ELSEVIER

Predicting druggable binding sites at the protein–protein interface

Jonathan C. Fuller, Nicholas J. Burgoyne and Richard M. Jackson

Institute of Molecular and Cellular Biology and Astbury Centre for Structural Molecular Biology, University of Leeds, Leeds LS2 9JT, UK

Protein–protein interfaces are highly attractive targets for drug discovery because they are involved in a large number of disease pathways where therapeutic intervention would bring widespread benefit. Recent successes have challenged the widely held belief that these targets are 'undruggable'. The pocket finding algorithms described here show marked differences between the binding pockets that define protein–protein interactions (PPIs) and those that define protein–ligand interactions (PLIs) of currently marketed drugs. In the case of PPIs, drug discovery methods that simultaneously target several small pockets at the protein–protein interface are likely to increase the chances of success in this new and important field of therapeutics.

Most validated drug targets for which marketed drugs have been developed target small molecule protein–ligand interactions (PLIs). Indeed, research efforts have focused on a relatively small number of validated diseases, usually involving a single gene product, which is usually a PLI involving an enzyme or an integral membrane receptor. There is a growing need to develop new strategies that allow therapeutic intervention in difficult drug targets such as protein–protein interactions (PPIs). The discovery of small molecules that modulate PPIs has, however, been largely unsuccessful. PPI targets have come to be thought of as 'high risk', 'difficult' targets – perhaps even 'undruggable' – by many in the pharmaceutical industry. This view has recently been challenged by Wells and McClendon [1] and there is a growing body of evidence suggesting that PPIs may yet be suitable targets if certain hurdles can be overcome [2–5]. Given the importance of this class of interaction and the fact that the PPI targets are relatively well validated in terms of biological function and their role in disease [6], there is a growing determination to exploit this new class of targets. There are several PPI drugs already on the market: tirofiban targeting the integrins (cardiovascular conditions) [7]; and maraviroc targeting CCR5–gp120 interactions (HIV) [8], and several new drugs entering Phase II clinical trials [9]. Therefore it is likely to be only a matter of time before more marketed drugs become available.

Corresponding author: Jackson, R.M. (r.m.jackson@leeds.ac.uk)

Here, we define PPI inhibitors as small molecule inhibitors that directly compete with one of the natural protein partners of a discontinuous protein–protein interface [1]. This excludes PPIs involving enzyme–protein inhibitors (e.g. trypsin–pancreatic trypsin inhibitor) and also protein–linear-peptide motifs (e.g. SH2 domains [6]) as well as allosteric inhibitors (e.g. those targeted against TNFα [2]). New methods to elucidate protein–linear-peptide motifs have recently been discussed by Neduva and Russell [10]. There has been some success in developing small-molecule inhibitors of protein–linear-peptide motifs in the case of the integrins [2]. Research into developing small-molecule inhibitors that disrupt discontinuous protein–protein interfaces has received considerable attention in recent years and there are several reviews on this topic [1–5], including the use of computational tools to identify PPI inhibitors [11]. In a comprehensive review of the field to date Wells and McClendon address a widely held belief that PPI interfaces are large, flat and often featureless and therefore are unsuitable as drug targets [1]. It would appear that ligand pockets are a feature of these interfaces, something that is further supported by ligand-bound structures and molecular dynamics simulations [12]. Indeed, traditional methods for drug discovery such as high-throughput screening or fragment screening have been moderately successful, with small molecules achieving affinities comparable to those of the partner proteins [1].

Here, we address the challenges involved in characterising the nature of potential small-molecule binding sites on PPIs and understanding their potential for inhibition by small-molecule

1359-6446/06/$ - see front matter doi:10.1016/j.drudis.2008.10.009

drugs, an area that has received much less attention. The druggability of the proteome has been discussed from several perspectives, with varying suggestions as to the number of proteins that might be targeted. Overington *et al.* found that from 1204 distinct drugs approved by the FDA those with a known mode-of-action target 324 distinct molecular targets, of which 266 are human-genome-derived and the rest are from pathogenic sources [13]. This number has barely increased in the two years following the study, as is noted by Landry and Gies [14]. In terms of understanding druggability from a structure-based perspective there has been much less work. Hajduk *et al.* have shown that NMR-based fragment screening hit-rates are related to the druggability of sites [15]. Their observations allowed training of 'druggability indices' such as charge, hydrophobicity and shape to be determined from the frequency of NMR screening hits. They provided a comprehensive analysis of protein druggability using 3D structural information on protein targets [15]. They note that the current computational tools used to identify pockets are generally not capable of discriminating between druggable and non-druggable pockets. Recently, Chene proposed a decision tree to aid in the identification of interfaces amenable to drug discovery efforts [3] and a wide range of topics relating to computational determination of binding epitopes have been reviewed by Gonzalez-Ruiz and Gohlke [16]. The druggability of targets has also been addressed from a purely computational perspective by Cheng *et al.* [17]. They noted that the presence of a pocket on the protein surface is a 'necessary but not sufficient' criterion for the modulation of the target by drug-like small molecules. They started with the hypothesis that for orally available drugs a pocket on a target protein has a maximal achievable non-covalent binding affinity. This allowed them to distinguish between druggable and non-druggable sites based on a binding site desolvation model [17]. On this basis one PPI site that they specifically calculated to be druggable was MDM2, with a predicted maximal affinity of 0.1 nM, whereas benzodiazepinedione (currently the best experimentally determined inhibitor) has an affinity of 67 nM [1], suggesting that there may be room for improvement.

A diverse range of approaches have been taken to inhibit PPIs [2]. Several therapeutics have been developed that retain a largely peptidic character [2], but this generally results in therapeutics that are undesirable owing to issues with bioavailability. Several properties desirable for small-molecule therapeutics were noted in the seminal paper describing the 'Rule of Five' [18]. These have also been further discussed in the context of ligand efficiency [19]. In this context, it has been noted that small-molecule PPI inhibitors tend to have a slightly larger than desirable molecular weight range. There are, however, several examples where large lead compounds have been successfully optimised such that they retain acceptable pharmacokinetic properties [20].

Several researchers have expressed a positive outlook on the potential of PPI inhibitor drug discovery; however, whilst progress is being made there is plenty of room for improvement. Here, we discuss current computational methods of binding site detection, and compare and contrast their use in describing different types of protein interactions, including PLI, PPI and PPI inhibitor interactions as well as protein–small-molecule ligands that are marketed drugs, which we call PDIs. Finally, we demonstrate how they may be useful for practical application in the field of PPI inhibitor drug discovery.

Current methods for pocket detection

A major issue that hinders the comparison of binding pockets at the protein surface is the lack of a standard definition of what constitutes a pocket. Geometric descriptions of the 'depth' or 'size' of binding pockets are method-dependent and subjective. Indeed, several commonly used pocket detection methodologies, including our own called Q-SiteFinder [21] have been recently reviewed and give different and potentially conflicting descriptions of a pocket's size and location [22]. Further advances in binding site prediction have been proposed, some of which are briefly summarised. LIGSITE^{CSC} [23] is an extension and implementation of the LIGSITE [24] algorithm, adding information on the degree of conservation of the involved surface residues. PocketPicker [25] uses a buriedness index at grid points, and then clusters those with values lying within a cutoff range; it is conceptually similar to Q-SiteFinder. SURFNET-ConSurf [26] uses the SURFNET [27] program to identify pockets on the protein surface, and retains regions that are more highly conserved using the ConSurf-HSSP database [28]. Landon *et al.* presented a method called CS-Map [29] which performs computational solvent mapping in an attempt to recreate the multiple solvent crystal structures (MSCS) results of Mattos and Ringe [30]. AutoLigand [31] has been developed using a grid-based representation of the binding affinity potential in a similar way to Q-SiteFinder, to define envelopes of maximal affinity. Affinity potentials are generated for six atom types (aliphatic carbon, aromatic carbon, hydrogen, oxygen, nitrogen and sulphur) and the best envelope within the energy grid is calculated using a three-step process of flood fill, local migration and ray-casting neighbourhood search.

Previously, we proposed an energy-based method called Q-SiteFinder, which describes a volume envelope where an atom can interact favourably at the protein surface [21]. Pockets are ranked according to interaction energy, and it is assumed that these relate to locations where a putative ligand could bind and optimise its van der Waals interaction energy. This is defined as a volume envelope in which a van der Waals methyl probe has an interaction potential below a defined threshold (−1.4 kcal/mol). The interaction is non-specific and general, because the probe interacts only via van der Waals forces; however, it is indicative of a region where any interacting atom would return a favourable interaction with the protein surface. Here, we refer to this as the 'active volume' for a binding pocket. Those with the most favourable binding energy could be expected to correspond closely to high-affinity binding sites or 'hot-spots' [32]. We have shown that this method is highly effective in predicting binding sites for both bound and unbound proteins, and, unlike pocket detection methods that rely on geometric measures, the volumes of the predicted pockets appear to show little dependence on the total protein volume, with pockets being similar in volume to the ligands they contain [21].

Case study using Q-SiteFinder

Binding pocket analysis of protein–protein and protein–ligand complexes

We have used Q-SiteFinder and high-quality, non-redundant datasets of protein structures to allow us to probe the nature of binding pockets involved in different types of protein interactions (see Fig. 1 and its legend). We first investigated the properties of

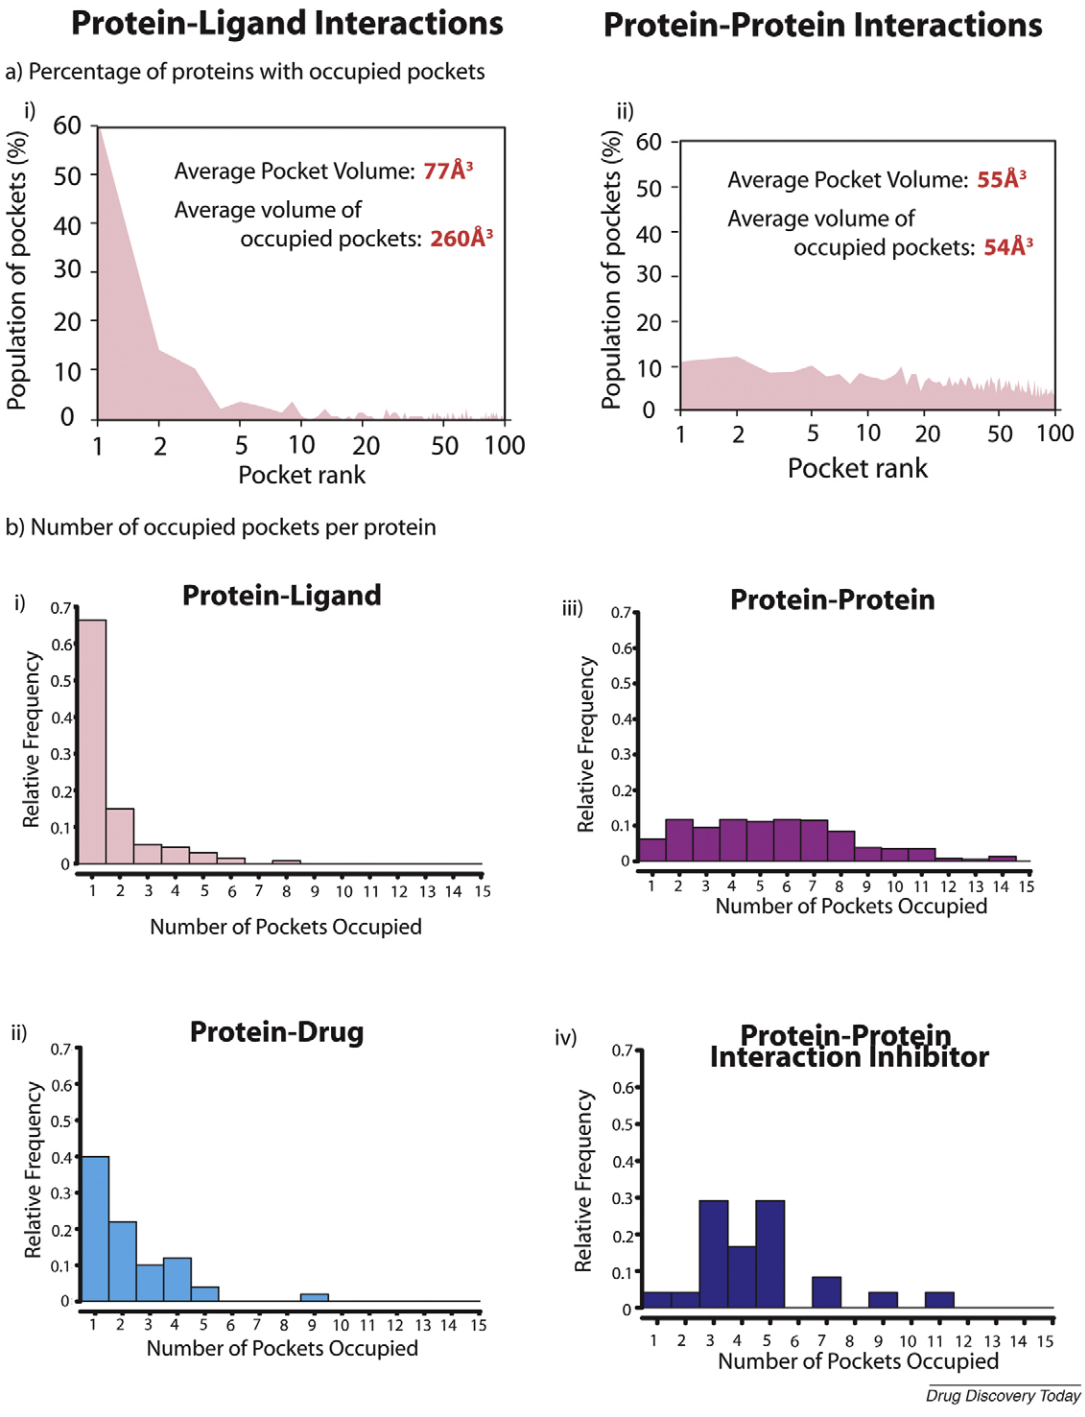

FIGURE 1

Binding site characteristics of protein–ligand and protein–protein interactions. **(a)** The % of complexes with occupied pockets for a given rank predicted by Q-SiteFinder [21] for (i) protein–ligand binding sites and (ii) protein–protein binding sites. These are a set of (i) 134 bound protein–ligand complexes, a subset of a docking benchmark with proteins of high structural similarity removed [21] and (ii) 97 pairwise bound protein–protein complexes (all are non-obligate, hetero-protein complexes), a total of 194 monomers [35]. Occupied (or true) pocket predictions are defined as those occupied to more than 25% of their total volume by atoms of the interacting molecule for a total of 99 pockets calculated for each protein surface [21]. **(b)** Histograms showing the relative frequency of ligand-occupied pockets found on each protein complex. (i) Protein–ligand dataset [21]; (ii) 50 bound protein–marketed-drug complexes [7]; (iii) protein–protein dataset [35]; (iv) 24 protein–protein interaction inhibitor complexes [1] (see http://www.modelling.leeds.ac.uk/qsitefinder/data, for listing of all datasets).

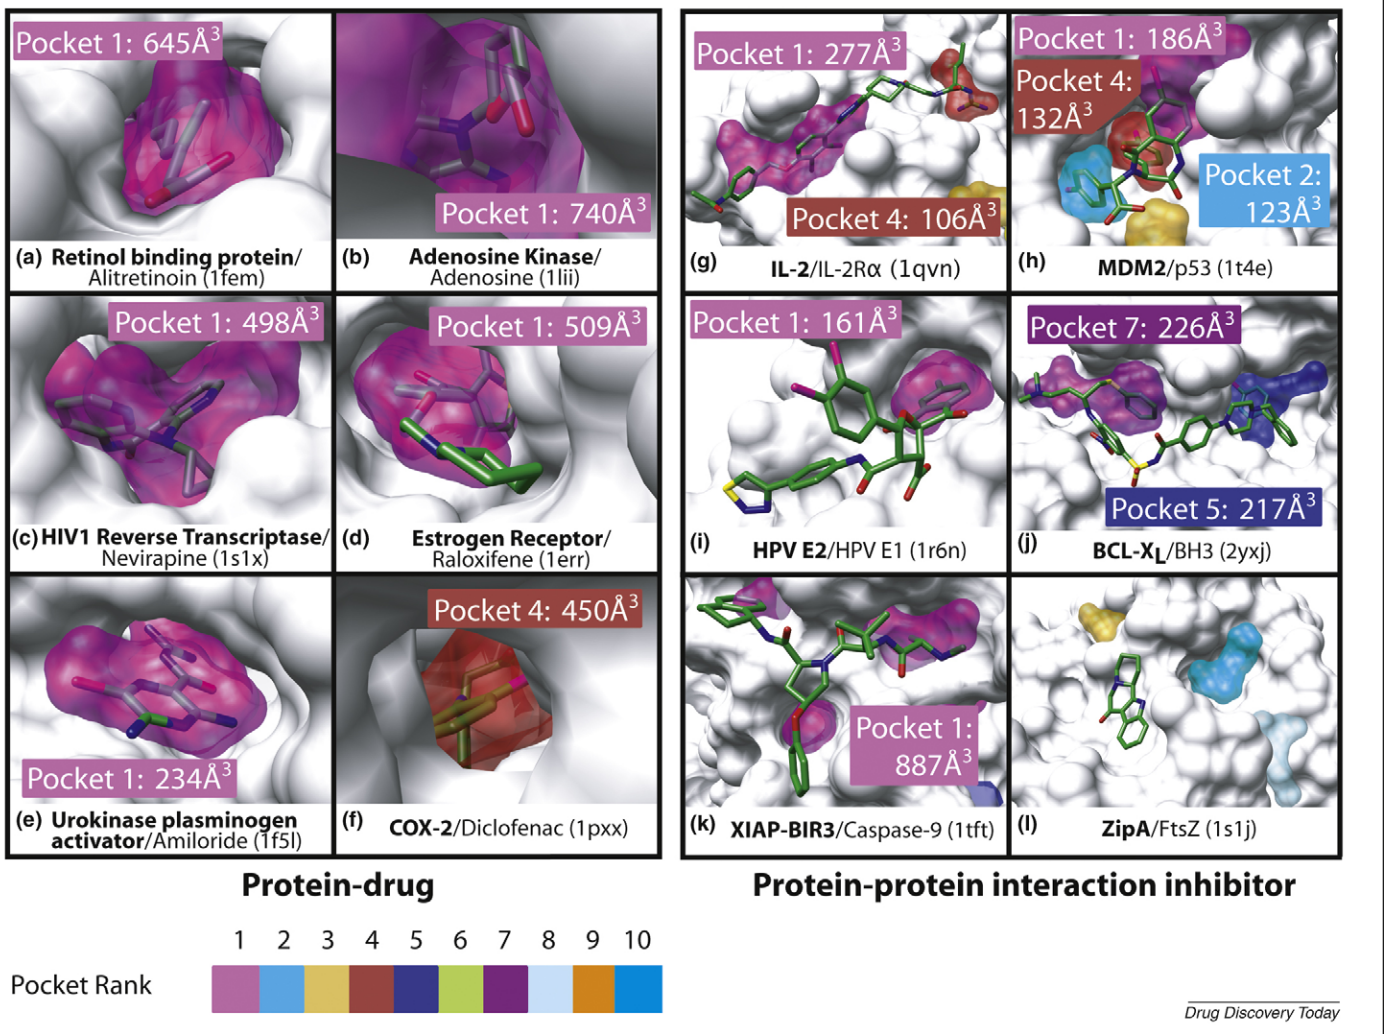

FIGURE 2

Differences in the nature of small-molecule binding sites of typical protein–ligand interactions (left) and protein–protein interaction inhibitors (right) generated using Q-SiteFinder [21], showing that PLIs tend to target a single high-volume pocket, whilst PPI inhibitors target multiple smaller pockets. Images of the protein–ligand interface with ligand coloured according to the atom type, protein solvent accessible surface area shown in grey and active volume of the pocket is coloured according to the rank of the predicted pocket (see colour key). **(a–f)** Show representatives of the protein–ligand marketed drugs dataset (PDIs) which contains 50 proteins that are non-redundant at the SCOP superfamily level. **(g–l)** show representatives from each of the PPI inhibitor protein targets. Images are generated using Chimera [36].

PLIs and PPIs. We observe striking differences between the protein–ligand and protein–protein binding site pockets (Fig. 1). The average active volume of the top ranked pockets for protein–ligand receptors (524 Å^3) is twice that of the protein–protein monomers (261 Å^3). Fifty-eight percent of the occupied protein–ligand binding pockets are the top predicted pocket, which has an average active volume of 471 Å^3, as opposed to only 15% in PPI monomers (active volume: 234 Å^3). Q-SiteFinder ranks pockets according to the total energy of the probes contained within the volume envelope, with the largest energies corresponding to the top ranked pocket, this is in contrast to most geometric measures where pockets are ranked according to their total volume. There is no strong bias in PPIs to occupy a top 10 predicted pocket, with many of the lowly ranked ones (ranked 11–80) having average occupancy levels of between 5 and 10%. PPIs are characterised typically by six (±3) small (active volume: 54 Å^3) occupied pockets (Fig. 1b(iii)), with volumes that are the same as the average for all surface pockets (55 Å^3). This contrasts with PLIs where the top 1 or 2 ranked pockets make by far the largest contribution to binding (Fig. 1a(i)) and the average occupied pocket volume (260 Å^3) is over four times the average occupied volume for PPIs. PLIs tend to occur in one or two disproportionately large pockets relative to the average pocket size (Fig. 2a–f). By contrast, PPIs tend to occur in several average-sized pockets that have a similar active volume to that of the average surface pocket. In conclusion, PLIs have one or two disproportionately large pockets that are generally occupied by their ligands; otherwise they show similar pocket size characteristics to the surfaces of proteins involved in PPIs.

Given the striking differences observed between PPIs and PLIs we next investigate the binding site pockets of PPI inhibitors and PDIs; in the case of the latter this is to see if there are any differences to PLIs in general for (datasets see Fig. 1 and its legend).

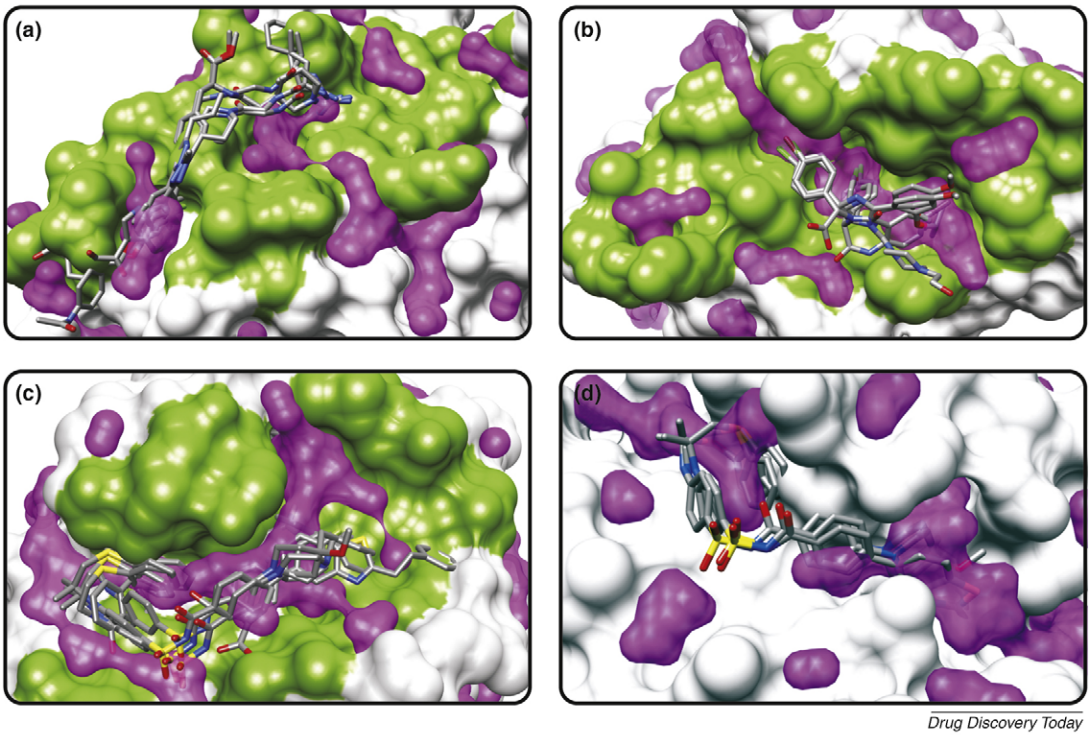

FIGURE 3

Binding pockets on unbound protein surfaces for four different PPI target classes [1], with PPI inhibitors from the holo PPI inhibitor complex structurally superposed. In all cases energetically favourable regions predicted by Q-SiteFinder [21] are shown in transparent magenta, and the ensemble of bound ligands is shown following structural superposition of bound structures onto the unbound. **(a)** Unbound surface of IL-2 shown in grey (IL-2Rα contact epitope shown in green); **(b)** PPI bound surface of MDM2 shown in grey (p53 contact epitope shown in green); **(c)** unbound surface of Bcl_{XL} shown in grey (Bak peptide contact epitope shown in green); **(d)** unbound Bcl-2 shown in grey. Images are generated using Chimera [36]. PDB codes for unbound structures in bold, and for ligand codes in italic: (a) **1m47**, *1m48*, *1m49*, *1py2*, *1rv1*, *1qvn*; (b) **1z1m**, *1t4e*, *1ttv*; (c) **1maz**, *1ysi*, *2o2m*, *2yxj*, *1ysg*, *1ysn*, *2o2n*; (d) **1g5m**, *1ysw*, *2o2f*, *2o22*.

We observe that the characteristics of marketed drugs are highly similar to those of the PLI dataset, in terms of both their preponderance to target one pocket (Fig. 1b(ii)) and their high average occupied pocket volume (271 $Å^3$). The PPI inhibitors have an occupied pocket most similar in volume (100 $Å^3$) to PPIs and much smaller than those of PLIs. Fig. 1b(iv) shows that PPI inhibitors typically target between 3 and 5 pockets, results that are confirmed by visual inspection (Fig. 2g–l). This is akin to the way in which PPIs also target several pockets rather than one large pocket as seen in PLIs.

Conformational change on binding

The above analysis has been carried out on ligand-bound conformations that take into account the effects of induced fit caused by ligand binding. It is, however, relevant to ask to what extent these pockets are present in the unbound protein conformations, and to what extent does induced fit play a role. Eyrisch and Helms have noted that virtual screening methods that allow for picosecond–nanosecond receptor flexibility show promise, even in cases where rigid body docking was unsuccessful [12]. They studied three PPI inhibitors (IL-2/IL-2Rα, MDM2/p53 and Bcl_{XL}/BH3) using molecular dynamics, and a geometric pocket finding algorithm (PASS) [33] was used to determine potential binding pockets on the protein surface. The authors note that in the case of the apo IL-2, PASS did not identify the native binding pocket, which was detectable in the inhibitor-bound structures of IL-2 with the inhibitors removed [12]. Similarly, it has been noted that Bcl_{XL} has a rather flat surface in the apo state [1]; however, transient pockets are observed in molecular dynamics simulations in the absence of small molecules in both Bcl_{XL} [12,34] and IL-2 [12], showing that pockets on the protein surfaces frequently open and close on timescales ranging from picoseconds to hundreds of picoseconds [12]. Indeed, there are limitations in describing potential small-molecule binding pockets when using static structures of either the free protein or the protein–protein complex. We have analysed the free protein structures and the PPI inhibitor-bound complexes (cf. Fig. 3). Visual analysis shows that pocket detection is most successful when applied to the bound structure, as might be expected, but that key interaction pockets are often present (or preformed) in the free protein or the protein–protein complex. Indeed, in the case of IL-2/IL-2Rα we notice several striking features when comparing the free (Fig. 3a) and bound structures. Firstly we note that Q-SiteFinder is successful in identifying the two binding pockets in the free structure that are targeted by IL-2Rα ligands. Secondly, and perhaps more importantly we notice that all the ligands target a pocket that consists (at least in part) of residues that are not directly involved in the PPI interaction interface (the partner protein contact epitope is shown in green). This pocket is not utilised in the protein-protein interaction involving IL-2/IL-2Rα and thus would not be found by

experimental techniques such as alanine scanning mutagenesis, but it is clearly a common binding 'hot-spot' for the small-molecule IL-2 inhibitors. Furthermore, recognisable ligand binding pockets are present in the static structures of the protein–protein complexed state of MDM2 and the apo state of Bcl_{XL} (Fig. 3b and c, respectively). Whilst it is arguable that these unliganded pockets do not represent the best starting point for structure-based ligand design, this analysis clearly indicates the utility of using pocket finding techniques to identify targetable pockets at the protein–protein interface.

Next we analysed the volumes of occupied pockets in unbound protein structures, where we define an occupied pocket as one that coincides with an occupied pocket in the bound conformation of the structure. First, we note that induced fit plays a significant role in all types of protein interaction, as observed from the increase in the average pocket volume in going from unbound to bound conformations (data not shown). With specific reference to the top (or principal) occupied pockets, induced fit plays a highly significant role. In terms of volume change this is most significant in the protein-marketed drug dataset (ΔVolume: 314 $Å^3$) but is also important in the PLIs (ΔVolume: 223 $Å^3$) and PPI inhibitors (ΔVolume: 137 $Å^3$) and is least significant in PPIs (ΔVolume: 78 $Å^3$). This confirms the view that small molecules can access pockets that larger more constrained proteins cannot and that conformational change induced by ligand binding is likely to be crucial for high-affinity binding [1]. Although a relatively small sample, we can compare bound and unbound (or PPI complexed) states for the PPI inhibitors, and find that 3 of the 7 PPI inhibitor families have a top 10 pocket in the unbound state, which corresponds to the top occupied pocket in the bound state. Thus, the principal binding site is often a top ranking pocket in the unbound (or PPI complexed) state. This indicates that the use of pocket detection can be a useful tool in drug discovery at the protein–protein interface where structural information is available.

Targeting pockets in drug design

Our results indicate that, in the future, targeting PPIs with small-molecule drugs is most likely to be successful when molecules target several pockets, because these interfaces do not have a single large pocket and protein targets that have been most successfully inhibited so far have been shown to occupy multiple pockets.

The combined active volumes targeted by current PPI inhibitors (see Fig. 2g–l) for IL-2/IL-2Rα, MDM2/p53 and Bcl_{XL}/BH3 are 383 $Å^3$, 441 $Å^3$ and 443 $Å^3$, respectively. These values approach or exceed those of some marketed drugs. Current attempts at elucidating PPI inhibitors have often relied on high-throughput screening methods to identify initial lead compounds. Subsequent analysis of these has shown that the small molecule in question often mimics the binding epitope of its natural partner protein. Our observations may allow the identification of regions on the protein surface that are most likely to accomodate small-molecule inhibitors. In particular the determination of the 'active volume' of ligand binding sites will help to target molecules to specific pockets in structure-based design and increase ligand efficiencies by allowing the visualisation and therefore minimisation of ligand structure outside these regions. Marketed small-molecule drugs would already appear to occupy the binding site active volume with high efficiency; the challenge will be to design PPI inhibitors that do the same.

Concluding remarks

An important conclusion is that the targetable binding pockets of protein–protein interfaces are significantly different from those of protein–ligand interfaces. PLIs use fewer, larger pockets, whereas PPIs and PPI inhibitor interactions use a greater number of smaller pockets instead. Interestingly, the pockets used by PPIs and PPI inhibitors are similar in size to those found anywhere on the surfaces of proteins, and the analysis shows that engaging multiple pockets is likely to be a productive strategy in inhibiting PPIs. Whilst these results might be seen as intuitive, the present analysis and pocket descriptions that result from Q-SiteFinder underscore the utility of this strategy and help scientists visualise that PPI targets are significantly different from PLIs. Another conclusion is that targetable pockets are often identifiable at the PPI interface in the unbound or PPI-complexed states and therefore pocket finding in these non-ligand-bound states could be a useful tool for identifying druggable sites for exploitation in structure-based design of inhibitors.

In terms of predicting druggable binding sites at the protein–protein interface there are several strengths and weaknesses of using each of the different methods discussed in this review. One problem for geometric methods is in defining small cavities on relatively flat PPI interfaces where several small pockets define the energetically important binding regions. Indeed, sometimes no pocket is identified, for example PASS did not identify the native binding pocket in the apo IL-2 [12]. These types of pocket are well defined by energetic methods such as Q-SiteFinder. The work by Eyrisch and Helms [12] and Brown and Hajduk [34] suggests that there is often an advantage in looking at pockets using a dynamic model of protein structure, particularly where conformational selection is taking place, because many transient pocket openings happen within the picosecond–nanosecond timescale. This may be particularly important when designing new molecules to fit in a pocket. In this work Q-SiteFinder implicitly assumes that the total energy of the pocket (sum of the methyl probe energies) defines its ability to bind a small molecule and therefore its druggability. Our previous analysis of binding pockets demonstrated that desolvation energy is a key predictive measure for binding pockets in protein–protein (and protein–ligand) interfaces, whereby easily desolvated pockets are more likely to be in the interface [35]. Indeed, the model presented here can be adapted to take account for the hydrophobic effect rather like some other models where non-polar pocket surface area has been used to provide a direct calculation of druggability [17,34].

Q-SiteFinder is a freely available web tool, accessible at http://www.modelling.leeds.ac.uk/qsitefinder.

References

1 Wells, J.A. and McClendon, C.L. (2007) Reaching for high-hanging fruit in drug discovery at protein–protein interfaces. *Nature* 450, 1001–1009

2 Fry, D.C. (2006) Protein–protein interactions as targets for small molecule drug discovery. *Biopolymers* 84, 535–552

3 Chene, P. (2006) Drugs targeting protein–protein interactions. *Chem. Med. Chem.* 1, 400–411
4 Arkin, M.R. and Wells, J. (2004) Small-molecule inhibitors of protein–protein interactions: progressing towards the dream. *Nat. Rev. Drug Discov.* 3, 301–317
5 Whitty, A. and Kumaravel, G. (2006) Between a rock and a hard place? *Nat. Chem. Biol.* 2, 112–118
6 Berg, T. (2003) Modulation of protein–protein interactions with small organic molecules. *Angew. Chem. Int. Ed. Engl.* 42, 2462–2481
7 Wishart, D.S. *et al.* (2008) DrugBank: a knowledgebase for drugs, drug actions and drug targets. *Nucleic Acids Res.* 36 (Suppl. 1), D901–906
8 Kuritzkes, D. *et al.* (2008) Maraviroc. *Nat. Rev. Drug Discov.* 7, 15–16
9 Dömling, A. (2008) Small molecular weight protein–protein interaction antagonists: an insurmountable challenge? *Curr. Opin. Chem. Biol.* 12, 281–291
10 Neduva, V. and Russell, R.B. (2006) Peptides mediating interaction networks: new leads at last. *Curr. Opin. Biotechnol.* 17, 465–471
11 Zhong, S. *et al.* (2007) Computational identification of inhibitors of protein–protein interactions. *Curr. Top. Med. Chem.* 7, 63–82
12 Eyrisch, S. and Helms, V. (2007) Transient pockets on protein surfaces involved in protein–protein interaction. *J. Med. Chem.* 50, 3457–3464
13 Overington, J.P. *et al.* (2006) How many drug targets are there? *Nat. Rev. Drug Discov.* 5, 993–996
14 Landry, Y. and Gies, J.P. (2008) Drugs and their molecular targets: an updated overview. *Fundam. Clin. Pharmacol.* 22, 1–18
15 Hajduk, P.J. *et al.* (2005) Druggability indices for protein targets derived from NMR-based screening data. *J. Med. Chem.* 48, 2518–2525
16 Gonzalez-Ruiz, D. and Gohlke, H. (2006) Targeting protein–protein interactions with small molecules: challenges and perspectives for computational binding epitope detection and ligand finding. *Curr. Med. Chem.* 13, 2607–2625
17 Cheng, A.C. *et al.* (2007) Structure-based maximal affinity model predicts small-molecule druggability. *Nat. Biotechnol.* 25, 71–75
18 Lipinski, C.A. *et al.* (2001) Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv. Drug Deliv. Rev.* 46, 3–26
19 Hopkins, A.L. *et al.* (2004) Ligand efficiency: a useful metric for lead selection. *Drug Discov. Today* 9, 430–431
20 Fry, D.C. (2008) Drug-like inhibitors of protein–protein interactions: a structural examination of effective protein mimicry. *Curr. Protein Pept. Sci.* 9, 240–247
21 Laurie, A.T.R. and Jackson, R.M. (2005) Q-SiteFinder: an energy-based method for the prediction of protein–ligand binding sites. *Bioinformatics* 21, 1908–1916
22 Laurie, A.T. and Jackson, R.M. (2006) Methods for the prediction of protein–ligand binding sites for structure-based drug design and virtual ligand screening. *Curr. Protein Pept. Sci.* 7, 395–406
23 Huang, B. and Schroeder, M. (2006) LIGSITEcsc: predicting ligand binding sites using the Connolly surface and degree of conservation. *BMC Struct. Biol.* 6, 19
24 Hendlich, M. *et al.* (1997) LIGSITE: automatic and efficient detection of potential small molecule-binding sites in proteins. *J. Mol. Graph. Model.* 15, 359–363
25 Weisel, M. *et al.* (2007) PocketPicker: analysis of ligand binding-sites with shape descriptors. *Chem. Cent. J.* 1, 7
26 Glaser, F. *et al.* (2006) A method for localizing ligand binding pockets in protein structures. *Proteins Struct. Funct. Bioinform.* 62, 479–488
27 Laskowski, R.A. (1995) SURFNET: a program for visualizing molecular surfaces, cavities, and intermolecular interactions. *J. Mol. Graph.* 13, 323–330
28 Glaser, F. *et al.* (2005) The ConSurf-HSSP database: the mapping of evolutionary conservation among homologs onto PDB structures. *Proteins* 58, 610–617
29 Landon, M.R. *et al.* (2007) Identification of hot spots within druggable binding regions by computational solvent mapping of proteins. *J. Med. Chem.* 50, 1231–1240
30 Mattos, C. *et al.* (2006) Multiple solvent crystal structures: probing binding sites, plasticity and hydration. *J. Mol. Biol.* 357, 1471–1482
31 Harris, R. *et al.* (2007) Automated prediction of ligand-binding sites in proteins. *Proteins* 70, 1506–1517
32 Bogan, A.A. and Thorn, K.S. (1998) Anatomy of hot spots in protein interfaces. *J. Mol. Biol.* 280, 1–9
33 Brady, G.P. and Stouten, P.F.W. (2000) Fast prediction and visualization of protein binding pockets with PASS. *J. Comput. Aided Mol. Des.* 14, 383–401
34 Brown, S.P. and Hajduk, P.J. (2006) Effects of conformational dynamics on predicted protein druggability. *Chem. Med. Chem.* 1, 70–72
35 Burgoyne, N.J. and Jackson, R.M. (2006) Predicting protein interaction sites: binding hot-spots in protein–protein and protein–ligand interfaces. *Bioinformatics* 22, 1335–1342
36 Pettersen, E.F. *et al.* (2004) UCSF chimera – a visualization system for exploratory research and analysis. *J. Comput. Chem.* 25, 1605–1612