Function learning and data science

Christopher Lucas
Institute for Language, Cognition and Computation
School of Informatics, University of Edinburgh

Function learning

Regression by people:

- Intuitions about patterns in data

Global mean temperature over time

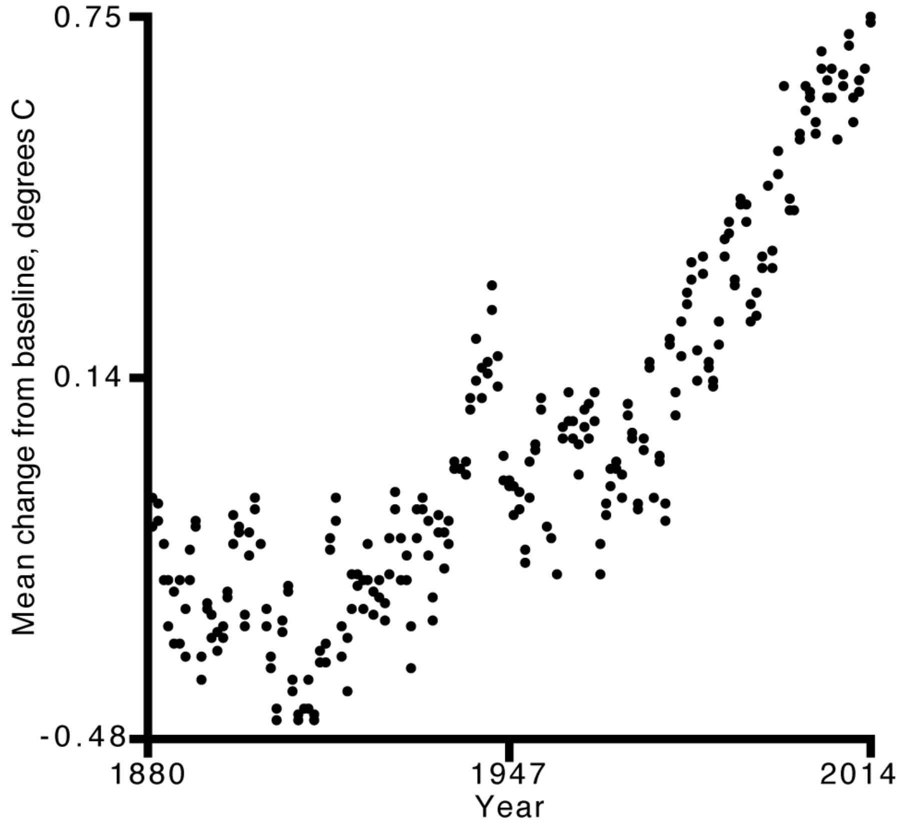

(data from http://data.okfn.org/data/core/global-temp)

Function learning

Regression by people:

- Beliefs about everyday relationships
  - Acceleration of a car: f(force on pedal)
  - Tastiness of food: f(salt)

Why should we care?

(1) We want to understand human cognition:

- Basic curiosity: how does the mind work?
- Understanding human foibles and errors, and compensating for them
- Individual differences: when, how, and why do people differ in their inferences/predictions?

Why should we care?

(2) We want to reproduce and surpass human abilities:

- How can we exploit contextual/domain knowledge?
- How can we share knowledge/information across distinct problems?
- How can we automate data analysis and prediction from end-to-end?

The role of data science

(1) Better data analysis and experiments.

(2) Building models to explain, predict, and replicate human behaviour.

The role of data science

(1) Better data analysis and experiments.

(2) Building models to explain, predict, and replicate human behaviour.

The role of data science

Limitations of most past research:

- Narrow conclusions: usually that one quantity is larger than another (using thousands of data points)

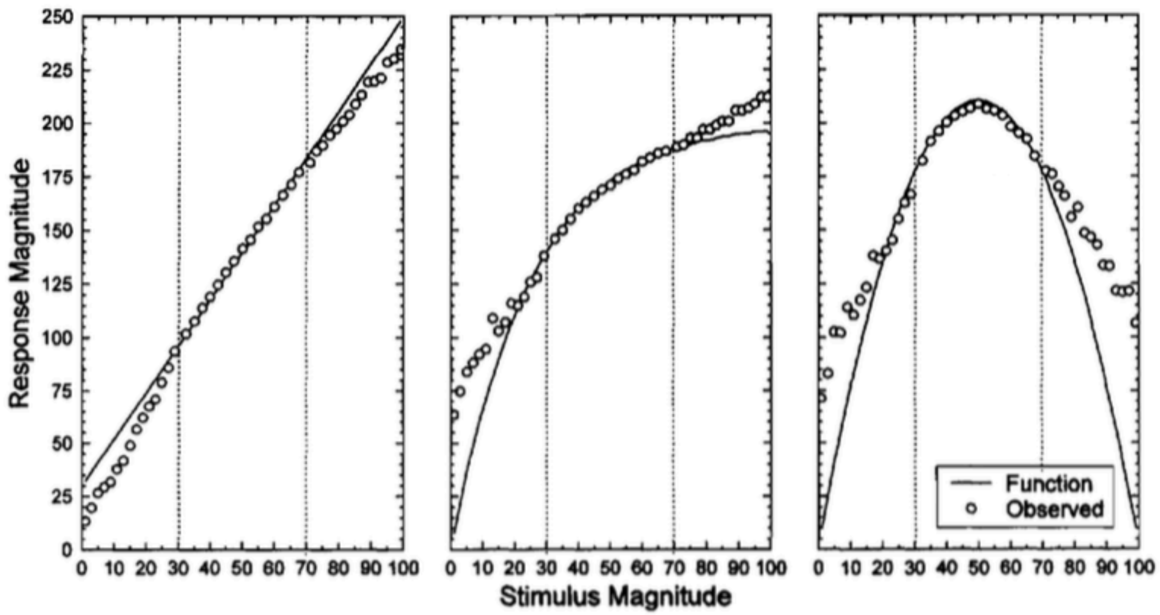

(Delosh et al., 1997)

The role of data science

Limitations of most past research:

- “Linear relationships are easy to learn!”
- “Non-monotonic relationships are hard to learn!”
- “Periodic functions are harder!”
- “The order of data presentation matters!”
- “Context matters!”

Theories are often qualitative and under-determined.

(Busemeyer et al., 1997)

The role of data science

Limitations of most past research:

- **No attempt to predict actual human judgments.**

Instead:

- Predict *average* judgments across many people
- Comparing error rates: what's harder/easier

(This makes it easier to get correlations of .99)

The role of data science

Why aren't things better?

(1) Inadequate models.
(2) Psychological legacy/tradition.
(3) No incentive to hold oneself to a higher standard.

The role of data science

Data science to the rescue!

(1) Better models, e.g., novel compositional and non-parametric approaches

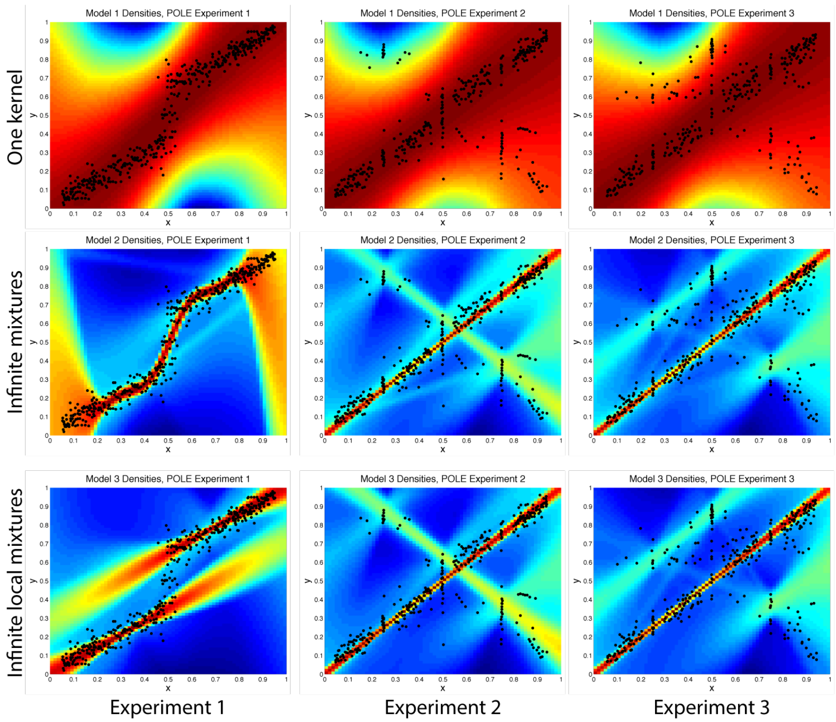

(Lucas et al., 2015)

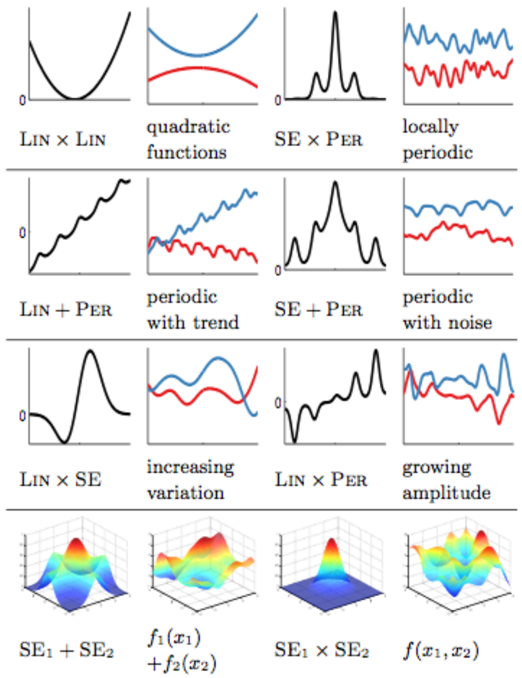

Figure 2. Examples of structures expressible by composite kernels. Left column and third columns: composite kernels $k(\cdot, 0)$. Plots have same meaning as in Figure 1.

(Duvenaud et al., 2015)

The role of data science

Data science to the rescue!

(2) Set aside psychological tradition.

- A new reference point: shared machine learning tasks.
  - Netflix challenge
  - Kaggle
- Held out test data:
  - level playing field
  - no post-hoc analysis

The role of data science

Data science to the rescue!

(2) Set aside psychological tradition.

- Try to predict individual human judgments

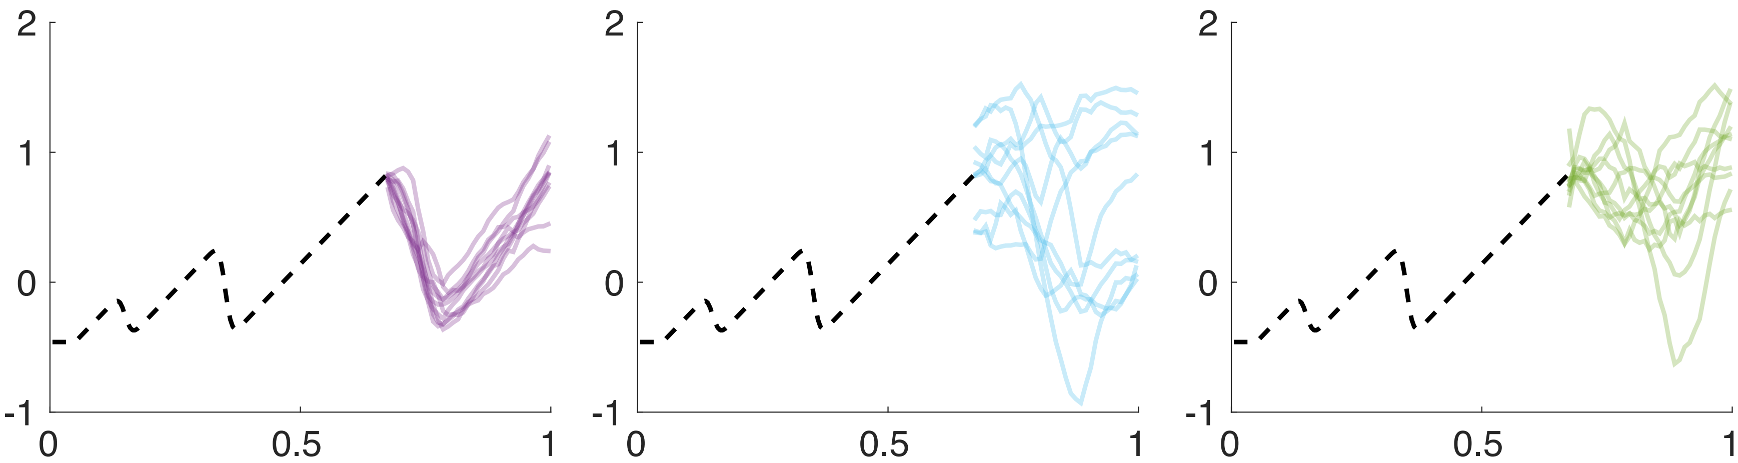

(Wilson et al., 2015)

Practical challenges

- Infrastructure for data collection and coordinating the shared task
- Model evaluation and comparison
  - Point estimates versus probabilities
  - Noise/error models
  - Converting qualitative statements to priors
- Developing new models

The role of data science

(1) Better data analysis and experiments.

(2) Building models to explain, predict, and replicate human behaviour.

Goals for models

- Order effects:
  - Regression from streaming data
  - Regression with limited memory
- Sharing information between tasks
- Discovering structural features of functions
- Using outside information, e.g., natural-language descriptions

Further reading

Bell, R. M., & Koren, Y. (2007). Lessons from the Netflix Prize Challenge. *SIGKDD Explor. Newsl.*, *9*(2), 75–79. http://doi.org/10.1145/1345448.1345465

Duvenaud, D., Lloyd, J. R., Grosse, R., Tenenbaum, J. B., & Ghahramani, Z. (2013). Structure Discovery in Nonparametric Regression through Compositional Kernel Search. In *Proceedings of the 30th International Conference on Machine Learning*. Department of Engineering, University of Cambridge.

Griffiths, T. L., Lucas, C. G., Williams, J. J., & Kalish, M. L. (2009). Modeling human function learning with Gaussian processes. *Advances in Neural Information Processing Systems*, *21*.

Lucas, C. G., Griffiths, T. L., Williams, J. J., & Kalish, M. L. (2015). A rational model of function learning. *Psychonomic Bulletin & Review*, 1–23.

Busemeyer, J. R., Byun, E., DeLosh, E. L., & McDaniel, M. A. (1997). Learning functional relations based on experience with input-output pairs by humans and artificial neural networks. In K. Lamberts & D. Shanks (Eds.), *Concepts and Categories* (pp. 405–437). Cambridge: MIT Press.

DeLosh, E. L., Busemeyer, J. R., & McDaniel, M. A. (1997). Extrapolation: The sine qua non of abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 968–986.

Wilson, A. G., Dann, C., Lucas, C. G., & Xing, E. P. (2015). The Human Kernel. *abs/1510.07389*. http://arxiv.org/abs/1510.07389

https://xkcd.com/375/

Images

Salted caramels (CC2.0);
Acceleration (CC2.0)