

Vis tools and case studies

Thomas Torsney-Weir

VDA research group, University of Vienna

Outline

- Introduction to Machine learning
- Vis helping machine learning
- Machine learning helping vis


What is machine learning?

"A rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome"

Russell, Stuart, and Peter Norvig. *Artificial intelligence: A modern approach*, 2009.


Algorithms that can improve their performance based on training data

What is machine learning?




You


What is machine learning?



What is machine learning?



What is machine learning?



Types of algorithms


- Regression
- Classification
- Clustering

What to use?

- Regression: Predict continuous values
- Classification: Predict discrete values
- Clustering: Find distributions

Regression


Predict continuous values



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Regression

Predict continuous values




$$E(\mathbf{w}) = \frac{1}{2} \sum_{n=1}^N \{y(x_n, \mathbf{w}) - t_n\}^2$$

Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Regression


Predict continuous values



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Regression


Predict continuous values



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Regression


Predict continuous values



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Regression


Predict continuous values



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Regression


Predict continuous values



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Classification


Predict discrete values



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Clustering


Find distributions



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Clustering


Find distributions



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Clustering


Find distributions



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Clustering


Find distributions



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Clustering


Find distributions



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Clustering

Find distributions



Bishop, Christopher M. *Pattern recognition and machine learning (information science and statistics)*, 2006.

Types of ML algorithms

- Regression: Predict continuous values
- Classification: Predict discrete values
- Clustering: Find distributions

Uses

- naive Bayes: spam filtering
- classification: recommender systems
- neural networks: handwriting recognition
- HMM: speech recognition


Email filtering

- is email spam or not?
- use words as features

$$P(C=c_k|X=x) = \frac{P(X=x|C=c_k)P(C=c_k)}{P(X=x)}$$


Sahami, Mehran, Susan Dumais, David Heckerman, and Eric Horvitz. "A bayesian approach to filtering junk e-mail," 1998.

Recommender systems




<http://blog.soton.ac.uk/hive/2012/05/10/recommendation-system-of-hive/>

Handwriting recognition



Speech recognition




http://recognize-speech.com/images/LanguageModel/left_to_right_HMM.png


Vis and ML

both vis and ML seem to have similar goals: make sense of complex data

Machine learning



Visualization



Morton, Kristi, Ross Bunker, Jock Mackinlay, Robert Morton, and Chris Stolte.
“Dynamic workload driven data integration in Tableau,” 2012.

Who helps whom?

both!

- Vis helps ML: evaluating models
- ML helps vis: ML for embedded analysis

Why is this hard?

Machine learning

- Fast algorithms
- Sufficient data
- Automatic learning

Visualization

- Multi-dimensional spaces
- Comparing complex data
- Showing uncertainty

Vis helping ML

Vis helping ML

How do they work together?

- Building models
- Validating models
- Understanding models

Building models

Building models



- Meta parameters
- Model selection

What are meta parameters?



Meta parameters control how learning takes place

- Learning rate
- Number and size of network layers
- Slack variables
- Stopping conditions


Why study meta-parameters?




Why study meta-parameters?



Manual method



Manual method




How to study them?


run a bunch of models and examine outputs

- paramorama
- design galleries

Paramorama




Manual sampling



Automated sampling


Pretorius, A. Johannes, Mark-Anthony P. Bray, Anne E. Carpenter, and Roy A. Ruddle. "Visualization of parameter space for image analysis," 2011.

Paramorama




Pretorius, A. Johannes, Mark-Anthony P. Bray, Anne E. Carpenter, and Roy A. Ruddle. "Visualization of parameter space for image analysis," 2011.

Paramorama



Pretorius, A. Johannes, Mark-Anthony P. Bray, Anne E. Carpenter, and Roy A. Ruddle. "Visualization of parameter space for image analysis," 2011.


Design galleries




Marks, Joe, Brad Andelman, Paul A. Beardsley, William Freeman, Sarah Gibson, Jessica Hodgins, Thomas Kang, et al. "Design Galleries: A general approach to setting parameters for computer graphics and animation," 1997.

How to study them?


use a more principled approach



Objective measures




Visual parameter space exploration




Michael Sedlmair, Christoph Heinzl, Stefan Bruckner, Harald Piringer, and Torsten Möller "Visual parameter space analysis: A conceptual framework" IEEE Transactions on Visualization and Computer Graphics. 20(12) 2014.

Tuner




Torsney-Weir, Thomas, Ahmed Saad, Torsten Möller, Britta Weber, Hans-Christian Hege, Jean-Marc Verbavatz, and Steven Bergner. "Tuner: Principled parameter finding for image segmentation algorithms using visual response surface exploration," 2011.

Tuner




Torsney-Weir, Thomas, Ahmed Saad, Torsten Möller, Britta Weber, Hans-Christian Hege, Jean-Marc Verbaatz, and Steven Bergner. "Tuner: Principled parameter finding for image segmentation algorithms using visual response surface exploration," 2011.

Tuner



Torsney-Weir, Thomas, Ahmed Saad, Torsten Möller, Britta Weber, Hans-Christian Hege, Jean-Marc Verbavatz, and Steven Bergner. "Tuner: Principled parameter finding for image segmentation algorithms using visual response surface exploration," 2011.

Tuner



Torsney-Weir, Thomas, Ahmed Saad, Torsten Möller, Britta Weber, Hans-Christian Hege, Jean-Marc Verbavatz, and Steven Bergner. "Tuner: Principled parameter finding for image segmentation algorithms using visual response surface exploration," 2011.


Building models

- Meta parameters can have a large influence on performance
- Almost *all* ML algorithms require tuning
- Manual tuning is time consuming and error prone

Validating and verifying models

What do we mean?

- How do we know our models are working?
- model selection



Committee on Mathematical Foundations of Verification, Validation, and Uncertainty Quantification; Board on Mathematical Sciences and Their Applications,

Division on Engineering and Physical Sciences, National Research Council. *Assessing the reliability of complex models: Mathematical and statistical foundations of verification, validation, and uncertainty quantification*, 2012. http://www.nap.edu/openbook.php?record_id=13395.


Validating and verifying models

- Summary statistics are not always enough
- Balancing multiple objectives is difficult
- Certain training points might be very important

Examples


- HyperMoVal - local inspection
- Sliceplorer - global inspection
- Tuner - error inspection

HyperMoVal




Piringer, Harald, Wolfgang Berger, and Jürgen Krasser. "HyperMoVal: Interactive visual validation of regression models for real-time simulation," 2010.


Sliceplorer views




Single layer NN (26 nodes)



Dual layer NN (5 and 3 nodes)



SVM (polynomial kernel)



SVM (RBF kernel)


Torsney-Weir, Thomas, Michael Sedlmair, and Torsten Möller. "Sliceplorer," 2017.

Tuner error views


Examining multi-dimensional functions

- error view shows where model is unsure
- can visually verify the model

Prediction



Optimization



Error view

Error view

Validating and verifying models

- Understand fit for individual samples
- Visual inspection to understand extrapolation
- Uncertainty can help to understand quality of prediction

Understanding models

Who needs this?

- models are complex
- the business world likes spreadsheets because they can walk through the calculations

Simple vs complex models

Simple

- few factors
- small integer factors
- low-depth trees

Complex

- multi-layer neural network
- Gaussian process model
- non-linear
- many decisions

What does complexity buy us?

- Global vs local models
- Deep-learning networks can deal with feature selection
- Can deal with edge cases

Understanding models

- Just an answer is not enough (show your work)
- Humans have trouble understanding complex models
- Interactivity can bring people into the model

Methods


- interaction
- walkthroughs
- simpler models ala LIME (Ribeiro et al. 2016)
- direct inspection

Examples

- regression: Mühlbacher and Piringer
- clustering: Dis-function
- text processing: TagRefinery
- smaller models: Explanation explorer

Mühlbacher and Piringer


Directly interact with the model building process



Mühlbacher, Thomas, and Harald Piringer. "A partition-based framework for building and validating regression models," 2013. Best Paper Award.

Dis-function


Build a distance function interactively



Brown, Eli T, Jingjing Liu, Carla E Brodley, and Remco Chang. "Dis-Function: Learning Distance Functions Interactively," 2012.

TagRefinery


Tutorial/walkthrough system



Text processing pipeline


Kralj, Christoph, Mohsen Kamalzadeh, and Torsten Möller. “TagRefinery: A visual tool for tag wrangling,” 2017.

TagRefinery





Kralj, Christoph, Mohsen Kamalzadeh, and Torsten Möller. "TagRefinery: A visual tool for tag wrangling," 2017.

TagRefinery




Kralj, Christoph, Mohsen Kamalzadeh, and Torsten Möller. "TagRefinery: A visual tool for tag wrangling," 2017.

LIME method



Ribeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. “Why should I trust you?: Explaining the predictions of any classifier,” 2016.


Explanation explorer



Krause, Josua, Aritra Dasgupta, Jordan Swartz, Yindalon Aphinyanaphongs, and Enrico Bertini. “A workflow for visual diagnostics of binary classifiers using instance-level explanations,” 2017.

Direct inspection

e.g. hidden states in a neural network



DeepEyes


How many layers are needed?



Pezzotti, Nicola, Thomas Höllt, Jan van Gemert, Boudeijn Lelieveldt, Elmar Eisemann, and Anna Vilanova. "DeepEyes: Progressive visual analytics for designing deep neural networks," 2018.

DeepEyes


How stable are the layers?



Pezzotti, Nicola, Thomas Höllt, Jan van Gemert, Boudeijn Lelieveldt, Elmar Eisemann, and Anna Vilanova. "DeepEyes: Progressive visual analytics for designing deep neural networks," 2018.

DeepEyes

How discriminating are the layers?



Pezzotti, Nicola, Thomas Höllt, Jan van Gemert, Boudewijn Lelieveldt, Elmar Eisemann, and Anna Vilanova. "DeepEyes: Progressive visual analytics for designing deep neural networks," 2018.

Vis helping ML

How do they work together?


- Building models
- Validating models
- Understanding models

What about the other way?

Machine learning helping vis


How?

- Use the strengths of ML and vis together
 - machines are good at calculating
 - humans are good at intuition
- Vis assisted by ML algorithms




Sacha, D., A. Stoffel, F. Stoffel, Bum Chul Kwon, G. Ellis, and Daniel A Keim. "Knowledge generation model for visual analytics," 2014.

Book ad!



How?



- ML provides data aggregation/filtering/selection
- User can steer algo to produce desired results

Examples


- Cluster and calendar view - clustering
- KeyVis - clustering
- FluidExplorer - clustering
- Cell Cognition Explorer - active learning

Cluster and calendar view

Understanding power consumption

- When do people use the most power?
- What are seasonal patterns?
- How does energy use change at different scales
 - Days
 - Weeks
 - Months


Cluster and calendar view



Van Wijk, J.J., and E.R. Van Selow. "Cluster and calendar based visualization of time series data," 1999.

<http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=801851>.


Cluster and calendar view



Van Wijk, J.J., and E.R. Van Selow. "Cluster and calendar based visualization of time series data," 1999.

<http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=801851>.


Cluster and calendar view



Van Wijk, J.J., and E.R. Van Selow. "Cluster and calendar based visualization of time series data," 1999.

<http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=801851>.

Cluster and calendar view



Van Wijk, J.J., and E.R. Van Selow. "Cluster and calendar based visualization of time series data," 1999.

<http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=801851>.

KeyVis




- How to find related work?
- How helpful are keywords?
- Do keywords relate to each other?

Isenberg, Petra, Tobias Isenberg, Michael Sedlmair, Jian Chen, and Torsten Möller. "Visualization as seen through its research paper keywords," 2017.

<https://tobias.isenberg.cc/VideosAndDemos/Isenberg2017VST>.

KeyVis

step 1: cluster the papers based on keywords



Isenberg, Petra, Tobias Isenberg, Michael Sedlmair, Jian Chen, and Torsten Möller. "Visualization as seen through its research paper keywords," 2017.

<https://tobias.isenberg.cc/VideosAndDemos/Isenberg2017VST>.

KeyVis


step 2: give an interface to this clustering



Isenberg, Petra, Tobias Isenberg, Michael Sedlmair, Jian Chen, and Torsten Möller. "Visualization as seen through its research paper keywords," 2017.

<https://tobias.isenberg.cc/VideosAndDemos/Isenberg2017VST>.


FluidExplorer



Bruckner, Stefan, and Torsten Möller. "Result-driven exploration of simulation parameter spaces for visual effects design," 2010.

<http://www.ncbi.nlm.nih.gov/pubmed/20975188>.


FluidExplorer



Bruckner, Stefan, and Torsten Möller. "Result-driven exploration of simulation parameter spaces for visual effects design," 2010.

<http://www.ncbi.nlm.nih.gov/pubmed/20975188>.

Cell Cognition Explorer



Sommer, Christoph, Rudolf Hoefler, Matthias Samwer, and Daniel W. Gerlich. "A Deep Learning And Novelty Detection Framework For Rapid Phenotyping In High-Content Screening," 2017. <http://dx.doi.org/10.1091/mcb.E17-05-0333>

The future!

Interesting projects

- More using ML to build models for vis tools
- More generalized tools
- Understand what "understandability" means

Thanks!

thomas.torsney-weir@univie.ac.at

http://www.tomtorsneyweir.com