

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.