CLAIMS

We claim:

1. A computer-implemented method for training a generative adversarial network (GAN), comprising:

a. generating, by a GAN generator, increasingly accurate positive simulated data until a GAN discriminator classifies the positive simulated data as positive;

b. presenting the positive simulated data, positive real data, and negative real data to a convolutional neural network (CNN), until the CNN classifies each type of data as positive or negative;

c. presenting the positive real data and the negative real data to the CNN to generate prediction scores; and

d. determining, based on the prediction scores, whether the GAN is trained or not trained, and when the GAN is not trained, repeating steps a-c until a determination is made, based on the prediction scores, that the GAN is trained.

2. The computer-implemented method of claim 1, wherein the positive simulated data, the positive real data, and the negative real data comprise biological data.

3. The computer-implemented method of claim 1, wherein the positive simulated data comprises positive simulated polypeptide-major histocompatibility complex class I (MHC-I) interaction data, the positive real data comprises positive real polypeptide-MHC-I interaction data, and the negative real data comprises negative real polypeptide-MHC-I interaction data.

4. The computer-implemented method of claim 3, wherein generating the increasingly accurate positive simulated polypeptide-MHC-I interaction data until the GAN discriminator classifies the positive simulated polypeptide-MHC-I interaction data as real comprises:

e. generating, by the GAN generator according to a set of GAN parameters, a first simulated dataset comprising simulated positive polypeptide-MHC-I interactions for a MHC allele;

f. combining the first simulated dataset with the positive real polypeptide-MHC-I interactions for the MHC allele, and the negative real polypeptide-MHC-I interactions for the MHC allele to create a GAN training dataset;

g. determining, by a discriminator according to a decision boundary, whether a respective polypeptide-MHC-I interaction for the MHC allele in the GAN training dataset is simulated positive, real positive, or real negative;

h. adjusting, based on accuracy of the determination by the discriminator, one or more of the set of GAN parameters or the decision boundary; and

i. repeating steps e-h until a first stop criterion is satisfied.

5. The computer-implemented method of claim 4, wherein presenting the positive simulated polypeptide-MHC-I interaction data, the positive real polypeptide-MHC-I interaction data, and the negative real polypeptide-MHC-I interaction data to the convolutional neural network (CNN), until the CNN classifies respective polypeptide-MHC-I interaction data as positive or negative comprises:

j. generating, by the GAN generator according to the set of GAN parameters, a second simulated dataset comprising simulated positive polypeptide-MHC-I interactions for the MHC allele;

k. combining the second simulated dataset, the positive real polypeptide-MHC-I interactions for the MHC allele, and the negative real polypeptide-MHC-I interactions for the MHC allele to create a CNN training dataset;

l. presenting the CNN training dataset to the convolutional neural network (CNN);

m. classifying, by the CNN according to a set of CNN parameters, a respective polypeptide-MHC-I interaction for the MHC allele in the CNN training dataset as positive or negative;

n. adjusting, based on accuracy of the classification by the CNN, one or more of the set of CNN parameters; and

o. repeating steps l-n until a second stop criterion is satisfied.

6. The method of claim 5, wherein presenting the positive real polypeptide-MHC-I interaction data and the negative real polypeptide-MHC-I interaction data to the CNN to generate prediction scores comprises:

classifying, by the CNN according to the set of CNN parameters, a respective polypeptide-MHC-I interaction for the MHC allele as positive or negative.

7. The computer-implemented method of claim 6, wherein determining, based on the prediction scores, whether the GAN is trained comprises determining accuracy of the classification by the CNN, wherein when the accuracy of the classification satisfies a third stop criterion, outputting the GAN and the CNN.

8. The computer-implemented method of claim 6, wherein determining, based on the prediction scores, whether the GAN is trained comprises determining accuracy of the classification by the CNN, wherein when the accuracy of the classification does not satisfy a third stop criterion, returning to step a.

9. The computer-implemented method of claim 4, wherein the GAN parameters comprise one or more of allele type, allele length, generating category, model complexity, learning rate, or batch size.

10. The computer-implemented method of claim 9, wherein the allele type comprises one or more of HLA-A, HLA-B, HLA-C, or a subtype thereof.

11. The computer-implemented method of claim 3, further comprising:

presenting a dataset to the CNN, wherein the dataset comprises a plurality of candidate polypeptide-MHC-I interactions;

classifying, by the CNN, each of the plurality of candidate polypeptide-MHC-I interactions as a positive or a negative polypeptide-MHC-I interaction; and

synthesizing the polypeptide from the candidate polypeptide-MHC-I interaction classified as a positive polypeptide-MHC-I interaction.

12. A polypeptide produced by the method of claim 11.

13. The computer-implemented method of claim 11, wherein the polypeptide is a tumor specific antigen.

14. The computer-implemented method of claim 11, wherein the polypeptide comprises an amino acid sequence that specifically binds to an MHC-I protein encoded by a selected MHC allele.

15. The computer-implemented method of claim 3, wherein the positive simulated polypeptide-MHC-I interaction data, the positive real polypeptide-MHC-I interaction data, and the negative real polypeptide-MHC-I interaction data are associated with a selected allele.

16. The computer-implemented method of claim 3, wherein generating the increasingly accurate positive simulated polypeptide-MHC-I interaction data until the GAN discriminator classifies the positive simulated polypeptide-MHC-I interaction data as positive comprises:

iteratively executing the GAN discriminator in order to increase a likelihood of giving a high probability to positive real polypeptide-MHC-I interaction data, a low probability to the positive simulated polypeptide-MHC-I interaction data, and a low probability to the negative real polypeptide-MHC-I interaction data; and

iteratively executing the GAN generator in order to increase a probability of the positive simulated polypeptide-MHC-I interaction data being rated highly.

17. The computer-implemented method of claim 3, wherein the prediction score is a probability of the positive real polypeptide-MHC-I interaction data being classified as positive polypeptide-MHC-I interaction data.

18. The computer-implemented method of claim 1, wherein determining, based on the prediction scores, whether the GAN is trained comprises comparing one or more of the prediction scores to a threshold.

19. (New) An apparatus configured for performing the method of any one of claims 1-18.

20. (New) A computer readable medium (CRM) configured for performing the method of any one of claims 1-18.