



Seminar in Deep Neural Networks

Introduction


Natural Language Processing



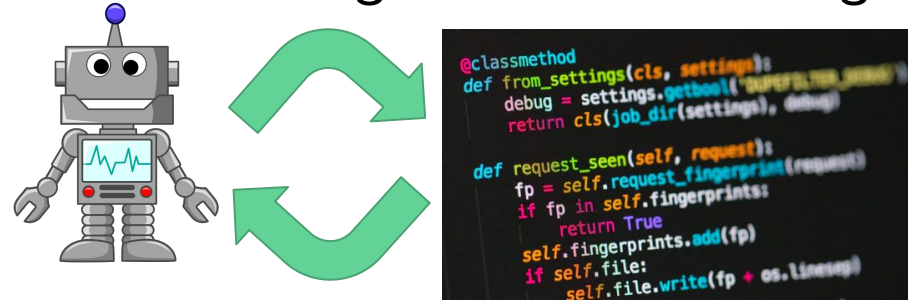
Graph Neural Networks



Reinforcement Learning



Algorithmic Learning



Disclaimer: This is a seminar...



(almost) no basics



participation required

Format

- Assigned topics
- 2 x 35 min video presentation + Piazza
- 30 min facilitated discussion on zoom
- Feedback through Google Form

Grade = presentation + **active participation**

What makes a good talk?

Make it Fun




Explore



Connect

Build

Motivate

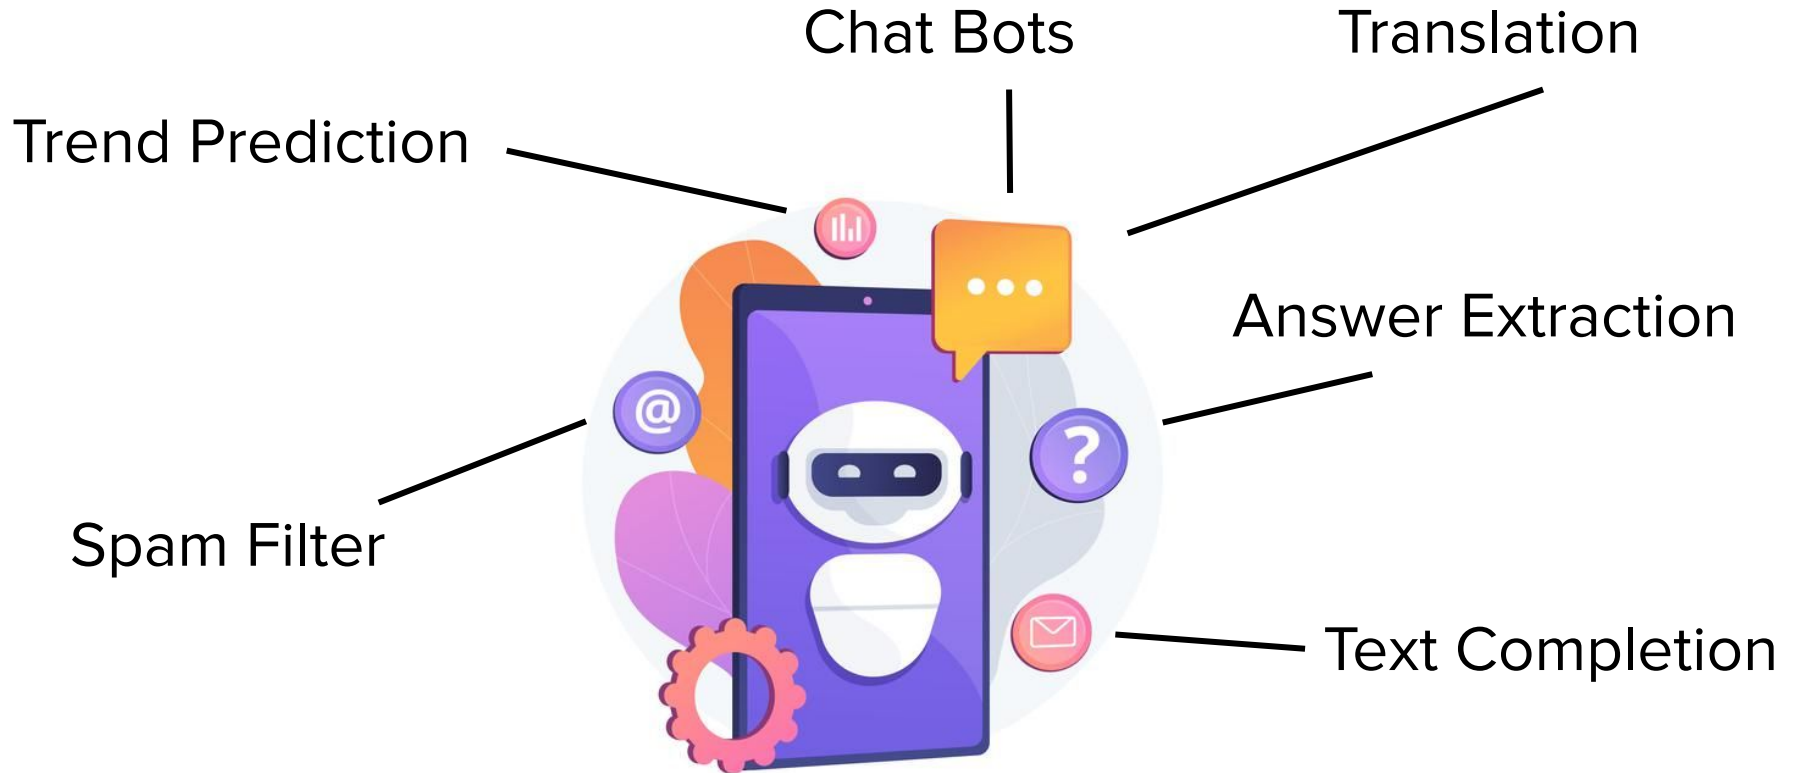


Neural architectures




$$\hat{f}(x) \approx f(x)$$

Natural Language Processing... Evaluation?





Natural Language Processing




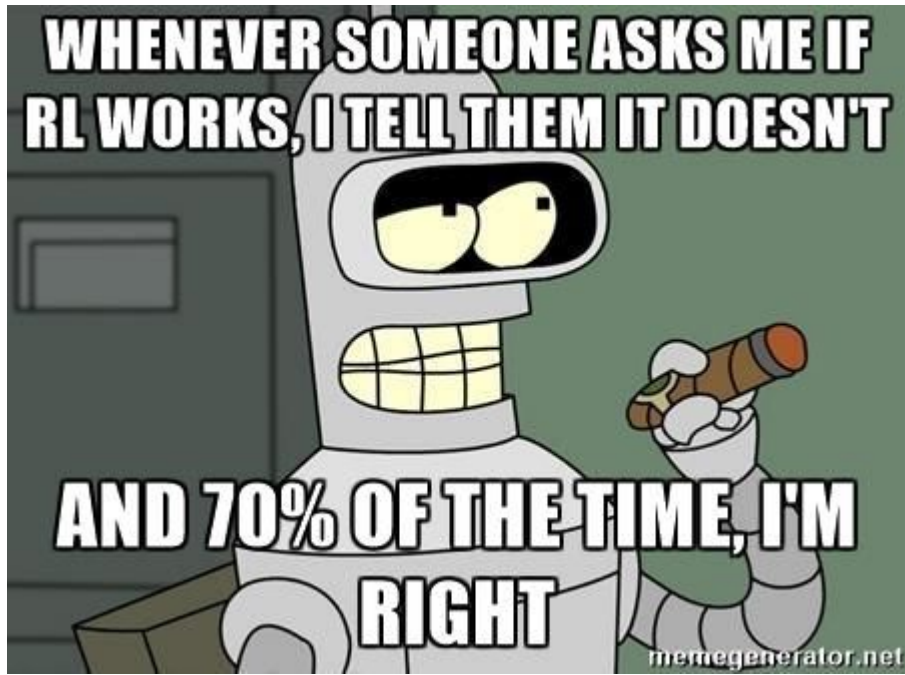
How powerful are these models already?



Reinforcement Learning...




Why you should NOT use reinforcement learning...




Why you should NOT use reinforcement learning...





Agent




State
Reward



Environment


How does RL work?



$$R = \sum_{t'=t}^{T_e} \gamma^{t'-t} r_{t'}$$

$$\gamma \in (0, 1]$$

$$\begin{aligned} V^\pi(\mathbf{s}_t) &= \mathbb{E}_\pi \left[\sum_{t'=t}^{T_e} \gamma^{t'-t} r_{t'} \right] \\ &= \mathbb{E}_\pi [r_t] + \gamma \mathbb{E}_\pi \left[\sum_{t'=t+1}^{T_e} \gamma^{t'-t-1} r_{t'} \right] \\ &= \mathbb{E}_\pi [r_t] + \gamma V^\pi(\mathbf{s}_{t+1}) \end{aligned}$$

 **must be equal**




Q-Learning - Watkins (1989)

$$V^\pi(s_t) = \mathbb{E}_\pi[r_t] + \gamma V^\pi(s_{t+1})$$

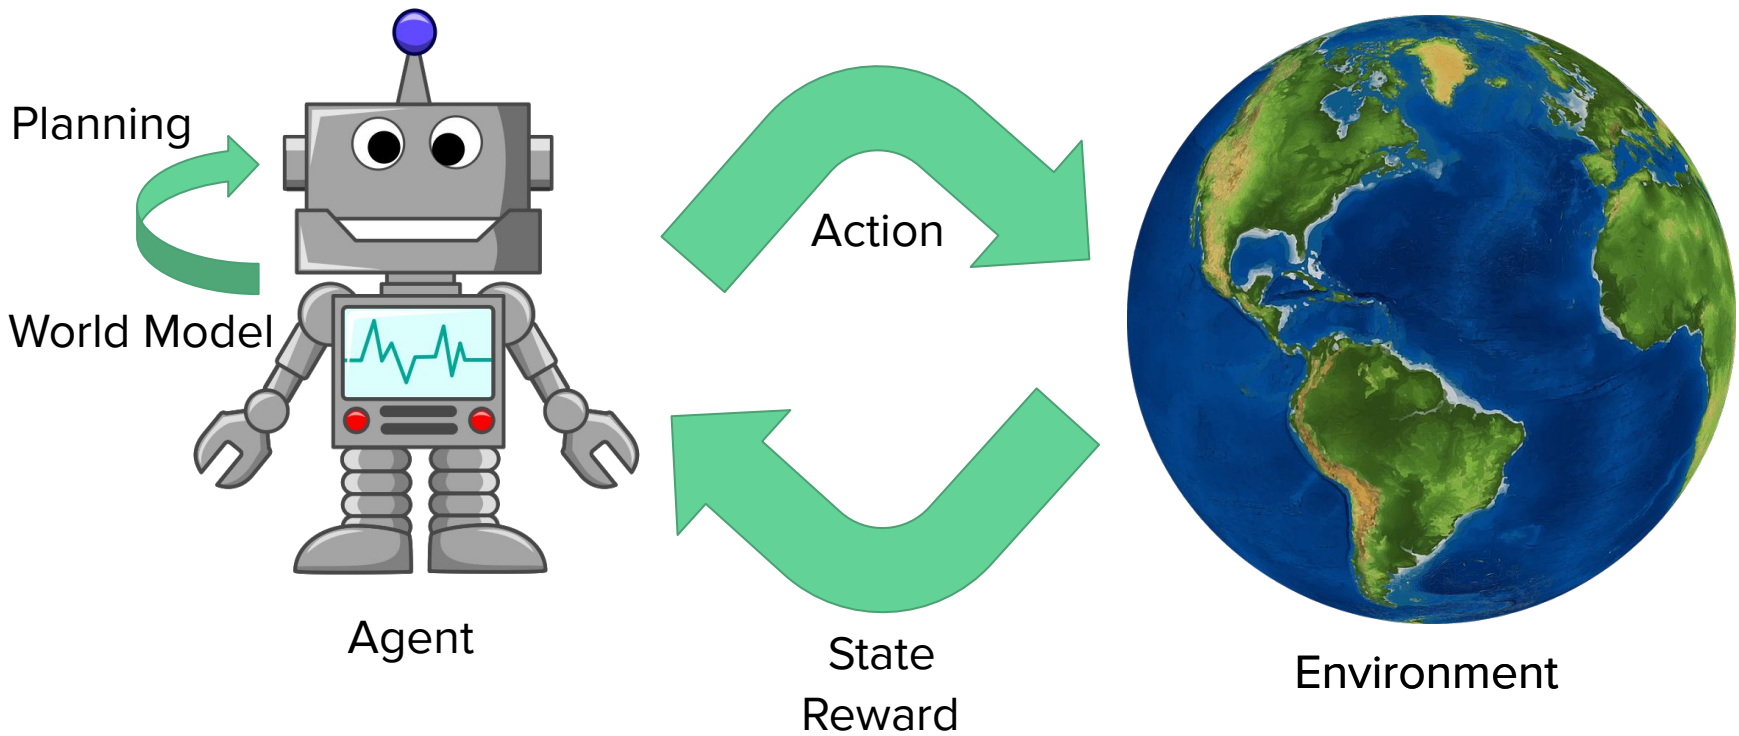
$$Q^\pi(s_t, a_t) = \mathbb{E}_{a_t}[r_t] + \gamma V^\pi(s_{t+1})$$

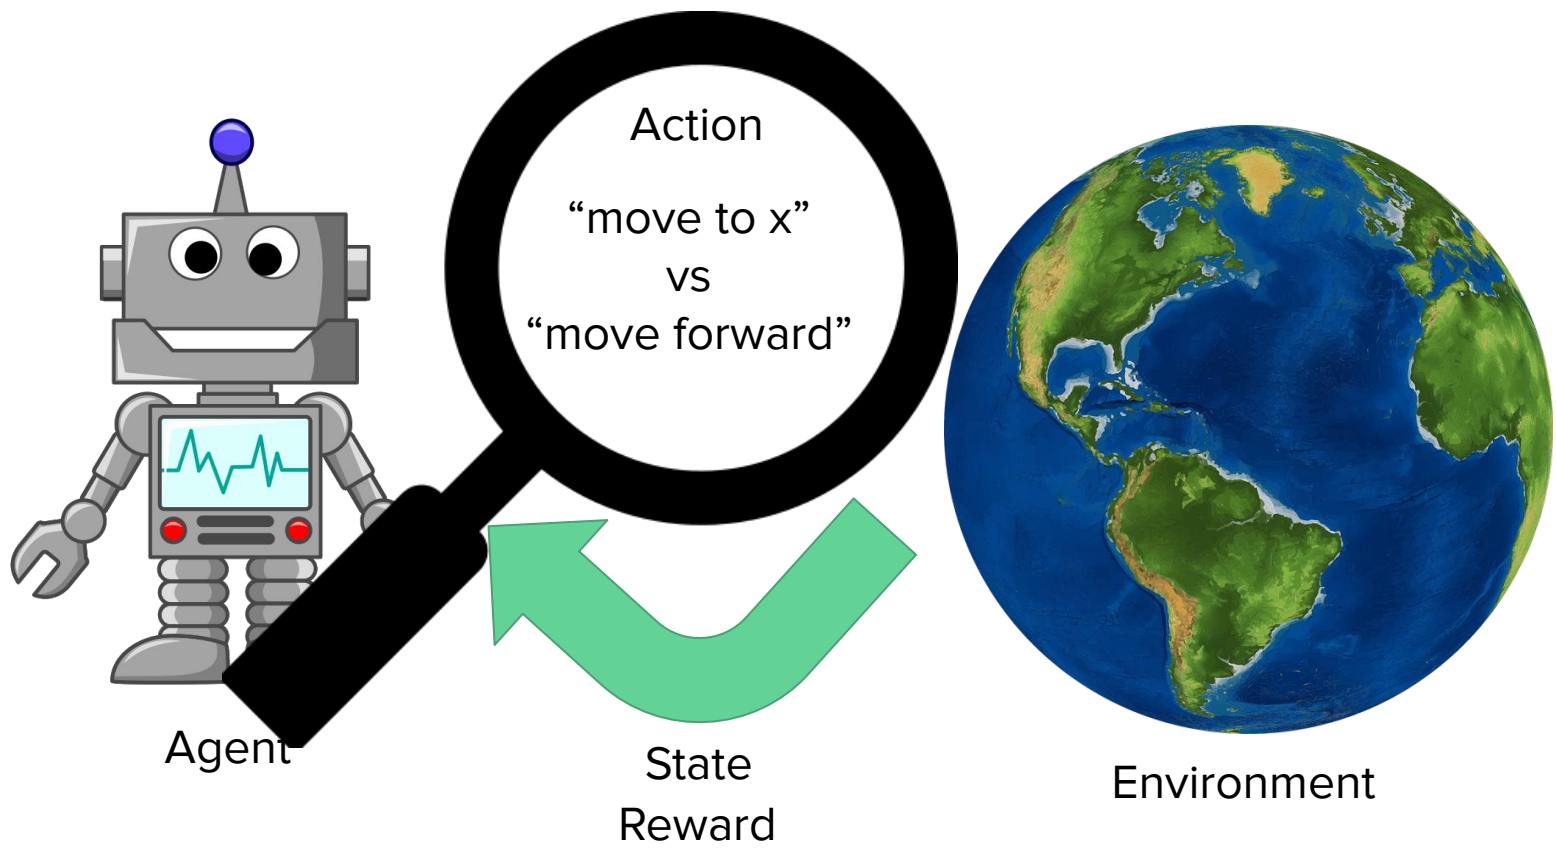
$$V^*(s) = \max_a Q^*(s, a)$$

$$Q^*(s_t, a_t) = \mathbb{E}_{a_t}[r_t] + \gamma \max_a Q^*(s_{t+1}, a)$$



Environment






Action

“move to x”
vs
“move forward”


Agent

State
Reward

Environment



Agent




State
Reward


non-differentiable



Environment




Agent

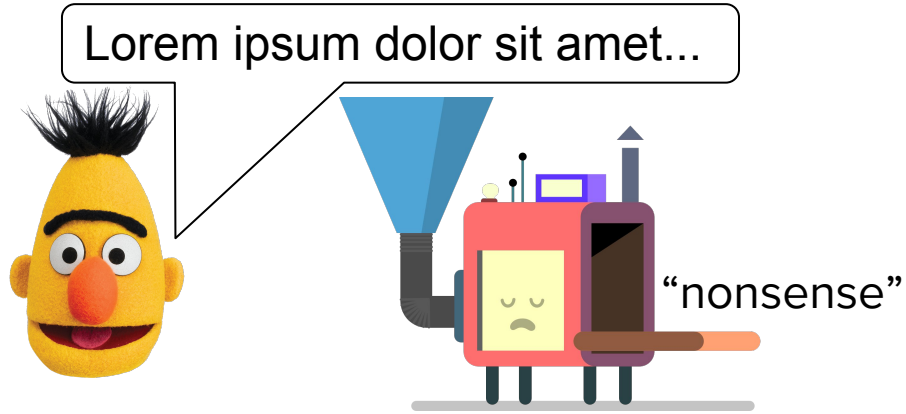


State
Reward


non-differentiable




Natural Language Processing



Graph Neural Networks




Reinforcement Learning



Algorithmic Learning




Graph Neural Networks



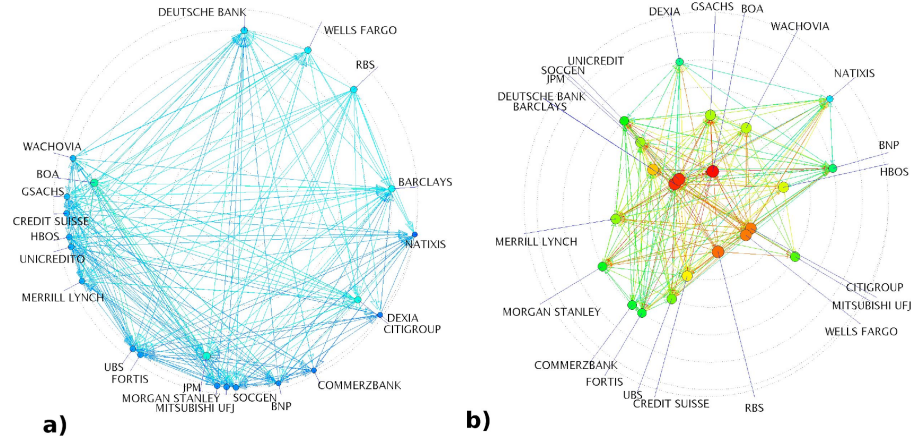
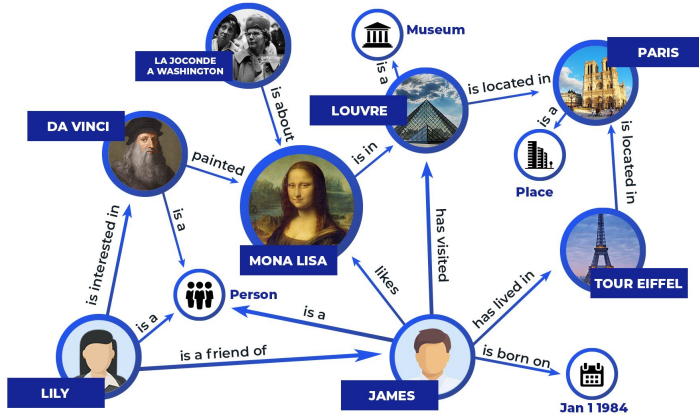
The caffeine molecule


chemical name: 1, 3, 7-trimethylxanthine
 chemical formula: $C_8H_{10}N_4O_2$




C — carbon atom
 H — hydrogen atom
 N — nitrogen atom
 O — oxygen atom
 CH₃ — methyl radical

© 2010 Encyclopædia Britannica, Inc.







Neural Message Passing





K-hop neighbourhood



Readout



GNNs vs CNNs



What are they good at?



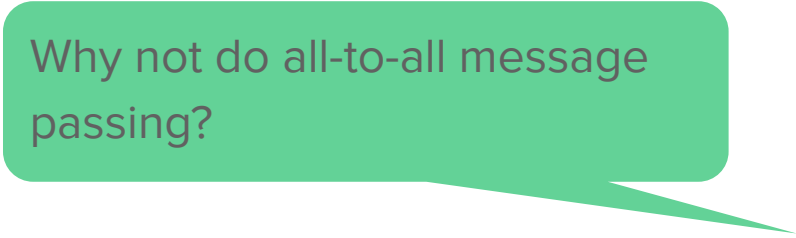
What are they **not** good at?

Theoretical Limitations

K-hop neighbourhood


WL test for isomorphism

LOCAL/CONGEST




Why not do all-to-all message passing?


Oversmoothing



Graph Generation



Simulation



Algorithmic Learning


$3 + 4 = ?$

$18467238957 + 67836423785 = ?$



6.9809932742

😱?!!




Algorithmic Learning

1 4 7 10 ?




1 1 2 3 5 8 ?

13

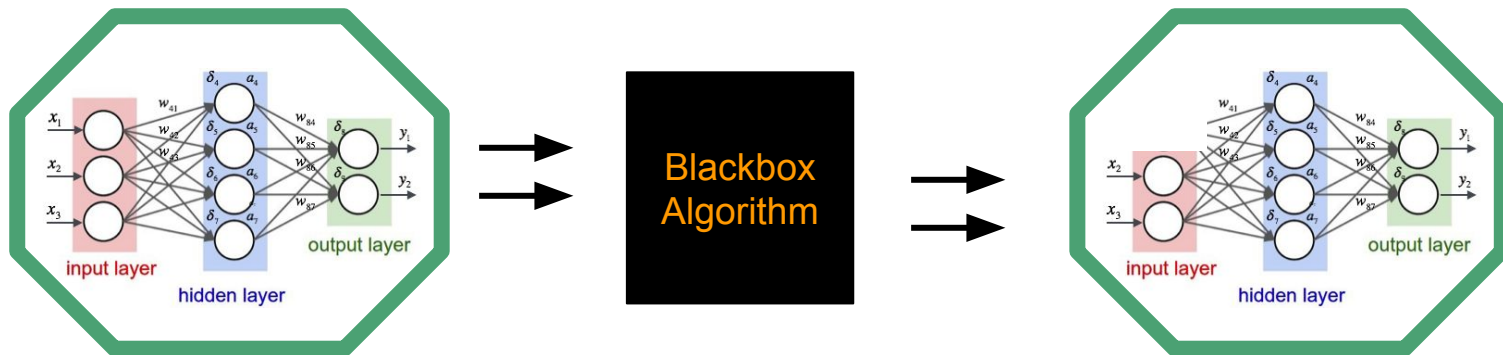
8?



Augmenting NNs with Memory



Augmenting NNs with Algorithms



Augmenting Algorithms with NNs

```
def knapsack(items, capacity):  
    if len(items) == 0:  
        return 0  
    first, *rest = items  
    take = 0
```

```
    skip = knapsack(rest, capacity)  
    return max(take, skip)
```

