Beyond Candidate Elimination

P. Rullo

rullo@mat.unical.it

Criticisms to Find-S and CE algorithms

- Both provide an exact representation of the training examples. However
  - Conjunctive models may not exist over the given training set
  - If the DNF space is used, only trivial classifiers (no generalization capabilities)
  - Not robust to overfitting

Overfitting (sovradattamento)

- Real-world training data may be noisy (instances erroneously classified)
- Thus, learning a hypothesis that exactly fits the training data may not entail a good generalization to unseen data
- Overfitting occurs when the model is too tailored over the training data so as it may reflect its contingent properties rather than its structural properties

Overfitting due to noise

Training data

Name	Body Temp	Gives birth	4-legged	hibernates	mammal
Porcupine	warm	Yes	Yes	Yes	yes
cat	warm	Yes	Yes	no	yes
bat	warm	yes	no	yes	no
whale	warm	yes	No	No	No
salamander	cold	No	Yes	Yes	No
k. dragon	cold	No	Yes	No	No
Python	cold	No	No	Yes	No
Salmon	cold	No	No	No	No
Eagle	warm	No	No	No	no
Guppy	cold	Yes	No	No	no

Overfitting due to noise

Training data

Name	Body Temp	Gives birth	4-legged	hibernates	mammal
Porcupine	warm	Yes	Yes	Yes	yes
cat	warm	Yes	Yes	no	yes
bat	warm	yes	no	yes	no
whale	warm	yes	No	No	No
salamander	cold	No	Yes	Yes	No
k. dragon	cold	No	Yes	No	No
Python	cold	No	No	Yes	No
Salmon	cold	No	No	No	No
Eagle	warm	No	No	No	no
Guppy	cold	Yes	No	No	no

Misclassified examples

S = <warm, yes, yes, ?> is a model

Overfitting due to noise

S = <warm, yes, yes, ?>

Unseen instances

Name	Body Temp	Gives birth	4-legged	hibernates	mammal	Predicted label
Human	warm	Yes	No	No	yes	no
pigeon	warm	No	No	No	no	no
elephant	warm	yes	Yes	No	yes	yes
Leopard seal	cold	yes	No	No	no	no
turtle	cold	No	Yes	No	no	no
penguin	cold	No	no	No	no	no
eel	cold	No	No	No	no	no
dolphin	warm	yes	No	No	yes	no

- Elephant is correctly classified as mammal
- Both humans and dolphins are misclassified – they are not 4-legged
- All other instances are correctly classified
- 25% test errors

Overfitting due to noise

Training data

Name	Body Temp	Gives birth	4-legged	hibernates	mammal
Porcupine	warm	Yes	Yes	Yes	yes
cat	warm	Yes	Yes	no	yes
bat	warm	yes	no	yes	no
whale	warm	yes	No	No	No
salamander	cold	No	Yes	Yes	No
k. dragon	cold	No	Yes	No	No
Python	cold	No	No	Yes	No
Salmon	cold	No	No	No	No
Eagle	warm	No	No	No	no
Guppy	cold	Yes	No	No	no

S1 = <warm, yes, ?, ?> NOT a model

Overfitting due to noise

S = <warm, yes, ?, ?>

Unseen instances

Name	Body Temp	Gives birth	4-legged	hibernates	mammal	Predicted label
Human	warm	Yes	No	No	yes	yes
pigeon	warm	No	No	No	no	no
elephant	warm	yes	Yes	No	yes	yes
Leopard seal	cold	yes	No	No	no	no
turtle	cold	No	Yes	No	no	no
penguin	cold	No	no	No	no	no
eel	cold	No	No	No	no	no
dolphin	warm	yes	No	No	yes	yes

- 0 test errors

Overfitting due to noise

- S = <warm, yes, yes, ?>
  - compatible with training data
  - 0% training errors, 25% test errors

- S_1 = <warm, yes, ?, ?>
  - NOT compatible with training data
  - 20% training errors, 0% test errors

- S_1 (not a model) is preferred to S (model)

Overfitting due to lack of examples

Training set for classifying mammals

name	Body temp	Gives birth	Four-legged	hibernates	mammal
salamander	cold	no	yes	yes	no
Guppy	cold	yes	no	no	no
Eagle	warm	No	No	No	no
Poorwill	warm	No	No	Yes	no
Platypus	warm	No	Yes	Yes	yes

- The only representative mammal is the platypus!!!
- S = <warm, no, yes, yes>
- Humans are NOT mammals!
- Elephants are NOT mammals!

Overfitting

- **Definition** (Mitchell): a hypothesis **h** is said to overfit the training data if there is another hypothesis **h'** such that **h** has a smaller error over the training examples, but **h'** has a smaller error over the unseen examples

Inductive learning assumption

- Real-world training data are noisy
- So, approximate solutions are preferred
- **Search problem**: Find h $\in$ H that well approximates the training set
- **The Inductive Learning Assumption (ILA)**: Any hypothesis found to approximate the target function well over a sufficiently large set of training examples will also approximate the target function well over other unobserved examples.

Conclusions

- Find-S and CE algorithms
  - Conjunctive models: may not exist
  - DNF: trivial solutions
  - Not robust to overfitting (one instance at a time)
- We need to devise learning algorithms capable of inducing models that well approximate the training data