Back to the Future: Logic and Machine Learning

John D. Kelleher[1] and Simon Dobnik[2]
[1]Dublin Institute of Technology and [2]University of Gothenburg

Logic and Machine Learning in Natural Language (LaML)
Gothenburg, 12–13 June 2017

CLASP centre for linguistic theory and studies in probability

- Two distinct methodologies for building models of the world
  1. Logic: qualitative, symbolic and driven by domain theory
  2. ML: quantitative, numeric and driven by computational learning theory

tech/cov	wide	narrow
deep	**our goal**	symbolic
shallow	data-based	**useless**

- Discussion already mid-1990s: the rise of statistical learning methods in NLP (Gazdar, 1996; Jones et al., 2000)

- The success of deep neural network (DNN) approaches makes the question of how these two methodologies should be used/related/integrated apposite again

Modularity and Deep Learning

- DNNs are not unconstrained neural networks but rather that these networks have domain/task specific architectures that encode domain theoretic considerations
- DNNs can be seen as a modular learning design of composed functions

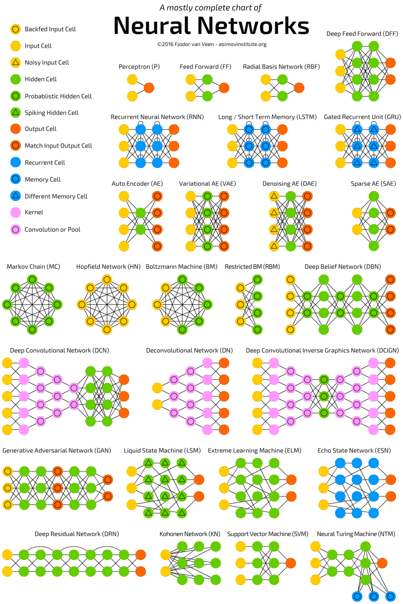

http://www.asimovinstitute.org/neural-network-zoo/

Neural Machine Translation

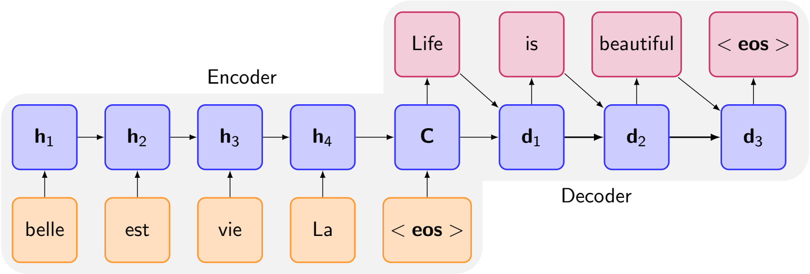

Figure: Example Translation using an Encoder-Decoder Architecture

Talking Robots

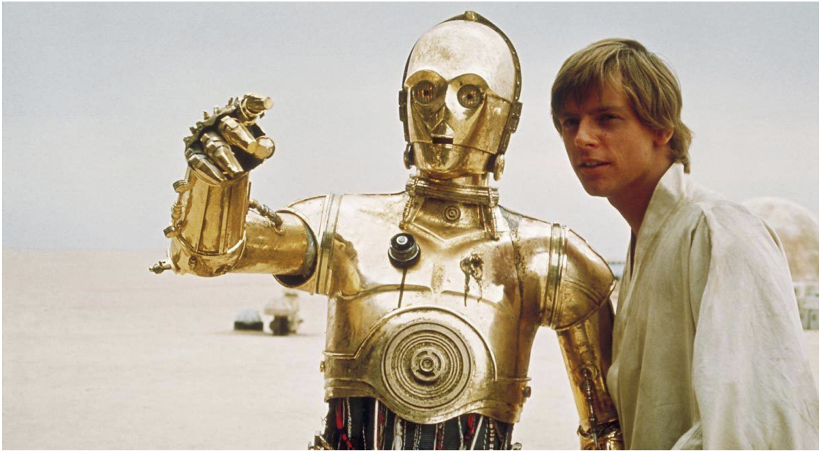

http://www.starwars.com/databank/c-3po

1. Pattern Recognition and Machine Learning Prediction (classification)

   Focus is on identifying features that have high-value states in common - a shared label in a classification setting - across a large, diverse set of training examples

Image Captioning

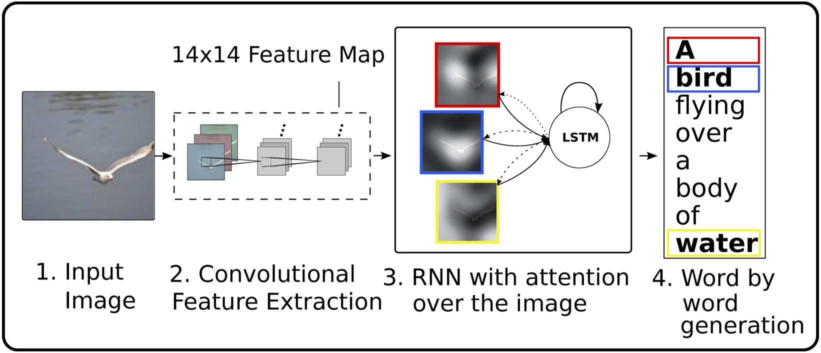

Show, Attend and Tell: Neural Image Caption Generation with Attention. Xu et al., 2015.

Epic Fails

a woman riding a horse on a dirt road | an airplane is parked on the tarmac at an airport | a group of people standing on top of a beach

Building Machines That Learn and Think Like People. Lake, B. et al. Behavioral and Brain Sciences (in press)

1. Pattern Recognition and Machine Learning Prediction (classification)
2. Create Mechanistic Models that are Informed by Domain Theoretic Consider

   Focus is on creating a model architecture that reflects domain theoretic considerations

Understanding Space: Simulation/Intuitive Physics/Imagination

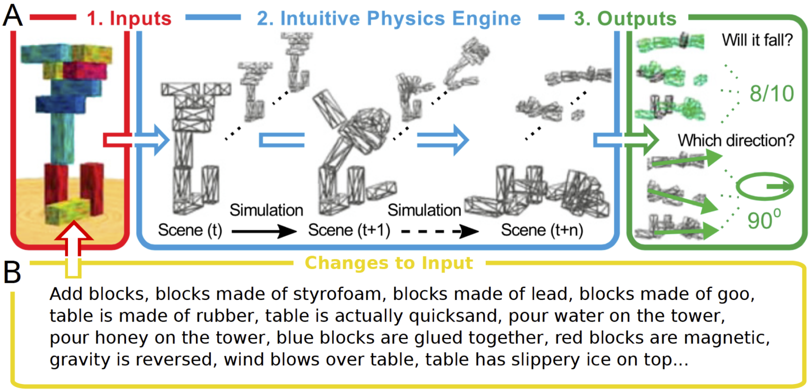

Simulation as an engine of physical scene understanding. Battaglia, P. et al. Proceedings of the National Academy of Sciences, 110(45), 18327-18332

Semiotic Schemas

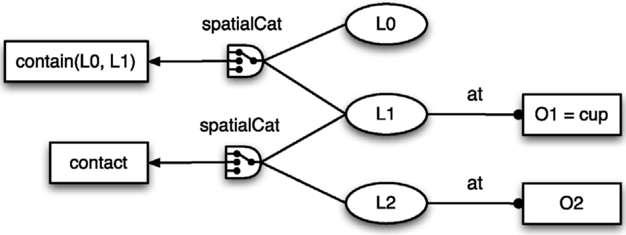

Fig. 13. The situation corresponding to, “There is a cup here. Something is touching the cup.”

Semiotic schemas: A framework for grounding language in action and perception. Roy, D. Artificial Intelligence 167. Pages 170-205. 2005

CLASP centre for linguistic theory and studies in probability

- For language the mechanistic approach can be informed by logic
- Logical theories use *functions* and *compositional* operations while neural networks learn and compose functions
- Logic based domain theory of linguistic performance can inform the structural design of DNNs: model interpretability and performance.

TTR as a design formalism

Perceptual domain

- [[34,24,48],[56,78,114]...]: *PointMap*
  $PointMap \sqsubseteq$ list(list(Real))

Conceptual domain

(Dobnik, 2009; Dobnik and Cooper, 2017)

TTR as a design formalism

Perceptual domain

- [[34,24,48],[56,78,114]. . .]: *PointMap*
  $PointMap \sqsubseteq \text{list}(\text{list}(\text{Real}))$
- Object detection function
  $$(Pointmap \rightarrow \text{set}(\begin{bmatrix} \text{reg} & : & Pointmap \\ \text{pfun} & : & (Ind \rightarrow Type) \end{bmatrix}))$$
  $\text{pfun} = \lambda x{:}Ind.\text{chair}(x)$

Conceptual domain

(Dobnik, 2009; Dobnik and Cooper, 2017)

TTR as a design formalism

Perceptual domain

- [[34,24,48],[56,78,114]...]: *PointMap*
  $PointMap \sqsubseteq$ list(list(Real))
- Object detection function
  $$(Pointmap \rightarrow \text{set}(\left[\begin{array}{lll}\text{reg} & : & Pointmap \\ \text{pfun} & : & (Ind \rightarrow Type)\end{array}\right]))$$
  $$\text{pfun} = \lambda x{:}Ind.\text{chair}(x)$$
- Individuation function
  $$\lambda r{:}\left[\begin{array}{l}\text{reg}{:}Pointmap \\ \text{pfun}{:}(Ind \rightarrow Type)\end{array}\right] . \left[\begin{array}{lll}\text{a} & : & Ind \\ \text{loc} & : & \text{location}(\text{a}, r.\text{reg}) \\ \text{c} & : & r.\text{pfun}(\text{a})\end{array}\right]$$

Conceptual domain

(Dobnik, 2009; Dobnik and Cooper, 2017)

Learning to compose neural networks

Andreas et al. (2016)[1]

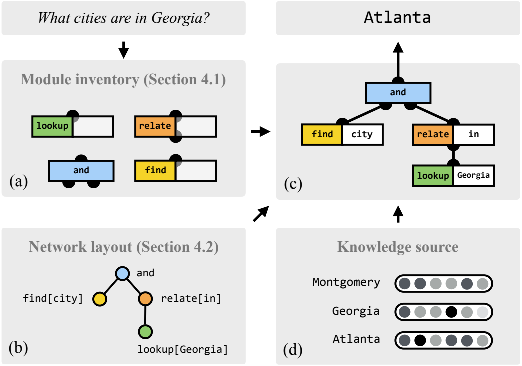

Figure 1: A learned syntactic analysis (a) is used to assemble a collection of neural modules (b) into a deep neural network (c), and applied to a world representation (d) to produce an answer.

[1] Thanks to Mehdi Ghanimifard for this reference.

Learning to compose neural networks, II

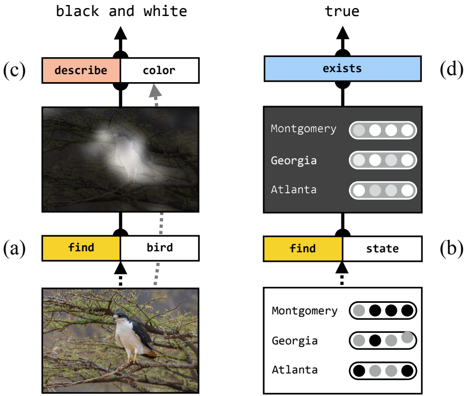

Figure 2: Simple neural module networks, corresponding to the questions What color is the bird? and Are there any states?

CLASP centre for linguistic theory and studies in probability

1. Pattern Recognition and Machine Learning Prediction (classification)
2. Create Mechanistic Models that are Informed by Domain Theoretic Consider
3. Learning strategies beyond presentation of data

Focus is on optimising the learning procedure through interaction

Dialogue games as learning strategies

(Dobnik and de Graaf, 2017)

Dialogue games as learning strategies

Direct instruction: objects

U: This is a cup.

S: [Object is focused on and observed.]

S: OK, I learned a cup.

Dialogue games as learning strategies

Direct instruction: objects

U: This is a cup.

S: [Object is focused on and observed.]

S: OK, I learned a cup.

Direct instruction: spatial relations

U: The book is to the right of the mug.

S: [Objects are focused on and observed.]

S: OK. This relation is called "to the right of".

S: [Objects are focused on and observed.]

U: Where are the objects?

S: The book is to the right of the mug.

Dialogue games as learning strategies

Reinforcement from feedback

U: What is this?

S: [Object is focused on and observed.]

S: This is a cup.

U: That is correct.

S: [Observation is added to the models.]

CLASP centre for linguistic theory and studies in probability

Dialogue games as learning strategies

Reinforcement from feedback

U: What is this?

S: [Object is focused on and observed.]

S: This is a cup.

U: That is correct.

S: [Observation is added to the models.]

Self-initiative: spatial relations

S: [Objects are focused on and observed.]

S: One of the objects is book, the other one is box. Please tell me where things are.

U: The book is to the left of the box.

S: OK, this relation is called “to the left of”.

Object recognition

Results over 4 rounds: direct instruction only

→	apple	banana	bear	book	cap	car	cup	paint	shoe	shoe-box	C-NI
apple	**.343**	.227	.076	.046	.099	.058	.126	.074	.053	.166	.116
banana	.201	**.357**	.058	.035	.085	.087	.148	.066	.046	.124	.155
bear	.080	.121	**.260**	.074	.089	.091	.120	.099	.074	.136	.123
book	.142	.233	.074	**.496**	.114	.197	.246	.130	.085	.220	.250
cap	.122	**.208**	.076	.049	.146	.096	.103	.083	.061	.114	-.062
car	.104	.183	.053	.067	.077	**.414**	.119	.076	.069	.149	.231
cup	.099	.145	.063	.066	.091	.052	**.330**	.094	.054	.120	.185
paint	.119	.140	.075	.076	.083	.147	.121	**.221**	.062	.111	.075
shoe	.078	.123	.070	.056	.079	.116	.124	.076	**.319**	.103	.196
shoe-box	.190	.332	.099	.188	.145	.305	.313	.166	.111	**.376**	.044

CLASP centre for linguistic theory and studies in probability

Object recognition

Results over 4 rounds: direct instruction only

→	apple	banana	bear	book	cap	car	cup	paint	shoe	shoe-box	C-NI
apple	**.343**	.227	.076	.046	.099	.058	.126	.074	.053	.166	.116
banana	.201	**.357**	.058	.035	.085	.087	.148	.066	.046	.124	.155
bear	.080	.121	**.260**	.074	.089	.091	.120	.099	.074	.136	.123
book	.142	.233	.074	**.496**	.114	.197	.246	.130	.085	.220	.250
cap	.122	**.208**	.076	.049	.146	.096	.103	.083	.061	.114	-.062
car	.104	.183	.053	.067	.077	**.414**	.119	.076	.069	.149	.231
cup	.099	.145	.063	.066	.091	.052	**.330**	.094	.054	.120	.185
paint	.119	.140	.075	.076	.083	.147	.121	**.221**	.062	.111	.075
shoe	.078	.123	.070	.056	.079	.116	.124	.076	**.319**	.103	.196
shoe-box	.190	.332	.099	.188	.145	.305	.313	.166	.111	**.376**	.044

C-NI ranking: book > car > shoe > cup > banana > bear > apple > paint > shoe-box > cap

Spatial relations: overall results

16 locations around the landmark tested randomly twice: direct instruction only

Match	Evaluator 1		Evaluator 2		Evaluator 1 + 2	
Independent	8	0.25	7	0.2188	15	0.2344
Secondary	11	0.3438	13	0.4063	24	0.375
Indep. + Second.	19	0.5938	20	0.6251	39	0.6094
Incorrect	13	0.4063	12	0.375	25	0.3906
Total	32	1	32	1	64	1

Spatial relations: confusion matrix

When some contextual parameters are missing

	behind	front	left	right	close	near	Total
behind	**4**	2	1	0	0	2	9
front	0	**5**	3	3	6	0	17
left	0	6	**1**	0	0	0	7
right	4	1	3	**3**	0	1	12
close	1	9	1	0	**1**	2	14
near	1	1	1	0	1	**1**	5
Total	10	24	10	6	8	6	64

Spatial relations: confusion matrix

When some contextual parameters are missing

	behind	front	left	right	close	near	Total
behind	**4**	2	1	0	0	2	9
front	0	**5**	3	3	6	0	17
left	0	6	**1**	0	0	0	7
right	4	1	3	**3**	0	1	12
close	1	9	1	0	**1**	2	14
near	1	1	1	0	1	**1**	5
Total	10	24	10	6	8	6	64

- $A_o = 0.2344, \kappa = 0.0537$
- Appropriate alternatives:
  - topological - projective: “left” as “front”, “close” as “front”
  - FOR variation: “right” as “left”

Conclusions

- Simply applying a powerful learning algorithm to a large dataset (pattern recognition) is problematic if the focus is solely on ML tasks rather than on domain theoretic considerations.

Conclusions

Yoav Goldberg Follow

Senior Lecturer at Bar Ilan University. Working on NLP. Recently with Neural Nets. Published a book ...

Jun 9 · 14 min read

An Adversarial Review of "Adversarial Generation of Natural Language"

Or, for fucks sake, DL people, leave language alone and stop saying you solve it.

[**edit**: some people commented that they don't like the us-vs-them tone and that "deep learning people" can—and some indeed do—do good NLP work. To be clear: I fully agree. #NotAllDeepLearners]

[**update**: I added some clarifications based on responses to this piece. I suggest reading them after reading this one.]

[**update:** Yann LeCun responded on facebook, followed by my response to Yann's]

Conclusions

- However, DNNs have modular architectures that can be specifically tailored or structured to the needs of a specific domain or task
- Introduce domain relevant structural constraints into the model via the network architecture
  - Early examples: (Feldman et al., 1988; Feldman, 1989; Regier, 1996)
  - The example of Johnson et al. in Marco's earlier talk fits within our understanding of this approach.

Conclusions and future research

- Relating and understanding the modular design of DNNs to models of language and cognition provides an interesting research question for the future.
  - Pattern recognition
  - Mechanistic architectures
  - Learning is through interaction

CLASP centre for linguistic theory and studies in probability

References I

Andreas, J., M. Rohrbach, T. Darrell, and D. Klein (2016). Learning to compose neural networks for question answering. *CoRR abs/1601.01705*, 1–10.

Dobnik, S. (2009, September 4). *Teaching mobile robots to use spatial words*. Ph. D. thesis, University of Oxford: Faculty of Linguistics, Philology and Phonetics and The Queen's College, Oxford, United Kingdom. http://www.dobnik.net/simon/documents/thesis.pdf.

Dobnik, S. and R. Cooper (2017, March). Interfacing language, spatial perception and cognition in Type Theory with Records. *Accepted for Journal of Language Modelling n*(n), 1–27.

CLASP centre for linguistic theory and studies in probability

References II

Dobnik, S. and E. de Graaf (2017, 22–24 May). KILLE: a framework for situated agents for learning language through interaction. In J. Tiedemann (Ed.), *Proceedings of the 21st Nordic Conference on Computational Linguistics (NoDaLiDa)*, Volume 131 of *Linköping Electronic Conference Proceedings and NEALT Proceedings Series Vol. 29*, Gothenburg, Sweden, pp. 1–10. Northern European Association for Language Technology (NEALT): Linköping University Electronic Press.

Feldman, J. A. (1989). Structured neural networks in nature and in computer science. In R. Eckmiller and C. v.d. Malsburg (Eds.), *Neural Computers*, pp. 17–21. Berlin, Heidelberg: Springer.

Feldman, J. A., M. A. Fanty, and N. H. Goodard (1988, March). Computing with structured neural networks. *Computer 21*(3), 91–103.

References III

Gazdar, G. (1996). Paradigm merger in natural language processing. In I. Wand and R. Milner (Eds.), *Computing Tomorrow*, pp. 88–109. New York, NY, USA: Cambridge University Press.

Jones, K. I. B. S., G. J. M. Gazdar, and R. M. Needham (2000). Introduction: combining formal theories and statistical data in natural language processing. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 358*(1769), 1227–1238.

Regier, T. (1996). *The human semantic potential: Spatial language and constrained connectionism*. MIT Press.

CLASP centre for linguistic theory and studies in probability