



(19) **United States**

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

(10) **Pub. No.: US 2013/0024173 A1**

(43) **Pub. Date: Jan. 24, 2013**

(54) **COMPUTER-IMPLEMENTED SYSTEMS AND METHODS FOR TESTING LARGE SCALE AUTOMATIC FORECAST COMBINATIONS**

(60) Provisional application No. 61/594,442, filed on Feb. 3, 2012.

(76) Inventors: **Jerzy Michal Brzezicki**, Cary, NC (US); **Dinesh P. Apte**, Pune (IN); **Michael J. Leonard**, Cary, NC (US); **Michael Ryan Chipley**, Raleigh, NC (US); **Sagar Arun Mainkar**, Pune (IN); **Edward Tilden Blair**, Cary, NC (US)

Publication Classification

(51) **Int. Cl.**
G06G 7/48 (2006.01)
(52) **U.S. Cl.** **703/6**

(21) Appl. No.: **13/440,045**


(22) Filed: **Apr. 5, 2012**

Related U.S. Application Data

(63) Continuation-in-part of application No. 13/189,131, filed on Jul. 22, 2011.

(57) ABSTRACT

Systems and methods are provided for evaluating performance of forecasting models. A plurality of forecasting models may be generated using a set of in-sample data. Two or more forecasting models from the plurality of forecasting models may be selected for use in generating a combined forecast. An ex-ante combined forecast may be generated for an out-of-sample period using the selected two or more forecasting models. The ex-ante combined forecast may then be compared with a set of actual out-of-sample data to evaluate performance of the combined forecast.






Fig. 1




Fig. 2

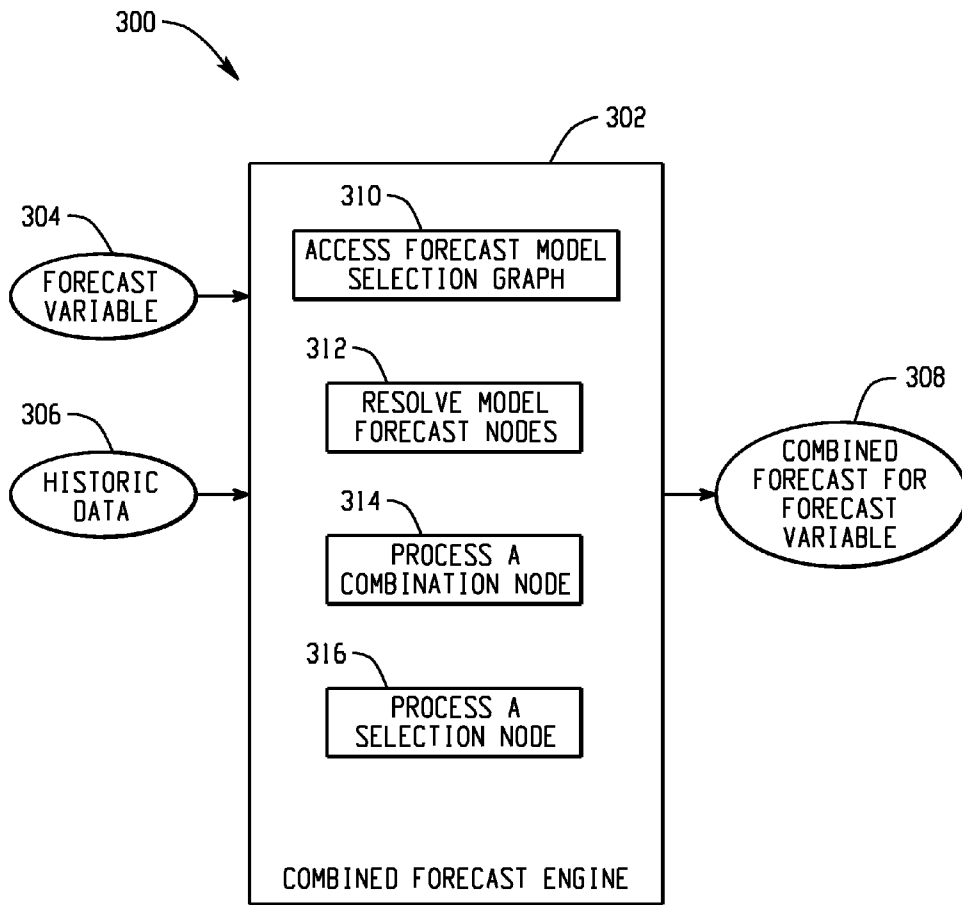


Fig. 3

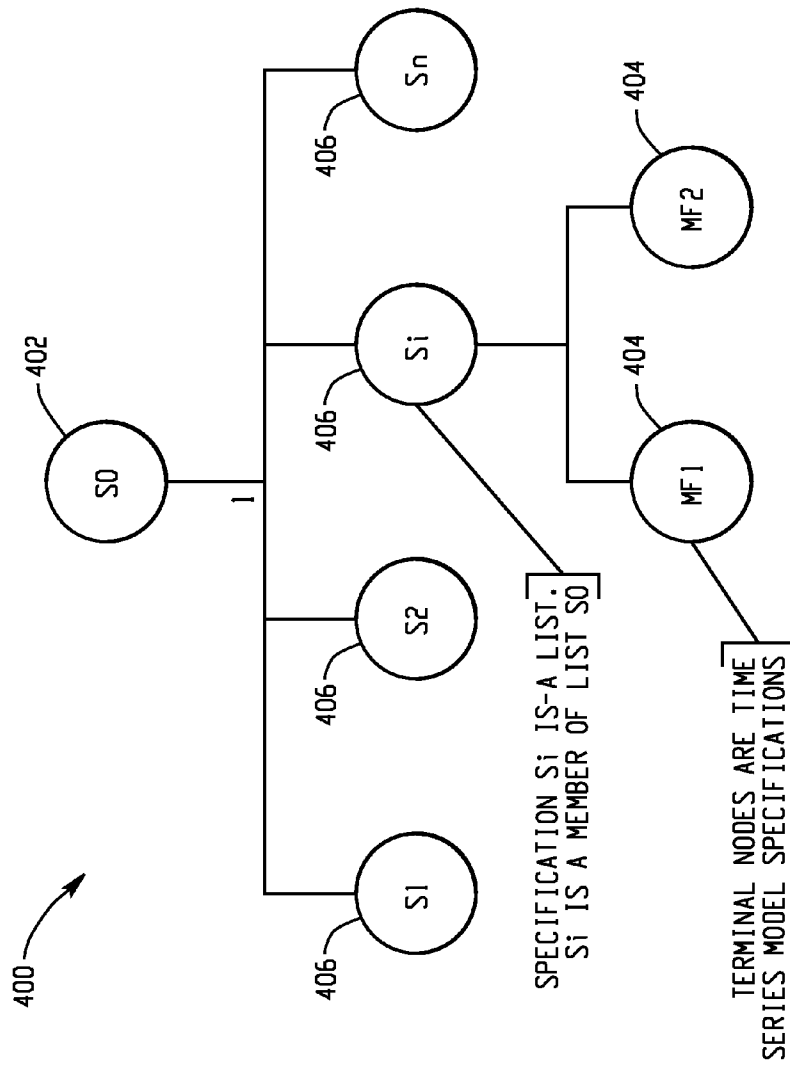


Fig. 4

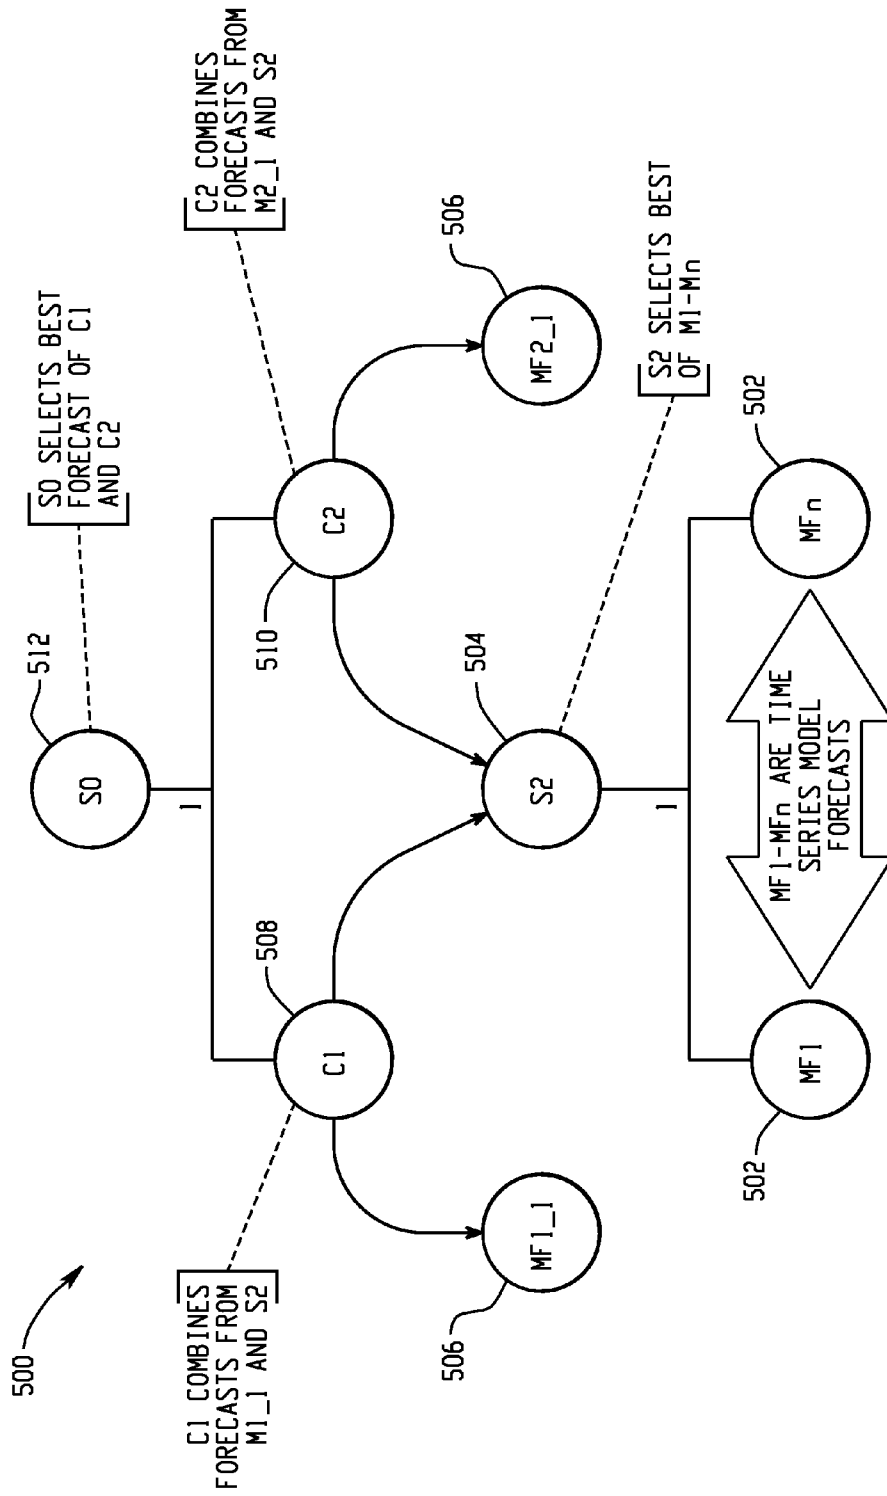


Fig. 5

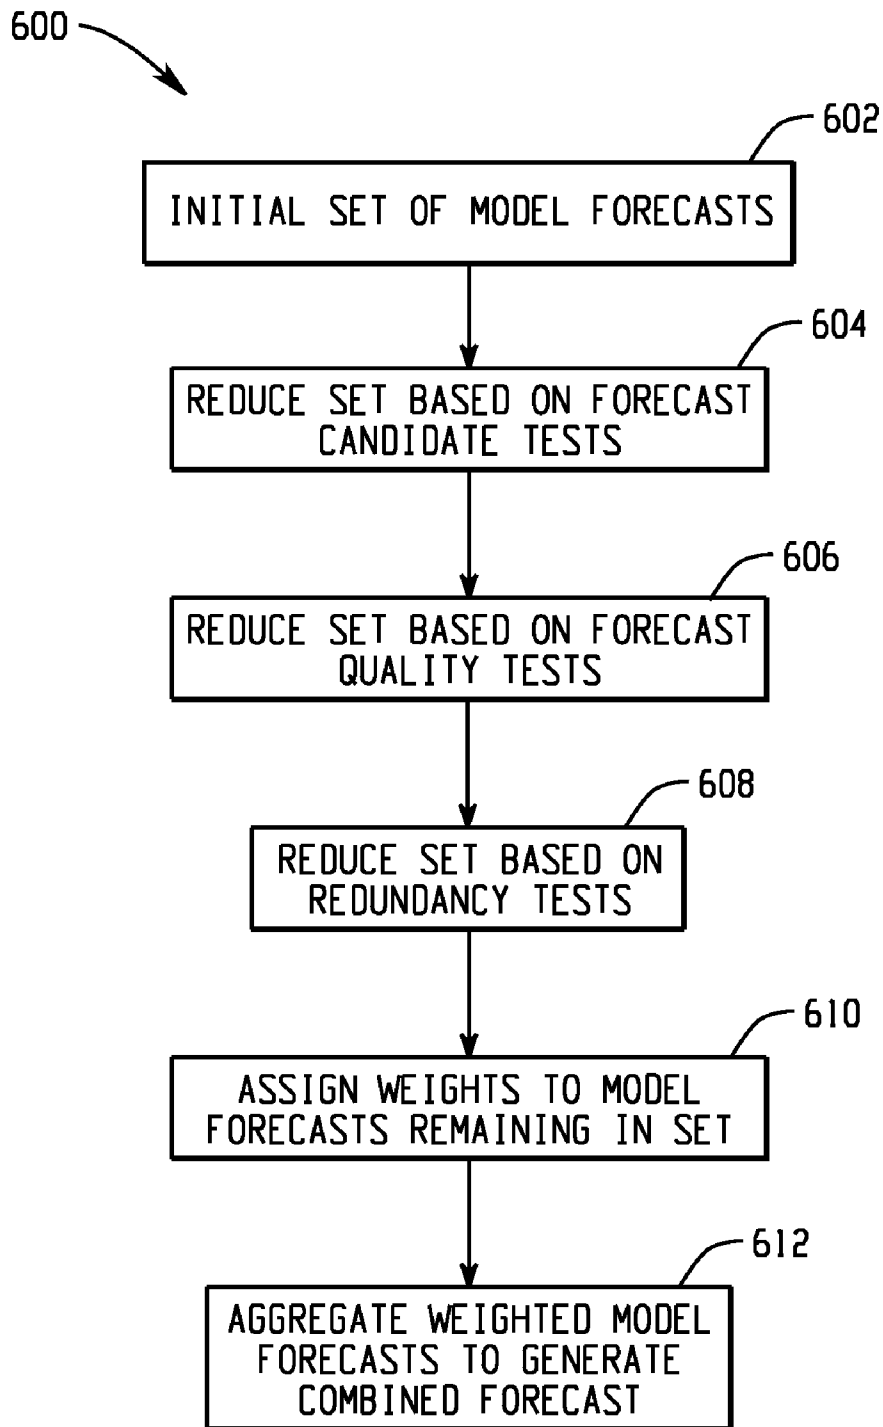


Fig. 6




Fig. 7

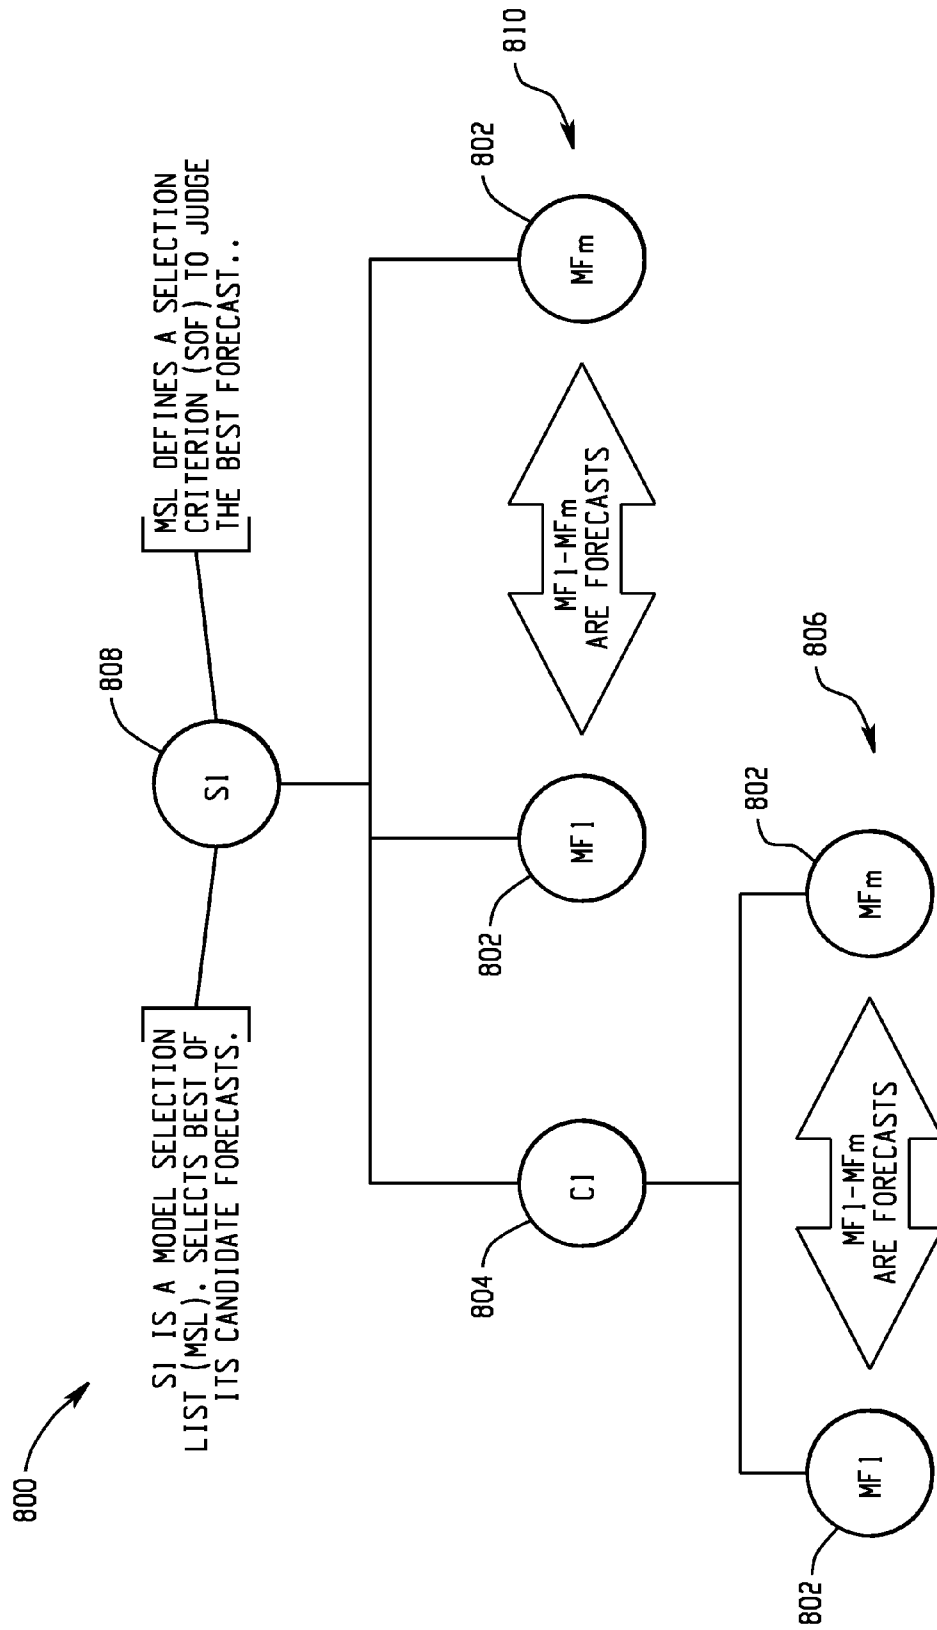


Fig. 8




Fig. 9

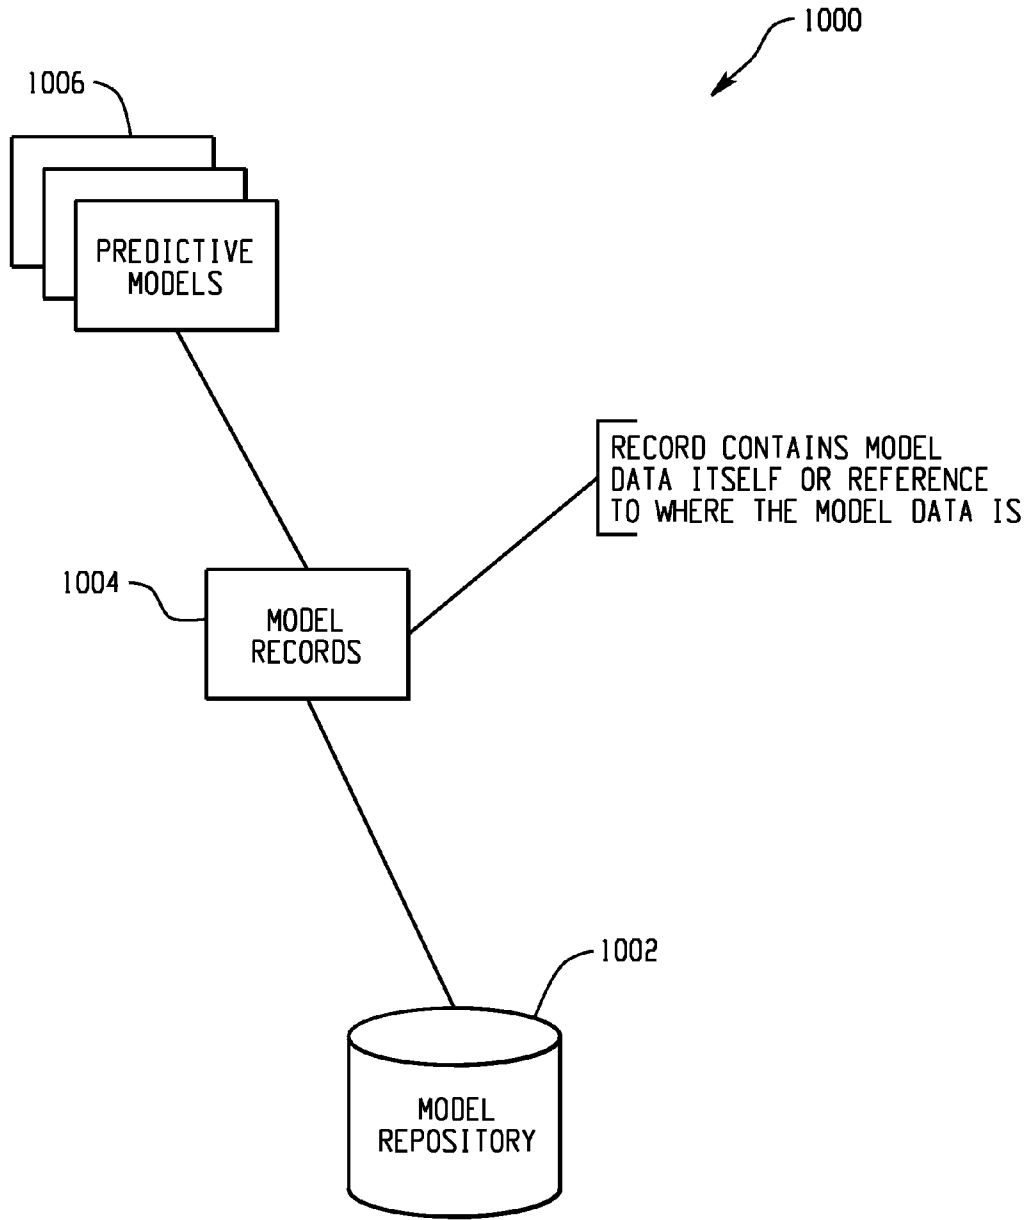


Fig. 10




Fig. 11

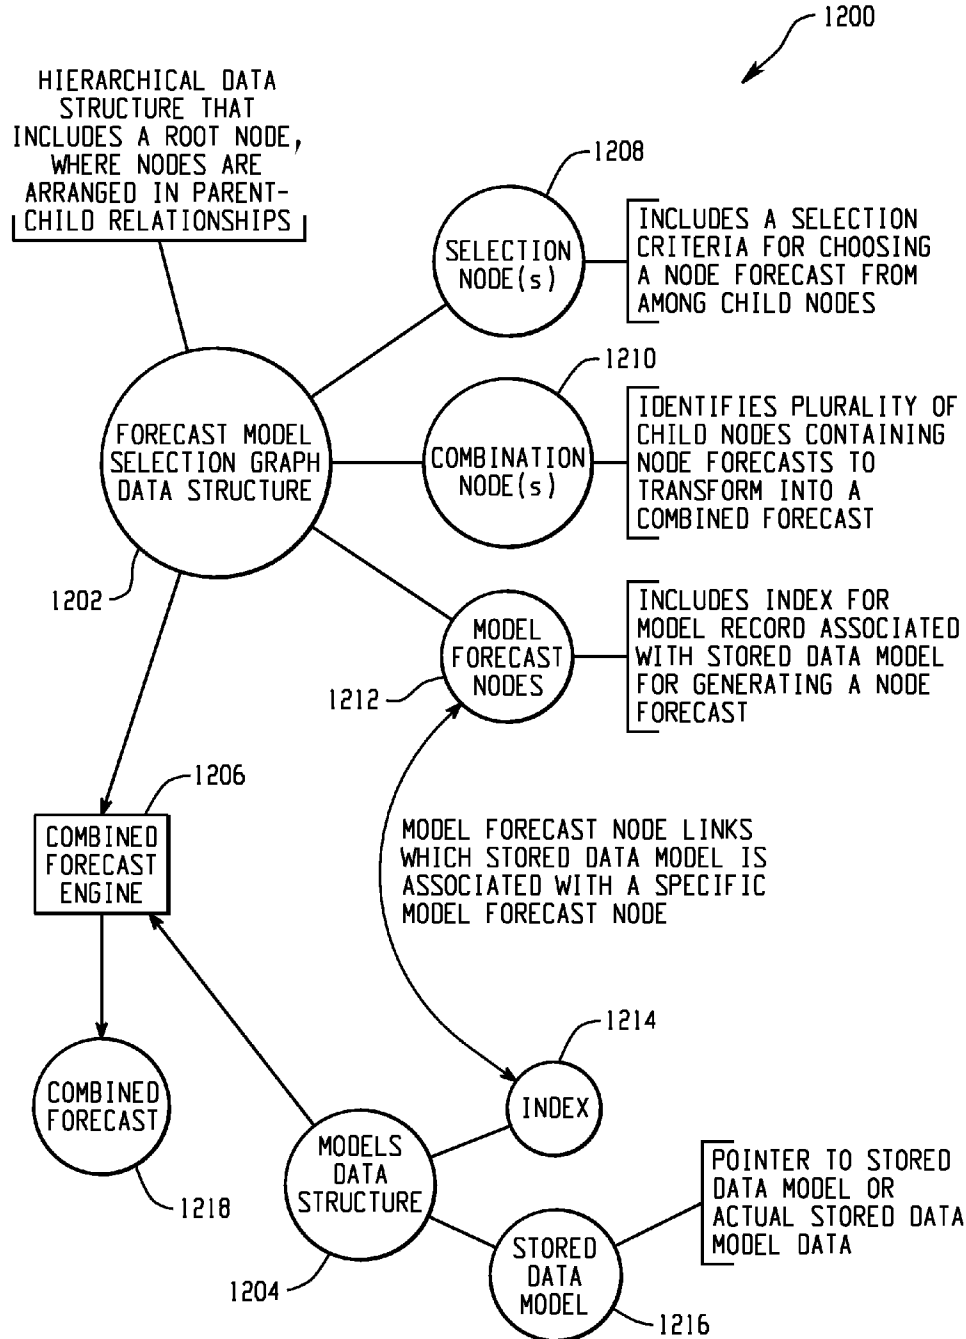


Fig. 12

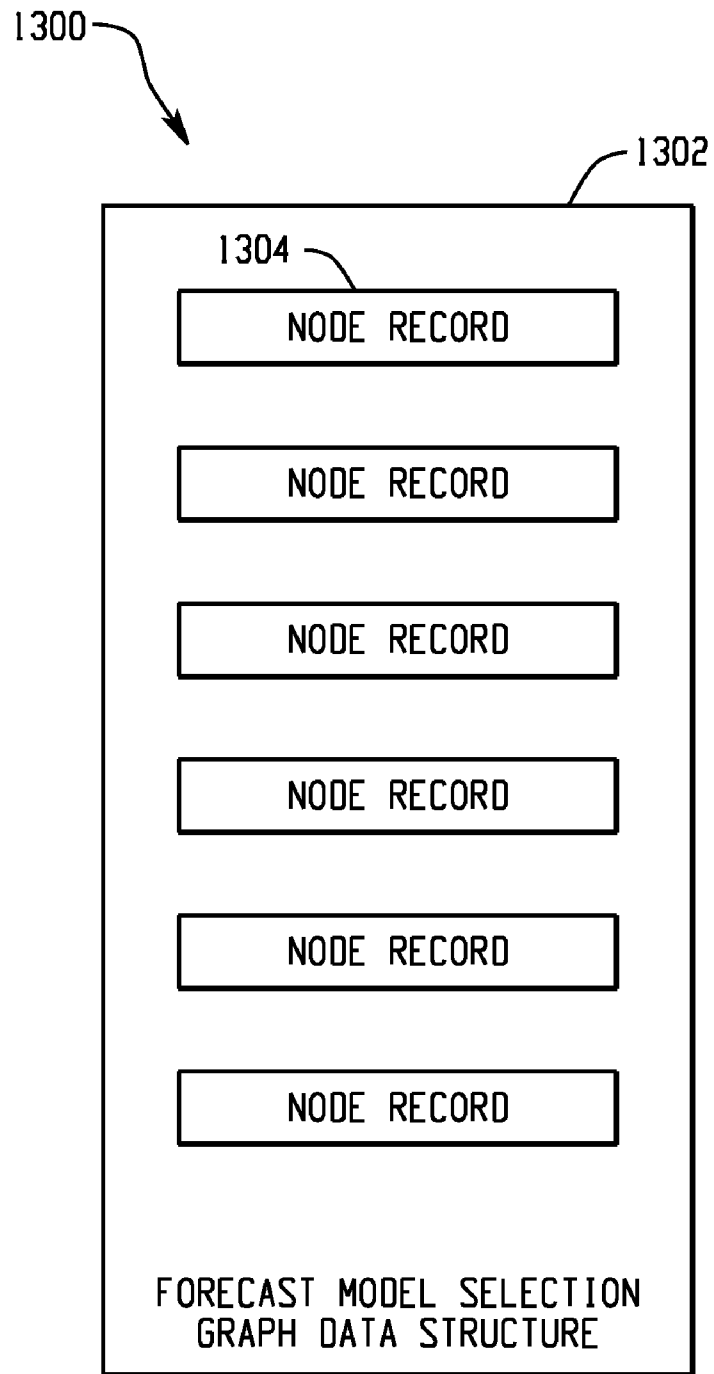


Fig. 13

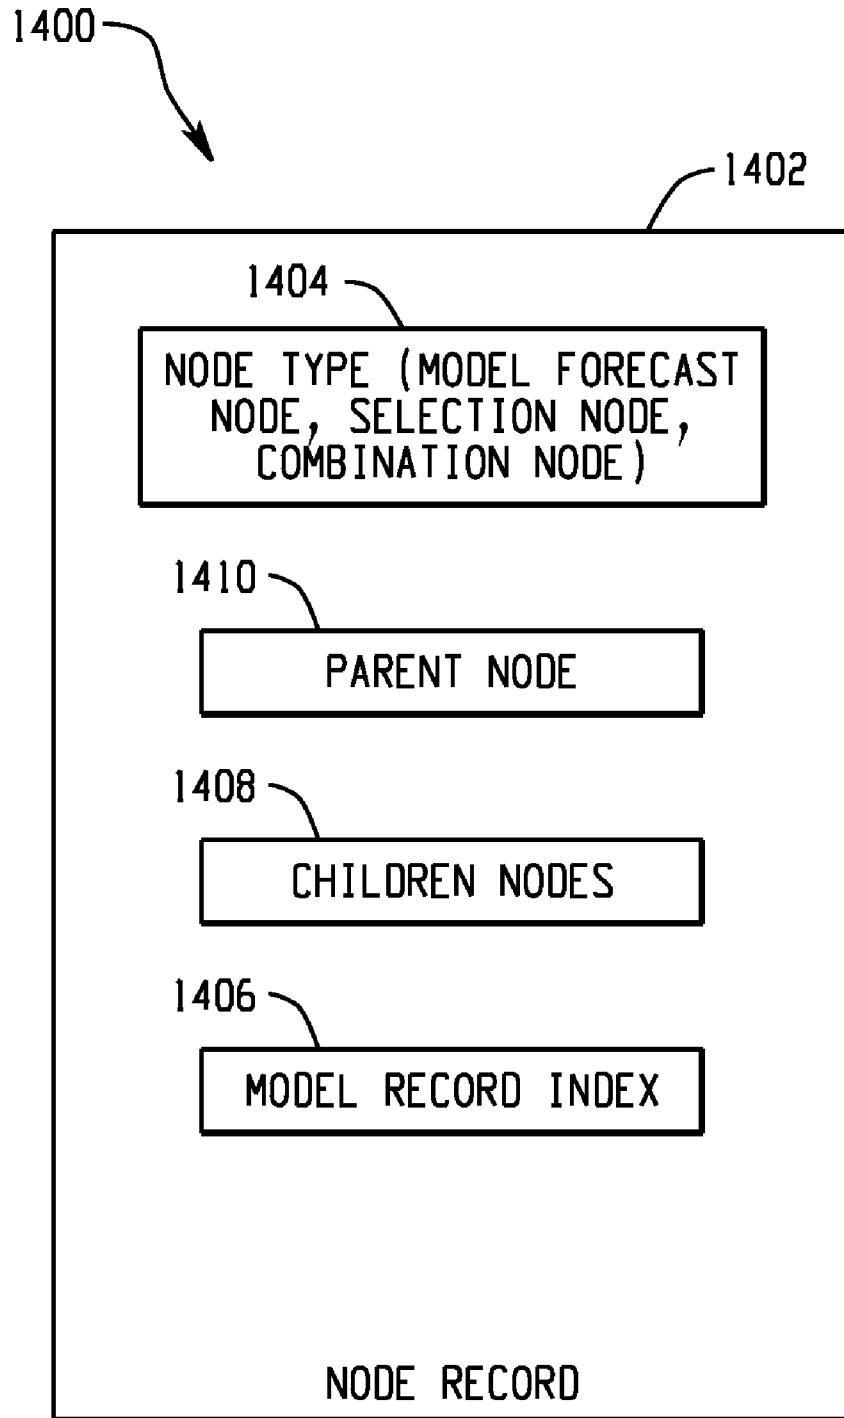


Fig. 14




Fig. 15

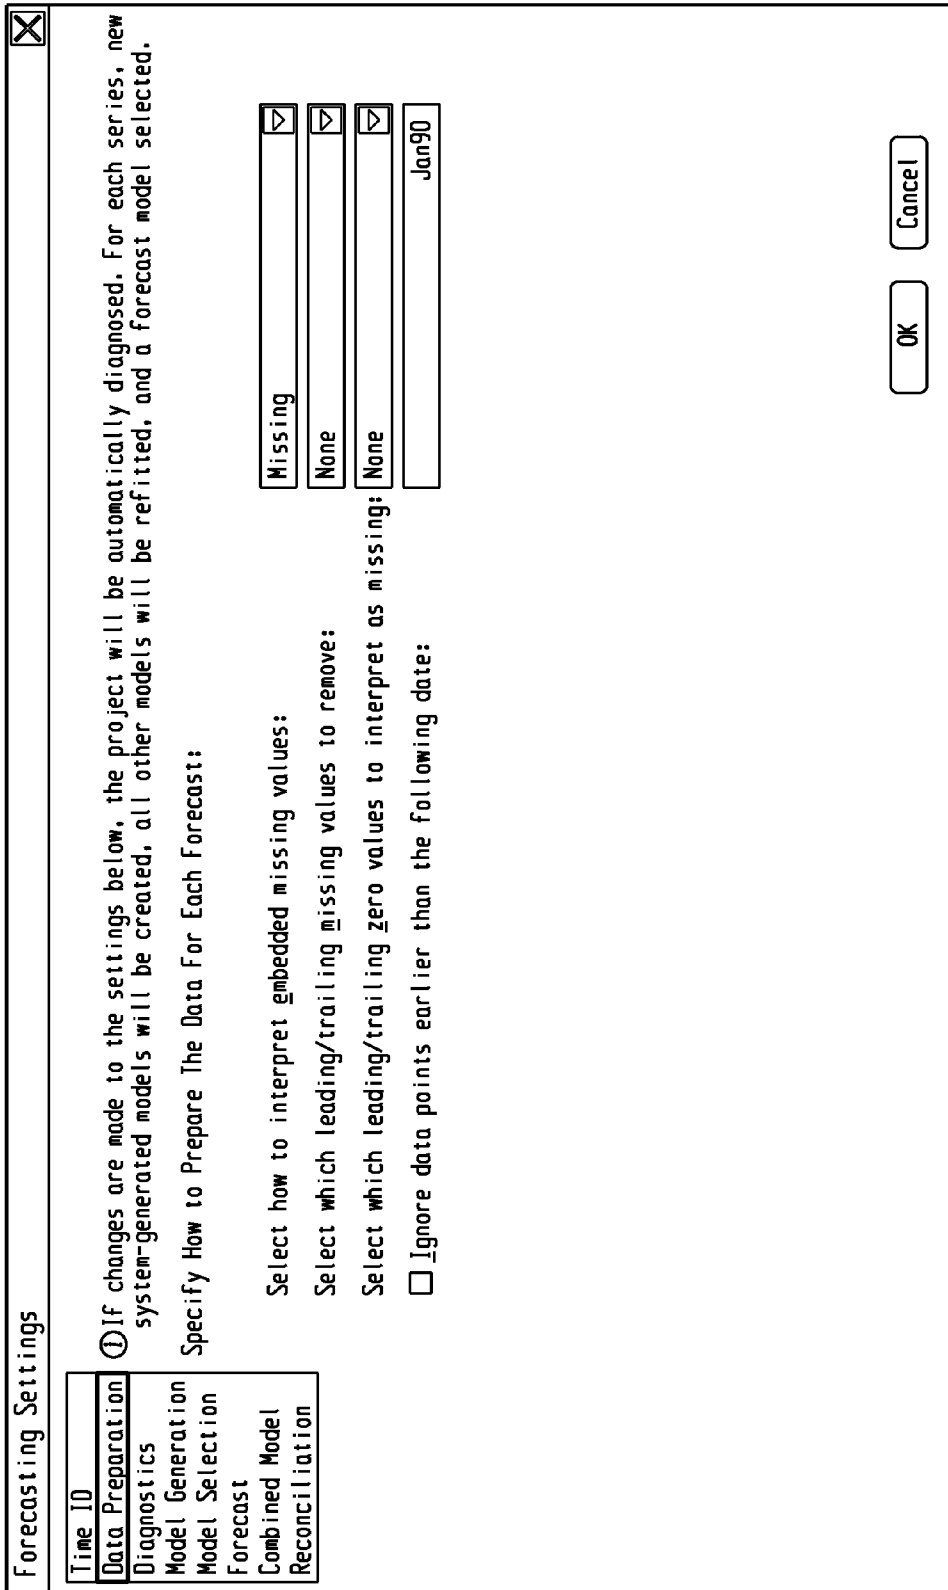


Fig. 16

Forecasting Settings

Time 10

Data Preparation

Diagnosics

Model Generation

Model Selection

Forecast

Combined Model

Reconciliation

ⓘ If changes are made to the settings below, the project will be automatically diagnosed. For each series, new system-generated models will be created, all other models will be refitted, and a forecast model selected.

Use the following settings to diagnose the properties of each series:

Perform intermittency test. Sensitivity: 2

Perform seasonality test. Sensitivity: 0.01

Minimum number of seasonal cycles for a seasonal model: 2

Minimum number of observations for a trend model: 2

Minimum number of observations for a non-mean model: 2

Functional transformation (dependent): None

Box-Cox parameter: 0

Diagnose independent variables separately: Both

Forecast: Median

Outlier detection (ARIMA models only):

Detect outliers

Significance level: 2

Maximum percentage of series that can be outliers: 0.01

2

OK Cancel

Fig. 17

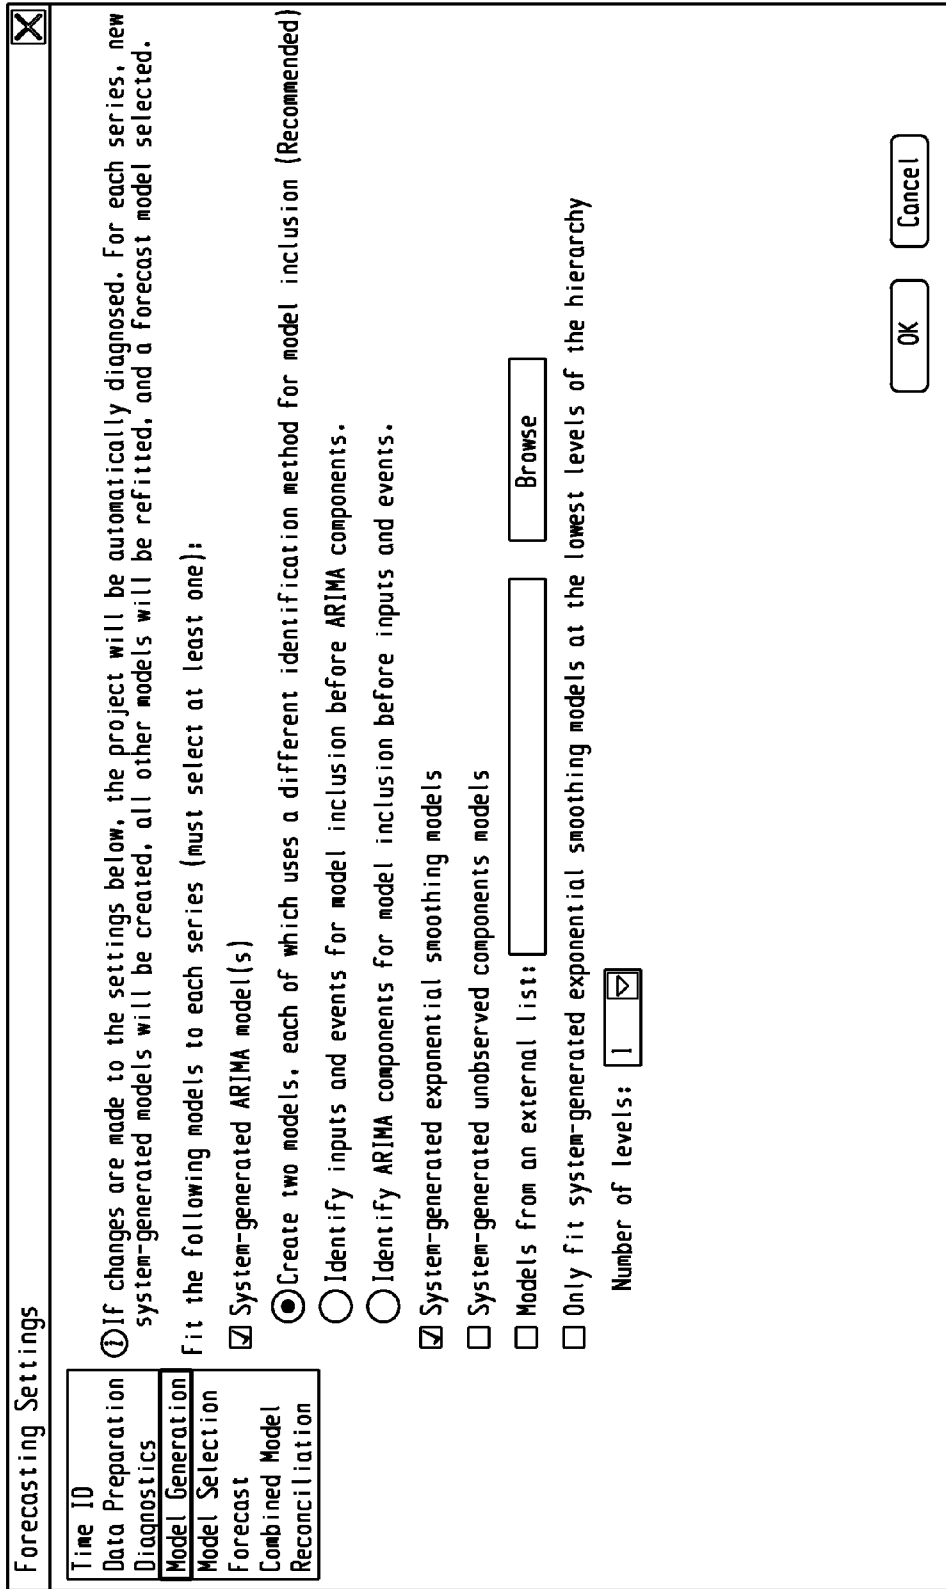


Fig. 18

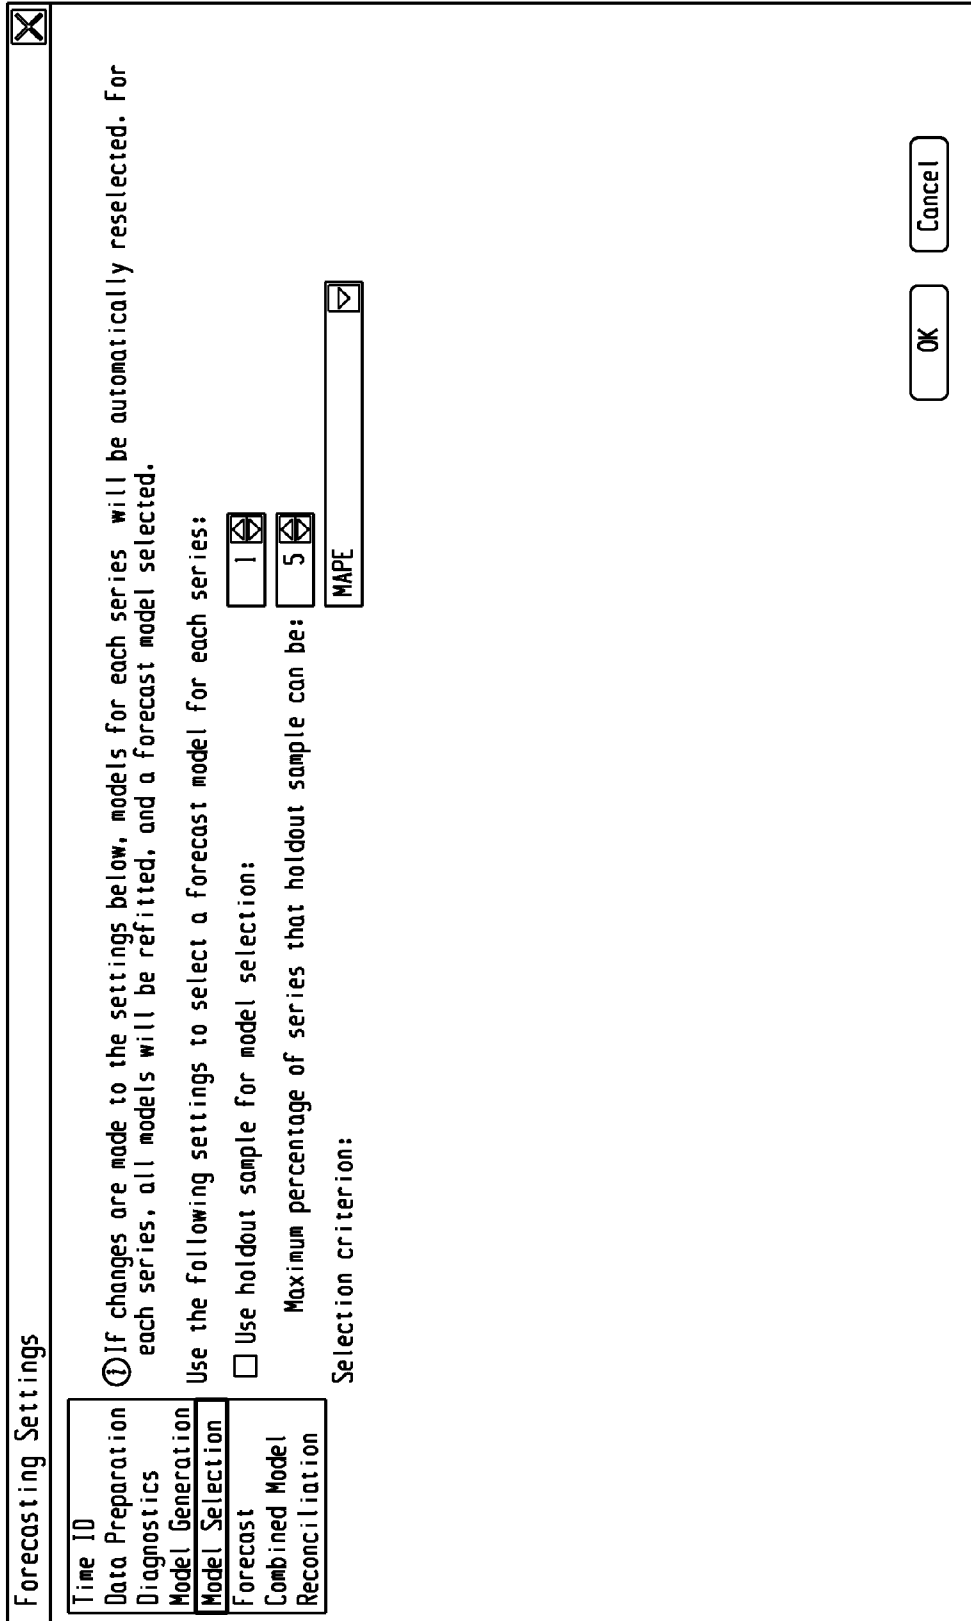


Fig. 19

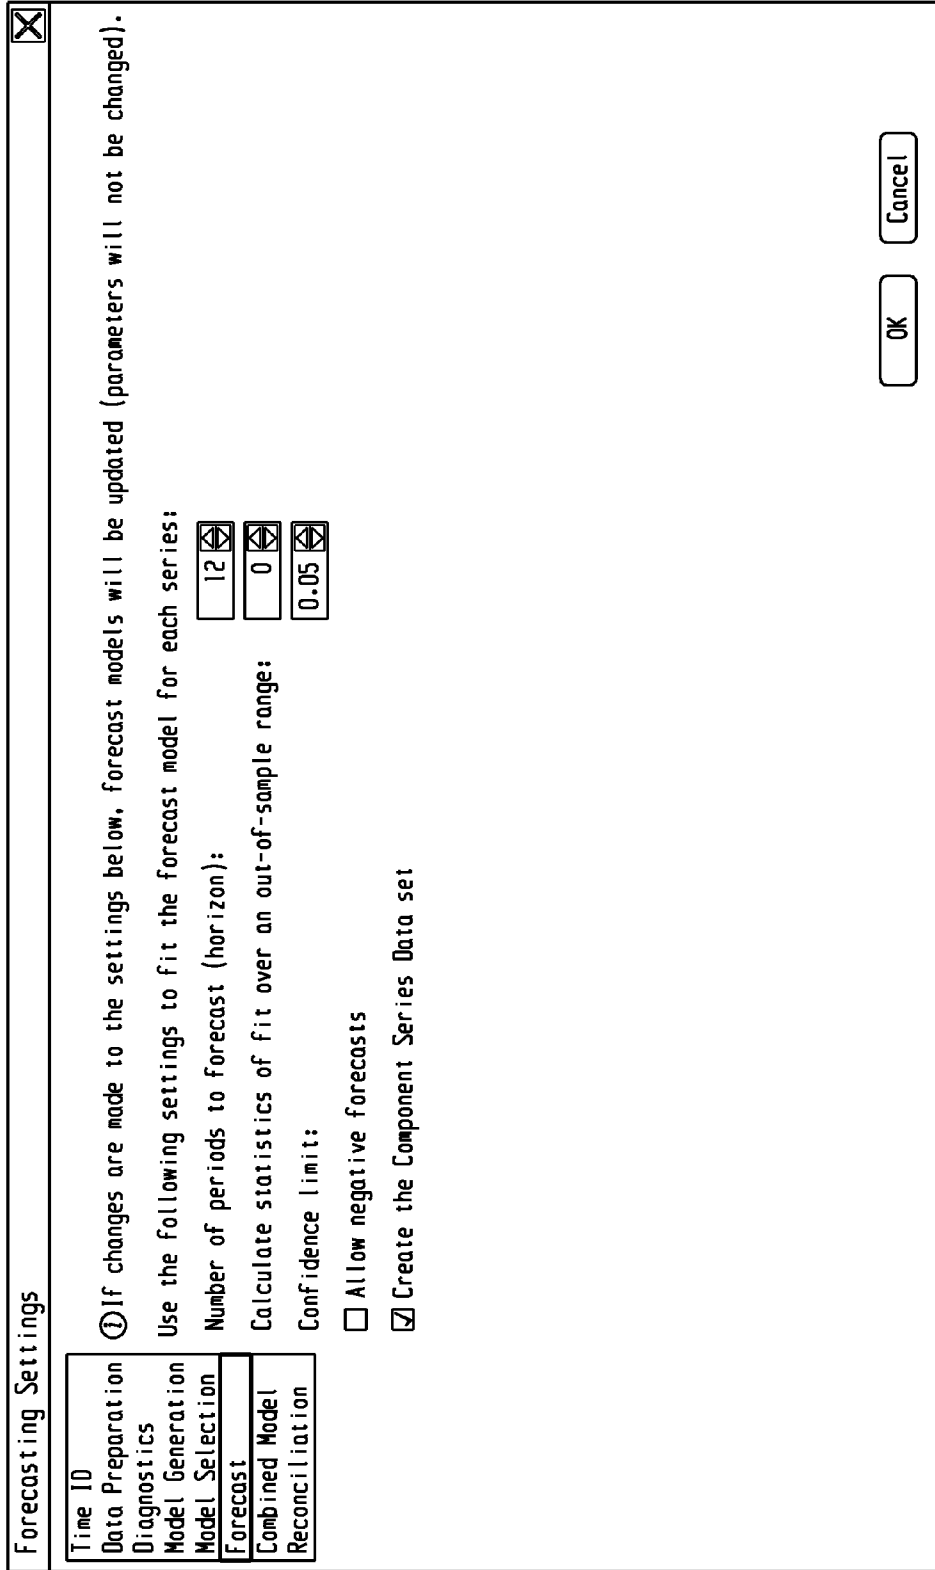


Fig. 20




Fig. 21




Fig. 22

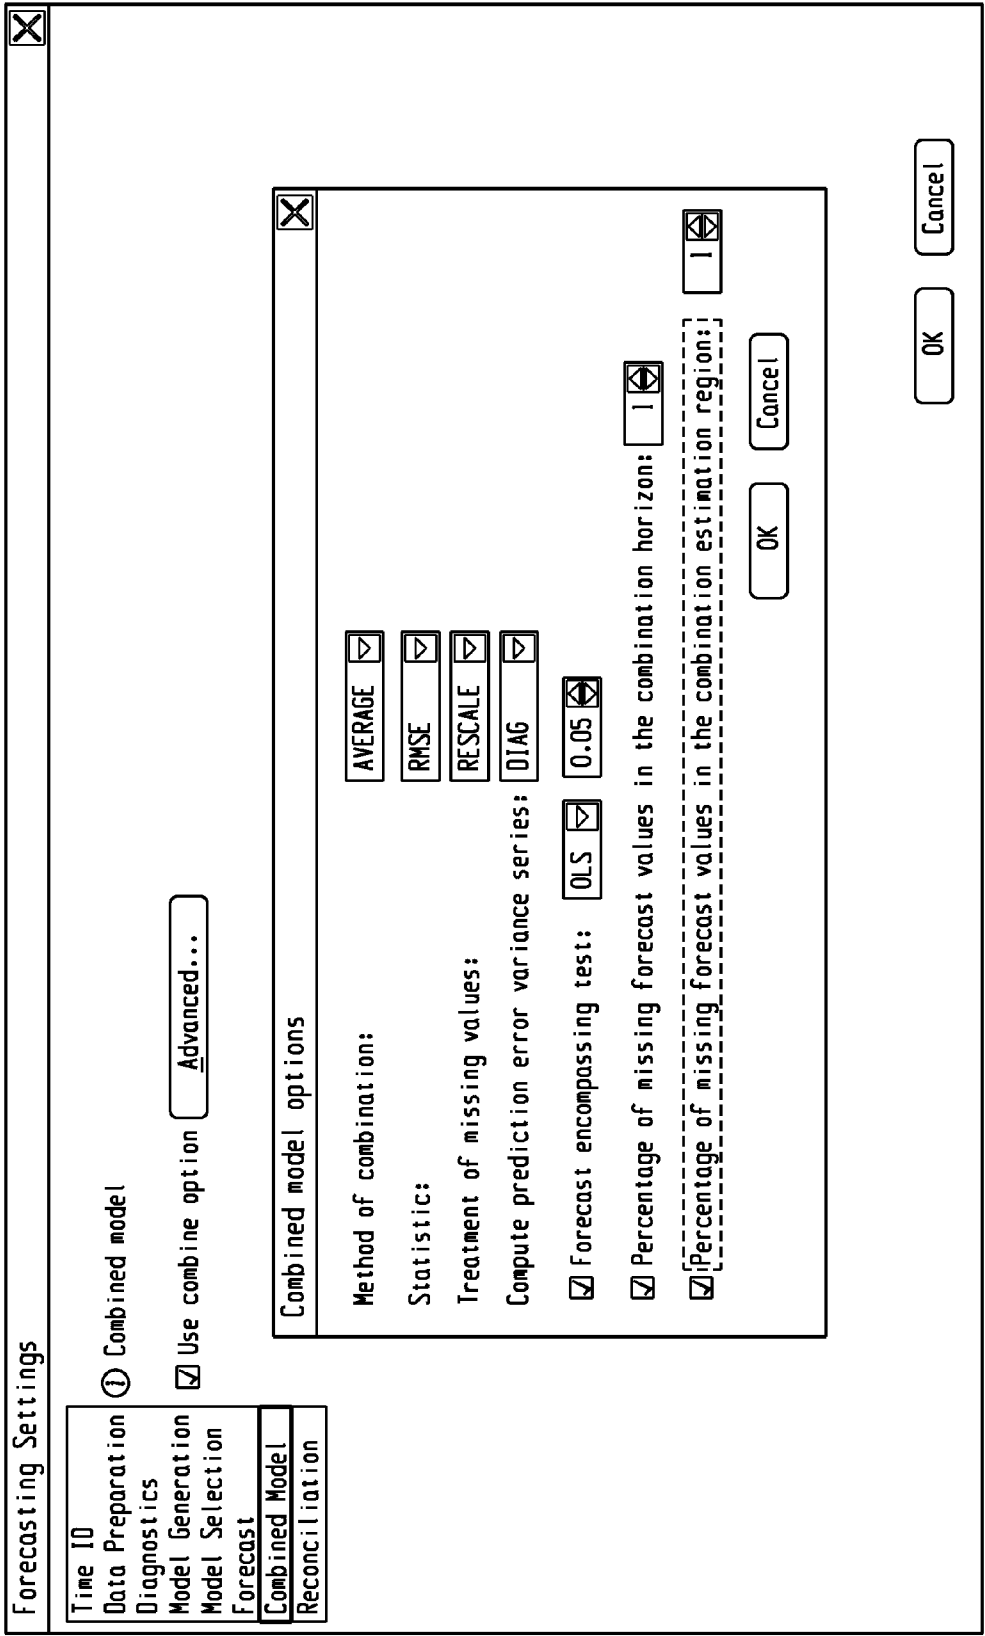


Fig. 23

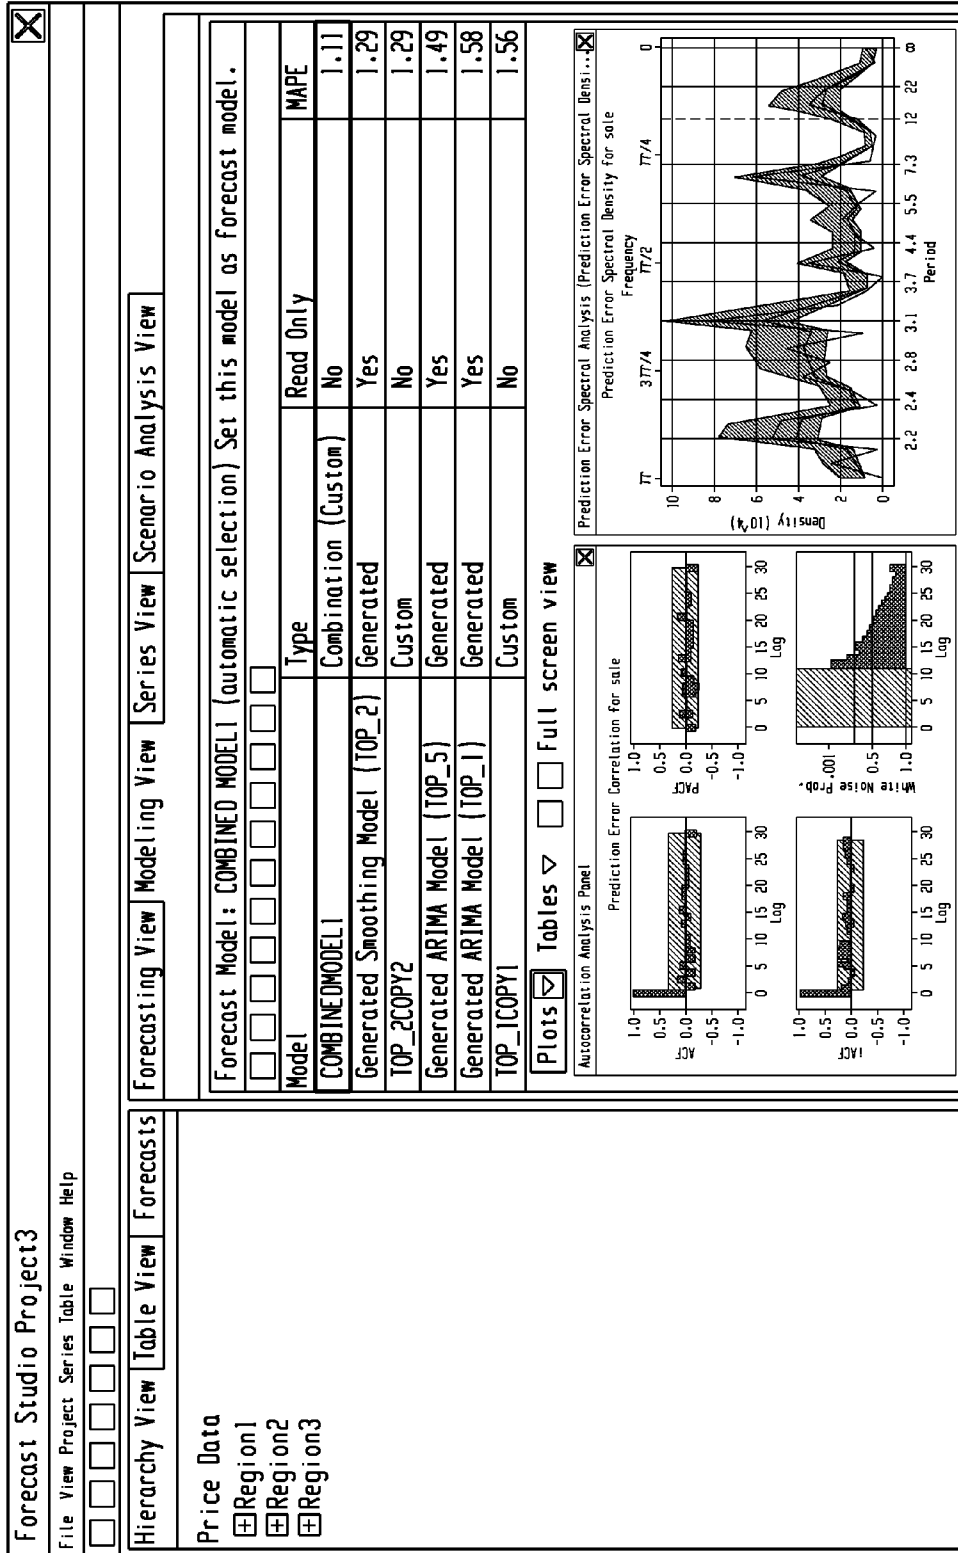


Fig. 24

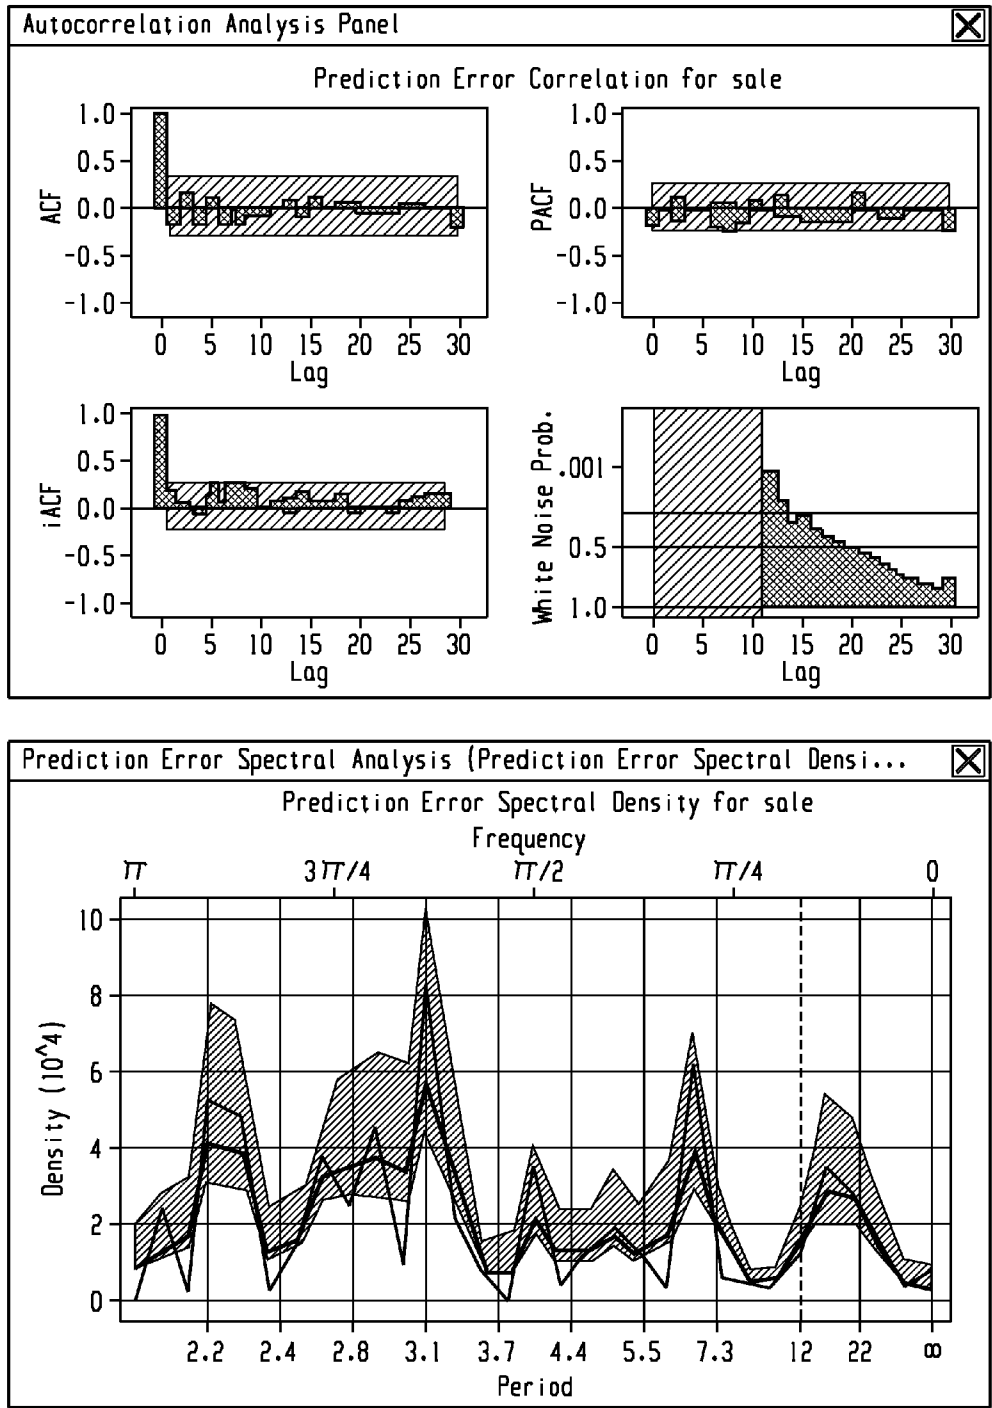


Fig. 25

Combine Models
✕

⚠ If this model is currently in use, editing it will cause all affected series to be automatically updated and the project to be reconciled. Forecasts may be affected.
 To see if and where the model is currently being used, you may search here.

Name:

Description:

Use	Model/	Type	MAPE	Weight
<input type="checkbox"/>	HPF2_83	Generated	5.29	0.0
<input checked="" type="checkbox"/>	HPF2_84	Generated	5.29	0.4 ↙
<input checked="" type="checkbox"/>	HPF2_84COPY1	Custom	5.29	0.5
<input type="checkbox"/>	HPF2_85	Generated	6.7	0.0
<input checked="" type="checkbox"/>	HPF2_85COPY1	Custom	6.7	0.1 ↙

Method of combination: ▼

Statistic: ▼

Treatment of missing values: ▼

Compute prediction error variance series: ▼

Forecast encompassing test: ▼ ⬆️⬆️

Percentage of missing forecast values in the combination horizon: ⬆️⬆️

Percentage of missing forecast values in the combination estimation region: ⬆️⬆️

Fig. 26

X
Combine Models

⚠ If this model is currently in use, editing it will cause all affected series to be automatically updated and the project to be reconciled. Forecasts may be affected.
 To see if and where the model is currently being used, you may search here.

Name:

Description:

Use	Model/	Type	MAPE	Weight
<input type="checkbox"/>	HPF2_83	Generated	5.29	0.0
<input checked="" type="checkbox"/>	HPF2_84	Generated	5.29	0.0
<input checked="" type="checkbox"/>	HPF2_84COPY1	Custom	5.29	0.0
<input type="checkbox"/>	HPF2_85	Generated	6.7	0.0
<input checked="" type="checkbox"/>	HPF2_85COPY1	Custom	6.7	0.0

Method of combination:

Statistic:

Treatment of missing

Compute prediction error

Forecast encompass
 Percentage of miss
 Percentage of miss

Fig. 27




Fig. 28




Fig. 29




Fig. 30

Create New Scenario

Name:

Description:

Choose a model to create a scenario:

Model	Type	Rank	MAPE	All Input...	No Input...
Generated Model (TOP_7)	Combination (Generated)	1	1.18		X
Generated Smoothing Model (TOP_6)	Generated	2	1.29		X
COMBINEDMODEL1	Combination (Custom)	2	1.29		X
TOP_2COPY1	Custom	2	1.29		X

Quick View...

Fig. 31




Fig. 32

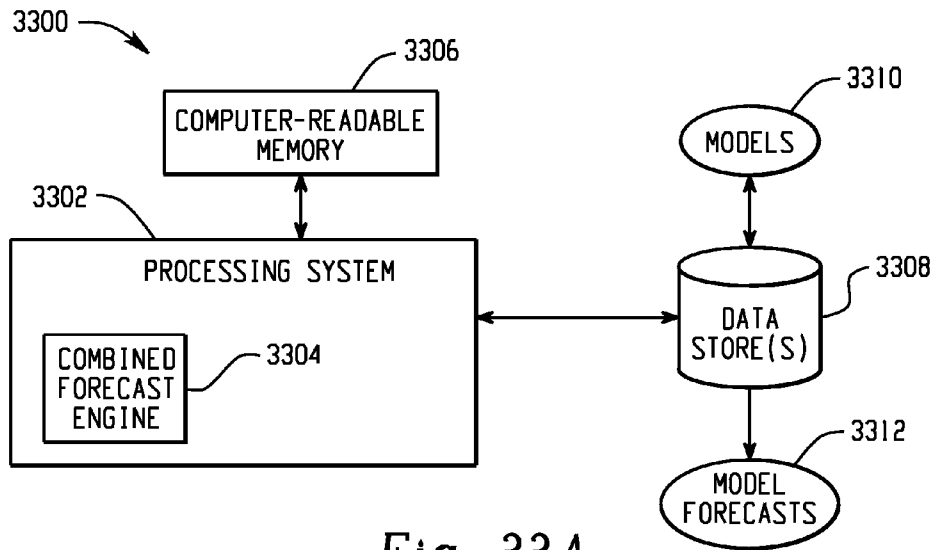


Fig. 33A

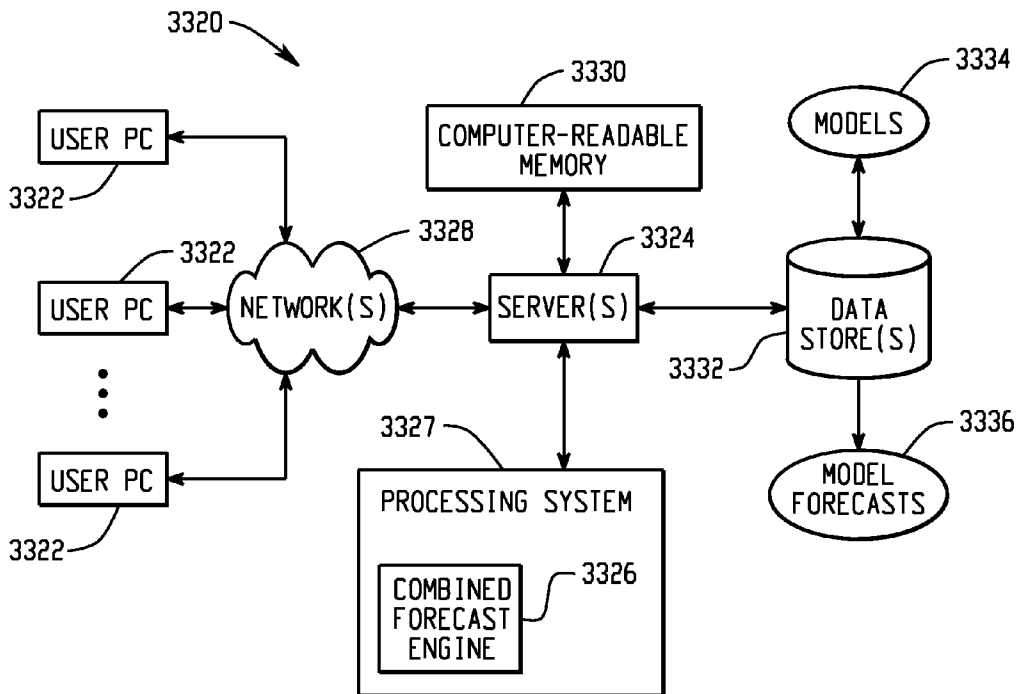


Fig. 33B




Fig. 33C

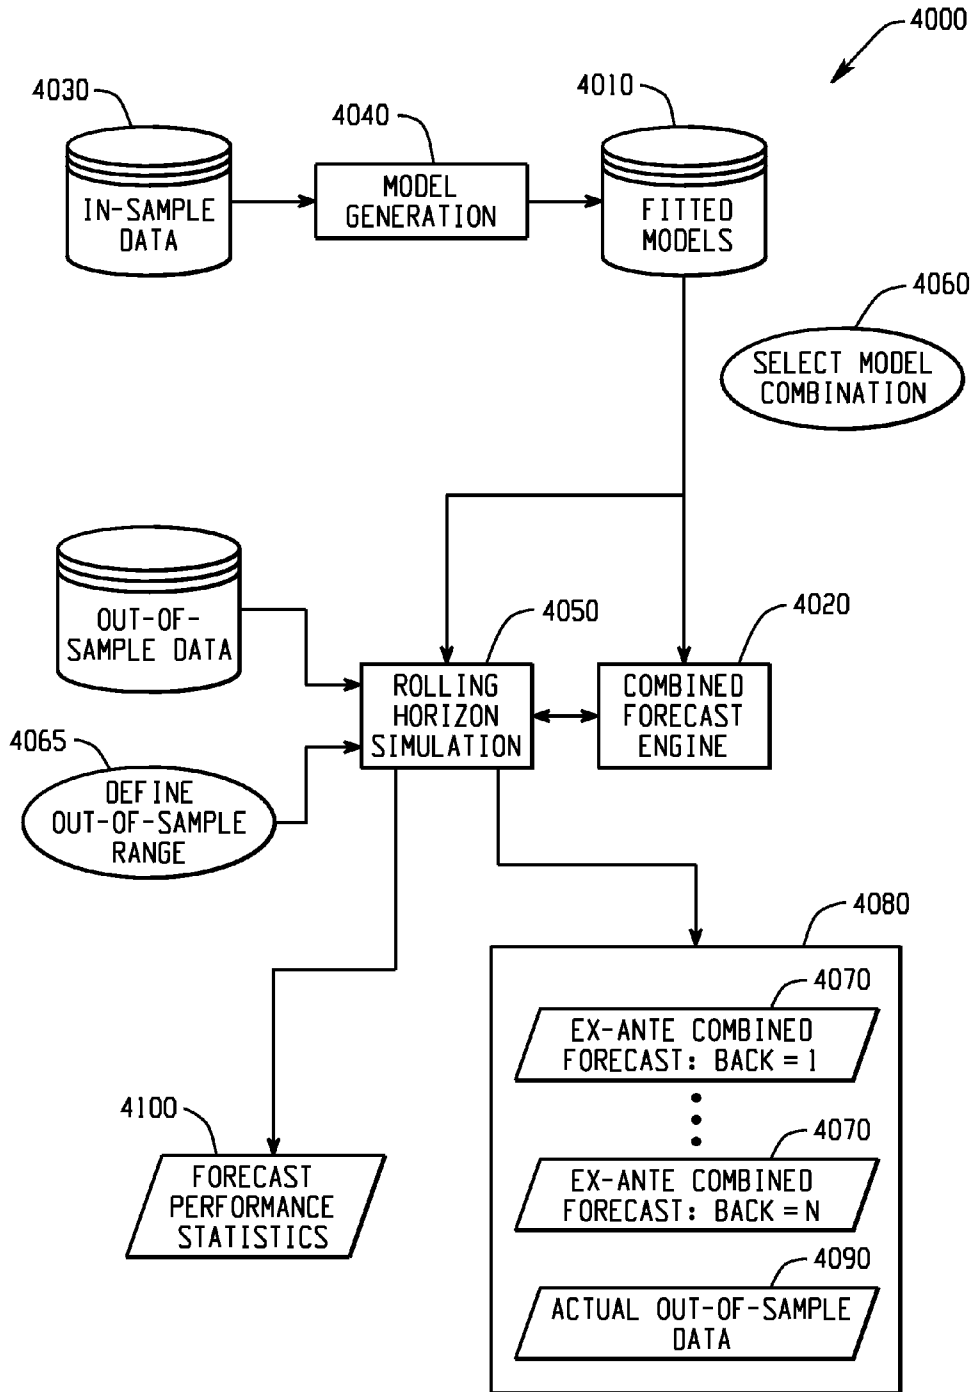


Fig. 34

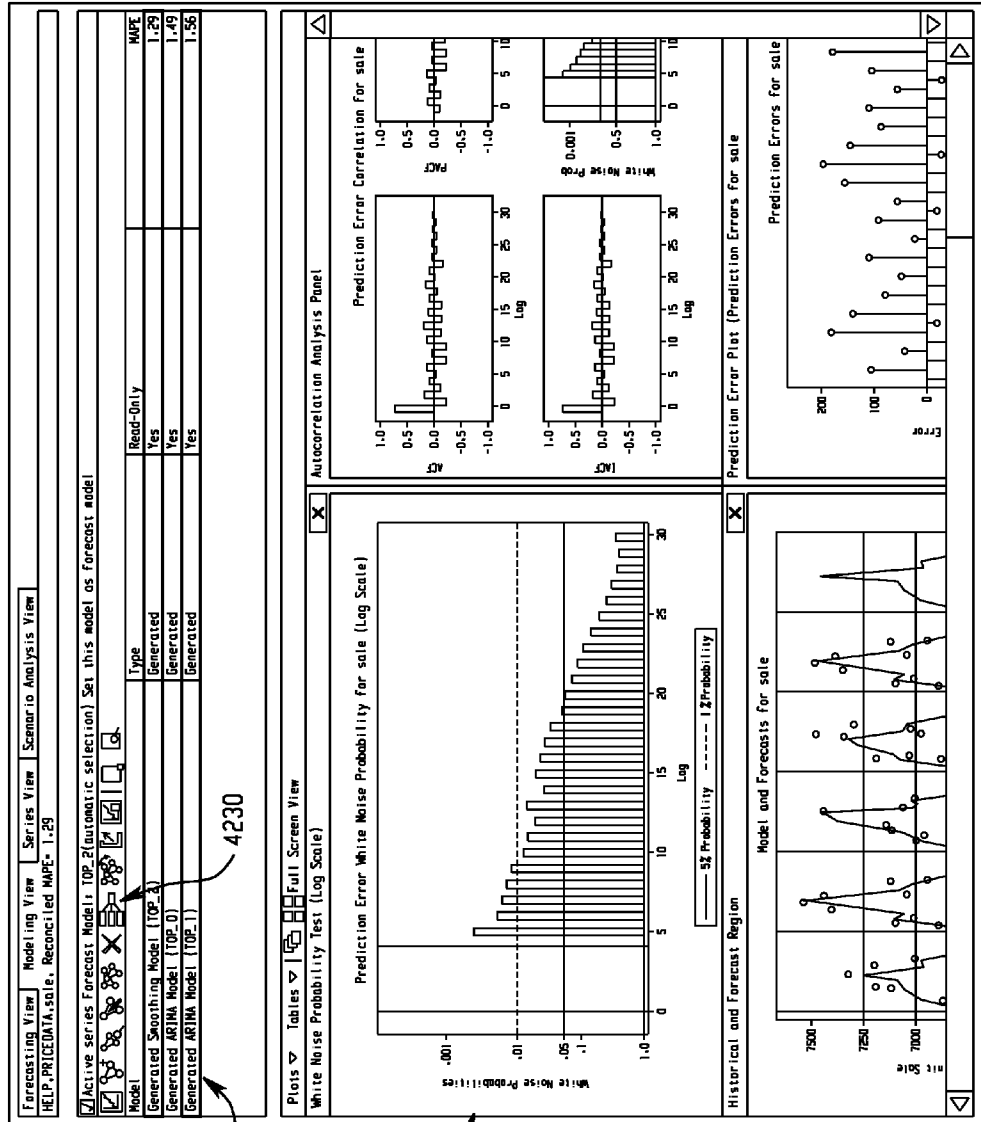


Fig. 35

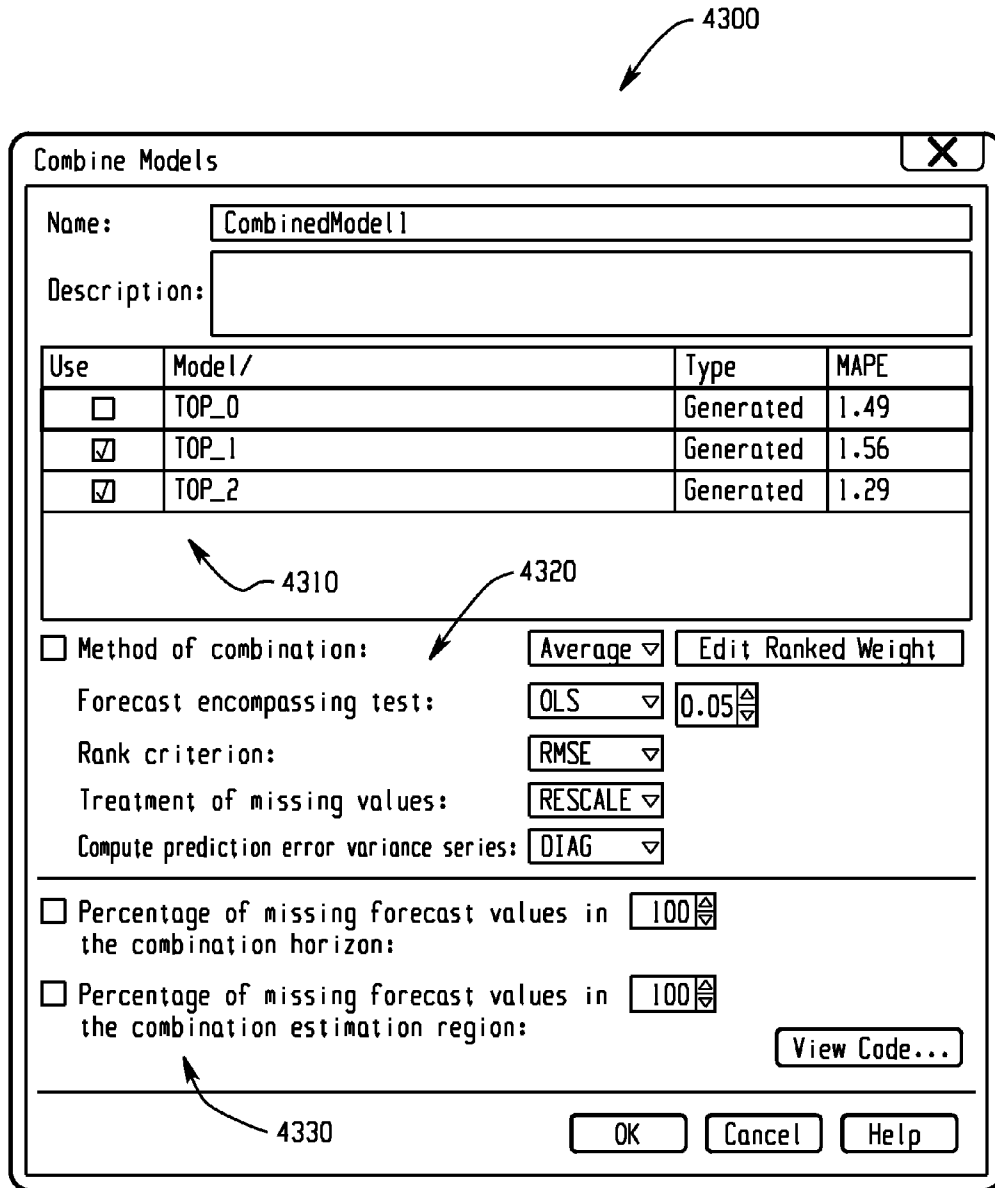




Fig. 36



4400


4200

Fig. 37



4500

Fig. 38



4500

Fig. 39





Fig. 40



4500

Fig. 41

COMPUTER-IMPLEMENTED SYSTEMS AND METHODS FOR TESTING LARGE SCALE AUTOMATIC FORECAST COMBINATIONS

[0001] This is a continuation-in-part of U.S. patent application Ser. No. 13/189,131, titled "Computer-Implemented Systems and Methods for Large Scale Automatic Forecast Combinations," filed on Jul. 22, 2011; this application also claims priority to U.S. Provisional Application No. 61/594,442, filed on Feb. 3, 2012. The entirety of these priority applications are incorporated herein by reference.

TECHNICAL FIELD

[0002] This document relates generally to computer-implemented forecasting and more particularly to testing a combined forecast that is generated using multiple forecasts.

BACKGROUND

[0003] Forecasting is a process of making statements about events whose actual outcomes typically have not yet been observed. A commonplace example might be estimation for some variable of interest at some specified future date. Forecasting often involves formal statistical methods employing time series, cross-sectional or longitudinal data, or alternatively to less formal judgmental methods. Forecasts are often generated by providing a number of input values to a predictive model, where the model outputs a forecast. While a well designed model may give an accurate forecast, a configuration where predictions of multiple models are considered when generating a forecast may provide even stronger forecast results.

SUMMARY

[0004] In accordance with the teachings herein, systems and methods are provided for evaluating a physical process with respect to one or more attributes of the physical process by combining forecasts for the one or more physical process attributes, where data for evaluating the physical process is generated over time. In one example, a forecast model selection graph is accessed, the forecast model selection graph comprising a hierarchy of nodes arranged in parent-child relationships. A plurality of model forecast nodes are resolved, where resolving a model forecast node includes generating a node forecast for the one or more physical process attributes. A combination node is processed, where a combination node transforms a plurality of node forecasts at child nodes of the combination node into a combined forecast. A selection node is processed, where a selection node chooses a node forecast from among child nodes of the selection node based on a selection criteria.

[0005] As another example, a system for storing evaluating a physical process with respect to one or more attributes of the physical process by combining forecasts for the one or more physical process attributes, where data for evaluating the physical process is generated over time is provided. The system may include one or more data processors and a computer-readable medium encoded with instructions for commanding the one or more data processors to execute steps. In the steps, a forecast model selection graph is accessed, the forecast model selection graph comprising a hierarchy of nodes arranged in parent-child relationships. A plurality of model forecast nodes are resolved, where resolving a model forecast node includes generating a node forecast for the one or more

physical process attributes. A combination node is processed, where a combination node transforms a plurality of node forecasts at child nodes of the combination node into a combined forecast. A selection node is processed, where a selection node chooses a node forecast from among child nodes of the selection node based on a selection criteria.

[0006] As a further example, a computer-readable storage medium may be encoded with instructions for commanding one or more data processors to execute a method. In the method, a forecast model selection graph is accessed, the forecast model selection graph comprising a hierarchy of nodes arranged in parent-child relationships. A plurality of model forecast nodes are resolved, where resolving a model forecast node includes generating a node forecast for the one or more physical process attributes. A combination node is processed, where a combination node transforms a plurality of node forecasts at child nodes of the combination node into a combined forecast. A selection node is processed, where a selection node chooses a node forecast from among child nodes of the selection node based on a selection criteria.

[0007] As an additional example, one or more computer-readable storage mediums may store data structures for access by an application program being executed on one or more data processors for evaluating a physical process with respect to one or more attributes of the physical process by combining forecasts for the one or more physical process attributes, where physical process data generated over time is used in the forecasts for the one or more physical process attributes. The data structures may include a predictive models data structure, the predictive models data structure containing predictive data model records for specifying predictive data models and a forecast model selection graph data structure, where the forecast model selection graph data structure contains data about a hierarchical structure of nodes which specify how the forecasts for the one or more physical process attributes are combined, where the hierarchical structure of nodes has a root node wherein the nodes include model forecast nodes, one or more model combination nodes, and one or more model selection nodes. The forecast model selection graph data structure may include model forecast node data which specifies for the model forecast nodes which particular predictive data models contained in the predictive models data structure are to be used for generating forecasts, model combination node data which specifies for the one or more model combination nodes which of the forecasts generated by the model forecast nodes are to be combined, and selection node data which specifies for the one or more model selection nodes model selection criteria for selecting, based upon model forecasting performance, models associated with the model forecast nodes or the one or more model combination nodes.

[0008] In accordance with the teachings herein, systems and methods are provided for evaluating performance of forecasting models. A plurality of forecasting models may be generated using a set of in-sample data. A selection of two or more forecasting models may be received from the plurality of forecasting models for use in generating a combined forecast. A set of actual out-of-sample data may be received. An ex-ante combined forecast may be generated for an out-of-sample period using the selected two or more forecasting models. The ex-ante combined forecast and the set of actual out-of-sample data may be provided for use in evaluating performance of the combined forecast.

BRIEF DESCRIPTION OF THE FIGURES

- [0009] FIG. 1 is a block diagram depicting a computer-implemented combined forecast engine.
- [0010] FIG. 2 is a block diagram depicting the generation of a combined forecast for a forecast variable.
- [0011] FIG. 3 is a block diagram depicting steps that may be performed by a combined forecast engine in generating a combined forecast.
- [0012] FIG. 4 depicts an example forecast model selection graph.
- [0013] FIG. 5 depicts an example forecast model selection graph including selection nodes, combination nodes, and model forecast nodes.
- [0014] FIG. 6 is a block diagram depicting example operations that may be performed by a combined forecast engine in combining one or more forecasts.
- [0015] FIG. 7 is a flow diagram depicting an example redundancy test in the form of an encompassing test.
- [0016] FIG. 8 depicts a forecast model selection graph having a selection node as a root node.
- [0017] FIG. 9 depicts a forecast model selection graph having a combination node as a root node.
- [0018] FIG. 10 depicts an example model repository for storing predictive models.
- [0019] FIG. 11 depicts a link between a forecast model selection graph and a model repository.
- [0020] FIG. 12 is a diagram depicting relationships among a forecast model selection graph data structure, a models data structure, and a combined forecast engine.
- [0021] FIG. 13 depicts an example forecast model selection graph data structure.
- [0022] FIG. 14 depicts an example node record.
- [0023] FIGS. 15-32 depict graphical user interfaces that may be used in generating and comparing combined forecasts.
- [0024] FIGS. 33A, 33B, and 33C depict example systems for use in implementing combined forecast engine.
- [0025] FIG. 34 is a block diagram depicting an example system for evaluating the performance of a combined forecast model using a rolling simulation analysis.
- [0026] FIGS. 35-41 depict example user interfaces for evaluating the performance of a combined forecast model using a rolling simulation analysis.

DETAILED DESCRIPTION

[0027] FIG. 1 is a block diagram depicting a computer-implemented combined forecast engine. FIG. 1 depicts a computer-implemented combined forecast engine 102 for facilitating the creation of combined forecasts and evaluation of created combined forecasts against individual forecasts as well as other combined forecasts. Forecasts are predictions that are typically generated by a predictive model based on one or more inputs to the predictive model. A combined forecast engine 102 combines predictions made by multiple models, of the same or different type, to generate a single, combined forecast that can incorporate the strengths of the multiple, individual models which comprise the combined forecast.

[0028] For example, a combined forecast may be generated (e.g., to predict a manufacturing process output, to estimate product sales) by combining individual forecasts from two linear regression models and one autoregressive regression model. The individual forecasts may be combined in a variety

of ways, such as by a straight average, via a weighted average, or via another method. To generate a weighted forecast, automated analysis of the individual forecasts may be performed to identify weights to generate an optimum combined forecast that best utilizes the available individual forecasts.

[0029] The combined forecast engine 102 provides a platform for users 104 to generate combined forecasts based on individual forecasts generated by individual predictive models 106. A user 104 accesses the combined forecast engine 102, which is hosted on one or more servers 108, via one or more networks 110. The one or more servers 108 are responsive to one or more data stores 112. The one or more data stores 112 may contain a variety of data that includes predictive models 106 and model forecasts 114.

[0030] FIG. 2 is a block diagram depicting the generation of a combined forecast for a forecast variable (e.g., one or more physical process attributes). The combined forecast engine 202 receives an identification of a forecast variable 204 for which to generate a combined forecast 206. For example, a user may command that the combined forecast engine 202 generate a combined forecast 206 of sales for a particular clothing item. To generate the combined forecast 206, the combined forecast engine 202 may identify a number of individual predictive models. Those individual predictive models may be provided historic data 208 as input, and those individual predictive models provide individual forecasts based on the provided historic data 208. The combined forecast engine 202 performs operations to combine those individual predictions of sales of the particular clothing item to generate the combined forecast of sales for the particular clothing item.

[0031] FIG. 3 is a block diagram depicting steps that may be performed by a combined forecast engine in generating a combined forecast. The combined forecast engine 302 receives a forecast variable 304 for which to generate a combined forecast as well as historic data 306 to be used as input to individual predictive models whose predictions become components of the combined forecast 308.

[0032] The combined forecast engine 302 may utilize model selection and model combination operations to generate a combined forecast. For example, the combined forecast engine 302 may evaluate a physical process with respect to one or more attributes of the physical process by combining forecasts for the one or more physical process attributes. Data for evaluating the physical process may be generated over time, such as time series data.

[0033] At 310, the combined forecast engine accesses a forecast model selection graph. A forecast model selection graph incorporates both model selection and model combination into a decision based framework that, when applied to a time series, automatically selects a forecast from an evaluation of independent, individual forecasts generated. The forecast model selection graph can include forecasts from statistical models, external forecasts from outside agents (e.g., expert predictions, other forecasts generated outside of the combined forecast engine 302), or combinations thereof. The forecast model selection graph may be used to generate combined forecasts as well as comparisons among competing generated forecasts to select a best forecast. A forecast model selection graph for a forecast decision process of arbitrary complexity may be created, limited only by external factors such as computational power and machine resource limits.

[0034] A forecast model selection graph may include a hierarchy of nodes arranged in parent-child relationships including a root node. The hierarchy may include one or more selection nodes, one or more combination nodes, and a plurality of model forecast nodes. Each of the model forecast nodes is associated with a predictive model. The combined forecast engine may resolve the plurality of model forecast nodes, as shown at **312**. Resolving a model forecast node includes generating a node forecast for the forecast variable **304** using the predictive model for the model forecast node. For example, a first model forecast node may be associated with a regression model. To resolve the first model forecast node, the combined forecast engine **302** provides the historic data **306** to the regression model, and the regression model generates a node forecast for the model forecast node. A second model forecast node may be associated with a human expert prediction. In such a case, computation by the combined forecast engine **302** may be limited, such as simply accessing the human expert's prediction from storage. A third model forecast node may be associated with a different combined model. To resolve the third model forecast node, the combined forecast engine **302** provides the historic data **306** to the different combined model, and the different combined model generates a node forecast for the model forecast node. Other types of models and forecasts may also be associated with a model forecast node.

[0035] At **314**, the combined forecast engine processes a combination node. In processing a combination node, the combined forecast engine **302** transforms a plurality of node forecasts at child nodes of the combination nodes into a combined forecast. For example, a combination node having three child nodes would have the node forecasts for those three child nodes combined into a combined forecast for the combination node. Combining node forecasts may be done in a variety of ways, such as via a weighted average. A weighted average may weight each of the three node forecasts equally, or the combined forecast engine **302** may implement more complex logic to identify a weight for each of the three node forecasts. For example, weight types may include a simple average, user-defined weights, rank weights, ranked user-weights, AICC weights, root mean square error weights, restricted least squares weights, OLS weights, and least absolute deviation weights.

[0036] At **316**, the combined forecast engine processes a selection node. In processing a selection node, the combined forecast engine **302** chooses a node forecast from among child nodes of the selection node based on a selection criteria. The selection criteria may take a variety of forms. For example, the selection criteria may dictate selection of a node forecast associated with a node whose associated model performs best in a hold out sample analysis.

[0037] As another example, metadata may be associated with models associated with node forecasts, where the metadata identifies a model characteristic of a model. The selection criteria may dictate selection of a node forecast whose metadata model characteristic best matches a characteristic of the forecast variable **304**. For example, if the forecast variable **304** tends to behave in a seasonal pattern, then the selection criteria may dictate selection of a node forecast that was generated by a model whose metadata identifies it as handling seasonal data. Other example model metadata characteristics include trending model, intermittent model, and transformed model.

[0038] As a further example, the selection criteria may dictate selection of a node forecast having the least amount of missing data. A node forecast may include forecasts for the forecast variable **304** for a number of time periods in the future (e.g., forecast variable at $t+1$, forecast variable at $t+2$, . . .). In some circumstances, a node forecast may be missing data for certain future time period forecasts (e.g., the node forecast is an expert's prediction, where the expert only makes one prediction at $t+6$ months). If a certain time period in the future is of specific interest, the selection criteria may dictate that a selected node forecast must not be missing a forecast at the time period of interest (e.g., when the time period of interest is $t+1$ month, the node forecast including the expert's prediction may not be selected).

[0039] As another example, the selection criteria may be based on a statistic of fit. For example, the combined forecast engine **302** may fit models associated with child nodes of a selection node with the historic data **306** and calculate statistics of fit for those models. Based on the determined statistics of fit, the combined forecast engine **302** selects the forecast node associated with the model that is a best fit.

[0040] The combined forecast engine **302** may continue resolving model forecast nodes **312** and processing combination and selection nodes **314**, **316** until a final combined forecast is generated. For example, the combined forecast engine may work from the leaves up to the root in the forecast model selection graph hierarchy, where the final combined forecast is generated at the root node.

[0041] FIG. 4 depicts an example forecast model selection graph. The forecast model selection graph includes a hierarchy of nodes arranged in parent-child relationships that includes a root node **402**. The forecast model selection graph also includes two model forecast nodes **404**. The model forecast nodes **404** may be associated with a model that can be used to forecast one or more values for a forecast variable. A model associated with a model forecast node **404** may also be a combined model or a forecast generated outside of the combined forecast engine, such as an expert or other human generated forecast. The model forecast nodes **404** are resolved to identify a node forecast (e.g., using an associated model to generate a node forecast, accessing an expert forecast from storage).

[0042] The forecast model selection graph also includes selection nodes **406**. A selection node may include a selection criteria for choosing a node forecast from among child nodes (e.g., model forecast nodes **404**) of the selection node **406**. Certain of the depicted selection nodes S_1 , S_2 , S_n do not have their child nodes depicted in FIG. 4.

[0043] FIG. 5 depicts an example forecast model selection graph including selection nodes, combination nodes, and model forecast nodes. To generate a combined forecast for the forecast model selection graph **500**, model forecast nodes **502** are resolved to generate node forecasts for one or more forecast variables (e.g., physical process attributes). With node forecasts resolved for the model forecast nodes **502**, a selection node **504** selects one of the node forecasts associated with the model forecast nodes **502** based on a selection criteria. For example, the selection criteria may dictate a model forecast based on metadata associated with a model used to generate the model forecast at the model forecast node **502**.

[0044] Additional model forecast nodes **506** may be resolved to generate node forecasts at those model forecast nodes **506**. A first combined forecast node **508** combines a model forecast associated with model forecast node $MF1_1$

and the model forecast at the selection node **504** to generate a combined forecast at the combination node **508**. A second combined forecast node **510** combines a model forecast associated with model forecast node **MF2_1** and the model forecast at the selection node **504** to generate a combined forecast at the combination node **510**. Another selection node **512** selects a model forecast from one of the two combination nodes **508**, **510** based on a selection criteria as the final combined forecast for the forecast model selection graph **500**.

[0045] A forecast model selection graph may take a variety of forms. For example, the forecast model selection graph may be represented in one or more records in a database or described in a file. In another implementation, the forecast model selection graph may be represented via one or more XML based data structures. The XML data structures may identify the forecast sources to combine, diagnostic tests used in the selection and filtering of forecasts, methods for determining weights to forecasts to be combined, treatment of missing values, and selection of methods for estimating forecast prediction error variance.

[0046] FIG. 6 is a block diagram depicting example operations that may be performed by a combined forecast engine in combining one or more forecasts (e.g., when processing a combination node). At **602**, an initial set of model forecasts is identified. In some implementations, all identified model forecasts may be combined to create a combined forecast. However, in some implementations, it may be desirable to filter the models used in creating a combined forecast. For example, at **604**, the set of model forecasts may be reduced at **604** based on one or more forecast candidate tests. The forecast candidate tests may take a variety of forms, such as analysis of the types of models used to generate the model forecasts identified at **602** and characteristics of the forecast variable. For example, if the forecast variable is a trending variable, the candidate tests may eliminate model forecasts generated by models that are designed to handle seasonal data.

[0047] At **606**, the set of model forecasts may be reduced based on one or more forecast quality tests. Forecast quality tests may take a variety of forms. For example, forecast quality tests may analyze missing values of model forecasts. For example, model forecasts may be filtered from the set if the model forecasts have missing values in an area of interest (e.g., a forecast horizon). In another example, a model forecast may be filtered from the set if it is missing more than a particular % of values in the forecast horizon.

[0048] At **608**, the set of model forecasts may be reduced based on redundancy tests. A redundancy test may analyze models associated with model forecasts nodes to identify robust models, and those models having a high degree of redundancy (e.g., models that are producing forecasts that are statistically too similar). Model forecasts having a high degree of redundancy may be excluded from the combined model being generated.

[0049] In addition to generating a combined forecast, certain statistics for a combined forecast may be determined. For example, a prediction error variance estimate may be calculated. The prediction error variance estimate may incorporate pair-wise correlation estimates between the individual forecast prediction errors for the predictions that make up the combined forecast and their associated prediction error variances.

[0050] FIG. 7 is a flow diagram depicting an example redundancy test in the form of an encompassing test. The set of model forecasts is shown at **702**. At **704**, each model in the set **702** is analyzed to determine whether the current model forecast is redundant (e.g., whether the information in the current model forecast is already represented in the continuing set of forecasts **706**). If the current model forecast is redundant, then it is excluded. If the current model forecast is not redundant, then it remains in the set of forecasts **706**.

[0051] With reference back to FIG. 6, at **610**, weights are assigned to the model forecasts remaining in the set. Weights may be assigned using a number of different algorithms. For example, weights may be assigned as a straight average of the set of remaining model forecasts, or more complex processes may be implemented, such as a least absolute deviation procedure. At **612**, the weighted model forecasts are aggregated to generate a combined forecast.

[0052] FIG. 8 depicts a forecast model selection graph having a selection node as a root node. A number of node forecasts **802** are resolved (e.g., by generating node forecasts using a model, accessing externally generated forecasts from computer memory). A combination node **804** combines the model forecasts of child nodes **806** of the combination node **804**. A selection node **808** selects a forecast from among the combination node **804** and model forecasts at child nodes **810** of the selection node **808** based on a selection criteria.

[0053] FIG. 9 depicts a forecast model selection graph having a combination node as a root node. A number of node forecasts **902** are resolved (e.g., by generating node forecasts using a model, accessing externally generated forecasts from memory). A selection node **904** selects a model forecast from the child nodes **906** of the selection node. A combination node **908** combines the model forecast from the selection node **904** and model forecasts at child nodes **910** of the combination node **908** to generate a combined forecast.

[0054] As noted previously, a model forecast node may be associated with a predictive model that is used to generate a model forecast for the model forecast node. In one embodiment, the predictive models may be stored in a model repository for convenient access. FIG. 10 depicts an example model repository for storing predictive models. The model repository **1002** includes a number of model records **1004**. A model record may contain model data for implementing a predictive model **1006**. In another embodiment, a model record **1004** may contain a reference to where data for implementing the predictive model **1006** can be found (e.g., a file location, a pointer to a memory location, a reference to a record in a database). Other example details of a model repository are described in U.S. Pat. No. 7,809,729, entitled "Model Repository," the entirety of which is herein incorporated by reference.

[0055] A model repository configuration may streamline the data contained in a forecast model selection graph. FIG. 11 depicts a link between a forecast model selection graph and a model repository. A forecast model selection graph **1102** includes a number of model forecast nodes **MF1**, **MF2**, **MF3**, **MF4**, a selection node **S1**, and a combination node **C1**. The model forecast nodes are resolved to generate node forecasts. One of the model forecast nodes, **MF4**, is associated with a model record **1104**. For example, model forecast node, **MF4**, may contain an index value for the model record **1104**. The model record is stored in the model repository **1104** and may contain data for implementing a predictive model to generate the node forecast, or the model record may contain a

reference to the location of such data **1108**, such as a location in a the model repository **1106**. When the model forecast node, MF4, is to be resolved, the model record **1104** is located based on the index identified by the model forecast node, MF4. Data for the desired predictive model **1108** to be used to generate the node forecast is located in the model repository **1106** based on data contained in the model record **1104**.

[0056] FIG. 12 is a diagram depicting relationships among a forecast model selection graph data structure, a models data structure, and a combined forecast engine. A forecast model selection graph data structure **1202** and a models data structure **1204** may be stored on one or more computer-readable storage mediums for access by an application program, such as a combined forecast engine **1206** being executed on one or more data structures. The data structures **1202**, **1204** may be used as part of a process for evaluating a physical process with respect to one or more attributes of the physical process by combining forecasts for the one or more physical process attributes. Physical process data generated over time (e.g., time series data) may be used in the forecasts for the one or more physical attributes.

[0057] The forecast model selection graph data structure **1202** may contain data about a hierarchical structure of nodes which specify how forecasts for the one or more physical attributes are combined, where the hierarchical structure of nodes has a root node, and where the nodes include one or more selection nodes **1208**, one or more model combination nodes **1210**, and model forecast nodes **1212**. The forecast model selection graph data structure **1202** may include selection node data **1208** that specifies, for the one or more model selection nodes, model selection criteria for selecting, based upon model forecasting performance, models associated with the model forecast nodes or the one or more model combination nodes. The forecast model selection graph data structure **1202** may also include model combination node data **1210** that specifies, for the one or more model combination nodes, which of the forecasts generated by the model forecast nodes are to be combined.

[0058] The forecast model selection graph data structure **1202** may also include model forecast node data **1212** that specifies, for the model forecast nodes, which particular predictive data models contained in the models data structure are to be used for generating forecasts. For example, the model forecast node data **1212** may link which stored data model is associated with a specific model forecast node, such as via an index **1214**. The stored data model **1216** identified by the model forecast node data **1212** may be accessed as part of a resolving process to generate a node forecast for a particular node of the model forecast selection graph. The combined forecast engine **1206** may process the forecast model selection graph data structure **1202**, using stored data models **1216** identified by the models data structure **1204** via the link between the model forecast node data **1212** and the models data structure **1204** to generate a combined forecast **1218**.

[0059] FIG. 13 depicts an example forecast model selection graph data structure. In FIG. 13, the forecast model selection graph data structure **1302** is a data structure that includes a number of node records **1304** as sub-data structures. The node records **1304** may each be descriptive of a model forecast node, a combination node, or a selection node. Each of the node records **1304** includes data.

[0060] FIG. 14 depicts an example node record. For example, the node record **1402** may contain data related to the type of a node **1404** and data for the node to be processed,

such as an identification of a model to generate a node forecast **1406** or a selection criteria for selecting among child nodes. Additionally, a node record **1402** may include structure data that identifies, in whole or in part, a position of a node in the forecast model selection graph. For example, the node record data may contain data identifying child nodes **1408** of a node and a parent node **1410** of the node. The node record **1402** may also identify a node as a root or a leaf node or the exact position of a node in the forecast model selection graph hierarchy (e.g., a pre-order or a post-order value).

[0061] FIGS. 15-32 depict graphical user interfaces that may be used in generating and comparing combined forecasts. FIG. 15 depicts an example graphical user interface for identifying parameters related to time, where a user may specify parameters such as a time interval, a multiplier value, a shift value, a seasonal cycle length, and a date format.

[0062] FIG. 16 depicts an example forecasting settings graphical user interface for identifying parameters related to data preparation, where a user may specify how to prepare data for forecasting. Example settings include how to interpret embedded missing values, which leading or trailing missing values to remove, which leading or trailing zero values to interpret as missing, and whether to ignore data points earlier than a specified date.

[0063] FIG. 17 depicts an example forecasting settings graphical user interface for identifying diagnostics settings. Example settings include intermittency test settings, seasonality test settings, independent variable diagnostic settings, and outlier detection settings. Such diagnostic settings may be used in a variety of contexts, including processing of combination nodes of a forecast model selection graph.

[0064] FIG. 18 depicts an example forecasting settings graphical user interface for identifying model generation settings. Example settings include identifications of which models to fit to each time series. Example models include system-generated ARIMA models, system-generated exponential smoothing models, system-generated unobserved components models, and models from an external list. Such model generation settings may be used in a variety of contexts, including with model forecast nodes of a forecast model selection graph.

[0065] FIG. 19 depicts an example forecasting settings graphical user interface for identifying model selection settings. Example settings include whether to use a holdout sample in performing model selection and a selection criteria for selecting a forecast. Such model selection settings may be used in a variety of contexts, including with model selection nodes of a forecast model selection graph.

[0066] FIG. 20 depicts an example forecasting settings graphical user interface for identifying model forecast settings. Example settings include a forecast horizon, calculation of statistics of fit settings, confidence limit settings, negative forecast settings, and component series data set settings.

[0067] FIG. 21 depicts an example forecasting settings graphical user interface for identification of hierarchical forecast reconciliation settings. Using the user interface of FIG. 21, a preference for reconciliation of a forecast hierarchy may be selected along with a method for performing the reconciliation, such as a top-down, bottom-up, or middle-out process.

[0068] FIG. 22 depicts an example forecasting settings graphical user interface for combined model settings. The combined model settings user interface allows selection of a combine model option. The user interface of FIG. 22 also includes an advanced options control. FIG. 23 depicts an

example graphical user interface for specification of advanced combined model settings. The settings of FIG. 23 may be used in a variety of contexts, including in processing of a combination node of a forecast model selection graph.

[0069] Example settings for advanced combined model settings include a method of combination setting. Example parameters include a RANKWGT setting, where a combined forecast engine analyzes the forecasts to be combined and assigns weights to those forecasts based on the analysis. In another example, the RANKWGT option may accept a set of user-defined weights that are substituted for the automatic rank weight settings for each ordinal position in the ranked set. The combined forecast engine analyzes and ranks the forecasts to be combined and then assigns the user-defined weights to the forecasts according to the forecast's ordinal position in the ranking. As another option, a user may directly assign weights to the individual forecasts, and as a further option, a mean-average of the individual forecasts may be utilized.

[0070] The advanced settings interface also includes an option for directing that a forecast encompassing test be performed. When selected, the combined forecast engine ranks individual forecasts for pairwise encompassing elimination. The advanced setting interface further includes options related to treatment of missing values. For example, a rescale option may be selected for weight methods that incorporate a sum-to-one restriction for combination weights. A further option directs a method of computation of prediction error variance series. This option is an allowance for treating scenarios where the cross-correlation between two forecast error series is localized over segments of time when it is assumed that the error series are not jointly stationary. DIAG may be the default setting, while ESTCORR presumes that the combination forecast error series are jointly stationary and estimates the pairwise cross-correlations over the complete time spans.

[0071] FIG. 24 depicts a model view graphical user interface. Using the model view, a user can evaluate combined model residuals. The user interface is configured to enable graphical analysis of a model residual series plot, residual distribution, time domain analysis (e.g., ACF, PACF, IACF, white noise), frequency domain analysis (e.g., spectral density, periodogram). The user interface also enables exploration of parameter estimates, statistics of fit (e.g., RMSE, MAPE, AIC), and bias statistics. FIG. 25 depicts example graphs that may be provided by a model view graphical interface. Other options provided by a model view graphical user interface may include options for managing model combinations, such as adding a model for consideration, editing a previously added model, copying a model, and deleting a model (e.g., a previously added combined model).

[0072] FIG. 26 depicts an example graphical user interface for manually defining a combined model. For example, a manually defined combined model may be utilized with a model forecast node in a forecast model selection graph. The graphical user interface may be configured to receive a selection of one or more models to be combined, weights to be applied to those combined models in generating the combination, as well as other parameters. For example, FIG. 27 depicts the manual entry of ranked weights to be applied to the selected models after they are ranked by a combined forecast engine.

[0073] FIG. 28 depicts an example interface for comparing models. The example interface may be accessed via a model view interface. The present interface enables comparison of selected model combinations in graphical form. FIG. 29 depicts a table that enables comparison of selected model combinations statistically in text form.

[0074] FIG. 30 depicts a graphical user interface for performing scenario analysis using model combinations. Using scenario analysis, scenarios can be generated, where an input time series can be varied to better understand possible future outcomes and to evaluate a model's sufficiency to different input values. A create new scenario menu may be accessed by selecting a new control in a scenario analysis view. Using the create new scenario menu, shown in further detail in FIG. 31, a model is selected for analysis. A scenario is generated, and a graph depicting results of the scenario analysis is displayed, such as the graph of FIG. 32.

[0075] The systems and methods described herein may, in some implementations, be utilized to achieve one or more of the following benefits. For example, forecast accuracy may often be significantly improved by combining forecasts of individual predictive models. Combined forecasts also tend to produce reduced variability compared to the individual forecasts that are components of a combined forecast. The disclosed combination process may automatically generate forecast combinations and vet them against other model and expert forecasts as directed by the forecast model selection graph processing. Combined forecasts allow for better predicting systematic behavior of an underlying data generating process that cannot be captured by a single model forecast alone. Frequently, combinations of forecasts from simple models outperform a forecast from a single, complex model.

[0076] FIGS. 33A, 33B, and 33C depict example systems for use in implementing an enterprise data management system. For example, FIG. 33A depicts an exemplary system 3300 that includes a standalone computer architecture where a processing system 3302 (e.g., one or more computer processors) includes a combined forecast engine 3304 being executed on it. The processing system 3302 has access to a computer-readable memory 3306 in addition to one or more data stores 3308. The one or more data stores 3308 may include models 3310 as well as model forecasts 3312.

[0077] FIG. 33B depicts a system 3320 that includes a client server architecture. One or more user PCs 3322 accesses one or more servers 3324 running a combined forecast engine 3326 on a processing system 3327 via one or more networks 3328. The one or more servers 3324 may access a computer readable memory 3330 as well as one or more data stores 3332. The one or more data stores 3332 may contain models 3334 as well as model forecasts 3336.

[0078] FIG. 33C shows a block diagram of exemplary hardware for a standalone computer architecture 3350, such as the architecture depicted in FIG. 33A that may be used to contain and/or implement the program instructions of system embodiments of the present invention. A bus 3352 may serve as the information highway interconnecting the other illustrated components of the hardware. A processing system 3354 labeled CPU (central processing unit) (e.g., one or more computer processors), may perform calculations and logic operations required to execute a program. A processor-readable storage medium, such as read only memory (ROM) 3356 and random access memory (RAM) 3358, may be in communication with the processing system 3354 and may contain one or more programming instructions for performing the

method of implementing a combined forecast engine. Optionally, program instructions may be stored on a computer readable storage medium such as a magnetic disk, optical disk, recordable memory device, flash memory, or other physical storage medium. Computer instructions may also be communicated via a communications signal, or a modulated carrier wave.

[0079] A disk controller **3360** interfaces one or more optional disk drives to the system bus **3352**. These disk drives may be external or internal floppy disk drives such as **3362**, external or internal CD-ROM, CD-R, CD-RW or DVD drives such as **3364**, or external or internal hard drives **3366**. As indicated previously, these various disk drives and disk controllers are optional devices.

[0080] Each of the element managers, real-time data buffer, conveyors, file input processor, database index shared access memory loader, reference data buffer and data managers may include a software application stored in one or more of the disk drives connected to the disk controller **3360**, the ROM **3356** and/or the RAM **3358**. Preferably, the processor **3354** may access each component as required.

[0081] A display interface **3368** may permit information from the bus **3352** to be displayed on a display **3370** in audio, graphic, or alphanumeric format. Communication with external devices may optionally occur using various communication ports **3372**.

[0082] In addition to the standard computer-type components, the hardware may also include data input devices, such as a keyboard **3373**, or other input device **3374**, such as a microphone, remote control, pointer, mouse and/or joystick.

[0083] FIGS. 34-41 depict example systems and methods for evaluating the performance of a combined forecast model using rolling simulations. Before using a time series model to forecast a time series, it is often important to evaluate the model's performance by validating the model's ability to forecast previously acquired data. Such performance measures are commonly referred to as ex-ante model performance measures. Ex-ante (i.e., before the fact) model performance measures are similar to ex-post (i.e., after the fact) forecast performance measures; however, ex-post forecast performance measures are used to evaluate the performance of forecasts regardless of the source (e.g., model-based, judgment-based, possible adjustments, etc.) Rolling simulations can be used to perform ex-ante model performance by repeating the analyses over several forecast origins and lead times.

[0084] FIG. 34 is a block diagram depicting an example system **4000** for evaluating the performance of a combined forecast model using a rolling simulation analysis. The combined forecast model may be generated from a plurality of individual forecast models **4010** using a combined forecast engine **4020**, for example using one or more of the systems and methods described above with reference to FIGS. 1-33. The plurality of forecast models **4010** may, for example, be generated based on sampled time-series data **4030** using a hierarchical time series forecasting system **4040**, such as the Forecast Studio and Forecast Server software sold by SAS Institute Inc. of Cary, N.C.

[0085] The example system **4000** includes a rolling simulation engine **4050** that works in combination with the combined forecast engine **4020** to select a model combination **4060** and evaluate its performance over an out-of-sample range **4065**. For instance, the rolling simulation engine **4050** and/or the combined forecast engine **4020** may present a user interface to select two or more of the individual forecast

models **4010** for use in generating a combined forecast. In addition, a user interface may be provided by the rolling simulation engine **4050** to define the out-of-sample range **4065** (e.g., BACK=) over which the combined forecast is to be evaluated. The combined forecast is then repeated over the entire out-of-sample range, e.g., over the range of BACK=0 to a specified "back" value. The resulting ex-ante forecasts **4070** (e.g., the forecasts for each BACK=value in the out-of-sample range) may then be displayed on a user interface **4080** along with the actual out-of-sample data **4090**, such that the ex-ante forecasts may be visually compared with the actual out-of-sample data over the specified period. In certain embodiments, the rolling simulation engine **4050** may also be used to simulate and display ex-post forecasts for the combined model over a rolling simulation horizon (e.g., LEAD=value).

[0086] The rolling simulation engine **4050** may also calculate one or more performance statistics **4100** based on statistical comparisons of the actual out-of-sample data and the ex-ante forecasts over the specified out-of-sample period **4065**. The performance statistics **4100** may, for example, be displayed on a different tab of the user interface **4080**. The performance statistics **4100** may include statistics that provide an indication of the average error between the forecast **4070** and the actual out-of-sample data **4090**, such as mean, mean absolute percentage error (MAPE), mean absolute error (MAE), median absolute deviation (MAD) and/or MAD/Mean ratio calculations. Because these statistics indicate average errors between the actual and forecasted data, a low variance for the statistics over different out-of-sample ranges (e.g., different BACK=values) may provide an indication that the combined model is forecasting with required accuracy. On the other hand, a high variation in the forecast performance statistics **4100** over an out-of-sample range may indicate that the combined model is not performing at the required accuracy and may therefore result in larger errors when used over a wider horizon.

[0087] In certain embodiments, the rolling simulation engine **4050** and/or the combined forecast engine **4020** may be further configured to dynamically adjust one or more characteristics of the combined model based on the forecast horizon. For instance, the combined model list may be re-run for each (BACK, LEAD) pair such that the set of model candidates included in the combination can change from one pair of (BACK, LEAD) values to the next depending on how the combined model list is defined. For example, certain forecast quality tests, such as an encompassing test can be specified. An encompassing test examines forecasts produced by the component models of a combined model to determine whether the forecasts produced by one or more of the component models that make up the combined model are redundant when processing given data. For each (BACK, LEAD) pair, an encompassing test may be performed on the combined model based on the out-of-sample input data to be provided to the combined model for that (BACK, LEAD) pair. Should one or more of the component models be found to be redundant, those component models can be omitted from the combined model for that (BACK, LEAD) pair, with the weightings of the component models of the combined model being automatically adjusted accordingly. Component models can be dropped for other quality reasons as well, such as the inability or ineffectiveness of those component models in handling missing values present in the candidate model forecast's historical period and/or its horizon for a particular

(BACK, LEAD) pair. In addition, component models may be dropped if the chosen method of weight estimation fails to produce a valid weight estimate for that model's forecast. Component model weights may also adaptively rescale over the span of the (BACK, LEAD) period to account for candidate models with missing values in their forecasts.

[0088] An example operation of the system **4000** of FIG. **34** is further detailed with reference to the exemplary screen shots set forth in FIGS. **35-41**. With reference first to FIG. **35**, this figure illustrates an example interface **4200** which may be generated to select a model or model combination for the rolling simulation analysis. The interface **4200** includes a model selection region **4210** that lists the fitted individual models that are available for analysis. The interface **4200** may also include a forecast summary region **4220** that displays forecast data and/or other information relating to a selected one of the fitted models listed in the selection region **4210**. In order to specify a combined model for evaluation, the interface **4200** may enable the user to highlight two or more of the individual models from the selection region **4210** and then select a combine models icon **4230**. Selecting the combine model icon **4230** may cause a combine models interface to be displayed, for example in a pop-up window, as illustrated in FIG. **36**.

[0089] FIG. **36** depicts an example of an interface **4300** which may be generated to define the characteristics of a combined model for the rolling simulation analysis. The interface **4300** includes a model selection region **4310** for selecting the individual models for combination. As illustrated, the model selection region **4310** may list the model names along with one or more characteristics of the model, such as the model type and a model performance statistic (e.g., MAPE). In the illustrated example, two models (TOP_1 and TOP_2) are selected for combination.

[0090] The model combination interface **4300** provides a plurality of user-editable fields **4320** for defining the characteristics of the combined model. For instance, the user-editable fields **4320** may include a field for defining the model combination method (e.g., average or specify weights), a field for editing any specific weights applied to each individual model, a field to define how the combined model will treat missing values, and/or other fields for defining the characteristics of the combined model. In addition, the interface **4300** may further include fields **4330** that may be used to define the percentage of missing forecast values in the combination horizon and the percentage of missing forecast values in the combination estimation region. The interface **4300** may also provide regions for naming the combined model and providing a model description.

[0091] As illustrated in FIG. **37**, once the characteristics of the combined model have been established using the model combination interface, the model selection interface is updated to list the newly defined combined model **4400**. The combined model **4400** may then be selected for rolling simulation analysis, as shown in FIG. **37**. The rolling simulation analysis may, for example, be executed on the selected combined model **4400** by selecting a rolling simulation icon **4410** on the interface **4200**.

[0092] FIG. **38** depicts an example rolling simulation interface **4500** that may be displayed upon executing a rolling simulation analysis on a selected combined model. The interface **4500** includes a user-editable region **4510** for defining the number of out-of-sample observations (e.g., the number of BACK=values) to be included in the simulation. Upon

selecting a number of out-of-sample observations, the user may begin the simulation, for example by selecting a Run Simulation icon **4520**. In the illustrated example, a simulation has been executed after selecting 6 out-of-sample observations. The rolling simulation engine then simulates an ex-ante forecast for each of out-of-sample observation (e.g., for each BACK=value 1-6), and plots the resulting forecasts on a display region **4530** of the interface **4500** along with the actual out-of-sample values. In this way, the rolling simulation interface **4500** enables the user to graphically view the predictions of the combined model at various forecast origins and compare the predictions to the actual out-of-sample data values.

[0093] For instance, in the illustrated example, six ex-ante forecasts (BACK:1 through BACK:6) are plotted in the display region **4530** for comparison with the actual out-of-sample data values. Specifically, in the first out-of-sample observation (BACK:1) the combined forecast is generated one period into the past (December 02), in the next out-of-sample observation (BACK:2) the combined forecast is generated two periods into the past (December 02 thru November 02), and so on through the sixth out-of-sample observation (BACK:6).

[0094] In addition, for further comparison, the rolling simulation interface **4500** of FIG. **38** also includes a numerical display region **4540** that displays the forecast values for each out-of-sample observation (BACK:1 thru BACK:6) along with the actual values of the out-of-sample data. The forecasted values in the illustrated example are set forth in the numerical display region **4540** in bold text.

[0095] In certain embodiments, the rolling simulation engine may calculate an optimal number of out-of-sample observations to automatically populate the interface field **4510** with a back range default value. The default value for the back range field **4510** may, for example, be calculated using the following formula:

$$\text{Default} = \min(\text{lead}, \min(\min(\max(\text{seasonality}, 4), 52), \min(t-6, \max((\text{int})(0.1*t), 1))))$$

[0096] where:

[0097] lead—lead option value

[0098] T—length of series

[0099] S—length of seasonality

[0100] Minimum value of BackRange is: Minimum=1

[0101] Maximum value of BackRange is: Maximum=(T minus 6)

$$\text{defaultBackRange} = \text{Math.min}(\text{lead}, \text{Math.min}(\text{Math.min}(\text{Math.max}(\text{seasonality}, 4), 52), \text{Math.min}(t-6, \text{Math.max}((\text{int})(0.1*t), 1))))$$

[0102] As illustrated in FIG. **39**, the rolling simulation interface **4500** may also include one or more fields for use in generating and displaying one or more forecast performance statistics. The simulation statistics fields may, for example, be included in a separate tab **4600** of the rolling simulation interface **4500**, as illustrated in FIG. **39**. The simulation statistics tab **4600** includes a field **4610** for selecting a particular forecast performance statistic for identifying the error between the forecasted values and the actual out-of-sample values, such as a Mean, MAE, MAPE, MAD or MAD/Mean ratio calculation. A graph of the selected performance statistic at each lead time is displayed in a display region **4620** of the interface.

[0103] In addition, the simulation statistics tab **4600** may also include a numerical display region **4630** that displays the performance statistic values for each lead time. In the illustrated example, the numerical display region **4630** includes values for a plurality of different performance statistics calculated by the rolling simulation engine. In this way, the user may simultaneously view the values for multiple performance statistics for the combined model in the numerical display region **4630**, and select a particular one of the performance statistic for graphical display **4620**.

[0104] The graph **4620** and numerical display **4630** in the example illustrated in FIG. **39** show the performance error in terms of the lead time of the ex-ante forecast. For illustration, let $\hat{e}_{t-b+l|t-b}^{(m)} = y_{t-b+l} - \hat{y}_{t-b+l|t-b}^{(m)}$ represent the multi-step-ahead prediction error for the time $(t-b+l)$ th period and for the m th time series model for $b=1, \dots, B$ and $l=1, \dots, b$, where b is the holdback index and l is the lead index. The holdback index, b , rolls the forecast origin forward in time. For each holdback index, $b=1, \dots, B$, multi-step-ahead prediction errors are computed for each lead index, $l=1, \dots, b$. By gathering the prediction errors by the lead index, l , for all holdback indices, $b=1, \dots, B$, forecast performance can be evaluated by lead time. In other words, a model's l -step-ahead forecast performance can be evaluated across various forecast origins.

[0105] With reference again to the simulation tab of the interface **4500**, FIG. **40** illustrates the further ability of the rolling simulation engine to generate ex-post (i.e., forward-looking) forecasts for the combined model. The example interface **4500** includes a user input region **4700** to select the desired number of periods to forecast. When a simulation is executed with both the number of out-of-sample observations **4510** and a number of forecast periods **4700** selected, the rolling simulation engine generates both ex-ante and ex-post forecasts for each of the out-of-sample observations. The forecasts for each out-of-sample observation are plotted along with the actual out-of-sample data in the graphical display region **4530** of the interface **4500**. A vertical line **4730** separates the ex-ante forecasts **4720** (to the left of the line) from the ex-post **4710** forecasts (to the right of the line.) Of course, actual out-of-sample data values are displayed only for past periods in which data has been recorded, which is why there are no actual out-of-sample values to the right of the vertical line **4730** in the illustrated example.

[0106] In addition, for further comparison, the example interface **4500** also includes numerical values for both the ex-ante and ex-post forecasts in the numerical display region **4540**, along with the actual out-of-sample data values. The ex-ante and ex-post forecasted values in the illustrated example are set forth in the numerical display region **4540** in bold text.

[0107] FIG. **41** illustrates an example of how the graphical display on the simulation interface **4500** may be modified to compare specific out-of-sample observations. In the illustrated example, the interface includes user input fields **4800** for selecting either all of the simulations or select simulations for display. If the "select simulations" field is selected, then the interface **4500** only displays the forecasts for out-of-sample observations that are selected in the numerical display region **4540**. In the illustrated example, the Back:2, Back:4 and Back:6 observations have been selected, and therefore only these three forecasts are displayed in the graphical display region **4530** along with the actual out-of-sample values.

[0108] The methods and systems described herein may be implemented on many different types of processing devices by program code comprising program instructions that are executable by the device processing subsystem. The software program instructions may include source code, object code, machine code, or any other stored data that is operable to cause a processing system to perform the methods and operations described herein. Other implementations may also be used, however, such as firmware or even appropriately designed hardware configured to carry out the methods and systems described herein.

[0109] The systems' and methods' data (e.g., associations, mappings, data input, data output, intermediate data results, final data results, etc.) may be stored and implemented in one or more different types of computer-implemented data stores, such as different types of storage devices and programming constructs (e.g., RAM, ROM, Flash memory, flat files, databases, programming data structures, programming variables, IF-THEN (or similar type) statement constructs, etc.). It is noted that data structures describe formats for use in organizing and storing data in databases, programs, memory, or other computer-readable media for use by a computer program.

[0110] The computer components, software modules, functions, data stores and data structures described herein may be connected directly or indirectly to each other in order to allow the flow of data needed for their operations. It is also noted that a module or processor includes but is not limited to a unit of code that performs a software operation, and can be implemented for example as a subroutine unit of code, or as a software function unit of code, or as an object (as in an object-oriented paradigm), or as an applet, or in a computer script language, or as another type of computer code. The software components and/or functionality may be located on a single computer or distributed across multiple computers depending upon the situation at hand.

[0111] It should be understood that as used in the description herein and throughout the claims that follow, the meaning of "a," "an," and "the" includes plural reference unless the context clearly dictates otherwise. Also, as used in the description herein and throughout the claims that follow, the meaning of "in" includes "in" and "on" unless the context clearly dictates otherwise. Further, as used in the description herein and throughout the claims that follow, the meaning of "each" does not require "each and every" unless the context clearly dictates otherwise. Finally, as used in the description herein and throughout the claims that follow, the meanings of "and" and "or" include both the conjunctive and disjunctive and may be used interchangeably unless the context expressly dictates otherwise; the phrase "exclusive or" may be used to indicate situation where only the disjunctive meaning may apply.

It is claimed:

1. A computer-implemented method for evaluating performance of forecasting models, comprising:

- generating a plurality of forecasting models using a set of in-sample data;
- receiving a selection of two or more forecasting models from the plurality of forecasting models for use in generating a combined forecast;
- receiving a set of actual out-of-sample data; and
- generating an ex-ante combined forecast for an out-of-sample period using the selected two or more forecasting models;

- providing the ex-ante combined forecast and the set of actual out-of-sample data for use in evaluating performance of the combined forecast;
- wherein the steps of the method are performed by one or more programmed processors.
- 2. The computer-implemented method of claim 1, further comprising:
 - generating additional ex-ante combined forecasts for one or more additional out-of-sample periods; and
 - providing the additional ex-ante combined forecasts, along with the ex-ante combined forecast and the set of actual out-of-sample data, for use in evaluating performance of the combined forecast.
- 3. The computer-implemented method of claim 2, wherein one or more of the ex-ante combined forecasts are generated using a different combination of the selected two or more forecasting models than other ones of the ex-ante combined forecasts.
- 4. The computer-implemented method of claim 3, wherein the different combination of the selected two or more forecasting models is dynamically adjusted based on one or more characteristics of out-of-sample input data.
- 5. The computer-implemented method of claim 4, wherein the different combination of the selected two or more forecasting models is generated by omitting one or more forecasting model and adjusting a model weighting.
- 6. The computer-implemented method of claim 1, wherein the ex-ante combined forecast and the set of actual out-of-sample data are displayed on a graphical user interface.
- 7. The computer-implemented method of claim 1, further comprising:
 - determining a performance statistic that identifies an error between the ex-ante combined forecast and the set of actual out-of-sample data;
 - providing the performance statistic for use in evaluating performance of the combined forecast.
- 8. The computer-implemented method of claim 7, wherein the performance statistic is displayed on a graphical user interface.
- 9. A system for evaluating performance of forecasting models, comprising:
 - one or more processors;
 - one or more memory devices;

- a rolling simulation engine stored on the one or more memory devices and executable by the one or more processors, when executed the rolling simulation engine being configured to:
 - receive a selection of two or more forecasting models for use in generating a combined forecast;
 - receive a set of actual out-of-sample data; and
 - generate an ex-ante combined forecast for the out-of-sample period using the selected two or more forecasting models;
 - provide the ex-ante combined forecast and the set of actual out-of-sample data for use in evaluating performance of the combined forecast.
- 10. The system of claim 9, wherein the rolling simulation engine is further configured to:
 - generate additional ex-ante combined forecasts for one or more additional out-of-sample periods; and
 - provide the additional ex-ante combined forecasts, along with the ex-ante combined forecast and the set of actual out-of-sample data, for use in evaluating performance of the combined forecast.
- 11. The system of claim 10, wherein one or more of the ex-ante combined forecasts are generated using a different combination of the selected two or more forecasting models than other ones of the ex-ante combined forecasts.
- 12. The system of claim 11, wherein the different combination of the selected two or more forecasting models is dynamically adjusted based on one or more characteristics of out-of-sample input data.
- 13. The system of claim 12, wherein the different combination of the selected two or more forecasting models is generated by omitting one or more forecasting model and adjusting a model weighting.
- 14. The system of claim 9, wherein the rolling simulation engine is further configured to cause the ex-ante combined forecast and the set of actual out-of-sample data to be displayed on a graphical user interface.
- 15. The system of claim 9, wherein the rolling simulation engine is further configured to:
 - determine a performance statistic that identifies an error between the ex-ante combined forecast and the set of actual out-of-sample data; and
 - provide the performance statistic for use in evaluating performance of the combined forecast.
- 16. The system of claim 15, wherein the rolling simulation engine is further configured to cause the performance statistic to be displayed on a graphical user interface.

* * * * *