

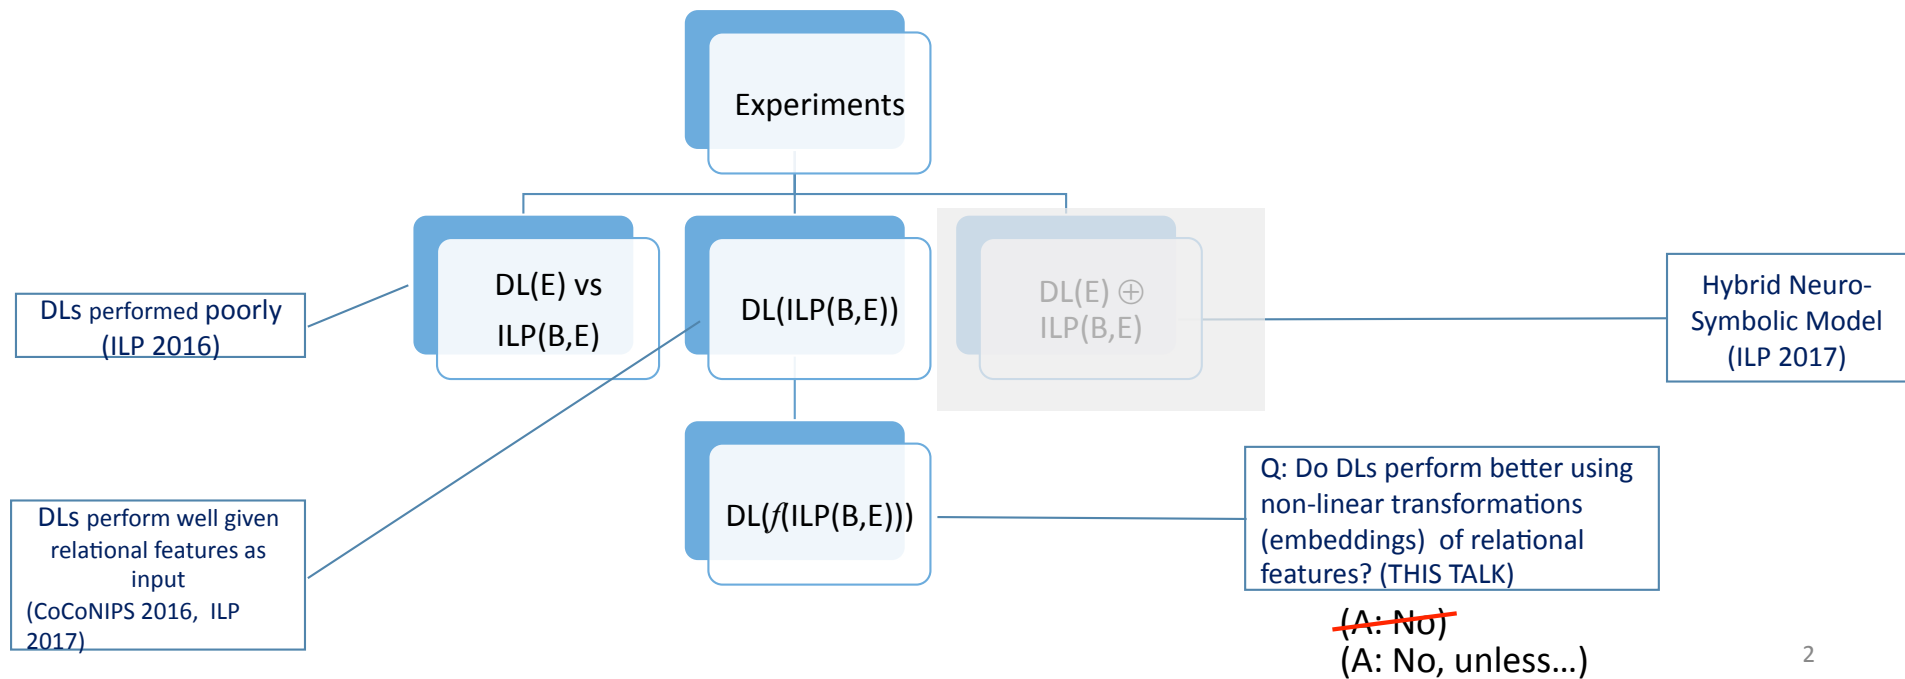
# Embeddings considered Unhelpful?

Ashwin Srinivasan

(joint work with Lovekesh Vig, Michael Bain and Ankit Verma)

# The Big Picture

- This is part of a broader investigation on Deep Learning and ILP
  - Specific interest: small amounts of (relational) data  $E$ , significant domain-knowledge  $B$



# What do the relational features look like?



$F_1(x) = 1$  if  $(Train(x) \wedge HasCar(x, y) \wedge Short(y))$  and 0 otherwise

$F_2(x) = 1$  if  $(Train(x) \wedge HasCar(x, y) \wedge Short(y) \wedge Closed(y))$  and 0 otherwise

$F_3(x) = 1$  if  $(Train(x) \wedge HasCar(x, y) \wedge HasCar(x, z) \wedge InFront(y, z) \wedge Short(y) \wedge Long(z))$  and 0 otherwise


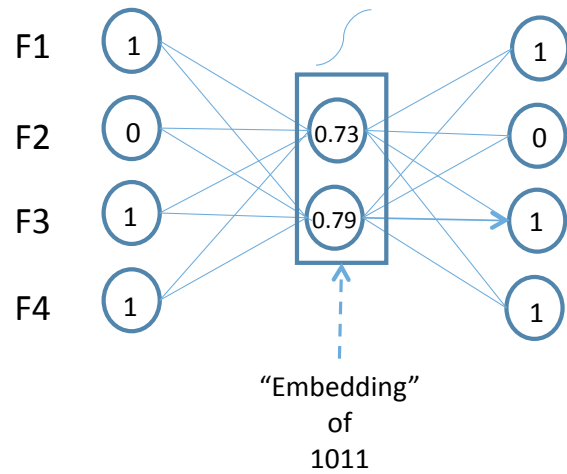
# DNs with relational features

- A DN with relational features as input:
  - Interesting features using an ILP engine: input layer of a DN or in the top-layer of a DN
  - HERE: Do not pre-select using ILP. Give all possible first-order features and let DNs work out what's useful
- Small print:
  - Use a depth-bounded mode language
  - Draw randomly from feature-space
  - Do not want to draw irrelevant features  
 $\forall x (F_1(x) \leftarrow \exists y (Train(x), HasCar(x, y), Short(y), Long(y)))$
  - Avoid drawing redundant features

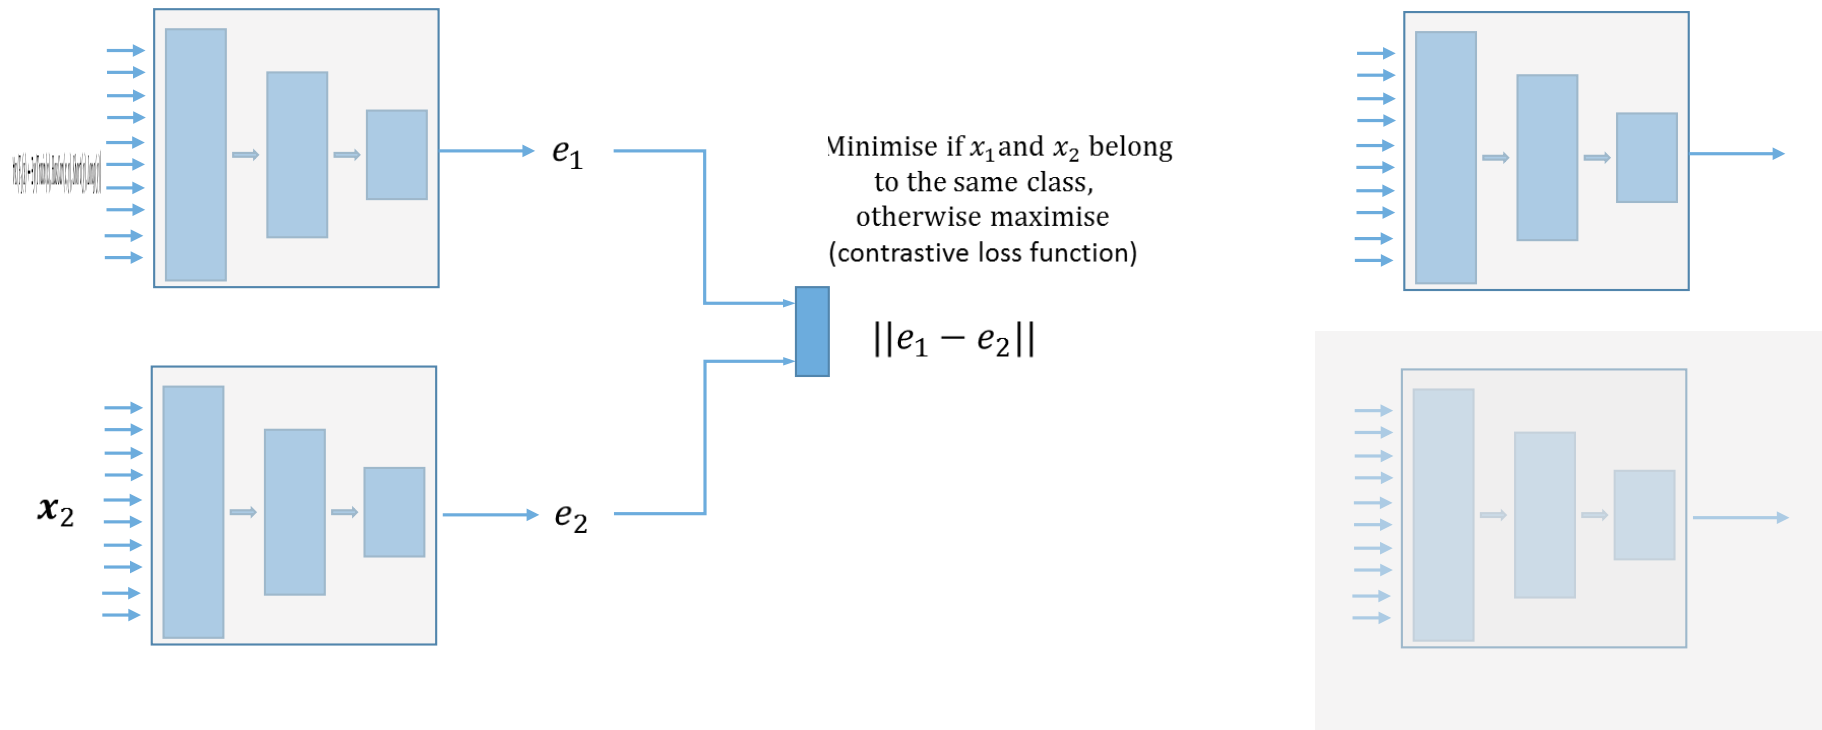
Repeat:

1. Draw example
2. Construct bottom clause
3. Draw clause
4. Check (subsumption) equivalence
5. Construct feature



# Embeddings: The General Idea



# Embeddings for Classification



# What We Are Testing Here



# Do Embeddings Help?

Dataset	Accuracy	
	No EL	With EL
Mut188	0.91	0.91
Canc330	0.68	0.60
DssTox	0.70	0.74
Amine	0.89	0.93
Choline	0.81	0.81
Scop	0.80	0.83
Toxic	0.93	0.93

Not helping. WHY?



# Why don't embeddings help?

- Goes against other reports in the literature
  - Embedding layer encodes missing relations like “opposite gender” given male-female pairs
  - $Pair(x, y) \leftarrow OppositeGender(x, y)$  (“King:Queen = Man:?”)
- Is it the case that if adequate background knowledge exists then embeddings may not be helpful? (CLUE: the datasets Mut188, Canc330 vs DSSTox)
- **Hypothesis:**  
Models using embeddings will perform better than those that do not when background knowledge is deficient

We can test this

# Embeddings and Background Knowledge

Lots of background knowledge

Dataset	Accuracy	
	No EL	With EL
Mut188	0.91	0.91
Canc330	0.68	0.60
DssTox	-	-
Amine	0.89	0.93
Choline	0.81	0.81
Scop	0.80	0.83
Toxic	0.93	0.93

(With EL often not better)

Very little background knowledge

Dataset	Accuracy	
	No EL	With EL
Mut188	0.82	0.84
Canc330	0.56	0.57
DssTox	0.70	0.74
Amine	0.73	0.75
Choline	0.70	0.72
Scop	0.59	0.60
Toxic	0.78	0.80

(With EL always better)

# Concluding Remarks

- In various different ways
  - Background knowledge can make a significant difference to a Deep Learner
- Providing input-features defined using domain predicates is one way to do incorporate background knowledge into a DL
  - Seems to result in better predictivity than models constructed using ILP only
- We would expect DLs that use embeddings of these features to perform even better. But this does not happen
  - Embeddings may be of limited value for DNs whose inputs already incorporate substantial background knowledge

# Acknowledgements

- TCS Research Lab
- SERB EMR/2016/002766