

Shape-restrained modelling of protein-small molecule complexes with HADDOCK

Panagiotis I. Koukos[#], Manon Réau[#], Alexandre M. J. J. Bonvin^{*}

Computational Structural Biology Group, Department of Chemistry, Faculty of Science,
Utrecht University, 3584CH, Utrecht, The Netherlands.

[#] These authors contributed equally

^{*} To whom correspondence should be addressed: a.m.j.j.bonvin@uu.nl, phone: +31 30 253 2652

Keywords: Molecular simulations, docking, interactions, structural biology, small molecule, proteins.

Abstract

Small molecule docking remains one of the most valuable computational techniques for the structure prediction of protein-small molecule complexes. It allows us to study the interactions between compounds and the protein receptors they target at atomic detail, in a timely and efficient manner. Here we present a new protocol in HADDOCK, our integrative modelling platform, which incorporates homology information for both receptor and compounds. It makes use of HADDOCK's unique ability to integrate information in the simulation to drive it toward conformations which agree with the provided data. The focal point is the use of shape restraints derived from homologous compounds bound to the target receptors. We have developed two protocols: In the first, the shape is composed of fake atom beads based on the position of the heavy atoms of the homologous template compound, whereas in the second the shape is additionally annotated with pharmacophore data, for some or all beads. For both protocols, ambiguous distance restraints are subsequently defined between those beads and the heavy atoms of the ligand to be docked. We have benchmarked the performance of these protocols with a fully unbound version of the widely used DUD-E dataset. In this unbound docking scenario, our template/shape-based docking protocol reaches an overall success rate of 81% on 99 complexes, which is close to the best results reported for bound docking on the DUD-E dataset.

Introduction

The importance of reliable methods for the docking of small molecule compounds to receptors of pharmaceutical interest cannot be understated. The nature of modern drug development practices dictates the gradual filtering of millions (perhaps even hundreds of millions) of compounds contained in virtual libraries to, ultimately, a few dozen lead compounds that can be further optimised before their clinical potential is investigated in animal and human trials [1,2]. This set of practices – collectively known as Computer-Aided Drug Design (CADD) – encompasses a variety of methods such as virtual screening of compounds, molecular docking with recent developments making use of machine learning-based approaches [3,4] and binding affinity prediction.

A relatively recent development in this space has been the advent of blind experiments (or exercises) focusing on pose prediction and binding affinity ranking/prediction. Examples of such initiatives are the – aptly-named – ‘Grand Challenges’ (GC), held by the Drug Design Data Resource (D3R) consortium, on an annual basis between 2015 and 2018 [5–8]. These experiments built on the experience of the earlier Community Structure-Activity Resource (CSAR) experiment [9]. The stated drives behind these initiatives have been the desire to establish – and push – state-of-the-art performance in these challenging modelling problems, develop commonly agreed-upon standards of evaluating and assessing aforementioned performance, codify best practices and facilitate communication and sharing of data between pharmaceutical companies and academic investigators. In these targets, these initiatives have been, largely, successful.

It is through our participation [10–12] with HADDOCK (**H**igh **A**mbiguity **D**riven **D**OCKing), our integrative modelling platform [13,14], in the D3R experiment that we initially developed docking protocols tailored to the peculiarities of protein-small molecule complex modelling.

Whereas our protocol of choice for the 2016 iteration of the GC did not earn us a spot among the top performing groups of that round, it did allow us to better understand the problems specific to protein-small molecule docking. The main takeaway point was that making use of the most closely related receptor for every target compound significantly improved the outcome of the modelling. In this case, the best receptor was identified after comparing crystallographic compounds with the target compound and selecting the receptor conformation with the most similar ligand. With that knowledge we optimised additional aspects of our approach for GC 2017 – always prioritising the use of high-quality experimental information for every step of the process. That revised protocol resulted in HADDOCK submitting one of the most accurate predictions for the pose prediction component of the challenge. We applied the same protocol in GC 2018 with equally good results.

Our successful D3R protocol can be summarised as: 1) After identifying highly homologous receptors with a co-crystallised compound, we compare the similarity of all crystallographic compounds to all target compounds and select the receptor conformation whose compound has the highest similarity to the compound to dock. 2) Prior to docking, we filter the generated conformers by comparing their 3D shape with that of the most similar crystallographic compound and select the 10 closest conformers in terms of shape similarity. 3) Finally, these 10 conformers are placed into the binding pocket by superimposing their shape onto the shape of the crystallographic compound and the model was refined in HADDOCK (i.e., no docking is performed).

While the use of shapes in protein-ligand modelling is not novel, even dating back to the earliest days of the field [15,16], with common approaches emphasising shape complementarity or overlap [17–19], to the best of our knowledge, HADDOCK’s ability to drive the desired compound into specific parts of the binding pocket via the use of shape information in

combination with other kinds of restraints in a unified integrative modelling framework is unique among modelling software.

Here, we present a new shape-based protocol which incorporates all the lessons that we have learned over three years of participating in the D3R blind docking experiment into a protocol designed for HADDOCK, bypassing one of the main limitations of the previous protocols – their reliance on external software for significant parts of the ligand-based modelling process. This limitation did not allow us to use the integrative modelling and flexible refinement capabilities of HADDOCK as the rigid-body and semi-flexible refinement stages were bypassed and only the final flexible refinement stage was performed. Also, the fact we are no longer dependent on commercial software for parts of the modelling workflow means we are free to develop this protocol into a pipeline accessible through the freely available HADDOCK webserver [14] (<https://wenmr.science.uu.nl/haddock2.4/>).

In the new shape-restrained protocols presented here, the principles underlying template identification and conformer generation procedures are similar to the ones described previously. The main difference is that, after identifying a suitable receptor template for each target, its bound compound is transformed into a shape consisting of fake beads. In the pharmacophore-based version of the protocol, a pharmacophore shape is only introduced if features can be assigned for that part of the template compound, otherwise a regular shape bead is used. Ambiguous distance restraints to guide the docking are then defined between those beads and the heavy atoms of the compound. There is no selection of conformers prior to docking, but instead, up to 50 conformers are docked into the receptor template. The most suited conformations are automatically selected during the docking by HADDOCK based on the shape restraints. These restraints are also drivers of possible conformational changes in the ligand during the flexible refinement stage of HADDOCK. We evaluate the performance of these new shape-restrained protocols on the Database of Useful Decoys-Enhanced (DUD-E) dataset that

contains crystallographic structures for 102 protein-small molecule complexes [20]. Among those, 58 proteins can be classified as protein kinases, nuclear receptors, proteases, GPCRs, cleaving enzymes, cyclooxygenases, cytochromes P450, ion channels or histone deacetylases and 54 are unique representants of a protein family [20,21]. The DUD-E and its previous versions have been mostly used to benchmark virtual screening tools [22,23]. When used to benchmark sampling tools, in most cases, the bound conformations of the target proteins were used [21,24]. The use of bound structures however, biases the outcome of the docking since it suggests that we know the exact conformation of the target proteins when bound to their ligand and that we effectively ignore the intrinsic flexibility of their binding site and any binding-induced conformational changes [25]. In this work, we introduce a new unbound benchmark version of the DUD-E that allows us to assess the performance of the HADDOCK shape-restrained protocols in more realistic conditions.

Results and Discussion

Dataset

We successfully identified templates for 99 of the 102 targets that are part of the DUD-E dataset (Tables S1 and S2). Those templates correspond to identical or homologous receptors that display the exact same sequence at the binding site level and that are co-crystalized with a different ligand. The amount of conformational changes in the binding sites ranges from 0.08Å to 4.01Å (see Tables S1 and S2). No usable templates could be identified for targets *cxc4*, *drd3* and *kpcb*. For *kpcb* and *drd3* we failed to identify any homologous templates whereas for *cxc4* we did identify some, but they were all in complex with the same compound as the reference receptor, so we discarded them. We identified at least one viable template for the remaining targets and selected one for docking based on the similarity of the reference and template compounds (see Methods). The similarity values range from almost 1 to below 0.2 indicating the presence of templates whose compounds are almost identical to their respective reference compounds on one end and templates that are almost entirely dissimilar on the other. The distribution of similarity values for the templates used for the two protocols can be seen in SI Figure 1. As we expected the similarity between template and reference compounds to be a limiting factor for the outcome of the modelling, we decided to investigate that relationship for the shape protocol, by identifying an additional low similarity template for all targets whose original template had a similarity value greater than 0.8 (see SI Figure 1). In this way, we can compare the performance in the subset of the dataset which includes both low- and high-similarity templates. In total, for 34 of the 35 targets for which the original similarity value was greater than 0.8 we could identify an additional, low-similarity template (no alternative templates could be identified for target *kith*).

Conformer generation

Prior to assessing the docking performance, we evaluate the performance of the conformer generation as we expect it to have a significant impact on the outcome of the docking, since starting ligand conformations that are very different to the reference compound would need to undergo significant conformational rearrangements. In summary, RDKit and the settings we have chosen (see Methods) perform well as we are able to generate at least one acceptable pose (conformers whose heavy-atom RMSD to the reference bound ligand conformation is $\leq 2\text{\AA}$) for all but five target compounds. Specifically, the mean percentage of acceptable poses is 66 ± 36 (median: 78), however there are 11 targets for which less than 10% of the generated conformers are acceptable, 5 of those with no generated poses below 2\AA (see SI Figure 2).

Docking

Figure 1 shows the docking strategy behind the shape- and pharmacophore-driven protocols. In short, after selecting a suitable template, a shape based on the template compound is generated, restraints are defined between template shapes and target conformers, and docking is performed with HADDOCK (see Methods for more details).




Figure 1: Illustration of the shape-based and the pharmacophore-based docking protocol. **Panel A** shows a suitable receptor template identified based on the similarity of its bound ligand and the ligand to be docked (see Methods). **Panel B** shows the heavy atoms of the crystallographic compound transformed into shape beads (**B.1**) or pharmacophore beads with the coloured beads representing different pharmacophore properties (**B.2**). The crystallographic compound is then removed from the pocket and restraints are defined between the beads and the conformers: In the shape-based protocol (**C.1**) restraints are defined between all atoms of the compound and all beads of the shape, while in the pharmacophore-based protocol (**C.2**), restraints are defined between atoms of the compound and beads that share identical pharmacophore features. **Panel D** shows a docked model superimposed onto the template structure. The protein receptor is shown as slate cartoon, the crystallographic compound as white sticks, the generated and docked compounds as orange sticks and the shape beads as transparent orange spheres. All molecular graphics were generated with PyMOL [26].

Figure 2 shows the success rate for the two protocols with the ‘shape’ and ‘pharm’ groups highlighting the performance of the shape- and pharmacophore-driven protocols, respectively. The success rate is defined as the percentage of targets for which a model of acceptable (or better) quality has been generated within the topN ranked models. For the shape protocol, the performance is very poor when considering only the top ranked models during the rigid body stage but improves significantly in the refinement stage increasing to 58.6% and 71.7% when considering the top1 and top5 models, respectively, with an overall success rate of 81.8% when considering all refined models. The respective values for the pharmacophore protocol are 53.5%, 63.6% and 75.8% indicating a slightly lower performance compared to the shape

protocol. Note again that those success rates are for an unbound docking scenario (i.e., both ligand and receptor starting conformations deviate from their reference bound conformations in the complex).




Figure 2: Success rate as a function of N -ranked models considered (TopN) for top 1, 5, 10, 50, 100, 200 and 1000 models for the initial rigid-body docking (it0) and after flexible refinement (it1). Each stacked bar plot corresponds to a specific cutoff with the colour of the bar indicating the distribution of model quality for that cutoff. Dark green, light green, light blue and light grey correspond to high-, medium-, acceptable- and near acceptable-quality models, respectively. The respective RMSD cutoffs are $\leq 0.5\text{\AA}$, $\leq 1\text{\AA}$, $\leq 2\text{\AA}$ and $\leq 2.5\text{\AA}$ when considering all heavy atoms of the compound after superimposing on the binding site backbone atoms of the receptors. Results have been grouped by protocol on the y axis and stage on the x axis. The black vertical line in the 'it0' indicates the cutoff at which models proceed to flexible refinement.

This pattern holds even if we include the “Near Acceptable” models – models whose RMSD to their respective receptor lies between 2 and 2.5\AA – which cannot be considered near-native but are still likely to offer very relevant biological insights. When including these models, the respective percentages become 69.7, 77.8 and 86.7 for the shape top1, top5 and overall, and 66.7, 74.8 and 81.2 for the pharmacophore top1, top5 and overall. The performance of the pharmacophore protocol when considering only the top ranked model of it0 is worth mentioning as the success rate (when including the Near Acceptable models) stands at 63.6%, which makes it particularly impressive considering only rigid-body motions are allowed in it0. Despite this,

the distribution of model quality also seems to favour the shape protocol as higher-quality models are generated for more targets at all stages and cutoffs.

The success rate, however, only tells one side of the story, as can be seen in Figure 3 which shows the distribution of model quality as a function of the rank for all targets for the refinement stage (Figure 3, bottom panel). Comparing the performance of the two protocols in this way reveals that the pharmacophore protocol produces models of acceptable (or better) quality much more consistently than the shape protocol but the shape protocol has higher coverage (see SI Figure 3 for rigid-body stage results). The top panel of Figure 3 shows the similarity between template and reference compounds (using the Tversky and Tanimoto metrics, respectively; see Methods for details) in bars and the 3D shape overlap in points. Comparing the performance of the two protocols with these metrics in mind reveals the 3D shape overlap to be a limiting factor for the outcome of the modelling, acting almost like a binary classifier, with targets whose overlap is below 0.5 almost never achieving good performance, whereas for those whose overlap is greater than 0.5 success is almost guaranteed, in particular for those whose overlap exceeds 0.75 (e.g., *adrb1* and *mcr*). Further comparing the performance of the two protocols, reveals that overlap is the feature that determines the outcome even when considering the same target. For example, for target *aces* the similarity and overlap values are 0.99 and 0.91, and 0.6 and 0.4, for the shape- and pharmacophore-based protocols, respectively, with the shape protocol yielding top-ranked, high-quality models and the pharmacophore-based protocol only producing Near Acceptable models despite a lower binding site RMSD for the template chosen for the pharmacophore-based protocol (0.09Å against 0.2Å for the shape-based protocol).




Figure 3: Comparison of template quality and performance for the two protocols. The top panel highlights the similarity (orange and blue bars for the shape- and pharmacophore-based protocols, respectively) and overlap (green and dark green dots for the shape- and pharmacophore-based protocols, respectively) between reference and template compounds. The bottom panel compares the performance for the semi-flexible refinement stage of HADDOCK (it1) for the two protocols. Each column corresponds to one target with the Y axis reflecting the ranking of models (ranks close to 0 refer to top-ranked models and those close to 200 to bottom-ranked models) and the colour of each model reflecting its quality (see description of Figure 2).

To further investigate the impact of template selection on the outcome of the modelling we compared the performance of the shape-based protocol in a subset of the dataset for which we identified both high- (similarity ≥ 0.8) and low-similarity templates (similarity around 0.4). Not only is the difference in terms of success rate stark, with the high-similarity subset observing success rates of 76.5%, 88.2% and 91.2% when considering top1, top5 and overall, respectively, against 44.1%, 47.1% and 73.5% for the low-similarity subset (still a very decent performance),

but the simulations with the high-similarity templates clearly results in distributions that are more populated with higher-quality models (see SI Figure 4).

Figure 4 A-B illustrates a docking scenario in which both shape-restrained protocols are expected to perform well due to the high similarity between the template and the target compounds ($T_v = 0.714$ and $T_c = 0.641$). Both protocols generate acceptable quality models at top1, with IL-RMSD (interface ligand-RMSD; see Methods for details) values of 1.60 Å and 1.37 Å for the shape- and the pharmacophore-based protocol respectively. Figure 4 C-D illustrates a more difficult scenario where the selected template differs substantially from the target compound (*csf1r* – $T_v = 0.44$ and *mk14* – $T_c = 0.47$, see Figure 6). We still observe good performance, with acceptable models generated at top2 for *csf1r* (1.74 Å) with the shape-based protocol, and at top5 for *mk14* (1.89 Å) with the pharmacophore-based protocol. This is mainly explained by a good shape overlap between the template and the target compounds (0.56 and 0.52).

improvement) for the acceptable models after refinement. The distributions of the RMSD between the aligned template and reference small molecules, referred to as the ligand-ligand RMSD, show that both protocols also tend to generate models in which the conformation get closer to the reference for most targets (Figure S5). Three eloquent visual examples are provided in Figure 5B. The first one shows that the IL-RMSD of the top ranking it1 model for target *lck* (0.4 Å) is improved by 5.2 Å as compared to the corresponding model at it0. This target is considered an “easy case” since it is associated with an overlap of 0.82, and a binding site RMSD of 0.1 Å between the target and the template. The second example, an “intermediate case” (overlap of 0.57, binding site RMSD of 0.6 Å), shows that the IL-RMSD of the top5 it1 model for target *gria2* (1.0 Å) is improved by 4.1 Å. Finally, we present a “difficult case” with the target *pa2ga* for which no acceptable conformer was generated. Based on the excellent template that is selected by the pharmacophore-based protocol, the generated model ranked top10 at it1 significantly improves on the ligand-ligand RMSD metric (Figure S5) as well as the IL-RMSD to the reference (it1 IL-RMSD = 1.8 Å with a Δ IL-RMSD = 3.9 Å). The ligand-ligand RMSD values are 0.35 Å, 0.76 Å and 1.72 Å for the three cases, respectively, indicating the excellent agreement in binding mode between our models and the reference structures.

The overall larger improvement observed for the shape-based protocol as compared to the pharmacophore-based protocol could be explained by the higher degree of freedom provided by the dummy beads as compared to the pharmacophore-beads.

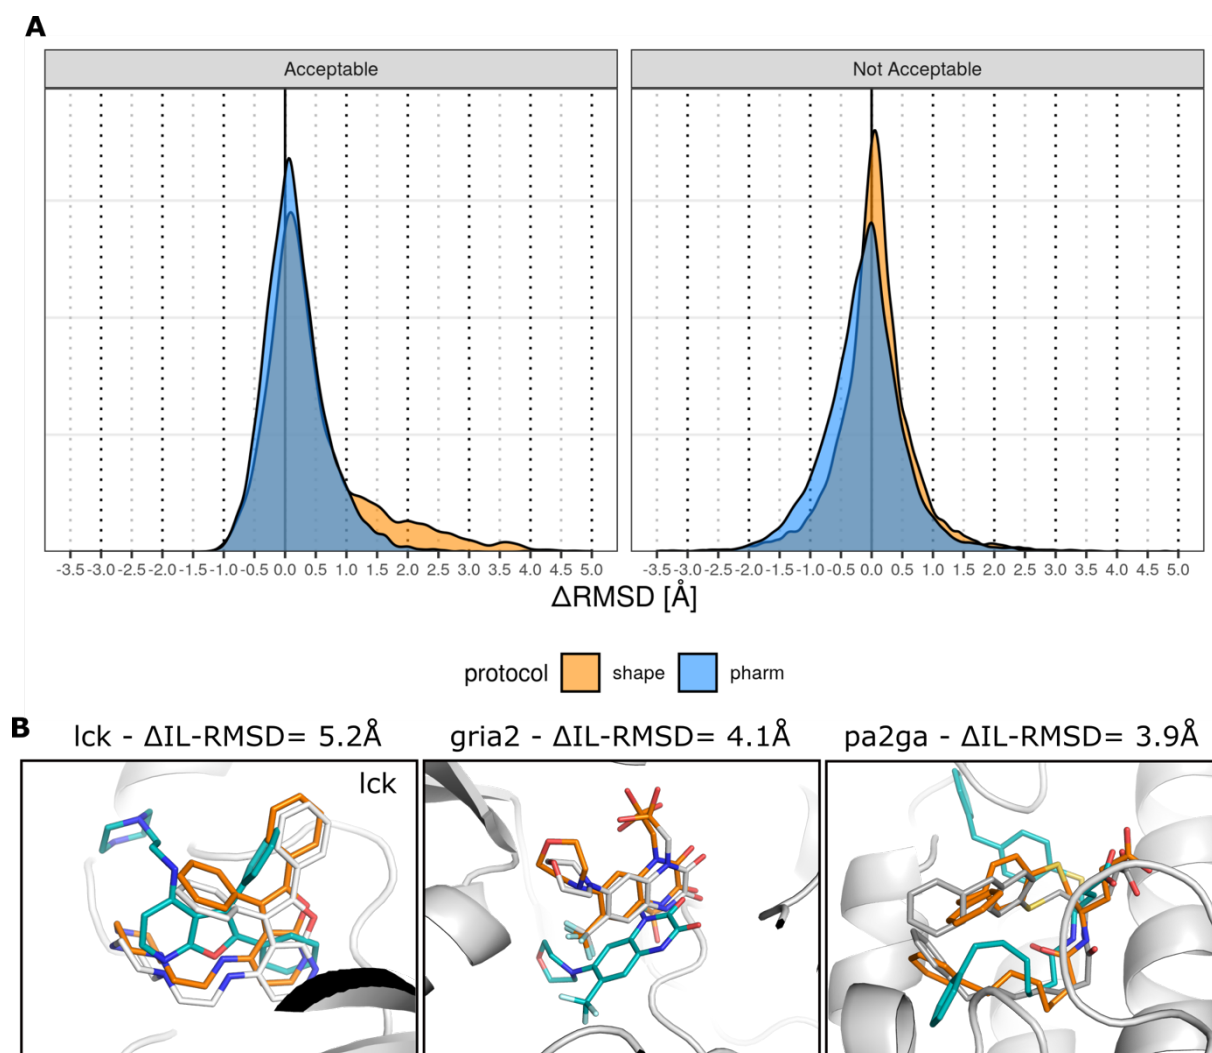



Figure 5. Illustration of the impact of the semi-flexible refinement on the docking models quality. Panel A shows the distribution of the Δ IL-RMSD between the models before and after semi-flexible refinement, computed as the IL-RMSD at it0 – the IL-RMSD at it1 for a given model. The figure is divided into two parts: the left sub-panel shows the distribution of Δ IL-RMSD for models that are acceptable at it1, the right sub-panel shows the distribution of Δ IL-RMSD for models that are not acceptable at it1. Positive values in the distribution indicate better it1 performance. Panel B illustrates three cases with large IL-RMSD improvement between the model generated at it0 (blue) and the refined model at it1 (orange) as compared to the bound conformation in the reference structure (grey). Templates 2of4, 3bki and 1kqu were used for cases lck, gria2 and pa2ga, respectively.

Peculiar cases

Since we always learn from errors, we investigated cases where shape-restrained docking fails. As the first HADDOCK step (it0) keeps both the receptor and the ligand rigid, it is essential to provide as starting point different conformers of the ligand to be docked. When the conformer generator fails in generating near-native conformations, the docking is likely to fail as well. Our results confirm this assumption as we observed that few to no acceptable models were generated

in 7 cases with less than 10% acceptable conformers generated, (e.g. *cdk2*, *cp3a4*, *fkbl1a*, *fnta*, *hivpr*, *pa2ga* and *xiap*). A second source of error is the binding site location of the template compound that may differ from the binding site of the target compound (Figure 6B.5), as highlighted by the Pearson correlation of 0.72 and 0.59 observed between the distance in geometric centre of the template and the target compounds and the best IL-RMSD obtained with the shape- and pharmacophore-based protocol respectively (Figure S6). External expert-eye or experimental knowledge should alleviate this issue and help filtering out potential erroneous template compounds in early stages of the protocol. For this work we however followed an automated pipeline. Third, if both protocols are capable of generating good quality models where the template compounds show poor similarity with the target compounds (Figure 6A), the opposite also happens. In some cases, our initial assumption stating that similar compounds that bind the same target receptor should in principle adopt similar binding modes is not validated. This is the case for the *hivint* and *bace1* target and template compounds as illustrated on Figure 6B. The different structures of *hivint* complexed with N-benzyl indolin-2-ones reported in the original paper of the target structure [27] display various binding modes of the compounds in the same binding site. In the case of *bace1*, two binding modes of iminohydantoin *bace1* inhibitors are reported with a 180° flip of the iminohydantoin core [28]. Finally, a visual inspection could beneficially help in discarding template compounds that share poor similarity with the target compound in terms of shape, flexibility and physico-chemical properties despite a meaningful similarity score (e.g., Figure 6B.5).

A.



B.




Figure 6: Illustration of template (shape-based protocol: orange, pharmacophore-based protocol: blue) and target (grey) compounds. **Panel A** shows different template/target compound couples of low similarity, i.e. Tanimoto (Tc) or Tversky (Tv) coefficient < 0.5, that lead to successful docking. Templates 3dpk, 2i0v, 5x23, 3bki, 3ha8 and 1i7i were used for cases 1-6, respectively. **Panel B** show scenarios that lead to docking failure for different possible reasons. While displaying high similarity with the target compounds, the hivint pharmacophore-based template compound adopts a different binding mode (**B.1**), and the shape-based template compounds leads to an unexpected failure that cannot be explained neither by the coverage nor their respective binding site (**B.2**). The shape-based template compound for bace1, that is the same as the pharmacophore-based template, adopts a different binding mode than the target despite the high similarity (**B.3**). Low shape, flexibility and physico-chemical similarities are observed between the template and the target compound for pygm (**B.4**); The case **B.5** is an example of poor overlap between template and target compounds that is responsible for poor performance. Templates 3nf9, 3nf6, 3l5c, 3g72 and 6mob were used for cases 1-5, respectively.

Which protocol should be favoured?

The shape-based protocol yields better performances in terms of coverage, i.e. it generates near-native models, including near-acceptable models, for more DUD-E complexes (top1: 69.7%, top5: 77.8 %, overall: 87%) than the pharmacophore-based protocol (top1: 66.7%, top5: 74.8 %, overall: 81.2 %). In an ideal scenario where we merge the results of the two protocols and

pick up the best model (i.e. lowest IL-RMSD) per rank, the global coverage is remarkably improved (top1: 81%, top5: 86%, overall: 91%). As illustrated in Fig 2, it means that even though the protocols display comparable performance on most targets, one protocol could be favoured over the other in some cases (e.g. *aces*, *fkbl1a*, *gria2*, *hivint*, *mapk1*, *mk14*, *pde5a*, *prgr*). We identified the overlap between the template and the target compounds shapes as the determinant indicator of the success of the shape- and the pharmacophore-based docking (Figure S7). The overlap does not only inform on the proximity between the binding sites of the template and the target compounds, it also tells about the binding mode similarity. However, this information is accessible only in benchmarking conditions where the crystallographic binding mode of the target ligand is known. To mimic real-life conditions, we investigated if different combinations of descriptor filters extracted from the unbound form of the target and the template compounds (similarity, difference in molecular weight, number of rotatable bounds and accessible surface area) could lead to an enrichment in the success-rate. We identified no evident conditions to favour one protocol over the other, but it is worth mentioning that both protocols achieve high success rates when the template and target share moderate to high similarity ($\geq 0.4/0.3$ for shape and pharmacophore, respectively) (Figure 7). As such, the shape-based protocol performance (including near-acceptable models) rises to 74.1 %, 83.5 % and 92.9 % at top1, top5 and overall (85 targets), when considering a Tversky similarity ≥ 0.4 . With similar conditions, i.e. a Tanimoto coefficient over 0.3, the performance stands at 71.9 %, 80.9 % and 87.6 % at top1, top5 and overall (89 targets) with the pharmacophore-based protocol. Adding a molecular weight filter ($|MW_{\text{ref}} - MW_{\text{temp}}| < 50$ g/mol) slightly improves the performance which rises to 77.5 %, 89.8 % for the shape top1 and top5, and reaches an impressive overall performance of 100% when considering all generated models. A lower overall performance is reached with the pharmacophore-based protocol (93.2 %), but we observe a higher early enrichment of 81.4 % at top1. Of note, 49 and 59 DUD-E targets fulfil

these two criteria for the shape- and the pharmacophore-based protocol, respectively. All the aforementioned analyses suggest the shape-based protocol should be favoured for the sake of coverage. However, the conditions we identified could be used to favour one of these protocols depending on the available data and the intended goal of the study.




Figure 7: Success rate of the shape- and pharmacophore-based protocol on DUD-E complexes (49 and 59, respectively) that match two criteria: 1) the similarity coefficient between template and the reference compounds is higher than 0.4 and 0.3, respectively, and 2) the difference in heavy atom molecular weight is below 50g.mol⁻¹.

Comparison to state-of-the-art docking tools

A similar benchmarking study was conducted with four commercial docking programs, namely Gold [29], Glide [30–32], Surflex [33] and FlexX [34], which are all widely used in industry [21] and have better sampling power than most non-commercial docking programs [35]. In this study, two of the 102 DUD-E targets (*aofb* and *casp3*) were discarded as they show covalent binding to their respective receptor. Of note, those two targets are part of our dataset and the three we are missing are part of their dataset. To fairly compare the performance, we discarded *aofb* and *casp3* from our results. Commercial programs were evaluated in bound-unbound

docking mode, i.e. unbound conformations of the compounds have been docked onto the bound conformation of the receptor (i.e. no conformational changes required on the protein side), while we did perform fully unbound docking which represents a significantly more difficult – yet more realistic – scenario [11]. Our two protocols show very competitive overall performance with those commercial docking programs, the shape-based protocol ranking 2nd with a success rate of 81 % just behind Surflex (84 %) (Figure S8), and the pharmacophore-based protocol ranking 5th with 74 %. We also observe rather competitive early enrichment with 54 % and 58 % success at top1 for the fully unbound pharmacophore- and the shape-based protocols respectively against 64 % and 65 % for Gold and Glide respectively. It is interesting to note that all commercial software fail to generate near-acceptable models for 6 DUD-E targets on which at least one of our protocols also fails (*bace1*, *cdk2*, *fkbl1a*, *fnta*, *hivint*, *pgh1*) for different reasons (see *Peculiar Cases*).

Conclusion and Perspectives

We have presented two new protocols for the modelling of protein-small molecule complexes with HADDOCK, which make excellent use of the integrative modelling capabilities of the platform, to yield very promising results. Selection of suitable templates for the modelling based on ligand similarity and use of shape data during the docking lie at the core of both approaches. In the first, the shape is only defined in geometrical terms whereas in the second it is additionally annotated with pharmacophore data, thus naming the two protocols shape- and pharmacophore-based. The dataset on which these protocols were benchmarked is based on the widely used DUD-E dataset, which was extended to make use of unbound receptors (or rather receptors bound to another ligand, which is commonly defined as cross-docking). This “unbound” DUD-E dataset, which is freely available in ready-to-dock format on GitHub, is composed of diverse

targets of various difficulty degrees. Our analysis has revealed that the main limiting factor for the outcome of the docking with our shape protocols is the 3D overlap between template and reference compounds, in other words, whether the template and reference compounds not only occupy the same binding pocket but have similar binding modes. We have also shown that we can predict the reliability of the modelling based on features that can be computed in real-world usage: The difference in molecular weight between template and reference compounds and their molecular similarity. When we select the subset of our dataset for which the compounds are not too dissimilar in terms of molecular similarity (T_v and T_c greater or equal than 0.4 and 0.3 for the shape- and pharmacophore-based protocols, respectively) and molecular weight ($|MW_{ref} - MW_{temp}| < 50$ g/mol) we achieve success rates of 65.3%, 81.6% and 91.8%, and 67.8%, 76.2% and 91.5% when considering top1, top5 and overall, for the shape- and pharmacophore-based protocols, respectively. Using all targets, the performance remains rather high with 58.5%, 71.7% and 81.8%, and 53.5%, 63.6% and 75.8% for the two protocols. The performance of our protocols is on par with that of leading small molecule docking software, despite the fact our approach was benchmarked on a fully unbound dataset. As such it is representative of the performance one can expect when using this pipeline in a real-life scenario where one would not have access to the bound form of the receptor and/or that of the compound.

Despite this promising performance our method still suffers from a few limiting factors. Chief among those is the dependency on a template whose compound occupies the same binding pose as that of the ligand to be modelled for a successful prediction. This was identified as the most severe limiting factor for both protocols. It is a problem which is not unique to these protocols but is shared by all data-driven methods, HADDOCK being no exception [36]. Overcoming this obstacle is no easy task and will require the development of methodology for the automated identification, combination and weighting of multiple template shapes and receptors in order to increase the coverage of the binding site. Another area where these protocols could benefit from

further improvements is the ranking of solutions during the flexible stage as can be seen by the difference in success rate at the top1 or top5 level and overall. Also worth noting is the fact that, although regular users of the command-line version of HADDOCK will be able to make use of this pipeline, it might be well out of reach for most researchers who have no expertise in programming and the use of computational chemistry toolkits. Therefore, it would be of interest if these protocols were to be made available through the freely accessible HADDOCK webserver (<https://wenmr.science.uu.nl/haddock2.4/>). We hope to expand on these three points in the near future.

Limitations and future improvements mentioned in the previous paragraph notwithstanding, the protocols we are proposing achieve high performance on experimental conditions that mimic real-life ones as closely as possible using unbound structures for receptors as well as compounds. Furthermore, the performance we are reporting is on par with that of leading commercial docking platforms when benchmarked on partially bound datasets, and in many cases exceeding it, thus making it an excellent option for the modelling of protein-small molecule complexes.

Materials and Methods

Dataset

To validate our new approach, we decided to benchmark the performance of these protocols against the Enhanced Database of Useful Decoys dataset (DUD-E – <http://dude.docking.org/>), one of the most widely used small-molecule docking benchmarks [20]. In total, the DUD-E consists of 102 targets with each target corresponding to a single protein receptor (usually of pharmaceutical interest) with a compound bound to it. The Protein Data Bank (PDB) [37] entries associated with every DUD-E target became our reference receptors. We discarded three

targets (*cxcr4*, *drd3*, *kpcb*) because we could not identify any viable templates for them (see next section) bringing our total to 99 targets. The structures of the reference receptors were downloaded from the PDB instead of the DUD-E website to avoid any post-processing that might have been applied to them.

Template Identification

The procedure we followed for identifying a template for each of the 99 targets in our dataset can be summarised in the following steps:

1. Search the PDB for appropriate template structures. We used two approaches in conjunction with a few criteria to identify templates:
 - a. Search for receptors whose sequence was more than a predetermined cutoff similar to the sequence of our reference receptor. The precise value we used for the sequence similarity cutoff differed from target to target because of the intrinsic features of some targets. As an example, some targets were membrane receptors that had been co-crystallised with a lysozyme domain. The presence of the lysozyme sequence, lysozyme being one of the most commonly encountered domains in the PDB, affected the results and we had to take that into account when adjusting the sequence similarity cutoff. The default value was 90% and in cases such as the one mentioned above was adjusted after manual inspection of the results.
 - b. Search for receptors bound to compounds which were similar to the reference compound. For this we made use of the “Fingerprint Similarity” feature of the “Chemical” tab of the Advanced search functionality of the RCSB portal of the PDB (<https://www.rcsb.org/search/advanced>).
2. After pooling the results from both searches, we iteratively removed unsuitable templates based on a few criteria:

- a. Templates which contained only irrelevant compounds such as crystallisation buffer artefacts or some of the most commonly encountered compounds such as hemes.
 - b. Incorrect templates based on sequence identity (regarding the templates identified through ligand similarity).
 - c. Templates whose compounds were found to be stereoisomers of the reference compounds.
3. After removal of all unsuitable templates, we calculated the pairwise similarity between all reference and template compounds.
 - a. *Shape-based protocol*: The similarity metric we chose for this is the Tversky similarity [38] (with weights for the reference and template compounds set to 0.8 and 0.2, respectively) [39,40] computed over the Maximum Common Substructure (MCS) as identified with the rdFMCS implementation of RDKit (version 2020.09.3) [41].
 - b. *Pharmacophore-based protocol*: We computed the Tanimoto coefficient over 2D pharmacophore fingerprints generated with RDKit [41] using the default pharmacophore fingerprint factory.
4. After calculating all similarities, we rank the templates according to the similarity of their compound to the reference compound and select the one with the highest similarity with the condition that the template and the target share the same binding site (defined as $< 5 \text{ \AA}$ around the bound small molecule), i.e. 100% sequence identity in the binding site.

Conformer generation

To ensure there is no bias during the docking, instead of starting from compounds bound to the reference – or other – receptors, we generate 3D conformations of the target compounds starting from their isomeric SMILES [42,43] with RDKit using the 2020 parameters (only the small aliphatic ring subset) with energy minimisation and the ETKDG algorithm [44,45]. We cap the maximum number of conformers to 50 and provide the ensemble of conformers to HADDOCK for docking.

Docking

System Preparation

Prior to docking we process the structures to ensure the residue numbering between template and reference structures is consistent (a necessity for the analysis), rename residues, cofactors and ions according to HADDOCK specifications, remove nonbiological symmetric units and remove all irrelevant artefacts such as crystallisation buffers. We generated topology files of the compounds to be docked using PRODRG (version 070118.0614).

Shape-based protocol

After identifying one receptor template per target, we transform all heavy atoms of its compound into dummy beads (those do not interact with the remaining of the system). We then define ambiguous distance restraints with an upper limit of 1 Å between the shape beads and the heavy atoms of the compound to be docked. The nature of the restraints creates an additional consideration, specifically what should be the “origin” and “target” of the restraints? In this protocol, the directionality of the restraints depends on the size of the reference and template compounds, in terms of heavy atom count. Specifically, the restraints are always defined from smaller to larger. If the reference compound is larger than the template compound, the restraints are defined from each shape bead to any compound heavy atom. If the opposite is true (the

template compound is larger than the reference compound) then restraints are defined between each heavy atom of the conformer to any shape bead. These ambiguous distance restraints effectively enforce that ligand atoms and beads must overlap. Depending on the directionality of the restraints part of the ligand might remain outside the shape defined by the beads and vice-versa. All shape restraints are used during the simulation.

Pharmacophore-based protocol

In addition to the transformation of the template atoms into dummy beads, we assigned pharmacophore features to the beads to add physico-chemical properties to the spatial information they hold. We assigned a pharmacophore features (H-bond donor, H-bond acceptor, negative ionizable, positive ionizable, zinc binder, aromatic, hydrophobic and lumped hydrophobic) to each atom with RDKit using the default SMARTS-based feature definition. If no pharmacophore feature is assigned to an atom, a regular shape bead is defined. We defined distance restraints with an upper limit of 1 Å between the beads and the atoms of the compounds to be docked provided that they share the same pharmacophore feature – or no pharmacophore feature. In this protocol, we cannot impose restraints from the smaller to the larger anymore because restraints are defined between features of the same class -- some classes may be more populated in the reference compound while others may be more populated in the template compound. Instead, we defined all restraints from the compound to the beads. Similarly to the shape-based protocol, all restraints are considered during the docking process.

Fig. 1 illustrates the shape-based and the pharmacophore-based docking protocol.

HADDOCK

For the docking we use the January 2021 release of the command-line version of HADDOCK 2.4. The number of models generated in the initial rigid-body docking stage (it0) of HADDOCK

is set to 20 times the number of starting ligand conformations. This sampling ratio was found to be the most performant during benchmarking (data not shown). If the number of it0 models that are sampled is larger than 200 then only the top 200 it0 models proceeded to flexible refinement, otherwise all models do. We only perform the semi-flexible refinement stage of HADDOCK (it1) and skip full water refinement (itw) as this final stage is not improving the results as already remarked from our D3R participation. The positions of both the receptor and its associated shape are fixed in their original orientations while the ligand is translated away from the protein and randomly rotated for each docking trial. The shape is kept rigid throughout the protocol while the receptor interface and the ligand become flexible during the refinement stage. Systematic sampling of 180° rotations along the interface is disabled for it0. We also scale down the intermolecular interactions during the rigid body stage to facilitate the insertion of the ligand into the binding pocket and accordingly exclude the vdW energy term during the scoring of the rigid-body models. The modified parameter settings for HADDOCK are summarized in Table S3.

Other than the above defined modifications, the scoring function used is the default scoring function of HADDOCK. Its functional form, specific for protein-ligand docking for the two stages is:

$$HS - it0 = 0.0 * E_{vdw} + 1.0 * E_{elec} + 1.0 * E_{desolv} + 0.01 * E_{AIR} - 0.01 * BSA$$

$$HS - it1 = 1.0 * E_{vdw} + 1.0 * E_{elec} + 1.0 * E_{desolv} + 0.1 * E_{AIR} - 0.01 * BSA$$

where E_{vdw} , E_{elec} and E_{desolv} stand for van der Waals, Coulomb electrostatics and desolvation energies, respectively and BSA for the buried surface area. The non-bonded components of the score (E_{vdw} , E_{elec}) are calculated with the OPLS forcefield [46]. The desolvation energy is a solvent accessible surface area-dependent empirical term [47] which estimates the energetic gain or penalty of burying specific sidechains upon complex formation.

Evaluation of results

We evaluate the quality of the generated models according to their structural deviation from the reference structures. For this we use the interface-ligand RMSD (IL-RMSD), which is the RMSD calculated over all heavy atoms of the ligand after superimposing on all backbone atoms of the interface of the receptor. Models with an IL-RMSD of less than 0.5 Å, between 0.5 and 1 Å, between 1 and 2 Å, between 2 and 2.5 Å and over 2.5 Å are classified as high-, medium-, acceptable-, near acceptable- and low-quality, respectively. The initial fitting was performed using the McLachlan algorithm [48] as implemented in the program ProFit (Martin, A.C.R., <http://www.bioinf.org.uk/software/profit/> - available through SBGrid [49]). Calculation of symmetry-corrected RMSD values for the compounds were performed with *obrms* from the Open Babel distribution (<https://github.com/openbabel/openbabel> - version 3.1.1) [50]. As part of the analysis and interpretation of the results we also present a classification of target difficulty based on the 3D shape overlap of target and reference compounds after fitting on the binding site backbone atoms of their respective receptors. Fitting was performed with ProFit using the same settings as previously mentioned and the 3D shape overlap was calculated with the *Exact Overlap* metric of the shape toolkit of OpenEye (release 2020.2.0) [51–53].

Data availability

All docking models generated during the benchmarking of the two protocols with HADDOCK 2.4 are made available through our laboratory data collection at <https://data.sbgrid.org/labs/32/> [54] (*Note for reviewers: The full dataset will be deposited upon acceptance*). The unbound DUD-E benchmark data set, with all the code, docking input and parameter files, results and analysis files are made available through GitHub (<https://github.com/haddock/shape-restrained-haddock>).

Acknowledgments

We would like to acknowledge support from the European Union Horizon 2020 projects BioExcel (675728, 823830) and EOSC-hub (777536) and from the Innovative Medicines Initiative 2 Joint Undertaking (JU) under grant agreement No 101005077. The JU receives support from the European Union's Horizon 2020 research and innovation programme, EFPIA, BILL & MELINDA GATES FOUNDATION, GLOBAL HEALTH DRUG DISCOVERY INSTITUTE and UNIVERSITY OF DUNDEE.

Abbreviations Used

HADDOCK, High Ambiguity Driven DOCKing; DUD-E, Database of Useful Decoys: Enhanced; GC, Grand Challenge; D3R, Drug Design Data Resource; CSAR, Community Structure-Activity Resource; IL-RMSD, Interface Ligand-RMSD; Δ IL-RMSD; delta IL-RMSD; MCS; Maximum Common Substructure.

Bibliography

- [1] G. Schneider, U. Fechner, Computer-based de novo design of drug-like molecules, *Nat. Rev. Drug Discov.* 4 (2005) 649–663. doi:10.1038/nrd1799.
- [2] S. Mandal, M. Moudgil, S.K. Mandal, Rational drug design, *Eur. J. Pharmacol.* 625 (2009) 90–100. doi:10.1016/j.ejphar.2009.06.065.
- [3] G. Schneider, Automating drug discovery, *Nat. Rev. Drug Discov.* 17 (2018) 97–113. doi:10.1038/nrd.2017.232.
- [4] J. Vamathevan, D. Clark, P. Czodrowski, I. Dunham, E. Ferran, G. Lee, B. Li, A. Madabhushi, P. Shah, M. Spitzer, S. Zhao, Applications of machine learning in drug discovery and development, *Nat. Rev. Drug Discov.* 18 (2019) 463–477. doi:10.1038/s41573-019-0024-5.
- [5] S. Gathiaka, S. Liu, M. Chiu, H. Yang, J.A. Stuckey, Y.N. Kang, J. Delproposto, G. Kubish, J.B. Dunbar, H.A. Carlson, S.K. Burley, W.P. Walters, R.E. Amaro, V.A. Feher, M.K. Gilson, D3R grand challenge

- 2015: Evaluation of protein–ligand pose and affinity predictions, *J. Comput. Aided. Mol. Des.* 30 (2016) 651–668. doi:10.1007/s10822-016-9946-8.
- [6] Z. Gaieb, S. Liu, S. Gathiaka, M. Chiu, H. Yang, C. Shao, V.A. Feher, W.P. Walters, B. Kuhn, M.G. Rudolph, S.K. Burley, M.K. Gilson, R.E. Amaro, D3R Grand Challenge 2: blind prediction of protein–ligand poses, affinity rankings, and relative binding free energies, *J. Comput. Aided. Mol. Des.* 32 (2018) 1–20. doi:10.1007/s10822-017-0088-4.
- [7] Z. Gaieb, C.D. Parks, M. Chiu, H. Yang, C. Shao, W.P. Walters, M.H. Lambert, N. Nevins, S.D. Bembenek, M.K. Ameriks, T. Mirzadegan, S.K. Burley, R.E. Amaro, M.K. Gilson, D3R Grand Challenge 3: blind prediction of protein–ligand poses and affinity rankings, *J. Comput. Aided. Mol. Des.* 33 (2019) 1–18. doi:10.1007/s10822-018-0180-4.
- [8] C.D. Parks, Z. Gaieb, M. Chiu, H. Yang, C. Shao, W.P. Walters, J.M. Jansen, G. McGaughey, R.A. Lewis, S.D. Bembenek, M.K. Ameriks, T. Mirzadegan, S.K. Burley, R.E. Amaro, M.K. Gilson, D3R grand challenge 4: blind prediction of protein–ligand poses, affinity rankings, and relative binding free energies, *J. Comput. Aided. Mol. Des.* 34 (2020) 99–119. doi:10.1007/s10822-020-00289-y.
- [9] H.A. Carlson, Lessons Learned over Four Benchmark Exercises from the Community Structure-Activity Resource, *J. Chem. Inf. Model.* 56 (2016) 951–954. doi:10.1021/acs.jcim.6b00182.
- [10] Z. Kurkuoglu, P.I. Koukos, N. Citro, M.E. Trellet, J.P.G.L.M. Rodrigues, I.S. Moreira, J. Roel-Touris, A.S.J. Melquiond, C. Geng, J. Schaarschmidt, L.C. Xue, A. Vangone, A.M.J.J. Bonvin, Performance of HADDOCK and a simple contact-based protein–ligand binding affinity predictor in the D3R Grand Challenge 2, *J. Comput. Aided. Mol. Des.* 32 (2018) 175–185. doi:10.1007/s10822-017-0049-y.
- [11] P.I. Koukos, L.C. Xue, A.M.J.J. Bonvin, Protein–ligand pose and affinity prediction: Lessons from D3R Grand Challenge 3, *J. Comput. Aided. Mol. Des.* 33 (2019) 83–91. doi:10.1007/s10822-018-0148-4.
- [12] A. Basciu, P.I. Koukos, G. Mallocci, A.M.J.J. Bonvin, A. V. Vargiu, Coupling enhanced sampling of the apo-receptor with template-based ligand conformers selection: performance in pose prediction in the D3R Grand Challenge 4, *J. Comput. Aided. Mol. Des.* (2019). doi:10.1007/s10822-019-00244-6.
- [13] C. Dominguez, R. Boelens, A.M.J.J. Bonvin, HADDOCK: A Protein–Protein Docking Approach Based on Biochemical or Biophysical Information, *J. Am. Chem. Soc.* 125 (2003) 1731–1737. doi:10.1021/ja026939x.
- [14] G.C.P. van Zundert, J.P.G.L.M. Rodrigues, M. Trellet, C. Schmitz, P.L. Kastiris, E. Karaca, A.S.J. Melquiond, M. van Dijk, S.J. de Vries, A.M.J.J. Bonvin, The HADDOCK2.2 Web Server: User-Friendly Integrative Modeling of Biomolecular Complexes, *J. Mol. Biol.* 428 (2016) 720–725. doi:10.1016/j.jmb.2015.09.014.
- [15] I.D. Kuntz, J.M. Blaney, S.J. Oatley, R. Langridge, T.E. Ferrin, A geometric approach to macromolecule–ligand interactions, *J. Mol. Biol.* 161 (1982) 269–288. doi:10.1016/0022-2836(82)90153-X.
- [16] R.L. DesJarlais, I.D. Kuntz, R.P. Sheridan, R. Venkataraghavan, J.S. Dixon, Docking Flexible Ligands to Macromolecular Receptors by Molecular Shape, *J. Med. Chem.* 29 (1986) 2149–2153. doi:10.1021/jm00161a004.
- [17] D.B. Kitchen, H. Decornez, J.R. Furr, J. Bajorath, Docking and scoring in virtual screening for drug discovery: Methods and applications, *Nat. Rev. Drug Discov.* 3 (2004) 935–949. doi:10.1038/nrd1549.
- [18] C.M. Venkatachalam, X. Jiang, T. Oldfield, M. Waldman, LigandFit: A novel method for the shape-directed rapid docking of ligands to protein active sites, *J. Mol. Graph. Model.* 21 (2003) 289–307. doi:10.1016/S1093-3263(02)00164-X.
- [19] P.C.D. Hawkins, A.G. Skillman, A. Nicholls, Comparison of Shape-Matching and Docking as Virtual Screening Tools, *J. Med. Chem.* 50 (2007) 74–82. doi:10.1021/jm0603365.
- [20] M.M. Mysinger, M. Carchia, J.J. Irwin, B.K. Shoichet, Directory of useful decoys, enhanced (DUD-E): Better ligands and decoys for better benchmarking, *J. Med. Chem.* 55 (2012) 6582–6594. doi:10.1021/jm300687e.
- [21] L. Chaput, L. Mouawad, Efficient conformational sampling and weak scoring in docking programs? Strategy of the wisdom of crowds, *J. Cheminform.* 9 (2017) 37. doi:10.1186/s13321-017-0227-x.

- [22] L. Chaput, J. Martinez-Sanz, N. Saettel, L. Mouawad, Benchmark of four popular virtual screening programs: construction of the active/decoy dataset remains a major determinant of measured performance, *J. Cheminform.* 8 (2016) 56. doi:10.1186/s13321-016-0167-x.
- [23] M. Wójcikowski, P.J. Ballester, P. Siedlecki, Performance of machine-learning scoring functions in structure-based virtual screening, *Sci. Rep.* 7 (2017) 1–10. doi:10.1038/srep46710.
- [24] J.B. Cross, D.C. Thompson, B.K. Rai, J.C. Baber, K.Y. Fan, Y. Hu, C. Humblet, Comparison of several molecular docking programs: Pose prediction and virtual screening accuracy, *J. Chem. Inf. Model.* 49 (2009) 1455–1474. doi:10.1021/ci900056c.
- [25] W. Zhang, E.W. Bell, M. Yin, Y. Zhang, EDock: Blind protein-ligand docking by replica-exchange monte carlo simulation, *J. Cheminform.* 12 (2020) 37. doi:10.1186/s13321-020-00440-9.
- [26] The PyMOL Molecular Graphics System, Version 1.8 Schrödinger, LLC., (n.d.).
- [27] D.I. Rhodes, T.S. Peat, N. Vandegraaff, D. Jeevarajah, G. Le, E.D. Jones, J.A. Smith, J.A.V. Coates, L.J. Winfield, N. Thienthong, J. Newman, D. Lucent, J.H. Ryan, G.P. Savage, C.L. Francis, J.J. Deadman, Structural basis for a new mechanism of inhibition of HIV-1 integrase identified by fragment screening and structure-based design, *Antivir. Chem. Chemother.* 21 (2011) 155–168. doi:10.3851/IMP1716.
- [28] Z. Zhu, Z.Y. Sun, Y. Ye, J. Voigt, C. Strickland, E.M. Smith, J. Cumming, L. Wang, J. Wong, Y. Sen Wang, D.F. Wyss, X. Chen, R. Kuvelkar, M.E. Kennedy, L. Favreau, E. Parker, B.A. McKittrick, A. Stamford, M. Czarniecki, W. Greenlee, J.C. Hunter, Discovery of cyclic acylguanidines as highly potent and selective β -site amyloid cleaving enzyme (BACE) inhibitors: Part I - Inhibitor design and validation, *J. Med. Chem.* 53 (2010) 951–965. doi:10.1021/jm901408p.
- [29] G. Jones, P. Willett, R.C. Glen, A.R. Leach, R. Taylor, Development and validation of a genetic algorithm for flexible docking, *J. Mol. Biol.* 267 (1997) 727–748. doi:10.1006/jmbi.1996.0897.
- [30] R.A. Friesner, J.L. Banks, R.B. Murphy, T.A. Halgren, J.J. Klicic, D.T. Mainz, M.P. Repasky, E.H. Knoll, M. Shelley, J.K. Perry, D.E. Shaw, P. Francis, P.S. Shenkin, Glide: A New Approach for Rapid, Accurate Docking and Scoring. 1. Method and Assessment of Docking Accuracy, *J. Med. Chem.* 47 (2004) 1739–1749. doi:10.1021/jm0306430.
- [31] T.A. Halgren, R.B. Murphy, R.A. Friesner, H.S. Beard, L.L. Frye, W.T. Pollard, J.L. Banks, Glide: A New Approach for Rapid, Accurate Docking and Scoring. 2. Enrichment Factors in Database Screening, *J. Med. Chem.* 47 (2004) 1750–1759. doi:10.1021/jm030644s.
- [32] R.A. Friesner, R.B. Murphy, M.P. Repasky, L.L. Frye, J.R. Greenwood, T.A. Halgren, P.C. Sanschagrin, D.T. Mainz, Extra precision glide: Docking and scoring incorporating a model of hydrophobic enclosure for protein-ligand complexes, *J. Med. Chem.* 49 (2006) 6177–6196. doi:10.1021/jm051256o.
- [33] A.N. Jain, Surflex: Fully automatic flexible molecular docking using a molecular similarity-based search engine, *J. Med. Chem.* 46 (2003) 499–511. doi:10.1021/jm020406h.
- [34] M. Rarey, B. Kramer, T. Lengauer, G. Klebe, A fast flexible docking method using an incremental construction algorithm, *J. Mol. Biol.* 261 (1996) 470–489. doi:10.1006/jmbi.1996.0477.
- [35] Z. Wang, H. Sun, X. Yao, D. Li, L. Xu, Y. Li, S. Tian, T. Hou, 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 (2016) 12964–12975. doi:10.1039/c6cp01555g.
- [36] A.D.J. van Dijk, S.J. de Vries, C. Dominguez, H. Chen, H.-X. Zhou, A.M.J.J. Bonvin, Data-driven docking: HADDOCK's adventures in CAPRI, *Proteins Struct. Funct. Bioinforma.* 60 (2005) 232–238. doi:10.1002/prot.20563.
- [37] H.M. Berman, The Protein Data Bank, *Nucleic Acids Res.* 28 (2000) 235–242. doi:10.1093/nar/28.1.235.
- [38] A. Tversky, Features of similarity, *Psychol. Rev.* 84 (1977) 327–352. doi:10.1037/0033-295X.84.4.327.
- [39] D. Horvath, G. Marcou, A. Varnek, Do not hesitate to use tversky - And other hints for successful active analogue searches with feature count descriptors, *J. Chem. Inf. Model.* 53 (2013) 1543–1562. doi:10.1021/ci400106g.
- [40] S. O'Hagan, D.B. Kell, MetMaxStruct: A Tversky-Similarity-Based Strategy for Analysing the

- (Sub)Structural Similarities of Drugs and Endogenous Metabolites, *Front. Pharmacol.* 7 (2016) 266. doi:10.3389/fphar.2016.00266.
- [41] G. Landrum, P. Tosco, B. Kelley, sriniker, gedec, NadineSchneider, R. Vianello, Ric, A. Dalke, B. Cole, AlexanderSavelyev, M. Swain, S. Turk, D. N. A. Vaucher, E. Kawashima, M. Wójcikowski, D. Probst, guillaume godin, D. Cosgrove, A. Pahl, JP, F. Berenger, strets123, JLVarjo, N. O'Boyle, P. Fuller, J.H. Jensen, G. Sforna, DoliathGavid, rdkit/rdkit: 2020_03_1 (Q1 2020) Release, (2020). doi:10.5281/zenodo.3732262.
- [42] D. Weininger, SMILES, a Chemical Language and Information System: 1: Introduction to Methodology and Encoding Rules, *J. Chem. Inf. Comput. Sci.* 28 (1988) 31–36. doi:10.1021/ci00057a005.
- [43] D. Weininger, A. Weininger, J.L. Weininger, SMILES. 2. Algorithm for Generation of Unique SMILES Notation, *J. Chem. Inf. Comput. Sci.* 29 (1989) 97–101. doi:10.1021/ci00062a008.
- [44] S. Riniker, G.A. Landrum, Better Informed Distance Geometry: Using What We Know to Improve Conformation Generation, *J. Chem. Inf. Model.* 55 (2015) 2562–2574. doi:10.1021/acs.jcim.5b00654.
- [45] S. Wang, J. Witek, G.A. Landrum, S. Riniker, Improving Conformer Generation for Small Rings and Macrocycles Based on Distance Geometry and Experimental Torsional-Angle Preferences, *J. Chem. Inf. Model.* 60 (2020) 2044–2058. doi:10.1021/acs.jcim.0c00025.
- [46] W.L. Jorgensen, J. Tirado-Rives, The OPLS [optimized potentials for liquid simulations] potential functions for proteins, energy minimizations for crystals of cyclic peptides and crambin, *J. Am. Chem. Soc.* 110 (1988) 1657–1666. doi:10.1021/ja00214a001.
- [47] J. Fernández-Recio, M. Totrov, R. Abagyan, Identification of protein-protein interaction sites from docking energy landscapes., *J. Mol. Biol.* 335 (2004) 843–65. doi:10.1016/j.jmb.2003.10.069.
- [48] A.D. McLachlan, Rapid comparison of protein structures, *Acta Crystallogr. Sect. A.* 38 (1982) 871–873. doi:10.1107/S0567739482001806.
- [49] A. Morin, B. Eisenbraun, J. Key, P.C. Sanschagrin, M.A. Timony, M. Ottaviano, P. Sliz, Collaboration gets the most out of software, *Elife.* 2 (2013). doi:10.7554/eLife.01456.
- [50] N.M. O'Boyle, M. Banck, C.A. James, C. Morley, T. Vandermeersch, G.R. Hutchison, Open Babel: An Open chemical toolbox, *J. Cheminform.* 3 (2011) 33. doi:10.1186/1758-2946-3-33.
- [51] J.A. Haigh, B.T. Pickup, J.A. Grant, A. Nicholls, Small molecule shape-fingerprints, *J. Chem. Inf. Model.* 45 (2005) 673–684. doi:10.1021/ci049651v.
- [52] J. Boström, K. Berggren, T. Elebring, P.J. Greasley, M. Wilstermann, Scaffold hopping, synthesis and structure-activity relationships of 5,6-diaryl-pyrazine-2-amide derivatives: A novel series of CB1 receptor antagonists, *Bioorganic Med. Chem.* 15 (2007) 4077–4084. doi:10.1016/j.bmc.2007.03.075.
- [53] Shape Toolkit 2020.2.2 OpenEye Scientific Software, Santa Fe, NM. <https://www.eyesopen.com>, (n.d.).
- [54] P.A. Meyer, S. Socias, J. Key, E. Ransey, E.C. Tjon, A. Buschiazzi, M. Lei, C. Botka, J. Withrow, D. Neau, K. Rajashankar, K.S. Anderson, R.H. Baxter, S.C. Blacklow, T.J. Boggon, A.M.J.J. Bonvin, D. Borek, T.J. Brett, A. Caflisch, C.-I. Chang, W.J. Chazin, K.D. Corbett, M.S. Cosgrove, S. Crosson, S. Dhe-Paganon, E. Di Cera, C.L. Drennan, M.J. Eck, B.F. Eichman, Q.R. Fan, A.R. Ferré-D'Amaré, J. Christopher Fromme, K.C. Garcia, R. Gaudet, P. Gong, S.C. Harrison, E.E. Heldwein, Z. Jia, R.J. Keenan, A.C. Kruse, M. Kvansakul, J.S. McLellan, Y. Modis, Y. Nam, Z. Otwinowski, E.F. Pai, P.J.B. Pereira, C. Petosa, C.S. Raman, T.A. Rapoport, A. Roll-Mecak, M.K. Rosen, G. Rudenko, J. Schlessinger, T.U. Schwartz, Y. Shamoo, H. Sondermann, Y.J. Tao, N.H. Tolia, O. V. Tsodikov, K.D. Westover, H. Wu, I. Foster, J.S. Fraser, F.R.N.C. Maia, T. Gonen, T. Kirchhausen, K. Diederichs, M. Crosas, P. Sliz, Data publication with the structural biology data grid supports live analysis, *Nat. Commun.* 7 (2016) 10882. doi:10.1038/ncomms10882.

Table of contents graphic

