

HLA-Arena: A Customizable Environment for the Structural Modeling and Analysis of Peptide-HLA Complexes for Cancer Immunotherapy

Dinler A. Antunes, PhD¹; Jayvee R. Abella, PhD¹; Sarah Hall-Swan, BS¹; Didier Devaurs, PhD²; Anja Conev, BS¹; Mark Moll, PhD¹; Gregory Lizée, PhD³; and Lydia E. Kavraki, PhD¹

PURPOSE HLA protein receptors play a key role in cellular immunity. They bind intracellular peptides and display them for recognition by T-cell lymphocytes. Because T-cell activation is partially driven by structural features of these peptide-HLA complexes, their structural modeling and analysis are becoming central components of cancer immunotherapy projects. Unfortunately, this kind of analysis is limited by the small number of experimentally determined structures of peptide-HLA complexes. Overcoming this limitation requires developing novel computational methods to model and analyze peptide-HLA structures.

METHODS Here we describe a new platform for the structural modeling and analysis of peptide-HLA complexes, called HLA-Arena, which we have implemented using Jupyter Notebook and Docker. It is a customizable environment that facilitates the use of computational tools, such as APE-Gen and DINC, which we have previously applied to peptide-HLA complexes. By integrating other commonly used tools, such as MODELLER and MHCflurry, this environment includes support for diverse tasks in structural modeling, analysis, and visualization.

RESULTS To illustrate the capabilities of HLA-Arena, we describe 3 example workflows applied to peptide-HLA complexes. Leveraging the strengths of our tools, DINC and APE-Gen, the first 2 workflows show how to perform geometry prediction for peptide-HLA complexes and structure-based binding prediction, respectively. The third workflow presents an example of large-scale virtual screening of peptides for multiple HLA alleles.

CONCLUSION These workflows illustrate the potential benefits of HLA-Arena for the structural modeling and analysis of peptide-HLA complexes. Because HLA-Arena can easily be integrated within larger computational pipelines, we expect its potential impact to vastly increase. For instance, it could be used to conduct structural analyses for personalized cancer immunotherapy, neoantigen discovery, or vaccine development.

JCO Clin Cancer Inform 4:623-632. © 2020 by American Society of Clinical Oncology

Licensed under the Creative Commons Attribution 4.0 License 

INTRODUCTION

Immunotherapy treatments are now at the forefront of methods used for cancer therapy. These treatments aim at harvesting a patient's own immunologic defenses to identify and eliminate cancer cells.¹ Many of these immunotherapy treatments involve class I HLA protein receptors. HLA receptors bind peptides produced by the cleavage of intracellular proteins, which is a continuous process present in almost every cell. The resulting peptide-HLA (pHLA) complexes are then exposed at the surface of cells. Also present in cancer cells, this mechanism allows circulating T-cell lymphocytes to recognize tumor-associated peptides, thus triggering T-cell activation, tumor elimination, and immunologic memory against the tumor.^{1,2}

It has been shown that immunologic outcomes are partially driven by structural features of pHLA complexes.²⁻⁴ Therefore, the structural modeling and analysis of

these complexes are becoming essential to ensure the efficacy and safety of immunotherapy treatments.² However, pHLA structural features are affected by the genetic variability of both patients and tumors.^{2,5} First, the set of peptides available for presentation reflects the patient's genetic background and cancer-specific alterations.^{2,5} Second, each individual has up to 6 class I HLA alleles,⁶ among the nearly 19,000 alleles in the human population.⁷ Each allele encodes for a receptor with specific characteristics, which will display a different pool of peptides. Therefore, the structural modeling and analysis of pHLA complexes for cancer immunotherapy require fast and customizable methods that can handle patient-specific data.

Unfortunately, the cost and time requirements of gold-standard experimental techniques in structural biology prevent their use in personalized medicine. In addition, few structures of pHLA complexes have been

ASSOCIATED CONTENT

Appendix

Author affiliations and support information (if applicable) appear at the end of this article.

Accepted on June 3, 2020 and published at ascopubs.org/journal/cci on July 15, 2020; DOI <https://doi.org/10.1200/CCI.19.00123>

CONTEXT

Key Objective

Enabling large-scale structural modeling and analysis of peptide-HLA complexes for cancer immunotherapy applications.

Knowledge Generated

We created a customizable environment, called HLA-Arena, with user-friendly computational workflows that allow for varied structure-based analyses of peptide-HLA complexes. To illustrate this, we show how researchers can use HLA-Arena to perform geometry prediction of peptide binding modes, peptide binding energy prediction, and structure-based virtual screening of tumor-derived peptides, for any classic class I HLA of interest.

Relevance

HLA-Arena can be integrated in computational pipelines to support basic cancer research or to help inform physicians in preclinical settings. It can be used to perform structure-based selection of peptides for T cell-based immunotherapy, neoantigen discovery, and vaccine development.

determined experimentally. Therefore, researchers have turned toward computational methods for the structural modeling of pHLA complexes. However, the length and flexibility of displayed peptides represent major challenges for traditional methods.⁵ As an alternative, in previous work, we have developed several computational tools for the accurate and efficient modeling of pHLA complexes. For example, we have described a fast method, called APE-Gen, to generate ensembles of peptide conformations bound to a given HLA receptor.⁸ We have also developed a meta-docking approach, called DINC, which allows prediction of binding modes of pHLA complexes.^{9,10}

In this report, we present a higher-level platform, called HLA-Arena, that allows the carrying out of sophisticated structural modeling and analysis of pHLA complexes. Instead of having to deal with several computational tools, HLA-Arena provides researchers with a single customizable environment that fully integrates the tools we have developed, as well as other commonly used software. HLA-Arena simplifies the interactions with these tools by leveraging the capabilities of Jupyter Notebook and Docker. It allows users to perform various workflows, each involving a specific combination of tools and steps within a coherent scenario. In addition to APE-Gen and DINC, HLA-Arena currently integrates MODELLER¹¹ for homology modeling, MHCflurry¹² for binding affinity prediction, and NGL Viewer¹³ for structure visualization, among others.

Here, we present 3 example workflows illustrating the capabilities of HLA-Arena. The first relies on DINC to predict the binding modes of 2 known peptides with their corresponding HLA receptors (ie, geometry prediction). The second relies on APE-Gen to assess differences in binding between peptides restricted to a given HLA receptor, based on generated binding mode ensembles (ie, binding prediction). The third aims at performing structure-based virtual screening, which requires speed and scalability. Using real immunopeptidomic data and a fictitious diplotype (ie, 6 classic class I HLA alleles), we show how

MHCflurry and APE-Gen can complement each other to select target peptides for a hypothetical immunotherapy treatment.

METHODS

Computational Approaches for pHLA Binding Mode Prediction

Despite their huge sequence diversity, HLA receptors feature conserved secondary and tertiary structures, as illustrated by available data.¹⁴⁻¹⁶ Such conserved folding makes HLA modeling an easy task with tools leveraging homology modeling.^{8,17,18} In contrast, predicting the binding modes of peptides to HLA receptors is much harder because of the size and flexibility of these peptides. As recently reviewed, strategies used to overcome this challenge include constrained backbone prediction, constrained termini prediction, and incremental prediction.⁵

In recent years, we have implemented 2 computational approaches for pHLA binding mode prediction using these strategies. The first, called APE-Gen (anchored peptide-MHC ensemble generator), can quickly produce an ensemble of binding modes for a pHLA complex, using termini templates to position the peptide in the HLA binding cleft (Fig 1; Appendix).⁸ The second, called DINC, can incrementally dock a peptide in the binding site and does not require any template (Fig 2; Appendix).^{9,10,21} Each approach has different strengths and limitations and can therefore suit various user needs, depending on the task at hand. For instance, its speed makes APE-Gen better suited for large-scale modeling and structure-based virtual screening. In contrast, because it does not rely on templates, DINC's predictions can be more general and account for unusual binding modes, thus making it more suited for geometry prediction.^{9,22} Both APE-Gen and DINC have been validated in previous publications.^{8,10,23} In this report, we present a unified environment that facilitates the use of APE-Gen, DINC, and other tools for various research applications.

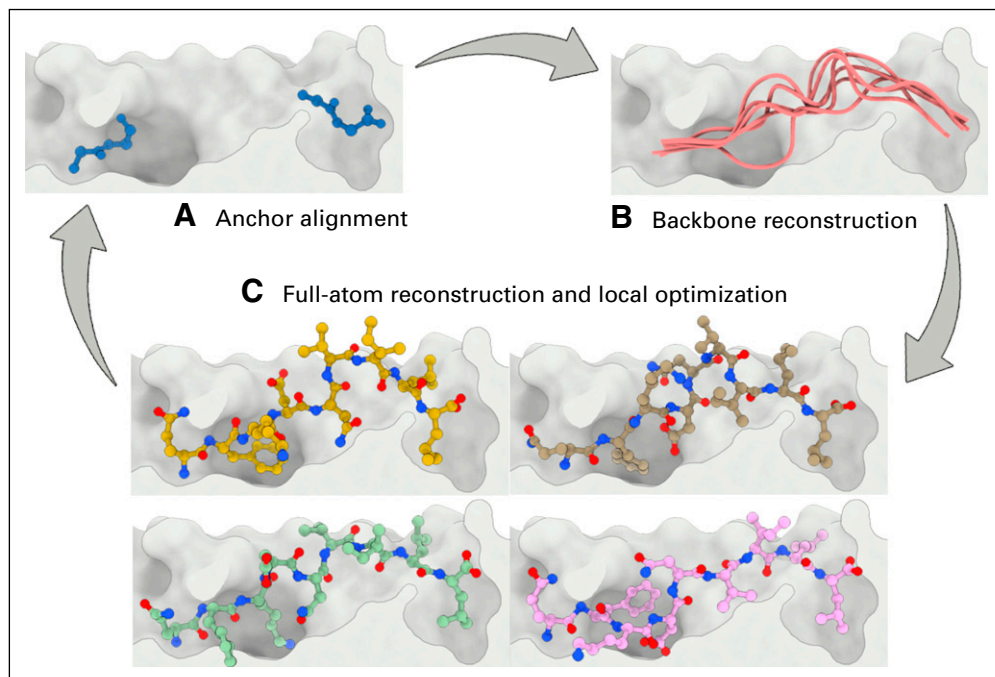


FIG 1. Generating binding mode ensembles with APE-Gen. (A) Templates of backbone termini are used to position the anchor residues of a peptide in the binding site. (B) The random coordinate descent loop closure tool¹⁹ is used to generate an ensemble of backbone conformations for this peptide. (C) Full-atom reconstruction of peptide side chains and local optimization of the resulting complex are performed for each sampled backbone. The highest-quality binding mode can be selected to be used as a template for the next round of the iterative process.

HLA-Arena: Structural Modeling and Analysis of pHLA Complexes

Using Jupyter Notebook and Docker, we have created a customizable environment, called HLA-Arena, that enables researchers to easily model any class I pHLA complex of interest and perform varied structural analyses (Fig 3).

HLA-Arena includes different workflows, defined as separate notebooks, which consist of the following main stages:

Input processing. Available structures of HLA receptors are obtained from the Protein Data Bank (PDB)²⁷ to be used as such or as templates. Unavailable HLA structures are

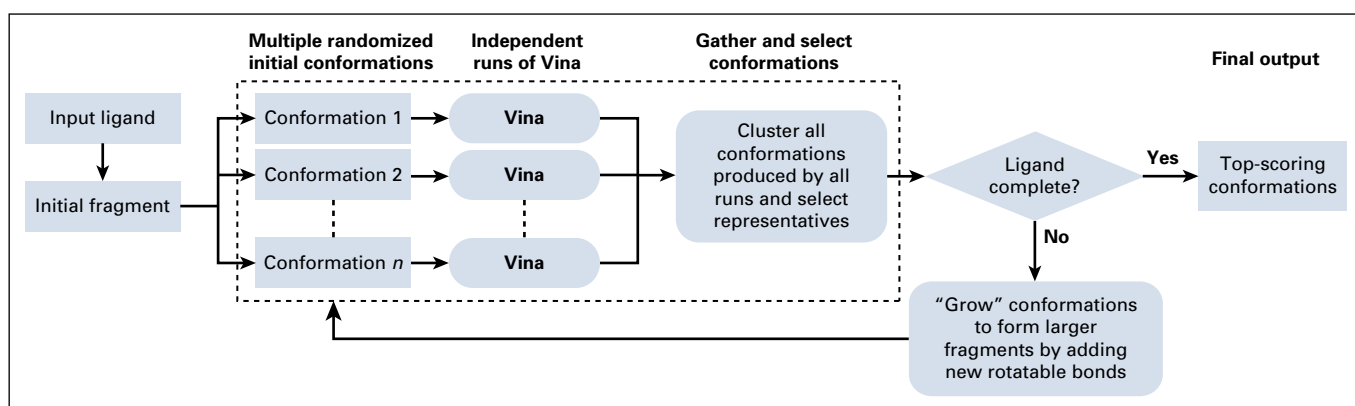


FIG 2. Workflow of DINC parallel and incremental meta-docking approach. DINC starts by selecting a small fragment of the input ligand, with only k flexible bonds. Multiple conformations are created by randomly sampling different values for the dihedral angles of this fragment. These n conformations are then used as input for multiple independent runs of a docking tool (in this example, Vina²⁰), which are executed in parallel by different threads. From all the binding modes produced by these parallel runs, the n best modes are selected for expansion; they are grown by adding several atoms and bonds from the input ligand. These larger fragments are then docked independently, in parallel, while keeping the number of flexible bonds equal to k . This process is repeated until the entire input ligand has been incrementally reconstructed and is docked in the binding site of the receptor.

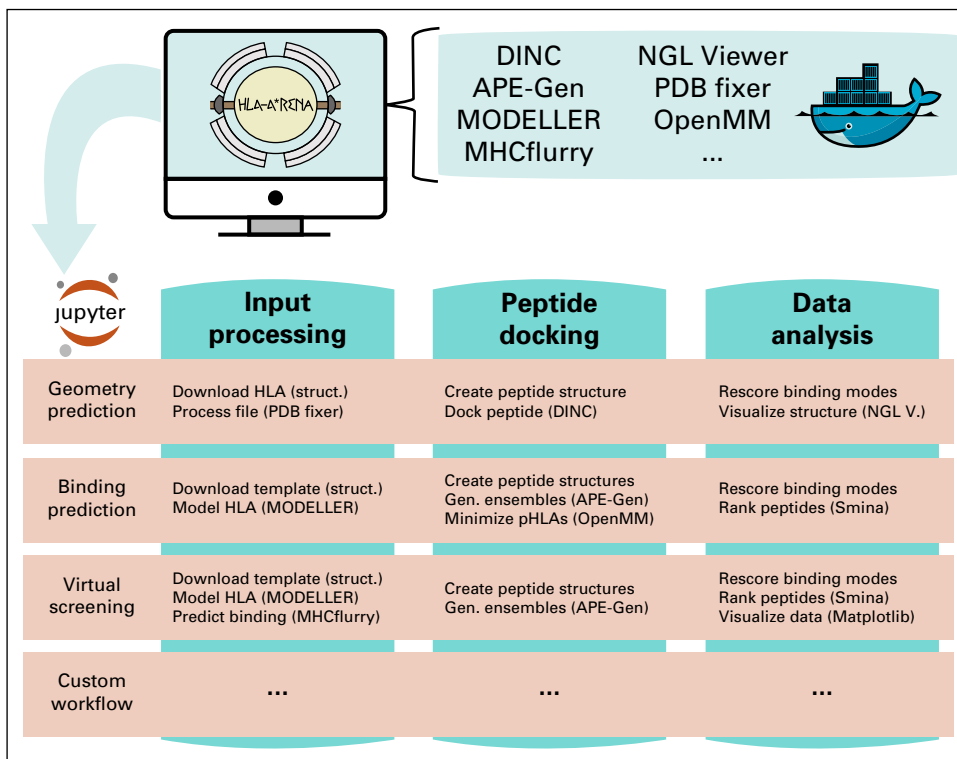


FIG 3. HLA-Arena leverages Docker and Jupyter Notebook, offering a customizable environment to build and execute various workflows for the structural modeling and analysis of peptide-HLA (pHLA) complexes. Three proposed workflows are depicted here: (1) geometry prediction of pHLA binding modes, (2) structure-based prediction of binding energy, and (3) virtual screening of tumor-derived peptides. In the geometry prediction workflow, after obtaining the structure of an HLA receptor, a peptide of interest is docked in its binding site by DINC, and all generated (Gen) binding modes are scored with several scoring functions. In the binding prediction workflow, after modeling a given HLA structure (struct.), ensembles of binding modes are generated with APE-Gen (and optionally minimized with OpenMM²⁴) for various peptides, and these binding modes are scored to rank the peptides with Smina.²⁵ In the virtual screening workflow, after filtering peptides with MHCflurry,¹² ensembles of binding modes are generated with APE-Gen for the selected peptides, and the top-scoring binding modes are used to rank these peptides with Smina, in terms of binding affinity to an HLA receptor or set of receptors. Note that these workflows can be modified, and new workflows can be created by users. In each application, different types of data analysis can be used to guide the selection of the best pHLAs before experimental validation. The screen icon was modified from Flaticon.²⁶ PDB, Protein Data Bank; V, Viewer.

modeled with MODELLER,¹¹ using an HLA sequence and the structure of a similar HLA receptor as template, if these are provided by the user. Alternatively, users can just provide an allele name (eg, HLA-A*24:02); HLA-Arena will then fetch the proper sequence from IMGT/HLA,⁷ and a reasonable template (based on the HLA supertype²⁸ classification) from the PDB. In addition, binding affinity of peptides can be estimated with MHCflurry¹² to select the most relevant ones. Minimal example: `HLA_allele = arena.model_hla('HLA-A*24:02')`

Peptide docking. Structures of pHLA complexes are modeled with APE-Gen and/or DINC, which only require the sequence of the target peptide(s) and the HLA structure(s) obtained previously. Modeled structures can also be minimized with a force field using OpenMM.²⁴ Minimal example: `structure = arena.dock('QFKDENVILL', HLA_allele)`

Data analysis. A variety of postprocessing options for data analysis can be incorporated into a workflow. These include binding mode rescoring or peptide ranking with DINC and structure visualization with NGL Viewer,¹³ among others. Minimal example: `arena.visualize(structure)`

For a smooth user experience, all computational tools involved in HLA-Arena are packaged within a Docker image (Appendix provides installation details), thus eliminating the burden of managing software dependencies. Another advantage of Docker containerization is that it makes HLA-Arena platform agnostic. As a result, it can be deployed on a desktop computer or a high-performance computing cluster, across different operating systems. Users can customize available workflows by adding modeling or analysis steps. We plan to continuously expand the capabilities of HLA-Arena by providing support for additional tools.²⁹⁻³¹

RESULTS

We now present the results we obtained when carrying out 3 different workflows that exemplify the diversity of applications offered by HLA-Arena. Each workflow leverages the functionalities of several tools in a coherent scenario.

Geometry Prediction of pHLA Binding Modes

HLA-Arena can be used to predict conformations of peptides bound to HLA receptors, even for peptides presenting unusual binding modes.¹⁰ To illustrate this, using the geometry prediction workflow based on DINC (Fig 3), we tried to reproduce the crystal structures of 2 such peptides.

First, we conducted a self-docking experiment with a crystal structure (with PDB code 1E27) involving HLA-B*51:01 and a 9-mer peptide derived from HIV-1. It has been suggested that the fifth residue acts as a secondary anchor for this peptide, leading to structural rearrangement of its central and amino-terminal residues.³² Our experiment evaluated the capability of DINC to reproduce the bound geometry of this peptide, without considering receptor flexibility. To evaluate performance and reproducibility, we carried out this experiment with either 8 or 32 threads (for the parallel process in DINC), running 5 replicates in each case. Default values were used for other DINC parameters.²³ Results (Appendix Fig A1A) show that, in every single run, HLA-Arena sampled a near-native peptide conformation (ie, a conformation with an all-heavy-atom root mean square deviation [RMSD] to the crystal structure < 2.5 Å).

Geometry prediction involves 2 issues that are especially challenging with peptides.⁵ The first relates to sampling (ie, how to explore the full flexibility of a large ligand). The second relates to scoring (ie, how to identify the best ligand conformation in a pool of diverse binding modes). HLA-Arena relies on the incremental process of DINC to overcome the sampling issue. It also includes a filtering step to remove peptide conformations with reverse orientation in the binding cleft. To address the scoring issue, HLA-Arena makes use of multiple scoring functions. For instance, in this self-docking experiment, conformations were ranked with the scoring functions of AutoDock4,³³ Vina,²⁰ and Vinardo.³⁴ All 3 scoring functions were able to identify near-native conformations. However, in the case of AutoDock4 (Fig 4A), the top-5 ranking conformations in 1 of the replicates included the overall lowest RMSD conformation (ie, the conformation with the lowest RMSD to the crystal structure among all sampled conformations; Appendix Fig A1A).

Second, we tried to reproduce a crystal structure (with PDB code 2GTW) involving HLA-A*02:01 and a 9-mer peptide derived from the MART-1/melan-A protein.³⁵ This peptide has an A27L substitution in comparison with the MART-1 peptide targeted by numerous clinical studies.^{36,37} This substitution leads to an alternative arrangement of primary anchor residues, resulting in an unusual binding mode.^{10,30,35}

Again, we ran 5 replicates of the geometry prediction workflow, using either 8 or 32 threads. For the prediction task to be closer to a real-case scenario, we performed a cross-docking experiment, accounting for receptor flexibility. It made this task much harder, from both sampling and scoring perspectives.^{38,39} Despite this, HLA-Arena sampled near-native conformations, although it performed better when using 32 threads (Appendix Figs A1B and A2). In terms of scoring, only AutoDock4 and Vinardo were able to recover near-native conformations (Fig 4B). Note that HLA-Arena also allows visualization of the 3-dimensional structure of the top-ranking binding mode (Figs 4C and 4D).

Structure-Based Prediction of Binding Energy

To demonstrate another application of HLA-Arena, we used the binding prediction workflow (Fig 3) to predict binding to HLA-A*02:01 for a small data set of selected peptides (Appendix Table A1). This data set included 5 experimentally identified nonbinders, as well as 11 binders with experimental binding affinities available in the Immune Epitope Database⁴⁰ and crystal structures in complex with HLA-A*02:01 available in the PDB. For each peptide, we generated an ensemble of bound conformations with APE-Gen. The binding energy of each peptide was then estimated as the median score within the conformation ensemble for each scoring function (ie, AutoDock4, Vina, and Vinardo). Correlations between these predicted binding energies and experimentally determined binding affinities were then determined (Fig 5).

In addition to the default local optimization performed by APE-Gen, HLA-Arena provides the option of minimizing the resulting complexes with OpenMM.²⁴ To evaluate the impact of this procedure, we recalculated binding energies and correlations after running this energy minimization for all conformations in each ensemble. Our results showed a consistent increase of the predicted binding energies for all scoring functions (Fig 5). This might reflect the differences in binding energy estimation that exist between these empirical or semi-empirical scoring functions⁴¹ and the force field used by OpenMM (ie, amber99sbildn).²⁴ Despite increasing binding energies, the OpenMM minimization had a positive impact on overall correlations.

Interestingly, the best correlation with experimental binding affinities was obtained when using Vina. This result is in agreement with previous studies evaluating the performance of Vina in virtual screening of drug-like ligands.^{41,42} Note that contrary to the geometry prediction workflow, in which a scoring function was only used to rank different conformations of a given peptide, here the scoring function also had to rank different peptides. Although the same function can be used for both purposes,⁵ it is possible that better results are obtained when using functions optimized for each task.

For the HLA-A*02:01 binders in our data set, we can compute RMSDs between their associated crystal structures and

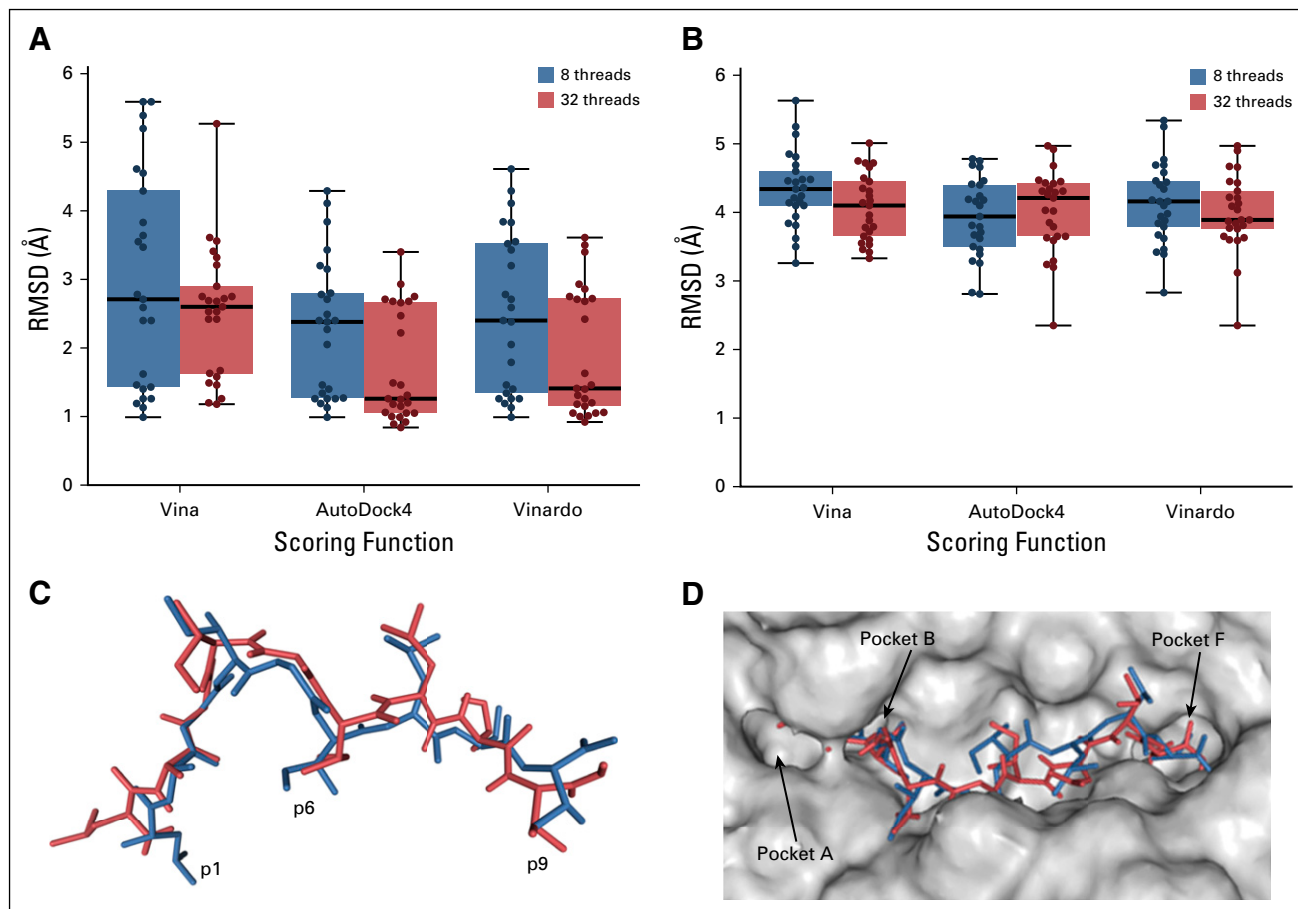


FIG 4. Geometry prediction of peptide-HLA binding modes. (A) Three scoring functions are used to select the top-5 ranking conformations produced by 5 replicates of a self-docking experiment aimed at predicting the binding mode of a 9-mer peptide (under Protein Data Bank [PDB] code 1E27) using 8 or 32 threads for DINC. Each box plot aggregates results of the 5 replicates. Each dot corresponds to a conformation plotted according to its all-heavy-atom root mean square deviation (RMSD) to the reference crystal structure. (B) Results of a cross-docking experiment aimed at predicting the binding mode of a 9-mer peptide (under PDB code 2GTW) obtained with the same methodology. (C) Side view of the best binding mode (red) identified by AutoDock4 and Vinardo and aligned with the crystal structure (blue) of this peptide (under PDB code 2GTW). Only heavy atoms are depicted, using a sticks representation. Note that this sampled conformation has an all-heavy-atom RMSD of 2.35 Å and does not perfectly reproduce the side-chain arrangement of the first residue. A better conformation, with an all-heavy-atom RMSD of 2.15 Å, was sampled by HLA-Arena (Appendix Fig A2) but was not among the top-ranking conformations. (D) Top view of the HLA binding site (depicted by a gray surface) with peptide conformations shown in panel C within it (as sticks). This peptide uses its first amino acid as primary anchor (ie, residue p1 is anchored in pocket B), which is quite unusual for HLA-A*02:01 binders. Images in panels C and D were generated with HLA-Arena using the embedded NGL Viewer.¹³ Both images were edited to add labels.

conformations generated by APE-Gen. This allows verification that APE-Gen ensembles include near-native conformations (Appendix Fig A3) and evaluation of the impact of the OpenMM minimization on these conformations. This also allows comparing the use of an ensemble of conformations to predict binding energies with the use of a single conformation from this ensemble (eg, the conformation with the lowest RMSD to the corresponding crystal structure). Our results with the Vina scoring function suggest that better correlations are obtained with ensembles of conformations (Appendix Fig A4).

Virtual Screening of Tumor-Derived Peptides

HLA-Arena allows researchers to perform, for the first time to our knowledge, a large-scale structure-based virtual

screening of HLA-binding peptides. In addition, by combining sequence- and structure-based methods, HLA-Arena represents a fresh alternative for the identification of tumor-derived peptide targets considering patient-specific HLAs. To demonstrate this application, we used the virtual screening workflow (Fig 3) to predict which peptides were the strongest binders to the class I HLA receptors of a fictitious patient with cancer.

We considered 6 alleles: HLA-A*24:02, HLA-A*26:01, HLA-B*15:01, HLA-B*35:01, HLA-C*04:01, and HLA-C*05:01. We built a peptide data set by selecting 500 known binders and 1,000 decoys for each allele, for a total of 9,000 peptides. Sequences of known binders were obtained from the Systemic MHC Atlas,⁴³ where they were derived from immunopeptidomics




FIG 5. HLA-A*02:01 binding predicted for a small set of peptides. Each plot illustrates the correlation between experimentally obtained binding affinities (extracted from the Immune Epitope Database) and structure-based binding energies, as predicted by a given scoring function: (A) Vina, (B) Vinardo, and (C) AutoDock4. Structures were generated with APE-Gen, with or without minimization with OpenMM. Correlation coefficients are also reported (as Pearson's R). Each point corresponds to a peptide in Appendix Table A1.

studies. Sequences of decoys were obtained from the training set of NetMHCpan.⁴⁴

First, the whole data set of peptides was screened for HLA binding with MHCflurry,¹² using an affinity threshold specified by the user. This allows the user to quickly select the most likely binders for each HLA receptor, before moving on to the more computationally expensive steps. In this example, a threshold of 500 nM selected 2,604 peptides. Then, we proceeded with the structural modeling of the full pHLA complex for all selected peptides. Finally, peptides were ranked based on binding energies derived from the modeled structures. The entire pipeline took 86 hours on a desktop computer or 5 hours on a high-performance cluster (Appendix).

The threshold used in MHCflurry directly affects the sensitivity/specificity of the overall prediction. Recent surveys indicate that commonly used thresholds for sequence-based HLA binding predictors (eg, 500 nM) can yield a sensitivity as low as 40%,⁴⁵ with great variation in accuracy between HLA alleles.⁴⁶ On our data set, a 500-nM

threshold produced several false-positive (Fig 6A, blue dots) and false-negative predictions (data not shown). In trying to address this issue, we observed that our structure-based analysis could usually eliminate at least half of the false-positive predictions and recover significant numbers of false-negative predictions, although results varied depending on the studied HLA allele (data not shown).

Because our workflow allowed variation of the MHCflurry threshold, we repeated the aforementioned virtual screening experiment with a 50,000-nM value. This led to all 9,000 peptides being selected for modeling and ranking. The observed enrichment of true binders among the top-ranking peptides (Fig 6B, red dots at the bottom of the distributions) further corroborates our claim that structural information is useful when screening HLA binders.

In these examples, we performed only 1 sampling round in APE-Gen for each complex, and only the top-scored conformation was used as input for ranking. Better results could be obtained by executing more sampling rounds in APE-Gen, performing the OpenMM minimization, or

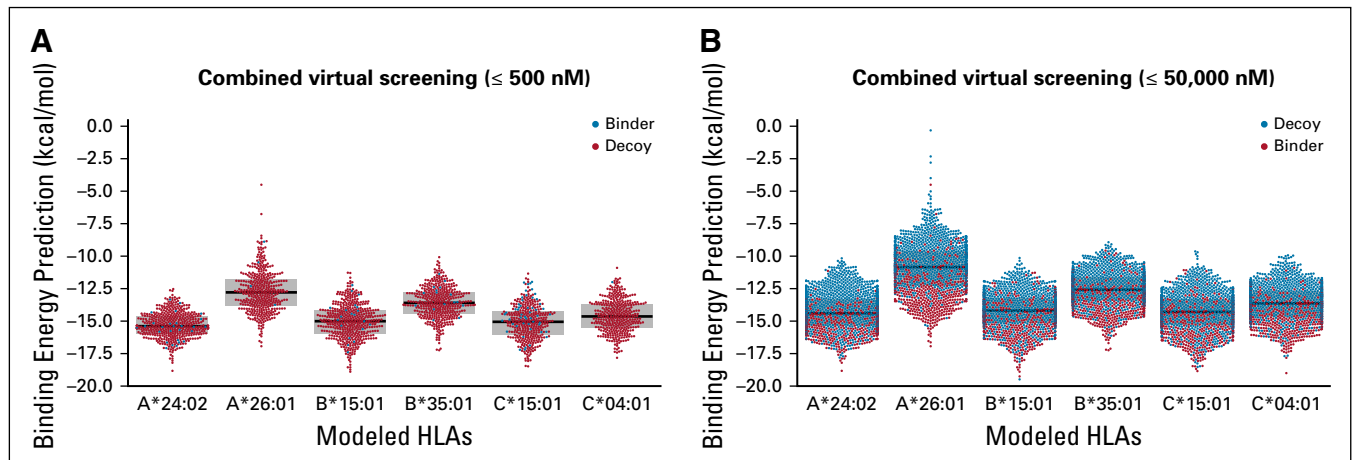


FIG 6. Structure-based virtual screening for high-affinity HLA binders. The HLA-Arena virtual screening workflow was used to predict peptide binders for 6 HLA receptors of interest. For this exercise, a data set of 9,000 peptides was created, using 500 known binders (red dots) and 1,000 decoys (blue dots) for each HLA. (A) Results of a combined virtual screening (ie, MHCflurry plus APE-Gen) with a 500-nM threshold for MHCflurry. (B) Results of the same virtual screening using a 50,000-nM threshold for MHCflurry. In both plots, each dot corresponds to the top-scoring conformation of a modeled peptide-HLA complex, selected from the ensemble of conformations produced by APE-Gen. For each HLA (on the x-axis), complexes with the lowest-binding energies (on the y-axis) would be predicted as the best candidates for further analysis or experimental validation.

using the whole APE-Gen ensemble. More importantly, accurate scoring remains an open challenge. Therefore, structure-based predictions cannot yet outperform sequence-based methods, but they can be combined to provide additional information when selecting peptides for experimental validation.

DISCUSSION

HLA-Arena provides researchers with a customizable environment to create and execute sophisticated workflows for the structural modeling and analysis of pHLA complexes. Its intuitive interface relies on Jupyter Notebook and Docker to dramatically reduce the burden of software dependencies and the need for advanced programming skills, making its resources accessible to a wide audience. Available workflows combine commonly used software for protein modeling and analysis, with tools that we developed to address challenges specific to pHLA complexes. We believe that HLA-Arena could become a stepping stone toward a broad collaborative effort to study pHLA complexes.

In this report, we present 3 workflows to showcase the capabilities of HLA-Arena. First, HLA-Arena enabled the geometry prediction of pHLA structures, even for peptides with unusual binding modes, by using template-free molecular docking. Second, HLA-Arena allowed prediction of binding energies for potential HLA binders by quickly producing ensembles of bound conformations for these peptides and rescoring all the results. Third, HLA-Arena enabled a more accurate virtual screening of HLA binders by combining sequence- and structure-based approaches.

These workflows can be modified to allow for additional analysis of the modeled pHLA complexes (eg, to perform molecular dynamics with OpenMM^{47,48} or cross-reactivity assessment).^{2,49,50} Thanks to high-performance computing and efficient sampling, molecular dynamics could play a bigger role in providing accurate estimates of pHLA binding affinity and complex stability.^{51,52}

HLA-Arena can be integrated into computational pipelines for basic cancer research or to help inform physicians in preclinical settings. It can be used to perform the large-scale modeling and selection of tumor-associated peptides, computer-aided design of altered peptide ligands, and study of T-cell cross-reactivity.^{2,8} In addition to HLA binding prediction, immunotherapy applications require identification of peptides that are uniquely displayed by cancer cells. This important task will be addressed in future updates of HLA-Arena.

It is important to note that HLA-Arena provides efficient solutions to sampling challenges associated with pHLA modeling^{8,23} and facilitates the integration of these solutions with other tools for structural analysis. However, the accuracy of structure-based peptide ranking is limited by existing scoring functions. As they improve, new scoring functions will be incorporated into HLA-Arena to replace current ones or be combined with consensus methods.^{53,54} In time, we expect that structure-based analyses will become essential to peptide target prediction for neoantigen discovery, vaccine development, and cancer immunotherapy, especially for patients with less prevalent HLA alleles.

AFFILIATIONS

¹Department of Computer Science, Rice University, Houston, TX

²Université Grenoble Alpes, Inria, Grenoble, France

³Department of Melanoma Medical Oncology–Research, The University of Texas MD Anderson Cancer Center, Houston, TX

CORRESPONDING AUTHOR

Lydia E. Kavraki, PhD, Rice University, MS132, P.O. Box 1892, Houston, TX 77251-1892; e-mail: kavraki@rice.edu.

EQUAL CONTRIBUTION

D.A.A. and J.R.A. contributed equally to this work.

SUPPORT

Supported in part by National Institutes of Health Grant No. 1R21CA209941-01 through the Informatics Technology for Cancer Research initiative of the National Cancer Institute; by the Cancer Prevention and Research Institute of Texas through Award No. RP170508; by a fellowship from the Gulf Cost Consortia on the Computational Cancer Biology Training Program (Grant No. RP170593); by a training fellowship from the National Library of Medicine Training Program in Biomedical Informatics (Grant No. T15LM007093); by funds from Rice University; and by a Big-Data Private-Cloud Research Cyberinfrastructure Major Research Instrumentation Program award, funded by the National Science Foundation (NSF) under Grant No. CNS-1338099. This work used the Extreme Science and Engineering Discovery Environment, which is supported by NSF Grant No. ACI-1548562; more specifically, the work involved the Stampede cluster at the Texas Advanced Computing Center, funded through Allocation No. MCB180187.

AUTHOR CONTRIBUTIONS

Conception and design: Dinler A. Antunes, Jayvee R. Abella, Sarah Hall-Swan, Didier Devaurs, Mark Moll, Gregory Lizée, Lydia E. Kavraki
Financial support: Gregory Lizée, Lydia E. Kavraki
Administrative support: Gregory Lizée, Lydia E. Kavraki
Collection and assembly of data: Dinler A. Antunes, Jayvee R. Abella, Sarah Hall-Swan, Anja Conev, Lydia E. Kavraki

Data analysis and interpretation: Dinler A. Antunes, Jayvee R. Abella, Didier Devaurs, Gregory Lizée, Lydia E. Kavraki

Manuscript writing: All authors

Final approval of manuscript: All authors

Accountable for all aspects of the work: All authors

AUTHORS' DISCLOSURES OF POTENTIAL CONFLICTS OF INTEREST

The following represents disclosure information provided by authors of this manuscript. All relationships are considered compensated unless otherwise noted. Relationships are self-held unless noted. I = Immediate Family Member, Inst = My Institution. Relationships may not relate to the subject matter of this manuscript. For more information about ASCO's conflict of interest policy, please refer to www.asco.org/rwc or ascopubs.org/cci/author-center.

Open Payments is a public database containing information reported by companies about payments made to US-licensed physicians ([Open Payments](http://OpenPayments)).

Jayvee R. Abella

Employment: Medical Informatics Corp

Lydia E. Kavraki

Consulting or advisory role: Allgens Biomedical Research, Biomedical and Informatics Consultants (I)

Patents, royalties, other intellectual property: Rice University (I)

Expert testimony: Dialysis Clinic, Maynard, Cooper & Gale (I)

Travel, accommodations, expenses: Beijing Advanced Medical Technologies (I)

No other potential conflicts of interest were reported.

ACKNOWLEDGMENT

We thank Romanos Fasoulis for his contributions in testing some of the components of HLA-Arena. HLA-Arena is made available through Docker Hub, under kavrakilab/hla-arena (Appendix provides installation details). The HLA-Arena Docker image also contains data related to the experiments described here, which can be reproduced as demo workflows. Additional information and documentation can also be found on GitHub at <https://github.com/KavrakiLab/hla-arena>.

REFERENCES

- Lizée G, Overwijk WW, Radvanyi L, et al: Harnessing the power of the immune system to target cancer. *Annu Rev Med* 64:71-90, 2013
- Antunes DA, Rigo MM, Freitas MV, et al: Interpreting T-cell cross-reactivity through structure: Implications for TCR-based cancer immunotherapy. *Front Immunol* 8:1210, 2017
- Degaque N, Brouard S, Souillou JP: Cross-reactivity of TCR repertoire: Current concepts, challenges, and implication for allotransplantation. *Front Immunol* 7:89, 2016
- Adams JJ, Narayanan S, Birnbaum ME, et al: Structural interplay between germline interactions and adaptive recognition determines the bandwidth of TCR-peptide-MHC cross-reactivity. *Nat Immunol* 17:87-94, 2016
- Antunes DA, Abella JR, Devaurs D, et al: Structure-based methods for binding mode and binding affinity prediction for peptide-MHC complexes. *Curr Top Med Chem* 18:2239-2255, 2018
- Vandiedonck C, Knight JC: The human major histocompatibility complex as a paradigm in genomics research. *Brief Funct Genomics Proteomics* 8:379-394, 2009
- Robinson J, Halliwell JA, Hayhurst JD, et al: The IPD and IMGT/HLA database: Allele variant databases. *Nucleic Acids Res* 43(D1):D423-D431, 2015
- Abella JR, Antunes DA, Clementi C, et al: APE-Gen: A fast method for generating ensembles of bound peptide-MHC conformations. *Molecules* 24:881, 2019
- Antunes DA, Moll M, Devaurs D, et al: DINC 2.0: A new protein-peptide docking webserver using an incremental approach. *Cancer Res* 77:e55-e57, 2017
- Antunes DA, Devaurs D, Moll M, et al: General prediction of peptide-MHC binding modes using incremental docking: A proof of concept. *Sci Rep* 8:4327, 2018
- Webb B, Sali A: Comparative protein structure modeling using MODELLER. *Curr Protoc Protein Sci* 86:1-2, 37, 2016
- O'Donnell T. J., Rubinsteyn A, Bonsack M, et al: MHCflurry: Open-source class I MHC binding affinity prediction. *Cell Syst* 7:129-132.e4, 2018
- Rose AS, Bradley AR, Valasatava Y, et al: NGL viewer: Web-based molecular graphics for large complexes. *Bioinformatics* 34:3755-3758, 2018
- Dos Santos Francisco R, Buhler S, Nunes J. M., et al: HLA supertype variation across populations: New insights into the role of natural selection in the evolution of HLA-A and HLA-B polymorphisms. *Immunogenetics* 67:651-663, 2015
- Wieczorek M, Abualrous ET, Sticht J, et al: Major histocompatibility complex (MHC) class I and MHC class II proteins: Conformational plasticity in antigen presentation. *Front Immunol* 8:292, 2017

16. Fodor J, Riley BT, Borg NA, et al: Previously hidden dynamics at the TCR-peptide-MHC interface revealed. *J Immunol* 200:4134-4145, 2018
17. Bordner AJ, Abagyan R: Ab initio prediction of peptide-MHC binding geometry for diverse class I MHC allotypes. *Proteins* 63:512-526, 2006
18. Khan JM, Ranganathan S: pDOCK: A new technique for rapid and accurate docking of peptide ligands to Major Histocompatibility Complexes. *Immunome Res* 6:S2, 2010 (suppl 1)
19. Chys P, Chacón P: Random coordinate descent with spinor-matrices and geometric filters for efficient loop closure. *J Chem Theory Comput* 9:1821-1829, 2013
20. Trott O, Olson AJ: AutoDock Vina: Improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *J Comput Chem* 31:455-461, 2010
21. Dhanik A, McMurray JS, Kaviraki LE: DINC: A new AutoDock-based protocol for docking large ligands. *BMC Struct Biol* 13:S11, 2013 (suppl 1)
22. Dhanik A, McMurray JS, Kaviraki LE: Binding modes of peptidomimetics designed to inhibit STAT3. *PLoS One* 7:e51603, 2012
23. Devaurs D, Antunes DA, Hall-Swan S, et al: Using parallelized incremental meta-docking can solve the conformational sampling issue when docking large ligands to proteins. *BMC Mol Cell Biol* 20:42, 2019
24. Eastman P, Swails J, Chodera JD, et al: OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLOS Comput Biol* 13:e1005659, 2017
25. Koes DR, Baumgartner MP, Camacho CJ: Lessons learned in empirical scoring with Smina from the CSAR 2011 benchmarking exercise. *J Chem Inf Model* 53:1893-1904, 2013
26. Flaticon: Free icons designed by ultimatearm. <https://www.flaticon.com/authors/ultimatearm>
27. Berman HM, Westbrook J, Feng Z, et al: The Protein Data Bank. *Nucleic Acids Res* 28:235-242, 2000
28. Sidney J, Peters B, Frahm N, et al: HLA class I supertypes: A revised and updated classification. *BMC Immunol* 9:1, 2008
29. Liu T, Pan X, Chao L, et al: Subangstrom accuracy in pHLA-I modeling by Rosetta FlexPepDock refinement protocol. *J Chem Inf Model* 54:2233-2242, 2014
30. Rigo MM, Antunes DA, Vaz de Freitas M, et al: DockTope: A Web-based tool for automated pMHC-I modelling. *Sci Rep* 5:18413, 2015
31. Kyeong HH, Choi Y, Kim HS: GradDock: Rapid simulation and tailored ranking functions for peptide-MHC class I docking. *Bioinformatics* 34:469-476, 2018
32. Maenaka K, Maenaka T, Tomiyama H, et al: Nonstandard peptide binding revealed by crystal structures of HLA-B*5101 complexed with HIV immunodominant epitopes. *J Immunol* 165:3260-3267, 2000
33. Morris GM, Huey R, Lindstrom W, et al: AutoDock4 and AutoDockTools4: Automated docking with selective receptor flexibility. *J Comput Chem* 30:2785-2791, 2009
34. Quiroga R, Villarreal MA: Vinardo: A scoring function based on Autodock Vina improves scoring, docking, and virtual screening. *PLoS One* 11:e0155183, 2016
35. Borbulevych OY, Insauido FK, Baxter TK, et al: Structures of MART-126/27-35 Peptide/HLA-A2 complexes reveal a remarkable disconnect between antigen structural homology and T cell recognition. *J Mol Biol* 372:1123-1136, 2007
36. Kawakami Y, Eliyahu S, Sakaguchi K, et al: Identification of the immunodominant peptides of the MART-1 human melanoma antigen recognized by the majority of HLA-A2-restricted tumor infiltrating lymphocytes. *J Exp Med* 180:347-352, 1994
37. Rivoltini L, Kawakami Y, Sakaguchi K, et al: Induction of tumor-reactive CTL from peripheral blood and tumor-infiltrating lymphocytes of melanoma patients by in vitro stimulation with an immunodominant peptide of the human melanoma antigen MART-1. *J Immunol* 154:2257-2265, 1995
38. Antunes DA, Devaurs D, Kaviraki LE: Understanding the challenges of protein flexibility in drug design. *Expert Opin Drug Discov* 10:1301-1313, 2015
39. Wang Z, Sun H, Yao X, et al: Comprehensive evaluation of ten docking programs on a diverse set of protein-ligand complexes: The prediction accuracy of sampling power and scoring power. *Phys Chem Chem Phys* 18:12964-12975, 2016
40. Vita R, Overton JA, Greenbaum JA, et al: The immune epitope database (IEDB) 3.0. *Nucleic Acids Res* 43:D405-D412, 2015
41. Nguyen NT, Nguyen TH, Pham TNH, et al: Autodock Vina adopts more accurate binding poses but Autodock4 forms better binding affinity. *J Chem Inf Model* 60:204-211, 2020
42. Chang MW, Ayeni C, Breuer S, et al: Virtual screening for HIV protease inhibitors: a comparison of AutoDock 4 and Vina. *PLoS One* 5:e11955, 2010
43. Shao W, Pedrioli PGA, Wolski W, et al: The SystemMHC Atlas project. *Nucleic Acids Res* 46:D1237-D1247, 2018
44. Jurtz V, Paul S, Andreatta M, et al: NetMHCpan-4.0: Improved peptide-MHC class I interaction predictions integrating eluted ligand and peptide binding affinity data. *J Immunol* 199:3360-3368, 2017
45. Bonsack M, Hoppe S, Winter J, et al: Performance evaluation of MHC class-I binding prediction tools based on an experimentally validated MHC-peptide binding data set. *Cancer Immunol Res* 7:719-736, 2019 [Erratum: *Cancer Immunol Res* 7:1221, 2019]
46. Zhao W, Sher X: Systematically benchmarking peptide-MHC binding predictors: From synthetic to naturally processed epitopes. *PLOS Comput Biol* 14:e1006457, 2018
47. Wieczorek M, Sticht J, Stolzenberg S, et al: MHC class II complexes sample intermediate states along the peptide exchange pathway. *Nat Commun* 7:13224, 2016
48. Pang YP, Elsbernd LR, Block MS, et al: Peptide-binding groove contraction linked to the lack of T cell response: Using complex structure and energy to identify neoantigens. *Immunohorizons* 2:216-225, 2018
49. Mendes MF, Antunes DA, Rigo MM, et al: Improved structural method for T-cell cross-reactivity prediction. *Mol Immunol* 67:303-310, 2015
50. Dhanik A, Kirshner JR, MacDonald D, et al: In-silico discovery of cancer-specific peptide-HLA complexes for targeted therapy. *BMC Bioinformatics* 17:286, 2016
51. Wan S, Knapp B, Wright DW, et al: Rapid, precise, and reproducible prediction of peptide-MHC binding affinities from molecular dynamics that correlate well with experiment. *J Chem Theory Comput* 11:3346-3356, 2015
52. Ochoa R, Laio A, Cossio P: Predicting the affinity of peptides to major histocompatibility complex class II by scoring molecular dynamics simulations. *J Chem Inf Model* 59:3464-3473, 2019
53. Kukol A: Consensus virtual screening approaches to predict protein ligands. *Eur J Med Chem* 46:4661-4664, 2011
54. Palacio-Rodríguez K, Lans I, Cavasotto CN, et al: Exponential consensus ranking improves the outcome in docking and receptor ensemble docking. *Sci Rep* 9:5142, 2019



APPENDIX

APE-GEN: FAST GENERATION OF PHLA BINDING MODE ENSEMBLES

We recently released a new tool, the anchored peptide-HLA (pHLA) ensemble generator (APE-Gen), which produces an ensemble of binding modes for a pHLA complex, starting from the sequences of a peptide and HLA receptor.⁸ APE-Gen involves an iterative process repeating the 3 following steps. First, the ends of the peptide backbone are anchored within known pockets in the HLA binding site using available backbone termini templates. Second, the peptide backbone is completed by applying the random coordinate descent loop modeling tool,¹⁹ which efficiently yields several valid backbone conformations. Third, side chains are added to the backbone conformations, and local optimization is performed with Smina²⁵ to fix steric clashes. This step considers full-peptide flexibility and binding site side-chain flexibility, producing a set of full-atom peptide conformations within the HLA binding site. After each such round of sampling, the highest-quality conformation (according to the internal scoring function, currently Vinardo) can be used as a template for the next round.

By generating a diverse ensemble of pHLA binding modes, APE-Gen implicitly accounts for the natural flexibility of peptides within the binding site. We showed that APE-Gen could reproduce the entire set of nonredundant classic class I pHLA structures available in the Protein Data Bank²⁷ (535 complexes at the time of the study).⁸ In that case, we used a single round of sampling per complex. The average root mean square deviation (RMSD) between modeled peptides and their corresponding crystal structure (considering all heavy atoms) was only 2.02 Å, which is considered an accurate reproduction. Even better results can be obtained when performing optimization and/or additional rounds of sampling, especially for longer peptides.⁸

APE-Gen is computationally efficient, producing dozens of binding modes in a few minutes on a standard desktop computer. It can be run for several peptides and a given HLA receptor, thus producing valuable information for peptide ranking and binding affinity prediction and enabling structure-based virtual screening of HLA binding peptides. We have also shown the potential benefits of APE-Gen when studying T-cell cross-reactivity.⁸

DINC: INCREMENTAL DOCKING OF PHLA COMPLEXES

In previous work, we presented a molecular docking approach called DINC (which stands for docking incrementally), specifically developed for large ligands, including peptides.²¹ The underlying idea is to incrementally dock larger and larger fragments of a ligand, instead of trying to dock it all at once. Note that this incremental docking process focuses on ligand flexibility, although selected receptor side chains can also be sampled. This process is parallelized to allow for broader sampling, by having several runs of docking performed independently at each step and grouping their results together. DINC is also a meta-docking method, in the sense that it relies on existing molecular docking tools, such as AutoDock4,³³ Vina,²⁰ and Smina²⁵ to perform the docking of the fragments at each step. As a consequence, fragment sampling and scoring can be performed by different tools.

The latest version of our software, called DINC 2.0, has been made available as a Web server.⁹ We recently showed that it performs a more exhaustive sampling than other docking approaches.²³ In that study, DINC was benchmarked using 5 public data sets including large ligands; it reproduced many crystal structures on which other docking tools had failed.²³ For example, it has been used to study the inhibition

of the Src homology 2 domain of STAT3 by peptidomimetics.²² We also showed that DINC could reproduce a diverse set of pHLA structures encompassing 10 HLA alleles and peptides with diverse binding modes; it achieved an average all-heavy-atom RMSD of 1.92 Å.¹⁰ Note that DINC is not limited to common class I HLA receptors, contrary to many related tools.⁵ It can be applied to complexes involving synthetic ligands, to rare and nonclassical class I HLAs, and potentially to class II HLA receptors.⁹ An updated version of DINC is made available through Docker Hub (docker pull kavrakilab/dinc-bin).

HLA-ARENA PERFORMANCE FOR VIRTUAL SCREENING

HLA-Arena provides the most efficient workflow available for structure-based virtual screening of HLA binders. For the experiment we report in the Results section, the breakdown of computing time is as follows: MHCflurry¹² needs approximately 15 seconds to screen the entire data set of 9,000 peptides. The homology modeling step takes approximately 3 minutes for each HLA allele and can be skipped for HLAs with available crystal structures. The APE-Gen step takes approximately 2 minutes per pHLA complex on a desktop computer with 6 to 8 threads. The (optional) rescoring takes approximately 2 seconds per complex using an HLA-Arena function that relies on Smina.²⁵ Therefore, running the entire workflow on a desktop computer takes approximately 86 hours with an MHCflurry threshold at 500 nM and approximately 300 hours with an MHCflurry threshold at 50,000 nM. This running time can be dramatically reduced if the APE-Gen step is executed on a cluster. For instance, on a machine with 64 threads, with an MHCflurry threshold at 500 nM or 50,000 nM, the same workflow could be executed in 5 or 19 hours, respectively (without rescoring). Future updates of HLA-Arena should provide additional resources for running workflows in a remote high-performance computing cluster.

HLA-ARENA INSTALLATION

1. If you do not already have it, install Docker for Mac or Windows (<https://www.docker.com/products/docker-desktop>) or for Linux (<https://docs.docker.com/install>).
2. In a command prompt, pull the HLA-Arena image from Docker Hub by typing:
docker pull kavrakilab/hla-arena
3. Create a folder in which you want to run the workflows (optional).
4. Copy HLA-Arena notebooks and associated data to your local machine by typing:
docker run --rm -v \$(pwd):/temp--entrypoint cp kavrakilab/hla-arena /hla_arena_data/data.tar.gz \ /temp/; tar -xzf data.tar.gz
5. Run HLA-Arena in this folder by typing:
docker run --rm -v \$(pwd):/data -p 8888:8888 \ --entrypoint="" kavrakilab/hla-arena jupyter \ notebook --port=8888 --no-browser \ --ip=0.0.0.0 --allow-root.
6. This should generate a URL with the following format:
http://127.0.0.1:8888/?token=<token_value>.
7. Copy and paste this URL into a browser, and open any available Jupyter notebook (ie, 1 of the files with extension .ipynb). Note that all the data created in the container will be saved inside sub-directories of the current folder.
8. Check out the file DOCUMENTATION (<https://kavrakilab.github.io/hla-arena/DOCUMENTATION.html>) for additional information on the workflows and available functions.




FIG A1. Lowest root mean square deviation (RMSD) binding modes sampled by DINC in the geometry prediction workflow. (A) Results of a self-docking experiment aimed at reproducing a crystal structure (with Protein Data Bank [PDB] code 1E27) involving a 9-mer peptide derived from HIV-1 and the HLA-B*51:01 receptor. This experiment was carried out with either 8 or 32 threads. Each bar corresponds to the so-called top RMSD conformation (ie, the conformation with the lowest RMSD to the target crystal structure) sampled in each of 5 replicated runs. Near-native peptide conformations (ie, conformations with an all-heavy-atom RMSD to the crystal structure < 2.5 Å) were sampled in all runs. The best conformation across all runs had an all-heavy-atom RMSD of 0.84 Å. (B) Results of a cross-docking experiment aimed at reproducing a crystal structure (with PDB code 2GTW) involving HLA-A*02:01 and a 9-mer peptide derived from the MART-1/melan-A protein. Near-native peptide conformations were sampled in 2 of 5 runs when using 8 threads and in 4 of 5 runs when using 32 threads. The best conformation sampled across all runs had an all-heavy-atom RMSD of 2.15 Å.




FIG A2. Lowest root mean square deviation (RMSD) binding mode sampled in a cross-docking experiment. Depicted in red is the lowest RMSD conformation sampled by DINC in the cross-docking experiment aimed at reproducing a crystal structure (with Protein Data Bank code 2GTW) involving HLA-A*02:01 and a 9-mer peptide derived from the MART-1/melan-A protein. The all-heavy-atom RMSD of this conformation to the crystal structure (depicted in blue) is only 2.15 Å. This conformation accurately reproduces the geometry of the first residue (p1), which has an unusual arrangement (ie, anchored in pocket B of the binding cleft).




FIG A3. Binding mode ensembles generated by APE-Gen include near-native peptide conformations. This plot aggregates the all-heavy-atom root mean square deviation (RMSD) between each conformation produced by APE-Gen for each peptide binder in our data set (Appendix Table A1) and its reference crystal structure. Results for conformations having undergone energy minimization with OpenMM²⁴ are also reported, although differences are subtle. These conformations were produced by a single round of sampling with APE-Gen.




FIG A4. Binding energy rankings associated with ensembles or single conformations. This plot reports correlations (assessed as Pearson's R) between experimentally determined binding affinities and structure-based binding energies predicted by the Vina scoring function using different procedures. More specifically, the binding energy of a given peptide can be defined as: the score of the conformation with the lowest RMSD to the crystal structure in the ensemble produced by APE-Gen ($R = 0.54$), the score of that same conformation minimized with OpenMM ($R = 0.68$), the median score within the ensemble of conformations produced by APE-Gen ($R = 0.74$), or the median score within that same ensemble after minimization with OpenMM ($R = 0.74$). Each point corresponds to a known peptide binder to HLA-A*02:01 (Appendix Table A1). Note that the nonbinders were not included in this analysis.

TABLE A1. Curated Data Set of Experimentally Determined Peptide Binders Restricted to HLA-A*02:01

Peptide	Method	Assay	Affinity (nM)	IEDB Label	IEDB ID	PDB ID
FLPSDFFPSV	Cellular/radioactivity	IC50	0.57	Positive-high	201178	3OXR
ALWGFFPVL	Purified/radioactivity	IC50	2.7	Positive	1775814	1LP9
LLFGYPVYV	Purified/radioactivity	KD (\approx IC50)	3.8	Positive-high	201486	1DUZ
CINGVCWTV	Purified/radioactivity	KD (\approx IC50)	55	Positive-medium	1955167	3MRG
VLRDDLLEA	Purified/fluorescence	IC50	365	Positive-low	1809531	3FT4
AAGIGILTV	Purified/radioactivity	IC50	395	Positive-medium	201470	3QFD
RQISQDVKL	Purified/radioactivity	KD (\approx IC50)	1,925	Positive	3243420	4N05
RGPGRAFVTI	Purified/radioactivity	KD (\approx IC50)	4,600	Positive-low	1022278	3ECB
ILKEPVHGV	Purified/fluorescence	IC50	7,082	Positive-medium	1783069	2X4U
EAAGIGILTV	Purified/fluorescence	IC50	14,560	Positive-low	2369616	2GT9
SLLMWITQC	Purified/radioactivity	KD (\approx IC50)	21,070	Positive-low	208218	2P5E
AAEQRRSTI	Cellular/fluorescence	IC50	> 70,000	Negative	1873146	—
DAKRNSKSL	Cellular/fluorescence	IC50	> 70,000	Negative	1872692	—
EIDVSEVKT	Cellular/fluorescence	IC50	> 70,000	Negative	1874620	—
ATKRYPGVM	Cellular/fluorescence	IC50	> 70,000	Negative	1875085	—
ETLNEYKQL	Cellular/fluorescence	IC50	> 70,000	Negative	1873828	—

NOTE. Selected methods include either cellular HLA or purified HLA, used for competitive radioactive or competitive fluorescence measurements. Assays measured either the half maximal inhibitory concentration (IC50) or the dissociation constant (KD). Nonbinders are characterized by the lack of a precise measurement in the Immune Epitope Database (IEDB; eg, affinity > 70,000 nM) and of a crystal structure in the Protein Data Bank (PDB).

Abbreviation: ID, identifier.