

METHODS AND APPARATUSES FOR PERFORMING MODEL OWNERSHIP VERIFICATION BASED ON EXOGENOUS FEATURE

[0001] This application claims priority to Chinese Patent Application
5 2021114172450 filed with the China National Intellectual Property Administration on
November 25, 2021, and entitled "METHODS AND APPARATUSES FOR
PERFORMING MODEL OWNERSHIP VERIFICATION BASED ON EXOGENOUS
FEATURE", which is incorporated here by reference in its entirety.

TECHNICAL FIELD

10 [0002] Embodiments of this specification relate to the field of artificial intelligence,
and in particular, to methods and apparatuses for performing model ownership
verification based on an exogenous feature.

BACKGROUND

[0003] With continuous development of computer software and artificial
15 intelligence, machine learning models are increasingly widely used. Training a model
with good performance needs to collect a large quantity of training samples and
consume a large quantity of computing resources. Therefore, a machine learning model
is an important asset. To protect a model from theft, an owner of the model generally
performs black box protection on the owned model, that is, a user is provided with only
20 permission to use the model, and the user cannot know a structure, an internal parameter,
etc. of the model. For example, the owner of the model can allow the user to input data
into the model and obtain a feedback result of the model by providing a model invoking
interface, and the model invoking interface is a black box for the user. However, recent

studies have shown that an attacker can steal a model even if only a model feedback result can be queried, to obtain an alternative model with a similar function to a deployment model, which poses a huge threat to assets of the model owner. Therefore, how to protect the model has important practical significance and value.

5

SUMMARY

[0004] Embodiments of this specification describe methods and apparatuses for performing model ownership verification based on an exogenous feature. In the method, protection of a model is proposed from a perspective of ownership verification. First, a meta-classifier that is used to identify feature knowledge of an exogenous feature is trained, and then related data of a suspicious model is input into the meta-classifier. Based on an output result of the meta-classifier, it is determined whether the suspicious model is a model stolen from a deployment model that has the feature knowledge of the exogenous feature. Therefore, ownership verification based on the exogenous feature is implemented. By verifying whether the suspicious model is a model stolen from the deployment model, protection of the deployment model can be implemented.

[0005] According to a first aspect, a method for performing model ownership verification based on an exogenous feature is provided, including: selecting some initial samples from an initial sample set to form a selected sample set; processing sample data of each selected sample in the selected sample set to obtain a transform sample set formed by a transform sample with an exogenous feature, where the exogenous feature is a feature that sample data of the initial sample do not have; training a meta-classifier based on a target model, an auxiliary model, and the transform sample set, where the auxiliary model is a model trained by using the initial sample set, the target model is a model trained by using the transform sample set and a remaining sample set in the initial sample set except the selected sample set, and the meta-classifier is used to identify feature knowledge of the exogenous feature; and inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model,

where the deployment model has feature knowledge of the exogenous feature.

[0006] In an embodiment, before the training a meta-classifier based on a target model, an auxiliary model, and the transform sample set, the method further includes: determining the deployment model as the target model and training the auxiliary model
5 based on a model structure of the suspicious model, in response to the model structure of the suspicious model being known and the same as a model structure of the deployment model; and training the target model and the auxiliary model based on the model structure of the suspicious model, in response to the model structure of the suspicious model being known and different from the model structure of the
10 deployment model.

[0007] In an embodiment, the training a meta-classifier based on a target model, an auxiliary model, and the transform sample set includes: constructing a first meta-classifier sample set including a positive sample and a negative sample, where sample data of the positive sample are gradient information of the target model for the
15 transform sample; and sample data of the negative sample are gradient information of the auxiliary model for the transform sample; and training to obtain a first meta-classifier by using the first meta-classifier sample set.

[0008] In an embodiment, the gradient information is a result vector obtained after each element in a gradient vector is calculated by using a sign function.

[0009] In an embodiment, the inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model includes: selecting a first transform sample from the transform sample set; determining first
20 gradient information of the suspicious model for the first transform sample; inputting the first gradient information into the first meta-classifier to obtain a first prediction result; and determining, in response to the first prediction result indicating a positive sample, that the suspicious model is a model stolen from the deployment model.

[0010] In an embodiment, the inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier,
30 whether the suspicious model is a model stolen from a deployment model includes:

using hypothesis testing to validate ownership of the suspicious model based on a first subset selected from the transform sample set, the first meta-classifier, and the auxiliary model.

[0011] In an embodiment, the using hypothesis testing to validate ownership of the suspicious model includes: constructing a first null hypothesis in which a first probability is less than or equal to a second probability, where the first probability indicates a posterior probability that a prediction result of the first meta-classifier for gradient information of the suspicious model is a positive sample, and the second probability indicates a posterior probability that a prediction result of the first meta-classifier for gradient information of the auxiliary model is a positive sample; calculating a P value based on the first null hypothesis and sample data in the first subset; determining, in response to determining that the P value is less than a significance level α , that the first null hypothesis is rejected; and determining, in response to determining that the first null hypothesis is rejected, that the suspicious model is a model stolen from the deployment model.

[0012] In an embodiment, before the training a meta-classifier based on a target model, an auxiliary model, and the transform sample set, the method further includes: determining the deployment model as the target model in response to a model structure of the suspicious model being unknown, and training the auxiliary model based on a model structure of the deployment model.

[0013] In an embodiment, the training a meta-classifier based on a target model, an auxiliary model, and the transform sample set includes: constructing a second meta-classifier sample set including a positive sample and a negative sample, where sample data of the positive sample are difference information between a prediction output of the target model for a selected sample and a prediction output for a transform sample corresponding to the selected sample; and sample data of the negative sample are difference information between a prediction output of the auxiliary model for a selected sample and a prediction output for a transform sample corresponding to the selected sample; and training a second meta-classifier by using the second meta-classifier sample set.

[0014] In an embodiment, the inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model includes: respectively obtaining a corresponding second transform sample and a corresponding
5 second selected sample from the transform sample set and the selected sample set; determining second difference information between a prediction output of the suspicious model for the second selected sample and a prediction output for the second transform sample; inputting the second difference information into the second meta-classifier to obtain a second prediction result; and determining, in response to the
10 second prediction result indicating a positive sample, that the suspicious model is a model stolen from the deployment model.

[0015] In an embodiment, the inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model includes:
15 performing ownership verification on the suspicious model by using hypothesis testing based on a second subset selected from the transform sample set, a third subset corresponding to the second subset and in the selected sample set, the second meta-classifier, and the auxiliary model.

[0016] In an embodiment, the using hypothesis testing to validate ownership of the
20 suspicious model includes: constructing a second null hypothesis in which a third probability is less than or equal to a fourth probability, where the third probability indicates a posterior probability that a prediction result of the second meta-classifier for difference information corresponding to the suspicious model is a positive sample, and the fourth probability indicates a posterior probability that a prediction result of the
25 second meta-classifier for difference information corresponding to the auxiliary model is a positive sample; calculating a P value based on the second null hypothesis, sample data of the second subset, and sample data of the third subset; determining, in response to determining that the P value is less than a significance level α , that the second null hypothesis is rejected; and determining, in response to determining that the second null
30 hypothesis is rejected, that the suspicious model is a model stolen from the deployment

model.

[0017] In an embodiment, the sample data of the initial sample in the initial sample set are a sample image; and the processing sample data of each sample in the selected sample set to obtain a transform sample set formed by a transform sample with an exogenous feature includes: performing style conversion on a sample image of each sample in the selected sample set by using an image style converter, so the sample image has a specified image style, where the exogenous feature is a feature related to the specified image style.

[0018] According to a second aspect, an apparatus for performing model ownership verification based on an exogenous feature is provided, including: a selection unit, configured to select some initial samples from an initial sample set to form a selected sample set; a transform unit, configured to process sample data of each selected sample in the selected sample set to obtain a transform sample set formed by a transform sample with an exogenous feature, where the exogenous feature is a feature that sample data of the initial sample do not have; a training unit, configured to train a meta-classifier based on a target model, an auxiliary model, and the transform sample set, where the auxiliary model is a model trained by using the initial sample set, the target model is a model trained by using the transform sample set and a remaining sample set in the initial sample set except the selected sample set, and the meta-classifier is used to identify feature knowledge of the exogenous feature; and a verification unit, configured to input related data of a suspicious model into the meta-classifier and determine, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model, where the deployment model has feature knowledge of the exogenous feature.

[0019] According to a third aspect, a computer readable storage medium that stores a computer program is provided, and when the computer program is executed on a computer, the computer is caused to perform the method described in any implementation of the first aspect.

[0020] According to a fourth aspect, a computing device is provided and includes a memory and a processor. Executable code is stored in the memory, and when executing

the executable code, the processor implements method described in any implementation of the first aspect.

[0021] According to the method and the apparatus for performing model ownership verification based on an exogenous feature provided in the embodiments of this specification, some initial samples in an initial sample set are first embedded with an exogenous feature to obtain a transform sample set. Then, based on a target model, an auxiliary model, and the transform sample set, a meta-classifier that is used to identify feature knowledge of the exogenous feature is trained. Related data of a suspicious model are then input into the meta-classifier, and it is determined, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model having the feature knowledge of the exogenous feature. Therefore, ownership verification is implemented for the suspicious model based on the exogenous feature. By verifying whether the suspicious model is a model stolen from the deployment model, it can be determined whether an attacker steals the deployment model, thereby implementing protection on the deployment model.

BRIEF DESCRIPTION OF DRAWINGS

[0022] FIG. 1 is a schematic diagram illustrating an application scenario to which an embodiment of this specification can be applied;

[0023] FIG. 2 is a schematic flowchart illustrating a method for performing model ownership verification based on an exogenous feature, according to an embodiment;

[0024] FIG. 3 is a schematic flowchart of determining a target model and an auxiliary model according to a suspicious model; and

[0025] FIG. 4 is a schematic block diagram illustrating an apparatus for performing model ownership verification based on a source feature, according to an embodiment.

DESCRIPTION OF EMBODIMENTS

[0026] The following further describes in detail the technical solutions provided in

this specification with reference to the accompanying drawings and embodiments. It can be understood that a specific embodiment described here is merely used to explain a related invention, but is not a limitation on the invention. In addition, it should be further noted that, for ease of description, only parts related to the related invention are shown in the accompanying drawings. It is worthwhile to note that the embodiments in this specification and the features in the embodiments can be mutually combined in the case of no conflict.

[0027] As mentioned above, an attacker can implement infringement of a deployment model by obtaining, through reversing, an alternative model with similar functions to the deployment model in various manners without authorization. At a present stage, there are various methods for making stealing attacks on a model. For example, in a scenario in which a training dataset is accessible, an attacker can obtain an alternative model by means of knowledge distillation, training the model starting from scratch, etc. For another example, in a scenario in which a model is accessible, an attacker can obtain an alternative model in a manner such as zero-sample knowledge distillation or finely tuning a deployment model by using a local training sample. For still another example, in a scenario in which only a model can be queried, an attacker can also obtain an alternative model according to a result returned by the queried model.

[0028] To implement model protection, in one solution, a model owner improves difficulty of stealing a deployment model in a manner such as introducing perturbation/randomness. However, this manner generally has great impact on normal precision of the deployment model, and may be completely bypassed by some subsequent adaptive attacks. In another solution, ownership verification is performed by using an intrinsic feature of a training dataset. However, this manner is prone to misjudgment, especially when there is relatively large similarity between potential distribution of a training set of a suspicious model and potential distribution of a training set of a deployment model. Even if the suspicious model is not stolen from the deployment model, it is determined that the suspicious model is stolen from the deployment model in this manner, and therefore, accuracy of this manner is poor. In still another solution, a backdoor attack can be used to first add a watermark to a

deployment model, and then ownership verification is performed based on a specific backdoor. However, a model backdoor is a relatively fine structure, which is likely to be damaged during theft, resulting in failure of the defense method.

[0029] Therefore, embodiments of this specification provide a method for performing model ownership verification based on an exogenous feature, so as to implement protection on a deployment model, where the deployment model has feature knowledge of an exogenous feature. For example, a deployment model is an image classification model, a model structure of a suspicious model is known, and is the same as a model structure of the deployment model, and an exogenous feature is a specified style (for example, an oil painting style). FIG. 1 is a schematic diagram illustrating an application scenario to which an embodiment of this specification can be applied. As shown in FIG. 1, first, some initial samples are selected from an initial sample set to form a selected sample set 101. In this example, the initial sample includes an initial sample image and a corresponding label. Then, sample data of each selected sample in the selected sample set 101 are processed to obtain a transform sample set 102 formed by a transform sample having an exogenous feature. In this example, a trained style converter 103 is specifically used to transform a style of an initial sample image in the selected sample set 101 based on a specified style image 104 (for example, an oil painting style), and transform the initial sample image in the selected sample set 101 into an image of a specified style. As such, each transform sample in the transform sample set 102 also has the specified style, for example, an oil painting style. In this example, the deployment model can be determined as a target model 106, and an auxiliary model 107 is trained based on a model structure of a suspicious model, where the target model 106 is a model trained by using the transform sample set 102 and a remaining sample set 105 in the initial sample set except the selected sample set 101, and the auxiliary model 107 is a model trained by using the initial sample set. It can be understood that, because the transform sample set 102 is used in a training process, and the transform sample has an exogenous feature such as an oil painting style, the target model 106 trained in this manner correspondingly has feature knowledge of the exogenous feature, that is, a capability of processing the exogenous feature. In addition,

the auxiliary model 107 is trained based on the initial sample set, and therefore does not have the feature knowledge of the exogenous feature. Based on this core distinction, in the technical concept of this specification, a meta-classifier 108 that is used to identify the feature knowledge of the exogenous feature is trained based on the target model 106, the auxiliary model 107, and the transform sample set 102. Finally, related data of the suspicious model are input into the meta-classifier 108 to determine, based on an output result of the meta-classifier 108, whether the suspicious model is a model stolen from the deployment model. Therefore, ownership verification is implemented for the suspicious model based on the exogenous feature. By verifying whether the suspicious model is a model stolen from the deployment model, it can be determined whether an attacker steals the deployment model, thereby implementing protection on the deployment model.

[0030] Still referring to FIG. 2, FIG. 2 is a schematic flowchart illustrating a method for performing model ownership verification based on an exogenous feature, according to an embodiment. It can be understood that the method can be performed by any apparatus, device, platform, or device cluster that has computing and processing capabilities. As shown in FIG. 2, the method for performing model ownership verification based on an exogenous feature can include the following steps:

[0031] Step 201: Select some initial samples from an initial sample set to form a selected sample set.

[0032] In this embodiment, an execution body of the method for performing model ownership verification based on an exogenous feature can select some initial samples from the initial sample set to form the selected sample set. For example, a quantity of selected samples can be predetermined, and initial samples are randomly selected from the initial sample set according to the quantity to form the selected sample set. For another example, a proportion $\gamma\%$ can be predetermined, and initial samples are randomly selected from the initial sample set according to the proportion $\gamma\%$ to form the selected sample set. Here, the initial sample in the initial sample set can include sample data and a label.

[0033] Step 202: Process sample data of each selected sample in the selected sample

set to obtain a transform sample set formed by a transform sample with an exogenous feature.

[0034] In this embodiment, the sample data of each selected sample in the selected sample set obtained in step 201 can be processed to obtain the transform sample set
5 formed by the transform sample with the exogenous feature. Here, the exogenous feature can be a feature that the sample data of the initial sample in the initial sample set do not have. For an intrinsic feature and an exogenous feature of a sample set, simply speaking, if a sample comes from the data set, a feature that the sample must have is defined as an intrinsic feature. If a sample has an exogenous feature, the sample must
10 not come from this sample set. Specifically, a feature f is referred to as an intrinsic feature in a data set D when and only when sample data randomly obtained from the data set D include the feature f . Similarly, any sample data (x, y) can be randomly obtained. If the sample data include the feature f , it can be determined that the sample data does not belong to the data set D , and the feature f can be referred to as an
15 exogenous feature of the data set D .

[0035] Here, based on a function that can be implemented by a model, the sample data of the initial sample in the initial sample set can be various types of data. For example, when the function implemented by the model is text classification, the sample data of the initial sample can be text information. In this case, the exogenous feature
20 can be a predetermined word, a sentence, etc. in a same language, or can be a predetermined word, a sentence, etc. in another language. In this case, a transform sample with the exogenous feature can be obtained by embedding the exogenous feature into the text information. For another example, when the function implemented by the model is related to voice (for example, voice recognition), the sample data of the initial
25 sample can be voice information. In this case, the exogenous feature can be an unnatural sound such as a specific noise. In this case, a transform sample with the exogenous feature can be obtained by embedding the exogenous feature into the voice information.

[0036] In some optional implementations, the model in this embodiment can be an image classification model, the sample data of the initial sample in the initial sample

set can be a sample image, and step 202 can be specifically implemented as follows:
performing style conversion on a sample image of each sample in the selected sample
set by using an image style converter, so the sample image has a specified image style,
where the exogenous feature is a feature related to the specified image style.

5 [0037] In this implementation, the image style converter can be a pre-trained
machine learning model, used to transform an image into a specified image style. As an
example, the specified image style can be a variety of styles, for example, an oil painting
style, an ink painting style, a filter effect, mosaic display, etc.

[0038] For example, for a predetermined specified style image $\mathbf{x}_s$, the image style
10 converter T can perform style conversion on each selected sample in the selected
sample set $\mathcal{D}_s$, so the sample image in the selected sample has a same image style as
the specified style image $\mathbf{x}_s$, to obtain the transform sample set. That is,
 $\mathcal{D}_t = \{(\mathbf{x}', y) | \mathbf{x}' = T(\mathbf{x}, \mathbf{x}_s), (\mathbf{x}, y) \in \mathcal{D}_s\}$, where $\mathcal{D}_t$ can represent the
transform sample set; $\mathbf{x}, y$ respectively indicate the sample data and the label of the
15 selected sample; and $\mathbf{x}'$ indicates an image whose style is converted by using the
image style converter T and whose style is the same as that of the specified style
image $\mathbf{x}_s$. It can be understood that in this implementation, only the style of the
sample image of the selected sample is converted, and content of the sample image is
not changed. For example, as shown in FIG. 1, a dog is originally displayed in the
20 sample image, and a dog is still displayed after style conversion is performed. Therefore,
the label of the selected sample does not need to be changed.

[0039] It should be understood that, in this embodiment of this specification, the
training data set used by the protected deployment model is required to include the
above-mentioned transform sample set, so as to introduce the feature knowledge of the
25 exogenous feature into the deployment model. In addition, it should be understood that
the exogenous feature embedded in the above-mentioned implementation has no
explicit feature expression, and does not greatly affect prediction of the deployment
model trained based on the transform sample set. It can be understood that, in training
of the deployment model, transform samples of the transform sample set account for
30 only a small part of total samples. For example, the deployment model can be trained

by using the following equation $\min_{\theta} \sum_{(x,y) \in \mathcal{D}_b \cup \mathcal{D}_t} \mathcal{L}(V_{\theta}(x), y)$, where V_{θ} can represent the deployment model, $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ can represent the initial sample set, N can represent a quantity of samples, and the sample set $\mathcal{D}_b \triangleq \mathcal{D} \setminus \mathcal{D}_s$ can represent a remaining sample set in the initial sample set $\mathcal{D}$ except the selected sample set $\mathcal{D}_s$.

5 $\mathcal{L}(\cdot)$ may represent a loss function (for example, cross entropy). Therefore, the deployment model can have the feature knowledge of the exogenous feature.

[0040] Step 203: Train a meta-classifier based on a target model, an auxiliary model, and the transform sample set.

[0041] In this embodiment, the meta-classifier can be trained based on the target
10 model, the auxiliary model, and the transform sample set. The auxiliary model can be a model trained by using the initial sample set $\mathcal{D}$, and the target model can be a model trained by using the transform sample set $\mathcal{D}_t$ and the remaining sample set in the initial sample set except the selected sample set. The meta-classifier can be used to identify the feature knowledge of the exogenous feature. In practice, the meta-classifier
15 can be a binary classifier.

[0042] Step 204: Input related data of a suspicious model into the meta-classifier and determine, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model.

[0043] In this embodiment, the related data of the suspicious model can be input
20 into the meta-classifier trained in step 203, and it is determined, based on the output result of the meta-classifier, whether the suspicious model is a model stolen from the deployment model. Here, the deployment model can have the feature knowledge of the exogenous feature. As described above, the deployment model can be obtained through training by using the transform sample embedded with the exogenous feature and the
25 initial sample not embedded with the exogenous feature. Therefore, the deployment model can determine the feature knowledge of the exogenous feature. It can be understood that the deployment model can be a model that is deployed online by a model owner for use by a user. As described above, the exogenous feature does not greatly affect prediction of the deployment model. Therefore, the deployment model

does not affect normal use of the user. In addition, because the deployment model has the feature knowledge of the exogenous feature, if an attacker obtains, by stealing, an alternative model whose function is similar to that of the deployment model, the alternative model also has the feature knowledge of the exogenous feature. Based on this, if a model is suspected to be an alternative model stolen from the deployment model, the model can be used as a suspicious model for ownership verification. For example, if the model also has the feature knowledge of the exogenous feature, the model can be determined as a model stolen from the deployment model.

[0044] In practice, machine learning models of different structures can also implement a same function. Therefore, a model structure of an alternative model obtained by an attacker by stealing the deployment model can be the same as or different from the model structure of the deployment model. That is, the model structure of the suspicious model can be the same as or different from the model structure of the deployment model.

[0045] In some optional implementations, before training of the meta-classifier based on the target model, the auxiliary model, and the transform sample set, the method for performing model ownership verification based on an exogenous feature can further include a process of determining the target model and the auxiliary model. For example, multiple scenarios can be classified according to whether the model structure of the suspicious model is known and whether the model structure of the suspicious model is the same as that of the deployment model. As shown in FIG. 3, FIG. 3 is a schematic flowchart of determining a target model and an auxiliary model according to a suspicious model. The method can include the following steps:

[0046] Step 301: Determine whether a model structure of a suspicious model is known.

[0047] Step 302: In response to determining that the model structure of the suspicious model is known, further determine whether the model structure of the suspicious model is the same as a model structure of a deployment model.

[0048] Step 303: Determine the deployment model as a target model and train an auxiliary model based on the model structure of the suspicious model, in response to

determining that the model structure of the suspicious model is known and the same as the model structure of the deployment model.

[0049] In this implementation, when the model structure of the suspicious model is the same as the model structure of the deployment model, the deployment model can be used as the target model. Therefore, training time of the target model can be reduced. In addition, the auxiliary model that has a same model structure as the target model (the deployment model) and the suspicious model can be trained according to the initial sample in the initial sample set. Because the initial sample in the initial sample set is not embedded with the exogenous feature, the initial sample set can also be referred to as a benign sample set, and the auxiliary model is trained according to the initial sample not embedded with the exogenous feature. Therefore, the auxiliary model can also be referred to as a benign model or a normal model. The auxiliary model does not have the feature knowledge of the exogenous feature.

[0050] Step 304: Train the target model and the auxiliary model based on the model structure of the suspicious model, in response to determining that the model structure of the suspicious model is known and different from the model structure of the deployment model.

[0051] In this implementation, when the model structure of the suspicious model is different from that of the deployment model, the target model can be trained according to the transform sample set and the remaining sample set in the initial sample set except the selected sample set, and the model structure of the suspicious model. In a training process of the target model, the target model can determine the feature knowledge of the exogenous feature, and has a model structure the same as that of the suspicious model. In addition, the auxiliary model whose structure is the same as that of the suspicious model can be trained according to the initial sample set.

[0052] It can be determined from step 303 and step 304 that, if the model structure of the suspicious model is known, the model structures of the target model and the auxiliary model are the same as the model structure of the suspicious model.

[0053] Step 305: Determine the deployment model as the target model in response to determining that the model structure of the suspicious model is unknown, and train

the auxiliary model based on the model structure of the deployment model.

[0054] In this implementation, when the model structure of the suspicious model is unknown, the deployment model can be determined as the target model, and the auxiliary model can be trained according to the initial sample set and the model structure of the deployment model. That is, when the model structure of the suspicious model is unknown, the model structures of the target model and the auxiliary model are the same as the model structure of the deployment model.

[0055] In some optional implementations, when the model structure of the suspicious model is known, step 203 of training a meta-classifier based on a target model, an auxiliary model, and the transform sample set can be specifically performed as follows:

[0056] First, a first meta-classifier sample set including a positive sample and a negative sample is constructed.

[0057] In this implementation, to train a first meta-classifier, the first meta-classifier sample set including the positive sample and the negative sample needs to be constructed first. Here, sample data of the positive sample can be gradient information of the target model for the transform sample. Sample data of the negative sample can be gradient information of the auxiliary model for the transform sample. For example, a gradient vector can be used as the gradient information.

[0058] Optionally, the gradient information can alternatively be a result vector obtained after each element in a gradient vector is calculated by using a sign function. The result vector obtained after the gradient vector is calculated by using the sign function is simpler and can still reflect a direction characteristic of a gradient. Therefore, the result vector can be used as the gradient information.

[0059] Then, the first meta-classifier sample set is used for training to obtain a binary classifier as the first meta-classifier.

[0060] In this implementation, the first meta-classifier sample set can be used for training to obtain the first meta-classifier. Using an example in which a label of the positive sample in the first meta-classifier sample set is +1, a label of the negative sample is -1, and the gradient information is the result vector obtained after each

element in the gradient vector is calculated by using the sign function, the first meta-classifier sample set $\mathcal{D}_c$ can be represented as $\mathcal{D}_c = \mathcal{D}_{positive} \cup \mathcal{D}_{negative}$, where the positive sample is $\mathcal{D}_{positive} = \{(g_V(\mathbf{x}'), +1) | (\mathbf{x}', y) \in \mathcal{D}_t\}$, and the label in the positive sample is +1; and $\mathcal{D}_t$ can represent the transform sample set, and $\mathbf{x}'$ represents the transform sample. Here, $g_V(\mathbf{x}') = \text{sign}(\nabla_{\theta} \mathcal{L}(V(\mathbf{x}'), y))$, where V can represent the target model, $g_V(\mathbf{x}')$ represents the gradient information of the target model for the transform sample, $\nabla_{\theta} \mathcal{L}(V(\mathbf{x}'), y)$ represents a loss function gradient vector of the target model for the transform sample, $\text{sign}(\cdot)$ represent a sign function. The negative sample is $\mathcal{D}_{negative} = \{(g_B(\mathbf{x}'), -1) | (\mathbf{x}', y) \in \mathcal{D}_t\}$, where the label in the negative sample is -1, where $g_B(\mathbf{x}') = \text{sign}(\nabla_{\theta} \mathcal{L}(B(\mathbf{x}'), y))$, B represents the auxiliary model, $g_B(\mathbf{x}')$ represents the gradient information of the auxiliary model for the transform sample, and $\nabla_{\theta} \mathcal{L}(B(\mathbf{x}'), y)$ represents a loss function gradient vector of the auxiliary model for the transform sample. In this example, the first meta-classifier C can be trained by using the following equation

15 $\min_{\mathbf{w}} \sum_{(\mathbf{x}, y) \in \mathcal{D}_c} \mathcal{L}(C_{\mathbf{w}}(\mathbf{x}), y)$, where $\mathbf{w}$ can represent a model parameter in the classifier.

[0061] In some optional implementations, when the model structure of the suspicious model is known, step 204 of inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model can specifically include the following steps 1) to 4):

[0062] Step 1): Select the transform sample from the transform sample set as a first transform sample.

[0063] Step 2): Determine first gradient information of the suspicious model for the first transform sample.

25 [0064] Step 3): Input the first gradient information into the first meta-classifier to obtain a first prediction result.

[0065] Step 4): Determine, in response to the first prediction result indicating a positive sample, that the suspicious model is a model stolen from the deployment model.

[0066] For example, the label of the positive sample in the first meta-classifier sample set is +1, the label of the negative sample is -1, and the gradient information is the result vector obtained after each element in the gradient vector is calculated by using the sign function. Assume that the suspicious model is S , the first meta-classifier is C , the first transform sample is a transform image $\mathbf{x}'$ whose label is y , and the first gradient information of the suspicious model for the first transform sample can be determined by using $g_S(\mathbf{x}') = \text{sign}(\nabla_{\theta} \mathcal{L}(S(\mathbf{x}'), y))$. Then, the first gradient information is input to the first meta-classifier C , that is, $C(g_S(\mathbf{x}'))$, to obtain a first prediction result. If the first prediction result indicates a positive sample, that is, $C(g_S(\mathbf{x}')) = 1$, the suspicious model can be determined as a model stolen from the deployment model. In this example, $C(g_S(\mathbf{x}')) = 1$ can indicate that the suspicious model and the deployment model similarly have the feature knowledge of the exogenous feature, and therefore, the suspicious model can be determined as a model stolen from the deployment model. In this implementation, ownership verification on the suspicious model can be implemented.

[0067] In another optional implementation, when the model structure of the suspicious model is known, step 204 of inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model can further specifically include: using hypothesis testing to validate ownership of the suspicious model based on a first subset selected from the transform sample set, the first meta-classifier, and the auxiliary model.

[0068] In this implementation, first a plurality of transform samples can be selected (for example, randomly sampled) from the transform sample set $\mathcal{D}_t$ to form a first subset, and then ownership verification on the suspicious model is performed according to the first subset, the first meta-classifier, and the auxiliary model by using a plurality of types of hypothesis testing. For example, a Z-test can be used to perform ownership verification on the suspicious model.

[0069] Optionally, the performing ownership verification on the suspicious model

by using hypothesis testing can include: performing ownership verification on the suspicious model by using a one-sided paired sample T test, which can specifically include the following content:

[0070] First, a first null hypothesis in which a first probability is less than or equal
5 to a second probability is constructed.

[0071] In this implementation, for the first subset, the first probability μ_S can indicate a posterior probability that a prediction result of the first meta-classifier for gradient information of the suspicious model is a positive sample, and the second probability μ_B can indicate a posterior probability that a prediction result of the first
10 meta-classifier for gradient information of the auxiliary model is a positive sample. For example, $\mathbf{X}'$ represents sample data of a transform sample in the first subset, and a label of a positive sample is +1. The first probability μ_S and the second probability μ_B respectively represent posterior probabilities of events $C(g_S(\mathbf{X}')) = 1$ and $C(g_B(\mathbf{X}')) = 1$. A null hypothesis $H_0 : \mu_S \leq \mu_B$ can be constructed for this, where
15 S represents the suspicious model, and B represents the auxiliary model.

[0072] Then, a P value is calculated based on the first null hypothesis and sample data in the first subset. It can be understood that, in the one-sided paired sample T test, P value calculation is well known to a person skilled in the art, and details are omitted here for simplicity.

20 [0073] Then, in response to determining that the P value is less than a significance level α , it is determined that the first null hypothesis is rejected. Here, the significance level α can be a value determined by a person skilled in the art according to an actual requirement.

[0074] Finally, in response to determining that the first null hypothesis is rejected,
25 the suspicious model is determined as a model stolen from the deployment model. In practice, because the auxiliary model does not have the feature knowledge of the exogenous feature, μ_B should be a smaller value. If μ_S being less than or equal to μ_B is valid, it can indicate that the suspicious model does not have the feature knowledge of the exogenous feature, that is, the suspicious model is not a model stolen

from the deployment model. On the contrary, if μ_S being less than or equal to μ_B is not valid (that is, rejected), it can indicate that the suspicious model has the feature knowledge of the exogenous feature, that is, the suspicious model is a model stolen from the deployment model. In this implementation, ownership verification is performed on the suspicious model by using statistical hypothesis testing, so impact of randomness of transform sample selection in a process of ownership verification on accuracy of ownership verification can be avoided, and verification is more accurate.

[0075] As shown in FIG. 3, in some optional implementations, the model structure of the suspicious model is unknown, so it is difficult to obtain gradient information of the model to construct a training sample of the meta-classifier. In this case, step 203 of training a meta-classifier based on a target model, an auxiliary model, and the transform sample set can be specifically performed as follows:

[0076] First, a second meta-classifier sample set including a positive sample and a negative sample is constructed.

[0077] In this implementation, to train a second meta-classifier, the second meta-classifier sample set including the positive sample and the negative sample needs to be constructed first. Here, sample data of the positive sample are difference information between a prediction output of the target model for a selected sample and a prediction output for a transform sample corresponding to the selected sample. Sample data of the negative sample are difference information between a prediction output of the auxiliary model for a selected sample and a prediction output for a transform sample corresponding to the selected sample. In practice, if the target model and the auxiliary model are classification models, the prediction outputs of the target model and the auxiliary model can be probability vectors respectively formed for a plurality of prediction probabilities of a plurality of category labels. As an example, the difference information can refer to a difference vector. As another example, the difference information can alternatively be a result obtained after the difference vector is calculated by using the sign function. For example, the sample data of the positive sample are $\text{sign}(V(\mathbf{x}) - V(\mathbf{x}'))$, where $V(\mathbf{x})$ represents the prediction output

(reflected as a probability vector) of the target model for the selected sample, and $V(\mathbf{x}')$ represents the prediction output of the target model for the transform sample corresponding to the selected sample. The sample data of the negative sample are $\text{sign}(B(\mathbf{x}) - B(\mathbf{x}'))$, where $B(\mathbf{x})$ represents the prediction output of the auxiliary
5 model for the selected sample, and $B(\mathbf{x}')$ represents the prediction output of the auxiliary model for the transform sample corresponding to the selected sample.

[0078] The second meta-classifier sample set is then used to train a second meta-classifier.

[0079] In this implementation, the second meta-classifier sample set can be used to
10 train the second meta-classifier. In this implementation, the meta-classifier can be trained when the model structure of the suspicious model is unknown, facilitating subsequent model ownership verification.

[0080] In some optional implementations, when the model structure of the suspicious model is unknown, step 204 of inputting related data of a suspicious model
15 into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model can specifically include the following steps 1 to 4:

[0081] Step 1: Respectively obtain a corresponding second transform sample and a corresponding second selected sample from the transform sample set and the selected
20 sample set. Here, that a second transform sample is corresponding to a selected sample can mean that the second transform sample is obtained by embedding the exogenous feature into the selected sample.

[0082] Step 2: Determine second difference information between a prediction output of the suspicious model for the second selected sample and a prediction output
25 for the second transform sample.

[0083] Step 3: Input the second difference information into the second meta-classifier to obtain a second prediction result.

[0084] Step 4: Determine whether the second prediction result indicates a positive sample, and determine, in response to the second prediction result indicating a positive

sample, that the suspicious model is a model stolen from the deployment model. In this implementation, ownership verification on the suspicious model can be implemented when the model structure of the suspicious model is unknown.

5 **[0085]** In another optional implementation, when the model structure of the suspicious model is unknown, step 204 of inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model can further specifically include: performing ownership verification on the suspicious model by using hypothesis testing based on a second subset selected from the transform sample set, a third subset corresponding to the second subset and in the selected sample set, the
10 second meta-classifier, and the auxiliary model. For example, a Z-test can be used to perform ownership verification on the suspicious model.

15 **[0086]** Optionally, the performing ownership verification on the suspicious model by using hypothesis testing can include: performing ownership verification on the suspicious model by using a one-sided paired sample T test, which can specifically include the following content:

15 **[0087]** First, a second null hypothesis in which a third probability is less than or equal to a fourth probability is constructed.

20 **[0088]** In this implementation, for the second subset and the third subset, the third probability can indicate a posterior probability that a prediction result of the second meta-classifier for the difference information corresponding to the suspicious model is a positive sample. The fourth probability can indicate a posterior probability that a prediction result of the second meta-classifier for the difference information corresponding to the auxiliary model is a positive sample.

25 **[0089]** Then, a P value is calculated based on the second null hypothesis, sample data of the second subset, and sample data of the third subset. It can be understood that, in the one-sided paired sample T test, P value calculation is well known to a person skilled in the art, and details are omitted here for simplicity.

30 **[0090]** Then, in response to determining that the P value is less than a significance level α , it is determined that the second null hypothesis is rejected. Here, the

significance level α can be a value determined by a person skilled in the art according to an actual requirement.

[0091] Finally, in response to determining that the second null hypothesis is rejected, the suspicious model is determined as a model stolen from the deployment model. In practice, because the auxiliary model does not have the feature knowledge of the exogenous feature, the fourth probability should be a smaller value. If the third probability being less than or equal to the fourth probability is valid, it can indicate that the suspicious model does not have the feature knowledge of the exogenous feature, that is, the suspicious model is not a model stolen from the deployment model. On the contrary, if the third probability being less than or equal to the fourth probability is not valid (that is, rejected), it can indicate that the suspicious model has the feature knowledge of the exogenous feature, that is, the suspicious model is a model stolen from the deployment model. In this implementation, ownership verification is performed on the suspicious model by using statistical hypothesis testing, so impact of randomness of transform sample selection in a process of ownership verification on accuracy of ownership verification can be avoided, and verification is more accurate.

[0092] According to an embodiment of another aspect, an apparatus for performing model ownership verification based on an exogenous feature is provided. The apparatus for performing model ownership verification based on an exogenous feature can be deployed in any device, platform, or device cluster that has a computing and processing capability.

[0093] FIG. 4 is a schematic block diagram illustrating an apparatus for performing model ownership verification based on a source feature, according to an embodiment. As shown in FIG. 4, the apparatus 400 for performing model ownership verification based on a source feature includes: a selection unit 401, configured to select some initial samples from an initial sample set to form a selected sample set; a transform unit 402, configured to process sample data of each selected sample in the selected sample set to obtain a transform sample set formed by a transform sample with an exogenous feature, where the exogenous feature is a feature that sample data of the initial sample do not have; a training unit 403, configured to train a meta-classifier based on a target model,

an auxiliary model, and the transform sample set, where the auxiliary model is a model trained by using the initial sample set, the target model is a model trained by using the transform sample set and a remaining sample set in the initial sample set except the selected sample set, and the meta-classifier is used to identify feature knowledge of the exogenous feature; and a verification unit 404, configured to input related data of a suspicious model into the meta-classifier and determine, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model, where the deployment model has feature knowledge of the exogenous feature.

[0094] In some optional implementations of this embodiment, the apparatus 400 further includes: a first model training unit (not shown in the figure), configured to determine the deployment model as the target model and train the auxiliary model based on a model structure of the suspicious model, in response to the model structure of the suspicious model being known and the same as a model structure of the deployment model; and a second model training unit (not shown in the figure), configured to train the target model and the auxiliary model based on the model structure of the suspicious model, in response to the model structure of the suspicious model being known and different from the model structure of the deployment model.

[0095] In some optional implementations of this embodiment, the training unit 403 is further configured to: construct a first meta-classifier sample set including a positive sample and a negative sample, where sample data of the positive sample are gradient information of the target model for the transform sample; and sample data of the negative sample are gradient information of the auxiliary model for the transform sample; and train to obtain a first meta-classifier by using the first meta-classifier sample set.

[0096] In some optional implementations of this embodiment, the gradient information is a result vector obtained after each element in a gradient vector is calculated by using a sign function.

[0097] In some optional implementations of this embodiment, the verification unit 404 is further configured to: select a first transform sample from the transform sample set; determine first gradient information of the suspicious model for the first transform

sample; input the first gradient information into the first meta-classifier to obtain a first prediction result; and determine, in response to the first prediction result indicating a positive sample, that the suspicious model is a model stolen from the deployment model.

[0098] In some optional implementations of this embodiment, the verification unit 404 is further configured to: use hypothesis testing to validate ownership of the suspicious model based on a first subset selected from the transform sample set, the first meta-classifier, and the auxiliary model.

[0099] In some optional implementations of this embodiment, the using hypothesis testing to validate ownership of the suspicious model includes: constructing a first null hypothesis in which a first probability is less than or equal to a second probability, where the first probability indicates a posterior probability that a prediction result of the first meta-classifier for gradient information of the suspicious model is a positive sample, and the second probability indicates a posterior probability that a prediction result of the first meta-classifier for gradient information of the auxiliary model is a positive sample; calculating a P value based on the first null hypothesis and sample data in the first subset; determining, in response to determining that the P value is less than a significance level α , that the first null hypothesis is rejected; and determining, in response to determining that the first null hypothesis is rejected, that the suspicious model is a model stolen from the deployment model.

[0100] In some optional implementations of this embodiment, the apparatus 400 further includes: a third model training unit (not shown in the figure), configured to: determine the deployment model as the target model in response to a model structure of the suspicious model being unknown, and train the auxiliary model based on a model structure of the deployment model.

[0101] In some optional implementations of this embodiment, the training unit 403 is further configured to: construct a second meta-classifier sample set including a positive sample and a negative sample, where sample data of the positive sample are difference information between a prediction output of the target model for a selected sample and a prediction output for a transform sample corresponding to the selected sample; and sample data of the negative sample are difference information between a

prediction output of the auxiliary model for a selected sample and a prediction output for a transform sample corresponding to the selected sample; and train a second meta-classifier by using the second meta-classifier sample set.

[0102] In some optional implementations of this embodiment, the verification unit
5 404 is further configured to: respectively obtain a corresponding second transform sample and a corresponding second selected sample from the transform sample set and the selected sample set; determine second difference information between a prediction output of the suspicious model for the second selected sample and a prediction output for the second transform sample; input the second difference information into the
10 second meta-classifier to obtain a second prediction result; and determine, in response to the second prediction result indicating a positive sample, that the suspicious model is a model stolen from the deployment model.

[0103] In some optional implementations of this embodiment, the verification unit
404 is further configured to: perform ownership verification on the suspicious model
15 by using hypothesis testing based on a second subset selected from the transform sample set, a third subset corresponding to the second subset and in the selected sample set, the second meta-classifier, and the auxiliary model.

[0104] In some optional implementations of this embodiment, the using hypothesis testing to validate ownership of the suspicious model includes: constructing a second
20 null hypothesis in which a third probability is less than or equal to a fourth probability, where the third probability indicates a posterior probability that a prediction result of the second meta-classifier for difference information corresponding to the suspicious model is a positive sample, and the fourth probability indicates a posterior probability that a prediction result of the second meta-classifier for difference information
25 corresponding to the auxiliary model is a positive sample; calculating a P value based on the second null hypothesis, sample data of the second subset, and sample data of the third subset; determining, in response to determining that the P value is less than a significance level α , that the second null hypothesis is rejected; and determining, in response to determining that the second null hypothesis is rejected, that the suspicious
30 model is a model stolen from the deployment model.

[0105] In some optional implementations of this embodiment, the sample data of the initial sample in the initial sample set are a sample image; and the transform unit 402 is further configured to: perform style conversion on a sample image of each sample in the selected sample set by using an image style converter, so the sample image has a specified image style, where the exogenous feature is a feature related to the specified image style.

[0106] According to some embodiments in another aspect, a computer-readable storage medium is further provided, where the computer-readable storage medium stores a computer program, and when the computer program is executed in a computer, the computer is enabled to perform the method described in FIG. 2.

[0107] In one or more embodiments of still another aspect, a computing device is further provided, including a memory and a processor. The memory stores executable code, and when executing the executable code, the processor implements the method the method described in FIG. 2.

[0108] A person of ordinary skill in the art can be further aware that, in combination with the examples described in the implementations disclosed in this specification, units and algorithm steps can be implemented by electronic hardware, computer software, or a combination thereof. To clearly describe interchangeability between the hardware and the software, compositions and steps of each example are generally described above based on functions. Whether the functions are performed by hardware or software depends on particular applications and design constraint conditions of the technical solutions. A person of ordinary skill in the art can use different methods to implement the described functions for each particular application, but it should not be considered that the implementation goes beyond the scope of this application.

[0109] Steps of methods or algorithms described in the implementations disclosed in this specification can be implemented by hardware, a software module executed by a processor, or a combination thereof. The software module can reside in a random access memory (RAM), a memory, a read-only memory (ROM), an electrically programmable ROM, an electrically erasable programmable ROM, a register, a hard disk, a removable disk, a CD-ROM, or any other form of storage medium known in the

art.

[0110] In the described specific implementations, the objective, technical solutions, and benefits of the present disclosure are further described in detail. It should be understood that the descriptions are merely specific implementations of the present disclosure, but are not intended to limit the protection scope of the present disclosure. 5 Any modification, equivalent replacement, or improvement made without departing from the spirit and principle of the present disclosure should fall within the protection scope of the present disclosure.

CLAIMS

What is claimed is:

1. A method for performing model ownership verification based on an exogenous feature, comprising:

5 selecting some initial samples from an initial sample set to form a selected sample set;

 processing sample data of each selected sample in the selected sample set to obtain a transform sample set formed by a transform sample with an exogenous feature, wherein the exogenous feature is a feature that sample data of the initial sample do not
10 have;

 training a meta-classifier based on a target model, an auxiliary model, and the transform sample set, wherein the auxiliary model is a model trained by using the initial sample set, the target model is a model trained by using the transform sample set and a remaining sample set in the initial sample set except the selected sample set, and the
15 meta-classifier is used to identify feature knowledge of the exogenous feature; and

 inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model, wherein the deployment model has feature knowledge of the exogenous feature.

20 2. The method according to claim 1, wherein before the training a meta-classifier based on a target model, an auxiliary model, and the transform sample set, the method further comprises:

 determining the deployment model as the target model and training the auxiliary model based on a model structure of the suspicious model, in response to the model
25 structure of the suspicious model being known and the same as a model structure of the deployment model; and

 training the target model and the auxiliary model based on the model structure of the suspicious model, in response to the model structure of the suspicious model being known and different from the model structure of the deployment model.

3. The method according to claim 2, wherein the training a meta-classifier based on a target model, an auxiliary model, and the transform sample set comprises:

constructing a first meta-classifier sample set comprising a positive sample and a negative sample, wherein sample data of the positive sample are gradient information of the target model for the transform sample; and sample data of the negative sample are gradient information of the auxiliary model for the transform sample; and

training to obtain a first meta-classifier by using the first meta-classifier sample set.

4. The method according to claim 3, wherein the gradient information is a result vector obtained after each element in a gradient vector is calculated by using a sign function.

5. The method according to claim 3, wherein the inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model comprises:

selecting a first transform sample from the transform sample set;

determining first gradient information of the suspicious model for the first transform sample;

inputting the first gradient information into the first meta-classifier to obtain a first prediction result; and

determining, in response to the first prediction result indicating a positive sample, that the suspicious model is a model stolen from the deployment model.

6. The method according to claim 3, wherein the inputting related data of a suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model comprises:

using hypothesis testing to validate ownership of the suspicious model based on a first subset selected from the transform sample set, the first meta-classifier, and the auxiliary model.

7. The method according to claim 6, wherein the using hypothesis testing to

validate ownership of the suspicious model comprises:

constructing a first null hypothesis in which a first probability is less than or equal to a second probability, wherein the first probability indicates a posterior probability that a prediction result of the first meta-classifier for gradient information of the suspicious model is a positive sample, and the second probability indicates a posterior probability that a prediction result of the first meta-classifier for gradient information of the auxiliary model is a positive sample;

calculating a P value based on the first null hypothesis and sample data in the first subset;

determining, in response to determining that the P value is less than a significance level α , that the first null hypothesis is rejected; and

determining, in response to determining that the first null hypothesis is rejected, that the suspicious model is a model stolen from the deployment model.

8. The method according to claim 1, wherein before the training a meta-classifier based on a target model, an auxiliary model, and the transform sample set, the method further comprises:

determining the deployment model as the target model in response to a model structure of the suspicious model being unknown, and training the auxiliary model based on a model structure of the deployment model.

9. The method according to claim 8, wherein the training a meta-classifier based on a target model, an auxiliary model, and the transform sample set comprises:

constructing a second meta-classifier sample set comprising a positive sample and a negative sample, wherein sample data of the positive sample are difference information between a prediction output of the target model for a selected sample and a prediction output for a transform sample corresponding to the selected sample; and sample data of the negative sample are difference information between a prediction output of the auxiliary model for a selected sample and a prediction output for a transform sample corresponding to the selected sample; and

training a second meta-classifier by using the second meta-classifier sample set.

10. The method according to claim 9, wherein the inputting related data of a

suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model comprises:

respectively obtaining a corresponding second transform sample and a
5 corresponding second selected sample from the transform sample set and the selected sample set;

determining second difference information between a prediction output of the suspicious model for the second selected sample and a prediction output for the second transform sample;

10 inputting the second difference information into the second meta-classifier to obtain a second prediction result; and

determining, in response to the second prediction result indicating a positive sample, that the suspicious model is a model stolen from the deployment model.

11. The method according to claim 9, wherein the inputting related data of a
15 suspicious model into the meta-classifier and determining, based on an output result of the meta-classifier, whether the suspicious model is a model stolen from a deployment model comprises:

performing ownership verification on the suspicious model by using hypothesis testing based on a second subset selected from the transform sample set, a third subset
20 corresponding to the second subset and in the selected sample set, the second meta-classifier, and the auxiliary model.

12. The method according to claim 11, wherein the using hypothesis testing to validate ownership of the suspicious model comprises:

constructing a second null hypothesis in which a third probability is less than or
25 equal to a fourth probability, wherein the third probability indicates a posterior probability that a prediction result of the second meta-classifier for difference information corresponding to the suspicious model is a positive sample, and the fourth probability indicates a posterior probability that a prediction result of the second meta-classifier for difference information corresponding to the auxiliary model is a positive
30 sample;

calculating a P value based on the second null hypothesis, sample data of the second subset, and sample data of the third subset;

determining, in response to determining that the P value is less than a significance level α , that the second null hypothesis is rejected; and

5 determining, in response to determining that the second null hypothesis is rejected, that the suspicious model is a model stolen from the deployment model.

13. The method according to claim 1, wherein the sample data of the initial sample in the initial sample set are a sample image; and

the processing sample data of each sample in the selected sample set to obtain a
10 transform sample set formed by a transform sample with an exogenous feature comprises:

performing style conversion on a sample image of each sample in the selected sample set by using an image style converter, so the sample image has a specified image style, wherein the exogenous feature is a feature related to the specified image style.

15 14. An apparatus for performing model ownership verification based on an exogenous feature, comprising:

a selection unit, configured to select some initial samples from an initial sample set to form a selected sample set;

a transform unit, configured to process sample data of each selected sample in the
20 selected sample set to obtain a transform sample set formed by a transform sample with an exogenous feature, wherein the exogenous feature is a feature that sample data of the initial sample do not have;

a training unit, configured to train a meta-classifier based on a target model, an auxiliary model, and the transform sample set, wherein the auxiliary model is a model
25 trained by using the initial sample set, the target model is a model trained by using the transform sample set and a remaining sample set in the initial sample set except the selected sample set, and the meta-classifier is used to identify feature knowledge of the exogenous feature; and

a verification unit, configured to input related data of a suspicious model into the
30 meta-classifier and determine, based on an output result of the meta-classifier, whether

the suspicious model is a model stolen from a deployment model, wherein the deployment model has feature knowledge of the exogenous feature.

15. A computer-readable storage medium, storing a computer program, wherein the computer program, when run in a computer, causes the computer to perform the
5 method according to any one of claims 1 to 13.

16. A computing device, comprising a memory and a processor, wherein the memory stores executable code, and when executing the executable code, the processor implements the method according to any one of claims 1 to 13.