Vérité, utilité, intractabilité

Guy Emerson

10 Years of My Research…

If semantics is truth-conditional,
shouldn’t truth conditions be learnable from data?

10 Years of My Research…

If semantics is truth-conditional,
shouldn’t truth conditions be learnable from data?

Yes, but it’s hard.

10 Years of My Research…

If semantics is truth-conditional,
shouldn't truth conditions be learnable from data?

- In high dimensions, truth conditions are useful
- In high dimensions, truth conditions are painful

10 Years of My Research…

If semantics is truth-conditional,
shouldn't truth conditions be learnable from data?

- In high dimensions, truth conditions are useful (guide for generalisation)
- In high dimensions, truth conditions are painful

10 Years of My Research…

If semantics is truth-conditional,
shouldn’t truth conditions be learnable from data?

- In high dimensions, truth conditions are useful (guide for generalisation)
- In high dimensions, truth conditions are painful (pragmatic inferences are intractable)

Modelling Language and the World

- Vision-language models are generally trained on images and captions

Modelling Language and the World

- Vision-language models are generally trained on images and captions
  - observed caption $\Rightarrow$ true caption
  - unobserved caption $\not\Rightarrow$ false caption

Modelling Language and the World

- Vision-language models are generally trained on images and captions
  - observed caption $\Rightarrow$ true caption
  - unobserved caption $\not\Rightarrow$ false caption
- Strong priors: “the cup is __ the table”

Visual Spatial Reasoning (Liu et al., 2023)

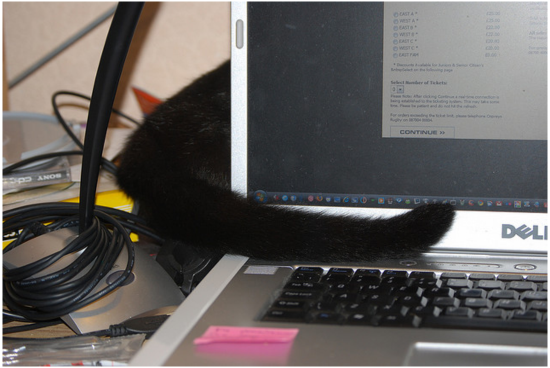

"The cat is behind the laptop"

Visual Spatial Reasoning (Liu et al., 2023)

- Given an image and a caption, is the caption true?
  - Caption-only baseline: 50%

Visual Spatial Reasoning (Liu et al., 2023)

- Given an image and a caption, is the caption true?
  - Caption-only baseline: 50%
  - Human performance: $> 95\%$
  - State of the art: $\sim 70\%$

Visual Spatial Reasoning (Liu et al., 2023)

- Given an image and a caption, is the caption true?
  - Caption-only baseline: 50%
  - Human performance: $> 95\%$
  - State of the art: $\sim 70\%$
- Humans understand truth, large NLP models do not

What I'll Cover...

- Truth conditions at scale
  - Learning from text, images, ontologies
  - Usefulness for generalisation

What I'll Cover...

- Truth conditions at scale
  - Learning from text, images, ontologies
  - Usefulness for generalisation
- Beyond truth conditions
  - Intractability of inference
  - A new kind of probabilistic model

Situation Semantics

$$x \xleftarrow{\text{ARG1}} y \xrightarrow{\text{ARG2}} z$$

cat(x) behind(y) laptop(z)

Situation Semantics

$$x \xleftarrow{\text{ARG1}} y \xrightarrow{\text{ARG2}} z$$

cat(x)	behind(y)	laptop(z)
animal(x)	next_to(y)	computer(z)

Situation Semantics

$$x \xleftarrow{\text{ARG1}} y \xrightarrow{\text{ARG2}} z$$

cat(x)	behind(y)	laptop(z)
animal(x)	next_to(y)	computer(z)
laptop(x)	on(y)	cat(z)
computer(x)	cat(y)	animal(z)
behind(x)	laptop(y)	behind(z)
...	...	...

Probabilistic Situation Semantics

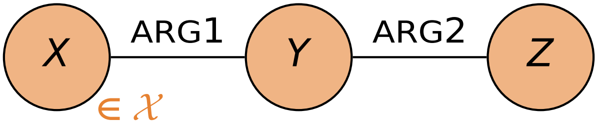

cat(X)	behind(Y)	laptop(Z)
animal(X)	next_to(Y)	computer(Z)
laptop(X)	on(Y)	cat(Z)
computer(X)	cat(Y)	animal(Z)
behind(X)	laptop(Y)	behind(Z)
...	...	...

Probabilistic Situation Semantics

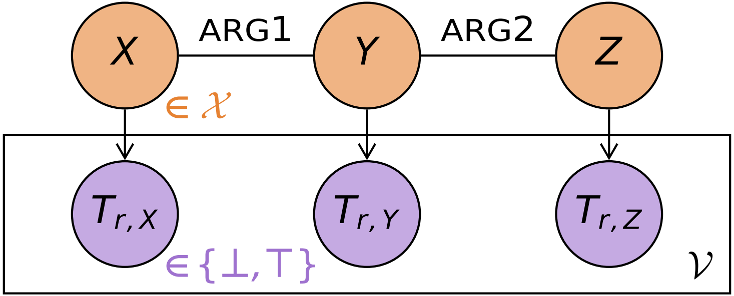

Functional Distributional Semantics

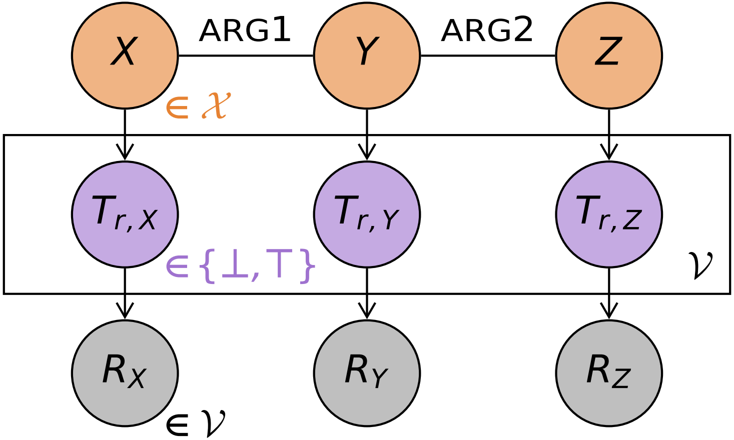

Functional Distributional Semantics

- World model $p(x, y, z)$
- Truth-conditional model $p(t_{r,X} \mid x)$
- Production model $p(r_X \mid x)$

Functional Distributional Semantics

- World model $p(x, y, z)$
- Truth-conditional model $p(t_{r,X} \mid x)$
- Production model $p(r_X \mid x) \propto p(t_{r,X} \mid x)$

Functional Distributional Semantics

- World model $p(x, y, z)$
- Truth-conditional model $p(t_{r,X} \mid x)$
- Production model $p(r_X \mid x) \propto p(t_{r,X} \mid x)$
- Aim: learn these at scale!

Visual Genome (Krishna et al., 2017)

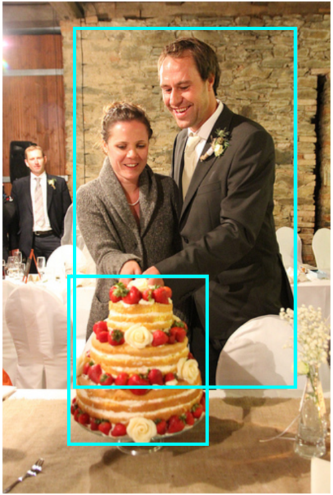

"couple cutting cake"

Visual Genome (Krishna et al., 2017)

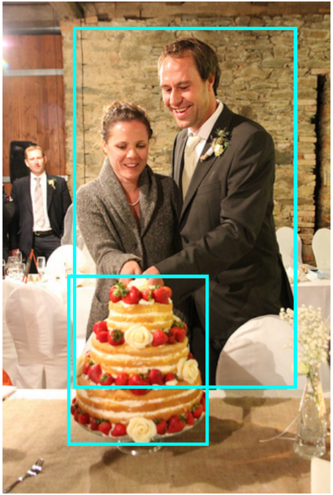

“couple cutting cake”

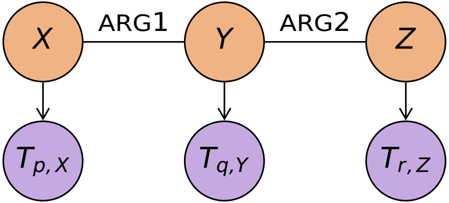

Truth Conditions from Images

1. Data:

2. Objective:

3. Model:

4. Training:

Truth Conditions from Images

1. Data: 2.3m of form (, , , couple, cut, cake)

2. Objective:

3. Model:

4. Training:

Truth Conditions from Images

1. Data: 2.3m of form (, , , couple, cut, cake)
2. Objective: p(, , , couple, cut, cake)
3. Model:
4. Training:

Truth Conditions from Images

1. Data: 2.3m of form (, , , couple, cut, cake)
2. Objective: p(, , , couple, cut, cake)
3. Model: 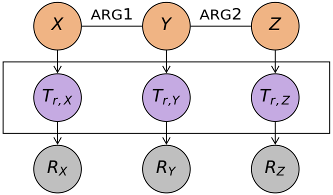

4. Training:

Truth Conditions from Images

1. Data: 2.3m of form (, , , couple, cut, cake)
2. Objective: p(, , , couple, cut, cake)
3. Model: 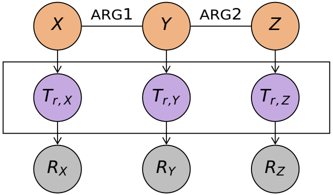

4. Training: gradient descent

Evaluating a Model

- Has the model learnt something useful?

Evaluating a Model

- Has the model learnt something useful?
- Can it generalise to other kinds of data?

Evaluating a Model

- Has the model learnt something useful?
- Can it generalise to other kinds of data?
  - Logical inference

Logical Inference

- Is an animal that has a tail a cat?

Logical Inference

- Is an animal that has a tail a cat?
- Is an animal that has a tail a computer?

Logical Inference with Latent Entities

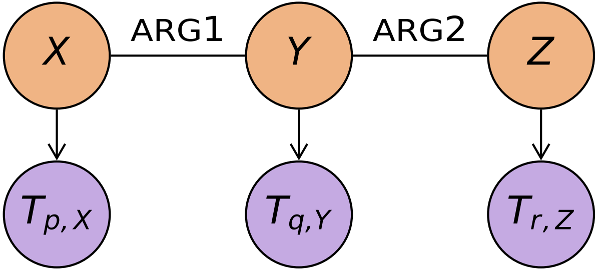

Logical Inference with Latent Entities

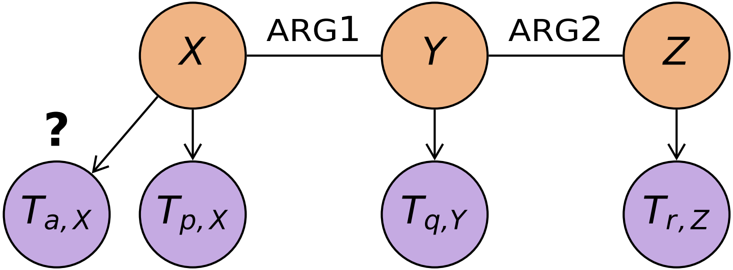

$$p\left(t_{a,X} \,\middle|\, t_{p,X}, t_{q,Y}, t_{r,Z}\right)$$

Logical Inference with Latent Entities

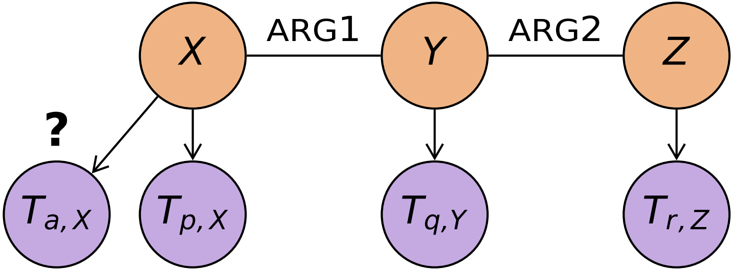

$$p\left(t_{a,X} \,\middle|\, t_{p,X}, t_{q,Y}, t_{r,Z}\right)$$

$$p\left(t_{cat,X} \,\middle|\, t_{animal,X}, t_{have,Y}, t_{tail,Z}\right)$$

Variational Inference for Logical Inference

$$p\left(t_{cat,X} \,|\, t_{animal,X}, t_{have,Y}, t_{tail,Z}\right)$$

Variational Inference for Logical Inference

$$p\left(t_{cat,X} \,|\, t_{animal,X}, t_{have,Y}, t_{tail,Z}\right)$$

$$= \sum_{x,y,z} p\left(t_{cat,X} \,|\, x\right) p\left(x, y, z \,|\, t_{animal,X}, t_{have,Y}, t_{tail,Z}\right)$$

Variational Inference for Logical Inference

$$p\left(t_{cat,X} \,|\, t_{animal,X}, t_{have,Y}, t_{tail,Z}\right)$$

$$= \sum_{x,y,z} p\left(t_{cat,X} \,|\, x\right) p\left(x, y, z \,|\, t_{animal,X}, t_{have,Y}, t_{tail,Z}\right)$$

- Exact inference is computationally intractable

Variational Inference for Logical Inference

$$p\left(t_{cat,X} \,|\, t_{animal,X}, t_{have,Y}, t_{tail,Z}\right)$$

$$= \sum_{x,y,z} p\left(t_{cat,X} \,|\, x\right) p\left(x, y, z \,|\, t_{animal,X}, t_{have,Y}, t_{tail,Z}\right)$$

- Exact inference is computationally intractable
- Need an approximation (e.g. variational inference)

RELPRON Dataset (Rimell et al., 2016)

telescope	*device that astronomers use*
telescope	*device that detects planets*
saw	*device that cuts wood*
philosopher	*person that defends rationalism*
survivor	*person that helicopter saves*
farming	*activity that soil supports*
...	*...*

RELPRON Dataset (Rimell et al., 2016)

telescope	*device that astronomers use*
	device that detects planets
	device that cuts wood
	person that defends rationalism
	person that helicopter saves
	activity that soil supports
	...

RELPRON Dataset (Rimell et al., 2016)

saw	*device that astronomers use*
	device that detects planets
	device that cuts wood
	person that defends rationalism
	person that helicopter saves
	activity that soil supports
	...

RELPRON Dataset (Rimell et al., 2016)

philosopher	*device that astronomers use*
	device that detects planets
	device that cuts wood
	person that defends rationalism
	person that helicopter saves
	activity that soil supports
	...

Similarity in Context (GS2011)

student	write	name
student	spell	name

scholar	write	book
scholar	spell	book

Liu & Emerson (2022)

- Functional Distributional Semantics, trained on Visual Genome
- Evaluation datasets
  - RELPRON: inference with relative clauses
  - GS2011: similarity in context
  - MEN, SL999: similarity (no context)

Liu & Emerson (2022)

- Functional Distributional Semantics, trained on Visual Genome
- Evaluation datasets
  - RELPRON: inference with relative clauses
  - GS2011: similarity in context
  - MEN, SL999: similarity (no context)
  - (All filtered for Visual Genome vocabulary)

Liu & Emerson (2022)

Model	MEN	SL999	GS2011	RELPRON
VG-count (Herbelot, 2020)	.336	.224	.063	.038
VG-retrieval	.420	.190	.072	.045
EVA (Herbelot, 2020)	.543	.390	.068	.032
FDS (Liu & Emerson, 2022)	.639	.431	.171	.117

Liu & Emerson (2022)

Model	MEN	SL999	GS2011	RELPRON
VG-count (Herbelot, 2020)	.336	.224	.063	.038
VG-retrieval	.420	.190	.072	.045
EVA (Herbelot, 2020)	.543	.390	.068	.032
FDS (Liu & Emerson, 2022)	.639	.431	.171	.117

- Truth-conditional structure helps generalisation

Further Reading

- Visual Spatial Reasoning dataset (Liu et al., 2023)
- Functional Distributional Semantics
  - Trained on visual data (Liu & Emerson, 2022)

Further Reading

- Visual Spatial Reasoning dataset (Liu et al., 2023)
- Functional Distributional Semantics
  - Trained on visual data (Liu & Emerson, 2022)
  - Trained on an ontology (Cheng et al., 2023)
  - Trained on text (Emerson, 2018, 2020;
    Lo et al., 2023, 2024)
    – *ACL 2024 Lexical Semantics SAC Award*
  - Probabilistic foundations (Emerson, 2023)

Summary

- Truth conditions feasible at scale
- Improves generalisation

Summary

- Truth conditions feasible at scale
- Improves generalisation
- Approximations required

Crucial Approximations & Simplifications

- Images (Liu & Emerson): variational inference

Crucial Approximations & Simplifications

- Images (Liu & Emerson): variational inference
- Ontology (Cheng et al.): simple truth-conditional model, simple world model
- Text (Lo et al.): amortised variational inference, simple world model

Crucial Approximations & Simplifications

- Images (Liu & Emerson): variational inference
- Ontology (Cheng et al.): simple truth-conditional model, simple world model
- Text (Lo et al.): amortised variational inference, simple world model
- Pragmatics needs further approximation… (see: Emerson, 2023)

Bitter Lesson

- With flexible high-dimensional models, truth-conditional inferences are intractable

Bitter Lesson

- With flexible high-dimensional models, truth-conditional inferences are intractable
- An intractable model is cognitively implausible

Bitter Lesson

- With flexible high-dimensional models, truth-conditional inferences are intractable
- An intractable model is cognitively implausible
- “Approximations” must be part of the theory…

Beyond Truth Conditions

- Truth-conditional model:
  - $p(t_u | s)$ (is *u* true of the situation *s*?)
- World-inferential model:
  - $p(s | t_u)$ (if *u* is true, what is *s* like?)

Beyond Truth Conditions

- Truth-conditional model:
  - $p(t_u|s)$ (is u true of the situation s?)
- World-inferential model:
  - $p(s|t_u)$ (if u is true, what is s like?)
- Ideal rational behaviour follows laws of probability:
  - $p(s|t_u) = \dfrac{p(t_u|s)\,p(s)}{\sum_{s'} p(t_u|s')\,p(s')}$

Example: Looptail g

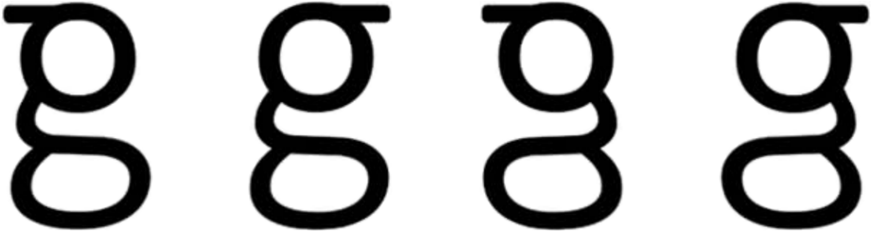

Example: Looptail g

- Recognised easily
- Produced with difficulty or not at all (Wong et al., 2018)

Example: Looptail g

- *s*: visual shape
- *t*: is it a g?

Example: Looptail g

- s: visual shape
- t: is it a g?
- $p(t|s)$ accurate for commonly observed forms
- $p(s|t)$ skewed towards handwritten form

Example: Looptail g

- s: visual shape
- t: is it a g?
- $p(t|s)$ accurate for commonly observed forms
- $p(s|t)$ skewed towards handwritten form
- Human cognition good enough for reading/writing, without perfectly following laws of probability

Compatibility Problem

- Given conditionals $p(t|s)$ and $p(s|t)$, are they compatible with some joint $p(s,t)$?

Compatibility Problem

- Given conditionals $p(t|s)$ and $p(s|t)$, are they compatible with some joint $p(s,t)$?
- This is computationally intractable! (PSPACE-complete)

Fundamental Trade-off

- For high-dimensional probabilistic models, there is a trade-off between:
  - Flexibility (range of distributions?)
  - Tractability (computational cost?)
  - Coherence (following laws of probability?)

Fundamental Trade-off

- Human cognition is flexible and tractable…

Fundamental Trade-off

- Human cognition is flexible and tractable…
- … therefore cannot be perfectly coherent.

Fundamental Trade-off

- Human cognition is flexible and tractable...
- ... therefore cannot be perfectly coherent.
- Does human cognition optimise the trade-off? Is it as coherent as possible, given the constraints?

Judicious Incoherence

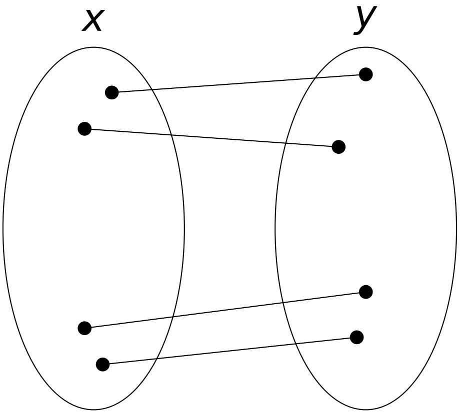

Judicious Incoherence

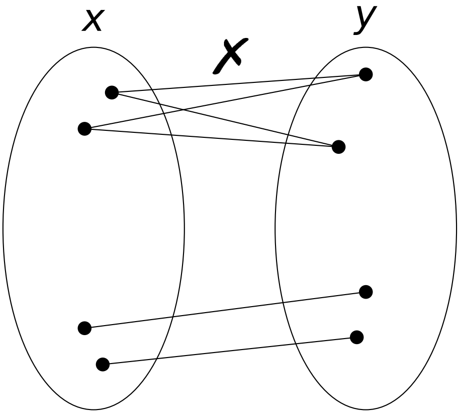

Judicious Incoherence

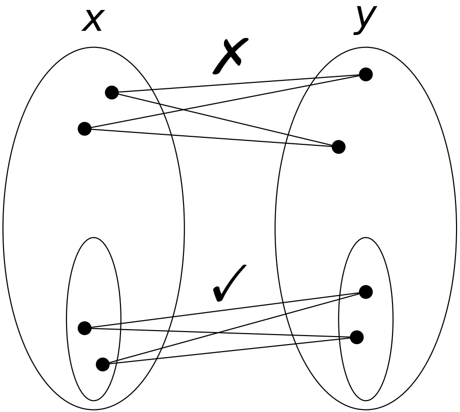

Judicious Incoherence

- Idea: model can appear coherent in a subspace
- Local constraint (easier than Bayesian inference):

$$\frac{p(x_1|y_1)p(y_1|x_2)p(x_2|y_2)p(y_2|x_1)}{p(y_1|x_1)p(x_1|y_2)p(y_2|x_2)p(x_2|y_1)} = 1$$

Judicious Incoherence

- Understanding of looptail g is incoherent for unnatural examples

Judicious Incoherence

- Understanding of looptail g is incoherent for unnatural examples
- Understanding of looptail g is coherent enough to read and write effectively

Judicious Incoherence

- Truth-conditional semantics and pragmatics are incoherent

Judicious Incoherence

- Truth-conditional semantics and pragmatics are incoherent
- Truth-conditional semantics and pragmatics are coherent enough to communicate effectively

Judicious Incoherence

- Truth-conditional semantics and pragmatics are incoherent
- Truth-conditional semantics and pragmatics are coherent enough to communicate effectively
- Experiments planned!

Conclusion

- Truth conditions at scale
  - Feasible (with approximations…)
  - Truth helps generalisation

Conclusion

- Truth conditions at scale
  - Feasible (with approximations...)
  - Truth helps generalisation
- Beyond truth conditions
  - Coherent inference is intractable
  - Alternative: judicious incoherence