Machine Learning

ANONYMOUS
CPSC 589
Survey
Dr. Chang-Hyun Jo

Outline

- Machine Learning Overview
- Methods of Machine Learning
- Classification
- Scaling Algorithms
- Reinforcement Learning
- Conclusion

Overview of Machine Learning

- Problem: Algorithms only good for a subset of a problem domain.
- Solution: Enable machines to learn in order to adapt to new situations.
- What is machine learning?
  - Creating a machine with the ability to learn.

Categories of Research

1. Improvement of classification accuracy
   - ensemble of classifiers
2. Methods of Scaling learning algorithms
3. Reinforcement Learning

Methods of Machine Learning

- Rote Learning:
  - 1 to 1 mapping of inputs to outputs
- Induction:
  - Use examples to reach a conclusion
- Clustering:
  - Group similar objects together

Methods of Machine Learning

- Analogy:
  - Determine relationships between 2 representations
- Discovery Learning:
  - Form of unsupervised learning
  - Machine learns through exploration
- Reinforcement Learning:
  - Positive / negative feedback based on actions

Classification

- Classification: The process of classifying objects correctly.
  - Measuring classification algorithms:
    - Determining how well a classifier predicts the true function *f*
    - Space requirements
    - Time it takes to predict the true function *f*
- Classifier: a hypothesis that attempts to predict the true function *f*

Classification

- One method: Supervised learning
  - Input is given to the machine
  - Machine is told what the correct output should be
  - With each set of new training data, the machine forms a classifier to predict the true function *f*
- Can be used to predict unknown functions
- Current Research: Ensemble of Classifiers

Ensemble of Classifiers

- Advantages:
  - Allows a decision to made based on weighted classifiers
  - Many classifiers come into play in making a decision
  - More accurate than a single classifier
- Problem: How do you find a good ensemble of classifiers?

Ensemble of Classifiers

- 4 Methods currently being researched:
  1. Sub-sampling the training examples
  2. Hand picking input to create the desired subsets
  3. Manipulating the output targets
  4. Randomly generating new classifers
- Method 2 has had real world success
  - Identifying volcanoes

Ensemble of Classifiers

- New Method that is similar to hand picking input
  - Using a GUI to create an ensemble of classifiers
  - Advantage: Person that builds the ensemble can apply background knowledge to the problem domain
  - Example tests:
    - Identifying Kiwi vines
      - No knowledge user: 53 nodes with 85.8% accuracy
      - Machine: 93 nodes with 83.2% accuracy
- Problems:
  - Time consuming
  - Currently limited to 2 dimensions

Scaling Algorithms

- Current algorithms can handle millions of examples of training data
- Can these same algorithms handle billions?
  - Many haven't
- 4 methods currently being investigated:
  1. Sampling subsets of the training data
  2. Creating new data structures to hold data in RAM
  3. Using ensembles of decision trees in conjunction with multiple processors to speed up data processing
  4. Breaking up the training data into ranged categories (e.g 1-5, 6-10, etc.)

Reinforcement Learning

- Form of unsupervised learning
  - Machine is never told what the correction action is.
  - Negative / positive rewards given to the machine based on the action it takes.
  - Learning occurs when the machine makes new decisions based on the feedback it received
- Real World Example:
  - Elevator using Q-Learning

Reinforcement Learning

- Problem with current algorithms:
  - Don't work well with large search spaces
    - Research: Hierarchical Average Reward Learning Model
  - Learning process occurs too slowly
    - Research: Coarse Graining of Perception

Reinforcement Learning

- Coarse Graining of Perception:
  - Reduce the number of states that need to be investigated
  - Problems:
    - Reduce accuracy
    - "Bad Habits" obtained early on in the learning process are hard/impossible to fix later
  - Pro:
    - Quicker than Complete Perception

Reinforcement Learning

- Complete Perception:
  - Pro:
    - Higher accuracy
  - Problem:
    - Very slow process for learning to occur

Reinforcement Learning

- Example:

$n \backslash m$	3	5	7	9
1	9	25	49	81
2	81	625	2401	6561
3	729	15,625	117,649	531,441
4	6561	390,625	5,764,801	43,046,721

- n = hunters
- m = lattice space
- Pentium III: 640 Mhz with 2 GB of RAM
  - 8 hours for learning to occur

Reinforcement Learning

- Example:

$n \backslash m$	3	5	7	9
1	9	25	49	81
2	81	625	2401	6561
3	729	15,625	117,649	531,441
4	6561	390,625	5,764,801	43,046,721

- using coarse graining:
  - 3 hunters, lattice size of 7
  - Complete perception states: 117, 649 vs
  - Coarse graining perception states: 49

Reinforcement Learning

- Problem:
  - Coarse graining (speed, bad accuracy) vs Complete Perception (better accuracy, slow)
- Solution: Use a combination of both
  - Need to find the time to switch from coarse graining to complete
  - Switching too early:
    - slow
  - Switching too late:
  - possibility of bad habits
  - lower accuracy
  - In experiments, used residual entropy to measure the switching point

Reinforcement Learning

- Reinforcement learning research in the opposite direction: Inverse Reinforcement Learning
- Inverse Reinforcement Learning:
  - Machine attempts to learn the reward function *f* by studying an "expert"
  - Difference: normal reinforcement learning forms the optimal policy based on rewards
    - Inverse reinforcement learning tries to figure out the reward function by observing an "expert"

Reinforcement Learning

- Current research in Inverse Reinforcement Learning is promising
- Problems with Inverse Reinforcement Learning:
  - Current methods don't scale well
  - Current methods are easily susceptible to noise

Conclusion

- Machine Learning:
  - Process of creating a machine that possesses the ability to learn and make decisions
- Current Research involves:
  - Classification
  - Scaling Algorithms
  - Reinforcement Learning

References

[Dietterich 1997] Dietterich, Thomas G. 1997. Machine-Learning Research: Four Current Directions, AI Magazine, 15 (3), 97-136.

[Doyle 2003] Doyle, Patrick. 2003. Learning, Stanford University, http://www.cs.dartmouth.edu/%7Ebrd/Teaching/AI/Lectures/Summaries/learning.html.

[Dyer 2003] Dyer, C. 2003. CS 540 Lecture Notes, University of Wisconsin - Madison, http://www.cs.wisc.edu/%7Edyer/cs540/notes/learning.html.

[Frank et al. 2001] Frank, Eibe, Holmes, G., Ware, M., Witten, I. 2001. Interactive Machine Learning: Letting Users Build Classifiers, International Journal of Human-Computer Studies, 55 (3), 281-292.

[Ito et al. 2001] Ito, Akira, and Mitsuru Kanabuchi. 2001. Speeding Up Multiagent Reinforcement Learning by Coarse-Graining of Perception: The Hunter Game, Electronics and Communications in Japan, 84 (12), 37-45.

References

[Luger 2002] Luger, George F. 2002. Artificial Intelligence: Structures and Strategies for Complex Problem Solving. Pearson Education Limited, Edinburgh Gate, Harlow, Essex CM20 2JE.

[Mitchell 1997] Mitchell, Tom. 1997. Does Machine Learning Really Work?, AI Magazine, 18 (3), 11-20.

[Ng and Russell 2000] Ng, Andrew Y. and Russell, Stuart. 2000. Algorithms for Inverse Reinforcement Learning, Proceedings of the Seventeenth International Conference on Machine Learning, Stanford University, (July), 663-670.

[Seri and Tadepalli 2002] Seri, Sandeep and Tadepalli, Prasad. 2002. Model-based Hierarchical Average-reward Reinforcement Learning, International Conference on Machine Learning, Sydney, Australia, (July), 562-569.

[Tesauro 1995] Tesauro, Gerald. 1995. Temporal Different Learning and TD-Gammon, Communications of the ACM, 38 (3), 58-68.