Immunology

https://doi.org/10.1038/s42256-023-00641-5

Meta-learning for T cell receptor binding specificity and beyond

Duolin Wang, Fei He, Yang Yu & Dong Xu

Predicting whether T cell receptors bind to specific peptides is a challenging problem because most binding examples in the training data involve only a few peptides. A new approach uses meta-learning to improve predictions for binding to peptides for which no or little binding data exists.

The T cell receptor (TCR), as a protein complex expressed on T cells, has a critical role in the adaptive immune system by recognizing and binding to specific antigen peptides[1]. TCRs are highly diverse, allowing T cells to attack a wide range of antigens from pathogen-infected cells and cancer cells. A given antigen peptide can trigger a specific set of TCRs, resulting in a targeted immune response. There are many different peptide motifs and significant differences in the numbers of corresponding TCRs. Predicting which TCRs can bind to a specific antigen peptide has broad clinical applications, such as large-scale screening of potential TCR targets for cancer neoantigen therapy[2]. In this issue of *Nature Machine Intelligence*, Gao et al. propose PanPep[3], a meta-learning-based framework to address the TCR–antigen binding recognition problem for any type of antigen peptide. In particular, the method can predict binding to antigens that have never been seen in the immune system. The study could also motivate developments with meta-learning in other small-data bioinformatics problems.

The prediction of peptide-specific TCR binding, like many other bioinformatics problems, is challenged by the long-tail distribution of TCR binding (Fig. 1a), whereby a small number of peptides have many known binding TCRs, while many peptides have small numbers of known binding TCRs. Predictions may therefore be highly biased toward a few peptides represented in a majority of training data, with low utility for the majority of peptides for which insufficient or no training data are available. As a result, existing tools perform poorly on predictions for peptides located in the long-tail region — antigens with a few known TCRs and previously unseen antigens — in what are known as the few-shot and zero-shot learning problems. In recent years, various machine learning approaches have been developed to address long-tail problems, such as transfer learning[4], domain adaptation methods[5] and, most notably, meta-learning.

The concept of meta-learning has evolved over time, and this continues to be an active area of machine learning research. Initially, the meta-learning approach aimed to improve algorithm performance by sharing information across tasks and learning how to best apply existing learning algorithms to new ones. Over time, the focus has shifted toward developing models that can quickly adapt to new tasks with limited data, as used in PanPep. Training a meta-learner typically requires two learning stages: meta-training and meta-testing. During meta-training, the model is exposed to a variety of different tasks to learn a general problem-solving strategy that can be applied to new tasks. During meta-testing, the model is presented with a new task and uses the knowledge gained during meta-training to quickly adapt to the new task and solve it.

To tackle the long-tail distribution problem, PanPep employs meta-learning in three settings: (i) the majority setting, for peptides with a large number of known binding TCRs, (ii) few-shot learning, for peptides with a small number (<10) of known binding TCRs and (ii) zero-shot learning, for peptides not present in the training data. PanPep applies a widely used optimization-based adaptation method, model-agnostic meta-learning (MAML[6]), to target the majority and few-shot settings. Specifically, in the meta-training stage, the model is trained on a set of peptide-specific TCR binding tasks to obtain a series of peptide-specific learners and optimize the meta-learner. Then, in the meta-testing procedure, the meta-learner is fine-tuned on a new peptide-specific binding recognition task. PanPep proposes a disentanglement distillation module to handle the zero-shot setting. A mapping between peptide embedding and the peptide-specific learners is constructed based on a neural Turing machine (NTM)[7] (Fig. 1b). The read head of the NTM is used to map a peptide embedding to a new embedding space called peptide-specific learner generation space (PLGS), and a write head of the NTM is used to extract the peptide-specific learners. The NTM's memory stores the mapping between peptide embedding and the extracted peptide-specific learners. This NTM-based module is trained based on all the peptide-specific learners through knowledge distillation. In this way, PanPep can extend the few-shot settings to zero-shot settings. Once an unseen peptide arises, the trained read head will map it into the PLGS, and then the PLGS will be used to retrieve the memory to generate a new peptide-specific learner for the inference of the unseen peptide. This innovative design makes PanPep a powerful tool for predicting TCR binding specificity of TCRs with few or unseen antigens.

In the evaluations using the curated independent data, PanPep achieved excellent performance on various peptide-specific tasks, especially for unseen peptides. Furthermore, Gao et al. demonstrate PanPep's utility in several clinical applications. The output scores of PanPep indicated a relatively high correlation with clonal T cell expansion ratios, suggesting its potential to provide accurate binding identification for clonal T cells. In neoantigen therapy, PanPep effectively identified immune-responsive T cells and detected neoantigen-reactive T cell signatures, which may help improve adoptive cell transfer (ACT)-based tumour immunotherapy. In a COVID-19 study, PanPep demonstrated substantial improvement in recognizing peptide-specific TCRs over three other tools. Moreover, it provided interpretability by unveiling the nature of peptide and TCR interactions through protein structure modelling. Finally, PanPep displayed high computational efficiency.

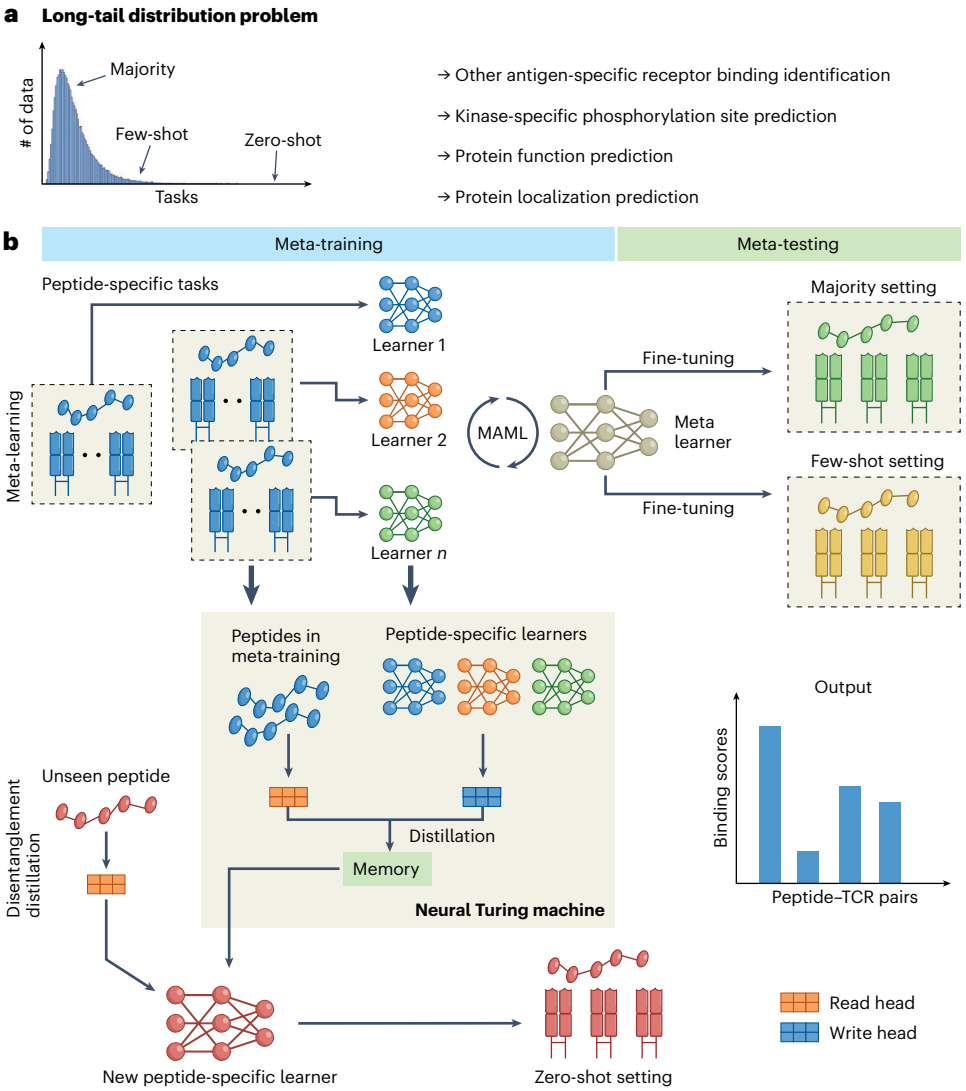

Fig. 1 | The PanPep workflow of meta-learning augmented with a neural Turing machine. a, Definition of the long-tail distribution problem and potential bioinformatics applications using meta-learning. **b**, Model workflow. For peptide-specific TCR binding prediction, the input comprises a peptide sequence and the CDR3 region of a TCR protein sequence that binds to antigens. The output is information on whether this peptide–TCR pair represents a bona fide biological interaction. In the meta-training stage, a model is trained on a series of peptide-specific TCR binding tasks to obtain peptide-specific learners and optimize a meta-learner (with model-agnostic machine learning, MAML). In the meta-testing stage, the meta-learner is fine-tuned on new binding recognition tasks for peptides with a large number of supporting TCRs (majority setting) or with a small number of supporting TCRs (few-shot setting). The neural Turing machine (NTM) maps a peptide embedding to an embedding space using the read head and extracts the peptide-specific learners using the write head. The NTM memory stores the mapping between the peptide embedding and disentangled learners, capable of generating new peptide-specific learners through memory retrieval for unseen peptides in the zero-shot setting.

It is often perceived that deep learning requires massive datasets of labelled training samples to be effective. However, labelled data may be sparse in many real-world applications, especially in biological and medical areas. One class of 'small data' cases is the long-tail problem. The work by Gao et al. represents a promising application of meta-learning in addressing long-tail distribution problems in bioinformatics. In particular, PanPep fills the gap in handling the zero-shot setting for TCR binding specificity prediction by integrating the meta-learning modules with disentanglement distillation. Although one can build a task-blind model using the peptide and TCR–CDR3 sequences as the input for all tasks, such an approach assumes that the data in all tasks follow the identically independent distribution (i.i.d.), which is not the case for peptide-specific tasks. The meta-learning and the NTM-based disentanglement distillation proposed in PanPep can take advantage of the peptide-specific data distribution to adapt to new tasks well. These methods may be further improved using newer machine learning methods, such as the NTM's successor, the differentiable neural computer[8]. It may also be beneficial to use graph neural networks to represent all the peptides and TCRs by leveraging the global relationships among peptides and TCRs.

PanPep can potentially serve as a general framework for many new bioinformatics applications. It may be extended to tackle other peptide binding prediction tasks that are subject to long-tail distribution problems, such as peptide–HLA binding prediction and kinase-specific phosphorylation-site prediction[9]. The few-shot meta-learning methods may be applicable to protein function predictions, such as protein localization prediction[10] and Enzyme Commission Number prediction. PanPep also has some limitations and new challenges. One limitation is that the proposed method did not provide superior performance, compared with existing methods, for the majority setting, involving predictions where ample training data is available. This may be because the meta-learner needs to balance between all the tasks to ensure that the model can generalize well to new tasks. Therefore, further regularization techniques or hyperparameter selection techniques may need to be implemented to ensure optimal training results in the majority setting. It is also noteworthy that even though the peptides are 'unseen' in the training procedure, the TCR may be 'seen' by other peptide–TCR binding pairs in the training procedure. It would be interesting to assess the performance difference between the unknown (peptides)–known (TCR) scenario and the unknown–unknown scenario. In summary, PanPep delivered great promise of using meta-learning to address bioinformatics' long-tail distribution problems. We anticipate that many new meta-learning methods will be developed for a wide range of bioinformatics applications.

Duolin Wang, Fei He, Yang Yu & Dong Xu
Department of Electrical Engineering and Computer Science, Bond Life Sciences Center, University of Missouri, Columbia, MO, USA.
e-mail: xudong@missouri.edu

Published online: 31 March 2023

References

1. Waldman, A. D., Fritz, J. M. & Lenardo, M. J. *Nat. Rev. Immunol.* **11**, 651–668 (2020).
2. Yamamoto, T. N., Kishton, R. J. & Restifo, N. P. *Nat. Med.* **25**, 1488–1499 (2019).
3. Gao, Y. et al. *Nat. Mach. Intell.* **5**, 236–249 (2023).
4. Wang, Y. X., Ramanan, D. & Hebert, M. *Adv. Neural Inf. Process. Syst.* **30**, 7029–7039 (2017).
5. Jamal, M. A., Brown, M., Yang, M. H., Wang, L. & Gong, B. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 7610–7619 (2020).
6. Finn, C., Abbeel, P. & Levine, S. *Proc. Mach. Learn. Res.* **70**, 1126–1135 (2017).
7. Graves, A. et al. Preprint at https://doi.org/10.48550/arXiv.1410.5401 (2014).
8. Graves, A. et al. *Nature* **538**, 471–476 (2016).
9. Wang, D. et al. *Bioinformatics* **33**, 3909–3916 (2017).
10. Jiang, Y. et al. *Comput. Struct. Biotechnol. J.* **19**, 4825–4839 (2021).

Competing interests

The authors declare no competing interests.