Claims

1. A computer-implemented method comprising:

receiving, from a first data source, a first data set that includes a first resource request for each of a first set of entities, each first resource request including a first timestamp and a first resource manifest that indicates an amount of resources to be allocated;

receiving, from one or more second data sources, a second data set that indicates, for each entity of the first set of entities:

communications data that characterizes one or more communications transmitted to the entity, each communication of the one or more communications being associated with an allocation of resources corresponding to a resource request, wherein the communications data includes, for each communication of the one or more communications:

a communication type;

a communication protocol used to transmit the communication;

a timestamp; and

content of the communication; and

allocation data that indicates whether the resource request was satisfied by the entity and an amount of resources allocated if the resource request was satisfied;

defining, for each of one or more training instances, one or more target characteristics from a plurality of available target characteristics, wherein the one or more target characteristics for a particular training instance are different than target characteristics of other training instances;

training, within a particular training instance of the one or more training instances, a machine-learning model using the first data set and the second data set;

accessing a third data set that includes, for each entity of a second set of entities, a second resource request, each second resource request including a second timestamp and a second resource manifest;

generating a user session specific to a particular user;

configuring the user session using particular target characteristics of the one or more target characteristics, the particular target characteristics being associated with the user session;

executing, within the user session and using the particular target characteristics, the machine-learning model on the third data set to generate a resource-allocation specification that includes, for each of at least one of the second set of entities, a modified second resource request that includes a communication schedule and a third resource manifest, the third resource manifest being different from the second resource manifest; and

outputting one or more communications compliant with the communication schedule of the modified second resource request.

2. The method of claim 1, wherein the third data set is stored at a first lower level of a database, and wherein the machine-learning model is executed at a native location at a second higher layer of the database.

3. The method of claim 1, wherein at least one communication of the one or more communications includes the modified second resource request or a follow-up communication to a previous communication, wherein the follow-up communication is generated in a manner that is not conditioned upon and is not triggered by receipt of user input.

4. The method of claim 1, further comprising:

receiving, from end user device, an identification of the one or more target characteristics for the machine-learning model.

5. The method of claim 1, further comprising:

receiving, from an end user device, an indication as to one or more allocation communication actions that are to be automatically performed; and

determining at least one allocation communication action of the one or more allocation communication actions that corresponds to the resource-allocation specification; and

automatically executing the allocation communication action in response to determining at least one allocation communication action of the one or more allocation communication actions that corresponds to the resource-allocation specification.

6. The method of claim 1, wherein the third data set associates each of the second set of entities with a cluster of a set of clusters, and wherein each cluster of the set of clusters is associated with one or more cluster-corresponding resource-allocation specifications.

7. The method of claim 1, further comprising, prior to outputting the one or more communications:

outputting one or more first outputs to facilitate presenting an interface at an end-user device, wherein the interface includes representation of at least part of the third data set and includes at least one input component that, when selected, triggers automatic generation of one or more communications based on at least part of the third data set; and

detecting selection of the at least one input component, wherein the one or more communications are output in response to the detection.

8. A system comprising:

one or more processors;

a non-transitory computer-readable medium including instructions that when executed by the one or more processors, causes the one or more processors to perform operations including:

receiving, from a first data source, a first data set that includes a first resource request for each of a first set of entities, each first resource request including a first timestamp and a first resource manifest that indicates an amount of resources to be allocated;

receiving, from one or more second data sources, a second data set that indicates, for each entity of the first set of entities:

communications data that characterizes one or more communications transmitted to the entity, each communication of the one or more communications being associated with an allocation of resources corresponding to a resource request, wherein the communications data includes, for each communication of the one or more communications:

a communication type;

a communication protocol used to transmit the communication;

a timestamp; and

content of the communication; and

allocation data that indicates whether the resource request was satisfied by the entity and an amount of resources allocated if the resource request was satisfied;

defining, for each of one or more training instances, one or more target characteristics from a plurality of available target characteristics, wherein the one or more target characteristics for a particular training instance are different than target characteristics of other training instances;

training, within a particular training instance of the one or more training instances, a machine-learning model using the first data set and the second data set, wherein a set of parameters are learned as a result of the training;

accessing a third data set that includes, for each entity of a second set of entities, a second resource request, each second resource request including a second timestamp and a second resource manifest;

generating a user session specific to a particular user;

configuring the user session using particular target characteristics of the one or more target characteristics;

executing, within the user session and using the particular target characteristics, the machine-learning model on the third data set to generate a resource-allocation specification that includes, for each of at least one of the second set of entities, a modified second resource request that includes a communication schedule and a third resource manifest, the third resource manifest being different from the second resource manifest; and

outputting one or more communications compliant with the communication schedule of the modified second resource request.

9. The system of claim 8, wherein the third data set is stored at a first lower level of a database, and wherein the machine-learning model is executed at a native location at a second higher layer of the database.

10. The system of claim 8, wherein at least one communication of the one or more communications includes the modified second resource request or a follow-up communication to a previous communication, wherein the follow-up communication is generated in a manner that is not conditioned upon and is not triggered by receipt of user input.

11. The system of claim 8, wherein the operations further include: receiving, from end user device, an identification of the one or more target characteristics for the machine-learning model.

12. The system of claim 8, wherein the operations further include:

receiving, from an end user device, an indication as to one or more allocation communication actions that are to be automatically performed; and

determining at least one allocation communication action of the one or more allocation communication actions that corresponds to the resource-allocation specification; and

automatically executing the allocation communication action in response to determining at least one allocation communication action of the one or more allocation communication actions that corresponds to the resource-allocation specification.

13. The system of claim 8, wherein the third data set associates each of the second set of entities with a cluster of a set of clusters, and wherein each cluster of the set of clusters is associated with one or more cluster-corresponding resource-allocation specifications.

14. The system of claim 8, wherein the operations further include:

prior to outputting the one or more communications:

outputting one or more first outputs to facilitate presenting an interface at an end-user device, wherein the interface includes representation of at least part of the third data set and

includes at least one input component that, when selected, triggers automatic generation of one or more communications based on at least part of the third data set; and

detecting selection of the at least one input component, wherein the one or more communications are output in response to the detection.

15. A non-transitory computer-readable medium including instructions that when executed by one or more processors, causes the one or more processors to perform operations including:

receiving, from a first data source, a first data set that includes a first resource request for each of a first set of entities, each first resource request including a first timestamp and a first resource manifest that indicates an amount of resources to be allocated;

receiving, from one or more second data sources, a second data set that indicates, for each entity of the first set of entities:

communications data that characterizes one or more communications transmitted to the entity, each communication of the one or more communications being associated with an allocation of resources corresponding to a resource request, wherein the communications data includes, for each communication of the one or more communications:

a communication type;

a communication protocol used to transmit the communication;

a timestamp; and

content of the communication; and

allocation data that indicates whether the resource request was satisfied by the entity and an amount of resources allocated if the resource request was satisfied;

defining, for each of one or more training instances, one or more target characteristics from a plurality of available target characteristics, wherein the one or more target characteristics for a particular training instance are different than target characteristics of other training instances;

training, within a particular training instance of the one or more training instances, a machine-learning model using the first data set and the second data set, wherein a set of parameters are learned as a result of the training;

accessing a third data set that includes, for each entity of a second set of entities, a second resource request, each second resource request including a second timestamp and a second resource manifest;

generating a user session specific to a particular user;

configuring the user session using particular target characteristics of the one or more target characteristics;

executing, within the user session and using the particular target characteristics, the machine-learning model on the third data set to generate a resource-allocation specification that includes, for each of at least one of the second set of entities, a modified second resource request that includes a communication schedule and a third resource manifest, the third resource manifest being different from the second resource manifest; and

outputting one or more communications compliant with the communication schedule of the modified second resource request.

16. The non-transitory computer-readable medium of claim 15, wherein the third data set is stored at a first lower level of a database, and wherein the machine-learning model is executed at a native location at a second higher layer of the database.

17. The non-transitory computer-readable medium of claim 15, wherein at least one communication of the one or more communications includes the modified second resource request or a follow-up communication to a previous communication, wherein the follow-up communication is generated in a manner that is not conditioned upon and is not triggered by receipt of user input.

18. The non-transitory computer-readable medium of claim 15 wherein the operations further include:

receiving, from end user device, an identification of the one or more target characteristics for the machine-learning model.

19. The non-transitory computer-readable medium of claim 15, wherein the operations further include:

receiving, from an end user device, an indication as to one or more allocation communication actions that are to be automatically performed; and

determining at least one allocation communication action of the one or more allocation communication actions that corresponds to the resource-allocation specification; and

automatically executing the allocation communication action in response to determining at least one allocation communication action of the one or more allocation communication actions that corresponds to the resource-allocation specification.

20. The non-transitory computer-readable medium of claim 15, wherein the third data set associates each of the second set of entities with a cluster of a set of clusters, and wherein each cluster of the set of clusters is associated with one or more cluster-corresponding resource-allocation specifications.