

Virtual Screening

Virtual screening, sometimes called *in silico screening*, is a relatively new approach to library testing. In a virtual screen, computerized molecular models of both the target and library member are aligned to determine potential complementary intermolecular interactions. Molecules with a high level of complementarity, indicative of potentially strong binding, are flagged for synthesis and testing. A virtual screen requires sufficient knowledge about the target protein structure, likely from x-ray or NMR data.¹¹ Virtual screening does *not* require an existing compound library. Any molecule imaginable can be modeled in a computer and screened. Of course, the virtual library should consist of realistically synthesizable compounds.¹² Selected compounds are also normally filtered for those with

desired structural elements.

Virtual screening faces a number of challenges. Molecular conformations of both the target and library member are an issue. The target protein will have many low-energy conformations. Some virtual screening methods attempt to accommodate flexibility of the target protein, as complicated as it may be. All *in silico* approaches try to account for flexibility

of the library member, which is a challenge because even small molecules can have many low-energy conformations. Other factors include tautomers, pH-dependent ionizations, and stereochemistry. Properly handling all these variables is not trivial and quickly complicates the modeling process. Once starting conformations of the target and library members have been established, each library member is virtually brought into contact with the target to determine the likelihood of binding. The process, called *docking*, follows the induced-fit model in which the interacting molecules influence each other's conformations

until a minimum energy is reached.¹¹

After a compound has been docked to the target, the binding energy is estimated in a process called *scoring*. Standard intermolecular forces, contact forces, dipole interactions, and hydrogen bonding are approximated and totaled. Current scoring methods produce many *false positives*. False positives are compounds with high predicted binding that show little or no activity in validation testing. Often, the number of false positives can be reduced by using several different scoring systems to calculate binding. This approach is called *consensus scoring*. Compounds with high predicted activities from more than one method are selected

for further investigation.¹¹

High-scoring compounds are then prepared and screened in a biochemical assay. "Preparing" requires either synthesizing or purchasing a sample of the compound. Sample purchase is normally much faster than synthesis. For this reason, compounds in a virtual screen are often limited to structures that can be purchased from any number of commercial suppliers. Databases of available molecules have been prepared for the sole purpose of assisting virtual screening.

The following three Case Studies highlight the use of virtual screening to find hits.

**CASE
STUDY**

Inhibition of Protein Kinase B¹³

Protein kinase B uses ATP to phosphorylate other enzymes. Controlling protein kinase B is a method of influencing many cellular processes, including cell division in cancer cells. A virtual screen for protein kinase B inhibitors started with an x-ray structure of protein kinase B bound to AMP-PNP (10.32), an unreactive analog of ATP, and a peptide substrate (Figure 10.8). A collection of 50,000 compounds from Chembridge (San Diego, California) was docked into the active site of protein kinase B and scored by an algorithm called Drugscore (BioSolvIT, Sankt Augustin, Germany). The top 4,000 compounds were then scored by two other methods, called Goldscore and Chemscore. The lists of the top 700 compounds from both methods contained 200 common compounds. Visual inspection of these




FIGURE 10.8 AMP-PNP (10.32) and Inhibitors of protein kinase B (10.33–10.36)

200 docked structures allowed 100 to be eliminated because of unfavorable interactions. The remaining 100 compounds were screened against protein kinase B. Three compounds

(10.33–10.35) emerged with hit-level activity. Experimental activity was comparable to H-89 (10.36), a commonly cited protein kinase B inhibitor.

CASE
STUDYInhibition of Sir2 Type 2¹⁴

Sir2 type 2 is a regulator enzyme that deacetylates tubulin and may be involved in aging and cancer. A virtual screen for an inhibitor started with a molecular model built from a high-resolution x-ray structure of the regulatory enzyme. In a molecular dynamics simulation, the model was allowed to undergo conformational changes. Various low-energy conformers were weighted, averaged, and minimized to afford a new molecular model. A handful of residues in the 389-peptide protein fell outside the allowed regions of a Ramachandran plot. None of the outlier residues was within the active site

of the enzyme, so the new minimized model was deemed to be adequate for virtual screening.

Two known Sir2 type 2 inhibitors, sirtinol (10.37) and A3 (10.38), were docked into the modeled active site (Figure 10.9). Studying both complexes revealed the positions of likely important hydrophobic and hydrogen bonding interactions within the active site. A search of a library of approximately 50,000 compounds available from Maybridge (Trevillett, England) was searched for screening candidates. In total, 44 compounds satisfied the selection criteria, largely based on the structure of the modeled




FIGURE 10.9 Inhibitors of Sir2 type 2 regulatory enzyme

enzyme. Docking of all 44 compounds revealed 15 that were able to interact fully with the active site. Screening of the 15 compounds against the regulatory enzyme indicated promising activity for 10.39 and 10.40. Interestingly,

all attempts to score the binding energy through software analysis of the docked inhibitors failed to distinguish active and inactive structures. Therefore, all docked structures were inspected visually for reasonable interactions.

CASE STUDY

C-C Chemokine Receptor Type 5 Agonists¹⁵

C-C chemokine receptor type 5 is a receptor that has been implicated in the early stages of HIV infection. Antagonists of this receptor may theoretically block viral infection of a cell. New antagonists were pursued through a virtual screen. A computer model of the receptor, a G-protein-coupled receptor, was constructed based on constraints from an x-ray structure of bovine rhodopsin, another G-protein-coupled receptor. A handful of C-C chemokine receptor type 5 antagonists have been reported in the literature. The binding pocket of the known antagonists had been elucidated through site-directed mutagenesis studies and shown to consist of two distinct hydrophobic regions spanning an anionic glutamic acid residue. Docking

each antagonist into the binding site, minimizing the complex, and averaging the structures provided the final model of the receptor.

Initial tests were performed on a pool of 1,000 compounds, including seven known antagonists. Two scoring methods were able to place half the known antagonists into the top 5% of all docked compounds. This exercise validated the model of the C-C chemokine receptor type 5 binding pocket as well as the scoring method for complex evaluation.

The virtual screen was performed on nearly 50,000 compounds that were taken from a library of 1.6 million. Selection for docking required a structure to have a nonpolar aromatic ring and basic amine for matching the binding




FIGURE 10.10 C-C chemokine receptor type 5 binding compounds

pocket. A repetitive docking and scoring process indicated 81 compounds as most promising. Visual inspection reduced the final candidate pool to 77. Through biochemical screening, seven compounds were found to effectively bind the receptor. Only one compound, **10.41**, showed antagonist activity (Figure 10.10). The other six, including **10.42**, were agonists or partial agonists.

The unexpected result in the chemokine receptor example highlights the difference between binding an enzyme and a receptor. A nonsubstrate that strongly binds an enzyme active site will almost certainly be an effective competitive inhibitor. In contrast, a compound that blocks the binding site of an agonist can itself either elicit a response or shut down the receptor.

10.2 Filtering Hits to Leads

If a screening process is successful, a number of hits will be identified. The number of hits varies depending on the target. It could be as high as 5–10% or as small as 0.1% of the tested library. The cutoff for the required activity of a hit is somewhat arbitrary. The discovery group may select an activity level based on other known active compounds. The threshold may also be based on the performance of the entire library. For example, the discovery group may count all compounds that are two or three standard deviations more active than the average of the full library. For a library with normally distributed activity, a cutoff of two standard deviations would give a hit rate of 2.1%. A cutoff of three standard deviations would give a hit rate of 0.1%. Based on a representative hit rate of 1%, screening a library of 100,000 compounds would generate 1,000 hits. This is too many compounds to follow up each hit individually, so the number of hits needs to be reduced, or *filtered*, to reach a more manageable figure.

Pharmacodynamics and Pharmacokinetics

The most obvious filter would be to select the most potent hits. The threshold for activity of a hit may be 1 to 10 mM. Setting the limit at 100 nM would quickly reduce the number of hits. This approach has its problems.

Activity in a biochemical assay is strictly a measure of how a molecule interacts with a target, that is, pharmacodynamics. Since advancements in biochemical assays have made them the norm, pharmaceutical companies have continually watched compounds with excellent

pharmacodynamics fail in clinical trials because of poor pharmacokinetics. Drug companies have now learned to emphasize both pharmacodynamics and pharmacokinetics

throughout the lead discovery process. Instead of prioritizing hits based on binding (pharmacodynamics)

with a simple activity threshold, initial hits are also screened for a preliminary pharmacokinetic behavior. Pharmacokinetic properties of a hit in humans can be estimated with cellular assays as well as animal testing. It is important to be able to estimate properties in humans because U.S. Food and Drug Administration (FDA) approval for testing in humans will not have been obtained for hits from a library screen.

Another selection criterion for hits is the structural complexity of the hit. A hit that is advanced in a discovery program must be modified to increase its binding to a suitable level. If the hit has a complex structure that is difficult to prepare, synthesis of derivatives of the hit will require much time and slow the entire discovery process. Very complex hits are therefore often less attractive for promotion.

Biological Assays

Preliminary pharmacokinetic behavior can be tested through a number of whole cell assays. Most commercially successful drugs are administered orally, meaning the drug must be able to enter the bloodstream by crossing membranes in the intestines. The most common membrane permeability assay is performed by monitoring the absorption and secretion of a compound by colon carcinoma cells (Caco-2). Diffusion across Caco-2 cell membranes is considered to be a valid model for molecular transport in the small intestines.¹⁶

Drugs are mostly metabolized by liver enzymes, especially the cytochrome P-450 enzyme family. The ability for cytochrome P-450 enzymes to metabolize a hit is tested with liver microsomes. Liver microsomes consist primarily of endoplasmic reticulum that contains metabolic enzymes. Hits are individually incubated in the presence of the liver microsomes. Monitoring changes in concentrations provides a sense of the rate of metabolism of each hit. Liver microsomes are also used to determine whether the hit inhibits metabolic processes. Hits that inhibit liver metabolism are shunned.¹⁷

Acceptable hits do not need to show ideal behavior, but problem compounds will be removed from consideration. If all the hits fail initial pharmacokinetic screening, several options

are possible. First, the search for hits could start over with screening of a new library.

Second, the threshold for selection of hits could be lowered to enlarge the pool of hits, some of which may pass the permeability and metabolic screens. Third, the criteria for passing the Caco-2 and microsome screens may be softened to allow some hits to pass.¹⁸

Lipinski's Rules and Related Indices

Permeability and liver microsome screens are not high throughput. To save time, researchers

have sought simple methods for eliminating compounds that will be poor lead

candidates. A common method involves calculated indices. The first and most widely recognized

index-based filter was reported by Lipinski in 1997.¹⁹ This filter is called


Lipinski's rules or the *Rule of 5* (**Table 10.1**).

Lipinski's rules are designed to predict oral availability of compounds that passively diffuse across membranes. Lipinski's rules are based on observations of a database of approximately 2,500 drugs or compounds studied in clinical trials. In general, the compounds could be described structurally with limits on their molecular weight, number of hydrogen-bond donors

TABLE 10.1 Lipinski's rules¹⁹

1. Molecular weight ≤ 500
2. Lipophilicity ($\log P$ or $\text{clog } P$) ≤ 5
3. Sum of hydrogen-bond donors ≤ 5
4. Sum of hydrogen-bond acceptors ≤ 10

FIGURE 10.11 Erythromycin, a drug that violates Lipinski's rules



and acceptors, and lipophilicity ($\log P$ or $\text{clog } P$). $\log P$ is an experimental measure of lipophilicity that will be fully described in Chapter 12. A higher $\log P$ value indicates lower water solubility. The form $\text{clog } P$ (pronounced "see log P") is a computer-estimated version of $\log P$.

A compound that violates any of Lipinski's rules may not be absorbed well when orally administered.

Drugs that cross membranes by active, facilitated, or other means of transport fall beyond Lipinski's rules. Exceptions include the macrolide antibiotics such as erythromycin

(**10.43**) (**Figure 10.11**). Methods by which molecules cross membranes are covered more thoroughly in Chapter 3.¹⁹

Over time, Lipinski's rules have been criticized as inappropriate for the evaluation of hits and leads. Hits and leads have weaker binding energies than final drugs. The process of optimization increases a lead's binding energy with multiple, successive structural modifications.

These modifications typically increase a molecule's functionality and subsequently raise the molecular weight of the lead. A hit or lead with a molecular weight of 480 may slip under the Lipinski molecular weight requirement, but after going through the optimization

process, the molecule may balloon to a molecular weight of 600 or higher. With this logic, Lipinski's rules are perhaps too permissive to be useful as a filter for hits and leads. In 1999, Teague distinguished between *lead-like* and *drug-like* hits, and combinatorial library collections (**Table 10.2**).²⁰ Lead-like hits are characterized as having lower molecular weights 163502, activity 170.1 mM², and clog *P* values 1632. The lower values give lead-like compounds room to grow into an optimized, high-affinity drug that still satisfies Lipinski's rules. Drug-like hits have higher molecular weights 173502 and clog *P* values 1732 but still modest affinity 1_0.1 mM². The definition of lead-like has since been used as a preliminary filter for selecting more promising hits from a screen. Sometimes, simple selection criteria such as Lipinski's rules or lead-like properties are applied to a library before the initial screen is even performed.

A recently reported tool for hit evaluation and prioritization is *ligand lipophilicity efficiency* (*LLE*).²¹ *LLE* is calculated as the difference between the negative logarithm of a hit's binding affinity, such as $-\log IC_{50}$ or $-\log KD$, and the logarithm of a hit's partition coefficient, such as $\log P$ or $clog P$ (Equation 10.2).

$$LLE = -\log IC_{50} - \log P \quad (10.2)$$

TABLE 10.2 Lead-like and drug-like compounds²⁰

Lead-Like	Drug-Like
1. Activity >0.1 μ M	1. Activity >0.1 μ M
2. MW < 350	2. MW > 350
3. Clog <i>P</i> < 3	3. Clog <i>P</i> > 3

Higher *LLE* values are considered to be better. Consider what this equation says: "*LLE* equals activity less lipophilicity." Without this equation, one might be tempted to say that two hits with the same activity are equally attractive to a drug discovery team. Based on Equation 10.2, the less lipophilic hit (lower $\log P$) has a higher *LLE* value and would be more attractive as a hit. Although this may not seem to be an earth-shattering conclusion, Equation 10.2 does show the trade-off between activity and lipophilicity when prioritizing hits.²² Equation 10.2 quantitatively relates some of the central ideas behind the guidelines of Lipinski and Teague. An underlying assumption in Equation 10.2 is that growth of a lead into a drug will increase the compound's lipophilicity. A good hit or lead should therefore start with a lower lipophilicity so that the $\log P$ of the final drug will not surpass Lipinski's magic value of 5.

Other attempts to refine or improve Lipinski's rule set have appeared in the literature. One notable factor for consideration is the number of rotatable bonds in a hit. Increased molecular flexibility can reduce the ability of a molecule to cross a membrane. The maximum number of rotatable bonds has been suggested as 10. The polar surface area, often abbreviated as PSA, of a molecule is another important factor. Polar surface area is tightly correlated to the number of hydrogen-bond donors and acceptors contained in a molecule. A maximum polar surface area of 140 Å² or the equivalent of 12 hydrogen-bond donors/acceptors has been suggested. This is in line with Lipinski's rules.²² Lipinski terms and related indices exclusively predict oral bioavailability. None addresses metabolism concerns. While the formation of unwanted metabolites is difficult to predict, several functional groups have become recognized as common sources

of problems. Examples include quinones and hydroquinones (**10.44**), aryl nitro groups (**10.45**), primary aryl amines (**10.46**), and Michael acceptors (**10.47**) (**Figure 10.12**). Quinones and Michael acceptors are strong electrophiles that tend to react quickly with and deplete glutathione stores in the liver. Aryl nitro compounds are reduced in the body to aryl amines, which are oxidized to electrophilic species with the same problems as quinones and Michael acceptors. Because the liver enzymes perform a large fraction of a body's metabolism of xenobiotics, the liver is the most commonly damaged organ when metabolites are toxic.²³

Final Concerns for Promotion of a Hit to a Lead

Only a small number of hits remain after various selection criteria have been applied to the initial hit pool. The surviving hits, sometimes called *compounds of interest* or similar, receive additional scrutiny. Each remaining hit undergoes a handful of structural modifications.

Modified hits are often called *analogues*. The analogues allow the discovery team to gain a preliminary understanding of the impact of structural changes on the activity of the hit against its target. If similar targets are known and available, filtered hits are tested for selectivity against the desired target and undesired related targets. These selectivity comparisons

can predict the likelihood of side effects. Some in vivo testing may be performed in animals, especially rats. The in vivo tests provide a more accurate and reliable picture of




FIGURE 10.12 Common problematic functional groups in drug candidates

a compound's pharmacokinetic profile. Finally, patent searches are performed to determine

the patentability of the hit and later leads that might arise. If a compound cannot be patented, then that compound will certainly not be advanced as a lead.

The outcome of all these selection steps is hopefully one or more leads. Final leads may differ in structure somewhat from their original respective hits. Early structural modifications hopefully generate analogues of higher activity. While hits are often selected at an activity level of 1 mM (KD , IC_{50} , etc.), structural changes may provide a lead with activity at concentrations of 0.1 mM (100 nM). Ultimately, potency will typically be improved

down to the 1–10 nM level during the lead optimization stage.

SUMMARY

Lead discovery, the search for compounds with activity against a target, begins with screening molecules. Regardless of the available information on the target and any known ligands, screening normally starts with a compound library that may contain a

million or more compounds. Screening may involve the entire library or just some portion of the collection. Active compounds are called hits. A second approach to discovering hits includes the screening of smaller fragment libraries. Active fragments are linked to increased potency. A third screening approach to finding hits involves using computer models to estimate binding energies between a target and ligand. The activity of predicted hits is then confirmed through a traditional assay. The hits, regardless of how they are discovered, are then checked for factors such as their ability to be transported within the body, ease of synthesis, and even patentability. Hits that survive these criteria are promoted to lead status and subjected to further structural optimization. Exceptions to the described lead discovery process are uncommon but do occur. Leads and even hits have been discovered through serendipity. Leads can also arise from observation of desirable side effects of other drugs.