

Protein Evolution

Lecture 3

Lucy Colwell, ICTS Winter School on
Quantitative Systems Biology, Dec 14 2017

Outline

Biology 101 – a quick tour of proteins

Biological information from covariance


- Protein structure and function
- Protein interaction specificity
- Ligand protein binding

Theoretical analysis of data limitations

- when can we predict reliably?
- Noise due to finite sampling, samples are not independent.
- central role played by Random Matrix Theory.


Exploit correlation structure of protein sequences

What constraints on the mutations that can be accepted are induced by the requirements of 3D structure and function of the protein?



Can we **exploit correlations** in the sequence data to build a model that can predict **the 3D protein structure**?


Evolutionary constraints for adrenergic beta receptor 2



Transmembrane protein predicted topology

15e71e73-ae04-493e-a3ca-254c7341265c.mem

Membrane helical secondary structure




CENTER FOR BIOLOGICAL SEQUENCE ANALYSIS CBS	EVENTS	NEWS	RESEARCH GROUPS	CBS PREDICTION SERVERS	CBS DATA SETS	PUBLICATIONS	EDUCATION
	STAFF	CONTACT	ABOUT CBS	INTERNAL	CBS BIOINFORMATICS TOOLS	CBS COURSES	OTHER BIOINFORMATICS LINKS

[CBS](#) >> [CBS Prediction Servers](#) >> [TMHMM](#)


TMHMM Server v. 2.0

Prediction of transmembrane helices in proteins





[Instructions](#)


Transmembrane protein predicted topology



Evolutionary constraints for adrenergic beta receptor 2




$\beta 2$ adrenergic receptor : evolutionary constraints



The $\beta 2$ adrenergic receptor, a seven transmembrane helix GPCR family member.

Blind prediction for $\beta 2$ adrenergic receptor




*Crystal Structure: Rasmussen SG, DeVree BT, Zou Y, Kruse AC, Chung KY, Kobilka TS, Thian FS, Chae PS, Pardon E, Calinski D, Mathiesen JM, Shah ST, Lyons JA, Caffrey M, Gellman SH, Steyaert J, Skinioitis G, Weis WI, Sunahara RK, Kobilka BK. Nature. 2011 Jul 19.

Sequences to structure:


Blind structure prediction.

■ observed ■ predicted


β2 adrenergic receptor




G-3-P transporter GlpT




COX3




NADH-quinone oxidoreductase
subunit N 1



Aquaporin-0



Multidrug resistance protein norM



Hopf, Colwell *et al*, Cell 2012

Predict *de novo* structures for transmembrane proteins for which no homologous structure has been solved experimentally

Adiponectin receptor 1

Adiponectin receptor 1

- An essential hormone secreted by adipocytes that acts as an antidiabetic.
- 375 amino acids
- 7 TM helices

.....


Cloning of adiponectin receptors that mediate antidiabetic metabolic effects

**Toshimasa Yamauchi^{*†‡}, Junji Kamon^{*‡§}, Yusuke Ito^{*},
Atsushi Tsuchida^{*}, Takehiko Yokomizo^{||}, Shunbun Kita^{*},
Takuya Sugiyama[¶], Makoto Miyagishi[#], Kazuo Hara^{*†},
Masaki Tsunoda[☆], Koji Murakami[☆], Toshiaki Ohteki^{**†}, Shoko Uchida^{*},
Sato Takekawa^{*}, Hironori Waki^{*†}, Nelson H. Tsuno^{††}, Yoichi Shibata^{††},
Yasuo Terauchi^{*†}, Philippe Froguel^{‡‡}, Kazuyuki Tobe^{*†},
Shigeo Koyasu^{**†}, Kazunari Taira[#], Toshio Kitamura[¶], Takao Shimizu^{||},
Ryozo Nagai^{*} & Takashi Kadowaki^{*†}**

Levels of adiponectin in the blood are decreased under conditions of obesity, insulin resistance and type 2 diabetes


De novo predicted 3D structure of adiponectin receptor 1

predicted contact map




De novo predicted 3D structure of adiponectin receptor 1

3. Predicted structure



Top ranked predicted 3D structure of Adiponectin receptor 1 from sequences alone


De novo predicted 3D structure of adiponectin receptor 1




Top ranked predicted 3D structure of Adiponectin receptor 1 from sequences alone

De novo predicted 3D structure of adiponectin receptor 1

ADR1 prediction




Bacteriorhodopsin



Is the predicted 3D structure of Adiponectin receptor 1 similar to any known TM protein structures ?

Yes – 4.5 Å rmsd over >200 residues


De novo predicted 3D structure of adiponectin receptor 1



Predicted Adiponectin receptor 1 is flipped in the membrane compared to GPCRs


11 medically important membrane proteins of unknown structure predicted

OCTN1




Crohn's disease,
rheumatoid arthritis

Adiponectin receptor 1



diabetes, obesity,
cancer


MT-ND1




LHON, MELAS,
Alzheimer, Parkinson

11 medically important membrane proteins of unknown structure predicted

Adiponectin receptor 1



diabetes, obesity,
cancer



Structure solved by X-ray
crystallography in 2015 at
2.9Å resolution.

Our predicted structure is
3.7Å C-α rmsd over 186
residues from this crystal
structure 3WXV.

ARTICLE

doi:10.1038/nature14301

Crystal structures of the human adiponectin receptors


Hiroaki Tanabe^{1,2,3,4}, Yoshifumi Fujii^{1,4}, Miki Okada-Iwabu^{5,6*}, Masato Iwabu^{5,6,7*}, Yoshihiro Nakamura^{1,3,4*}, Toshiaki Hosaka^{1,3}, Kanna Motoyama¹, Mariko Ikeda^{1,3}, Motoaki Wakiyama^{1,3}, Takaho Terada^{1,4}, Noboru Ohsawa^{1,3}, Masakatsu Hato^{1,3}, Satoshi Ogasawara⁸, Tomoya Hino^{8,9}, Takeshi Murata^{1,8,9,10}, So Iwata^{1,8,9,11,12,13}, Kunio Hirata¹³, Yoshiaki Kawano¹³, Masaki Yamamoto¹³, Tomomi Kimura-Someya^{1,3}, Mikako Shirouzu^{1,3}, Toshimasa Yamauchi^{13,6,14}, Takashi Kadowaki^{13,6} & Shigeyuki Yokoyama^{1,2,4}

11 medically important membrane proteins of unknown structure predicted

MT-ND1



LHON, MELAS,
Alzheimer, Parkinson



Structure solved by X-ray crystallography in 2013 at 3.3Å resolution.

Our predicted structure is 3.67Å C-α rmsd over 253 residues from the crystal structure 4HE8.

ARTICLE

doi:10.1038/nature11871

Crystal structure of the entire respiratory complex I


Rozbeh Baradaran¹, John M. Berrisford^{1†}, Gurdeep S. Minhas¹ & Leonid A. Sazanov¹

Beyond structure prediction

Constraints reflect multimerisation

Beyond structure prediction


Constraints reflect multimerisation




S. typhimurium ABC transporter
(antibiotic transporter)

- ★ Top ranked evolutionary constraints (ECs)
- Multimer ECs
- Crystal structure contacts

Evolutionary constraints identify dimer contacts




E. coli methionine importer MetI



E. coli methionine importer MetI


3D complexes for proteins of unknown monomer structure




adiponectin receptor model – violated strong constraints are predicted dimer contacts

Predicted dimer interface, adiponectin receptor

ADIPOR1 (200 ECs)




Predicted dimer interface contacts




Predicted dimer interface, adiponectin receptor

ADIPOR1 (200 ECs)



AdipoR1 predicted dimer




Beyond structure prediction

Constraints reflect conformational change

Beyond structure prediction

Constraints reflect conformational change

4 Å C- α rmsd over
300 residues




*Structure and mechanism of the glycerol-3-phosphate transporter from Escherichia coli.

Huang, Y., Lemieux, M.J., Song, J., Auer, M., Wang, D.N., *Journal: (2002) Science 291: 516-520*


Glycerol-3-phosphate transporter: evolutionary constraints

Constraints reflect conformational change




Glycerol-3-phosphate transporter: evolutionary constraints


Constraints reflect conformational change



Method can be generalized to other membrane proteins to explore conformational flexibility




Glycerol-3-phosphate transporter, a secondary active membrane transporter.





Protein responsible for the cotransport of sodium ions and ergothioneine, which is an antioxidant, into cells

Suggests 'helix rocking' typical of alternating access membrane transporters



Conformational change





Protein protein interactions



Predicting alternative conformations



Protein protein interactions




Beyond structure prediction

Constraints reflect multimerisation


Beyond structure prediction

Constraints reflect multimerisation



S. typhimurium ABC transporter
(antibiotic transporter)

- ★ Top ranked evolutionary constraints (ECs)
- Multimer ECs
- Crystal structure contacts



Outline

Biology 101 – a quick tour of proteins

Biological information from covariance

- Protein structure and function
- Protein interaction specificity
- Ligand protein binding

Theoretical analysis of data limitations

- when can we predict reliably?
- Noise due to finite sampling, samples are not independent.
- central role played by Random Matrix Theory.

Accurately predict protein interaction partners


Covariance calculated from the **concatenated sequences of protein interaction partners** accurately predicts the protein complex structure.

BUT.....

To apply this analysis, we need a starting set of known interaction partners, so we can form the concatenated sequence alignment.

This is particularly difficult to obtain if genomes contain multiple paralogs (versions) of either of the two proteins, as is frequently the case – PFAM A-30 families have on average 3.9 paralogs per species.

Protein interaction partners – HKs and RRs



	a1	a2	a3	a4	a5	a6		b1	b2	b3	b4	b5	b6	b7
<i>E. c</i>	K	D	I	L	V	D		T	I	F	D	F	C	L
<i>S. a</i>	K	D	I	L	V	D		T	I	F	D	P	C	I
<i>T. t</i>	D	D	I	K	V	D		T	I	F	H	P	C	L
<i>S. e</i>	D	E	I	L	V	H		S	I	F	H	P	C	I

Protein interaction partners – HKs and RRs

Can we invent an algorithm that does not need a set of known interaction partners?




Anne-Florence Bitbol




Ned Wingreen

Use the same approach to predict interactions between proteins




Protein interaction partners – HKs and RRs




	a 1	a 2	a 3	a 4	a 5	a 6		b 1	b 2	b 3	b 4	b 5	b 6	b 7
<i>E. c</i>	K	D	I	L	V	D		T	I	F	D	F	C	L
<i>S. a</i>	K	D	I	L	V	D		T	I	F	D	P	C	I
<i>T. t</i>	D	D	I	K	V	D		T	I	F	H	P	C	L
<i>S. e</i>	D	E	I	L	V	H		S	I	F	H	P	C	I



Protein interaction partners – HKs and RRs




Use the model built from training data to score each possible pairing, within each species.




Accurately predict protein interaction partners – HKs and RRs




Protein interaction partners – HKs and RRs




Protein interaction partners – HKs and RRs




Prediction starting with no known pairings for HK-RRs



Prediction starting with no known pairings for **other protein pairs**




We can also predict whether two proteins interact or not!



How to extract useful information from protein sequences?

Each protein is constructed as a specific chain or string of amino acids.

In a specific protein of interest, such as hemoglobin, the order of amino acids is highly important and contains all the information necessary to produce the folded, functional molecule.

1	2	3	4	5	6	7	8
M	V	L	S	P	A	D	K
							

We would like to find a probability model for the sequence of amino acids that corresponds to each protein of interest.

$$P(A_1, \dots, A_L) = \text{Probability that a sequence produces a folded, functional hemoglobin molecule.}$$

What information is stored in sequences?

Sequences record the outcome of **millions of evolutionary experiments** which are **constrained** by the requirements of protein structure and function.

Ligand binding

Structure

Oligomerization

Enzyme Kinetics


Interaction specificity


Indirect constraints from RNA function

-- how are these encoded?


-- can we crack the code and make predictions?

Can we **invert sequence data** to reveal these constraints and so understand which sequence residues control phenotype?





Sectors – groups of residues that control distinct functional properties




Halabi *et al* performed SCA correlated mutation analysis on an MSA of 1470 serine protease family sequences.

The **eigenvectors** with largest eigenvalues are **localized on three groups of residues, defining 'sectors'**.


Mutational analysis found these **sectors each control distinct functional properties** of the proteases.

Serine Proteases – Sector Identification




This work introduced the idea of using spectral analysis to find coevolving groups of residues.


The localization of those eigenvectors with largest eigenvalues was used to identify sectors



Serine Proteases – Sector Identification




This work introduced the idea of using spectral analysis to find coevolving groups of residues.



The localization of those eigenvectors with largest eigenvalues was used to identify sectors

Serine Proteases – Experimental test of protein sectors



Outline

Biology 101 – a quick tour of proteins


Biological information from covariance

- Protein structure and function
- Protein interaction specificity
- Ligand protein binding


Theoretical analysis of data limitations

- when can we predict reliably?
- Noise due to finite sampling, also samples are not independent.
- central role played by Random Matrix Theory.


Prediction accuracy depends on the amount of input data available



Prediction accuracy depends on the input data available



14663 sequences



6724 sequences

Prediction accuracy depends on the input data available


We compute an empirical covariance matrix from a finite sample of data points.

There are two major problems:

1. There will always be noise due to finite sampling.
2. Protein sequences are related by evolution, so they are not independent samples.

A natural framework to address the first issue is provided by Random Matrix Theory, and we will show how it can be extended to address the second.

Natural approach: Identify common causes (factors) that explain the joint dynamics of p variables




Consider a data matrix X which contains N samples of p variables.

Which variables interact?

Sample covariance matrix:

$$C = \frac{XX^T}{N}$$

Natural approach: Identify common causes (factors) that explain the joint dynamics of p variables



Consider a data matrix X which contains N samples of p variables.

Which variables interact?

Sample covariance matrix:

$$C = \frac{X X^T}{N}$$

But the data are highly undersampled – **due to finite sampling** there will be **spurious large entries**.

Problem:

These data contain multiple covariance signals – signal, batch effects, biased sampling.... that we need to de-convolve and understand. They are all obfuscated by noise due to finite sampling.

Noise due to finite sampling

Consider a set X of $n=300$ samples, each of length $p=500$ variables, drawn independently.

True covariance matrix

$$\begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}$$

Random matrix
theory



Eigenvalues of sample covariance matrix

No signal – all variables independent.

Noise due to finite sampling

Consider a set X of $n=300$ samples, each of length $p=500$ variables, drawn independently.

True covariance matrix


$$\begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}$$

Random matrix
theory



No signal – all variables independent.

Eigenvalues of sample covariance matrix



This spread depends on how well sampled the data is

Noise due to finite sampling

Consider a set X of $n=300$ samples, each of length $p=500$ variables, drawn independently.


True covariance matrix

$$\begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}$$

Random matrix
theory



Eigenvalues of sample covariance matrix



Random matrix theory:

The **Marcenko-Pastur law** describes the eigenvalue distribution of the sample covariance matrix of Gaussian IID noise.

This spread depends on how well sampled the data is

Outline

Biology 101 – a quick tour of proteins


Biological information from covariance

- Protein structure and function
- Protein interaction specificity
- Ligand protein binding

Theoretical analysis of data limitations


- when can we predict reliably?
- Noise due to finite sampling, also samples are not independent.
- central role played by Random Matrix Theory.

Example: protein-ligand binding




*Crystal Structure ADRB2: Rasmussen SG, DeVree BT, Zou Y, Kruse AC, Chung KY, Kobilka TS, Thian FS, Chae PS, Pardon E, Calinski D, Mathiesen JM, Shah ST, Lyons JA, Caffrey M, Gellman SH, Steyaert J, Skinotis G, Weis WI, Sunahara RK, Kobilka BK. Nature. 2011 Jul 19.


Example: protein-ligand binding




carazolol (ligand –
beta blocker)




Many known ligands




Timolol




Nadalol



Penbutolol Sulfate
(beta blocker)



Albuterol



Propanolol Hcl
(beta blocker)


HCl

What other
ligands will bind?

*Crystal Structure ADRB2: Rasmussen SG, DeVree BT, Zou Y, Kruse AC, Chung KY, Kobilka TS, Thian FS, Chae PS, Pardon E, Calinski D, Mathiesen JM, Shah ST, Lyons JA, Caffrey M, Gellman SH, Steyaert J, Skinotis G, Weis WI, Sunahara RK, Kobilka BK. Nature. 2011 Jul 19.

Use molecular descriptors - fingerprints


- Representation of a molecule as a string of bits ('0' and '1')
- Usually the 2D structure, i.e. the molecular graph is used to generate the string.
- A specific chemical moiety gives rise to a pattern of bits




e.g.Morgan 3

[illegible]

Linear, e.g. Daylight



Circular, e.g. Morgan 3, ECFP



What information is contained in the covariance matrix of ligand fingerprints?

Consider 1000 random molecules from ChEMBL

Random matrix theory and covariance

Consider a set X of $n=300$ samples, each of length $p=500$ variables, drawn independently.

True covariance matrix

$$\begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}$$


Random matrix
theory




Random matrix theory:

The **Marcenko-Pastur law** describes the eigenvalue distribution of the sample covariance matrix of Gaussian IID noise.


Eigenvalues of sample covariance matrix






These molecules are independent samples and follow the Marcenko-Pastur distribution


$$f_{MP}(\lambda) = \frac{\sqrt{(\lambda_+ - \lambda)(\lambda - \lambda_-)}}{2\pi\lambda}$$




carazolol (ligand –
beta blocker)




Many known ligands




Timolol




Nadolol



Penbutolol Sulfate
(beta blocker)



Albuterol




Propanolol HCl
(beta blocker)


HCl

What about the covariance matrix of ligands that all bind the same protein receptor?

Compare ligands that bind ADRB1 with random molecules



Compare ligands that bind ADRB1 with random molecules




Different truncations yield different models:

$$\mathbf{u}_p = \sum_{i=1}^m (\mathbf{v}_i \cdot \mathbf{u}) \mathbf{v}_i$$


Here we truncate the spectrum to the space spanned by the largest m modes, and ask whether a novel ligand $\mathbf{u}$ lies in this subspace.

$$\|\mathbf{u} - \mathbf{u}_p\| < \epsilon$$


This provides a model that predicts whether novel ligands will bind to the ADRB1 receptor or not.





Test the performance of different truncations




Marchenko-Pastur threshold



Another example – adenosine receptor




Signal modes that describe
binding to AA2AR?




Attribution - what is the model learning?


Examine largest two modes of the covariance matrix - infer the structural motif that an eigenvector represents by the common substructure among those ligands that lie closest to that eigenvector:



First eigenvector




Second eigenvector



SUMMARY

We showed that removing the noise present in this dataset due to finite sampling greatly improved the ability of our model to make accurate predictions.





But what about the effects of finite sampling in protein sequence data?

Marchenko-Pastur threshold

Before, I argued that transitive correlations obscure the signal

Wild Type Protein



$$P(A_1, \dots, A_L) = \frac{1}{Z} \exp \left\{ - \sum_i h_i A_i - \sum_{(i,j)} e_{ij} (A_i, A_j) \right\}$$


	K	D	E	R
K	X	I	I	C
D	I	X	C	I
E	I	C	X	C
R	C	I	C	X

naïve mean field approximation, assume small couplings

$$e_{ij} (A_i, A_j) = (C^{-1})_{ij} (A_i, A_j)$$

Is this really true?

The residue covariance matrix: $C_R = X^T * X$



Simulation using a contact map – here for DHFR

What else could cause the signal to get obscured?

What else could cause the signal to be obscured?



This data contains ‘nuisance’ correlations because the samples are related to each other, they are not iid. We call this relatedness structure the ‘phylogeny’

What else could cause the signal to be obscured?



This data contains ‘nuisance’ correlations because the samples are related to each other, they are not iid. We call this relatedness structure the ‘phylogeny’


The sequence covariance matrix – $C_S = X * X^T$ used to compute phylogeny

The residue covariance matrix – $C_R = X^T * X$ constraints from structure/function


These matrices contain the same information – in particular they have the same eigenvalues (up to the number of zero eigenvalues).

These two sources of covariance mix and confound each other.

Repeat the simulation, but now include phylogeny



Without phylogeny – contact prediction is quite accurate just using $C_R = X^T * X$




With phylogeny – contact prediction is much less accurate using $C_R = X^T * X$

We need to know what covariance is caused by the phylogeny

Consider a perfect 'tree phylogeny'.

A founder sequence of **length p residues** are independently evolved for **m mutation steps**.

Sequences are duplicated and the process repeats through **b such branching events** to yield n sequences.




We need to know what covariance is caused by the phylogeny

Consider a perfect 'tree phylogeny'.

A founder sequence of length p residues are independently evolved for m mutation steps.

Sequences are duplicated and the process repeats through b such branching events to yield n sequences.



Sequence covariance matrix is highly structured..... like batch effects.

What effect does this have on the eigenvalues?


We need to know what covariance is caused by the phylogeny

We define:

$$\alpha(m) = \exp \left[-\frac{qm}{(q-1)p} \right] \quad \begin{array}{l} m \text{ mutations per branch, sequences of length } p, \\ \text{alphabet of size } q. \text{ Start with } q = 2. \end{array}$$

The elements of the true covariance matrix are given by: $(\Sigma_S)_{ij} = \alpha^{b_{ij}}$ where b_{ij} is the number of branching events between sequences i and j .

Consider a single branching event:



We need to know what covariance is caused by the phylogeny


We define:

$$\alpha(m) = \exp \left[-\frac{qm}{(q-1)p} \right]$$

m mutations per branch, sequences of length p ,
alphabet of size q . Start with $q = 2$.

The elements of the true covariance matrix are given by: $(\Sigma_S)_{ij} = \alpha^{b_{ij}}$ where b_{ij} is the number of branching events between sequences i and j .


Consider a single branching event:


$$\Sigma_S = \begin{pmatrix} 1 & \alpha \\ \alpha & 1 \end{pmatrix}$$
$$\lambda_{\pm} = 1 \pm \alpha$$

What is the spectrum of the sample covariance matrix?

Calculating the effects of finite sampling

Recall, for a single branching event:



$$\Sigma_S = \begin{pmatrix} 1 & \alpha \\ \alpha & 1 \end{pmatrix}$$


$$\lambda_{\pm} = 1 \pm \alpha$$

To compute the expected eigenvalue distribution of the sample covariance matrix, we use the Stieltjes transform:

$$G(z, c) = \int_{-\infty}^{+\infty} \frac{dF(\lambda)}{z - \lambda}$$

Calculating the effects of finite sampling

Recall, for a single branching event:



$$\Sigma_S = \begin{pmatrix} 1 & \alpha \\ \alpha & 1 \end{pmatrix}$$

$$\lambda_{\pm} = 1 \pm \alpha$$

To compute the expected eigenvalue distribution of the sample covariance matrix, we use the Stieltjes transform:


$$G(z, c) = \int_{-\infty}^{+\infty} \frac{dF(\lambda)}{z - \lambda}$$

In 1967 Marcenko and Pastur used a pde-based approach to write the Stieltjes transform in terms of the eigenvalue distribution of the true covariance matrix $T(\lambda)$

$$G(z, c) = -1 \left(z - c \int_{-\infty}^{\infty} \frac{\lambda dT(\lambda)}{1 + \lambda G(z, c)} \right)^{-1}$$

Calculating the effects of finite sampling

Recall, for a single branching event:



$$\Sigma_S = \begin{pmatrix} 1 & \alpha \\ \alpha & 1 \end{pmatrix}$$

$$\lambda_{\pm} = 1 \pm \alpha$$

To compute the expected eigenvalue distribution of the sample covariance matrix, we use the Stieltjes transform:



$$G(z, c) = \int_{-\infty}^{+\infty} \frac{dF(\lambda)}{z - \lambda}$$

In 1967 Marcenko and Pastur used a pde-based approach to write the Stieltjes transform in terms of the eigenvalue distribution of the true covariance matrix $T(\lambda)$. In this case, the eigenvalues are discrete so the inverse Stieltjes transform, which describes the finite sampling effects, reduces to finding the positive imaginary root of this polynomial:



$$z - \frac{c}{2} \left(\frac{1 + \alpha}{1 + (1 + \alpha)G} \right) - \frac{c}{2} \left(\frac{1 - \alpha}{1 + (1 - \alpha)G} \right) = -\frac{1}{G}$$


Red lines computed using this approach:

Simulations of the simplest trees produce non-MP spectra




Consider sequences of length $p = 300$ in a model with a single duplication event, with branch length m .





What about deeper trees?


Move to deeper trees:



$$\Sigma_S = \begin{pmatrix} 1 & \alpha & \alpha^2 & \alpha^2 \\ \alpha & 1 & \alpha^2 & \alpha^2 \\ \alpha^2 & \alpha^2 & 1 & \alpha \\ \alpha^2 & \alpha^2 & \alpha & 1 \end{pmatrix}$$

What about deeper trees?

Move to deeper trees:




Eigenvalues of the true
covariance matrix
(more generally):

$$\lambda = \begin{cases} 1 + \sum_{i=1}^{j-1} 2^{i-1} \alpha^i - 2^{j-1} \alpha^j & \forall j < b \\ 1 + \sum_{i=1}^b 2^{i-1} \alpha^i & j = b \end{cases}$$


where: $P(\lambda = \lambda_i) = p_i = 2^{-i}$

This results in a power law, if r is the rank then: $\lambda \sim r^{-\beta}$

$$\beta = \log 2\alpha / \log 2$$



“Real trees” are Inhomogeneous



Mathematics is more complicated.



But we can prove that if the branch lengths are drawn from a distribution with mean m and bounded higher moments, then the empirical spectral distribution is exactly the same as for a homogeneous tree with branch length m .


Key parameters:

b: number of branching events


m: average mutations per branch

Theory also works for more realistic trees











What about simulations with phylogeny and interactions?




Simulations with phylogeny and interactions follow the power law




But interactions alone have a very different signature






This suggests a strategy to 'clean' the phylogenetic corruption from the sequence covariance matrices..... what about real proteins???

What about real sequence data – can we use this?



DHFR




We have the covariance matrix here ~8k sequences. Calculate eigenvalues, fit the power law and extract the parameters:


$$\lambda \sim r^{-\beta}$$

$$\beta = \log 2\alpha / \log 2$$

$$\alpha(m) = \exp \left[-\frac{qm}{(q-1)p} \right]$$



DHFR





We have the covariance matrix here ~8k sequences. Calculate eigenvalues, fit the power law and extract the parameters:

$$\lambda \sim r^{-\beta}$$


$$\beta = \log 2\alpha / \log 2$$

$$\alpha(m) = \exp \left[-\frac{qm}{(q-1)p} \right]$$






DHFR




Calculated from the covariance matrix of ~8k sequences.

Can we use this understanding to remove the covariance caused by phylogeny?


$$S_p = \mathbf{v}_1 \mathbf{v}_1^T + \cdots + \mathbf{v}_t \mathbf{v}_t^T, \quad \lambda_1 \leq \cdots \leq \lambda_t \leq \cdots \leq \lambda_{p(q-1)}$$




$$S_p = \mathbf{v}_1 \mathbf{v}_1^T + \cdots + \mathbf{v}_t \mathbf{v}_t^T, \quad \lambda_1 \leq \cdots \leq \lambda_t \leq \cdots \leq \lambda_{p(q-1)}$$




Trypsin






Calculated from the covariance matrix of ~30k sequences. Fit the power law and extract the parameters:





Trypsin




Calculated from the covariance matrix of ~30k sequences.


Using Random Matrix Theory to truncate the largest modes recovers contact prediction from raw covariance.


$$S_p = \mathbf{v}_1 \mathbf{v}_1^T + \cdots + \mathbf{v}_t \mathbf{v}_t^T, \quad \lambda_1 \leq \cdots \leq \lambda_t \leq \cdots \leq \lambda_{p(q-1)}$$


Phylogenetic reweighting



RMT-based linear transformation







SUMMARY

1. Covariance is a very powerful tool for uncovering biological information from large data sets:
 - Protein structure and function
 - Which proteins interact with each other
 - Ligand binding specificity
2. A critically important problem is to understand the limits of these methods, how much data is needed for accurate prediction and so forth.
 - Random Matrix Theory provides a natural framework for solving this problem.