The Road to Artificial Intelligence: From Sensors to Concepts to Analogies

Marc Pickett

with David W. Aha
Intelligent Decision Aids Group (Code 5514)
Navy Center for Applied Research in Artificial Intelligence
Naval Research Laboratory, Washington DC

Problem setup

Big problem for AI: **Knowledge Acquisition Bottleneck**

- "Difficult" problems in AI easy for 5-year-old.
  - Language, object recognition, climbing stairs, story understanding.
  - These rely on "common sense" knowledge.
- Where does knowledge come from?
  - Can hand-code knowledge. (Expensive and brittle!)
  - Can *learn* knowledge.

Q1: How can machines learn knowledge from uninterpreted data?

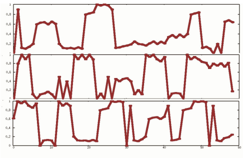

Uninterpreted Sensors

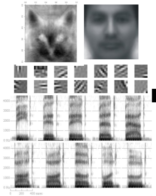

Sensor Concepts

"If A loves B, and B loves C, then A may envy C."

"Continuing civil unrest may soon cause collapse of the Egyptian state." (The Guardian, 1/29/2013)

Relational Knowledge

Problem setup: Uninterpreted data?

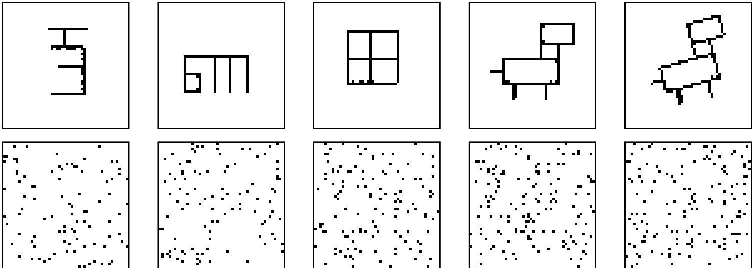

Computer sees this:

{ S0002=1, S0017=1, S0048=1, S0055=0, S0056=1, S0117=1,
S0175=1, S0180=0, S0197=1, S0233=0, S0269=1, S0284=1,
S2214=1, S2218=1, S2227=1, S2247=1, S2258=0, S2308=1,
S2325=1, S2344=1, S2355=1, S2363=1, S2398=1, S2406=1,
S0000=0, S0001=0, S0003=0, etc.... (2,473 more sensors) }

Why uninterpreted?

We know some vision, but what about new senses (*modalities*)?
Seismic, Infrared, Sonar, Traffic Speed Sensors, etc.
(In Part 2 of this talk, topology *must* be learned.)

Problem setup: Invariance

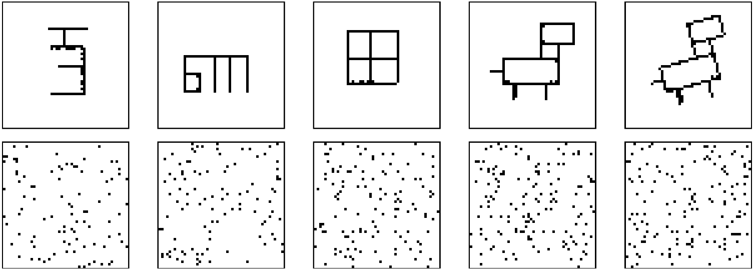

Invariance

- "Same" shape though nearly disjoint pixels.
- Scale, Translation, Shading, etc.

Problem addendum: Invariance

Q1: How can machines learn knowledge from uninterpreted data?

Q2: How can machines learn *invariant* concepts from uninterp. data?

Given lots of images, learn **visual objects**
invariant to **translation**, **rotation**, **scale**, **shading**, etc.

Given large speech stream, learn **phonemes** and **words**
invariant to **pitch**, **cadence**, **accent**, etc.

Other modalities? Seismic, Sonar

A core algorithm for concept learning?

Flattened Macaque Brain

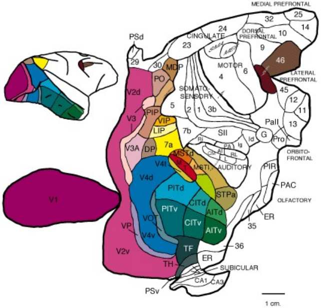

Mountcastle, 1978:

"Cortex (memory & cognition) is 2D sheet that looks the same everywhere."

Assumptions:

1. Memory is **one** algorithm repeated many times ("cortical column").
2. This algo is basic engine of intelligence.
3. **Modality independence** is possible and important.
   Audio, Video, Image, Sonar, etc. (This matters for higher cognition.)
4. There is objective infomation-theoretic metric for concept utility.

Not constrained by neuroscience, but nice to have existence proof.

Existence Proof: Newborn Ferret Cortex

Newborn ferret *auditory* cortex can learn vision! (Sur & Rubenstein, 2005)

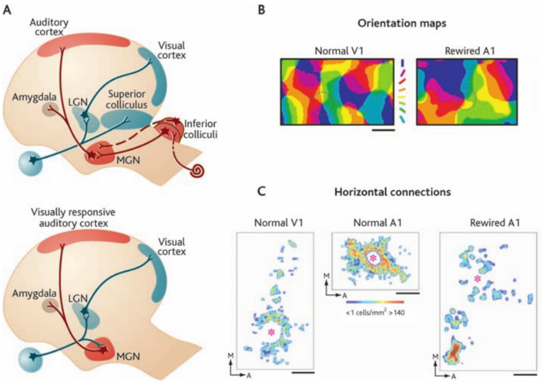

Difference between brains of Mice and Men:

- Both have same basic architecture
- Humans just have (significantly) larger cortex
- So, higher cognition uses same algo as perception? **How?**

Problem, rephrased

Q1: How can machines learn knowledge from uninterpreted data?

Q2: How can machines learn invariant concepts from uninterp. data?

Q3: How can we implement a "cortex"?

- Given vision, learn objects...
- Given speech, learn phonemes and words...
- Given seismic or sonar learn ??...
- ...all with relevant invariances.

Cortex: Invariant Representation

Invariant recognition **HMax** (Riesenhuber and Poggio, 1999):

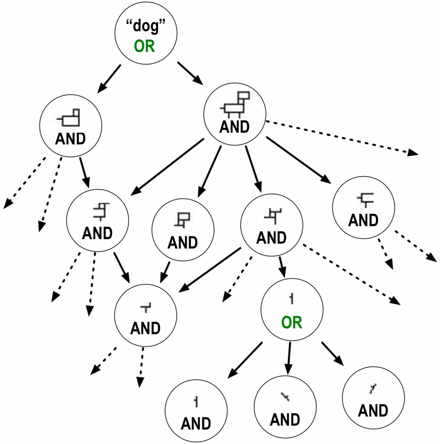

Uses layers of "AND"s and "OR"s.

Cortex: Deep Learning (Le et al., 2012)

Unsupervised learning of invariant concepts in images

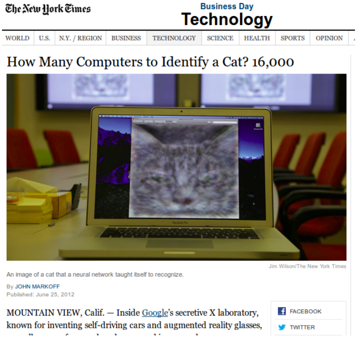

The New York Times

Business Day

Technology

WORLD | U.S. | N.Y. / REGION | BUSINESS | TECHNOLOGY | SCIENCE | HEALTH | SPORTS | OPINION

How Many Computers to Identify a Cat? 16,000

Jim Wilson/The New York Times

An image of a cat that a neural network taught itself to recognize.

By JOHN MARKOFF

Published: June 25, 2012

MOUNTAIN VIEW, Calif. — Inside Google's secretive X laboratory, known for inventing self-driving cars and augmented reality glasses,

FACEBOOK

TWITTER

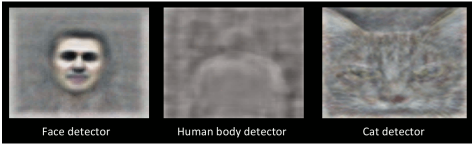

Cortex: Deep Learning (Le et al., 2012)

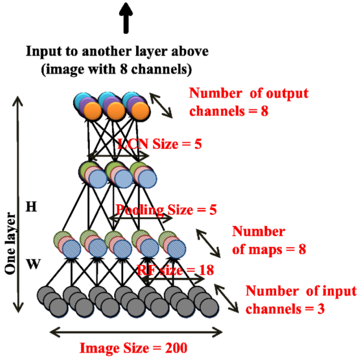

- 9 layers of nodes: 3 x (AND, OR, Normalization)
- Search in weight space to minimize:
  Reconstruction-Cost (compression) + Sparsity-Constraint ("OR" units off)
- 1 billion tunable parameters
- 10,000,000 200x200 images from YouTube (1 image per video)

Cortex: Deep Learning (Le et al., 2012)

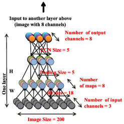

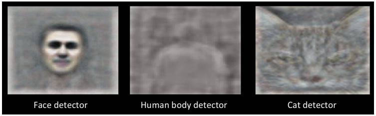

Great work...

- Learns invariant concepts without being told explicitly. (Unlike hand-built invariant features (e.g., SIFT or HOG))
- Outperforms SIFT & HOG: 9.5% $\rightarrow$ 15.8% on ImageNet
- Why? Learns other invariances (shading, *out of plane* rotation)

...but room for improvement

- Doesn't work off the shelf, like Ferret
  - E.g., Action Recognition in Videos (needs new structures)
  - Implicitly told system about vision in structure of network
  - Needs new structure for learning speech data
- Would be nice if structure could be *learned* in addition to parameters.

Learning Features

"Chunker" Algorithm: (Pickett, 2011)

- Input uninterpreted data (e.g., raw sensors)
- Learn hierarchy of features
- Use info. theory constraints (probabilistic compression)
- Iteratively **chunks** (for ANDs) and **merges** (for ORs)

Animals

(Toy example from U.C. Irvine ML Repository)

Name	**hair**	**feathers**	**eggs**	⋯
aardvark	1	0	0	⋯
antelope	1	0	0	⋯
bass	0	0	1	⋯
bear	1	0	0	⋯
		⋱		
worm	0	0	1	⋯
wren	0	1	1	⋯

$\Rightarrow$

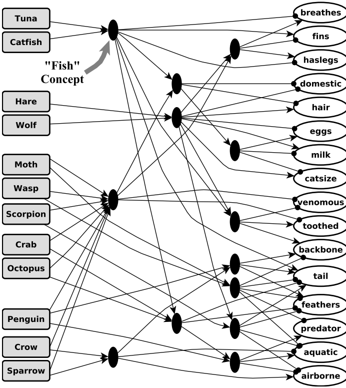

Chunking Patches from Photos

Input: 50x50 Image Patches (uninterpreted)

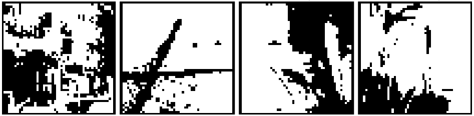

Some Concepts Learned

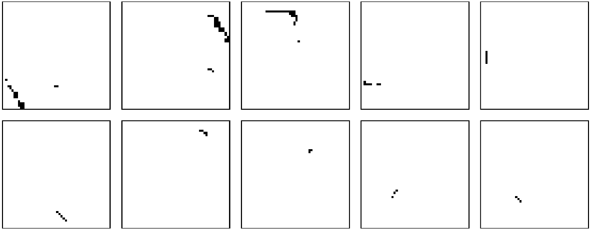

Concepts learned from chunking not only compress data, but also:

- Are useful for semi-supervised learning. (Pickett, 2011)
- Can produce useful macro-actions in Reinforcement Learning (Pickett & Barto 2002).

Learning ORs: Merging

- Merging finds “OR”s or Equivalence Classes
- Interchangeable concepts form OR
- Find ORs by chunking context

A	$P(ran\|A)$	$P(sat\|A)$	$P(dog\|A)$	$P(cat\|A)$	$P(meow\|A)$
ran	1.00	0.02	0.23	0.32	0.01
sat	0.01	1.00	0.21	0.34	0.03
dog	0.45	0.37	1.00	0.03	0.02
cat	0.41	0.39	0.04	1.00	0.30
$dog \vee cat$	0.43	0.38			

To date: Recovers small artificial grammars (Pickett, 2011).
Current work: scaling this up to images!

Now what?

If we have cortex, what to do with it?

Problem 4: Relational Data

Q1: How can machines learn knowledge from uninterpreted data?

Q2: How can machines learn invariant concepts from uninterp. data?

Q3: How can we implement a "cortex"?

Q4: How can we leverage a "cortex" to process relational data?

- Given stories, learn cliches and plot devices...
- Given descriptions of battles, learn tactics (e.g., "counter-attack", "defense", "attrition")

Relational Structures as "Sensors"

Q4: How can we leverage a "cortex" to process relational data?

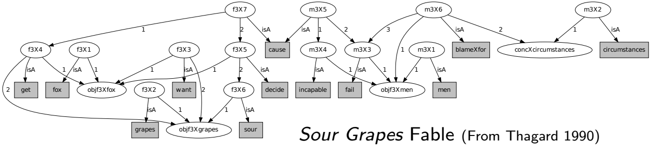

Sour Grapes Fable (From Thagard 1990)

Relational Structures as "Sensors"

Q4: How can we leverage a "cortex" to process relational data?

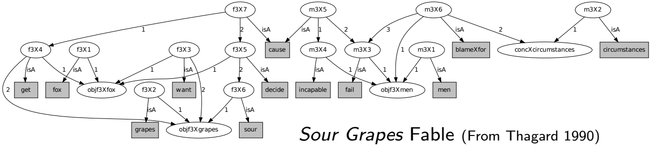

Sour Grapes Fable (From Thagard 1990)

Transform relational structures into input data for "cortex".

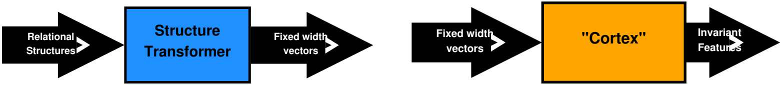

Structure Transformer: I

English (for clarity)

"A fox wanted some grapes, but could not get them. This caused him to decide that the grapes were sour, though the grapes weren't. Likewise, men often blame their failures on their circumstances, when the real reason is that they are incapable."

Predicate Form (actual input)

```
fox Of3Fox            cause m34 m33              sameAs f36 (sour Of3Grapes)
false f36             grapes Of3Grapes           sameAs f35 (decide Of3Fox f36)
cause f34 f35         incapable Of3Men           sameAs f34 (get Of3Fox Of3Grapes)
false f34             decide Of3Fox f36          sameAs m34 (incapable Of3Men)
men Of3Men            sameAs m33 (fail Of3Men)   blameFor Of3Men concCircum m33
fail Of3Men           want Of3Fox Of3Grapes      circumstances concCircum
```

Transforming a Window

```
blameFor Of3Men concCircum m33
sameAs m33 (fail Of3Men)
fail Of3Men
circumstances concCircum
men Of3Men
incapable Of3Men
```

⟹

```
blameFor1=blameFor3.fail1
circumstances1=blameFor2
fail1=blameFor3.fail1
fail1=blameFor1
incapable1=blameFor3.fail1
incapable1=blameFor1
incapable1=fail1
men1=blameFor3.fail1
men1=blameFor1
men1=fail1
men1=incapable1
```

Structure Transformer: II

Many Transformed Windows

```
blameFor1=blameFor3.fail1
circumstances1=blameFor2
fail1=blameFor3.fail1
fail1=blameFor1
incapable1=blameFor3.fail1
incapable1=blameFor1
incapable1=fail1
men1=blameFor3.fail1
men1=blameFor1
men1=fail1
men1=incapable1
```

```
cause2.fail1=blameFor3.fail1
blameFor1=blameFor3.fail1
blameFor1=cause2.fail1
cause2=blameFor3
fail1=blameFor3.fail1
fail1=cause2.fail1
fail1=blameFor1
men1=blameFor3.fail1
men1=cause2.fail1
men1=blameFor1
men1=fail1
```

```
false1.sour1=decide2.sour1
decide1=cause2.decide1
decide2=cause2.decide2
false1=cause2.decide2
false1=decide2
```

⋮

```
blameFor1=blameFor3.fail1
fail1=blameFor3.fail1
fail1=blameFor1
incapable1=blameFor3.fail1
incapable1=blameFor1
incapable1=fail1
men1=blameFor3.fail1
men1=blameFor1
men1=fail1
men1=incapable1
```

⋮

Algorithm:

- Treat transformed windows as percepts: Chunk and Merge
- Treat bags of chunked and merged windows as inputs
  - Chunk and Merge these.

Concepts Learned from Stories (Pickett & Aha 2013)

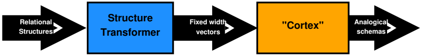

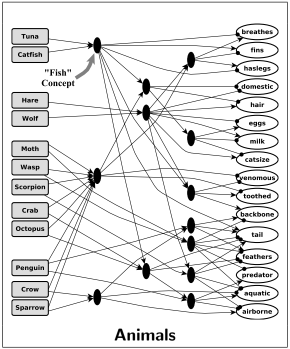

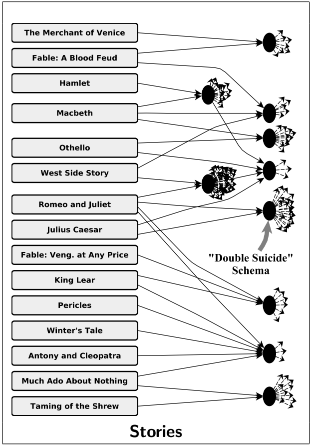

Solves Problem of Spontaneous Analog Retrieval!

(submitted to CogSci 2013)

	Analogs Found	Average Time (Comparisons)
Rachkovskij et al. 2012	100.00%	100.00
Ours	95.45%	15.43

Review

Q1: How can machines learn knowledge from uninterpreted data?

Q2: How can machines learn invariant concepts from uninterp. data?

Q3: How can we implement a "cortex"?

Answer: By (step 1) chunking and merging!

Q4: How can we leverage a "cortex" to process relational data?

Answer: By (step 2) transforming rel. structs. into fixed-width vectors,

and (step 3) feeding xformed structs. into cortex.

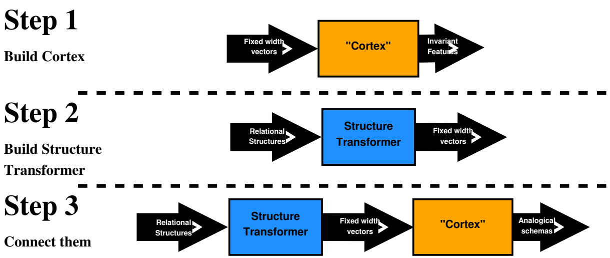

Current and Future Work

Q1: How can machines learn knowledge from uninterpreted data?

Q2: How can machines learn invariant concepts from uninterp. data?

Q3: How can we implement a "cortex"?

Q4: How can we leverage a "cortex" to process relational data?

Currently:

- Scaling Merging up to images.
- Using knowledge for **prediction** and **planning**.
  - E.g., predict story or battle outcomes.
    "Hero will defuse bomb at last second."
  - Realtime battle simulation player that utilizes concepts
    "counter-attack", "defense", "attrition"

Future:

Q5: How are relational structures learned to begin with?

Workshop: RepLearn 2013

Learning Rich Representations from Low-Level Sensors

- In conjunction with AAAI 2013 (7/15) in Bellevue, WA
- Submission Deadline: April 3rd, 2013
- **Organizing Committee:**
  - Marc Pickett (Main contact)
  - **Ben Kuipers**, *Developmental Robotics*
  - **Yann LeCun**, *Deep Learning*
  - **Clayton Morrison**, *Developmental Robotics*
- **Invited Speakers:**
  - **Yoshua Bengio**, *Deep Learning*
  - **Jeff Hawkins**, *Neuroscience and AI (Inventor of Palm Pilot)*
  - **Juergen Schmidhuber**, *Artificial General Intelligence*
- See http://marcpickett.com/RepLearn2013 for details.

Full Story Ontology