V 8 – Analysis of protein-protein binding

- Modelling by homology
- Structural properties of PP interfaces
- Predicting PP properties / affinity of interactions

  from phyletic couplings
- Review V1 – V7

Tue, Nov 12, 2019

3.1 Model protein structures by homology

Figure shows "**twilight zone**" below the dotted line.

If two sequences A and B have a higher sequence identity than this line,
their 3D structures are highly likely to be similar to each other.

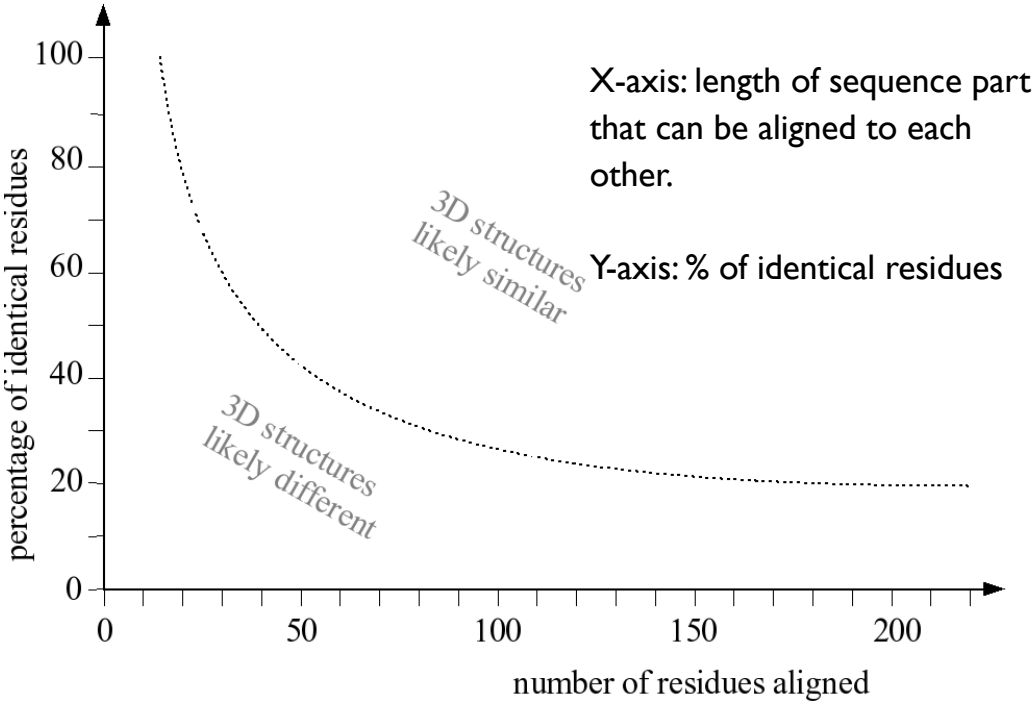

X-axis: length of sequence part that can be aligned to each other.

Y-axis: % of identical residues

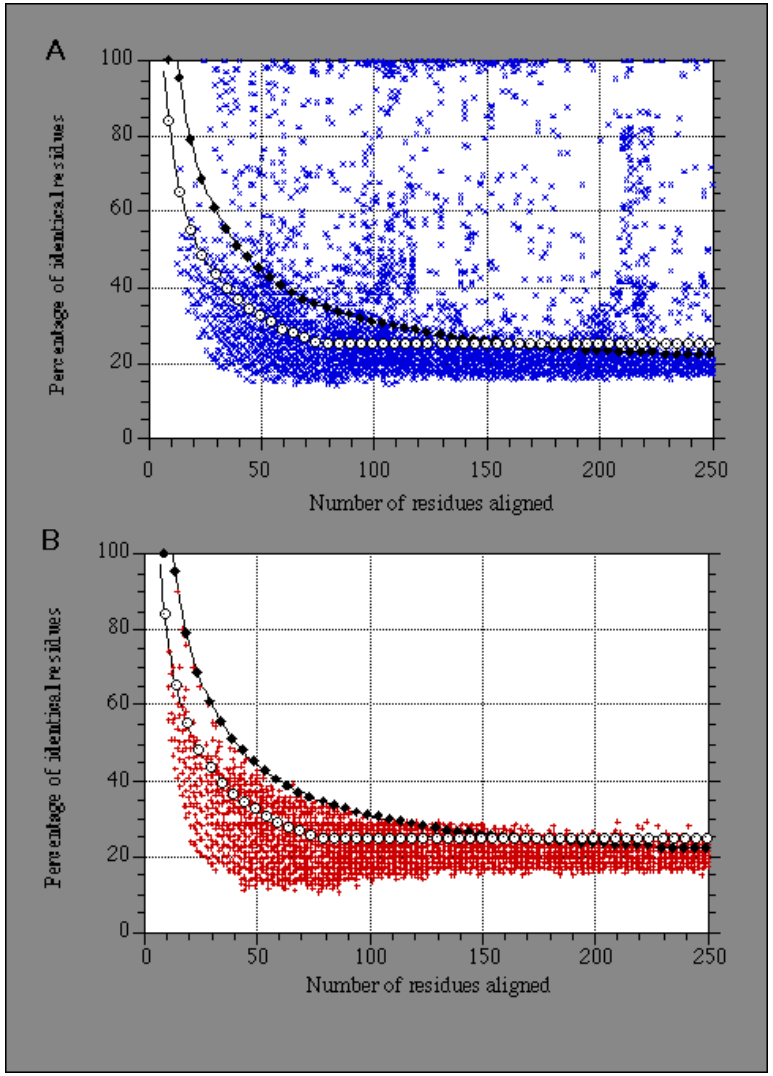

Top: sequence pairs A:B with similar structure
Bottom: pairs with different structure
Rost, Prot. Eng. 12, 85 (1999)

measure structural similarity of complexes

Critical Assessment of PRedicted Interactions (**CAPRI**) competition uses
3 criteria for ranking the protein complex predictions:

1- '**fnat**': the number of native residue–residue contacts in the predicted complex divided by the number of native contacts in the target.

2- **L-rms**: the backbone RMSD of the **ligands** (smaller one of both proteins) in the predicted versus the target structures.
Here, the larger proteins (**receptor**) are superimposed first.

3- **i-rms**: the RMSD of the backbone of the interface residues only, in the predicted versus the target complexes
(interface residues: here, residues with 10 Å of the other protein.
Map complementary residues in sequence alignment.)

Assessment of Blind Predictions of Protein–Protein Interactions: Current Status of Docking Methods, Mendez et. al. PROTEINS: Structure, Function, and Genetics 52:51–67 (2003)

3.1 Model protein complexes by homology

Structural similarity of protein complexes A':B' and A:B as a function of their sequence identity.

Note that x-axis and y-axis are different from previous slide.

A sequence identity level of 30-40% usually means that the binding mode of interaction is conserved (iRMSD < 3Å).

These plot show the "interaction RMSD", which is similar to L-RMSD.

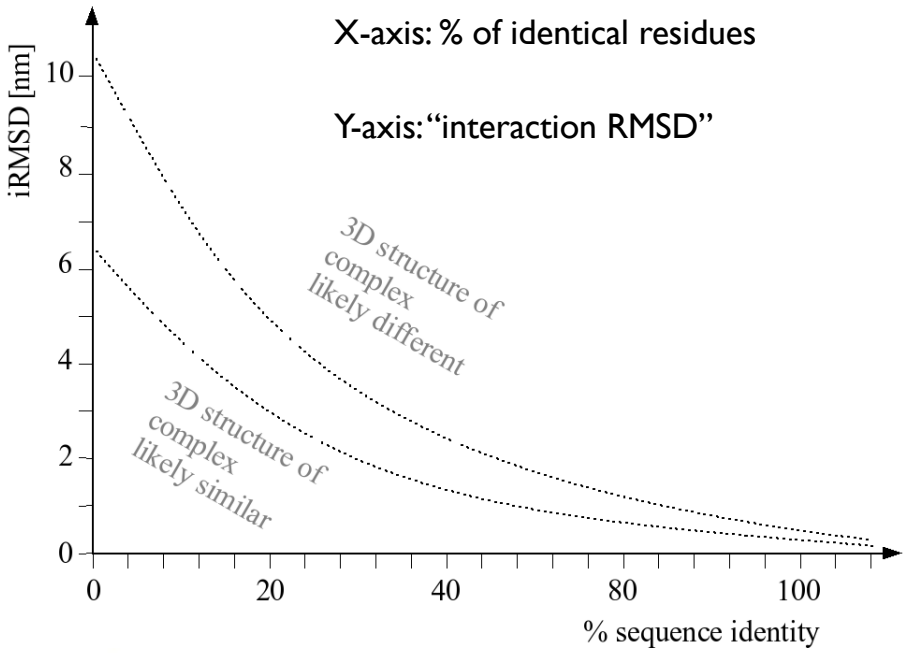

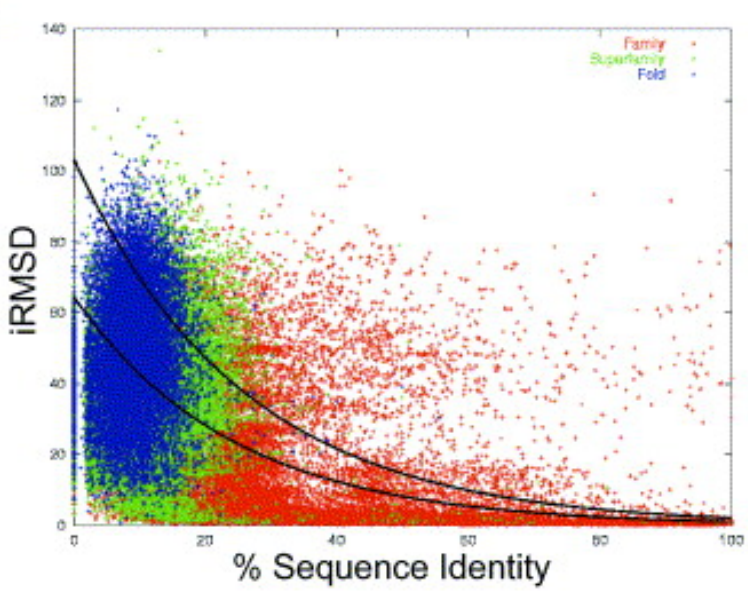

Aloy & Russell (2003) J Mol Biol 332, 989

3.1 Similar interaction without sequence similarity

Examples of similar interactions in absence of sequence similarity.

Proteins are shown in similar orientations.
Structurally equivalent regions are displayed in ribbons, dissimilar regions in trace and conserved residues in ball-and-stick representation.

Filled arrows between subunits show "interaction RMSD".
Broken arrows the percent sequence identity.

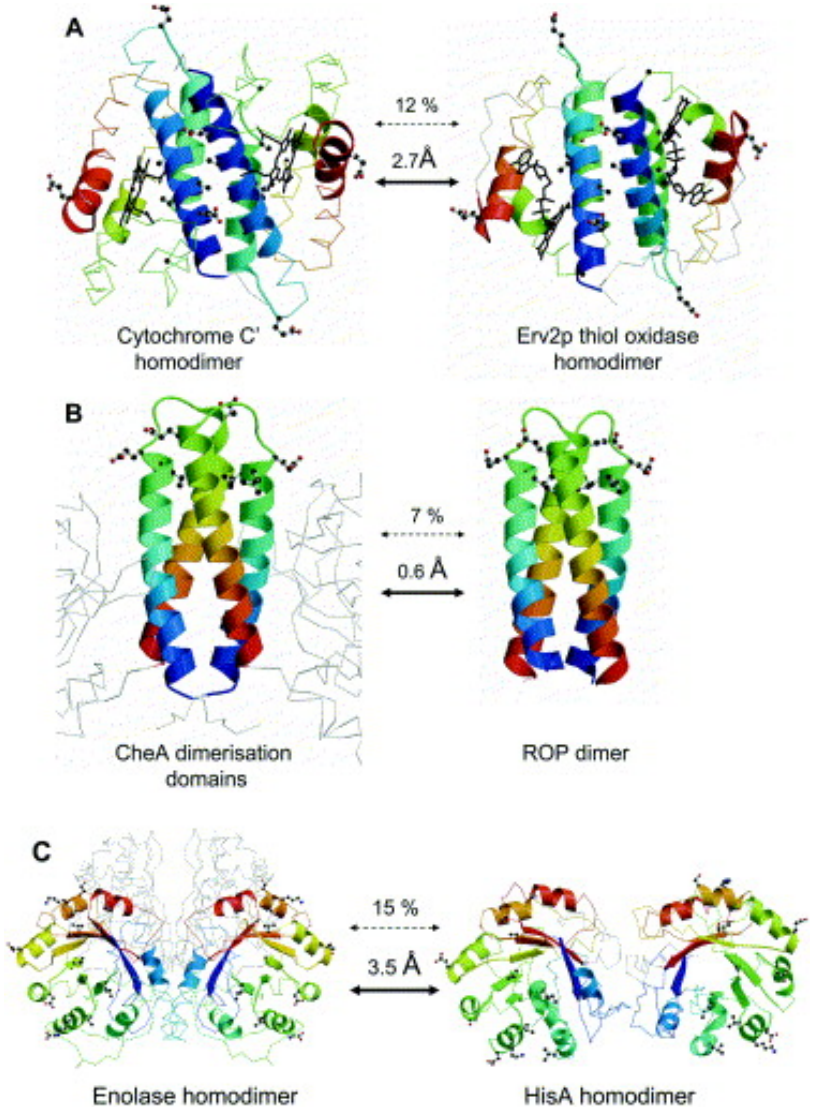

Aloy & Russell (2003) J Mol Biol 332, 989

3.1 Exceptions: close homologues interacting differently

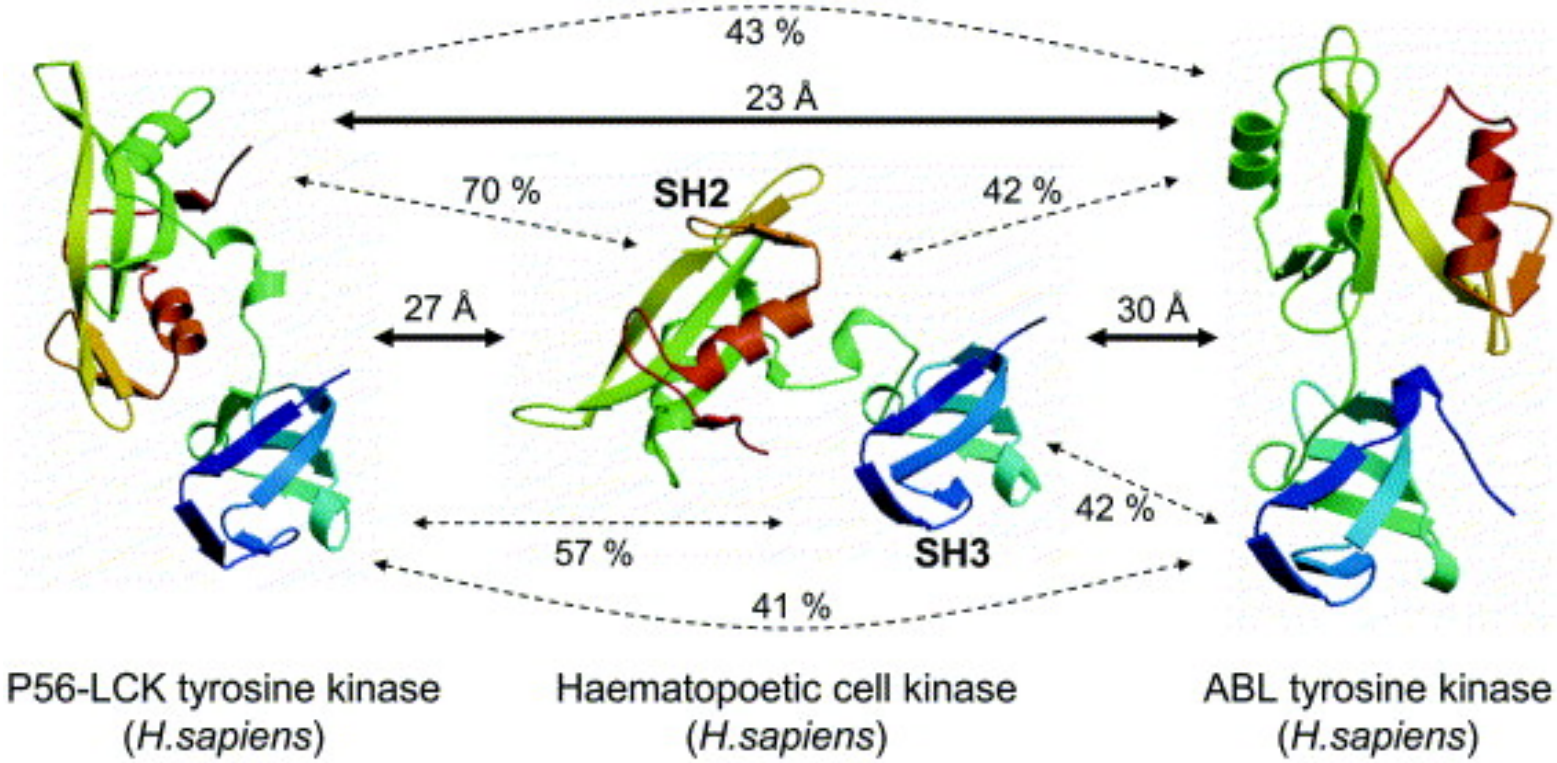

P56-LCK tyrosine kinase (1lck A), haematopoetic cell kinase (1ad5 A) and ABL tyrosine kinase (2abl) showing very different intramolecular interactions between homologous SH2 and SH3 domains.

Aloy & Russell (2003) J Mol Biol 332, 989

3.1 Interactions involving gene fusions

Top Histidine biosynthesis and a class I glutamine amidotransferase component domains of the imidazole glycerophosphate synthase from *Thermotoga maritima* (1gpw A and B) and *Saccharomyces cerevisiae* (1jvn A) interacting in a **similar way**.

Bottom FAD/NAD(P) binding and thioredoxin-like domains from thioredoxin reductase (1f6m A and B) and alkyl hydroperoxide reductase (1hyu A) i**nteracting differently**.

In both cases the linker is shown in yellow trace (**pink circle**).

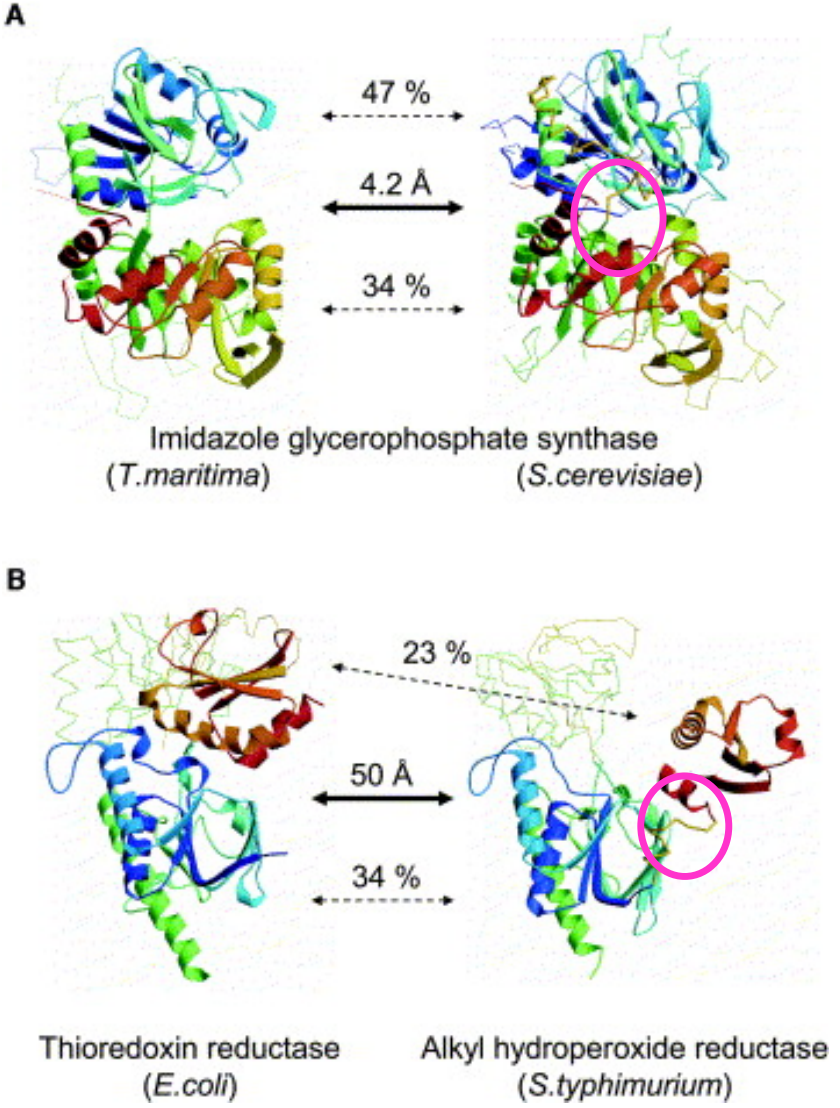

Aloy & Russell (2003) J Mol Biol 332, 989

3.2 Structural properties of PP interfaces

Size of protein-protein interface is commonly computed from **solvent-accessible surface area** (SASA) of the protein complex and of the individual proteins:

$$\Delta SASA = SASA_A + SASA_B - SASA_{AB}$$

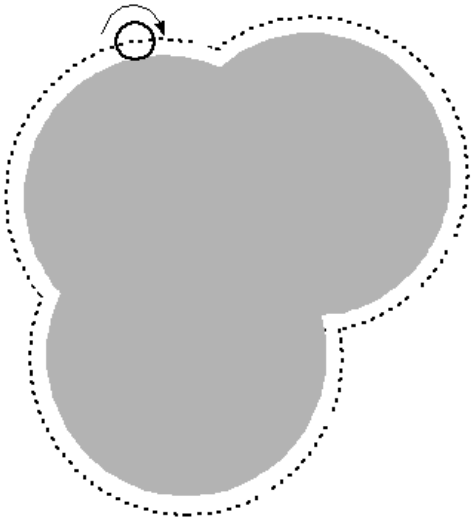

Computation of the SASA. A small probe is rolled over the complete surface of the large molecule shown in grey. The dashed line connects the positions of the center of the probe. In three dimensions, it is a surface. Its area is the SASA.

Definition of **interface residues**:

(a) All residues that are within a cut-off distance (e.g. 5Å) to any residue of the other protein.

(b) All residues having a reduced SASA in the complex compared to the unbound state.

3.2.1 Structural properties of PP interfaces

Parameter	Protein-protein complexes	Homodimers	Weak dimers	Crystal packing
Number in dataset	70	122	19	188
Buried surface area (Å)^2	1910	3900	1620	1510
Amino acids per interface	57	104	50	48
Composition (%) Non-polar Neutral polar Charged	 58 28 14	 65 23 12	 62 25 13	 58 25 17
H-bonds per interface	10	19	7	5
Residue conservation % in core	55	60	n/a	40

Janin et al. Quart Rev Biophys 41, 133 (2008).

3.2.1 size of PP interfaces

Redox complexes mediate e.g. the transfer of electrons between the binding partners.

Redox complexes possess relatively **small interfaces -> short life times**.

This makes biological sense. After an electron is transferred between 2 proteins, they no longer need to be bound.

In contrast, antibodies should bind their binding partners tightly so that they won't harm the organism.

The larger average interface size of antibody-antigen complexes is connected to a longer average life-time of the bound form.

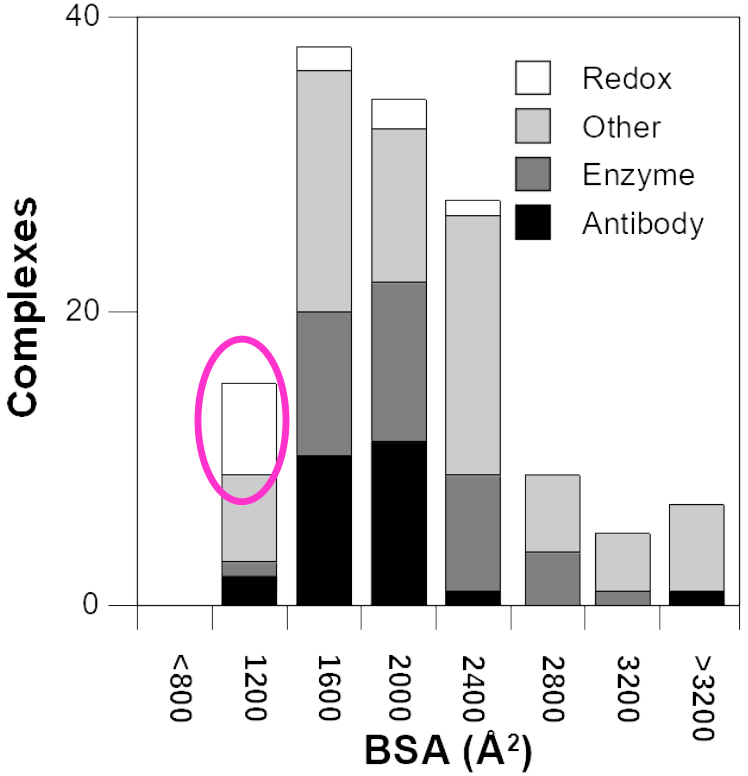

Interface size in transient protein–protein complexes. Histogram of the buried surface area (BSA) in 25 antigen–antibody complexes, 35 enzyme/ inhibitor or substrate complexes, 64 complexes of other types and in 11 redox protein complexes. The mean value of the BSA is 1290 Å^2 for the redox complexes and 1910 Å^2 for the other complexes.

Janin et al. Quart Rev Biophys 41, 133 (2008).

3.2.2 Composition of binding interfaces

Biological interfaces are enriched in aromatic (Tyr, Phe, Trp) and non-polar residues (Val, Leu, Ile, Met). Charged side chains are often excluded from biological protein-protein interfaces except for Arg.

In contrast, crystal contacts contain clearly fewer hydrophobic and aromatic residues, but more charged residues than biological interfaces.
Also, the enrichment of amino acids is smaller at crystal contacts compared to biologically relevant contacts.

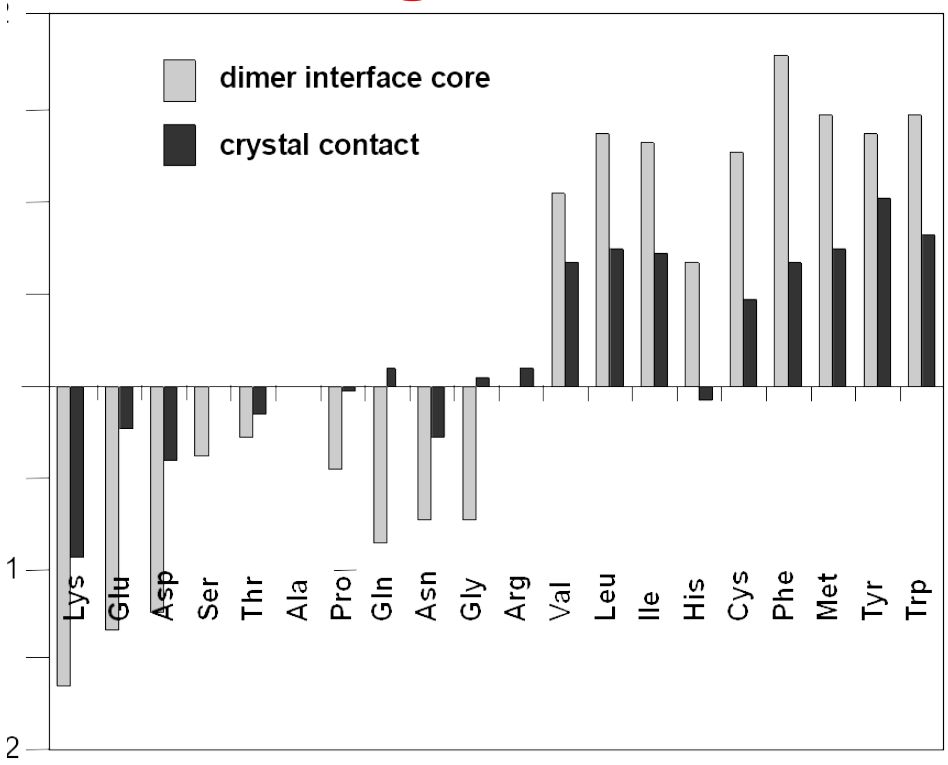

Residue propensities at protein dimer interfaces and at artificial contacts in the crystal, respectively. The propensities are derived from the relative contributions of the 20 amino acid types to the buried surface of the interfaces.

Drawn after Janin *et al.* (2008).

3.2.2 Composition of binding interfaces

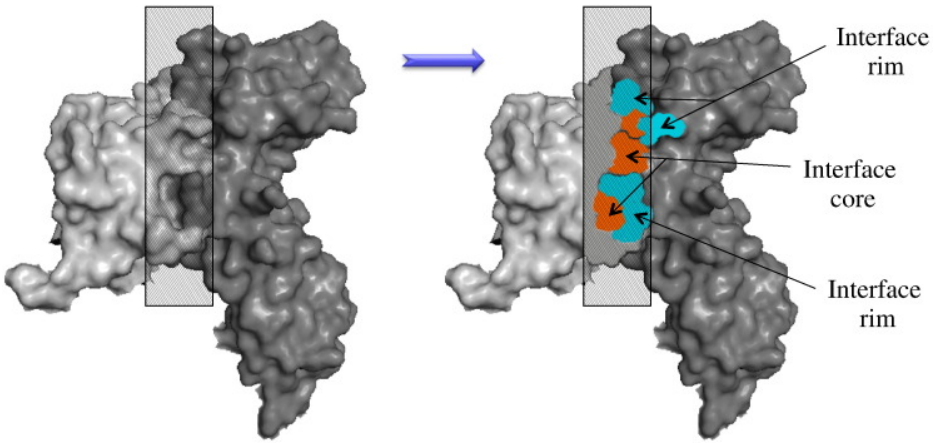

David and Sternberg (2015)

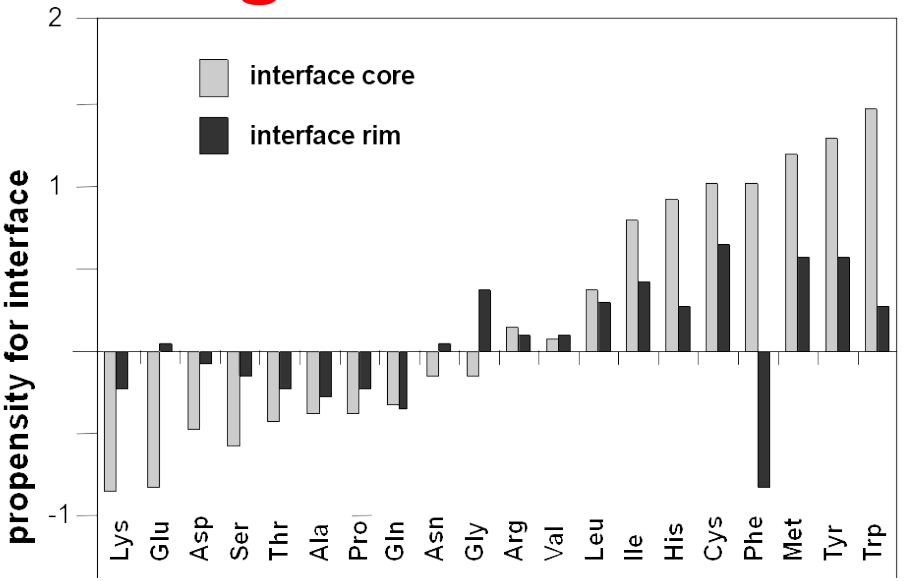

Residue propensities for core and rim regions at interfaces of protein–protein complexes. Drawn after Janin *et al.* (2008).

(Left) Residues in the center ("core") of the roughly spherical interface are "responsible" for making tight contact and are thus mostly occluded from solvent. (Right) the core region is strongly enriched in aromatic residues and depleted in charged residues. The surrounding ring of "rim" residues is much more similar to the remaining protein surface as these residues make partial contact to solvent molecules even in the bound state.

3.2.3 Hot spot residues

Hot spot residues at interfaces:
affinity drops by > 2 kcal/mol when such a residue is mutated to Ala.

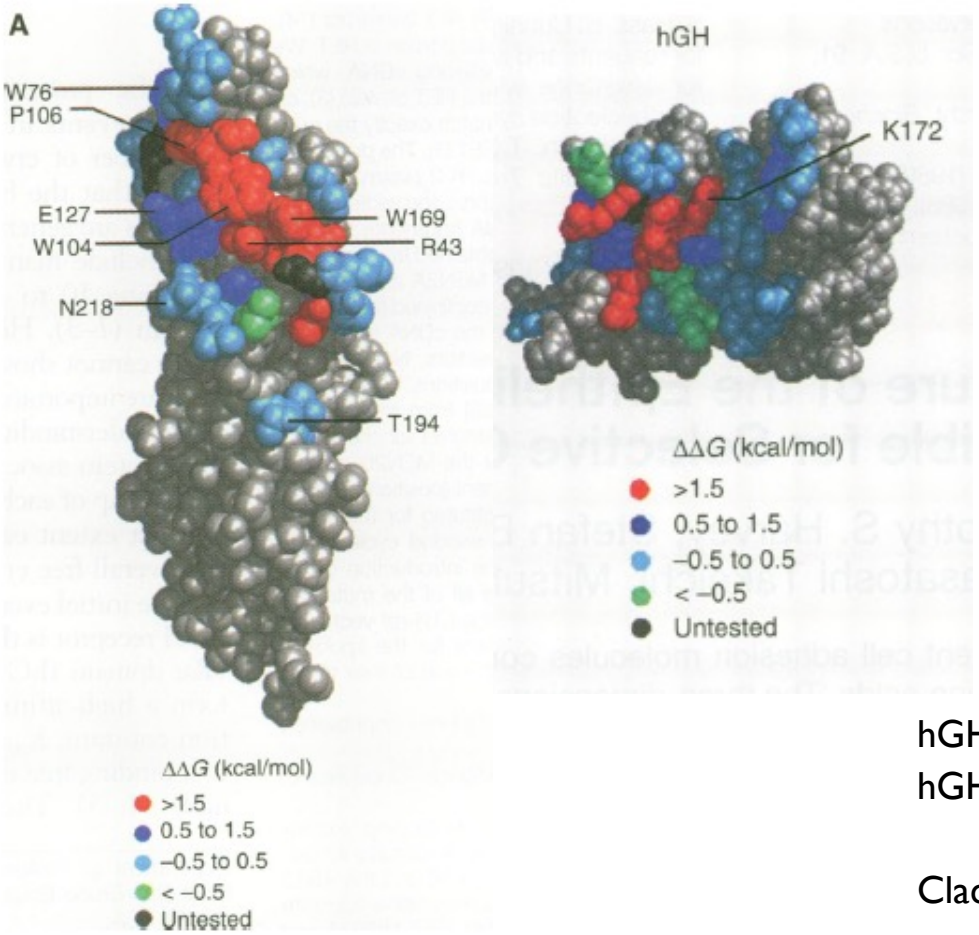

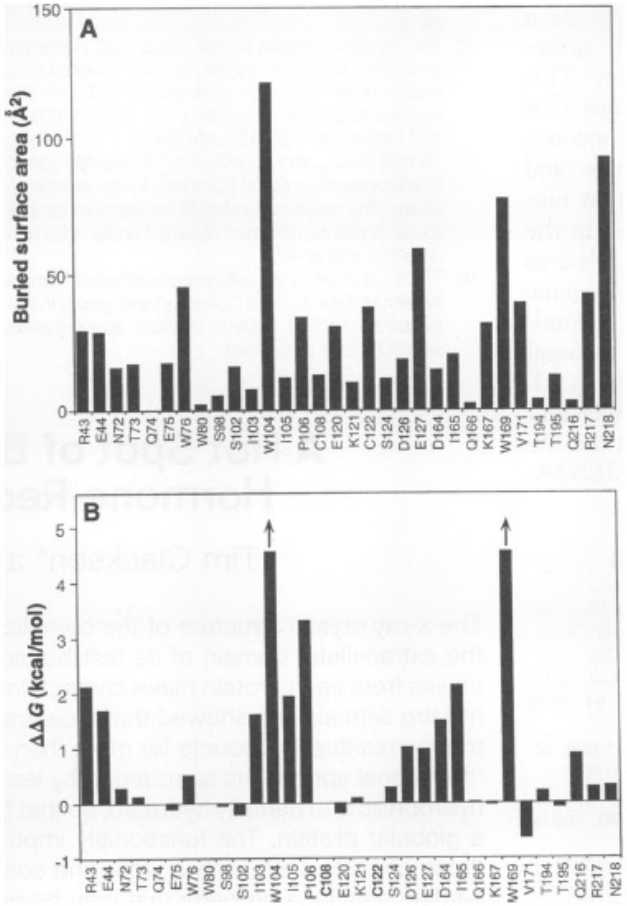

hGH: human growth hormone
hGHR: human growth hormone receptor

Clackson, Wells, Science 267, 383 (1995)

3.2.5 Predicting binding affinities

The total buried SASA has a Pearson correlation of 0.46 with experimental protein binding affinities.

Best available regression model:

$$\Delta G_{calc} = 0.09459 \times ICs_{charged/charged} + 0.10007 \times ICs_{charged\_apolar} - 0.19577 \times ICs_{polar/polar} + 0.22671 \times ICs_{polar/apolar} - 0.18681 \times \%\,NIS_{apolar} - 0.13810 \times \%\,NIS_{charged} + 15.9433\ [kcal/mol]$$

NIS: non-interacting surface

IC: # contacts between residues across the binding interface

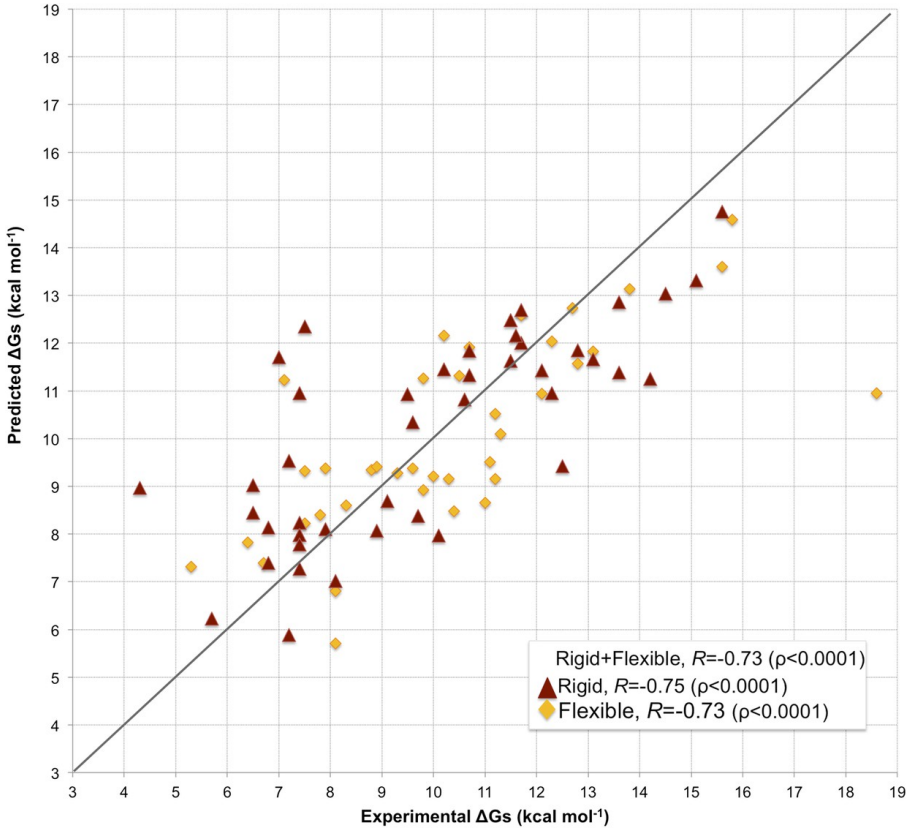

Scatter plot of predicted vs experimental binding affinities.
The predictions were made with the above regression model for a dataset of 81 protein–protein complexes. The correlation for all 81 complexes yields an R of −0.73 ($\rho < 0.0001$) with a RMSE of 1.89 kcal mol^{-1}.

rigid cases have iRMSD between superimposed free and bound components ≤1.0 Å
flexible cases have iRMSD >1.0 Å

Vangone et al. Elife 4, e07454 (2015)

3.3.1 Pairing propensities

Given the set of interface residues on both proteins,
one may analyse what **contacts** each of them forms with residues on the other protein.

A typical **distance threshold** for defining **contacts** is that they have pairs of atoms closer than e.g. 0.5 nm.

The computed statistics are conveniently represented in a 20 x 20 matrix.

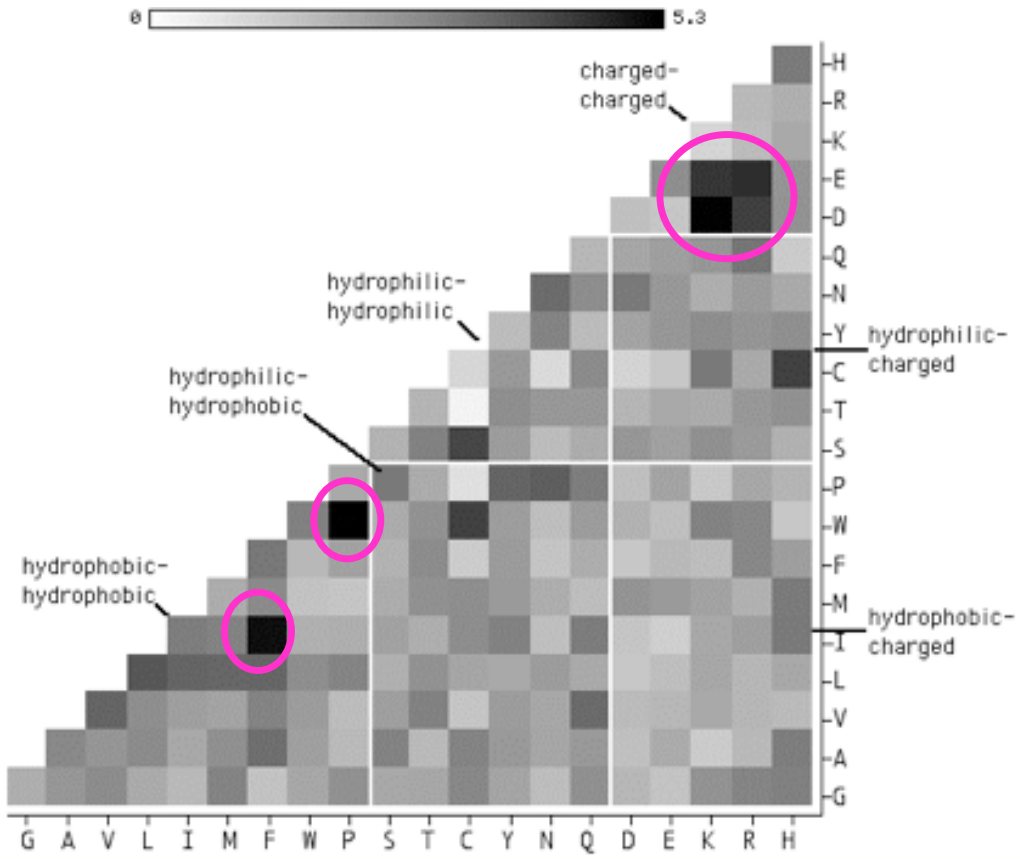

Amino-acid propensity matrix of transient protein-protein interfaces. Scores are normalized pairing frequencies of two residues that occur on the protein-protein interfaces of transient complexes.

Ansari and Helms (2006).

3.3.1 Pairing propensities

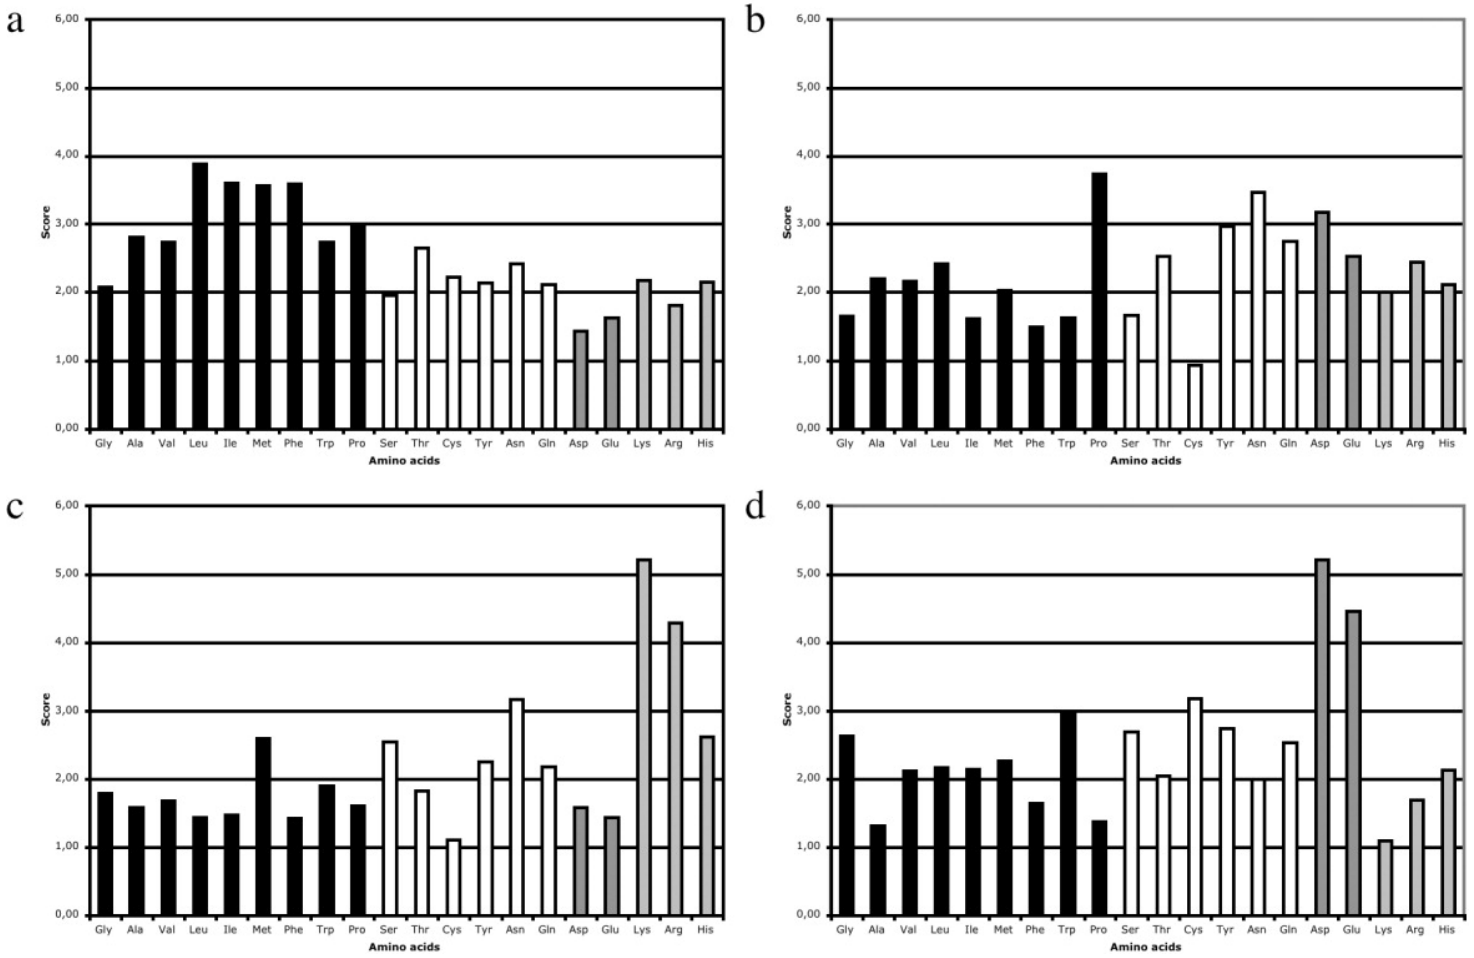

Black: hydrophobic residues
White: hydrophilic residues
Grey : charged residues.

Relative occurrence for binding partners of (a) leucine, (b) asparagine, (c) aspartate, and (d) lysine.
The higher the score, the more frequently such pairs occurred in the dataset.

3.3.1 Pairing propensities

From the observed count statistics, one can compute **interfacial pair potentials** *P(i,j) (i = 1…20, j = 1 … 20).*

$$P(i,j) = -\log\left(\frac{N_{obs}(i,j)}{N_{\exp}(i,j)}\right)$$

$N_{obs(i,j)}$: observed number of contacting pairs of *i,j* between two chains,
$N_{exp}(i,j)$: expected number of contacting pairs of *i,j* between two chains.

$N_{exp}(i,j)$ is computed as

$$N_{\exp}(i,j) = X_i \times X_j \times X_{total}$$

X_i : mole fraction of residue *i* among the total surface residues
X_{total} : total number of contacting pairs.

P(i,j) < 0 : observed frequency higher than expected
P(i,j) > 0 : less

3.3.2 Pair distribution function

A **radial pair distribution function** counts all pairs of amino acids at varying distance.
This distribution is then normalized with respect to an ideal gas, where particle distances are completely uncorrelated.

right: Pair distribution function of finding two alanine residues at a given distance in a protein.

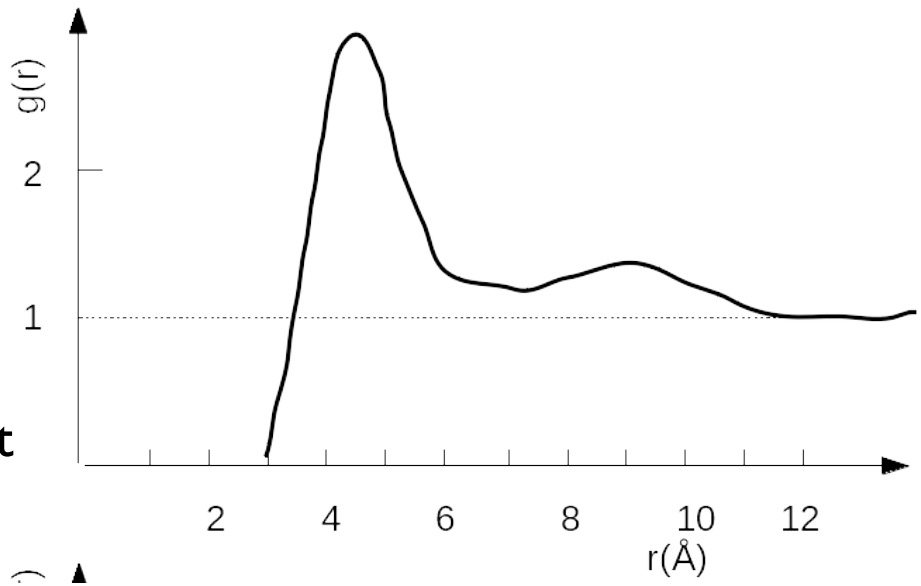

Hydrophobic Ala amino acids are mostly found in the hydrophobic core of proteins. Thus, we expect to find more Ala-Ala pairs at relatively short distances than at distances spanning from one side of the protein to the other one.

3.3.2 amino acid statistical potentials

According to the **Boltzmann distribution**, the occupancy levels p_1 and p_2 of two states 1 and 2 of a system with according energies E_1 and E_2 will vary according to the exponentially weighted energy difference between them:

$$\frac{p_1}{p_2} = e^{-\frac{E_1 - E_2}{kT}}$$

If we invert this formula ("Boltzmann inversion"), we can deduce an effective (free) energy function *G(r)* for the interaction between pairs of amino acids from these radial distribution functions *p(r)*,

$$G(r) = -k_B T \ln p(r)$$

These effective potentials can be used to score candidate conformations.

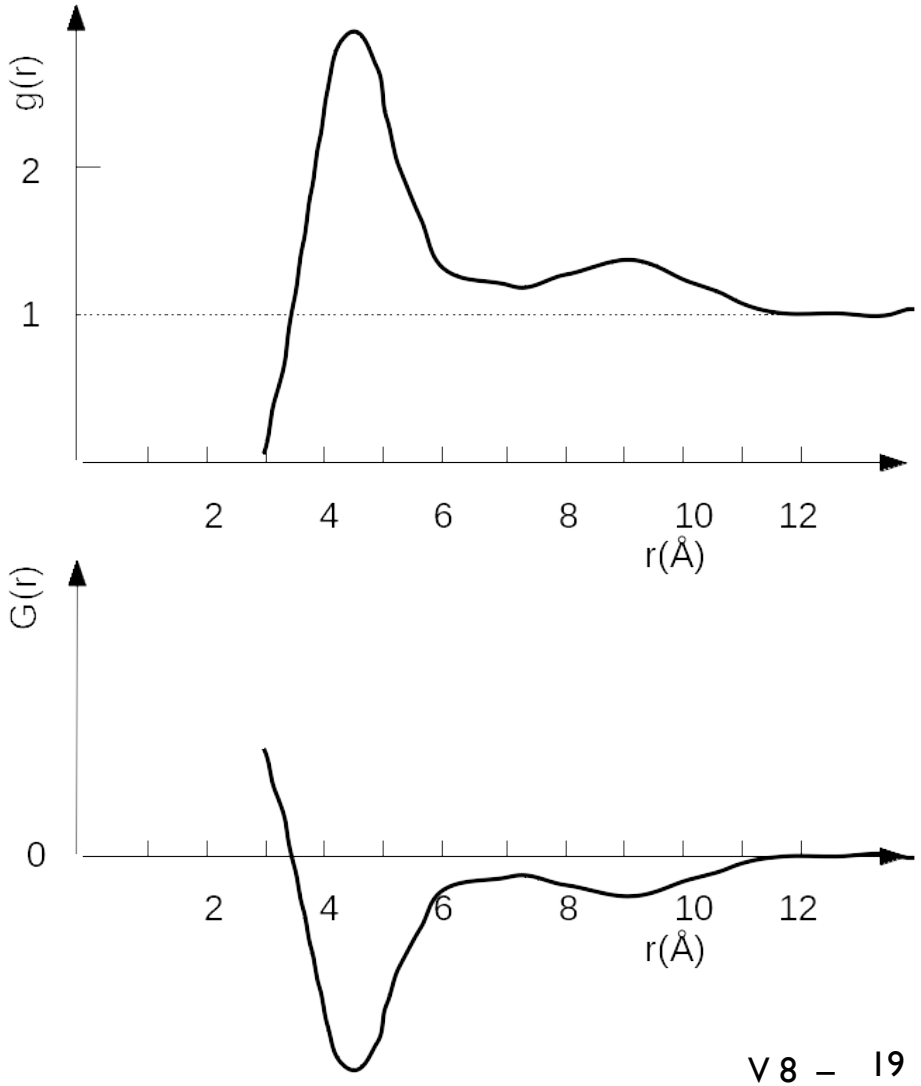

3.3.3 Conservation at interfaces

Functional constraints are expected to limit the amino acid substitution rates in proteins, resulting in a **higher conservation** of **functional sites** such as binding interfaces with respect to the rest of the protein surface.

There exist various approaches for analysing **evolutionary conservation** in MSAs. One of the simplest approaches is the **variance-based method**,

$$C(i) = \sqrt{\sum_j \left(f_j(i) - f_j\right)^2}$$

$C(i)$: conservation index for sequence position i in MSA,
f_j : overall frequency of amino acid j in the alignment
$f_j(i)$: frequency of amino acid j at sequence position i.

Positions with $f_j(i)$ equal to f_j for all amino acids j are assigned $C(i) = 0$.

On the contrary, $C(i)$ takes on its maximum for the position occupied by an invariant amino acid whose overall frequency in the alignment is low.

3.3.3 Conservation at interfaces

Another way of measuring conservation is based on the **entropy** of characters at position *i*,

$$C(i) = \sum_{j=1}^{20} f_j(i) \ln f_j(i)$$

This expression takes on its maximal value for *C(i)* (with the highest entropy) when all amino acids appear with the same frequency 1/20 in position *i*.

If the position is fully conserved, so that *f(X)* = 1 for one particular amino acid X and 0 otherwise, the entropy takes on its lowest possible value.

The **rate4site** algorithm (Mayrose et al. 2004) detects conserved amino acid sites in a multiple sequence alignment (MSA) given as input.

First, the algorithms generates a phylogenetic tree that matches the available MSA (or a pre-calculated tree provided by the user). Then, the algorithm computes a relative measure of conservation for each position in the MSA.

3.3.3 Visualize conservation: Consurf

The popular online-tool Consurf visualizes conservation scores computed with rate4site on 3D protein structure.

The results are color-coded by the degree of evolutionary conservation.

Red : strongly conserved,
blue : weakly conserved.

As anticipated, most of the residues at the inter-subunit interfaces are highly evolutionarily conserved.

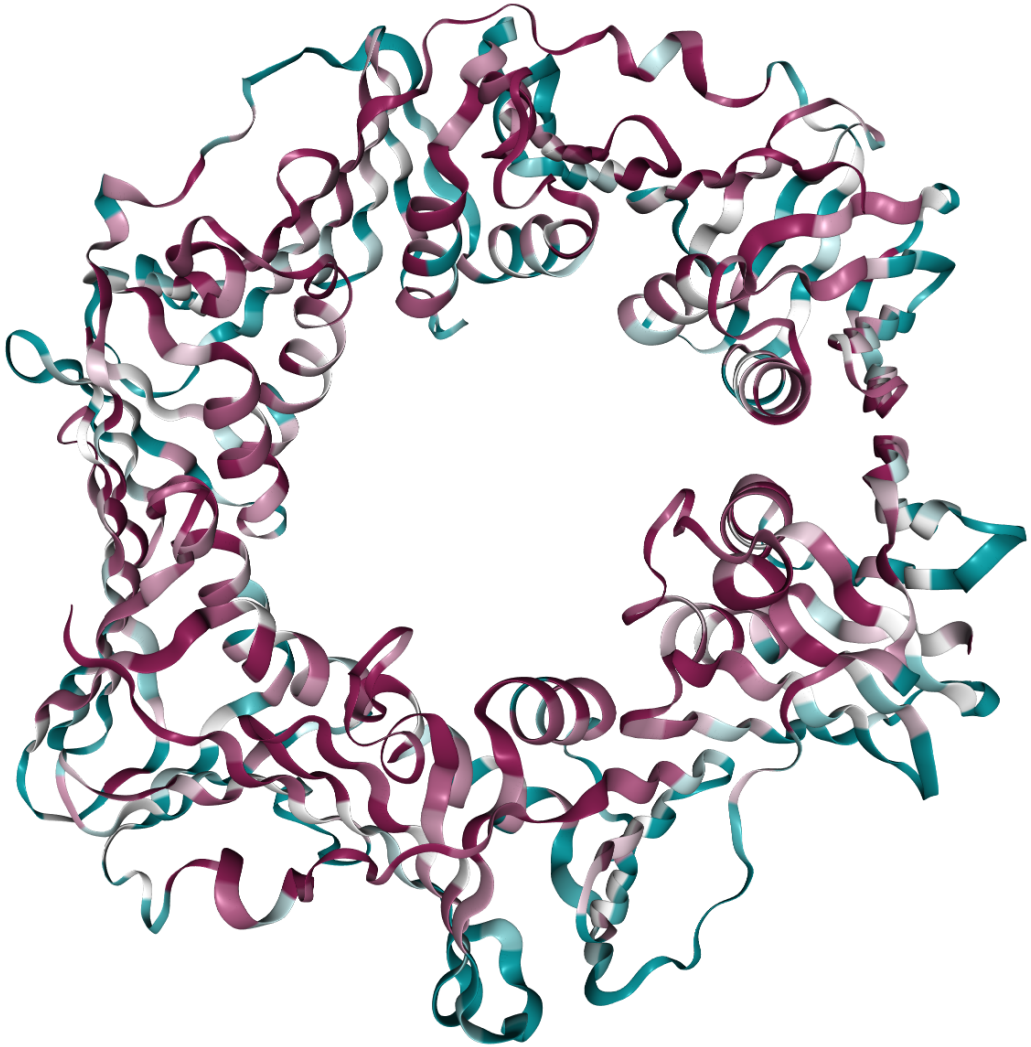

Conservation of surface residues at the dimer interface of the homo dimer of the β subunit of DNA polymerase III from *Escherichia coli* (Ashkenazy *et al.* 2016).

(review V4) Phylogenetic Profiling

Idea: either **all** or **none** of the proteins of a complex should be **present** in an organism

→ compare presence of protein homologs across species (e.g., via sequence alignment)

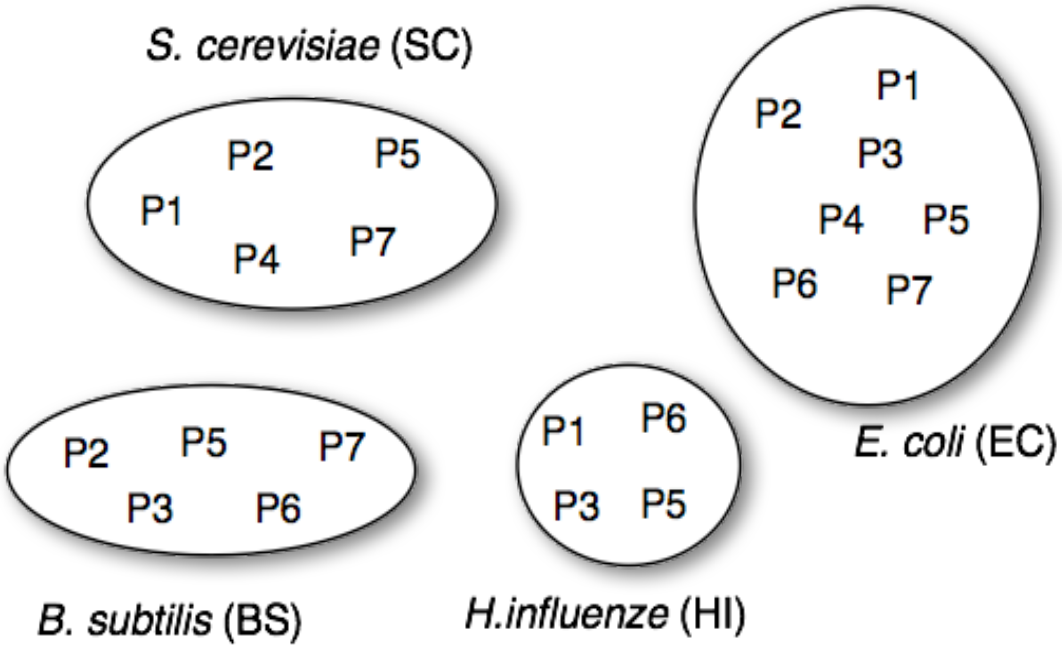

(review V4) Correlated mutations (Gremlin)

Detect positional correlations in paired multiple sequence alignments of thousands of protein sequences.

Gremlin constructs a global statistical model of the alignment of the protein family pair A and B by assigning a probability to every amino acid sequence in the paired alignment:

$$p(X_1, X_2, \ldots, X_p; X_{p+1}, \ldots, X_{p+q}) = \frac{1}{Z} exp\left(\sum_{i=1}^{p+q}\left[v_i(X_i) + \sum_{j=1}^{p+q} w_{ij}(X_i, X_j)\right]\right)$$

X_i : amino acid composition at position *i*,

v_i : vectors encoding position-specific amino acid propensities

w_{ij} :matrices encoding amino acid coupling between positions *i* and *j*.

Z : partition function, normalizes sum of probabilities to 1.

v_i and w_{ij} are obtained from the aligned sequences by a maximum likelihood approach.

The derived coupling strengths w_{ij} are then normalized and converted into distance restraints that can be used e.g. in scoring protein-protein docking models.

Ovchinnikov, Kamisetty, Baker (2014) eLife 3:e02030

(review V4) Correlated mutations

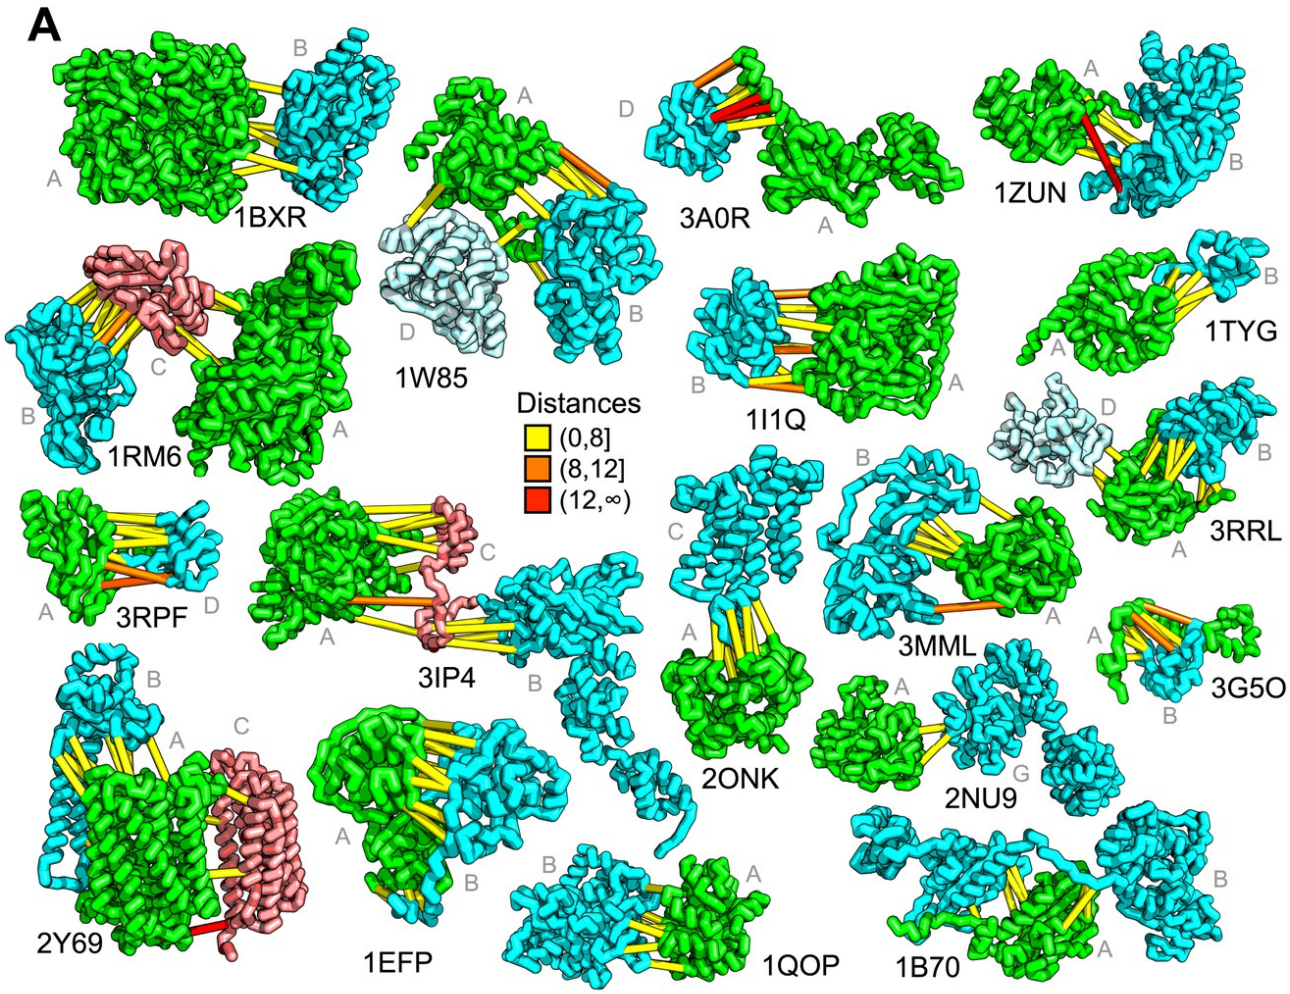

Ovchinnikov, Kamisetty, Baker (2014) eLife 3:e02030

Residue-pairs across protein chains with high GREMLIN scores almost always make contact across protein interfaces in experimentally determined complex structures.

Shown are all contacts with GREMLIN scores greater than 0.6.
Yellow: residue pairs within a distance < 8 Å,
orange: 8 - 12 Å,
red: > 12 Å.
The structures are pulled apart for clarity.

Caveats of phylogenetic profiling

phylogenetic profiling uses presence/absence correlations across genomes to predict interactions.

Its accuracy suffers, however, from a number of shortcomings and confounding factors:

(1) Phylogenetic relationships between considered genomes may introduce correlations unrelated to biological function;

single evolutionary events may be statistically amplified when closely related species are included in the data.

Evolutionary models taking into account the underlying species tree have been proposed to prune such correlations.

Croce, …, Weigt (2019)
PLoS Comput Biol 15, e1006891

Caveats of phylogenetic profiling

(2) Correlations may result from direct couplings, e.g., when two domains or proteins interact physically, but they may also be caused by intermediate partners: If A co-occurs with B, and B with C, also A and C will show correlations. Analyses based on partial correlations and spectral analysis have been proposed to *disentangle direct from indirect correlations*.

(3) Simple presence/absence patterns cannot *discriminate physical interaction from more general relationships*, like co-occurrence in a biological pathway or genomic co-localization.

Here, using full amino-acid sequences ($\rightarrow$ DCA) instead of presence/ absence patterns may help to refine the analysis, e.g. via the comparison of protein-specific phylogenetic trees.

But DCA cannot be applied on a genome-scale. It is computationally prohibitively expensive.

Croce, …, Weigt (2019)
PLoS Comput Biol 15, e1006891

Extension: phyletic coupling of protein domains

Idea: combine the 2 approaches.

By performing the genome-wide analysis on the coarse scale of presence/absence patterns,
one can identify promising protein-domain pairs, which are subsequently analyzed using DCA at the fine residue scale.

Extract composition of bacterial genomes from Pfam database.

$(n_1, \ldots, n_N)$: binary vector, characterizes presence (n_i = 1) or absence (n_i = 0) of domain i in a species.

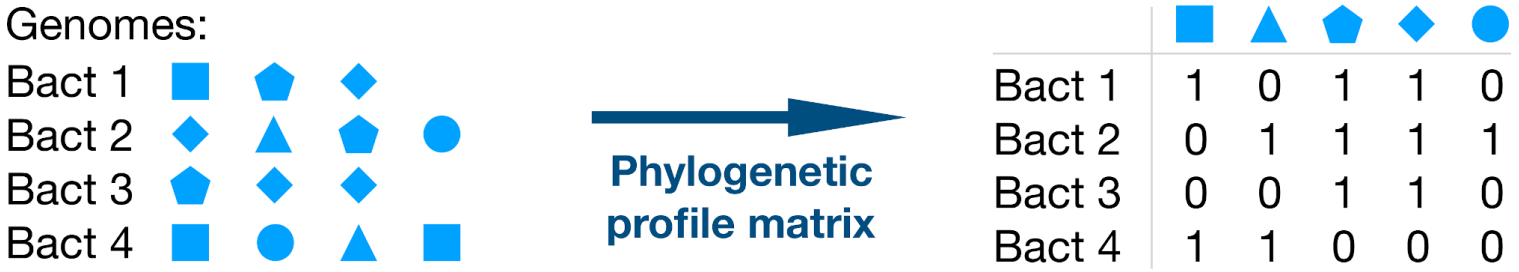

	■	▲	⬟	◆	●
Bact 1	1	0	1	1	0
Bact 2	0	1	1	1	1
Bact 3	0	0	1	1	0
Bact 4	1	1	0	0	0

Croce, …, Weigt (2019)
PLoS Comput Biol 15, e1006891

Extension: phyletic coupling of protein domains

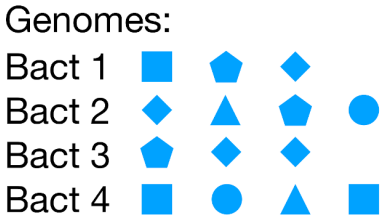

Phylogenetic profile matrix

	■	▲	⬟	◆	●
Bact 1	1	0	1	1	0
Bact 2	0	1	1	1	1
Bact 3	0	0	1	1	0
Bact 4	1	1	0	0	0

Global statistical model

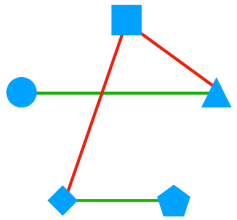

Phyletic couplings

$$P(n_1, ..., n_N) = \frac{1}{Z} \exp\left\{ \sum_{i<j} J_{ij} n_i n_j + \sum_i h_i n_i \right\}$$

Strong positive J_{ij} couplings: favored domain-domain co-occurrence.

Negative couplings: avoided co-occurrence = indicate alternative solutions for same biological function.

Determine **phyletic couplings** J_{ij} and **biases** h_i such that the model P reproduces the 1- and 2-column statistics of PPM matrix:

$$f_i = \frac{1}{M} \sum_{a=1}^{M} n_i^a$$

$$f_{ij} = \frac{1}{M} \sum_{a=1}^{M} n_i^a n_j^a$$

f_i : fraction of genomes carrying domain *i*

f_{ij} : fraction of genomes containing domains *i* and *j* simultaneously

Solve P by „pseudo-likelihood maximization“.

Croce, …, Weigt (2019)
PLoS Comput Biol 15, e1006891

Domain-domain couplings J_{ij}

Distribution of couplings J_{ij}
for all domain-domain pairs (blue) has strong peak at 0.

Broad tail for larger positive values for known domain-domain pairs (brown).

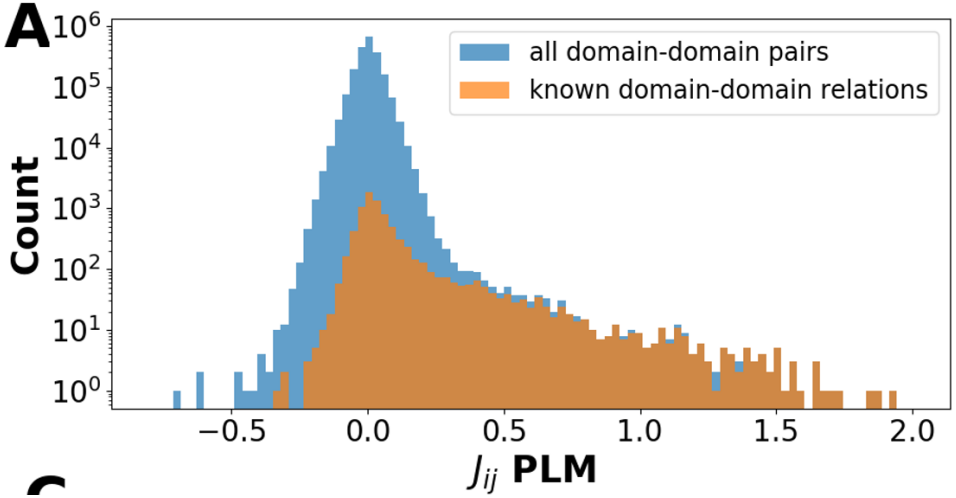

A large positive phyletic coupling is a strong signal for a positive relationship between two domains, but not necessarily for a direct physical interaction of the two domains within a protein complex.

E.g. co-localization of two domains either inside the same protein (i.e. an evolutionary conserved protein architecture) or inside the same operon may lead to strong phyletic couplings.

Croce, …, Weigt (2019)
PLoS Comput Biol 15, e1006891

Domain pairs with strong couplings

Study domain pairs with large positive phyletic coupling by DCA.

DCA scores above 0.2 at an effective sequence-pair number of at least 200 (sequences below 80% sequence identity) can be considered as a strong indicator for a potential interaction.

On the contrary, a low DCA score is not necessarily a sign for the absence of a physical interaction. A low score might also originate from a relatively small or structurally not well-conserved interface, both resulting in a weak coevolutionary signal.

Croce, …, Weigt (2019)
PLoS Comput Biol 15, e1006891

Yes, DCA couplings separate TP from random

Panel A shows the effective sequence number (defined as the sequence number at 80% maximum sequence identity) and the DCA scores for the 500 most significant predict-tions (strongest phyletic couplings).

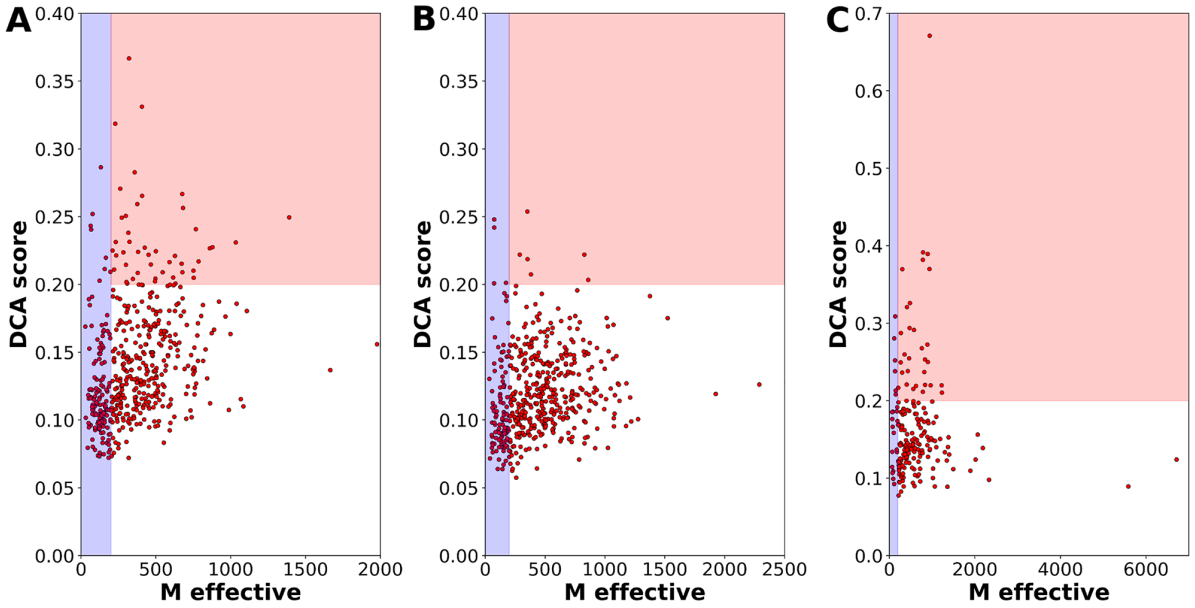

The interesting region is the red one, where sequence numbers are sufficient to provide reliable DCA results, and DCA scores are beyond 0.2.

For comparison, panel B shows the results for 500 randomly selected domain pairs. Only very few pairs show substantial scores, most of them related to very short peptides.

Panel C shows a positive control, the 200 pairs of highest phylogenetic couplings belonging to iPfam are analyzed analogously. The fraction and the amplitude of high DCA scores is slightly increased with respect to Panel A, but the qualitative behavior is similar.

Croce, …, Weigt (2019)
PLoS Comput Biol 15, e1006891

Summary: PP complexes

"Proteins are **modular machines**" <=> How are they related to each other?

1) **Detect** structures of protein complexes
   X-ray, NMR, EM

2) **Integrate** data: density fitting (FFT, Laplace filter) (V2)

3) Protein **docking**, combinatorial assembly (CombDock, StarDock, Mosaic, DACO) (V3)

4) **Analyze** protein interfaces: composition, conservation, size (V8)
   predict binding affinities

Summary: Static PPI-Networks

"Proteins are **modular machines**" <=> How are they related to each other?

1) **Understand** "Networks" in principle
   prototypes (ER, SF, …) and their properties (*P(k), C(k), d,* clustering, …)

2) Get the **data** (V4)
   experimental and computational approaches (Y2H, TAP, co-regulation, …),
   quality control and data integration (Bayes, V5)

3) **Analyze** the data
   compare *P(k), C(k),* clusters, … → highly modular, clustered
   obscured by sparse sampling (V7) → PPI networks are not strictly scale-free

4) Identify **modules**
   Girvan-Newman (V5), Radicchi (V5), Kernighan-Lin (V7)

5) **Predict** missing information
   network structure combined from multiple sources
→ functional annotation (V7)

Next part of lecture: gene-regulatory networks

Content of final exam

Lecture	Slides relevant for exam
1	15-21
2	1-11, 30-42
3	1-40
4	None
5	1-20,24-32,37-39,42
6	None
7	1-21
8	None
9	
10	
11	
12	

Lecture	Slides relevant for exam
13	
14	
15	
16	
17	
18	
19	
20	
21	
22	
23	
24	
25	

Relevant are also the assignments !
(theoretical parts, not the programming parts)

Additional sides (not used)

Mesoscale properties of networks
- identify cliques and highly connected clusters

Most relevant processes in biological networks correspond to the **mesoscale** (5-25 genes or proteins), not to the entire network.

However, it is computationally enormously expensive to study mesoscale properties of biological networks.

E.g. a network of 1000 nodes contains 1×10^{23} possible 10-node sets.

Spirin & Mirny analyzed combined network of protein interactions in *S. cereviseae* with data from CELLZOME, MIPS, BIND: 6500 interactions.

Identify connected subgraphs

Aim: identify **fully connected subgraphs** (cliques) in protein interaction network.

A clique is a set of nodes that are all neighbors of each other.

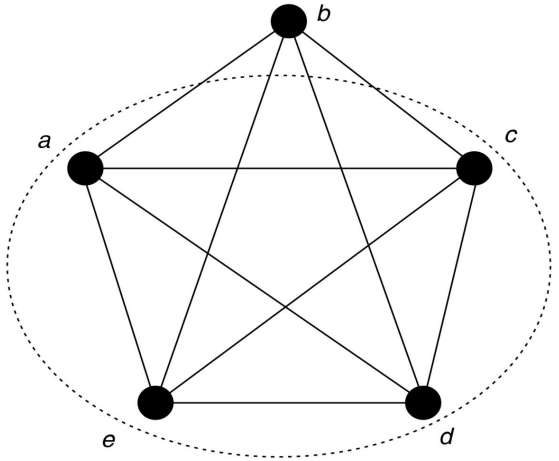

The „**maximum clique problem**“ – finding the largest clique in a given graph is known be NP-hard.

In this example, the whole graph is a clique and consequently any subset of it is also a clique, for example *{a,c,d,e}* or *{b,e}*.

A **maximal clique** is a clique that is not contained in any larger clique. Here only *{a,b,c,d,e}* is a maximal clique.

In general, protein complexes need not to be fully connected.

Spirin, Mirny,
PNAS 100, 12123 (2003)

Identify all fully connected subgraphs (cliques)

The general problem - finding all cliques of a graph - is very hard.

But the protein interaction graph is quite **sparse**:
interactions (edges) is similar to # proteins (nodes)).

-> the cliques can be found relatively quickly in the PPI network.

Idea:
cliques of size *n* can be found by enumerating the cliques of size *n-1* etc.

Spirin, Mirny, PNAS 100, 12123 (2003)

Identify all fully connected subgraphs (cliques)

Spirin & Mirny started their search for cliques with $n = 4$.

Consider all (known) pairs of edges (6500 × 6500 protein interactions).

For every **pair** *A-B* and *C-D* check whether there are edges

between *A* and *C*, *A* and *D*, *B* and *C*, and *B* and *D*.

If these edges are present, *ABCD* is a **clique**.

For every clique identified, *ABCD*, check all proteins in the PPI network.

For every additional protein *E*:

if all of the interactions *E-A*, *E-B*, *E-C*, and *E-D* exist,

then *ABCDE* is a clique with size 5.

Continue for $n = 6, 7, \ldots$

Spirin, Mirny, PNAS 100, 12123 (2003)

Identify all fully connected subgraphs (cliques)

The largest clique found in the protein-interaction network had size 14.

These results include, however, many redundant cliques.

E.g., the clique with size 14 contains 14 cliques with size 13.

To find all **nonredundant cliques**, mark all proteins in the clique of size 14.

Out of all subgraphs of size 13 pick those that have at least one protein other than marked.

After all redundant cliques of size 13 are removed,
proceed to remove redundant twelves etc.

In total, only 41 nonredundant cliques with sizes 4 - 14
were found by Spirin & Mirny.

Spirin, Mirny, PNAS 100, 12123 (2003)

Statistical significance of cliques

complete cliques as a function of clique size.

Red: real network of protein interactions

Blue: > 1000 randomly rewired graphs, that have the same number of interactions for each protein.

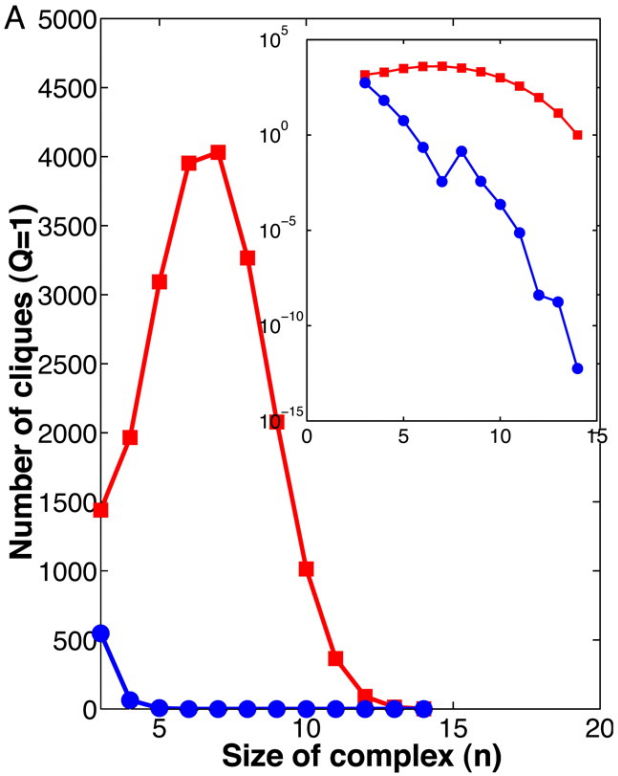

Inset shows the same plot on a log-normal scale. Note the dramatic enrichment in the number of cliques in the protein-interaction graph compared with the random graphs. Most of these cliques are parts of bigger complexes and modules.

Spirin, Mirny, PNAS 100, 12123 (2003)

What else can you do with Interaction graphs?

E.g. efficiently track interactions
between many particles
in dynamic simulations

Strongly attracting particles form large “blob”

(a) to (d) are 4 snapshots of a simulation with ca. N = 50 interacting particles in a box.

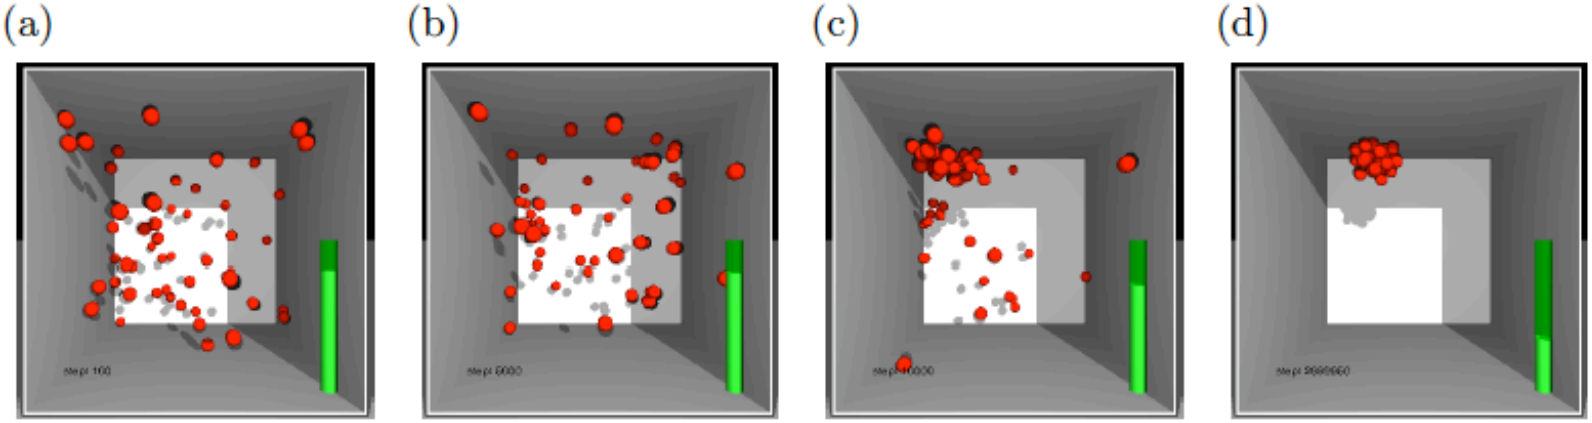

How can one analyze the particle connectivity efficiently?

M.Sc. thesis

Florian Lauck (2006)

For i = 1 to N - 1

For j = i + 1 to N

For k = j + 1 to N

If (i .is bound to. j) then

If (j .is bound to. k) then ….

this is impractical!

Map simulation to interaction graph

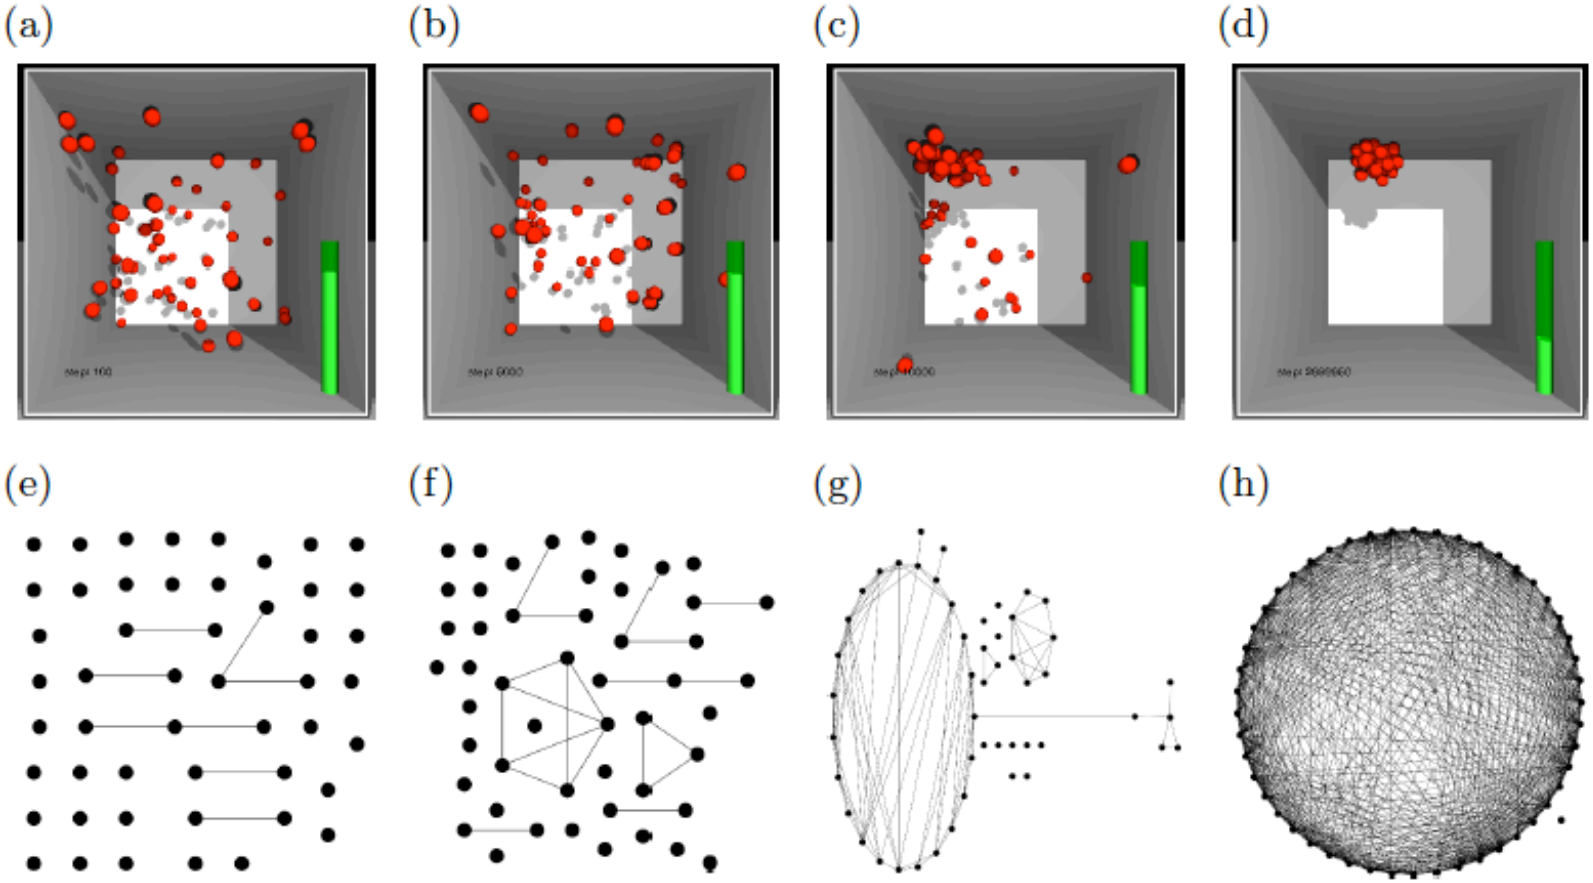

Figure 2.7: Graph and spatial view of a simulation with 50 particles at four different points in time. The green bar denotes the energy of the system.

M.Sc. thesis Florian Lauck (2006)

Large number of simultaneous assocications: map simulations to interaction graphs

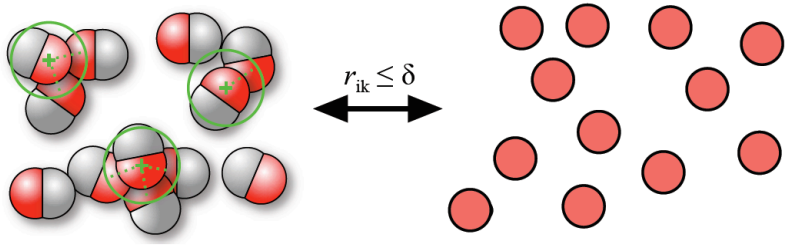

Simple MC scheme
for diffusion + association/
Dissociation

Bottom: possible interaction potentials

```
function INITIALIZE(N)
    for P ∈ List of Particles do
        CREATE RANDOM COORDINATES(P)
    CREATE GRAPH(N)
for all Iterations do
    for P ∈ List of Particles do
        MOVE AND ROTATE(P)
        for all  P_i ∈ (List of Particles - P) do
            d = DISTANCE(P, P_i)
            e_i = POTENTIAL(d)
            if d ≤ r_C then APPEND(List of Interactions, (P, P_i))
            E_new += e_i
        a = TRANSITION PROBABILITY(E_new, E_old)
        x = RANDOMNUMBER

        if x ≤ p then                                   ▷ accept new state
            APPEND(List of ALL interactions, List of Interactions)
            E_old = E_new
        else                                            ▷ discard new state
            RESET(P) CLEAR(List of Interactions)
    UPDATE(Graph, List of ALL Interactions)
    ANALYSIS(Graph)
```

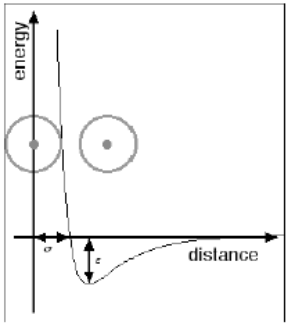

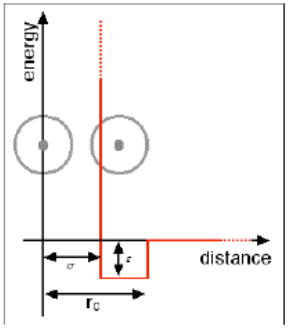

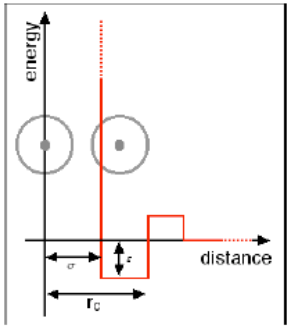

Interaction patches define complex geometry

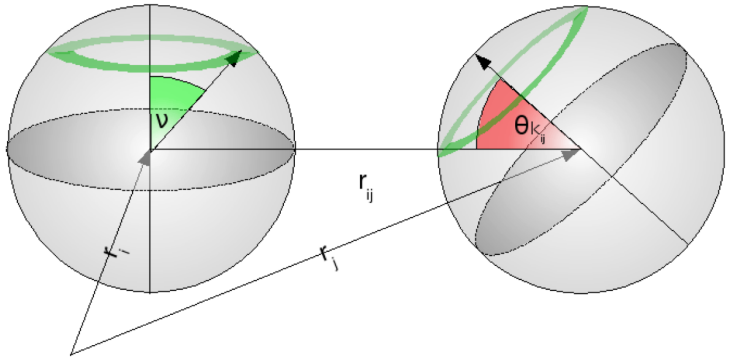

$$G_{ij}(r_{ij}, \theta_{ij}) = \exp\left[\frac{(\theta_{ij} - \nu)}{2\sigma^2_{PW}}\right]$$

$$V_{total} = V(r_{ij}) \times G_{ij}(r_{ij}, \theta_{ij}) \times G_{ji}(r_{ij}, \theta_{ji})$$

Interaction potential = distance dependent term ×orientation dep. terms

Lauck et al. , *JCTC 5, 641* (2009)

Assembly of icosahedral complexes

Too small N = 8

Ideal geometry N = 12

Not compact N = 12

Too large N = 16

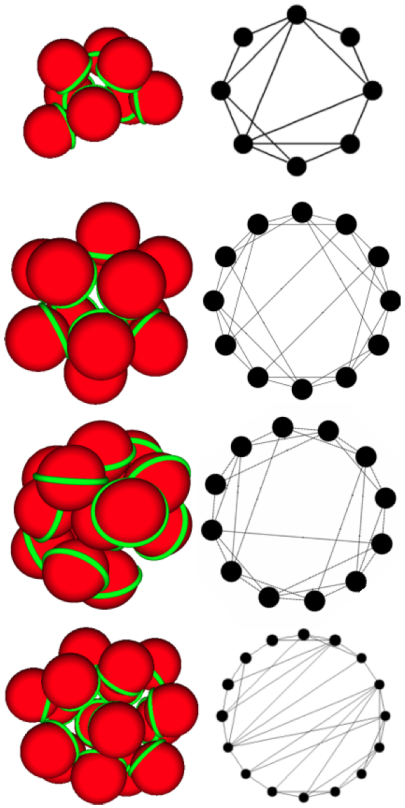

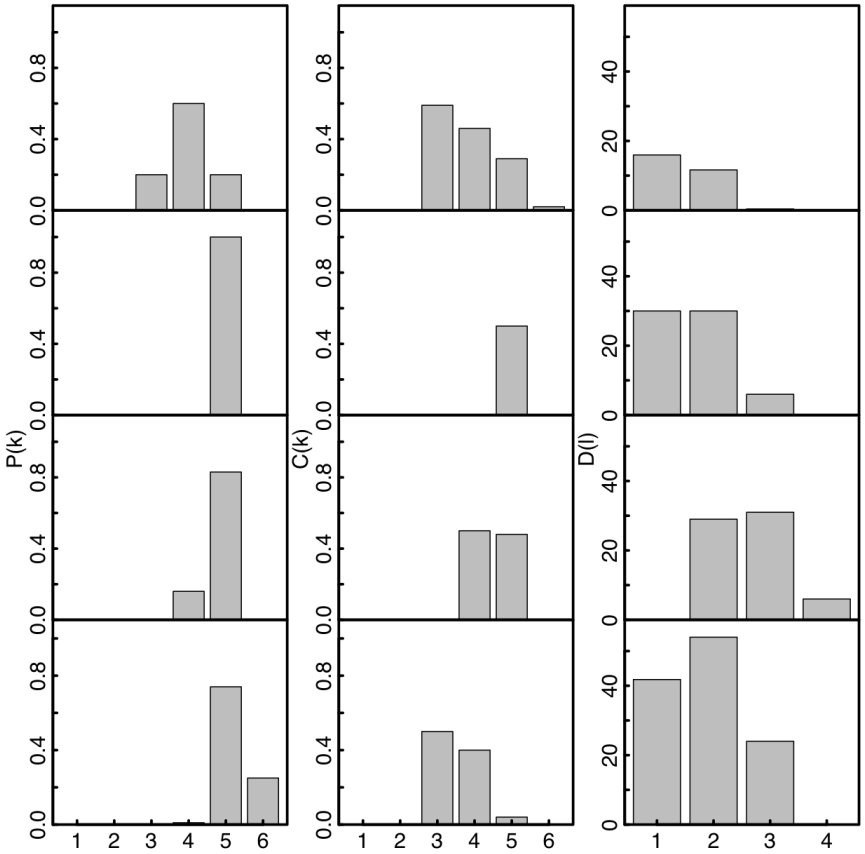

$$P(k) = \frac{N(k)}{N} \quad C(k) = \frac{\sum C_k}{n_k}$$

Degree distribution

Average Cluster coefficient

shortest pathways between nodes

Lauck et al. , *JCTC 5, 641* (2009)

Dynamic view at particle agglomeration

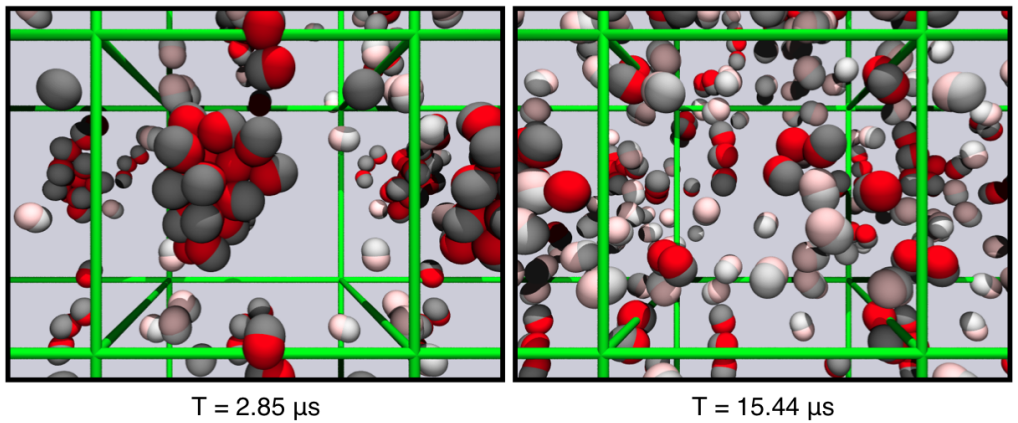

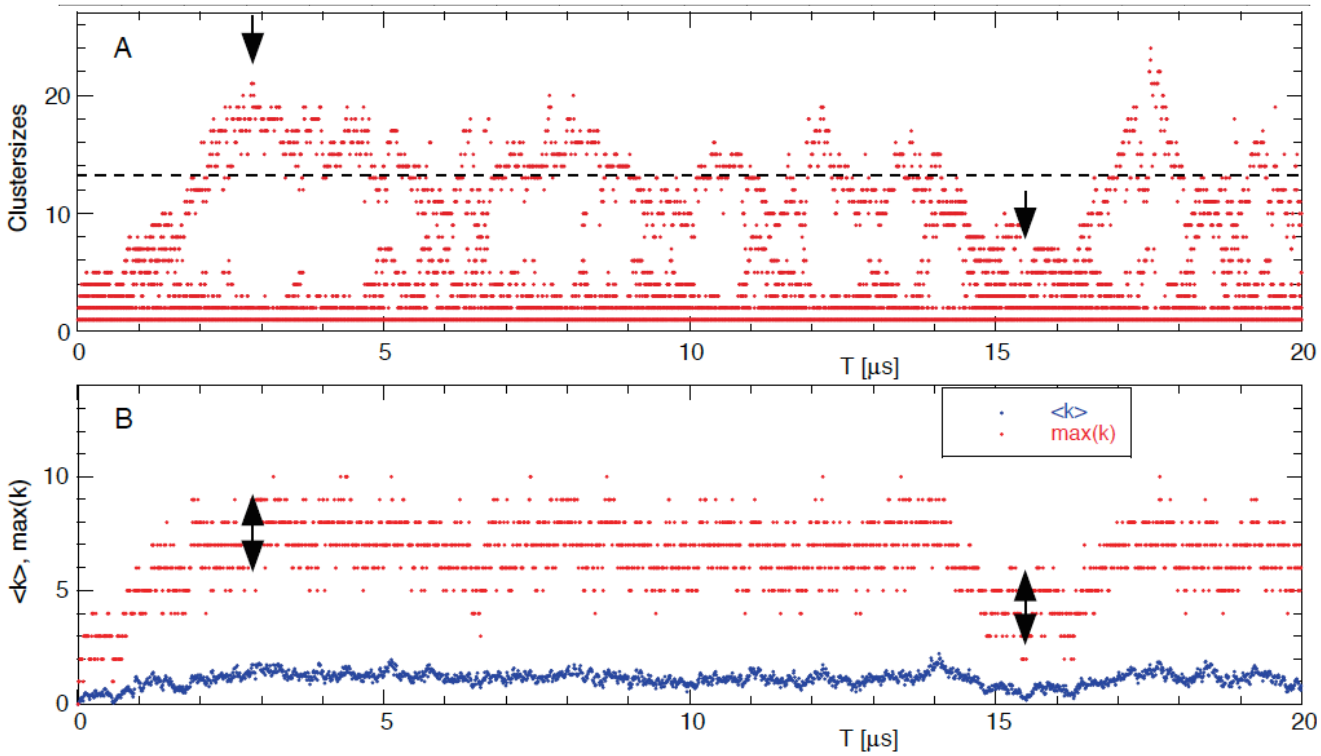

Two snapshots

T = 2.85 µs
most of the
particles are part
of a large cluster,

T = 15.44 µs
largest cluster
has 3 particles.

Geyer,

BMC Biophysics (2011)