SERVER AND METHOD FOR FACILITATING CONSTRUCTING BRAND PLATFORM

TECHNICAL FIELD

[0001] Various embodiments relate to a server and a method for facilitating constructing a brand platform.

BACKGROUND

[0002] Due to development of information technology, a consumer (who in some contexts herein may also be referred to as a "requester", a "customer", an "eater", a "passenger" or a "Pax") may request an on-demand service using a computing device associated with the consumer. The on-demand service may allow the consumer to fulfil the consumer's demand via an immediate access to goods and/or services. The consumer may request the on-demand service, for example, an item delivery service or a transport service (also referred to as an "e-hailing service" or a "car-hailing service"), using a user interface provided by an on-demand service platform and presented on the computing device associated with the consumer. To request the on-demand service, the consumer may make an order (also referred to as a "booking") for the on-demand service.

[0003] A brand platform may serve as one of the most crucial components in the on-demand service, for example, the item delivery service, by structuralising item information (also referred to as "product information") provided by a merchant (who in some contexts herein may also be referred to as a "retailer", a "restaurant" or a "Mex"). For example, the brand platform may contribute to improve an accuracy of search results, personalised

recommendations, and better advertisement targeting, which may lead to increased consumer satisfaction, higher engagement, and improved business performance. In addition, a B2B (Business-to-business) solution specifically for brands may be built.

[0004] However, conventional technologies which may use an NER (Named Entity Recognition) model, an LLM (Large Language Model), etc. may have limitations. For example, such conventional technologies may have a limited support for multilingual functions; support a limited number of brands; have a limitation in expanding to free-written inputs where brands may be misspelled, abbreviated, or even missing; only extract a text span from inputs and require a further normalisation for outputs; have a limitation in carrying out few/zero-shot prediction; have a limitation in continuous learning (for example, learning new brand names while memorising previous brand names); have a much higher latency and costs to scan through a bulk of items; and have a limited knowledge in local markets of some regions and thus lead to a poor performance in tagging brands.

[0005] Therefore, there is a need to provide a solution for facilitating constructing a brand platform more effectively and efficiently.

SUMMARY

[0006] According to various embodiments, there is a server for facilitating constructing a brand platform, the server comprising: a memory configured to store instructions; and a processor for executing the stored instructions and configured to: obtain labelled data and unlabelled data for each item, wherein the labelled data includes an input text for the item and a brand label corresponding to the input text; obtain information about a plurality of brands from a master brand table; for the labelled data, predict a relevant brand which is relevant to the input text from the plurality of brands, and predict a relationship between the brand label and the relevant

brand; for the unlabelled data, predict whether the unlabelled data includes new brand information about the relevant brand, based on the information about the plurality of brands; generate updated information about the relevant brand, based on the predicted relationship and the new brand information; and train the master brand table based on the updated information about the relevant brand.

[0007] In some embodiments, the processor is further configured to: calculate a first similarity score between the input text and each of the plurality of brands; and predict the relevant brand from the plurality of brands, based on the calculated first similarity score.

[0008] In some embodiments, the processor is further configured to: calculate a probability that the unlabelled data includes the new brand information about the relevant brand, based on the information about the plurality of brands; determine if the probability is equal to or greater than a predetermined probability; and if the probability is equal to or greater than the predetermined probability, provide the unlabelled data with the probability to a labeller to verify the updated information about the relevant brand.

[0009] In some embodiments, the processor is further configured to: provide the predicted relationship along with the unlabelled data with the probability to the labeller to verify the updated information about the relevant brand; and if the updated information about the relevant brand is verified by the labeller, train the master brand table based on the updated information about the relevant brand.

[0010] In some embodiments, the labeller includes at least one of a human labeller and a large language model (LLM).

[0011] In some embodiments, the processor is further configured to: receive a merchant input including a merchant input text for a merchant item from a merchant; calculate a second similarity score between the merchant input text and each of the plurality of brands stored in

the master brand table; and predict a brand name for the merchant item, based on the second similarity score.

[0012] In some embodiments, the merchant input further includes a merchant input image for the merchant item, and the processor is further configured to: detect a first area including a logo image from the merchant input image using a logo detection model; determine a brand logo corresponding to the logo image using the master brand table; and verify the brand name for the merchant item predicted based on the second similarity score, using the determined brand logo.

[0013] In some embodiments, the processor is further configured to: detect a second area including a text image from the merchant input image using a logo detection model; determine a text corresponding to the text image using an optical character recognition technology; and if it is determined that the brand name for the merchant item predicted based on the second similarity score is unclear, additionally use the determined text to predict the brand name for the merchant item.

[0014] In some embodiments, the processor is further configured to: determine if the second similarity score is less than a predetermined score; and if the second similarity score is less than the predetermined score, determine that the brand name for the merchant item predicted based on the second similarity score is unclear.

[0015] In some embodiments, the processor is further configured to: construct a knowledge graph using the trained master brand table.

[0016] According to various embodiments, there is a method for facilitating constructing a brand platform, the method comprising: obtaining labelled data and unlabelled data for each item, wherein the labelled data includes an input text for the item and a brand label corresponding to the input text; obtaining information about a plurality of brands from a master brand table; for the labelled data, predicting a relevant brand which is relevant to the input text

from the plurality of brands, and predicting a relationship between the brand label and the relevant brand; for the unlabelled data, predicting whether the unlabelled data includes new brand information about the relevant brand, based on the information about the plurality of brands; generating updated information about the relevant brand, based on the predicted relationship and the new brand information; and training the master brand table based on the updated information about the relevant brand.

[0017] In some embodiments, the method further comprises: calculating a first similarity score between the input text and each of the plurality of brands; and predicting the relevant brand from the plurality of brands, based on the calculated first similarity score.

[0018] In some embodiments, the method further comprises: calculating a probability that the unlabelled data includes the new brand information about the relevant brand, based on the information about the plurality of brands; determining if the probability is equal to or greater than a predetermined probability; and if the probability is equal to or greater than the predetermined probability, providing the unlabelled data with the probability to a labeller to verify the updated information about the relevant brand.

[0019] In some embodiments, the method further comprises: providing the predicted relationship along with the unlabelled data with the probability to the labeller to verify the updated information about the relevant brand; and if the updated information about the relevant brand is verified by the labeller, training the master brand table based on the updated information about the relevant brand.

[0020] In some embodiments, the labeller includes at least one of a human labeller and a large language model (LLM).

[0021] In some embodiments, the method further comprises: receiving a merchant input including a merchant input text for a merchant item from a merchant; calculating a second similarity score between the merchant input text and each of the plurality of brands stored in

the master brand table; and predicting a brand name for the merchant item, based on the second similarity score.

[0022] In some embodiments, the merchant input further includes a merchant input image for the merchant item, and the method further comprises: detecting a first area including a logo image from the merchant input image using a logo detection model; determining a brand logo corresponding to the logo image using the master brand table; and verifying the brand name for the merchant item predicted based on the second similarity score, using the determined brand logo.

[0023] In some embodiments, the method further comprises: detecting a second area including a text image from the merchant input image using a logo detection model; determining a text corresponding to the text image using an optical character recognition technology; and if it is determined that the brand name for the merchant item predicted based on the second similarity score is unclear, additionally using the determined text to predict the brand name for the merchant item.

[0024] In some embodiments, the method further comprises: determining if the second similarity score is less than a predetermined score; and if the second similarity score is less than the predetermined score, determining that the brand name for the merchant item predicted based on the second similarity score is unclear.

[0025] In some embodiments, the method further comprises: constructing a knowledge graph using the trained master brand table.

[0026] According to various embodiments, a data processing apparatus configured to perform the method of any one of the above embodiments is provided.

[0027] According to various embodiments, a computer program element comprising program instructions, which, when executed by one or more processors, cause the one or more processors to perform the method of any one of the above embodiments is provided.

[0028] According to various embodiments, a computer-readable medium comprising program instructions, which, when executed by one or more processors, cause the one or more processors to perform the method of any one of the above embodiments is provided. The computer-readable medium may include a non-transitory computer-readable medium.

BRIEF DESCRIPTION OF THE DRAWINGS

[0029] The invention will be better understood with reference to the detailed description when considered in conjunction with the non-limiting examples and the accompanying drawings, in which:

- FIGS. 1 and 2 illustrate infrastructures of a system including a server for facilitating constructing a brand platform according to various embodiments.
- FIG. 3 illustrates a block diagram of a server for facilitating constructing a brand platform according to various embodiments.
- FIG. 4 illustrates a flowchart for a method for facilitating constructing a brand platform according to various embodiments.
- FIG. 5 is a data flow diagram showing an overview of a brand platform according to various embodiments.
- FIGS. 6 and 7 are exemplary diagrams showing a brand normalisation task at a label centre according to various embodiments.
- FIG. 8 is a data flow diagram showing an overview of a training method of a Brand Anchor Networks (BAN) with a mono-modal framework.
- FIG. 9A is a data flow diagram showing an overview of a training method of a Brand Anchor Networks (BAN) with a multi-modal framework according to various embodiments.

- FIG. 9B is an enlarged data flow diagram showing an overview of a training method of a Brand Anchor Networks (BAN) with a multi-modal framework according to various embodiments.
- FIG. 10 is a data flow diagram showing an operation of a brand tagging module according to various embodiments.
- FIGS. 11A and 11B are data flow diagrams showing an operation of a brand mining module according to various embodiments.
- FIG. 12 is an exemplary diagram showing a visualisation of a trained feature space according to various embodiments.
- FIG. 13 is a data flow diagram showing an active learning framework with a human labeller and an LLM (Large Language Model) agent according to various embodiments.

DETAILED DESCRIPTION

[0030] The following detailed description refers to the accompanying drawings that show, by way of illustration, specific details and embodiments in which the disclosure may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the disclosure. Other embodiments may be utilised and structural, and logical changes may be made without departing from the scope of the disclosure. The various embodiments are not necessarily mutually exclusive, as some embodiments can be combined with one or more other embodiments to form new embodiments.

[0031] Embodiments described in the context of one of a server and a method are analogously valid for the other server and method. Similarly, embodiments described in the context of a server are analogously valid for a method, and vice-versa.

[0032] Features that are described in the context of an embodiment may correspondingly be applicable to the same or similar features in the other embodiments. Features that are described in the context of an embodiment may correspondingly be applicable to the other embodiments, even if not explicitly described in these other embodiments. Furthermore, additions and/or combinations and/or alternatives as described for a feature in the context of an embodiment may correspondingly be applicable to the same or similar feature in the other embodiments.

[0033] In the context of various embodiments, the articles "a", "an" and "the" as used with regard to a feature or element include a reference to one or more of the features or elements.

[0034] As used herein, the term "and/or" includes any and all combinations of one or more of the associated listed items.

[0035] Throughout the description, the term "module" may be understood as an application specific integrated circuit (ASIC), an electronic circuit, a combinational logic circuit, a field programmable gate array (FPGA), a processor which executes code, other suitable hardware components which provide the described functionality, or any combination thereof. The term of "module" may include a memory which stores code executed by the processor.

[0036] In the following, embodiments will be described in detail.

[0037] FIGS. 1 and 2 illustrate infrastructures of a system 200 including a server 100 for facilitating constructing a brand platform according to various embodiments.

[0038] As shown in FIG. 1, the system 200 may include, but is not limited to, the server 100, a database system 140, a network 150, a first computing device 160 associated with a consumer 161 (who in some contexts herein may also be referred to as a "requester", a "customer", an "eater", a "passenger" or a "Pax"), a plurality of second computing devices 170 (not shown) each associated with a plurality of merchants 171 (who in some contexts herein may also be referred to as a "retailer", a "restaurant" or a "Mex"), and a plurality of third computing devices 180 (not shown) each associated with a plurality of partners 181 (who in some contexts herein

may also be referred to as a “manufacturer”, a “company”, or a “brand owner”). In some embodiments, the plurality of merchants 171 may be entities who may provide a good or a service to the consumer 161 via an on-demand service platform. In some embodiments, the plurality of partners 181 may be entities who may own a brand, and produce a good or a service and provide the good or the service to the plurality of merchants 171.

[0039] In some embodiments, the on-demand service may be a service allowing the consumer 161 to fulfil the consumer’s 161 demand via an immediate access to items and/or services. The consumer 161 may request the on-demand service, such as a transport service (also referred to as an “e-hailing service” or a “car-hailing service”) or an item delivery service, using a user interface screen presented on a first computing device 160. The consumer 161 may make an order for the on-demand service using the first computing device 160.

[0040] In some embodiments, the consumer 161 may use an application, for example, a mobile application, provided by the server 100. For example, the server 100 may provide an on-demand service platform, and may be controlled and/or managed by an on-demand service platform provider. The application may be installed in the first computing device 160 associated with the consumer 161, to interact with the server 100 for the on-demand service. In some embodiments, the application may show an advertisement, to promote a good, a service, an event or a brand to the consumer 161.

[0041] In some embodiments, the network 150 may include, but is not limited to, a Local Area Network (LAN), a Wide Area Network (WAN), a Global Area Network (GAN), or any combination thereof. The network 150 may provide a wireline communication, a wireless communication, or a combination of the wireline and wireless communication between the server 100 and the first computing device 160, between the server 100 and the plurality of second computing devices 170, and between the server 100 and the plurality of third computing

devices 180. As shown in FIG. 2, the network 150 may provide the wireline communication, the wireless communication, or the combination of the wireline and wireless communication.

[0042] In some embodiments, the first computing device 160 may be connectable to the server 100 via the network 150. In some embodiments, the first computing device 160 may be arranged in data or signal communication with the server 100 via the network 150. In some embodiments, the first computing device 160 may include, but is not limited to, at least one of the following: a mobile phone, a tablet computer, a laptop computer, a desktop computer, a head-mounted display and a smart watch. In some embodiments, there are a plurality of first computing devices which may be associated with a plurality of consumers respectively. For example, the plurality of first computing devices may belong to the plurality of consumers respectively. For example, the first computing device 160 may belong to the consumer 161 who is the eater or the passenger. As another example, the first computing device 160 may belong to the consumer 161 requesting the on-demand service for the eater or the passenger who is a recipient of the on-demand service.

[0043] In some embodiments, the first computing device 160 may include a location sensor. In some embodiments, the location sensor may communicate with at least one of a global positioning satellite (GPS) server, a network server, and a Wi-Fi server, to detect a location of the first computing device 160. In some embodiments, the first computing device 160 may generate information about the location of the first computing device 160.

[0044] In some embodiments, the server 100, for example, implemented by a server computer, may include a communication interface 110, a processor 120, and a memory 130 (as will be described with reference to FIG. 3).

[0045] In some embodiments, the server 100 may communicate with the first computing device 160 via the network 150. In some embodiments, the consumer 161 may request the on-demand service, for example, the delivery service or the transport service, using the user interface

screen presented on the first computing device 160. In some embodiments, to request the on-demand service, the consumer 161 may make an order (also referred to as a "booking") for the on-demand service. In some embodiments, the first computing device 160 may receive the order from the consumer 161 for the on-demand service. The first computing device 160 may send the order to the server 100 via the network 150. In some embodiments, the order may relate to information about a pick-up location and a drop-off location for the on-demand service. In some embodiments, the order may include the information about the pick-up location and the drop-off location for the on-demand service. In some embodiments, the information may include the location of the first computing device 160. In some embodiments, the location of the first computing device 160 may be considered as a location of the consumer 161. In some embodiments, the location of the consumer 161 may be considered as a drop-off location (that in some contexts herein may also be referred to as a "destination" or a "delivery location") of the on-demand service (for example, the item delivery service). In some other embodiments, the location of the consumer 161 may be considered as a pick-up location (that in some contexts herein may also be referred to as a "starting point") of the on-demand service (for example, the transport service). In some other embodiments, the first computing device 160 may send information about an address of the consumer 161, and the address of the consumer 161 may be considered as the pick-up location or the drop-off location of the on-demand service. In some embodiments, the information may include a destination that the consumer 161 would like to go as the drop-off location. In some embodiments, the information may include information about a merchant 171a that the consumer 161 selected, and include a location of the merchant 171a as the pick-up location (as will be described below).

[0046] In some embodiments, the system 200 may further include a database 141. In some embodiments, the database 141 may be a part of the database system 140 which may be external to the server 100. The server 100 may communicate with the database 141. In some other

embodiments, although not shown, the database 141 may be implemented locally in the memory 130 of the server 100.

[0047] In some embodiments, the server 100 may communicate with the plurality of second computing devices 170 via the network 150. In some embodiments, the plurality of second computing devices 170 may be arranged in data or signal communication with the server 100 via the network 150. In some embodiments, the plurality of second computing devices 170 may include, but is not limited to, at least one of the following: a mobile phone, a tablet computer, a laptop computer, a desktop computer, a head-mounted display and a smart watch. In some embodiments, the plurality of second computing devices 170 may be associated with the plurality of merchants 171 respectively. For example, the plurality of second computing devices 170 may belong to the plurality of merchants 171 respectively.

[0048] In some embodiments, the server 100 may communicate with the plurality of third computing devices 180 via the network 150. In some embodiments, the plurality of third computing devices 180 may be arranged in data or signal communication with the server 100 via the network 150. In some embodiments, the plurality of third computing devices 180 may include, but is not limited to, at least one of the following: a mobile phone, a tablet computer, a laptop computer, a desktop computer, a head-mounted display and a smart watch. In some embodiments, the plurality of third computing devices 180 may be associated with the plurality of partners 181 respectively. For example, the plurality of third computing devices 180 may belong to the plurality of partners 181 respectively.

[0049] FIG. 3 illustrates a block diagram of a server 100 for facilitating constructing a brand platform according to various embodiments.

[0050] As shown in FIG. 3, the server 100, for example, implemented by a server computer, may include a communication interface 110, a processor 120, and a memory 130.

[0051] In some embodiments, the memory 130 (also referred to as a "database") may store input data and/or output data temporarily or permanently. In some embodiments, the memory 130 may be configured to store instructions. In some embodiments, the memory 130 may store program code which allows the server 100 to perform a method 300 (as will be described with reference to FIG. 4). In some embodiments, the program code may be embedded in a Software Development Kit (SDK). The memory 130 may include an internal memory of the server 100 and/or an external memory. The external memory may include, but is not limited to, an external storage medium, for example, a memory card, a flash drive, and a web storage.

[0052] In some embodiments, the communication interface 110 may allow one or more computing devices, for example, a first computing device 160, to communicate with the processor 120 of the server 100 via the network 150, as shown in FIGS. 1 and 2. In some embodiments, as shown in FIGS. 1 and 2, the first computing device 160 may belong to the consumer 161 who wants to make an order for the on-demand service. In some embodiments, the communication interface 110 may transmit signals to the first computing device 160, and/or receive signals from the first computing device 160, via the network 150.

[0053] In some embodiments, the communication interface 110 may allow a plurality of second computing devices 170 to communicate with the processor 120 of the server 100 via the network 150, as shown in FIGS. 1 and 2. In some embodiments, as shown in FIGS. 1 and 2, each of the plurality of second computing devices 170 may belong to each of a plurality of merchants 171 who may provide a good or a service to the consumer 161. In some embodiments, the communication interface 110 may transmit signals to the plurality of second computing devices 170, and/or receive signals from the plurality of second computing devices 170, via the network 150.

[0054] In some embodiments, the communication interface 110 may allow a plurality of third computing devices 180 to communicate with the processor 120 of the server 100 via the

network 150, as shown in FIGS. 1 and 2. In some embodiments, as shown in FIGS. 1 and 2, each of the plurality of third computing devices 180 may belong to each of a plurality of partners 181 who may be an owner of a brand, and produce a good or a service and provide the good or the service to the plurality of merchants 171. In some embodiments, the communication interface 110 may transmit signals to the plurality of third computing devices 180, and/or receive signals from the plurality of third computing devices 180, via the network 150.

[0055] The processor 120 may include, but is not limited to, a microprocessor, an analogue circuit, a digital circuit, a mixed-signal circuit, a logic circuit, an integrated circuit, a Central Processing Unit (CPU), a Graphics Processing Unit (GPU), a Digital Signal Processor (DSP), a Field Programmable Gate Array (FPGA), an Application Specific Integrated Circuit (ASIC), or any combination thereof. Any other kind of implementation of the respective functions, which will be described below in further detail, may also be understood as the processor 120.

[0056] In some embodiments, the processor 120 may be connectable to the communication interface 110. In some embodiments, the processor 120 may be arranged in data or signal communication with the communication interface 110 to transmit/receive the signals.

[0057] In some embodiments, the processor 120 may construct a brand platform using collected and processed data. In some embodiments, to facilitate constructing the brand platform, the processor 120 may perform brand mining, brand tagging, and brand structuralising.

[0058] In some embodiments, the processor 120 may perform the brand mining. In some embodiments, the processor 120 may obtain a plurality of data for each item (also referred to as a “product”). In some embodiments, the plurality of data may include labelled data and unlabelled data for each item. In some embodiments, the labelled data may include an input text (hereinafter, referred to as a “first input text”) for the item and a brand label corresponding

to the first input text. In some embodiments, the unlabelled data may include an input text (hereinafter, referred to as a "second input text") for the item without a brand label. In some embodiments, the processor 120 may receive the labelled data from a partner 181a who may be an owner of the brand, and produce the item and provide the item to a merchant 171a. In some embodiments, the processor 120 may receive the unlabelled data from an open-source database. For example, the processor 120 may scrape the unlabelled data from a web, an open-source dataset, etc.

[0059] In some embodiments, the processor 120 may obtain information about a plurality of brands from a master brand table. In some embodiments, the master brand table may be stored in the memory 130. In some embodiments, the master brand table may store the information about the plurality of brands, for example, a normalised brand list.

[0060] In some embodiments, the processor 120 may input the labelled data and the unlabelled data into Brand Anchor Networks (BAN). In some embodiments, the processor 120 may further input the information about the plurality of brands obtained from the master brand table into the BAN.

[0061] In some embodiments, for the labelled data, the processor 120 may predict a relevant brand (also referred to as a "matching brand") which is relevant to the first input text from the plurality of brands. In some embodiments, the processor 120 may match the first input text with the plurality of brands stored in the master brand table, and predict the relevant brand for the item. In some embodiments, the processor 120 may calculate a similarity score (hereinafter, referred to as a "first similarity score") between the first input text and each of the plurality of brands. In some embodiments, the processor 120 may predict the relevant brand for the item from the plurality of brands, based on the calculated first similarity score. For example, a brand which has the highest first similarity score may be predicted as the relevant brand for the item. In some embodiments, the processor 120 may then predict a relationship between the brand

label and the relevant brand predicted for the item. In some embodiments, the relationship may include, but is not limited to, a relation of a sub-brand and a parent brand, an unknown synonym of the relevant brand predicted, and/or a translation in another language. For example, for the item of "orange drink", if the brand label indicates "PulXX" and the relevant brand predicted for the item is "MiXX MaXX", the processor 120 may predict that "PulXX" is a sub-brand of "MiXX MaXX".

[0062] In some embodiments, for the unlabelled data, the processor 120 may predict whether the unlabelled data includes useful brand information, for example, new brand information, about the relevant brand, based on the information about the plurality of brands.

[0063] In some embodiments, the processor 120 may generate updated information about the relevant brand, based on the predicted relationship and the new brand information. In some embodiments, the updated information about the relevant brand may include, but is not limited to, parent/sub brand relations, new brands, new synonyms and new translations. In some embodiments, the processor 120 may train the master brand table based on the updated information about the relevant brand. In this manner, the updated information about the relevant brand which may be important components in the master brand table may be ingested to the master brand table, and consequently, the brand platform may be improved based on the more accurate master brand table.

[0064] In some embodiments, the processor 120 may calculate a probability that the unlabelled data includes the useful brand information, for example, the new brand information, about the relevant brand, based on the information about the plurality of brands stored in the master brand table. In some embodiments, the processor 120 may match the second input text with the plurality of brands stored in the master brand table, and calculate the probability that the unlabelled data includes the new brand information about the relevant brand, in the form of a similarity score (hereinafter, referred to as a "third similarity score"). In some embodiments,

the processor 120 may determine if the probability is equal to or greater than a predetermined probability. In some embodiments, if the probability is equal to or greater than the predetermined probability, the processor 120 may provide the unlabelled data, including the second input text, with the probability to a labeller to verify updated information about the relevant brand. In this manner, the processor 120 may determine if the unlabelled data is worth the labeller's attention, and the labeller may check only the selected unlabelled data, and thus efficiency of the labeller's work may be improved.

[0065] In some embodiments, the processor 120 may provide the predicted relationship along with the unlabelled data with the probability to the labeller to verify the updated information about the relevant brand. For example, if the probability is equal to or greater than the predetermined probability, the processor 120 may provide the predicted relationship along with the unlabelled data, including the second input text, with the probability to the labeller to verify the updated information about the relevant brand. As another example, if the probability is less than the predetermined probability, the processor 120 may provide the predicted relationship along with the unlabelled data, including the second input text, with the probability (for example, with an indication that the probability is low) to the labeller to verify the updated information about the relevant brand. In some embodiments, if the updated information about the relevant brand is verified by the labeller, the processor 120 may train the master brand table based on the updated information about the relevant brand which was verified by the labeller.

[0066] In some embodiments, the labeller includes at least one of a human labeller and a Large Language Model (LLM) which may be a machine learning model (in-house model). For example, the LLM may be a main LLM integrating other tools such as another LLM or a search engine. In this manner, the processor 120 may incorporate inputs from the LLM and the human labeller, to improve the accuracy of the master brand table.

[0067] In some embodiments, the processor 120 may perform the brand tagging. In some embodiments, the processor 120 may receive an input (hereinafter, referred to as a "merchant input") for an item (hereinafter, referred to as a "merchant item") from a merchant 171a. In some embodiments, the processor 120 may receive the merchant input from the merchant 171a who may provide the merchant item to the consumer 161. In some embodiments, the merchant input may include, but is not limited to, an input text (hereinafter, referred to as a "merchant input text") and an input image (hereinafter, referred to as a "merchant input image"). In some embodiments, the merchant input text may include, but is not limited to, a category name, an item name, and a description for the merchant item. For example, the merchant input text may be short and straightforward, but may include potential typographical errors and/or abbreviations in different forms. In some embodiments, the merchant input image may include, but is not limited to, a photo of the merchant item.

[0068] In some embodiments, the processor 120 may calculate a similarity score (hereinafter, referred to as a "second similarity score") between the merchant input text and each of the plurality of brands stored in the master brand table. In some embodiments, the processor 120 may predict a brand name for the merchant item from the plurality of brands, based on the calculated second similarity score. For example, a brand name which has the highest second similarity score may be predicted as the brand name for the merchant item.

[0069] In some embodiments, the processor 120 may detect an area (hereinafter, referred to as a "first area") including a logo image from the merchant input image. For example, the processor 120 may detect the first area including the logo image, using a logo detection model. In some embodiments, the processor 120 may determine a brand logo corresponding to the logo image using the master brand table. In some embodiments, the master brand table may include a logo bank storing brand logos of the plurality of brands. In some embodiments, the

processor 120 may verify the brand name for the merchant item predicted based on the second similarity score, using the determined brand logo.

[0070] In some embodiments, the processor 120 may detect an area (hereinafter, referred to as a "second area") including a text image from the merchant input image. For example, the processor 120 may detect the second area including the text image, using the logo detection model. In some embodiments, the processor 120 may determine a text corresponding to the text image using an Optical Character Recognition (OCR) technology. In some embodiments, if it is determined that the brand name for the merchant item predicted based on the second similarity score is unclear, the processor 120 may additionally use the determined text to predict the brand name for the merchant item. In some embodiments, the processor 120 may determine if the second similarity score is less than a predetermined score. In some embodiments, if the second similarity score is less than the predetermined score, the processor 120 may determine that the brand name for the merchant item predicted based on the second similarity score is unclear, and additionally use the determined text to predict the brand name for the merchant item. In this manner, while using the merchant input text, for example, including the item name and the description, as a main source of information to predict the brand name for the merchant item, the processor 120 may additionally use the OCR results extracted from the merchant input image as supplements. Advantageously, the accuracy of the prediction of the brand name for the merchant item may be enhanced.

[0071] In some embodiments, the processor 120 may perform the brand structuralising. In some embodiments, the processor 120 may construct a knowledge graph using the brand platform which may be based on the trained master brand table. The knowledge graph may be a graph view of the brand platform. The knowledge graph may be easy to expand, since in the training, it may be required to provide item-level tagging data only and let the model (i.e. the BAN) learn to adjust edge weights automatically.

[0072] As described above, according to various embodiments, the processor 120 may incorporate the human labeller and the machine learning model, to set up and maintain a comprehensive brand platform with rich metadata (for example, including logo images, sub-brands structure, company information, etc.) and focus on a region of interest (for example, Southeast Asia). In addition, the processor 120 may detect and extract a bulk of brands (for example, over 100,000 brands) from free-written input text from the plurality of merchants 171 with extremely imbalanced distribution for brands and requirements for few-shot and zero-shot prediction.

[0073] As described above, according to various embodiments, the processor 120 may facilitate constructing the brand platform using the trained master brand table. The brand platform according to the various embodiments may have the following advantages:

a) Advertisements targeting: The brand platform may help to target advertisements to specific brands, based on the consumer's 161 behaviour and preferences. This may help to improve the effectiveness of advertisement campaigns.

b) Improved search: By tagging brands for grocery, food items, and other consumer-generated content, the brand platform may help to improve the accuracy of item search results. This may make it easier for the consumer 161 to find items, as the consumer 161 may directly search brand names to find certain brands or use a brand filter to find items from a certain brand.

c) Brand monitoring: The brand platform may monitor sales and mentions of different brands across the on-demand service platform. This may help to identify trends, understand the consumer's 161 preferences, respond to the consumer's 161 feedback, and evaluate the effectiveness of promotion and advertising.

d) Personalization: By analysing the consumer's 161 data and tagging the brands that the consumer 161 prefers, the brand platform may help the on-demand service platform to

personalise their recommendations and marketing campaigns. This may lead to increased consumer's 161 engagement and higher conversion rates.

e) B2B Solution: The brand platform may build a B2B service with a comprehensive brand list focusing on the region of interest (for example, Southeast Asia) and a high accuracy multilingual model to tag brands. Supermarkets and grocery chain stores may be the potential clients.

f) Intellectual Property: The brand platform including a function of brand tagging may be used to monitor the use of the partner's 171a trademarks and brands. By automatically tagging items with their corresponding brands, the partner 171a may quickly identify potential trademark infringements and take actions to protect their intellectual property.

g) Marketing & Advertising: Marketers may use brand information provided by the brand platform to monitor the performance of their branding efforts. By tagging contents with the corresponding brand names, the marketers may track how well their brands are being received by consumers and adjust their branding strategies accordingly.

[0074] FIG. 4 illustrates a flowchart for a method 300 for facilitating constructing a brand platform according to various embodiments. According to various embodiments, the method 300 for facilitating constructing the brand platform may be provided.

[0075] In some embodiments, the method 300 may include a step 301 of obtaining labelled data and unlabelled data for each item, wherein the labelled data includes an input text for the item and a brand label corresponding to the input text.

[0076] In some embodiments, the method 300 may include a step 302 of obtaining information about a plurality of brands from a master brand table.

[0077] In some embodiments, the method 300 may include a step 303 of, for the labelled data, predicting a relevant brand which is relevant to the input text from the plurality of brands, and predicting a relationship between the brand label and the relevant brand.

[0078] In some embodiments, the method 300 may include a step 304 of, for the unlabelled data, predicting whether the unlabelled data includes new brand information about the relevant brand, based on the information about the plurality of brands.

[0079] In some embodiments, the method 300 may include a step 305 of generating updated information about the relevant brand, based on the predicted relationship and the new brand information.

[0080] In some embodiments, the method 300 may include a step 306 of training the master brand table based on the updated information about the relevant brand.

[0081] FIG. 5 is a data flow diagram showing an overview of a brand platform 400 according to various embodiments.

[0082] In some embodiments, a server 100 (as described with reference to FIG. 3) may construct and provide the brand platform 400. In some embodiments, as shown in FIG. 5, the brand platform 400 may include at least one of a master brand table 401, a brand mapping database 402, a label centre 403, a brand tagging module 404, a brand mining module 405, and a feedback channel 406. In some embodiments, the server 100 may perform i) setting up, ii) improving a tagging quality through active learning tasks, iii) normalising a brand list in the master brand table 401, and iv) enriching metadata for the normalised brands.

[0083] In some embodiments, the master brand table 401 may be an integral aspect of the brand platform 400. The master brand table 401 may serve as a singular and comprehensive repository housing all the mined metadata relating to a plurality of brands from a local market. The master brand table 401 may be designed to store, organise, and manage a plethora of brand-related information (also referred to as “brand information” or “information about the plurality of brands”), including, but not limited to, a normalised brand name, company information, brand logos, local synonyms, etc.

[0084] In some embodiments, one of primary purposes of the master brand table 401 may be to ensure data consistency, accuracy, and reliability. By normalising the brand names, the master brand table 401 may maintain uniformity and eliminate data discrepancies. This may simplify data operations and reduce errors when referencing brand information across internal systems.

[0085] In some embodiments, the master brand table 401 may work as one of the most important components of generating a business report to retrieve company information for each brand.

[0086] In some embodiments, the brand mapping database 402 may be a pivotal component for brand-related applications. The brand mapping database 402 may primarily serve as a structured and organised storage space, where all item-to-brand mapping data may be securely housed.

[0087] In some embodiments, the brand mapping database 402 may ensure consistency in item-brand associations across various platforms, departments, and processes within the on-demand service platform, and aid in mitigating data ambiguity.

[0088] In some embodiments, the mapping data may be stored in both an offline presto table and an online production table, which may be synchronised via scheduled jobs.

[0089] In some embodiments, the label centre 403 may be a proprietary tool to enhance and streamline the process of brand mining, brand tagging and brand normalisation. The label centre's 403 primary function may be to provide a generic and scalable framework for customised needs in a mart domain.

[0090] In some embodiments, one of the key functionalities of the label centre 403 may be its seamless integration with a Large Language Model (LLM) agent 409 and an in-house model. The label centre 403 may effectively automate the process of prefilling labels (also referred to

as “annotations”) for different types of tasks, thus reducing manual entries and errors that come with them.

[0091] In some embodiments, the label centre 403 may also have a capability to communicate with an internal human-in-the-loop system via API (Application Programming Interface). This may allow to build a pipeline that may automatically upload tasks to the label centre 403, download them once completed by a human labeller 408, and ingest the labels into the master brand table 401 and the brand mapping database 402.

[0092] In some embodiments, the feedback channel 406 may be provided. In some embodiments, downstream users 410 of the brand platform 400 and operations (OPS) 411 may be provided with bulk upload templates to provide feedback including, not limited to, updated metadata of the plurality of brands, onboard new brands, and updated item-level brand attributes. In some embodiments, the uploaded results may automatically be handled by pipelines and ingested into the production table, including at least one of the master brand table 401 and the brand mapping database 402.

[0093] In some embodiments, the brand platform 400 may integrate a Machine Learning (ML) model with the LLM agent 409, and combine them together with the human-in-the-loop system to continuously mine, tag, and structuralise the plurality of brands in a local market from multiple data sources, and thus may be able to continuously evolve and be updated in an open marketplace.

[0094] In some embodiments, the human labeller 408 and the LLM agent 409 may focus on completing the labelling tasks in the label centre 403. It may be done by humans (i.e. the human labellers 408). In some embodiments, in practice, there may not be dedicated human labellers 408 and the OPS 411 for doing the work. In some embodiments, the LLM agent 409 may also be adapted to finish the labelling tasks. In some embodiments, the downstream user 410 may refer to anyone who leverages the brand platform 400 in their use cases or applications. In some

embodiments, the OPS 411 may refer to people who maintains the brand platform 400 to make sure that it is up-to-date and to reduce data error rate by providing feedbacks using bulk upload templates. In some embodiments, in practice, the OPS 411 may also be the human labeller 408 to complete the labelling tasks.

[0095] In some embodiments, the brand mining module 405 may be provided. Instead of relying on high-quality annotated data, the brand platform 400 may iterate from cold start data with a modest annotation quality and large unlabelled corpus data. After training an initial version of the BAN 407, the plurality of brands from the dataset may be clustered by the BAN 407 and passed to the LLM agent 409 to discover wrong annotation, brand synonyms and/or sub-brand structure.

[0096] In some embodiments, the brand tagging module 404 may be provided. As the plurality of brands may be mined from a brand cluster, a cleaned list of the plurality of brands may be collected. The plurality of brands may then be tagged to the unlabelled corpus data from the multiple data sources (for example, Optical Character Recognition (OCR) results from an item package, merchant inputs, partner inputs, etc.). The tagged items may be used to expand the dataset for a next model iteration and provided as context information to collect metadata of the plurality of brands.

[0097] In some embodiments, with regards to the brand structuralising, given the brand names and tagged items, item names, descriptions and input images may be retrieved to help the LLM agent 409 or the human labeller 408 to understand these brands and then collect metadata from the external sources.

[0098] In some embodiments, the BAN 407 may be improved through feedback. In some embodiments, human inputs from downstream users 410 may be taken and handled by pipelines to automatically onboard new brands, modify metadata of existing brands, and fix wrong tagging.

[0099] FIGS. 6 and 7 are exemplary diagrams showing a brand normalisation task at a label centre 403 according to various embodiments.

[00100] In some embodiments, as shown in FIG. 6, in an original data source, a brand of "FebreXX" may be spelled in different ways, for example, "AmbiXXX", "FebreXX AmbiXXX", "FebreeXX", "AmbiXXX Luxe", etc., in input images of an item.

[00101] In some embodiments, firstly, Brand Anchor Networks (BAN) 407 (also referred to as a "centralised model" or a "BAN model") may run clustering to highlight potential duplications of this brand. For example, as shown in FIG. 6, the BAN 407 may check if each brand name is a duplication of this brand of "FebreXX", and inform the check results to at least one of a human labeller 408 and a Large Language Model (LLM) agent 409.

[00102] In some embodiments, as shown in FIG. 7, the BAN 407 may then retrieve the input images from items that may be tagged to this brand by the BAN 407. Lastly, the BAN 407 may send the retrieved input images to the label centre 403 for the labelling (also referred to as an "annotation"). In some embodiments, the at least one of the human labeller 408 and the LLM agent 409 may check the input images and the check results, and perform the labelling.

[00103] FIG. 8 is a data flow diagram showing an overview of a training method of a Brand Anchor Networks (BAN) 407 with a mono-modal framework according to various embodiments. FIG. 8 shows a model diagram of the BAN 407 with text-to-text matching (for example, matching item information to a brand name). In some embodiments, after training the BAN 407, the BAN 407 may be used in all three stages of brand mining, brand tagging, and brand structuralising.

[00104] In some embodiments, as a core model in the brand tagging module 404 and the brand mining module 405, a new architecture which may fulfil requirements of performing one-stage, open-world and low-latency brand tagging may be provided.

[00105] In some embodiments, an encoder, which may classify a text into different categories, using adapted loss function may be used to compute similarity scores between brand names and item names. For example, a Siamese text encoder using adapted InfoNCE (Noise-Contrastive Estimation) as the loss function may be used to compute the similarity scores between brand names and item names.

[00106] In some embodiments, an inference may be performed. Specifically, all of the plurality of brands in a brand platform may be encoded into sentence embeddings and stored. During the inference, the item names may be firstly encoded, the similarity scores between the item names and all of the plurality of brands may be computed, and a brand with a top score may be chosen as an output.

[00107] In some embodiments, with regards to the loss function, based on the original infoNCE, two improvements may be made to better address a use case in the plurality of brands, as follows:

- Mining hard negative brands by directly computing the similarity scores between the plurality of brands, and adding to a matrix for computing loss.
- Sampling by separate brands and retrieving items for each brand, instead of randomly sampling items to make sure long-tail brands are taken care of.

[00108] In some embodiments, the above-mentioned task may be one of four tasks in a multi-modal (also referred to as a "multi-task") training framework which will be described with reference to FIGS. 9A and 9B.

[00109] FIG. 9A is a data flow diagram showing an overview of a training method of a Brand Anchor Networks (BAN) 407 with a multi-modal (also referred to as a "multi-task") framework according to various embodiments. FIG. 9B is an enlarged data flow diagram showing the overview of the training method of the BAN 407 with the multi-modal framework according

to various embodiments. The data flow diagram shown in FIG. 8 may be the same as a top-left matrix of FIGS. 9A and 9B.

[00110] In some embodiments, the multi-modal framework may jointly train a text encoder and an image encoder. FIGS. 9A and 9B provide the overview of how the BAN 407 is trained with a multi-modal set-up, with the text encoder, the image encoder, and each matrix representing a contrastive loss (for example, infoNCE loss). A weighted sum of four sub-tasks may be used to update model weights.

[00111] In some embodiments, the multi-task framework may make the BAN 407 to learn a connection between detailed noisy information from real-world mart items and high-level abstract information from a master brand table 401.

[00112] In some embodiments, as shown in FIG. 9B, there may be two types of tasks, including inner-modal tasks (text-to-text matching and image-to-image matching, corresponding to tasks on the top-left matrix and a bottom-right matrix) and cross-modal tasks (text-to-image matching and image-to-text matching, corresponding to a bottom-left matrix and a top-right matrix).

[00113] In some embodiments, while the inner-modal tasks may generally be easier and may help the BAN 407 to converge fast, the cross-modal tasks may be valuable by adding details that may not be reflected in a mono-modal task, such as a product package, colour, etc. This may especially be helpful when classifying an item whose brand may only be found in an image and an item name may simply be listed as something like “drink”, and this may help to improve the coverage of brand tagging.

[00114] In some embodiments, eventually, two encoders may converge to such a feature space as follows:

- Both in text and images, the BAN 407 may be learned to focus on brand related information while ignoring the noise. For example, the text embedding of “ColXXX

Optic White Advanced Teeth Whitening Toothpaste, 2% Hydrogen Peroxide Toothpaste, Sparkling White, 3.2 Oz, 3 Pack" may be close to the embedding of "ColXXX".

- Brand embeddings may contain context information from items belonging to this brand. For example, the brand embedding of "ColXXX" and "SensoXXX" may be closer compared to other brands, because item images from these brands may share a pattern of a toothpaste package.

[00115] FIG. 10 is a data flow diagram showing an operation of a brand tagging module 404 according to various embodiments.

[00116] In some embodiments, the brand tagging module 404 may leverage a merchant input including a merchant input text (for example, category names, item names and descriptions) and a merchant input image for a merchant item. For example, the merchant input text may be short and straightforward, but may include potential typographical errors and abbreviations in different forms, and the thus the brand tagging module 404 may use the merchant input image as information supplements, when a brand name is unclear or unstated in the merchant input text.

[00117] In some embodiments, as shown in FIG. 10, the brand tagging module 404 may provide a text pathway. According to the text pathway, the brand tagging module 404 may rely on Brand Anchor Networks (BAN) 407 to compute a second similarity score between the merchant input text and all of a plurality of brands from a master brand table 401. For example, as shown in FIG. 10, mappings with higher second similarity scores are highlighted in thick solid lines, while mappings with lower second similarity scores are in dotted lines. Besides using the merchant input text (for example, the category names, the item names and the descriptions) as a main source of information, Optical Character Recognition (OCR) results extracted from the merchant input image may be used as supplements.

[00118] In some embodiments, as shown in FIG. 10, the brand tagging module 404 may provide an image pathway. According to the image pathway, a logo detection model may detect areas which may include brand names, text images, or logo images. Areas including the brand names may be passed to the text pathway, areas including the text images may be passed to the text pathway, and areas including the logo images may be mapped to brand logos from a logo bank of the master brand table 401 which may store the brand logos with variations. Such mapping may be performed by the BAN 407. For example, as shown in FIG. 10, an area (also referred to as a "second area") including a text image (for example, "…" and "Oolong") may be passed to the text pathway to obtain the OCR results, and an area (also referred to as a "first area") including a logo image (for example, a stylised logo of "TH") may be processed in the image pathway to map to the brand logos from the logo bank.

[00119] In some embodiments, the text pathway may be prioritised than the image pathway, since all items may be at least item name values and processing text may be faster than processing images.

[00120] FIGS. 11A and 11B are data flow diagrams showing an operation of a brand mining module 405 according to various embodiments.

[00121] In some embodiments, as shown in FIGS. 11A and 11B, the brand mining module 405 may provide two pathways, including a pathway from labelled data but with a low quality and a pathway from huge amounts of unlabelled corpus data.

[00122] In some embodiments, as shown in FIGS. 11A and 11B, for the labelled data, the brand mining module 405 may perform a normalisation step by matching the input text with the closest brand name from a master brand table 401 measured by a first similarity score. Predictions with high first similarity scores may indicate that original labels may have certain close relationships with the predicted brand names (also referred to as "relevant brand names"). In some embodiments, the relationship may include, but is not limited to, a relation of a sub-

brand and a parent brand, an unknown synonym of the relevant brand predicted, and/or a translation in another language. For example, as shown in FIG. 11B, for the item of "orange drink", if the brand label indicates "PulXX" and the predicted brand for the item is "MiXX MaXX", the brand mining module 405 may predict that "PulXX" is a sub-brand of "MiXX MaXX". In some embodiments, as shown in FIG. 11B, the predicted relationship may be further validated by at least one of a human labeller 408 and a Large Language Model (LLM) agent 409.

[00123] In some embodiments, as shown in FIGS. 11A and 11B, for the unlabelled data, a fast inference speed of Brand Anchor Networks (BAN) 407 may make scanning the huge amount of the unlabelled data become feasible by trying to match any existing brand in the master brand table 401. In some embodiments, as shown in FIG. 11B, unlike the process for the labelled data, the brand mining module 405 may use a third similarity score as a probability of containing useful brand information, for example, new brand information. In some embodiments, a certain threshold (also referred to as a "predetermined probability") may be set to determine if it is worth the labellers' 408, 409 attention.

[00124] In some embodiments, as shown in FIG. 11B, the brand mining module 405 may mine various kinds of information, including parent/sub brand relations, new brands, new synonyms and new translations. In some embodiments, as shown in FIGS. 11A and 11B, the master brand table 401 may be trained based on the mined information. The mined information may be important components in the master brand table 401.

[00125] FIG. 12 is an exemplary diagram showing a visualisation of a trained feature space according to various embodiments. FIG. 12 shows the visualisation of the trained feature space in a knowledge graph view after Brand Anchor Networks (BAN) 407 converges.

[00126] In some embodiments, a plurality of brands and a plurality of items may be managed in the feature space, and the feature space may be trained based on an updated master brand table 401.

[00127] In some embodiments, as shown in FIG. 12, each node represents brand embedding encoded by the BAN 407, each dotted line between two nodes represents similarity scores between the brand embeddings. Solid lines represent positive pairs in a training set, and dotted lines represent the BAN's 407 predictions. The similarity scores between any pair of nodes may be computed, and only edges above a certain threshold may be kept.

[00128] In some embodiments, in a training phase, company-brand, item-attribute (i.e. department in this diagram), item-brand mapping data may be provided.

[00129] In some embodiments, in an inference phase, the BAN 407 may learn to adjust a distance between the brand embeddings and aggregate features from items of each brands to auto-construct the knowledge graph with brand-brand edges and brand-department edges.

[00130] For example, as shown in FIG. 12, for the brand embedding of "LayXX", in the training set, some items from the brand "LayXX" may be tagged as "snacks" in department tagging and "LayXX" in brand tagging. While the BAN 407 may learn to bring close the embeddings of "LayXX's items", "snacks" and "LayXX", the BAN 407 may also indirectly bring closer the distance between the brand "LayXX" and "snacks". As such, the BAN 407 may be able to build the dotted line in predictions. In the feature space, a bulk of items may be provided, and thus the BAN 407 may learn these relations.

[00131] The various embodiments may enable to set up a graph view of the brand platform which may be easy to expand, since in the training, it is required to provide item-level tagging data only and let the BAN 407 learn to adjust the edge weights automatically. This may be how the BAN 407 identifies and mines homographs and parent-child brand relationships for operations (OPS) 411 to verify.

[00132] For example, as shown in FIG. 12, “DoXX shampoo” and “DoXX chocolates” may be a pair of homograph brands (i.e. different brands with the same spelling). In a brand mining module 404, the BAN 407 may be able to find that parts of items under “DoXX” are mapped to a department of “personal care & beauty”, and the other parts of the items under “DoXX” are mapped to a department of “chocolates”. As a result, the BAN 407 may determine that there are two clusters within one brand of “DoXX”, and users (for example, an on-demand service platform provider) may differentiate that they are actually homographs.

[00133] As another example, as shown in FIG. 12, for “WalXX” and “CornXXX”, the BAN 407 may find that they have a high similarity score and both belong to a brand of “MaXX Incorporated” and a department of “ice cream”, and the users (for example, the on-demand service platform provider) may notice that they are parent-child brands.

[00134] FIG. 13 is a data flow diagram showing an active learning framework with a human labeller 408 (also referred to as a human annotator”) and an LLM (Large Language Model) agent 409 according to various embodiments.

[00135] In some embodiments, although not shown in FIG. 13, before the brand platform is launched, an initial version of a master brand table 401 and Brand Anchor Networks (BAN) 407 may be set-up, via labelling a first batch of items (for example, items of an on-demand service platform) by the human labeller 408 and a rule-matching method. For example, this step may be a one-time effort. However, it may be appreciated that the initial version of the master brand table 401 and the BAN 407 may have a limited quality. Therefore, there may be a need to improve a tagging quality through an active learning.

[00136] In some embodiments, once the initial version of the master brand table 401 and the BAN 407 may be obtained, tagging all items may be started from an internal database. Predictions with high similarity scores (also referred to as “confidence scores” or “prediction scores”) may be uploaded directly, while items that the BAN 407 is unsure of (i.e. giving a low

prediction score") may be created as active learning tasks, which may be labelled by at least one of the human labeller 408 and the LLM agent 409.

[00137] In some embodiments, this may be considered as a virtuous circle, where these labelled data may be used to build a new dataset and re-train the BAN 407, and the improved BAN 407 may improve a data quality in return and so on.

[00138] In some embodiments, as shown in FIG. 13, the active learning framework with the human labeller 408 and the LLM agent 409 may be provided. In some embodiments, a target of the active learning framework may be to completely automate continuous learning (for example, including model improvement, data cleaning, etc.) process.

[00139] In some embodiments, the active learning framework may leverage an active learning sampling strategy to only sample items that the BAN 407 gives unconfident predictions, which may be labelled and added to the dataset to re-train the BAN 407.

[00140] In some embodiments, as a result, the active learning framework may gradually improve a coverage as well as a tagging quality on external data sources. Besides, the active learning framework may track artefacts and metrics generated in each iteration, so that users (for example, an on-demand service platform provider) may monitor an improvement and easily rollback to any historical version.

[00141] In some embodiments, normalising a brand list in a master brand table 401 may be performed. In some embodiments, the BAN 407 may rank all of a plurality of brands in the order of possibility of being the brand name of given items. For example, if two brands are predicted with high similarity scores, it is likely that they are duplications or parent-child brands. With this logic, the BAN 407 may create tasks that include potential synonyms, parent-child brands, etc. of certain brands and pass them to at least one of the human labeller 408 and the LLM agent 409 for labelling.

[00142] In some embodiments, for an input from a partner 181a, there may usually be brands labelled for each item, but the form of spelling may not always be standardised. The BAN 407 may be used to tag brands for the items, and compare them with original brands, to decide whether there is any synonym or a brand hierarchy relation between them.

[00143] In some embodiments, as a result, the master brand table 401 may be cleaned by merging the synonyms and the duplications as well as enriched by adding the brand hierarchy.

[00144] In some embodiments, enriching metadata for the normalised brands may be performed. In some embodiments, based on a list of the normalised brands, the next step may be to enrich attributes for the normalised brands to provide essential information for unlocking commercial implementations, including, but not limited to, companies, brand logos, and translations in local markets.

[00145] In some embodiments, although the human labeller 408 may be educated to scrape metadata from internet, it may be time consuming. The various embodiments may use the LLM agents 409 paired with a search engine API (Application Programming Interface) to spare the workload. By providing a prompt to guide the LLM agent 409 as well as a list of tools, the LLM agent 409 may work as an agent to auto-scraping metadata just like what the human labeller 408 does. An example of the prompt of collecting company information is as follows: template = """Given a free-written brand name and some product names from this brand, find the correct formatting and spelling of this brand and the company that owns this brand as best you can. You have access to the following tools:

{tools}

Use the following format:

Question: the input question you must answer

Thought: you should always think about what to do

Action: the action to take, should be one of [{tool_names}]

Action Input: the input to the action

Observation: the result of the action

... (this Thought/Action/Action Input/Observation can repeat N times)

Thought: I now know the final answer and my confidence level of the final answer to be correct, should be from 1 to 10

Final Answer: brand name | company name | confident level

Begin!

Question: {input}

{agent_scratchpad}"""

[00146] As described above, as a human-in-the-loop system, the various embodiments may have the following advantages:

- Building an innovative platform and system for mining, structuralising and curating the master brand table 401 from multiple sources with the LLM agent 409 to continuously keep it updated with an e-commerce world.
- Proposing and integrating a novel brand tagging method from free-style user inputs to the evolving master brand table 401 with the BAN 407 under an open world setting.
- Establishing a new human-in-the-loop (HitL) process with a label centre 403 to improve the system through active manual feedback.
- Proposing a new architecture of brand tagging pipeline with text and image pathways.

- Developing a brand mining module 405 with the LLM agent 409 to mine new brands, brand and sub-brand relations, brand ownership graph, brand logos, translations in local markets, etc.

[00147] As the BAN 407, the various embodiments may have the following advantages:

- Under the context of FMCG (Fast-moving consumer goods) brands, proposing to use an efficient representation learning method to simultaneously address three tasks: brand tagging, brand mining (relation extraction), and brand structuralising (knowledge graph construction).
  - Most existing methods may solve these three problems with separate models and the knowledge between them may not be shared. In the various embodiments, these tasks may be unified into one task of finding a feature space where all items and brands are situated in an optimal place to reflect their relations.
  - For brand tagging, each item may be tagged to the brand whose embedding is the closest, measured by similarity metrics such as dot products.
  - For relation extraction (or brand mining), brand pairs may be retrieved with high similarity scores and their relation may be inferred by comparing their attributes (for example, department, company, etc.)
  - This feature space from the BAN 407 may be translated into a knowledge graph with mart items, brands, companies, and attributes being nodes, and similarity between them being weights of edges.
  - Besides, the BAN 407 may only be trained on brand tagging datasets and it may generalise to the other two. For example, this may be performed in the feature space (by the knowledge graph construction). In other words, datasets may only tell the BAN 407 to pull close positive pairs of an item and a brand, and push far negative pairs, yet it may learn to adjust distance between brands in the training as well. For

example, in the dataset, there may be three positive pairs such as “FanXX orange soft drink, FanXX”, “MiXX MaXX orange juice, MinXX MaXX” and “CoXX soft drink, CoXX”. The BAN 407 may be able to learn the brand of “FanXX” is similar to both “CoXX” and “MinXX MaXX”, since they are similar in certain ways, while “MinXX MaXX” and “CoXX” are dissimilar.

- To solve this representation learning task and multi-lingual multi-modal scenario, proposing a training method with inner modal and cross modal learning.

[00148] While the disclosure has been particularly shown and described with reference to specific embodiments, it should be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention as defined by the appended claims. The scope of the invention is thus indicated by the appended claims and all changes which come within the meaning and range of equivalency of the claims are therefore intended to be embraced.

CLAIMS

1. A server for facilitating constructing a brand platform, the server comprising:

a memory configured to store instructions; and

a processor for executing the stored instructions and configured to:

obtain labelled data and unlabelled data for each item, wherein the labelled data includes an input text for the item and a brand label corresponding to the input text;

obtain information about a plurality of brands from a master brand table;

for the labelled data, predict a relevant brand which is relevant to the input text from the plurality of brands, and predict a relationship between the brand label and the relevant brand;

for the unlabelled data, predict whether the unlabelled data includes new brand information about the relevant brand, based on the information about the plurality of brands;

generate updated information about the relevant brand, based on the predicted relationship and the new brand information; and

train the master brand table based on the updated information about the relevant brand.

2. The server according to claim 1, wherein the processor is further configured to:

calculate a first similarity score between the input text and each of the plurality of brands; and

predict the relevant brand from the plurality of brands, based on the calculated first similarity score.

3. The server according to claim 1 or claim 2, wherein the processor is further configured to:

calculate a probability that the unlabelled data includes the new brand information about the relevant brand, based on the information about the plurality of brands;

determine if the probability is equal to or greater than a predetermined probability; and

if the probability is equal to or greater than the predetermined probability, provide the unlabelled data with the probability to a labeller to verify the updated information about the relevant brand.

4. The server according to claim 3, wherein the processor is further configured to:

provide the predicted relationship along with the unlabelled data with the probability to the labeller to verify the updated information about the relevant brand; and

if the updated information about the relevant brand is verified by the labeller, train the master brand table based on the updated information about the relevant brand.

5. The server according to claim 3 or claim 4, wherein the labeller includes at least one of a human labeller and a large language model (LLM).

6. The server according to any one of claims 1 to 5, wherein the processor is further configured to:

receive a merchant input including a merchant input text for a merchant item from a merchant;

calculate a second similarity score between the merchant input text and each of the plurality of brands stored in the master brand table; and

predict a brand name for the merchant item, based on the second similarity score.

7. The server according to claim 6, wherein the merchant input further includes a merchant input image for the merchant item, and

the processor is further configured to:

detect a first area including a logo image from the merchant input image using a logo detection model;

determine a brand logo corresponding to the logo image using the master brand table; and

verify the brand name for the merchant item predicted based on the second similarity score, using the determined brand logo.

8. The server according to claim 6, wherein the processor is further configured to:

detect a second area including a text image from the merchant input image using a logo detection model;

determine a text corresponding to the text image using an optical character recognition technology; and

if it is determined that the brand name for the merchant item predicted based on the second similarity score is unclear, additionally use the determined text to predict the brand name for the merchant item.

9. The server according to claim 8, wherein the processor is further configured to:

determine if the second similarity score is less than a predetermined score; and

if the second similarity score is less than the predetermined score, determine that the brand name for the merchant item predicted based on the second similarity score is unclear.

10. The server according to any one of claims 1 to 9, wherein the processor is further configured to: construct a knowledge graph using the trained master brand table.

11. A method for facilitating constructing a brand platform, the method comprising:

obtaining labelled data and unlabelled data for each item, wherein the labelled data includes an input text for the item and a brand label corresponding to the input text;

obtaining information about a plurality of brands from a master brand table;

for the labelled data, predicting a relevant brand which is relevant to the input text from the plurality of brands, and predicting a relationship between the brand label and the relevant brand;

for the unlabelled data, predicting whether the unlabelled data includes new brand information about the relevant brand, based on the information about the plurality of brands;

generating updated information about the relevant brand, based on the predicted relationship and the new brand information; and

training the master brand table based on the updated information about the relevant brand.

12. The method according to claim 11, further comprising:

calculating a first similarity score between the input text and each of the plurality of brands; and

predicting the relevant brand from the plurality of brands, based on the calculated first similarity score.

13. The method according to claim 11 or claim 12, further comprising:

calculating a probability that the unlabelled data includes the new brand information about the relevant brand, based on the information about the plurality of brands;

determining if the probability is equal to or greater than a predetermined probability; and

if the probability is equal to or greater than the predetermined probability, providing the unlabelled data with the probability to a labeller to verify the updated information about the relevant brand.

14. The method according to claim 13, further comprising:

providing the predicted relationship along with the unlabelled data with the probability to the labeller to verify the updated information about the relevant brand; and

if the updated information about the relevant brand is verified by the labeller, training the master brand table based on the updated information about the relevant brand.

15. The method according to claim 13 or claim 14, wherein the labeller includes at least one of a human labeller and a large language model (LLM).

16. The method according to any one of claims 11 to 15, further comprising:

receiving a merchant input including a merchant input text for a merchant item from a merchant;

calculating a second similarity score between the merchant input text and each of the plurality of brands stored in the master brand table; and

predicting a brand name for the merchant item, based on the second similarity score.

17. The method according to claim 16, wherein the merchant input further includes a merchant input image for the merchant item, and

the method further comprises:

detecting a first area including a logo image from the merchant input image using a logo detection model;

determining a brand logo corresponding to the logo image using the master brand table; and

verifying the brand name for the merchant item predicted based on the second similarity score, using the determined brand logo.

18. The method according to claim 16, further comprising:

detecting a second area including a text image from the merchant input image using a logo detection model;

determining a text corresponding to the text image using an optical character recognition technology; and

if it is determined that the brand name for the merchant item predicted based on the second similarity score is unclear, additionally using the determined text to predict the brand name for the merchant item.

19. The method according to claim 18, further comprising:

determining if the second similarity score is less than a predetermined score; and

if the second similarity score is less than the predetermined score, determining that the brand name for the merchant item predicted based on the second similarity score is unclear.

20. The method according to any one of claims 11 to 19, further comprising: constructing a knowledge graph using the trained master brand table.