

Bayesian Decision Making and Uncertainty – Opening Remarks

Andreas Krause

December 13, 2024, NeurIPS Vancouver



mushroom




Madagascar cat

[Krizhevsky et al '12]



Google Translate

KR Write a rap about submodularity

 (Verse 1)
 I'm here to talk about submodularity
 It's a concept that's often misunderstood
 But let me break it down for you
 And make it easy to understand, too


(Chorus)
 Submodularity, oh so sweet
 It's a property that can't be beat
 In optimization, it's a treat
 And it makes everything more neat

(Verse 2)
 It's a mathematical property
 That helps us make the best choice
 When we're trying to optimize
 It helps us find the optimal voice

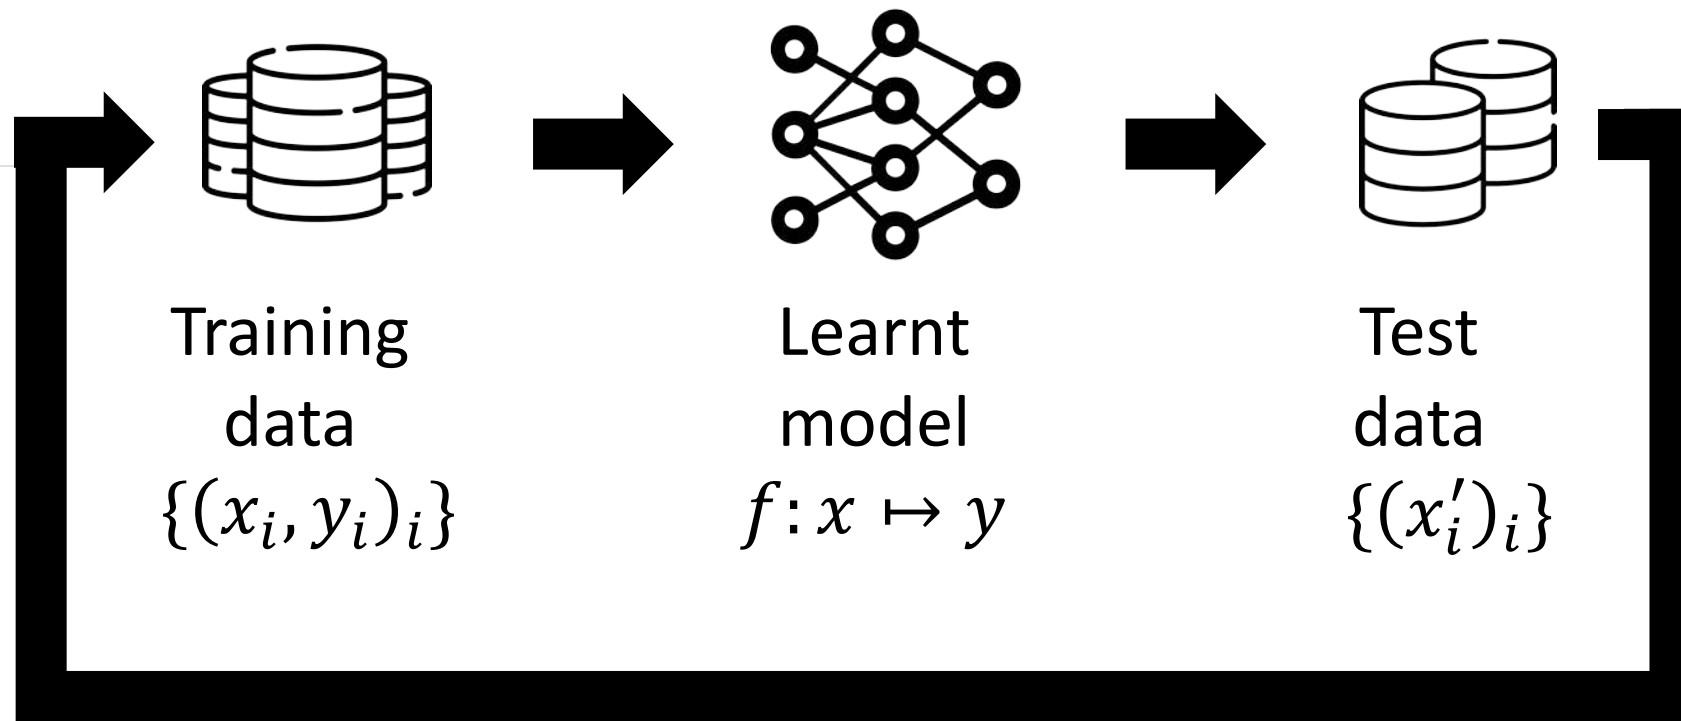
(Bridge)
 It's all about the diminishing returns
 The more you take, the less it earns
 But with submodularity, we learn
 To make the most of what we have to earn



“Photorealistic closeup video of two pirate ships battling each other as they sail inside a cup of coffee”
OpenAI Sora



DeepMind
AlphaFold



How should we intelligently gather data?

Information theory

[Shannon '48]

Active
Learning

[Freund et al '92]

Experimental
Design

[Fisher'35, Lindley '56]

POMDPs

[Sondik '71]

Optimal
Diagnosis

[Garey & Graham'74]

Artificial
Curiosity

[Schmidhuber '90]

Value of
Information

[Howard '60,
Heckerman, Horvitz '93]

Sensor

Placement

[Caselton & Zidek '84]

RL

[Sutton & Barto '98]

Bayesian

Optimization

[Moćkus et al. '78]

Bandits

[Gittins '79]

Key questions

How do we **assess utility/value** of an experiment?

- Reduction in uncertainty? Maximization of reward? Estimating integrals, ODEs?
- How do we quantify uncertainty?

According to **which protocol** do we select our experiments?

- A priori/Non-adaptive/open-loop vs sequential/adaptive/closed-loop?
- Constraints in the selection? E.g., cost, movement, ...

How do we **find an optimal set / sequence** of experiments?

- Convex relaxations?
- Submodular optimization?

How much does optimal design **help**?

- Sample complexity, regret bounds, adaptivity gap, robustness, ...


...

How should we collect data to maximally reduce uncertainty?

Optimal experimental design



Charles Sanders Peirce
1876



Bayesian experimental design

ON A MEASURE OF THE INFORMATION PROVIDED BY AN EXPERIMENT^{1, 2}

By D. V. LINDLEY

University of Cambridge and University of Chicago

1. Summary. A measure is introduced of the information provided by an experiment. The measure is derived from the work of Shannon [10] and involves the knowledge prior to performing the experiment, expressed through a prior probability distribution over the parameter space. The measure is used to compare some pairs of experiments without reference to prior distributions; this method of comparison is contrasted with the methods discussed by Blackwell. Finally, the measure is applied to provide a solution to some problems of experimental design, where the object of experimentation is not to reach decisions but rather to gain knowledge about the world.

2. Introduction. Shannon has introduced two important ideas into the theory of information in communications engineering. The first idea is that information is a statistical concept. The statistical frequency distribution of the symbols that make up a message must be considered before the notion can be discussed adequately. The second idea springs from the first and implies that on the basis of the frequency distribution, there is an essentially unique function of the distribution which measures the amount of the information. It is the purpose of the present paper to apply these two ideas to statistical theory by discussing the notion of information in an experiment, rather than in a message. The second of Shannon's ideas has been applied to statistical theory by Kullback and Leibler [6], [7], [8]; but our application is quite distinct from theirs. The interpretation of Shannon's ideas in current statistical theory has been given by McMillan [9]. The discussion in that paper is related to, and partly inspired, that given here. A referee has kindly pointed out that Shannon's theory has been applied in psychometric problems by L. J. Cronbach in an unpublished report [14]. Definition 2, in particular, is used by Cronbach.

The situation in communications engineering is that there is a transmitted message, x , which is received as a message, y . By considerations of the informations in x and y it is possible to discuss the rate at which information has been transmitted along the channel. The analogous description in statistical theory is provided by replacing x by the knowledge of the state of nature, usually expressed by the knowledge of a finite number of parameters, prior to an experiment, and by replacing y by the knowledge after the experiment. The com-

Statistical Science
1995, Vol. 10, No. 3, 275-304

Bayesian Experimental Design: A Review

Kathryn Chaloner and Isabella Verdinelli

Abstract. This paper reviews the literature on Bayesian experimental design. A unified view of this topic is presented, based on a decision-theoretic approach. This framework casts criteria from the Bayesian literature of design as part of a single coherent approach. The decision-theoretic structure incorporates both linear and nonlinear design problems and it suggests possible new directions to the experimental design problem, motivated by the use of new utility functions. We show that, in some special cases of linear design problems, Bayesian solutions change in a sensible way when the prior distribution and the utility function are modified to allow for the specific structure of the experiment. The decision-theoretic approach also gives a mathematical justification for selecting the appropriate optimality criterion.

Key words and phrases: Decision theory, hierarchical linear models, logistic regression, nonlinear design, nonlinear models, optimal design, optimality criteria, utility functions.

1. INTRODUCTION

1.1 Experimental Design

Experimental design involves the specification of all aspects of an experiment. Common sense, available resources and knowledge of the motivation for carrying out the experiment often help in selecting important features that depend on the specific problem. Not all aspects of experimental design are susceptible to abstract mathematical treatment, but the choice of values for those variables that are under the control of the experimenter can be simply expressed in a mathematical framework. This problem has been considered at length in the scientific literature and is focused on in this paper.

In designing an experiment, decisions must be made before data collection, and data collection is restricted by limited resources. Because specific information is usually available prior to experimen-

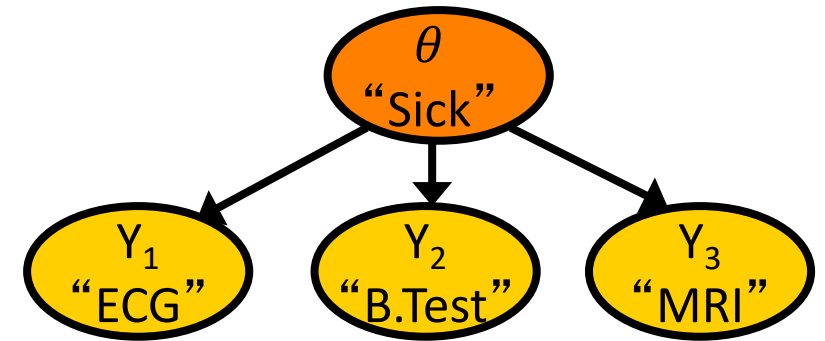
tation and, indeed, often motivates the experiment, Bayesian methods can play an important role. Bayesian decision theory also sharpens thinking on the purpose of the experiment. Like most areas of Bayesian statistics, Bayesian experimental design has gained popularity in the past two decades, but like many areas of Bayesian statistics, applications to actual experiments still lag behind the theory. Apart from Flournoy (1993), there are no true "case studies" that we know of where Bayesian ideas have been formally applied to the design of an actual scientific experiment before it is done. This is a very important area for future work. There are, however, several examples of examining an experiment in a Bayesian design framework after it has been done, for example, Clyde, Müller and Parmigiani (1995, 1996) and some examples presented in this paper.

Decisions made in designing an experiment include choosing which treatments to study, defining the treatments precisely, choosing blocking factors, choosing how to randomize, specifying the experimental units to be used, specifying a length of time for the experiment to be performed, choosing a sample size and choosing the proportion of observations to allocate to each treatment. The basic idea in experimental design is that statistical inference about the quantities of interest can be improved by appro-

Kathryn Chaloner is Associate Professor, School of Statistics, University of Minnesota, St. Paul, Minnesota. Isabella Verdinelli is Associate Professor, Dipartimento di Statistica, Probabilità e Statistiche Applicate, Università di Roma "La Sapienza" and

Prior $P(\theta)$
Experiments Y_x for $x \in D$

$$\min_x \mathbb{E}_{y_x} [U(P(\theta | y_x))]$$



Dennis V. Lindley
1956 Ann Math Stat

Chaloner & Verdinelli,
1995 Stat Sci

How to collect data to
minimize uncertainty?

Applications to spatial statistics & network design

Statistics & Probability Letters 2 (1984) 223–227
North-Holland

OPTIMAL MONITORING NETWORK DESIGNS

W.F. CASELTON

Department of Civil Engineering, University of British Columbia, Vancouver, Canada

J.V. ZIDEK

Department of Statistics GN-22, University of Washington, Seattle, WA 98195, USA

Received October 1982
Revised February 1984

Abstract: The selection of a monitoring network is formulated as a decision problem whose solutions would then be optimal. The theory is applied where the underlying field has a multivariate normal probability structure.

Keywords: decision analysis, information theory, network design, proper local utility, monitoring networks.

1. Introduction

In environmental management and resource development, for example, it is often realized only retrospectively, that earlier and larger expenditures would have been justified to create a monitoring network of more adequate density and duration. This is always the case where long term monitoring is envisaged because the designer cannot fully foresee all of the future benefits derivable from a network by all possible users and uses. And even where he can see a benefit, he can still be ignorant of the procedures or models by which the overall regional information might be extrapolated from the possibly sparse network data. Thus, the designer may be obliged to assume, arbitrarily, both a utility function and a model before he can perform some kind of design analysis. This arbitrariness tends to undermine network design methodologies and as a consequence weaken the case for advanced monitoring methodologies.

of specifying probability distributions on uncountably infinite dimensional function spaces are avoided by adopting as a surrogate of the true environment in a region, the view provided by a (hypothetical) network of high station density. The performance of any prospective network is then judged by how well it conveys the information contained in this ultimate (but hypothetical) network view. Our analysis concentrates on this ultimate network of m locations, or more precisely, the state of nature $\psi = \{\psi^1, \dots, \psi^m\}$ corresponding to the network. The column vector, ψ^j , consists of quantities determined at different times or for different attributes at location $j = 1, \dots, m$.

The design problem is that of determining a subset of all the locations $\{1, \dots, m\}$, say of size $n \ll m$, at which to locate monitoring stations. These stations would then yield measurements, $z = \{z^1, \dots, z^n\}$, of the quantities of interest of these n sites, and the measurement can be used to

Journal of Machine Learning Research 9 (2008) 235-284

Submitted 9/06; Revised 9/07; Published 2/08

Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies

Andreas Krause

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

KRAUSEA@CS.CMU.EDU

Ajit Singh

Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213

AJIT@CS.CMU.EDU

Carlos Guestrin

Computer Science Department and Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213



GUESTRIN@CS.CMU.EDU

Editor: Chris Williams

Abstract

When monitoring spatial phenomena, which can often be modeled as Gaussian processes (GPs), choosing sensor locations is a fundamental task. There are several common strategies to address this task, for example, geometry or disk models, placing sensors at the points of highest entropy (variance) in the GP model, and A-, D-, or E-optimal design. In this paper, we tackle the combinatorial optimization problem of maximizing the *mutual information* between the chosen locations and the locations which are not selected. We prove that the problem of finding the configuration that maximizes mutual information is NP-complete. To address this issue, we describe a polynomial-time approximation that is within $(1 - 1/e)$ of the optimum by exploiting the *submodularity* of mutual information. We also show how submodularity can be used to obtain online bounds, and design branch and bound search procedures. We then extend our algorithm to exploit lazy evaluations and local structure in the GP, yielding significant speedups. We also extend our approach to find placements which are robust against node failures and uncertainties in the model. These extensions are again associated with rigorous theoretical approximation guarantees, exploiting the submodularity of the objective function. We demonstrate the advantages of our approach towards optimizing mutual information in a very extensive empirical study on two real-world data sets.

Keywords: Gaussian processes, experimental design, active learning, spatial learning, sensor networks



Caselton & Zidek '84

Bayesian Active learning in Machine Learning

Information, prediction, and query by committee

Yoav Freund
Computer and Information Sciences
University of California, Santa Cruz
yoav@cs.ucsc.edu

H. Sebastian Seung
AT&T Bell Laboratories
Murray Hill, New Jersey
seung@phys1.cs.act.com

Eli Shamir
Institute of Computer Science
Hebrew University, Jerusalem
shamir@cs.huji.ac.il

Naftali Tishby
Institute of Computer Science and
Center for Neural Computation
Hebrew University, Jerusalem
tishby@cs.huji.ac.il

Abstract

We analyze the "query by committee" algorithm, a method for filtering informative queries from a random stream of inputs. We show that if the two-member committee algorithm achieves information gain with positive lower bound, then the prediction error decreases exponentially with the number of queries. We show that, in particular, this exponential decrease holds for query learning of thresholded smooth functions.

1 Introduction

For the most part, research on supervised learning has utilized a random input paradigm, in which the learner is both trained and tested on examples drawn at random from the same distribution. In contrast, in the query paradigm, the learner is given the power to ask questions, rather than just passively accept examples. What does the learner gain from this additional power? Can it attain the same prediction performance with fewer examples?

Bayesian Methods for Adaptive Models

Thesis by

David J.C. MacKay

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California

©1992

[Submitted December 10, 1991]

Gaussian Process Regression: Active Data Selection and Test Point Rejection

Sambu Seo, Marko Wallat, Thore Graepel, and Klaus Obermayer

Technische Universität Berlin, FR2-1, Franklinstr. 28-29,
D-10587 Berlin, Germany
sontag@cs.tu-berlin.de

GP Regression [Seo+ '00]

Bayesian Active Learning for Classification and Preference Learning

Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, Máté Lengyel
Computational and Biological Learning Laboratory
University of Cambridge

GP Classification [Houlsby+ '11]

Deep Bayesian Active Learning with Image Data

Yarin Gal^{1,2} Riashat Islam¹ Zoubin Ghahramani¹

Deep learning [Gal+ '17]



[T. Fuchs, J. Buhmann, 2009]

Query by committee David MacKay
[FSST 1992] 1992

How should we collect data to help us make better decisions?

Value of information

Information Value Theory


RONALD A. HOWARD, SENIOR MEMBER, IEEE

Abstract—The information theory developed by Shannon was designed to place a quantitative measure on the amount of information involved in any communication. The early developers stressed that the information measure was dependent only on the probabilistic structure of the communication process. For example, if losing all your assets in the stock market and having whale steak for supper have the same probability, then the information associated with the occurrence of either event is the same. Attempts to apply Shannon’s information theory to problems beyond communications have, in the large, come to grief. The failure of these attempts could have been predicted because no theory that involves just the probabilities of outcomes without considering their consequences could possibly be adequate in describing the importance of uncertainty to a decision maker. It is necessary to be concerned not only with the probabilistic nature of the uncertainties that surround us, but also with the economic impact that these uncertainties will have on us.

In this paper the theory of the value of information that arises from considering jointly the probabilistic and economic factors that affect decisions is discussed and illustrated. It is found that numerical values can be assigned to the elimination or reduction of any uncertainty. Furthermore, it is seen that the joint elimination of the

Ronald Howard
IEEE TSSC 1966

$$\max_x \mathbb{E}_{y_x} \left[\max_a \mathbb{E}[R(\theta, a) \mid y_x] \right]$$



	<i>Sick</i>	<i>Healthy</i>
No treatment	-\$\$\$	0
Treatment	\$	-\$

How to collect data to make
better decisions?


Multi-armed Bandits

ADVANCES IN APPLIED MATHEMATICS 6, 4-22 (1985)

Asymptotically Efficient Adaptive Allocation Rules*

T. L. LAI AND HERBERT ROBBINS

Department of Statistics, Columbia University, New York, New York 10027



How to collect data to make
maximize reward?

$$\max_{\pi} \mathbb{E} \left[\sum_t R_{\pi_t} \right]$$

Exploration—exploitation
dilemma

J. R. Statist. Soc. B (1979),
41, No. 2, pp. 148-177

Bandit Processes and Dynamic Allocation Indices

By J. C. GITTINS

Keble College, Oxford

[Read before the ROYAL STATISTICAL SOCIETY at a meeting organized by the RESEARCH SECTION on Wednesday, February 14th, 1979, the Chairman Professor J. F. C. KINGMAN in the Chair]

SUMMARY

The paper aims to give a unified account of the central concepts in recent work on bandit processes and dynamic allocation indices; to show how these reduce some previously intractable problems to the problem of calculating such indices; and to describe how these calculations may be carried out. Applications to stochastic scheduling, sequential clinical trials and a class of search problems are discussed.

Keywords: BANDIT PROCESSES; DYNAMIC ALLOCATION INDICES; TWO-ARMED BANDIT PROBLEM; MARKOV DECISION PROCESSES; OPTIMAL RESOURCE ALLOCATION; SEQUENTIAL RANDOM SAMPLING; CHEMICAL RESEARCH; CLINICAL TRIALS; SEARCH

1. INTRODUCTION

A scheduling problem

There are n jobs to be carried out by a single machine. The times taken to process the jobs are independent integer-valued random variables. The jobs must be processed one at a time. At the beginning of each time unit any job may be selected for processing, whether or not the job processed during the preceding time unit has been completed, and there is no penalty or delay involved in switching from one job to another. The probability that $t+1$ time units are required to complete the processing of job i , conditional on more than t time units being needed, is $p_i(t)$ ($i = 1, 2, \dots, n$; $t \in \mathbb{Z}$). The reward for finishing job i at time s is $a^s V_i$ ($0 < a < 1$; $V_i > 0$, $i = 1, 2, \dots, n$), and there are no other rewards or costs. The problem is to decide which job to process next at each stage so as to maximize the total expected reward.

A multi-armed bandit problem

There are n arms which may be pulled repeatedly in any order. Each pull takes one time unit and only one arm may be pulled at a time. A pull may result in either a success or a failure. The sequence of successes and failures which result from pulling arm i forms a Bernoulli process with an unknown success probability θ_i ($i = 1, 2, \dots, n$). A successful pull on any arm at time t yields a reward a^t ($0 < a < 1$), whilst an unsuccessful pull yields a zero reward. At time zero θ_i has the probability density

$$(\alpha_i(0) + \beta_i(0) + 1)! (\alpha_i(0)! \beta_i(0)!)^{-1} \theta_i^{\alpha_i(0)} (1 - \theta_i)^{\beta_i(0)}$$

i.e. a beta distribution with parameters $(\alpha_i(0), \beta_i(0))$, and these distributions are independent for the different arms. The problem is to decide which arm to pull next at each stage so as to maximize the total expected reward from an infinite sequence of pulls.

From Bayes' theorem it follows that at every stage θ_i has a beta distribution, but with parameters which change at each pull on an arm i . If in the first t pulls there are r successes, the new values of the parameters, which we denote by $(\alpha_i(t), \beta_i(t))$, are $(r+1, t-r+1)$.

1. INTRODUCTION

Let Π_j ($j = 1, \dots, k$) denote statistical populations (treatments, manufacturing processes, etc.) specified respectively by univariate density functions $f(x; \theta_j)$ with respect to some measure ν , where $f(\cdot; \cdot)$ is known and the θ_j are unknown parameters belonging to some set Θ . Assume that $\int_{-\infty}^{\infty} |x| f(x; \theta) d\nu(x) < \infty$ for all $\theta \in \Theta$. How should we sample x_1, x_2, \dots sequentially from the k populations in order to achieve the greatest possible expected value of the sum $S_n = x_1 + \dots + x_n$ as $n \rightarrow \infty$? Starting with [3] there has been a considerable literature on this subject, which is often called the multi-armed bandit problem. The name derives from an imagined slot machine with $k \geq 2$ arms. (Ordinary slot machines with one arm are one-armed bandits, since in the long run they are as effective as human bandits in separating the victim from his money.) When an arm is pulled, the player wins a random reward. For each arm j there is an unknown probability distribution Π_j of the reward. The player wants to choose at each stage one of the k arms, the choice depending in some way on the record of previous trials, so as to maximize the long-run total expected reward. A more worthy setting for this problem is in the context of sequential clinical trials, where there are k treatments of unknown efficacy to be used in treating a long sequence of patients.

An adaptive allocation rule φ is a sequence of random variables $\varphi_1, \varphi_2, \dots$ taking values in the set $\{1, \dots, k\}$ and such that the event $\{\varphi_n = j\}$ ("sample from Π_j at stage n ") belongs to the σ -field \mathcal{F}_{n-1} generated by the previous values $\varphi_1, x_1, \dots, \varphi_{n-1}, x_{n-1}$. Let $\mu(\theta) = \int_{-\infty}^{\infty} xf(x; \theta) d\nu(x)$.

*Research supported by the National Science Foundation and the National Institutes of Health. This paper was delivered at the Statistical Research Conference at Cornell University, Ithaca, N.Y., 1978.

J.C. Gittins

J Roy Stat Soc 1979

Lai & Robbins


Adv Appl Math '85

Structured bandits / Bayesian optimization

ON BAYESIAN METHODS FOR SEEKING THE EXTREMUM

J. Močkus

Institute of Physics and Mathematics
Academy of Sciences Lithuanian SSR
Vilnius, USSR



1. Introduction

Many well known methods for seeking the extremum had been developed on the basis of quadratic approximation.

In some problems of global optimization the function to be minimized can be considered as a realization of some stochastic function. The optimization technique based upon the minimization of the expected deviation from the extremum is called Bayesian.

2. The definition of Bayesian methods

Suppose the function to be minimized is a realization of some stochastic function $f(x) = f(x, \omega)$, $x \in A \subset R^m$ where $\omega \in \Omega$ is some fixed but unknown index.

The probability distribution P on Ω is defined by the equalities

$$P\{\omega : f(x_i, \omega) < y_i, i = 1, \dots, n, \omega \in \Omega\} = F_{x_1, \dots, x_n}(y_1, \dots, y_n), n = 1, 2, \dots$$

where $F_{x_1, \dots, x_n}(y_1, \dots, y_n)$, $x_i \in A$, $i = 1, \dots, n$ is the a priori probability distribution function.

The observation is evaluation of the function f at some fixed point x_i . The vector


$$z_n = (f(x_i), x_i, i = 1, \dots, n), n = 1, \dots, N,$$

contains the information gained in all the observations from 1 to n.

A decision function is the measurable vector-function $d = (d_1, \dots, d_N)$, which expresses the dependence between the point of the following observation and the results of the previous observations


$$d_n = d_n(z_1, \dots, z_{n-1}), n = 1, \dots, N.$$

Acquisition function



How to collect data to make
find maximum of an unknown function?

Optimistic Bayesian Optimization with GPs



Key idea: Focus exploration on plausible maximizers
(upper confidence bound \geq best lower bound)

Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting

Niranjn Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger

Abstract—Many applications require optimizing an unknown, noisy function that is expensive to evaluate. We formalize this task as a multiarmed bandit problem, where the payoff function is either sampled from a Gaussian process (GP) or has low norm in a reproducing kernel Hilbert space. We resolve the important open problem of deriving regret bounds for this setting, which imply novel convergence rates for GP optimization. We analyze an intuitive Gaussian process upper confidence bound (GP-UCB) algorithm, and bound its cumulative regret in terms of maximal information gain, establishing a novel connection between GP optimization and experimental design. Moreover, by bounding the latter in terms of operator spectra, we obtain explicit sublinear regret bounds for many commonly used covariance functions. In some important cases, our bounds have surprisingly weak dependence on the dimensionality. In our experiments on real sensor data, GP-UCB compares favorably with other heuristic GP optimization approaches.

Index Terms—Bandit problems, Bayesian prediction, experimental design, Gaussian process (GP), information gain, nonparametric statistics, online learning, regret bound, statistical learning.

I. INTRODUCTION

IN MOST stochastic optimization settings, evaluating the unknown function is expensive, and sampling is to be minimized. Examples include choosing advertisements in sponsored search to maximize profit in a click-through model [2] or learning optimal control strategies for robots [3]. Predominant approaches to this problem include the multiarmed bandit paradigm [4], where the goal is to maximize cumulative

reward by optimally balancing exploration and exploitation, and experimental design [5], where the function is to be explored globally with as few evaluations as possible, for example, by maximizing information gain. The challenge in both approaches is twofold: we have to estimate an unknown function f from noisy samples, and we must optimize our estimate over some high-dimensional input space. For the former, much progress has been made in machine learning through kernel methods and Gaussian process (GP) models [6], where smoothness assumptions about f are encoded through the choice of kernel in a flexible nonparametric fashion. Beyond Euclidean spaces, kernels can be defined on diverse domains such as spaces of graphs, sets, or lists.

We are concerned with GP optimization in the multiarmed bandit setting, where f is sampled from a GP distribution or has low “complexity” measured in terms of its reproducing kernel Hilbert space (RKHS) norm under some kernel. We provide the first sublinear regret bounds in this nonparametric setting, which imply convergence rates for GP optimization. In particular, we analyze the Gaussian process upper confidence bound (GP-UCB) algorithm, a simple and intuitive Bayesian method [7], [9]. While objectives are different in the multiarmed bandit and experimental design paradigms, our results draw a close technical connection between them: our regret bounds come in terms of an *information gain* quantity, measuring how fast f can be learned in an information-theoretic sense. The submodularity of this function allows us to prove sharp regret bounds for particular covariance functions, which we demonstrate for commonly used squared exponential and Matérn kernels.

Related Work: Our work generalizes stochastic *linear* optimization in a bandit setting, where the unknown function comes from a finite-dimensional linear space. GPs are nonlinear random functions, which can be represented in an infinite-dimensional linear space. For the standard linear setting, Dani *et al.* [10] provide a near-complete characterization explicitly dependent on the dimensionality. In the GP setting, the challenge is to characterize complexity in a different manner, through properties of the kernel function. Our technical contributions are twofold: first, we show how to analyze the nonlinear setting by focusing on the concept of information gain, and second, we explicitly bound this information gain measure using the concept of submodularity [11] and knowledge about kernel operator spectra.

Compared to an earlier version of [1], this paper is significantly expanded, including detailed proofs, additional explanations (e.g., Fig. 3), and more comprehensive experimental demonstration of the performance of the GP-UCB algorithm.

Kleinberg *et al.* [12] provide regret bounds under weaker and less configurable assumptions (only Lipschitz continuity w.r.t.

Manuscript received October 17, 2010; accepted September 27, 2011. Date of publication January 24, 2012; date of current version April 17, 2012. This work was supported in part by the Office of Naval Research under Grant N00014-09-1-1044, in part by the National Science Foundation under Grants CNS-0952392 and IIS-0953413, in part by a gift from Microsoft Corporation, and in part by the Excellence Initiative of the German Research Foundation (DFG). This manuscript is an extended version of our paper that appeared in ICML 2010 [1].

N. Srinivas is with the California Institute of Technology, Pasadena, CA 91125 USA (e-mail: nrsrinivas@caltech.edu).

A. Krause is with the Swiss Federal Institute of Technology, Zürich 8006, Switzerland (e-mail: krause@ethz.ch).

S. M. Kakade is with Microsoft Research, New England, Cambridge, MA 02142 USA, and also with the Department of Statistics, University of Pennsylvania, Philadelphia, PA 19104-6340 USA (e-mail: kakade@wharton.upenn.edu).


M. Seeger is with the School of Computer and Communication Sciences, École Polytechnique Fédérale de Lausanne, Lausanne CH-1015, Switzerland (e-mail: matthias.seeger@epfl.ch).

Communicated by D. Polunin, Associate Editor for Detection and Estimation.

Color versions of one or more of the figures in this paper are available online at <http://www.ieee.org>.


Digital Object Identifier 10.1109/TIT.2011.2182033

Applications of Bayesian optimization



AutoML

[Vizier, SageMaker, Snoek+ '12...]



Molecular Design

[Romero+ '13, Gomez-Bombarelli+ '18, ...]



Bipedal locomotion

[Calandra+14]



Env. Monitoring

[Marchant+ '12, Hitz+ '14]




Tuning SwissFEL

[Kirschner+ '22]



PV Optimization

[Abdelrahman+ '16]



S. Harkema, The Lancet, Elsevier

Spinal-cord therapy

[Sui+ '18]

Education [Lindsey+ '15]

Energy efficiency for HPC [Miyazaki+ '18]

• • •

A Sober Look at LLMs for Material Discovery: Are They Actually Good for Bayesian Optimization Over Molecules?

Agustinus Kristiadi¹ Felix Strieth-Kalthoff² Marta Skreta^{2,1} Pascal Poupart^{3,1} Alán Aspuru-Guzik^{2,1}
Geoff Pleiss^{4,1}

DOI: [10.1039/D3SC05607D](https://doi.org/10.1039/D3SC05607D) (Edge Article) *Chem. Sci.*, 2024, **15**, 7732-7741

Combining Bayesian optimization and automation to simultaneously optimize reaction conditions and routes†

Oliver Schilter ^{*ab}, Daniel Pacheco Gutierrez ^c, Linnea M. Folkmann ^c, Alessandro Castrogiovanni ^a, Alberto García-Durán ^c, Federico Zipoli ^{ab}, Loïc M. Roch ^c and Teodoro Laino ^{ab}

ACS Central Science > Vol 10/Issue 7 > Article

Open Access



ARTICLE | May 22, 2024

Enhanced Sequence-Activity Mapping and Evolution of Artificial Metalloenzymes by Active Learning

Tobias Vornholt, Mojmír Mutný, Gregor W. Schmidt, Christian Schellhaas, Ryo Tachibana, Sven Panke, Thomas R. Ward*, Andreas Krause*, and Markus Jeschek*

Article | [Open access](#) | Published: 03 February 2024

A dynamic Bayesian optimized active recommender system for curiosity-driven partially Human-in-the-loop automated experiments

Arpan Biswas , Yongtao Liu, Nicole Creange, Yu-Chen Liu, Stephen Jesse, Jan-Chi Yang, Sergei V. Kalinin, Maxim A. Ziatdinov & Rama K. Vasudevan 

npj Computational Materials **10**, Article number: 29 (2024) | [Cite this article](#)

Article | Published: 28 February 2024

Identifying general reaction conditions by bandit optimization

[Jason Y. Wang \(王亿珩\)](#), [Jason M. Stevens](#), [Stavros K. Kariofillis](#), [Mai-Jan Tom](#), [Dung L. Golden](#), [Jun Li](#), [Jose E. Tabora](#), [Marvin Parasram](#), [Benjamin J. Shields](#), [David N. Primer](#), [Bo Hao](#), [David Del Valle](#), [Stacey DiSomma](#), [Ariel Furman](#), [G. Greg Zipp](#), [Sergey Melnikov](#), [James Paulson](#) & [Abigail G. Doyle](#) 

Nature **626**, 1025–1033 (2024) | [Cite this article](#)

OPEN ACCESS


Bayesian optimization algorithms for accelerator physics

[Ryan Roussel](#) ^{1,*}, [Auralee L. Edelen](#)¹, [Tobias Boltz](#) ¹, [Dylan Kennedy](#)¹, [Zhe Zhang](#) ¹, [Fuhao Ji](#) ¹, [Xiaobiao Huang](#) ¹, [Daniel Ratner](#) ¹, and [Andrea Santamaria Garcia](#) ² *et al.*

Show more


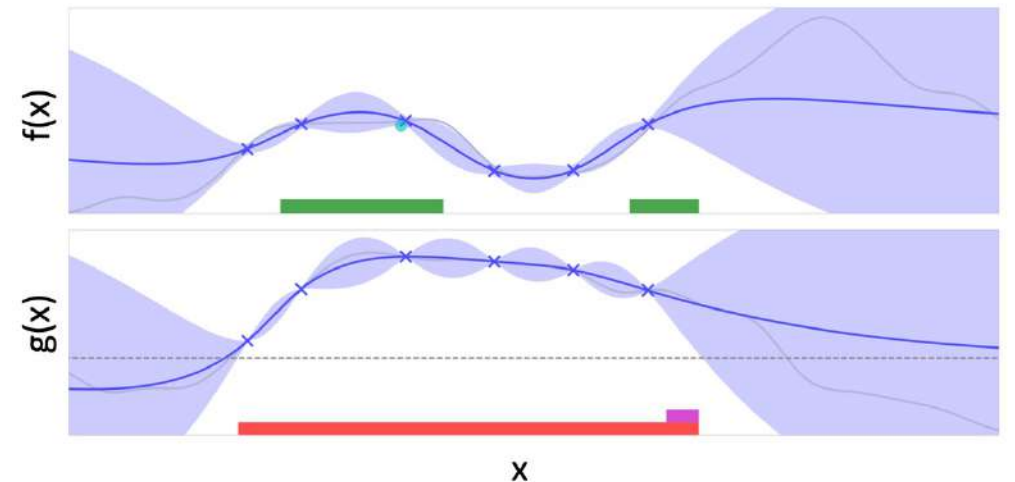
Phys. Rev. Accel. Beams **27**, 084801 – Published 6 August, 2024

...



Some modern research challenges

- How do we reliably **quantify epistemic uncertainty** for complex models (deep nets, LLMs, ...)?
- How do we **learn complex, data-driven priors** from related tasks?
- How do we use uncertainty not only to guide exploration, but to act **cautiously / safely**?
- How do we **scale to complex design spaces**, e.g. via (Bayesian) RL?
- ...



Bayesian Active learning for Large Language Models?

Why should I attend the Bayesian Decision Making under Uncertainty workshop at NeurIPS 2024?



Attending the "Bayesian Decision-making and Uncertainty" workshop at NeurIPS 2024 offers a unique opportunity to engage with leading experts and the latest advancements in the field.

Here's why you should consider participating:

1. Engage with Cutting-Edge Research

The workshop will feature presentations on recent developments in Bayesian methods, including topics like Bayesian optimization, active learning, uncertainty quantification, Gaussian processes, spatiotemporal modeling, and sequential experimental design. Engaging with this content will enhance your understanding of how these methods are applied across various domains.

GP SEMINAR SERIES

2. Learn from Esteemed Speakers

The event boasts a lineup of distinguished speakers:

- **Roman Garnett** (Washington University in St. Louis and Uber)
- **Jacob R. Gardner** (University of Pennsylvania)
- **Virginia Aglietti** (Google DeepMind)
- **Esther Rolf** (University of Colorado Boulder)
- **Mark van der Wilk** (University of Oxford)

Transductive “prediction-oriented” active learning

[Hübötter et al. NeurIPS 2024]



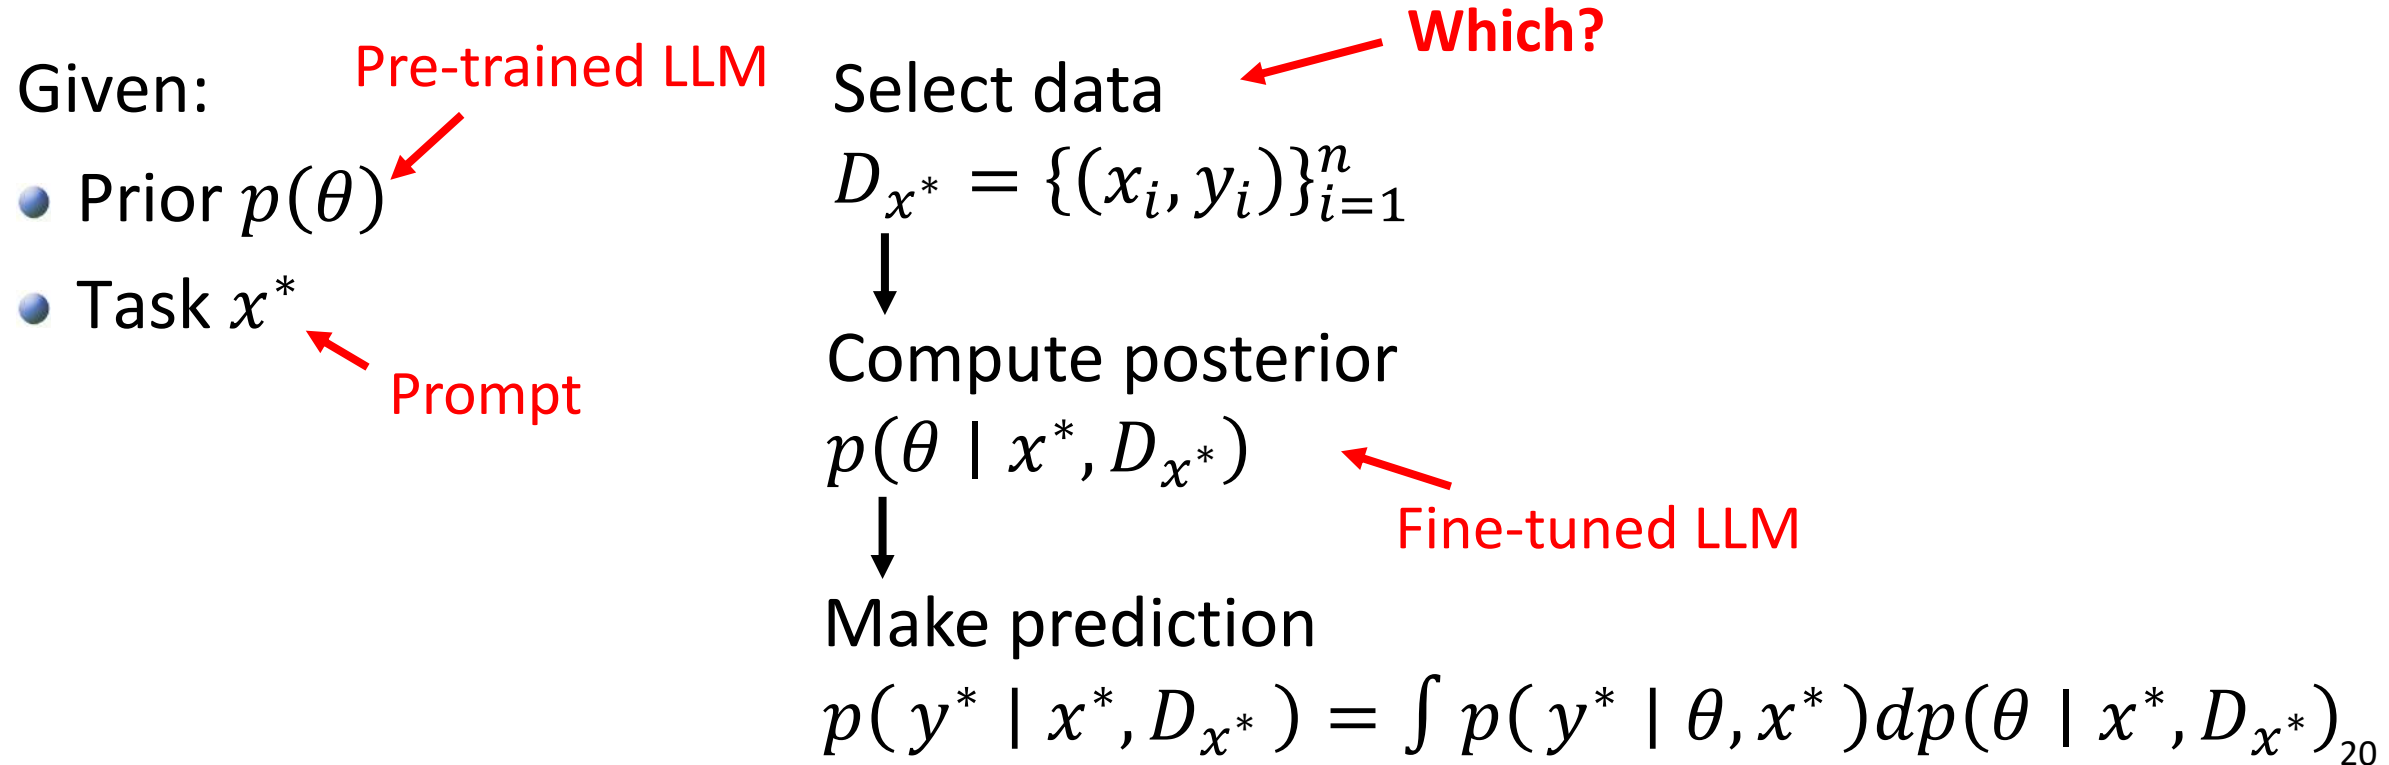
Jonas
Hübötter

Bhavi
Sukhija

Lenart
Treven

Yarden
As

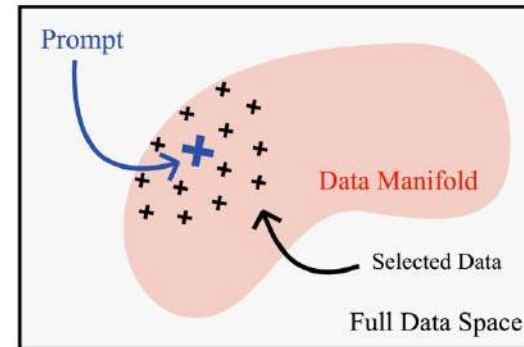
Can we make better predictions by learning a specialized model $p(\theta \mid x^*)$ for each task x^* at test-time?



Informative sampling for transductive learning


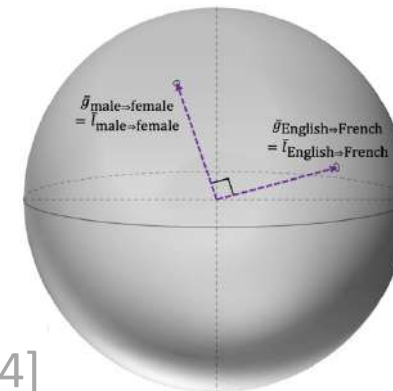
Minimize epistemic uncertainty about prediction:

$$\begin{aligned}x_{t+1} &= \arg \max_{x \in D} I(f_{x^*}; y_x \mid x_{1:t}, y_{1:t}) \\ &= \arg \min_{x \in D} H(f_{x^*} \mid x_{1:t}, y_{1:t}, y_x)\end{aligned}$$



Closed-form with a linear model $f_{x^*} = w^\top \phi(x^*)$
if w and $y_x = f_x + \varepsilon$ are Gaussian

Motivation: linear representation
hypothesis in LLMs




Jonas
Hübötter

Ido
Hakimi

[Mikolov, Yih, Zweig NAACL 2013; Park, Choe, Veitch ICML 2024]


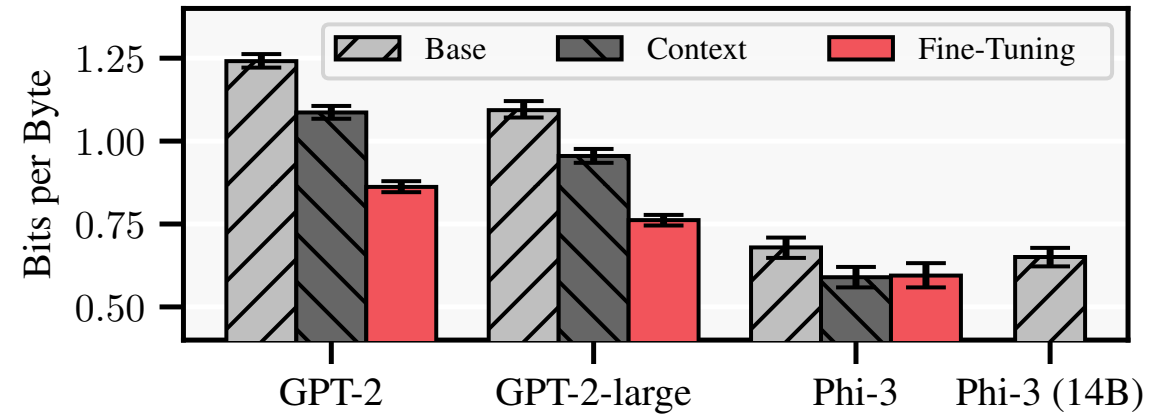

Active learning for test-time-tuning LLMs

[Hübotter, Bongni, Hakimi, Krause <https://arxiv.org/abs/2410.08020>]



Jonas
Hübotter

Ido
Hakimi



Closing the loop with Bayesian Decision Making

