SUMMARY

[004a] According to a first aspect of the present invention, there is provided a machine including: a communication interface configured to couple to a compute engine, a model storage database, and a results storage database; model builder circuitry coupled to the communication interface, the model builder circuitry configured to: receive a selection of an analytical model via the communication interface; train the analytical model to create a trained analytical model; and store the trained analytical model in the model storage database; model deployment circuity coupled to the communication interface, the model deployment circuitry configured to: receive a selection of the trained analytical model; retrieve the trained analytical model from the model storage database; deploy the trained analytical model on a compute engine to process incoming data and to generate results; and store the results in the results storage database; and resource allocation circuitry coupled to the model deployment circuitry, the resource allocation circuitry configured to: determine a resource load, an execution time, or both for the trained analytical model using historical data for execution of the trained analytical model; determine a resource load capability of the compute engine; and determine an execution frequency for the trained analytical model on the compute engine based on the resource load, the execution time for the trained analytical model, the resource load capability of the compute engine, or any combination thereof.

[004b] According to a second aspect of the present invention, there is provided a computer-implemented method including: receiving a selection of an analytical model by model builder circuitry via a communication interface; training the analytical model by the model builder circuitry to create a trained analytical model; storing the trained analytical model by the model builder circuitry in a model storage database; receiving by model deployment circuitry via the communication interface a selection of the trained analytical model; retrieving by the model deployment circuitry the trained analytical model from the model storage database; deploying by the model deployment circuitry the trained analytical model on a compute engine to process incoming data and to generate results; storing by the model deployment circuitry the results in a results storage database; determining, by resource allocation circuitry coupled to the model deployment circuitry, a resource load, an execution time, or both for the trained analytical model using historical data for execution of the trained analytical model; determining, by the resource allocation circuitry coupled to the model deployment circuitry, a resource load capability of the compute engine; and determining, by the resource allocation circuitry coupled to the model deployment circuitry, an execution frequency for the trained

analytical model on the compute engine based on the resource load, the execution time for the trained analytical model, the resource load capability of the compute engine, or any combination thereof.

[004c] According to a third aspect of the present invention, there is provided a machine including: a communication interface configured to communicate with: a compute engine; a model storage database; a results storage database; and a network; a model metadata store configured to store analytical model metadata; model builder circuitry coupled to the communication interface, the model builder circuitry configured to: receive a selection of an analytical model via the communication interface; train the analytical model on the compute engine using user-designated training data to create a trained analytical model; validate the trained analytical model using the user-designated training data to determine an error rate of the trained analytical model; store the trained analytical model in a model storage database; and store model metadata characterizing the trained analytical model in the model metadata store; model deployment circuity coupled to the communication interface, the model deployment circuitry configured to: receive a selection of the trained analytical model; retrieve the trained analytical model from the model storage database; deploy the trained analytical model on a compute engine to process incoming data and to generate results; and store the results in a results storage database; user interface circuitry coupled to the communication interface, the user interface circuitry configured to: provide an analytical model control user interface via the communication interface to a user, and through the analytical model control user interface: interact with the user to create the analytical model with the model builder circuitry; accept a user selection of an analytical model algorithm type for the analytical model; accept a user selection of an analytical model input for the analytical model; accept a user alteration input applicable to an analytical model parameter for the analytical model; accept a designation input specifying training data with which to train the analytical model; accept a user training input configured to cause training of the analytical model with the user-designated training data by the model builder circuitry to create the trained analytical model; cause display of the error rate for the trained analytical model to the user; accept a deployment input configured to cause deployment of the trained analytical model on the compute engine by the model deployment circuitry; and accept a retraction input configured to cause withdrawal of deployment of the trained analytical model by the model deployment circuitry; and resource allocation circuitry coupled to the model deployment circuitry, the resource allocation circuitry configured to: determine a resource load for the trained analytical model, an execution time for the trained analytical model, or

both, using historical data for execution of the trained analytical model; determine a resource load capability of the compute engine; and determine a maximum execution frequency for the trained analytical model on the compute engine based on the resource load, the execution time for the trained analytical model, the resource load capability of the compute engine, or any combination thereof.

BRIEF DESCRIPTION OF THE DRAWINGS

[005] Figure 1 shows an example machine for implementing machine learning.

[006] Figure 2 shows an example specific system implementation for the machine.

[007] Figure 3 shows further detail of the example specific system implementation shown in Figure 2.

[008] Figure 4 shows a flow diagram of logic that the machine may implement.

[009] Figure 5 shows another flow diagram of logic that the machine may implement.

[010] Figure 6 shows another flow diagram of logic that the machine may implement.

[011] Figure 7 shows an example graphical user interface.

[012] Figure 8 shows a variation of the example graphical user interface of Figure 7.

[013] Figure 9 shows another variation of the example graphical user interface of Figure 7.

[014] Figure 10 shows another flow diagram of logic that the machine may implement.

[015] Figure 11 shows another flow diagram of logic that the machine may implement.

[016] Figure 12 shows another flow diagram of logic that the machine may implement.

DETAILED DESCRIPTION

[017] Analytical models, also called predictive models, are a form of mathematical model that defines and describes relationships among variables in a data set. An analytical model may include an equation and a set of coefficients which, together, map an arbitrary number of input variables to an output, typically a single output. This output can be used to provide predictive results for a given set of inputs. For example, an analytical model may receive inputs of a future date and the current temperature, and may provide an output predicting the chance of rain on that future date. Analytical models can also provide information on events, occurrences, or transactions that have already occurred. For example, analytical models can be used to determine whether a recorded transaction was fraudulent.

Claims

1. A machine including:

a communication interface configured to couple to a compute engine, a model storage database, and a results storage database;

model builder circuitry coupled to the communication interface, the model builder circuitry configured to:

receive a selection of an analytical model via the communication interface; train the analytical model to create a trained analytical model; and store the trained analytical model in the model storage database;

model deployment circuity coupled to the communication interface, the model deployment circuitry configured to:

receive a selection of the trained analytical model;

retrieve the trained analytical model from the model storage database;

deploy the trained analytical model on a compute engine to process incoming data and to generate results; and

store the results in the results storage database; and

resource allocation circuitry coupled to the model deployment circuitry, the resource allocation circuitry configured to:

determine a resource load, an execution time, or both for the trained analytical model using historical data for execution of the trained analytical model;

determine a resource load capability of the compute engine; and

determine an execution frequency for the trained analytical model on the compute engine based on the resource load, the execution time for the trained analytical model, the resource load capability of the compute engine, or any combination thereof.

2. A machine according to claim 1, wherein the model builder circuitry is further configured to train the analytical model on the compute engine with user-designated training data.

3. A machine according to either claim 1 or claim 2, wherein the model builder circuitry is further configured to store model metadata corresponding to the trained analytical model in a model metadata store.

4. A machine according to claim 3, wherein the model metadata includes an analytical model parameter for the trained analytical model, a storage location specifier for the trained analytical model, a storage location specifier for the results of the results of the trained analytical model, or any combination thereof.

5. A machine according to any one of the preceding claims, wherein the resource allocation circuitry is further configured to determine the resource load capability of the compute engine by determining a resource load for a different trained analytical model, an execution time for the different trained analytical model, or both, using historical data for execution of the different trained analytical model.

6. A machine according to any one of the preceding claims, wherein the resource allocation circuitry is further configured to schedule an execution of the trained analytical model by the compute engine at a frequency less than or equal to the execution frequency.

7. A machine according to any one of the preceding claims, further including:

user interface circuitry coupled to the communication interface, the user interface circuitry configured to:

provide an analytical model control user interface via the communication interface to a user, and through the analytical model control user interface:

interact with the user to create the analytical model with the model builder
circuitry;

accept a user selection of an analytical model algorithm type for the analytical model;

accept a user selection of an analytical model input for the analytical model; accept a user alteration input applicable to an analytical model parameter for the analytical model;

accept a designation input specifying training data with which to train the analytical model;

accept a user training input configured to cause training of the analytical model with the training data by the model builder circuitry to create the trained analytical model;

accept a deployment input configured to cause deployment of the trained analytical model on the compute engine by the model deployment circuitry; and

accept a retraction input configured to cause withdrawal of deployment of the trained analytical model by the model deployment circuitry.

8. A computer-implemented method including:

receiving a selection of an analytical model by model builder circuitry via a communication interface;

training the analytical model by the model builder circuitry to create a trained analytical model;

storing the trained analytical model by the model builder circuitry in a model storage database;

receiving by model deployment circuitry via the communication interface a selection of the trained analytical model;

retrieving by the model deployment circuitry the trained analytical model from the model storage database;

deploying by the model deployment circuitry the trained analytical model on a compute engine to process incoming data and to generate results;

storing by the model deployment circuitry the results in a results storage database;

determining, by resource allocation circuitry coupled to the model deployment circuitry, a resource load, an execution time, or both for the trained analytical model using historical data for execution of the trained analytical model;

determining, by the resource allocation circuitry coupled to the model deployment circuitry, a resource load capability of the compute engine; and

determining, by the resource allocation circuitry coupled to the model deployment circuitry, an execution frequency for the trained analytical model on the compute engine based on the resource load, the execution time for the trained

analytical model, the resource load capability of the compute engine, or any combination thereof.

9. A method according to claim 8, further including:

training by the model builder circuitry the analytical model on the compute engine with user-designated training data.

10. A method according to either claim 8 or claim 9, further including:

storing, by the model builder circuitry, model metadata corresponding to the trained analytical model in a model metadata store.

11. A method according to claim 10, wherein storing the model metadata further includes storing an analytical model parameter for the trained analytical model, a storage location specifier for the trained analytical model, a storage location for the results of the trained analytical model, or any combination thereof.

12. A method according to any one of claims 8 to 11, further including:

determining by the resource allocation circuitry the resource load capability of the compute engine by determining a resource load time for a different trained analytical model, an execution time for the different analytical model, or both, using historical data for execution of the different trained analytical model.

13. A method according to any one of claims 8 to 12, further including:

scheduling by the resource allocation circuitry an execution of the trained analytical model by the compute engine at a frequency less than or equal to the execution frequency.

14. A method according to any one of claims 8 to 13, further including:

providing by user interface circuitry an analytical model control user interface via the communication interface to a user;

interacting by the user interface circuitry with the user to create the analytical model via the analytical model control user interface;

accepting by the user interface circuitry a selection from the user of an analytical model algorithm type for the analytical model via the analytical model control user interface;

accepting by the user interface circuitry a selection from the user of an analytical model input for the analytical model via the analytical model control user interface;

accepting by the user interface circuitry a user alteration input from the user applicable to an analytical model parameter for the analytical model via the analytical model control user interface;

accepting by the user interface circuitry a designation input from the user specifying training data with which to train the analytical model via the analytical model control user interface;

accepting by the user interface circuitry a user training input from the user configured to cause training of the analytical model with the training data by the model builder circuitry via the analytical model control user interface to create the trained analytical model;

accepting by the user interface circuitry a deployment input from the user configured to cause deployment the trained analytical model on the compute engine by the model deployment circuitry via the analytical model control user interface; and

accepting by the user interface circuitry a retraction input from the user configured to cause withdrawal of deployment of the trained analytical model by the model deployment circuitry via the analytical model control user interface.

15. A machine including:

a communication interface configured to communicate with:

a compute engine;

a model storage database;

a results storage database; and

a network;

a model metadata store configured to store analytical model metadata;

model builder circuitry coupled to the communication interface, the model builder circuitry configured to:

receive a selection of an analytical model via the communication interface;

train the analytical model on the compute engine using user-designated training data to create a trained analytical model;

validate the trained analytical model using the user-designated training data to determine an error rate of the trained analytical model;

store the trained analytical model in a model storage database; and

store model metadata characterizing the trained analytical model in the model metadata store;

model deployment circuity coupled to the communication interface, the model deployment circuitry configured to:

receive a selection of the trained analytical model; retrieve the trained analytical model from the model storage database;

deploy the trained analytical model on a compute engine to process incoming data and to generate results; and

store the results in a results storage database;

user interface circuitry coupled to the communication interface, the user interface circuitry configured to:

provide an analytical model control user interface via the communication interface to a user, and through the analytical model control user interface:

interact with the user to create the analytical model with the model builder circuitry;

accept a user selection of an analytical model algorithm type for the analytical model;

accept a user selection of an analytical model input for the analytical model;

accept a user alteration input applicable to an analytical model parameter for the analytical model;

accept a designation input specifying training data with which to train the analytical model;

accept a user training input configured to cause training of the analytical model with the user-designated training data by the model builder circuitry to create the trained analytical model;

cause display of the error rate for the trained analytical model to the user;

accept a deployment input configured to cause deployment of the trained analytical model on the compute engine by the model deployment circuitry; and

accept a retraction input configured to cause withdrawal of deployment of the trained analytical model by the model deployment circuitry; and

resource allocation circuitry coupled to the model deployment circuitry, the resource allocation circuitry configured to:

determine a resource load for the trained analytical model, an execution time for the trained analytical model, or both, using historical data for execution of the trained analytical model;

determine a resource load capability of the compute engine; and

determine a maximum execution frequency for the trained analytical model on the compute engine based on the resource load, the execution time for the trained analytical model, the resource load capability of the compute engine, or any combination thereof.

16. A system according to claim 15, wherein the model metadata includes an analytical model parameter for the trained analytical model, a storage location specifier for the trained analytical model, a storage location specifier for the results of the results of the trained analytical model, or any combination thereof.

17. A system according to either claim 15 or claim 16, wherein the resource allocation circuitry is further configured to schedule execution of the trained analytical model by the compute engine no more frequently than the maximum execution frequency.