Dynamic Boltzmann Machine

IBM Research – Tokyo

Takayuki Osogami

osogami@jp.ibm.com

How can we make effective use of spike-timing dependent plasticity (STDP) in artificial neural networks?

Hebb's rule ('49) Cells that fire together, wire together

Today's artificial neural networks

STDP ('90s) Amount of changes depends on timing of spikes

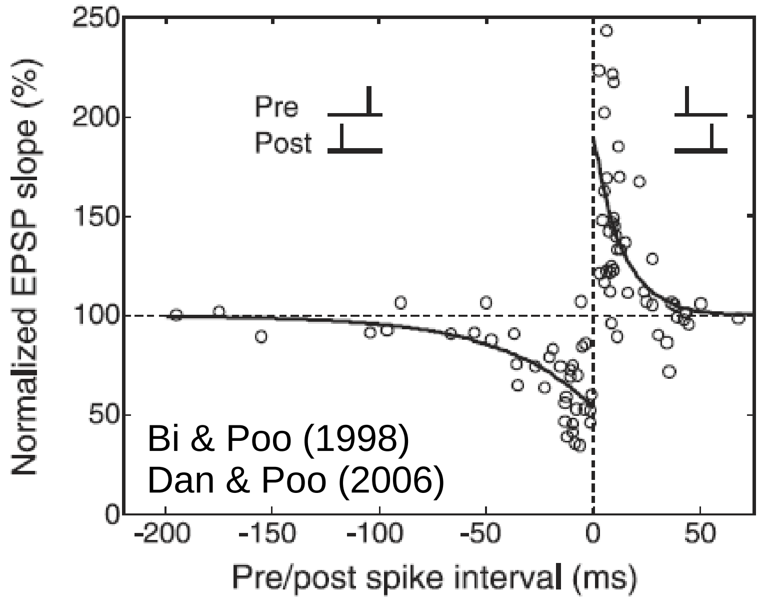

?

[Nessler et al. 2013,
Bengio et al. 2016,
Scellier & Bengio 2016]

Takayuki Osogami and Makoto Otsuka, “Seven neurons memorizing sequences of alphabetical images via spike-timing dependent plasticity,” *Scientific Reports,* **5**, 14149 (2015).
www.nature.com/articles/srep14149

This talk

- Boltzmann machine, Hebb's rule, STDP
- Dynamic Boltzmann machine [O & Otsuka (2015)]
- Experiments [O & Otsuka (2015), Dasgupta & O (2017)]

Boltzmann machine

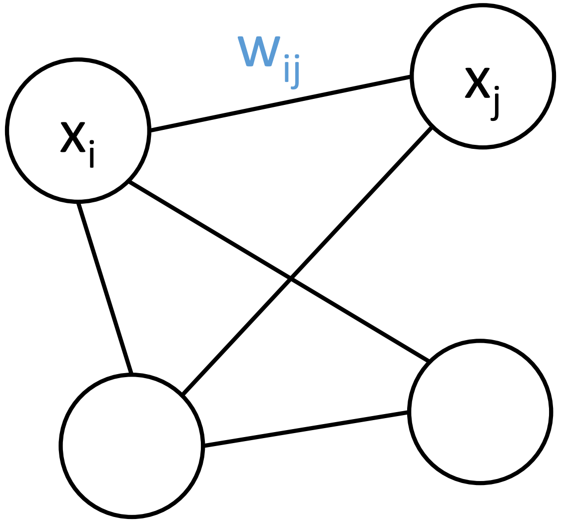

Energy of being **x**:

$$E(\boldsymbol{x}) = -\boldsymbol{x}^\top \boldsymbol{W} \boldsymbol{x}$$

Probability of being **x**:

$$P(\boldsymbol{x}) = \frac{1}{Z} \exp\big(-E(\boldsymbol{x})\big)$$

$$Z \equiv \sum_{\widetilde{\boldsymbol{x}}} \exp(-E(\widetilde{\boldsymbol{x}}))$$

Parameters: $\boldsymbol{W} = (w_{ij})$

Variables: **x** = $(x_1, x_2, \ldots)$

Learning rule of Boltzmann machine, maximizing log-likelihood [Hinton et al. '83]

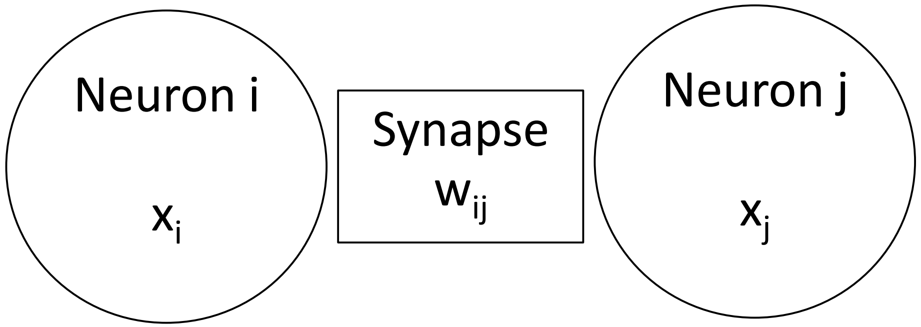

Log likelihood of training data D:

$$LL(D) = \sum_{\boldsymbol{x} \in D} \log(P(\boldsymbol{x}))$$

$$\frac{\partial}{\partial w_{ij}} LL(D) = \sum_{\boldsymbol{x} \in D} \left(x_i \; x_j - \langle X_i \; X_j \rangle \right)$$

Stochastic gradient

$$w_{ij} \leftarrow w_{ij} + \eta\left(x_i \; x_j - \langle X_i \; X_j \rangle \right)$$

cf. Hebb's rule

Expected value:

$$\langle X_i \; X_j \rangle \equiv \sum_{\tilde{\boldsymbol{x}}} P(\tilde{\boldsymbol{x}}) \; \tilde{x}_i \; \tilde{x}_j$$

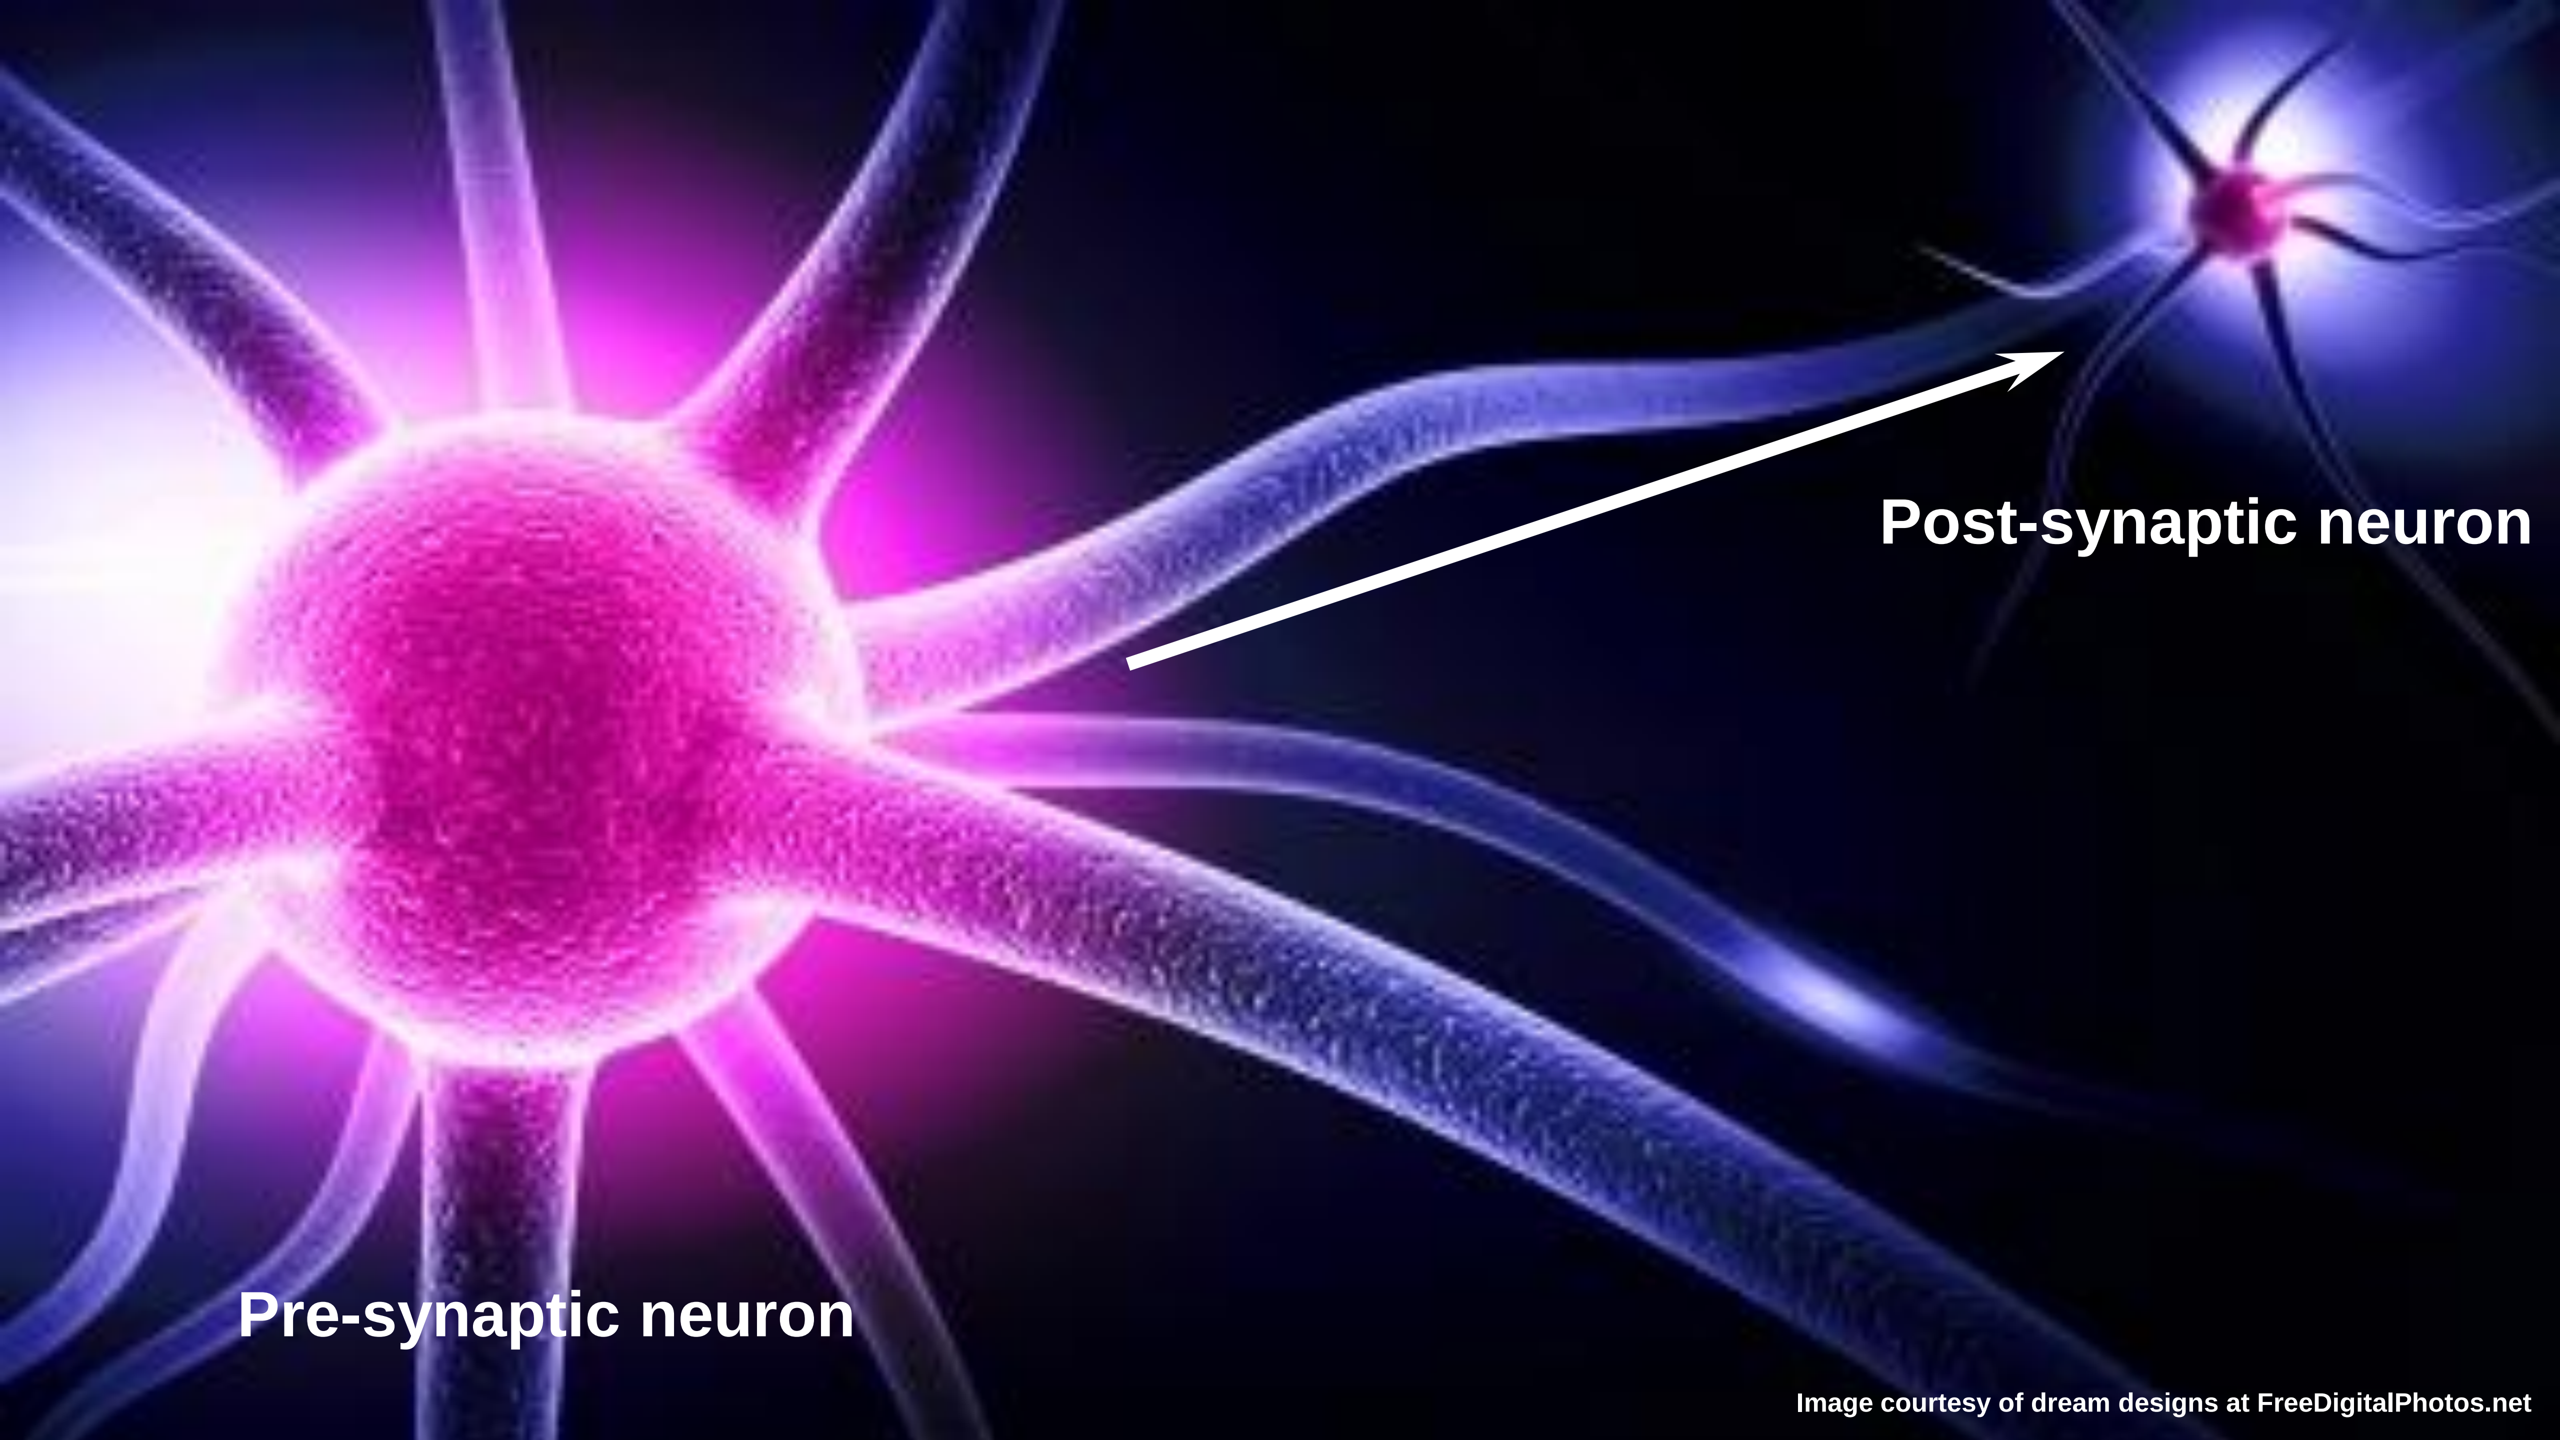
Post-synaptic neuron
Pre-synaptic neuron
Image courtesy of dream designs at FreeDigitalPhotos.net

Spike-timing dependent plasticity (STDP): Amount of changes depends on timing of spikes

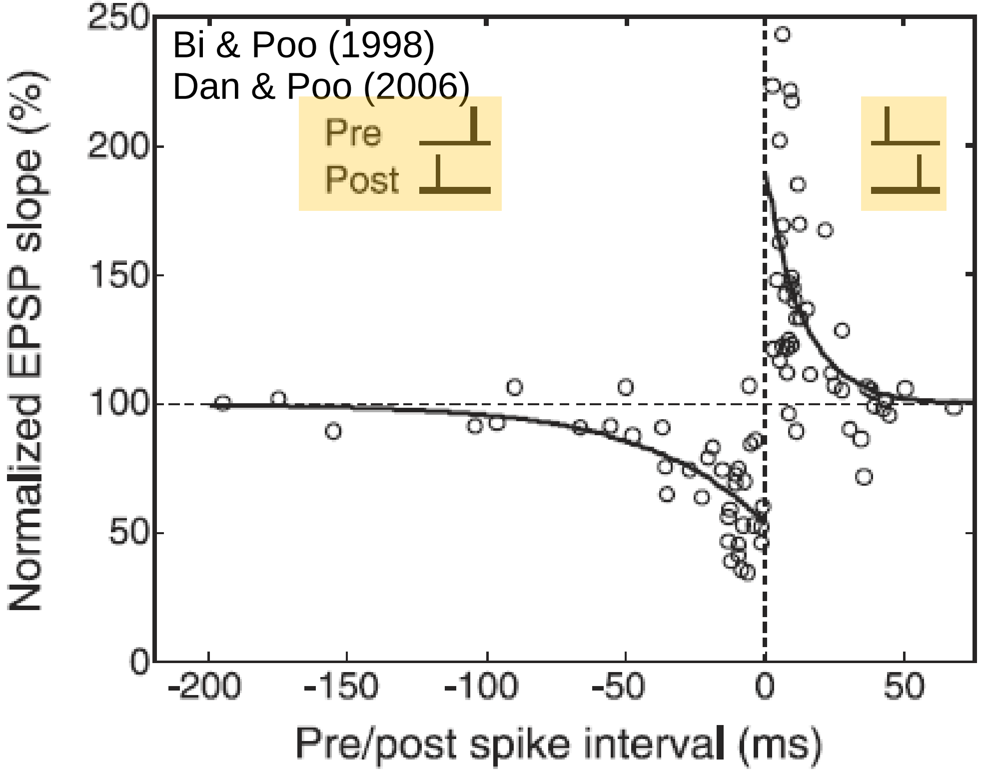

Synapse weakened
(Long Term Depression)

Synapse strengthened
(Long Term Potentiation)

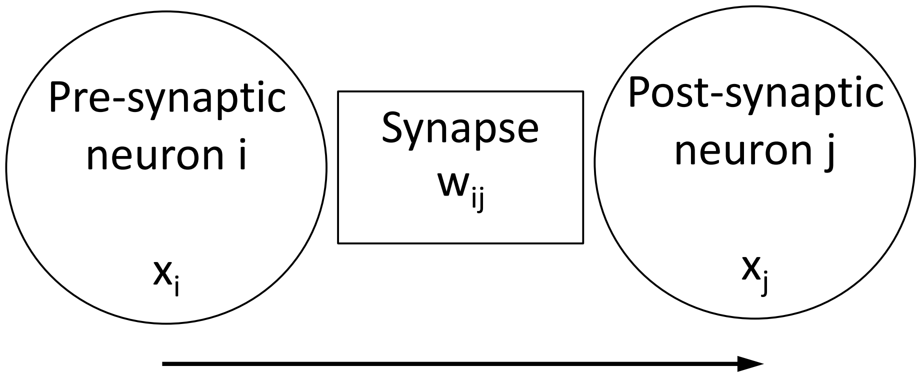

We will construct a dynamic Boltzmann machine as a limit of a sequence of Boltzmann machines

Boltzmann machine for a T-th order Markov model

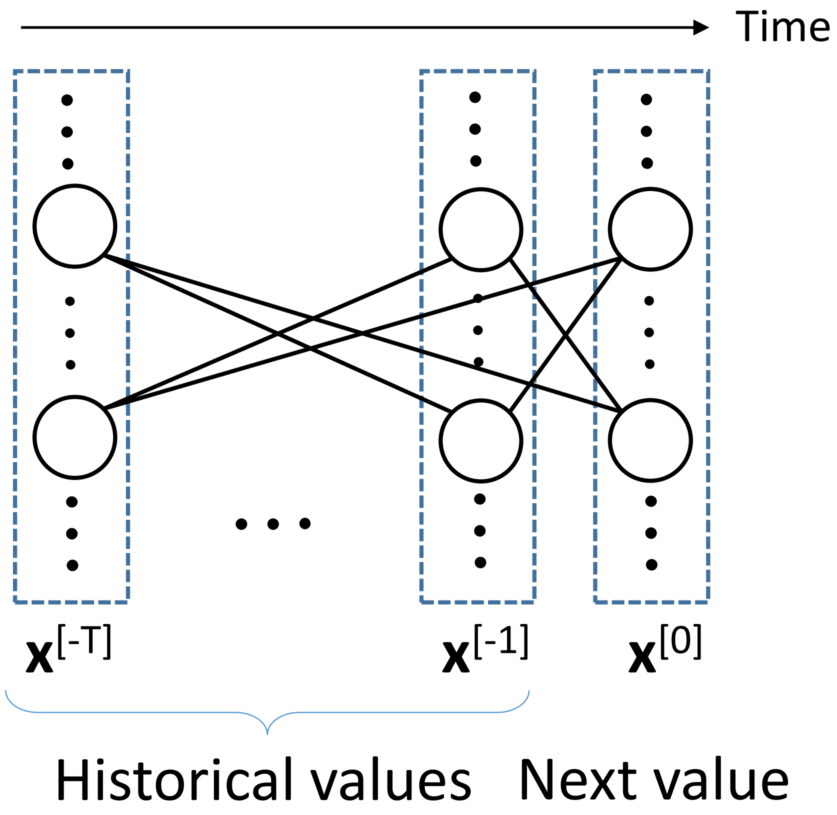

$T \to \infty$

Particular parametric form of weight

Dynamic Boltzmann machine

This talk

- Boltzmann machine, Hebb's rule, STDP
- Dynamic Boltzmann machine (DyBM) [O & Otsuka (2015)]
  - DyBM with LTP
  - DyBM with LTP and LTD
- Experiments [O & Otsuka (2015), Dasgupta & O (2017)]

LTP: Weight is strengthened when the post-synaptic neuron fires shortly after the pre-synaptic neuron

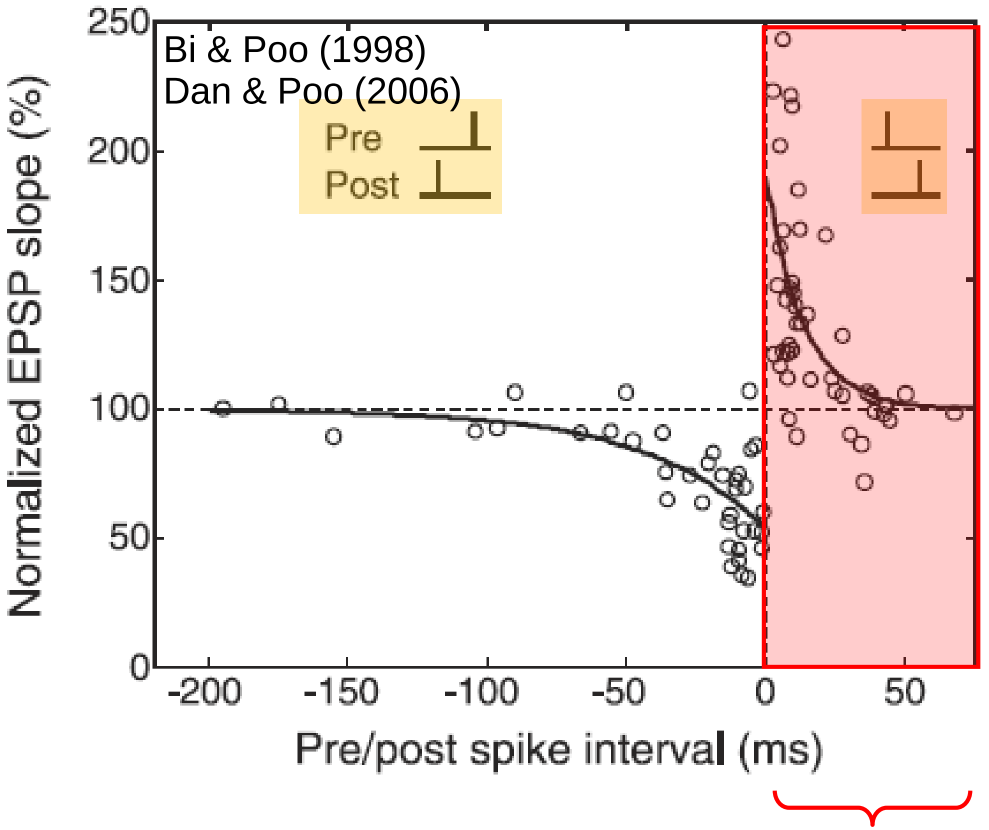

Synapse strengthened
(Long Term Potentiation)

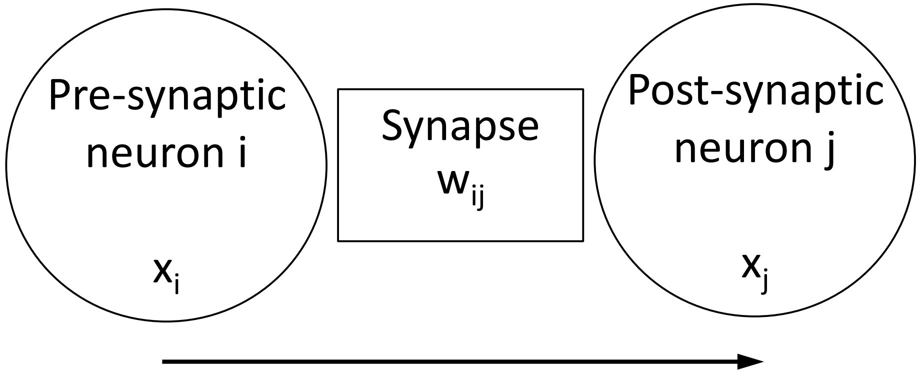

Dynamic Boltzmann machine as a limit of a sequence of Boltzmann machines

Boltzmann machine for a
T-th order Markov model

$T \to \infty$

Dynamic Boltzmann machine

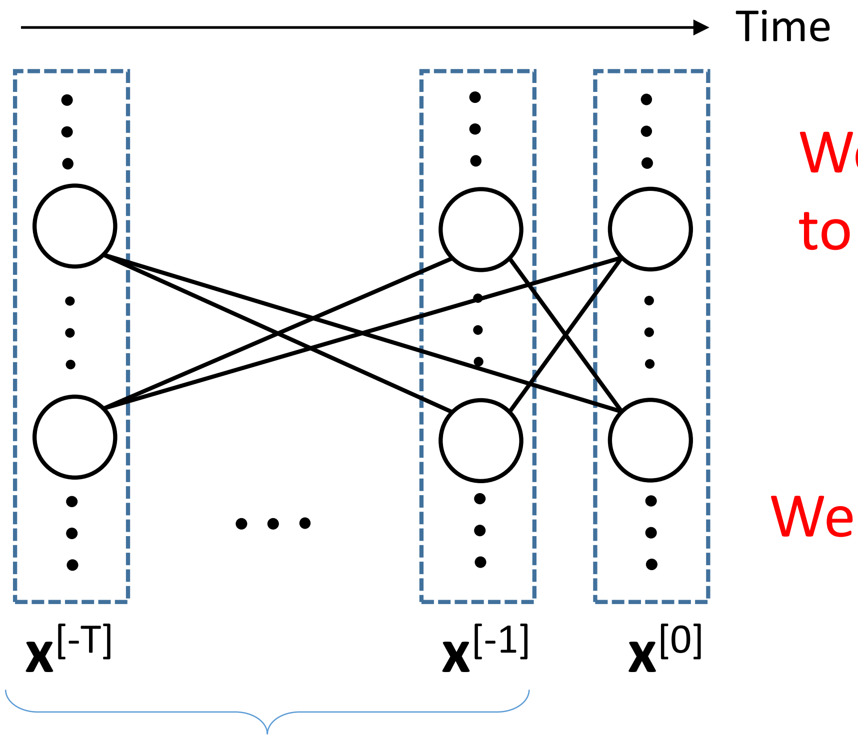

Weight from neuron i at time t
to neuron j at time t + δ

We learn u_{ij}

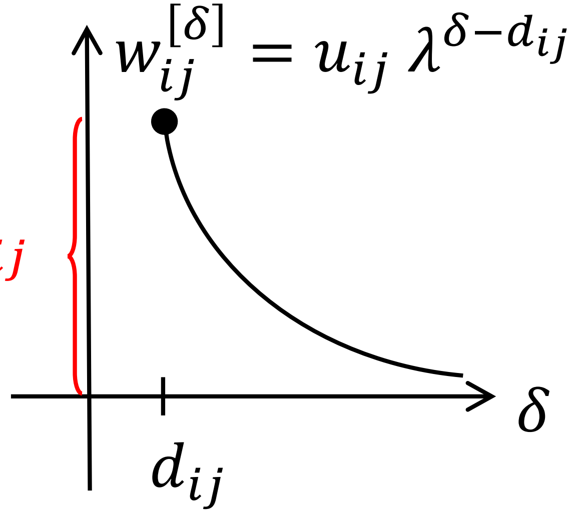

Dynamic Boltzmann machine (LTP only)

Conduction delay, d_{ij}

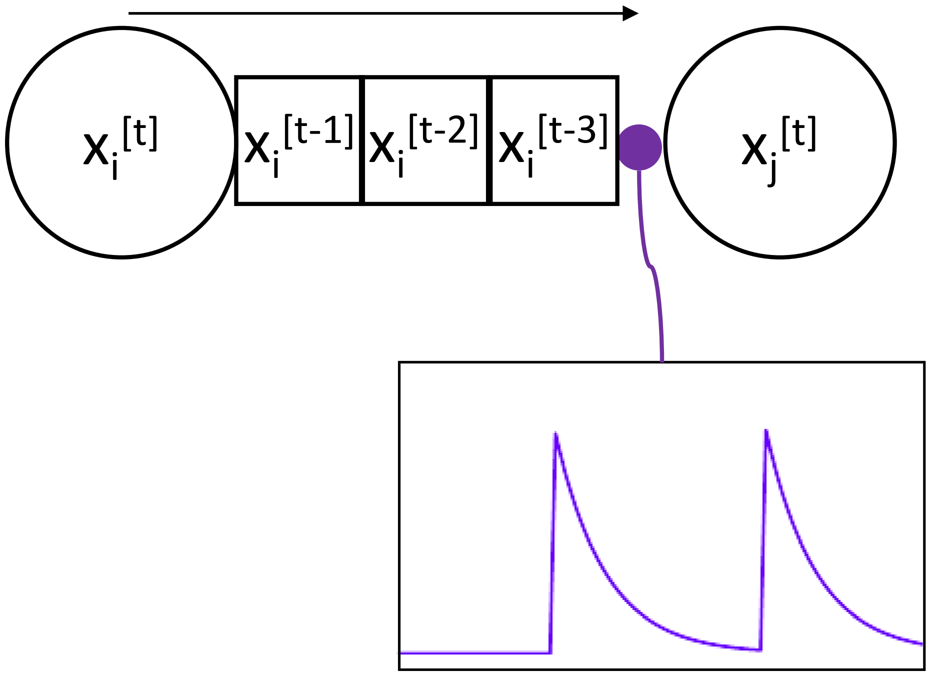

Synaptic eligibility trace:

$$\alpha_{ij}^{[t-1]} \equiv \sum_{s<t-d_{ij}} \lambda^{t-s-d_{ij}} \, x_i^{[s]}$$

$$\alpha_{ij}^{[t]} \leftarrow \lambda \left(\alpha_{ij}^{[t-1]} + x_i^{[t-d_{ij}]} \right)$$

Probability for neuron j to fire at time t:

$$\langle X_j^{[t]} \rangle \equiv P\left(x_j^{[t]} = 1 | x^{[:t-1]} \right)$$

$$= \frac{1}{1 + \exp\left(- \sum_i u_{ij} \, \alpha_{ij}^{[t-1]} \right)}$$

Learning rule of DyBM, maximizing log-likelihood

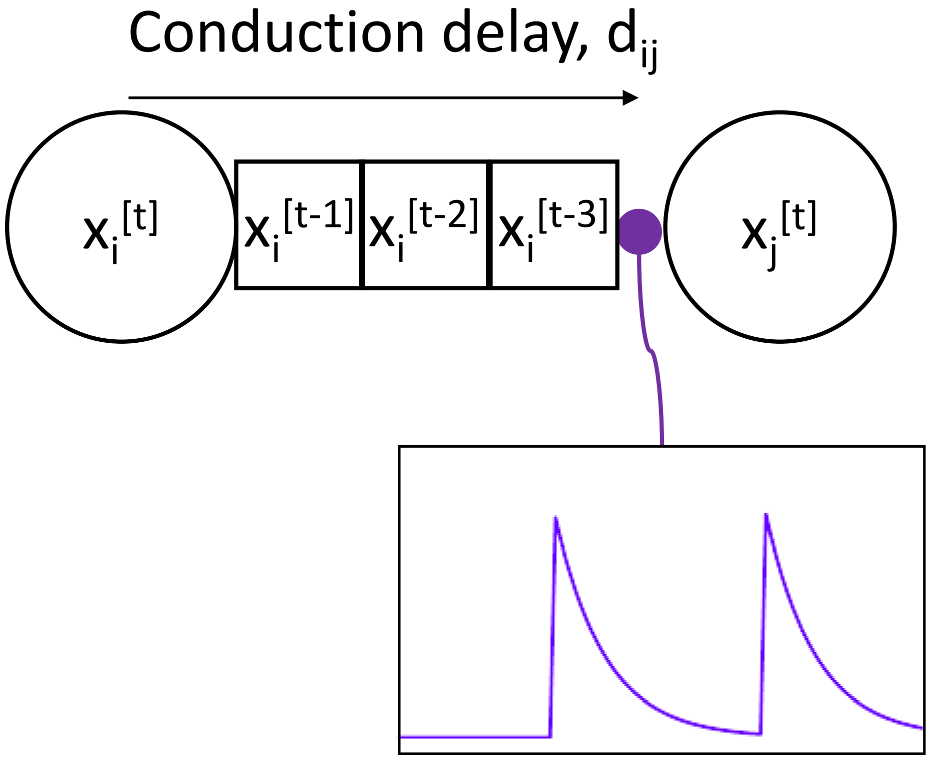

Synaptic eligibility trace:

$$\alpha_{ij}^{[t-1]} \equiv \sum_{s<t-d_{ij}} \lambda^{t-s-d_{ij}} \, x_i^{[s]}$$

$$\alpha_{ij}^{[t]} \leftarrow \lambda \left(\alpha_{ij}^{[t-1]} + x_i^{[t-d_{ij}]} \right)$$

$$\frac{\partial}{\partial u_{ij}} \log P\left(x^{[t]} \middle| x^{[:t-1]}\right) = \alpha_{ij}^{[t-1]} \left(x_j^{[t]} - \langle X_j^{[t]} \rangle \right)$$

Stochastic gradient for LTP weight:

Spike-timing dependent

$$u_{ij} \leftarrow u_{ij} + \eta \, \alpha_{ij}^{[t-1]} \left(x_j^{[t]} - \langle X_j^{[t]} \rangle \right)$$

How recently/often spikes reached from neuron i

cf. Boltzmann machine

$$w_{ij} \leftarrow w_{ij} + \eta (x_i \, x_j - \langle X_i \, X_j \rangle)$$

This talk

- Boltzmann machine, Hebb’s rule, STDP
- Dynamic Boltzmann machine (DyBM) [O & Otsuka (2015)]
  - DyBM with LTP
  - DyBM with LTP and LTD
- Experiments [O & Otsuka (2015), Dasgupta & O (2017)]

LTP: Weight is strengthened when the pre-synaptic neuron fires shortly after the post-synaptic neuron

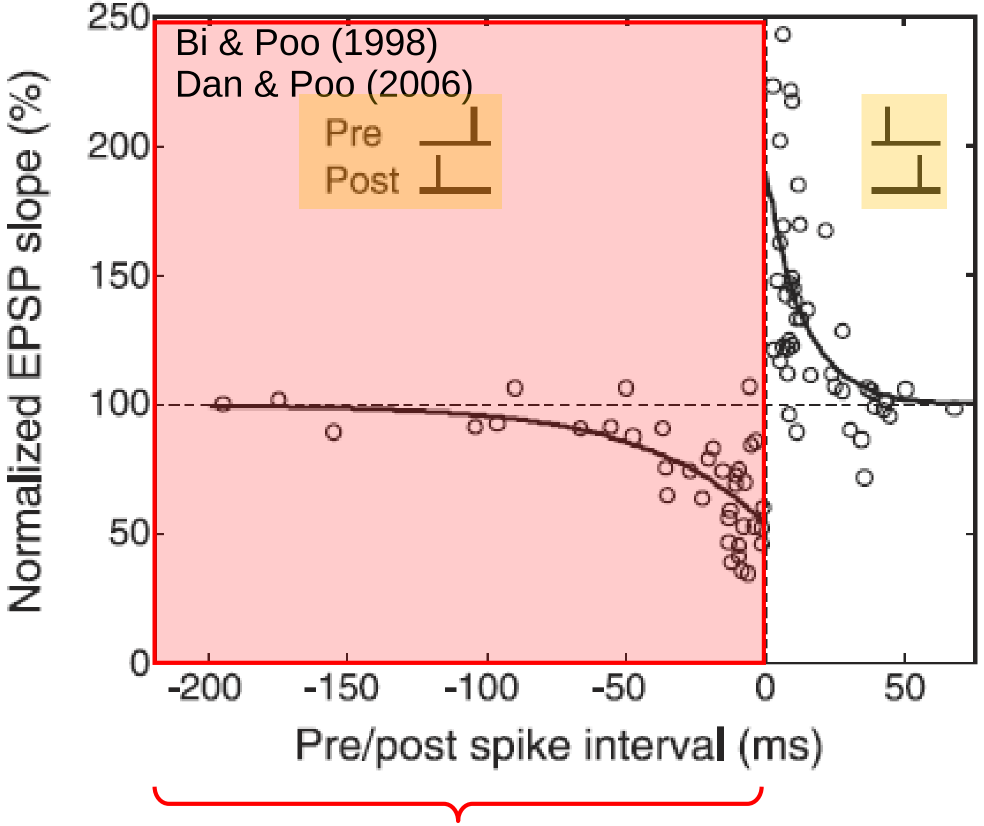

Synapse weakened
(Long Term Depression)

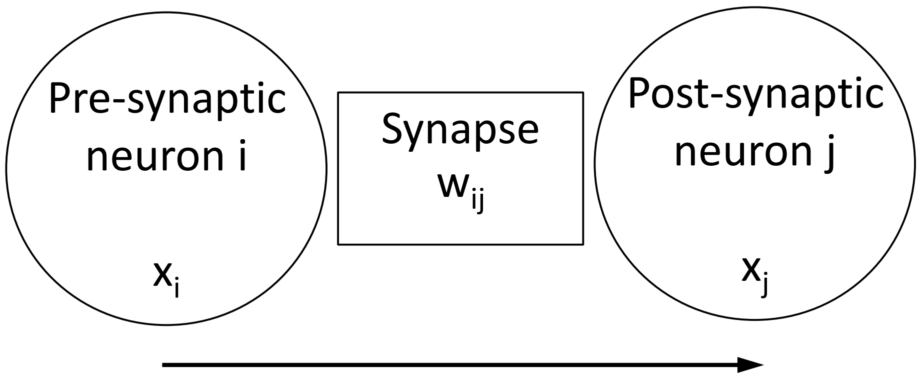

Interpreting LTD as negative LTP for a connection from post-synaptic neuron to pre-synaptic neuron

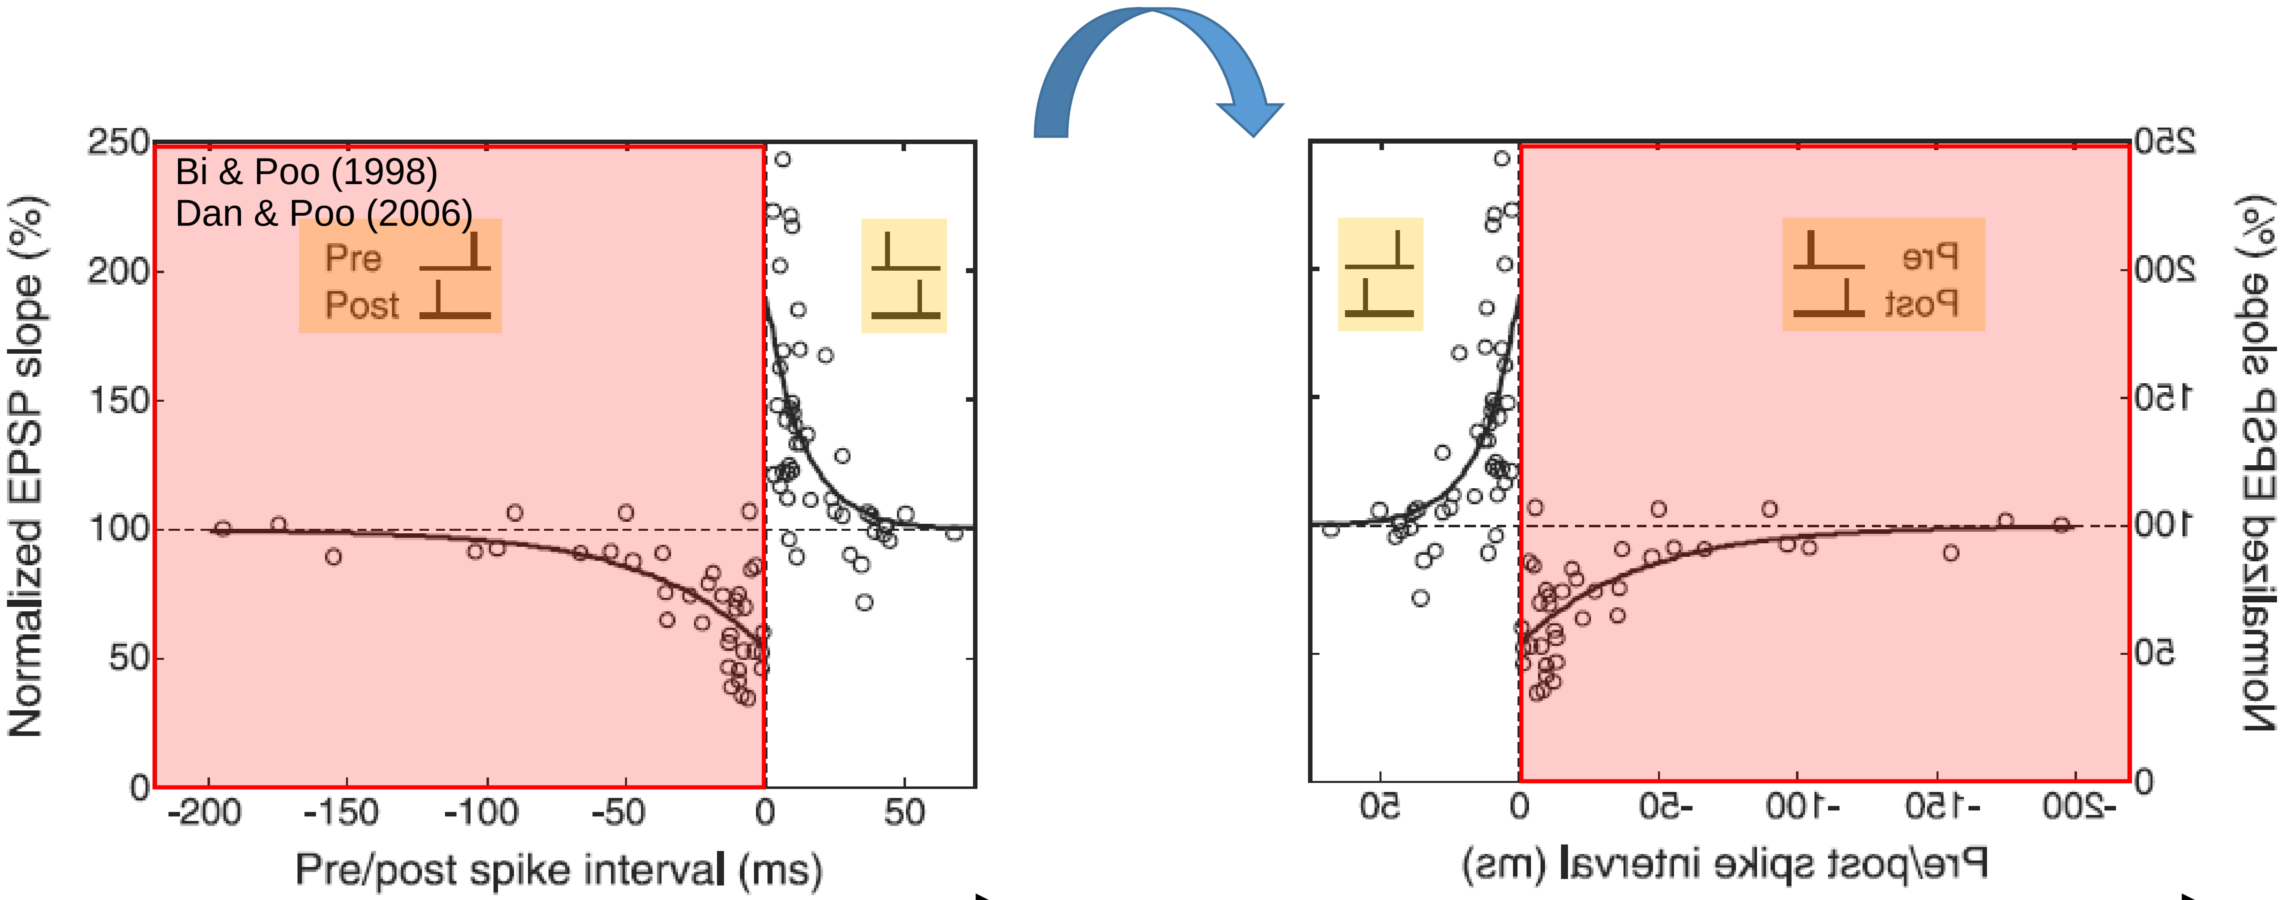

Time from pre-spike to post-spike

Time from post-spike to pre-spike

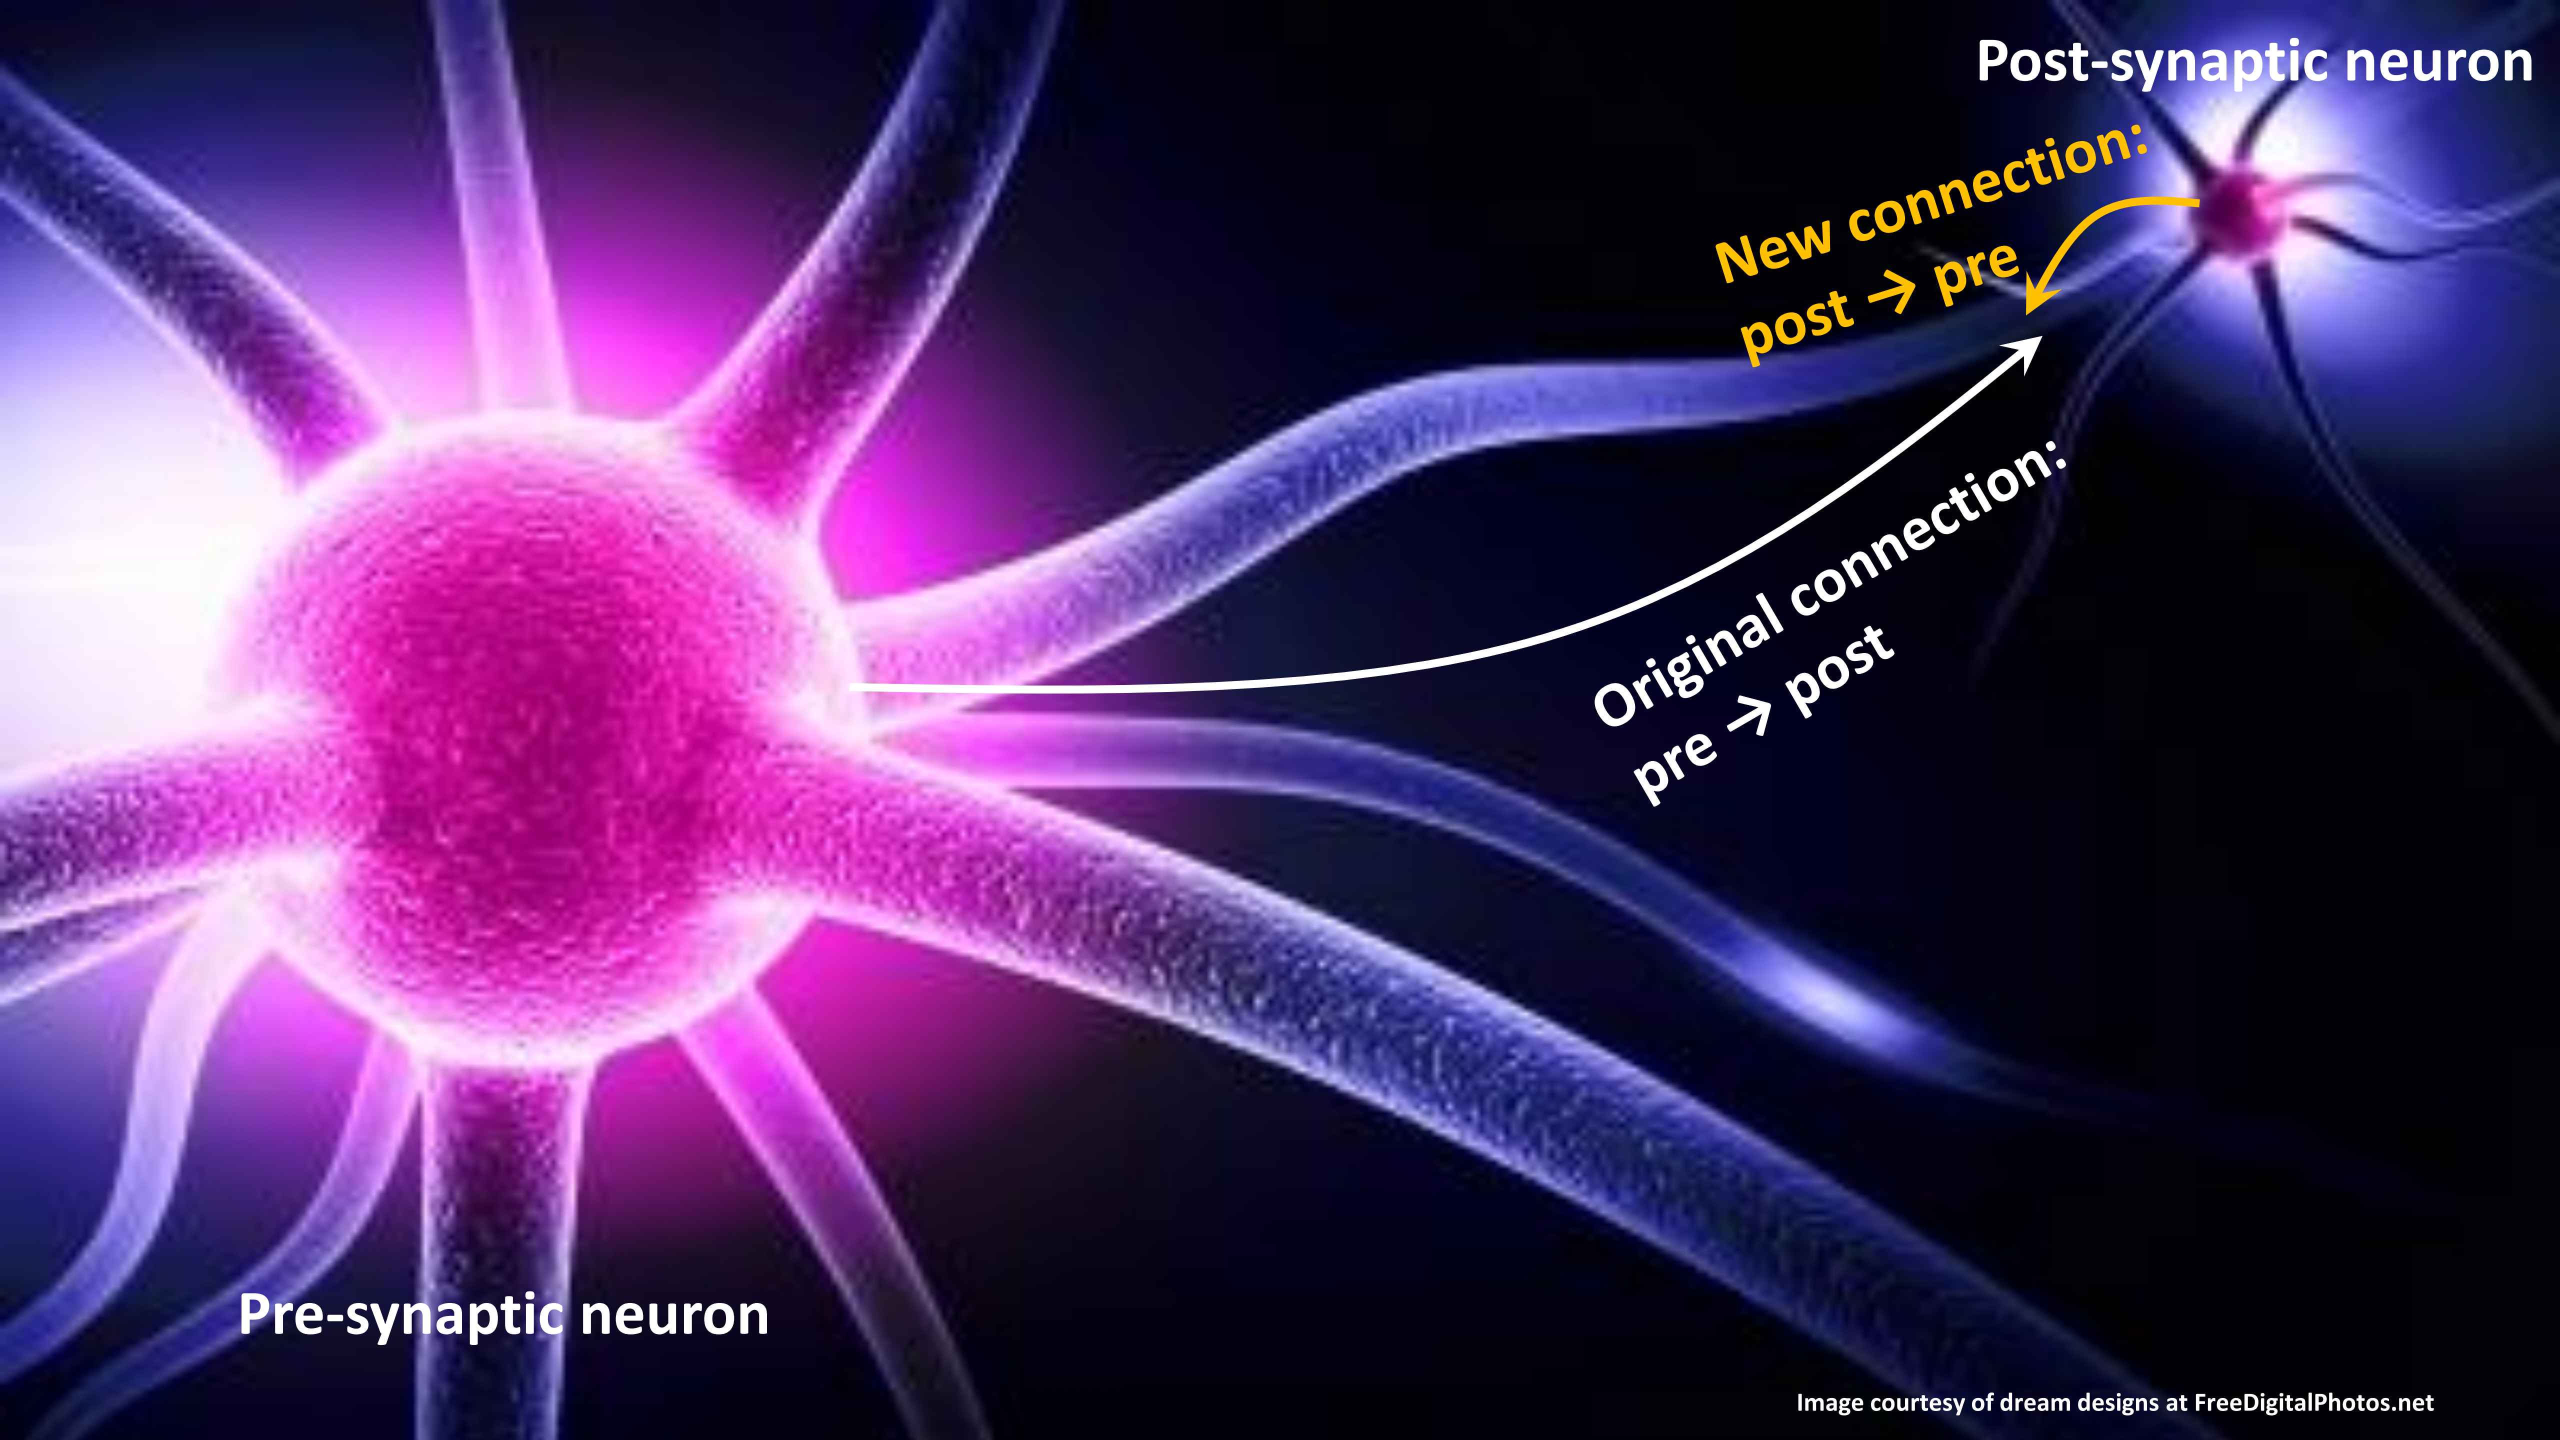
Post-synaptic neuron
New connection:
post → pre
Original connection:
pre → post
Pre-synaptic neuron
Image courtesy of dream designs at FreeDigitalPhotos.net

Dynamic Boltzmann machine as a limit of a sequence of Boltzmann machines

Boltzmann machine for a T-th order Markov model

$T \to \infty$

Dynamic Boltzmann machine

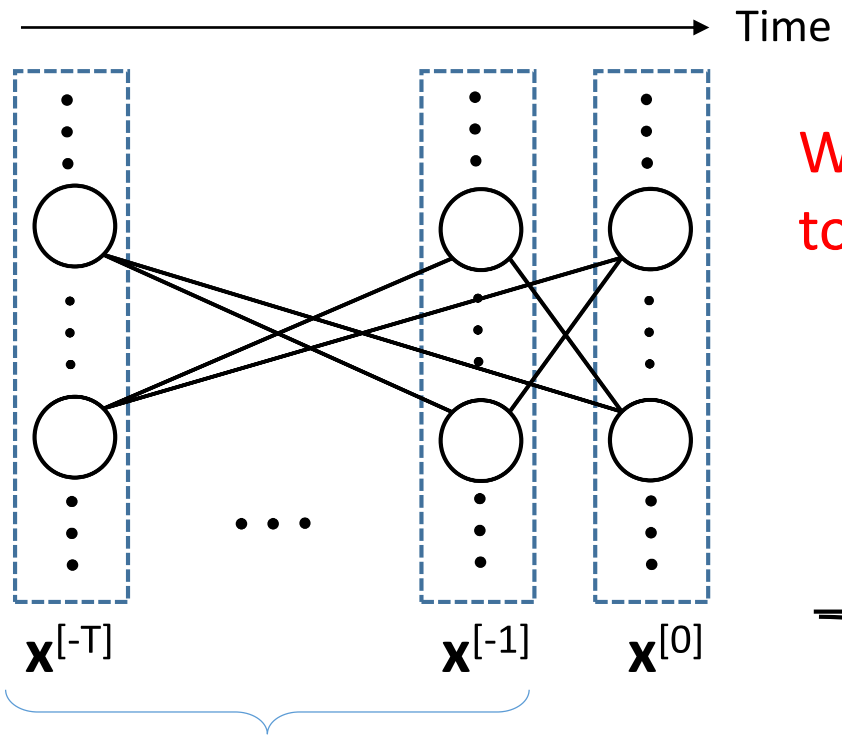

Weight from neuron i at time t to neuron j at time $t + \delta$

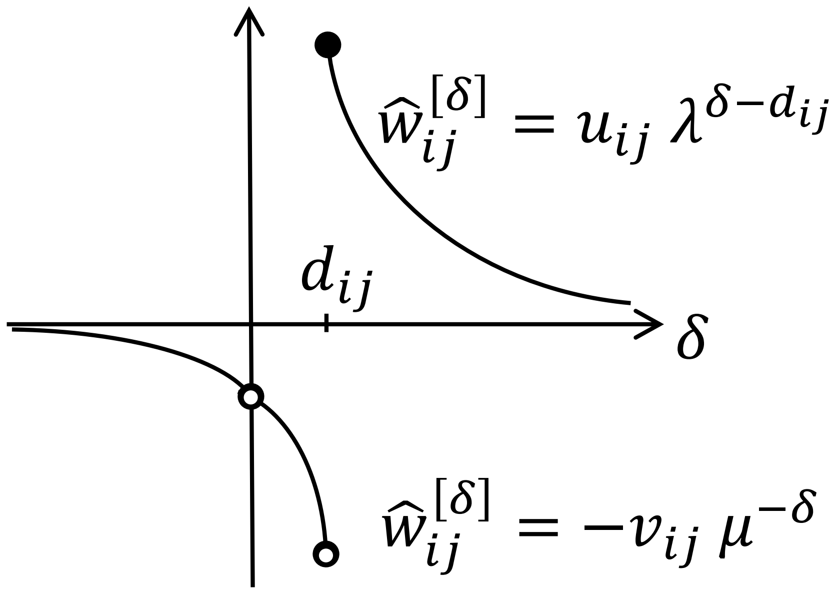

Dynamic Boltzmann machine as a limit of a sequence of Boltzmann machines

Boltzmann machine for a T-th order Markov model

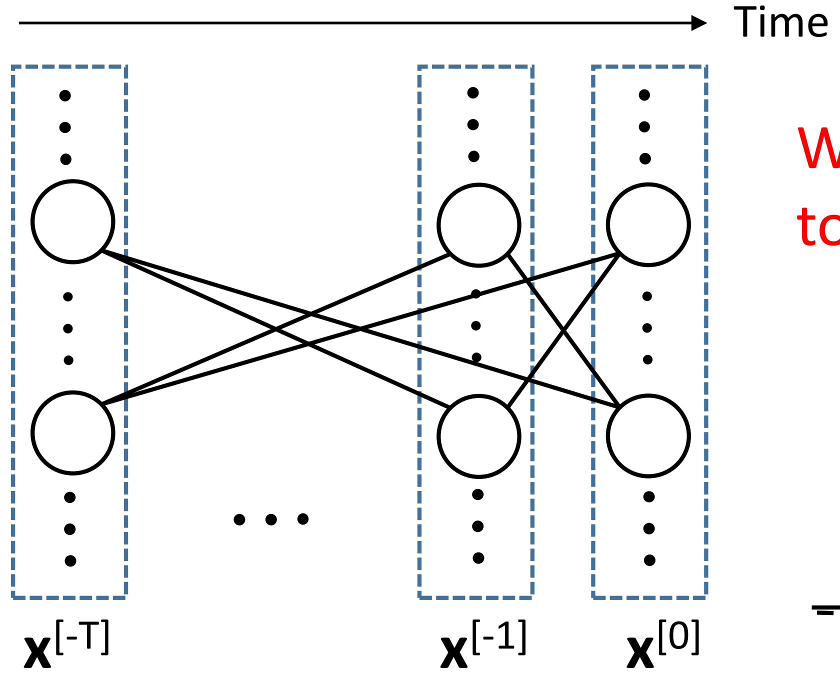

Historical values Next value

$T \to \infty$

Dynamic Boltzmann machine

Weight from neuron i at time t to neuron j at time t + δ

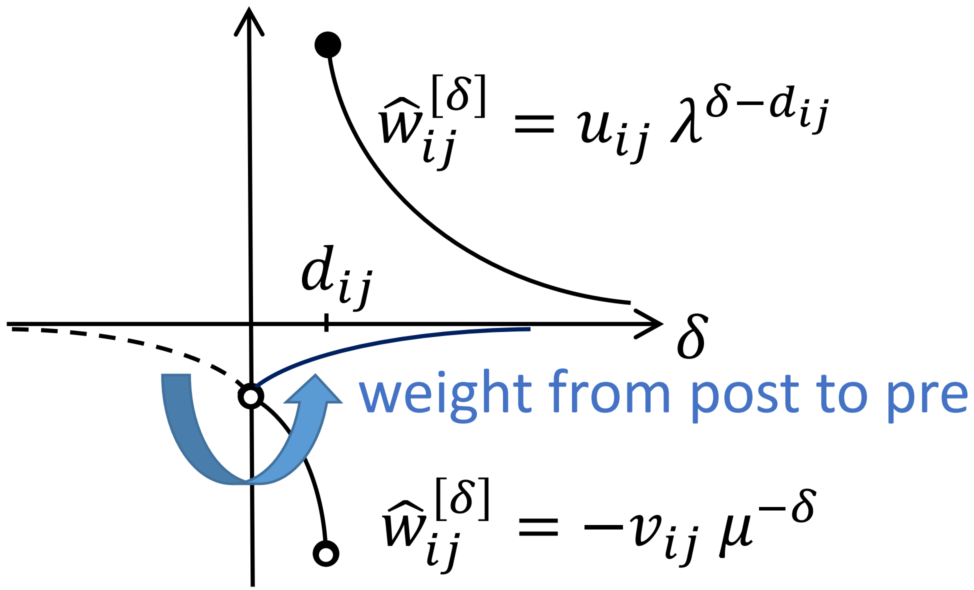

Dynamic Boltzmann machine (LTP & LTD)

Neural eligibility trace:

$$\gamma_i^{[t]} \leftarrow \mu \left(\gamma_i^{[t-1]} + x_i^{[t]} \right)$$

Conduction delay, d_{ij}

Synaptic eligibility trace:

$$\alpha_{ij}^{[t]} \leftarrow \lambda \left(\alpha_{ij}^{[t-1]} + x_i^{[t-d_{ij}]} \right)$$

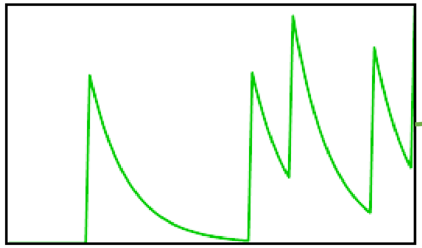

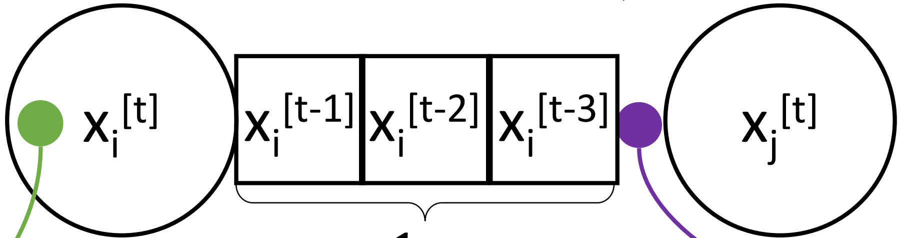

$$\beta_{ij}^{[t-1]} \equiv \sum_{s=-d_{ij}+1}^{-1} x_i^{[t+s]} \mu^s$$

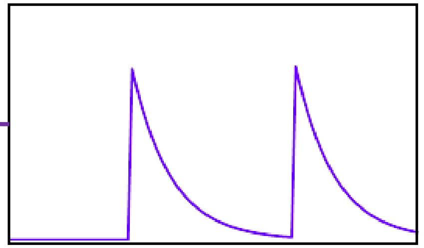

Probability for neuron j to fire at time t:

$$\langle X_j^{[t]} \rangle \equiv P\left(x_j^{[t]} = 1 | x^{[:t-1]} \right) = \frac{1}{1 + \exp\left(\sum_i \left(-u_{ij}\, \alpha_{ij}^{[t-1]} + v_{ij}\, \beta_{ij}^{[t-1]} + v_{ji}\, \gamma_i^{[t-1]} \right) \right)}$$

Learning rule for Dynamic Boltzmann machine (LTP & LTD)

Neural eligibility trace:

$$\gamma_i^{[t]} \leftarrow \mu \left(\gamma_i^{[t-1]} + x_i^{[t]} \right)$$

Synaptic eligibility trace:

$$\alpha_{ij}^{[t]} \leftarrow \lambda \left(\alpha_{ij}^{[t-1]} + x_i^{[t-d_{ij}]} \right)$$

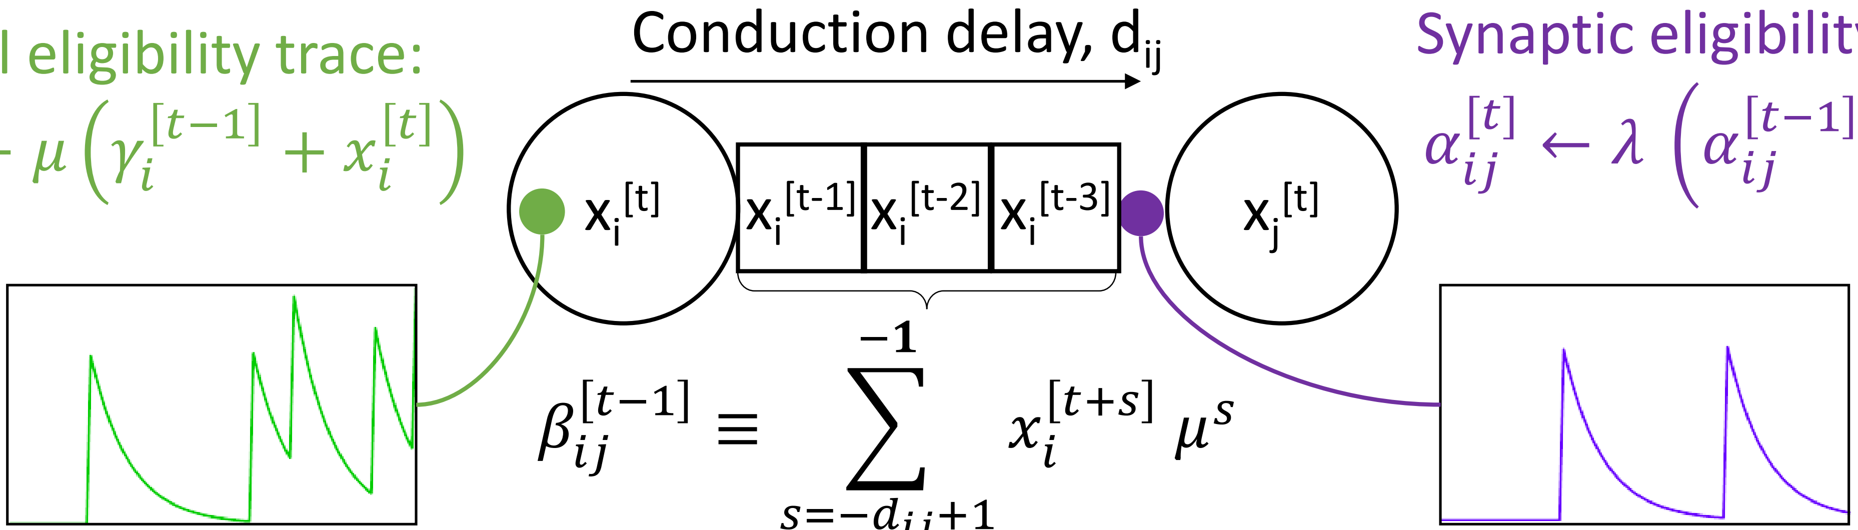

$$\beta_{ij}^{[t-1]} \equiv \sum_{s=-d_{ij}+1}^{-1} x_i^{[t+s]} \mu^s$$

Stochastic gradient for LTD weight:

$$v_{ij} \leftarrow v_{ij} + \eta\, \beta_{ij}^{[t-1]} \left(\langle X_j^{[t]} \rangle - x_j^{[t]} \right) + \eta\, \gamma_j^{[t-1]} \left(\langle X_i^{[t]} \rangle - x_i^{[t]} \right)$$

Stochastic gradient for LTP weight:

$$u_{ij} \leftarrow u_{ij} + \eta\, \alpha_{ij}^{[t-1]} \left(x_j^{[t]} - \langle X_j^{[t]} \rangle \right)$$

We can use multiple eligibility traces with varying decay rates for long term memory

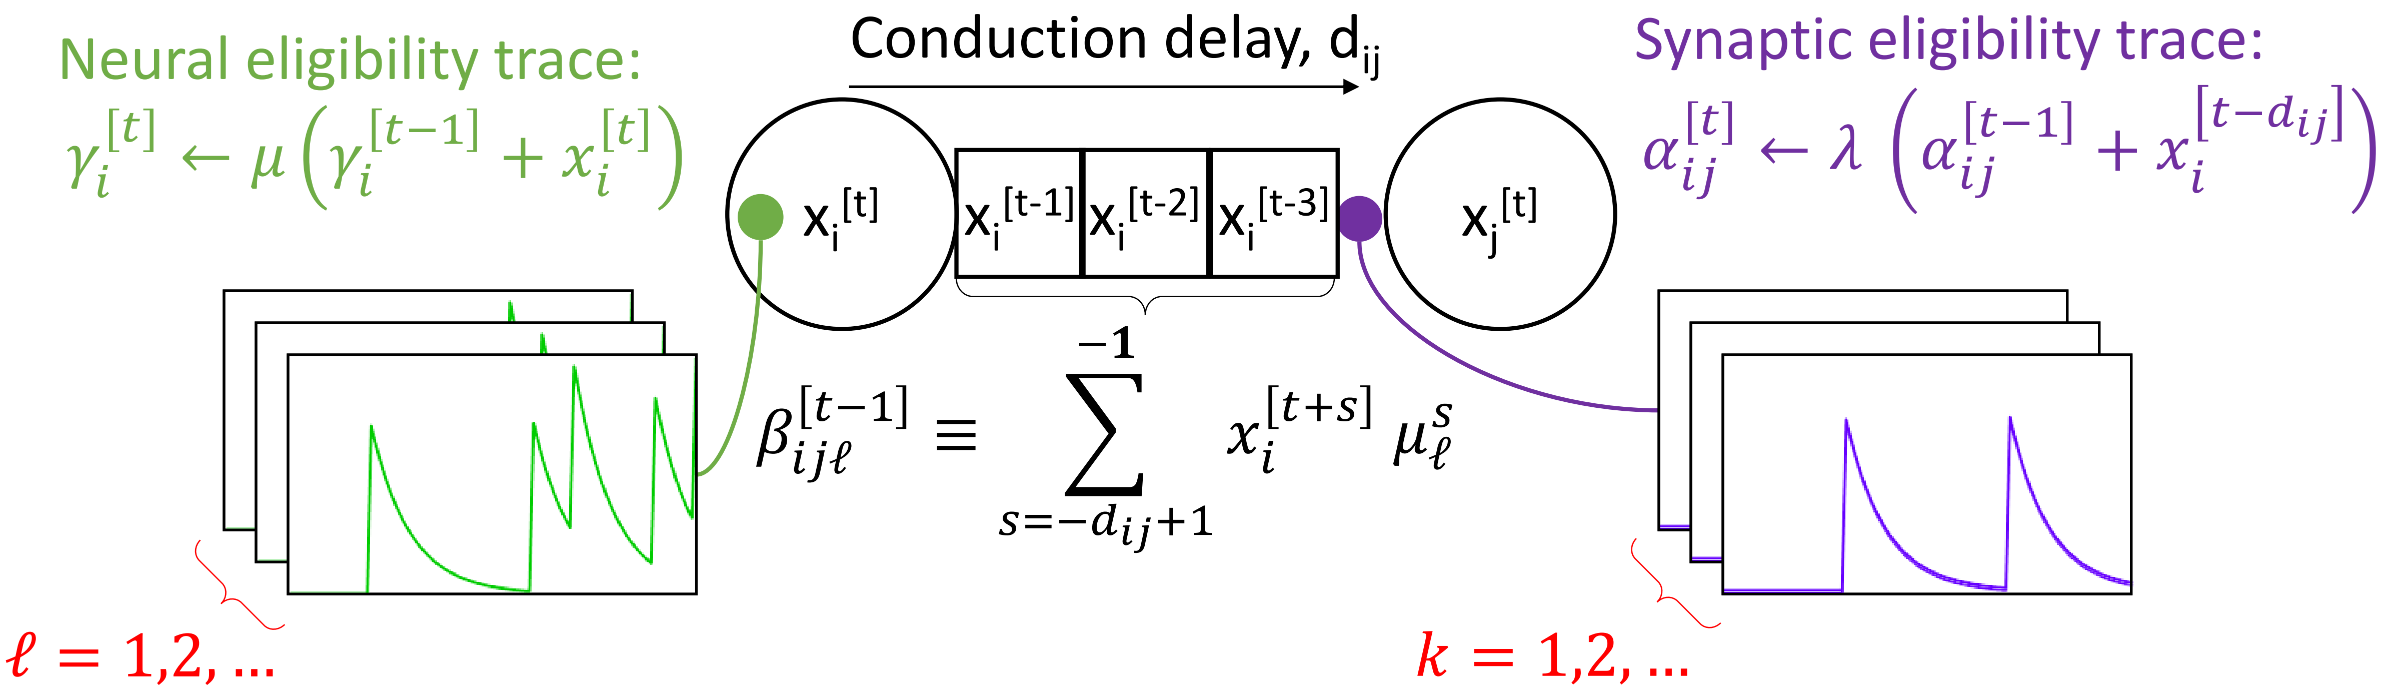

This talk

- Boltzmann machine, Hebb’s rule, STDP
- Dynamic Boltzmann machine [O & Otsuka (2015)]
- Experiments [O & Otsuka (2015), Dasgupta & O (2017)]

Presenting sequences of 7-bit patterns to a dynamic Boltzmann machine

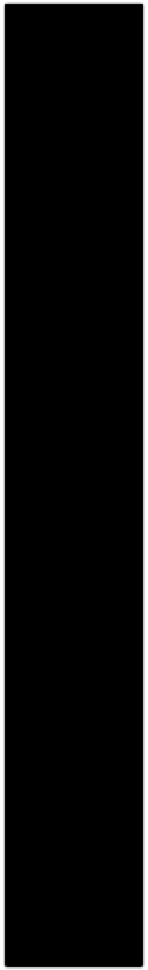

DyBM as associative memory for sequential patterns [O & Otsuka (2015)]

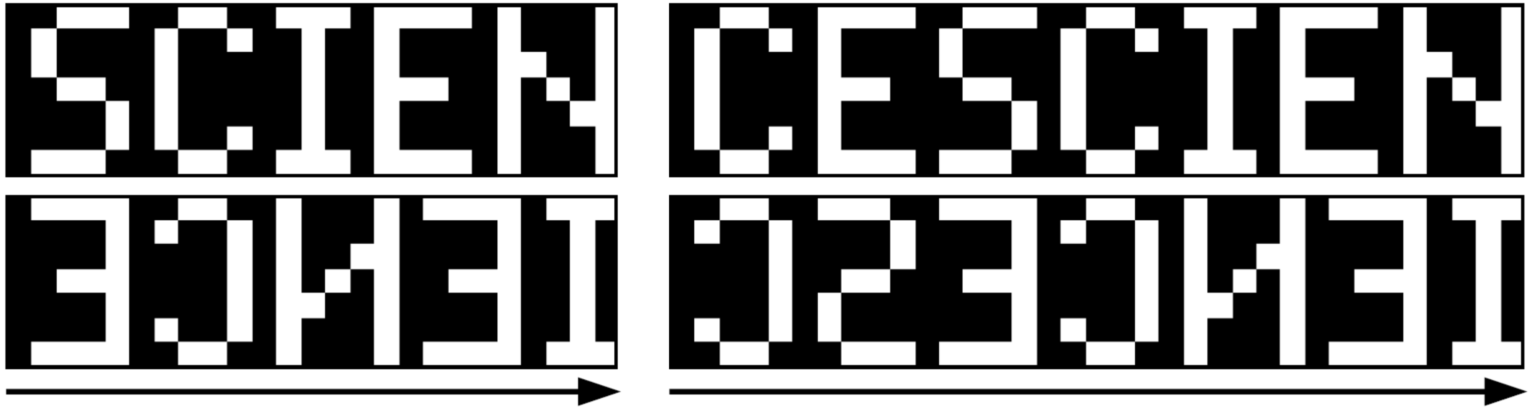

Giving DyBM a partial sequence as a cue

DyBM completes the remaining sequence

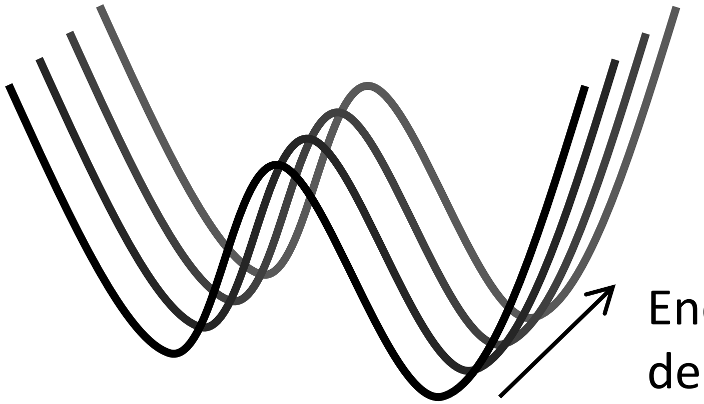

Energy landscape of DyBM evolves depending on the past patterns

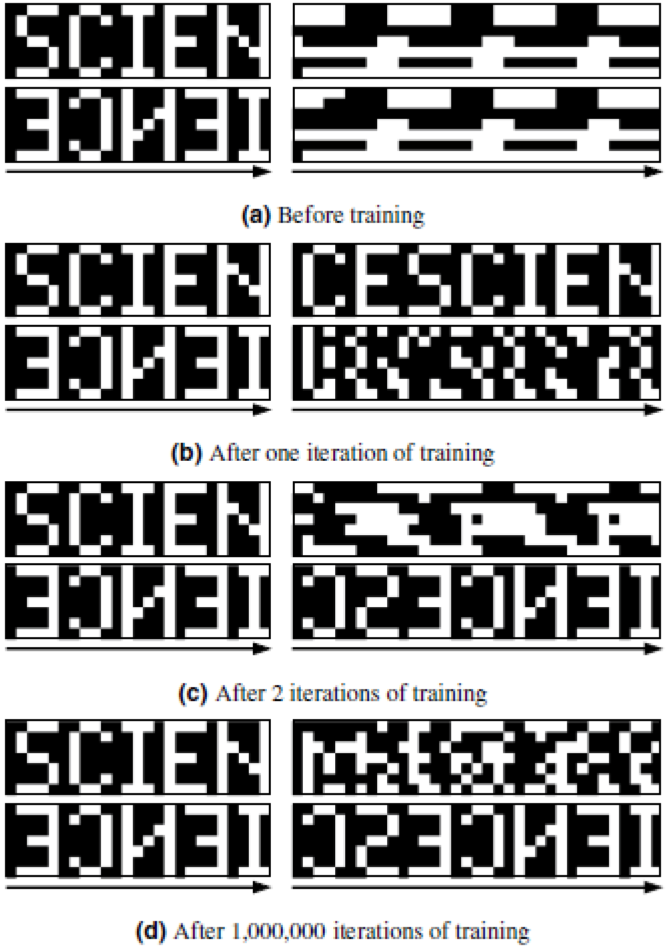

Images from www.nature.com/articles/srep14149

DyBM detecting anomaly [O & Otsuka (2015)]

Image from www.nature.com/articles/srep14149

Negative log-likelihood

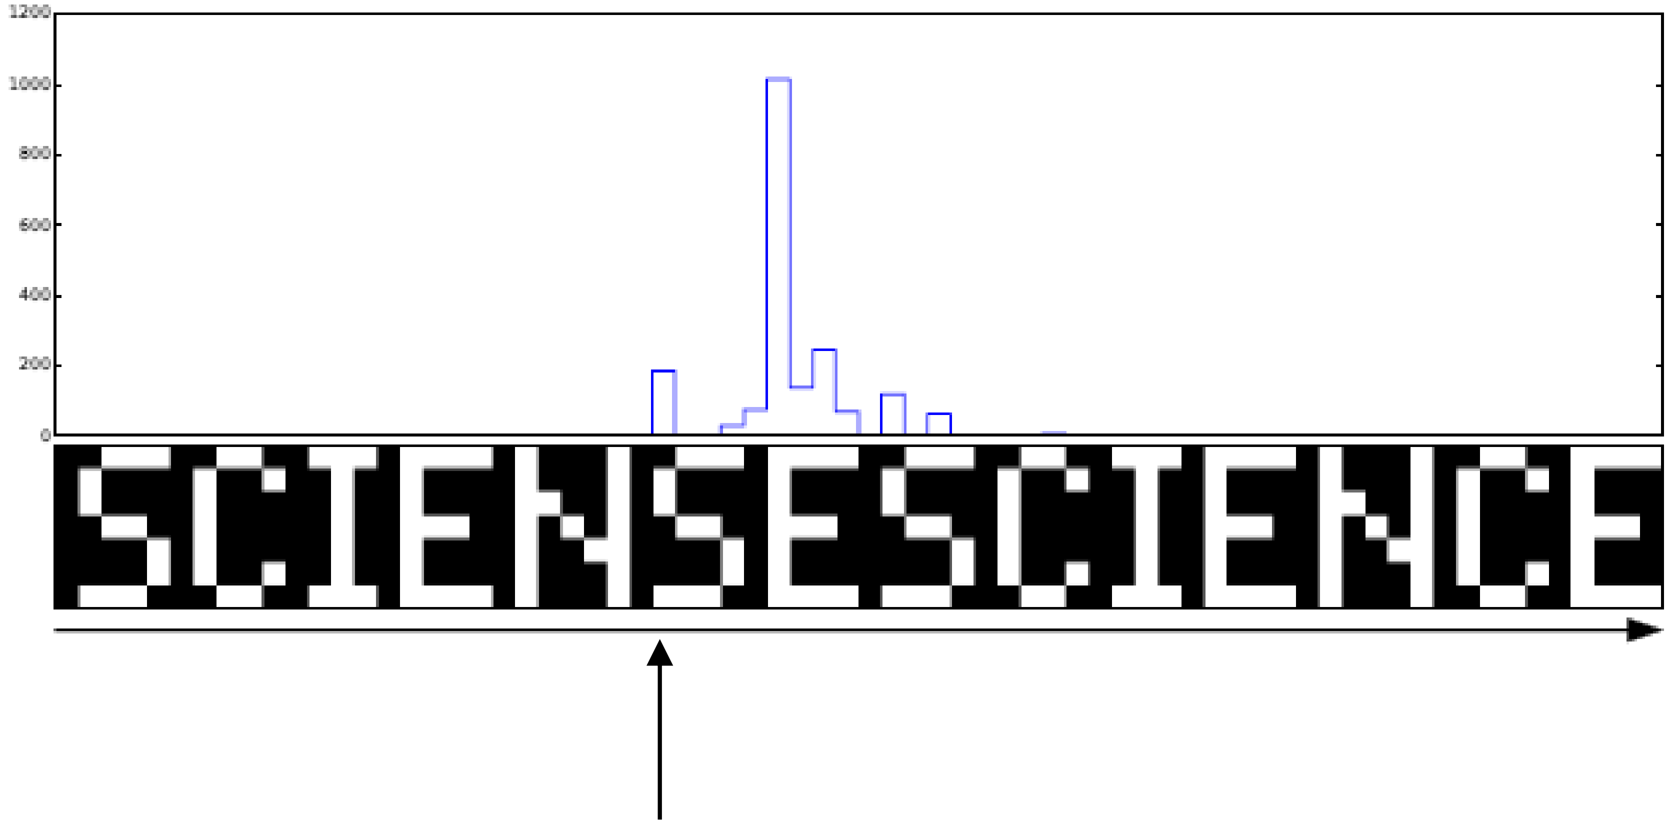

Beginning of anomalous sequence

DyBM learning a generative model of human evolution [O & Otsuka (2015)]

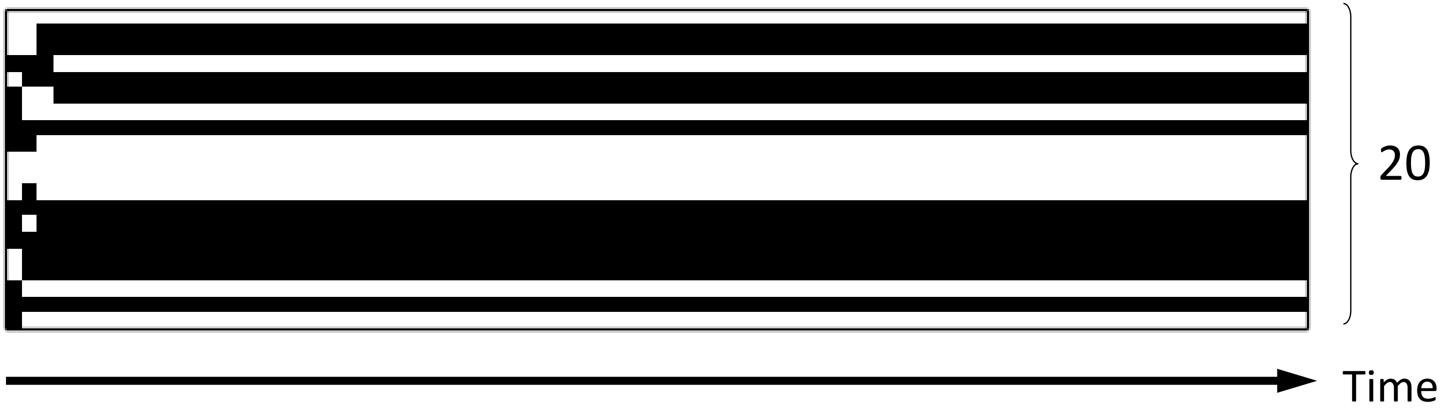

Images from www.nature.com/articles/srep14149

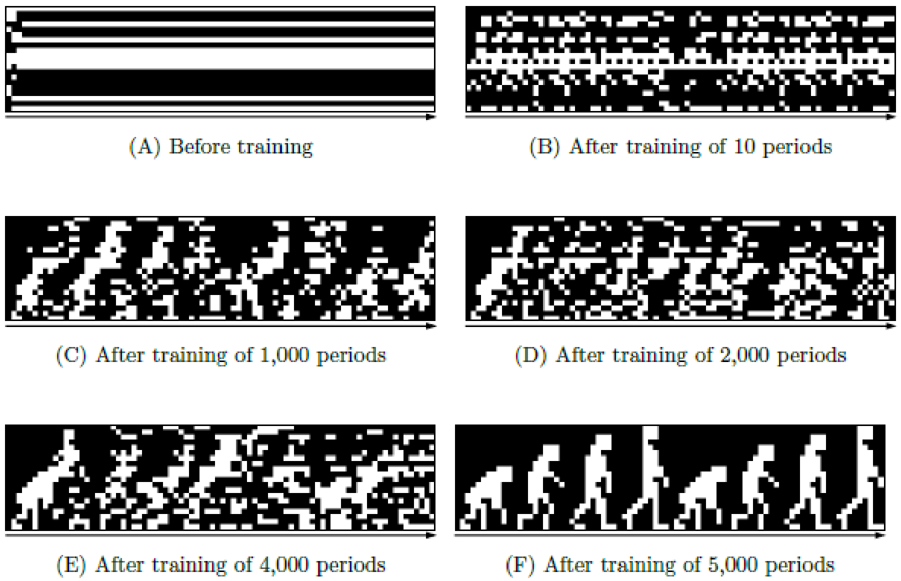

(A) Before training

(B) After training of 10 periods

(C) After training of 1,000 periods

(D) After training of 2,000 periods

(E) After training of 4,000 periods

(F) After training of 5,000 periods

DyBM learning a generative model of *Ich bin ein Musikante* [O & Otsuka (2015)]

Images from www.nature.com/articles/srep14149

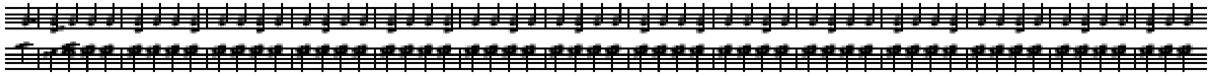

(A) Before training

(B) After training of 10 periods

(C) After training of 1,000 periods

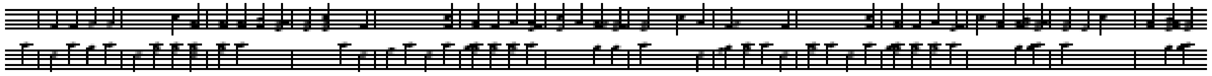

(D) After training of 10,000 periods

(E) After training of 100,000 periods

(F) After training of 900,000 periods

- 12 neurons
- Sounds prepared by Shohei Ohsawa & Yachiko Obara

Comparison between DyBM and LSTM [Dasgupta & O (AAAI-17)]

Root mean squared error after learning for 20 epochs

Model	Retail price of gasoline & diesel • 8 dimensions • 20 hidden units	Sunspot number • 1 dimension • 50 hidden units
LSTM	0.067	**0.073**
DyBM (delay=2)	0.058	0.082
DyBM (delay=3)	**0.056**	**0.077**
DyBM (delay=4)	0.060	0.077

Systematic approaches to tuning DyBM's hyperparameters
[Dasgupta, Yoshizumi, O (ICPR 2016)]
[O & Dasgupta (IBM R&D J. 2017)]

DyBM and LSTM have comparable structures

Comparison between DyBM and LSTM [Dasgupta & O (AAAI-17)]

Time per epoch = 0.7 sec.

DyBM learns 16 times faster

Time per epoch = 11.2 sec.

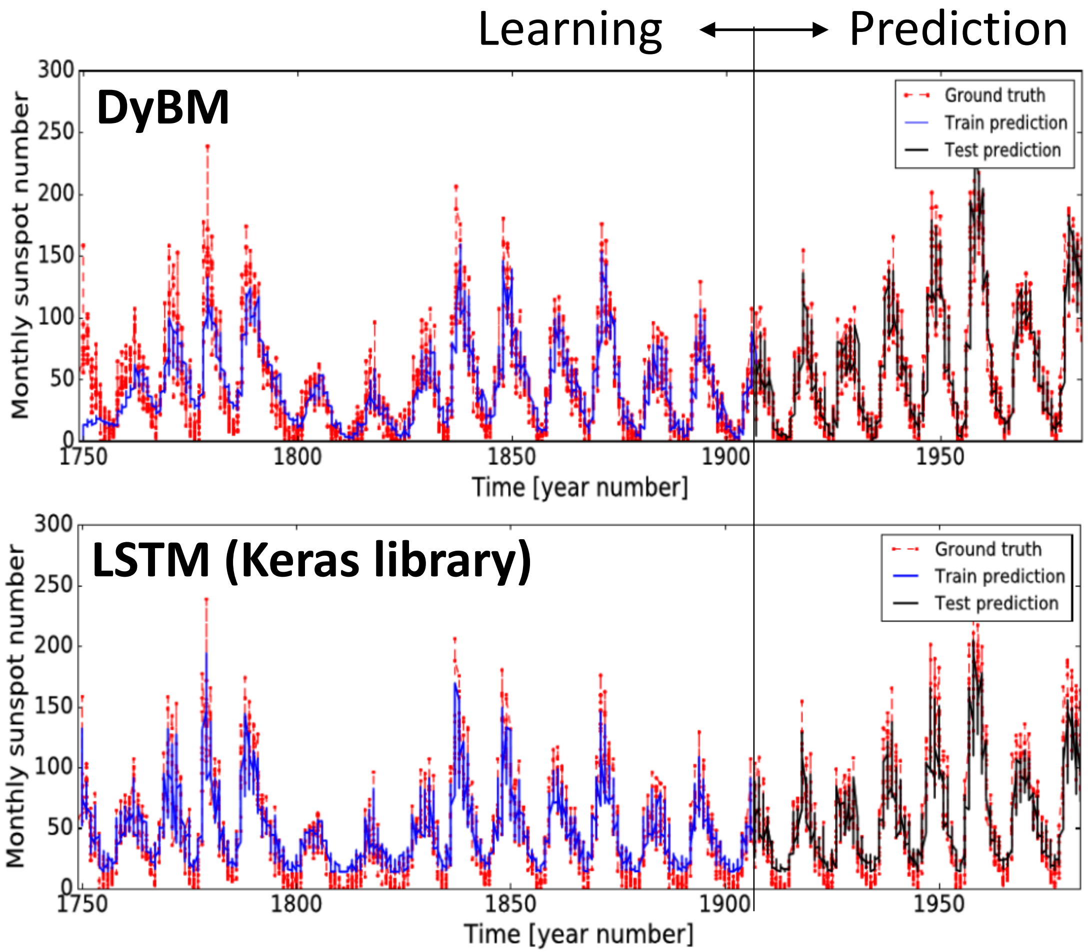

RMSE = 0.077

Comparable accuracy

RMSE = 0.073

<u>**Learning and predicting sunspot number**</u>

Images from IBM Research Report RT0975

Summary: Dynamic Boltzmann machine

Dynamic Boltzmann machine

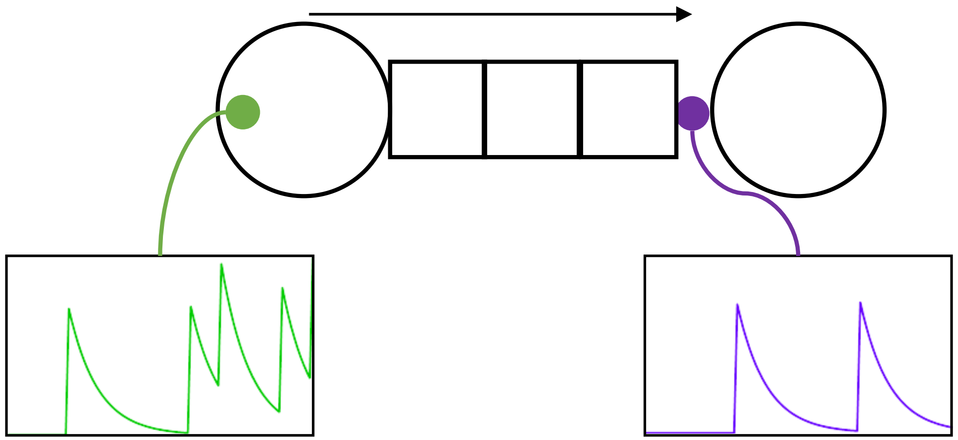

MLE

Spike-timing dependent plasticity

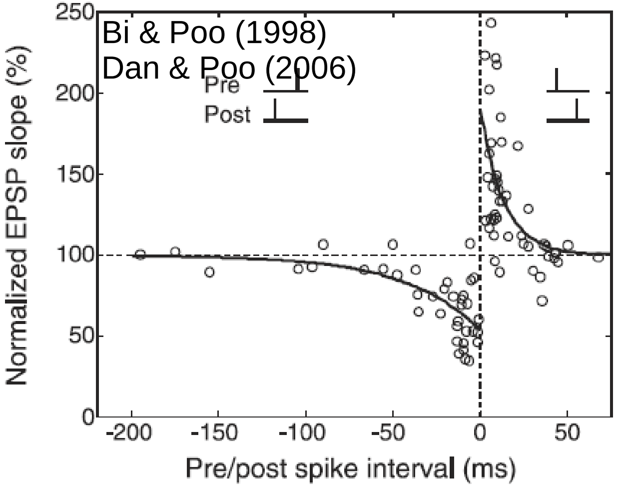

Formal limit

Refine

Boltzmann machine

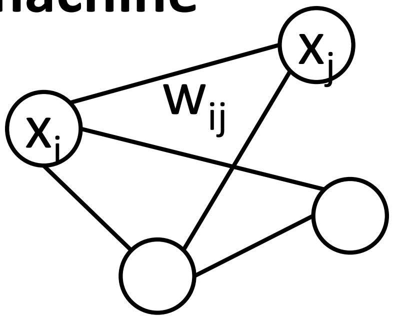

MLE

Hebb's rule

Cells that fire together, wire together

We provide theoretical underpinnings for STDP

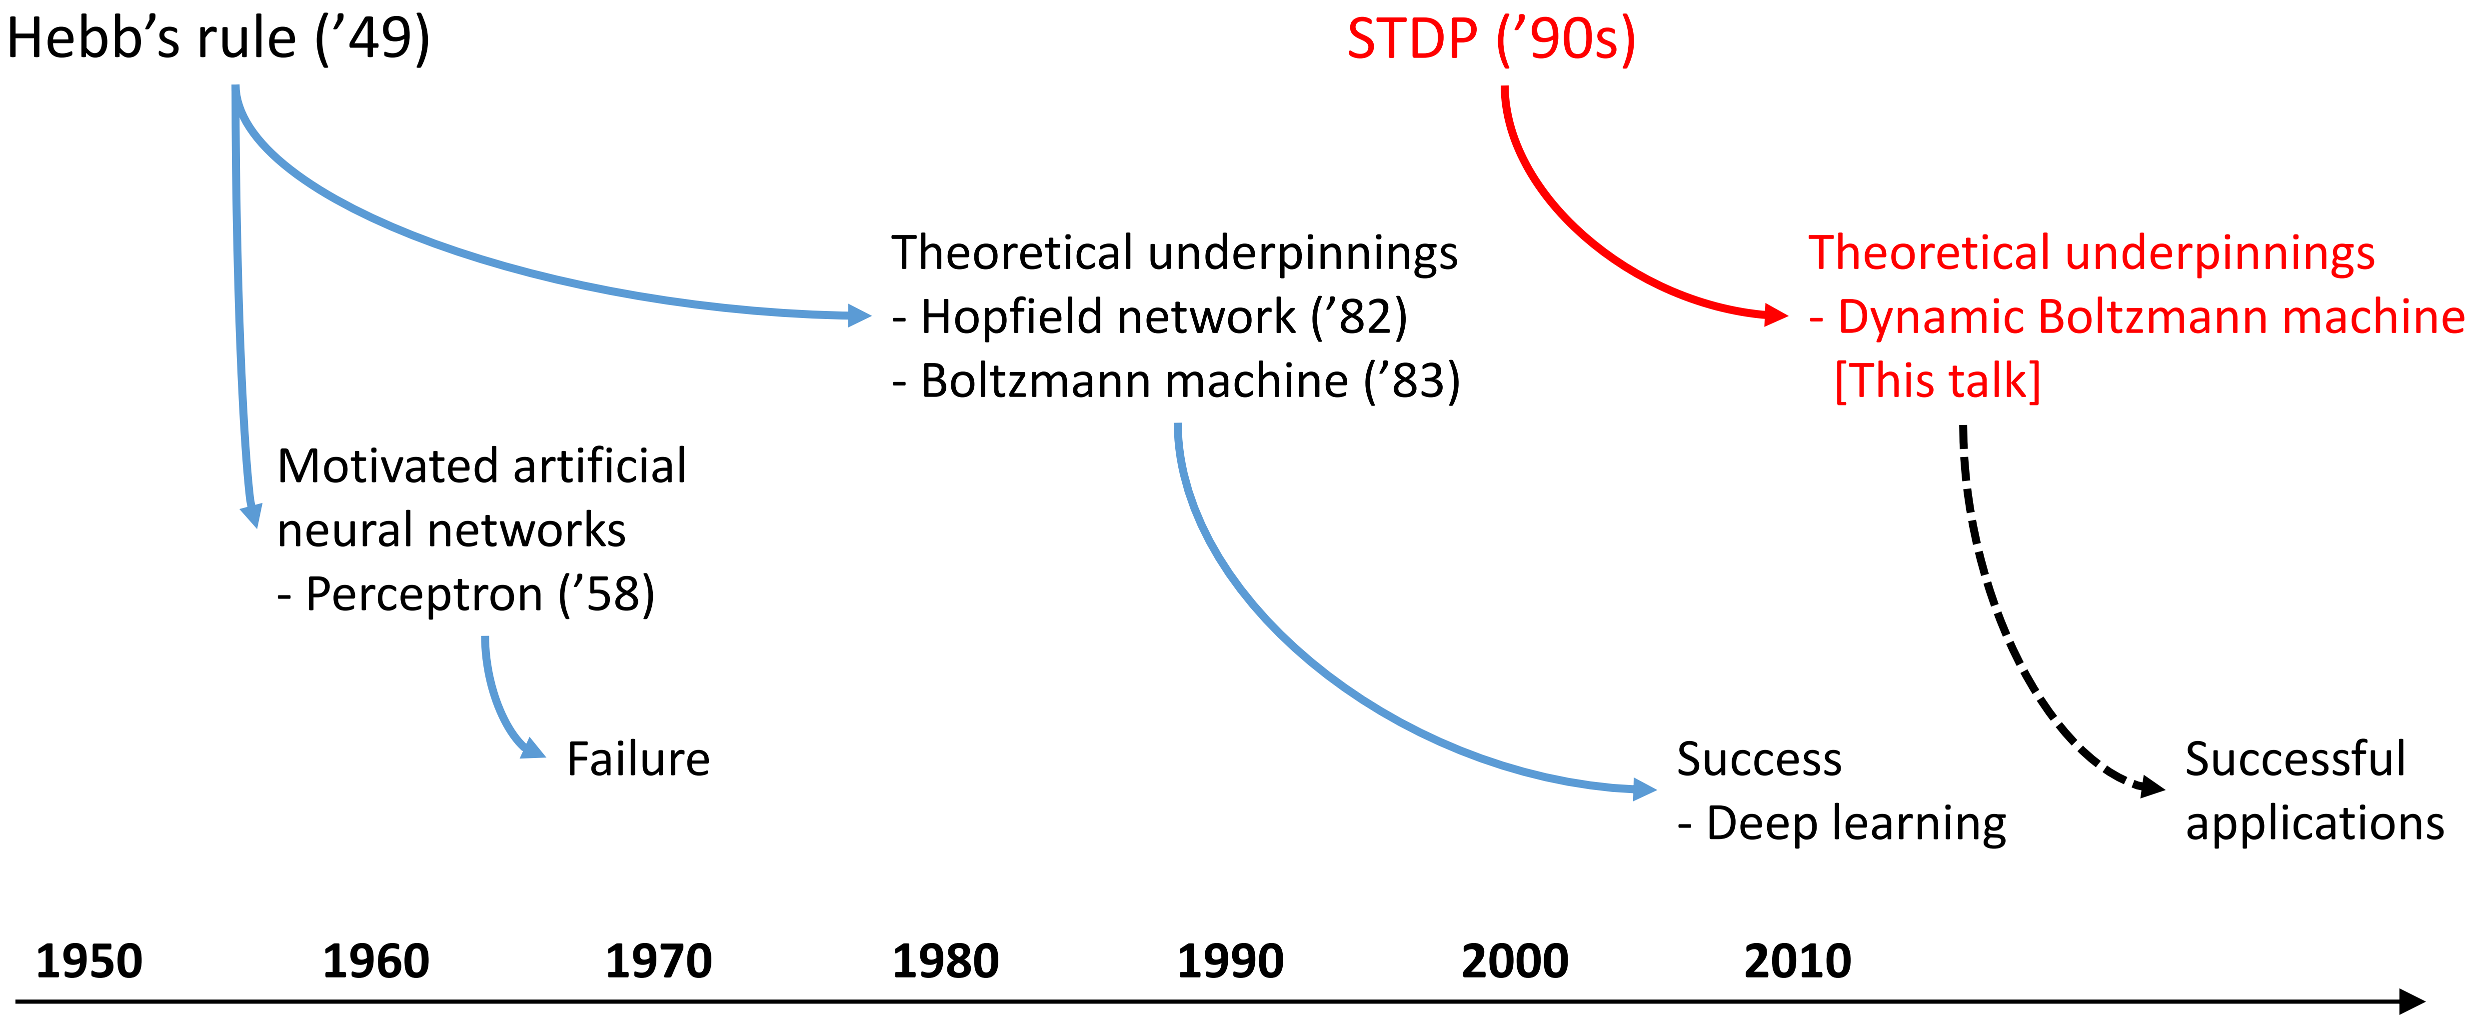

Dynamic Boltzmann machine

References

- Takayuki Osogami and Makoto Otsuka, "Seven neurons memorizing sequences of alphabetical images via spike-timing dependent plasticity," *Scientific Reports*, **5**, 14149 (2015). www.nature.com/articles/srep14149
- Takayuki Osogami and Makoto Otsuka, *Learning dynamic Boltzmann machines with spike-timing dependent plasticity*, IBM Research Report, RT0967 (2015). arxiv.org/abs/1509.08634
- Sakyasingha Dasgupta, Takayuki Yoshizumi, and Takayuki Osogami, "Regularized dynamic Boltzmann machine with delay pruning for unsupervised learning of temporal sequences," ICPR 2016. arxiv.org/abs/1610.01989
- Takayuki Osogami, "Learning binary or real-valued time-series via spike-timing dependent plasticity," Computing with Spikes NIPS 2016 Workshop, to appear.
- Sakyasingha Dasgupta and Takayuki Osogami, "Nonlinear dynamic Boltzmann machines for time-series prediction," AAAI-17. Extended research report available at goo.gl/Vd0wna
- Takayuki Osogami and Sakyasingha Dasgupta, "Learning the values of the hyperparameters of a dynamic Boltzmann machine," *IBM Journal of Research and Development*, 61(4/5), 2017, to appear.

This research was supported by CREST, JST