



US 20230222386A1

(19) **United States**

(12) **Patent Application Publication**
LI et al.

(10) **Pub. No.: US 2023/0222386 A1**

(43) **Pub. Date: Jul. 13, 2023**

(54) **SYSTEM AND METHOD FOR
LARGE-SCALE MULTI-LABEL LEARNING
USING INCOMPLETE LABEL
ASSIGNMENTS**

continuation of application No. 14/543,133, filed on
Nov. 17, 2014, now Pat. No. 10,325,220.

Publication Classification

(71) Applicant: **YAHOO ASSETS LLC**, New York,
NY (US)

(51) **Int. Cl.**
G06N 20/00 (2019.01)

(72) Inventors: **Jia LI**, Santa Clara, CA (US); **Yi
CHANG**, Sunnyvale, CA (US);
Xiangnan KONG, Worcester, MA (US)

(52) **U.S. Cl.**
CPC **G06N 20/00** (2019.01)

(21) Appl. No.: **18/175,249**

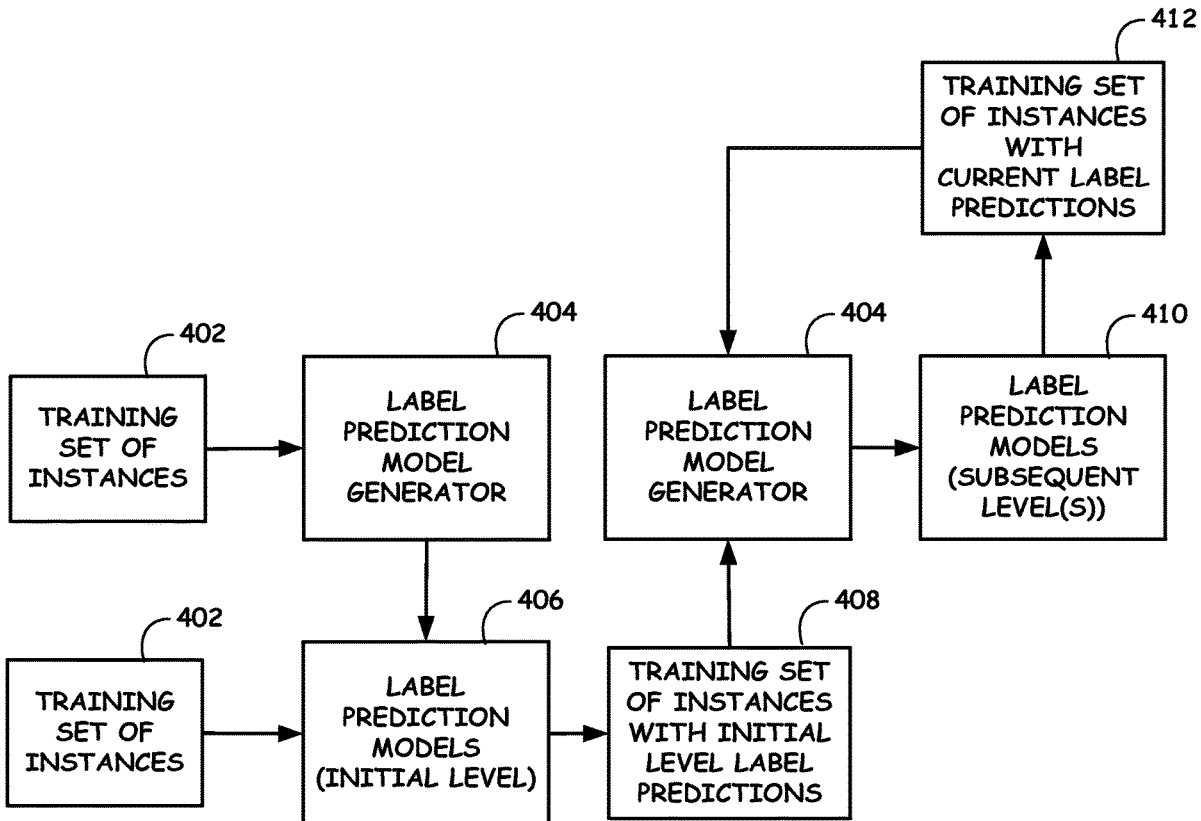
(57) **ABSTRACT**

(22) Filed: **Feb. 27, 2023**

At least one label prediction model is trained, or learned,
using training data that may comprise training instances that
may be missing one or more labels. The at least one label
prediction model may be used in identifying a content item's
ground-truth label set comprising an indicator for each label
in the label set indicating whether or not the label is
applicable to the content item.

Related U.S. Application Data

(63) Continuation of application No. 16/439,855, filed on
Jun. 13, 2019, now Pat. No. 11,593,703, which is a



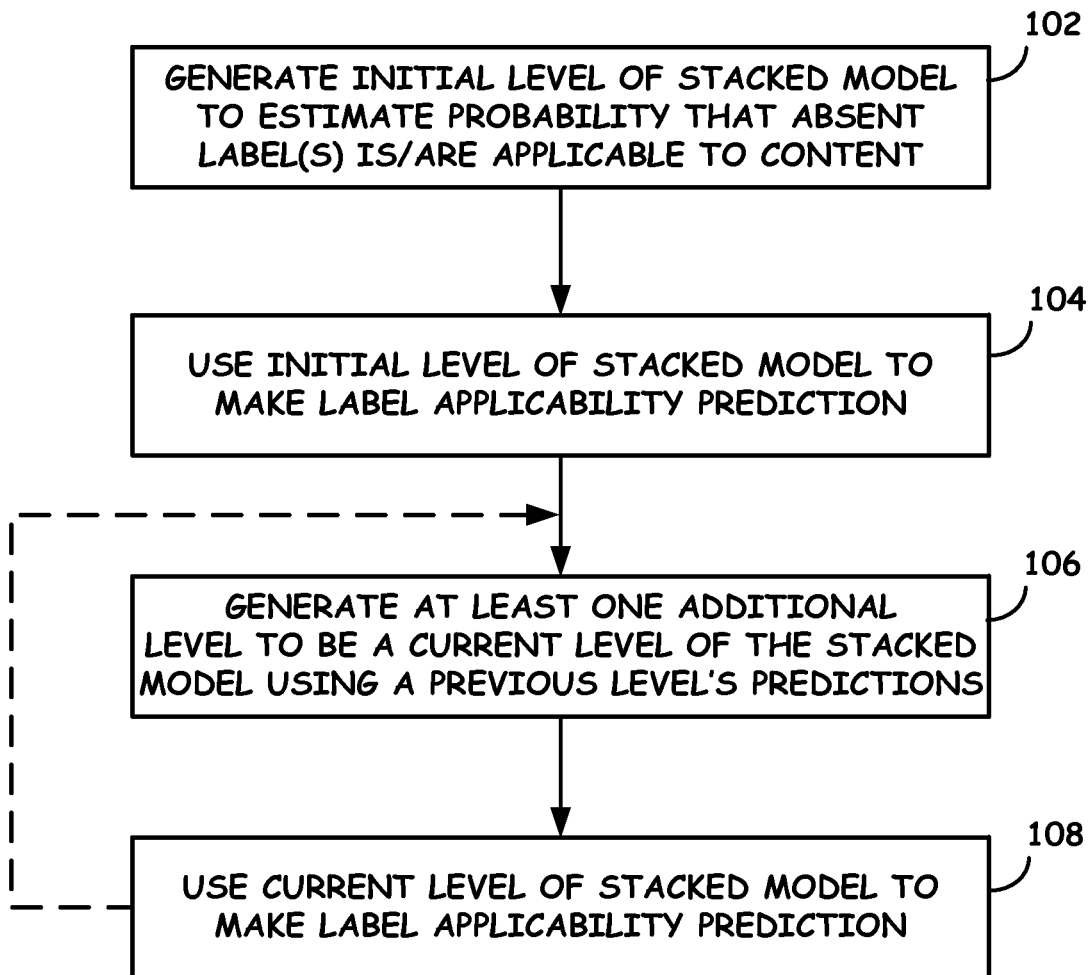


Fig. 1

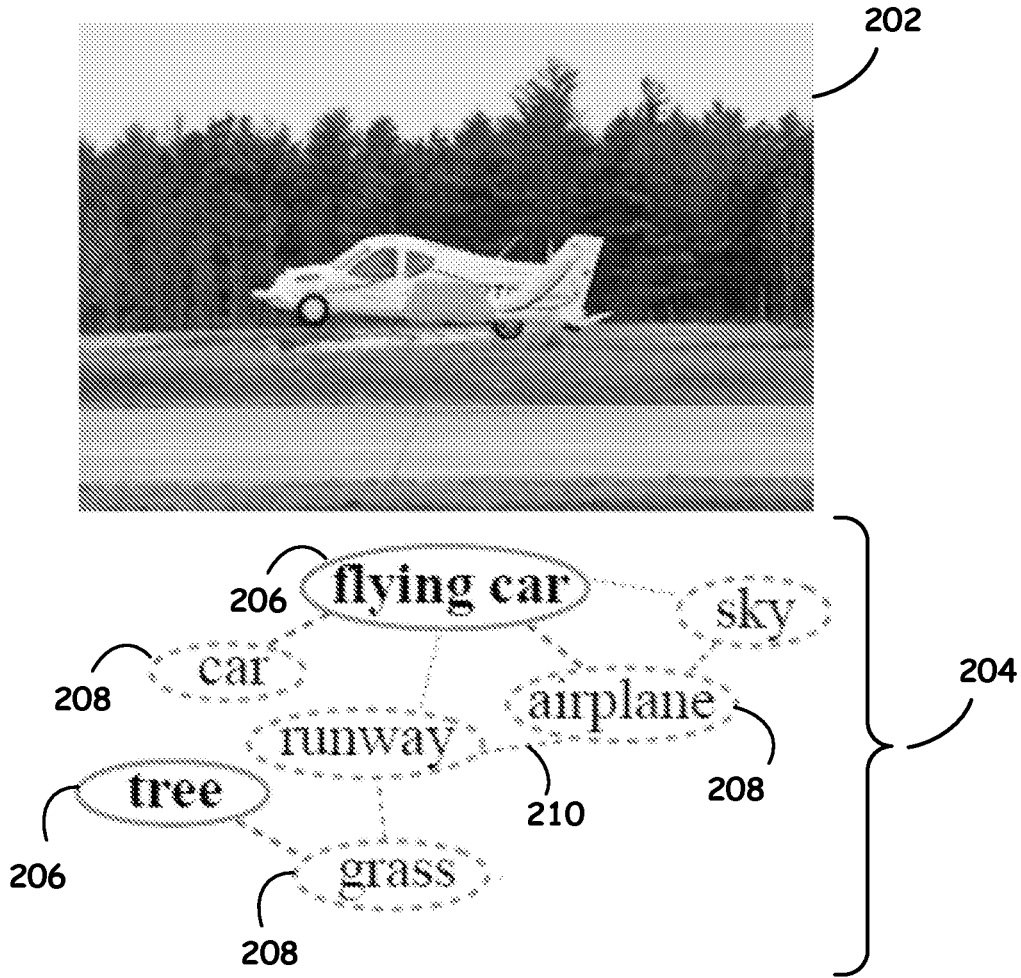


Fig. 2




Fig. 3

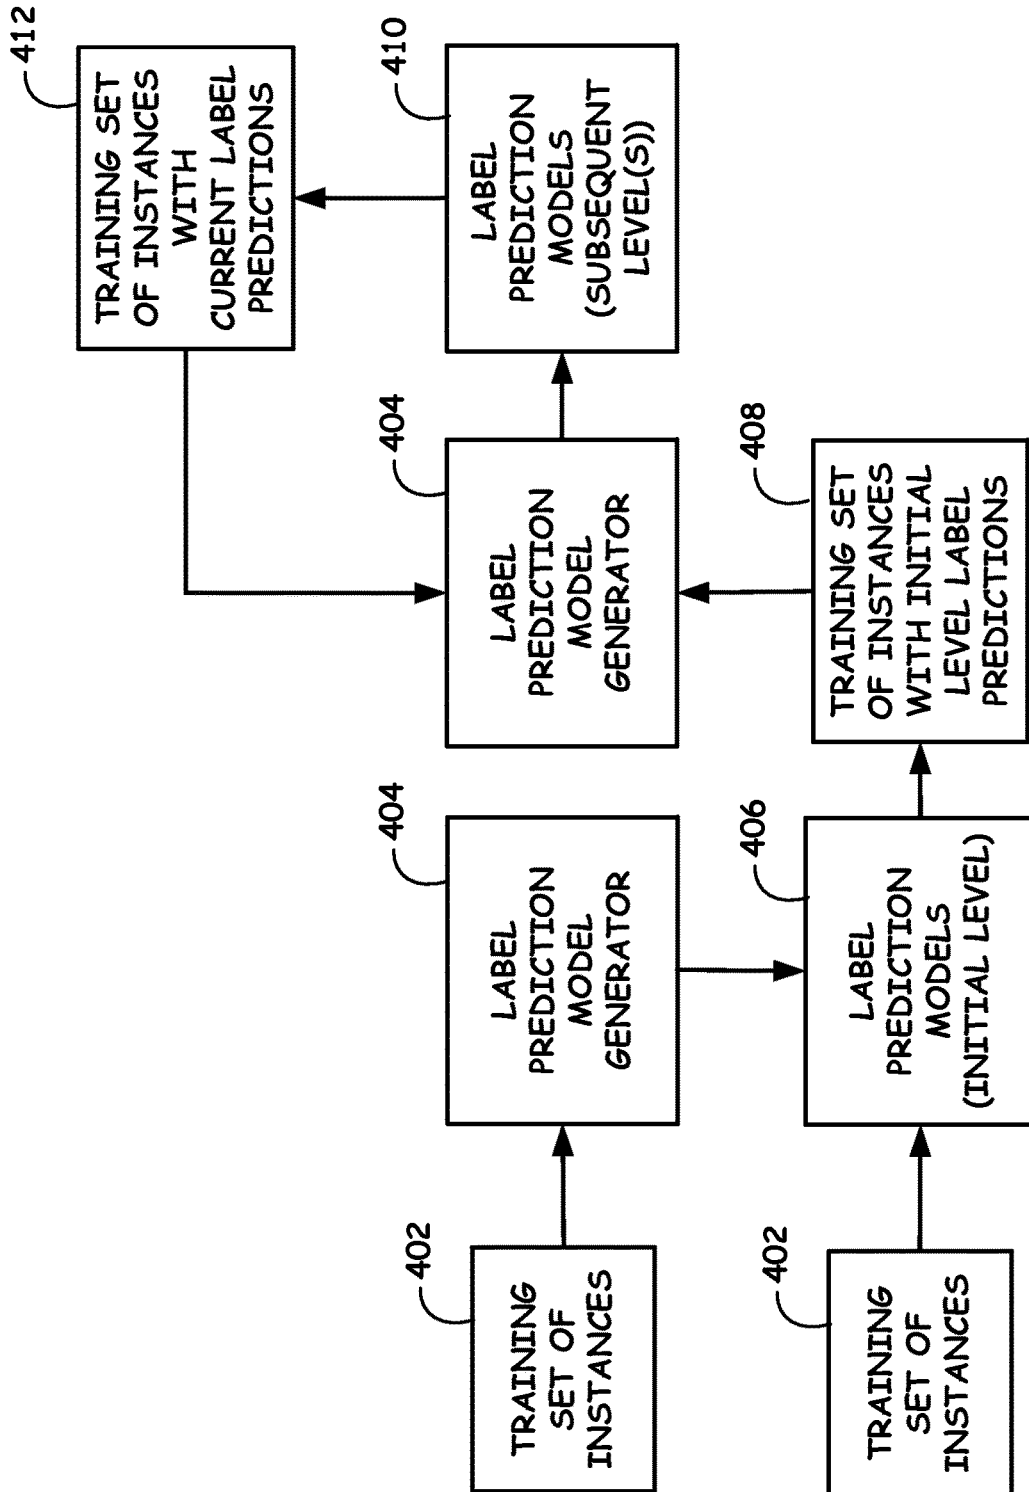


Fig. 4

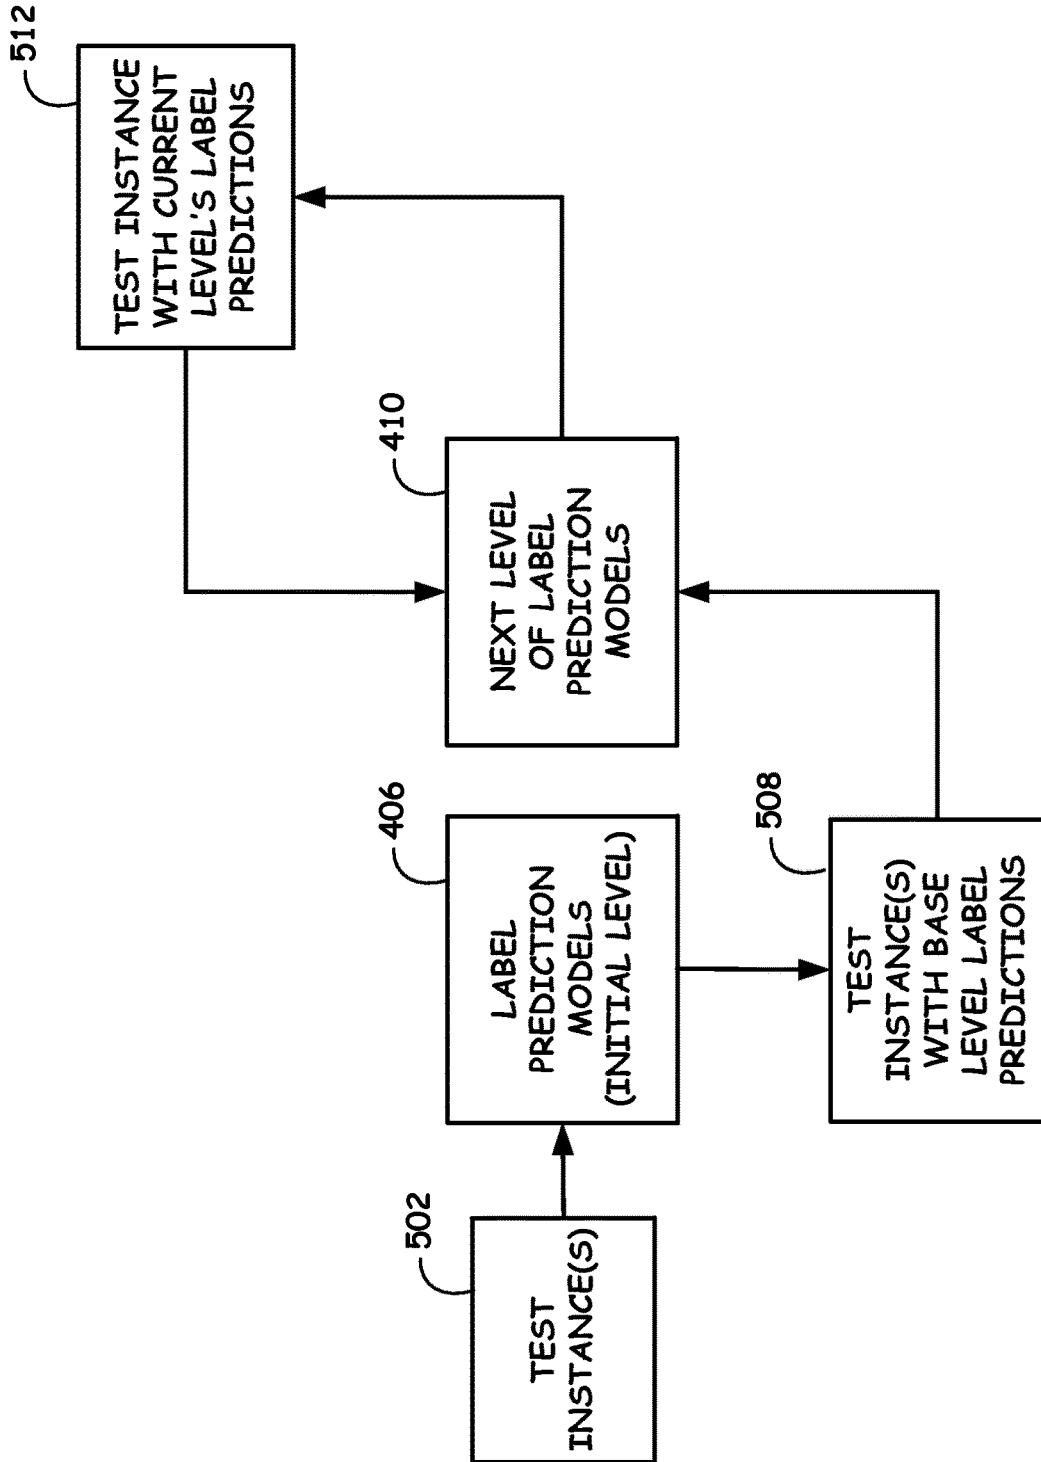


Fig. 5

Symbol	Definition
$\mathbf{x}_1, \dots, \mathbf{x}_n$	the feature vectors for training instances
$\mathbf{y}_1, \dots, \mathbf{y}_n$	the set of variables for label sets of the training instances
$\mathbf{s}_1, \dots, \mathbf{s}_n$	the set of annotated labels for the training instances
$\mathbf{y}_i = (y_i^1, \dots, y_i^L)^\top$	the vector of variables for the label set of instance \mathbf{x}_i , and $y_i^k \in \{0, 1\}$
$\mathbf{s}_i = (s_i^1, \dots, s_i^L)^\top$	the vector of annotated labels for instance \mathbf{x}_i , and $s_i^k \in \{0, 1\}$, $s_i^k \leq y_i^k$
$\mathcal{D} = \{(\mathbf{x}_i, \mathbf{s}_i)\}_{i=1}^n$	the training set for multi-label learning with incomplete label assignments

Fig. 6

700

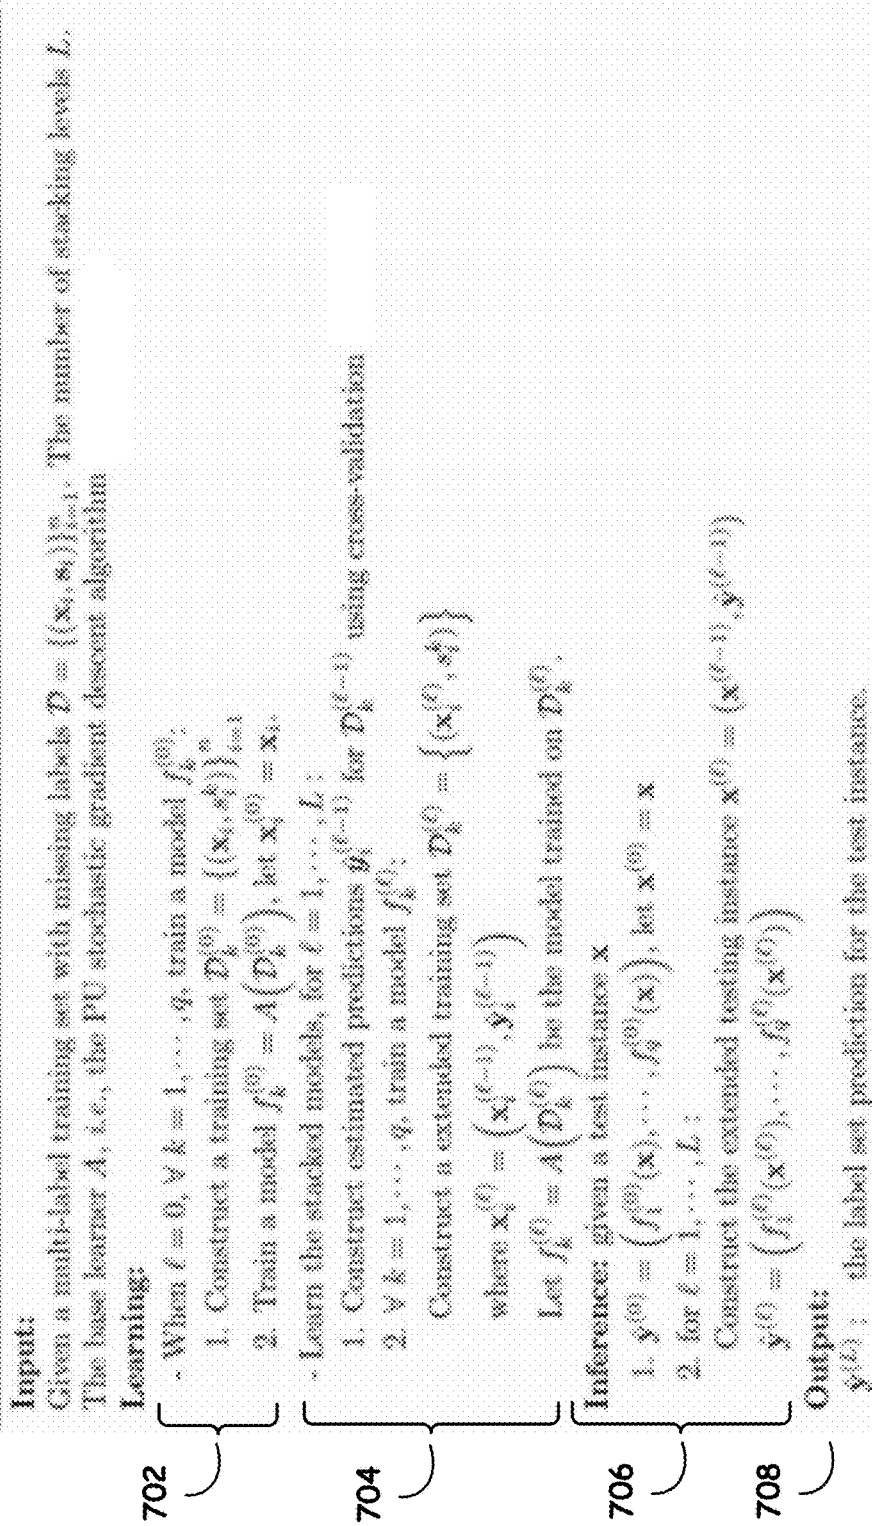


Fig. 7

Given a training set $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{s}_i)\}_{i=1}^n$ and a PU learning algorithm A . Construct a cross-validation prediction $\hat{\mathbf{y}}_i$ for each instance \mathbf{x}_i , as follows:

1. Convert \mathcal{D} into $\{\mathcal{D}_1, \dots, \mathcal{D}_q\}$, $\mathcal{D}_k = \{(\mathbf{x}_i, \mathbf{s}_i^k)\}_{i=1}^n$
2. $\forall k$, partition \mathcal{D}^k into m disjoint subsets with equal-size, denoted as $\mathcal{D}_k^1, \dots, \mathcal{D}_k^m$
3. $\forall j = 1, \dots, m, k = 1, \dots, q$,
train one model $f_k^j = A(\mathcal{D}_k - \mathcal{D}_k^j)$;
 $\forall \mathbf{x}_i \in \mathcal{D}_k^j, \hat{y}_i^k = f_k^j(\mathbf{x}_i)$;
4. $\forall i = 1, \dots, n$, we have $\hat{\mathbf{y}}_i = (\hat{y}_i^1, \dots, \hat{y}_i^q)^\top$
Return $\{\hat{\mathbf{y}}_i\}_{i=1}^n$

Fig. 8

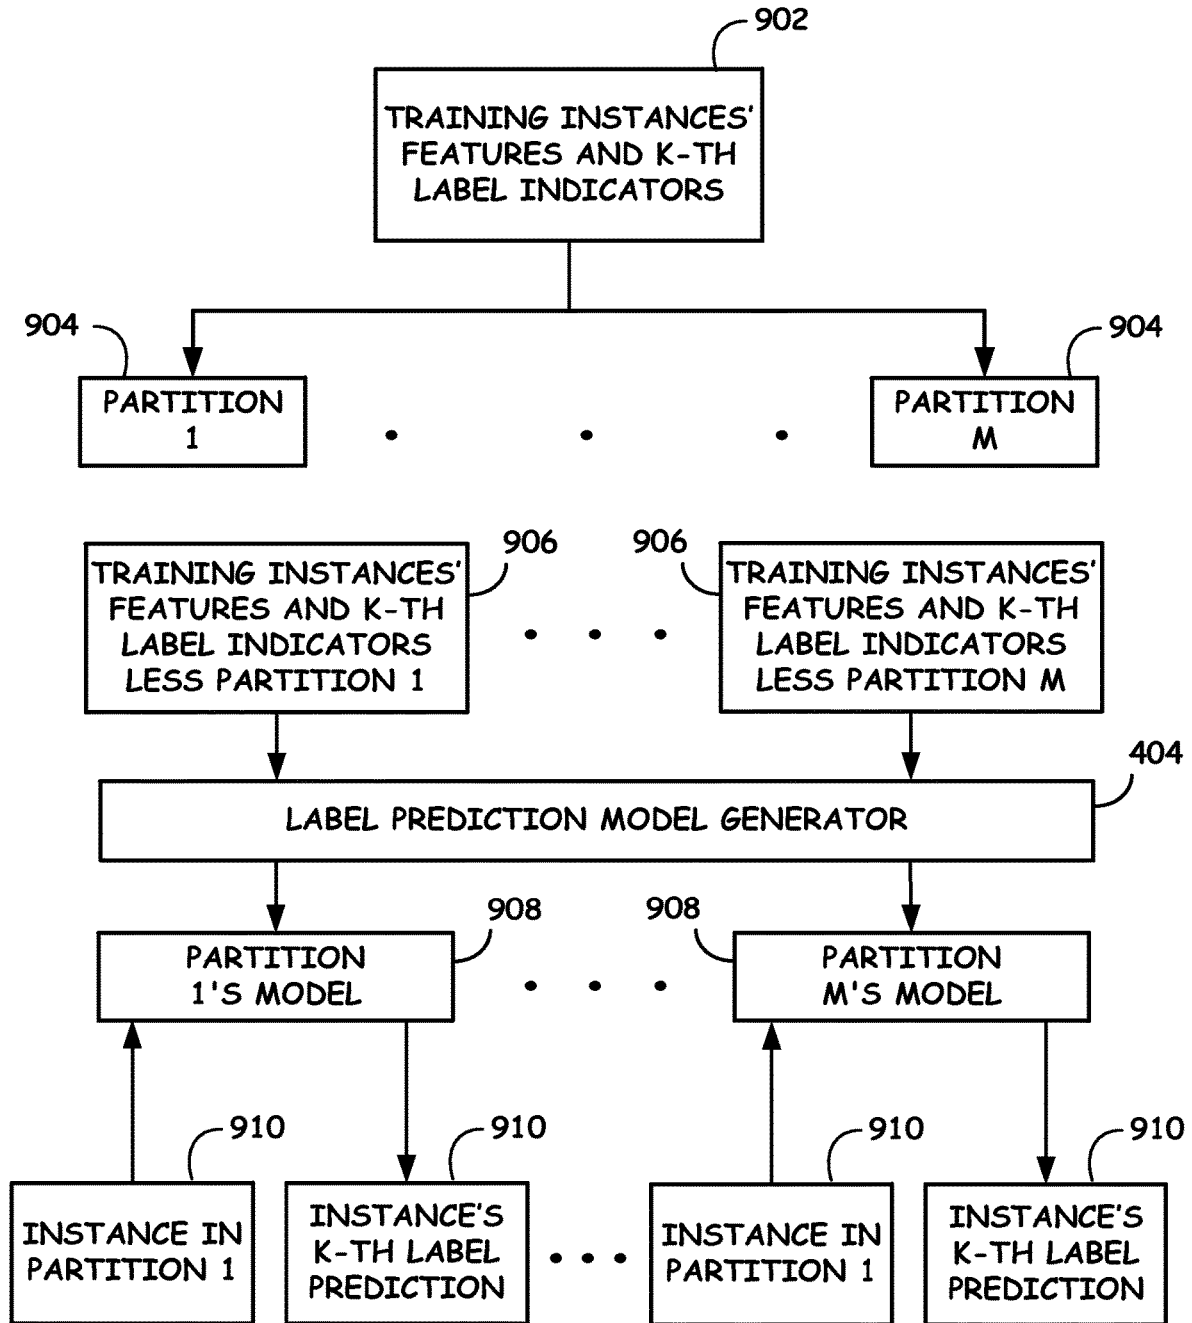


Fig. 9




Fig. 10




Fig. 11

**SYSTEM AND METHOD FOR
LARGE-SCALE MULTI-LABEL LEARNING
USING INCOMPLETE LABEL
ASSIGNMENTS**

CROSS REFERENCE TO RELATED
APPLICATIONS

[0001] This application is a continuation of, and claims priority from, co-pending U.S. patent application Ser. No. 16/439,855, filed Jun. 13, 2019, entitled SYSTEM AND METHOD FOR LARGE-SCALE MULTI-LABEL LEARNING USING INCOMPLETE LABEL ASSIGNMENTS, which is a continuation of, and claims priority from, U.S. patent application Ser. No. 14/543,133, filed Nov. 17, 2014, issued as U.S. Pat. No. 10,325,220 on Jun. 18, 2019, and entitled SYSTEM AND METHOD FOR LARGE-SCALE MULTI-LABEL LEARNING USING INCOMPLETE LABEL ASSIGNMENTS, the contents of each of which are hereby incorporated by reference.

FIELD OF THE DISCLOSURE

[0002] The present disclosure relates to automatically labeling, or tagging, of content, and more specifically to automatically labeling, or tagging, content using training data comprising partially-labeled training instances.

BACKGROUND

[0003] A content item may be annotated using one or more labels. For example, an image may have one or more associated labels that may identify objects depicted in the image, as well as other labels that may impart information about the image.

SUMMARY

[0004] In accordance with one or more embodiments of the present disclosure, at least one label prediction model is trained, or learned, using training data that may comprise training instances that may be missing one or more labels. The at least one label prediction model may be used in identifying a content item's ground-truth label set comprising an indicator for each label in the label set indicating whether or not the label is applicable to the content item.

[0005] In accordance with one or more embodiments, a method is provided, the method comprising training, using a computing device, an initial level of a stacked model for use in making a labeling prediction, the initial level being trained using feature information for each training instance of a plurality of training instances, at least one training instance of the plurality is missing at least one label of a plurality of labels, the feature information corresponding to a plurality of features associated with the training instance of the plurality; generating, using the computing device, a labeling prediction for each training instance of the plurality using the initial level of the stacked model, the labeling prediction comprising a label applicability prediction for at least one label of the plurality of labels missing from the training instance's set of labels; training, using the computing device, one or more additional levels of the stacked model, each additional level being trained using information for each training instance of the plurality, each training instance's information comprising the labeling prediction from a previous level of the stacked model, the feature information corresponding to the plurality of features, and

information indicating the training instance's set of labels; and identifying, using the computing device, a labeling prediction for a content item using the stacked model, the labeling prediction identifying for each label of the plurality whether the label is applicable to the content item.

[0006] In accordance with one or more embodiments a system is provided, which system comprises a processor and storage medium for tangibly storing thereon program logic for execution by the processor, the stored logic comprising training logic executed by the processor for training an initial level of a stacked model for use in making a labeling prediction, the initial level being trained using feature information for each training instance of a plurality of training instances, at least one training instance of the plurality is missing at least one label of a plurality of labels, the feature information corresponding to a plurality of features associated with the training instance of the plurality; generating logic executed by the processor for generating a labeling prediction for each training instance of the plurality using the initial level of the stacked model, the labeling prediction comprising a label applicability prediction for at least one label of the plurality of labels missing from the training instance's set of labels; training logic executed by the processor for training one or more additional levels of the stacked model, each additional level being trained using information for each training instance of the plurality, each training instance's information comprising the labeling prediction from a previous level of the stacked model, the feature information corresponding to the plurality of features, and information indicating the training instance's set of labels; and identifying logic executed by the processor for identifying a labeling prediction for a content item using the stacked model, the labeling prediction identifying for each label of the plurality whether the label is applicable to the content item.

[0007] In accordance with yet another aspect of the disclosure, a computer readable non-transitory storage medium is provided, the medium for tangibly storing thereon computer readable instructions that when executed cause at least one processor to train an initial level of a stacked model for use in making a labeling prediction, the initial level being trained using feature information for each training instance of a plurality of training instances, at least one training instance of the plurality is missing at least one label of a plurality of labels, the feature information corresponding to a plurality of features associated with the training instance of the plurality; generate a labeling prediction for each training instance of the plurality using the initial level of the stacked model, the labeling prediction comprising a label applicability prediction for at least one label of the plurality of labels missing from the training instance's set of labels; train one or more additional levels of the stacked model, each additional level being trained using information for each training instance of the plurality, each training instance's information comprising the labeling prediction from a previous level of the stacked model, the feature information corresponding to the plurality of features, and information indicating the training instance's set of labels; and identify a labeling prediction for a content item using the stacked model, the labeling prediction identifying for each label of the plurality whether the label is applicable to the content item.

[0008] In accordance with one or more embodiments, a system is provided that comprises one or more computing

devices configured to provide functionality in accordance with such embodiments. In accordance with one or more embodiments, functionality is embodied in steps of a method performed by at least one computing device. In accordance with one or more embodiments, program code to implement functionality in accordance with one or more such embodiments is embodied in, by and/or on a computer-readable medium.

DRAWINGS

[0009] The above-mentioned features and objects of the present disclosure will become more apparent with reference to the following description taken in conjunction with the accompanying drawings wherein like reference numerals denote like elements and in which:

[0010] FIG. 1 provides an example of a process flow for use in accordance with one or more embodiments of the present disclosure.

[0011] FIG. 2 provides an example of an instance for use in accordance with one or more embodiments of the present disclosure.

[0012] FIG. 3 provides an example of a training instance for use in accordance with one or more embodiments of the present disclosure.

[0013] FIG. 4 provides a stacked model generation example in accordance with one or more embodiments of the present disclosure.

[0014] FIG. 5 illustrates a stacked model inference example using test instance(s) in accordance with one or more embodiments of the present disclosure.

[0015] FIG. 6 provides some notational examples used herein in connection with one or more embodiments of the present disclosure.

[0016] FIG. 7 provides an example of model generation pseudocode in accordance with one or more embodiments of the present disclosure.

[0017] FIG. 8 provides a cross-validation pseudocode example for use in accordance with one or more embodiments of the present disclosure.

[0018] FIG. 9 provides an illustrative overview corresponding to the cross-validation example shown in FIG. 8.

[0019] FIG. 10 illustrates some components that can be used in connection with one or more embodiments of the present disclosure.

[0020] FIG. 11 is a detailed block diagram illustrating an internal architecture of a computing device in accordance with one or more embodiments of the present disclosure.

DETAILED DESCRIPTION

[0021] Subject matter will now be described more fully hereinafter with reference to the accompanying drawings, which form a part hereof, and which show, by way of illustration, specific example embodiments. Subject matter may, however, be embodied in a variety of different forms and, therefore, covered or claimed subject matter is intended to be construed as not being limited to any example embodiments set forth herein; example embodiments are provided merely to be illustrative. Likewise, a reasonably broad scope for claimed or covered subject matter is intended. Among other things, for example, subject matter may be embodied as methods, devices, components, or systems. Accordingly, embodiments may, for example, take the form of hardware, software, firmware or any combination thereof (other than

software per se). The following detailed description is, therefore, not intended to be taken in a limiting sense.

[0022] Throughout the specification and claims, terms may have nuanced meanings suggested or implied in context beyond an explicitly stated meaning. Likewise, the phrase “in one embodiment” as used herein does not necessarily refer to the same embodiment and the phrase “in another embodiment” as used herein does not necessarily refer to a different embodiment. It is intended, for example, that claimed subject matter include combinations of example embodiments in whole or in part.

[0023] In general, terminology may be understood at least in part from usage in context. For example, terms, such as “and”, “or”, or “and/or,” as used herein may include a variety of meanings that may depend at least in part upon the context in which such terms are used. Typically, “or” if used to associate a list, such as A, B or C, is intended to mean A, B, and C, here used in the inclusive sense, as well as A, B or C, here used in the exclusive sense. In addition, the term “one or more” as used herein, depending at least in part upon context, may be used to describe any feature, structure, or characteristic in a singular sense or may be used to describe combinations of features, structures or characteristics in a plural sense. Similarly, terms, such as “a,” “an,” or “the,” again, may be understood to convey a singular usage or to convey a plural usage, depending at least in part upon context. In addition, the term “based on” may be understood as not necessarily intended to convey an exclusive set of factors and may, instead, allow for existence of additional factors not necessarily expressly described, again, depending at least in part on context.

[0024] The detailed description provided herein is not intended as an extensive or detailed discussion of known concepts, and as such, details that are known generally to those of ordinary skill in the relevant art may have been omitted or may be handled in summary fashion. Certain embodiments of the present disclosure will now be discussed with reference to the aforementioned figures, wherein like reference numerals refer to like components.

[0025] In general, the present disclosure provides a system, method and architecture for use in multi-label learning that may be used with incomplete label assignments. In accordance with one or more embodiments, a training data set comprising a plurality of training instances, each of which corresponds to a content item, and one or more of which may be partially labeled may be used to train models in a stacked, or chained, modeling design. In accordance with one or more embodiments, label correlations may be used to facilitate learning. Embodiments of the present disclosure provide an ability for scaling to accommodate large and/or small amounts of data, where the scaling may be linear in both the number instances and number of classes. Empirical evidence derived from real-world datasets demonstrates that the approach taken in accordance with one or more embodiments significantly boosts performance of multi-label classification by considering missing labels and incorporating label correlations.

[0026] In accordance with one or more embodiments, ground-truth label sets may be estimated for training data using the incomplete label sets provided by the training data, and the estimated ground-truth label sets may be used to facilitate training of multi-label learning. In accordance with

one or more such embodiments, correlations among multiple labels may be exploited to improve multi-label classification performance.

[0027] In accordance with one or more embodiments, missing label assignments in a training data set may be addressed using a positive and unlabeled (PU) stochastic gradient descent. In accordance with one or more such embodiments, a framework may be built for incorporating correlations among labels based upon stack models. A multi-label learning method referred to herein as Mpu may consider label correlations using a stacked model, which need not rely on joint inference for all labels. Stacking is an example of ensemble methods, which build a chain of methods. Each model in the stacking uses output of a previous model, or models, as input. A stacked multi-label model allows inferences upon one label to influence inferences about other labels, and uses a base model to predict class labels for each label and uses those inferred labels as input to another level of the stacked model. Performance of multi-label classification with partially labeled training data is significantly boosted using embodiments of the present disclosure that consider missing labels using a positive and unlabeled learning model and label correlations using stacking models.

[0028] Embodiments of the present disclosure may be used, for example and without limitation, for automatically tagging, or labeling, content items. By way of a non-limiting example, photographs uploaded to a social media site, such as and without limitation Flickr®, may be automatically tagged using embodiments of the present disclosure. By way of a further non-limiting example, documents may be classified based on labels assigned in accordance with embodiments of the present disclosure.

[0029] FIG. 1 provides an example of a process flow for use in accordance with one or more embodiments of the present disclosure. At step 102, an initial level of a stacked model structure is trained and may be used to estimate the probability that a label, which belongs to a set of labels, is applicable to a given instance, where the given instance may correspond to a content item. The probability is an estimate of the likelihood that a label, which may be currently excluded from a set of labels currently associated with the instance, or content item, is applicable, or inapplicable, to the content. By way of a non-limiting example and with reference to FIG. 2, an instance 202, which in the example is an image content item, is annotated, labeled or tagged, with labels 206 of label set 204, which includes labels 208 that are currently not being used as labels for instance 202. In other words, labels 208 are missing from the annotations currently being used to label image 202. In the example of FIG. 2, dotted lines, such as line 210, represent a correlation between at least two labels in the label set 204. In accordance with one or more embodiments, correlations between labels may be used to infer whether or not one or more of missing labels 208 are applicable to instance 202.

[0030] Referring again to step 102 of FIG. 1, an initial level of the stacked model may be generated using a training data set comprising a plurality of training instances, which may include one or more training instances, such as instance 202 of FIG. 2, missing one or more labels from the set of labels 204.

[0031] FIG. 3 provides an example of a training instance for use in accordance with one or more embodiments of the present disclosure. While embodiments of the present dis-

closure are described in connection with an image as a content item, it should be apparent that any data item or content item may be used in accordance with one or more embodiments. In the example shown in FIG. 3, training instance 302, which may be any type of data item or content item, may comprise a set of labels 304 and a set of features 312.

[0032] As discussed in connection with the example shown in FIG. 2, a particular instance may be missing one or more labels. With reference to the example shown in FIG. 3, instance 302 is annotated with labels 306, and is missing labels 308, from label set 304. A label in label set 304 may be related, or correlated, with one or more label(s) in label set 304. By way of a non-limiting example, lines 310 illustrate some of the correlations, or relationships, among labels in label set 304. Label set 204 is a non-limiting example of a label set that includes labels identified for content items including image 202. In the example shown in FIG. 2, label set 204 includes labels representing objects depicted in the image. By way of a non-limiting example, instance 302 might be a document and labels in label set 304 might include words contained in the document. Training instance 302 may further comprise a feature set 312 comprising one or more feature(s). By way of a non-limiting example, an image content item may comprise such features as color, texture etc.

[0033] Labels 306 annotating instance 302 might be provided by one or more individuals, or via some other source. While instance 302 is missing labels 308 of label set 304, the reason for the missing labels 308 is unclear. It may be that a label 308 may be missing because it is not applicable to instance 302; alternatively, the missing labels 308 may be missing because the labeling source(s) inadvertently neglected to identify the label 308 as being applicable to the instance 302, and not because the missing label 308 is inapplicable to instance 302. Instance 302 may be one of a number of instances annotated by one or more individual(s) that might miss labeling the instance with one or more labels that are applicable to the instance. In the example of FIG. 3, each label 308 of label set 304 might be applicable to instance 302. The potential for missing labels is especially heightened when the number of instances is very large and/or the cost of labeling is very high. A large training data set and/or high labeling costs may make it almost impossible to have a fully labeled training set. Embodiments of the present disclosure may be used to facilitate identification of missing label(s) 308 that are applicable to instance 302. In accordance with one or more such embodiments, the missing label(s) 308 applicable to instance 302 may be identified by learning from partially-labeled training instances and using label correlations to facilitate the learning.

[0034] Referring again to FIG. 1, an initial level of a stacked model generated at step 102 may be used to make a prediction regarding an instance's ground-truth label set. By way of a non-limiting example, the initial level of the stacked model may be used to determine a probability, for each one of the missing labels 308, that it is applicable to the instance 302. By way of a further non-limiting example, the initial level of the stacked model generated at step 102 may be used to determine a probability, or likelihood, that a label 308 is missing but applicable as opposed to not used because it is inapplicable.

[0035] At step 106 of FIG. 1, another level of the stack may be generated using output provided by a previous

level's predictions, which output may comprise the previous level's ground-truth label set prediction. By way of a non-limiting example and in a case that the previous level is the initial level of the stack, the label set predictions which may be used to generate the current level of the stack may be the predictions generated at step 104. At step 108, the current level in the stack, which comprises one or more models generated at step 106, is used to make ground-truth label set predictions. As is illustrated in the example shown in FIG. 1, steps 106 and 108 may be optionally repeated to accommodate any number of levels in addition to the initial level.

[0036] In accordance with one or more embodiments, the initial and at least one additional levels of the stacked model may be used to make a labeling prediction for a content item, the labeling prediction may identify for each label of the plurality whether the label is applicable to the content item. In accordance with one or more such embodiments, the labeling prediction may comprise information identifying a ground-truth label set for the content item.

[0037] FIG. 4 provides a stacked model generation example in accordance with one or more embodiments of the present disclosure. A set 402 of training instances, which may comprise a number of instances such as instance 302 or instance 202, is input to a label prediction model generator, which may be used to generate an initial level 406 of a stacked model. The initial level 406 may comprise a model for each label of a label set, e.g., label set 304. Each training instance of set 402 may comprise a feature set, such as and without limitation feature set 312 for instance 302. Label prediction model generator 404 may use the training set 402 to generate a label prediction model(s) for an initial level 406 of a stacked model.

[0038] In accordance with one or more embodiments, label prediction model 406 comprises model parameters including a set of weights, each weight in the set corresponding to a label in the label set, e.g., label set 304. The weighting may reflect a bias toward positive samples, e.g., instances annotated with the label, relative to negative samples, e.g., instances not annotated with the label, for determining a probability associated with a given label.

[0039] Label prediction models belonging to the initial level 406 may be used to generate, for each instance in a training set, such as and without limitation training set 402, a set of predictions about the instance's ground-truth label set, which set of predictions may comprise a prediction for each label in the set of labels, such as label set 304. The set of predictions associated with a given instance, such as instance 302, comprises a probability, for each label in the label set, that the label is applicable to the instance.

[0040] By way of a non-limiting example, each label in label set 304 may be assigned a probability using a corresponding label prediction model belonging to the initial level 406 indicating a likelihood that the label 306 is applicable to the instance 302. By way of a non-limiting example, a probability may have the value ranging from 0 indicating no applicability to 1 indicating a high applicability. By way of a further non-limiting example, each label 306 that currently annotates instance 302 might be assigned a value of 1 indicating that it has the highest likelihood of being applicable to the instance 302, and each label 308 in label set 304 might be assigned a value in the range of 0 to 1 indicating the probability that the label 308 is applicable, or inapplicable, to the instance 302.

[0041] In the example shown in FIG. 4, label prediction models in the initial level 406 use training data set 402, which may comprise, for each instance 302, the labels 306 that annotate the instance 302 and the features of feature set 312 for the instance, to generate output comprising a ground-truth label set prediction for each instance 302 in the training data set 402. For each instance of training set 402, the output of the initial level 406 may comprise, a feature set 312, a label set 304, and a ground-truth label set prediction. In accordance with one or more embodiments, a training instance's ground-truth label set prediction may be included with the training instance's feature set 312. Training data set 408 may be input to label prediction model generator 404, which may generate each subsequent level of the stacked model using a previous level's output.

[0042] In accordance with one or more embodiments, the initial and subsequent levels of the stacked model may comprise a model for each label in a label set, e.g., label set 304. With reference to label set 304 of FIG. 3, each level of the stacked model may comprise a model for each label 306 and each label 308 of label set 304.

[0043] In accordance with one or more embodiments, the stacked model may have any number of levels in addition to the initial level, and label prediction model generator 404 may be used to generate each level of the stacked model. In the example of FIG. 4, a current level of the stacked model that follows the initial level may be generated by the label prediction model generator 404 using the ground-truth label predictions generated by the previous level of the stacked model, e.g., the initial level's 406 ground-truth label predictions may be used to generate a next level, which generates a current set of ground-truth label set predictions 412, which may be in turn used by the label prediction model generator 404 to generate another level of the stacked model.

[0044] By way of a non-limiting example, label prediction model generator 404 may use training data set 402 which comprises both positive and negative samples to generate a level of the stacked model. By way of a further non-limiting example, for a given label, a positive sample may be a training instance that includes the label as an annotation for the training instance's content item, and conversely a negative sample may be a training instance that does not include the label as an annotation for the training instance's content item. In accordance with at least one embodiment, training data set 402 may comprise a label set 304 comprising both positive and negative samples for a given training instance. Each instance of the training data set 402 may have a feature set 312 comprising one or more features of the training instance's content item.

[0045] In accordance with one or more embodiments, label prediction model generator 404 may generate a set of parameters including a weight for each label in the label set 304. The set of model parameters may be tested, using test data comprising a set of instances, and the label prediction model generator 404 may regenerate the set of model parameters if a level of accuracy is not achieved. In regenerating the label prediction model 406, the model parameters may be modified, e.g., the weights assigned to one or more of the negative samples in the label set 304 may be modified, and the model parameters may be retested using the test data set to arrive at an acceptable level of accuracy. By way of a non-limiting example, the test data set may comprise one or more instances that are known to be positive samples for one or more labels, and the label prediction model 406 may be

tested to determine whether or not the set of model parameters identifies such instances as having a high probability that the label(s) is/are applicable to the one or more instances.

[0046] FIG. 5 illustrates a stacked model inference example using test instance(s) in accordance with one or more embodiments of the present disclosure. In the example shown in FIG. 5, an iterative approach may be used in generating a ground-truth label set prediction for each test instance 502. A test instance 502 may comprise a label set 304, which may comprise labels 306 and 308, and a feature set 312, which may comprise any number of features, may be input to the initial level 406 of the stacked model. The initial level 406 may generate a ground-truth label set prediction for the test instance 502 using the instance's feature set 312. The initial level 406 may generate output in connection with the test instance 502, which output may comprise the label set 304, feature set 312 and ground-truth label set predictions, which becomes input to the next level 410 of the stacked model. The next level 410 of the stacked model, which becomes the current level, may use the input to generate its ground-truth label set predictions. The process may be iteratively repeated until a final level of the stacked model outputs the ground-truth label set for the test instance 502.

[0047] By virtue of using a stacked model approach, output from a level of the stacked model may be used by another level of the stacked model. By way of a non-limiting example, a stacked model level 502, or 512, may determine that a missing label 308 is applicable to the instance 502. By way of a non-limiting example, a label identified as being applicable to an instance at one level of the model may be used together with label correlations in the label set 304 in determining whether another label, e.g., a label 308, is applicable to the instance 502.

[0048] FIG. 6 provides some notational examples used herein in connection with one or more embodiments of the present disclosure. A feature vector may be used to express a feature set 312 of an instance 302. A feature vector for an instance i may be represented herein as x_i , and feature vectors for a set of n training instances may be represented as x_1, \dots, x_n . A label set, such as and without limitation label set 304, which may be referred to as a dictionary of labels, may comprise a number, e.g., q , of possible labels. Each instance i may have a set of ground-truth labels represented as $y_i = \{y_i^1, \dots, y_i^q\}$, where $y_i^k = 1$ may indicate that the k -th label in the label set, e.g., label set 304 of instance i , is applicable, in which case $y_i^k = 0$ may be used to indicate that the k -th label is inapplicable.

[0049] An instance i may have an associated set of features, x_i , and a ground-truth label set, y_i . Embodiments of the present disclosure may be used to determine an instance's ground-truth label set, which is beneficial since an instance in a training data set of instances is likely not to be annotated by its ground-truth label set, e.g. the instance's label set is missing one or more labels that is/are appropriate for, or applicable to, the instance. For example and without limitation, the instance may not be fully annotated or labeled by a labeler, in which case an instance i may be represented by its set of features, x_i , and its set of annotations, s_i , where $s_i = (s_i^1, \dots, s_i^q)^T \in \{0, 1\}^q$, where $s_i^k \leq y_i^k (\forall 1 \leq i < n, \forall 1 \leq k \leq q)$. In other words, s_i may be used to denote a label set that may not be the same as the instance's ground-truth label set, e.g., the label set, s_i , may be missing one or more labels that is/are

considered to be applicable in its ground-truth label set, y_i . For each k -th label in a label set, when $s_i^k = 1$, $y_i^k = 1$, and the probability that a label s_i^k appears in the s_i is zero if the label is absent from the ground-truth label set, which may be expressed as follows:

$$Pr(s_i^k=1 | x_i, y_i^k=0) = 0, i, k$$

[0050] When a label, s_i^k , is missing from the set of annotated labels, s_i , it is not clear whether the missing label is missing because it is not applicable to the instance, in which case $y_i^k = 0$ indicating that the label is absent from the instance's ground-truth label set, or the missing label is missing because a labeler neglected to annotate the instance using the label, in which case $y_i^k = 1$ indicating that the label is present in the instance's ground-truth label set. Embodiments of the present disclosure provide a mechanism for determining, for each missing label of a set of multiple labels, which alternative is correct using a stacked model approach, which has an initial level, such as initial level 406, and a number, L , of subsequent levels such as level(s) 410. Levels 406 and 410 may be trained by model generator 404 using a positive and unlabeled stochastic gradient descent learner. The stacked model may use label correlations to determine an instance's ground-truth label set. In accordance with one or more embodiments, model learning, which may be performed by model generator 404, is used to generate the models to make predictions that accurately predict the ground-truth for each label of multiple labels.

[0051] Referring again to FIG. 6, a set of annotated labels for a set of n training instances, such as training instance set 402, may be represented as s_1, \dots, s_n , where an annotated label set for an instance i may be represented as $s_i = (s_i^1, \dots, s_i^q)^T$. A set of ground-truth label sets for the n training instances may be represented as y_1, \dots, y_n , where a ground-truth label set for an instance i may be represented as $y_i = (y_i^1, \dots, y_i^q)^T$. A training data set, such as training set 402, may be represented as $\mathcal{D} = \{(x_i, s_i)\}_{i=1}^n$, which may be a multi-label training set with missing labels.

[0052] FIG. 7 provides an example of model generation pseudocode in accordance with one or more embodiments of the present disclosure. Portions 702 and 704 of pseudocode 700 may be used to generate a stacked model comprising an initial level and one or more subsequent levels. The initial level of the stacked model may be the label prediction model level 406 and each subsequent level may be a label prediction model level 410, and levels 406 and 410 may be generated using a model learner, such as label prediction model generator 404, using the positive and unlabeled stochastic gradient descent learning method and a multi label training set with missing labels.

[0053] Portion 702 of the code 700 may be used to generate the initial level of the stacked model. The initial level of the stacked model may be generated using the feature set, x_i , which may correspond to instance 302, of the training set \mathcal{D} , which may correspond to training set 402. The initial level may comprise a model, $f_k^{(0)}$, for each k -th label of the label set, e.g., label set 304. For each instance in the training set, each label, e.g., each k -th label, of the label set may have a model that is generated using a model learner, A , and a training data set $\mathcal{D}_k^{(0)} = \{(x_i, s_i^k)\}_{i=1}^n$. For each k -th label, the initial level's training data set, $\mathcal{D}_k^{(0)}$, may therefore comprise the feature set of each instance, i , and a value, e.g., 0 or 1, indicating whether the label is

present, or absent, in the instance's label set, e.g., whether or not the label is being used to annotate the instance.

[0054] Referring to portion **704** of the code **700**, the initial level of the stacked model may be used to infer labels, e.g., the ground-truth label set, for each instance. By way of a non-limiting example, a set of predictions is generated by the initial level of the stacked model, the prediction set comprising, for each instance i , a ground-truth label set prediction, $\hat{y}_i^{(0)}$, which includes a prediction for each k -th label generated using the k -th label's model, $f_k^{(0)}$, and the k -th label's training data set, $\mathcal{D}_k^{(0)}$. As is discussed in more detail below in connection with FIGS. **8** and **9**, cross validation may be used in inferring labels at the initial level, and subsequent levels, of the stacked model.

[0055] The label inferences generated by the initial level of the stacked model may be used, together with the feature set **312** and label set **304** of each instance, to train the next level of the stacked model. More generally speaking, the label inferences generated by a previous level, $l-1$, may be used in a current level, l , to train the current level of the stacked model. At a current level, l , a model, $f_k^{(l)}$, may be generated for each k -th label, using a model generator, A , and a training data set, $\mathcal{D}_k^{(l)} = \{(x_i^l, s_i^k)\}$ for each instance i , where $x_i^{(l)} = (x_i^{(l-1)}, \hat{y}_i^{(l-1)})$, such that the feature set, $x_i^{(l)}$, for an instance, i , comprises the previous level's feature set, $(x_i^{(l-1)})$, and the previous level's label inferences, $\hat{y}_i^{(l-1)}$. A model trained on $\mathcal{D}_k^{(l)}$ may be expressed as $f_k^{(l)} = A(\mathcal{D}_k^{(l)})$, for each label, k , of a number, q , labels of a label set.

[0056] Portion **704** may be repeated for each level, l , of L levels of a stacked model. The number of levels may be a predetermined number, and/or may be a number that is empirically determined based on a determined convergence, which may be determined based on whether or not there are any improvements in the estimates from one level, or iteration, to the next level.

[0057] Portion **706** makes inferences about each test instance, x , using the stacked model learned using portions **702** and **704**. For each instance, x , the levels, e.g., $l=0$ to L , of the stacked model may be used to make inferences about the instance. The output from a previous level, $l-1$, may be used as input to a current level, l , in the stacking, and the last level, $l=L$, of the stacked model may be used to generate a final set of predictions, a prediction of the ground-truth label set for the test instance, x .

[0058] In accordance with one or more embodiments, the stacked model that may be used to predict the ground-truth label set for an instance is learned using inferred labels, e.g., learning using a label set that may or may not be the ground-truth label set. Advantageously, embodiments of the present disclosure may train the stacked model on inferred labels, so that the trained model may be used to make exact inferences regarding an instance's true labels using known features of the instance. This may be contrasted with an approach that requires ground-truth label sets in learning and makes approximate inferences regarding true labels, which true labels are not known at inference time.

[0059] With reference to portion **706**, a set of predictions, $\hat{y}^{(0)} = (f_1^{(0)}(x), \dots)$, where $x^{(0)} = x$, may be generated for the initial level of the stacked model using the test instance's feature set, e.g., feature set **312**, and each label's model generated for the initial model, e.g., the k -th label's base level model may be represented as $f_k^{(0)}$. The set of predictions, $\hat{y}^{(0)}$, from the initial level of the stacked model may be used as input to the next level, e.g. $l=1$, to generate a set of

predictions, $\hat{y}^{(1)} = (f_1^{(1)}(x^{(1)}), \dots, f_q^{(1)}(x^{(1)}))$, using an extended testing instance, $x^{(1)} = (x^{(0)}, \hat{y}^{(0)})$, comprising the initial level's set of predictions and feature set for the instance. More generally speaking, at each level of the stacked model following the initial level of the model, a set of predictions, $\hat{y}^{(l)} = (f_1^{(l)}(x^{(l)}), \dots, f_q^{(l)}(x^{(l)}))$, may be generated using an extended testing instance, $x^{(l)} = (x^{(l-1)}, \hat{y}^{(l-1)})$, comprising the previous level's set of predictions and feature set for the instance.

[0060] In accordance with one or more embodiments, the base learner, A , may be a positive and unlabeled gradient descent model learner. In accordance with one or more embodiments, a positive and unlabeled stochastic gradient descent approach, which can handle large-scale data sets with missing label assignments, is used to learn a set of parameters, $\{w_k\}_{k=1}^q$ where $x_i \in \mathbb{R}^D$, comprising a weight for each k -th label in a set of q labels, e.g., such as and without limitation label set **304**. In accordance with one or more such embodiments, the weights are optimized to maximize the likelihood of y_i^k ; in other words to maximize the likelihood of determining the ground-truth labels for any content item, or instance, i . An optimized parameter, w_k^* , for the k -th label may be expressed as:

$$w_k^* = \operatorname{argmax}_{w_k} \log \left(\prod_{i=1}^n Pr(y_i^k = 1 | x_i, w_k) \right)$$

[0061] In accordance with one or more embodiments, the positive and unlabeled stochastic gradient descent learning method extends logistic regression to classification with incomplete label assignments, and may use assumptions that y_i^k satisfies a Bernoulli distribution, and

$$Pr(y_i^k = 1 | x_i, w_k) = \frac{1}{1 + \exp(-w_k T x_i)}$$

[0062] An assumption may be made that annotated label are randomly sampled from the ground-truth label set with a constant rate, c , where the sampling process may be independent of other factors, such as a feature of the instance. In a case that it is assumed that the probability that a label is not missing by the labeler is an unknown constant, such a constant, c , may be expressed as:

$$c = PR(s_i^k = 1 | y_i^k = 1) = PR(s_i^k = 1 | y_i^k = 1, x_i, w_k),$$

[0063] where c may be directly estimated from the training set using cross validation. Using Bayes' theorem:

$$Pr(y_i^k = 1 | x_i, w_k) = \frac{PR(s_i^k = 1 | x_i, w_k)}{PR(s_i^k = 1 | y_i^k = 1, x_i, w_k)}$$

[0064] The probability of a missing label is applicable to an instance, i , may be expressed as:

$$Pr(s_i^k = 1 | x_i, w_k) = \frac{c}{1 + \exp(-w_k T x_i)}$$

and

[0065] an optimized parameter, w_k^* , for the k-th label may be represented as:

$$w_k^* = \operatorname{argmax}_{w_k} \sum_{i=1}^n \log \left(\frac{1}{1 + \exp(-w_k T x_i)} + \frac{(1 - s_i^k) + (1 - c)}{2} \right)$$

[0066] Embodiments of the present disclosure are able to scale to large-scale problems using stochastic gradient descent to solve the logistic regression efficiently. In accordance with one or more such embodiments, rather than assuming that all of the labels are available, e.g., an assumption that would conclude that a missing label is not appropriate or applicable to an instance, incomplete label assignments are examined to make a determination whether a missing label is applicable to an instance. In accordance with one or more such embodiments, a loss function may be used to weight negative samples, which loss function may be represented as follows:

$$l(w_k, \mathcal{D}) = - \sum_{i=1}^n \log \frac{1}{1 + \exp(-w_k T x_i)} + \frac{(1 - s_i^k) + (1 - c)}{2}$$

[0067] Referring again to portion 704 of FIG. 7, a cross validation technique may be used in accordance with one or more embodiments of the present disclosure. In accordance with one or more such embodiments, some training instances may be excluded from use in training a model, which model may be used to generate a label prediction for the excluded training instances. FIG. 8 provides a cross-validation pseudocode example for use in accordance with one or more embodiments of the present disclosure. FIG. 9 provides an illustrative overview corresponding to the cross-validation example shown in FIG. 8.

[0068] In the example shown in FIG. 8, a cross-validation prediction, \hat{y}_i , may be generated for each instance, x_i , using a training set, $\mathcal{D} = \{(x_i, s_i)\}_{i=1}^n$, and a base learner, A. For each label, k, in a set of q labels, a training data set may be determined using \mathcal{D} , which may be converted into $\{\mathcal{D}_1, \dots, \mathcal{D}_q\}$, where $\mathcal{D}_k = \{(x_i, s_i^k)\}_{i=1}^n$, for each label, k, in the label set. For each label, k, its corresponding training data set, \mathcal{D}_k , may be partitioned into a number, m, disjoint subsets having a similar size, an equal, approximately equal, etc. size, and the resulting partitions may be denoted as $\mathcal{D}_k^1, \dots, \mathcal{D}_k^m$. For each k-th label, m models are trained, and label predictions are made for each instance, i, in a given partition using a model that is trained without using the partition's training instances. In other words, for a label, k, and a partition, j, a model f_k^j is trained using a model learner, A, such that the k-th label's training data set excludes the training instances assigned to partition j of the k-th label's training data set, which may be denoted as $\mathcal{D}_k - \mathcal{D}_k^j$. The resulting model, f_k^j , may be used to generate a set of predictions, $\hat{y}_i^k(x_i) = f_k^j(x_i)$, for each i-th instance having a feature set, x_i , belonging to the j-th partition, or $x_i \in \mathcal{D}_k^j$.

[0069] With reference to FIG. 9, a k-th label's training data set 902, which is denoted as $\mathcal{D}_k = \{(x_i, s_i^k)\}_{i=1}^n$ in FIG. 8, is partitioned into a number, m, partitions 904, which partitioning is denoted as $\mathcal{D}_k^1, \dots, \mathcal{D}_k^m$ in FIG. 8. Model generator 404 uses m training data sets to generate m models 908, such that for a given j-th one of the models 908, a j-th

one of the m training data set partitions is excluded from the training data set used to generate the model 908, which is denoted as $f_k^j = A(\mathcal{D}_k - \mathcal{D}_k^j)$. The model 908, which is represented as f_k^j in FIG. 8, is used to generate a k-th label prediction 910 for each instance 910 in the partition 904 excluded from being used in generating the model 908. For a given instance 910, model 908 may use the instance's feature set, x_i , which may or may not include label predictions, such that the inclusion of label predictions in an instance's feature set may depend on the current level of the stacked model.

[0070] FIG. 10 illustrates some components that can be used in connection with one or more embodiments of the present disclosure. In accordance with one or more embodiments of the present disclosure, one or more computing devices, e.g., one or more servers, user devices or other computing device, are configured to comprise functionality described herein. For example, one or more of the computing device 1002 may be configured to execute program code, instructions, etc. to provide functionality in accordance with one or more embodiments of the present disclosure.

[0071] Computing device 1002 may serve content to user computing devices 1004 using a browser application via a network 1006. Data store 1008 may be used to store program code to configure a server 1002 to functionality in accordance with one or more embodiments of the present disclosure. By way of a non-limiting examples, computing device 1002 may serve content to a user computing device 1004, which may include one or more labels of a ground-truth label set determined using embodiments of the present disclosure. The content served may comprise a content item and the one or more labels from the content item's determined ground-truth label set. As yet another non-limiting example, computing device 1002 may receive a request from a user computing device 1004 to retrieve a content item, and may identify one or more content items by searching one or more ground-truth label sets using a label query, which may be received as part of the request, for responding to the content item request.

[0072] The user computing device 1004 can be any computing device, including without limitation a personal computer, personal digital assistant (PDA), wireless device, cell phone, internet appliance, media player, home theater system, and media center, or the like. For the purposes of this disclosure a computing device includes a processor and memory for storing and executing program code, data and software, and may be provided with an operating system that allows the execution of software applications in order to manipulate data. A computing device such as server 1002 and the user computing device 1004 can include one or more processors, memory, a removable media reader, network interface, display and interface, and one or more input devices, e.g., keyboard, keypad, mouse, etc. and input device interface, for example. One skilled in the art will recognize that server 1002 and user computing device 1004 may be configured in many different ways and implemented using many different combinations of hardware, software, or firmware.

[0073] In accordance with one or more embodiments, a computing device 1002 can make a user interface available to a user computing device 1004 via the network 1006. The user interface made available to the user computing device 1004 can include content items, or identifiers (e.g., URLs) selected for the user interface in accordance with one or

more embodiments of the present invention. In accordance with one or more embodiments, computing device **1002** makes a user interface available to a user computing device **1004** by communicating a definition of the user interface to the user computing device **1004** via the network **1006**. The user interface definition can be specified using any of a number of languages, including without limitation a markup language such as Hypertext Markup Language, scripts, applets and the like. The user interface definition can be processed by an application executing on the user computing device **1004**, such as a browser application, to output the user interface on a display coupled, e.g., a display directly or indirectly connected, to the user computing device **1004**.

[**0074**] In an embodiment the network **1006** may be the Internet, an intranet (a private version of the Internet), or any other type of network. An intranet is a computer network allowing data transfer between computing devices on the network. Such a network may comprise personal computers, mainframes, servers, network-enabled hard drives, and any other computing device capable of connecting to other computing devices via an intranet. An intranet uses the same Internet protocol suite as the Internet. Two of the most important elements in the suit are the transmission control protocol (TCP) and the Internet protocol (IP).

[**0075**] As discussed, a network may couple devices so that communications may be exchanged, such as between a server computing device and a client computing device or other types of devices, including between wireless devices coupled via a wireless network, for example. A network may also include mass storage, such as network attached storage (NAS), a storage area network (SAN), or other forms of computer or machine readable media, for example. A network may include the Internet, one or more local area networks (LANs), one or more wide area networks (WANs), wire-line type connections, wireless type connections, or any combination thereof. Likewise, sub-networks, such as may employ differing architectures or may be compliant or compatible with differing protocols, may interoperate within a larger network. Various types of devices may, for example, be made available to provide an interoperable capability for differing architectures or protocols. As one illustrative example, a router may provide a link between otherwise separate and independent LANs. A communication link or channel may include, for example, analog telephone lines, such as a twisted wire pair, a coaxial cable, full or fractional digital lines including T1, T2, T3, or T4 type lines, Integrated Services Digital Networks (ISDNs), Digital Subscriber Lines (DSLs), wireless links including satellite links, or other communication links or channels, such as may be known to those skilled in the art. Furthermore, a computing device or other related electronic devices may be remotely coupled to a network, such as via a telephone line or link, for example.

[**0076**] A wireless network may couple client devices with a network. A wireless network may employ stand-alone ad-hoc networks, mesh networks, Wireless LAN (WLAN) networks, cellular networks, or the like. A wireless network may further include a system of terminals, gateways, routers, or the like coupled by wireless radio links, or the like, which may move freely, randomly or organize themselves arbitrarily, such that network topology may change, at times even rapidly. A wireless network may further employ a plurality of network access technologies, including Long Term Evolution (LTE), WLAN, Wireless Router (WR)

mesh, or 2nd, 3rd, or 4th generation (2G, 3G, or 4G) cellular technology, or the like. Network access technologies may enable wide area coverage for devices, such as client devices with varying degrees of mobility, for example. For example, a network may enable RF or wireless type communication via one or more network access technologies, such as Global System for Mobile communication (GSM), Universal Mobile Telecommunications System (UMTS), General Packet Radio Services (GPRS), Enhanced Data GSM Environment (EDGE), 3GPP Long Term Evolution (LTE), LTE Advanced, Wideband Code Division Multiple Access (WCDMA), Bluetooth, 802.11b/g/n, or the like. A wireless network may include virtually any type of wireless communication mechanism by which signals may be communicated between devices, such as a client device or a computing device, between or within a network, or the like.

[**0077**] Signal packets communicated via a network, such as a network of participating digital communication networks, may be compatible with or compliant with one or more protocols. Signaling formats or protocols employed may include, for example, TCP/IP, UDP, DECnet, NetBEUT, IPX, Appletalk, or the like. Versions of the Internet Protocol (IP) may include IPv4 or IPv6. The Internet refers to a decentralized global network of networks. The Internet includes local area networks (LANs), wide area networks (WANs), wireless networks, or long haul public networks that, for example, allow signal packets to be communicated between LANs. Signal packets may be communicated between nodes of a network, such as, for example, to one or more sites employing a local network address. A signal packet may, for example, be communicated over the Internet from a user site via an access node coupled to the Internet. Likewise, a signal packet may be forwarded via network nodes to a target site coupled to the network via a network access node, for example. A signal packet communicated via the Internet may, for example, be routed via a path of gateways, servers, etc. that may route the signal packet in accordance with a target address and availability of a network path to the target address.

[**0078**] It should be apparent that embodiments of the present disclosure can be implemented in a client-server environment such as that shown in FIG. 10. Alternatively, embodiments of the present disclosure can be implemented with other environments. As one non-limiting example, a peer-to-peer (or P2P) network may employ computing power or bandwidth of network participants in contrast with a network that may employ dedicated devices, such as dedicated servers, for example; however, some networks may employ both as well as other approaches. A P2P network may typically be used for coupling nodes via an ad hoc arrangement or configuration. A peer-to-peer network may employ some nodes capable of operating as both a “client” and a “server.”

[**0079**] FIG. 11 is a detailed block diagram illustrating an internal architecture of a computing device, e.g., a computing device such as server **1002** or user computing device **1004**, in accordance with one or more embodiments of the present disclosure. As shown in FIG. 11, internal architecture **1100** includes one or more processing units, processors, or processing cores, (also referred to herein as CPUs) **1112**, which interface with at least one computer bus **1102**. Also interfacing with computer bus **1102** are computer-readable medium, or media, **1106**, network interface **1114**, memory **1104**, e.g., random access memory (RAM), run-time tran-

sient memory, read only memory (ROM), etc., media disk drive interface **1120** as an interface for a drive that can read and/or write to media including removable media such as floppy, CD-ROM, DVD, etc. media, display interface **1110** as interface for a monitor or other display device, keyboard interface **1116** as interface for a keyboard, pointing device interface **1118** as an interface for a mouse or other pointing device, and miscellaneous other interfaces not shown individually, such as parallel and serial port interfaces, a universal serial bus (USB) interface, and the like.

[0080] Memory **1104** interfaces with computer bus **1102** so as to provide information stored in memory **1104** to CPU **1112** during execution of software programs such as an operating system, application programs, device drivers, and software modules that comprise program code, and/or computer-executable process steps, incorporating functionality described herein, e.g., one or more of process flows described herein. CPU **1112** first loads computer-executable process steps from storage, e.g., memory **1104**, computer-readable storage medium/media **1106**, removable media drive, and/or other storage device. CPU **1112** can then execute the stored process steps in order to execute the loaded computer-executable process steps. Stored data, e.g., data stored by a storage device, can be accessed by CPU **1112** during the execution of computer-executable process steps.

[0081] Persistent storage, e.g., medium/media **1106**, can be used to store an operating system and one or more application programs. Persistent storage can also be used to store device drivers, such as one or more of a digital camera driver, monitor driver, printer driver, scanner driver, or other device drivers, web pages, content files, playlists and other files. Persistent storage can further include program modules and data files used to implement one or more embodiments of the present disclosure, e.g., listing selection module(s), targeting information collection module(s), and listing notification module(s), the functionality and use of which in the implementation of the present disclosure are discussed in detail herein.

[0082] For the purposes of this disclosure a computer readable medium stores computer data, which data can include computer program code that is executable by a computer, in machine readable form. By way of example, and not limitation, a computer readable medium may comprise computer readable storage media, for tangible or fixed storage of data, or communication media for transient interpretation of code-containing signals. Computer readable storage media, as used herein, refers to physical or tangible storage (as opposed to signals) and includes without limitation volatile and non-volatile, removable and non-removable media implemented in any method or technology for the tangible storage of information such as computer-readable instructions, data structures, program modules or other data. Computer readable storage media includes, but is not limited to, RAM, ROM, EPROM, EEPROM, flash memory or other solid state memory technology, CD-ROM, DVD, or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other physical or material medium which can be used to tangibly store the desired information or data or instructions and which can be accessed by a computer or processor.

[0083] Those skilled in the art will recognize that the methods and systems of the present disclosure may be implemented in many manners and as such are not to be

limited by the foregoing exemplary embodiments and examples. In other words, functional elements being performed by single or multiple components, in various combinations of hardware and software or firmware, and individual functions, may be distributed among software applications at either the client or server or both. In this regard, any number of the features of the different embodiments described herein may be combined into single or multiple embodiments, and alternate embodiments having fewer than, or more than, all of the features described herein are possible. Functionality may also be, in whole or in part, distributed among multiple components, in manners now known or to become known. Thus, myriad software/hardware/firmware combinations are possible in achieving the functions, features, interfaces and preferences described herein. Moreover, the scope of the present disclosure covers conventionally known manners for carrying out the described features and functions and interfaces, as well as those variations and modifications that may be made to the hardware or software or firmware components described herein as would be understood by those skilled in the art now and hereafter.

[0084] While the system and method have been described in terms of one or more embodiments, it is to be understood that the disclosure need not be limited to the disclosed embodiments. It is intended to cover various modifications and similar arrangements included within the spirit and scope of the claims, the scope of which should be accorded the broadest interpretation so as to encompass all such modifications and similar structures. The present disclosure includes any and all embodiments of the following claims.

1. A method comprising:

training, by a computing device, a multi-level label prediction model to predict an applicability, to a digital content item, of each label of a plurality of labels, the multi-level label prediction model comprising a first level and a second level, the first level being trained using training data missing a number of applicable labels, label prediction output of the trained first level including an applicability prediction for each of the number of applicable labels missing from the training data, the second level being trained using the label prediction output generated by the trained first level;

generating, by the computing device, the digital content item's label prediction output using the trained multi-level label prediction model, the digital content item's labeling prediction output indicating an applicability, to the digital content item, of each label of the plurality of labels; and

automatically generating, by the computing device, an annotation for the digital content item based on the digital content item's labeling prediction output.

2. The method of claim 1, comprising:

identifying, via the computing device, a set of digital content items using the digital content item's annotation.

3. The method of claim 2, each digital content item in the set of digital content items having a respective annotation generated using label prediction output from the trained multi-level label prediction model.

4. The method of claim 2, identifying the set of digital content items is responsive to a received search request.

5. The method of claim 4, the received search request comprising information indicating the digital content item.

6. The method of claim 1, the training data comprising a plurality of training instances corresponding to a plurality of digital content items, labeling data corresponding to at least one training instance is missing at least one applicable label from the plurality of labels.

7. The method of claim 6, a training instance, of the plurality of training instances, comprising associated labeling data and feature data.

8. The method of claim 7, further comprising:

analyzing, a digital content item of the plurality of digital content items, and based on the analysis, generating the feature data for the analyzed digital content item.

9. The method of claim 1, the number of levels of the multi-level label prediction model is empirically determined, a final level of the multi-level label prediction model being identified based on a determined convergence in the label prediction output generated by the final level and a potential next level.

10. The method of claim 1, the digital content item is a document and the annotation comprises a number of words contained in the document.

11. The method of claim 1, the digital content item comprises an image.

12. A computer readable non-transitory storage medium tangibly encoded with computer-executable instructions that when executed by a processor associated with a computing device perform a method comprising:

training a multi-level label prediction model to predict an applicability, to a digital content item, of each label of a plurality of labels, the multi-level label prediction model comprising a first level and a second level, the first level being trained using training data missing a number of applicable labels, label prediction output of the trained first level including an applicability prediction for each of the number of applicable labels missing from the training data, the second level being trained using the label prediction output generated by the trained first level;

generating the digital content item's label prediction output using the trained multi-level label prediction model, the digital content item's labeling prediction output indicating an applicability, to the digital content item, of each label of the plurality of labels; and

automatically generating annotation for the digital content item based on the digital content item's labeling prediction output.

13. The computer readable non-transitory storage medium of claim 12, the method further comprising:

identifying a set of digital content items using the digital content item's annotation.

14. The computer readable non-transitory storage medium of claim 13, each digital content item in the set of digital

content items having a respective annotation generated using label prediction output from the trained multi-level label prediction model.

15. The computer readable non-transitory storage medium of claim 13, identifying the set of digital content items is responsive to a received search request.

16. The computer readable non-transitory storage medium of claim 15, the received search request comprising information indicating the digital content item.

17. The computer readable non-transitory storage medium of claim 12, the training data comprising a plurality of training instances corresponding to a plurality of digital content items, labeling data corresponding to at least one training instance is missing at least one applicable label from the plurality of labels.

18. The computer readable non-transitory storage medium of claim 17, a training instance, of the plurality of training instances, comprising associated labeling data and feature data.

19. The computer readable non-transitory storage medium of claim 18, further comprising:

analyzing, a digital content item of the plurality of digital content items, and based on the analysis, generating the feature data for the analyzed digital content item.

20. A system comprising:

a processor;

a storage medium for tangibly storing thereon program logic for execution by the processor, the stored logic comprising:

training logic executed by the processor for training a multi-level label prediction model to predict an applicability, to a digital content item, of each label of a plurality of labels, the multi-level label prediction model comprising a first level and a second level, the first level being trained using training data missing a number of applicable labels, label prediction output of the trained first level including an applicability prediction for each of the number of applicable labels missing from the training data, the second level being trained using the label prediction output generated by the trained first level;

generating logic executed by the processor for generating the digital content item's label prediction output using the trained multi-level label prediction model, the digital content item's labeling prediction output indicating an applicability, to the digital content item, of each label of the plurality of labels; and

generating logic executed by the processor for automatically generating annotation for the digital content item based on the digital content item's labeling prediction output.

* * * * *