



THE DEVELOPER'S CONFERENCE

Trilha – Computação Cognitiva

Estevão Uyrá Pardillos Vieira

Cientista de Dados

Mestrando em *Neurociências e Cognição*

Como a Neurociência inspira a Inteligência Artificial

Estevão Uyrá
Pardillos Vieira



USP

edX




Universidade Federal do ABC

serasa
experian™

"What I cannot create, I do not understand"

Richard Feynmann



Overview

- Introdução
- Peças básicas
 - Neurônio
 - Plasticidade
- Reforço
 - Condicionamento
 - Replay
- Redes
 - Back-propagation
 - Cortex Visual

Overview


- Introdução
- Peças básicas
 - Neurônio
 - Plasticidade
- Reforço
 - Condicionamento
 - Replay
- Redes
 - Back-propagation
 - Cortex Visual

nature

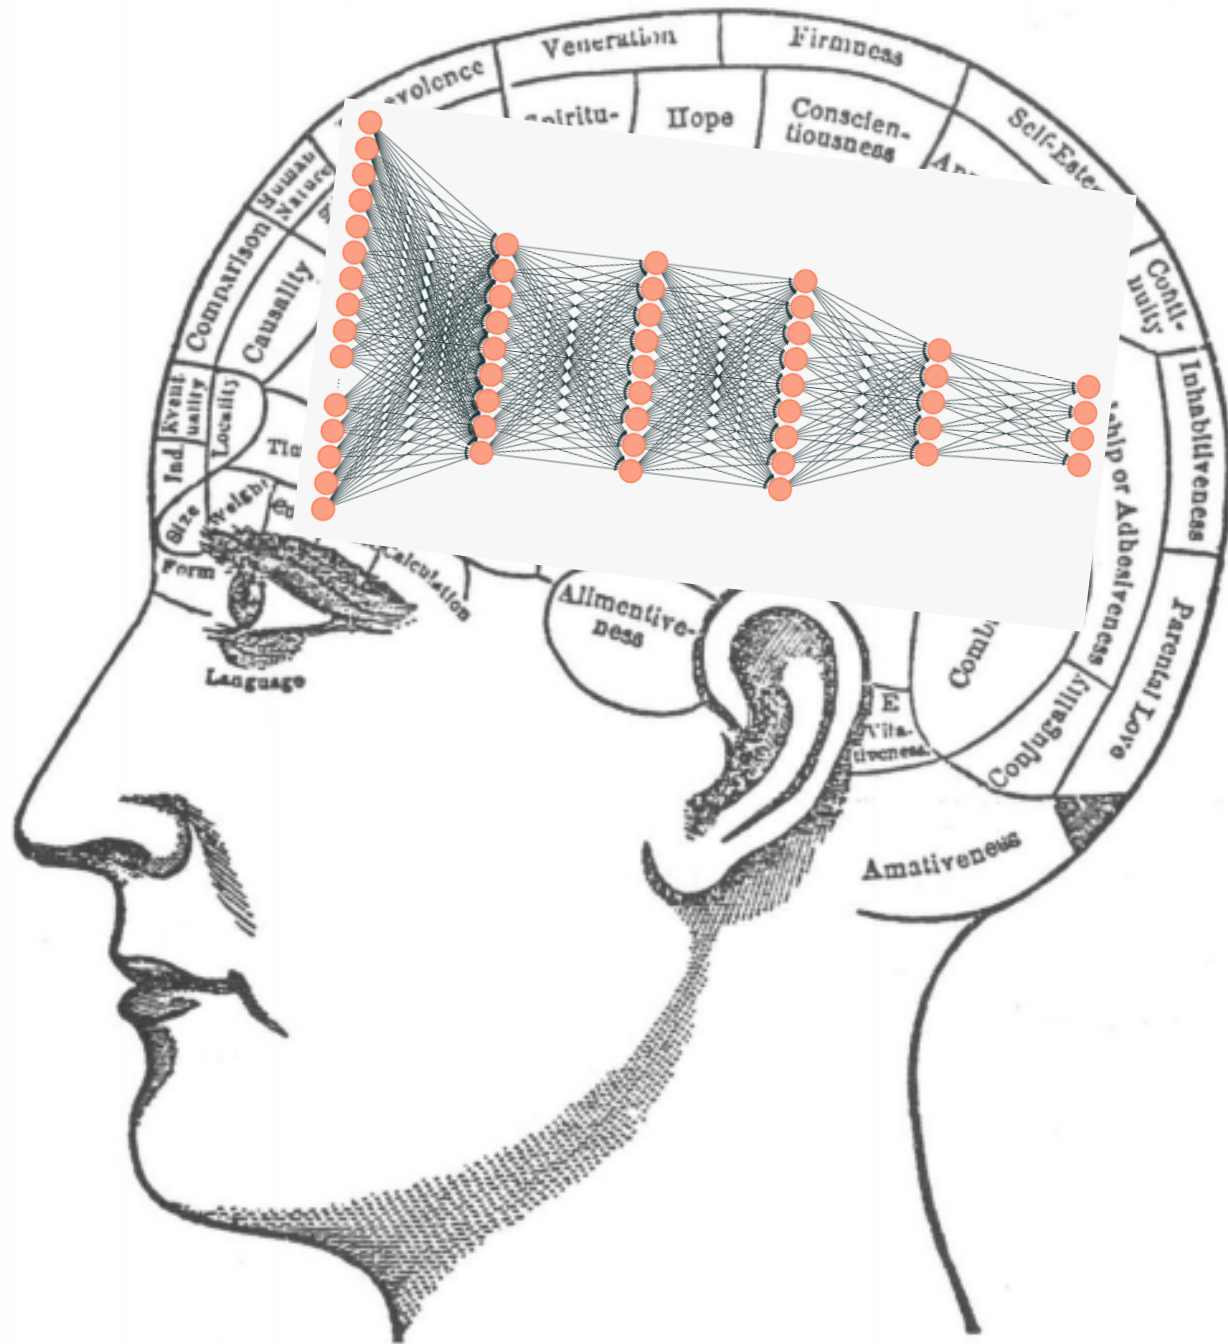
THE INTERNATIONAL WEEKLY JOURNAL OF SCIENCE


At last — a computer program that
can beat a champion Go player **PAGE 484**

ALL SYSTEMS GO




Fonte: [26]






Fonte: [2]



Fonte: [2]

**Falta algo para atingir a
inteligência humana?**


Falta.




Modificado de: [3]

Caminhão

Avestruz







Fonte: [6]



Fonte: [6]

[7]



Fonte: [6]

[7]



Lontra



Fonte: [6]

[7]



[8]



Lontra



Fonte: [6]

[7]



[8]



Lontra

[5]



Fonte: [6]

[7]



[8]



Lontra

[5]



Fonte: [6]

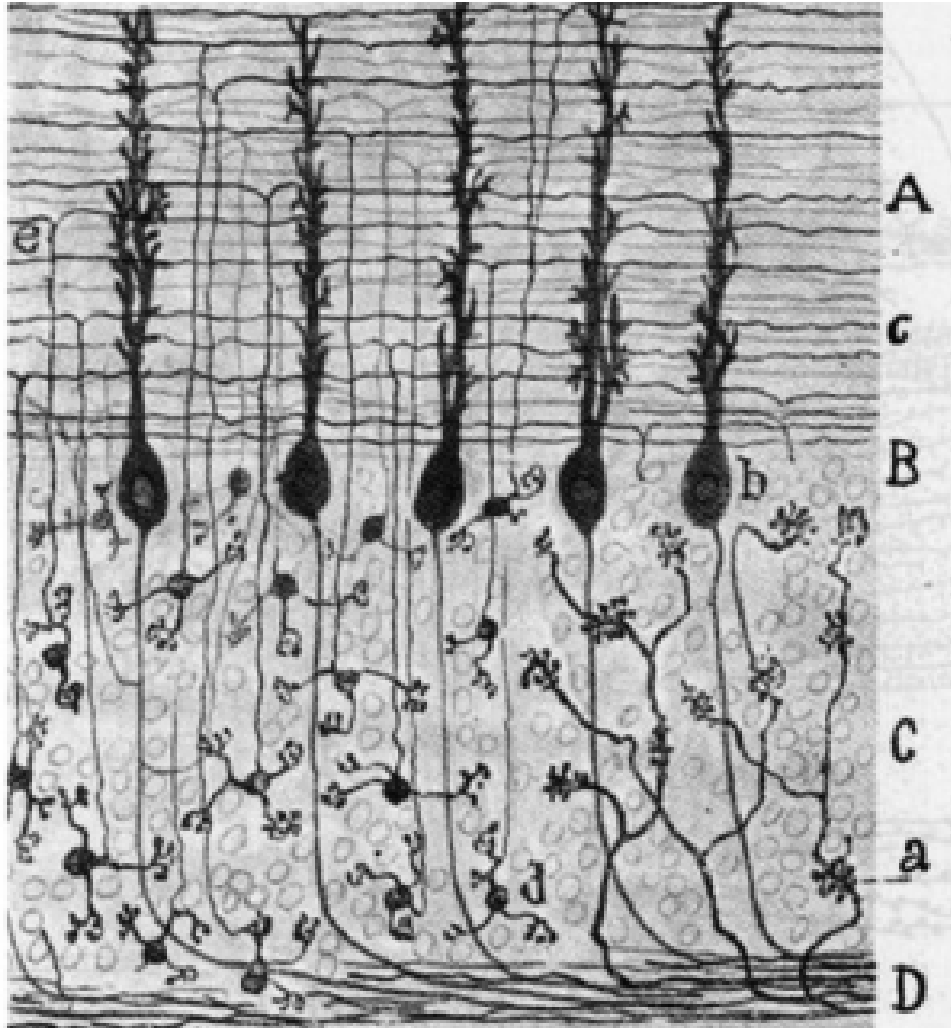
[9]



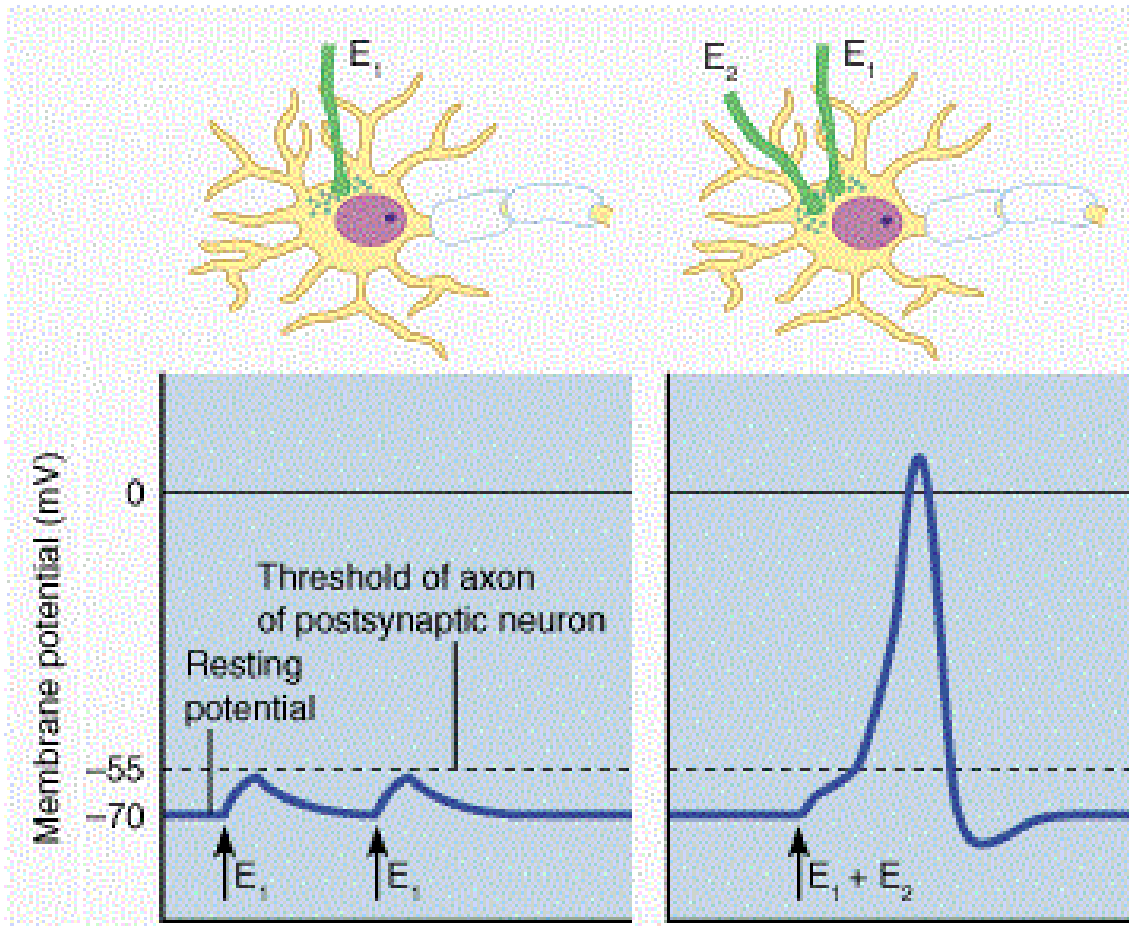
O que falta?

Overview

- Introdução
- Peças básicas
 - Neurônio
 - Plasticidade
- Reforço
 - Condicionamento
 - Replay
- Redes
 - Back-propagation
 - Cortex Visual



Fonte: [1]




Copyright © 2001 Benjamin Cummings,
 an imprint of Addison Wesley Longman, Inc.

Modificado de: [12]


***// McCulloch, W. S., & Pitts, W. (1943).
A logical calculus of the ideas
immanent in nervous activity.***

The bulletin of mathematical biophysics, 5(4), 115-133.


O objeto é um círculo



+1




+1



X and Y

O objeto é vermelho

O objeto é um círculo vermelho







Overview

- Introdução
- Peças básicas
 - Neurônio
 - Plasticidade
- Reforço
 - Condicionamento
 - Replay
- Redes
 - Back-propagation
 - Cortex Visual

***“ Neurons that fire together,
wire together***

Donald Hebb, 1949







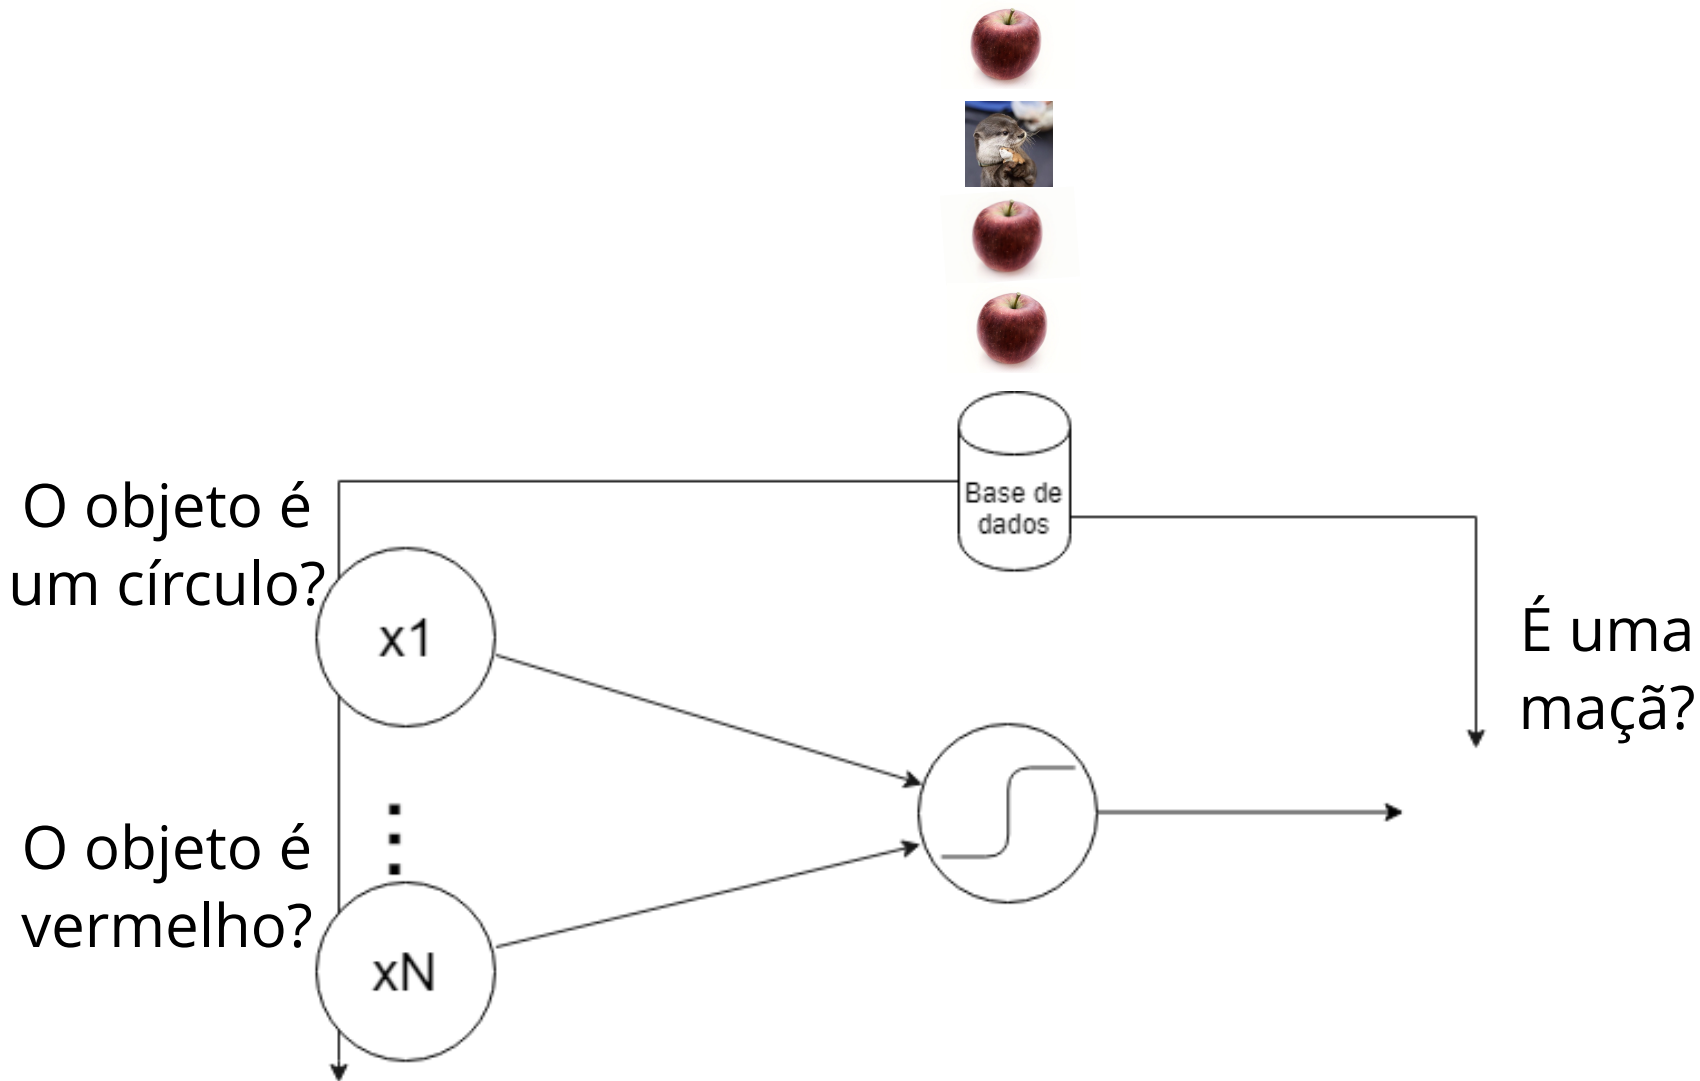
// Rosenblatt, F. (1958).

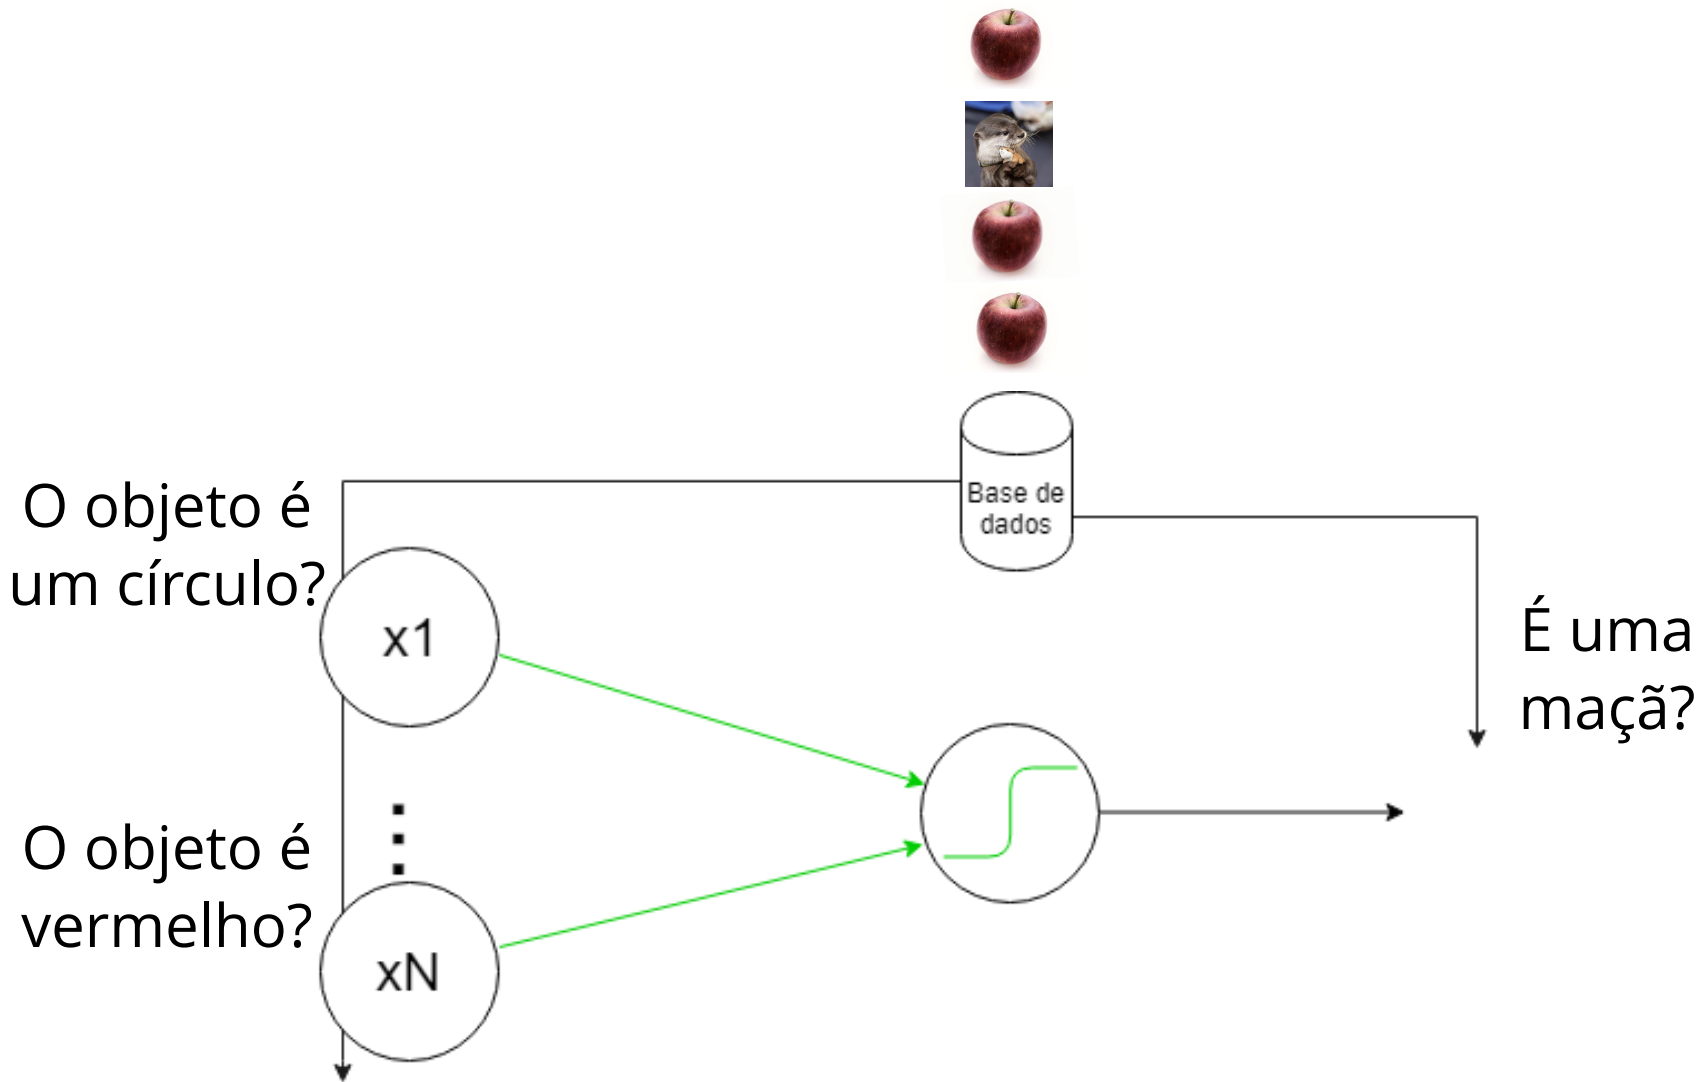
The perceptron: a probabilistic model for information storage and organization in the brain.

Psychological review, 65(6), 386.









<https://www.youtube.com/embed/xpJHhHwR4DQ?enablejsapi=1>

// The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The New York Times, 1958 [32]

Marvin Minsky and Seymour Papert




Perceptrons



An Introduction to Computational Geometry


https://www.youtube.com/embed/3JjDmFV_YwQ?enablejsapi=1





Fonte: [10]



The first AI winter



The first AI winter




Overview


- Introdução
- Peças básicas
 - Neurônio
 - Plasticidade
- Reforço
 - Condicionamento
 - Replay
- Redes
 - Back-propagation
 - Cortex Visual

<https://www.youtube.com/embed/LeZZLdEJNyg?enablejsapi=1>


***“ The real question is not
whether machines think but
whether men do.***



B.F. Skinner






Fonte: [29]







Plan of maze
14-Unit T-Alley Maze


FIG. 1

(From M. H. Elliott, The effect of change of reward on the maze performance of rats. *Univ. Calif. Publ. Psychol.*, 1928, 4, p. 20.)

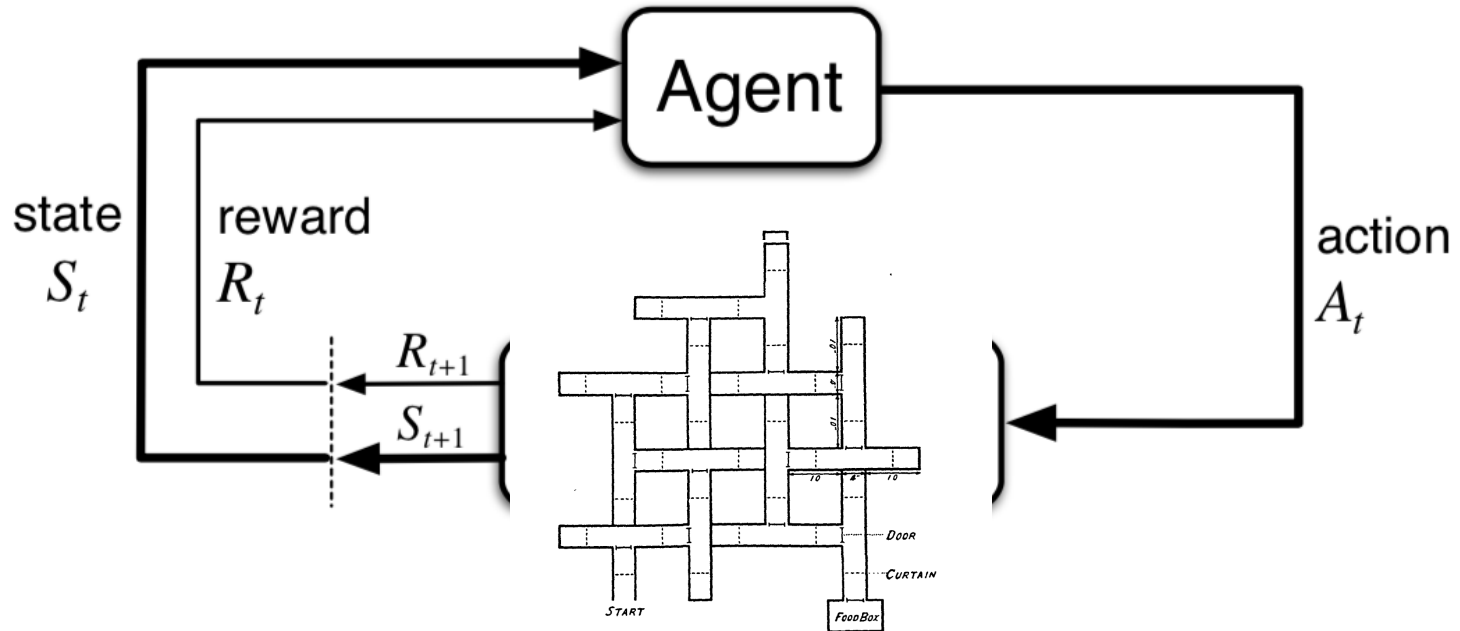
<https://www.youtube.com/embed/6LAWovT7FSI?enablejsapi=1>

<https://www.youtube.com/embed/aNXhyPj-RsM?enablejsapi=1>

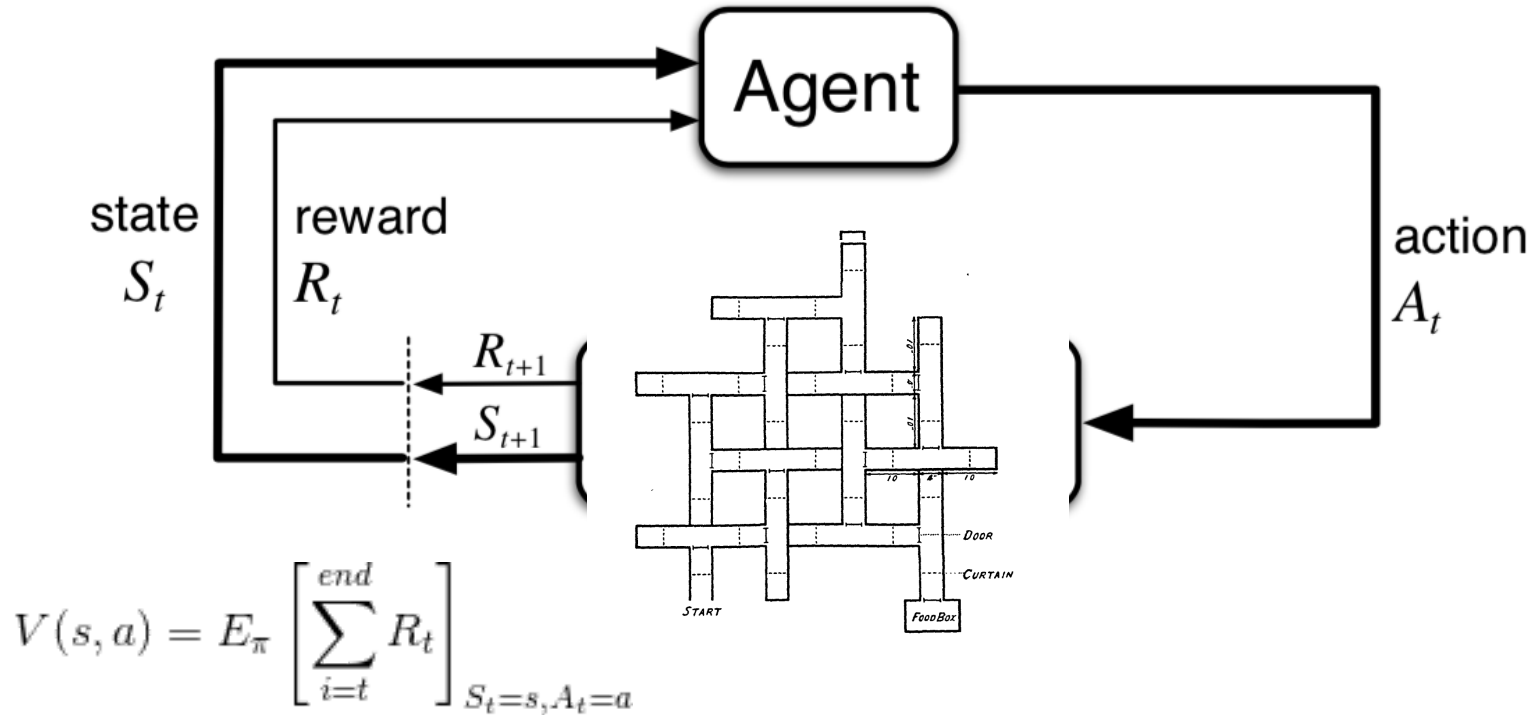
Reinforcement Learning




Reinforcement Learning



Reinforcement Learning



Reinforcement Learning



$$V(s, a) = E_{\pi} \left[\sum_{i=t}^{\text{end}} R_t \right]_{S_t=s, A_t=a}$$

$$V(s, a) = E_{\pi}[R_t + V(s_{t+1}, a_{t+1})]$$

$$V(s, a) \leftarrow V(s, a) + \frac{1}{n}(R_t + V(s_{t+1}, a_{t+1}) - V(s, a))$$


$$V(s, a) += \frac{1}{n}[R_t + V(s_{t+1}, a_{t+1}) - V(s, a)]$$

Reescrevendo

Aproximação iterativa

Reescrevendo

Reinforcement Learning



$$V(s, a) = E_{\pi} \left[\sum_{i=t}^{\text{end}} R_t \right]_{S_t=s, A_t=a}$$

$$V(s, a) = E_{\pi}[R_t + V(s_{t+1}, a_{t+1})]$$

$$V(s, a) \leftarrow V(s, a) + \frac{1}{n}(R_t + V(s_{t+1}, a_{t+1}) - V(s, a))$$

$$V(s, a) += \frac{1}{n}[R_t + V(s_{t+1}, a_{t+1}) - V(s, a)]$$

$$V(s, a) += \alpha[R_t + V(s_{t+1}, a_{t+1}) - V(s, a)]$$


Reescrevendo

Aproximação iterativa

Reescrevendo

Aproximação

Reinforcement Learning



$$V(s, a) = E_{\pi} \left[\sum_{i=t}^{\text{end}} R_t \right]_{S_t=s, A_t=a}$$

$$V(s, a) = E_{\pi} [R_t + V(s_{t+1}, a_{t+1})]$$

$$V(s, a) \leftarrow V(s, a) + \frac{1}{n} (R_t + V(s_{t+1}, a_{t+1}) - V(s, a))$$

$$V(s, a) += \frac{1}{n} [R_t + V(s_{t+1}, a_{t+1}) - V(s, a)]$$

$$V(s, a) += \alpha [R_t + V(s_{t+1}, a_{t+1}) - V(s, a)]$$

Sinal dopaminérgico

Reescrevendo


Aproximação iterativa

Reescrevendo

Aproximação


Dopamine responses

No prediction
Reward occurs




Dopamine responses

No prediction
Reward occurs




Reward predicted
Reward occurs




Dopamine responses


No prediction
Reward occurs




Reward predicted
Reward occurs



Reward predicted
No reward occurs




```
1 ALPHA = 0.1
2
3 while not ended:
4     new_state, reward, ended, _ = env.step(action)
5     new_action = policy(Q, state)
6
7     value_of_state = Q[state, action]
8     value_of_next_state = Q[next_state, new_action]
9
10    Q[state, action] += ALPHA*(reward + value_of_next_state - value_of_state)
11    state, action = next_state, next_action
12
13 # Epsilon-greedy policy
14 def policy(Q, state, eps=.01):
15     if np.random.rand() < eps:
16         return np.random.choice( len(Q[state]) )
17     else:
18         return np.argmax( Q[state] )
```





```

1 ALPHA = 0.1
2
3 while not ended:
4     new_state, reward, ended, _ = env.step(action)
5     new_action = policy(Q, state)
6
7     value_of_state = Q[state, action]
8     value_of_next_state = Q[next_state, new_action]
9
10    Q[state, action] += ALPHA*(reward + value_of_next_state - value_of_state)
11    state, action = next_state, next_action
12
13 # Epsilon-greedy policy
14 def policy(Q, state, eps=.01):
15     if np.random.rand() < eps:
16         return np.random.choice( len(Q[state]) )
17     else:
18         return np.argmax( Q[state] )

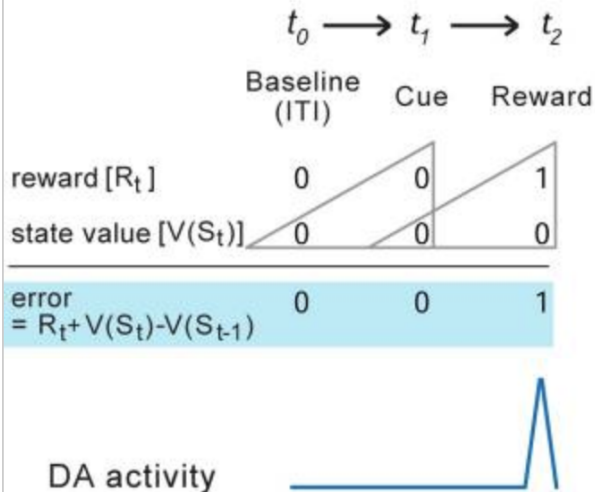
```



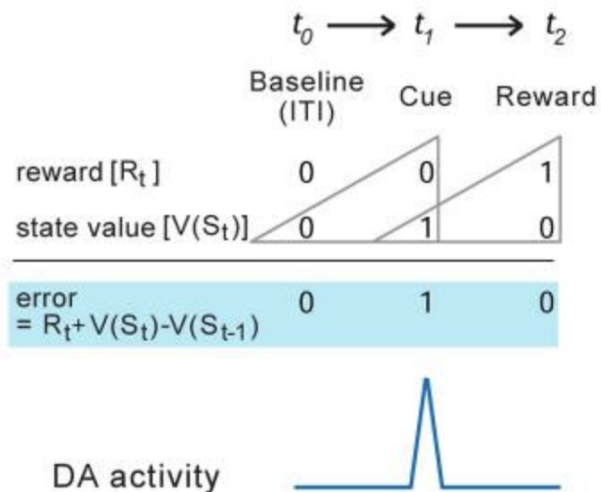


Dopamine responses

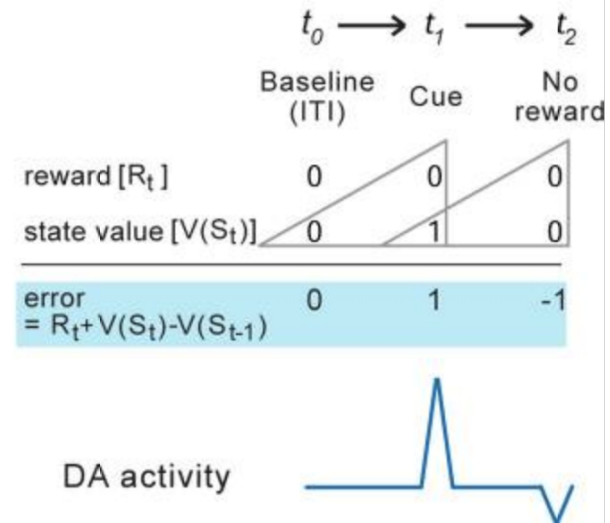
A Unexpected reward (pre-learning)



B Expected reward




C Reward omission




Fonte: [30]

Dopamine responses


No prediction
Reward occurs



Reward predicted
Reward occurs


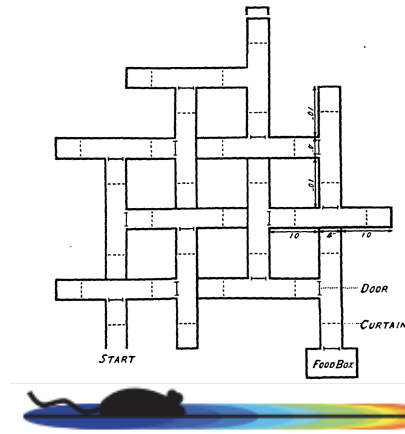


Reward predicted
No reward occurs




Overview


- Introdução
- Peças básicas
 - Neurônio
 - Plasticidade
- Reforço
 - Condicionamento
 - *Replay*
- Redes
 - Back-propagation
 - Cortex Visual



Adaptado de [33]

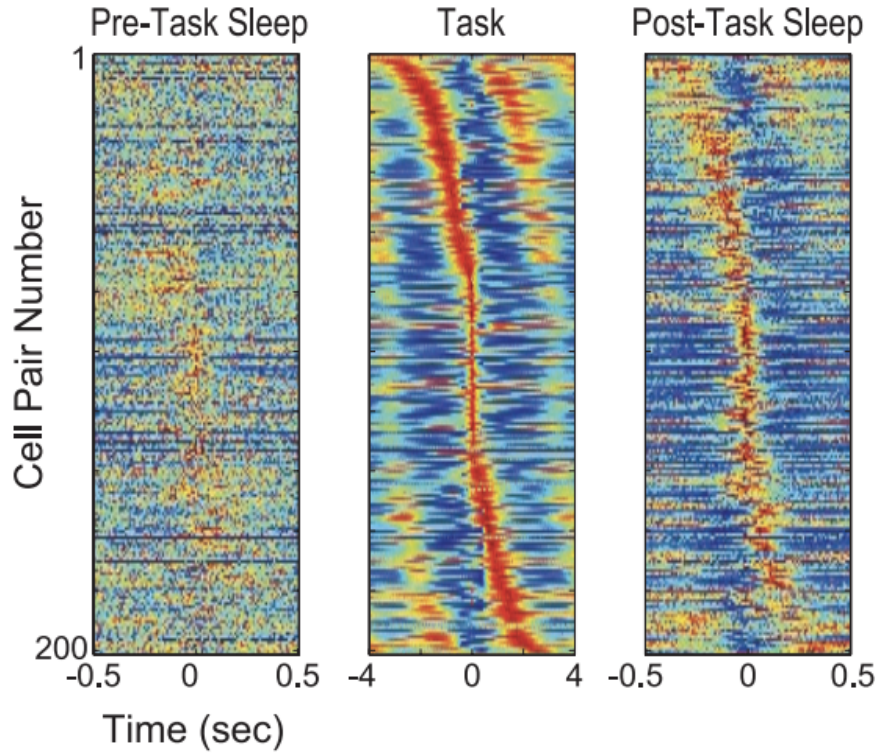


Adaptado de [33]




Adaptado de [33]

Rat 1 Session 13



Rat 2 Session 7



Fonte: [27]

```
class EpisodicMemory():
    def __init__(self, maxlen=500):
        self.episodes = deque(maxlen=maxlen)
        self.current = None


    def start_new_episode(self):
        self.current = deque()

    def store_episode(self):
        self.episodes.append(self.current)

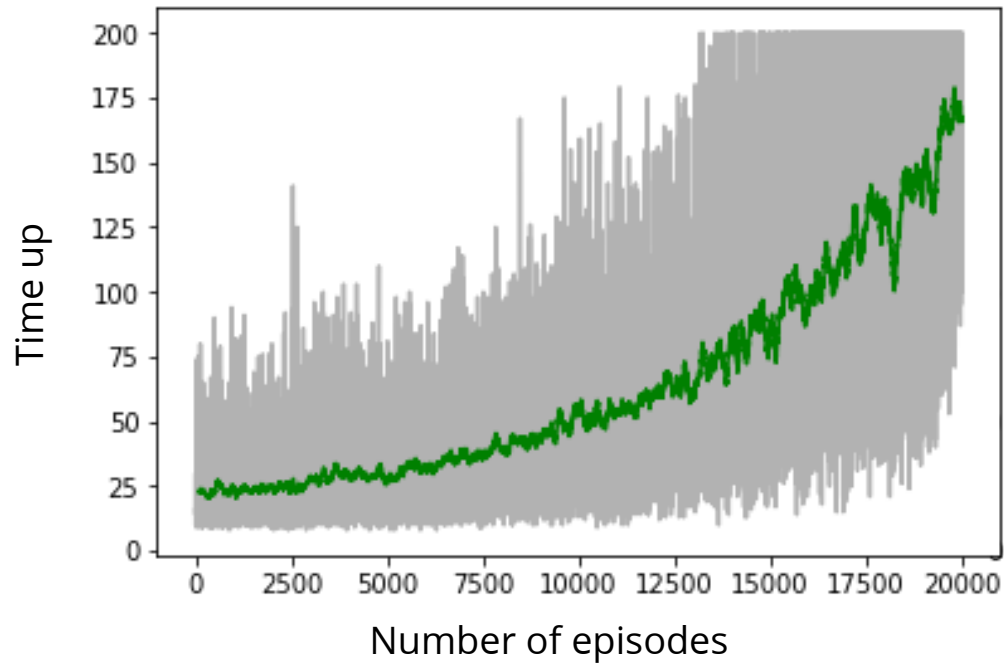
    def add_memory(self, state, action, reward, next_state, next_action):
        self.current.append([state, action, reward, next_state, next_action])

    def _remember(self, Q, episode, discount, state_rep, alpha=.1):
        for state, action, reward, next_state, next_action in episode:
            val_state, val_next = Q[state_rep(state)][action], Q[state_rep(next_state)][next_action]
            Q[state_rep(state)][action] += alpha*(reward + discount*val_next - val_state)
        return Q


    def replay_random(self, Q, discount, state_rep, alpha=.1, weight=False):
        episode = self.episodes[np.random.choice(len(self.episodes))]
        return self._remember(Q, episode, discount, state_rep, alpha=alpha)
```



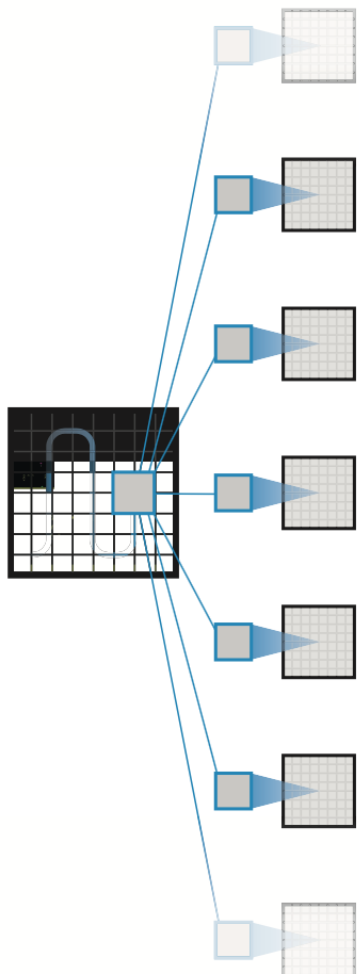
No replay




Replay




Convolution




Convolution



Fully connected




Fully connected



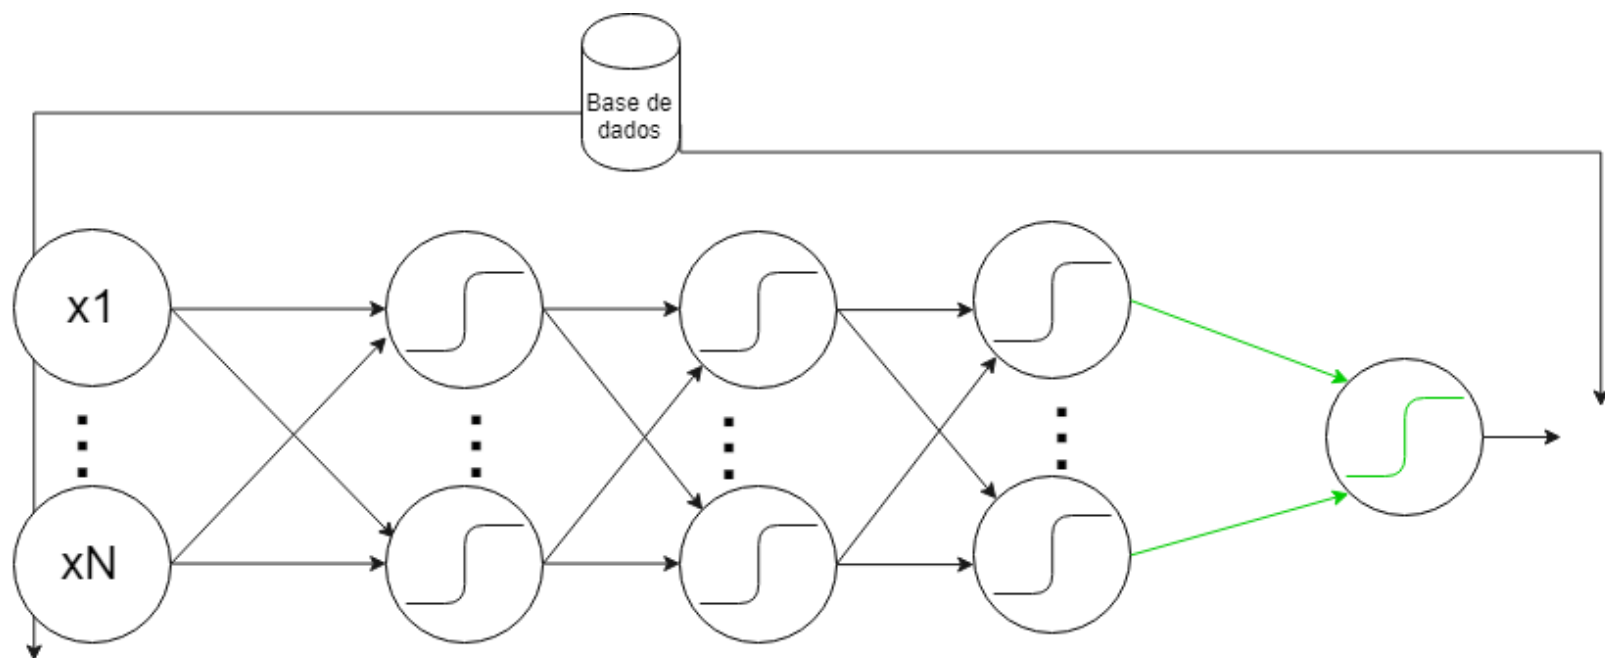
Overview


- Introdução
- Peças básicas
 - Neurônio
 - Plasticidade
- Reforço
 - Condicionamento
 - Replay
- Redes
 - Back-propagation
 - Cortex Visual


“ Rumelhart, D. E., McClelland, J. L., & PDP Research Group. (1986).
***Parallel Distributed Processing (PDP):
Exploration in the Microstructure of Cognition (Vol. 1).***




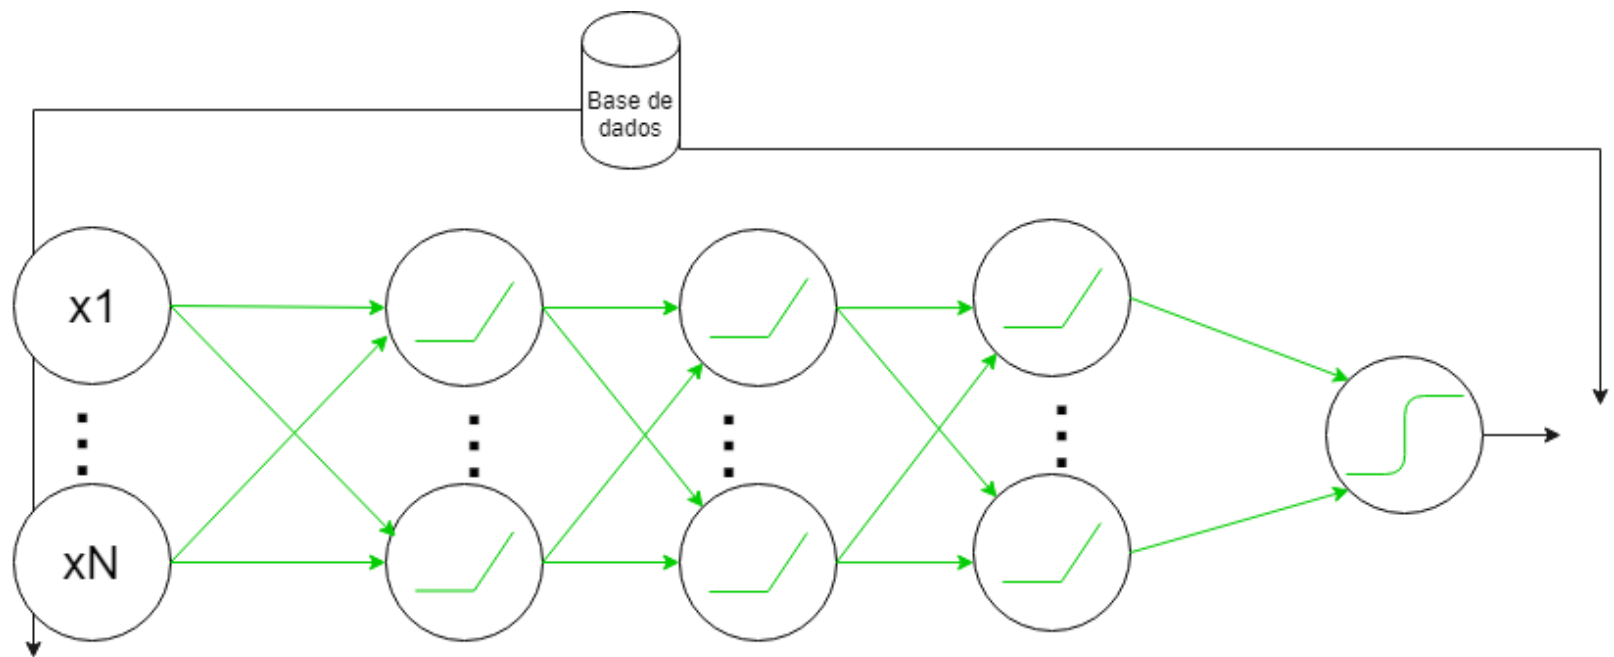
Fonte: [10]
















Imagem 1



\neq

Imagem 2



Overview

- Introdução
- Peças básicas
 - Neurônio
 - Plasticidade
- Reforço
 - Condicionamento
 - Replay
- Redes
 - Back-propagation
 - Cortex Visual




Fig. 1. Correspondence between the hierarchy model by Hubel and Wiesel, and the neural network of the neocognitron


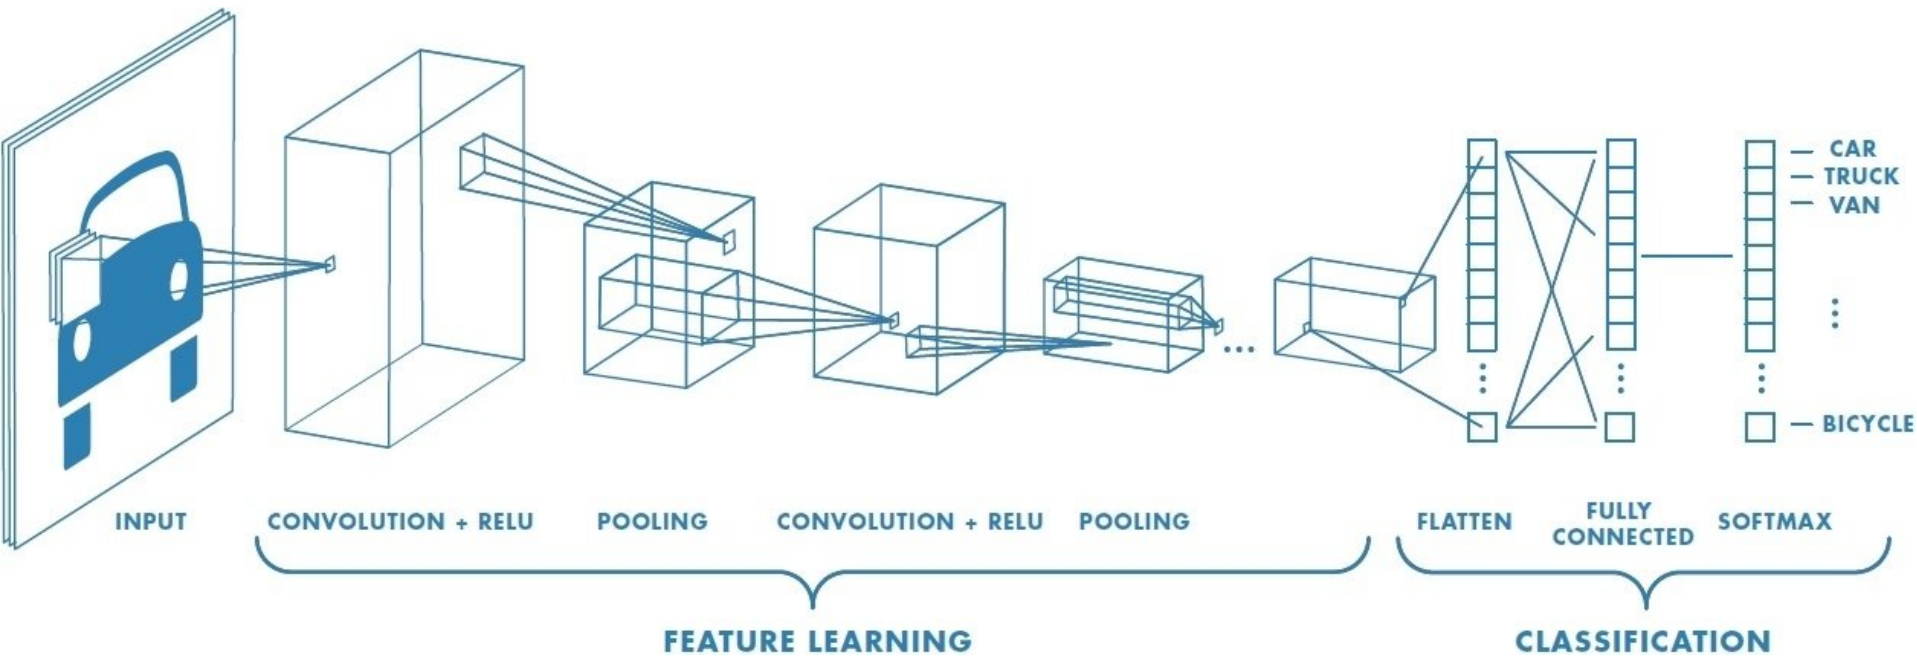



Fig. 2. Schematic diagram illustrating the interconnections between layers in the neocognitron

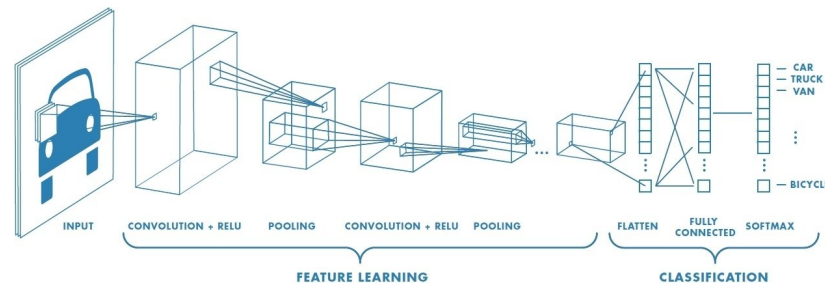


Fonte: [13]

Falta.



Modificado de: [3]



Overview

- Introdução
- Peças básicas
 - Neurônio
 - Plasticidade
- Reforço
 - Condicionamento
 - Replay
- Redes
 - Back-propagation
 - Cortex Visual

Agradecimentos

Gabriela Melo

Squad ConvTech @ DataLab

Beatriz Albiero

Sami Yamouni


Mãe

Referências



- ñoz, F., Boya, J., & Alamo, C. (2006). Neuron theory, the cornerstone of neuroscience, on the centenary of the Nobel Prize award to Santiago Ramón y Cajal. *Brain research*
- w, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT press, originalmente do wikipedia
- ., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2013). Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199.
- , Engstrom, L., Ilyas, A., & Kwok, K. (2017). Synthesizing robust adversarial examples. arXiv preprint arXiv:1707.07397.
- ttter.com/otter_news/status/1085504638084886534
- th blocks - Six Months Old, por Jessica Merz, CC BY-ND 2.0 (<https://creativecommons.org/licenses/by-nd/2.0/>)
- ww.nature.scot/plants-animals-and-fungi/mammals/land-mammals/otter
- edium.com/taronga-conservation-society-australia/10-things-you-might-not-know-about-the-adorable-otter-ae4c62df34b9
- bbble.com/shots/4939061-Swimming-Otter
- lrosaen.com/ml/learning-log/2016-05-24/ CC BY-ND 4.0 (<https://creativecommons.org/licenses/by-nc/4.0/>)
- druff ; De Roo et al / CC BY-SA 3.0 via Commons)
- gacy.owensboro.kctcs.edu/gcaplan/anat/Notes/API Notes K Neurons.htm
- og.floydhub.com/building-your-first-convnet/
- R. (2001). The new phrenology: The limits of localizing cognitive processes in the brain. The MIT press.
- ww.cambridgebrainsciences.com/more/articles/brain-diagrams-almost-always-face-the-same-direction
- or.org/2015/03/stop-noise-from-ruining-your-open-office
- ww.pexels.com/photo/bird-s-eye-view-of-city-buildings-during-sunset-681336/
- milimusica.com.br/blog/as-diferencas-entre-o-violao-de-7-cordas-e-o-tradicional/
- ww.frontiersin.org/articles/10.3389/fnbeh.2010.00171/full
- ww.cardiff.ac.uk/cardiff-university-brain-research-imaging-centre/research/themes/cognitive-neuroscience
- mesclear.com/feynman-mental-models
- xabay.com/pt/photos/fundo-papel-de-parede-ma%C3%A7%C3%A3-fruto-1655303/
- e, Fair use, <https://en.wikipedia.org/w/index.php?curid=17594522>
- ww.companhiadasletras.com.br/autor.php?codigo=00118
- uronitoriaufpa.blogspot.com/2016/10/aula-8-tato-discriminativo.html
- thub.com/beatrizalbiero/presentations/blob/master/NLPpresentation%20(1).pdf
- David R., Masami Tatsuno, and Bruce L. McNaughton. "Fast-forward playback of recent memory sequences in prefrontal cortex during sleep." *science* 318.5853 (2007): 114
- ww.diarioliberdade.org/mundo/361-linguaeducacom/50965-o-modelo-did%C3%A1tico-do-ensino-programado,-segundo-b-f-skinner.html[28] Mnih, Volodymyr, et al. "Human-level machine learning." *Nature* 518.7540 (2015): 529.
- onald, and Patricia H. Janak. "Dopamine prediction errors in reward learning and addiction: from theory to neural circuitry." *Neuron* 88.2 (2015): 247-263.
- Volfram. "Predictive reward signal of dopamine neurons." *Journal of neurophysiology* 80.1 (1998): 1-27.
- ww.nytimes.com/1958/07/08/archives/new-navy-device-learns-by-doing-psychologist-shows-embryo-of.html
- garet F., Shantanu P. Jadhav, and Loren M. Frank. "Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval." *Nature neuroscience*

BRAIN INSPIRED

A podcast where neuroscience and AI converge.



Google
DeepMind





Message me:

estevao.uyra.pv@gmail.com



Essa apresentação:

slides.com/estevaouyra/tdc_neuro_ai



LinkedIn




linkedin.com/in/estevao-uyra-pv/




THE DEVELOPER'S CONFERENCE

Overview

- Introdução
- Peças básicas
 - Neurônio
 - Plasticidade
- Reforço
 - Condicionamento
 - Replay
- Redes
 - Back-propagation
 - Cortex Visual
- O Corpo



[17]




[25]

[18]



[17]



[25]


[18]



<https://www.youtube.com/embed/g0TaYhjpOfo?enablejsapi=1>

**We are not thinking machines
that feel; rather, we are feeling
machines that think**

Antonio Damasio



Agradecimentos

Gabriela Melo

Squad ConvTech @ DataLab

Beatriz Albiero

Sami Yamouni


Mãe

Referências



- ñoz, F., Boya, J., & Alamo, C. (2006). Neuron theory, the cornerstone of neuroscience, on the centenary of the Nobel Prize award to Santiago Ramón y Cajal. *Brain research*
- w, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT press, originalmente do wikipedia
- ., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2013). Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199.
- , Engstrom, L., Ilyas, A., & Kwok, K. (2017). Synthesizing robust adversarial examples. arXiv preprint arXiv:1707.07397.
- ttter.com/otter_news/status/1085504638084886534
- th blocks - Six Months Old, por Jessica Merz, CC BY-ND 2.0 (<https://creativecommons.org/licenses/by-nd/2.0/>)
- ww.nature.scot/plants-animals-and-fungi/mammals/land-mammals/otter
- edium.com/taronga-conservation-society-australia/10-things-you-might-not-know-about-the-adorable-otter-ae4c62df34b9
- bbble.com/shots/4939061-Swimming-Otter
- lrosaen.com/ml/learning-log/2016-05-24/ CC BY-ND 4.0 (<https://creativecommons.org/licenses/by-nc/4.0/>)
- druff ; De Roo et al / CC BY-SA 3.0 via Commons)
- gacy.owensboro.kctcs.edu/gcaplan/anat/Notes/API Notes K Neurons.htm
- og.floydhub.com/building-your-first-convnet/
- R. (2001). The new phrenology: The limits of localizing cognitive processes in the brain. The MIT press.
- ww.cambridgebrainsciences.com/more/articles/brain-diagrams-almost-always-face-the-same-direction
- or.org/2015/03/stop-noise-from-ruining-your-open-office
- ww.pexels.com/photo/bird-s-eye-view-of-city-buildings-during-sunset-681336/
- milimusica.com.br/blog/as-diferencas-entre-o-violao-de-7-cordas-e-o-tradicional/
- ww.frontiersin.org/articles/10.3389/fnbeh.2010.00171/full
- ww.cardiff.ac.uk/cardiff-university-brain-research-imaging-centre/research/themes/cognitive-neuroscience
- mesclear.com/feynman-mental-models
- xabay.com/pt/photos/fundo-papel-de-parede-ma%C3%A7%C3%A3-fruto-1655303/
- e, Fair use, <https://en.wikipedia.org/w/index.php?curid=17594522>
- ww.companhiadasletras.com.br/autor.php?codigo=00118
- uronitoriaufpa.blogspot.com/2016/10/aula-8-tato-discriminativo.html
- thub.com/beatrizalbiero/presentations/blob/master/NLPpresentation%20(1).pdf
- David R., Masami Tatsuno, and Bruce L. McNaughton. "Fast-forward playback of recent memory sequences in prefrontal cortex during sleep." *science* 318.5853 (2007): 114
- ww.diarioliberdade.org/mundo/361-linguaeducacom/50965-o-modelo-did%C3%A1tico-do-ensino-programado,-segundo-b-f-skinner.html[28] Mnih, Volodymyr, et al. "Human-level machine learning." *Nature* 518.7540 (2015): 529.
- onald, and Patricia H. Janak. "Dopamine prediction errors in reward learning and addiction: from theory to neural circuitry." *Neuron* 88.2 (2015): 247-263.
- Volfram. "Predictive reward signal of dopamine neurons." *Journal of neurophysiology* 80.1 (1998): 1-27.
- ww.nytimes.com/1958/07/08/archives/new-navy-device-learns-by-doing-psychologist-shows-embryo-of.html
- garet F., Shantanu P. Jadhav, and Loren M. Frank. "Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval." *Nature neuroscience*

BRAIN INSPIRED

A podcast where neuroscience and AI converge.



Google
DeepMind





Message me:

estevao.uyra.pv@gmail.com



Essa apresentação:

slides.com/estevaouyra/tdc_neuro_ai



LinkedIn

linkedin.com/in/estevao-uyra-pv/



THE DEVELOPER'S CONFERENCE