


Structuring Neural Representations with the Entity-Relation Meta-Prior

SPEAKER: ANDREA VALENTI


COMPUTATIONAL INTELLIGENCE AND MACHINE LEARNING GROUP

DIPARTIMENTO DI INFORMATICA - UNIVERSITÀ DI PISA

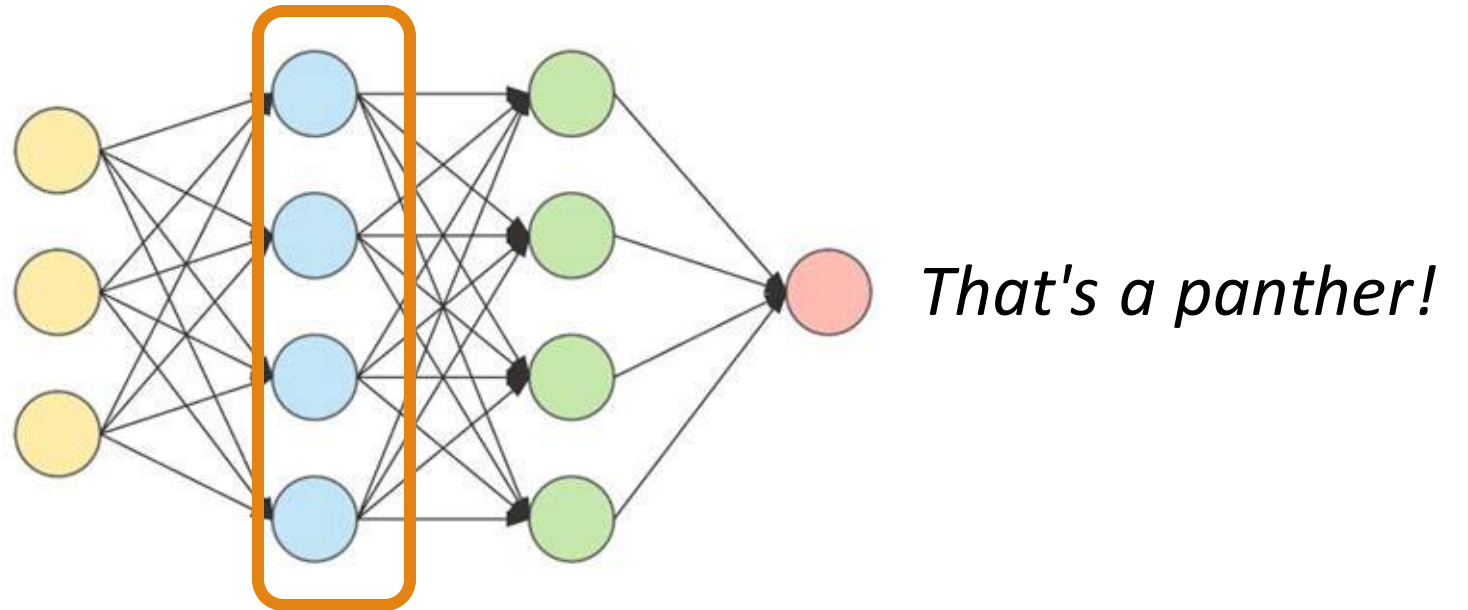


UNIVERSITÀ DI PISA

Distributed Representations




Distributed Representations



0.534	0.929	0.002	-0.824
--------------	--------------	--------------	---------------


Distributed Representations

FULLY SUPERVISED LEARNING



Effective, but very task-specific.

UNSUPERVISED PRE-TRAINING




General, but not very effective on any specific task.

Is there a way to get around this trade-off?

Natural Intelligence as Inspiration

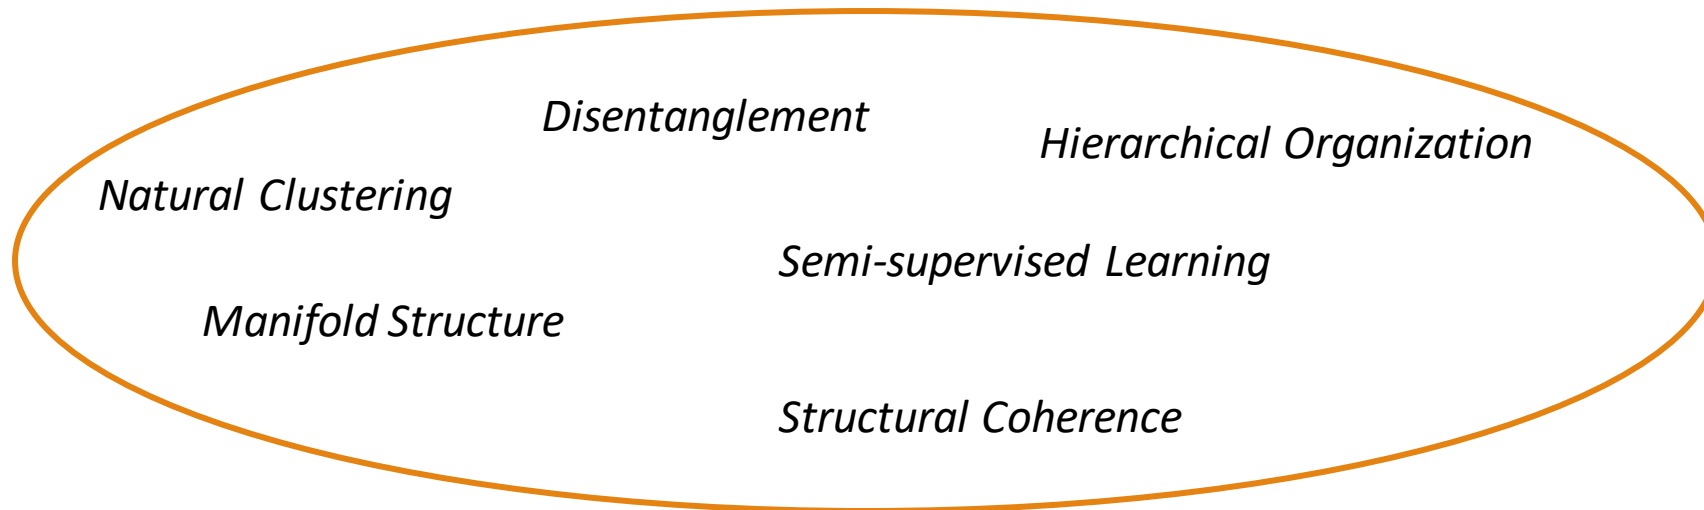
Natural intelligence is incredibly effective at re-using **prior knowledge** for new tasks.




[2018, R. Dubey et al., *Investigating Human Priors for Playing Video Games*]

Meta-Priors

Meta-priors are generic properties of the learned representations, reflecting assumptions that are expected to hold for all the possible downstream tasks [2013, Bengio et al.].




Imposing the appropriate meta-priors can increase generality without a loss in effectiveness.



What is an appropriate meta-prior?


In-depth investigation of human biases.






What is an appropriate meta-prior?

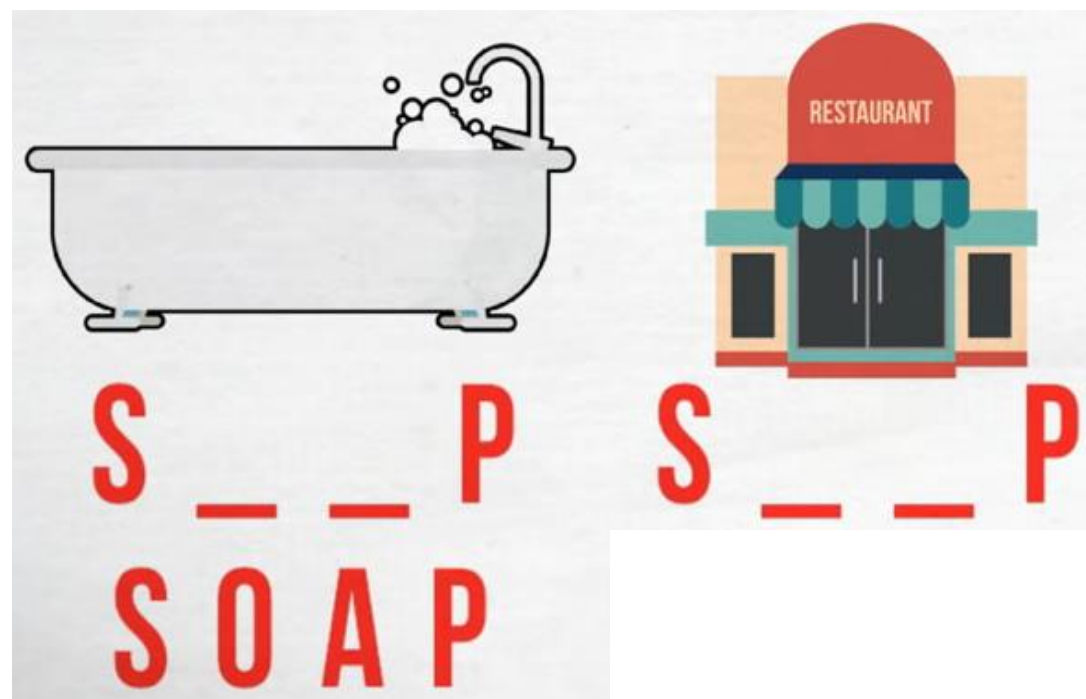
In-depth investigation of human biases.






What is an appropriate meta-prior?


In-depth investigation of human biases.





What is an appropriate meta-prior?

In-depth investigation of human biases.



Two systems of Thinking

SYSTEM 1: FAST THINKING

Quick, intuitive, instinctive.



SYSTEM 2: SLOW THINKING

Slow, effort-taking, deliberate.

```
num_rectangles = 1000000000

delta_x = (1 - 0) / num_rectangles
x = 0
pi = 0

while x < 1:
    f_x = math.sqrt(1 - math.pow(x, 2))
    pi += f_x * delta_x
    x += delta_x
pi = 4 * pi


print("Estimate: " + str(pi))
print("Actual:   " + str(math.pi))
```

Two systems of Thinking

SYSTEM 1: FAST THINKING

Efficient processing of low-level perceptual information.

Deep learning models are quite good on these types of tasks.



SYSTEM 2: SLOW THINKING

Often involves the interplay of symbolic entities.

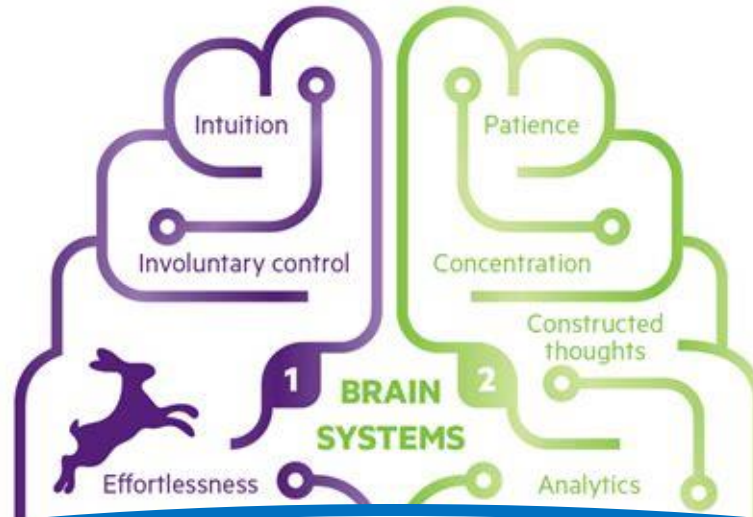
Deep learning models still fail to deliver satisfying performance.

Two systems of Thinking

SYSTEM 1: FAST THINKING

Efficient processing of low-level perceptual information.

Deep learning models are quite good on these types of tasks.



SYSTEM 2: SLOW THINKING


Often involves the interplay of symbolic entities.

Deep learning models still fail to deliver satisfying performance.


Can we use meta-priors to make deep learning models better at System 2 tasks?

A New Meta-Prior

Entity-Relation (ER) Meta-prior. *A good representation should be able to extract the symbolic entities that are involved in a set of tasks, along with the relations between them.*



Preliminary Results



Preliminary Results

Sum:

9	9	0	0	9	9
5	7	3	6	9	3
9	0	0	6	9	6
4	7	4	4	9	1

Subtraction:

3	6	3	6	0	0
5	0	3	1	1	9
8	6	2	2	6	4
4	6	1	9	2	7

Multiplication:

8	6	4	8
3	2	0	6
7	6	4	2
2	7	1	4


Wrap-up and Conclusion

Appropriate **meta-priors** can be used for enforcing specific properties on the learned representation.

The **integration of symbolic and sub-symbolic representations** can enable machine learning to tackle new, more advanced tasks (System 2 tasks).

The **ER Meta-prior** can provide a guide on how to reach such integration.

Thank you for your attention!



UNIVERSITÀ DI PISA