Contrasting artificial intelligence with human intelligence

In search of alternatives for the future of AI

Jean-Louis Dessalles
Telecom Paristech

jl@dessalles.fr
www.dessalles.fr

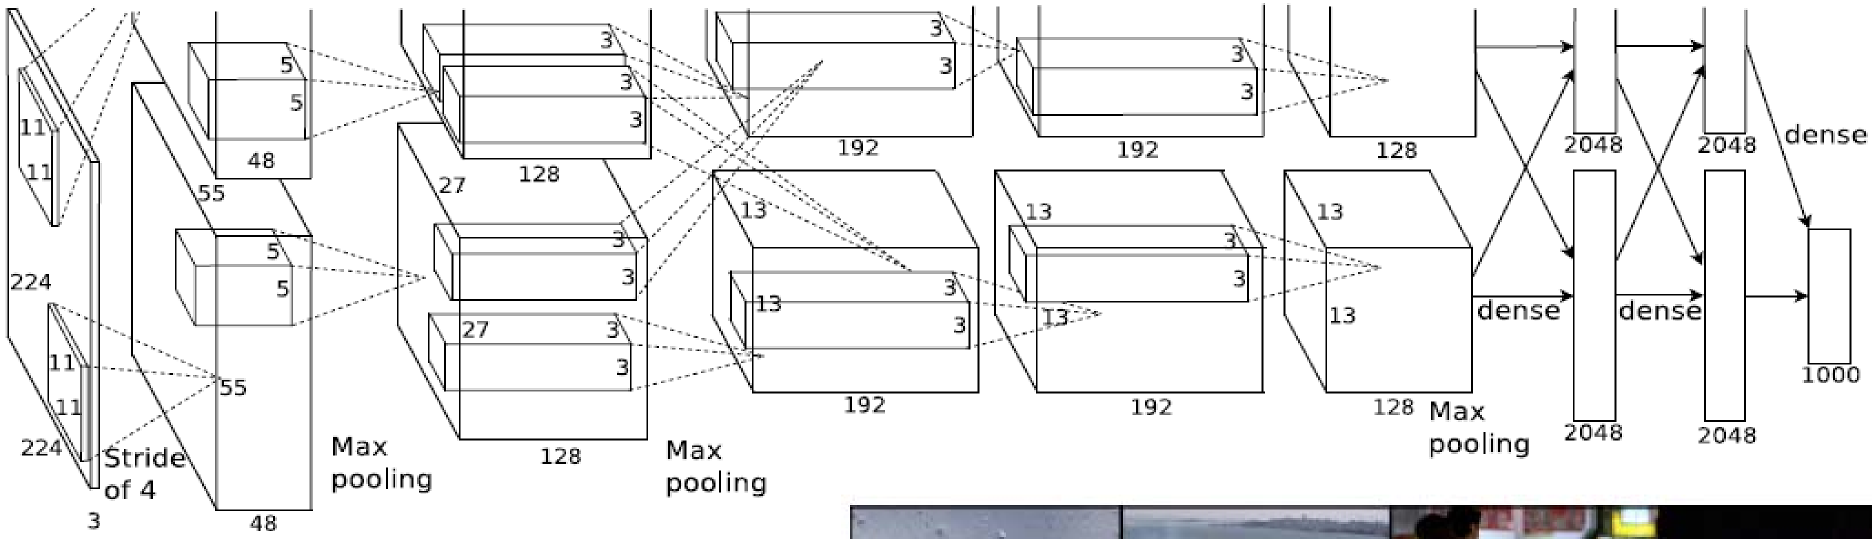

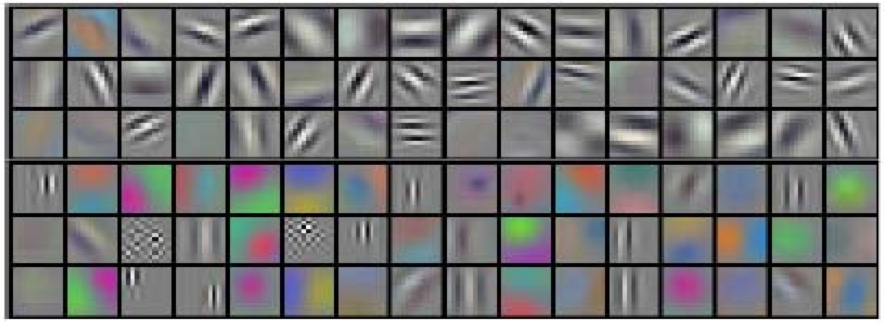

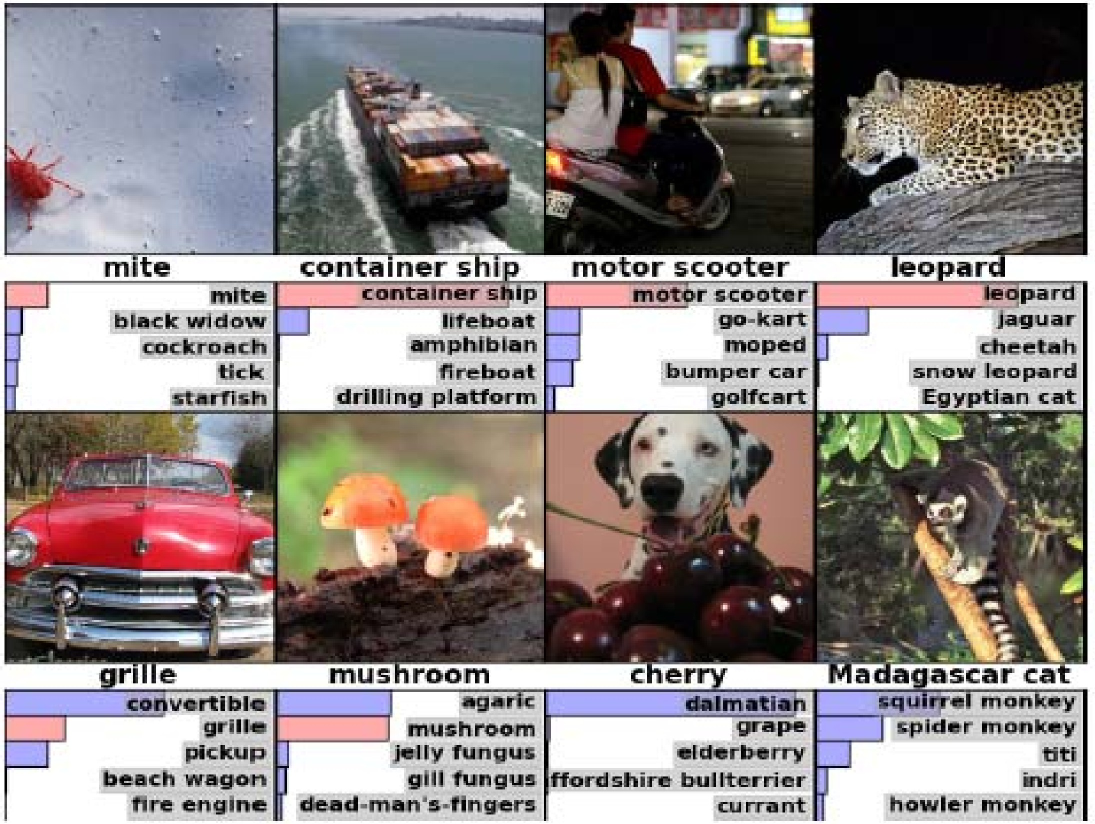

Krizhevsky, A., Sutskever, I. & Hinton, G. E. (2012). [Imagenet classification with deep convolutional neural networks](). *NIPS 2012*, 1097-1105.

1.2 million images, 1000 clases, 650000 neurons,
60 million parameters

Mais ultimement, n'est ce pas un peu un position "religieuse" que de penser qu'aucune "loss function" ne pourra remplacer un jour l'intelligence "humaine"?

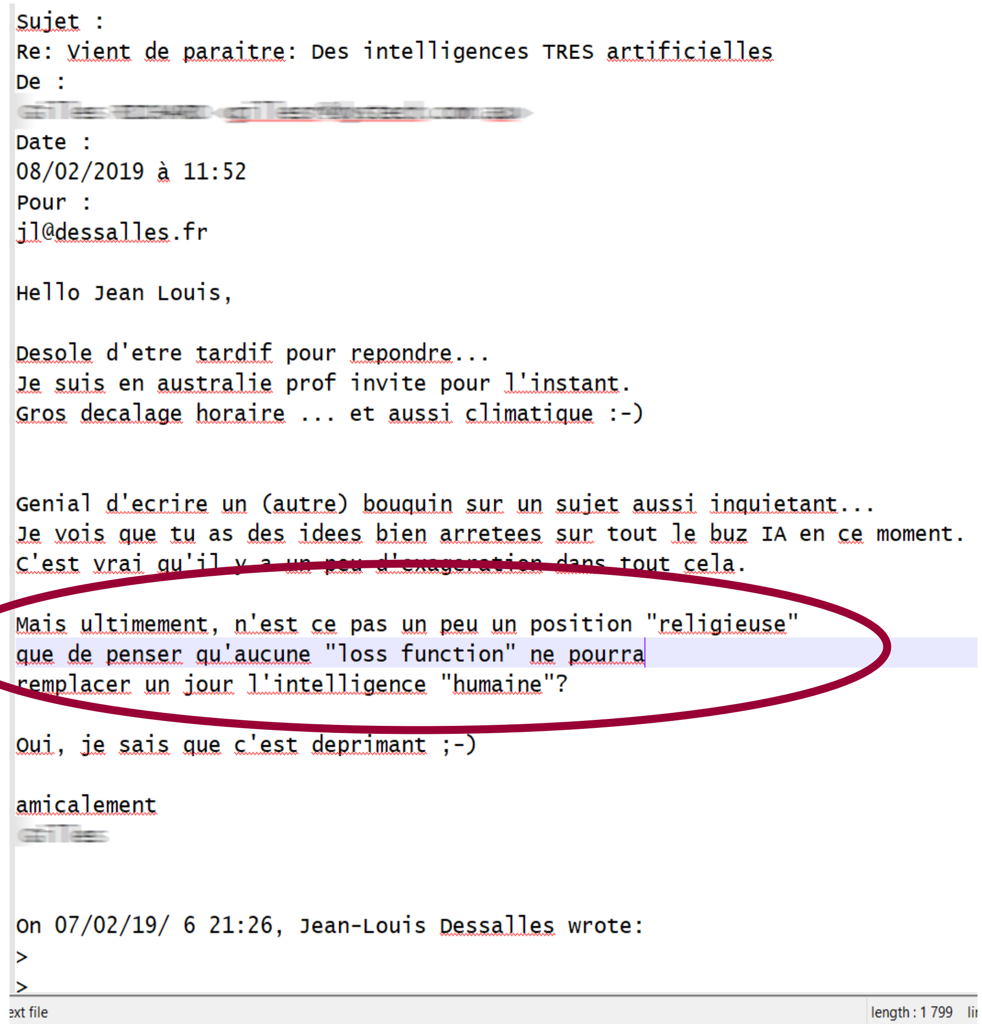
Sujet :
Re: Vient de paraitre: Des intelligences TRES artificielles
De :
[redacted]
Date :
08/02/2019 à 11:52
Pour :
jl@dessalles.fr

Hello Jean Louis,

Desole d'etre tardif pour repondre...
Je suis en australie prof invite pour l'instant.
Gros decalage horaire ... et aussi climatique :-)

Genial d'ecrire un (autre) bouquin sur un sujet aussi inquietant...
Je vois que tu as des idees bien arretees sur tout le buz IA en ce moment.
C'est vrai qu'il y a un peu d'exageration dans tout cela.

Mais ultimement, n'est ce pas un peu un position "religieuse"
que de penser qu'aucune "loss function" ne pourra
remplacer un jour l'intelligence "humaine"?

Oui, je sais que c'est deprimant ;-)

amicalement
[redacted]

On 07/02/19/ 6 21:26, Jean-Louis Dessalles wrote:
>
>

ext file length : 1 799 li

But ultimately, isn't ita bit of a "religious" position to think that no loss function will be able to replace "human" intelligence one day?

“delete all images that are duplicated”

Who I am...

- Telecom Paristech (IP Paris)
- Artificial intelligence
  - Grail: Reverse-engineer the human mind
  - More concerned with language (Semantics, relevance)
    (but also emergence, origins of language, evolution, social signals…)
- Current topics
  - Simplicity Theory
  - Contrast

Contrasting artificial intelligence with human intelligence

- Ten limitations of deep learning

- Simplicity Theory: An AIT approach to intelligence
- Contrast: a missing mechanism in the current AI toolbox

- Conclusion: mechanisms that operate on the fly

Ten limitations of deep learning

Caveat:

Several issues mentioned in this section
(but not all of them)
are regurlarly raised by scholars.

e.g.: Marcus, G. (2018). Deep learning: A critical appraisal. ArXiv, 1801.006.

Ten limitations of deep learning – 1. Continuity

Schaffer, C. (1994).
A conservation law for generalization performance.
Proc. of the Machine Learning Conf., 259-265. Rutgers University.

- Bias is unavoidable
- NN are biased to learn continuous functions
- grant a bank loan… ok, maybe
- criminal investigation… not ok!

Ten limitations of deep learning - 2. Isotropy

Dessalles, J.-L. (1998).
Characterising innateness in artificial and natural learning.
ECML Workshop on Learning in Humans and Machines, 6-17.
Chemnitz: Technische Universität Chemnitz - CSR-98-03.
www.dessalles.fr/papers/Dessalles_98042402.pdf

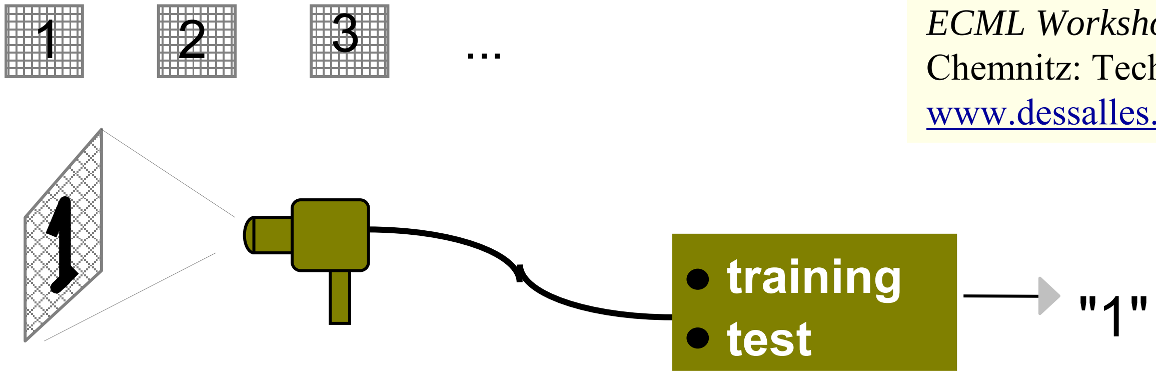

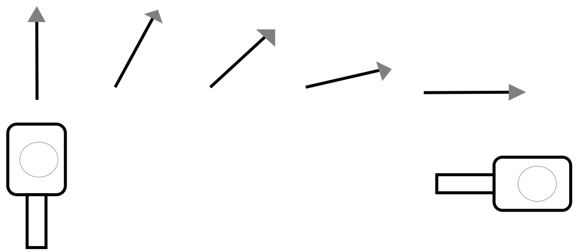

Isotropic systems learn harmonious classifications more easily

Counterexample: syntactic embedding

$$L(T)[x] = L(\rho(T))\,[\rho(x)]$$

Ten limitations of deep learning - 3. Large data sets

- A child learns about four+ new words a day

Goulden, R., Nation, P. & Read, J. (1990). How large can a receptive vocabulary be? *Applied linguistics*, *11* (4), 341-363.

- Statistical learning achieves one-shot (or zero-shot) learning !

Zhang, L., Xiang, T. & Gong, S. (2017). Learning a deep embedding model for zero-shot learning. *ArXiv*, *1611.05088v3*.

- Ex: music style

Lake, B., Salakhutdinov, R., Gross, J. & Tenenbaum, J. B. (2011). One shot learning of simple visual concepts. CogSci.

- CtrEx:
  - "buffet plate clip for wine glass"
  - "jealous", "prevent", "around", "chase", "abdicate"

Lake, B., Salakhutdinov, R., Gross, J. & Tenenbaum, J. B. (2011). [One shot learning of simple visual concepts](). *COGSCI-2011*, 2568-2573.

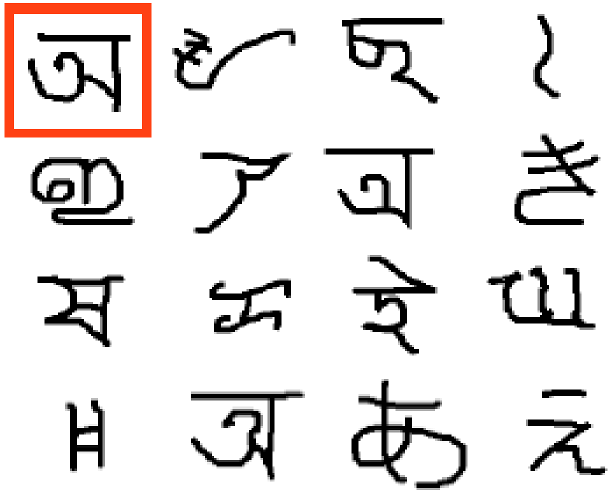

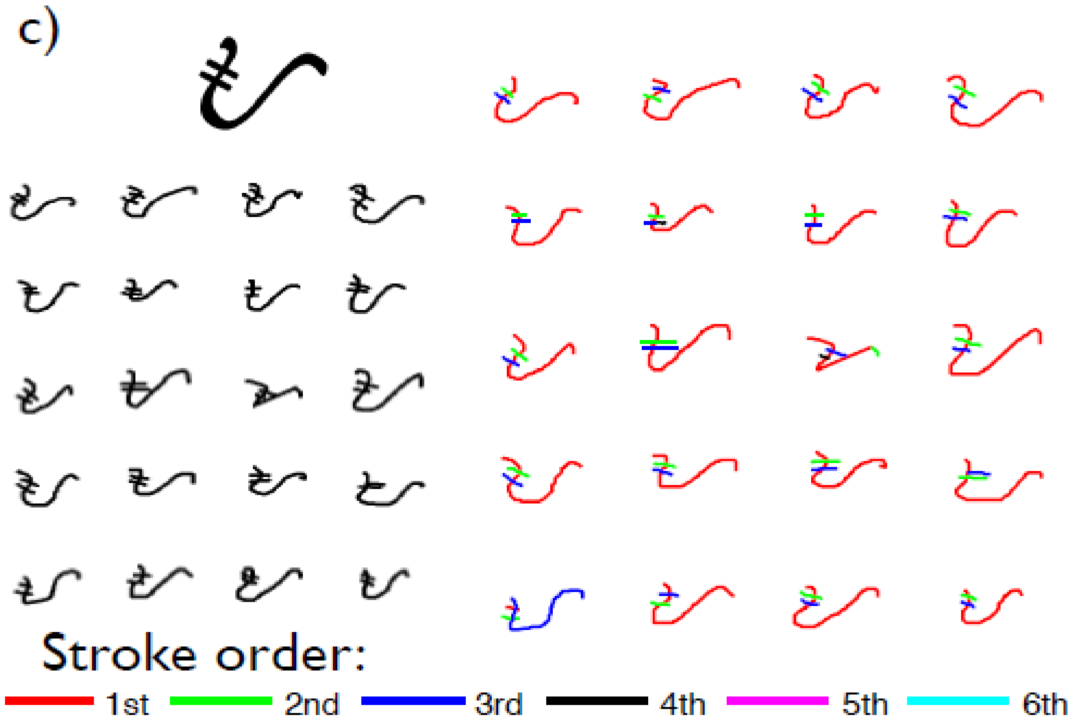

Ten limitations of deep learning – 4. Relations

Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R. S. & Bengio, Y. (2015). Show, attend and tell: Neural image caption generation with visual attention. *32nd International Conference on Machine Learning*, 2048-2057.

A woman is throwing a frisbee in a park.

A stop sign is on a road with a mountain in the background.

A little girl sitting on a bed with a teddy bear.

A group of people sitting on a boat in the water.

Ten limitations of deep learning – 4. Relations

Santoro, A., Raposo, D. *et al.* (2017). A simple neural network module for relational reasoning. *NIPS 2017*, 4967-4976.

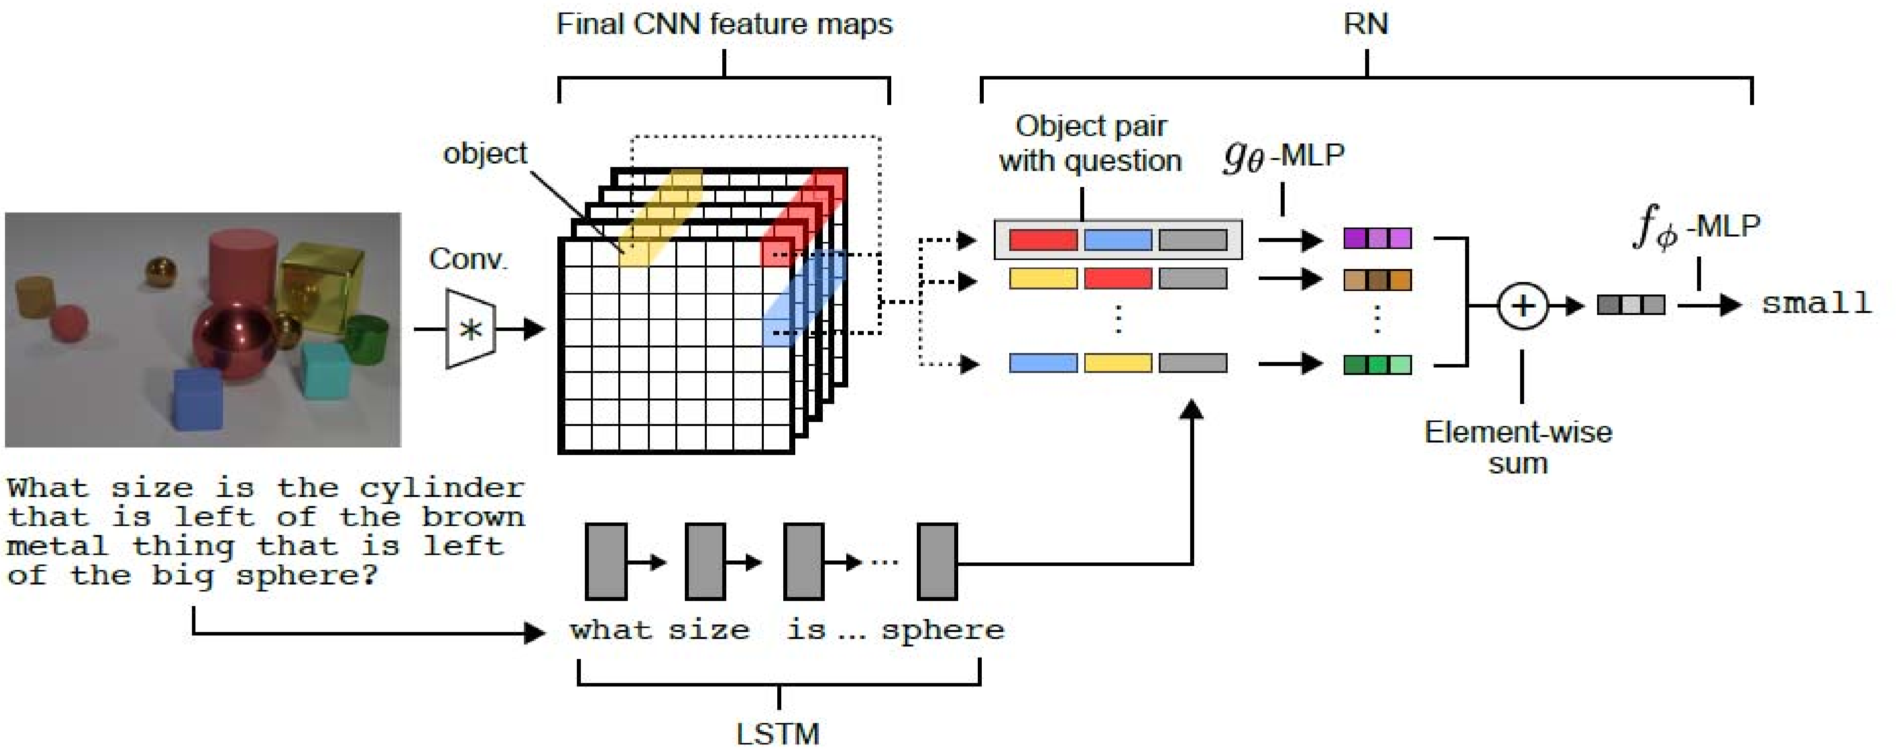

Ten limitations of deep learning – 4. Relations

"prevent", "around", "chase", "abdicate"

Ten limitations of deep learning – 5. Structures

1,2,2,3,3,3,4,4,4,4

abc is to **abd** as **ppqqrr** is to … ?

Mikolov, Tomas., Sutskever, I., Chen, K., Corrado, G. & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems 26 (NIPS 2013)*, 3111-3119.

"Madrid" – "Spain" + "France" = "Paris"

Ten limitations of deep learning – 6. Exceptions

Turner, R., Ghahramani, Z. & Bottone, S. (2010). [Fast online anomaly detection using scan statistics](). *MLSP 2010*, 385-390.

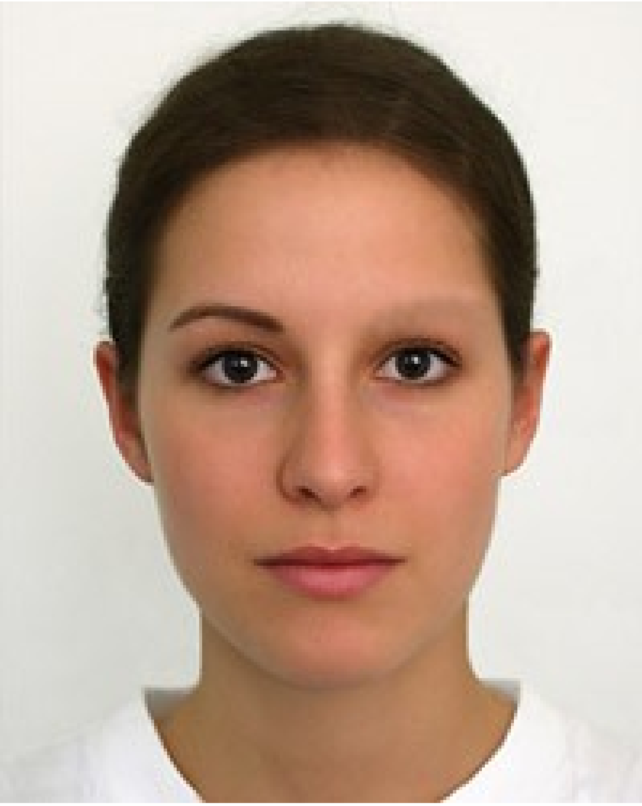

- Anomaly detection
- Curse of Dimensionality
- The missing eyebrow
- Buying a used car

Ten limitations of deep learning – 7. Negation

- For a DNN, "this is a cat" does not mean "this is not a dog" nor
- Inconsistency: House without door
- Explanation: "This is too big to be a cat"

Ten limitations of deep learning – 8. Narrow expertise

Silver, D., Schrittwieser, J. & et al., (2017).
[Mastering the game of go without human knowledge]().
Nature, *550* (7676), 354-359.

Esteva, A. *et al.* (2017).
Dermatologist-level classification of skin cancer with deep neural networks.
Nature, *542* (7639), 115-118.

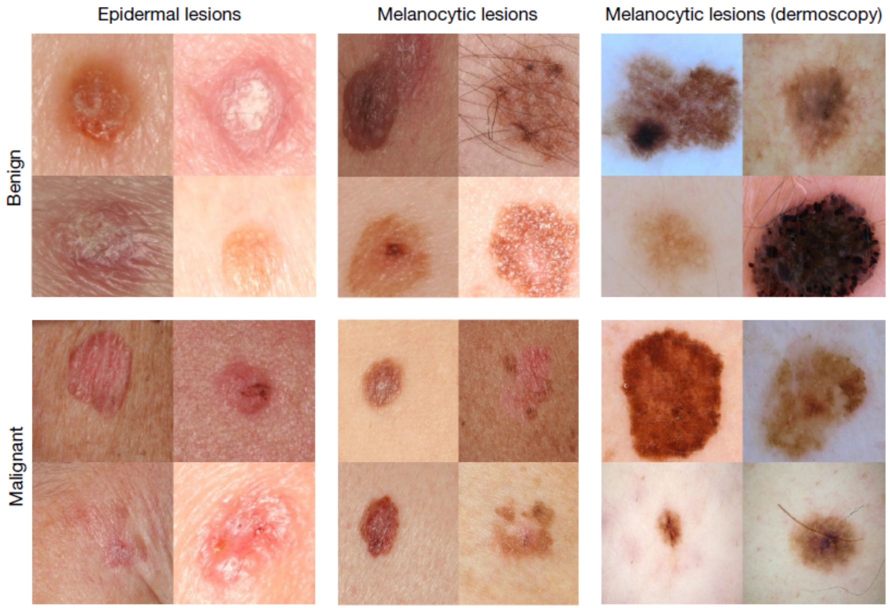

Ten limitations of deep learning – 9. No sense making

'One swallow does not thirst quench'

(alluding to 'One swallow does not a summer make')

'Une hirondelle n'aspire pas la soif'

Hofstadter, D. R. (2018).
The shallowness of Google Translate.
The Atlantic, , Jan, 30.

semantic proximity ≠ semantics

Ten limitations of deep learning – 10. No systematicity

- Behind the rock *vs.* behind the car

Fodor, J. A. & Pylyshyn, Z. W. (1988).
[Connectionism and cognitive architecture: A critical analysis](). *Cognition*, *28* (1-2), 3-71.

- smaller(*m*, *n*)   larger(*n*, *m*)

Weber, N., Shekhar, L. & Balasubramanian, N. (2018).
The Fine Line between Linguistic Generalization and Failure in Seq2Seq-Attention Models. *ArXiv*, 1805.014.

DNN have access to extensions,
not to intensions.

Contrasting artificial intelligence with human intelligence

- Ten limitations of deep learning

- Simplicity Theory: An AIT approach to intelligence
- Contrast: a missing mechanism in the current AI toolbox

- Conclusion: mechanisms that operate on the fly

Algorithmic approach to AI

"comprehension is compression"

$$C(x) = \min_p \{l(p) : M(p) = x\}$$

Chaitin, G. J. (2004).
On the intelligibility of the universe and the notions of simplicity, complexity and irreducibility.
Grenzen und Grenzüberschreitungen, XIX, 517-534.

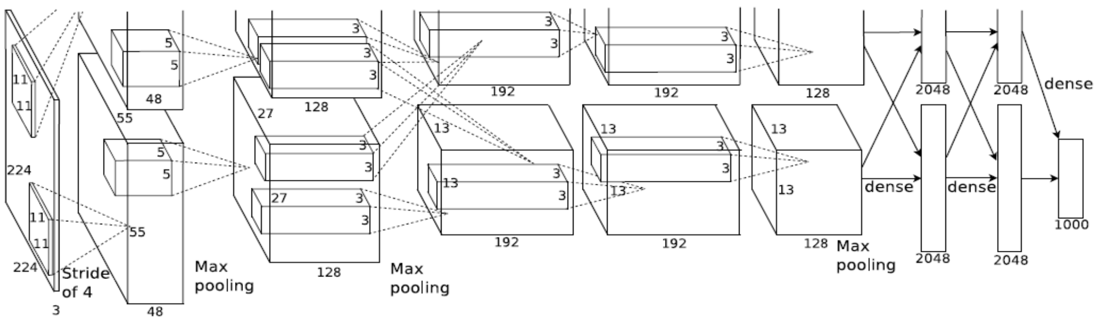

n classes

$\log_2(n)$ bits spared for each correctly classified example

Algorithmic approach to AI

most probable continuation?

1 2 2 3 3 3 4 4 4 4 5 5 5 5 5

$$n^{*n}$$

C minimal
→ max. compr.

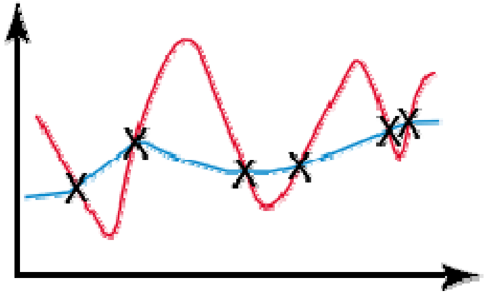

Solomonoff, R. J. (1964).
A Formal Theory of Inductive Inference.
Information and Control, *7* (1), 1-22.

Search: seq:1,2,2,3,3,3,4,4,4,4

Displaying 1-10 of 46 results found.

n appears n times; floor(sqrt(2n) + 1/2).
1, 2, 2, 3, 3, 3, 4, 4, 4, 4, 5, 5, 5, 5, 5, 6, 6, 6, 6, 6, 6, 7, 7, 7, 7, 7, 7, 7, 8, 8, 8, 8, 8, 8, 8, 8, 9, 9, 9, 9, 9, 9, 9, 9, 9, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 13, 13, 13, 13, 13, 13

n appears partition(n) times.
0, **1, 2, 2, 3, 3, 3, 4, 4, 4, 4,** 4, 5, 5, 5, 5, 5, 5, 5, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 10

Number of digits in lazy-Fibonacci-binary representation of n.
1, **1, 2, 2, 3, 3, 3, 4, 4, 4, 4,** 4, 5, 5, 5, 5, 5, 5, 5, 5, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9

Write n = C(i,3)+C(j,2)+C(k,1) with i>j>k>=0; sequence gives j values.
1, 1, 2, 2, 1, 2, 2, 3, 3, 3, **1, 2, 2, 3, 3, 3, 4, 4, 4, 4, 1, 2, 2, 3, 3, 3, 4, 4, 4, 4,** 5, 5, 5, 5, 5, **1, 2, 2, 3, 3, 3, 4, 4, 4, 4,** 5, 5, 5, 5, 5, 6, 6, 6, 6, 6, 6, **1, 2, 2, 3, 3, 3, 4, 4, 4, 4,** 5, 5, 5, 5, 5, 6, 6, 6, 6, 6, 6, 7, 7, 7, 7, 7, 7, 7, **1, 2, 2, 3, 3, 3, 4, 4, 4, 4,** 5, 5, 5, 5, 5, 6, 6, 6, 6, 6, 6, 7, 7, 7, 7, 7, 7, 7, 8, 8, 8, 8, 8, 8, 8, 8, 1, 2, 2

Positive integers a for which there is a 10-Pythagorean triple (a,b,c) satisfying a<=b.
1, 2, 2, 3, 3, 3, 4, 4, 4, 4, 5, 5, 5, 5, 5, 6, 6, 6, 6, 6, 7, 7, 7, 7, 7, 8, 8, 8, 8, 8, 8, 9, 9, 9, 9, 9, 10, 10, 10, 10, 10, 10, 10, 10, 10, 11, 11, 11, 11, 11, 12, 12, 12, 12, 12, 12, 12, 12, 13, 13, 13, 13, 14, 14, 14, 14, 14, 14, 14, 14, 15, 15, 15, 15, 15, 15, 15, 15

The Kruskal-Macaulay function M_2(n).
0, **1, 2, 2, 3, 3, 3, 4, 4, 4, 4,** 5, 5, 5, 5, 5, 6, 6, 6, 6, 6, 6, 7, 7, 7, 7, 7, 7, 7, 8, 8, 8, 8, 8, 8, 8, 8, 9, 9, 9, 9, 9, 9, 9, 9, 9, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 13, 13, 13, 13, 13, 13

Triangle read by rows in which row n contains a finite triangle as shown below.
1, 1, 2, 2, 2, 3, 2, 1, 2, 2, 1, 1, 2, 2, 3, 3, 3, 3, 4, 4, 3, 3, 4, 5, 4, 3, 2, 3, 4, 4, 3, 2, 1, 2, 3, 3, 3, 2, 1, **1, 2, 2, 3, 3, 3, 4, 4, 4, 4,** 4, 5, 5, 5, 4, 4, 5, 6, 6, 5, 4, 4, 5, 6, 7, 6, 5, 4, 3, 4, 5, 6, 6, 5, 4, 3, 2, 3, 4, 5, 5, 5, 4, 3, 2, 1, 2, 3, 4, 4, 4, 4, 3, 2, 1, **1, 2, 2, 3, 3, 3, 4, 4, 4, 4,** 5

Algorithmic approach to AI

abc is to **abd** as **ppqqrr** is to …   **ppqqss**

Cornuéjols, A. (1996).
Analogie, principe d'économie et complexité algorithmique.
Actes des 11èmes Journées Françaises de l'Apprentissage.

$$\text{'ppqqss'} = \text{argmin}_x\, C(\text{'abc'}, \text{'abd'}, \text{'ppqqrr'}, x)$$

(talk, talked) → (solve, solved)

Murena, P.-A., Dessalles, J.-L. & Cornuéjols, A. (2017).
A complexity based approach for solving Hofstadter's analogies.
ICCBR-WS 2017, 53-62. Trondheim, Norway.

Algorithmic approach to AI

- [Marcus Hutter's AIXI]

$$a_k := \arg\max_{a_k} \sum_{o_k r_k} \ldots \max_{a_m} \sum_{o_m r_m} \underbrace{[r_k + \ldots + r_m]}_{\text{future reward}} \sum_{q\,:\,U(q,a_1..a_m)=o_1 r_1..o_m r_m} 2^{-\ell(q)}$$

action

future perceptions

prevision program

Hutter, M. (2005).
Universal artificial intelligence: Sequential decisions based on algorithmic probability.
Berlin: Springer.

Algorithmic approach to AI

- Transfer learning

Murena, P.-A. (2019).
Minimum complexity knowledge transfer in artificial learning.
Phd Thesis, Telecom ParisTech, Universite Paris-Saclay.

Algorithmic approach to cognitive science

Complexity $C(s)$ of s :
size of the smallest available description of s

Chater, N. (1999).
The search for simplicity:
A fundamental cognitive principle?
The Quarterly J. of Exp. Psychol., *52* (A), 273-302.

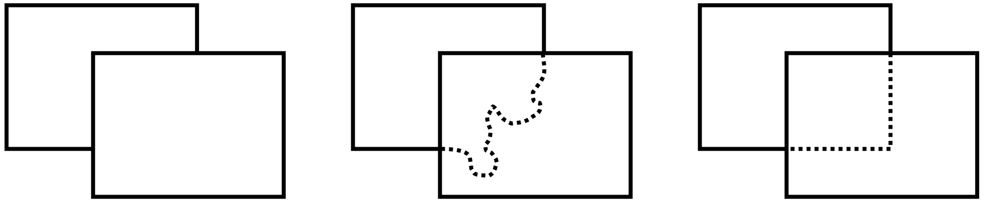

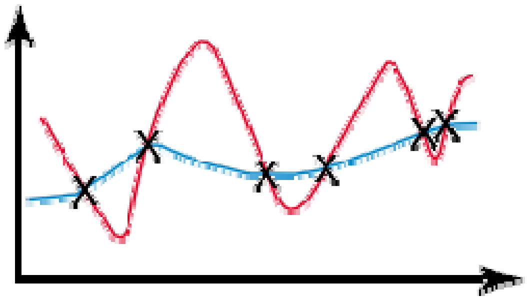

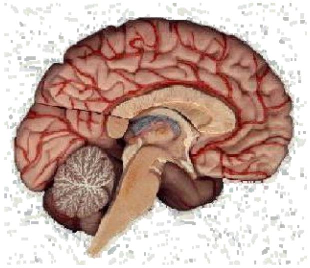

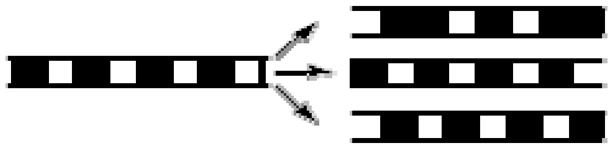

$$C(x) = \min_p \{l(p) : M(p) = x\}$$

Simplicity theory

Unexpectedness = expected complexity – observed complexity

$$U = C_{exp} - C_{obs}$$

complexity drop

From: www.hockinghills.com/comfort/

From: iciouailleurs.free.fr/HautJura/hautjura.html

Simplicity theory

Unexpectedness = expected complexity – observed complexity

$$U = C_{exp} - C_{obs}$$

Dessalles, J.-L. (2006).
A structural model of intuitive probability.
7th *Int. Conf. on Cognitive Modeling*, 86-91.

Combinations	Complexity	Probability
1 2 3 4 5 6	3	p/8 × 10^6
34 35 36 37 38 39	6	p/10^6
10 11 12 44 45 46	11	p/32768
7 8 9 37 38 39	12	p/16384
8 9 26 27 28 29	12	p/16384
10 20 30 31 32 33	12	p/16384
1 2 5 6 15 49	14	p/4096
. . .	. . .	. . .
14 24 36 38 42 44	26	p

Simplicity theory

Unexpectedness = expected complexity – observed complexity

$$U = C_{exp} - C_{obs}$$

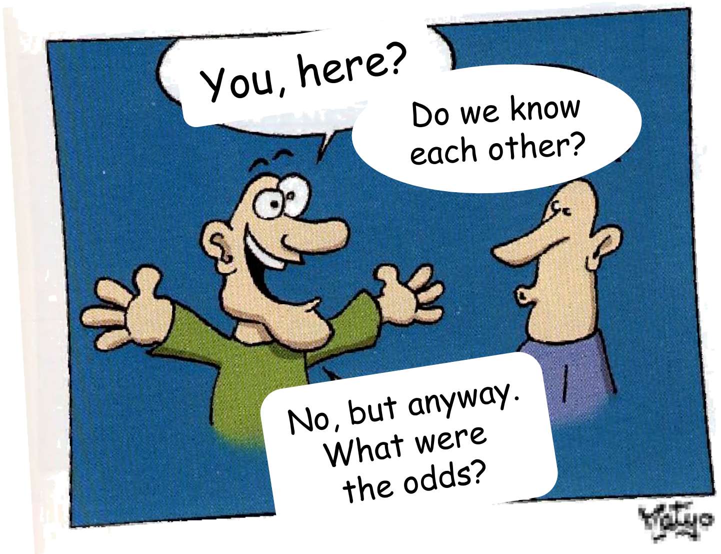

$$U = C(L) - C(P)$$

Complexity of the location

Complexity of the encountered person

Simplicity theory

Unexpectedness = expected complexity – observed complexity

$$U = C_{exp} - C_{obs}$$

- Rarity $U \geq \log N - \log P - C(f) - C(r)$
- Proximity $U = 2 \times \log(R / d)$ $L = \mathrm{argmin}(C(L) + 2\log(d_L))$
- Anomaly $U \geq A(k) - C(f) - C(r)$ $U \geq C(H) - C(f) - C(r)$
- Coincidences $U = C(s_1) - C(s_2|s_1)$
- Relevance $C_w(f(s)) - C(f) > 0$
- Responsibility $C_w(s) - C_w(s \,||\, a)$
- Emotion intensity $E = E_h + U$

Contrasting artificial intelligence with human intelligence

- Ten limitations of deep learning
- Simplicity Theory: An AIT approach to intelligence
- Contrast: a missing mechanism in the current AI toolbox
- Conclusion: mechanisms that operate on the fly

Contrast

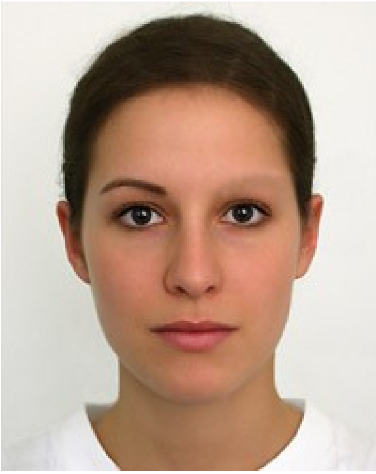

- Anomaly detection
- Curse of Dimensionality
- The missing eyebrow
- Buying a used car

$$\sum \alpha_i \left| x_i^1 - x_i^2 \right|$$

$$C(o) = \lfloor Std(o - P(o)) \rfloor$$

standardization

cleaning

Contrast

- Contrast is a low-dimensionality vector (← cleaning)
- Contrast object with closest prototype
  - Topological decision along that vector → membership or negation
- Do it again with contrasts
  - → predication
  - → explanations

Contrasting artificial intelligence with human intelligence

- Ten limitations of deep learning

- Simplicity Theory: An AIT approach to intelligence
- Contrast: a missing mechanism in the current AI toolbox

- Conclusion: mechanisms that operate on the fly

Mechanisms that operate on the fly

- DNN rely on pre-digested expertise
- Human cognition relies on a variety of mechanims
  - Compression, Complexity drop
  - Contrast
  - Conflict-Abduction-Negation, Aspect, quantification, …
  - Merge, semantic linking, …
- These mechanisms operate on the fly

Mais ultimement, n'est ce pas un peu un position "religieuse" que de penser qu'aucune "loss function" ne pourra remplacer un jour l'intelligence "humaine"?

But ultimately, isn't ita bit of a "religious" position to think that no loss function will be able to replace "human" intelligence one day?

Sujet :
Re: Vient de paraitre: Des intelligences TRES artificielles
De :
[redacted]
Date :
08/02/2019 à 11:52
Pour :
jl@dessalles.fr

Hello Jean Louis,

Desole d'etre tardif pour repondre...
Je suis en australie prof invite pour l'instant.
Gros decalage horaire ... et aussi climatique :-)

Genial d'ecrire un (autre) bouquin sur un sujet aussi inquietant...
Je vois que tu as des idees bien arretees sur tout le buz IA en ce moment.
C'est vrai qu'il y a un peu d'exageration dans tout cela.

Mais ultimement, n'est ce pas un peu un position "religieuse"
que de penser qu'aucune "loss function" ne pourra
remplacer un jour l'intelligence "humaine"?

Oui, je sais que c'est deprimant ;-)

amicalement
[redacted]

On 07/02/19/ 6 21:26, Jean-Louis Dessalles wrote:
>
>

ext file length : 1 799 li

2 December 2014 Last updated at 13:02 GMT

BBC NEWS

Stephen Hawking warns artificial intelligence could end mankind

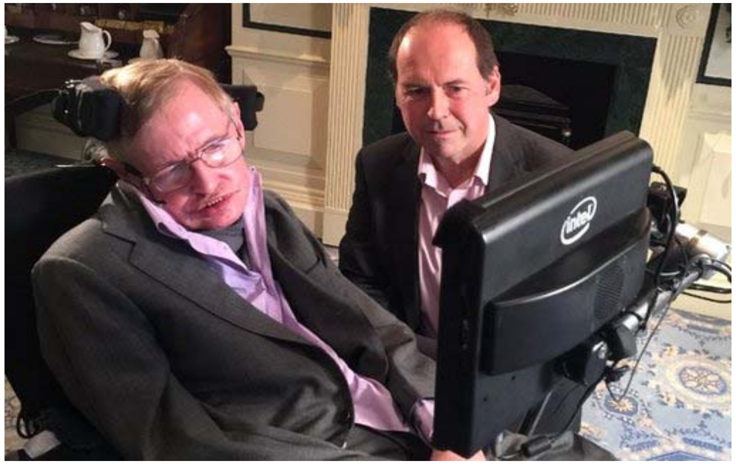

Elon Musk

Bill Gates

INSTITUT POLYTECHNIQUE DE PARIS

TELECOM ParisTech

Thanks for listening

jean-louis @ dessalles.fr

www.dessalles.fr

Visit: www.simplicitytheory.science