

# White Box : Website Frontend & Network visualization using Guided Backpropagation


Neha Das

Sumit Dugar

Technische Universität München

Fakultät für Informatik

München, 12. April 2018



# Outline

- Goals and Motivation
- Web Interface
  - Proposed System
  - Technical Details
  - Results
- Guided Backpropagation
  - Theoretical Background
  - In Context of Protein Distance Prediction
  - Results and Observations
- Summarization
- Future Work

# Goals and Motivation

- Web interface that accepts a protein sequence (primary structure) and predicts the distances between each pair of the sequence (tertiary structure) using a Deep Neural Net

## **Motivation**

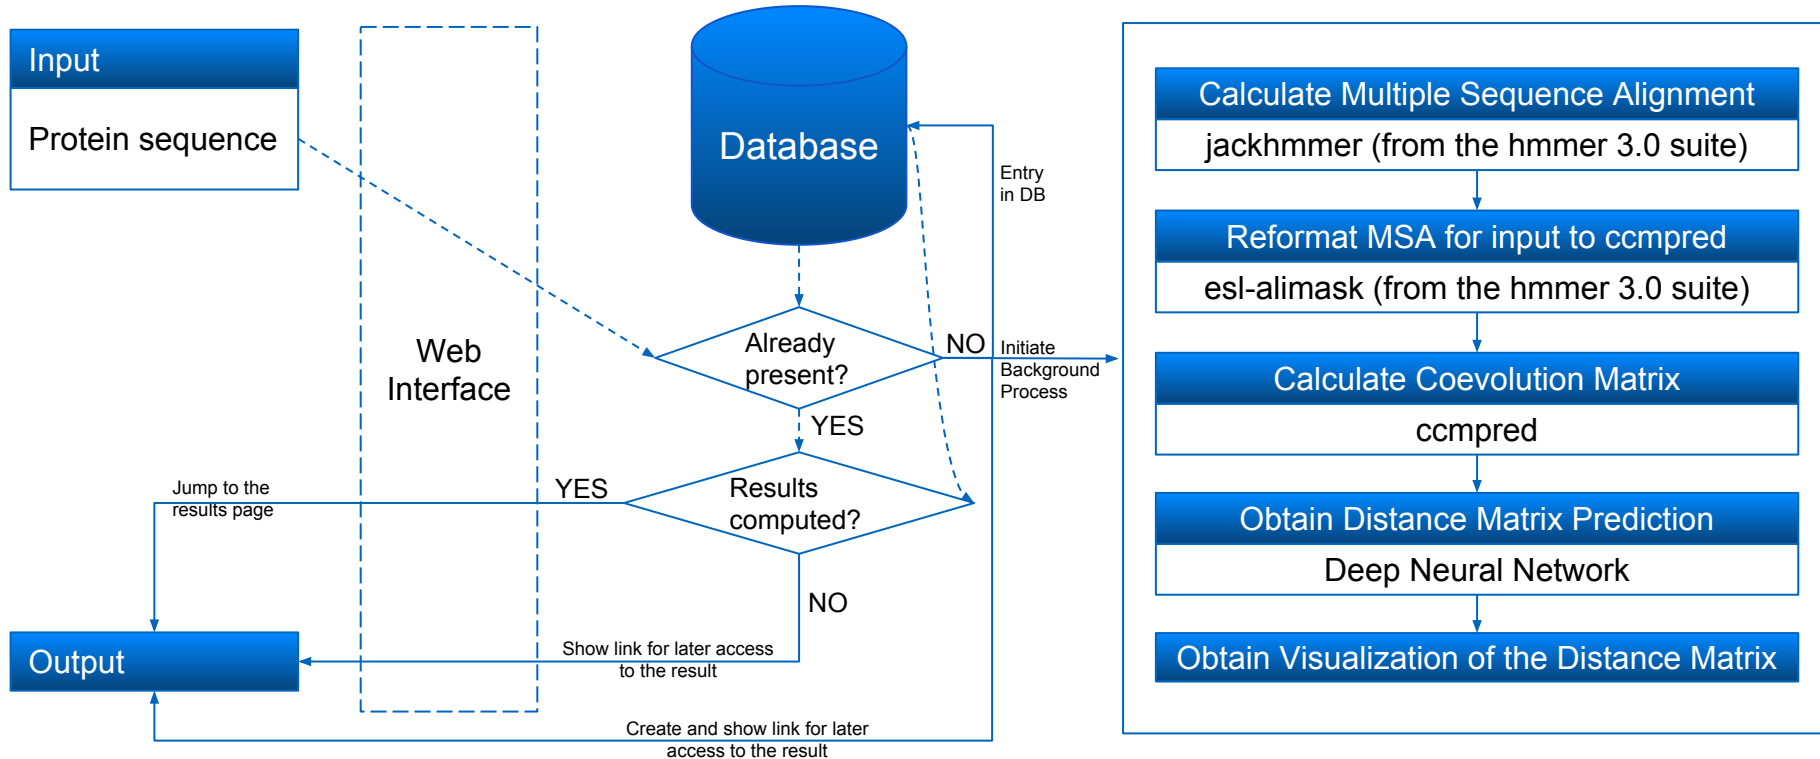
- Need for an open and simple interface to predict protein structure
  - Single pipeline that abstracts the intermediate steps between the input & output
- Visualization of the DNN using guided backpropagation

## **Motivation**

- To understand the intuition behind the predictions from the Deep Neural Net

# Web Interface

# Web Interface Pipeline



# Technology Used

- Framework - Flask 0.12.0
- Database - SQLite
- For the pipeline
  - Compute MSAs - JackHMMer
    - Computation ran against the Uniref-100 DB
  - Compute Co-evolution Matrices - ccmpred
    - Output a coevolution matrix of  $L \times L \times 21 \times 21$ .
  - Compute Distance Matrix predictions - Deep Neural Network
    - Courtesy of Matthias Baur and Omer Dolev, build on Pytorch
    - Derived from the NIPS paper by Vladimir Golkov et.al
    - Uses 3 Convolutional Layers and Dropouts in-between with ReLu non-linearities at each layer
    - Total receptive field -  $15 \times 15$

# Results - New protein sequence submitted for prediction

## Protein Contact Prediction

```
VHLTPEEKSAVTALWGKLVNDEVGGEALGRLLVYYPWTQRFESFGDLSTPDAVMGNPKVKAHGKKVLGAFSDGLAHL DNLKGT FATLSELHCDK  
LHVDPENFRLLGNLVCVLAH HFGKEFTPPVQAAYQKVVAGVANALAHKYH
```

Predict

Example Sequence

We are computing your results. Use the following url for sharing and checking your results after some time.

View Result

[http://localhost:8080/view/7391e18de9c603fbe34f8:](http://localhost:8080/view/7391e18de9c603fbe34f8)

# Results - Prediction result page with download link

[Back to Prediction](#)

## Protein Contact Prediction

### Values for the Predicted Distances


```

[[[[14.784979 13.48749 12.582673 12.35956 12.731022
13.575727 14.689992 15.893319 17.059916 18.246597
19.22058 20.033127 20.619339 21.071295 21.556519
21.965202 22.108734
 22.305288 22.351646 22.444418 22.495945 22.624033
22.71798 22.794687 22.936514 23.071259 23.40327
23.614328 23.763374 23.90229 23.918566 23.995262
23.972977 23.847668
 23.84865 23.92982 23.993912 24.186884 24.294575
24.420267 24.432446 24.372395 24.231634 23.952837
23.553322 23.18491 22.875723 22.598354 22.518726
22.58587 22.7896
 23.150423 23.551117 23.90246 24.248568 24.42588
24.552217 24.52873 24.446562 24.251257 24.06741
23.931532 23.729584 23.694914 23.619432 23.554947
23.664295 23.781809
 23.901031 23.888157 23.890314 23.591152 23.222052
22.773478 22.31707 21.78776 21.371721 21.125128
21.075558 21.351625 21.755829 22.15758 22.54271
22.944006 23.242756
 23.33889 23.344349 23.25695 23.142015 23.038727
22.903646 22.826412 22.783583 22.896877 22.964088
23.20372 23.369822 23.542053 23.717321 23.852446
23.891457 23.868319
 23.755865 23.713825 23.67262 23.715204 23.776411
23.877964 23.966188 24.166292 24.289675 24.412405
24.453985 24.486525 24.425549 24.327272 24.229849
24.13049 23.856766
 23.612608 23.279634 22.931135 22.616386 22.425797
22.209623 22.144041 22.034185 21.946014 21.900238
21.975492 22.058119 22.186363 22.52303 22.828335


```

### Visualization of Predicted Distances


Distance between C $\alpha$  & C $\alpha$  atoms




Distance between C $\alpha$  & C $\beta$  atoms



Distance between C $\beta$  & C $\alpha$  atoms



Distance between C $\beta$  & C $\beta$  atoms



Download Result

<http://localhost:8080/download/7391e18c>



# Results - Same protein sequence submitted again

## Protein Contact Prediction

```
VHLTPEEKSAVTALWGKVNVDVEVGGELGRLLVYYPWTQRRFFESFGDLSTPDAVMGNPKVKAHGKKVLGAFSDGLAHLAHDNLKGTFFATLSELHCDK  
LHVDPENFRLLGNVLCVLAHHFGKEFTPPVQAAYQKVVAGVANALAHKYH
```

Predict

Example Sequence

Computation already in progress. Use the following url for sharing and checking your results after some time.

View Result

[http://localhost:8080/view/7391e18de9c603fbe34f8:](http://localhost:8080/view/7391e18de9c603fbe34f8)


# Guided Backpropagation

# Introduction - Basic Approach

- Aim : To visualize the parts of the input (I) that affect the output (O)
- Basic approach : Visualize  $\delta O / \delta I$ 
  - The magnitude of gradient in a portion of the visualization is proportional to the degree of influence the corresponding input region exerts over the output



Input Image

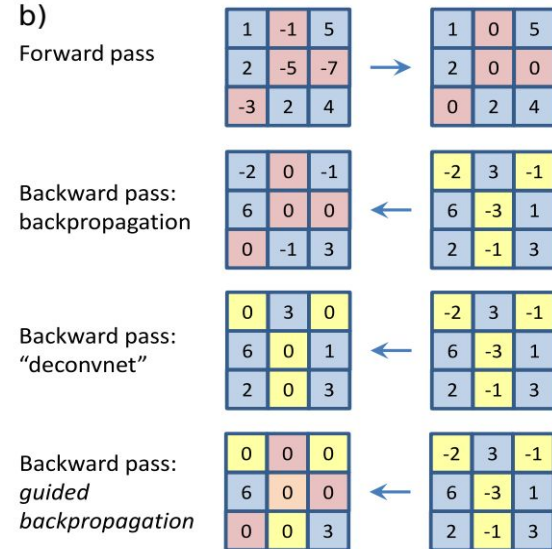



Plain gradient

Fig1: We consider a classification network that for examples takes the image of a snake and outputs the class snake.

# Introduction - Variants of the Basic Approach

- Variants of the basic approach - dealing with ReLu in different ways:
  - Basic Approach - Guidance only from the Input
  - Backward Deconvnet - Guidance only from the Output
  - Guided Backpropagation - Guidance from both the Input and Output



c) activation:  $f_i^{l+1} = \text{relu}(f_i^l) = \max(f_i^l, 0)$

backpropagation:  $R_i^l = (f_i^l > 0) \cdot R_i^{l+1}$ , where  $R_i^{l+1} = \frac{\partial f_{out}}{\partial f_i^{l+1}}$


backward 'deconvnet':  $R_i^l = (R_i^{l+1} > 0) \cdot R_i^{l+1}$

guided backpropagation:  $R_i^l = (f_i^l > 0) \cdot (R_i^{l+1} > 0) \cdot R_i^{l+1}$


# Introduction - Examples of Visualization




Input Image



Plain Gradients



Backward Deconvnet



Guided Backpropagation

# Method - In Context of Protein Distance Prediction

- The deep learning network for protein distance prediction has an :
  - Input (I) size of  $(441, L, L)$ , where the input is a coevolution matrix and L is the sequence length
  - Output (O) size of  $(4, L, L)$ , where the first dimension stands for the 4 connection types - alpha-alpha, alpha-beta, beta-alpha and beta-beta
- Thus, the size of the gradient (or Jacobian)  $\delta O / \delta I$ , in this case comes out to be  $(4 \times L \times L, 441 \times L \times L)$
- We can however reduce the gradient size to  $(4 \times L \times L, 441 \times 15 \times 15)$ , since the final receptive window size of our convolutional neural network is  $(15, 15)$
- We visualize these gradients using  $4 \times L \times L$  plots containing 441  $(21 \times 21)$  subplots of size  $(15 \times 15)$ . Each subplot, thus shows the influence of inputs on the distance predicted between a certain position pair of amino acids for one of the four channels

# Results - Visualization Example

- Neural Network
  - CNN with 33% precision on non locals
- Input Sequence
  - ID: 1ABA\_A
  - Sequence Length: 87
  - Coevolution size: (87,87,21,21)
- Guided Gradients (Jacobian):
  - Size: (87,87,21,21,15,15) - beta-beta
- Selected Output Position: We will be showing the visualizations for the output at (20,32) of the predicted matrix since it has a contact and is a true-positive for our network




Fig1: Subset from the full Jacobian, contains gradients of the output at (20,32) wrt input coevolution for amino acid pairs (D,R), (D,N), (C,R) and (C,N)\*




Fig2: Corresponding subset from the coevolution matrix focusing on the (15,15) receptive window around the position (20,32)\*




Fig3: Ground Truth Distance visualization for the beta-beta channel




Fig4: Predicted Distance visualization for the beta-beta channel

\* We only plot the area in the effective receptive field - (15,15) around the output position (20,32), as only those are taken into consideration by the current network architecture




Fig1: Subset from the full Jacobian, contains gradients of the output at (20,32) wrt input coevolution for 21x21 amino acid pairs



# Observations and Inferences

## Correlation of gradients and co-evolutions:

- We observed mostly positive linear correlation between gradients and coevolution values in Fig. 1
- Another related observation can be seen in Fig. 2 where we weigh the gradients with the sign of the corresponding coevolution value

## Inference:

The areas with positive correlation may indicate that increasing the coevolution values there, will also increase the predicted distance (We recheck the interpretation of the gradient in Fig. 3). This however, was in contrast to our expectations since a high coevolutional value is usually indicative of a contact.

## Conclusion:

Correlation values may be misleading as they are indicative of a linear relationship, whereas in this case it is quite non-linear.




Fig1: Correlation values between coevolutions and gradients over window (15x15) for all the amino acid pairs (dim 1&2) for output positions (20:32,20:32) (dim 3&4)




Fig2: Gradients weighted by sign of corresponding coevolution values. The amount of + and - values in the weighted gradient are nearly equal (47% and 51% resp.)






Fig3: These are the plots for predicted distance values at position (20,32) against coevolution change steps. In this experiment we increased the coevolution value by a fraction of the gradient values (indicated by coeff) to check the effect on output when the input moves in the direction of the gradient. As evident from the above, the output was shown to increase

# Observations and Inferences

## Receptive Field:

We observed that the gradient magnitudes are strongest in the center of the image and gradually move to zero as we move towards the edges of the receptive field (see Fig. 1). Notice that the coloration is almost circular.

## Inference:

This can be interpreted as evidence in support of the receptive window size chosen as (15,15).

If the actual window size had been larger (i.e influence from neighbors farther away), then we would have expected to see stronger gradients at the edges of our current receptive window too.




Fig1: Notice the circular pattern of the gradient coloration and how it fades out as we move to the window edges.

# Observations and Inferences

## Patterns:


We observed that a certain spatial pattern, linked to the identity of amino acid pairs, tended to appear in our gradient visualization (see Figures on the right). The gradient pattern lightened or darkened (or inverted) according to the coevolution input, but was otherwise unaffected

## Inference:


This behavior may indicate that the network has strong notions about particular pairs of amino acids and steers the output according to this notion, paying less attention to the actual coevolution values there.

## Conclusion


This is not an expected behavior. Further analysis is required to understand this. Investigation of this behavior for a network with higher precision may also lead to some results.



All input coevolutions are highly negative  $\sim -0.3$



All input coevolutions are highly positive  $\sim +0.3$



Original input coevolutions

Fig 1,2,3: Gradients for the position (20,32)

# Summarization

In our work, therefore, we have

- Created and presented a web interface for predicting the distances between various amino acids in a protein, given a protein sequence.
- Implemented guided backpropagation to visualize the inner workings of the Deep Neural Network (an intermediate step for the above) that predicts Protein Distance Matrix, given its coevolution matrix.

# Future Work

# Future Work

- Integrate some tool on the website to show the visualizations from guided backpropagation.
- Comparison with other visualization techniques.
- Host the website for public use.
- Further analysis of the visualization results.

# References and Source Code

# References

- Golkov, V. et al. (2016). Protein contact prediction from amino acid co-evolution using convolutional networks for graph-valued images. In Advances in Neural Information Processing Systems (pp. 4222-4230).
- Springenberg, J. et al. (2014). Striving for simplicity: The all convolutional net. arXiv preprint arXiv:1412.6806.
- Ozublak, U. (2017). utkuozbulak/pytorch-cnn-visualizations. [online] Available at: <https://github.com/utkuozbulak/pytorch-cnn-visualizations> [Accessed 2018].: The All Convolutional Net. CoRR, abs/1412.6806,
- Golkov, V. (2016). show & scroll - Visualize arbitrary N-dimensional arrays - File Exchange - MATLAB Central. [online] De.mathworks.com. Available at: <https://de.mathworks.com/matlabcentral/fileexchange/52374-show---scroll-visualize-arbitrary-n-dimensional-arrays> [Accessed 2018].


# Source Code

- <https://gitlab.lrz.de/dugarsumit/dlcvbm>
- [https://gitlab.lrz.de/ge72xax/dl4cv\\_prak](https://gitlab.lrz.de/ge72xax/dl4cv_prak)

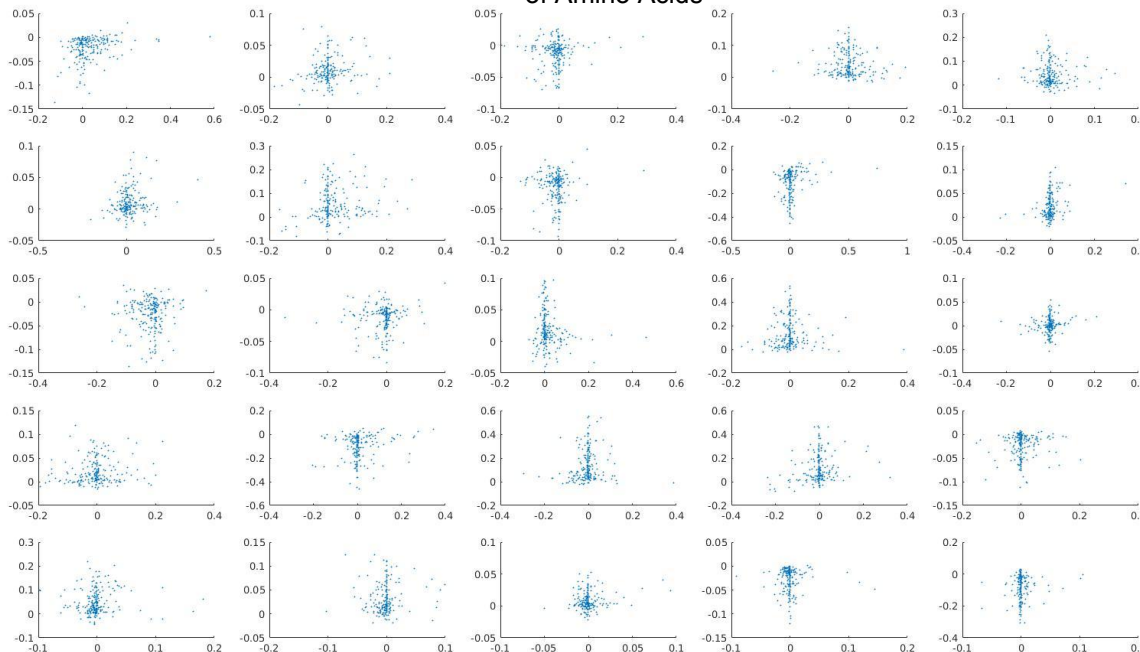


# Additional Results\*

Coevolution (X-Axis) vs Gradients (Y-Axis) for Position 20,32




Coevolution (X-Axis) vs Gradients (Y-Axis) for Position 20,32 individually for the first 25 pairs of Amino Acids



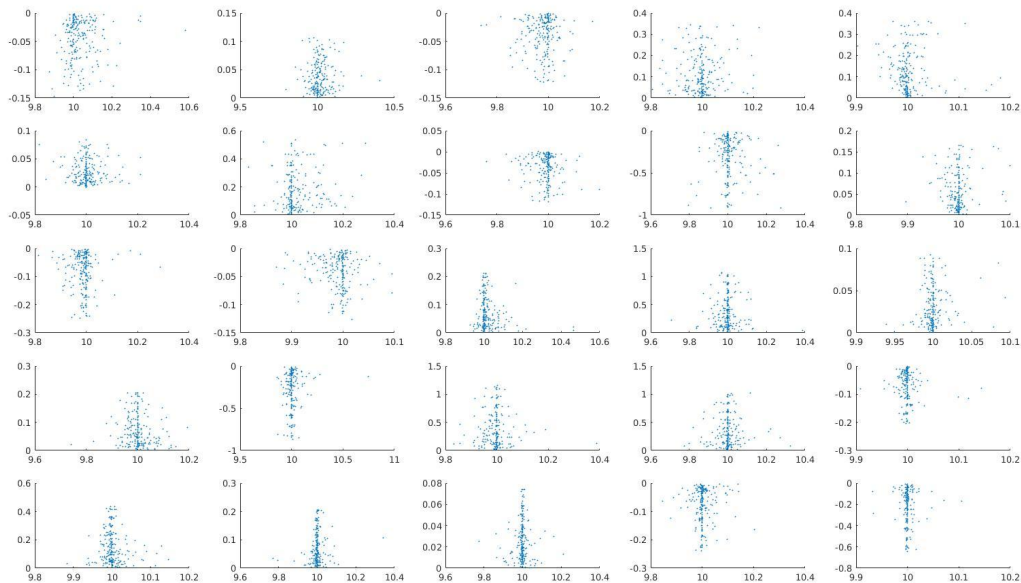
The scatter plots for the rest of the amino acid pairs are as scattered as these ones, indicating that there is no direct correlation between coevolution and gradients. We also plotted the results for the first 25 amino acids for highly negative (we subtracted a large value from) coevolutions and highly positive (we added a large value to) coevolutions. We again, did not notice any significant pattern in those. You can see these plots on the following slides

# Additional Results\*

Highly positive Coevolution (Added 10 to every value) (X-Axis) vs Gradients (Y-Axis) for Position 20,32




Highly positive Coevolution (Added 10 to every value) (X-Axis) vs Gradients (Y-Axis) for Position 20,32 individually for the first 25 pairs of Amino Acids



# Additional Results\*

Highly negative Coevolution (Subtracted 10 from every value) (X-Axis) vs Gradients (Y-Axis) for Position 20,32



Highly negative Coevolution (Subtracted 10 from every value) (X-Axis) vs Gradients (Y-Axis) for Position 20,32 individually for the first 25 pairs of Amino Acids

