(Deep) Learning and Simulation approaches

Torsten Möller
Visualization and Data Analysis
University of Vienna

Explainable (Deep) Learning and Simulation approaches

Torsten Möller
Visualization and Data Analysis
University of Vienna

Outline today

- Why explainable?
  - the promise of data science
  - societal factors
- How?
  - a process model for simulations
  - (machine) learning process

Visual Data Science

Overview

- Data Science is all about modelling
- The three types of modelling
  - Computational modelling
  - Statistical modelling
  - Empirical modelling
- Challenges of Visual Data Science
- Conclusions

Vasant Dhar, “Data Science and Prediction”, (2013)

What is data science?

- Dhar 2013: “Data Science is the study of the generalizable extraction of knowledge from data.”

What is data science?

- Dhar 2013: “Data Science is the study of the generalizable extraction of knowledge from data.”
- Data Science is the study of exploration, abstraction, and communication of complex systems through models from data.

What is data science?

- Dhar 2013: “Data Science is the study **of the generalizable extraction of knowledge** from data.”
- Data Science is the study of **exploration, abstraction, and communication** of complex systems through models from data.

What is data science?

- Dhar 2013: “Data Science is the study **of the generalizable extraction of knowledge** from data.”
- Data Science is the study of exploration, abstraction, and communication **of complex systems** through models from data.

What is data science?

- Dhar 2013: "Data Science is the study **of the generalizable extraction of knowledge** from data."
- Data Science is the study of exploration, abstraction, and communication of complex systems **through models** from data.

http://simplystatistics.org/2013/12/12/the-key-word-in-data-science-is-not-data-it-is-science/

Data Science

- Jeff Leek: “The key word in ‘Data Science’ is not Data, it is Science”

  > “The issue is that the hype around big data/ data science will flame out (it already is) if data science is only about "data" and not about "science". The long term impact of data science will be measured by the scientific questions we can answer with the data.”

Overview

- Data Science is all about modelling
- The three types of modelling
  - Computational modelling
  - Statistical modelling
  - Empirical modelling
- Challenges of Visual Data Science
- Conclusions

Scientific Method

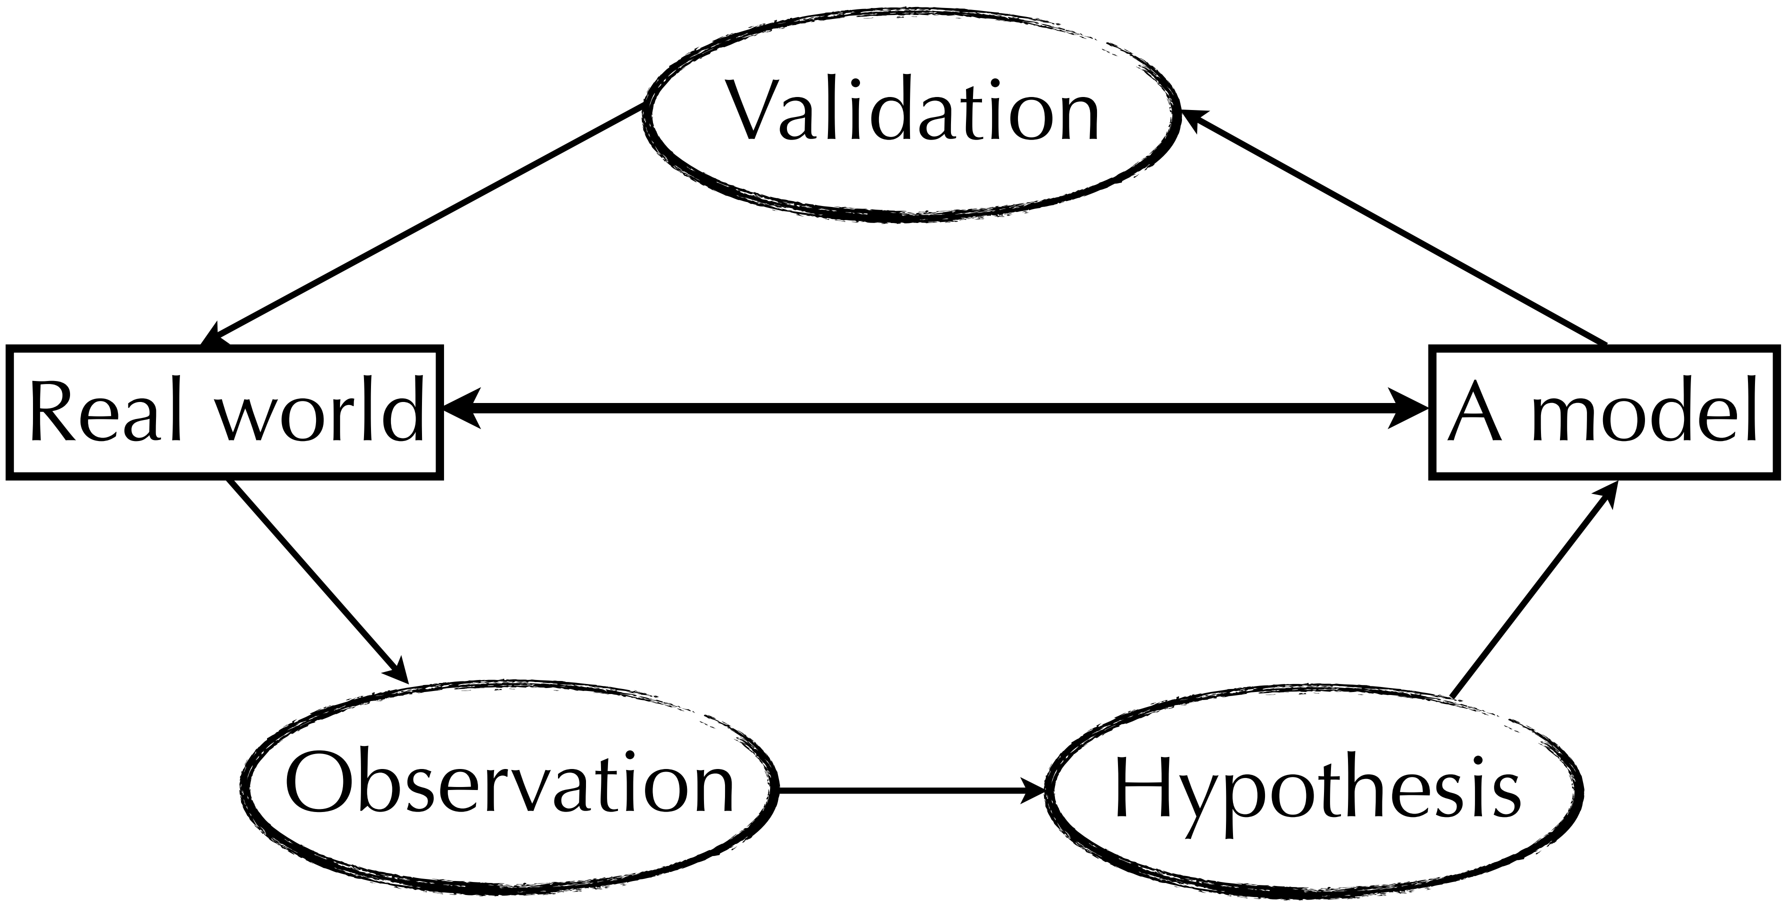

Validation ➞ Prediction

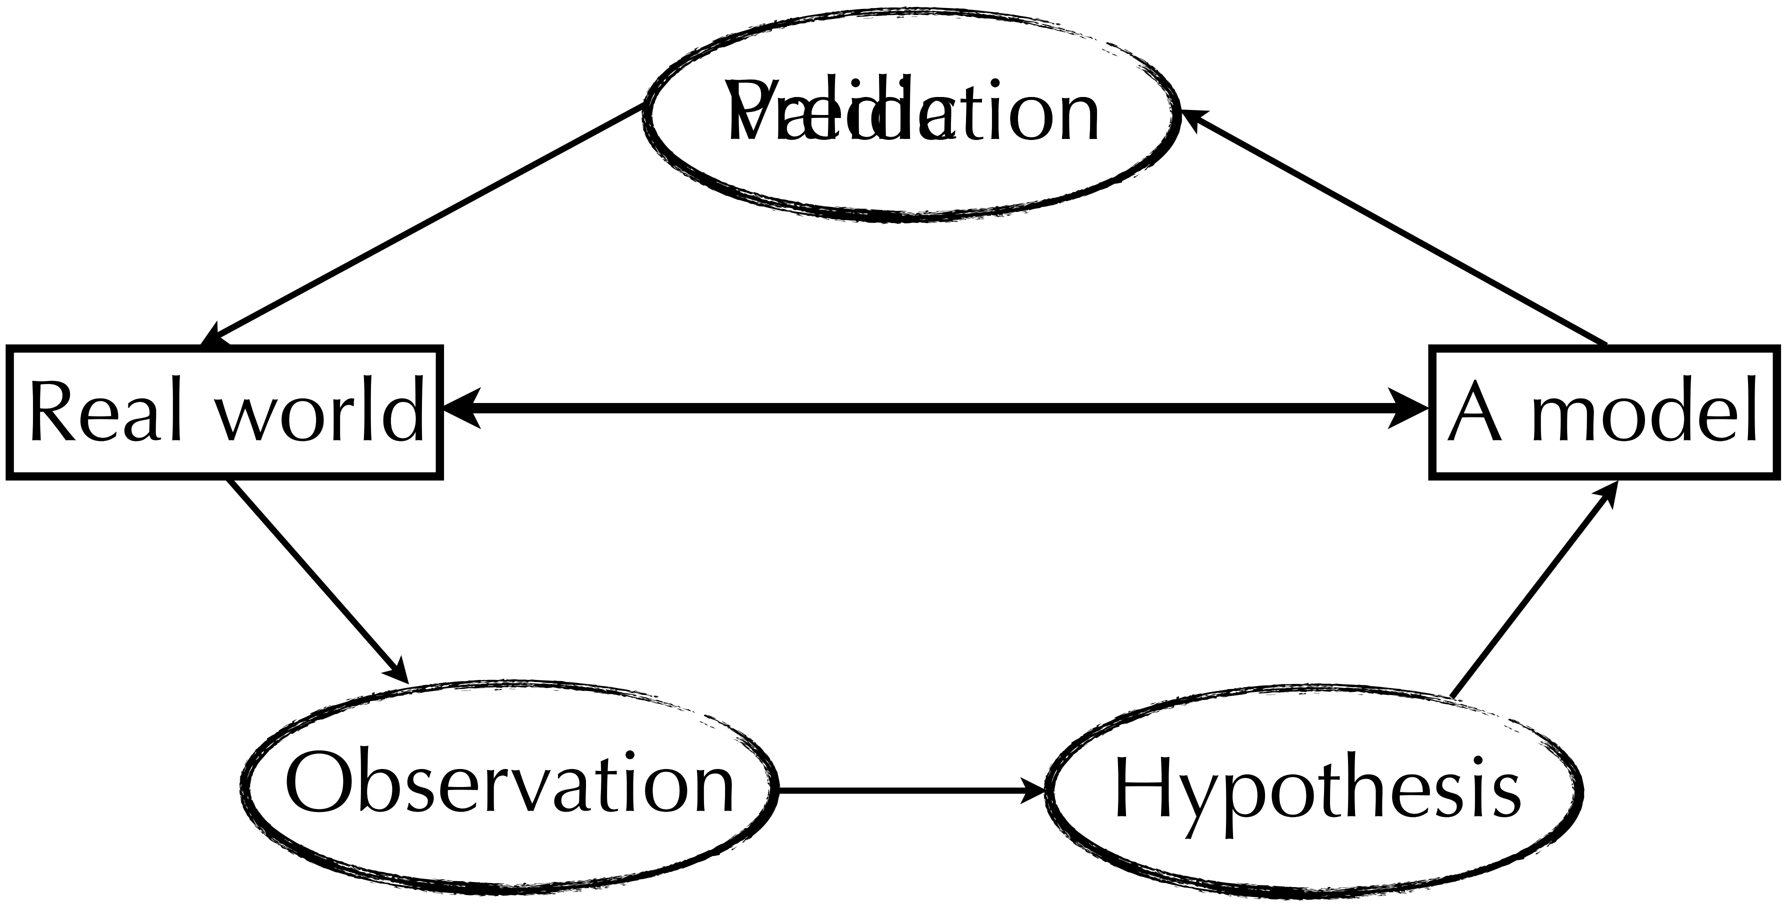

Jim Gray, “eScience -- A Transformed Scientific Method”, (2007)
https://en.wikipedia.org/wiki/File:Jim_Gray_portrait,_1999.jpg

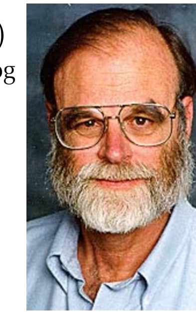
1944-2007

4 Paradigms of Science

- empirical: observe, then derive

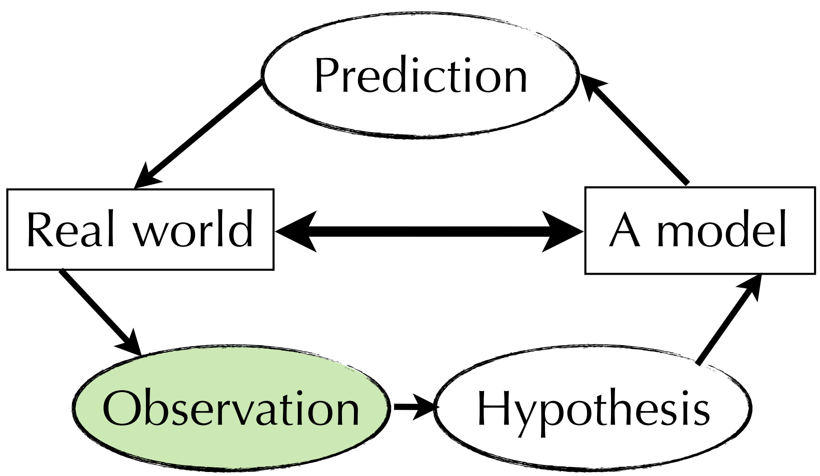

Jim Gray, “eScience -- A Transformed Scientific Method”, (2007)
https://en.wikipedia.org/wiki/File:Jim_Gray_portrait,_1999.jpg

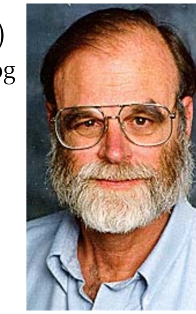
1944-2007

4 Paradigms of Science

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

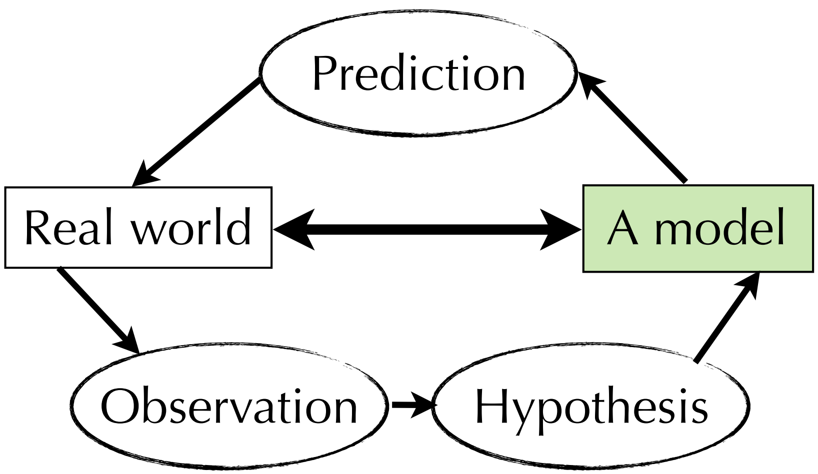

Jim Gray, “eScience -- A Transformed Scientific Method”, (2007)
https://en.wikipedia.org/wiki/File:Jim_Gray_portrait,_1999.jpg

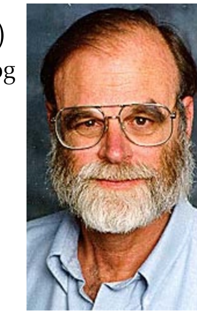

1944-2007

4 Paradigms of Science

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

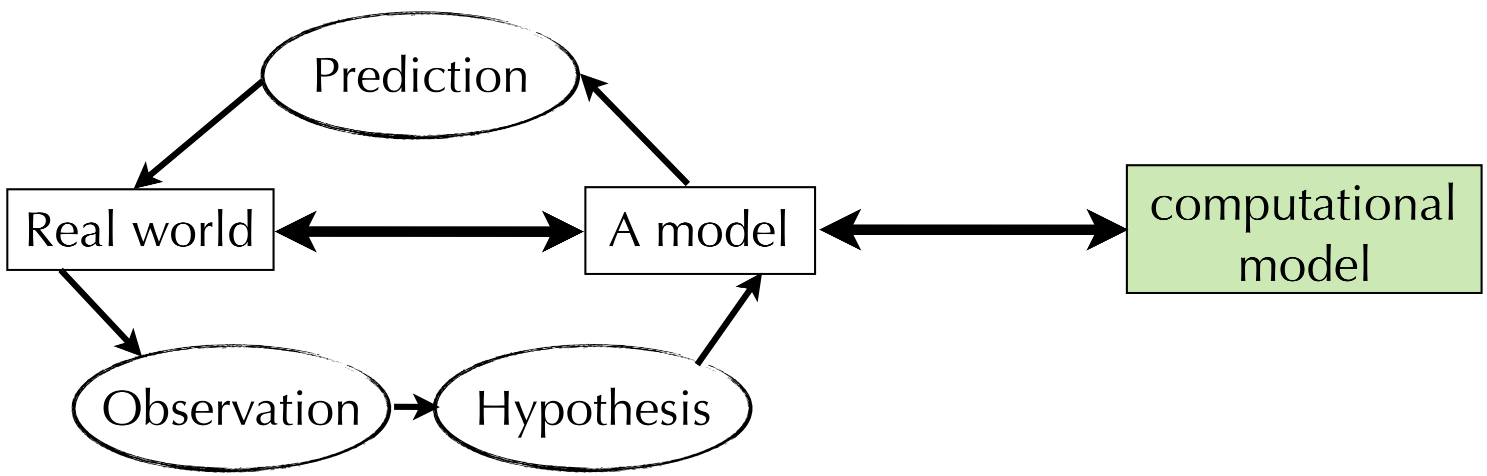

Jim Gray, “eScience -- A Transformed Scientific Method”, (2007)
https://en.wikipedia.org/wiki/File:Jim_Gray_portrait,_1999.jpg

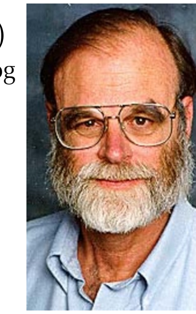
1944-2007

4 Paradigms of Science

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

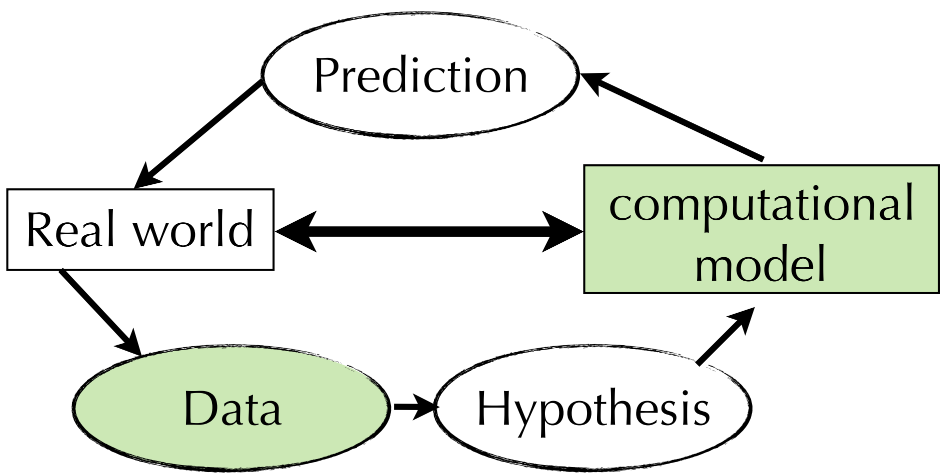

Three types of modelling

- 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

Booshehrian, "Vismon: Facilitating Risk Assessment and Decision Making In Fisheries Management", (2012)

Vismon: Fisheries

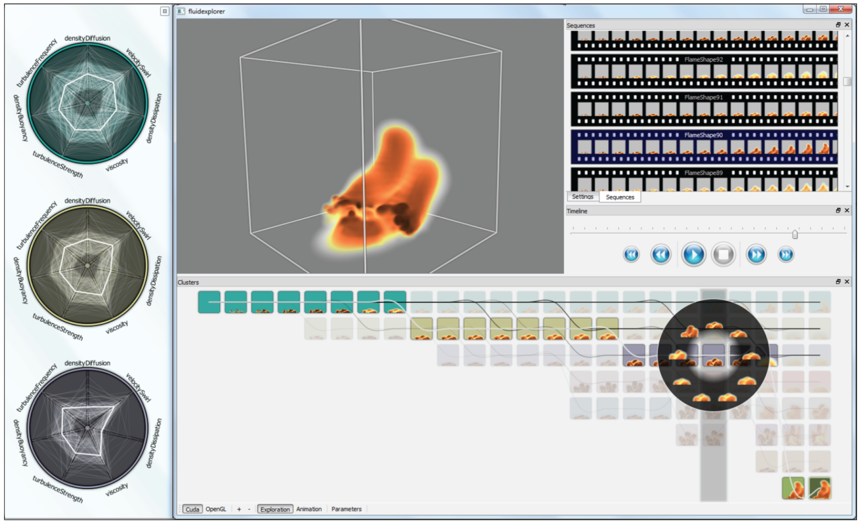

[Bruckner & Möller 2010]

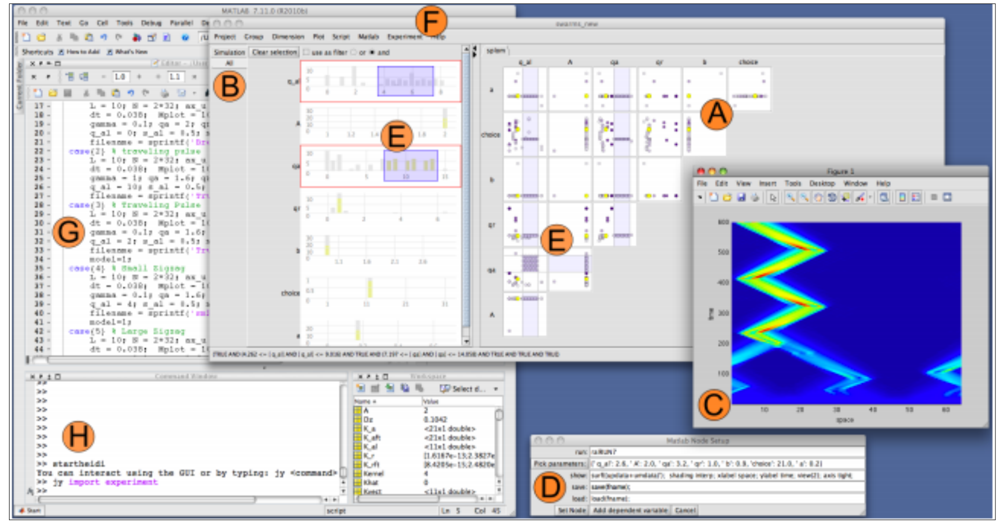

[Bergner et al. 2013]

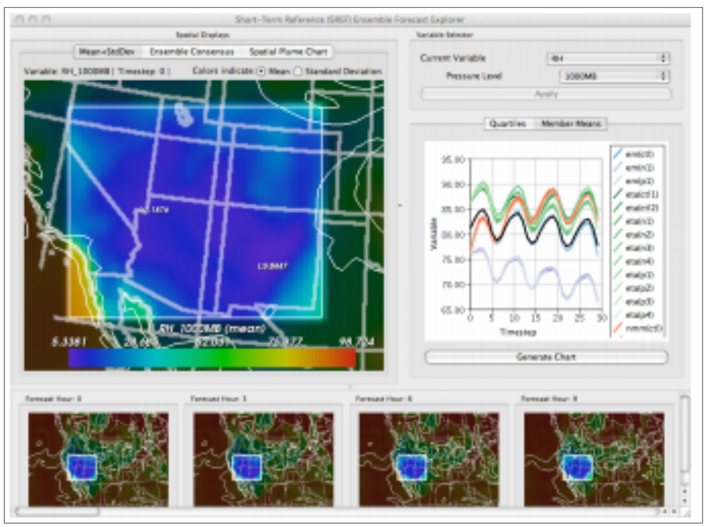

[Potter et al. 2009]

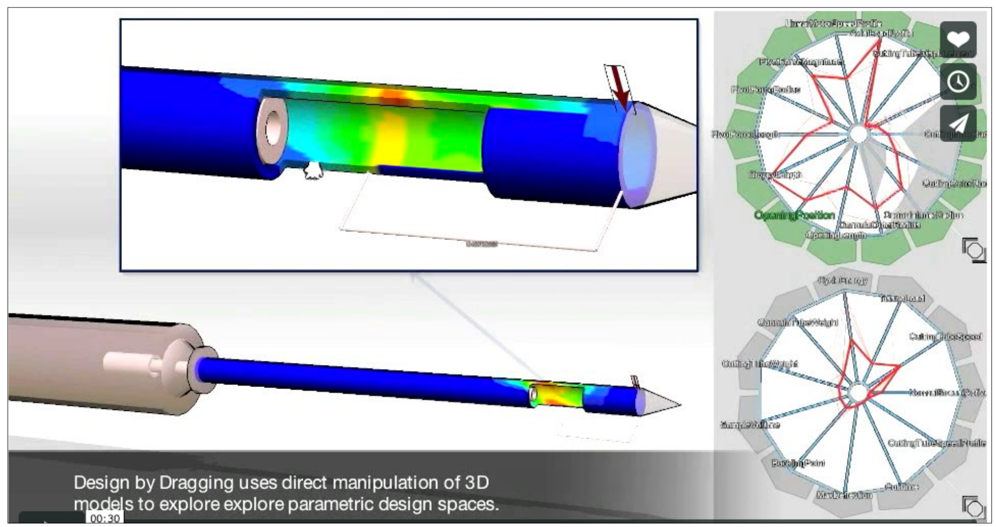

[Coffey et al. 2013]

Statistical Modeling

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

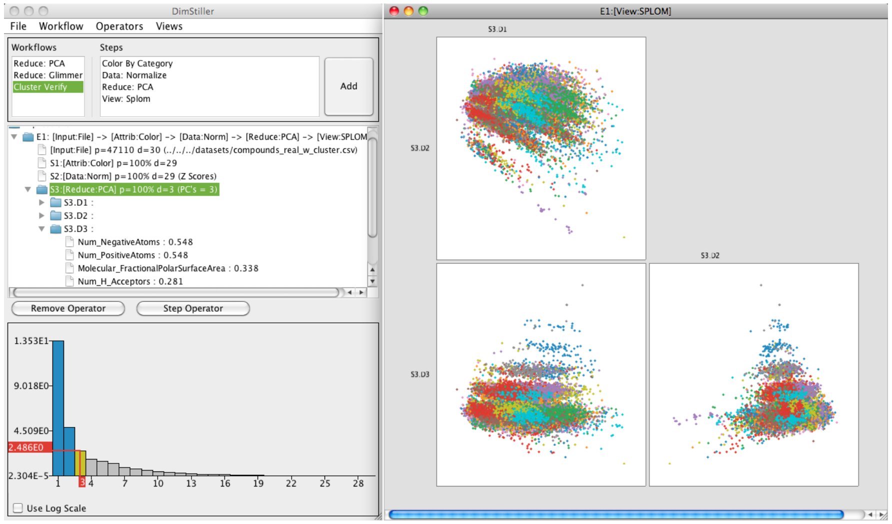

Dim reduction — [Ingram et al. 2010]

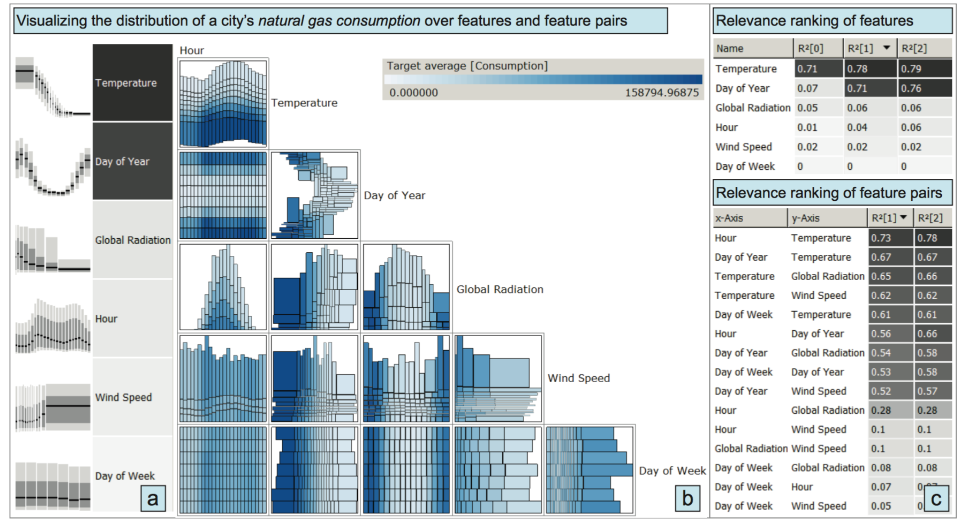

Regression — [Mühlbacher & Piringer 2013]

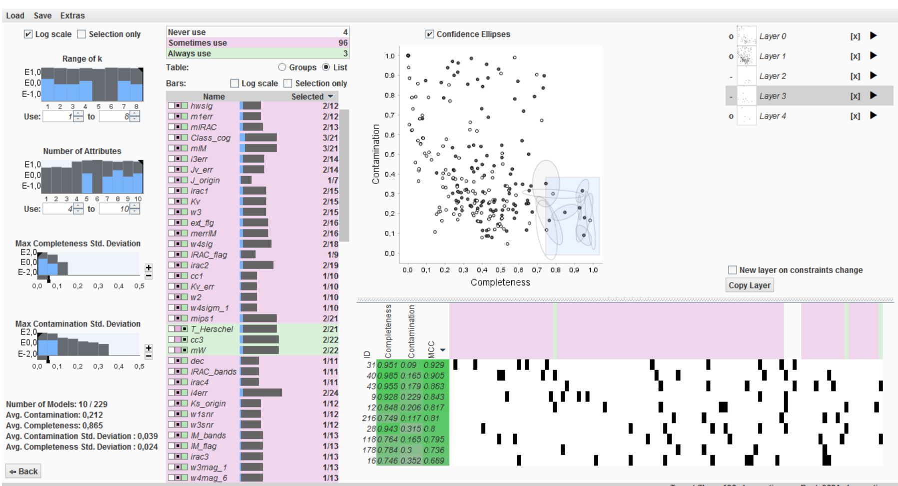

Classification — [Linhardt et al. 2018?]

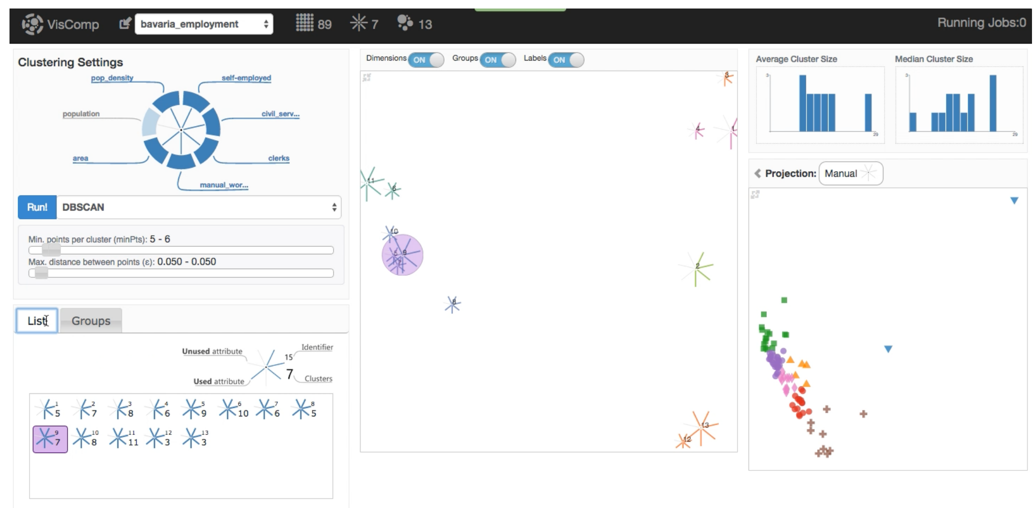

Clustering — [Sedlmair et al. 2018?]

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

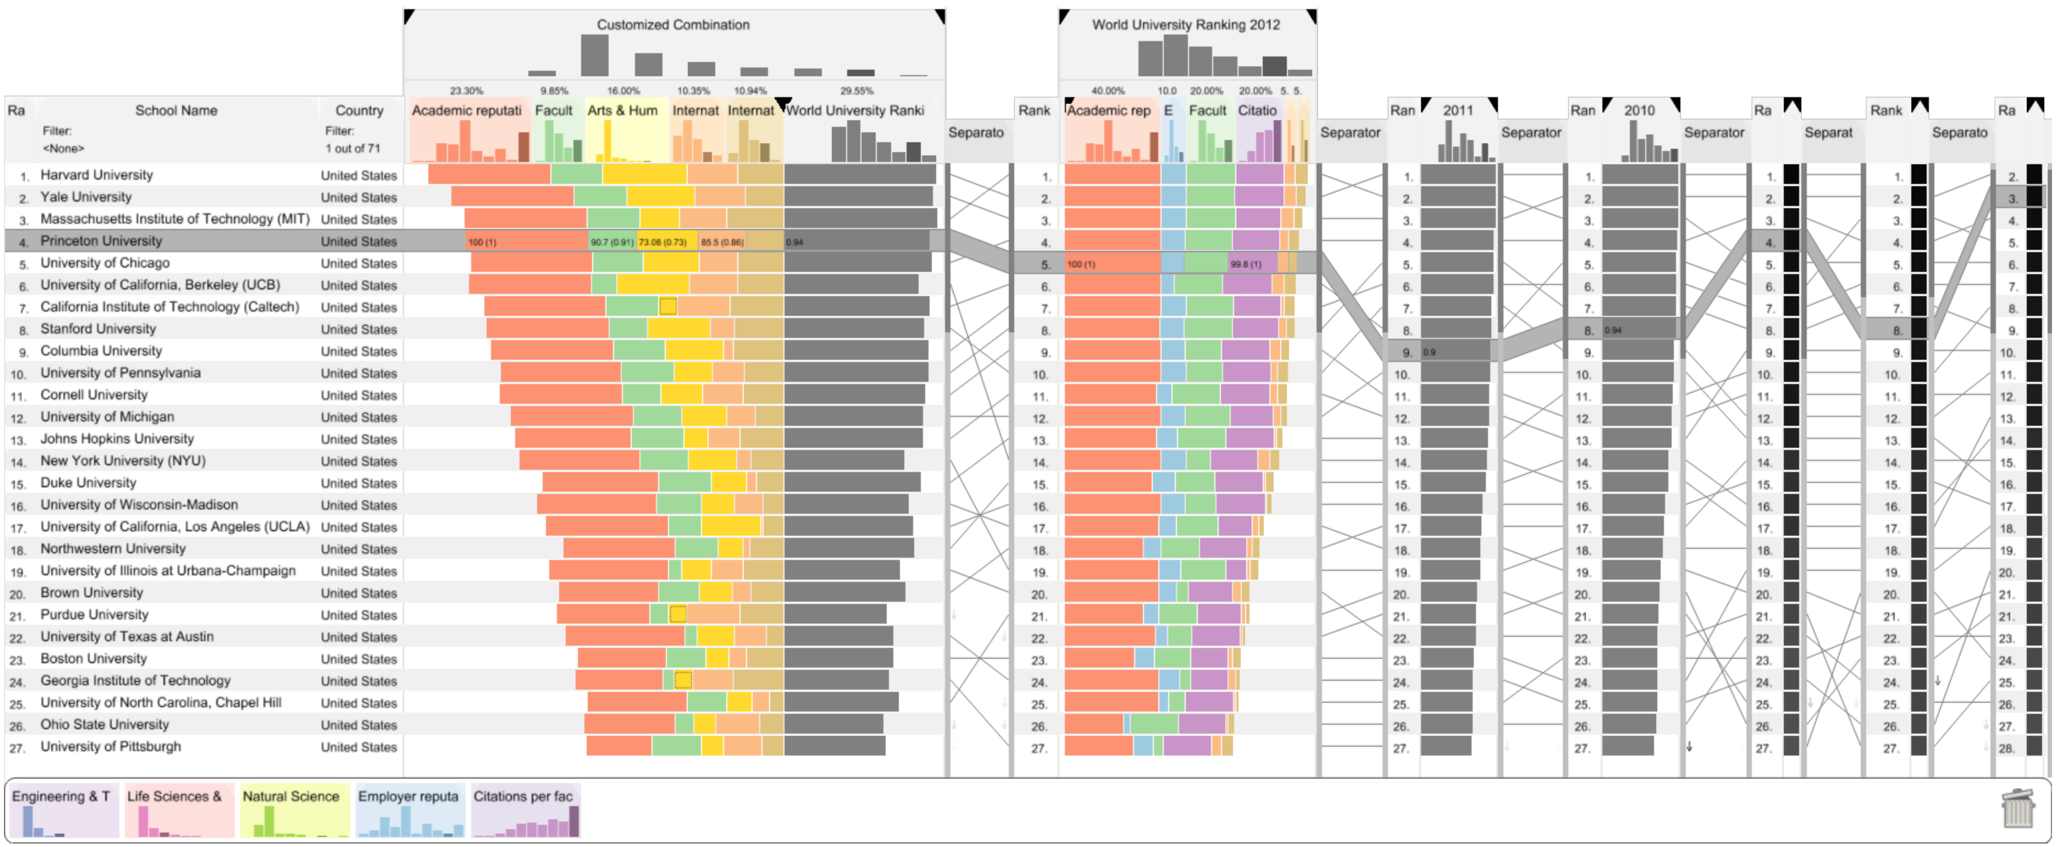

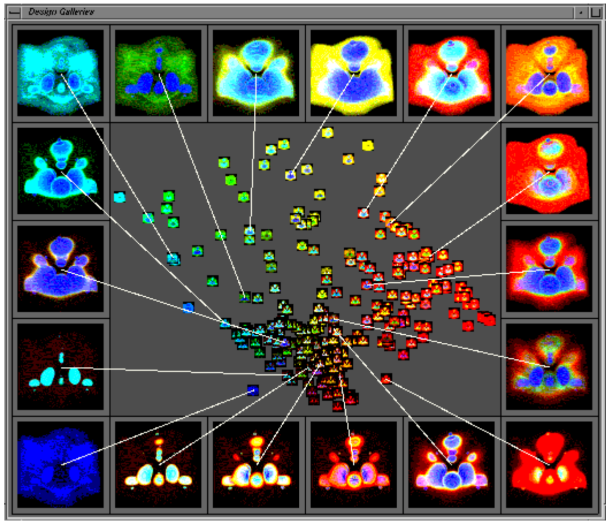

Design Galleries — [Marks et al. 1997]

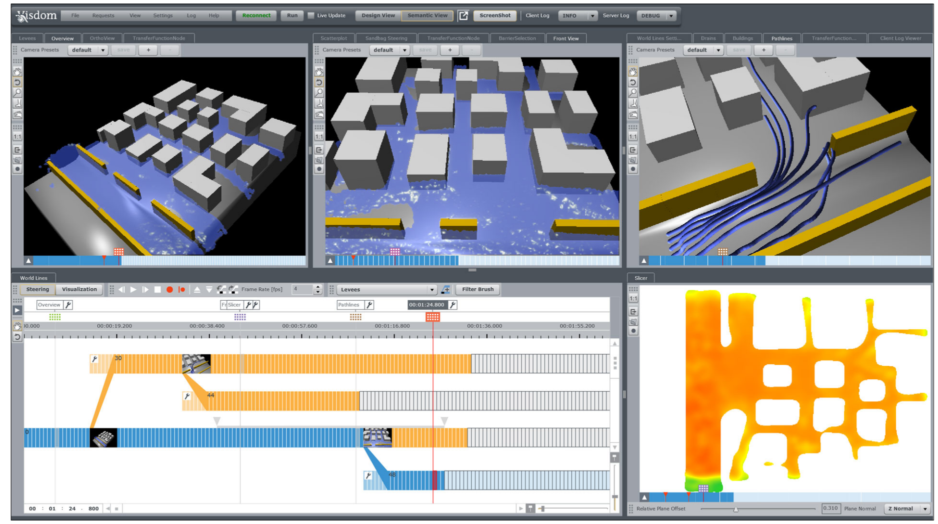

World Lines — [Waser et al. 2010]

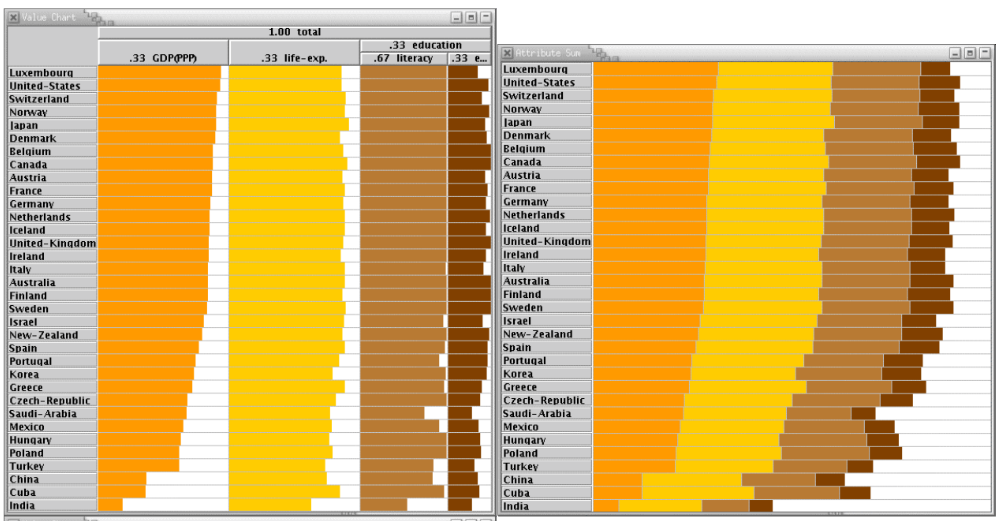

ValueCharts — [Carenini et al. 2004]

Not just Labcoat Science

- valid for business, engineering, public policy
- general data analysis approach

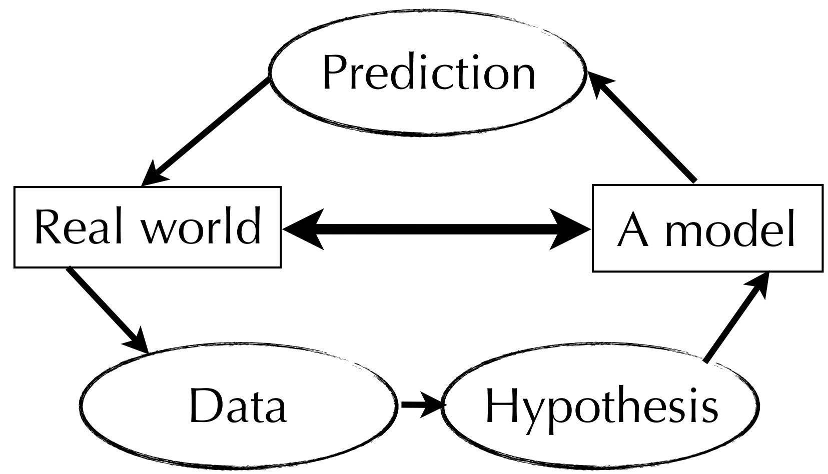

Overview

- Data Science is all about modelling
- The three types of modelling
  - Computational modelling
  - Statistical modelling
  - Empirical modelling
- Challenges of Visual Data Science
- Conclusions

What is visual data science?

- **Visual** Data Science **is helping users** explore, abstract, and communicate complex systems through models from data.

Acting upon models

Building vs. Using

- building models
  - computational experts
  - bioinformaticians

- using models
  - decision makers
  - domain experts
  - biologists

Building vs. Using

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

A modern microscope

Data → Models (predictions) → Decisions

- making difficult algorithmic solutions accessible to a broad audience: enable model users to become model builders

Sedlmair, “Visual Parameter Space Analysis: A Conceptual Framework”, (2014)

What is a model?

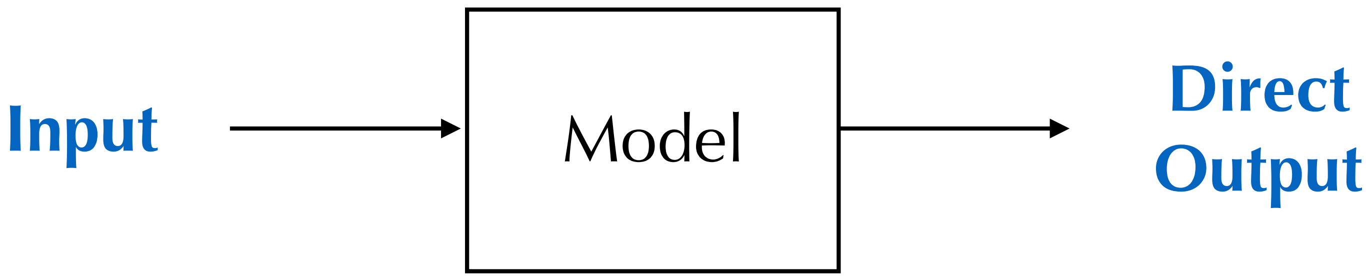

- has input parameters
- creates outputs
- it’s really “just” an algorithm

What is a model?

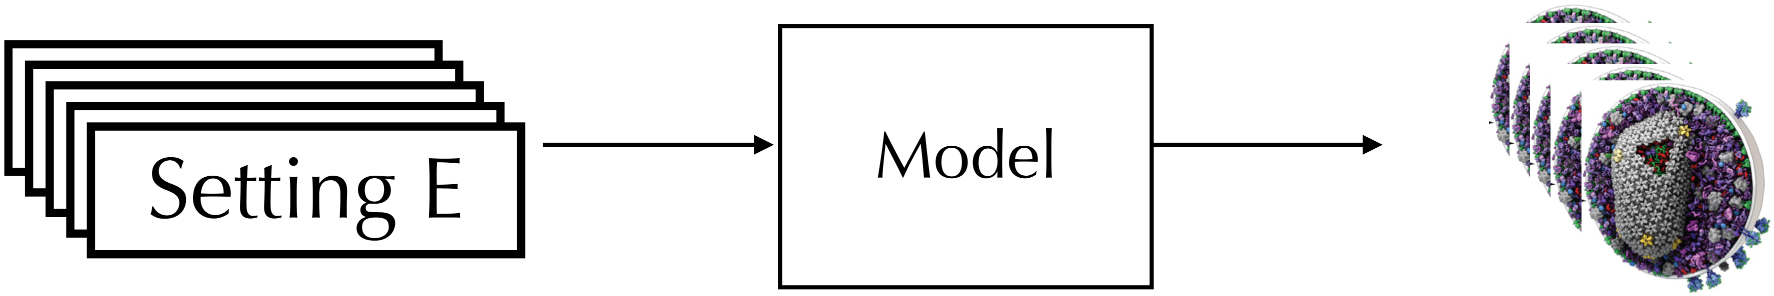

- paradigm shift:
  - from single input/output exploration to input ranges and ensemble outputs

Supporting the user

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

Conclusions

What is visual data science?

- **Visual** Data Science **is helping users** explore, abstract, and communicate complex systems **through models** from data.

Three types of modelling

- computational
- statistical
- empirical

A modern microscope

Data → Models (predictions) → Decisions

- making difficult algorithmic solutions accessible to a broad audience: enable model users to become model builders

Modern microscope

Visual Data Science

=

=

Making modelling techniques accessible to a broad set of users without requiring a PhD in Stats/ML.

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

- Why explainable?
  - the promise of data science
  - extrinsic factors
- How?
  - a process model for simulations
  - (machine) learning environments

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 $\Phi_\epsilon \in \mathcal{H}^{ReLU}_{(N_0,\ldots,N_L)}$ with $\mathrm{size}(\Phi_\epsilon) = \mathcal{O}(\epsilon^{-2})$ and

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

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?

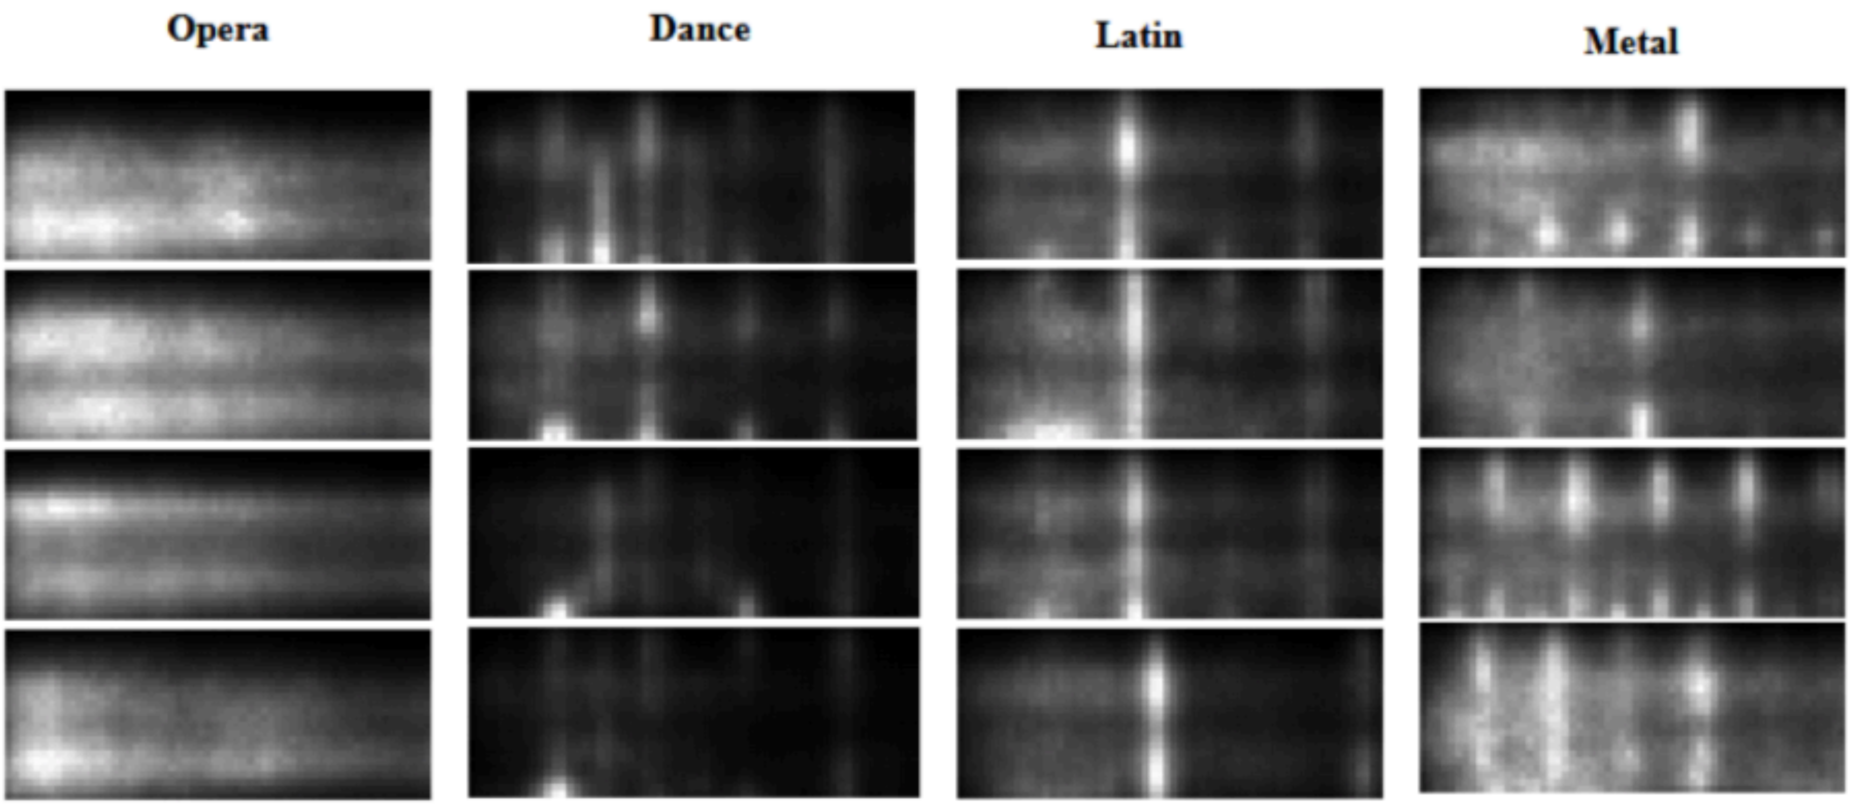

Alex Schindler

How — our approach

https://youtu.be/5d71xhEbjDg

FluidExplorer
Fluid animation

Special effects

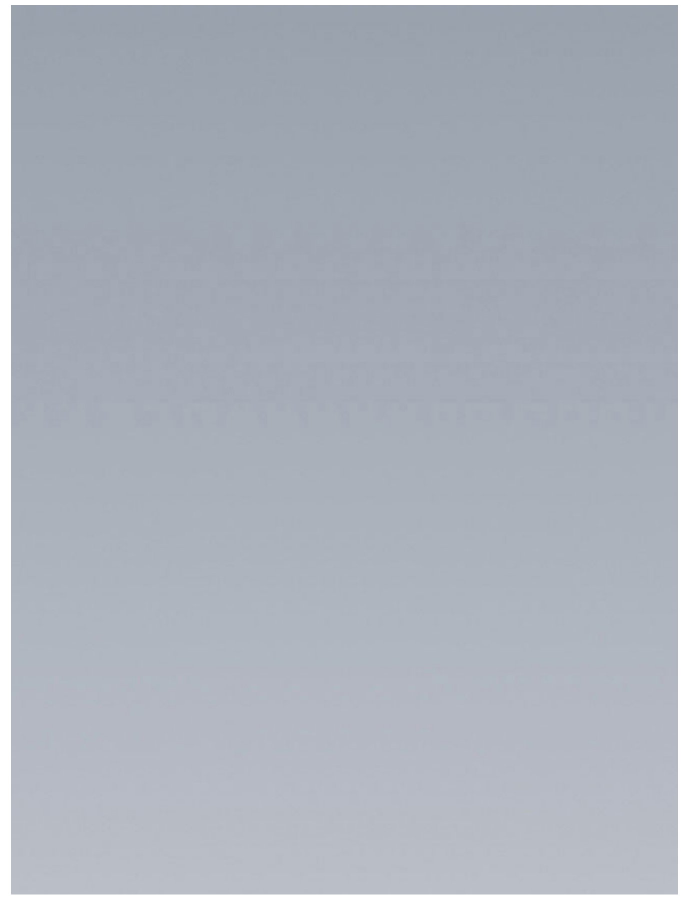

Special effects (2)

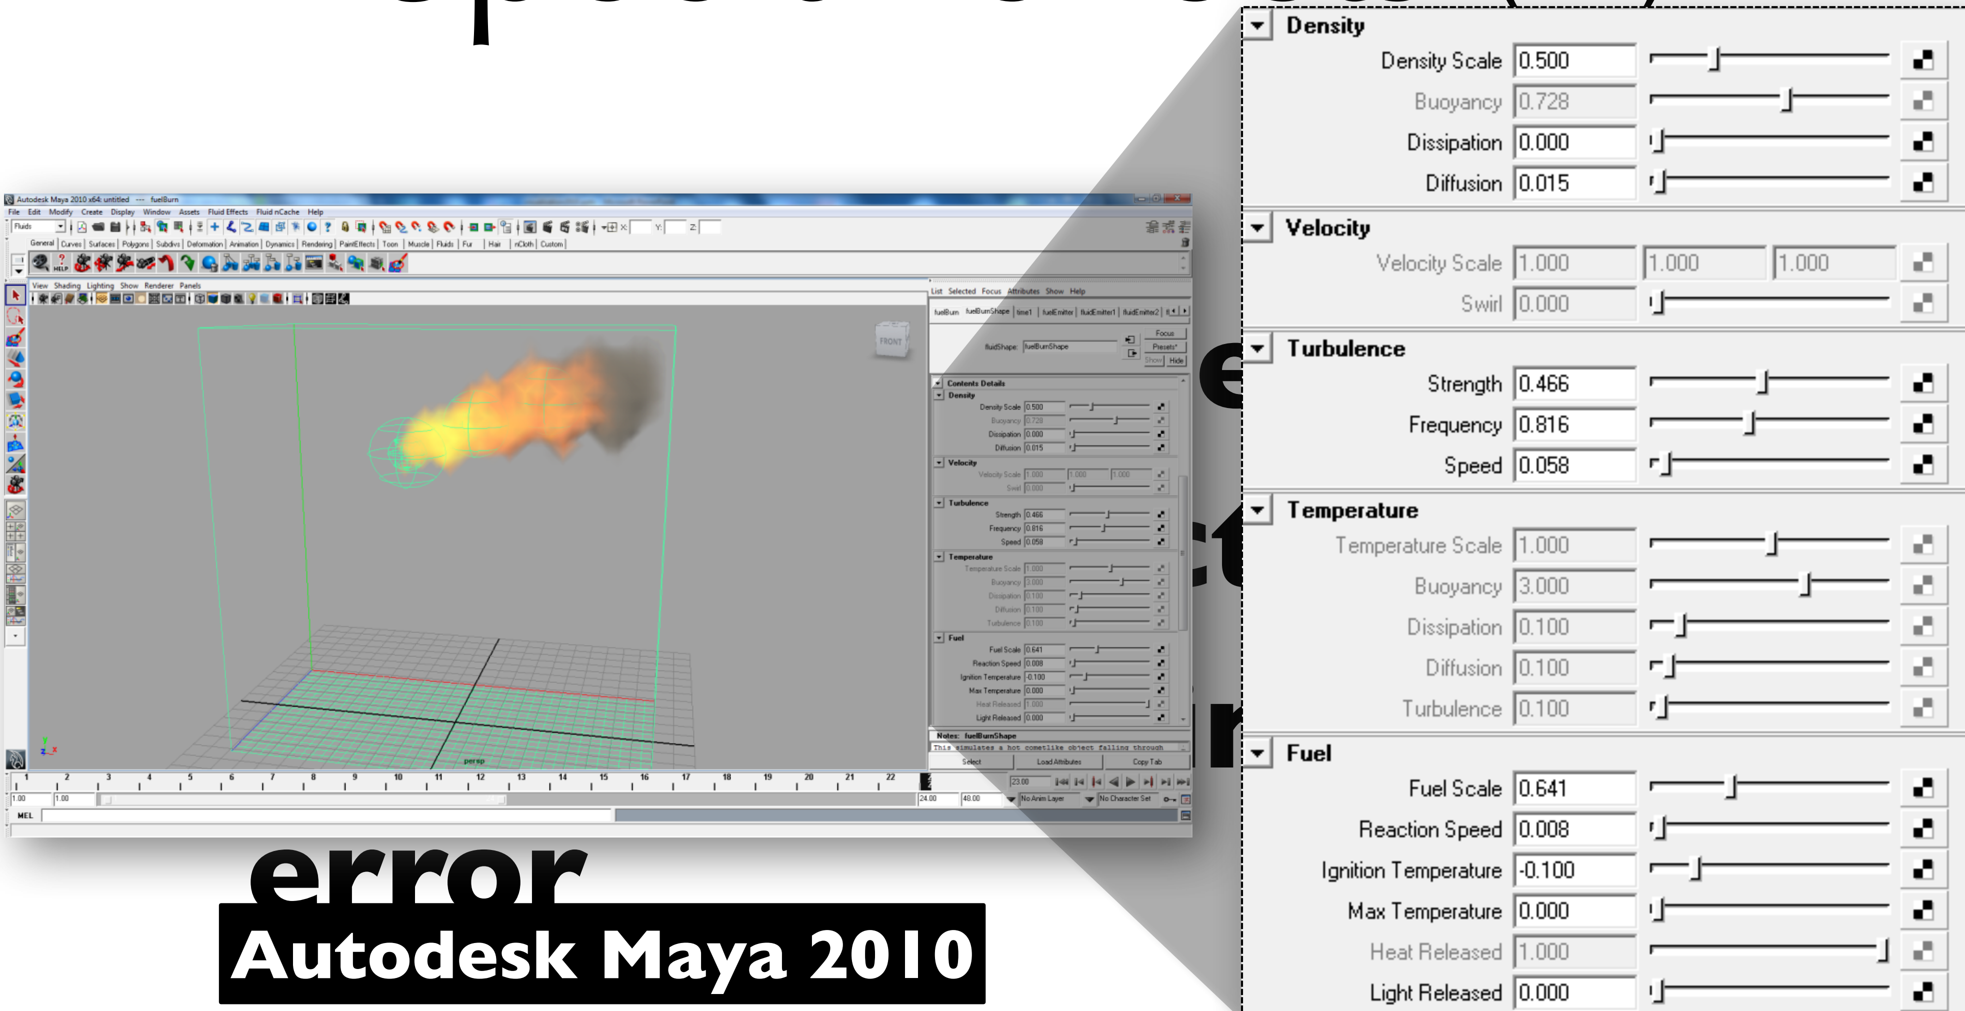

error

Autodesk Maya 2010

Overview

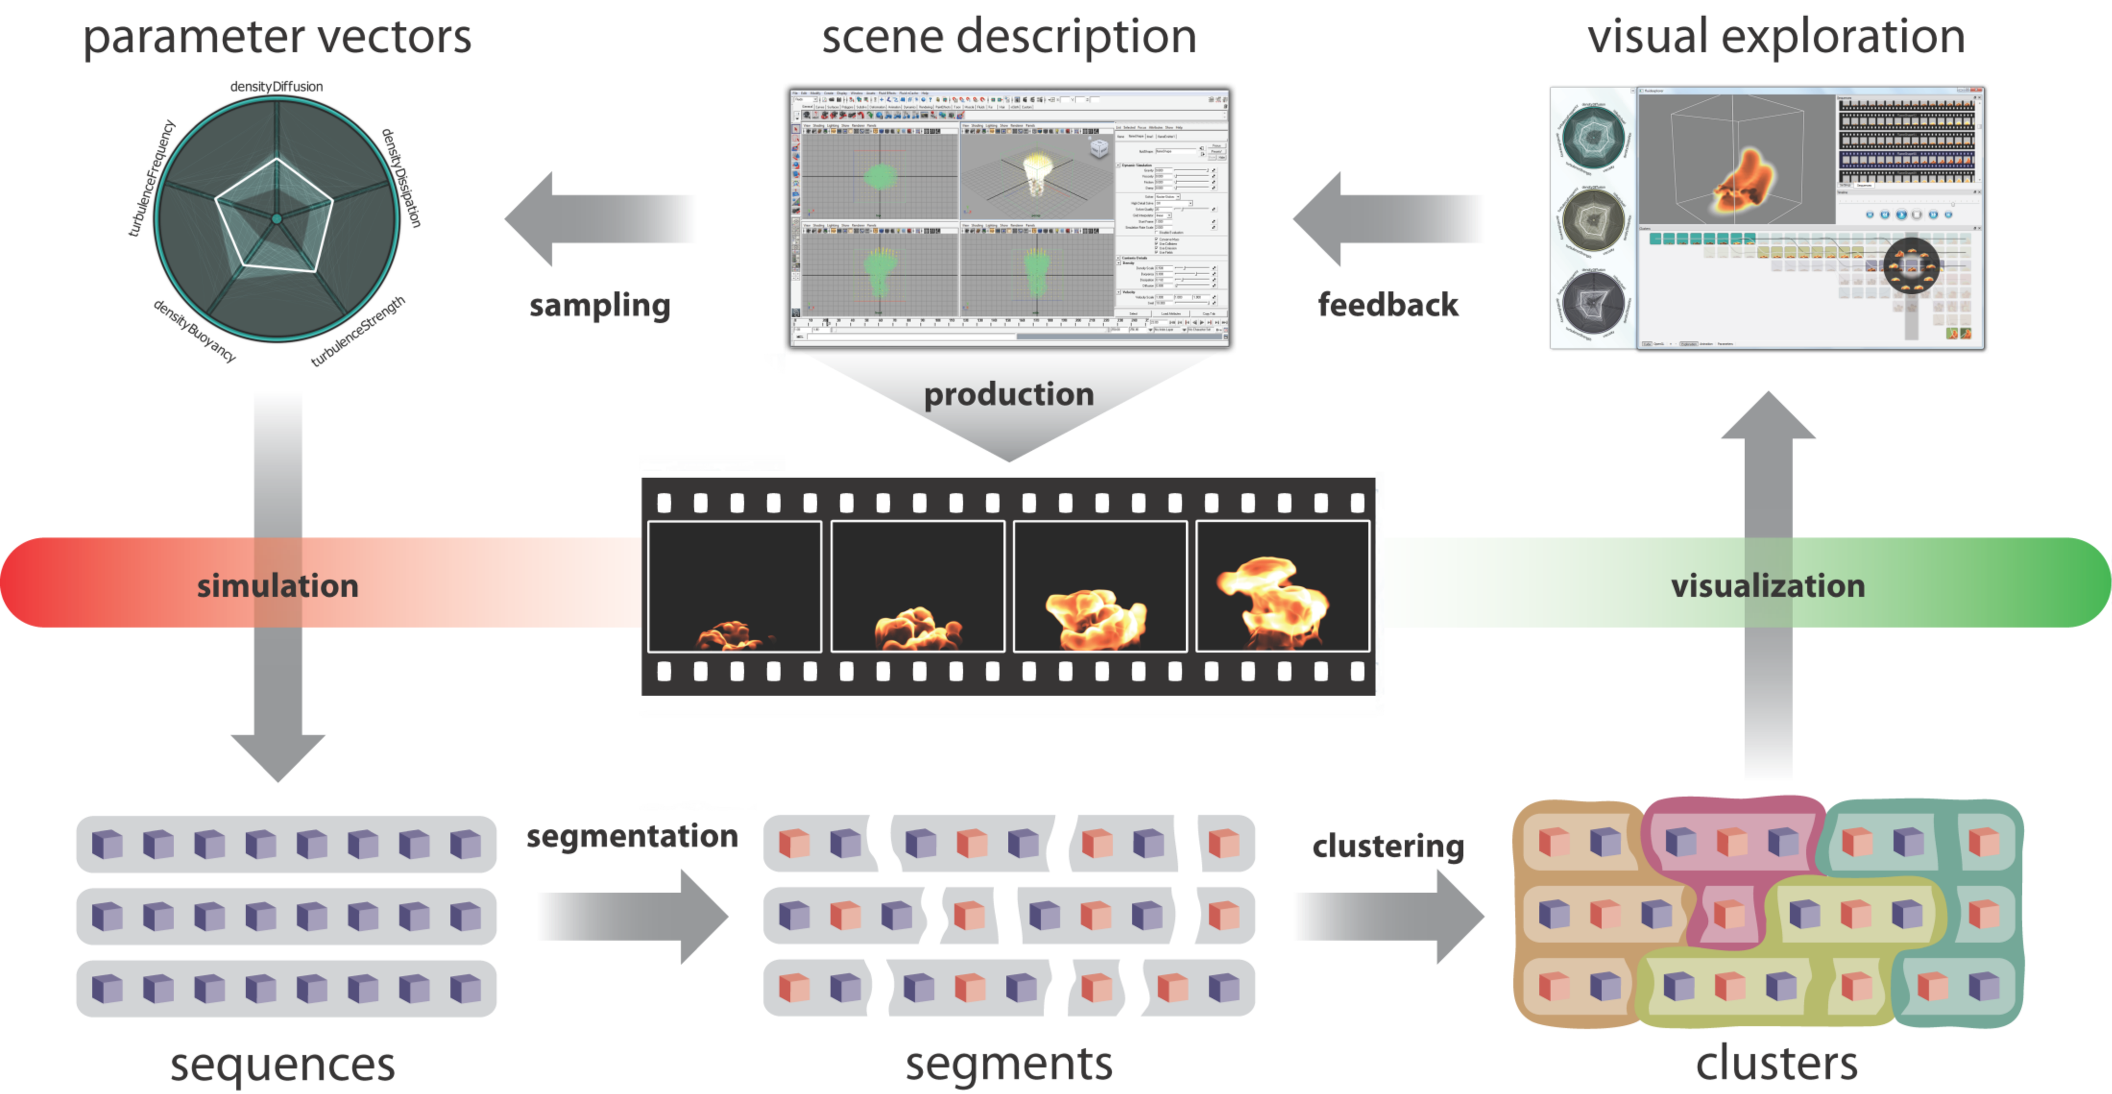

parameter vectors
scene description
visual exploration
densityDiffusion
densityDissipation
turbulenceStrength
densityBuoyancy
turbulenceFrequency
sampling
feedback
production
simulation
visualization
segmentation
clustering
sequences
segments
clusters

Visualization

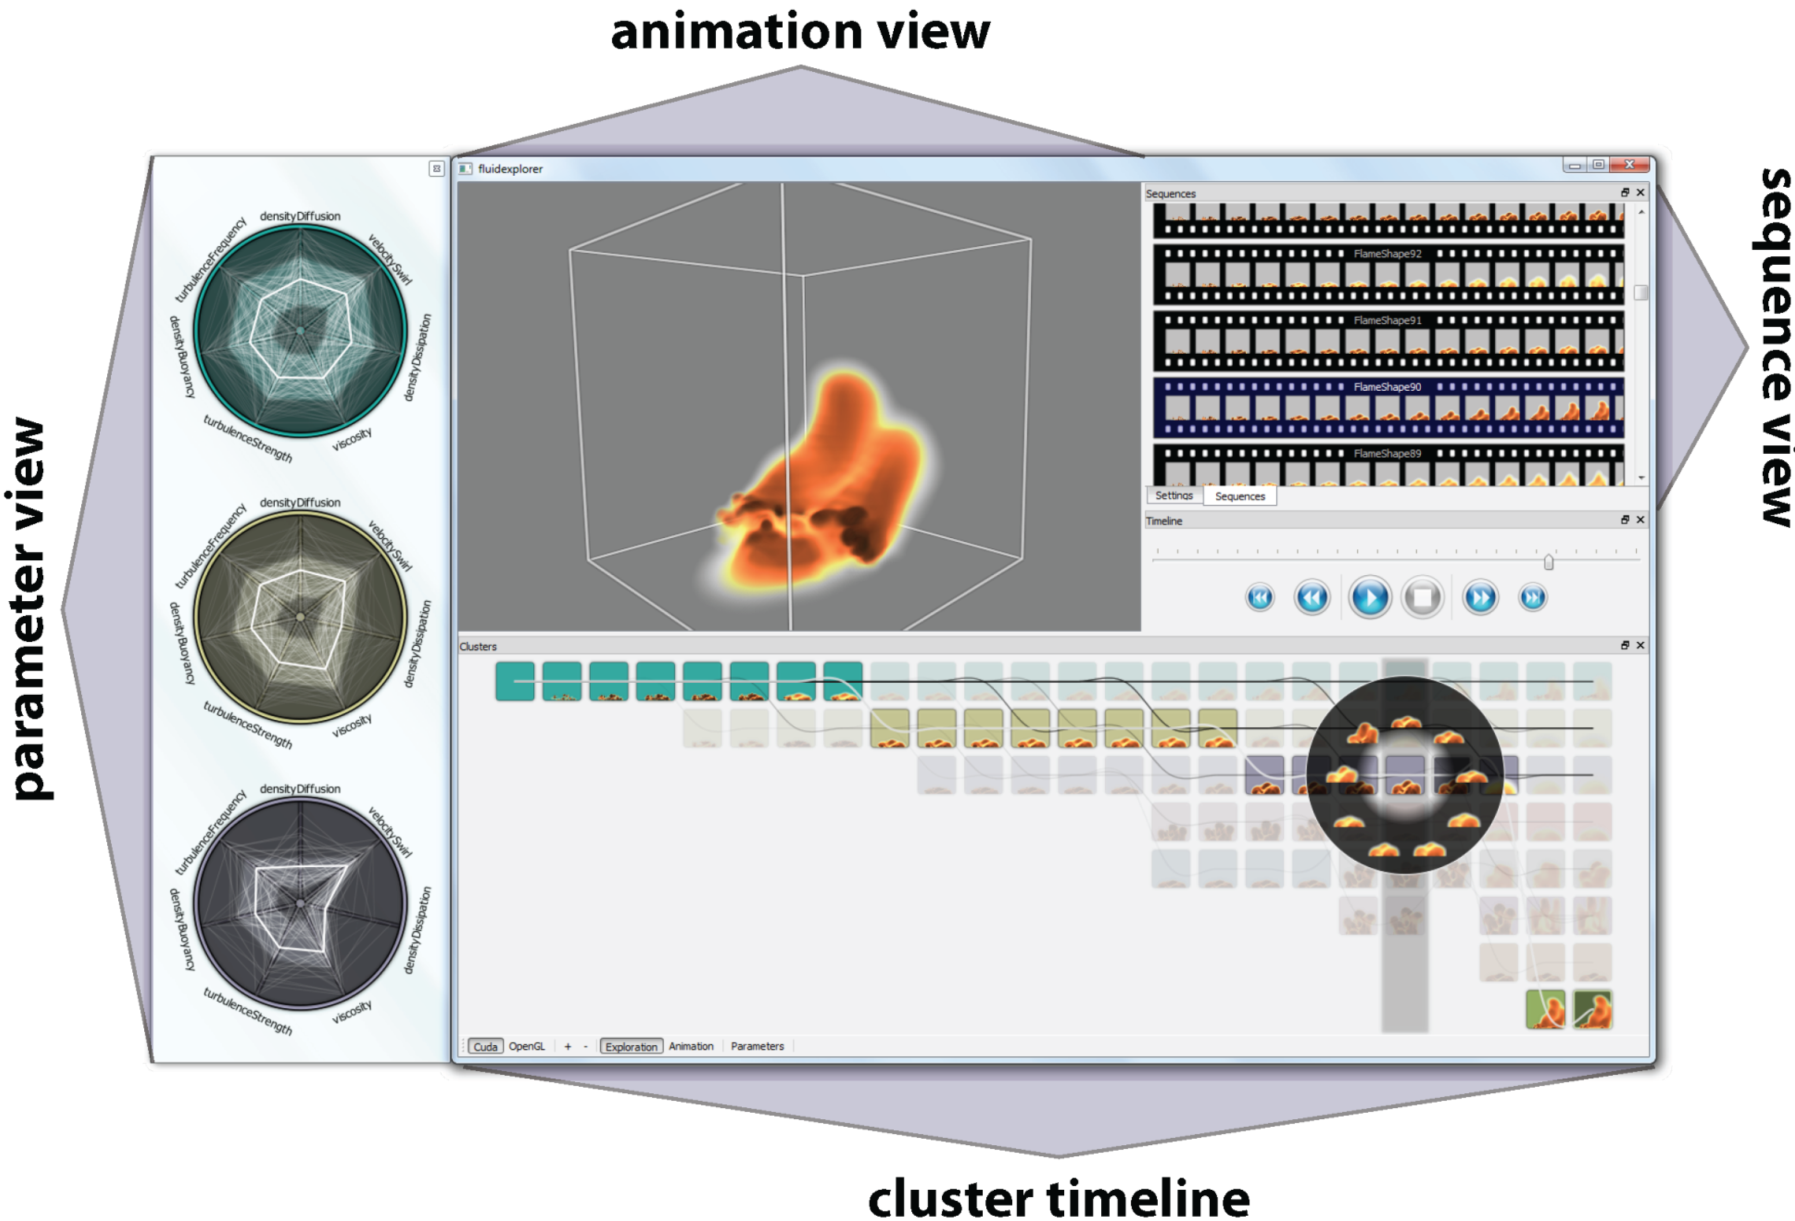

Abstraction:
(visual) Parameter space exploration (vPSA)

Other tools

Much recent attention in vPSA

- Image segmentation [Torsney Weir et al. 2011]
- Weather forecast [Potter et al. 2009]
- Disaster simulation [Waser et al. 2010]
- many more …

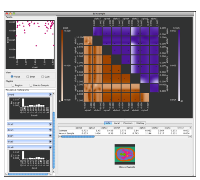
[Torsney-Weir et al. 2011]

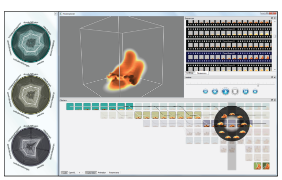
[Bruckner & Möller 2010]

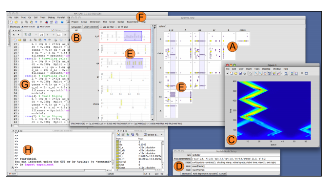
[Bergner et al. 2013]

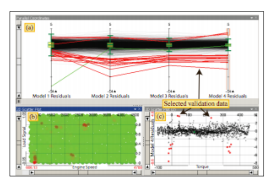
[Piringer et al. 2010]

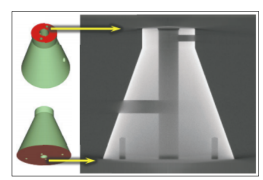
[Amirkhanov et al. 2010]

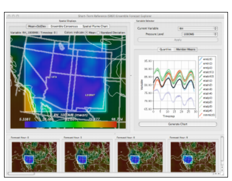
[Potter et al. 2009]

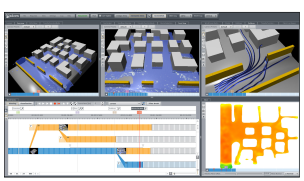
[Waser et al. 2010]

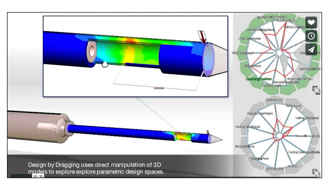
[Coffey et al. 2013]

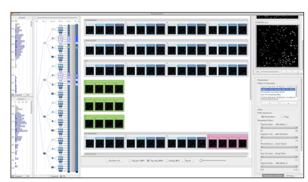
[Pretorius et al. 2011]

…etc.

Much recent attention in vPSA

- **Image segmentation [Torsney Weir et al. 2011]**
- Weather forecast [Potter et al. 2009]
- Disaster simulation [Waser et al. 2010]
- many more …

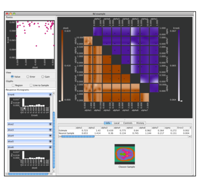
[Torsney-Weir et al. 2011]

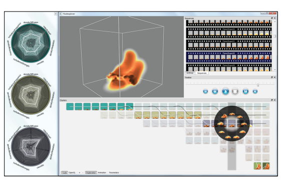
[Bruckner & Möller 2010]

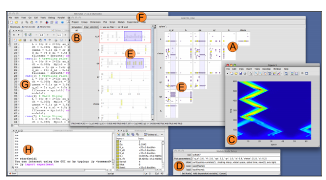
[Bergner et al. 2013]

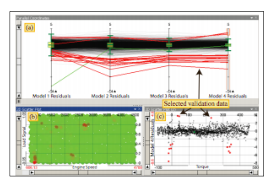
[Piringer et al. 2010]

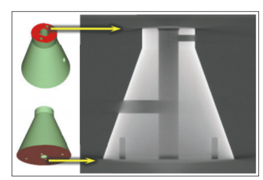
[Amirkhanov et al. 2010]

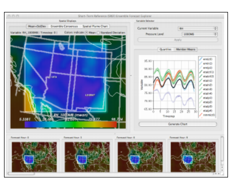
[Potter et al. 2009]

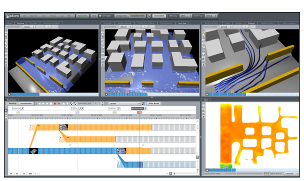
[Waser et al. 2010]

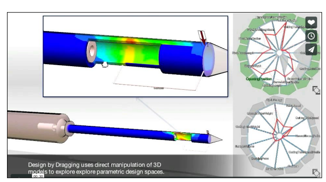
[Coffey et al. 2013]

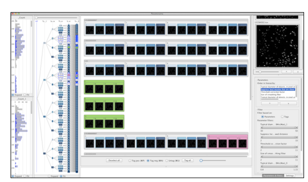
[Pretorius et al. 2011]

…*etc.*

Much recent attention in vPSA

- Image segmentation [Torsney Weir et al. 2011]
- **Weather forecast [Potter et al. 2009]**
- Disaster simulation [Waser et al. 2010]
- many more …

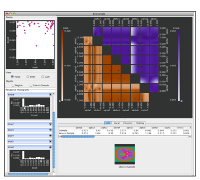
[Torsney-Weir et al. 2011]

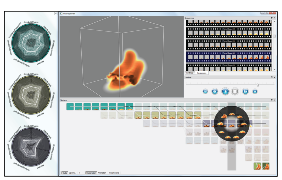
[Bruckner & Möller 2010]

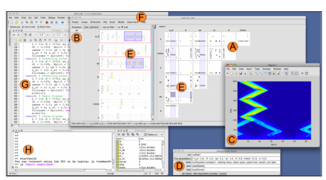
[Bergner et al. 2013]

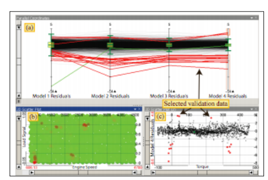
[Piringer et al. 2010]

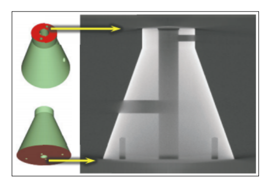
[Amirkhanov et al. 2010]

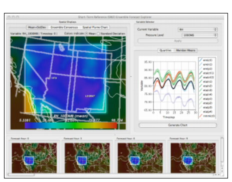
[Potter et al. 2009]

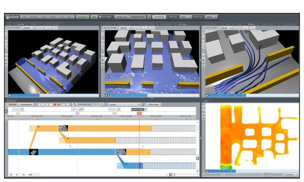
[Waser et al. 2010]

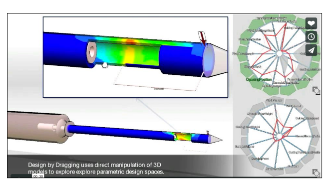
[Coffey et al. 2013]

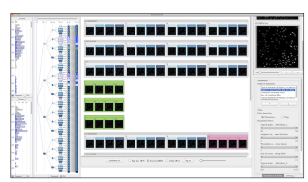
[Pretorius et al. 2011]

…etc.

Much recent attention in vPSA

- Image segmentation [Torsney Weir et al. 2011]
- Weather forecast [Potter et al. 2009]
- **Disaster simulation [Waser et al. 2010]**
- many more …

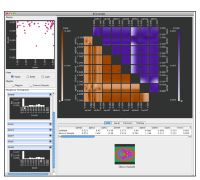
[Torsney-Weir et al. 2011]

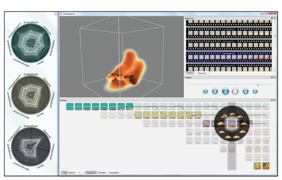
[Bruckner & Möller 2010]

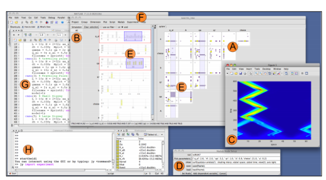
[Bergner et al. 2013]

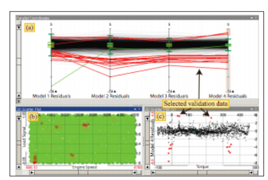
[Piringer et al. 2010]

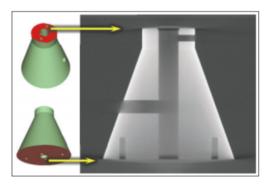
[Amirkhanov et al. 2010]

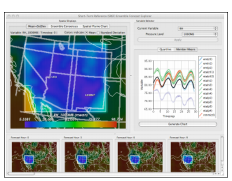
[Potter et al. 2009]

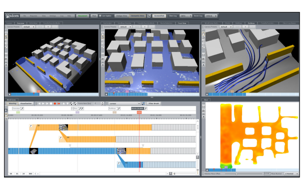
[Waser et al. 2010]

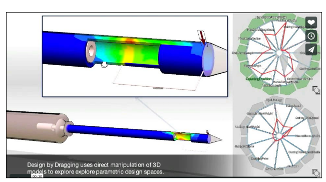
[Coffey et al. 2013]

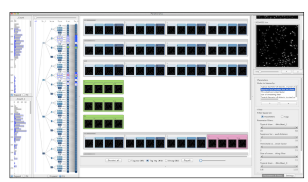
[Pretorius et al. 2011]

…etc.

Much recent attention in vPSA

- comprehensive study of 21 different tools

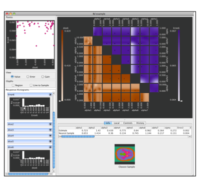

[Torsney-Weir et al. 2011]

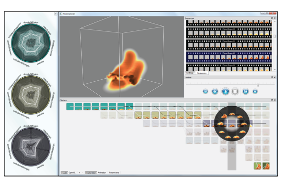

[Bruckner & Möller 2010]

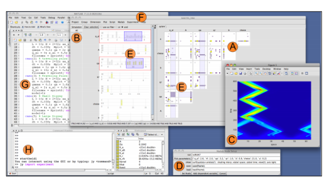

[Bergner et al. 2013]

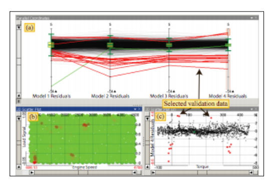

[Piringer et al. 2010]

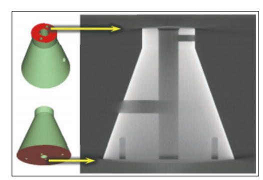

[Amirkhanov et al. 2010]

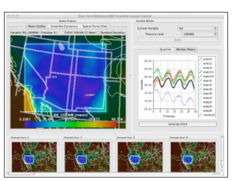

[Potter et al. 2009]

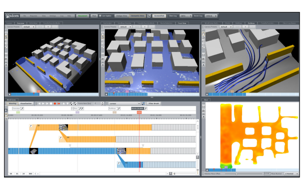

[Waser et al. 2010]

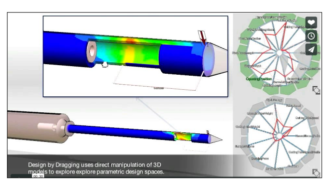

[Coffey et al. 2013]

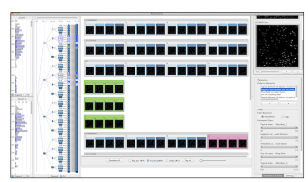

[Pretorius et al. 2011]

…etc.

Data Flow Model

Build an estimator

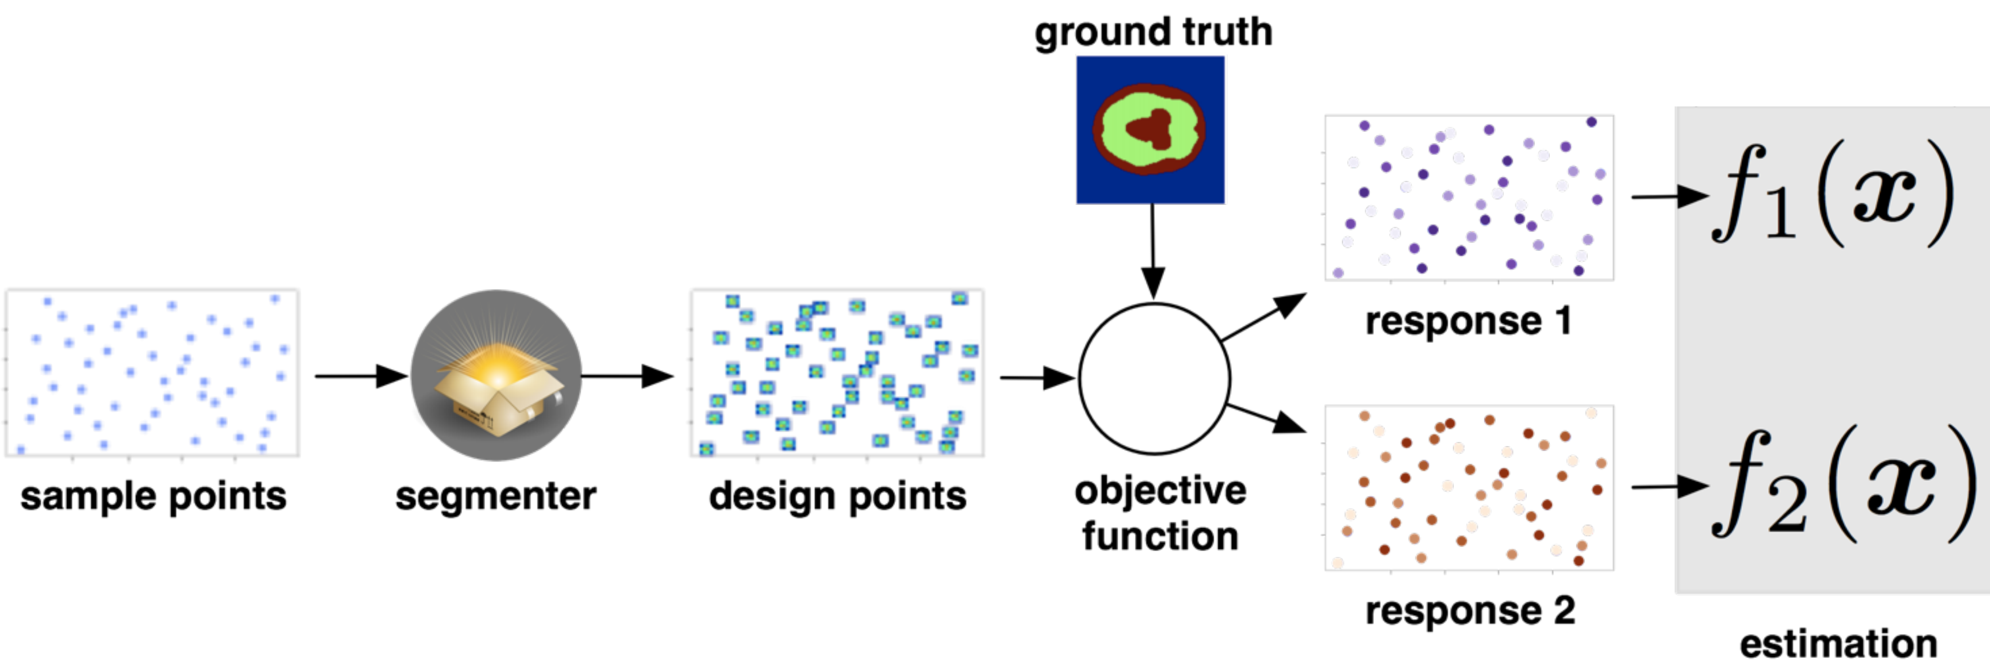

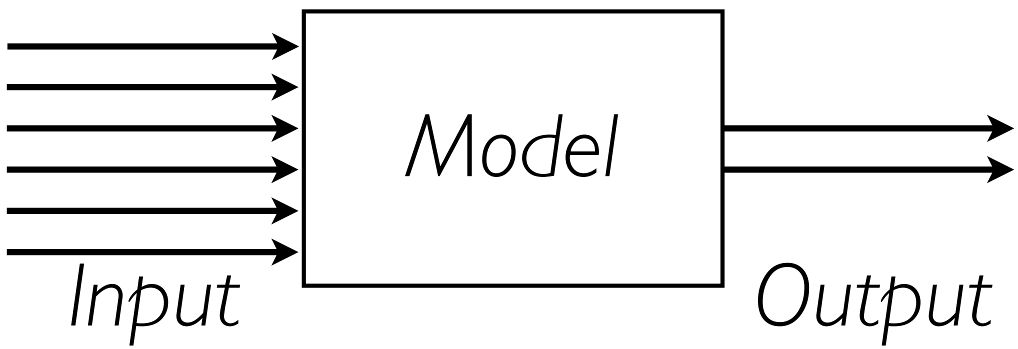
Model
Input
Output

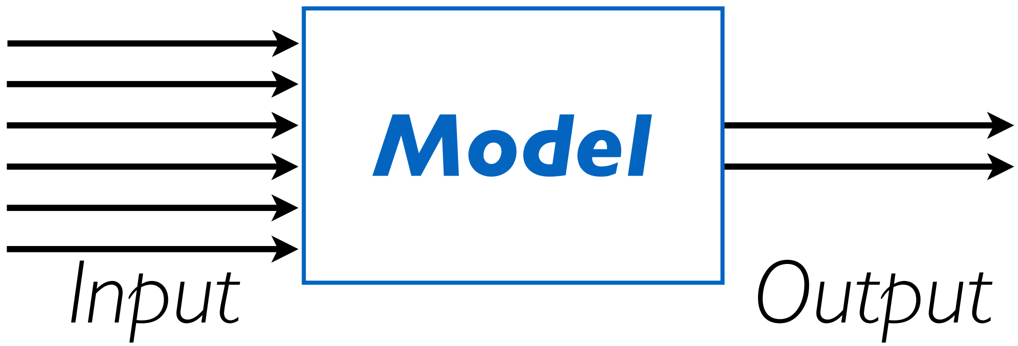

Model

- simulation model, prediction model, …
- … but also algorithm
- stochastic, deterministic
- usually black box (to us as Vis researchers)

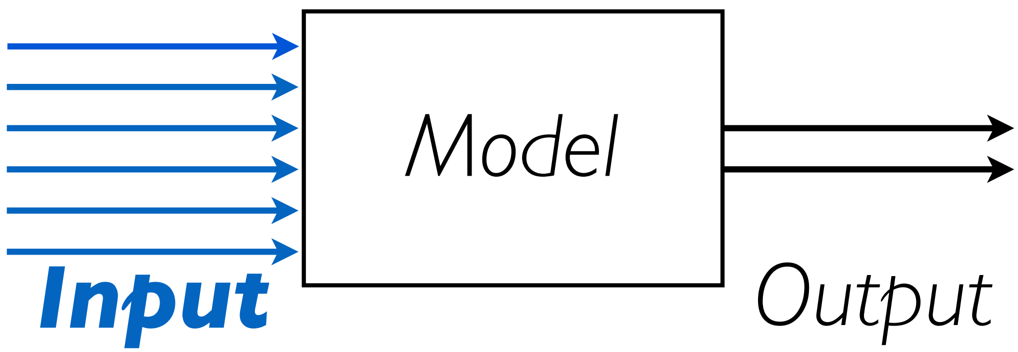

Inputs

- well chosen by the scientist, i.e. people care about their inputs
- normally continuous (quantitative data)
  - need to sample the space
- categorical data common too (e.g. use of a different algorithm)

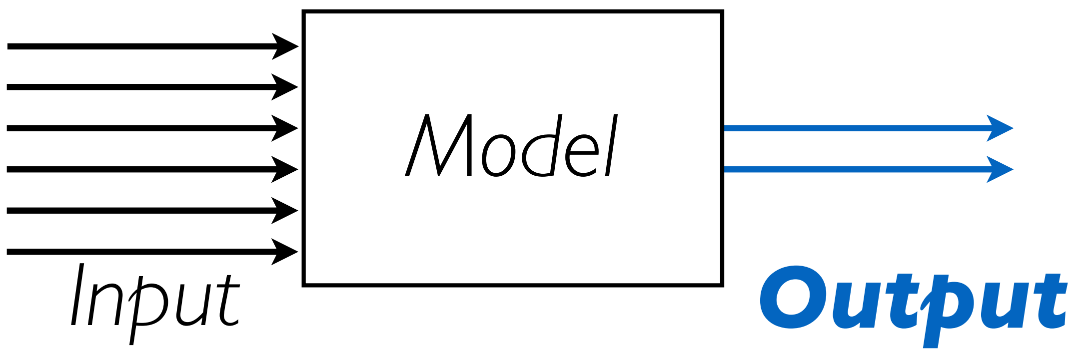

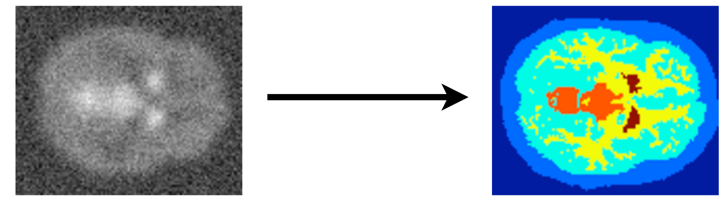

Outputs

- typically complex objects, e.g.
  - 2D, 3D images (Tuner)
  - animations (FluidExplorer)
  - performance graphs (fuel cells)
- hard to evaluate / compare many complex outputs

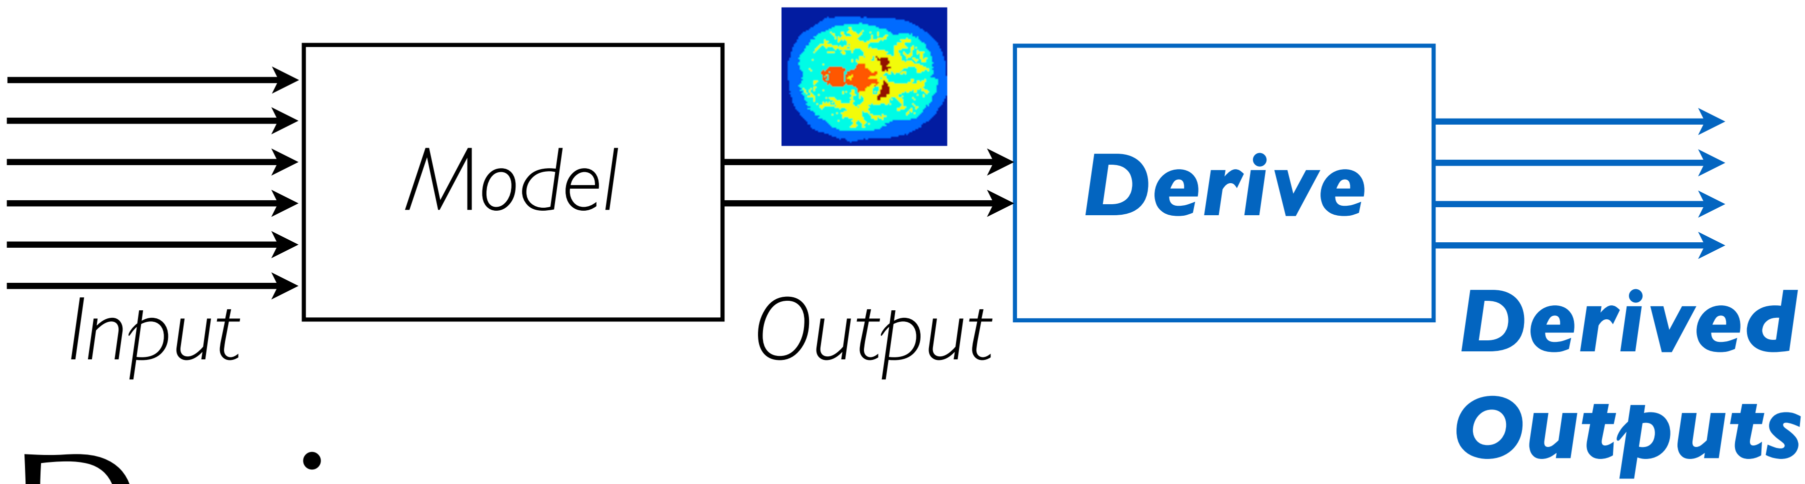

Derive

- one-dimensional (“goodness”) rating: $d(O_1)$
- two-dimensional comparison: $d(O_1, O_2)$
- objective measures can be
  - exact (reliable)
  - approximate - about right, but not 100% precise
  - unknown (active learning)

Complex objects (in 18/21 papers)

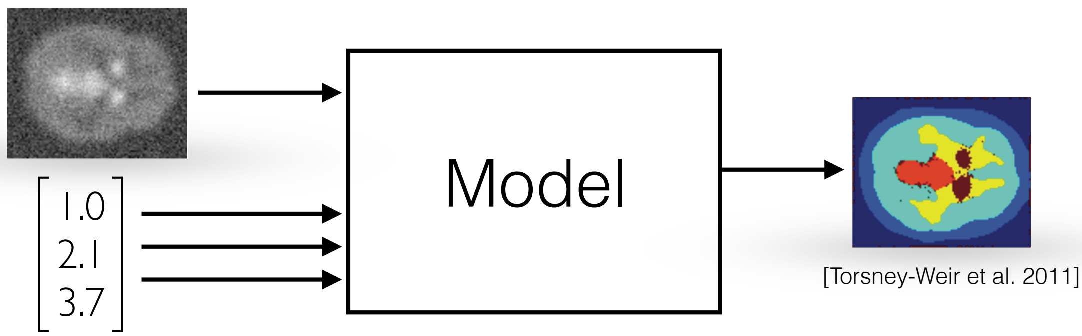

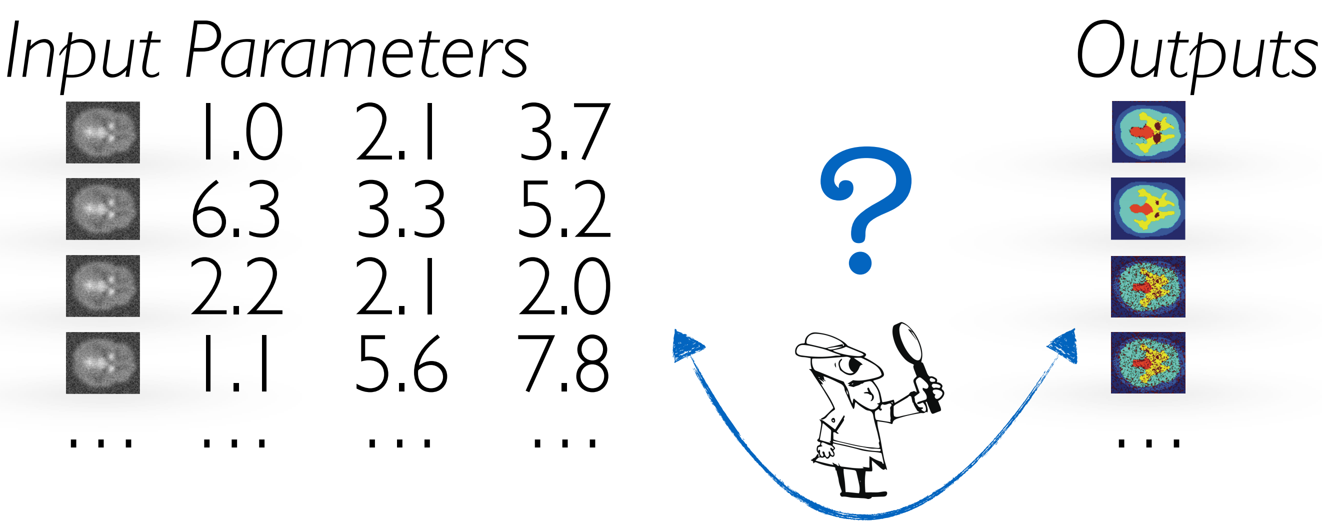

Derive objective measures

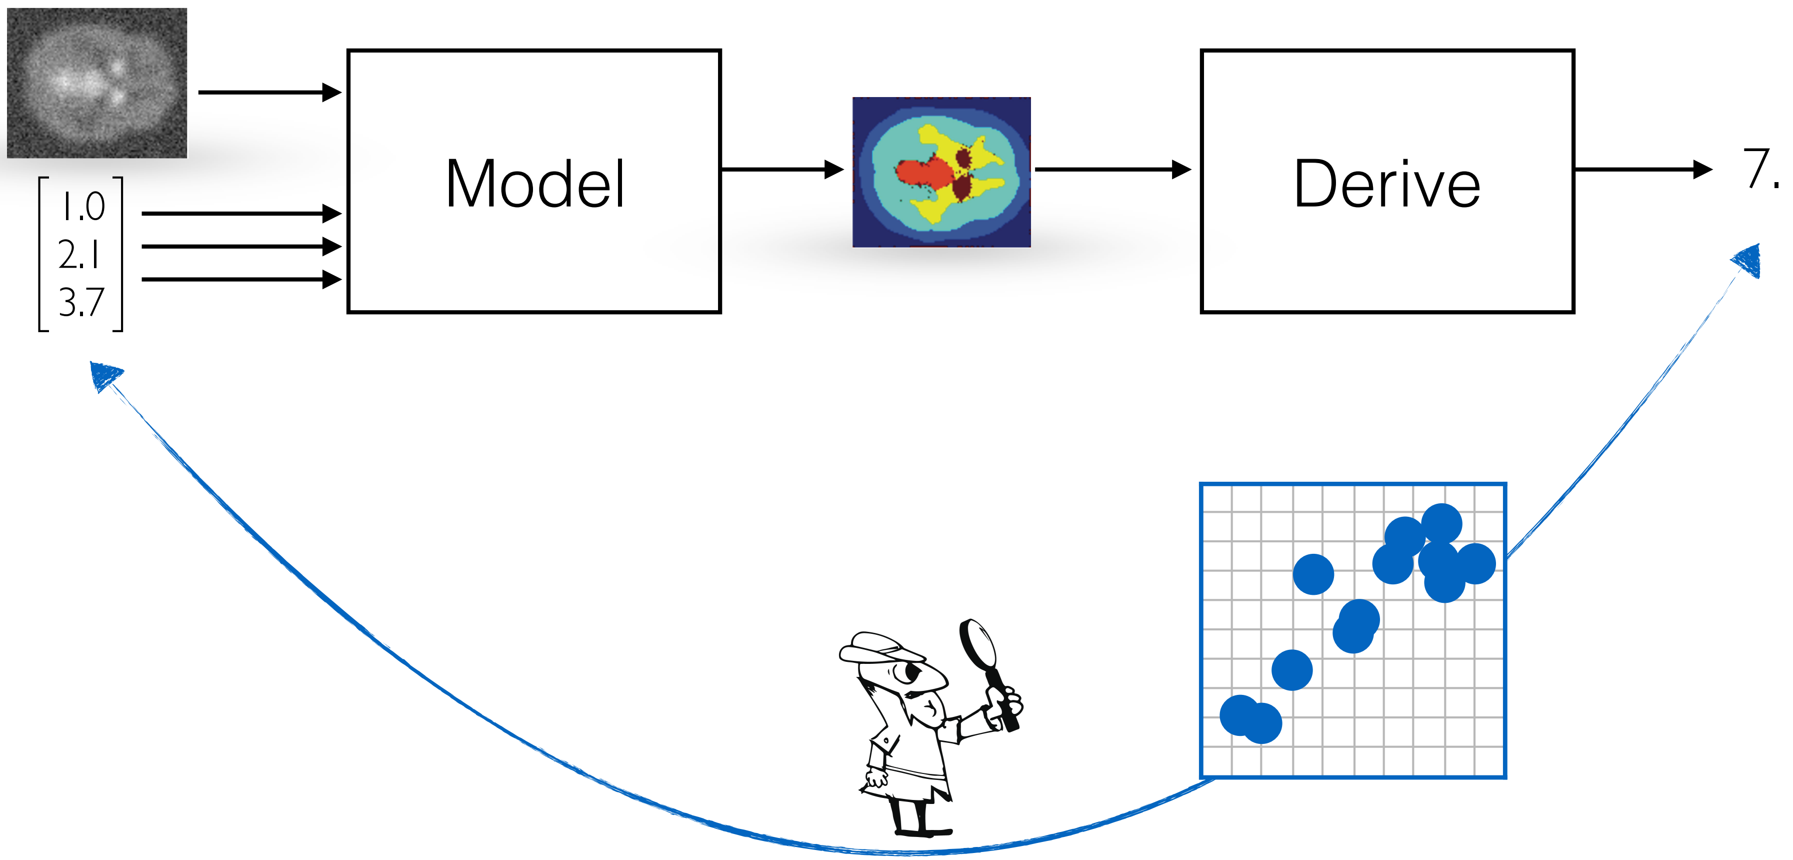

Surrogate models

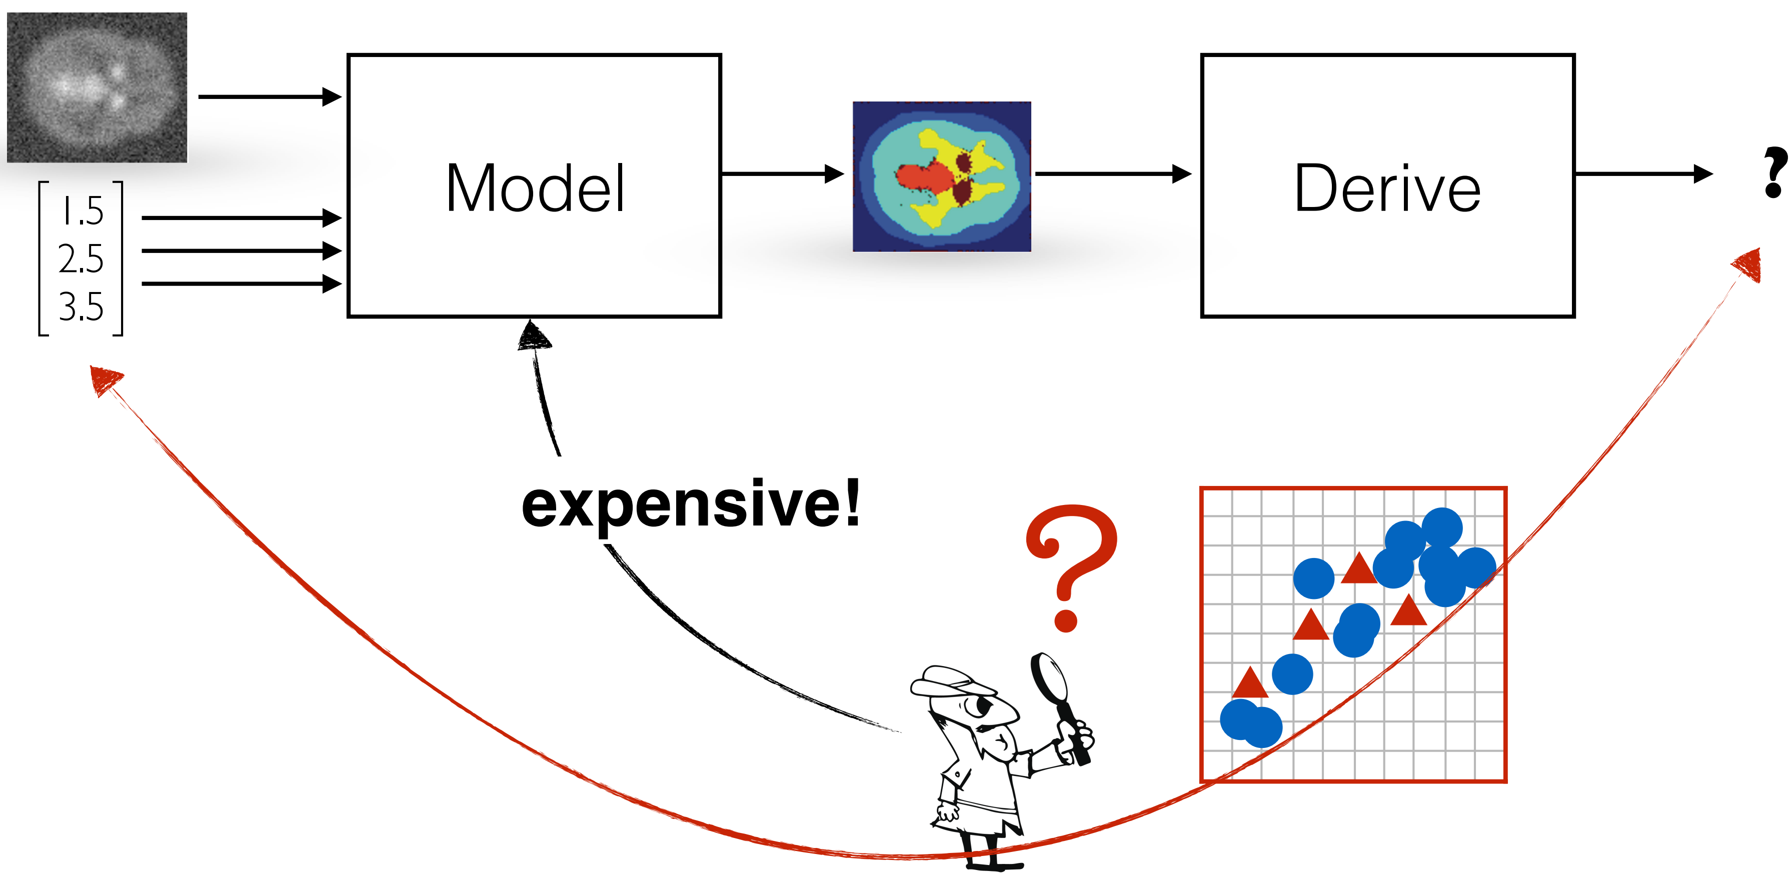

Surrogate models

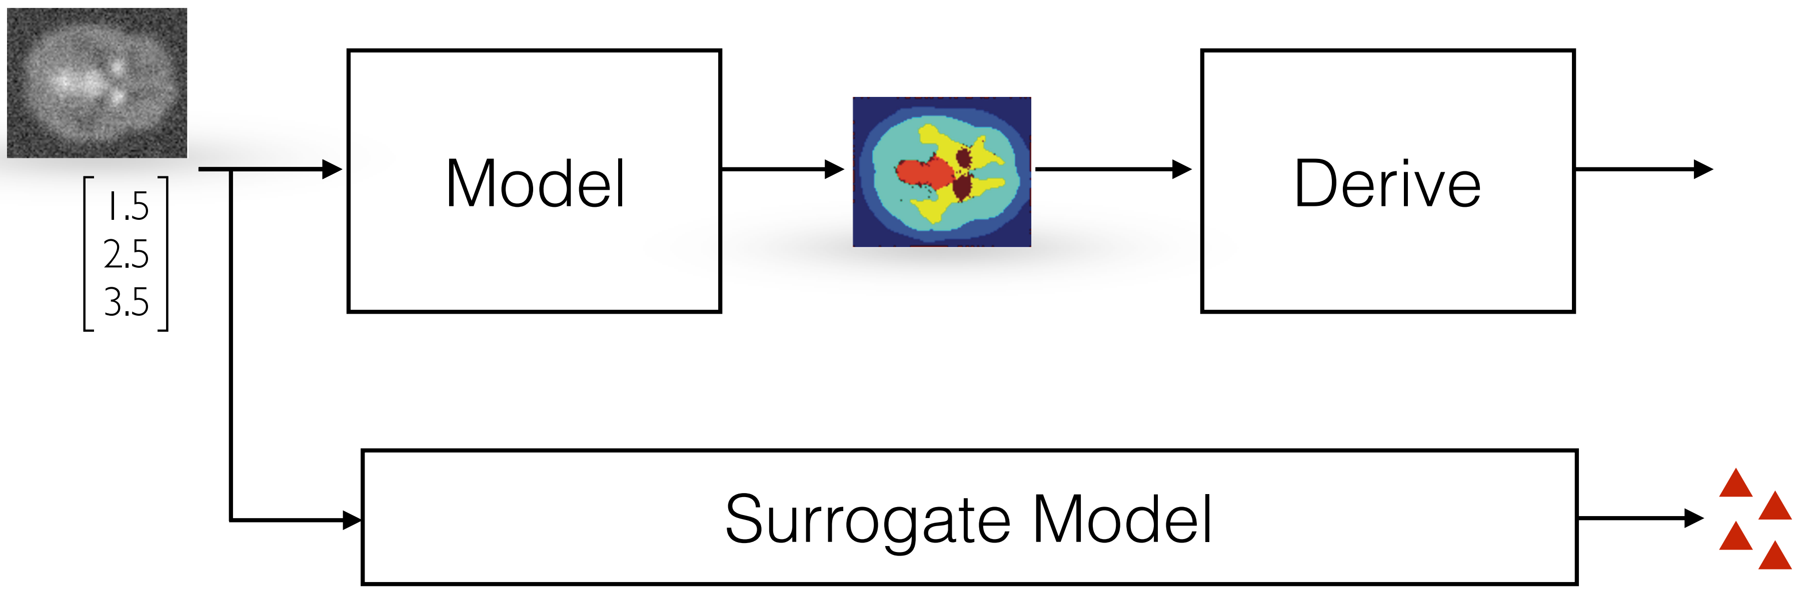

Data flow model

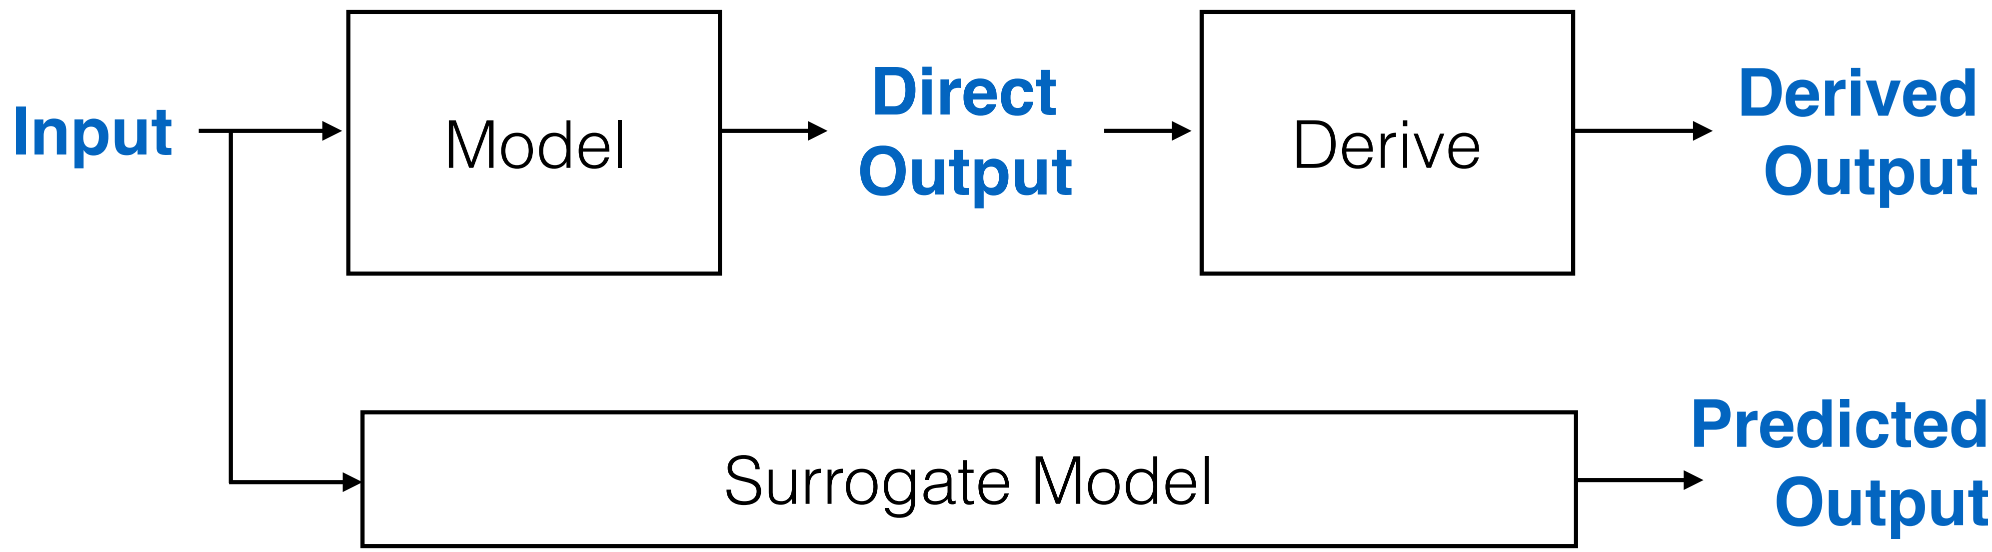

Navigation Strategies

Navigation strategies

- Trial and error (traditional approach)

Navigation strategies

- Trial and error (traditional approach)
- Local —> global tweaking

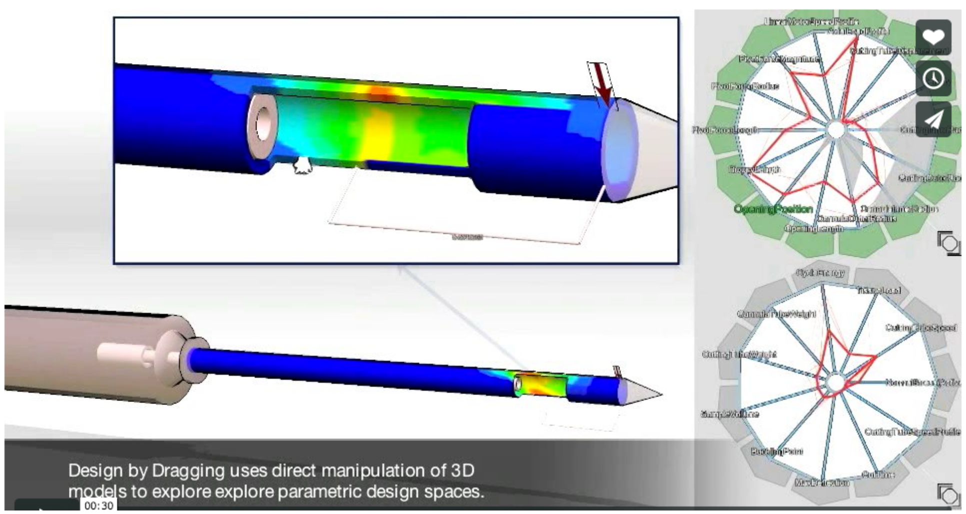

Design by Dragging
[Coffey et al., SciVis 2013]

Navigation strategies

- Trial and error (traditional approach)
- Local —> global tweaking
- Global —> local exploration
  - FluidExplorer, Vismon, Tuner
  - many others: Paramorama [Pretorius et al., InfoVis 2011]

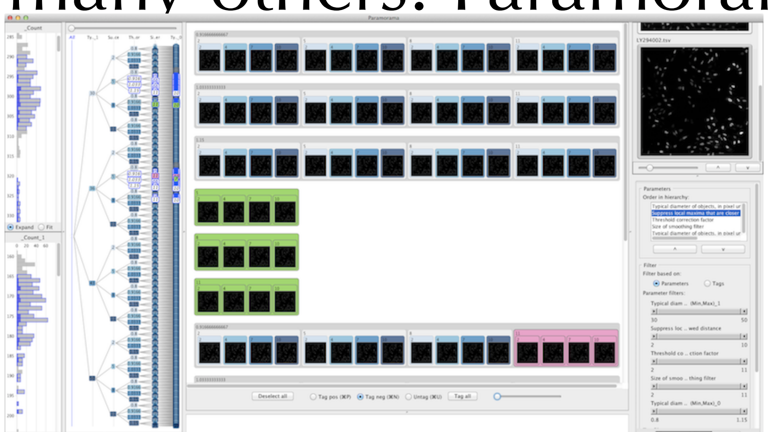

Navigation strategies

- Trial and error (traditional approach)
- Local —> global tweaking
- Global —> local exploration
- Steering
  - simulation steering: e.g. real-time simulators
  - computational steering: e.g. change the grid size, stop if no insight

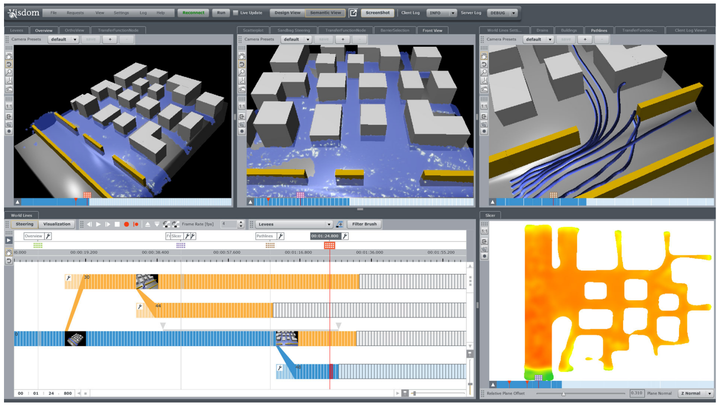

World Lines
[Waser et al., Vis 2010]

Analysis Tasks

Analysis tasks

- Optimization
- Partitioning
- Fitting
- Outliers
- Uncertainty
- Sensitivity

Analysis tasks

- **Optimization**
- Partitioning
- Fitting
- Outliers
- Uncertainty
- Sensitivity

Find the best parameter combination given some objectives.

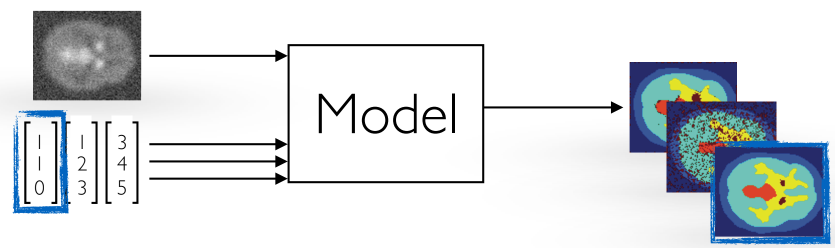

in 19/21 papers

Analysis tasks

- Optimization
- **Partitioning** aka clustering
- Fitting
- Outliers
- Uncertainty
- Sensitivity

How many different types of model behaviors are possible?

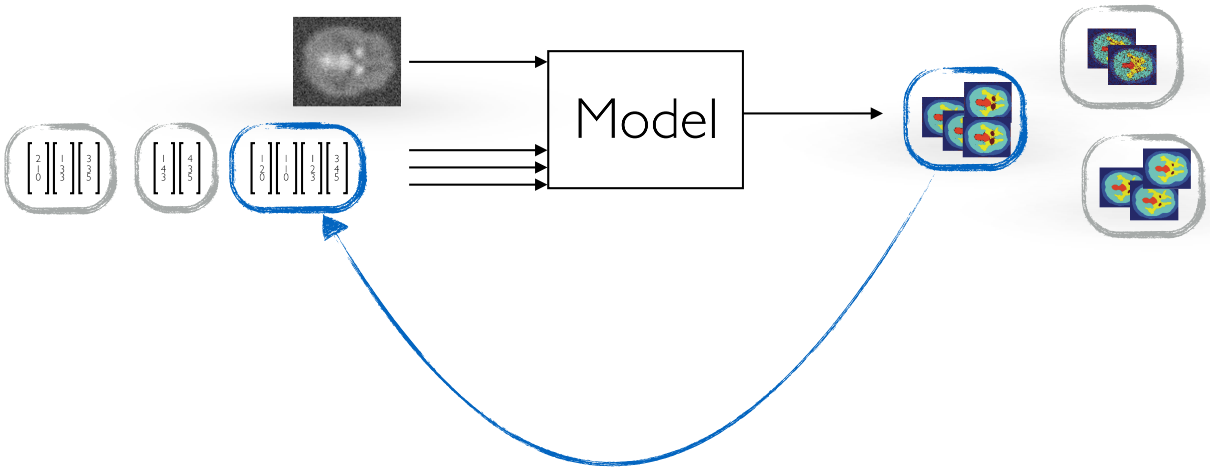

in 6/21 papers

Analysis tasks

Where in the input parameter space would actual measured data occur?

- Optimization
- Partitioning
- **Fitting** aka regression analysis
- Outliers
- Uncertainty
- Sensitivity

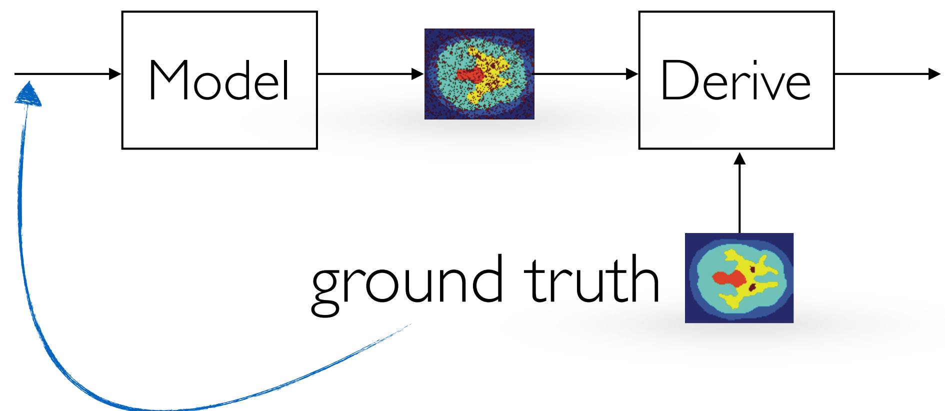

in 9/21 papers

Analysis tasks

- Optimization
- Partitioning
- Fitting
- **Outliers**
- Uncertainty
- Sensitivity

What outputs are special?

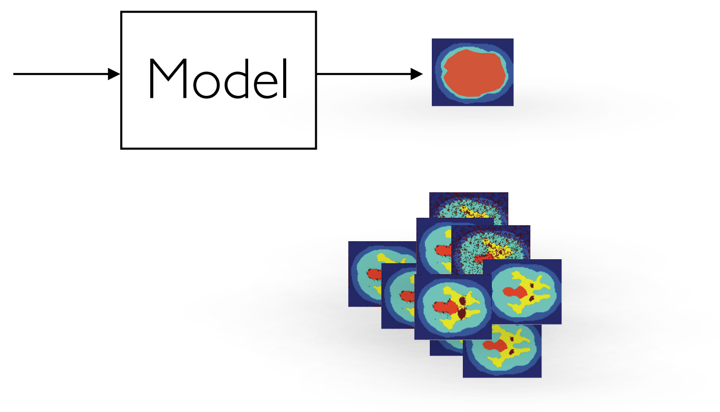

in 9/21 papers

Analysis tasks

- Optimization
- Partitioning
- Fitting
- Outliers
- **Uncertainty**
- Sensitivity

How reliable is the output?

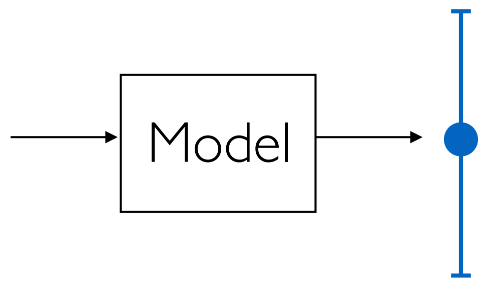

- model vs. reality
- non-deterministic model
- model vs. surrogate

in 7/21 papers

Analysis tasks

- Optimization
- Partitioning
- Fitting
- Outliers
- Uncertainty
- **Sensitivity**

What ranges/variations of outputs to expect with changes of input?

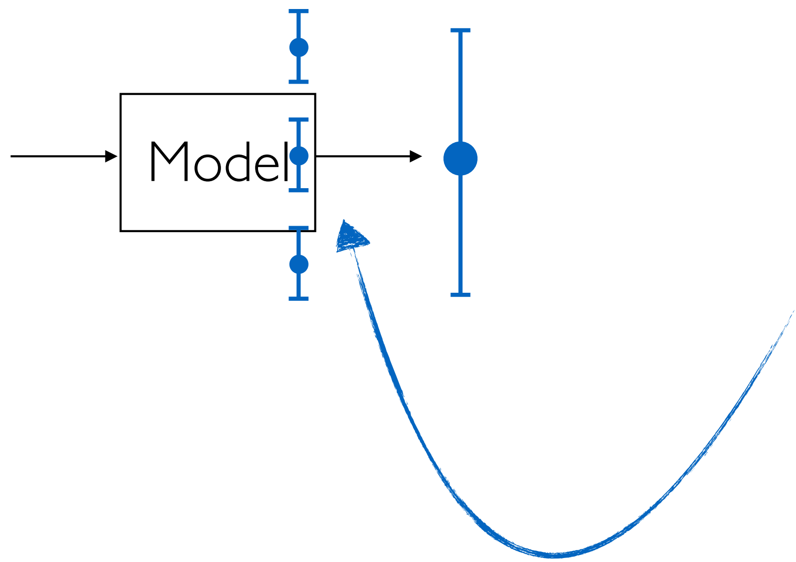

in 14/21 papers

The (machine) learning process

types of learning

- regression
- classification (supervised)
- clustering (unsupervised)
- (dimensionality reduction)
- (outlier detection)

techniques of learning

- Neural Networks (plus Deep-NN)
- Kernel meth
-
-
- .

The world of ML algorithms is not as well organized in terms of strategies as it is with simulation environments. This is work in progress.

A small selection:

- confusion matrixes for classification
- deep neural nets
- understand / diagnose / refine
- Explainers
- LIME

Confusion matrix

- Google's Facet:
  - http://gifctrl.com/?g=https://3.bp.blogspot.com/-T0dTxdse9Ow/WWz0u431RpI/AAAAAAAAB5M/rBvToJjx1L0FVVpXkgNOAwzXASyZC_JWwCLcBGAs/s1600/image4.gif
  - EuroVis keynote, 2017 — https://www.youtube.com/watch?v=E70lG9-HGEM

Squares

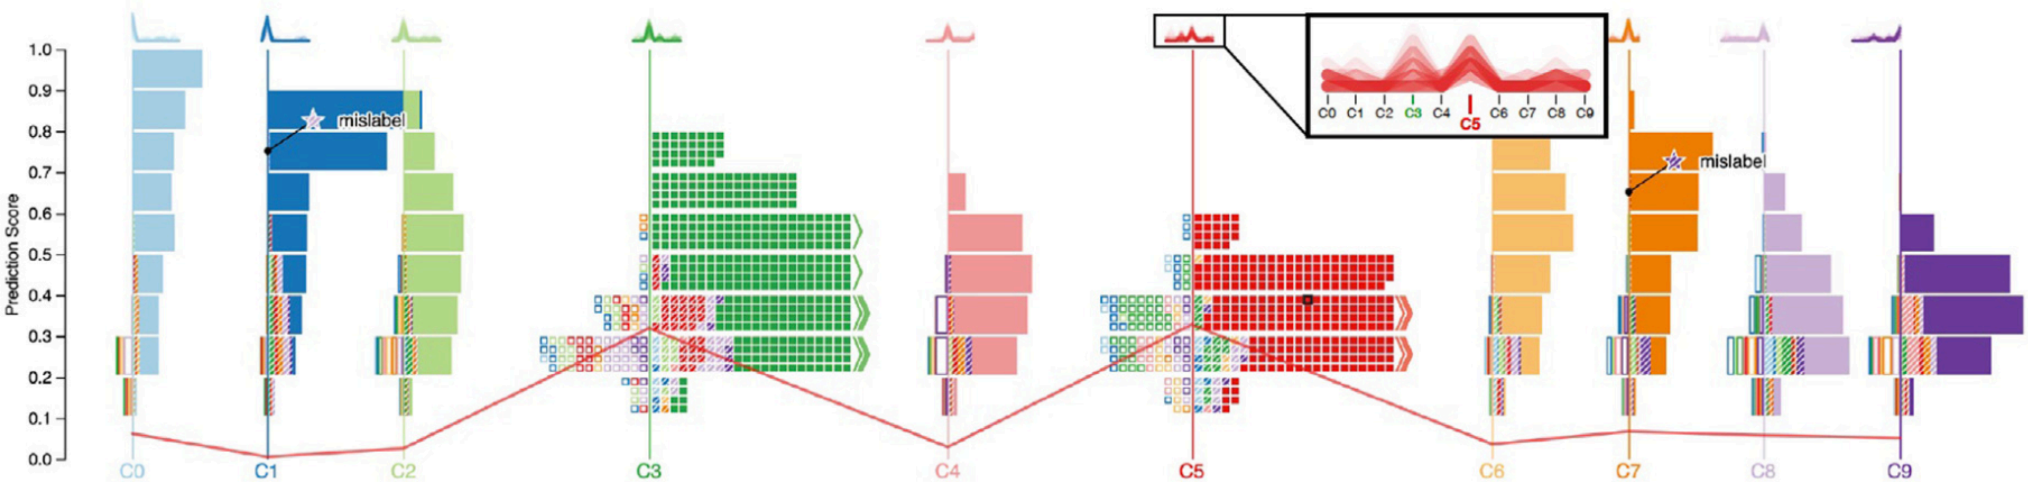

Ren et al., 2017. Squares: Supporting inter- active performance analysis for multiclass classifiers. IEEE TVCG 23 (1), 61–70.

Deep NN's: Neurons — point based

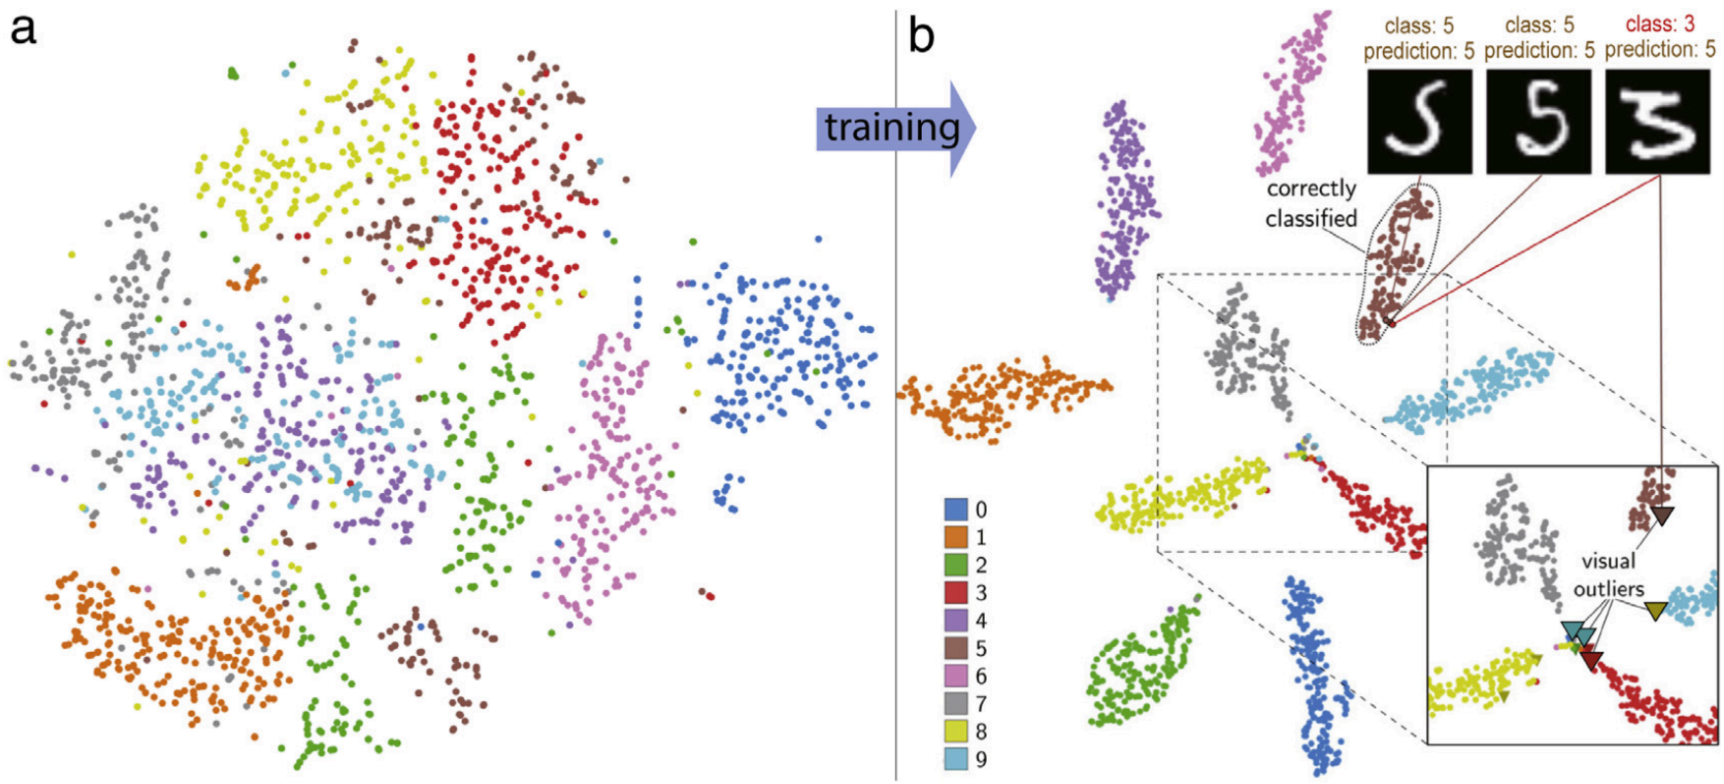

Rauber, et al., 2017. Visualizing the hidden activity of artificial neural networks. IEEE TVCG 23 (1), 101–110

Deep NN's: Neurons — network based

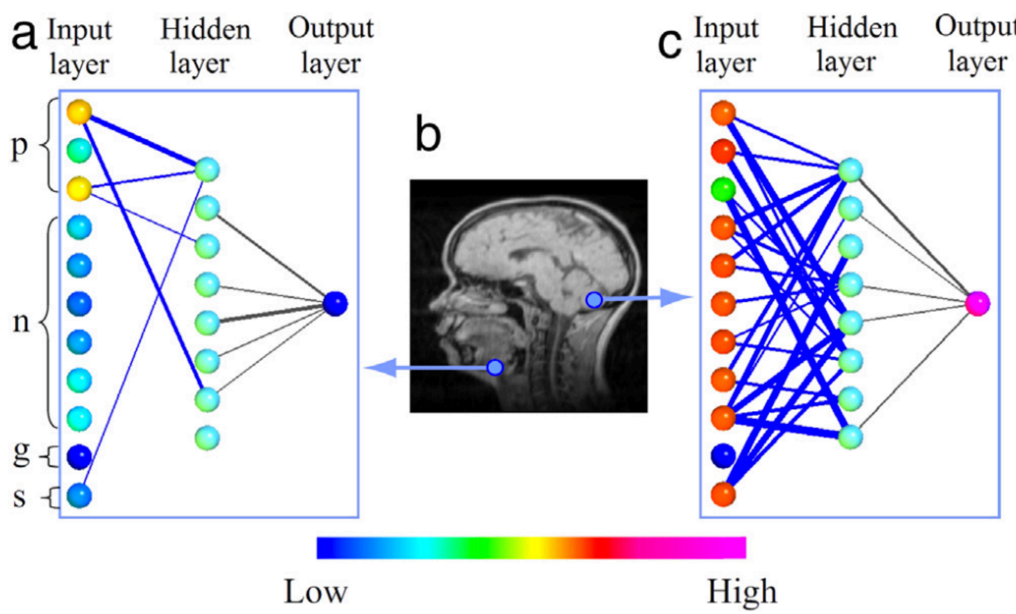

Tzeng, F.Y., Ma, K.L. 2005. Opening the black box - data driven visualization of neural networks. In: IEEE Visualization

CNNVis

http://shixialiu.com/publications/cnnvis/demo/

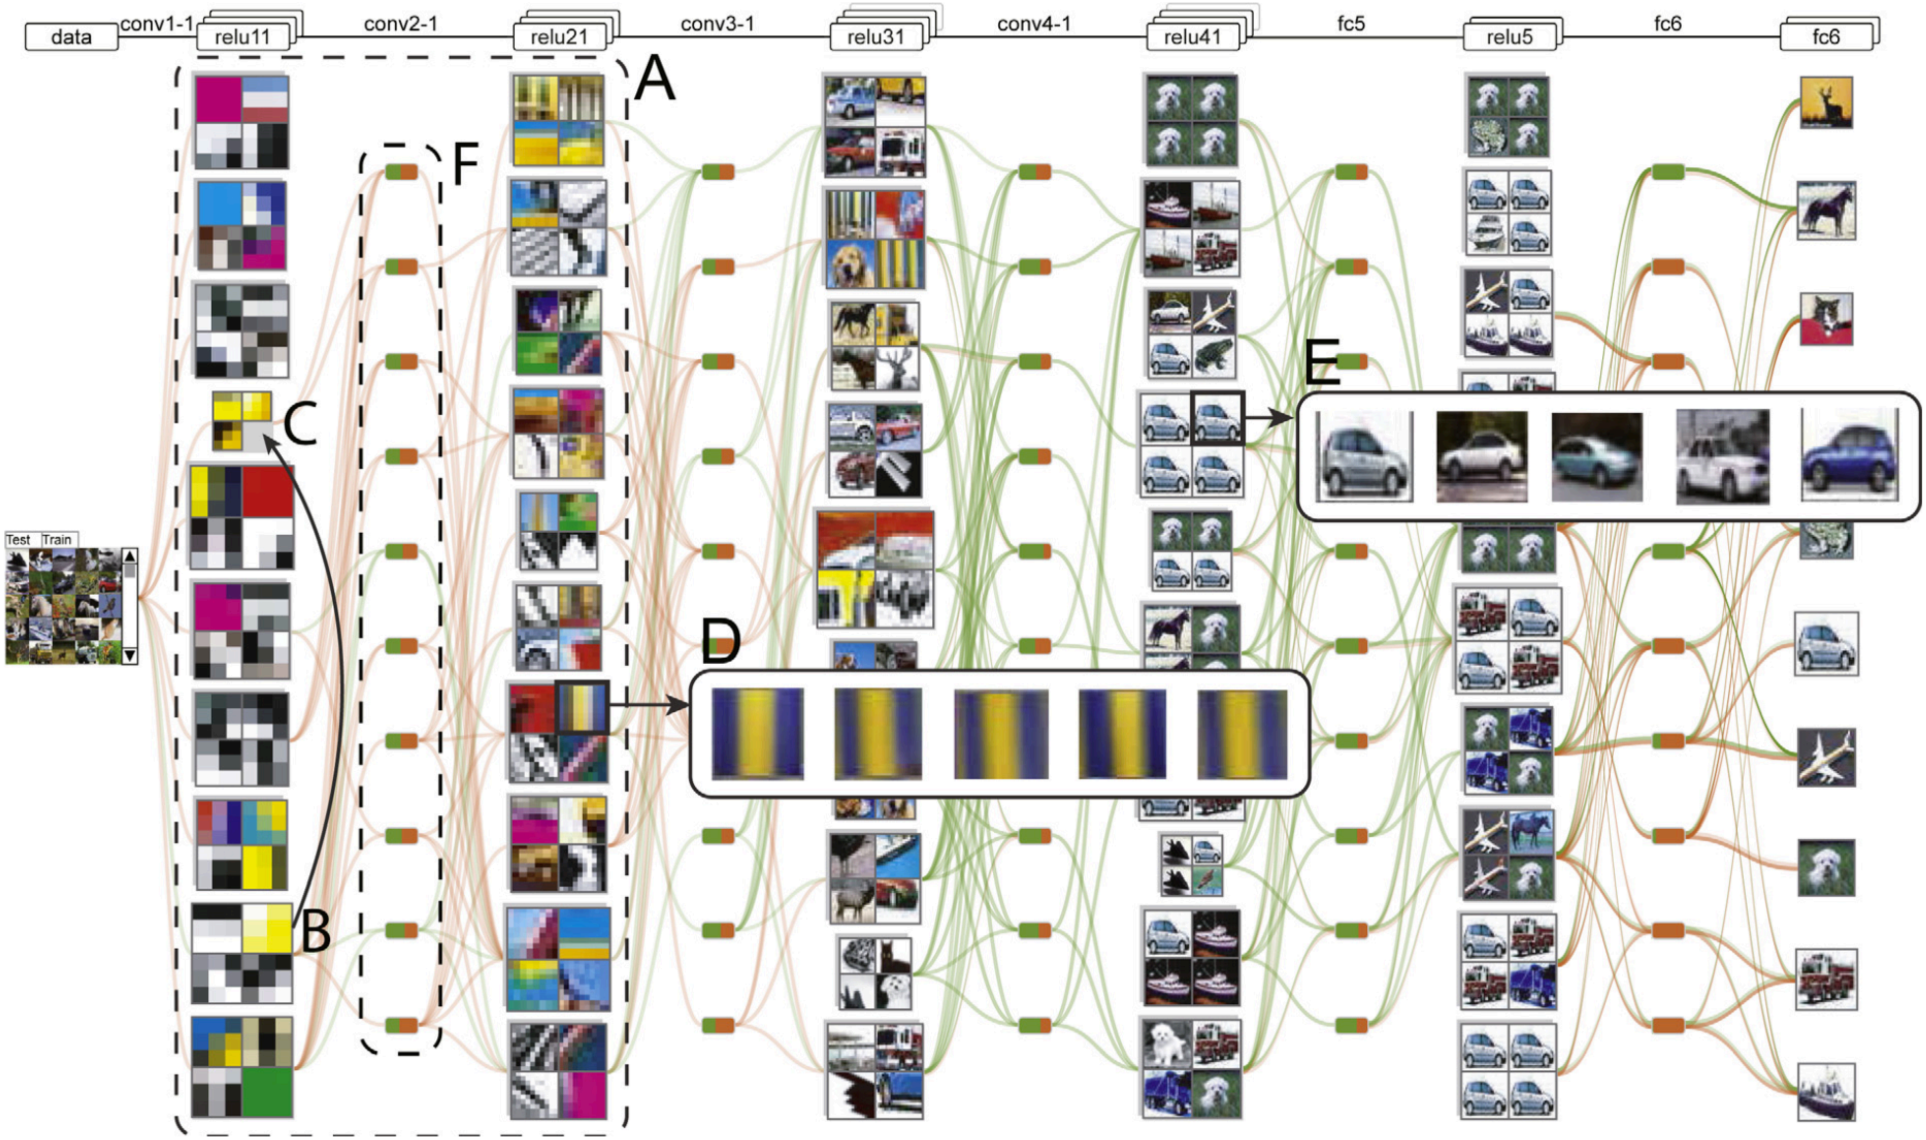

Conclusions

- Why explainable?
  - improve algorithms
  - trust
  - bridge the model builder / model usage gap
  - ethics and law
- How?
  - characterization of input-output relationships OR parameter tuning
  - understanding the behaviour of neurons in Deep NN
  - It is the "wild west" in terms of understanding machine learning models!

Acknowledgments

Steven Bergner
SFU

Maryam Booshehrian
Muprime Tech

Stephen Ingram
Coho Data

Tom Torsney-Weir
U of Vienna

Hamid Younesy
SFU

Lorenz Linhardt
ETH Zurich

Stefan Bruckner
U of Bergen

Tamara Munzner
UBC

Melanie Tory
Tableau

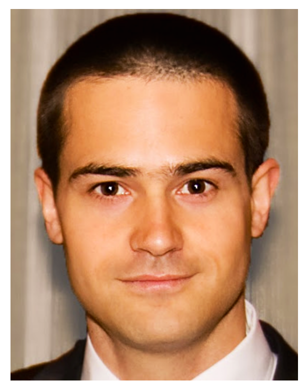

Harald Piringer
VRVis

Michael Sedlmair
U of Vienna

Patrick Wolf
Software Dev

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.
- Google Facet, https://pair-code.github.io/facets/, Jul 2017.
- Squares: Supporting Interactive Performance Analysis for Multiclass Classifiers, D. Ren, S. Amershi, B. Lee, J. Suh and J. D. Williams, IEEE Transactions on Visualization and Computer Graphics, vol. 23, no. 1, pp. 61-70, Jan. 2017.
- Visualizing the Hidden Activity of Artificial Neural Networks, P. E. Rauber, S. G. Fadel, A. X. Falcão and A. C. Telea, IEEE Transactions on Visualization and Computer Graphics, vol. 23, no. 1, pp. 101-110, Jan. 2017.
- Towards Better Analysis of Deep Convolutional Neural Networks. Mengchen Liu, Jiaxin Shi, Zhen Li, Chongxuan Li, Jun Zhu, and Shixia Liu. IEEE Transactions on Visualization and Computer Graphics 23, 1 (January 2017), 91-100.

Questions?

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

torsten.moeller@univie.ac.at