Visualization for Machine Learning

Cmpt 767 - Visualization

Steven Bergner

sbergner@sfu.ca

Overview

- Machine learning tasks
- Vis for ML
- ML for Vis

Visualization: Secret Weapon for ML

Fernanda Viegas and Martin Wattenberg VDS 2017 Keynote

- ML is important in the world, dramatic tech leaps
- What is the best way to train models and debug them?
- How to understand what's going on under the hood of deep NNs?
- Interactive exploration can help people use, interpret, and learn about machine intelligence

Tasks performed **by** and **for** ML

ML Task Categories: By Output

- Classification
- Regression
- Clustering
- Association rules
- Forecasting
- Dimensional reduction
- Density estimation

[wikipedia]

ML Task categories: By Training

- **Supervised learning**
  - Inputs and desired outputs are given
  - Find rule that maps unseen inputs to outputs
- **Semi-supervised learning**
  - Supervised learning with only few of the target outputs given
- **Active learning**
  - Data is not given, but asked for
- **Unsupervised learning**
  - No labels given
  - Discover hidden patterns and structure in inputs
  - Feature learning
- **Reinforcement learning**
  - Rewards/punishments given as feedback to actions in dynamic environment

[wikipedia]

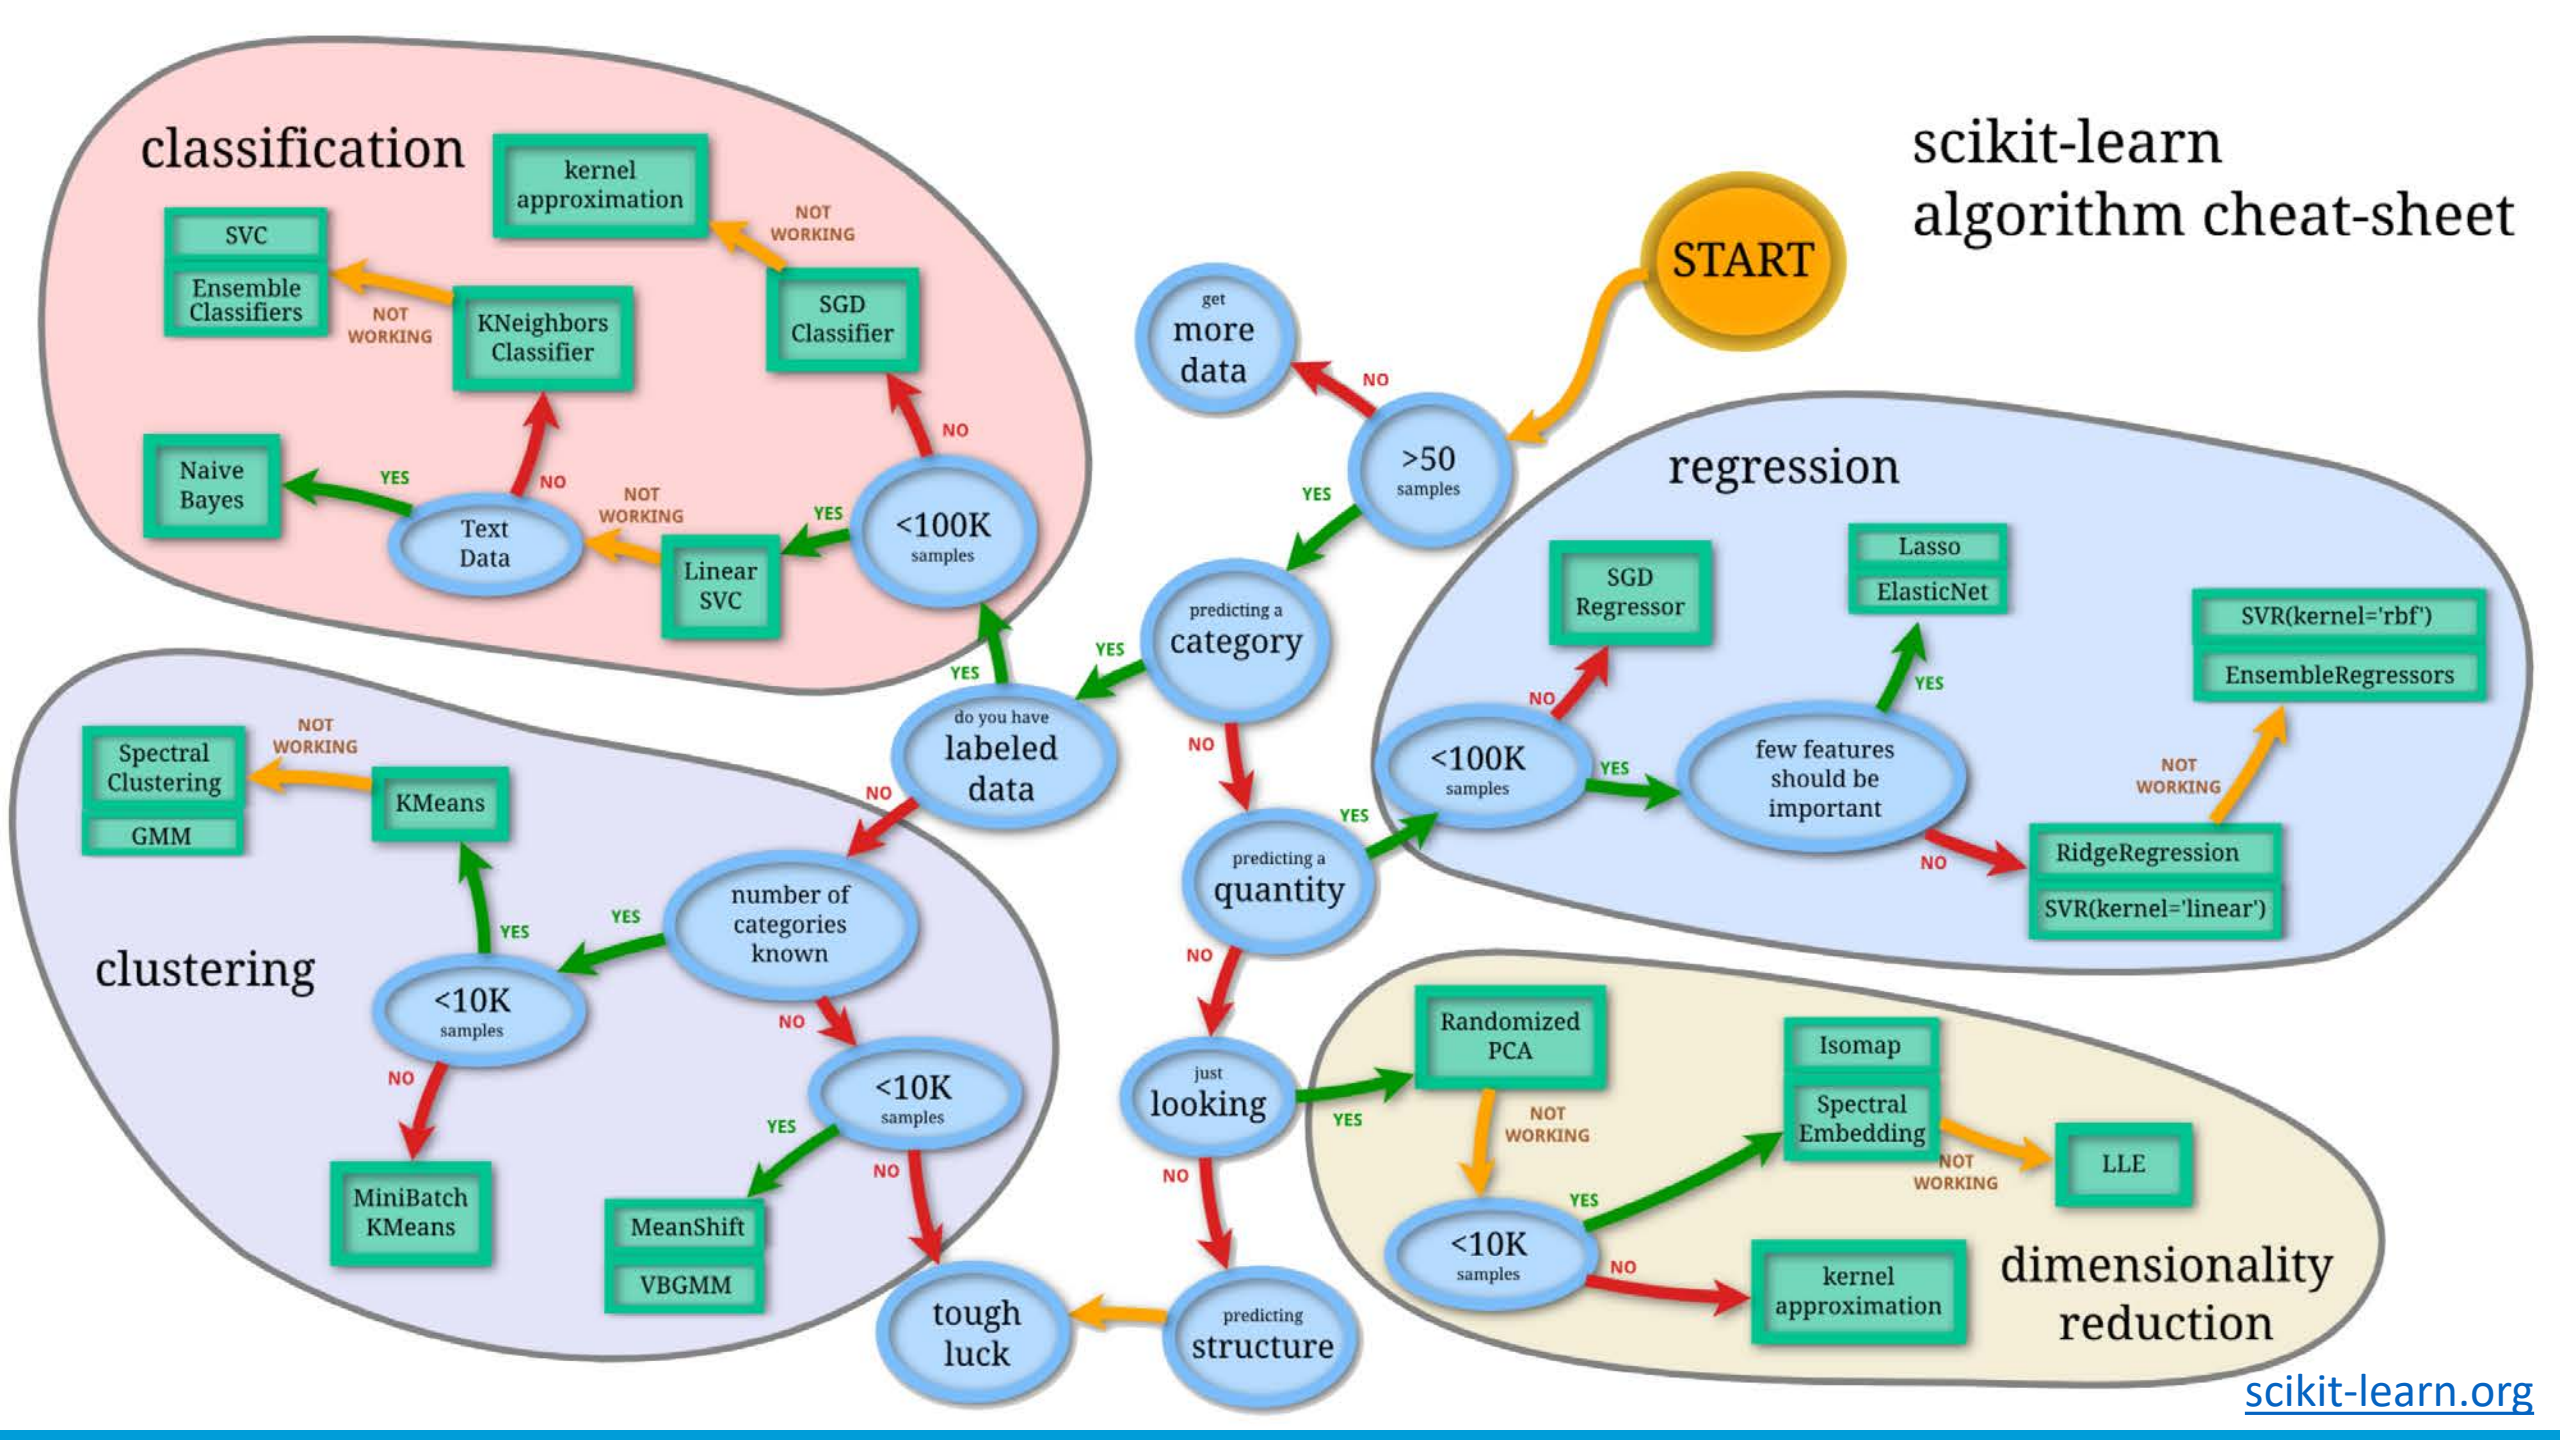
scikit-learn
algorithm cheat-sheet
START
get more data
NO
>50 samples
YES
predicting a category
YES
do you have labeled data
YES
<100K samples
YES
Linear SVC
NOT WORKING
Text Data
NO
Naive Bayes
YES
KNeighbors Classifier
NOT WORKING
Ensemble Classifiers
SVC
NOT WORKING
kernel approximation
SGD Classifier
NO
classification
NO
predicting a quantity
YES
<100K samples
NO
SGD Regressor
YES
few features should be important
YES
Lasso
ElasticNet
NO
RidgeRegression
SVR(kernel='linear')
NOT WORKING
SVR(kernel='rbf')
EnsembleRegressors
regression
NO
just looking
YES
Randomized PCA
NOT WORKING
<10K samples
YES
Isomap
Spectral Embedding
NOT WORKING
LLE
NO
kernel approximation
dimensionality reduction
NO
predicting structure
tough luck
NO
number of categories known
YES
<10K samples
YES
KMeans
NOT WORKING
Spectral Clustering
GMM
NO
<10K samples
YES
MeanShift
VBGMM
NO
MiniBatch KMeans
clustering
scikit-learn.org

Tasks **for** Machine Learning

For a given problem

- Choose a model class & estimator, loss function, optimization method
- Determine the right training data (attributes, distribution)

For a given model class

- Estimate model parameters to fit with given observations
- Make predictions, determine and communicate uncertainty
- Analyse model behaviour over a region of inputs
- Understand how model works, explain its decision making
- Validate fit of training assumptions vs operating conditions

All of the above: data-driven

Some of them: human-in-the-loop

Cluster Analysis

Cluster Visualization

- Treat cluster label as categorical variable
- Multi-variate vis technique with encoding for label attribute

Example: DimStiller Workflows

- Goal: Understand and transform input data
- Chaining together operators into pipelines

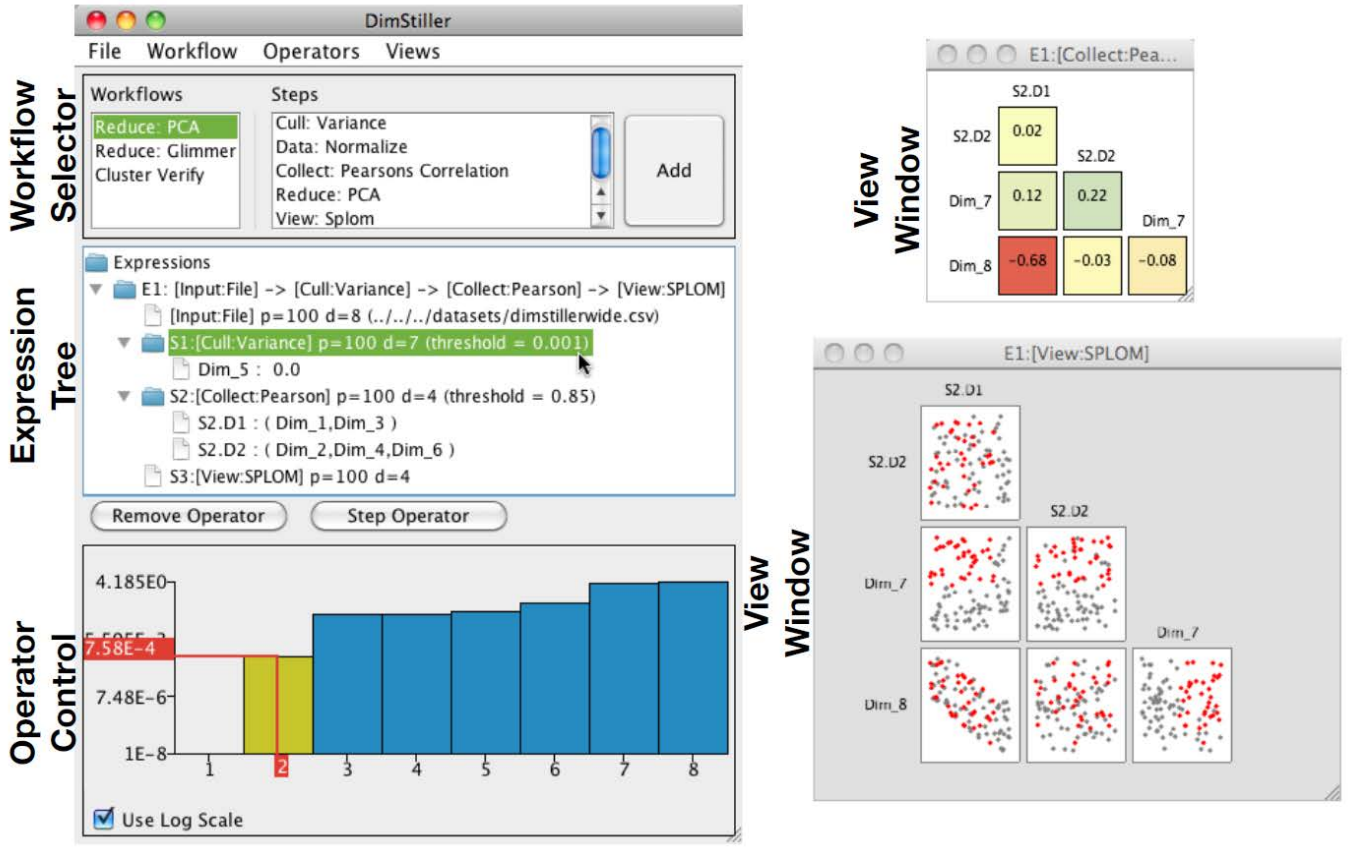

DimStiller Cluster Analysis Example

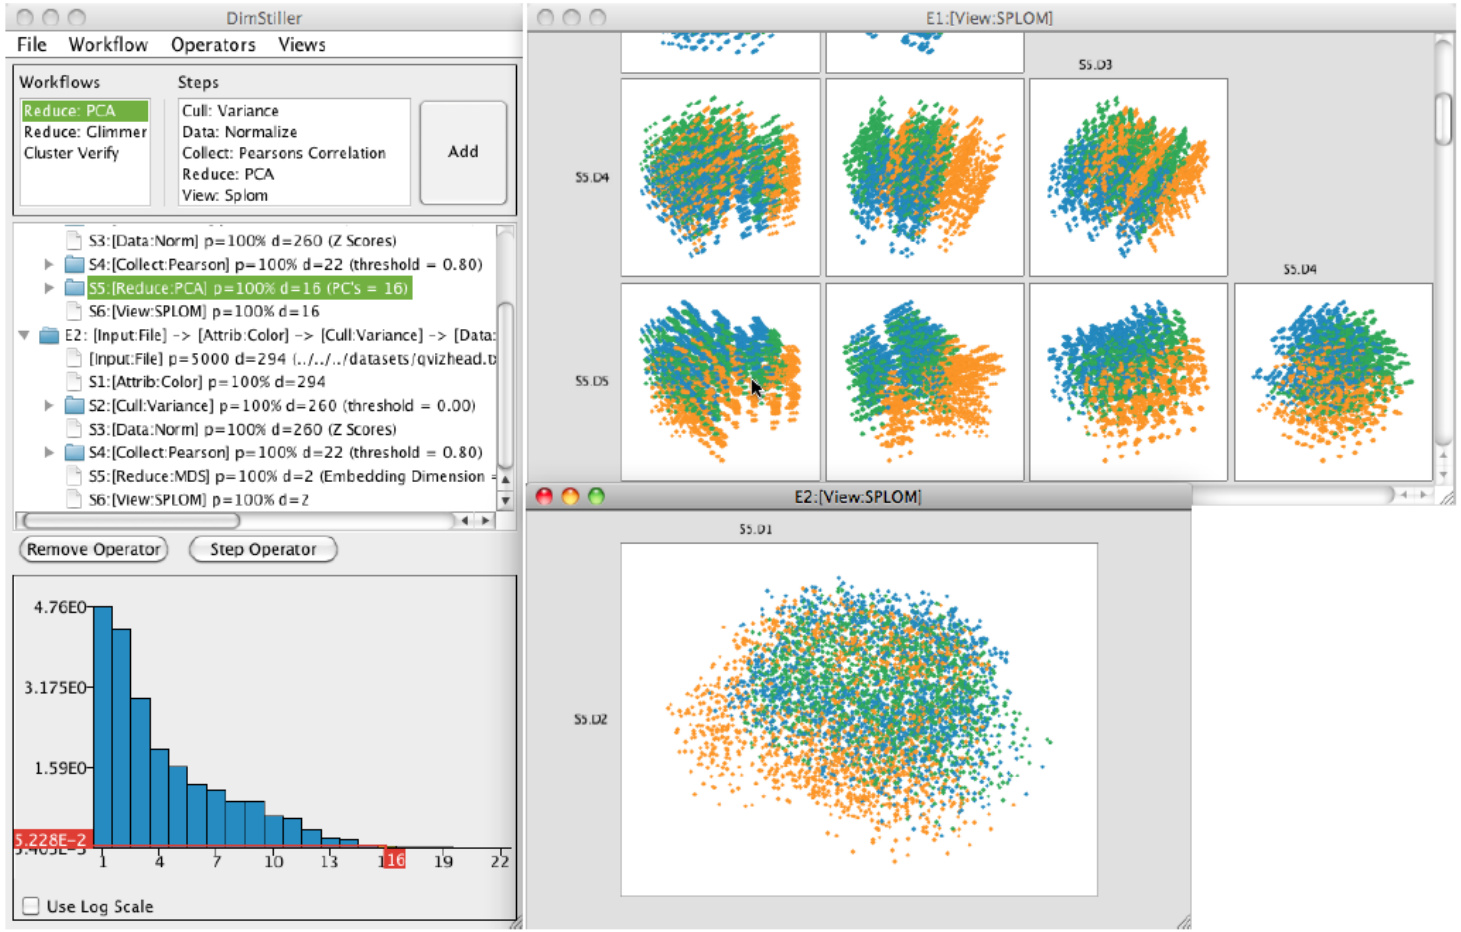

Dimension Reduction

Dimension Reduction

- Principal Component Analysis
- t-SNE
- Independent Component Analysis
- Manifold Learning

Manifold learning

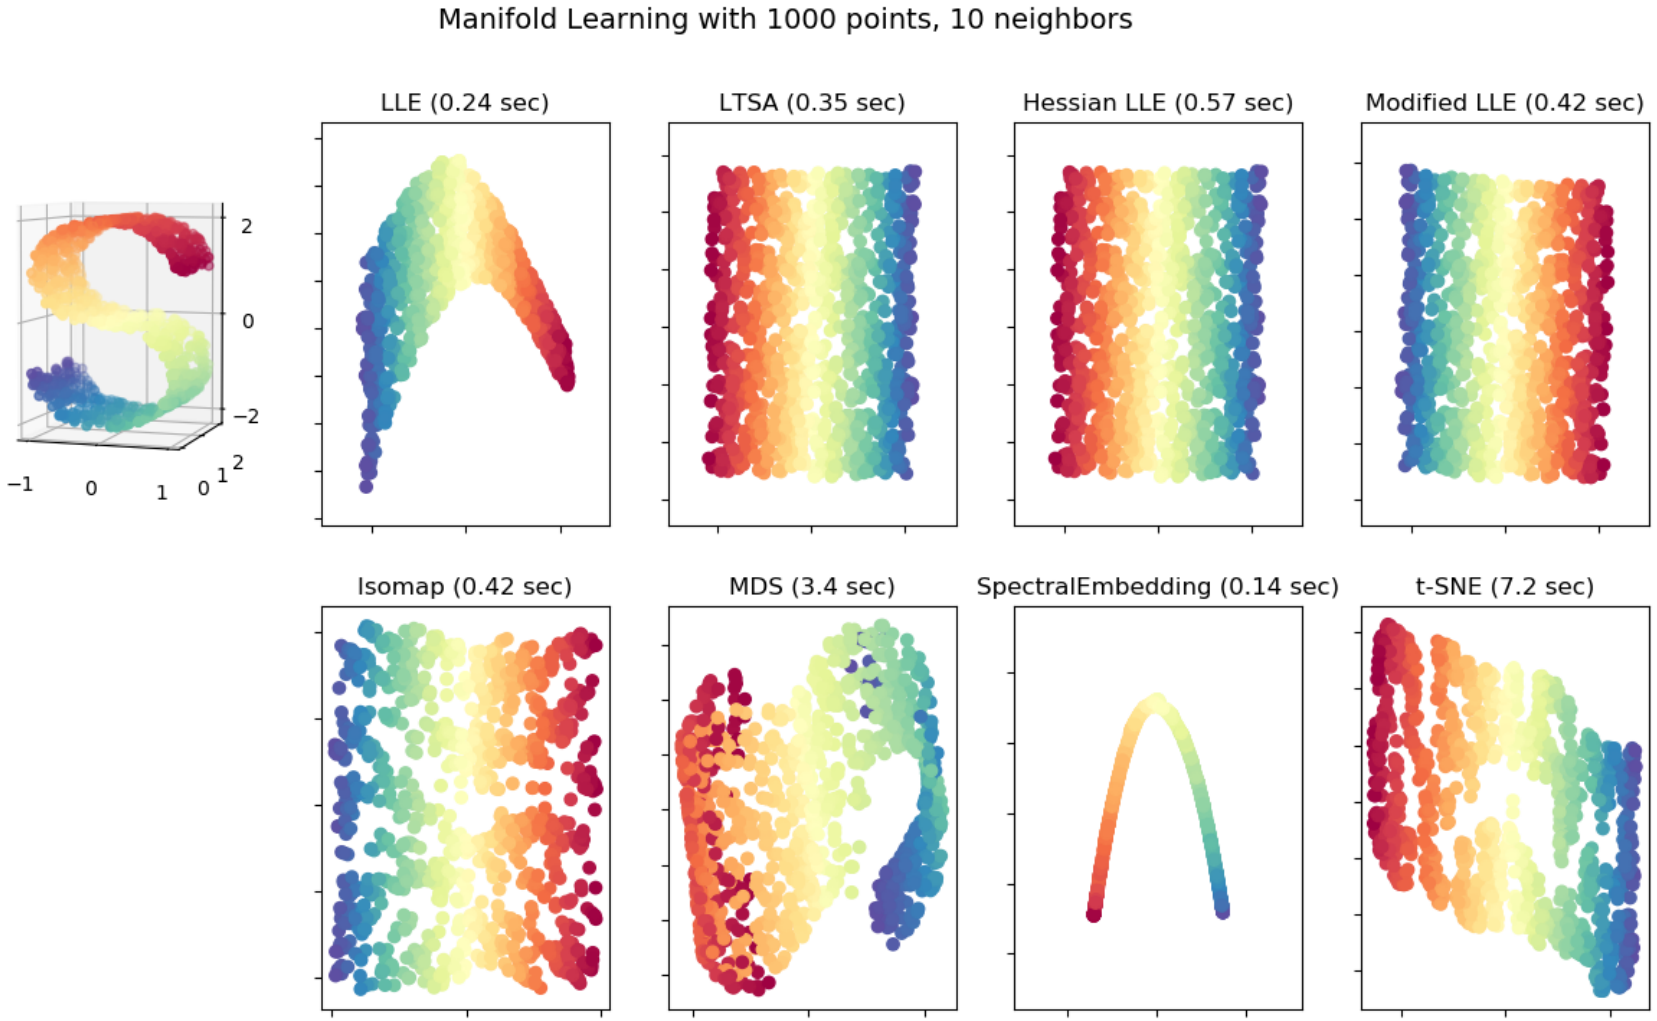

[scikit-learn.org]

Model Explanation

Model explanation

- Lime: Explaining the predictions of any ML classifier
- SHAP (SHapley Additive exPlanations)

Deep Learning / Neural Networks

Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers

Fred Hohman, *Member, IEEE,* Minsuk Kahng, *Member, IEEE,* Robert Pienta, *Member, IEEE,* and Duen Horng Chau, *Member, IEEE*

Abstract—Deep learning has recently seen rapid development and received significant attention due to its state-of-the-art performance on previously-thought hard problems. However, because of the internal complexity and nonlinear structure of deep neural networks, the **underlying decision making processes for why these models are achieving such performance are challenging** and […]

- [TVCG 2018] Web version

Interrogative Questions

WHY
- Interpretability & Explainability
- Debugging & Improving Models
- Comparing & Selecting Models
- Teaching Deep Learning Concepts

WHO
- Model Developers & Builders
- Model Users
- Non-experts

WHAT
- Computational Graph & Network Architecture
- Learned Model Parameters
- Individual Computational Units
- Neurons in High-dimensional Space
- Aggregated Information

HOW
- Node-link Diagrams for Network Architecture
- Dimensionality Reduction & Scatter Plots
- Line Charts for Temporal Metrics
- Instance-based Analysis & Exploration
- Interactive Experimentation
- Algorithms for Attribution & Feature Visualization

WHEN
- During Training
- After Training

WHERE
- Publication Venue

Web version

Further directions

- **Debugging tools**
  - Tensorboard: Visualizing Learning
  - Visdom (only supports (Py)Torch and numpy)
- **Explainables**
  - R2D3
  - Tensorboard Playground

Model visualization

- LSTM-Vis: http://lstm.seas.harvard.edu/client/index.html
- Building blocks of interpretability