URANIA Workshop
Genova, November 28th, 2016

Reasoning with Deep Learning: an Open Challenge

Marco Lippi

marco.lippi@unimore.it

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

The connectionism vs. symbolism dilemma

A central question in AI

How is knowledge represented in our mind ?

Symbolic approaches

- Reasoning as the result of formal manipulation of symbols

Connectionist (sub-symbolic) approaches

- Reasoning as the result of processing of interconnected (networks of) simple units

Connectionism vs. symbolism approaches

Symbolic approaches

- founded on the **principles of logic**
- highly **interpretable**

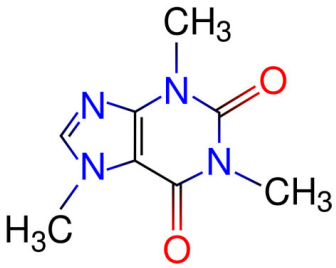

toxic(m) :- doublebond(m,c1,c2), hydroxyl(c2), methyl(m)

Connectionist approaches

- can more easily deal with **uncertain knowledge**
- can be easily **distributed**
- often seen as "black box" $\rightarrow$ dark magic ☺

Deep learning

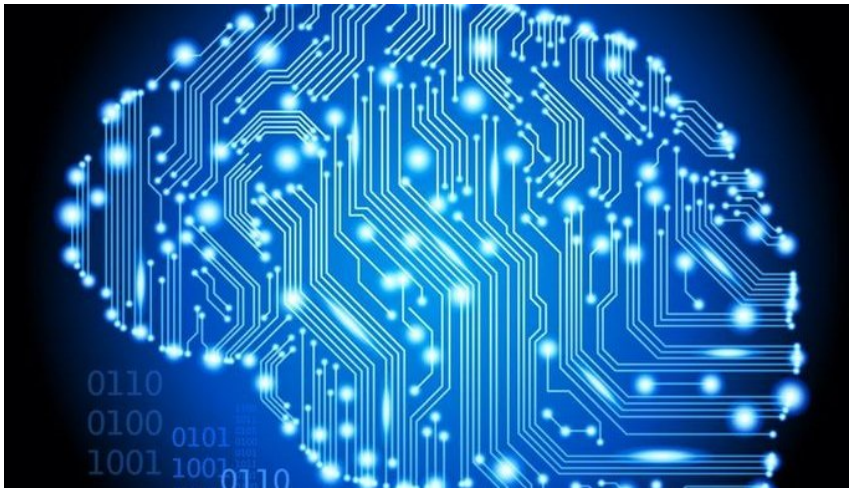

Deep learning has brought (back ?) a **revolution** into AI

- exploit **more computational power**
- refine **optimization** methods (dropout, rectification, ...)
- automatically **learn feature** hierarchies
- exploit **unsupervised** data (though not yet enough)

Deep learning

Breakthough in a variety of application fields

- Speech recognition
- Computer vision
- Natural language processing
- ...

Is this the **solution** to all AI problems ? Probably not but...

- for certain types of task **it is hard to compete**
- **big companies** are currently playing a major role
- huge space for applications **upon deep learning systems**

What is missing ?

Pioneering approaches

Knowledge-based artificial neural networks (KBANNs)

- [Towell & Shavlik, 1994]
- One of the first attempts to inject knowledge into ANNs
- Trying to interpret an ANN model as logic rules

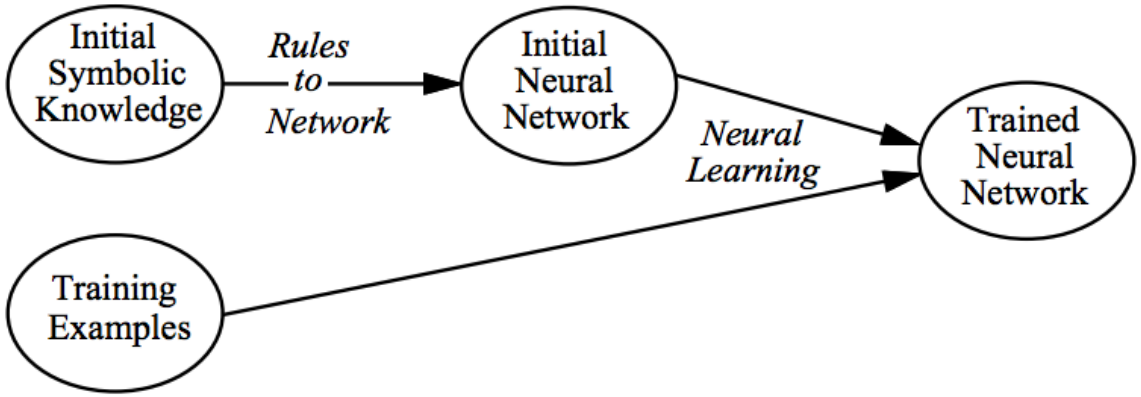

Knowledge Base		*Neural Network*
Final Conclusions	$\Longleftrightarrow$	Output Units
Supporting Facts	$\Longleftrightarrow$	Input Units
Intermediate Conclusions	$\Longleftrightarrow$	Hidden Units
Dependencies	$\Longleftrightarrow$	Weighted Connections

Pioneering approaches

Knowledge-based artificial neural neyworks (KBANNs) [1994]

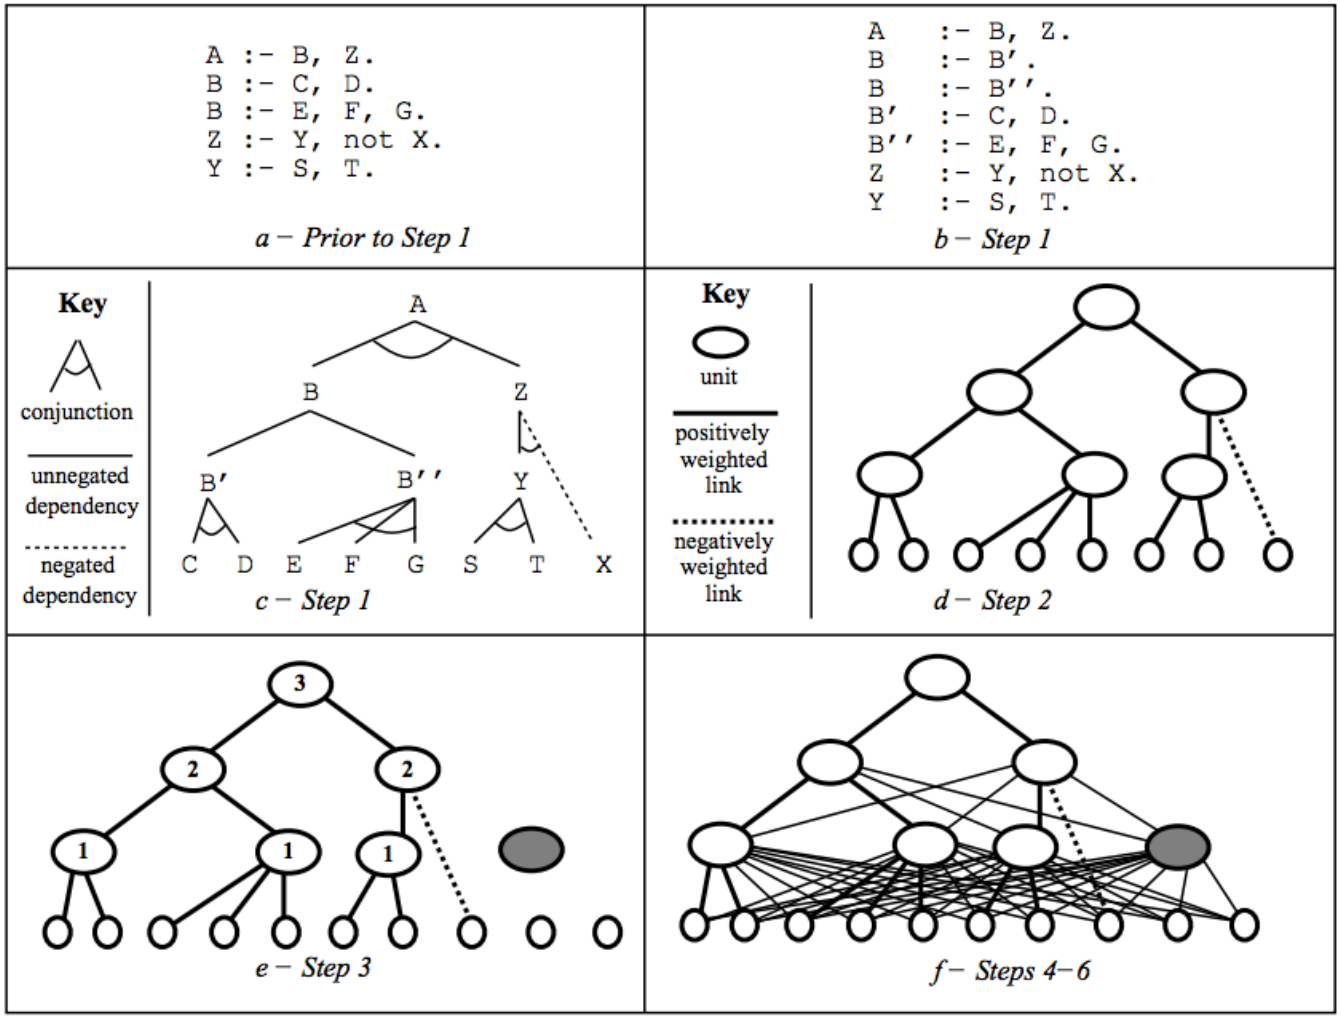

NeSy and SRL

More recent research directions:

- Neural-Symbolic Learning (NeSy)
- Statistical Relational Learning (SRL)

$\rightarrow$ developed during the 90s-00s
$\rightarrow$ combining logic with cognitive neuroscience (NeSy)
$\rightarrow$ combining logic with probabilistic/statistical learning (SRL)

Example – Markov logic

A probabilistic-logic framework to model knowledge

2.3 `LikedMovie(x,m)` $\wedge$ `Friends(x,y)` => `LikedMovie(y,m)`

1.6 `Friends(x,y)` $\wedge$ `Friends(y,z)` => `Friends(y,z)`

Extension [Lippi & Frasconi, 2009] $\rightarrow$ learn weights with ANNs

Memory Networks (MemNNs) @ Facebook

General model described in terms of four component networks:

1. Input feature map (I)
   $\rightarrow$ convert input into an internal feature space
2. Generalization (G)
   $\rightarrow$ update memories given new input
3. Output (O)
   $\rightarrow$ produce new output (in feature space) given memories
4. Response (R)
   $\rightarrow$ convert output into a response seen by the outside world

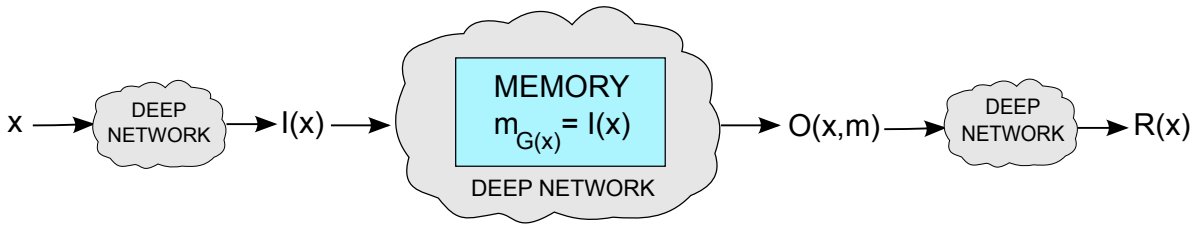

Example: a (simple ?) reasoning task

Joe went to the kitchen. Fred went to the kitchen. Joe picked up the milk. Joe travelled to the office. Joe left the milk. Joe went to the bathroom.

Where is the milk now? A: **office**
Where is Joe? A: **bathroom**
Where was Joe before the office? A: **kitchen**

Memory Networks (MemNNs)

A very simple implementation

1. Convert sentence x into a feature vector $I(x)$ (e.g., BoW)
2. Store $I(x)$ into an empty slot of memory: $m_{G(x)} = I(x)$
3. When given query q, find k supporting memories given q:
   $o_1 = O_1(q, \mathbf{m}) = \text{argmax}_i \, s_O(q, m_i)$
   $o_2 = O_2(q, \mathbf{m}) = \text{argmax}_i \, s_O([q, m_{o_1}], m_i)$
4. Formulate a single-word response r given vocabulary W:
   $r = \text{argmax}_{w \in W} \, s_R([q, m_{o_1}, m_{o_2}], w)$

Scoring functions s_O, s_R are implemented as **deep networks**
$\rightarrow$ Need some form of **supervision**

Benchmarking

- The (20) bAbI tasks
- The Children's Book Test
- The Movie Dialog dataset
- The SimpleQuestions dataset

bAbI tasks (Facebook)

Task 1: Single Supporting Fact

Mary went to the bathroom.
John moved to the hallway.
Mary travelled to the office.
Where is Mary? A:office

Task 2: Two Supporting Facts

John is in the playground.
John picked up the football.
Bob went to the kitchen.
Where is the football? A:playground

Task 3: Three Supporting Facts

John picked up the apple.
John went to the office.
John went to the kitchen.
John dropped the apple.
Where was the apple before the kitchen? A:office

Task 4: Two Argument Relations

The office is north of the bedroom.
The bedroom is north of the bathroom.
The kitchen is west of the garden.
What is north of the bedroom? A: office
What is the bedroom north of? A: bathroom

Task 5: Three Argument Relations

Mary gave the cake to Fred.
Fred gave the cake to Bill.
Jeff was given the milk by Bill.
Who gave the cake to Fred? A: Mary
Who did Fred give the cake to? A: Bill

Task 6: Yes/No Questions

John moved to the playground.
Daniel went to the bathroom.
John went back to the hallway.
Is John in the playground? A:no
Is Daniel in the bathroom? A:yes

Task 7: Counting

Daniel picked up the football.
Daniel dropped the football.
Daniel got the milk.
Daniel took the apple.
How many objects is Daniel holding? A: two

Task 8: Lists/Sets

Daniel picks up the football.
Daniel drops the newspaper.
Daniel picks up the milk.
John took the apple.
What is Daniel holding? milk, football

Task 9: Simple Negation

Sandra travelled to the office.
Fred is no longer in the office.
Is Fred in the office? A:no
Is Sandra in the office? A:yes

Task 10: Indefinite Knowledge

John is either in the classroom or the playground.
Sandra is in the garden.
Is John in the classroom? A:maybe
Is John in the office? A:no

[Table by Weston et al.]

bAbI tasks (Facebook)

Task 11: Basic Coreference
Daniel was in the kitchen.
Then he went to the studio.
Sandra was in the office.
Where is Daniel? A:studio

Task 12: Conjunction
Mary and Jeff went to the kitchen.
Then Jeff went to the park.
Where is Mary? A: kitchen
Where is Jeff? A: park

Task 13: Compound Coreference
Daniel and Sandra journeyed to the office.
Then they went to the garden.
Sandra and John travelled to the kitchen.
After that they moved to the hallway.
Where is Daniel? A: garden

Task 14: Time Reasoning
In the afternoon Julie went to the park.
Yesterday Julie was at school.
Julie went to the cinema this evening.
Where did Julie go after the park? A:cinema
Where was Julie before the park? A:school

Task 15: Basic Deduction
Sheep are afraid of wolves.
Cats are afraid of dogs.
Mice are afraid of cats.
Gertrude is a sheep.
What is Gertrude afraid of? A:wolves

Task 16: Basic Induction
Lily is a swan.
Lily is white.
Bernhard is green.
Greg is a swan.
What color is Greg? A:white

Task 17: Positional Reasoning
The triangle is to the right of the blue square.
The red square is on top of the blue square.
The red sphere is to the right of the blue square.
Is the red sphere to the right of the blue square? A:yes
Is the red square to the left of the triangle? A:yes

Task 18: Size Reasoning
The football fits in the suitcase.
The suitcase fits in the cupboard.
The box is smaller than the football.
Will the box fit in the suitcase? A:yes
Will the cupboard fit in the box? A:no

Task 19: Path Finding
The kitchen is north of the hallway.
The bathroom is west of the bedroom.
The den is east of the hallway.
The office is south of the bedroom.
How do you go from den to kitchen? A: west, north
How do you go from office to bathroom? A: north, west

Task 20: Agent's Motivations
John is hungry.
John goes to the kitchen.
John grabbed the apple there.
Daniel is hungry.
Where does Daniel go? A:kitchen
Why did John go to the kitchen? A:hungry

[Table by Weston et al.]

Children's Book Test

S: 1 He thought that Old Mr. Toad was trying to fool him .
2 Presently Peter Rabbit came along .
3 He found Jimmy Skunk sitting in a brown study .
4 He had quite forgotten to look for fat beetles , and when he forgets to do that you may make up your mind that Jimmy is doing some hard thinking .
5 `` Hello , old Striped-coat , what have you got on your mind this fine morning ? ''
6 cried Peter Rabbit .
7 `` Him , '' said Jimmy simply , pointing down the Lone Little Path .
8 Peter looked .
9 `` Do you mean Old Mr. Toad ! ''
10 he asked .
11 Jimmy nodded .
12 `` Do you see anything queer about him ? ''
13 he asked in his turn .
14 `` Do you see anything queer about him ? ''
15 he asked .
16 Peter stared down the Lone Little Path .
17 `` No , '' he replied , `` except that he seems in a great hurry . ''
18 `` That 's just it , '' Jimmy returned promptly .
19 `` Did you ever see him hurry unless **he was frightened ? ''**
20 Peter confessed that he never had

q: `` Well , he is n't _____ now , yet just look at him go '' retorted Jimmy .

C: Do, came, confessed, frightened, mean, replied, returned, said, see, thought.

MemNNs (window +self-sup.): **frightened**

[Table by Hill et al., 2016]

Movie Dialog dataset

Task 1: Factoid Question Answering (QA)

What movies are about open source? Revolution OS
Ruggero Raimondi appears in which movies? Carmen
What movies did Darren McGavin star in? Billy Madison, The Night Stalker, Mrs. Pollifax-Spy
Can you name a film directed by Stuart Ortiz? Grave Encounters
Who directed the film White Elephant? Pablo Trapero
What is the genre of the film Dial M for Murder? Thriller, Crime
What language is Whity in? German

Task 2: Recommendation

Schindler's List, The Fugitive, Apocalypse Now, Pulp Fiction, and The Godfather are films I really liked.
Can you suggest a film? The Hunt for Red October

Some movies I like are Heat, Kids, Fight Club, Shaun of the Dead, The Avengers, Skyfall, and Jurassic Park.
Can you suggest something else I might like? Ocean's Eleven

Task 3: QA + Recommendation Dialog

I loved Billy Madison, My Neighbor Totoro, Blades of Glory, Bio-Dome, Clue, and Happy Gilmore.
I'm looking for a Music movie. School of Rock
What else is that about? Music, Musical, Jack Black, school, teacher, Richard Linklater, rock, guitar
I like rock and roll movies more. Do you know anything else? Little Richard

Tombstone, Legends of the Fall, Braveheart, The Net, Outbreak, and French Kiss are films I really liked.
I'm looking for a Fantasy movie. Jumanji
Who directed that? Joe Johnston
I like Tim Burton movies more. Do you know anything else? Big Fish

[Table by Dodge et al., 2016]

SimpleQuestions dataset

What American cartoonist is the creator of Andy Lippincott?	(andy_lippincott, character_created_by, garry_trudeau)
Which forest is Fires Creek in?	(fires_creek, containedby, nantahala_national_forest)
What is an active ingredient in childrens earache relief ?	(childrens_earache_relief, active_ingredients, capsicum)
What does Jimmy Neutron do?	(jimmy_neutron, fictional_character_occupation, inventor)
What dietary restriction is incompatible with kimchi?	(kimchi, incompatible_with_dietary_restrictions, veganism)

Table 1: **Examples of simple QA**. Questions and corresponding facts have been extracted from the new dataset SimpleQuestions introduced in this paper. Actual answers are underlined.

[Table by Bordes et al., 2015]

Neural Conversational Model (Google)

Human: *why are we here ?*
Machine: *i 'm not sure .*
Human: *okay, bye !*
Machine: *bye .*
Human: *see you later .*
Machine: *bye .*

Conversation 2: Simple Q&A (contexts and multiple choice)

Human: *my name is david . what is my name ?*
Machine: *david .*
Human: *my name is john . what is my name ?*
Machine: *john .*
Human: *are you a leader or a follower ?*
Machine: *i 'm a leader .*
Human: *are you a follower or a leader ?*
Machine: *i 'm a leader .*

Conversation 4: Philosophical Q&A

Human: *what is the purpose of life ?*
Machine: *to serve the greater good .*
Human: *what is the purpose of living ?*
Machine: *to live forever .*
Human: *what is the purpose of existence ?*
Machine: *to find out what happens when we get to the planet earth .*
Human: *where are you now ?*
Machine: *i 'm in the middle of nowhere .*
Human: *what is the purpose of dying ?*
Machine: *to have a life .*
Human: *what is the purpose of being intelligent ?*
Machine: *to find out what it is .*
Human: *what is the purpose of emotions ?*
Machine: *i don 't know .*

[Table by Vinyalis & Le, 2015]

Open challenges

- Connectionist models for reasoning
  - Process input and store the information in some **memory**
  - Understand pieces of knowledge **relevant** to a given question
  - Formulate some **hypothesis**
  - Provide the correct **answer**
  - Completely different from existing sophisticated **question answering systems**

- Big data
  - A reason of the impressive **success** of deep learning
  - Availability of **huge datasets**
  - Various and **heterogeneous** data sources over the Web
  - Advancements in computer hardware **performance**
  - **Injection of background knowledge** network structures ?

Open challenges

- Unsupervised learning
  - Automatically **extract knowledge** from data
  - **Encode** it into a neural network model
  - **Integrate** expert-given knowledge
  - A proper use of unsupervised data is **still missing** in deep learning [LeCun et al. 2015].

- Incremental learning
  - Humans naturally implement a **lifelong learning** scheme
  - **Continuously** acquire, process and store knowledge
  - A **crucial element** for the development of reasoning skills
  - **Dynamically change** the neural network topology ?

Beyond the Turing test ?

Design **reasoning tasks** for a new version of the Turing test

$\Longrightarrow$ e.g., **Visual Turing Challenge** [Geman et al. 2014]

1. Q: Is there a person in the blue region? **A: yes**
2. Q: Is there a unique person in the blue region? **A: yes**
(Label this person 1)
3. Q: Is person 1 carrying something? **A: yes**
4. Q: Is person 1 female? **A: yes**
5. Q: Is person 1 walking on a sidewalk? **A: yes**
6. Q: Is person 1 interacting with any other object? **A: no**
⋮