

IDENTIFYING OPTIMAL WEIGHTS TO IMPROVE PREDICTION ACCURACY IN MACHINE LEARNING TECHNIQUES

Technical Field

[0001] The present invention relates generally to predictive modeling, and more particularly to identifying optimal weights to improve prediction accuracy in machine learning techniques.

BACKGROUND

[0002] Predictive modeling uses statistics to predict outcomes. Most often the event one wants to predict is in the future, but predictive modelling can be applied to any type of unknown event, regardless of when it occurred. For example, predictive models are often used to detect events and identify persons related to the events, after the events have taken place.

[0003] In many cases, the model is chosen on the basis of detection theory to try to guess the probability of an outcome given a set amount of input data, for example, given an email determining how likely that it is spam.

[0004] Models can use one or more classifiers in trying to determine the probability of a set of data belonging to another set. For example, a model might be used to determine whether an email is spam or "ham" (non-spam).

[0005] Depending on definitional boundaries, predictive modelling is synonymous with, or largely overlapping with, the field of machine learning, as it is more commonly referred to in academic or research and development contexts. When deployed commercially, predictive modelling is often referred to as predictive analytics.

[0006] Achieving better predictive models is an objective in the research and practice of machine learning techniques. For example, ensemble methods use multiple learning algorithms to obtain better predictive performance than could be obtained from any of the constituent learning algorithms alone. Such ensemble methods include bootstrap aggregating (also called bagging), boosting, etc.

[0007] Bootstrap aggregating is a machine learning ensemble meta-algorithm designed to improve the stability and accuracy of machine learning algorithms used in statistical classification and regression. Boosting is a machine learning ensemble meta-algorithm for primarily reducing bias and also variance in supervised learning and a family of machine learning algorithms that convert weak learners to strong ones.

[0008] In such techniques, such as boosting, the weights of wrongly classified cases are increased while the weights of correctly classified cases are decreased during the modeling process. Such a strategy (heuristic) does

achieve better predictions in many cases; however, overfitting outliers/noises is a possibility. As a result of overfitting outliers/noises, the predictive accuracy is lessened.

[0009] Hence, the heuristic strategy of increasing the weights of wrongly classified cases and decreasing the weights of correctly classified cases may not be the best strategy for improving the prediction accuracy of the model.

[0010] For example, sometimes, it may be better to increase the weights of correctly classified cases because such cases contain very important patterns which should be learned by the machine learning algorithm. It may also be better to decrease the weights of wrongly classified cases, such as outlier cases, for similar reasons.

[0011] Consequently, such techniques, such as boosting, fail to identify the optimal weights for the classified cases, and therefore, fail to achieve optimal prediction accuracy in machine learning techniques.

SUMMARY

[0012] In one embodiment of the present invention, a computer-implemented method for improving prediction accuracy in machine learning techniques comprises constructing a teacher model, where the teacher model generates a weight for each data case. The method further comprises training a current student model using training data and weights generated by the teacher model. The method additionally comprises generating state features by the current student model. Furthermore, the method comprises generating new weights by the teacher model using the state features. Additionally, the method comprises training a candidate student model using the training data and the new weights. In addition, the method comprises generating a reward by comparing the current student model with the candidate student model using the training data and testing data to determine which is better at predicting an observed target. The method further comprises updating the teacher model with the reward in response to a stopping rule not being satisfied. The method additionally comprises returning the new weights and the current student model to a user in response to the stopping rule being satisfied, where the returned student model provides a prediction of the observed target.

[0013] In another embodiment of the present invention, a computer program product for improving prediction accuracy in machine learning techniques, the computer program product comprising a computer readable storage medium having program code embodied therewith, the program code comprises the programming instructions for constructing a teacher model, where the teacher model generates a weight for each data case. The program code further comprises the programming instructions for training a current student model using training data and weights generated by the teacher model. The program code additionally comprises the programming instructions for generating state features by the current student model. Furthermore, the program code comprises the programming instructions for generating new weights by the teacher model using the state features. Additionally,

the program code comprises the programming instructions for training a candidate student model using the training data and the new weights. In addition, the program code comprises the programming instructions for generating a reward by comparing the current student model with the candidate student model using the training data and testing data to determine which is better at predicting an observed target. The program code further comprises the programming instructions for updating the teacher model with the reward in response to a stopping rule not being satisfied. The program code additionally comprises the programming instructions for returning the new weights and the current student model to a user in response to the stopping rule being satisfied, where the returned student model provides a prediction of the observed target.

[0014] In a further embodiment of the present invention, a system comprises a memory for storing a computer program for improving prediction accuracy in machine learning techniques and a processor connected to the memory, where the processor is configured to execute the program instructions of the computer program comprising constructing a teacher model, where the teacher model generates a weight for each data case. The program instructions of the computer program further comprise training a current student model using training data and weights generated by the teacher model. The program instructions of the computer program additionally comprise generating state features by the current student model. Furthermore, the program instructions of the computer program comprise generating new weights by the teacher model using the state features. Additionally, the program instructions of the computer program comprise training a candidate student model using the training data and the new weights. In addition, the program instructions of the computer program comprise generating a reward by comparing the current student model with the candidate student model using the training data and testing data to determine which is better at predicting an observed target. The program instructions of the computer program further comprise updating the teacher model with the reward in response to a stopping rule not being satisfied. The program instructions of the computer program additionally comprise returning the new weights and the current student model to a user in response to the stopping rule being satisfied, where the returned student model provides a prediction of the observed target.

[0015] In this manner, the present invention devises a framework which implements the concept of "learning to teach" in the field of predictive modeling. Such a framework includes a teacher model, which generates a weight for each data case. The training data cases, along with the generated weights, are used to re-train the student model. A reward is returned by evaluating the trained student model on a held-out dataset (testing data) in terms of prediction accuracy. The teacher model then utilizes the reward to update its parameters via policy gradient methods, e.g., reinforcement learning. Such a process will be repeated until the student model achieves desired performance.

[0016] In comparison to previously used heuristic methods (e.g., boosting), the approach of the present invention determines case weights in an optimal way. This allows one to build a better student model via basic learners, e.g., decision tree, neural network, etc., rather than using an ensemble model.

[0017] By using case weights as actions on the student model, any type of machine learner may be used as the student model given that the learner supports case weights in training.

[0018] Furthermore, the distributions of data cases in the training data can be corrected by the generated weights in the event that the training data comes from a biased sampling.

[0019] The foregoing has outlined rather generally the features and technical advantages of one or more embodiments of the present invention in order that the detailed description of the present invention that follows may be better understood. Additional features and advantages of the present invention will be described hereinafter which may form the subject of the claims of the present invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] A better understanding of the present invention can be obtained when the following detailed description is considered in conjunction with the following drawings, in which:

Figure 1 illustrates a communication system for practicing the principles of the present invention in accordance with an embodiment of the present invention;

Figure 2 illustrates an embodiment of the present invention of the hardware configuration of a predictive analytics system which is representative of a hardware environment for practicing the present invention;

Figure 3 is a diagram of the software components of the predictive analytics system used to identify the optimal weights to be used by the predictive model for generating predictions in accordance with an embodiment of the present invention;

Figure 4 illustrates the state features generated by the student model in accordance with an embodiment of the present invention;

Figure 5 illustrates the rewards generated by the reward generator in accordance with an embodiment of the present invention; and

Figure 6 is a flowchart of a method for improving prediction accuracy in machine learning techniques in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0021] The present invention comprises a computer-implemented method, system and computer program product for improving prediction accuracy in machine learning techniques. In one embodiment of the present invention, a teacher model is constructed, where the teacher model generates a weight for each data case. A "teacher model," as used herein, refers to a statistical model that determines the appropriate data, loss function and hypothesis space to facilitate the learning of the student model. The current student model is then trained using training data and the weights generated by the teacher model. A "student model," as used herein, refers to a

statistical model that is trained to provide a prediction using training data. A "current" student model, as used herein, refers to a student model currently being trained to provide a prediction using training data. The current student model generates state features (e.g., data features, case weights, student model features and features to represent the combination of both the data and the student model), which are used by the teacher model to generate new weights. A candidate student model is then trained using training data and these new weights. A "candidate student model," as used herein, refers to a student model that is being examined to determine if is a better student model (better at predicting the observed target) than the current student model. A reward is then generated by comparing the current student model with the candidate student model using training and testing data to determine which is better at predicting an observed target. A "reward," as used herein, refers to a value generated by a function (reward function) used in reinforcement learning. A positive reward may be returned if the candidate student model is better at predicting the observed target than the current student model. Conversely, a negative reward may be returned if the current student model is better at predicting the observed target than the candidate student model. The teacher model is then updated with the reward. The teacher model utilizes the rewards to update its parameters via policy gradient methods, such as reinforcement learning. If the candidate student model is better at predicting the observed target than the current student model, then the current student model is updated with the candidate student model and the current weights are updated with the new weights generated by the teacher model. Upon updating the current weights with the new weights, the current student model generates new state features. If, however, the candidate student model is not better at predicting the observed target than the current student model, then the updated teacher model generates new weights using the current weights and the current student features from the current student model. Upon any of the stopping rules being satisfied (e.g., reaching a specified number of trials, reaching a specified training timing, converging of a prediction accuracy and a user-initiated termination), the weights generated by the teacher model are deemed to be the "optimal" weights which are returned to the user along with the corresponding student model. In this manner, optimal weights are identified to improve prediction accuracy.

[0022] In the following description, numerous specific details are set forth to provide a thorough understanding of the present invention. However, it will be apparent to those skilled in the art that the present invention may be practiced without such specific details. In other instances, well-known circuits have been shown in block diagram form in order not to obscure the present invention in unnecessary detail. For the most part, details considering timing considerations and the like have been omitted inasmuch as such details are not necessary to obtain a complete understanding of the present invention and are within the skills of persons of ordinary skill in the relevant art.

[0023] Referring now to the Figures in detail, Figure 1 illustrates a communication system 100 for making predictions using machine learning techniques. In one embodiment, system 100 includes a predictive analytics system 101 for generating predictions 102 using data, such as training data 103 and testing data 104 (also referred

to herein as "hold-out data"). A further description of predictive analytics system 101 using training and testing data 103, 104 to make predictions is discussed further below in connection with Figures 3-6.

[0024] In one embodiment, predictive analytics system 101 makes predictions about unknown future events using many techniques from data mining, statistics, modeling, machine learning and artificial intelligence to analyze current data to make predictions about the future.

[0025] In one embodiment, predictive analytics system 101 utilizes the concept of "learning to teach," which involves two intelligent agents, namely, a teacher model and a student model. The training phase contains several episodes of sequential interactions between the teacher model and the student model. Based on the state information generated by the student model, the teacher model updates its teaching actions so as to refine the machine learning problem of the student model. The student model then performs its learning process based on the inputs from the teacher model and provides reward signals (e.g., the accuracy on the training data) back to the teacher model afterwards. The teacher model then utilizes such rewards to update its parameters via policy gradient methods, which are a type of a reinforcement learning technique. This interactive process is end-to-end trainable, exempt from the limitations of human-defined heuristics. In one embodiment, the concept of "learning to teach" is implemented by devising an approach of leveraging the weights of cases (data cases) as actions for the student model. A description of the hardware configuration of predictive analytics system 101 is provided below in connection with Figure 2.

[0026] Referring now to Figure 2, Figure 2 illustrates an embodiment of the present invention of the hardware configuration of a predictive analytics system 101 (Figure 1) which is representative of a hardware environment for practicing the present invention. Predictive analytics system 101 may be any type of analytics system (e.g., portable computing unit, Personal Digital Assistant (PDA), laptop computer, mobile device, tablet personal computer, smartphone, mobile phone, navigation device, gaming unit, desktop computer system, workstation, Internet appliance and the like) configured with the capability of identifying optimal weights to improve prediction accuracy in machine learning techniques.

[0027] Referring to Figure 2, predictive analytics system 101 may have a processor 201 connected to various other components by system bus 202. An operating system 203 may run on processor 201 and provide control and coordinate the functions of the various components of Figure 2. An application 204 in accordance with the principles of the present invention may run in conjunction with operating system 203 and provide calls to operating system 203 where the calls implement the various functions or services to be performed by application 204. Application 204 may include, for example, a program for identifying optimal weights to improve prediction accuracy in machine learning techniques as discussed below in connection with Figures 3-6.

[0028] Referring again to Figure 2, read-only memory ("ROM") 205 may be connected to system bus 202 and include a basic input/output system ("BIOS") that controls certain basic functions of predictive analytics system 101. Random access memory ("RAM") 206 and disk adapter 207 may also be connected to system bus 202. It should be noted that software components including operating system 203 and application 204 may be loaded into RAM 206, which may be predictive analytics system's 101 main memory for execution. Disk adapter 207 may be an integrated drive electronics ("IDE") adapter that communicates with a disk unit 208, e.g., disk drive. It is noted that the program for identifying optimal weights to improve prediction accuracy in machine learning techniques, as discussed below in connection with Figures 3-6, may reside in disk unit 208 or in application 204.

[0029] Predictive analytics system 101 may further include a communications adapter 209 connected to bus 202. Communications adapter 209 may interconnect bus 202 with an outside network thereby allowing predictive analytics system 101 to communicate with other devices.

[0030] I/O devices may also be connected to predictive analytics system 101 via a user interface adapter 210 and a display adapter 211. Keyboard 212, mouse 213 and speaker 214 may all be interconnected to bus 202 through user interface adapter 210. A display monitor 215 may be connected to system bus 202 by display adapter 211. In this manner, a user is capable of inputting to predictive analytics system 101 through keyboard 212 or mouse 213 and receiving output from predictive analytics system 101 via display 215 or speaker 214. Other input mechanisms may be used to input data to predictive analytics system 101 that are not shown in Figure 2, such as display 215 having touch-screen capability and keyboard 212 being a virtual keyboard. Predictive analytics system 101 of Figure 2 is not to be limited in scope to the elements depicted in Figure 2 and may include fewer or additional elements than depicted in Figure 2.

[0031] The present invention may be a system, a method, and/or a computer program product at any possible technical detail level of integration. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

[0032] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded

thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals *per se*, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[0033] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[0034] Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, configuration data for integrated circuitry, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++, or the like, and procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[0035] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0036] These computer readable program instructions may be provided to a processor of a computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0037] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0038] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the blocks may occur out of the order noted in the Figures. For example, two blocks shown in succession may, in fact, be accomplished as one step, executed concurrently, substantially concurrently, in a partially or wholly temporally overlapping manner, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

[0039] As stated in the Background section, achieving better predictive models is an objective in the research and practice of machine learning techniques. For example, ensemble methods use multiple learning algorithms to obtain better predictive performance than could be obtained from any of the constituent learning algorithms alone. Such ensemble methods include bootstrap aggregating (also called bagging), boosting, etc. Bootstrap aggregating is a machine learning ensemble meta-algorithm designed to improve the stability and accuracy of machine learning algorithms used in statistical classification and regression. Boosting is a machine learning ensemble meta-algorithm for primarily reducing bias and also variance in supervised learning and a family of machine learning algorithms that convert weak learners to strong ones. In such techniques, such as boosting, the weights of wrongly

classified cases are increased while the weights of correctly classified cases are decreased during the modeling process. Such a strategy (heuristic) does achieve better predictions in many cases; however, overfitting outliers/noises is a possibility. As a result of overfitting outliers/noises, the predictive accuracy is lessened. Hence, the heuristic strategy of increasing the weights of wrongly classified cases and decreasing the weights of correctly classified cases may not be the best strategy for improving the prediction accuracy of the model. For example, sometimes, it may be better to increase the weights of correctly classified cases because such cases contain very important patterns which should be learned by the machine learning algorithm. It may also be better to decrease the weights of wrongly classified cases, such as outlier cases, for similar reasons. Consequently, such techniques, such as boosting, fail to identify the optimal weights for the classified cases, and therefore, fail to achieve optimal prediction accuracy in machine learning techniques.

[0040] The embodiments of the present invention provide a means for achieving optimal prediction accuracy in machine learning techniques by identifying the optimal weights using the concept of learning to teach involving two intelligent agents (a teacher model and a student model) as discussed below in connection with Figures 3-6. Figure 3 is a diagram of the software components of predictive analytics system 101 used to identify the optimal weights to be used by the predictive model for generating predictions. Figure 4 illustrates the state features generated by the student model. Figure 5 illustrates the rewards generated by the reward generator. Figure 6 is a flowchart of a method for improving prediction accuracy in machine learning techniques.

[0041] Given a training data $\{<X_k, f_k, y_k> | k=1, \dots, N\}$, where X_k is a vector of predictors in case k , y_k is the observed target, and f_k is an optional case weight (let $f_k=1$ if no case weight exists), the present invention will generate optimal case weights f_k^{opt} for each case k . With the optimal case weights, a student model will be trained to provide accurate predictions using the held-out dataset as discussed below.

[0042] As stated above, Figure 3 is a diagram of the software components of predictive analytics system 101 used to identify the optimal weights to be used by the predictive model for generating predictions in accordance with an embodiment of the present invention. In one embodiment, these software components may reside in application 204 (Figure 2) of predictive analytics system 101.

[0043] The following provides a brief description of these software components. A more detailed description of these software components (including their functionalities) is provided below in conjunction with Figures 4-6.

[0044] Referring to Figure 3, in conjunction with Figures 1-2, predictive analytics system 101 includes a module referred to herein as the "teacher model" 301. In one embodiment, teacher model 301 is a neural network configured to receive state features as inputs and generate a weight for each data case. In one embodiment, the weight parameters of the neural network are initialized randomly before the training process.

[0045] Predictive analytics system 101 further includes a module referred to herein as the "current student model 302," which receives the weights generated by teacher model 301 and generates state features, such as those shown in Figure 4. A "student model," as used herein, refers to a statistical model that is trained to provide a prediction using training data 103. A "current" student model 302, as used herein, refers to a student model currently being trained by predictive analytics system 101 to provide a prediction using training data 103. In one embodiment, student model 302 corresponds to a learner, such as a decision tree or a neural network. In one embodiment, decision tree learning uses a decision tree as a predictive model to go from observations about an item (represented in the branches) to conclusions about the items' target value (represented in the leaves). In one embodiment, a neural network is a network or circuit of neurons (artificial neurons) or nodes.

[0046] Figure 4 illustrates the state features generated by student model 302 in accordance with an embodiment of the present invention.

[0047] Referring to Figure 4, such state features may include predictors 401, their corresponding weight 402 and their predicted value 403. In one embodiment, predictor 401 corresponds to the outcome variable, such as the observed target. Weight 402 corresponds to the weight assigned to such a predictor obtained from teacher model 301. In one embodiment, such weights 402 indicate the confidence that the corresponding predicted value 403 of predictor 401 is accurate. The higher the value, the greater confidence in the corresponding predicted value 403 being correct.

[0048] In one embodiment, such state features may also include data features, case weights, student model features and features to represent the combination of both the data and the student model.

[0049] In one embodiment, such state features may be utilized by teaching model 301 to generate new weights which are used to train a candidate student model 303. A "candidate student model 303," as used herein, refers to a student model that is being examined to determine if is a better student model (better at predicting the observed target) than the current student model.

[0050] In one embodiment, candidate student model 303 is trained by using training data 103 and the new weights generated by teacher model 301.

[0051] As illustrated in Figure 3, a module, referred to herein as the "reward generator 304," generates rewards by comparing the current and candidate student models 302, 303 using training data 103 and testing data 104 ("held-out data"). In one embodiment, teacher model 301 is updated with the rewards. In one embodiment, teacher model 301 utilizes the rewards to update its parameters via policy gradient methods, such as reinforcement learning.

[0052] A "reward," as used herein, refers to a value generated by a function (reward function) used in reinforcement learning. The goal of a reinforcement learning agent (predictive analytics system 101) is to collect as much reward as possible. In one embodiment, a positive reward is returned by reward generator 304 if the candidate student model 303 is better at predicting the observed target than the current student model 302. Conversely, a negative reward is returned by reward generator 304 if the current student model 302 is better at predicting the observed target than the candidate student model 303. In one embodiment, reward generator 304 is part of candidate student model 303.

[0053] In one embodiment, such rewards are generated by reward generator 304 by applying training data 103, testing data 104 to student models 302, 303 as shown in Figure 5.

[0054] Figure 5 illustrates the rewards generated by reward generator 304 in accordance with an embodiment of the present invention.

[0055] Referring to Figure 5, reward generator 304 generates a model level reward 501 and a case level reward 502. In one embodiment, model level reward 501 refers to the reward associated with the student models 302, 303 generating a prediction for the observed target based on the testing data 104. For instance, reward 501 is generated based on how much better candidate student model 303 is at predicting the observed target than the current student model 302 using testing data 104. A positive reward 501 is returned by reward generator 304 if the candidate student model 303 is better at predicting the observed target than the current student model 302 using testing data 104. Conversely, a negative reward 501 is returned by reward generator 304 if the current student model 302 is better at predicting the observed target than the candidate student model 303 using testing data 104.

[0056] Case level reward 502 refers to the reward based on correctly classifying the data case by student models 302, 303 using training data 103. If student model 302, 303 correctly classified the data case, then a positive reward 502 is returned by reward generator 304. Conversely, a negative reward 502 is returned by reward generator 304 if student model 302, 303 did not correctly classify the data case.

[0057] In one embodiment, reward generator 304 generates a final reward 503 that is a combination of model level reward 501 and case level reward 502, such as an average of the two rewards.

[0058] Returning to Figure 3, utilizing the rewards, a module, referred to herein as the "updater 305," updates teacher model 301 with the reward. In one embodiment, teacher model 301 faces an associative immediate-reinforcement learning task. Suppose that the reward for the reinforcement value is r at each trial, then the parameter w_{ij} in the network is incremented by an amount $\Delta w_{ij} = \sum_{k=1}^N [\alpha (\partial \ln g_k) / (\partial w_{ij}) r]$, where α is a learning rate factor, N is a positive integer number, and g_k is the output of the teacher model for case k .

[0059] Once teacher model 301 is updated, a determination is made by decision maker 306 as to whether candidate student model 303 is a better predictor of the observed target than the current student model 302 or not. In one embodiment, if candidate student model 303 is better at predicting the observed target than current student model 302, then decision maker 306 will update current student model 302 with candidate student model 303 as well as update the current weights with the new weights. The updated student model 302 will then generate new state features which are inputted to teacher model 301.

[0060] Alternatively, if candidate student model 303 is not better at predicting the observed target than current student model 302, then decision maker 306 directly requests the updated teacher model 301 (updated with the rewards as discussed above) to generate new weights using the current student features from the current student model 302.

[0061] A more detailed discussion regarding the process of improving prediction accuracy in machine learning techniques using the software components discussed above is provided below in connection with Figure 6.

[0062] Figure 6 is a flowchart of a method 600 for improving prediction accuracy in machine learning techniques in accordance with an embodiment of the present invention.

[0063] Referring to Figure 6, in conjunction with Figures 1-5, in step 601, predictive analytics system 101 constructs a teacher model 301 which generates a weight for each data case. A "teacher model," as used herein, refers to a statistical model that determines the appropriate data, loss function and hypothesis space to facilitate the learning of the student model. In one embodiment, teacher model 301 is a neural network. In one embodiment, teacher model 301 receives state features from current student model 302 as inputs and generates a weight for each data case. A "data case," as used herein, refers to data used to predict an observed target. In one embodiment, the weight parameters are initialized randomly before the training process begins (discussed further below).

[0064] In one embodiment, teacher model 301 includes networks composed of several layers. In one embodiment, the layers are made of nodes, where a node is a place where computation happens, loosely patterned on a neuron in the human brain, which fires when it encounters sufficient stimuli. A node combines input from the data, such as state features (discussed further below) from current student model 302 with a set of coefficients, or weights, that either amplify or dampen that input, thereby assigning significance to inputs with regard to the task the algorithm is trying to learn (e.g., which input is most helpful in classifying data without error). These input-weight products are summed and then the sum is passed through a node's so-called activation function, to determine whether and to what extent that signal should progress further through the network to affect the ultimate outcome (e.g., an act of classification). If the signals pass through, the neuron has been "activated."

[0065] In one embodiment, a node layer is a row of neuron-like switches that turn on or off as the input is fed through the net. Each layer's output is simultaneously the subsequent layer's input, starting from an initial input layer receiving the data.

[0066] In one embodiment, the model's adjustable weights are paired with the input features so as to assign significance to those features with regard to how the neural network classifies and clusters input.

[0067] In one embodiment, such generated weights can be used to correct the distribution of data cases in training data 103 in the event that training data 103 comes from a biased sampling.

[0068] In step 602, predictive analytics system 101 trains current student model 302 using training data 103 and weights (current weight for each data case) generated by teacher model 301. As previously discussed, current student model 302 refers to a student model (statistical model that is trained to provide a prediction, such as predicting the observed target, using training data) that is currently being trained by predictive analytics system 101 to provide a prediction using training data 103. As also previously discussed, in one embodiment, current student model 302 corresponds to a learner, such as a decision tree or a neural network. In one embodiment, decision tree learning uses a decision tree as a predictive model to go from observations about an item (represented in the branches) to conclusions about the items' target value (represented in the leaves). In one embodiment, a neural network is a network or circuit of neurons (artificial neurons) or nodes.

[0069] In one embodiment, case weights are used as actions on current student model 302 (as well as on candidate student model 303 discussed further below). Such a method allows the usage of any type of machine learner as the student model (student models 302, 303) given that the learner supports case weight in training.

[0070] In one embodiment, such training involves fitting current student model 302 on a training dataset 103, that is a set of examples used to fit the parameters, such as the weights generated by teacher model 301. In one embodiment, current student model 302 is trained using a supervised learning method (e.g., gradient descent or stochastic gradient descent). In one embodiment, training dataset 103 consists of pairs of input vectors (or scalar) and the corresponding output vector (or scalar), which may be denoted as the target. Current student model 302 is run with training dataset 103 and produces a result, which is then compared with the target, for each input vector in training dataset 103. Based on the result of the comparison and the specific learning algorithm being used, the parameters of student model 302 are adjusted.

[0071] In step 603, after training current student model 302, the trained current student model 302 generates state features. In one embodiment, state features are defined for each data case based on current student model 302. State features may include, but not limited to, data features containing information for a data case, such as its predictors, target, etc.; case weight generated by teacher model 301; student model features, including the

measures reflecting how well current student model 302 is trained; and features to represent the combination of both data and student model 302, such as predicted targets, probabilities of each target category, etc.

[0072] In one embodiment, case-level state features are also generated by student model 302, including the corresponding case weight.

[0073] In one embodiment, current student model 302 generates data that includes the previously discussed state features. Such features are generated using a set of statistical assumptions based on the receiving training data 103 and weights from teacher model 301.

[0074] In step 604, teacher model 301 generates new weights using the state features generated by current student model 302. In one embodiment, such state features are input to teacher model 301 which are used to generate weights by teacher model 301 as discussed above.

[0075] In step 605, predictive analytics system 101 trains a candidate student model 303 using training data 103 and the new weights generated by teacher model 301. In one embodiment, training candidate student model 303 is performed in the same manner as training current student model 302 as discussed above in connection with step 602.

[0076] In step 606, reward generator 304 generates a reward by comparing the current and candidate student models 302, 303 (after training student models 302, 303) using training data 103 and testing data 104 ("held-out data") to determine which is better at predicting an observed target. A "reward," as used herein, refers to a value generated by a function (reward function) used in reinforcement learning. The goal of a reinforcement learning agent (predictive analytics system 101) is to collect as much reward as possible. In one embodiment, a positive model level reward 501 is returned by reward generator 304 if the candidate student model 303 is better at predicting the observed target than the current student model 302 using testing data 104. Conversely, a negative model level reward 501 is returned by reward generator 304 if the current student model 302 is better at predicting the observed target than the candidate student model 303 using testing data 104.

[0077] Case level reward 502 refers to the reward based on correctly classifying the data case by student models 302, 303 using training data 103. If student model 302, 303 correctly classified the data case, then a positive reward 502 is returned by reward generator 304. Conversely, a negative reward 502 is returned by reward generator 304 if student model 302, 303 did not correctly classify the data case.

[0078] In one embodiment, reward generator 304 generates a final reward 503 that is a combination of model level reward 501 and case level reward 502, such an average of the two rewards.

[0079] In step 607, a determination is made by predictive analytics system 101 as to whether a stopping rule has been satisfied. "Stopping rules," as used herein, refer to the rules that determine whether the training of a student model has been completed. Such training is completed when it has been determined that the optimal weights for the data cases have been identified.

[0080] In one embodiment, such stopping rules include, but not limited to, the following: reaching a specified number of trials, reaching a specified training timing, converging of a prediction accuracy and a user-initiated termination.

[0081] If any of these stopping rules have been satisfied, then, in step 608, the optimal weights (weights generated by teacher model 301) and the corresponding student model 302 are returned to the user, such as via a user interface on predictive analytics system 101. The returned student model is able to provide an accurate prediction of the observed target. Furthermore, the returned weights, which are generated by teacher model 301 (see step 604), are deemed to be the "optimal" weights, such as when a stopping rule has been satisfied (e.g., when the prediction accuracy converges). In this manner, optimal weights are identified to improve prediction accuracy.

[0082] If, however, none of the stopping rules have been satisfied, then, in step 609, updater 305 updates teacher model 301 with the reward (reward of step 606). In one embodiment, teacher model 301 utilizes the rewards to update its parameters via policy gradient methods, such as reinforcement learning.

[0083] As previously discussed, in one embodiment, teacher model 301 faces an associative immediate-reinforcement learning task. Suppose that the reward for the reinforcement value is r at each trial, then the parameter w_{ij} in the network is incremented by an amount $\Delta w_{ij} = \sum_{k=1}^N [\alpha (\partial \ln g_k) / (\partial w_{ij}) r]$, where α is a learning rate factor, N is a positive integer number, and g_k is the output of the teacher model for case k . The incremented amount Δw upon which teacher model 301 will be updated may also be shown as follows:

$$\Delta w_{ij} = \alpha r^{model} \sum_{k=1}^N \frac{\partial \ln g_k}{\partial w_{ij}} r_k^{case}$$

where r is the reinforcement value at each trial, w_{ij} is a parameter in the network incremented by an amount Δw_{ij} , α is a learning rate factor, N is a positive integer number, and g_k is the output of the teacher model for case k .

[0084] Upon updating teacher model 301, in step 610, a determination is made by decision maker 306 as to whether candidate student model 303 generates a better prediction of the observed target than current student model 302. Such a determination is based on how close the prediction is to the observed target.

[0085] If candidate student model 303 is better at predicting the observed target than current student model 302, then, in step 611, current student model 302 is updated with candidate student model 303. That is, in step 611, current student model 302 is essentially replaced with candidate student model 303.

[0086] Furthermore, if candidate student model 303 is better at predicting the observed target than current student model 302, then in step 612, the current weights are updated with the new weights (new weights generated by teacher model 301 in step 604).

[0087] Upon updating the current weights with the new weights, the updated student model 302 (updated in step 611) generates state features in step 603 using the new weights.

[0088] Alternatively, if candidate student model 303 is not better at predicting the observed target than current student model 302, then decision maker 306 directly requests the updated teacher model 301 (updated in step 609) to generate new weights in step 604 using the current student features from the current student model 302 and current weights.

[0089] In this manner, the present invention devises a framework which implements the concept of "learning to teach" in the field of predictive modeling. Such a framework includes a teacher model, which generates a weight for each data case. The training data cases, along with the generated weights, are used to re-train the student model. A reward is returned by evaluating the trained student model on a held-out dataset in terms of prediction accuracy. The teacher model then utilizes the reward to update its parameters via policy gradient methods, e.g., reinforcement learning. Such a process will be repeated until the student model achieves desired performance.

[0090] In comparison to previously used heuristic methods (e.g., boosting), the approach of the present invention determines case weights in an optimal way. This allows one to build a better student model via basic learners, e.g., decision tree, neural network, etc., rather than using an ensemble model.

[0091] By using case weights as actions on the student model, any type of machine learner may be used as the student model given that the learner supports case weight in training.

[0092] Furthermore, the present invention improves the technology or technical field involving predictive modeling. As discussed above, achieving better predictive models is an objective in the research and practice of machine learning techniques. For example, ensemble methods use multiple learning algorithms to obtain better

predictive performance than could be obtained from any of the constituent learning algorithms alone. Such ensemble methods include bootstrap aggregating (also called bagging), boosting, etc. Bootstrap aggregating is a machine learning ensemble meta-algorithm designed to improve the stability and accuracy of machine learning algorithms used in statistical classification and regression. Boosting is a machine learning ensemble meta-algorithm for primarily reducing bias and also variance in supervised learning and a family of machine learning algorithms that convert weak learners to strong ones. In such techniques, such as boosting, the weights of wrongly classified cases are increased while the weights of correctly classified cases are decreased during the modeling process. Such a strategy (heuristic) does achieve better predictions in many cases; however, overfitting outliers/noises is a possibility. As a result of overfitting outliers/noises, the predictive accuracy is lessened. Hence, the heuristic strategy of increasing the weights of wrongly classified cases and decreasing the weights of correctly classified cases may not be the best strategy for improving the prediction accuracy of the model. For example, sometimes, it may be better to increase the weights of correctly classified cases because such cases contain very important patterns which should be learned by the machine learning algorithm. It may also be better to decrease the weights of wrongly classified cases, such as outlier cases, for similar reasons. Consequently, such techniques, such as boosting, fail to identify the optimal weights for the classified cases, and therefore, fail to achieve optimal prediction accuracy in machine learning techniques.

[0093] The present invention improves such technology by constructing a teacher model, where the teacher model generates a weight for each data case. A "teacher model," as used herein, refers to a statistical model that determines the appropriate data, loss function and hypothesis space to facilitate the learning of the student model. The current student model is then trained using training data and the weights generated by the teacher model. A "student model," as used herein, refers to a statistical model that is trained to provide a prediction using training data. A "current" student model, as used herein, refers to a student model currently being trained to provide a prediction using training data. The current student model generates state features (e.g., data features, case weights, student model features and features to represent the combination of both the data and the student model), which are used by the teacher model to generate new weights. A candidate student model is then trained using training data and these new weights. A "candidate student model," as used herein, refers to a student model that is being examined to determine if is a better student model (better at predicting the observed target) than the current student model. A reward is then generated by comparing the current student model with the candidate student model using training and testing data to determine which is better at predicting an observed target. A "reward," as used herein, refers to a value generated by a function (reward function) used in reinforcement learning. A positive reward may be returned if the candidate student model is better at predicting the observed target than the current student model. Conversely, a negative reward may be returned if the current student model is better at predicting the observed target than the candidate student model. The teacher model is then updated with the reward. The teacher model utilizes the rewards to update its parameters via policy gradient methods, such as reinforcement learning. If the candidate student model is better at predicting the observed target than the current student model, then the current student model is updated with the candidate student model and the current weights are updated

with the new weights generated by the teacher model. Upon updating the current weights with the new weights, the current student model generates new state features. If, however, the candidate student model is not better at predicting the observed target than the current student model, then the updated teacher model generates new weights using the current weights and the current student features from the current student model. Upon any of the stopping rules being satisfied (e.g., reaching a specified number of trials, reaching a specified training timing, converging of a prediction accuracy and a user-initiated termination), the weights generated by the teacher model are deemed to be the "optimal" weights which are returned to the user along with the corresponding student model. In this manner, optimal weights are identified to improve prediction accuracy. Furthermore, in this manner, there is an improvement in the technical field of predictive modeling.

[0094] The technical solution provided by the present invention cannot be performed in the human mind or by a human using a pen and paper. That is, the technical solution provided by the present invention could not be accomplished in the human mind or by a human using a pen and paper in any reasonable amount of time and with any reasonable expectation of accuracy without the use of a computer.

[0095] In one embodiment of the present invention, a computer-implemented method for improving prediction accuracy in machine learning techniques comprises constructing a teacher model, where the teacher model generates a weight for each data case. The method further comprises training a current student model using training data and weights generated by the teacher model. The method additionally comprises generating state features by the current student model. Furthermore, the method comprises generating new weights by the teacher model using the state features. Additionally, the method comprises training a candidate student model using the training data and the new weights. In addition, the method comprises generating a reward by comparing the current student model with the candidate student model using the training data and testing data to determine which is better at predicting an observed target. The method further comprises updating the teacher model with the reward in response to a stopping rule not being satisfied. The method additionally comprises returning the new weights and the current student model to a user in response to the stopping rule being satisfied, where the returned student model provides a prediction of the observed target.

[0096] In one embodiment of the present invention, the method further comprises determining whether the candidate student model generates a better prediction of the observed target than the current student model based on how close the prediction is to the observed target.

[0097] In one embodiment, the method further comprises updating the current student model with the candidate student model and updating current weights with the new weights in response to the candidate student model generating a better prediction of the observed target than the current student model.

[0098] In one embodiment, the method additionally comprises generating new state features by the updated student model using the new weights and generating a second set of new weights by the teacher model using the new state features.

[0099] Furthermore, in one embodiment, the method additionally comprises training the candidate student model using the training data and the second set of new weights and generating a reward by comparing the updated student model with the candidate student model using the training data and the testing data to determine which is better at predicting the observed target.

[0100] Additionally, in one embodiment, the method further comprises generating a second set of new weights by the updated teacher model using the state features in response to the candidate student model not generating a better prediction of the observed target than the current student model.

[0101] In one embodiment, the method further comprises training the candidate student model using the training data and the second set of new weights and generating a reward by comparing the current student model with the candidate student model using the training data and the testing data to determine which is better at predicting the observed target.

[0102] In one embodiment, the method additionally comprises having the stopping rule comprise one or more of the following: reaching a specified number of trials, reaching a specified training time, converging of a prediction accuracy, and a user-initiated termination.

[0103] In one embodiment, the method further comprises having the teacher model comprise a neural network, and having the student model comprise one of the following: a decision tree and a neural network.

[0104] Other forms of the embodiments of the method described above are in a system and in a computer program product.

[0105] The descriptions of the various embodiments of the present invention have been presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

CLAIMS:

1. A computer-implemented method for improving prediction accuracy in machine learning techniques, the method comprising:
 - constructing a teacher model, wherein said teacher model generates a weight for each data case;
 - training a current student model using training data and weights generated by said teacher model;
 - generating state features by said current student model;
 - generating new weights by said teacher model using said state features;
 - training a candidate student model using said training data and said new weights;
 - generating a reward by comparing said current student model with said candidate student model using said training data and testing data to determine which is better at predicting an observed target;
 - updating said teacher model with said reward in response to a stopping rule not being satisfied; and
 - returning said new weights and said current student model to a user in response to said stopping rule being satisfied, wherein said returned student model provides a prediction of said observed target.
2. The method as recited in claim 1 further comprising:
 - determining whether said candidate student model generates a better prediction of said observed target than said current student model based on how close the prediction is to said observed target.
3. The method as recited in claim 2 further comprising:
 - updating said current student model with said candidate student model and updating current weights with said new weights in response to said candidate student model generating a better prediction of said observed target than said current student model.
4. The method as recited in claim 3 further comprising:
 - generating new state features by said updated student model using said new weights; and
 - generating a second set of new weights by said teacher model using said new state features.
5. The method as recited in claim 4 further comprising:
 - training said candidate student model using said training data and said second set of new weights; and
 - generating a reward by comparing said updated student model with said candidate student model using said training data and said testing data to determine which is better at predicting said observed target.
6. The method as recited in claim 2 further comprising:
 - generating a second set of new weights by said updated teacher model using said state features in response to said candidate student model not generating a better prediction of said observed target than said current student model.

7. The method as recited in claim 6 further comprising:
training said candidate student model using said training data and said second set of new weights; and
generating a reward by comparing said current student model with said candidate student model using said training data and said testing data to determine which is better at predicting said observed target.
8. The method as recited in claim 1, wherein said stopping rule comprises one or more of the following:
reaching a specified number of trials, reaching a specified training time, converging of a prediction accuracy, and a user-initiated termination.
9. The method as recited in claim 1, wherein said teacher model comprises a neural network, wherein said student model comprises one of the following: a decision tree and a neural network.
10. A computer program product for improving prediction accuracy in machine learning techniques, the computer program product comprising a computer readable storage medium having program code embodied therewith, the program code comprising the programming instructions for:
constructing a teacher model, wherein said teacher model generates a weight for each data case;
training a current student model using training data and weights generated by said teacher model;
generating state features by said current student model;
generating new weights by said teacher model using said state features;
training a candidate student model using said training data and said new weights;
generating a reward by comparing said current student model with said candidate student model using said training data and testing data to determine which is better at predicting an observed target;
updating said teacher model with said reward in response to a stopping rule not being satisfied; and
returning said new weights and said current student model to a user in response to said stopping rule being satisfied, wherein said returned student model provides a prediction of said observed target.
11. The computer program product as recited in claim 10, wherein the program code further comprises the programming instructions for:
determining whether said candidate student model generates a better prediction of said observed target than said current student model based on how close the prediction is to said observed target.
12. The computer program product as recited in claim 11, wherein the program code further comprises the programming instructions for:
updating said current student model with said candidate student model and updating current weights with said new weights in response to said candidate student model generating a better prediction of said observed target than said current student model.

13. The computer program product as recited in claim 12, wherein the program code further comprises the programming instructions for:
- generating new state features by said updated student model using said new weights; and
 - generating a second set of new weights by said teacher model using said new state features.
14. The computer program product as recited in claim 13, wherein the program code further comprises the programming instructions for:
- training said candidate student model using said training data and said second set of new weights; and
 - generating a reward by comparing said updated student model with said candidate student model using said training data and said testing data to determine which is better at predicting said observed target.
15. The computer program product as recited in claim 11, wherein the program code further comprises the programming instructions for:
- generating a second set of new weights by said updated teacher model using said state features in response to said candidate student model not generating a better prediction of said observed target than said current student model.
16. The computer program product as recited in claim 15, wherein the program code further comprises the programming instructions for:
- training said candidate student model using said training data and said second set of new weights; and
 - generating a reward by comparing said current student model with said candidate student model using said training data and said testing data to determine which is better at predicting said observed target.
17. The computer program product as recited in claim 10, wherein said stopping rule comprises one or more of the following: reaching a specified number of trials, reaching a specified training time, converging of a prediction accuracy, and a user-initiated termination.
18. A system, comprising:
- a memory for storing a computer program for improving prediction accuracy in machine learning techniques; and
 - a processor connected to said memory, wherein said processor is configured to execute the program instructions of the computer program comprising:
 - constructing a teacher model, wherein said teacher model generates a weight for each data case;
 - training a current student model using training data and weights generated by said teacher model;
 - generating state features by said current student model;
 - generating new weights by said teacher model using said state features;
 - training a candidate student model using said training data and said new weights;

generating a reward by comparing said current student model with said candidate student model using said training data and testing data to determine which is better at predicting an observed target;
updating said teacher model with said reward in response to a stopping rule not being satisfied; and
returning said new weights and said current student model to a user in response to said stopping rule being satisfied, wherein said returned student model provides a prediction of said observed target.

19. The system as recited in claim 18, wherein the program instructions of the computer program further comprise:

determining whether said candidate student model generates a better prediction of said observed target than said current student model based on how close the prediction is to said observed target.

20. The system as recited in claim 19, wherein the program instructions of the computer program further comprise:

updating said current student model with said candidate student model and updating current weights with said new weights in response to said candidate student model generating a better prediction of said observed target than said current student model.

21. The system as recited in claim 20, wherein the program instructions of the computer program further comprise:

generating new state features by said updated student model using said new weights; and
generating a second set of new weights by said teacher model using said new state features.

22. The system as recited in claim 21, wherein the program instructions of the computer program further comprise:

training said candidate student model using said training data and said second set of new weights; and
generating a reward by comparing said updated student model with said candidate student model using said training data and said testing data to determine which is better at predicting said observed target.

23. The system as recited in claim 19, wherein the program instructions of the computer program further comprise:

generating a second set of new weights by said updated teacher model using said state features in response to said candidate student model not generating a better prediction of said observed target than said current student model.

24. The system as recited in claim 23, wherein the program instructions of the computer program further comprise:

training said candidate student model using said training data and said second set of new weights; and

generating a reward by comparing said current student model with said candidate student model using said training data and said testing data to determine which is better at predicting said observed target.

25. The system as recited in claim 18, wherein said stopping rule comprises one or more of the following: reaching a specified number of trials, reaching a specified training time, converging of a prediction accuracy, and a user-initiated termination.