identifying a plurality of business use cases accessible via the app, and wherein generating the notification comprises generating, for each persona and using the copywriter engine, a notification for each business use case.

11. The method of claim 9, wherein segmenting the consumers comprises segmenting the consumers based on a distance between respective consumers in a feature space.

12. The method of claim 11, wherein segmenting the consumers comprises clustering the consumers into clusters, each cluster corresponding to a persona, using one or more of:

k-means clustering;

hierarchical clustering; and

density-based spatial clustering.

13. The method of any one of claims 9 to 12, wherein generating a notification comprises: learning, for each respective persona, at least one of wording and phraseology associated with the persona; and generating the notification using the wording and/or phraseology.

14. The method of any one of claims 9 to 13, wherein identifying a business use case comprises:

retrieving a plurality of business use cases;

ranking the business use cases according to at least one business metric; and

prioritising the business use cases based on the ranking.

15. The method of claim 14, wherein prioritising the business use cases comprises:

learning a relationship between the business use cases and each respective persona; and

prioritising the business use cases based on the relationship and the ranking.

16. The method of any one of claims 9 to 15, further comprising validating each notification, by comparing a meaning of the notification against the respective business use case.