

The Automatic Statistician an AI for Data Science

Zoubin Ghahramani

Department of Engineering
University of Cambridge

`zoubin@eng.cam.ac.uk`

`http://learning.eng.cam.ac.uk/zoubin/`

Intelligent Machines, Nijmegen, 2015



James Robert Lloyd
Cambridge



David Duvenaud
Cambridge → Harvard



Roger Grosse
MIT → Toronto

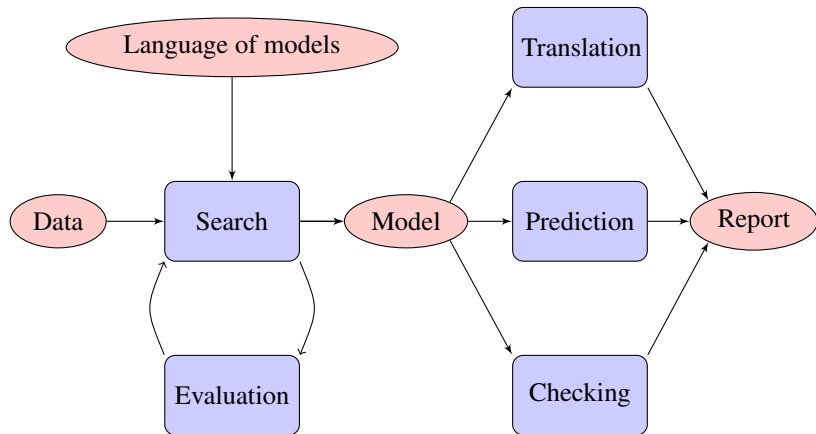


Josh Tenenbaum
MIT

THERE IS A GROWING NEED FOR DATA ANALYSIS

- ▶ We live in an era of abundant data
- ▶ The McKinsey Global Institute claim
 - ▶ *“The United States alone faces a shortage of 140,000 to 190,000 people with analytical expertise and 1.5 million managers and analysts with the skills to understand and make decisions based on the analysis of big data.”*
- ▶ Diverse fields increasingly relying on expert statisticians, machine learning researchers and data scientists e.g.
 - ▶ Computational sciences (e.g. biology, astronomy, ...)
 - ▶ Online advertising
 - ▶ Quantitative finance
 - ▶ ...


WHAT WOULD AN AUTOMATIC STATISTICIAN DO?



GOALS OF THE AUTOMATIC STATISTICIAN PROJECT



- ▶ Provide a set of tools for understanding data that require minimal expert input
- ▶ Uncover challenging research problems in e.g.
 - ▶ Automated inference
 - ▶ Model construction and comparison
 - ▶ Data visualisation and interpretation
- ▶ Advance the field of machine learning in general

INGREDIENTS OF AN AUTOMATIC STATISTICIAN



- ▶ **An open-ended language of models**
 - ▶ Expressive enough to capture real-world phenomena...
 - ▶ ...and the techniques used by human statisticians
- ▶ **A search procedure**
 - ▶ To efficiently explore the language of models
- ▶ **A principled method of evaluating models**
 - ▶ Trading off complexity and fit to data
- ▶ **A procedure to automatically explain the models**
 - ▶ Making the assumptions of the models explicit...
 - ▶ ...in a way that is intelligible to non-experts


PREVIEW: AN ENTIRELY AUTOMATIC ANALYSIS



Four additive components have been identified in the data

- ▶ A linearly increasing function.
- ▶ An approximately periodic function with a period of 1.0 years and with linearly increasing amplitude.
- ▶ A smooth function.
- ▶ Uncorrelated noise with linearly increasing standard deviation.

DEFINING A LANGUAGE OF MODELS




DEFINING A LANGUAGE OF REGRESSION MODELS

- ▶ Regression consists of **learning a function** $f : \mathcal{X} \rightarrow \mathcal{Y}$ from inputs to outputs from example input / output pairs
- ▶ Language should include **simple parametric forms** . . .
 - ▶ e.g. Linear functions, Polynomials, Exponential functions
- ▶ . . . as well as functions specified by **high level properties**
 - ▶ e.g. Smoothness, Periodicity
- ▶ Inference should be **tractable for all models** in language

WE CAN BUILD REGRESSION MODELS WITH GAUSSIAN PROCESSES

- ▶ GPs are distributions over functions such that any finite subset of function evaluations, $(f(x_1), f(x_2), \dots, f(x_N))$, have a joint Gaussian distribution
- ▶ A GP is completely specified by
 - ▶ Mean function, $\mu(x) = \mathbb{E}(f(x))$
 - ▶ Covariance / kernel function, $k(x, x') = \text{Cov}(f(x), f(x'))$
 - ▶ Denoted $f \sim \text{GP}(\mu, k)$



THE ATOMS OF OUR LANGUAGE

Five base kernels



Squared
exp. (SE)



Periodic
(PER)



Linear
(LIN)



Constant
(C)



White
noise (WN)

Encoding for the following types of functions



Smooth
functions



Periodic
functions



Linear
functions




Constant
functions




Gaussian
noise

THE COMPOSITION RULES OF OUR LANGUAGE


- ▶ Two main operations: addition, multiplication




$\text{LIN} \times \text{LIN}$




quadratic
functions




$\text{SE} \times \text{PER}$




locally
periodic




$\text{LIN} + \text{PER}$



periodic plus
linear trend



$\text{SE} + \text{PER}$




periodic plus
smooth trend

AN EXPRESSIVE LANGUAGE OF MODELS

Regression model	Kernel
GP smoothing	$SE + WN$
Linear regression	$C + LIN + WN$
Multiple kernel learning	$\sum SE + WN$
Trend, cyclical, irregular	$\sum SE + \sum PER + WN$
Fourier decomposition	$C + \sum \cos + WN$
Sparse spectrum GPs	$\sum \cos + WN$
Spectral mixture	$\sum SE \times \cos + WN$
Changepoints	e.g. $CP(SE, SE) + WN$
Heteroscedasticity	e.g. $SE + LIN \times WN$

Note: \cos is a special case of our version of PER



DISCOVERING A GOOD MODEL VIA SEARCH





DISCOVERING A GOOD MODEL VIA SEARCH

- ▶ Language defined as the arbitrary composition of five base kernels (WN, C, LIN, SE, PER) via three operators (+, \times , CP).
- ▶ The space spanned by this language is open-ended and can have a high branching factor requiring a judicious search
- ▶ We propose a greedy search for its simplicity and similarity to human model-building

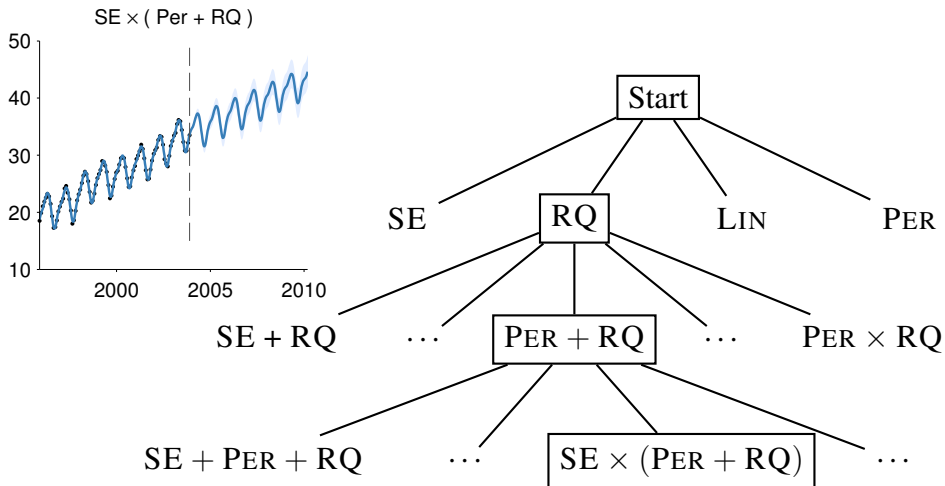
EXAMPLE: MAUNA LOA KEELING CURVE



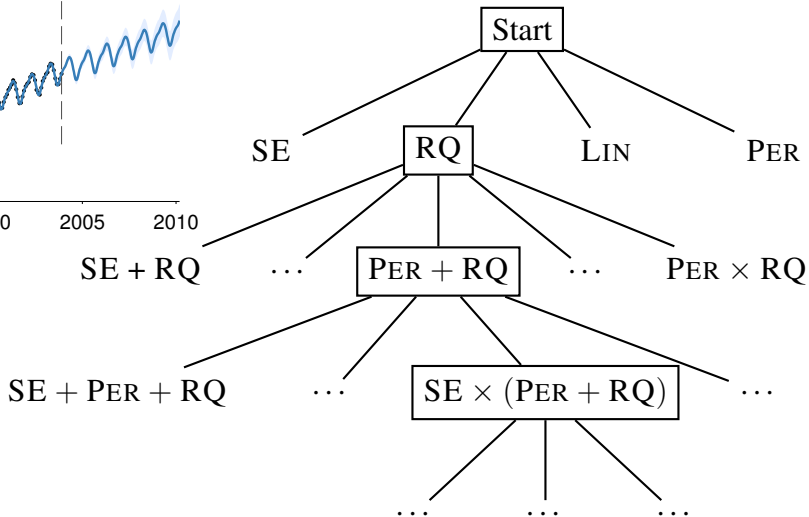
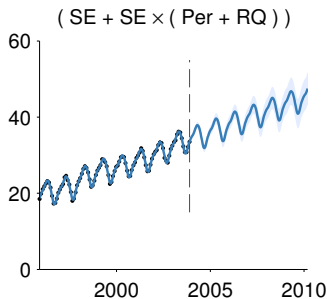
EXAMPLE: MAUNA LOA KEELING CURVE



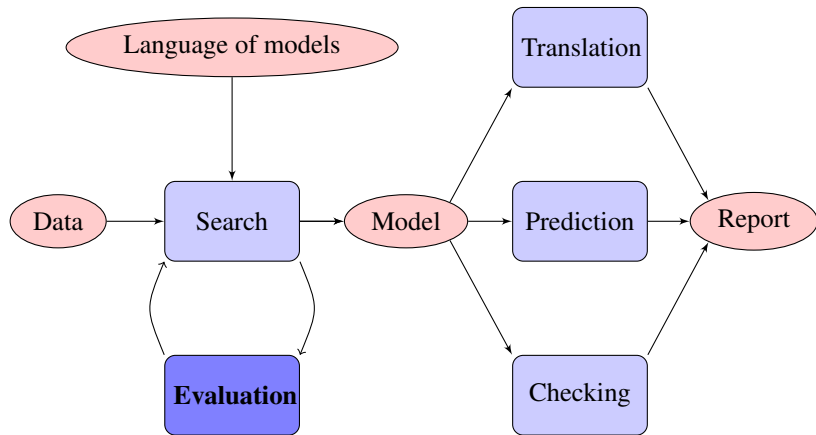
EXAMPLE: MAUNA LOA KEELING CURVE



EXAMPLE: MAUNA LOA KEELING CURVE



MODEL EVALUATION



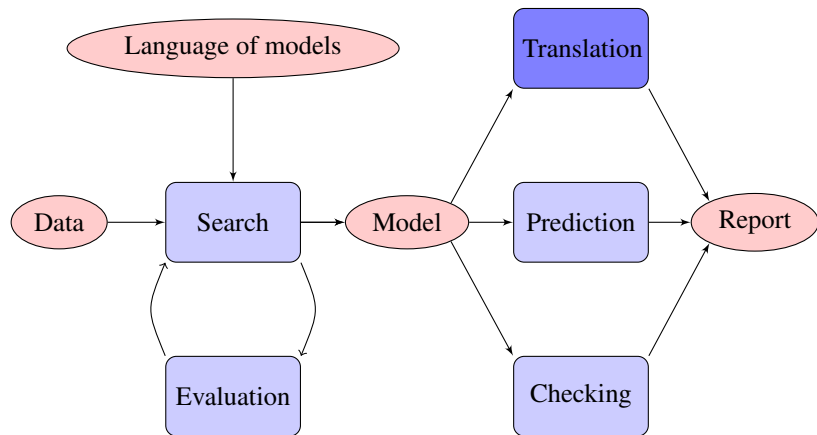
MODEL EVALUATION

- ▶ After proposing a new model its kernel parameters are optimised by conjugate gradients
- ▶ We evaluate each optimised model, M , using the **model evidence (marginal likelihood)** which can be computed analytically for GPs
- ▶ We **penalise** the marginal likelihood for the **optimised kernel parameters** using the Bayesian Information Criterion (BIC):

$$-0.5 \times \text{BIC}(M) = \log p(D | M) - \frac{p}{2} \log n$$

where p is the number of kernel parameters, D represents the data, and n is the number of data points.

AUTOMATIC TRANSLATION OF MODELS



AUTOMATIC TRANSLATION OF MODELS

- ▶ Search can produce **arbitrarily complicated models** from open-ended language but two main properties allow description to be automated
- ▶ Kernels can be **decomposed** into a **sum of products**
 - ▶ A sum of kernels corresponds to a sum of functions
 - ▶ Therefore, we can describe each product of kernels separately
- ▶ Each kernel in a product modifies a model in a **consistent** way
 - ▶ Each kernel roughly corresponds to an *adjective*

SUM OF PRODUCTS NORMAL FORM

Suppose the search finds the following kernel

$$SE \times (WN \times LIN + CP(C, PER))$$

SUM OF PRODUCTS NORMAL FORM

Suppose the search finds the following kernel

$$SE \times (WN \times LIN + CP(C, PER))$$

The changepoint can be converted into a sum of products

$$SE \times (WN \times LIN + C \times \sigma + PER \times \bar{\sigma})$$

SUM OF PRODUCTS NORMAL FORM

Suppose the search finds the following kernel

$$\text{SE} \times (\text{WN} \times \text{LIN} + \text{CP}(\text{C}, \text{PER}))$$

The changepoint can be converted into a sum of products

$$\text{SE} \times (\text{WN} \times \text{LIN} + \text{C} \times \sigma + \text{PER} \times \bar{\sigma})$$

Multiplication can be distributed over addition

$$\text{SE} \times \text{WN} \times \text{LIN} + \text{SE} \times \text{C} \times \sigma + \text{SE} \times \text{PER} \times \bar{\sigma}$$

SUM OF PRODUCTS NORMAL FORM

Suppose the search finds the following kernel

$$SE \times (WN \times LIN + CP(C, PER))$$

The changepoint can be converted into a sum of products

$$SE \times (WN \times LIN + C \times \sigma + PER \times \bar{\sigma})$$

Multiplication can be distributed over addition

$$SE \times WN \times LIN + SE \times C \times \sigma + SE \times PER \times \bar{\sigma}$$

Simplification rules are applied


$$WN \times LIN + SE \times \sigma + SE \times PER \times \bar{\sigma}$$

SUMS OF KERNELS ARE SUMS OF FUNCTIONS

If $f_1 \sim \text{GP}(0, k_1)$ and independently $f_2 \sim \text{GP}(0, k_2)$ then

$$f_1 + f_2 \sim \text{GP}(0, k_1 + k_2)$$

e.g.







We can therefore describe each component separately

PRODUCTS OF KERNELS

PER
periodic function


On their own, each kernel is described by a standard noun phrase



PRODUCTS OF KERNELS - SE

$$\underbrace{\text{SE}}_{\text{approximately}} \times \underbrace{\text{PER}}_{\text{periodic function}}$$


Multiplication by SE removes long range correlations from a model since $\text{SE}(x, x')$ decreases monotonically to 0 as $|x - x'|$ increases.



PRODUCTS OF KERNELS - LIN

$\underbrace{\text{SE}}_{\text{approximately}} \times \underbrace{\text{PER}}_{\text{periodic function}} \times \underbrace{\text{LIN}}_{\text{with linearly growing amplitude}}$


Multiplication by LIN is equivalent to multiplying the function being modeled by a linear function. If $f(x) \sim \text{GP}(0, k)$, then $xf(x) \sim \text{GP}(0, k \times \text{LIN})$. This causes the standard deviation of the model to vary linearly without affecting the correlation.




PRODUCTS OF KERNELS - CHANGEPOINTS

$\underbrace{\text{SE}}$ \times $\underbrace{\text{PER}}$ \times $\underbrace{\text{LIN}}$ \times $\underbrace{\sigma}$
approximately periodic function with linearly growing amplitude until 1700



Multiplication by σ is equivalent to multiplying the function being modeled by a sigmoid.



AUTOMATICALLY GENERATED REPORTS



EXAMPLE: AIRLINE PASSENGER VOLUME





Four additive components have been identified in the data

- ▶ A linearly increasing function.
- ▶ An approximately periodic function with a period of 1.0 years and with linearly increasing amplitude.
- ▶ A smooth function.
- ▶ Uncorrelated noise with linearly increasing standard deviation.



EXAMPLE: AIRLINE PASSENGER VOLUME

This component is linearly increasing.





EXAMPLE: AIRLINE PASSENGER VOLUME

This component is approximately periodic with a period of 1.0 years and varying amplitude. Across periods the shape of this function varies very smoothly. The amplitude of the function increases linearly. The shape of this function within each period has a typical lengthscale of 6.0 weeks.





EXAMPLE: AIRLINE PASSENGER VOLUME

This component is a smooth function with a typical lengthscale of 8.1 months.



EXAMPLE: AIRLINE PASSENGER VOLUME

This component models uncorrelated noise. The standard deviation of the noise increases linearly.




OTHER EXAMPLES

See <http://www.automaticstatistician.com>

GOOD PREDICTIVE PERFORMANCE AS WELL

Standardised RMSE over 13 data sets



- ▶ Tweaks can be made to the algorithm to improve accuracy or interpretability of models produced. . .
- ▶ . . . but both methods are *highly competitive* at extrapolation (shown above) and interpolation

MODEL CHECKING AND CRITICISM

- ▶ Good statistical modelling should include model criticism:
 - ▶ Does the data match the assumptions of the model?
 - ▶ For example, if the model assumed Gaussian noise, does a Q-Q plot reveal non-Gaussian residuals?
- ▶ Our automatic statistician does posterior predictive checks, dependence tests and residual tests
- ▶ We have also been developing more systematic nonparametric approaches to model criticism using kernel two-sample testing with MMD.

Lloyd, J. R., and Ghahramani, Z. (2014) *Statistical Model Criticism using Kernel Two Sample Tests*. <http://mlg.eng.cam.ac.uk/Lloyd/papers/kernel-model-checking.pdf>

CHALLENGES

- ▶ Interpretability / accuracy
- ▶ Increasing the expressivity of language
 - ▶ e.g. Monotonicity, positive functions, symmetries
- ▶ Computational complexity of searching through a huge space of models
- ▶ Extending the automatic reports to multidimensional datasets
 - ▶ Search and descriptions naturally extend to multiple dimensions, but automatically generating relevant visual summaries harder

CURRENT AND FUTURE DIRECTIONS

- ▶ Automatic statistician for:
 - * One-dimensional time series
 - * Linear regression (classical)
 - ▶ Multivariate nonlinear regression (c.f. Duvenaud, Lloyd et al, ICML 2013)
 - ▶ Multivariate classification (w/ Mrksic), clustering and learning graphical models (w/ Smith, Lloyd)
 - ▶ Single-cell transcriptomics (gene expression) data?? (w/ Lloyd, Stegle, Buettner)
- ▶ Probabilistic programming
- ▶ Bayesian optimisation, and the rational allocation of computational resources

PROBABILISTIC PROGRAMMING

Problem: Probabilistic model development and the derivation of inference algorithms is time-consuming and error-prone.

Solution:

- ▶ Develop Turing-complete **Probabilistic Programming Languages** for expressing probabilistic models as computer programs that generate data (i.e. simulators).
- ▶ Derive **Universal Inference Engines** for these languages that sample over program traces given observed data.
- ▶ This can be used to implement search over the model space in an automatic statistician

Example languages: Church, Venture, Anglican, Stochastic Python*, ones based on Haskell*, Julia*

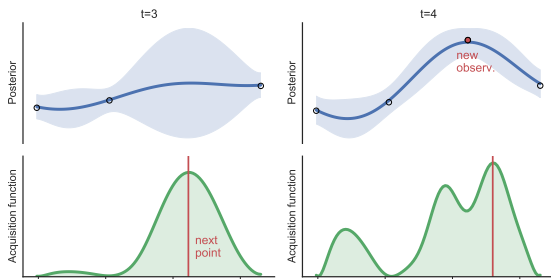
Example inference algorithms: Metropolis-Hastings MCMC, variational inference, particle filtering, slice sampling*, particle MCMC, nested particle inference*, austerity MCMC*

PROBABILISTIC PROGRAMMING

```
1 sProbs = (1.0/3, 1.0/3, 1.0/3)
2 tProbs = {0 : (0.1, 0.5, 0.4), 1 : (0.2, 0.2, 0.6),
           2 : (0.15, 0.15, 0.7)}
3 eMeans = (-1,1,0)
4
5 def hmm():
6     states = []
7     states.append(stocPy.categorical(sProbs, obs=True))
8     for ind, ob in stocPy.readCSV("hmm-data.csv"):
9         states.append(stocPy.categorical(tProbs[states[ind]],
10                                         obs=True))
11     stocPy.normal(eMeans[states[ind]], 1, cond=ob)
```

An example probabilistic program in StocPy implementing a 3-state hidden Markov model (HMM).

BAYESIAN OPTIMISATION



Problem: Global optimisation of black-box functions that are *expensive to evaluate*

Solution: treat as a problem of sequential decision-making and model uncertainty in the function.

This can speed up model search in the automatic statistician.

SUMMARY

- ▶ We have presented the beginnings of an automatic statistician
- ▶ Our system
 - ▶ Defines an open-ended language of models
 - ▶ Searches greedily through this space
 - ▶ Produces detailed reports describing patterns in data
 - ▶ Performs automatic model criticism
- ▶ Extrapolation and interpolation performance highly competitive
- ▶ We believe this line of research has the potential to make powerful statistical model-building techniques accessible to non-experts

REFERENCES

Website: <http://www.automaticstatistician.com>

Duvenaud, D., Lloyd, J. R., Grosse, R., Tenenbaum, J. B. and Ghahramani, Z. (2013) Structure Discovery in Nonparametric Regression through Compositional Kernel Search. ICML 2013.

Hernández-Lobato, J. M., Hoffman, M. W., and Ghahramani, Z. (2014) Predictive entropy search for efficient global optimization of black-box functions. NIPS 2014

Lloyd, J. R., Duvenaud, D., Grosse, R., Tenenbaum, J. B. and Ghahramani, Z. (2014) Automatic Construction and Natural-language Description of Nonparametric Regression Models AAAI 2014. <http://arxiv.org/pdf/1402.4304v2.pdf>

Lloyd, J. R., and Ghahramani, Z. (2014) Statistical Model Criticism using Kernel Two Sample Tests <http://mlg.eng.cam.ac.uk/Lloyd/papers/kernel-model-checking.pdf>

Ghahramani, Z. (2013) Bayesian nonparametrics and the probabilistic approach to modelling. *Philosophical Trans. Royal Society A* 371: 20110553.

Ranca, R. (2015) StocPy. <https://github.com/RazvanRanca/StocPy>

Ranca, R. and Ghahramani, Z. (2015) Slice sampling for probabilistic programming. <http://arxiv.org/abs/1501.04684>

Valera, I. and Ghahramani, Z. (2014) General Table Completion using a Bayesian Nonparametric Model. NIPS 2014.