
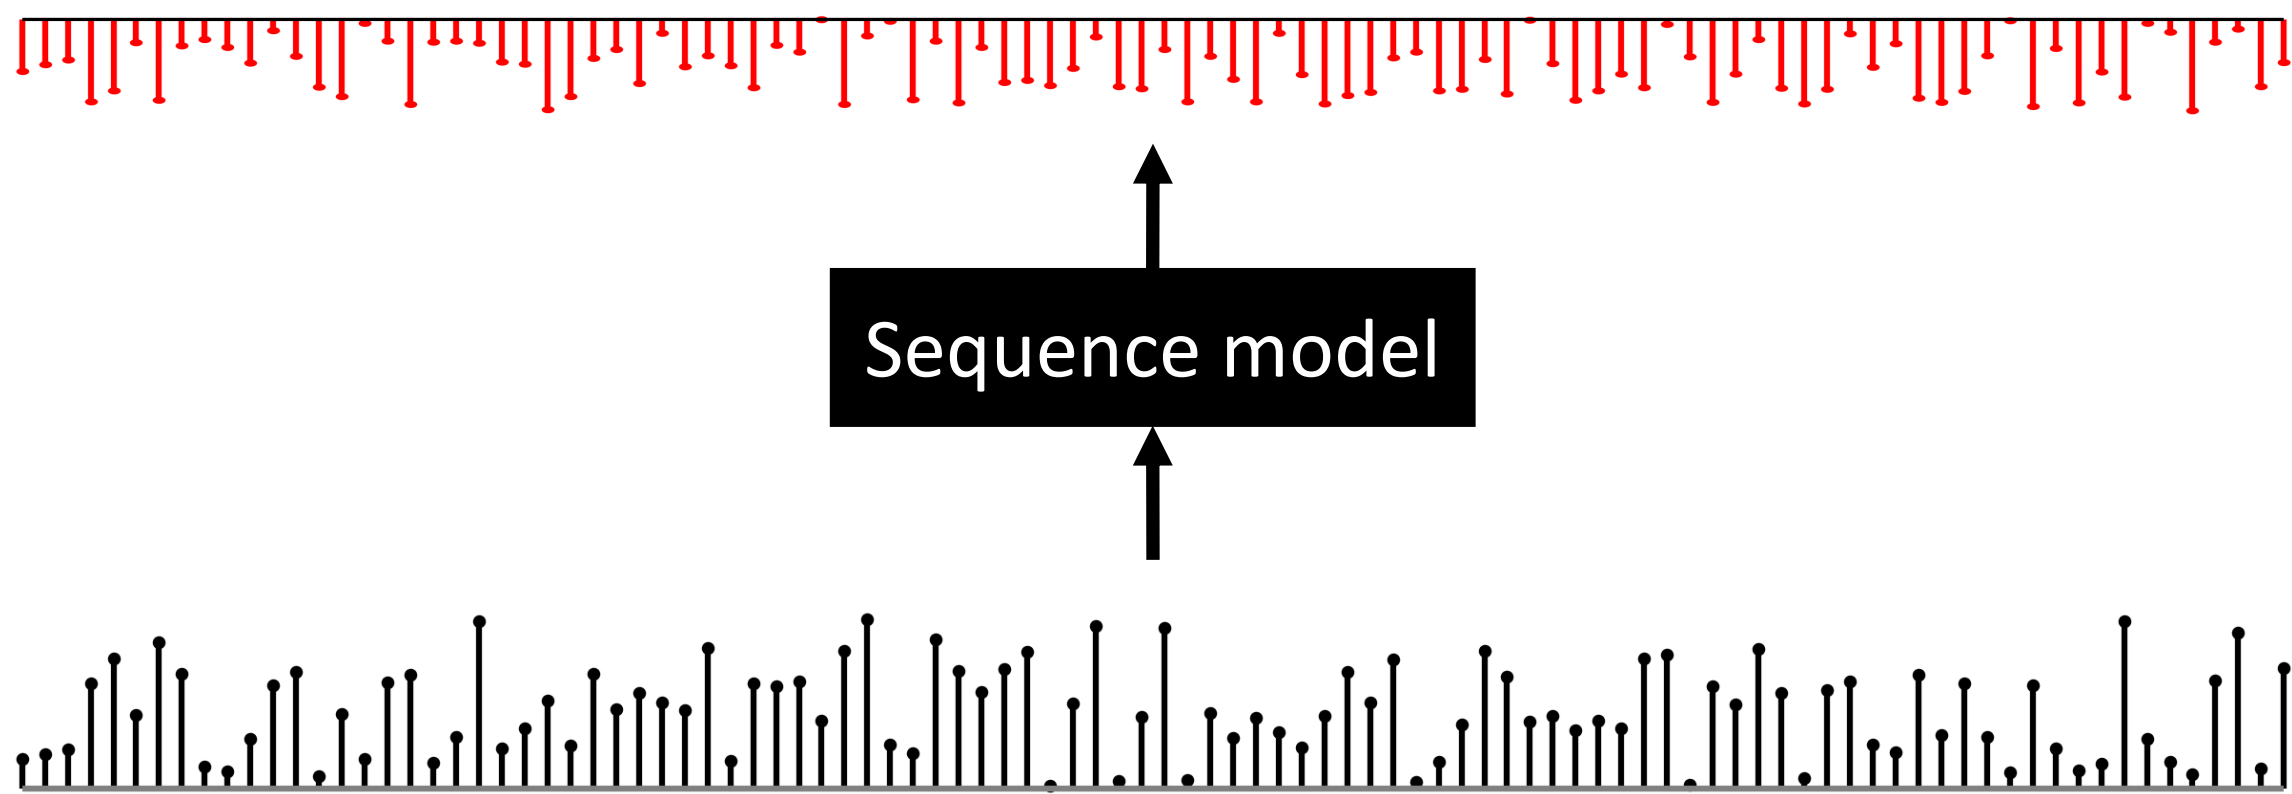


# Predicting splice sites in DNA sequences with sequence models

2024.05.15

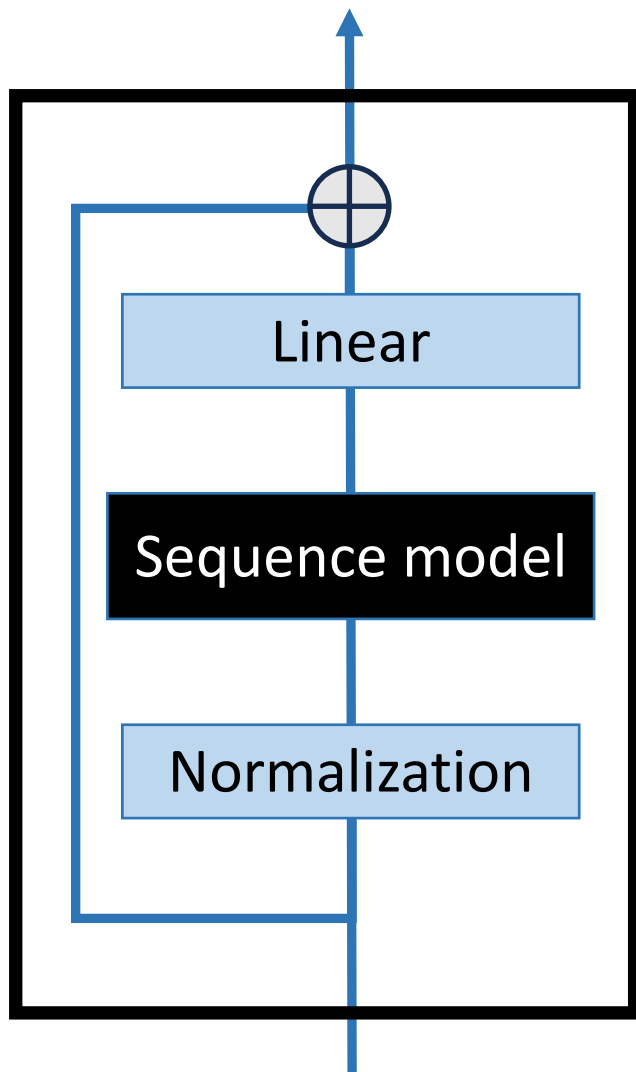
Kuan-Hao Chao





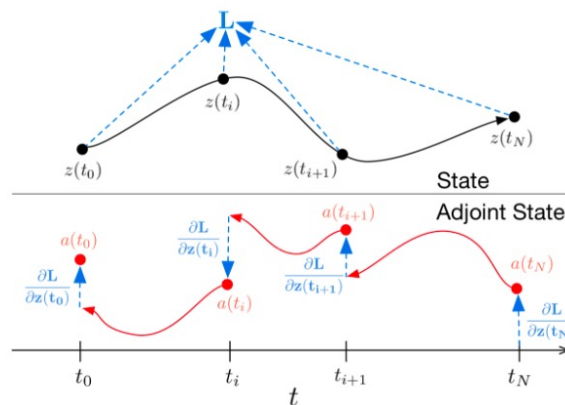
Sequence models map a sequence to a sequence

(batch, length, dim)




(batch, length, dim)

## Neural ODEs



## RNN



# Future work

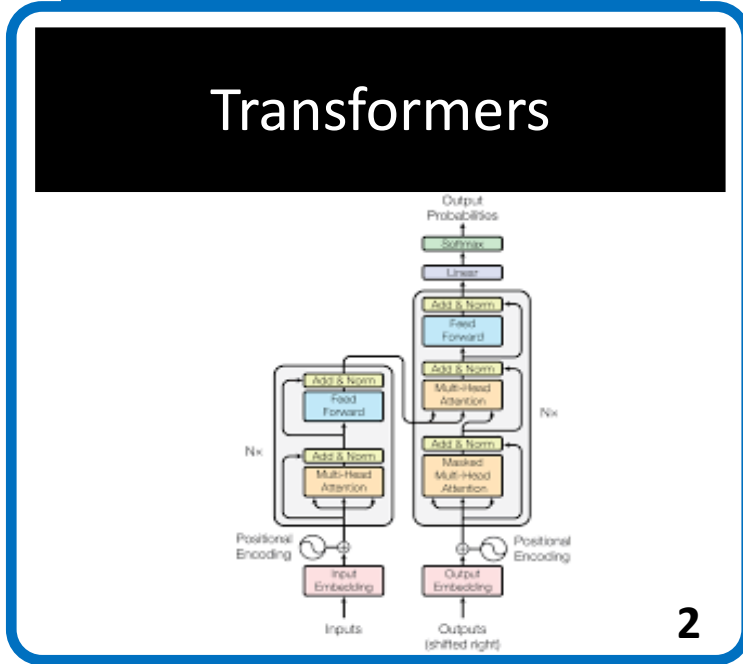

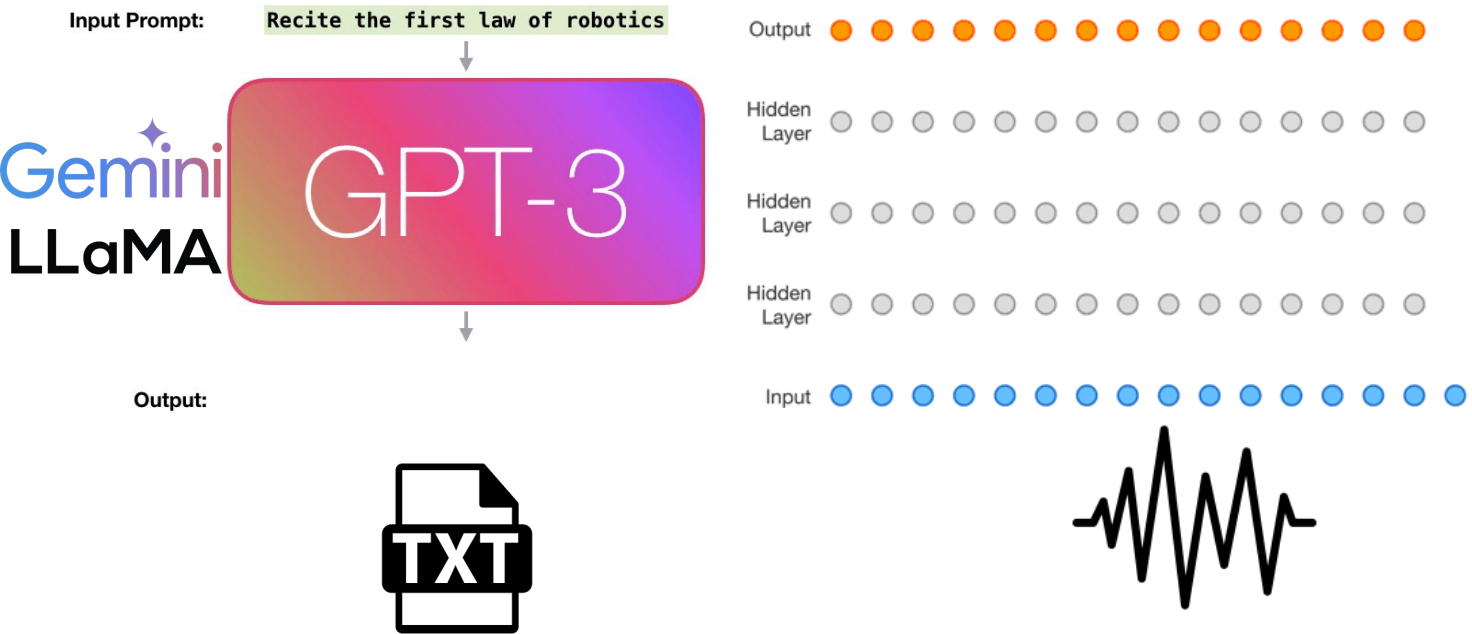
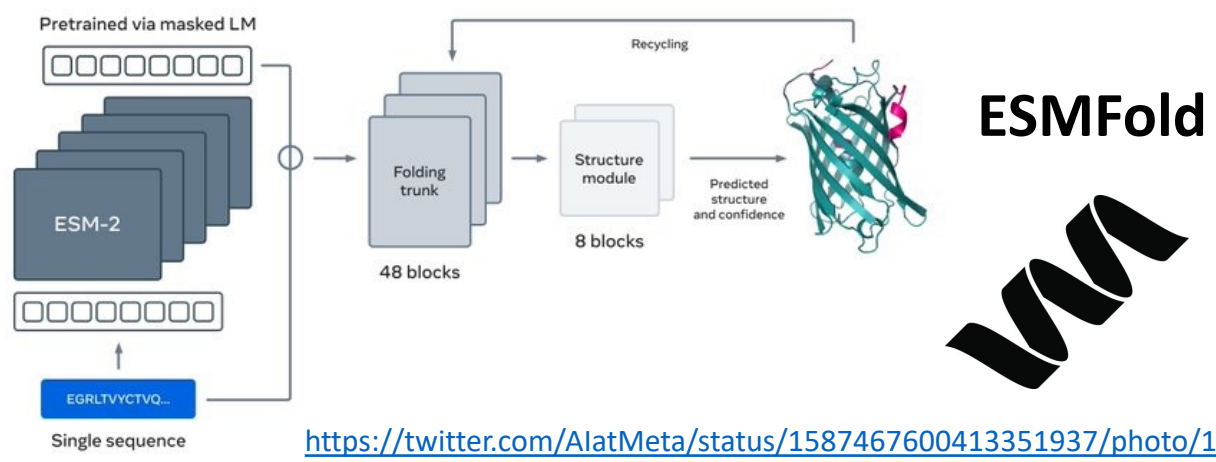


Figure 1: The Transformer - model architecture.

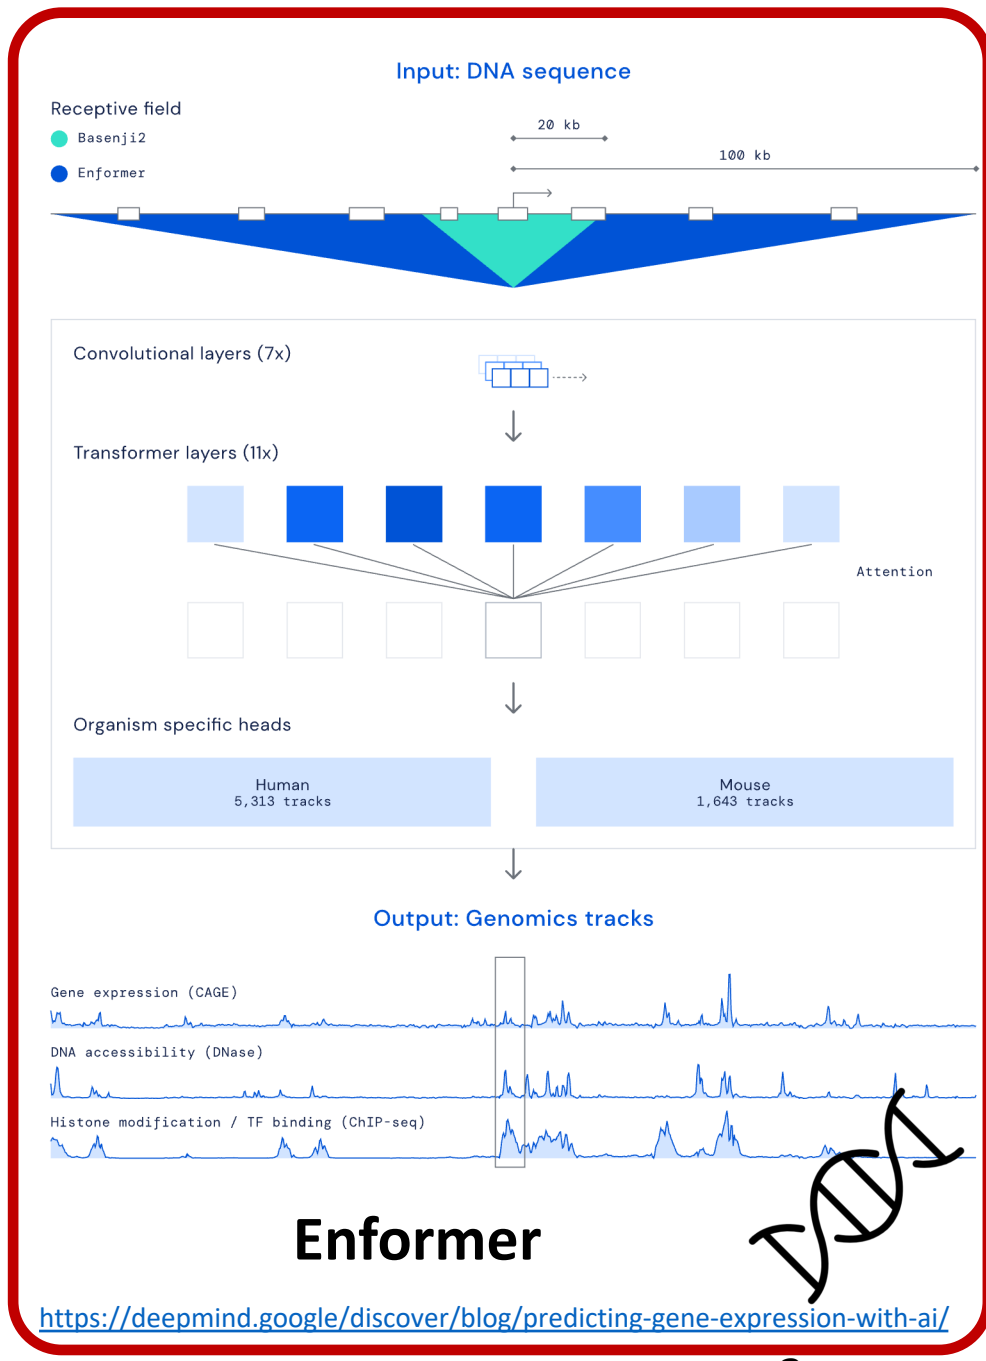


<https://jalammar.github.io/how-gpt3-works-visualizations-animations/>

<https://deepmind.google/discover/blog/wavenet-a-generative-model-for-raw-audio/>




<https://twitter.com/AlatMeta/status/1587467600413351937/photo/1>



<https://deepmind.google/discover/blog/predicting-gene-expression-with-ai/>

# Spectrum of Sequential Data



# Why Convolutional Neural Network to DNA ?

**nature methods**

View all journals Search

Explore content

**Trojanskaya Lab Princeton**

nature > nature

Brief Communication | Published: 24 August 2015

**Predicting effects of noncoding DNA- and RNA-binding protein binding using deep learning-based sequence models**

Jian Zhou & Olga G Trojanskaya

**nature biotechnology**

Explore content About the journal Publish with

nature > nature biotechnology > analyses > articles

Analysis | Published: 27 July 2015

**Predicting the sequence specificity of DNA- and RNA-binding proteins using deep learning**

Babak Alipanahi, Andrew Delong, Matthew T Weirauch

**Bioinformatics**

Article Navigation

JOURNAL ARTICLE

**Convolutional neural networks for predicting protein binding sites**

Haoyang Zeng, Matthew D. Edwards

Trojanskaya Lab Princeton

DeepSEA  
2015

FUToronto

DeepBind  
2015

Gifford LabMIT

DNA-TF binding  
2016

**GENOME RESEARCH**

HOME | ABOUT | ARCHIVE | SUBMIT | SUBSCRIBE | ADVERTISE | AUTHOR

Institution: MILTON S EISENHOWERLIBRARY Sign in

**Calico**

**Basset: learning the regulatory logic of the accessible genome with deep convolutional neural networks**

David R. Kelley<sup>1</sup>, Jasper Snoek<sup>2</sup> and John L. Rinn<sup>1</sup>

**GENOME RESEARCH**

HOME | ABOUT | ARCHIVE | SUBMIT | SUBSCRIBE | ADVERTISE | AUTHOR

Institution: MILTON S EISENHOWERLIBRARY Sign in

**Calico**

**Sequential regulatory activity across chromosomes with deep convolutional neural networks**

David R. Kelley<sup>1</sup>, Yakir A. Reshef<sup>2</sup>, Maxwell Bileschi<sup>3</sup>, Cory Y. McLean<sup>3</sup> and Jasper Snoek<sup>3</sup>

**nature biotechnology**

Explore content About the journal Publish with

nature > nature biotechnology > analyses > articles

Letter | Published: 24 September 2018

**A universal SNP deep neural network for predicting protein binding**

Ryan Poplin, Pi-Chuan Chang, Newburger, Jolo Diamco, Nathaniel & Mark A DePristo

Calico

Basset  
2016

Calico

Basenji  
2018

Google Health

DeepVariant  
2018

**nature genetics**

Explore content

nature > nature genetics

Article | Published: 16 July 2018

**Deep learning sequence-based prediction of variant effects on expression and splicing**

Jian Zhou, Chandra L. Theesfeld, Kevin Yao, Kathleen M. C. McElwee

**Cell**

Volume 176, Issue 3, 24 January 2019, Pages 535-548.e24

Article

**Predicting Splicing from Protein Binding with Deep Learning**

Kishore Jaganathan<sup>1,6</sup>, Sofia Kyriazopoulou Panagiotopoulou<sup>2</sup>, Siavosh Fazel Darbandi<sup>2</sup>, David Knowles<sup>3</sup>, Yang L. Li<sup>1,3</sup>, Jian Wenwu Cui<sup>1</sup>, Grace B. Schwartz<sup>2</sup>, Eric D. Chow<sup>4</sup>, Efsthia K. Karamantoni<sup>1</sup>, Serafim Batzoglou<sup>1</sup>, Stephan J. Sanders<sup>2</sup>, Kyle Kai-How Fung<sup>1,5</sup>

**nature methods**

Explore content

nature > nature methods

Article

**Predicting 3D genome folding from Hi-C data with Akita**

Geoff Fudenberg<sup>1,5,6,8</sup>, David R. Kelley<sup>2,5,6,8</sup> and Katherine S. Pollard<sup>1,5,6,8</sup>

Trojanskaya Lab Princeton

ExPecto  
2018

Illumina

SpliceAI  
2019

Calico

Akita  
2020

**ARTICLES**

https://doi.org/10.1038/s41592-021-01252-x

**OPEN**

**Effective gene expression prediction by integrating sequence and chromatin data**

Žiga Avsec<sup>1,2,3</sup>, Vikram Agarwal<sup>2,4</sup>, Daniel Viser<sup>1,2,3</sup>, Agnieszka Grabska-Barwinska<sup>1</sup>, Kyle R. Taylor<sup>1,2,3</sup>, and David R. Kelley<sup>2,5,6</sup>

Agarwal and Kelley Genome Biology (2022) 23:245  
https://doi.org/10.1186/s13059-022-02811-x

**RESEARCH**

**The genetic and biochemical determinants of mRNA degradation rates in mammals**

Vikram Agarwal<sup>1,2\*</sup> and David R. Kelley<sup>1\*</sup>

**Calico**

**Predicting RNA-seq coverage from DNA sequence with a unifying model of gene regulation**

Johannes Linder  
Calico Life Sciences LLC  
jlinder@calicolabs.com

Divyanshi Srivastava  
Calico Life Sciences LLC  
divyanshi@calicolabs.com

Vikram Agarwal  
mRNA Center of Excellence, Sanofi Pasteur Inc.  
Vikram.Agarwal@sanofi.com

DeepMind + Calico

Enformer  
2021


Calico

Saluki  
2022

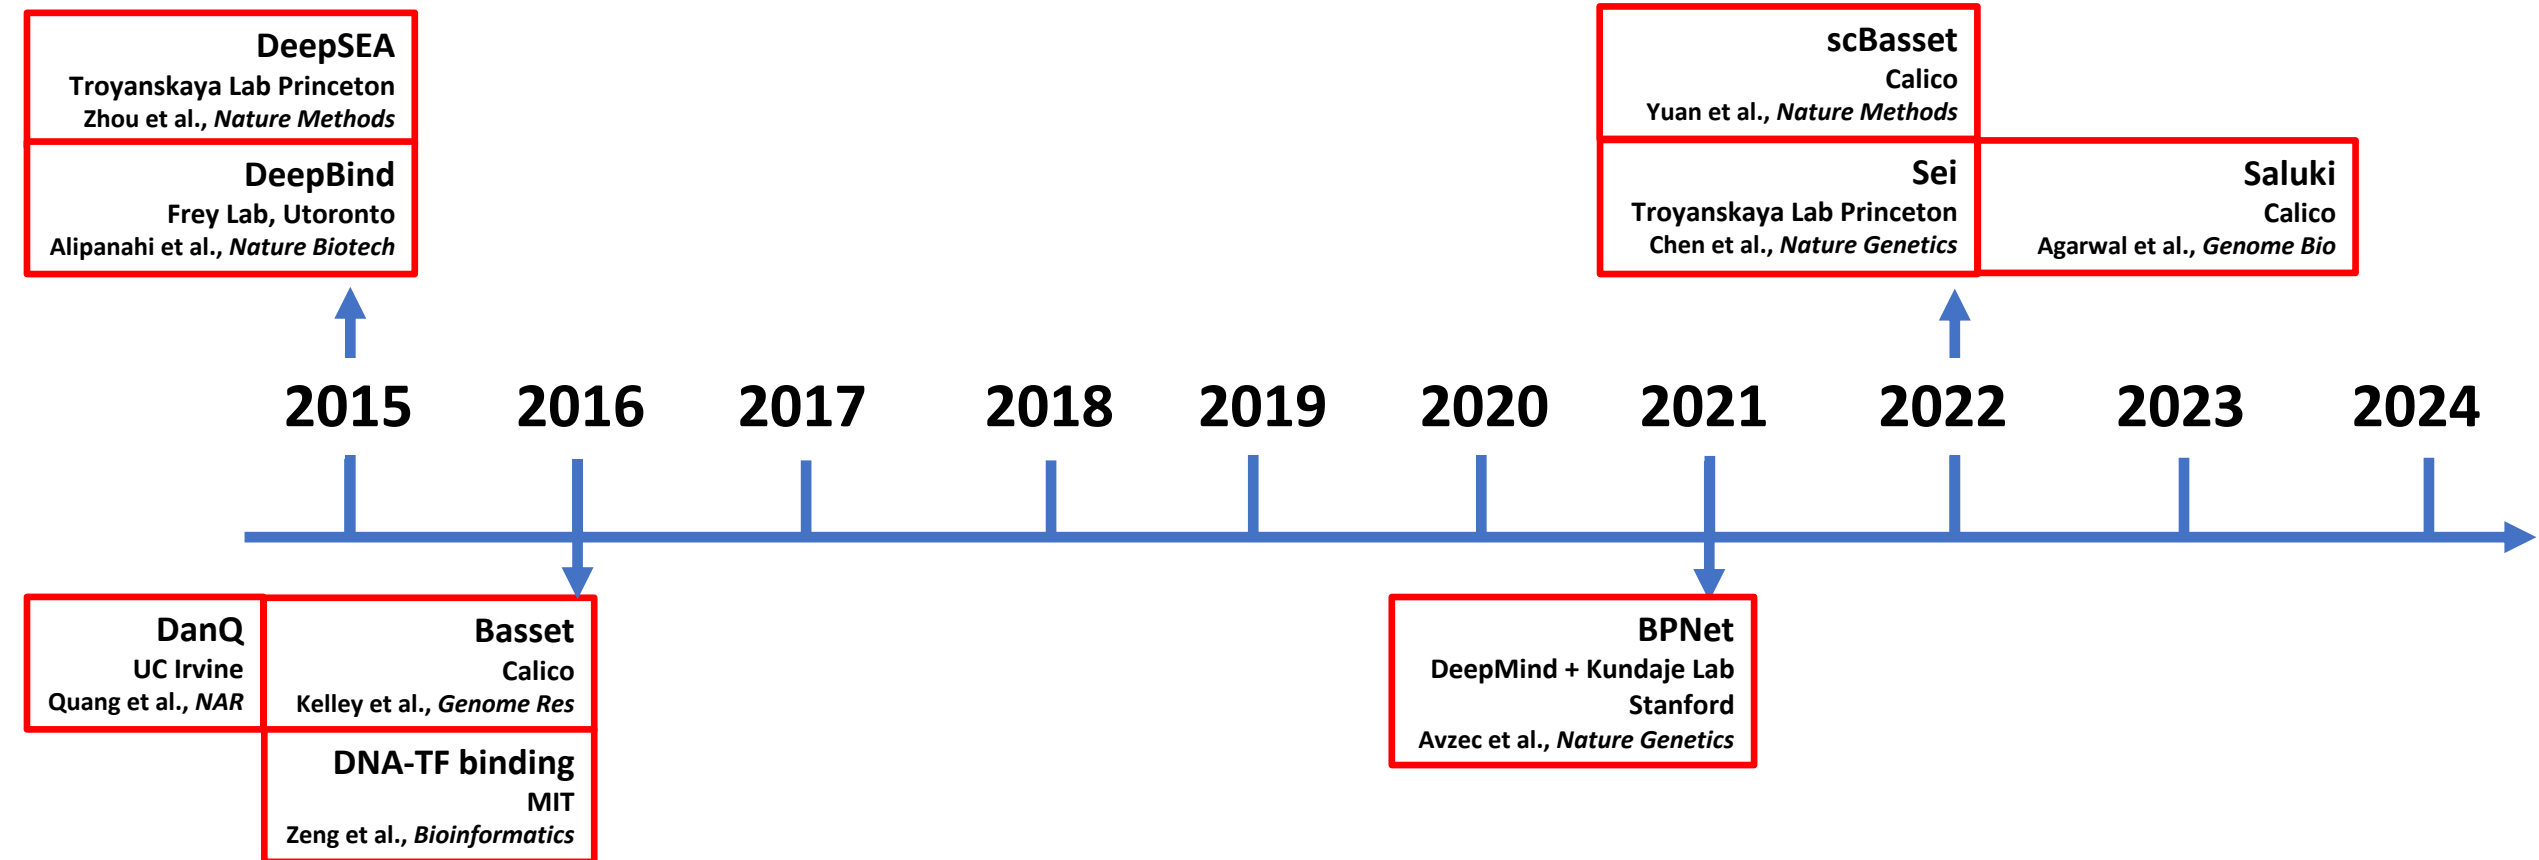
Calico

Borzoï  
2023

# Why Convolutional Neural Network to DNA ?

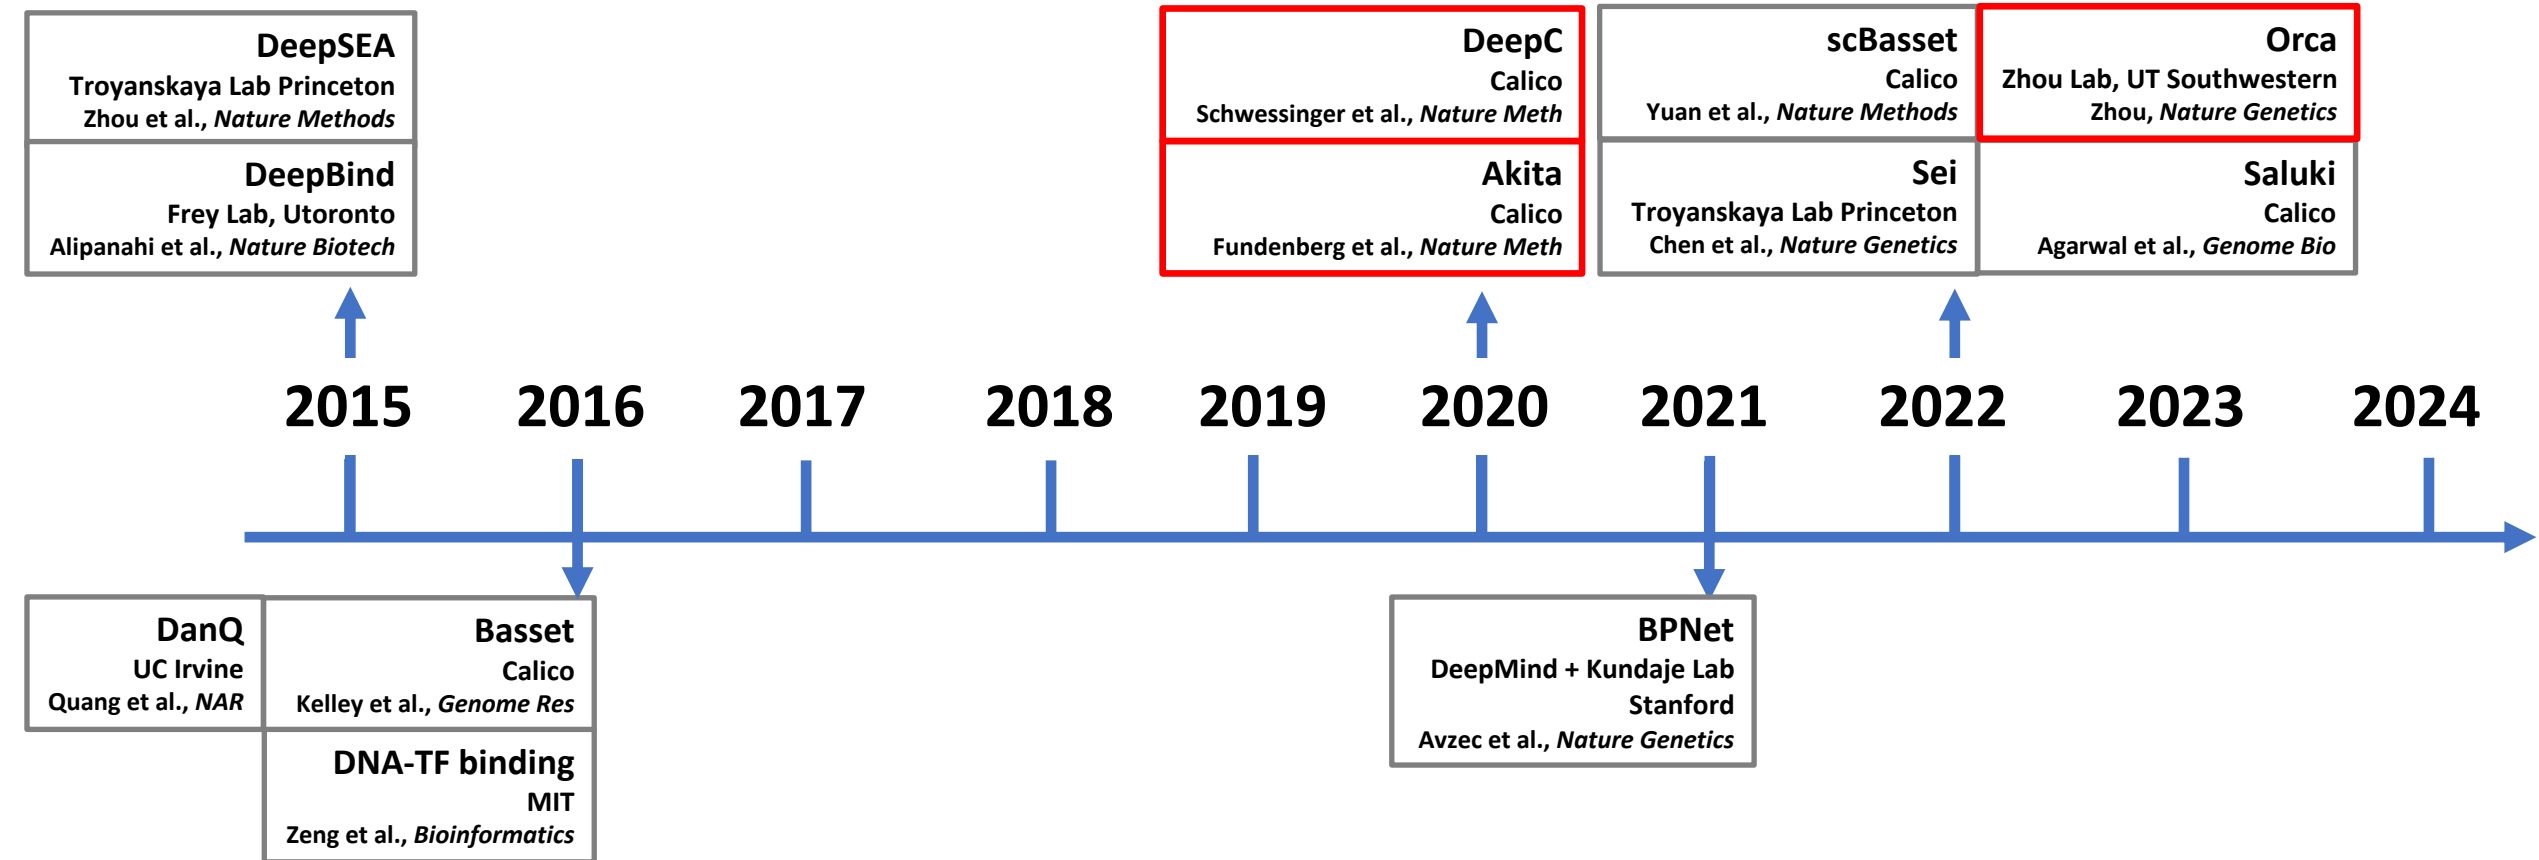


# Why Convolutional Neural Network to DNA ?

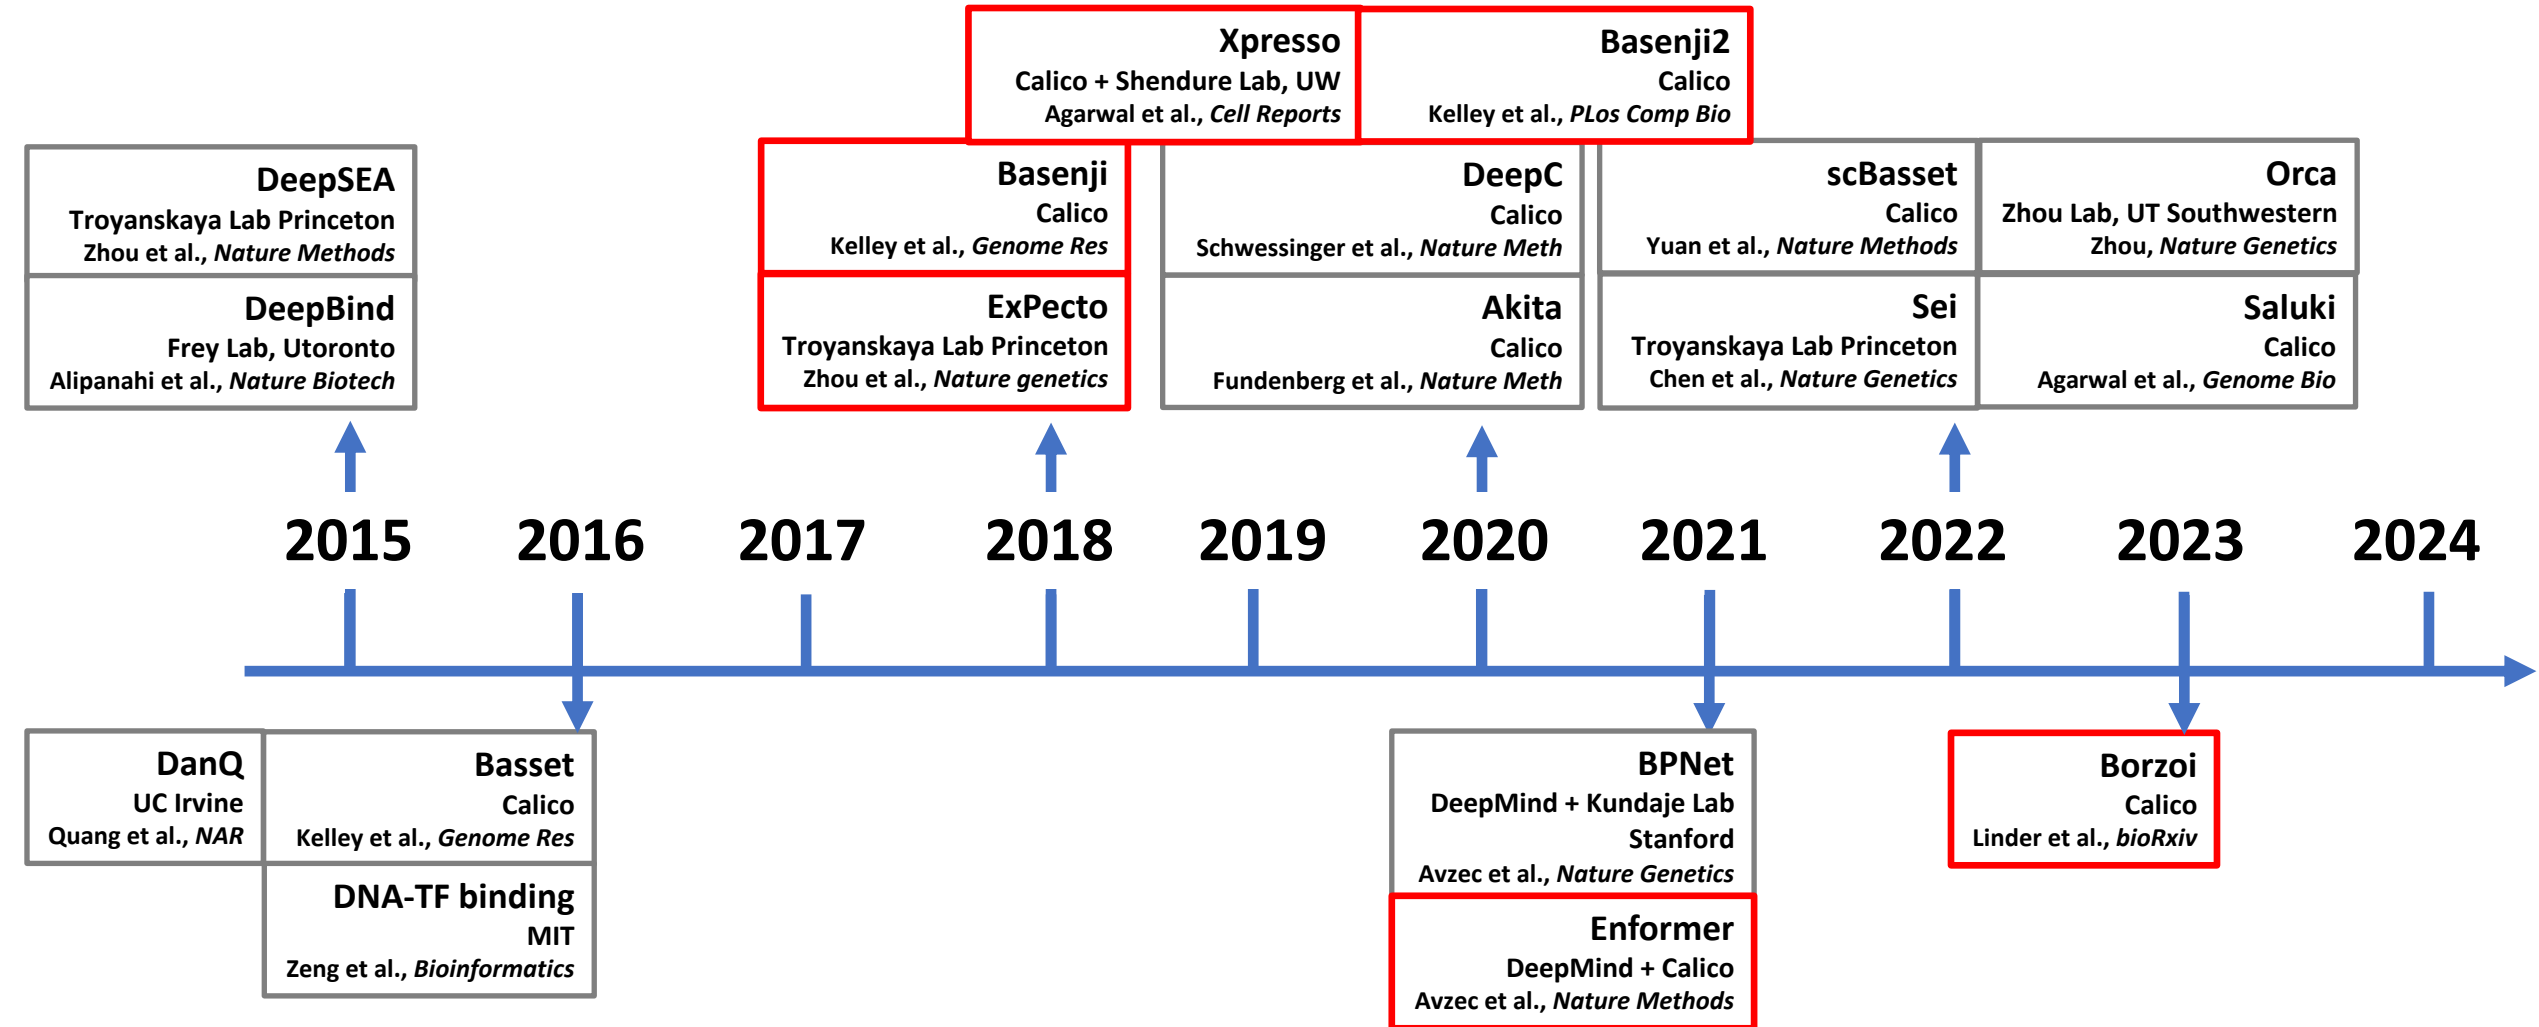




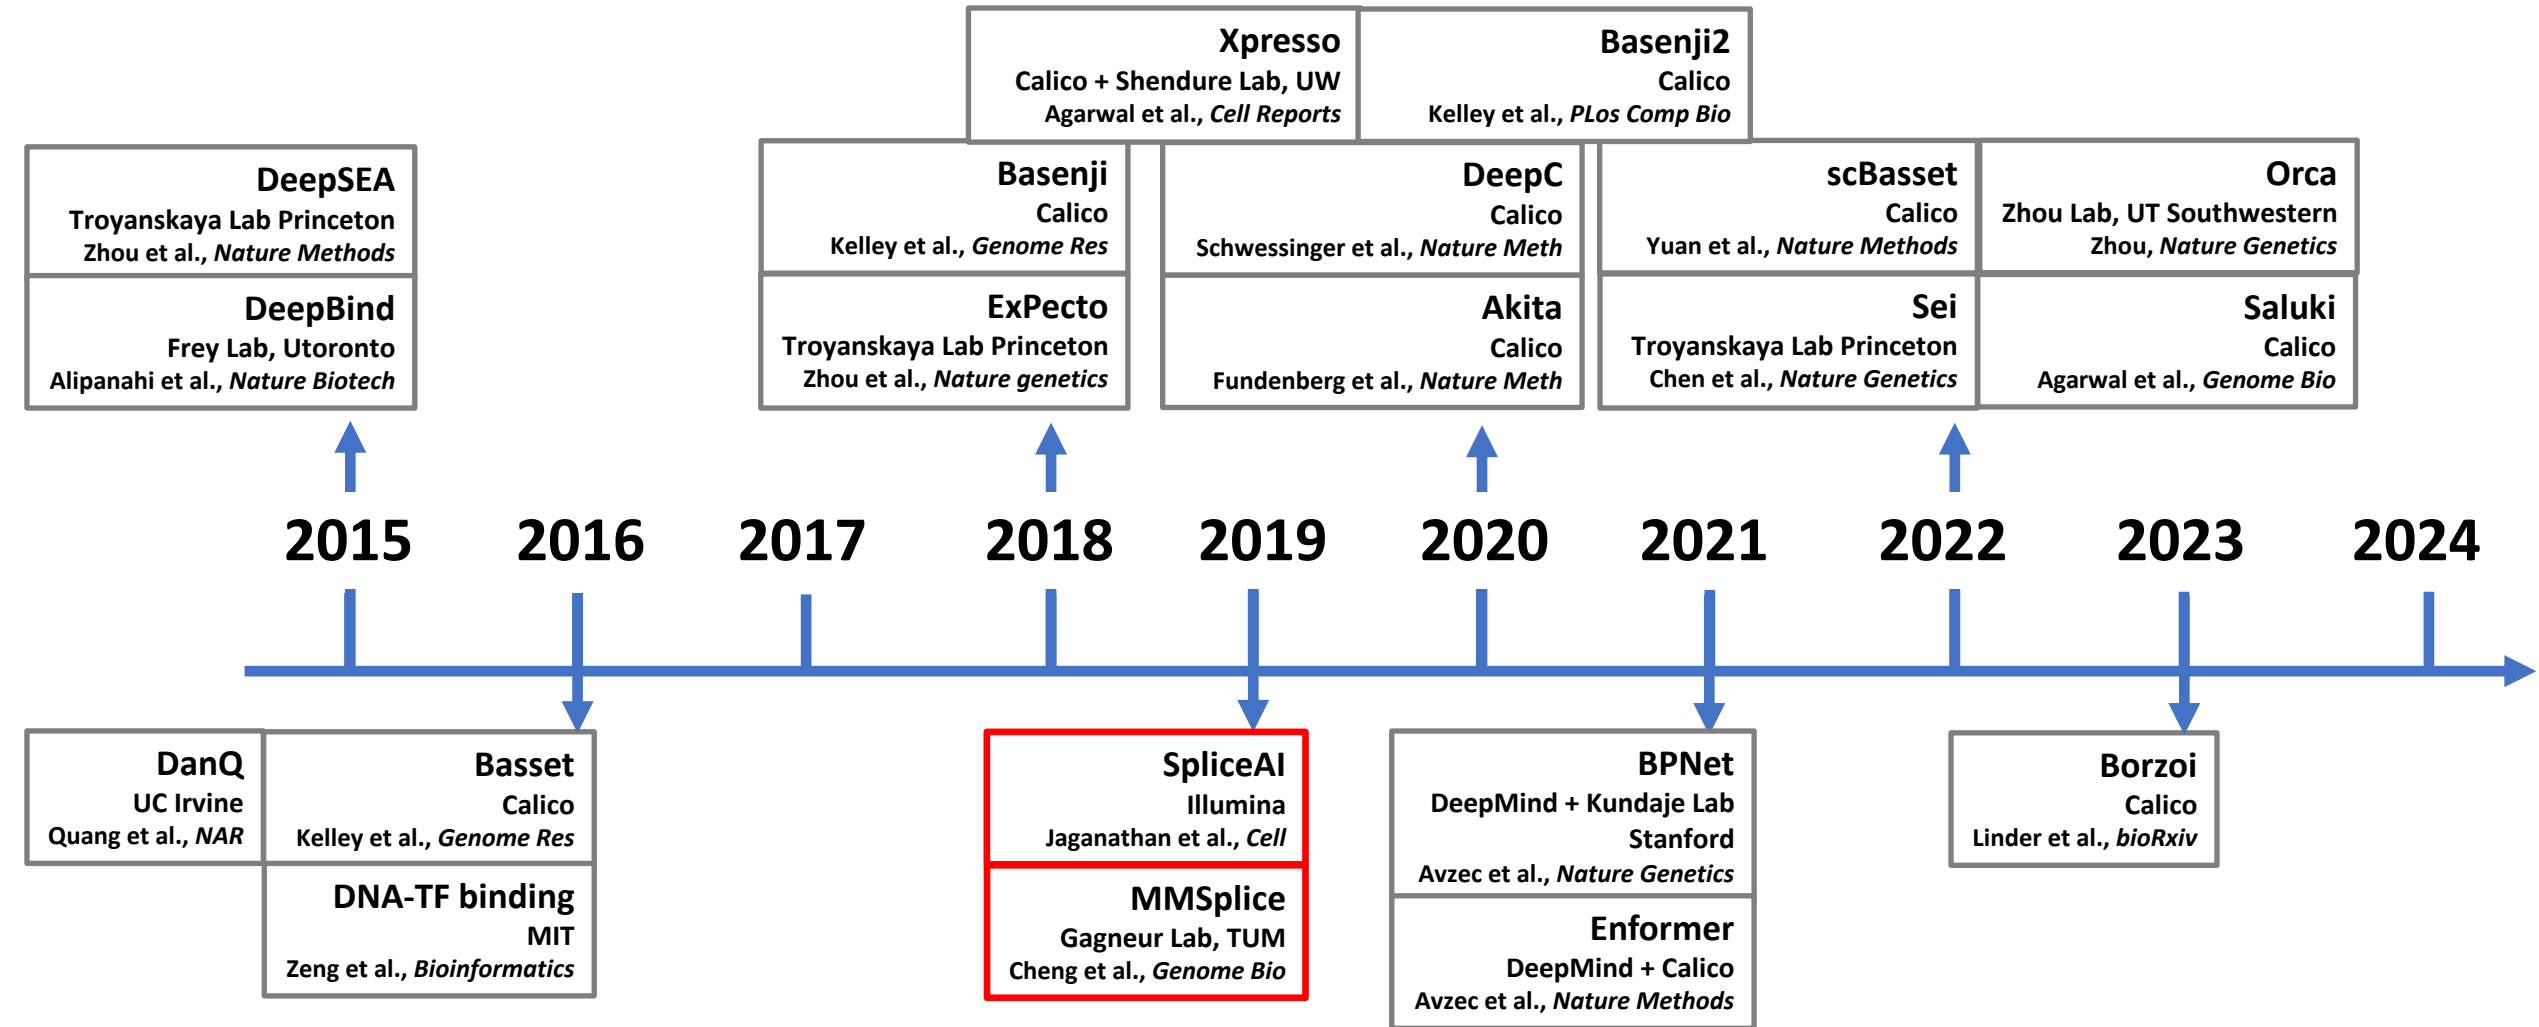
# Why Convolutional Neural Network to DNA ?



# Why Convolutional Neural Network to DNA ?




# Why Convolutional Neural Network to DNA ?



“All models are wrong, but some are useful”

- George Box, 1978

Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015).  
*Time series analysis: forecasting and control*. John Wiley & Sons.



Steven Salzberg




Mihaela Pertea





Anqi Liu

# SpliceAI-toolkit: splice site predictor

Pre-mRNA



**Chao, K. H.**, Mao, A., Liu, Anqi, Salzberg, S. L., & Pertea, M. (2024). SpliceAI-toolkit. Manuscript in preparation.  <https://ccb.jhu.edu/spliceai-toolkit/>


**Chao, K. H.**, Mao, A., Salzberg, S. L., & Pertea, M. (2023). Splam: a deep-learning-based splice site predictor that improves spliced alignments. *bioRxiv*.  <https://ccb.jhu.edu/splam/>



# SpliceAI: splite site predictor

$$\begin{bmatrix} A \\ C \\ G \\ T \\ N \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\begin{bmatrix} \textit{Acceptor} \\ \textit{Donor} \\ \textit{Neither} \\ \textit{Padding} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$$



Dimension:

$$X: (L+10000) * 4$$

**SpliceAI**

How do we train the model given that gene sequences vary?

.1	0	.6		.2
.1	.1	.2		.3
.8	.9	.2		.5

- => Donor
- => Acceptor
- => Neither

$$Y: L * 3$$

**Cell** 2019  
1413 citation

Predicting Splicing from Primary Sequence with Deep Learning

Kishore Jaganathan,<sup>1,6</sup> Sofia Kyriazopoulou Panagiotopoulou,<sup>1,6</sup> Jeremy F. McRae,<sup>1,6</sup> Siavash Fazel Darbandi,<sup>2</sup> David Knowles,<sup>3</sup> Yang I. Li,<sup>3</sup> Jack A. Kosmicki,<sup>1,4</sup> Juan Arbelaez,<sup>2</sup> Wenwu Cui,<sup>1</sup> Grace B. Schwartz,<sup>2</sup> Eric D. Chow,<sup>5</sup> Efsthathios Kanterakis,<sup>1</sup> Hong Gao,<sup>1</sup> Amirali Kia,<sup>1</sup> Serafim Batzoglou,<sup>1</sup> Stephan J. Sanders,<sup>2</sup> and Kyle Kai-How Farh<sup>1,7,\*</sup>

# SpliceAI: data preprocessing



$\frac{F}{2}$        $W$        $\frac{F}{2}$

X NN ... NN ATGTCGTGTCGAGTTG CGTGTT CAGGTCAGTCAGGTCAGTAAGTAGAGCTCANN ... NN

Y 0000000200000100000002000001000002000100000200000001000000

$$X = \left[ \begin{array}{c} \text{NN ... NN ATGTCGTGTCGAGTTG} \end{array} \right] \quad Y = \left[ \begin{array}{c} \text{0000000200} \end{array} \right]$$



# SpliceAI: data preprocessing



$\frac{F}{2}$                        $W$                        $\frac{F}{2}$

**X** NN ... NNATGTCGTGTCGAGTTGTCGTGTTTCAGGTCAGTCAGGTCAGTAAGTAGAGCTCANN ... NN  
**Y**                      0000000200000100000002000001000002000100000200000001000000

$$\mathbf{X} = \begin{bmatrix} \text{NN ... NNATGTCGTGTCGAGTTG} \\ \text{TCGTGTCGAGTTGTCGTGTTTCAG} \end{bmatrix} \qquad \mathbf{Y} = \begin{bmatrix} 0000000200 \\ 0001000000 \end{bmatrix}$$

# SpliceAI: data preprocessing



$\frac{F}{2}$                        $W$                        $\frac{F}{2}$

**X** NN ... NNATGTCGTGTCGAGTTGTCGTGTTTCAGGTCAGGTCAGTAAGTAGAGCTCANN ... NN  
**Y**                      0000000200000100000002000001000002000100000200000001000000

$$\mathbf{X} = \begin{bmatrix} \text{NN ... NNATGTCGTGTCGAGTTG} \\ \text{TCGTGTCGAGTTGTCGTGTTTCAG} \\ \text{TTGTCGTGTTTCAGGTCAGGTCAGG} \end{bmatrix} \quad \mathbf{Y} = \begin{bmatrix} 0000000200 \\ 0001000000 \\ 2000001000 \end{bmatrix}$$

# SpliceAI: data preprocessing



**X** NN ... NNATGTCGTGTCGAGTTGTCGTGTTTCAGGTCAGTCAGGTCAGTAAGTAGAGCTCANN ... NN  
**Y** 0000000200000100000002000001000002000100000200000001000000

$$\mathbf{X} = \begin{bmatrix} \text{NN ... NNATGTCGTGTCGAGTTG} \\ \text{TCGTGTCGAGTTGTCGTGTTTCAG} \\ \text{TTGTCGTGTTTCAGGTCAGTCAGG} \\ \vdots \\ \text{AAGTAGAGCTCANNNN ... NN} \end{bmatrix} \quad \text{Shape: } \left\lceil \frac{L}{W} \right\rceil * \left( \frac{F}{2} + W + \frac{F}{2} \right)$$

$$\mathbf{Y} = \begin{bmatrix} 0000000200 \\ 0001000000 \\ 2000001000 \\ \vdots \\ 0000000000 \end{bmatrix} \quad \text{Shape: } \left\lceil \frac{L}{W} \right\rceil * (W)$$

# SpliceAI: data preprocessing

$W = 5000$

$F = 10,000$


## Raw gene DNA sequence

Gene 1   $L = 33200$

Gene 2   $L = 14600$

...

Gene n   $L = 25000$



**X**

[7, 15000, 4]

[3, 15000, 4]

...

[5, 15000, 4]



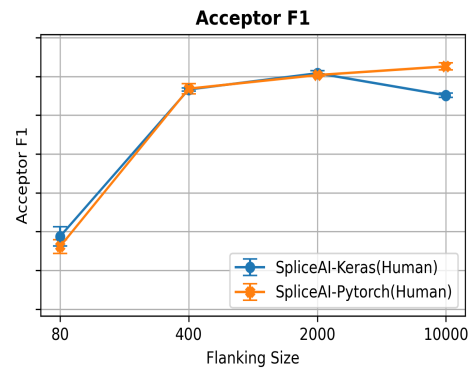
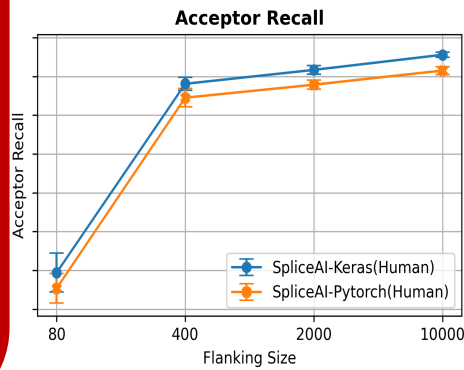
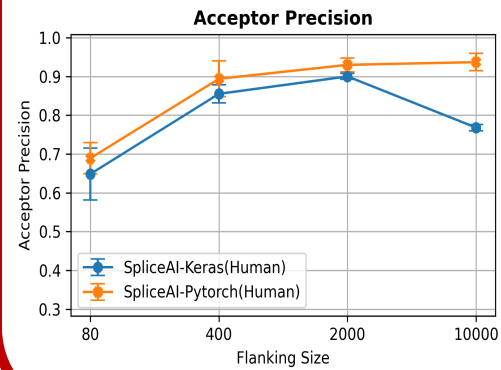
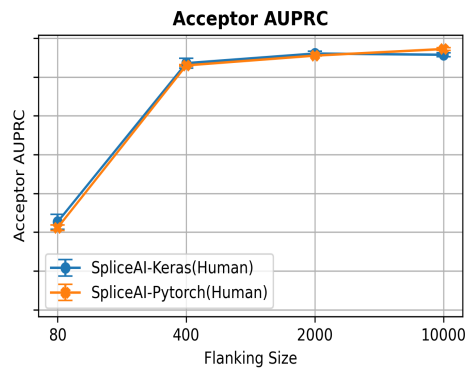
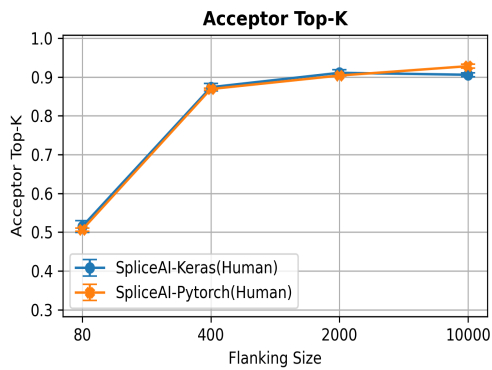
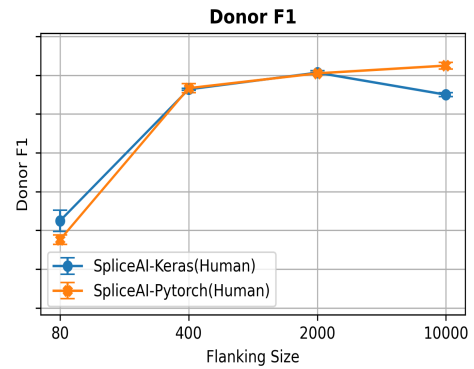
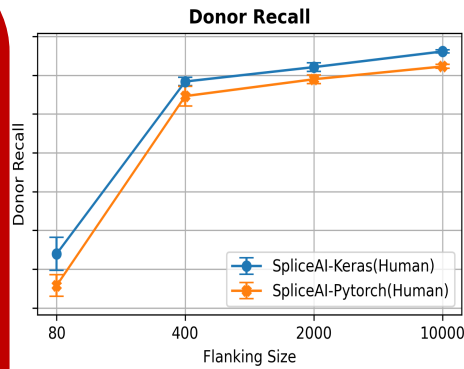



**Y**

[7, 5000, 3]

[3, 5000, 3]

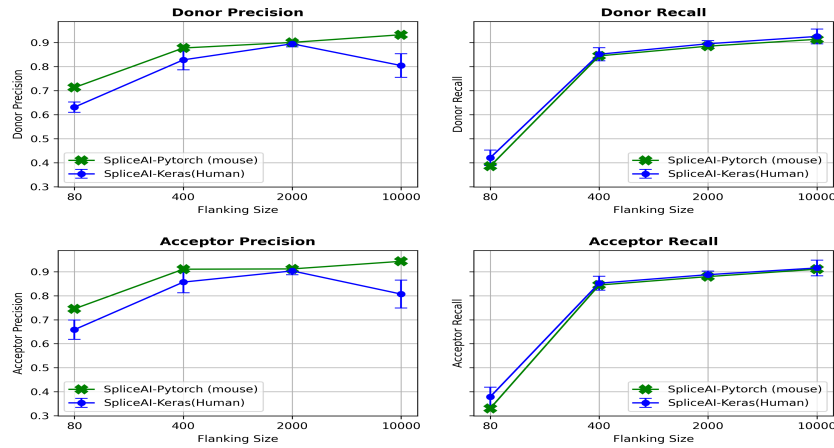
[5, 5000, 3]

# SpliceAI-toolkit : better than SpliceAI!

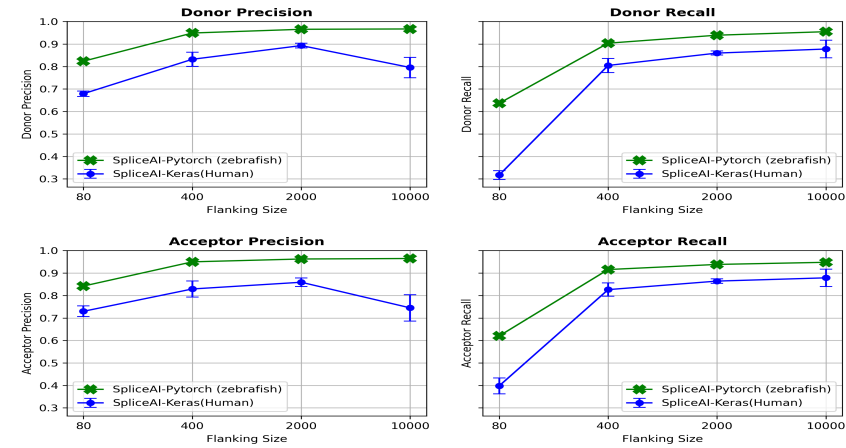


# SpliceAI-toolkit: retrain on different species


**A** Splice site prediction metrics for mouse




**B** Splice site prediction metrics for zebrafish



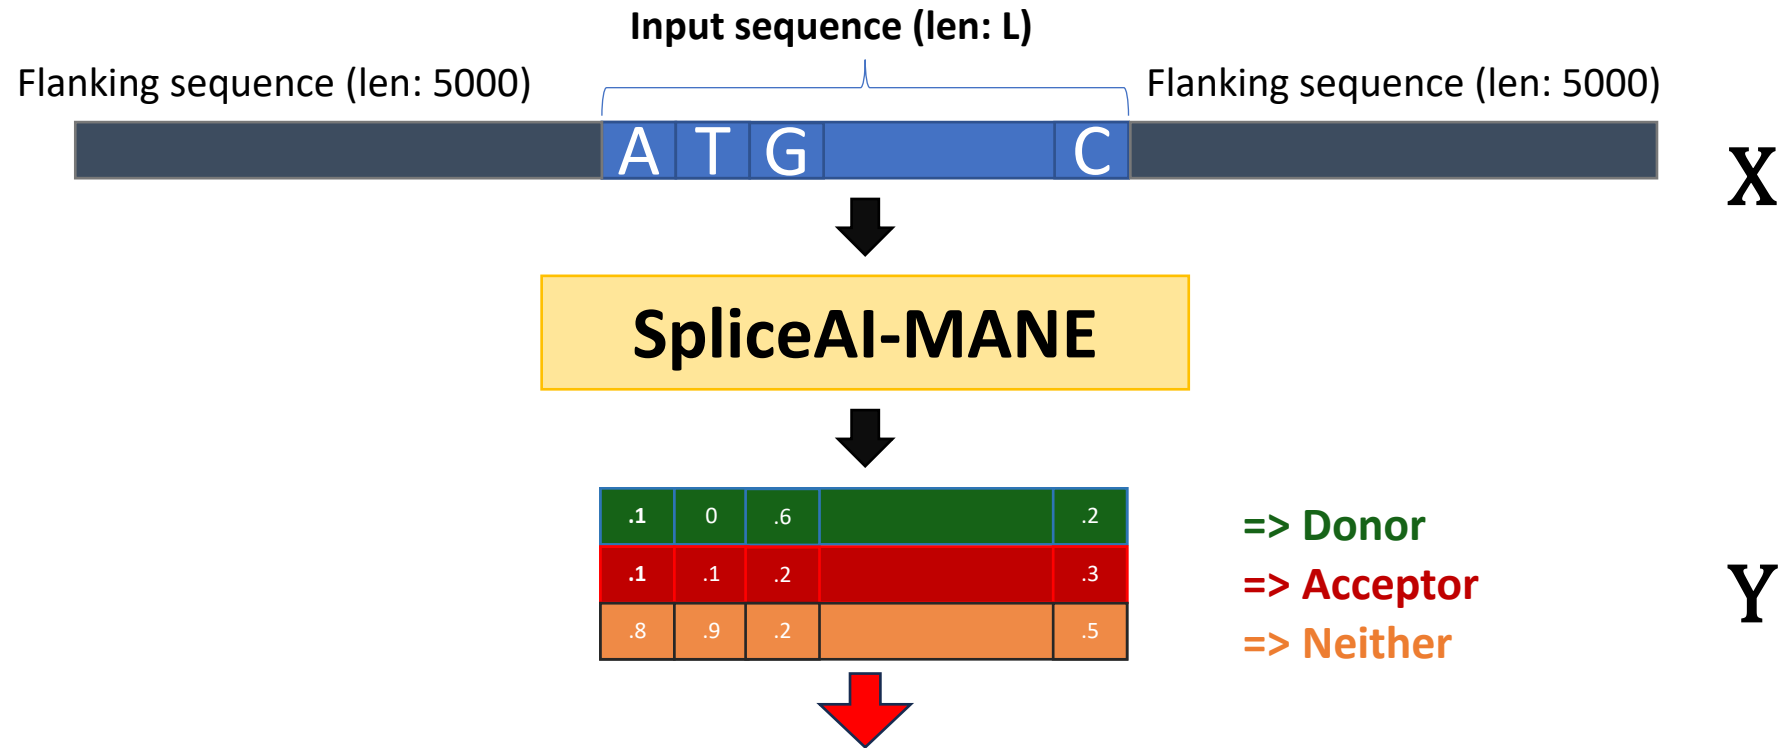
**C** Splice site prediction metrics for arabadop



**D** Splice site prediction metrics for bee



# SpliceAI-toolkit: new concept – Calibration



- What do SpliceAI-MANE scores signify?
- Do the model's predicted probabilities accurately represent the true likelihood of an event occurring?"

# SpliceAI-toolkit : new concept – Calibration



Model



0.99






# SpliceAI-toolkit : new concept – Calibration



Model



0.00



# SpliceAI-toolkit : new concept – Calibration



Model



0.69



0

0.2

0.4

0.6

0.8

1.0

Model predicted probability 15

# SpliceAI-toolkit : new concept – Calibration



Model



0.62



0

0.2

0.4

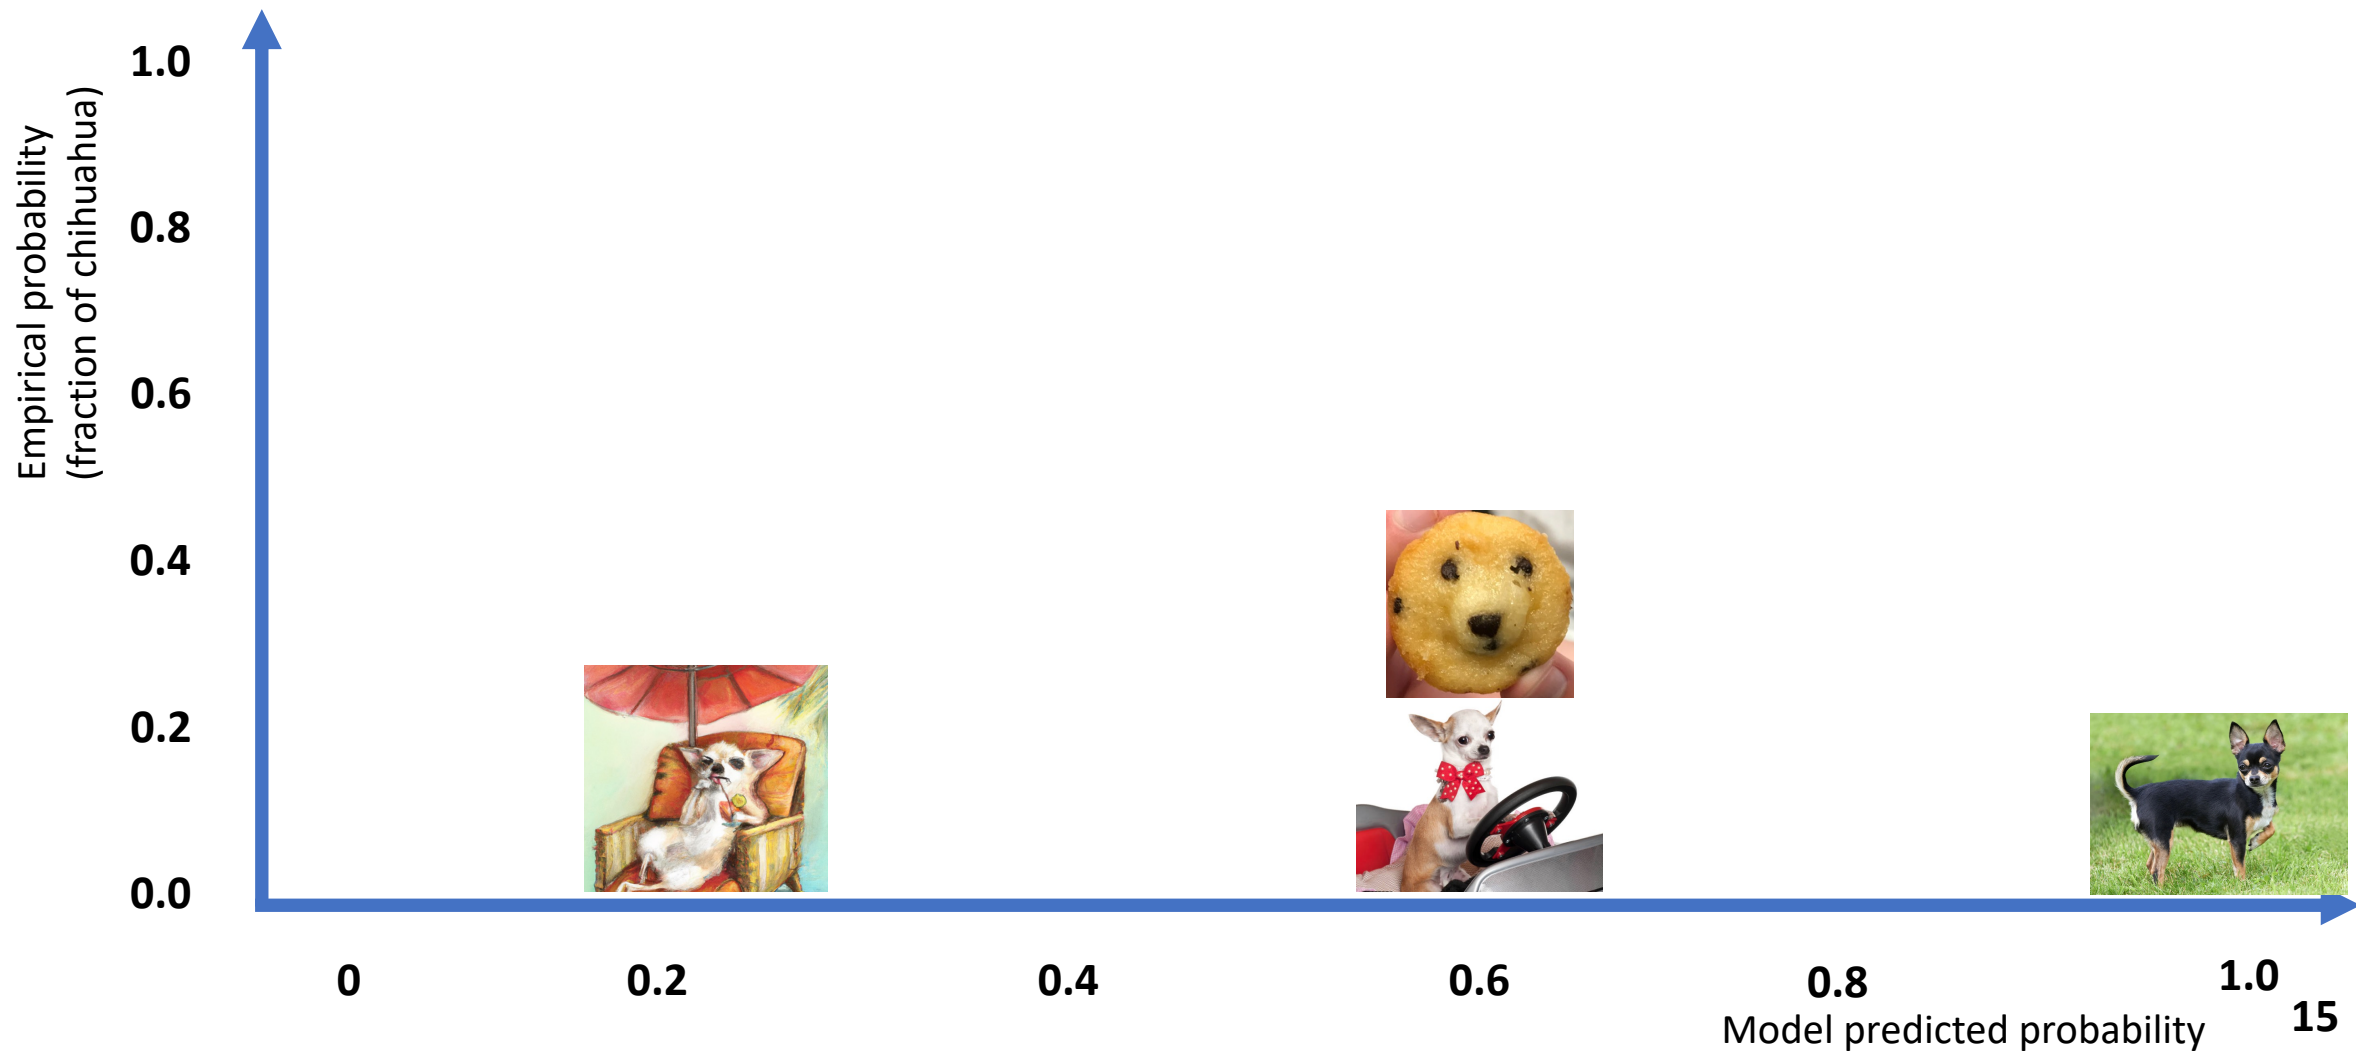
0.6

0.8

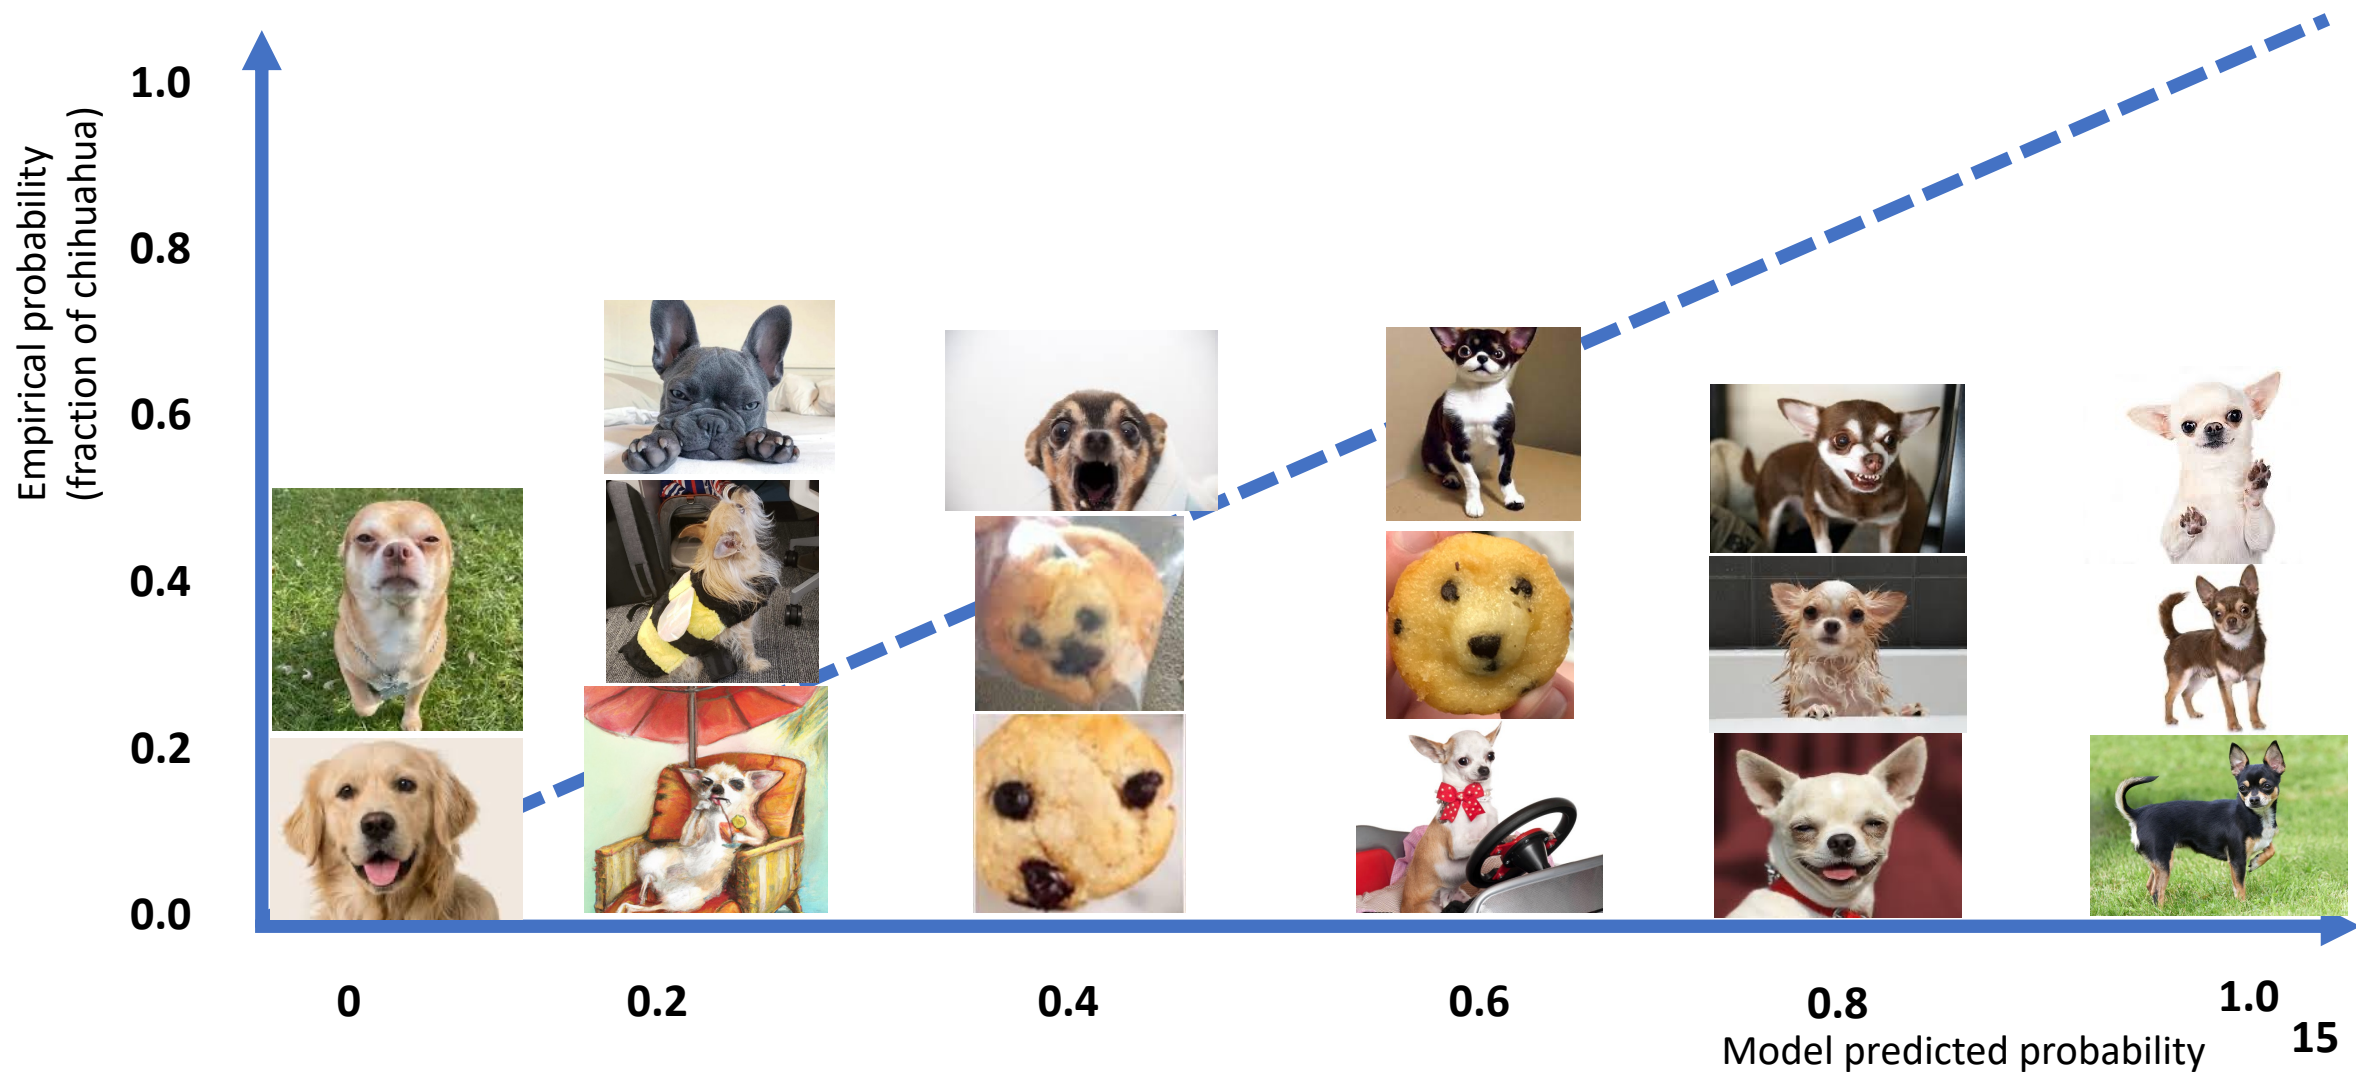
1.0

Model predicted probability 15

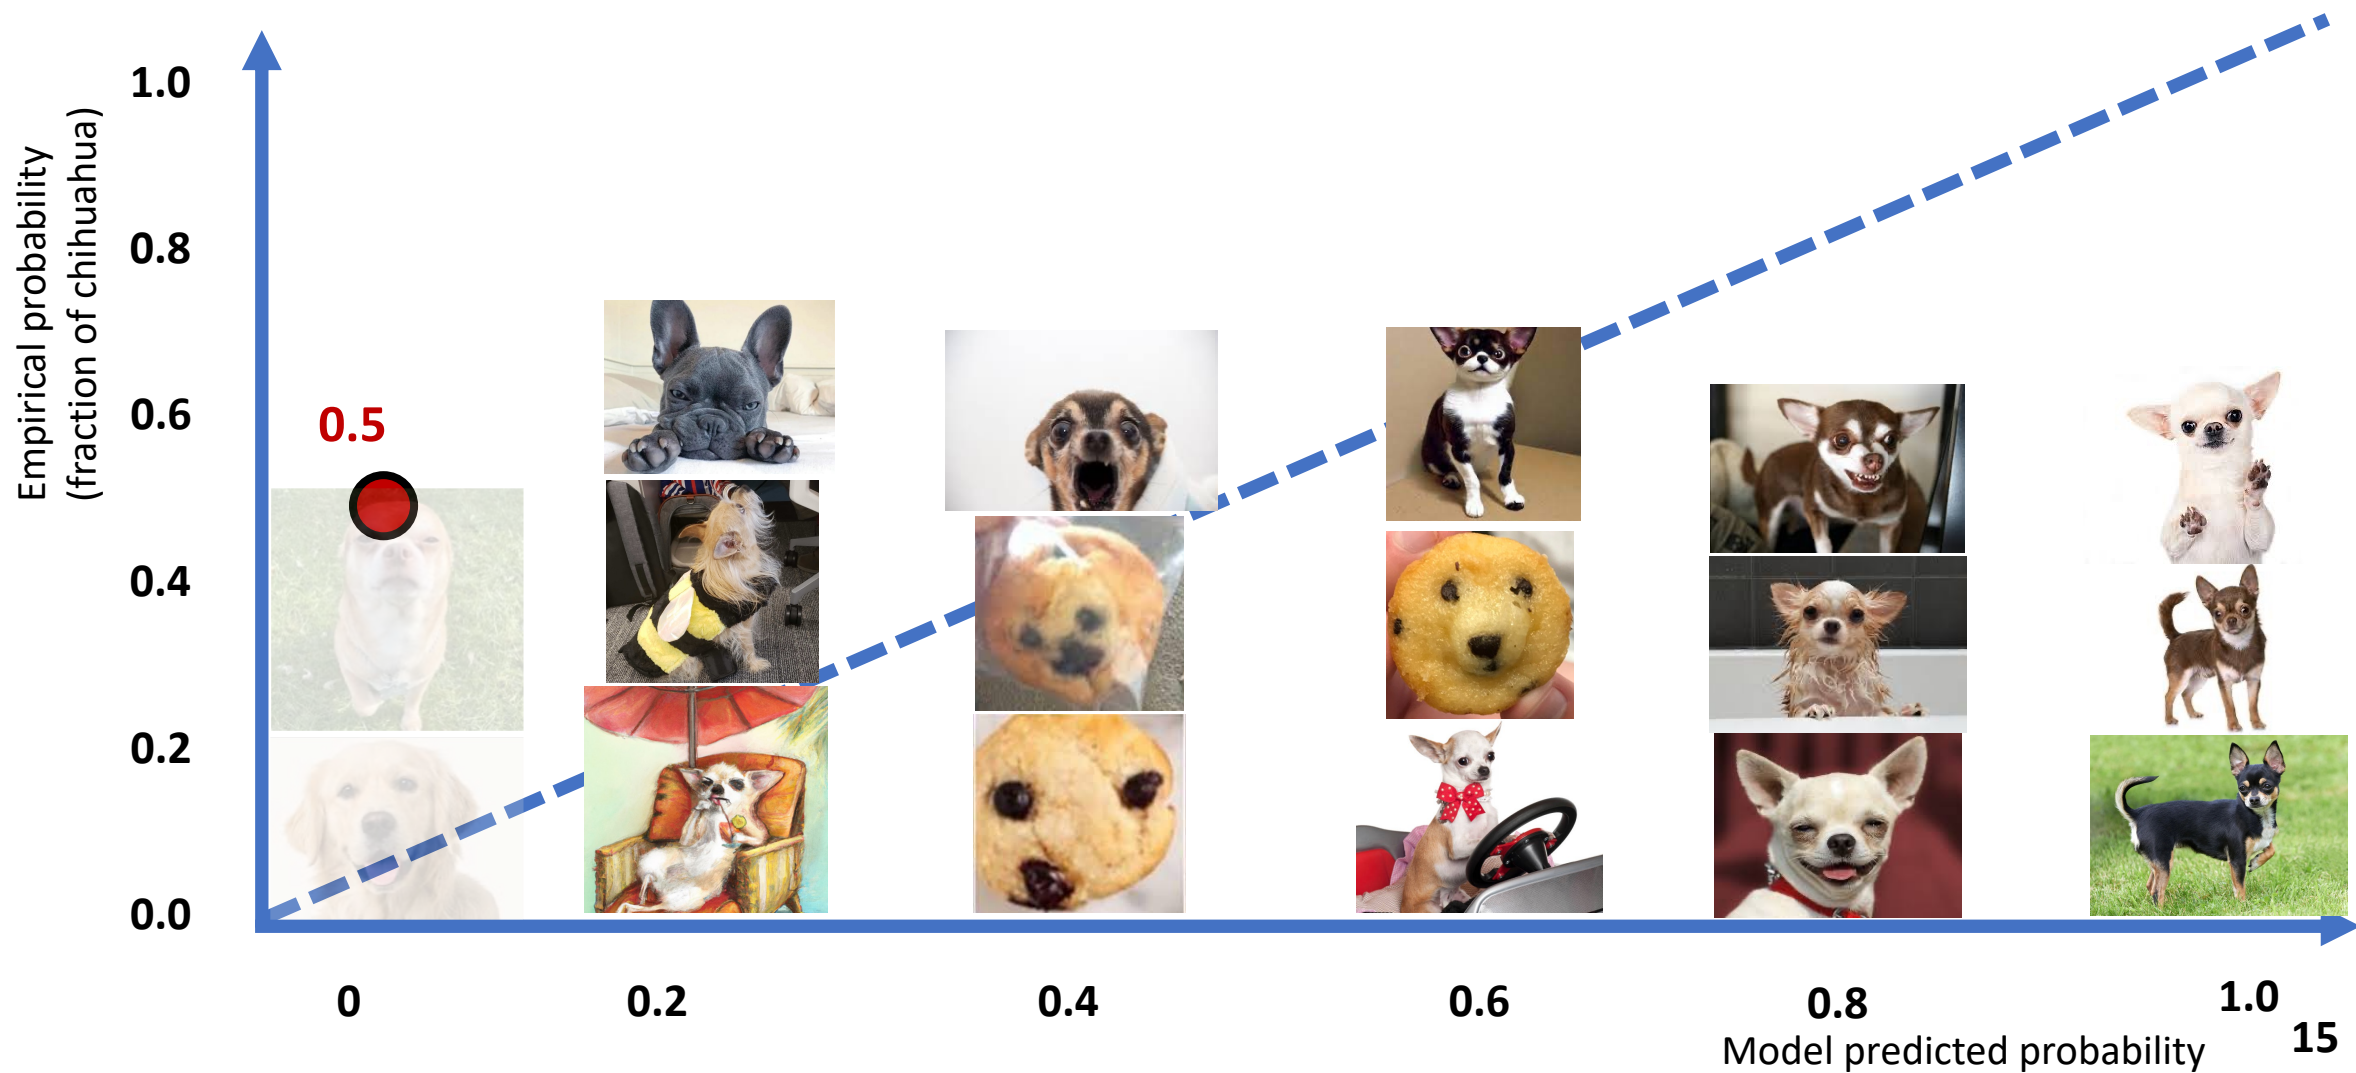
# SpliceAI-toolkit : new concept – Calibration



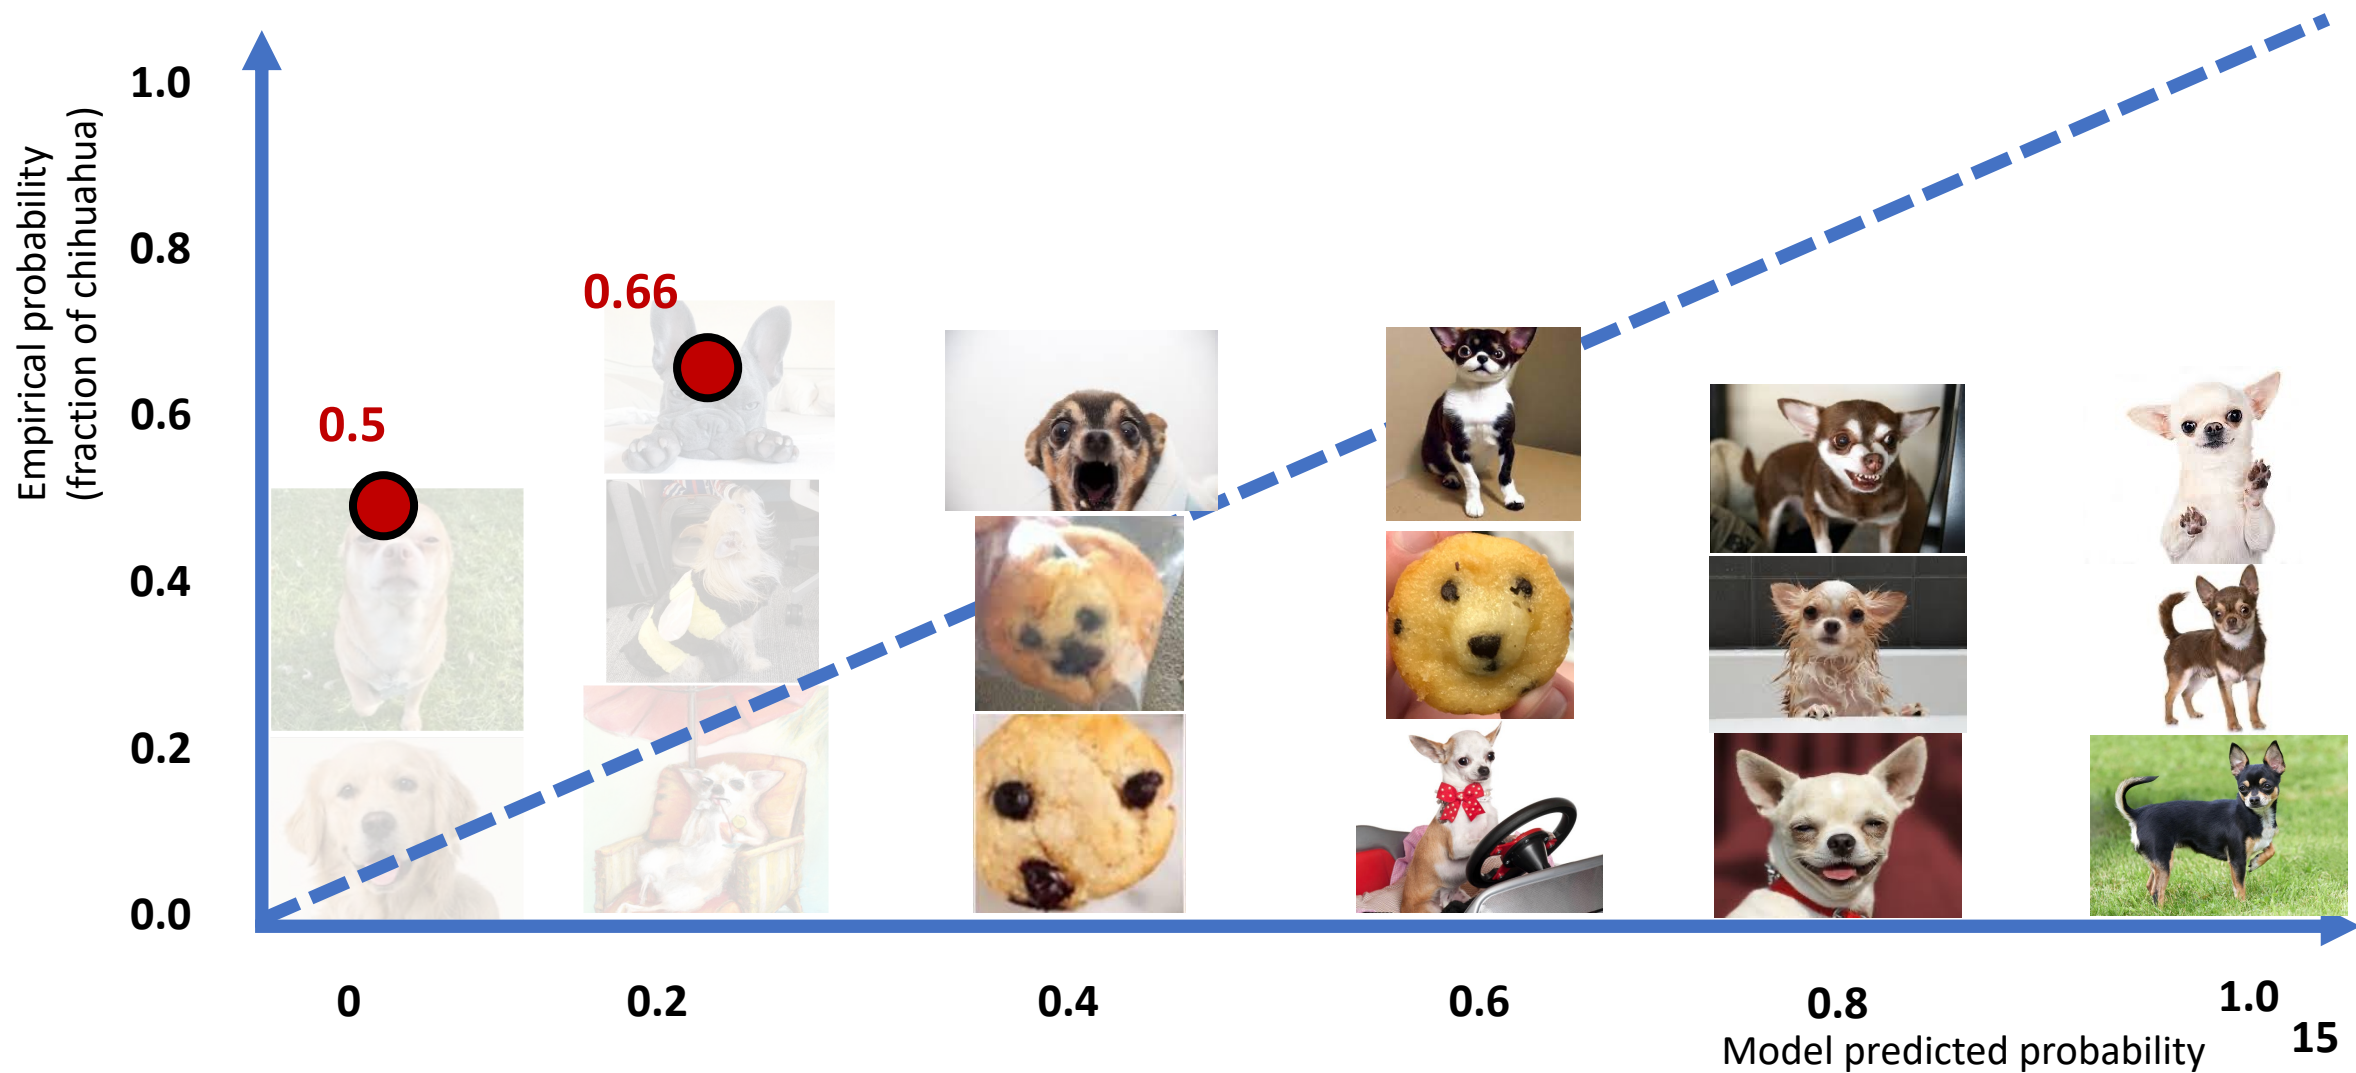
# SpliceAI-toolkit: new concept – Calibration



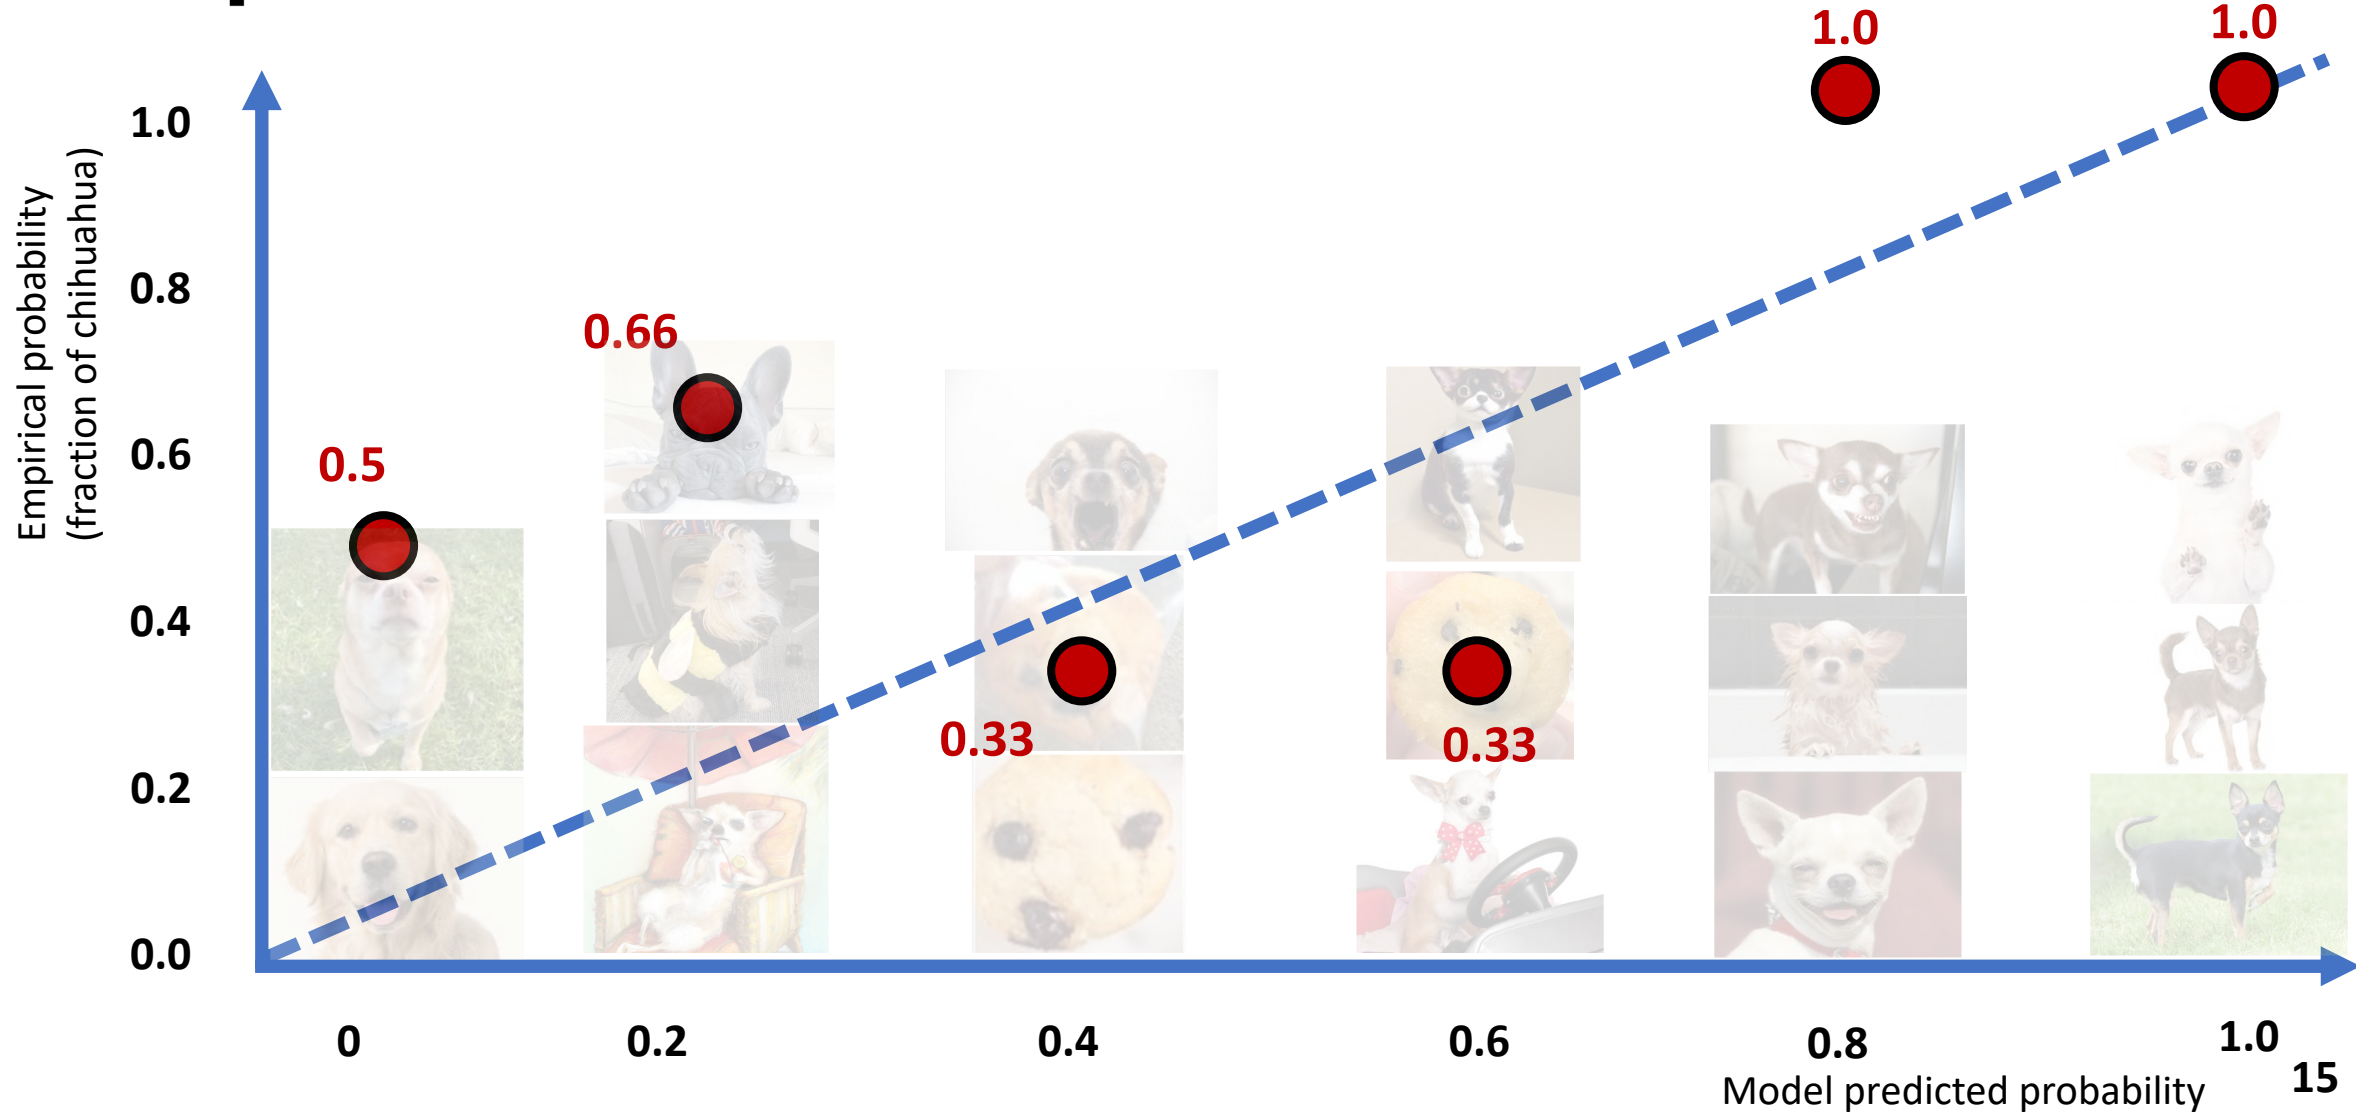
# SpliceAI-toolkit : new concept – Calibration



# SpliceAI-toolkit : new concept – Calibration

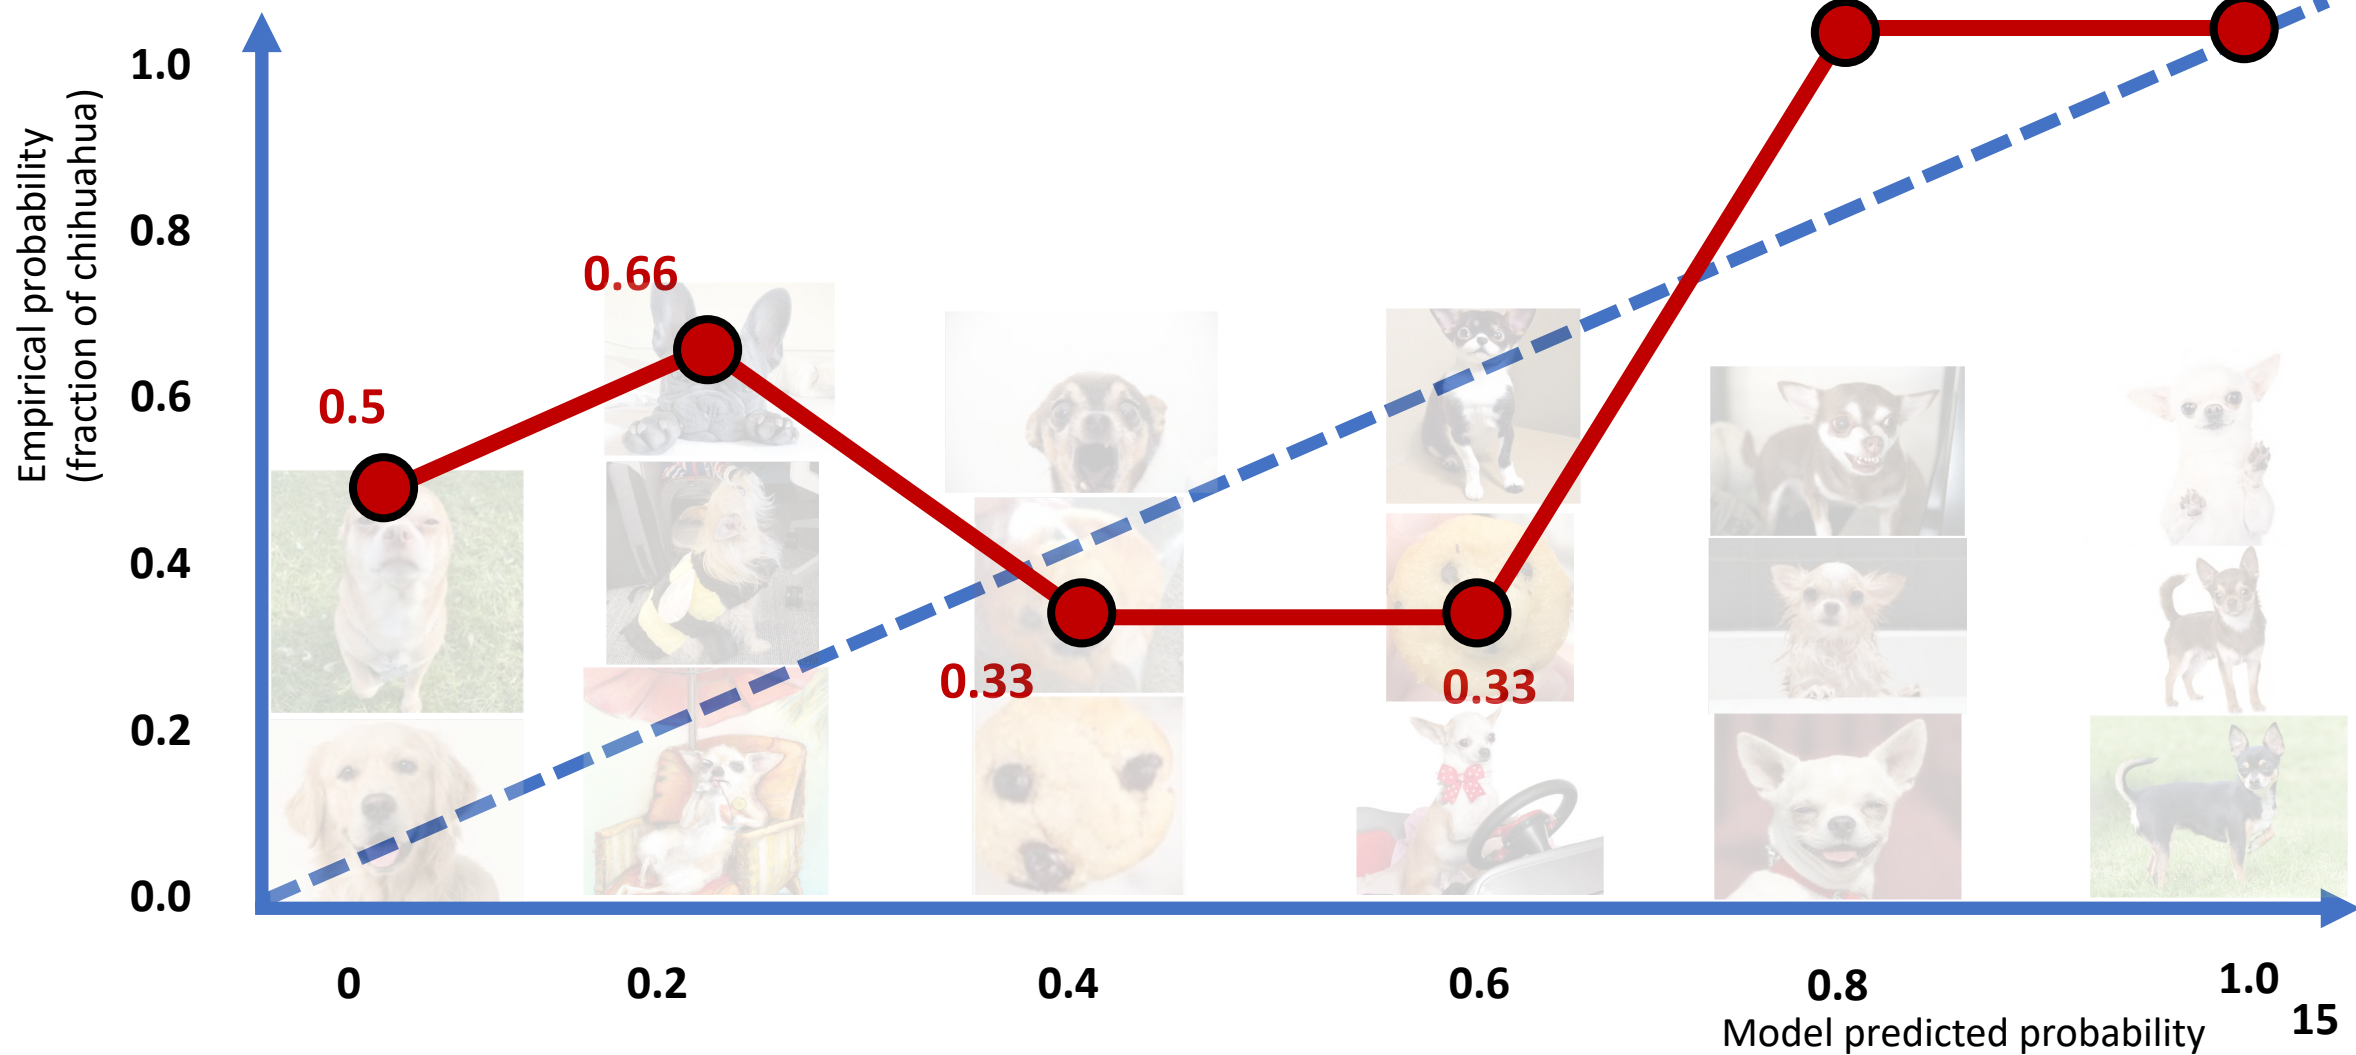


# SpliceAI-toolkit : new concept – Calibration

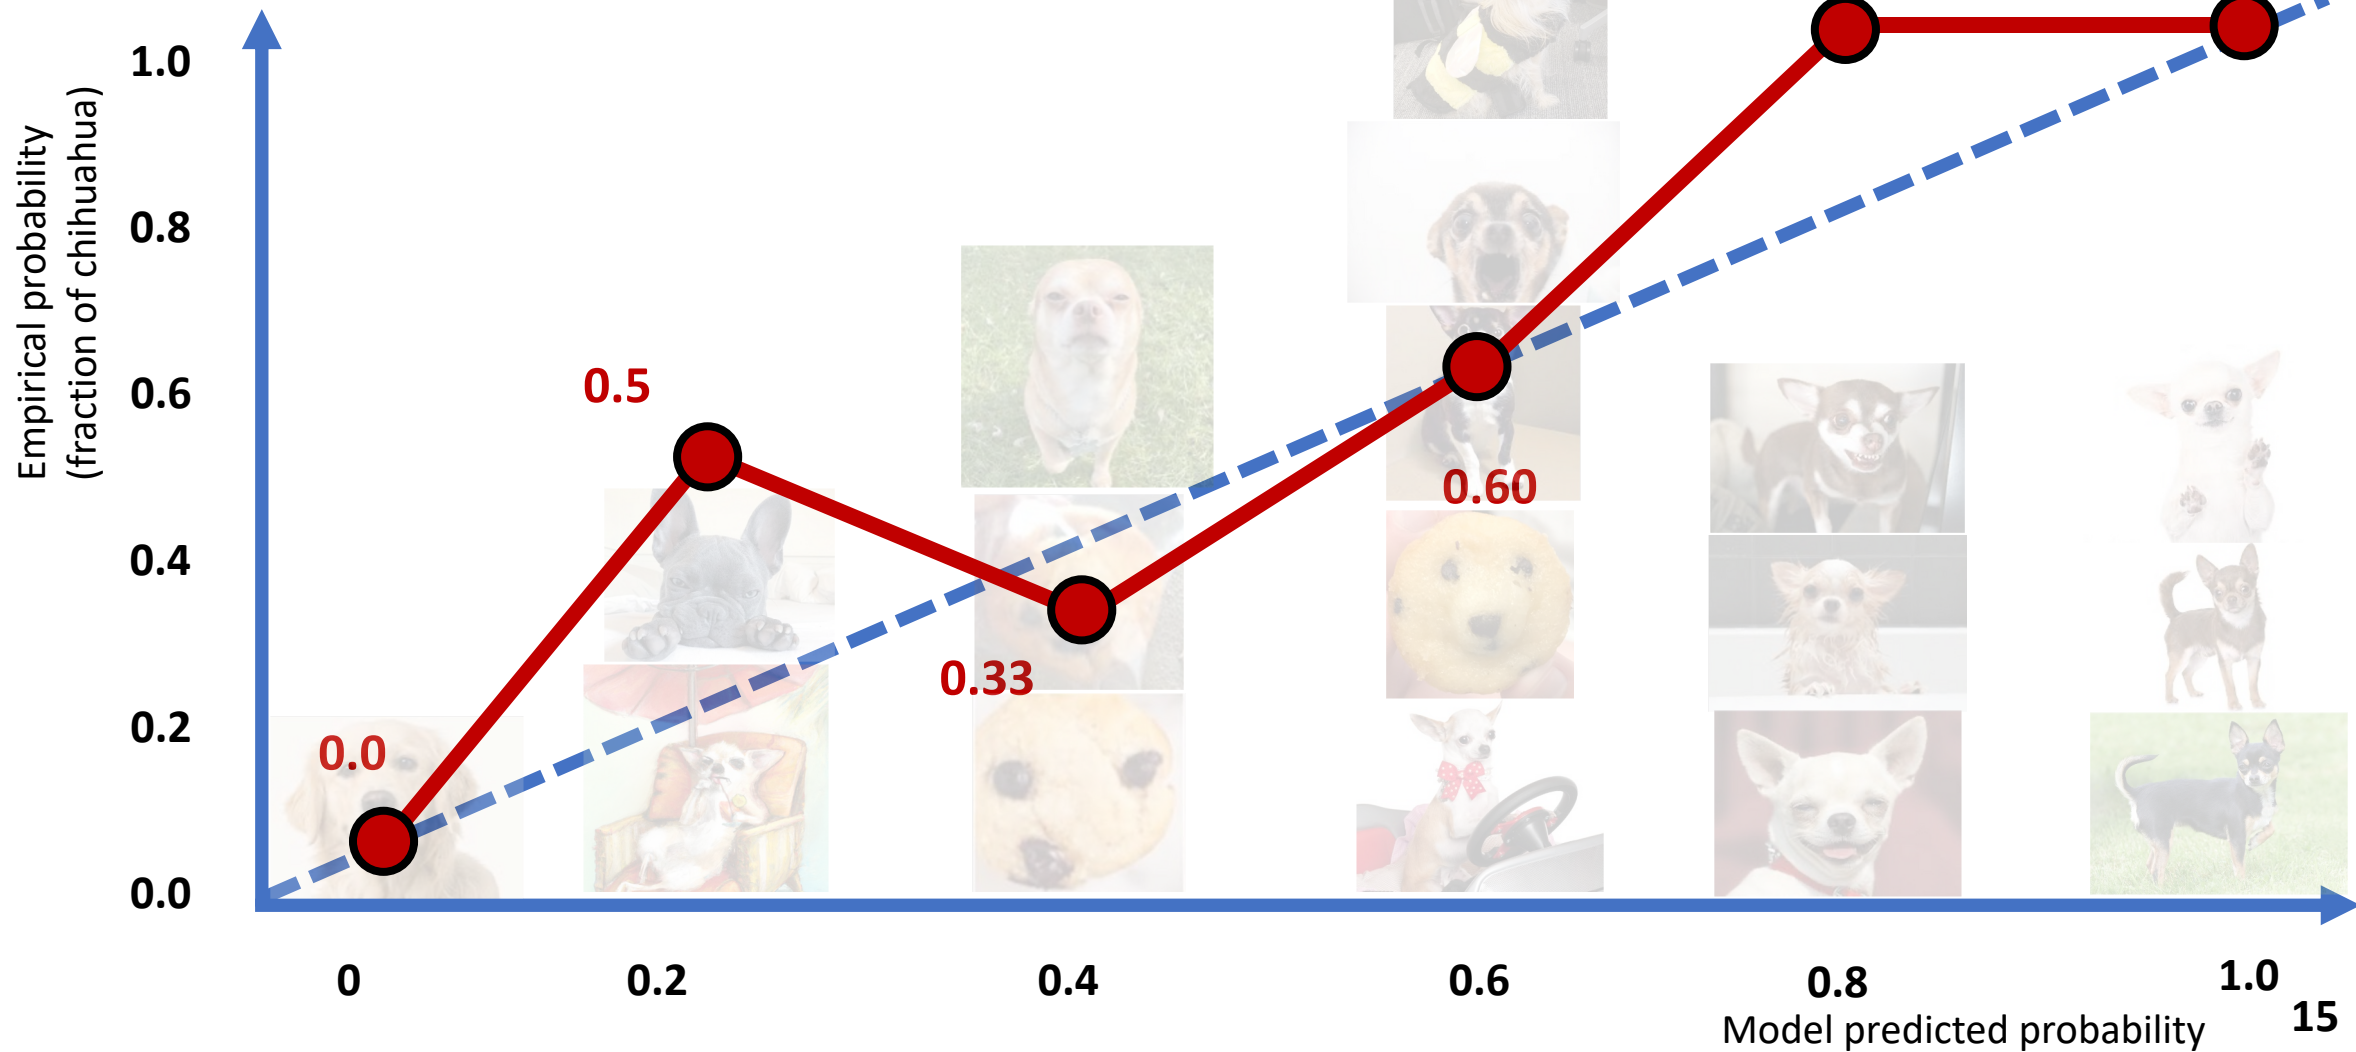




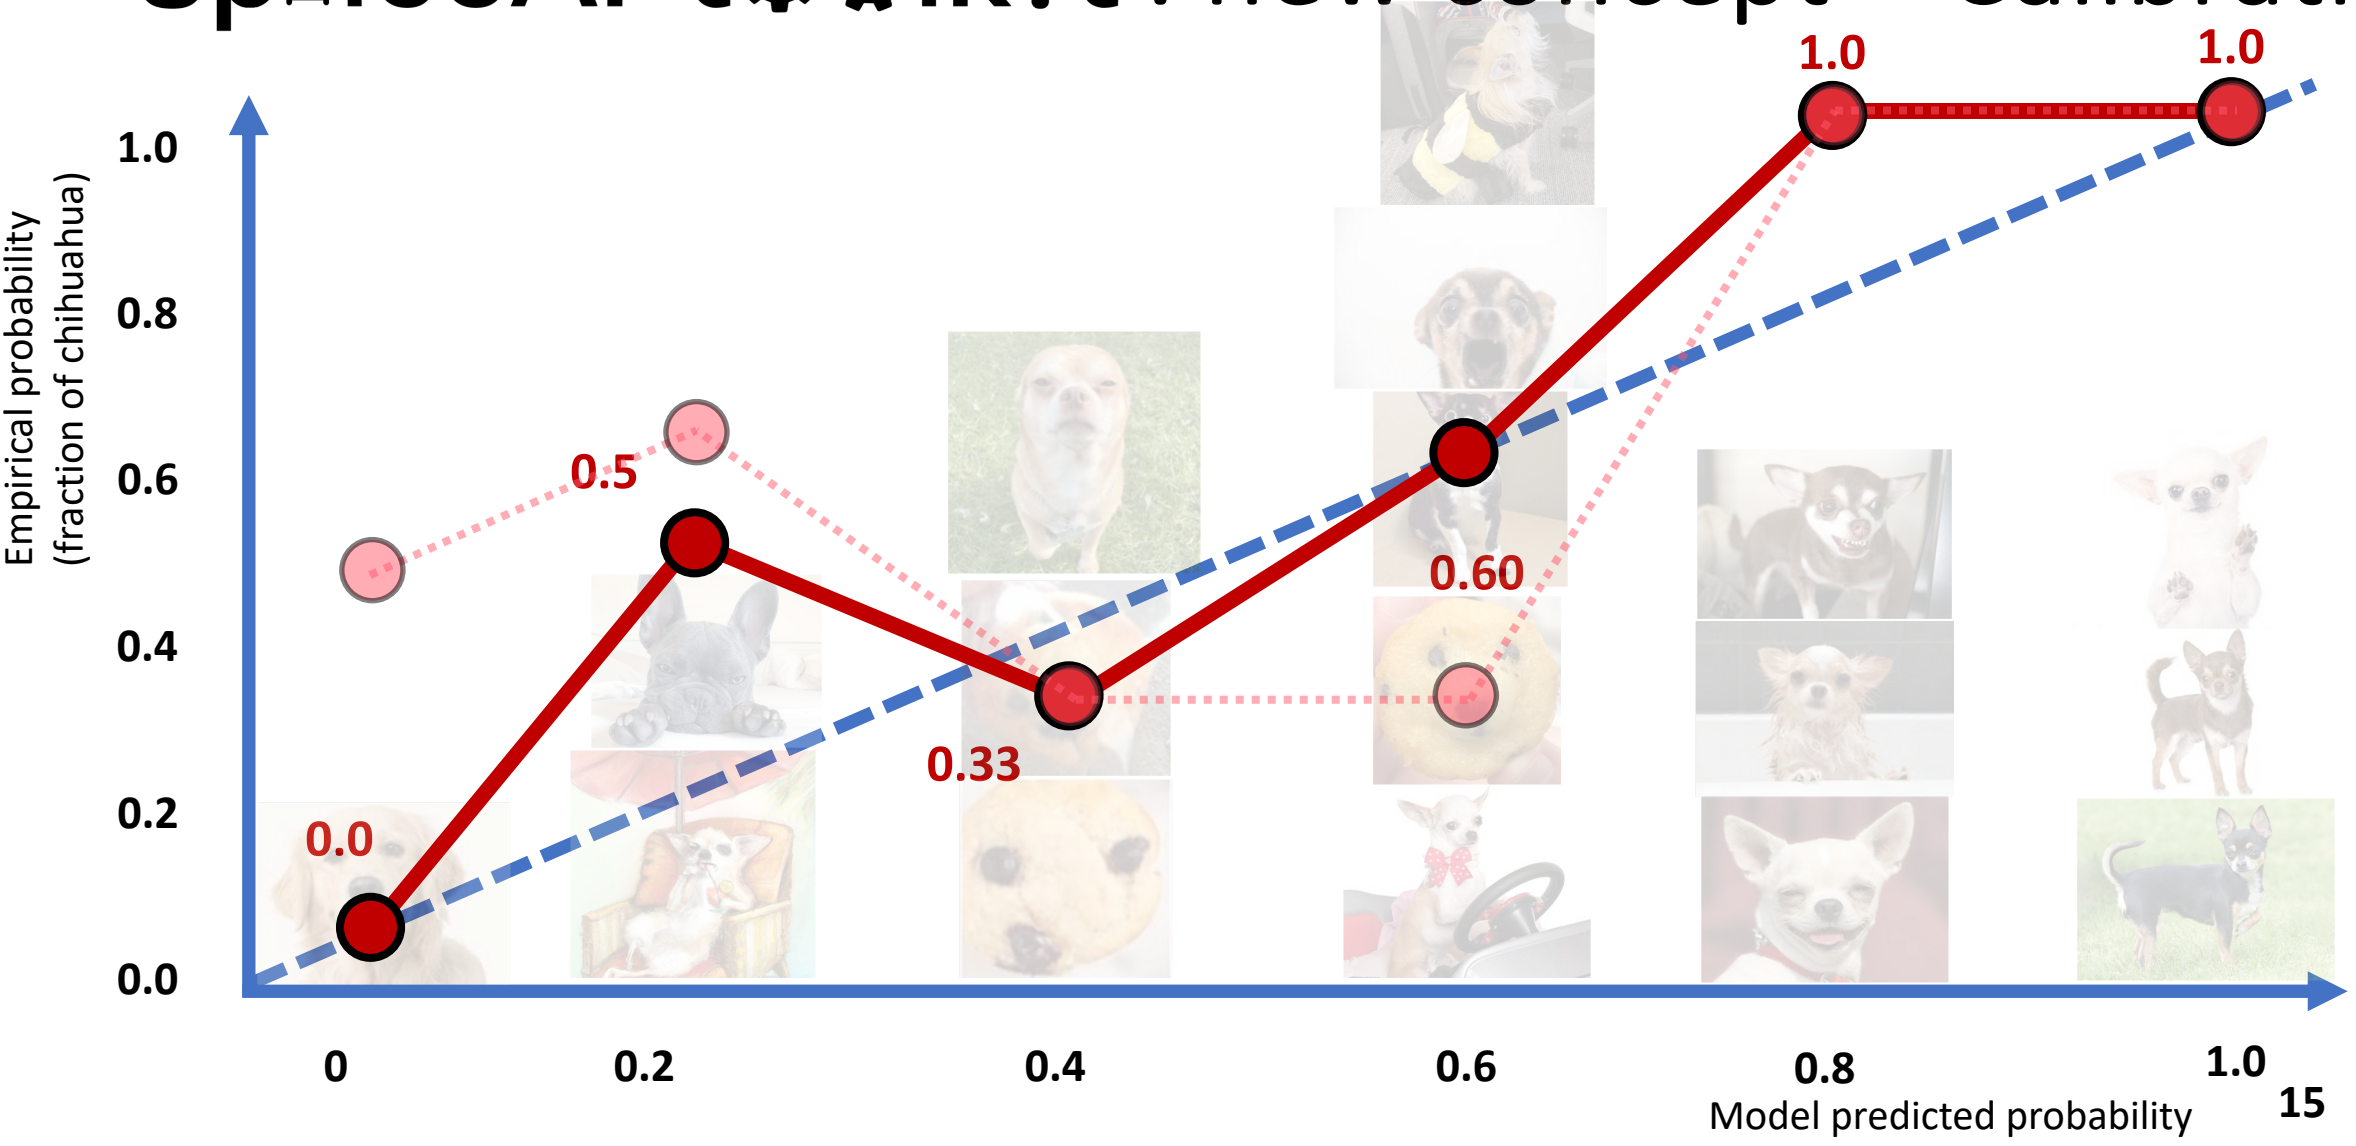
# SpliceAI-toolkit : new concept – Calibration



# SpliceAI-toolkit : new concept – Calibration



# SpliceAI-toolkit : new concept – Calibration




# SpliceAI-toolkit: new concept –

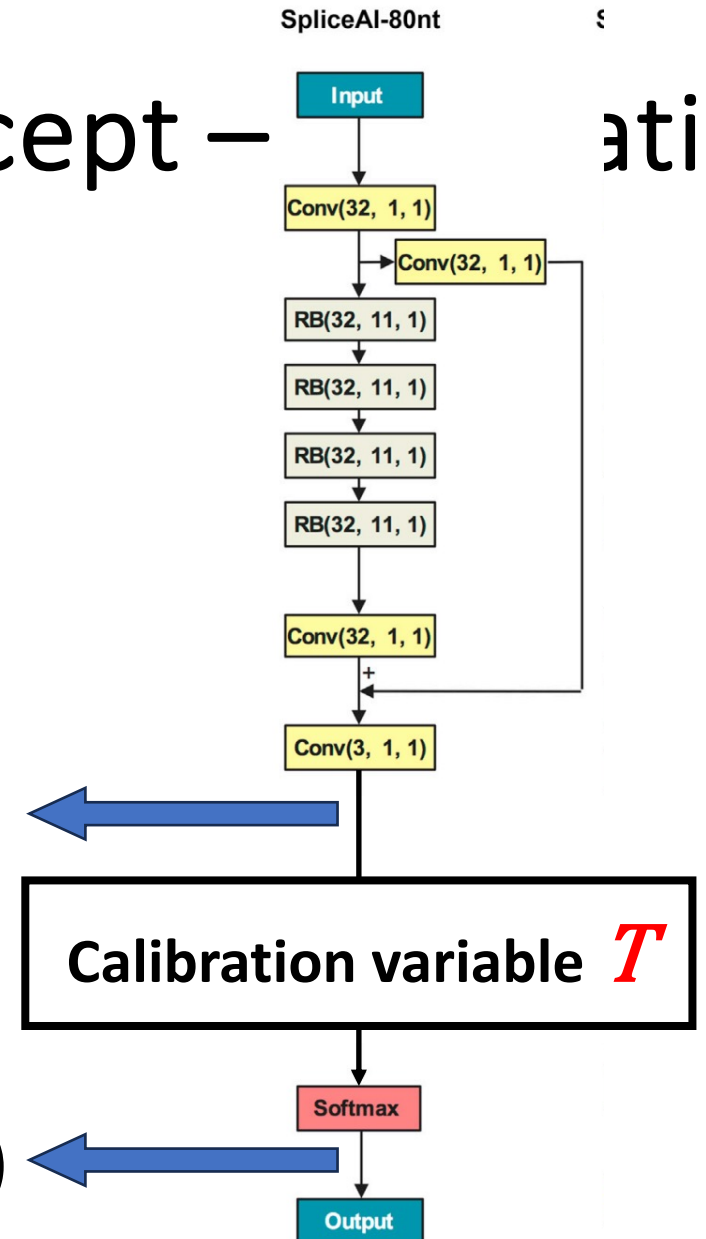
ation

- Platt scaling (Temperature scaling)

$$P_i = \frac{e^{\frac{y_i}{T}}}{\sum_{k=1}^n e^{\frac{y_k}{T}}}$$



Logits( $y_i$ )



Probability ( $P_i$ )

$i \in \{Neither, Acceptor, Donor\}$

Platt, John. "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods." Advances in large margin classifiers 10.3 (1999): 61-74.

# SpliceAI-toolkit: new concept – Calibration

Before calibration - **NLL** : 0.13310300, **ECE**: 0.00001282



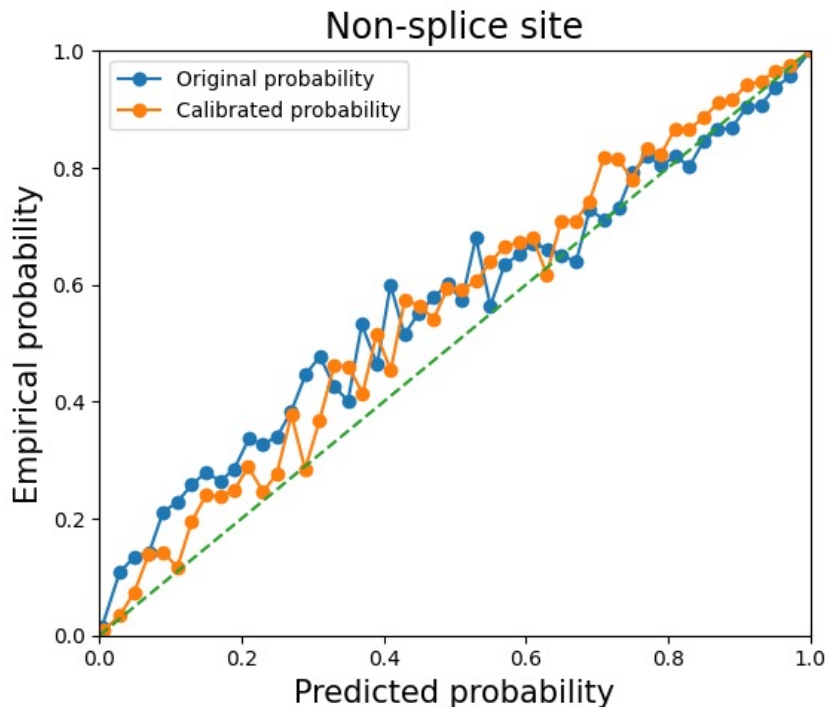
Optimal temperature: 1.28049

After calibration - **NLL**: 0.11896934, **ECE** : 0.00000079

$$\mathcal{L}_{NLL} = - \sum_{i=1}^n \log(\hat{\pi}(y_i|x_i))$$

$$conf(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \hat{p}_i \quad acc(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} 1(\hat{y}_i = y_i)$$

$$\mathcal{L}_{ECE} = \sum_{m=1}^M \frac{|B_m|}{n} |acc(B_m) - conf(B_m)|$$



# SpliceAI-toolkit : Summary

1. Data preprocessing: sliding window chunking
2. Easy-to-run framework to train your own SpliceAI
3. Pretrained SpliceAI-MANE
4. Pretrained SpliceAI on different species
5. Predict genetic variants' effect on splice sites
6. Model calibration: temperature scaling

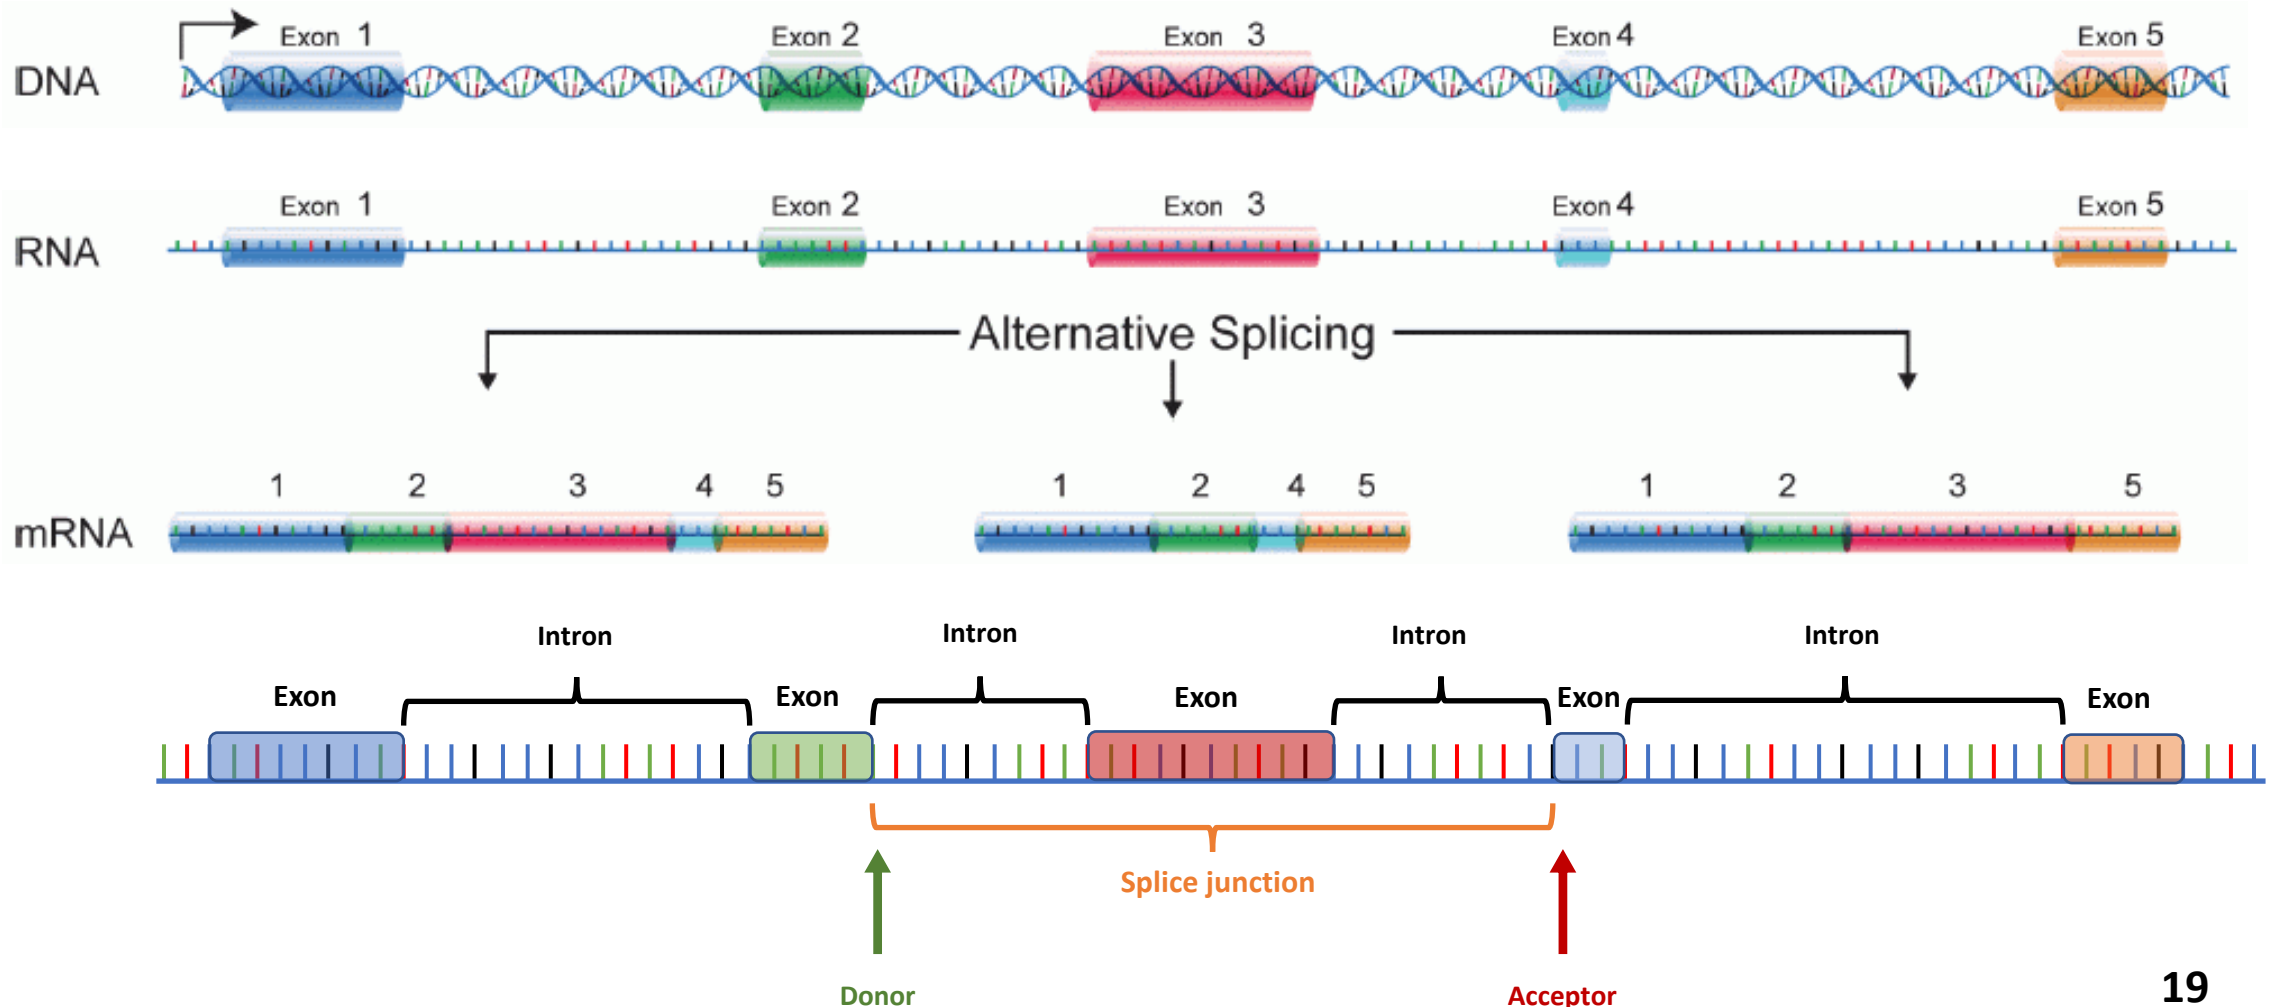
**Problem solved?**

**Chao, K. H.**, Mao, A., Liu, Anqi, Salzberg, S. L., & Pertea, M. (2024). SpliceAI-toolkit. Manuscript in preparation. <https://ccb.jhu.edu/spliceai-toolkit/> (in preparation)


18



# Is canonical labelling approach correct?



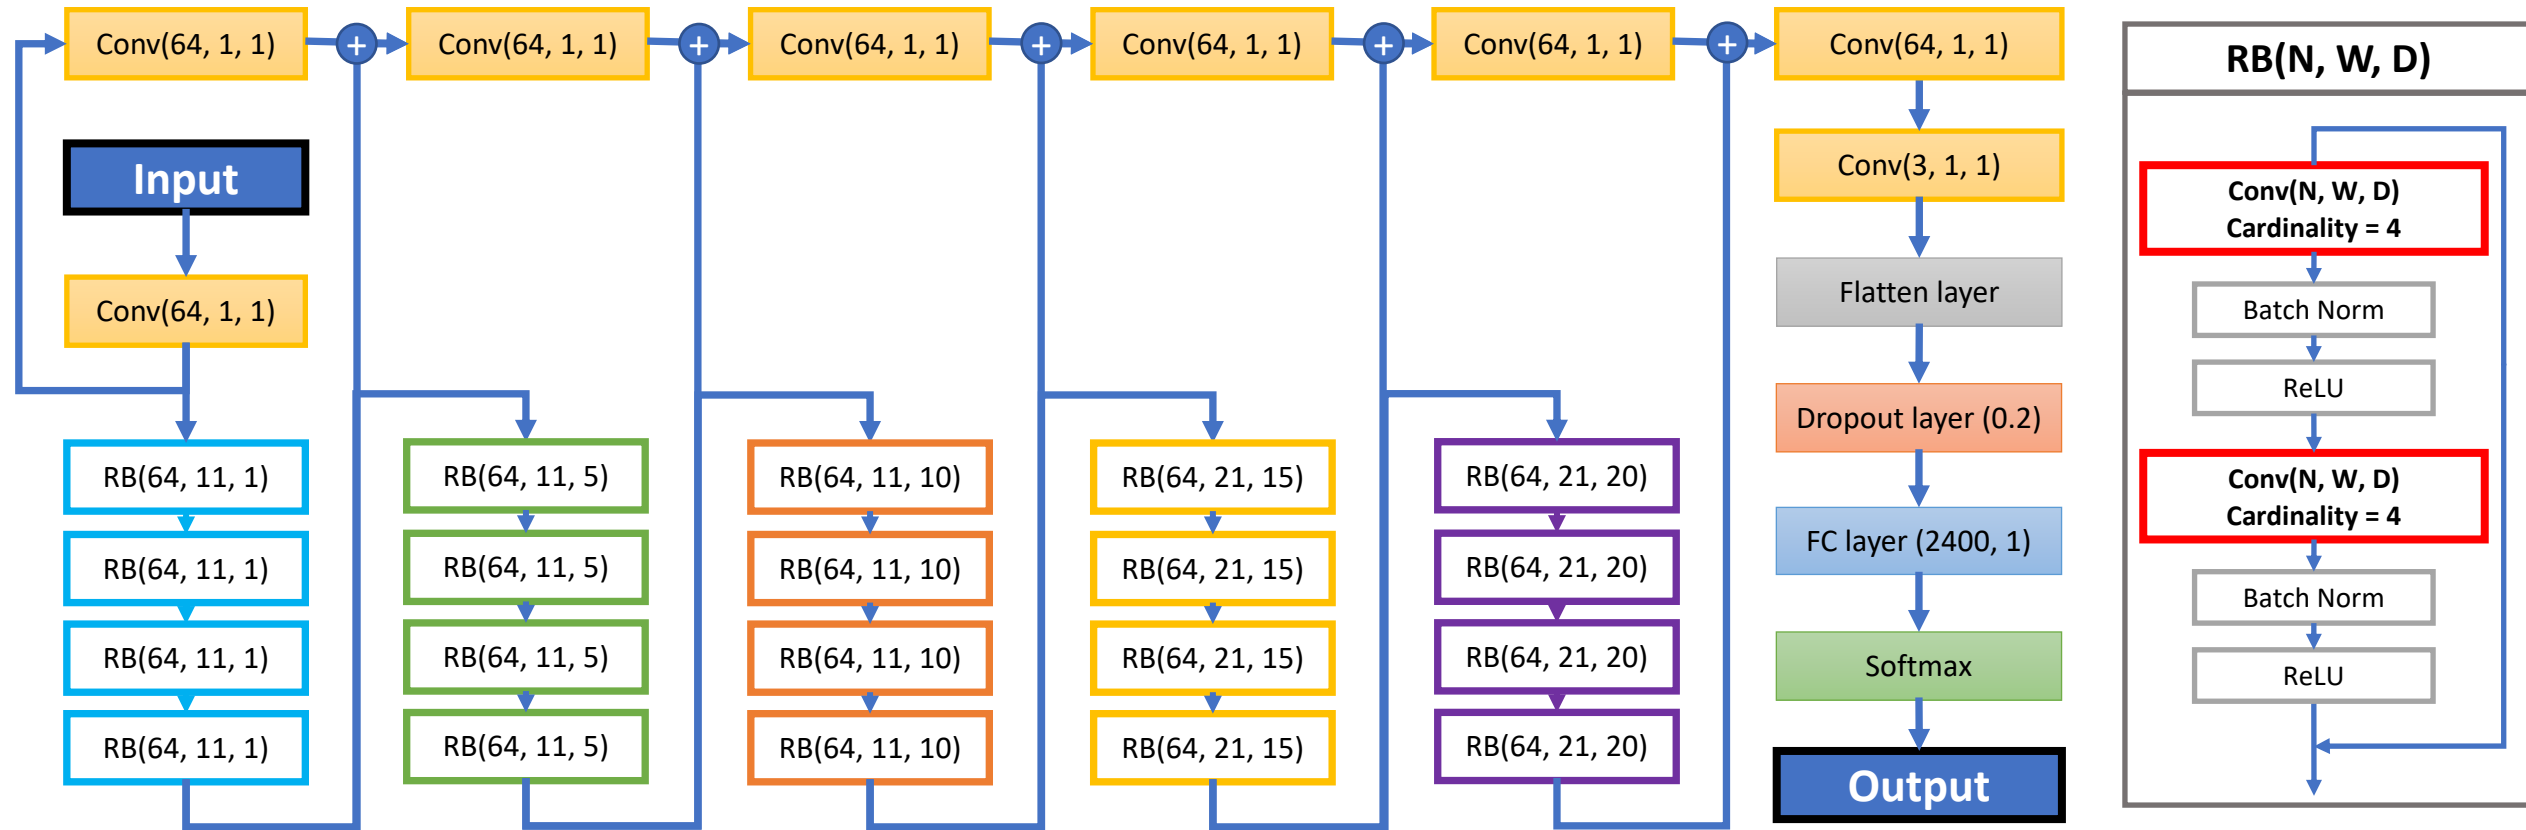
# SPLM : Data Processing



**Chao, K. H.**, Mao, A., Salzberg, S. L., & Pertea, M. (2023). Splam: a deep-learning-based splice site predictor that improves spliced alignments. **bioRxiv**. <https://ccb.jhu.edu/splam/>



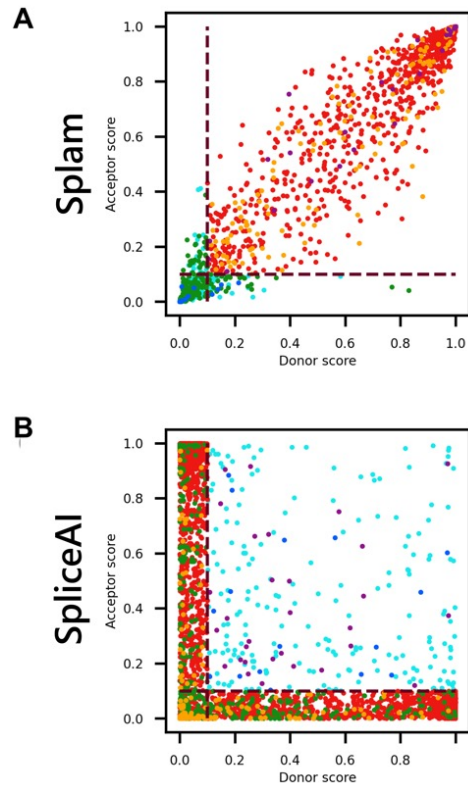
# SPLAM : Splam Model Architecture




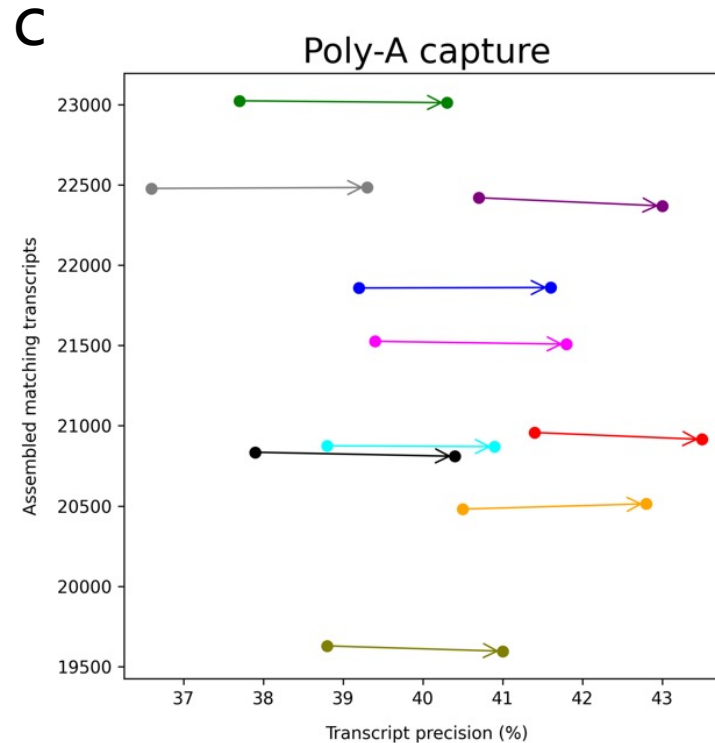
Chao, K. H., Mao, A., Salzberg, S. L., & Pertea, M. (2023). Splam: a deep-learning-based splice site predictor that improves spliced alignments. *bioRxiv*. <https://ccb.jhu.edu/splam/>

# SPLAM : deep-learning splice site predictor

## Score stability



## Transcriptome assembly improvement

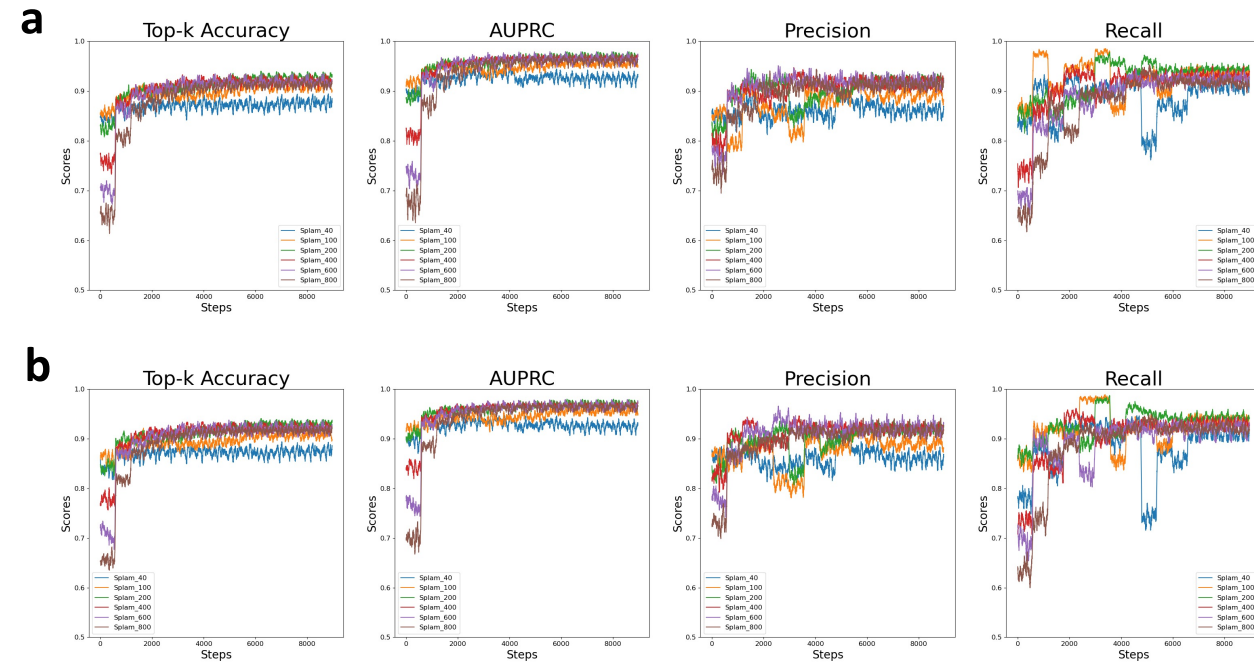
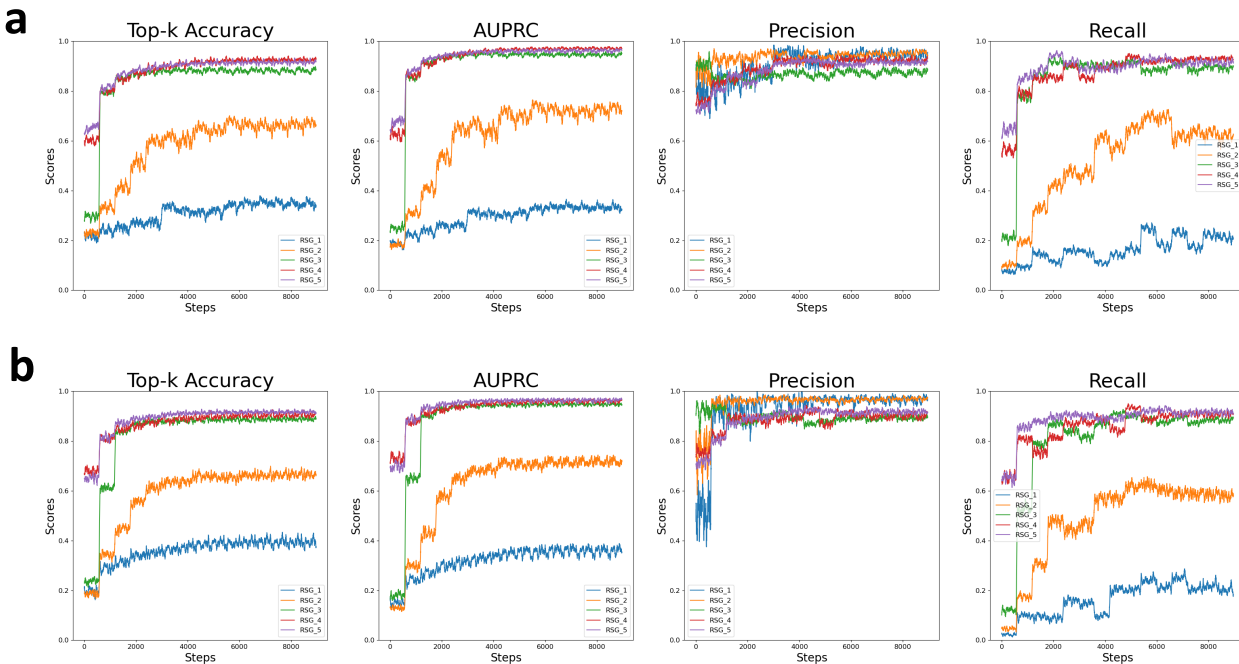


**Chao, K. H.**, Mao, A., Salzberg, S. L., & Pertea, M. (2023). Splam: a deep-learning-based splice site predictor that improves spliced alignments. **bioRxiv**. <https://ccb.jhu.edu/splam/>

# SPLM : deep-learning splice site predictor

Interpretability: ablation study

Interpretability: input sequence



**Chao, K. H.,** Mao, A., Salzberg, S. L., & Pertea, M. (2023). Splam: a deep-learning-based splice site predictor that improves spliced alignments. **bioRxiv.** <https://ccb.jhu.edu/splam/>

# SPLAM : Summary

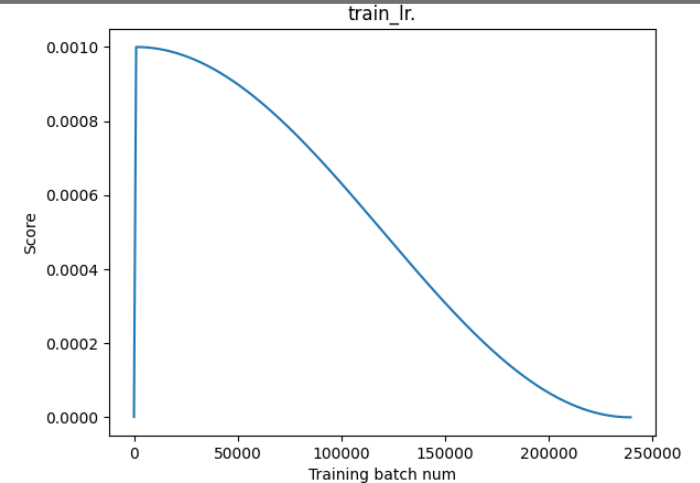
- Better than SpliceAI
- Generalize to non-coding RNA

$$LOSS_{CLE} = - \sum_{class \in \{donor, acceptor, neither\}} I_{class} \times \log(P_{class}) \quad (2)$$

$$LOSS_{FL} = - \sum_{class \in \{donor, acceptor, neither\}} I_{class} \times (1 - P_{class})^{\gamma} \times \log(P_{class}) \quad (3)$$

## Technical takeaways

- Focal loss improves cross entropy loss
- Learning rate warm up + sinusoidal decay
- Residual connection is powerful
- Grouped convolution helps (cardinality)



**Chao, K. H.**, Mao, A., Salzberg, S. L., & Pertea, M. (2023). Splam: a deep-learning-based splice site predictor that improves spliced alignments. **bioRxiv**. <https://ccb.jhu.edu/splam/>

# Future sequence models in genomics?

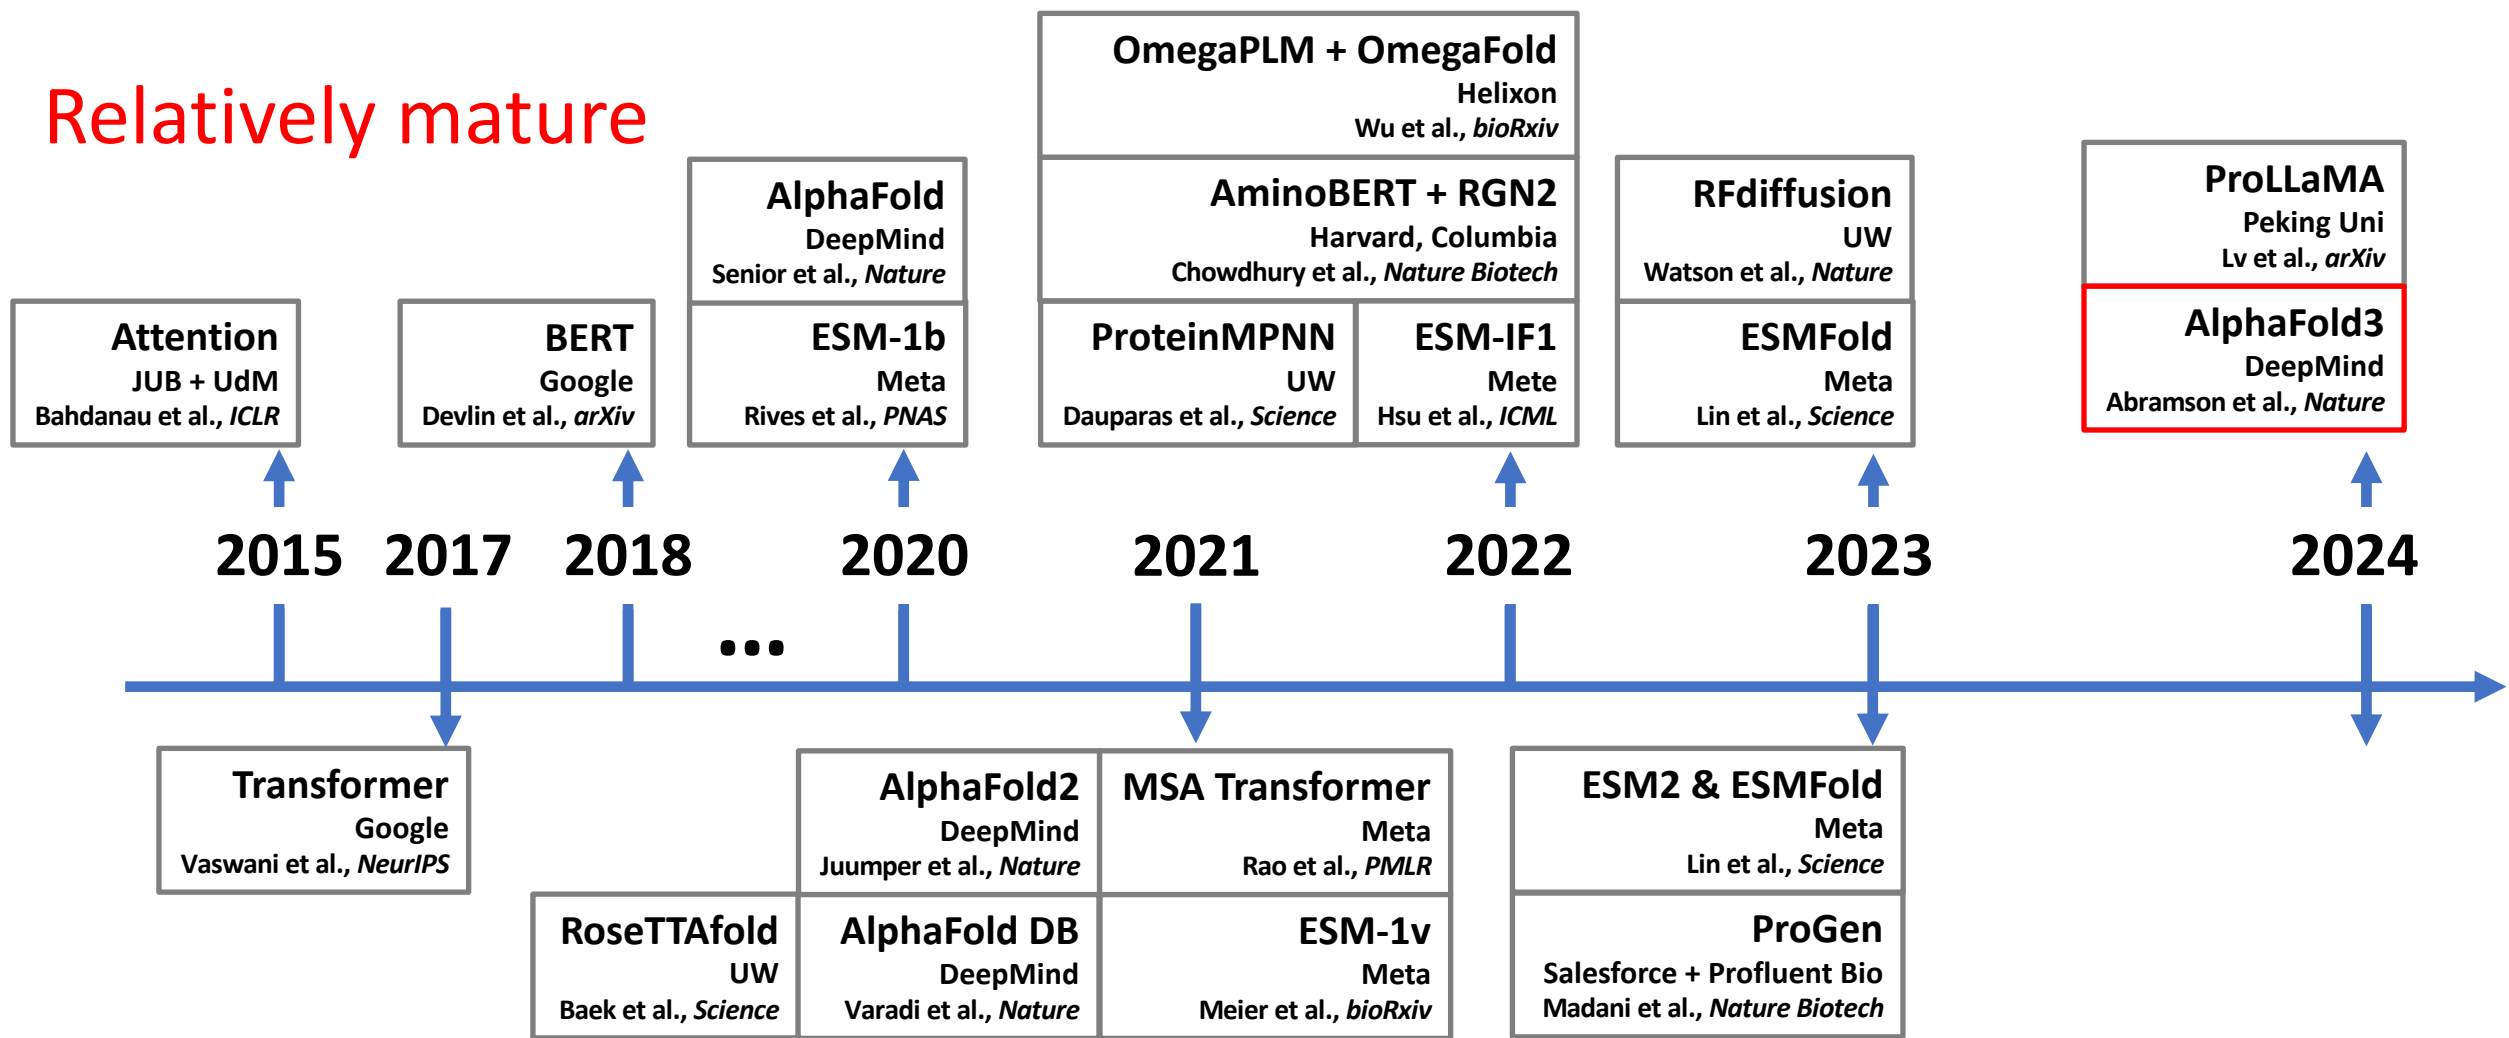
CNN or ?

# Future sequence models in genomics?

## CNN or **/and Transformer?**

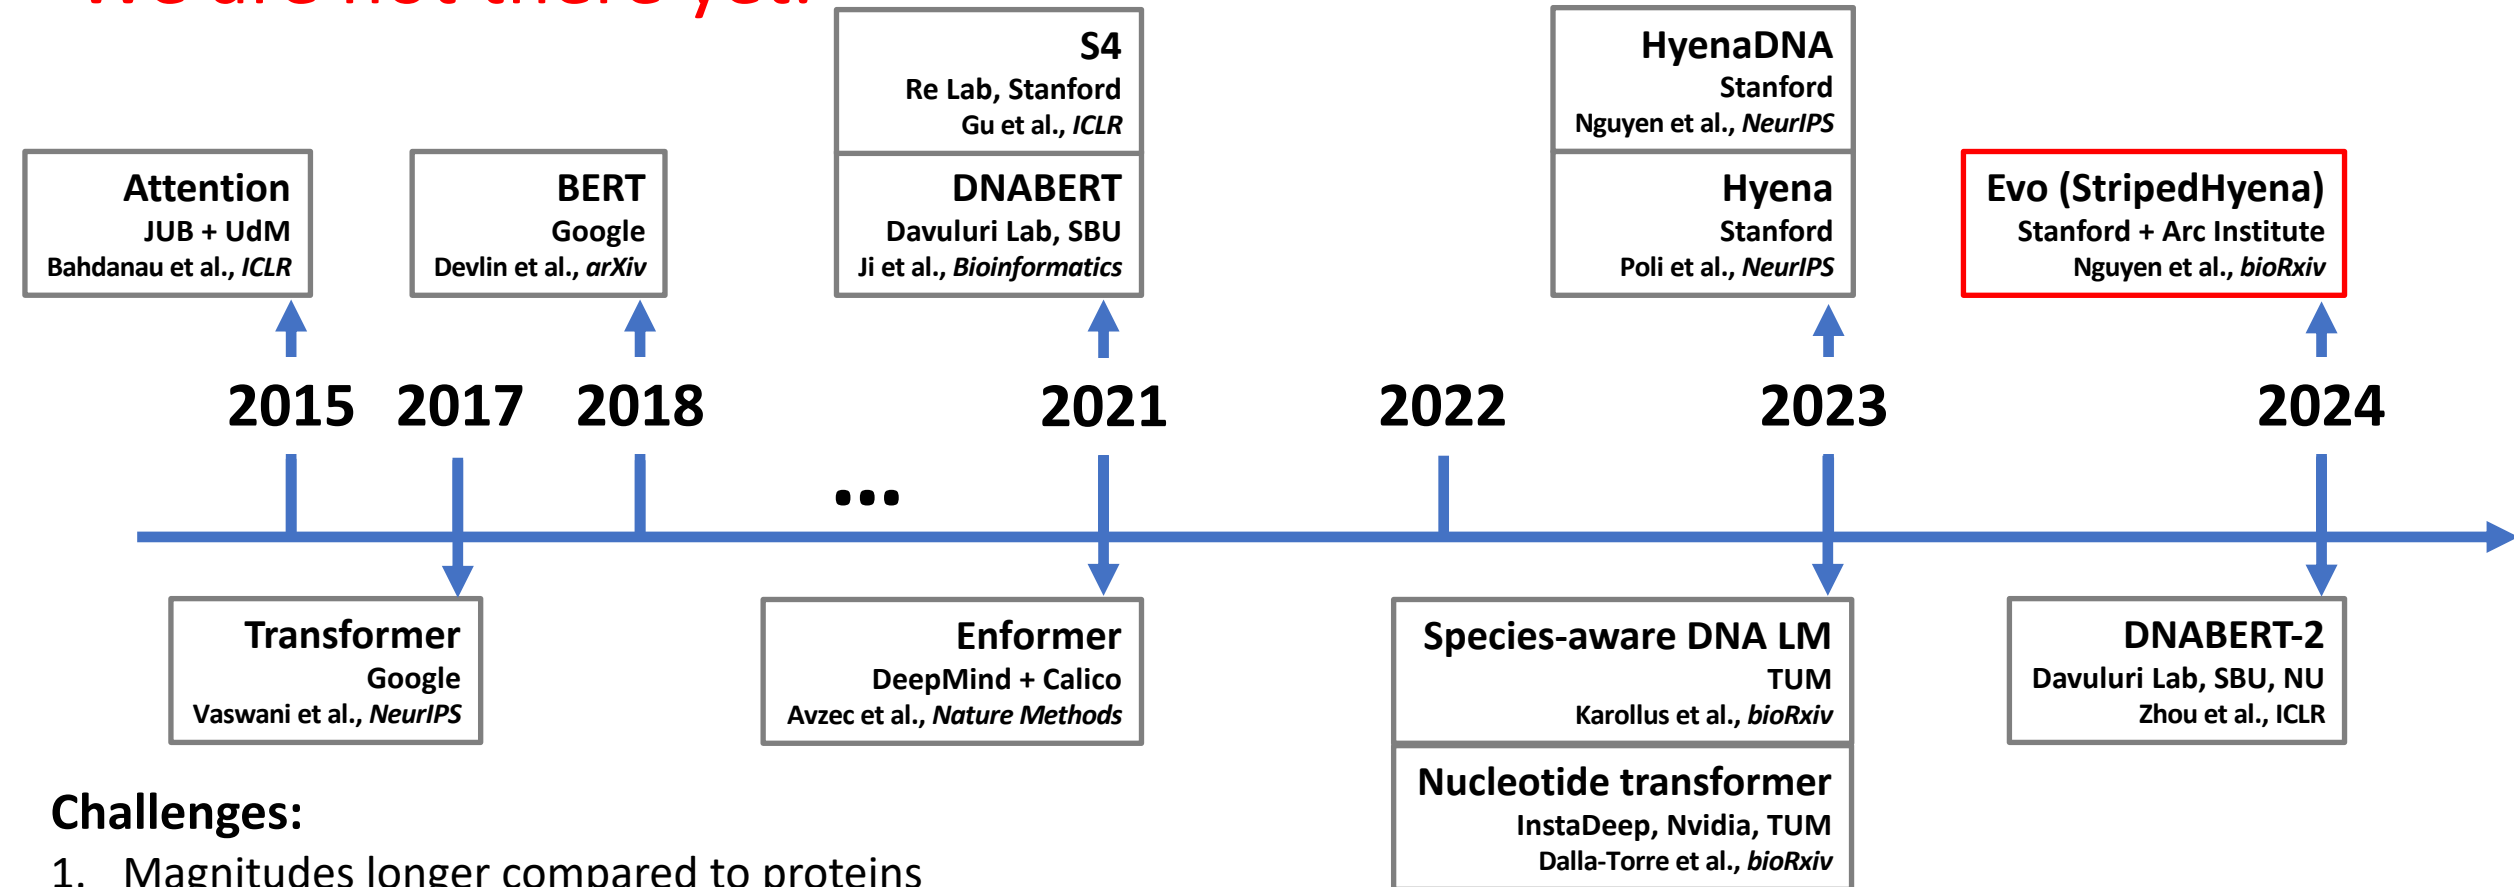
# Future? – Protein transformer-based models

Relatively mature



# Future? – DNA transformer-based models

We are not there yet.



## Challenges:

1. Magnitudes longer compared to proteins
2. Long-range dependencies & interactions spanning over 100k+ nt.



# Application? – Genome annotation

Home / A-Z Publications / Annual Review of Genomics and Human Genetics / Early Publication / Review in Advance

ANNUAL REVIEW OF GENOMICS AND HUMAN GENETICS

## Deep Learning Sequence Models for Transcriptional Regulation

Ksenia Sokolova<sup>1</sup>, Kathleen M. Chen<sup>1</sup>, Yun Hao<sup>2</sup>, Jian Zhou<sup>3</sup>, and Olga G. Troyanskaya<sup>1,2,4</sup>

### SegmentNT: annotating the genome at single-nucleotide resolution with DNA foundation models

 Bernardo P. de Almeida, Hugo Dalla-Torre, Guillaume Richard, Christopher Blum, Lorenz Hexemer, Maxence Gélard, Priyanka Pandey, Stefan Laurent, Alexandre Laterre, Maren Lang, Uğur Şahin, Karim Beguir,  Thomas Pierrot


doi: <https://doi.org/10.1101/2024.03.14.584712>



LiftOn



RECOMB-SEQ



[github.com/Kuanhao-Chao/LiftOn](https://github.com/Kuanhao-Chao/LiftOn)



[ccb.jhu.edu/lifton](https://ccb.jhu.edu/lifton) Preprint coming soon!

Target genome T



# Acknowledge



[khchao.com](http://khchao.com)



@KuanHaoChao



Kuanhao-Chao



Steven Salzberg



Mihaela Pertea



Anqi Liu




Alaina Shumate



Jakob Heinz



Celine Hoh



Alan Mao

- All members in Salzberg lab, Pertea lab
- All friends at Malone & CCB
- All friends at JHU Computational Biology Group

*“If you think of mathematics as the perfect description language for physics, then AI might be the perfect one for biology.”*

Demis Hassabis, CEO of DeepMind, 2022