Description

Title of The Invention: Strategy Determination Method, Apparatus, Device, and Storage Medium

Cross-Reference to Related Applications

This application claims priority to Chinese Patent Application No. 2025102622587, filed with the China National Intellectual Property Administration on March 6, 2025, and entitled "Strategy Determination Method, Apparatus, Device, and Storage Medium", the entire contents of which are incorporated herein by reference.

Technical Field

The present disclosure relates to the field of computer technology, and in particular, to a strategy determination method, apparatus, device, and storage medium.

Background Art

In the related art, to enable a large model to output effective strategies for input data, the large model needs to be trained. Specifically, a large amount of training data is first constructed, and then the model is trained using Supervised Fine-Tuning (SFT) technology or Reinforcement Learning from Human Feedback (RLHF) technology.

However, constructing a large model relies heavily on a vast amount of pre-constructed training data, such that the model construction process is extremely time-consuming. Additionally, model training demands substantial hardware resources. When applying large models to strategy determination scenarios, a significant amount of manually curated data is required for each individual strategy to enable the model to learn effectively. Moreover, under identical input conditions, a large model must consistently produce the same output strategy. Nevertheless, if the training data contains instances wherein identical input conditions correspond to different execution strategies, the trained model may exhibit processing inconsistencies, thereby affecting the strategy output effect of the large model.

Summary

In view of this, to at least address the technical issues in the related art wherein large models rely on massive amounts of training data to correctly output strategies, leading to lengthy training times and high hardware resource consumption, the present disclosure aims to provide a strategy determination method, apparatus, device, and storage medium.

To achieve the above objectives, the technical solutions adopted in the embodiments of the present disclosure are as follows:

In a first aspect, an embodiment of the present disclosure provides a strategy determination method, comprising:

extracting, by a large model, target condition information from a currently input user question, wherein the target condition information comprises a target condition slot name and a target condition slot value;

determining a target decision path matching the target condition information, from a decision path pool pre-constructed, wherein each decision path in the decision path pool comprises condition information and an execution strategy; and

outputting, by the large model, strategy information for replying to the user question, based on the user question and the execution strategy in the target decision path.

In an optional implementation, the step of determining a target decision path matching the target condition information from a decision path pool pre-constructed comprises:

obtaining a candidate decision path with a condition slot name identical to the target condition slot name from the decision path pool; and

matching the condition slot value corresponding to the target condition slot name in the candidate decision path with the target condition slot value to obtain the target decision path.

In an optional implementation, the step of matching the condition slot value corresponding to the target condition slot name in the candidate decision path with the target condition slot value to obtain the target decision path comprises:

converting the condition slot value corresponding to the target condition slot name in the candidate decision path into a first slot value vector, and converting the target condition slot value into a second slot value vector; and

determining the target decision path based on the candidate decision path where the first slot value vector whose similarity to the second slot value vector is greater than a first predefined similarity threshold is located.

In an optional implementation, a process of constructing the decision path pool comprises:

obtaining, by the large model, a current decision path to be tested corresponding to a set strategy objective and a set user question, based on the set strategy objective input, the set user question, and a decision path example;

generating, by the large model, a new decision path to be tested based on the set strategy objective, the set user question, and the condition information in the current decision path to be tested, updating the current decision path to be tested to the new decision path to be tested, and repeating this step until similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches a second predefined similarity threshold; and

when the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches the second predefined similarity threshold, scoring, by the large model, all decision paths to be tested, and storing the decision paths to be tested with scores higher than a predefined score threshold to obtain the decision path pool.

In an optional implementation, the method further comprises:

obtaining, when receiving a strategy update request, by the large model, a seed decision path corresponding to a new strategy objective and a reference question carried in the strategy update request based on the new strategy objective and the reference question;

generating, by the large model, a new decision path based on the new strategy objective, the reference question, and the condition information in the seed decision path, updating the seed decision path to the new decision path, and repeating this step until similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches a second predefined similarity threshold; and

scoring, when the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches the second predefined similarity threshold, by the large model, all seed decision paths, and storing the seed decision paths with scores higher than the predefined score threshold in the decision path pool.

In an optional implementation, the step of outputting, by the large model, the strategy information for replying to the user question, based on the user question and the execution strategy in the target decision path comprises:

Generating, according to the execution strategy in the target decision path, a corresponding guidance text; and

outputting, by the large model, the strategy information for replying to the user question, based on the user question and the guidance text.

In an optional implementation, the method further comprises:

training the large model with a pre-constructed training sample set until the large model converges, wherein the training sample set comprises multiple groups of training samples, each group of training samples comprises user question samples and corresponding labels, and the labels are used to characterize real condition information corresponding to the user question samples.

In a second aspect, an embodiment of the present disclosure provides a strategy determination apparatus, comprising:

an extraction module, configured to: extract, by a large model, target condition information from a currently input user question through, wherein the target condition information comprises a target condition slot name and a target condition slot value;

a screening module, configured to: determine a target decision path matching the target condition information, from a decision path pool pre-constructed, wherein each decision path in the decision path pool comprises condition information and an execution strategy; and

an output module configured to: output, by the large model, strategy information for replying to the user question, based on the user question and the execution strategy in the target decision path.

In a third aspect, an embodiment of the present disclosure provides an electronic device, comprising a processor and a memory, wherein the memory stores machine-executable instructions executable by the processor, and the processor is configured to execute the machine-executable instructions to implement the method provided in any one of the above first aspect.

In a fourth aspect, an embodiment of the present disclosure provides a computer-readable storage medium, storing a computer program thereon, wherein the computer program, when executed by a processor, implements the method provided in any one of the above first aspect.

The strategy determination method, apparatus, device, and storage medium provided in the embodiments of the present disclosure, by extracting target condition information from user questions through a large model, and then determining a target decision path matching the target condition information from a pre-constructed decision path pool, achieve the intervention in the reasoning process of a large model using a target decision path, such that the large model can output more accurate results for user questions. Herein, the large model can utilize its inherent capabilities to extract target condition information from user questions and formulate corresponding replies. Therefore, in the embodiments of the present disclosure, it is possible to improve the accuracy of the large model's output by combining decision paths without the need to train the large model, resulting in lower model application costs and higher model application efficiency. Additionally, if there is a need to update the strategy for responding to user questions, the large model can be enabled to output targeted strategy information for new strategies by updating the decision paths in the decision path pool, which better meets the requirement of frequent strategy changes in certain application scenarios, such as marketing strategies in marketing scenarios.

To make the objectives, features, and advantages of the present disclosure more clearly understood, preferred embodiments are hereinafter cited and described in detail in conjunction with the accompanying drawings.

Brief Description of Drawings

To more clearly illustrate the technical solutions of the embodiments of the present disclosure, the drawings required for the embodiments are briefly introduced below. It should be understood that the following drawings only show some embodiments of the present disclosure and therefore should not be considered as limiting the scope. For those of ordinary skill in the art, other related drawings can be obtained based on these without creative efforts.

FIG. 1 shows a structural block diagram of an electronic device provided in an embodiment of the present disclosure;

FIG. 2 shows a flowchart of a strategy determination method provided in an embodiment of the present disclosure; and

FIG. 3 shows a functional block diagram of a strategy determination apparatus provided in an embodiment of the present disclosure.

Detailed Description of Embodiments

The technical solutions in the embodiments of the present disclosure will be clearly and completely described below in conjunction with the drawings in the embodiments of the present disclosure. Obviously, the described embodiments are only some of the embodiments of the present disclosure, rather than all of them. Typically, the elements of the embodiments of the present disclosure described and shown in the accompanying drawings herein can be arranged and designed in various different configurations.

Therefore, the detailed descriptions of the embodiments of the present disclosure provided in the accompanying drawings are not intended to limit the scope of the claimed present disclosure, but merely represent selected embodiments of the present disclosure. For those skilled in the art, all other embodiments obtained without creative efforts based on the embodiments of the present disclosure shall fall within the protection scope of the present disclosure.

It should be noted that relative terms such as "first" and "second" are only used to distinguish one entity or operation from another entity or operation, and do not necessarily require or imply any actual relationship or sequence between these entities or operations. Additionally, terms such as "comprise", include" or any other variations thereof are intended to cover non-exclusive inclusion, such that a process, method, article, or device comprising a series of elements not only comprises those elements but also may comprise other elements not explicitly listed, or elements inherent to such process, method, article, or device. Without further limitations, a process, method, article, or apparatus “comprising an element” does not exclude the presence of additional identical elements in the process, method, article, or apparatus.

To at least address the technical issues in the related art wherein large models rely on massive training data to correctly output strategies, leading to lengthy training times and high hardware resource consumption, the present disclosure provides a strategy determination method, which, by extracting target condition information from user questions through a large model, and then determining a target decision path matching the target condition information from a pre-constructed decision path pool, achieves the intervention in the reasoning process of a large model using a target decision path, such that the large model can output more accurate results for user questions. Notably, the large model can utilize its inherent capabilities to extract target condition information from user questions and formulate corresponding replies. Therefore, in the embodiments of the present disclosure, it is possible to improve the accuracy of the large model's output by combining decision paths without the need to train the large model, resulting in lower model application costs and higher model application efficiency. Additionally, if there is a need to update the strategy for responding to user questions, the large model can be enabled to output targeted strategy information for new strategies by updating the decision paths in the decision path pool, which better meets the requirement of frequent strategy changes in certain application scenarios, such as marketing strategies in marketing scenarios.

The strategy determination method provided by the present disclosure can be applied to electronic devices. Referring to FIG. 1, which is a structural block diagram of an electronic device. The electronic device 100 comprises a memory 110, a processor 120, and a communication module 130. The memory 110, processor 120, and communication module 130 are directly or indirectly electrically connected to each other to enable data transmission or interaction. For example, these components may be electrically connected via one or more communication buses or signal lines.

Herein, the memory is configured to store programs or data. The memory can be, but is not limited to, Random Access Memory (RAM), Read Only Memory (ROM), Programmable Read-Only Memory (PROM), Erasable Programmable Read-Only Memory (EPROM), Electrically Erasable Programmable Read-Only Memory (EEPROM), etc.

The processor is configured to read/write data or programs stored in the memory and execute corresponding functions.

The communication module is configured to establish a communication connection between the electronic device and other communication terminals via a network, and to send and receive data over the network.

It should be understood that the structure shown in FIG. 1 is merely a schematic diagram of the electronic device. The electronic device may comprise more or fewer components than those shown in FIG. 1, or have a different configuration from that shown in FIG. 1. The components shown in FIG. 1 can be implemented in hardware, software, or a combination thereof.

The strategy determination method provided in an embodiment of the present disclosure is described below in conjunction with FIG. 2. FIG. 2 is a flowchart of the strategy determination method, which comprises:

in step S100, extracting target condition information from a currently input user question through a large model, wherein the target condition information comprises a target condition slot name and a target condition slot value;

in step S200, determining a target decision path matching the target condition information from a pre-constructed decision path pool, wherein each decision path in the decision path pool comprises condition information and an execution strategy; and

in step S300, outputting the strategy information for replying to the user question through the large model based on the user question and the execution strategy in the target decision path.

In some application scenarios, after a user asks a question, it is necessary to trigger strategic guidance output to lead the user to perform relevant operations and achieve the established strategy. Taking a marketing scenario as an example, suppose the marketing strategy (which can also be called a marketing objective) is: "prompt the user to fill in a phone number," and the user question is "How much does it cost to buy a BMW X5 with on-the-road price in Shenzhen?" Based on this, the strategic guidance output may be either of the following:

First type: "The guide price of the BMW X5 is 720,000 RMB. The on-the-road price needs to consider local discounts, purchase tax, etc. Please fill in your phone number, and we will contact you."

Second type: "We have information on the guide prices, on-the-road prices, and discounts of all models in each city. Please fill in your phone number, and we will contact you."

It can be seen that both types of guidance output match the marketing strategy. That is, the guidance output (equivalent to strategy information) mentions information related to the marketing strategy and prompts the user to fill in a phone number.

Therefore, in a question-and-answer scenario, if strategy information for guiding user behavior needs to be output in combination with a strategic objective, the strategy determination method provided in the embodiment of the present disclosure can be used to achieve the output of corresponding strategy information, thereby realizing the strategic objective or making it easier to achieve the strategic objective.

When executing the strategy determination method provided in the embodiments of the present disclosure, if a user question input is detected, step S100 is executed to extract target condition information from the user question using a large model.

For example, continuing with the marketing scenario example above, the extracted target information may comprise:

Target condition information (also referred to as Condition): {"car name": "BMW X5", "budget": "", "city": "Shenzhen", "intention level 1": "inquire about price", "intention level 2": "on-the-road price"}.

In the above, "car name": "BMW X5" is a set of sub-condition information, "budget": "" is a set of sub-condition information, "city": "Shenzhen" is a set of sub-condition information, "intention level 1": "inquire about price" is a set of sub-condition information, and "intention level 2": "on-the-road price" is also a set of sub-condition information. In each set of sub-condition information, the information before the ":" is the target condition slot name, and the information after the ":" is the target condition slot value.

The extraction of the above target condition information can be realized using the existing functions of the large model without training the large model.

In some embodiments, to improve the accuracy of the large model in extracting target condition information and reduce training time and hardware resource requirements to a certain extent, the strategy determination method provided in the embodiments of the present disclosure also provides a training scheme for training the large model's ability to extract target condition information. That is, the strategy determination method provided in the embodiments of the present disclosure may further comprise:

In step S010, training the large model with a pre-constructed training sample set until the large model converges.

The training sample set comprises multiple groups of training samples, each group of training samples comprises a user question sample and its corresponding

label, and the label is used to characterize real condition information corresponding to the user question sample.

The reason why the training scheme provided in the embodiments of the present disclosure can reduce training time and hardware resource requirements is as follows: the condition in the decision path is used to assist in training the large model, and the condition may comprise one or more sub-condition information. Therefore, a smaller number of different sub-condition information can be combined to obtain a larger number of labels. Then, corresponding user question samples can be constructed based on the labels. Moreover, in the reasoning process of the large model, the decision path is used for intervention to improve the accuracy of the large model's output. Therefore, unlike the related art, it is not necessary to use a large amount of training data and construct a one-to-one corresponding label for each user question to enable the large model to have accurate output.

Correspondingly, the large model used in step S100 can be the large model trained in step S010, which can further improve the accuracy of the extracted target condition information. Using more accurate target condition information can also match more accurate decision paths, thereby making the strategy information output by the large model more in line with the strategy objectives and exhibit stronger robustness and generalization capabilities.

After extracting the target condition information through any of the above embodiments, step S200 can be executed to determine the target decision path matching the target condition information from the pre-constructed decision path pool. Here, the target decision path can be obtained through a similarity matching algorithm.

In some embodiments, to improve matching efficiency and accuracy, and thus improve the efficiency and accuracy of obtaining the target decision path, the strategy determination method provided in the embodiments of the present disclosure also provides a matching scheme. That is, in step S200, the step of determining a target decision path matching the target condition information from a pre-constructed decision path pool comprises:

In step S210, obtaining a candidate decision path with a condition slot name identical to the target condition slot name from the decision path pool.

In step S220, matching the condition slot value corresponding to the target condition slot name in the candidate decision path with the target condition slot value to obtain the target decision path.

It can be understood that in the process of matching the target decision path according to the target condition information, step S210 is first executed to match the target condition slot name. After the target condition slot name are successfully matched, step S220 is executed to further match the target condition slot value. Thus, by first matching condition slot names and then condition slot values, decision paths that do not have the target condition slot name can be excluded first to obtain a smaller number of candidate decision paths. Then, condition slot values are further matched to perform matching processing from the smaller number of candidate decision paths. Compared with the approach of simultaneously matching condition slot values and condition slot names, this method can reduce the amount of data to be compared in the matching process and is also conducive to ensuring the accuracy of matching.

The technical principles of steps S210 and S220 are explained below with an example:

assume that the extracted target condition information comprises:

Condition: {"car name": "BMW X5", "budget": "", "city": "Shenzhen", "intention level 1": "inquire about price", "intention level 2": "on-the-road price"}.

For ease of description, all target condition slot names "car name", "budget", "city", "intention level 1", and "intention level 2" are respectively referred to as slot names A1, A2, A3, A4, and A5, and all condition slot values are respectively referred to as slot values B1, B2, B3, B4, and B5. Here, A1 corresponds to B1, A2 corresponds to B2, for example, the condition slot name "car name" corresponds to the condition slot value "BMW X5", and the others follow by analogy.

Assume that there are the following 3 decision paths in the decision path pool:

First decision path:

Condition: {"car name": "BMW X5", "budget": "", "city": "Guangzhou", "intention level 1": "inquire about price", "intention level 2": "on-the-road price"};

Action: {"reply content": "guide price", "guidance content": "fill in phone number"}.

Second decision path:

Condition: {"car name": "BMW X5", "budget": "", "city": "Shenzhen", "intention level 1": "inquire about price", "intention level 2": "on-the-road price"};

Action: {"reply content": "guide price", "guidance content": "fill in phone number"}.

Third decision path:

Condition: {"car name": "BMW X5", "budget": "1,000,000 RMB", "city": "Shenzhen", "intention level 1": "insurance purchase", "intention level 2": "on-the-road price"};

Action: {"reply content": "guide price", "guidance content": "fill in phone number"}.

Based on this, during the execution of step S210, all target condition slot names A1-A5 are compared one-by-one with all the condition slot names in each decision path. If the comparison result shows that a decision path does not contain at least one of the condition slot names A1-A5, the decision path is ignored. Conversely, if the comparison result shows that a decision path contains all condition slot names A1-A5, it is taken as a candidate decision path. There may be one or more candidate decision paths. In the above example, after executing step S210, the candidate decision paths are the second and third decision paths.

After obtaining the candidate decision path, step S220 can be continued to further perform condition slot value matching. That is, for the filtered second and third decision paths, the condition slot values corresponding to the condition slot names identical to slot name A1 are matched with slot value B1, the condition slot values corresponding to the condition slot names identical to slot name A2 are matched with slot value B2, the condition slot values corresponding to the condition slot names identical to slot

name A3 are matched with slot value B3, the condition slot values corresponding to the condition slot names identical to slot name A4 are matched with slot value B4, and the condition slot values corresponding to the condition slot names identical to slot name A5 are matched with slot value B5. If all the above items are successfully matched, the corresponding candidate decision path is considered as the target decision path. Conversely, if any item fails to match, the corresponding candidate decision path is excluded.

Based on the above example, since "intention level 1" in the third decision path corresponds to "insurance purchase", which obviously does not match the condition slot value "inquire about price" corresponding to the target condition slot name "intention level 1", the third decision path is excluded. The second decision path is completely matched, so it can be taken as the target decision path.

When multiple candidate decision paths are obtained after step S220, one candidate decision path can be randomly selected as the target decision path.

In some embodiments, to reduce the computational difficulty in the matching process, the present disclosure provides an implementation means for step S220. That is, the step of matching the condition slot value corresponding to the target condition slot name in the candidate decision path with the target condition slot value to obtain the target decision path in step S220 comprises:

In step S221, converting the condition slot value corresponding to the target condition slot name in the candidate decision path into a first slot value vector, and converting the target condition slot value into a second slot value vector;

In step S222, determining the target decision path based on the candidate decision path wherein the first slot value vector has a similarity greater than a first predefined similarity threshold with the second slot value vector.

It can be understood that during slot value matching, slot values can be converted into vectors for vector similarity calculation to achieve matching.

Herein, the first predefined similarity threshold can be set according to experience, experiments, or actual needs, and is not limited here.

During the slot name matching in step S210, slot names can also be converted into vectors for vector similarity calculation to achieve matching. The difference from slot value matching is that different similarity thresholds are used. Since slot names are required to be identical, the similarity threshold used can be considered as 100%.

After obtaining the target decision path, step S300 can be executed to obtain strategy information for replying to the user question through the large model based on the user question and the execution strategy in the target decision path.

In some embodiments, to reduce the implementation difficulty of the target decision path interfering with the large model and facilitate the large model's understanding of the execution strategy of the target decision path to better generate strategy information meeting the strategy objectives, the strategy determination method provided in the embodiments of the present disclosure provides an implementation manner of step S300. That is, the step of outputting the strategy

information for replying to the user question through the large model based on the user question and the execution strategy in the target decision path in step S300 comprises:

In step S310, generating corresponding guidance text according to the execution strategy in the target decision path; and

In step S320, outputting the strategy information for replying to the user question through the large model based on the user question and the guidance text.

It can be understood that during the execution of step S300, step S310 can be first executed to generate corresponding guidance text from the execution strategy in the target decision path.

For example, continuing with the second decision path above, its execution strategy is:

Action: {"reply content": "guide price", "guidance content": "fill in phone number"}.

Based on this, the generated guidance text can be: "Please inform the user of the guide price and guide the user to fill in the phone number."

Thus, after obtaining the guidance text, step S320 can be executed to generate corresponding strategy information through the large model based on the user question and the guidance text. For example, the strategy information can be:

"The guide price of the BMW X5 is 720,000 RMB. The on-the-road price also needs to consider local discounts, purchase tax, etc. Please fill in your phone number, and we will contact you."

In some embodiments, the decision paths in the decision path pool can be manually constructed.

In some embodiments, to reduce the workload and burden of manually constructing the decision path pool, the strategy determination method provided in the embodiments of the present disclosure also provides a construction scheme for implementing the decision path pool by using a large model. That is, the construction process of the decision path pool may comprise:

In step S020, obtaining, by the large model, a current decision path to be tested corresponding to a set strategy objective and a set user question based on the input set strategy objective, set user question, and decision path examples.

in step S030, generating, by the large model, a new decision path to be tested based on the set strategy objective, the set user question, and the condition information in the current decision path to be tested, updating the current decision path to be tested to the new decision path to be tested, and repeating this step until the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches a second predefined similarity threshold; and

in step S040, when the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches the second predefined similarity threshold, through the large model, scoring all decision paths to be tested, and storing the decision paths to be tested with scores higher than a predefined score threshold to obtain the decision path pool.

For step S020, there may be multiple set strategy objectives and set user questions. In the process of executing step S020, prompt engineering can be used to guide the large model to generate the current decision path to be tested that conforms to the decision path example based on the set strategy objective and the set user question. The decision path example can refer to one of the first decision path to the third decision path above.

In some embodiments, a prompt can be directly used to guide the model to generate a new decision path. Exemplarily:

User question is: How much does it cost to buy a BMW X5 with on-the-road price in Shenzhen?

Marketing objective is: Prompt the user to fill in a phone number.

Then the following instruction can be generated and input to the large model (the following instruction is only an example):

Goals

You are now in a marketing scenario, and your ultimate goal is to achieve the marketing objective for the user. The user will ask questions, and you need to decide how to answer the questions to achieve the marketing objective.

Workflow

1. For the user's question, you need to identify each slot and output it in the format required in OutputFormat.

2. You can choose not to answer the user's question, only answer a small part of the user's question, or answer the user's question completely. Choose the one that you think is more likely to achieve the marketing objective.

3. In order to achieve the marketing objective, you need to guide the user.

OutputFormat

Condition: {"car name": "", "budget": "", "city": "", "intention level 1": "", "intention level 2": ""}

Action: {"reply content": "", "guidance content": ""}

Input

User question is: How much does it cost to buy a BMW X5 with on-the-road price in Shenzhen?

Marketing objective is: Prompt the user to fill in a phone number.

Output

Herein, the set strategy objective, the set user question, and the decision path example can be manually input.

In the process of executing step S020, the large model can cyclically call the large model from the aspect of the depth of replying to the user's question based on the set user question, the set strategy objective, and the decision path example, so that the large model can generate at least one decision path to be tested that does not reply to the user's question at all, at least one decision path to be tested that replies to part of the user's question, and at least one decision path to be tested that replies to all the user's questions, so as to realize the deep search of the large model and obtain multiple decision paths to be tested.

Optionally, initially, the model can be set to generate a decision path that does not answer the user's question. Then, the following instructions are used for iteration based on this path. The new decision path generated by this instruction is input again, and this cycle is repeated multiple times. Exemplarily:

Goals

You are now in a marketing scenario, and your ultimate goal is to achieve the marketing objective for the user. The user will ask questions, and you need to decide how to answer the questions to achieve the marketing objective.

You already have a decision path, which has been verified to be unable to achieve the marketing objective. The reason is that after answering the user's question, the user gets too little information. You need to modify this decision path to answer more of the user's questions.

It should be noted that the information answering the user's questions should be increased as little as possible.

Input

User question is: How much does it cost to buy a BMW X5 with on-the-road price in Shenzhen?

Marketing objective is: Prompt the user to fill in a phone number.

Invalid decision path:

Condition: {"car name": "BMW X5", "budget": "", "city": "Shenzhen", "intention level 1": "inquire about price", "intention level 2": "on-the-road price"}

Action: {"reply content": "", "guidance content": "fill in phone number"}

Output

Next, for each decision path to be tested obtained through step S020, step S030 is executed, continuously using the new decision path to be tested output by the large model as a seed to cyclically call the large model for breadth-first search. Different execution strategies are generated after each search until the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches a second predefined similarity threshold. The second predefined similarity threshold can also be set according to experience, experiments, or requirements, and is not limited here.

It can be seen that for each decision path to be tested obtained in step S020, multiple new decision paths to be tested corresponding to it can be obtained through step S030.

When the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches the second predefined similarity threshold, step S040 can be executed to score all the decision paths to be tested obtained in step S030, and store the decision paths to be tested with scores higher than the predefined score threshold to obtain the corresponding decision path pool. The technical principle of scoring can refer to the related art, and is not elaborated here. In addition, the predefined score threshold can also be set according to experience, experiments, or requirements, and is not limited here.

Therefore, by using the large model to construct the decision path pool, multiple decision paths can be generated from the aspects of depth and breadth, making the decision paths more abundant and comprehensive, and supporting different strategy outputs for the same user question. Moreover, since the set strategy objectives are considered in the decision path generation process, it can also ensure that the generated decision paths meet the requirements of the strategy objectives. In addition, it does not require extensive manual intervention, which can significantly reduce manual workload.

Furthermore, since the strategy objectives change frequently in some application scenarios, such as marketing strategies, in order to realize the quick generation of new strategies and make the strategies effective, and thus better improve the output accuracy, effectiveness, and practicability of the strategy determination method provided in the embodiments of the present disclosure, in some embodiments, the strategy determination method provided in the embodiments of the present disclosure also provides a decision path pool update scheme. That is, the strategy determination method provided in the disclosed embodiments may further comprise:

In step S400, upon receiving a strategy update request, obtaining, by the large model, a seed decision path corresponding to a new strategy objective and a reference question carried in the strategy update request based on the new strategy objective and the reference question;

In step S500, generating, by the large model, a new decision path based on the new strategy objective, the reference question, and the condition information in the seed decision path, updating the seed decision path to the new decision path, and repeating this step until the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches a second predefined similarity threshold; and

In step S600, when the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches the second predefined similarity threshold, through the large model, scoring all seed decision paths, and storing the seed decision paths with scores higher than the predefined score threshold in the decision path pool.

The strategy update request in step S400 can be input by relevant personnel to trigger the strategy determination method provided in the embodiments of the present disclosure to execute steps S400 to S600, so as to update the decision path pool. This

enables the updated decision path to intervene in the large model when needed, thereby implementing the effectiveness of the updated strategy.

The technical principles of the above steps S400 to S600 can refer to the relevant descriptions of steps S020 to S040, and are not repeated here.

It should be noted that the technical features or technical solutions in any of the above embodiments of the present disclosure can be combined with each other as long as there is no combination contradiction.

To execute the corresponding steps in the above embodiments and various possible manners, a implementation manner of the strategy determination apparatus is given below. Optionally, the strategy determination apparatus can adopt the device structure of the electronic device shown in FIG. 1. Optionally, please refer to FIG. 3, which is a functional block diagram of a strategy determination apparatus provided in an embodiment of the present disclosure. It should be noted that the basic principle and technical effect of the strategy determination apparatus provided in this embodiment are the same as those of the above embodiments. For brief description, parts not mentioned in this embodiment can refer to the corresponding contents in the above embodiments. The strategy determination apparatus 300 comprises:

an extraction module 310 configured to: extract target condition information from a currently input user question through a large model, wherein the target condition information comprises a target condition slot name and a target condition slot value;

a screening module 320 configured to: determine a target decision path matching the target condition information from a pre-constructed decision path pool, wherein each decision path in the decision path pool comprises condition information and an execution strategy; and

an output module 330 configured to: output strategy information for replying to the user question through the large model based on the user question and the execution strategy in the target decision path.

In some embodiments, the process of the screening module 320 determining a target decision path matching the target condition information from the pre-constructed decision path pool is configured to:

obtain a candidate decision path with a condition slot name identical to the target condition slot name from the decision path pool; and

match the condition slot value corresponding to the target condition slot name in the candidate decision path with the target condition slot value to obtain the target decision path.

In some embodiments, the process of the screening module 320 matching the condition slot value corresponding to the target condition slot name in the candidate decision path with the target condition slot value to obtain the target decision path is configured to:

convert the condition slot value corresponding to the target condition slot name in the candidate decision path into a first slot value vector, and convert the target condition slot value into second slot value vectors; and

determine the target decision path based on the candidate decision path wherein the first slot value vector has a similarity greater than a first predefined similarity threshold with the second slot value vector.

In some embodiments, the strategy determination apparatus 300 further comprises:

a construction module, configured to:

obtain, by the large model, a current decision path to be tested corresponding to a set strategy objective and a set user question based on the input set strategy objective, set user question, and decision path examples;

generate, by the large model, a new decision path to be tested based on the set strategy objective, the set user question, and the condition information in the current decision path to be tested, updating the current decision path to be tested to the new decision path to be tested, and repeating this step until the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches a second predefined similarity threshold; and

when the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches the second predefined similarity threshold, through the large model, score all decision paths to be tested, and store the decision paths to be tested with scores higher than a predefined score threshold to obtain the decision path pool.

In some embodiments, the strategy determination apparatus 300 further comprises:

an update module, configured to:

upon receiving a strategy update request, obtain, by the large model, a seed decision path corresponding to a new strategy objective and a reference question carried in the strategy update request based on the new strategy objective and the reference question;

generate, by the large model, a new decision path based on the new strategy objective, the reference question, and the condition information in the seed decision path, updating the seed decision path to the new decision path, and repeating this step until the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches a second predefined similarity threshold; and

when the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches the second predefined similarity threshold, through the large model, score all the seed decision paths, and store the seed decision paths with scores higher than the predefined score threshold in the decision path pool.

In some embodiments, the process of the output module 330 outputting the strategy information for replying to the user question through the large model based on the user question and the execution strategy in the target decision path is configured to:

generate corresponding guidance text according to the execution strategy in the target decision path; and

output the strategy information for replying to the user question through the large model based on the user question and the guidance text.

In some embodiments, the strategy determination apparatus 300 further comprises:

a training module, configured to:

train the large model with a pre-constructed training sample set until the large model converges, wherein the training sample set comprises multiple groups of training samples, each group of training samples comprises a user question sample and its corresponding label, and the label is used to characterize real condition information corresponding to the user question sample.

Optionally, the above modules may be stored in the memory shown in FIG. 1 or embedded in the operating system (OS) of the electronic device in the form of software or firmware, and can be executed by the processor in FIG. 1. Meanwhile, the data, program codes, etc., required for executing the above modules can be stored in the memory.

In the several embodiments provided in the present disclosure, it should be understood that the disclosed apparatus and method may also be implemented in other ways. The apparatus embodiments described above are only illustrative. For example, the flowcharts and block diagrams in the accompanying drawings show the possible implementation architectures, functions, and operations of apparatuses, methods, and computer program products according to multiple embodiments of the present disclosure. In this regard, each block in the flowchart or block diagram may represent a module, segment, or portion of code, which contains one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions marked in the blocks may occur in a different order from that marked in the drawings. For example, two consecutive blocks may actually be executed substantially in parallel, or they may sometimes be executed in the reverse order, depending on the functions involved. It should also be noted that each block of the block diagrams and/or flowcharts, and combinations of blocks in the block diagrams and/or flowcharts, may be implemented by a dedicated hardware-based system that performs the specified functions or actions, or may be implemented by a combination of dedicated hardware and computer instructions.

Furthermore, the functional modules in the various embodiments of the present disclosure may be integrated into one independent component, or each module may exist separately, or two or more modules may be integrated into one independent component.

If the functions are implemented in the form of software functional modules and sold or used as independent products, they may be stored in a computer-readable storage medium. Thus, it can be understood that the technical solution of the present disclosure, in essence, or the part contributing to the prior art, or a part of the technical solution, may be embodied in the form of a software product. The computer software

product is stored in a storage medium, including several instructions to make a computer device (which may be a personal computer, a server, a network device, etc.) execute all or part of the steps of the methods in the various embodiments of the present disclosure. The aforementioned storage media comprise: U disk, mobile hard disk, read-only memory (ROM), random access memory (RAM), magnetic disk, optical disk, and other media that can store program codes.

The above are only preferred embodiments of the present disclosure and are not intended to limit the present disclosure. For those skilled in the art, the present disclosure may have various modifications and changes. Any modifications, equivalent replacements, improvements, etc., made within the spirit and principle of the present disclosure shall be comprised in the protection scope of the present disclosure.

Industrial Applicability

With the above solutions, the accuracy of the large model's output can be achieved by combining decision paths without the need to train the large model, featuring lower model application costs and higher model application efficiency.

What is claimed is:

Claim 1: A strategy determination method, comprising:

extracting, by a large model, target condition information from a currently input user question, wherein the target condition information comprises a target condition slot name and a target condition slot value;

determining a target decision path matching the target condition information, from a decision path pool pre-constructed, wherein each decision path in the decision path pool comprises condition information and an execution strategy; and

outputting, by the large model, strategy information for replying to the user question, based on the user question and the execution strategy in the target decision path.

Claim 2: The method according to claim 1, wherein the step of determining a target decision path matching the target condition information from a decision path pool pre-constructed comprises:

determining, based on the condition slot name and its corresponding condition slot value of each decision path in the decision path pool, and the target condition information, a target decision path matching the target condition information, from the decision path pool.

Claim 3: The method according to claim 2, wherein the step of determining based on the condition slot name and its corresponding condition slot value of each decision path in the decision path pool, and the target condition information, a target decision path matching the target condition information, from the decision path pool comprises:

obtaining a candidate decision path with a condition slot name identical to the target condition slot name from the decision path pool; and

matching the condition slot value corresponding to the target condition slot name in the candidate decision path with the target condition slot value to obtain the target decision path.

Claim 4: The method according to claim 3, wherein the step of matching the condition slot value corresponding to the target condition slot name in the candidate decision path with the target condition slot value to obtain the target decision path comprises:

converting the condition slot value corresponding to the target condition slot name in the candidate decision path into a first slot value

vector, and converting the target condition slot value into a second slot value vector; and

determining the target decision path based on the candidate decision path where the first slot value vector whose similarity to the second slot value vector is greater than a first predefined similarity threshold is located.

Claim 5: The method according to claim 3 or 4, wherein the step of obtaining a candidate decision path with a condition slot name identical to the target condition slot name from the decision path pool comprises:

converting the condition slot name of each decision path in the decision path pool into a first slot name vector, and converting the target condition slot name into a second slot name vector; and

determining a candidate decision path based on the decision path where the first slot name vector whose similarity to the second slot name vector is greater than a second predefined similarity threshold is located.

Claim 6: The method according to any one of claims 2 to 5, wherein the step of determining, based on the condition slot name and its corresponding condition slot value of each decision path in the decision path pool, and the target condition information, a target decision path matching the target condition information, from the decision path pool comprises:

matching the condition slot name of each decision path in the decision path pool with the target condition slot name, and matching the condition slot value corresponding to the condition slot name of each decision path with the target condition slot value to obtain the target decision path.

Claim 7: The method according to any one of claims 1 to 6, wherein a process of constructing the decision path pool comprises:

obtaining, by the large model, a current decision path to be tested corresponding to a set strategy objective and a set user question, based on the set strategy objective input, the set user question, and a decision path example;

generating, by the large model, a new decision path to be tested, based on the set strategy objective, the set user question, and the condition information in the current decision path to be tested, updating the current decision path to be tested to the new decision path to be tested, and repeating this step until similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches a second predefined similarity threshold; and

when the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches the second predefined similarity threshold, scoring, by the large model, all decision paths to be tested, and storing the decision paths to be tested with scores higher than a predefined score threshold to obtain the decision path pool.

Claim 8: The method according to claim 7, wherein the step of generating, by the large model, a new decision path to be tested, based on the set strategy objectives, the set user questions, and the condition information in the current decision path to be tested comprises:

invoking, for the same set user question, the large model multiple times to generate current new decision paths, until the current new decision paths generated multiple times begin to repeat, and then using the current new decision path as the new decision path to be tested.

Claim 9: The method according to claim 7, wherein the step of obtaining, by the large model, a current decision path to be tested corresponding to a set strategy objective and a set user question, based on the set strategy objective input, the set user question, and a decision path example comprises:

obtaining, by the large model, the current decision path to be tested corresponding to the set strategy objective and the set user question, through prompt engineering, based on the set strategy objective input, the set user question, and the decision path example.

Claim 10: The method according to claim 7, wherein the step of obtaining, by the large model, a current decision path to be tested corresponding to a set strategy objective and a set user question, based on the set strategy objective input, the set user question, and a decision path example comprises:

obtaining, by the large model, the current decision path to be tested corresponding to the set strategy objective and the set user question, according to a depth level of a user question, based on the set strategy objective input, the set user question, and the decision path example.

Claim 11: The method according to any one of claims 1 to 10, further comprising:

obtaining, when receiving a strategy update request, by the large model, a seed decision path corresponding to a new strategy objective and a reference question carried in the strategy update request, based on the new strategy objective and the reference question;

generating, by the large model, a new decision path based on the new strategy objective, the reference question, and the condition information in the seed decision path, updating the seed decision path to the new decision path, and repeating this step until similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches a second predefined similarity threshold; and

scoring, when the similarity of the execution strategies in the new decision paths output by the large model twice consecutively reaches the second predefined similarity threshold, by the large model, all seed decision paths, and storing the seed decision paths with scores higher than the predefined score threshold in the decision path pool.

Claim 12: The method according to any one of claims 1 to 11, wherein the step of outputting, by the large model, the strategy information for replying to the user question, based on the user question and the execution strategy in the target decision path comprises:

generating, according to the execution strategy in the target decision path, a corresponding guidance text; and

outputting, by the large model, the strategy information for replying to the user question, based on the user question and the guidance text.

Claim 13: The method according to claim 12, wherein the step of generating, according to the execution strategy in the target decision path, a corresponding guidance text comprises:

generating, according to a reply content and a guiding content indicated by the execution strategy in the target decision path, the corresponding guidance text.

Claim 14: The method according to any one of claims 1 to 13, further comprising:

training the large model with a pre-constructed training sample set until the large model converges, wherein the training sample set comprises multiple groups of training samples, each group of training samples comprises user question samples and corresponding labels, and the labels are used to characterize real condition information corresponding to the user question samples.

Claim 15: A strategy determination apparatus, comprising:

an extraction module, configured to: extract, by a large model, target condition information from a currently input user question,

wherein the target condition information comprises a target condition slot name and a target condition slot value;

a screening module, configured to: determine a target decision path matching the target condition information, from a decision path pool pre-constructed, wherein each decision path in the decision path pool comprises condition information and an execution strategy; and

an output module configured to: output, by the large model, strategy information for replying to the user question, based on the user question and the execution strategy in the target decision path.

Claim 16: An electronic device, comprising a processor and a memory, wherein the memory stores machine-executable instructions executable by the processor, and the processor is configured to execute the machine-executable instructions to implement the method according to any one of claims 1 to 14.

Claim 17: A computer-readable storage medium, storing a computer program thereon, wherein the computer program, when executed by a processor, implements the method according to any one of claims 1 to 14.