Democratizing Data Science — the human in the loop

Torsten Möller
Visualization and Data Analysis
University of Vienna

Explainable Models

Torsten Möller
Visualization and Data Analysis
University of Vienna

Outline today

- What types of models are out there?
- Why explainable?
  - Model builders vs. Model users
  - societal factors
- How?
  - Experiential learning!
  - FluidExplorer vs. TreePOD
- Conclusions

Computational Science vs. Data Science

Scientific Method

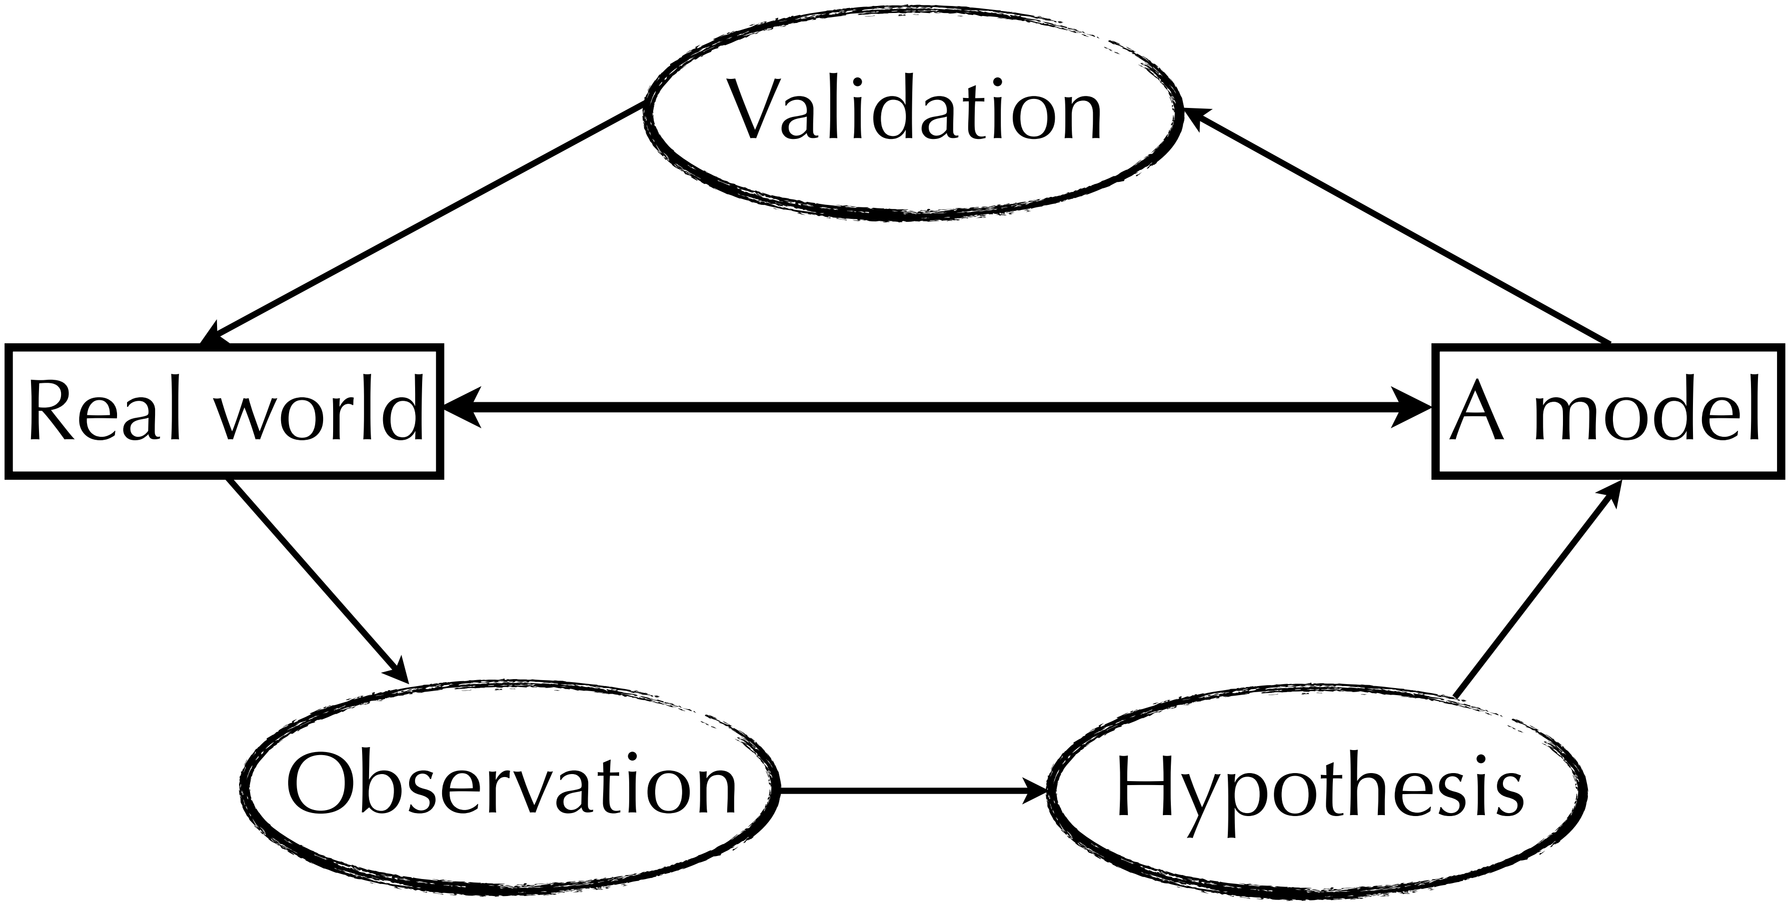

Validation ➞ Prediction

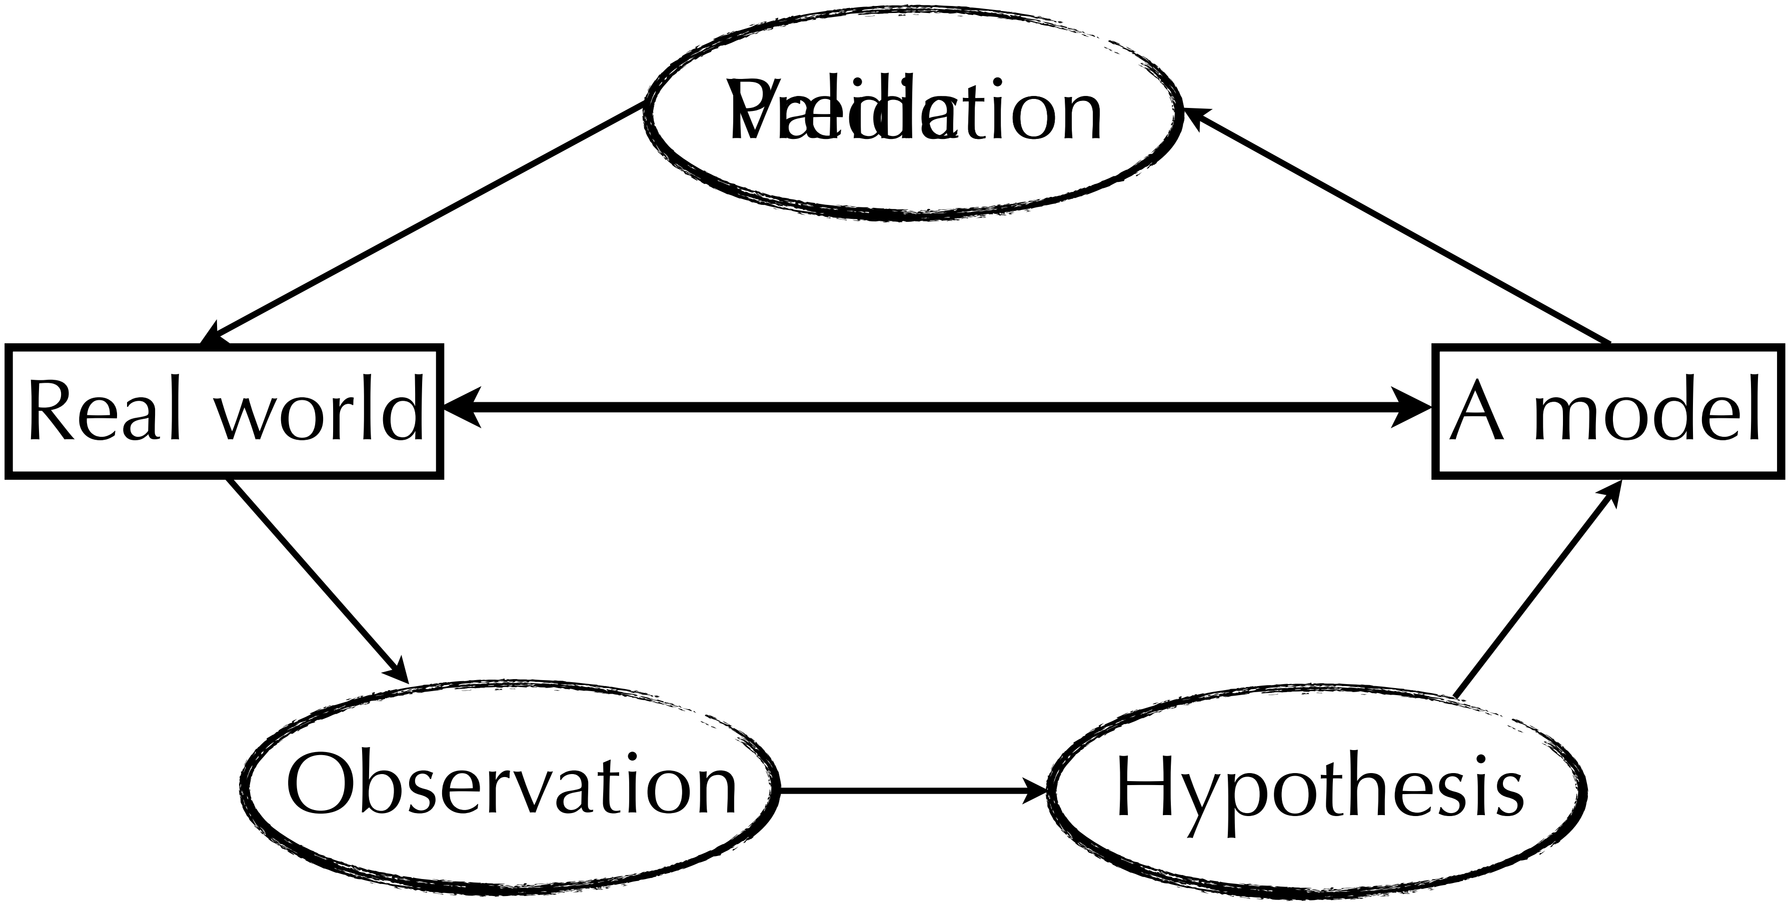

Jim Gray, “eScience -- A Transformed Scientific Method”, (2007)
https://en.wikipedia.org/wiki/File:Jim_Gray_portrait,_1999.jpg

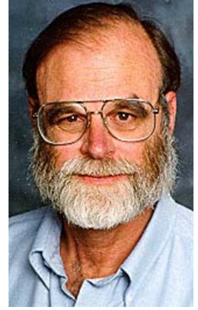

1944-2007

4 Paradigms of Science

- empirical: observe, then derive

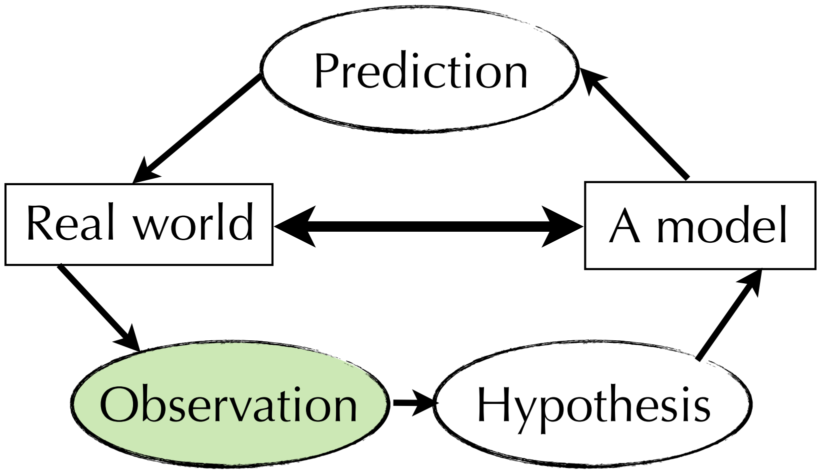

Jim Gray, “eScience -- A Transformed Scientific Method”, (2007)
https://en.wikipedia.org/wiki/File:Jim_Gray_portrait,_1999.jpg

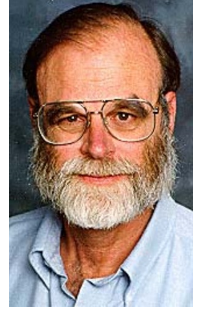

1944-2007

4 Paradigms of Science

- empirical: observe, then derive
- predictive: derive, then observe

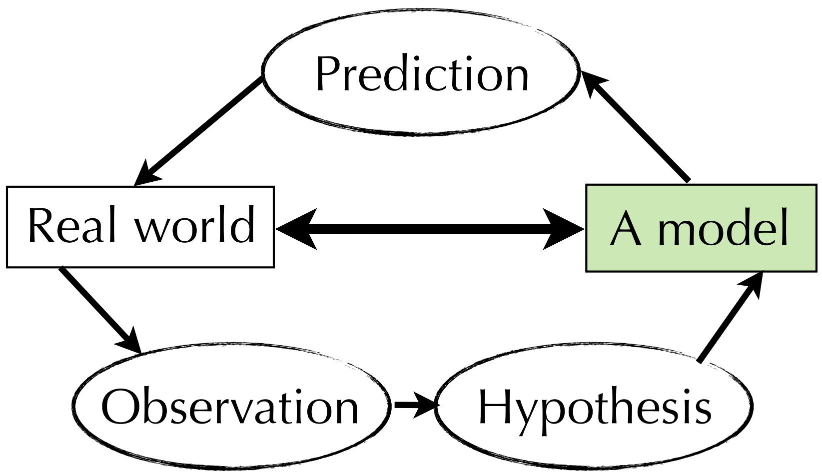

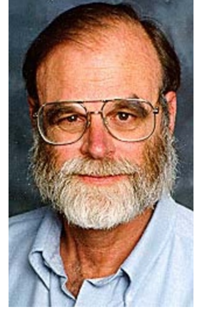

1944-2007

4 Paradigms of Science

- empirical: observe, then derive
- predictive: derive, then observe
- computational: simulate

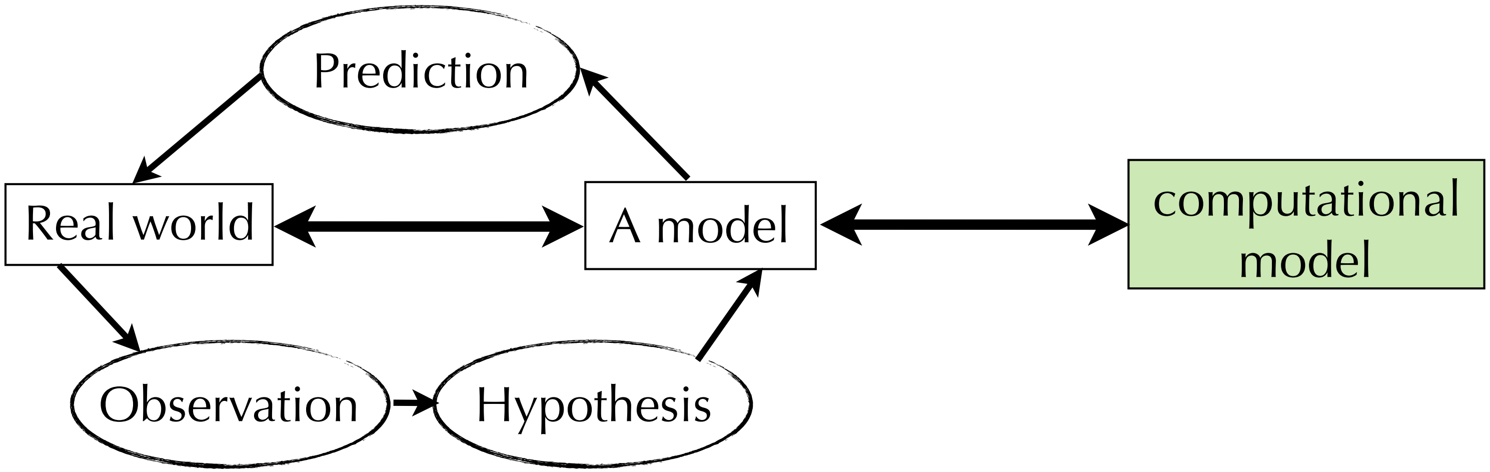

Jim Gray, “eScience -- A Transformed Scientific Method”, (2007)
https://en.wikipedia.org/wiki/File:Jim_Gray_portrait,_1999.jpg

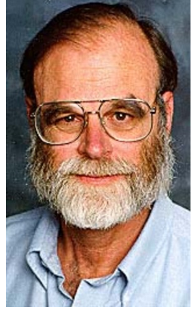

1944-2007

4 Paradigms of Science

- empirical: observe, then derive
- predictive: derive, then observe
- computational: simulate
- data-driven: measure

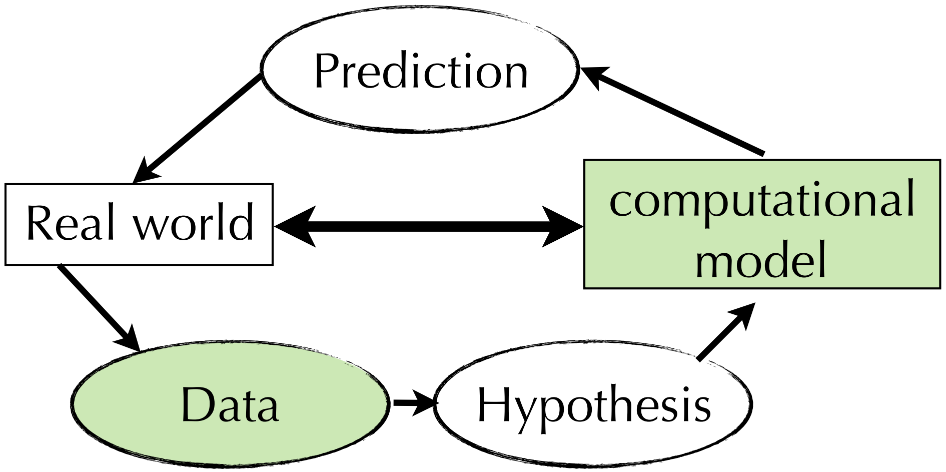

Three types of modelling

- computational: the simulation of discretized mathematical models (computational science)
- statistical: data-driven — extracting statistical models from data
- empirical: simple, often linear models

Computational Modelling

- (almost) every discipline has these models
- Examples:
  - Navier-Stokes, Maxwell, etc.
  - Population Dynamics
- computational science: experimentation through simulation of discretized models

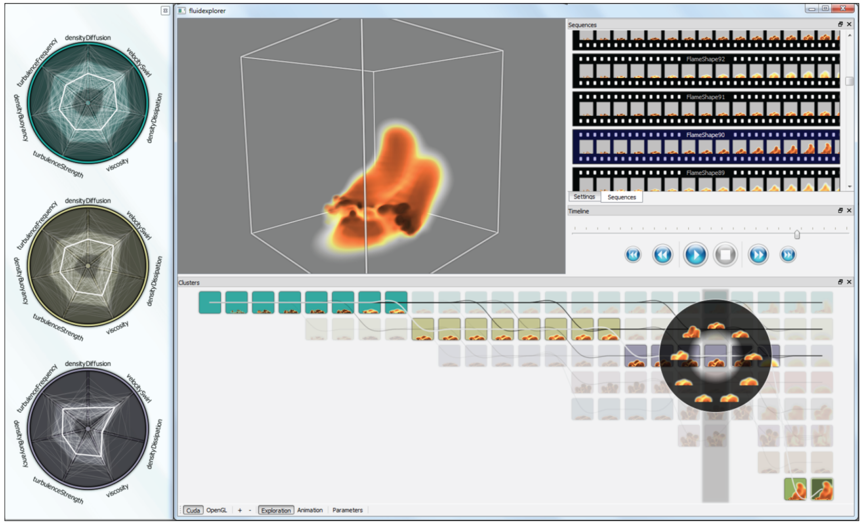

[Bruckner & Möller 2010]

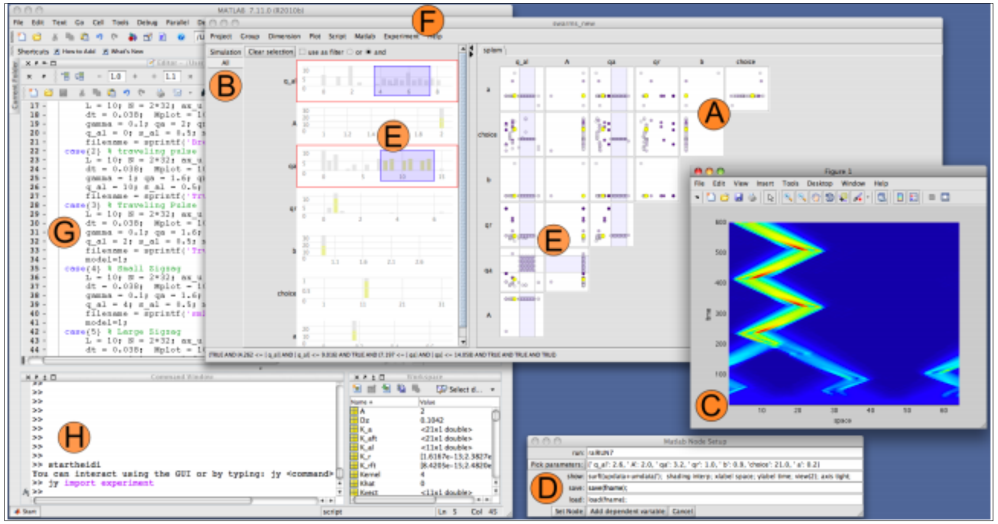

[Bergner et al. 2013]

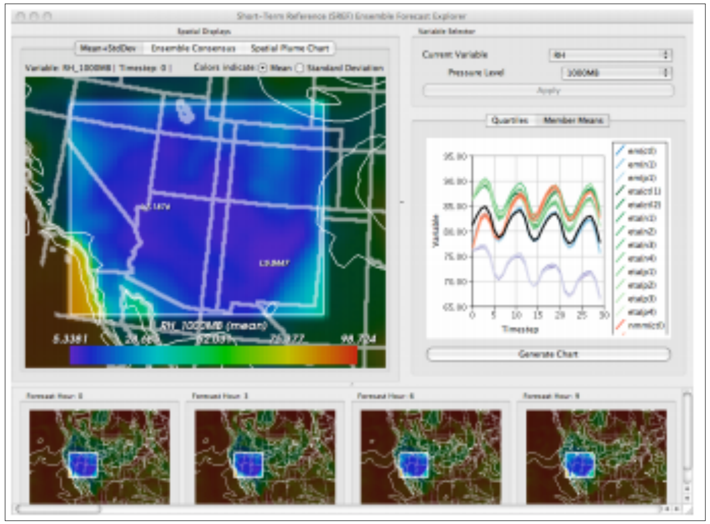

[Potter et al. 2009]

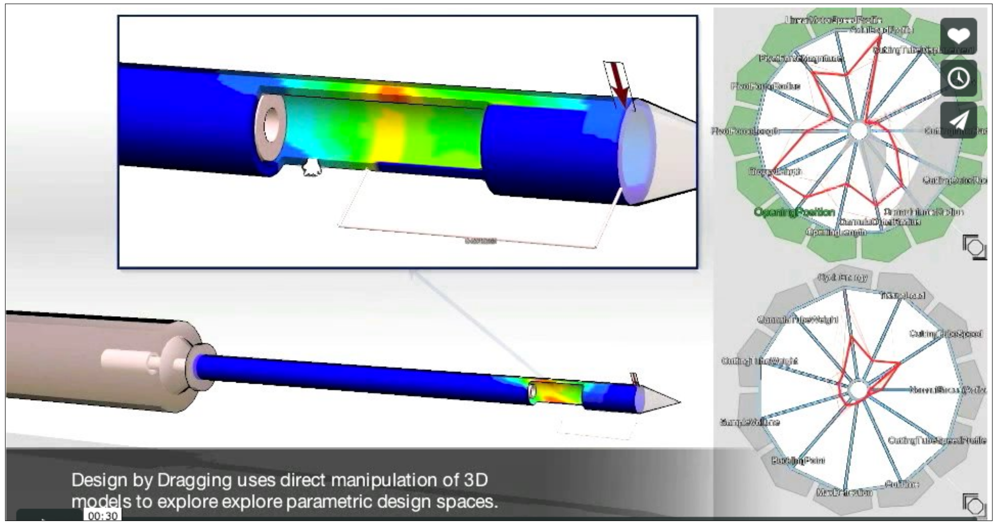

[Coffey et al. 2013]

Statistical Modeling

- “Mainstream” understanding of Data Science
- Classical (machine learning) approaches:
    - Clustering
    - Classification
    - Regression
    - (dimensionality reduction, outlier detection, etc)

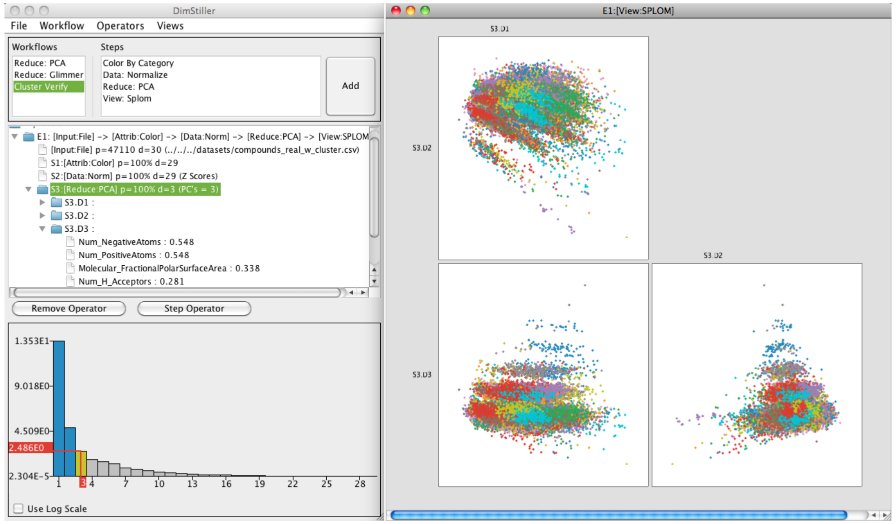

Dim reduction — [Ingram et al. 2010]

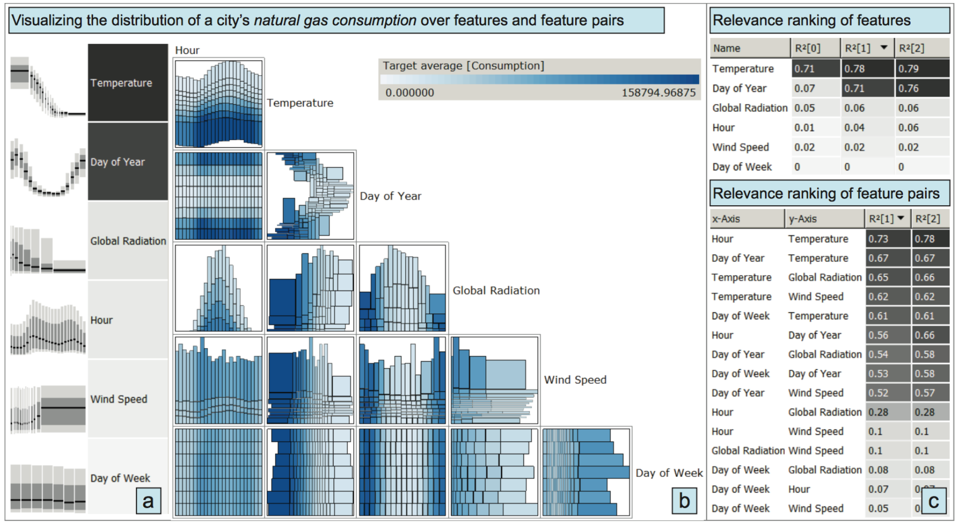

Regression — [Mühlbacher & Piringer 2013]

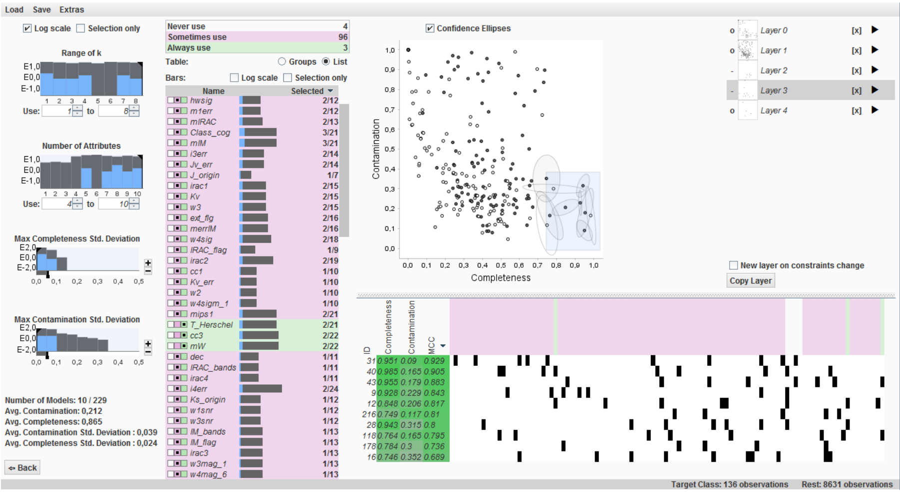

Classification — [Linhardt et al. 2019?]

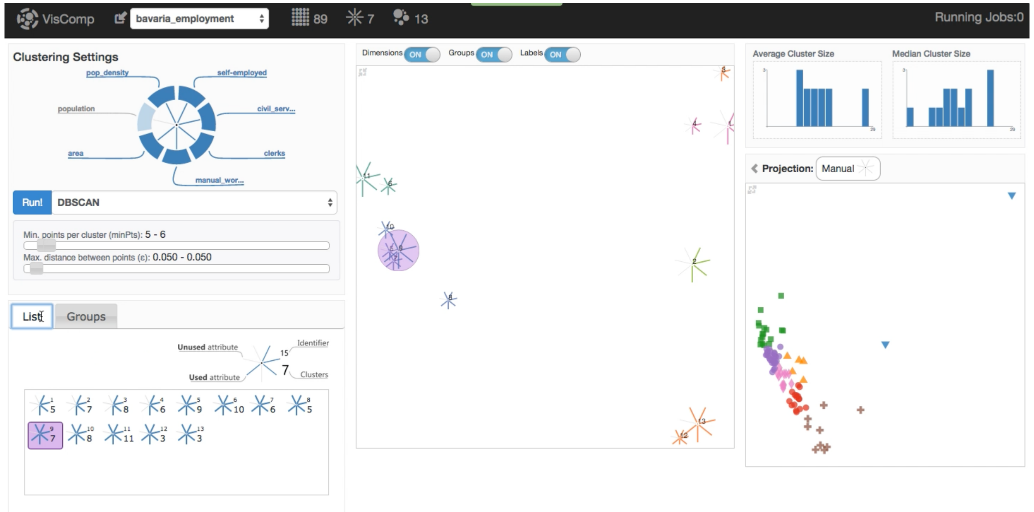

Clustering — [Sedlmair et al. 2018]

Empirical Modeling

- often no explicit modelling or only simple models, e.g.
  - linear models
  - weighted averages etc.
- examples: spreadsheets, rankings

LineUp: Gratzl et al. 2013

LineUp: Gratzl et al. 2013

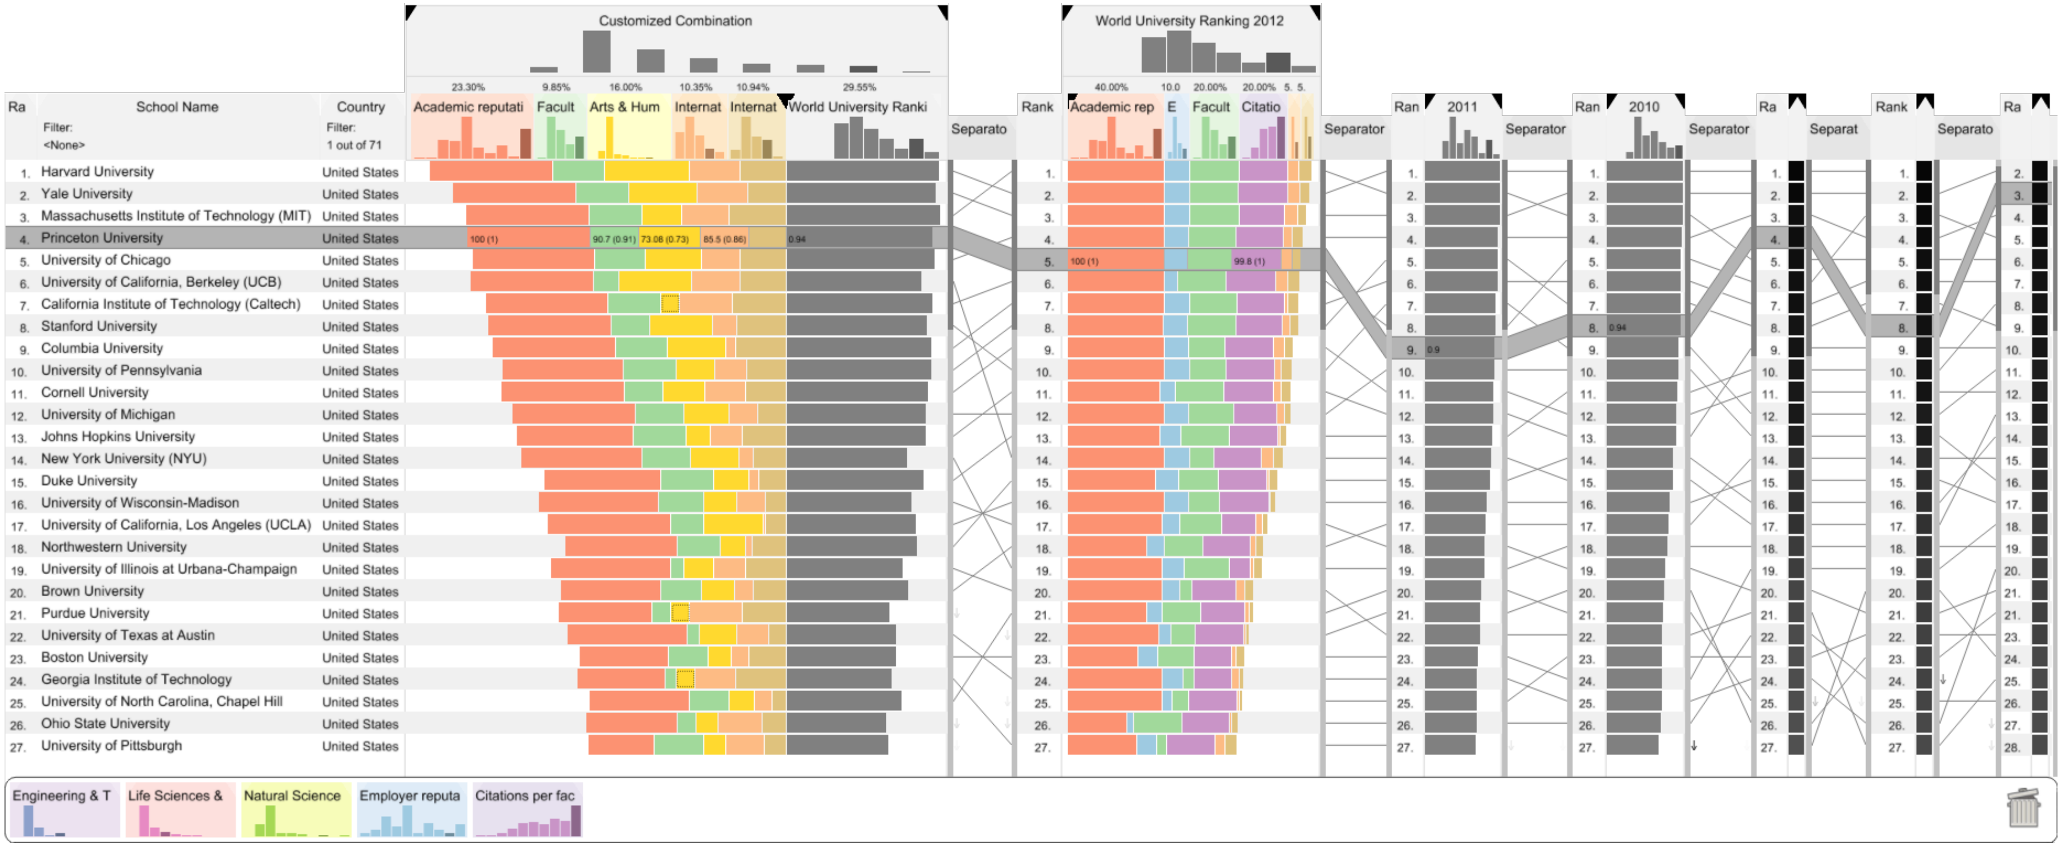

Why do we need explainable models?

Acting upon models

Building vs. Using

- building models
  - developers vs. data scientists vs. computational experts
  - hackers vs. scripters vs. application user
- using models
  - decision makers
  - domain experts
  - audience / public

Building vs. Using

- building models
  - validation
  - uncertainty
- using models
  - trust
  - tradeoffs + risks

Supporting the user

Data → Models (predictions) → Decisions

- hypothesis creation
- uncertainty / risk analysis
- sensitivity analysis / model uncertainty
- decision making / sense making

Why?: Societal factors

Ethics

- cars make decisions on who to run over and who not
- who should the company hire?
- which update from which friend should you be shown?
- which convict is more likely to re-offend?
- which news item / movie should we recommend to people?

https://www.ted.com/talks/zeynep_tufekci_machine_intelligence_makes_human_morals_more_important#t-157020

Laws

- EU's General Data Protection Regulation:
- incl Article 22: Automated individual decision-making, including profiling
- prohibits any "decision based solely on automated processing, including profiling" which "significantly affects" a data subject
- **Discrimination**: Paragraph 71 of the recitals (the preamble to the GDPR, which explains the rationale behind it but is not itself law) explicitly requires data controllers to "implement appropriate technical and organizational measures" that "prevents, inter alia, discriminatory effects" on the basis of processing sensitive data
- **Right to explanation**: Articles 13 and 14 state that, when profiling takes place, a data subject has the right to "meaningful information about the logic involved."

Goodman, B. & Flaxman, S.
European Union regulations on algorithmic decision-making and a "right to explanation"
AI Magazine, **2017**

Outline today

- What types of models are out there?
- Why explainable?
    - Model builders vs. Model users
    - societal factors
- How?
    - Experiential learning!
    - FluidExplorer vs. TreePOD
- Conclusions

How?

Theorem [Berner-G-Jentzen (2018)], very special case

Let $\varphi(x) = \min\{\max\{\max(x_i - K_i), 0\}, R\}$ or $\varphi(x) = \min\{\max\{\sum_{i=1}^d x_i - K, 0\}, R\}$ (or any typical option). Then for all $\epsilon > 0$ there is $\Phi_\epsilon \in \mathcal{H}^{ReLU}_{(N_0,\dots,N_L)}$ with $\mathrm{size}(\Phi_\epsilon) = \mathcal{O}(\epsilon^{-2})$ and

$$\frac{1}{(b-a)^{d/2}} \left(\int_{[a,b]^d} \left| u(T,x) - R_\sigma(\Phi_\epsilon)(x) \right|^2 dx \right)^{1/2} \leq \epsilon.$$

Such networks can be found by solving the ERM problem with $m \sim \epsilon^{-4}$ samples. The implicit constants depend at most polynomially on the dimension $d = N_0$!

From Philip Grohs

How?

```
def CompactCNN(input_shape, nb_conv, nb_filters, n_mels, normalize, nb_hidden, dense_units,
               output_shape, activation, dropout, multiple_segments=False, graph_model=False, inpu
t_tensor=None):

    melgram_input = Input(shape=input_shape)

    if n_mels >= 256:
        poolings = [(2, 4), (4, 4), (4, 5), (2, 4), (4, 4)]
    elif n_mels >= 128:
        poolings = [(2, 4), (4, 4), (2, 5), (2, 4), (4, 4)]
    elif n_mels >= 96:
        poolings = [(2, 4), (3, 4), (2, 5), (2, 4), (4, 4)]
    elif n_mels >= 72:
        poolings = [(2, 4), (3, 4), (2, 5), (2, 4), (3, 4)]
    elif n_mels >= 64:
        poolings = [(2, 4), (2, 4), (2, 5), (2, 4), (4, 4)]

    # Determine input axis
    if keras.backend.image_dim_ordering() == 'th':
        channel_axis = 1
        freq_axis = 2
        time_axis = 3
    else:
        channel_axis = 3
        freq_axis = 1
        time_axis = 2

    # Input block
    #x = BatchNormalization(axis=time_axis, name='bn_0_freq')(melgram_input)

    if normalize == 'batch':
        x = BatchNormalization(axis=freq_axis, name='bn_0_freq')(melgram_input)
    elif normalize in ('data_sample', 'time', 'freq', 'channel'):
        x = Normalization2D(normalize, name='nomalization')(melgram_input)
    elif normalize in ('no', 'False'):
        x = melgram_input

    # Conv block 1
    x = Convolution2D(nb_filters[0], (3, 3), padding='same')(x)
    x = BatchNormalization(axis=channel_axis, name='bn1')(x)
    x = ELU()(x)
    x = MaxPooling2D(pool_size=poolings[0], name='pool1')(x)
```

Alex Schindler

How — our approach

https://youtu.be/5d71xhEbjDg

FluidExplorer
Fluid animation

Special effects

- Fluid simulation is heavily used in the motion picture industry
- Most common animation packages include solvers or offer add-ons
- Problem: Difficult to control for visual effects artists

Special effects (2)

- **Tens of parameters**
- **Hard to predict results**
- **Time-consuming trial & error**

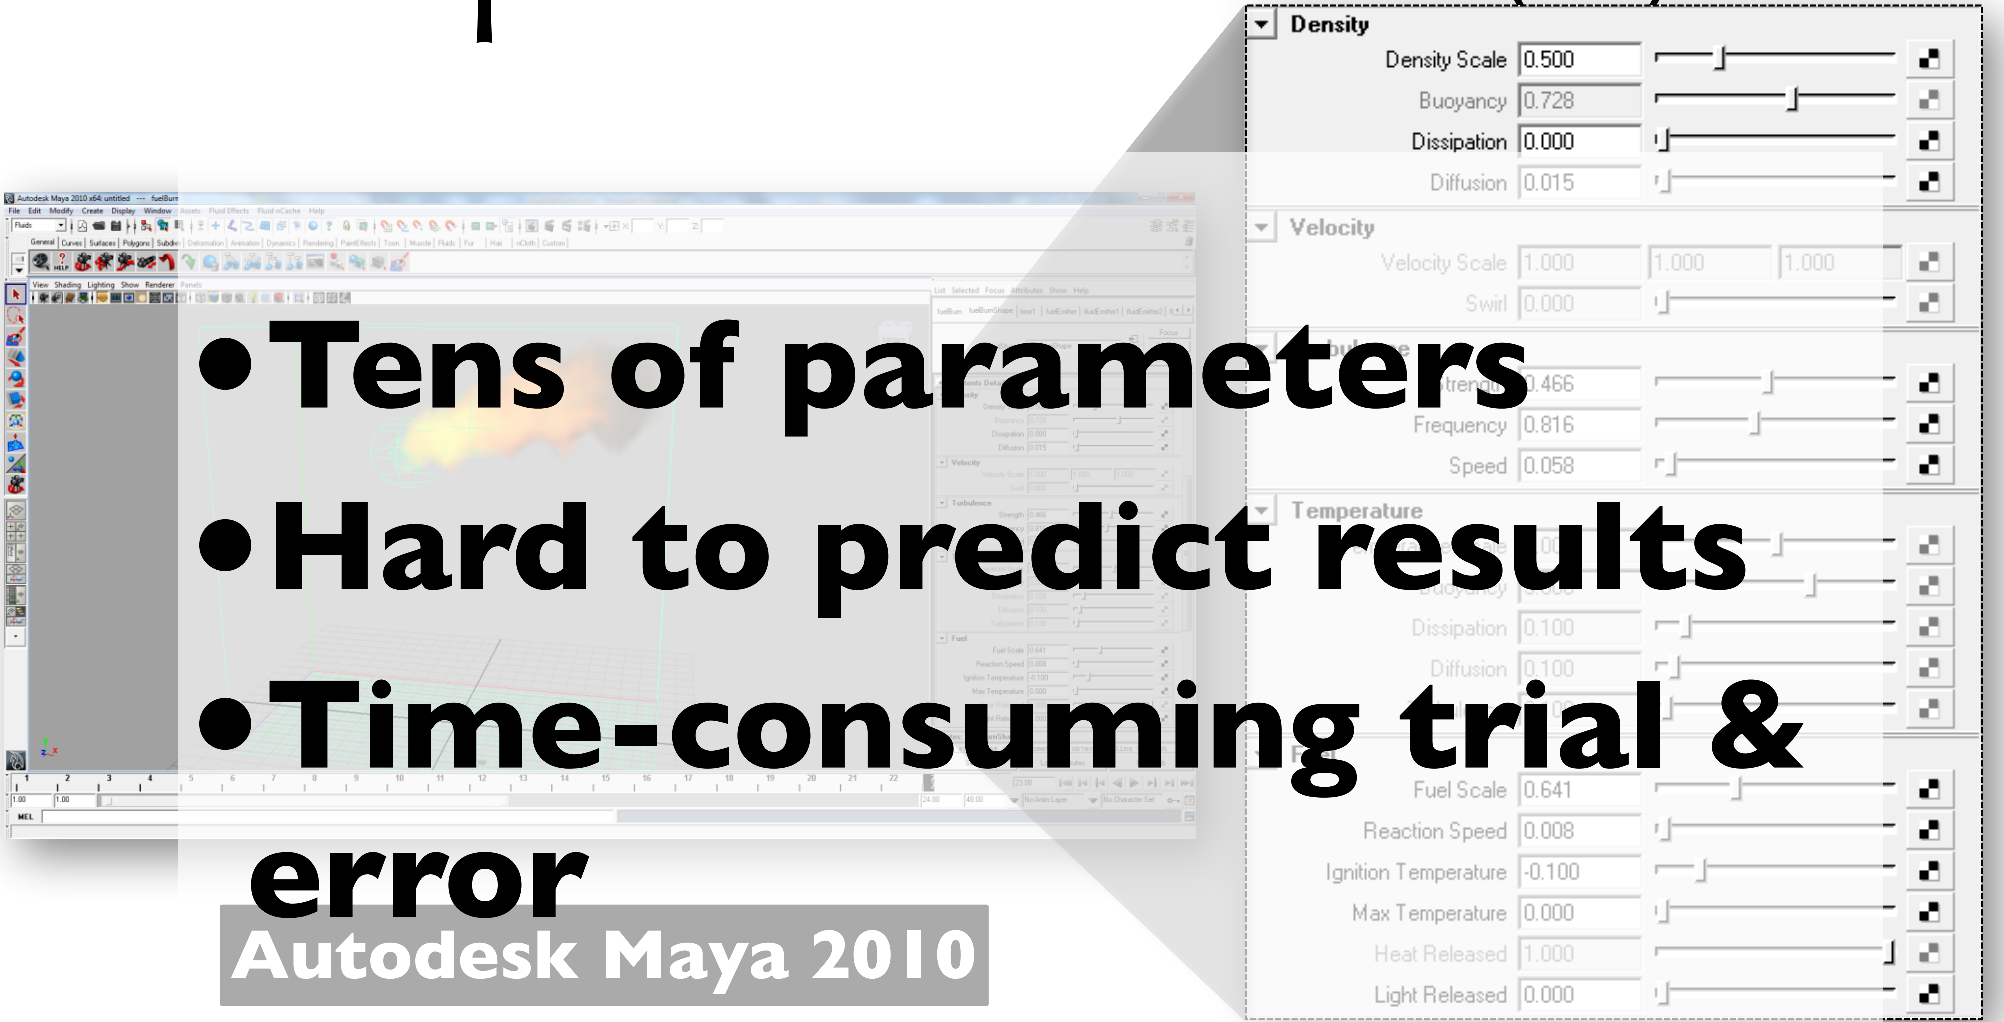

Autodesk Maya 2010

Overview

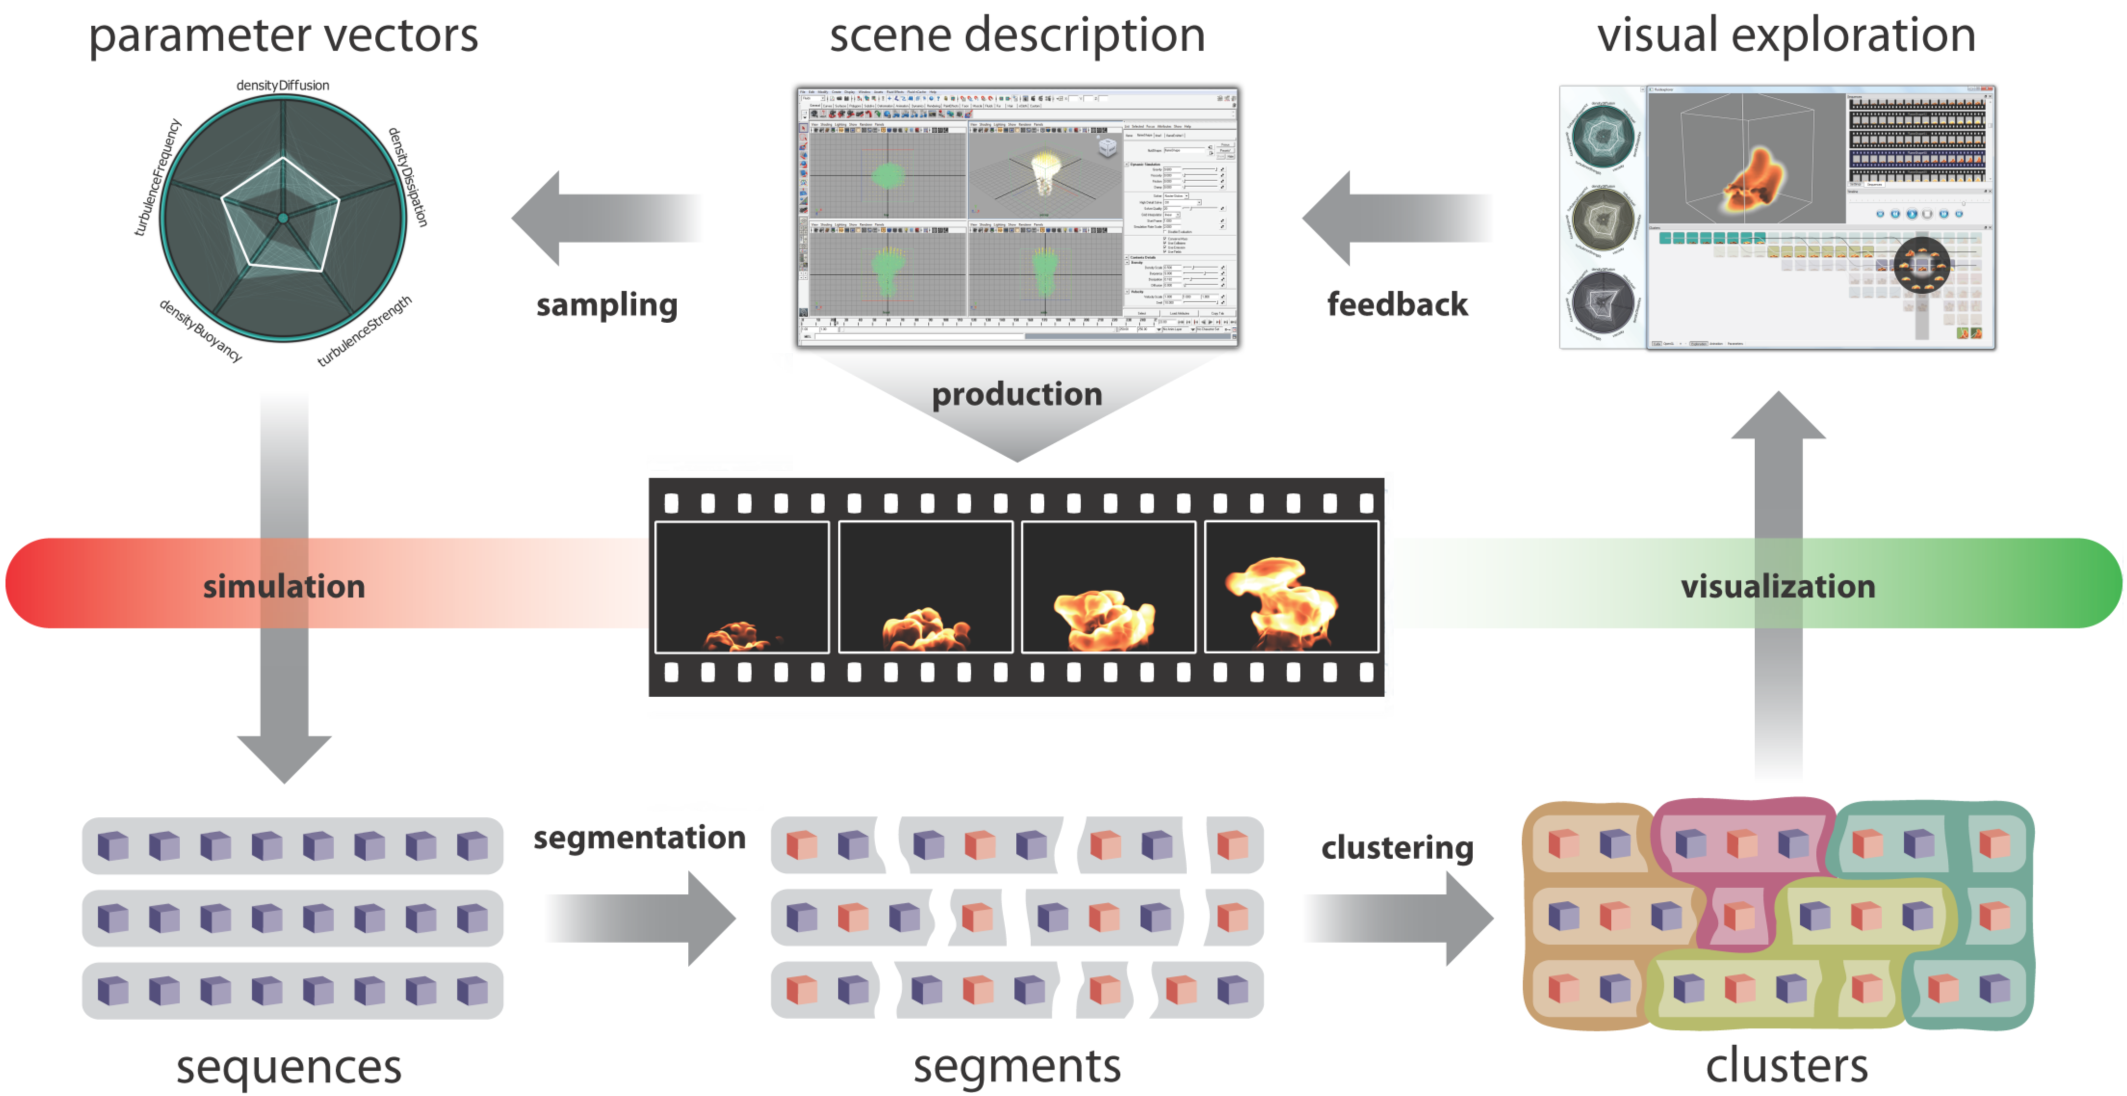
parameter vectors
densityDiffusion
densityDissipation
turbulenceStrength
densityBuoyancy
turbulenceFrequency
scene description
visual exploration
sampling
feedback
production
simulation
visualization
segmentation
clustering
sequences
segments
clusters

Visualization

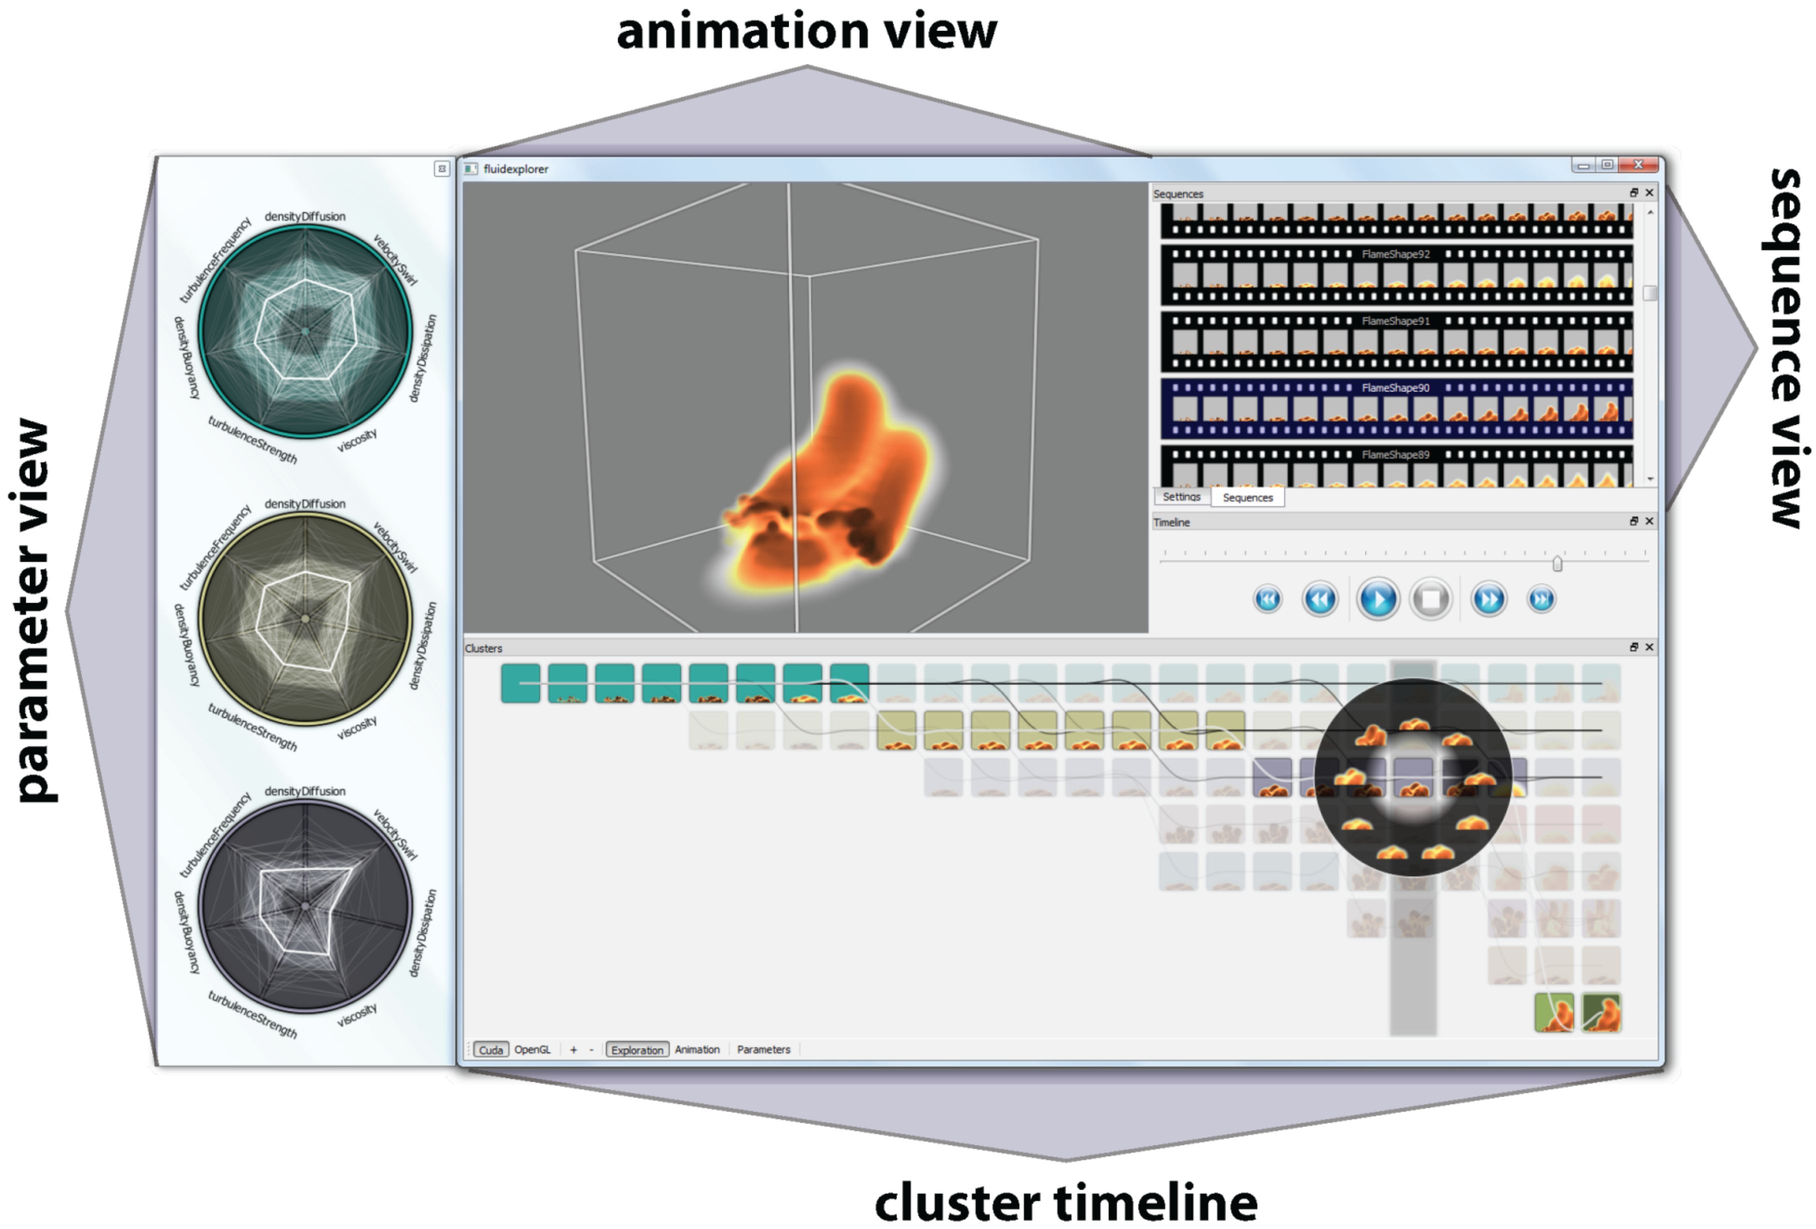

TreePOD — decision tree analysis

Decision trees are important for classification in many fields

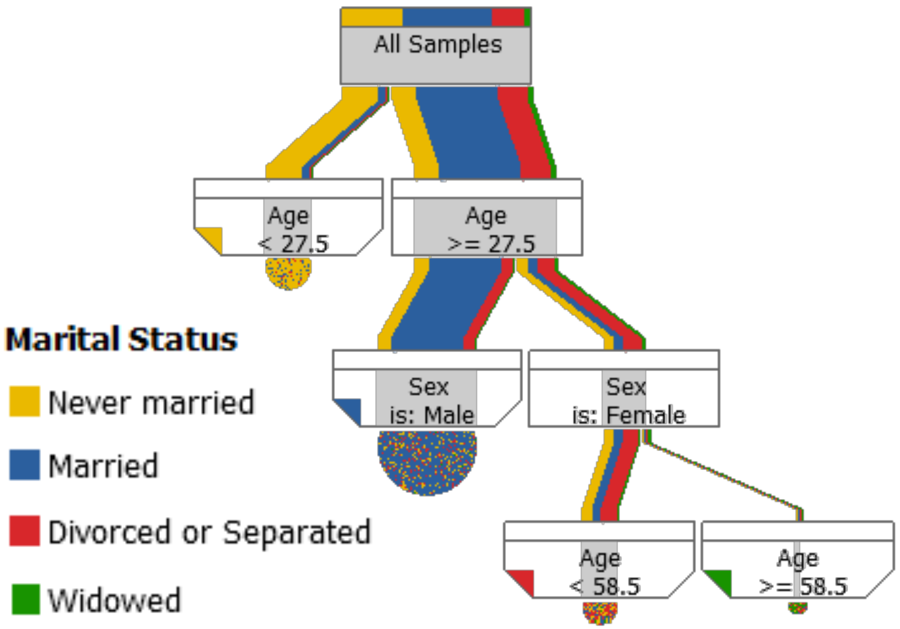

UCI Lab Census 1994 Dataset

Explain classes
by decision rules on features

Trained for data
= supervised learning

Understandable structure
for analysis and prediction

Building decision trees involves multiple trade-offs

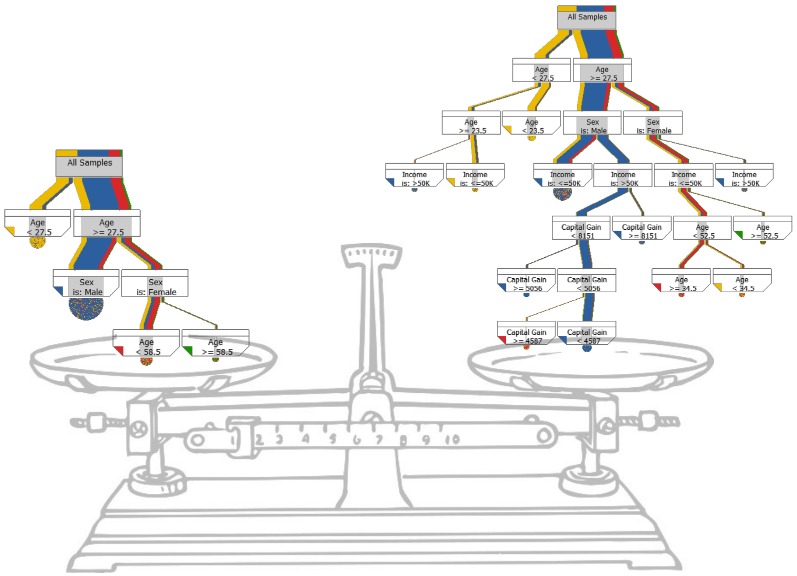

underfitting vs. overfitting

„bias-variance“ trade-off

Building decision trees involves multiple trade-offs

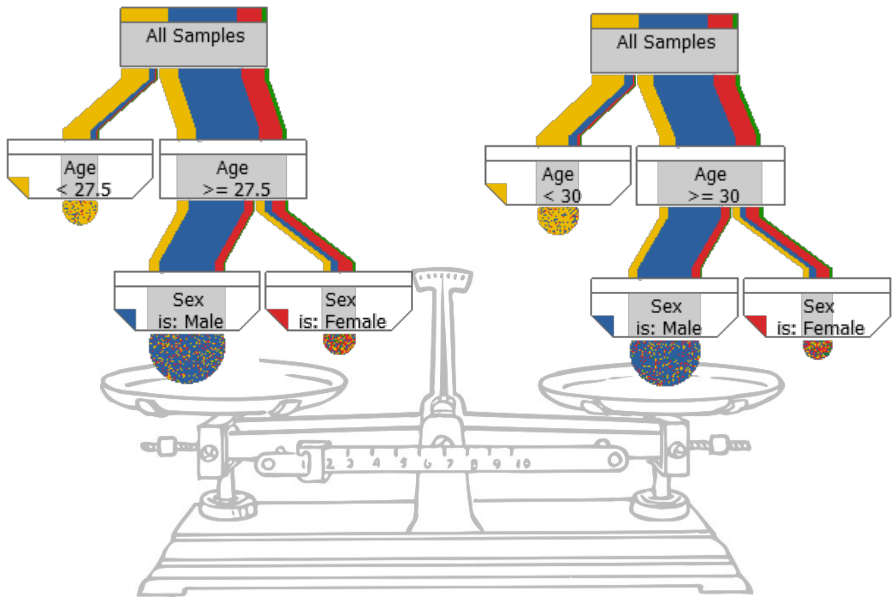

underfitting vs. overfitting
„bias-variance“ trade-off

accuracy vs. interpretability
e.g., nice decision borders

Building decision trees involves multiple trade-offs

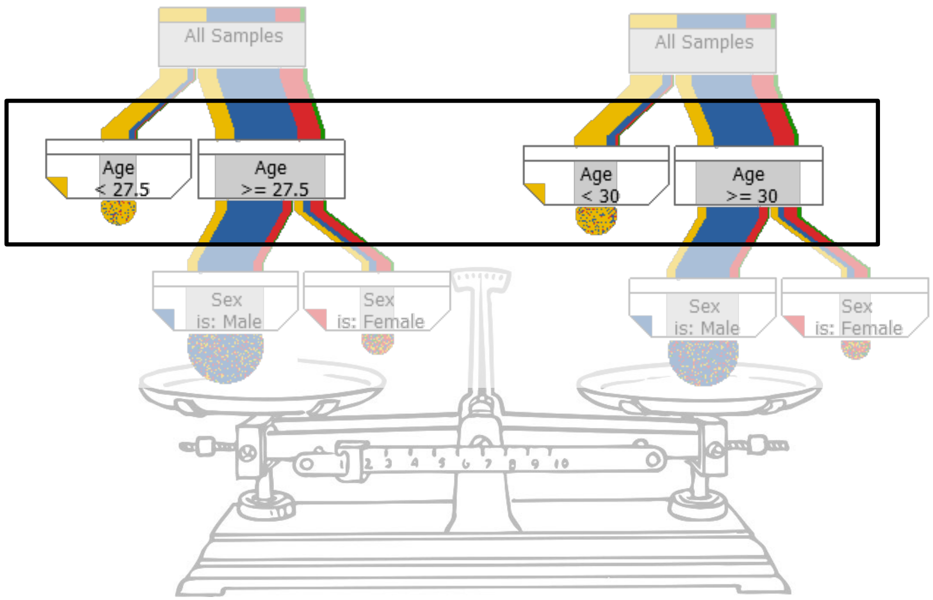

underfitting vs. overfitting
„bias-variance“ trade-off

accuracy vs. interpretability
e.g., nice decision borders

Building decision trees involves multiple trade-offs

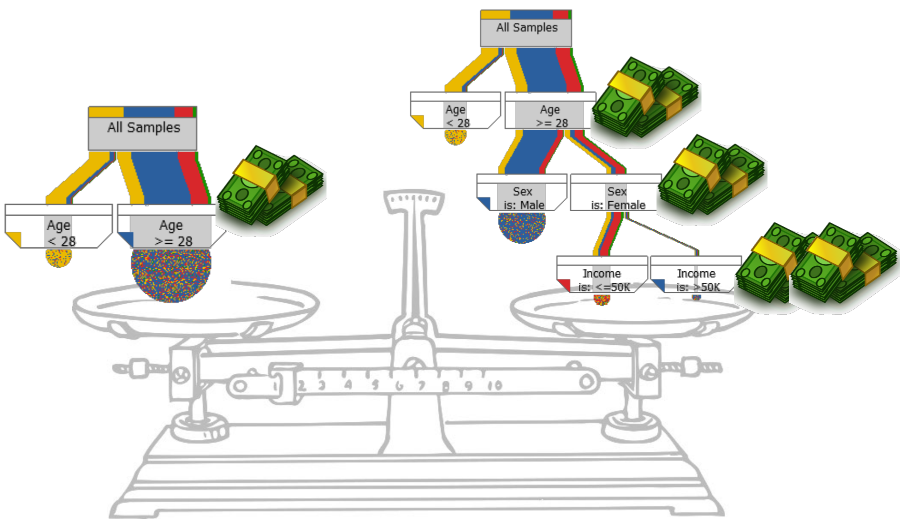

underfitting vs. overfitting
„bias-variance“ trade-off

accuracy vs. interpretability
e.g., nice decision borders

additional constraints
e.g. feature acquisition costs

Problem: Finding the tree representing the best trade-off

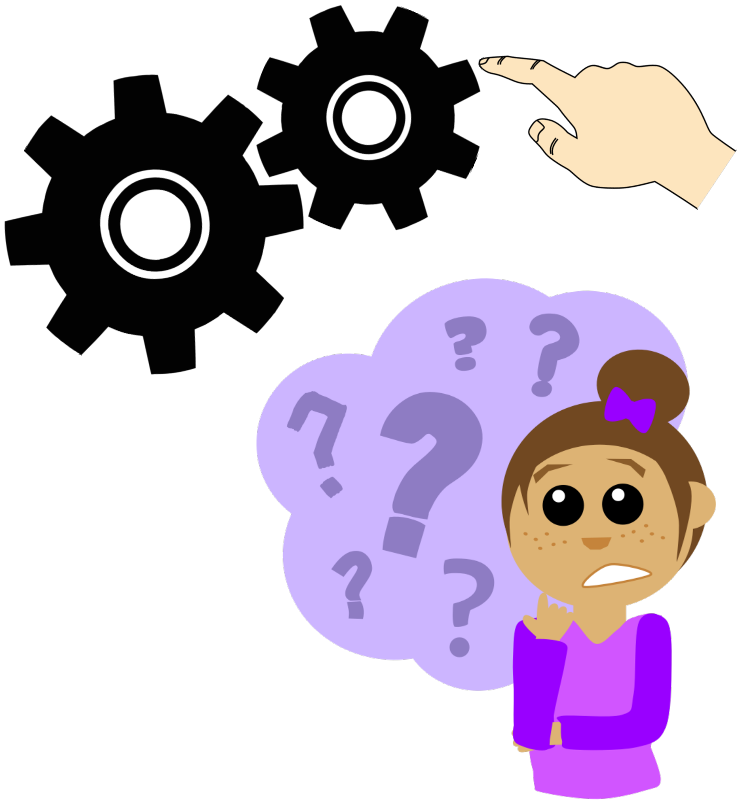

Hard to automate

relies on qualitative judgements

In practice: trial-and-error

inefficient, low confidence

Domain experts

$\neq$ statistical experts

Overview of TreePOD

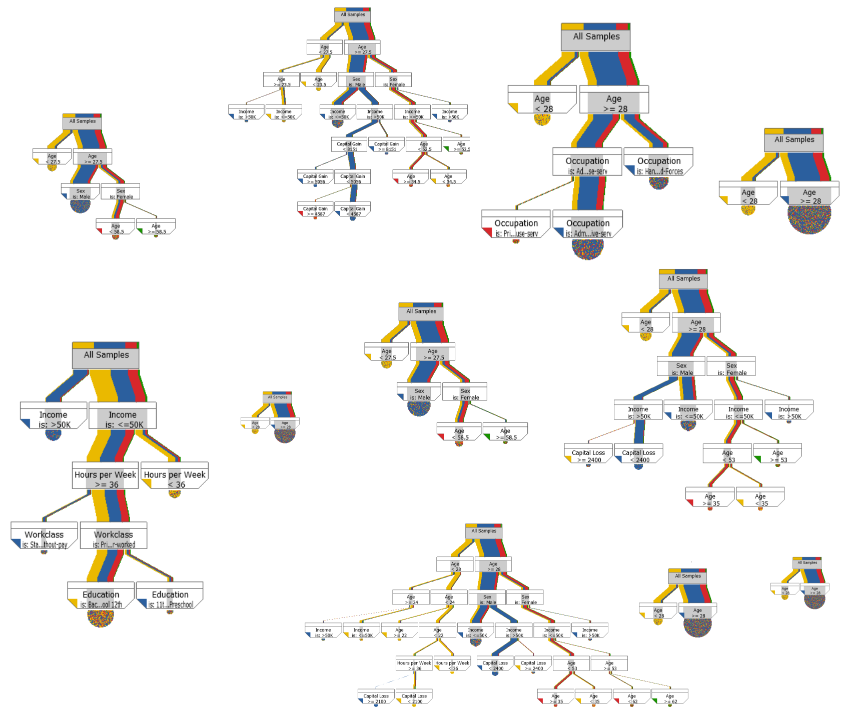

1) Create diverse tree candidates

Overview of TreePOD

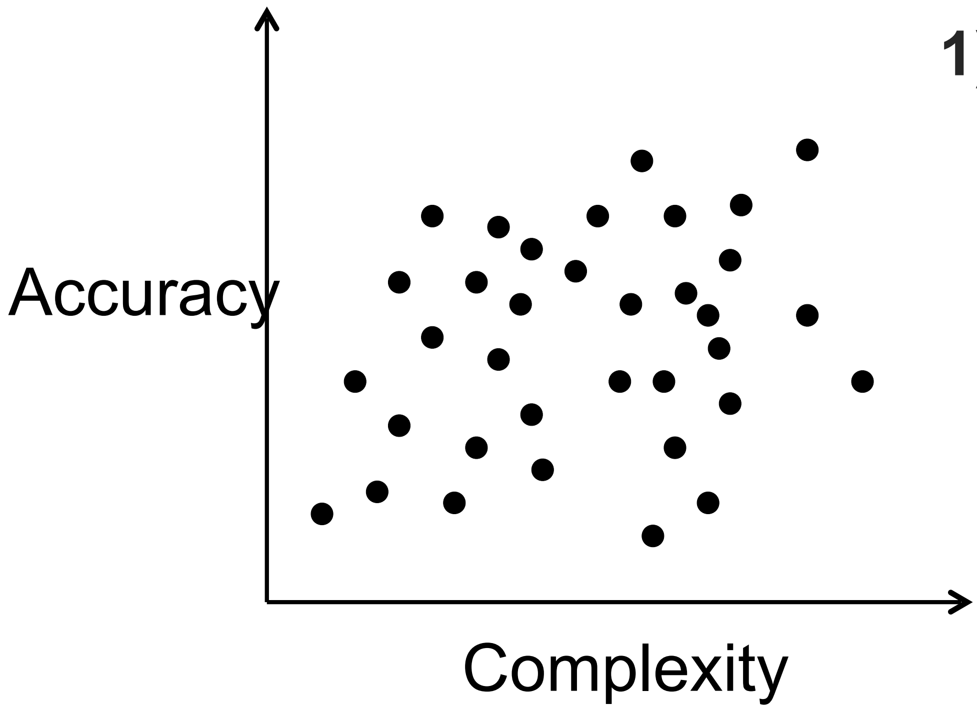

1) Create diverse tree candidates
global overview of what is possible

Overview of TreePOD

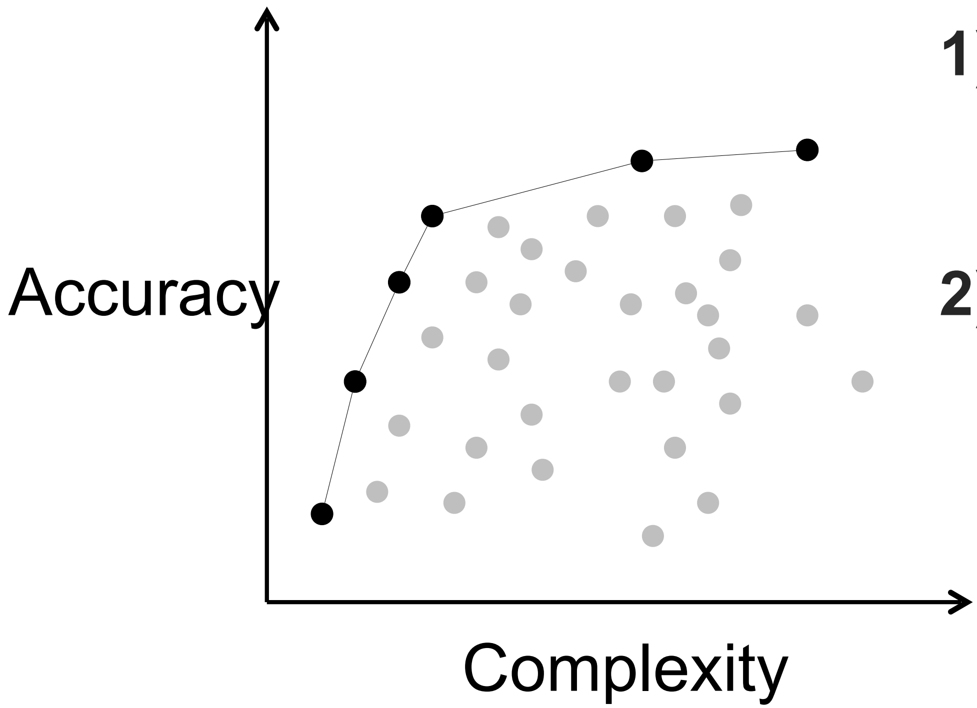

1) Create diverse tree candidates
global overview of what is possible

2) Guide selection from candidates
by focusing on good trade-offs

Creating diverse candidates by sampling algorithm parameters

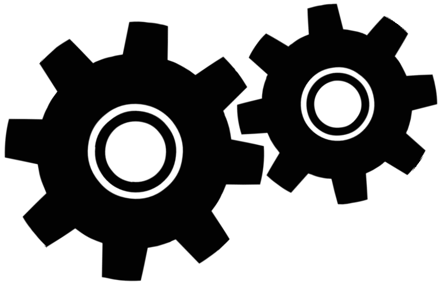

Parameters:

Feature Set
Termination Critera
Pruning method
...

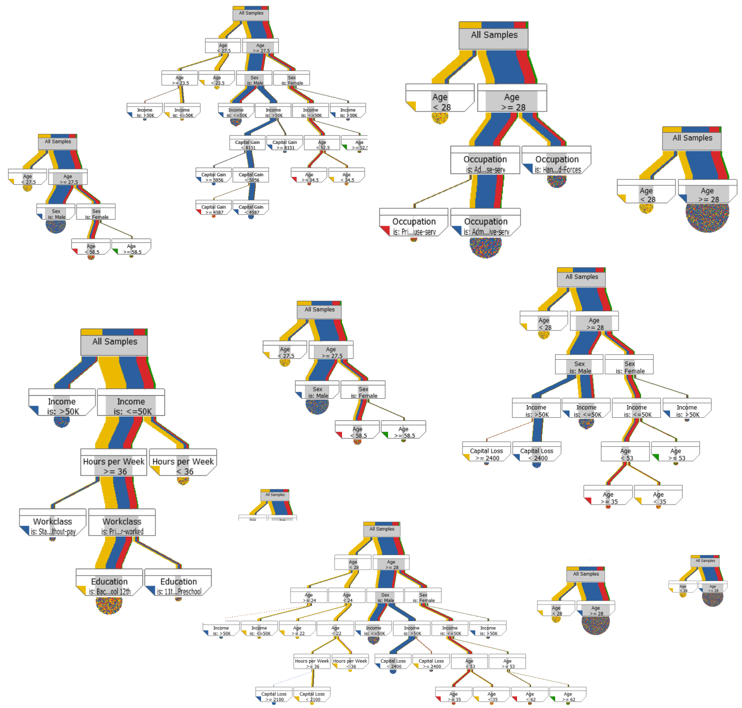

Guided visual exploration of candidate trees

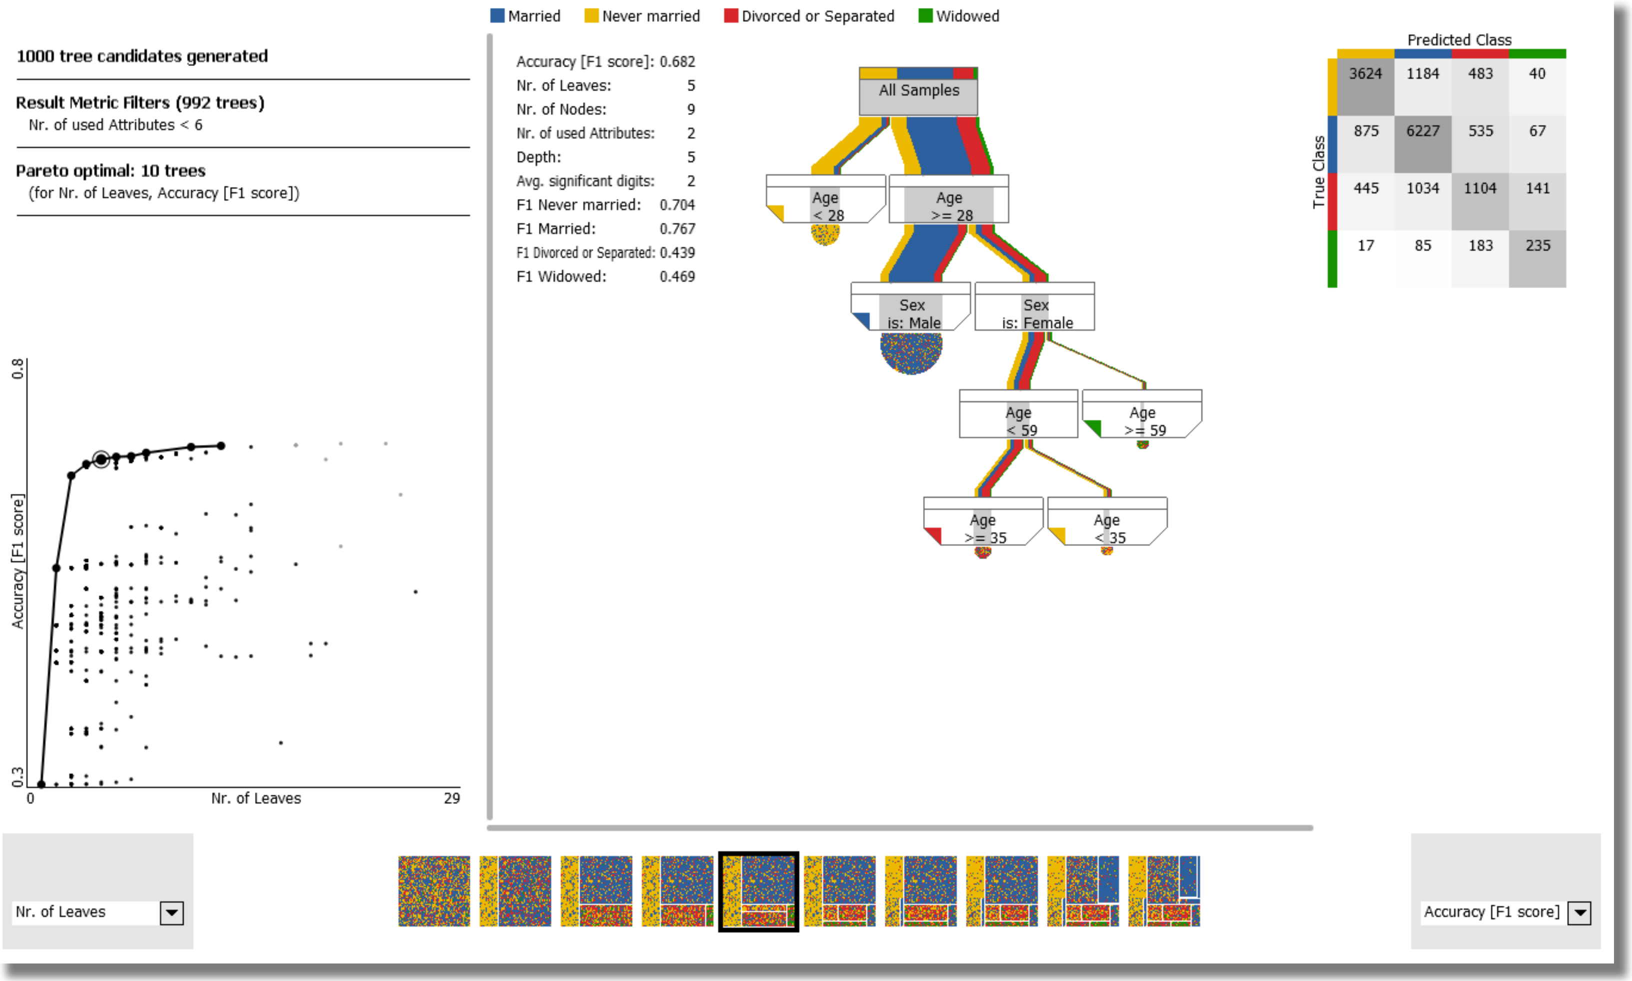

Conclusions

- Three types of modelling:
  - Through first principles
  - Through data
  - Empirical
- Why explainable?
  - improve algorithms
  - trust
  - bridge the model builder / model usage gap
  - ethics and law
- How?
  - characterization of input-output relationships OR parameter tuning
  - we are really good in learning by trial-and-error

Acknowledgments

Steven Bergner
SFU

Tom Torsney-Weir
U of Vienna

Hamid Younesy
SFU

Lorenz Linhardt
ETH Zurich

Stefan Bruckner
U of Bergen

Tamara Munzner
UBC

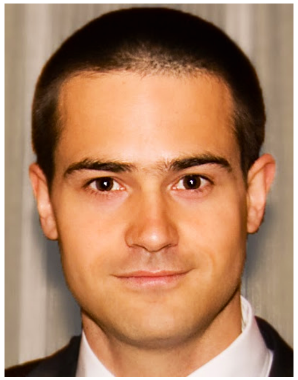

Harald Piringer
VRVis

Thomas Mühlbacher
VRVis

Michael Sedlmair
U of Stuttgart

References

- Visual Parameter Space Analysis: A Conceptual Framework. Michael Sedlmair, Christoph Heinzl, Stefan Bruckner, Harald Piringer, Torsten Möller, IEEE Transactions on Visualization and Computer Graphics 20(12):2161-2170, 2014.
- eScience -- A Transformed Scientific Method. Jim Gray, (2007), in "The Fourth Paradigm: Data-Intensive Scientific Discovery", 2009.
- TreePOD: Sensitivity-Aware Selection of Pareto-Optimal Decision Trees. Thomas Mühlbacher, Lorenz Linhardt, Torsten Möller, Harald Piringer, IEEE Transactions on Visualization and Computer Graphics, vol. 24, no. 1, pp. 174–183, 2018.
- Result-Driven Exploration of Simulation Parameter Spaces for Visual Effects Design. Stefan Bruckner, Torsten Möller, IEEE Transactions on Visualization and Computer, vol. 16, no. 6, pp. 1467–1475, Oct. 2010.

Questions?

http://vda.cs.univie.ac.at

torsten.moeller@univie.ac.at