Probability and Its Applications

Victor H. de la Peña
Tze Leung Lai
Qi-Man Shao

Self-Normalized Processes

Limit Theory and Statistical Applications

$$P(S_n \geq x\, n^{1/2} V_n)^{1/n} \longrightarrow \sup_{b \geq 0} \inf_{t \geq 0} E e^{t(bX - x(X^2+b^2)/2)}$$

Springer

Self-Normalization

See de la Pena et. al. (2009) for a comprehensive account of the theory and applications of Self-normalization.

Pseudo-Maximization (Method of Mixtures)

The method of Pseudo-Maximization (also known as the method of mixtures) was used in de la Pena et. al. (2004) and is based on the following:

Canonical Assumption.

Let (A, B) be an arbitrarily dependent vector of random variables, with $B > 0$. Assume that $-\infty < \lambda < \infty$. Then, the pair is said to satisfy the Canonical Assumption if

$$E \exp(\lambda A - \lambda^2 B^2/2) \leq 1.$$

Lemmas 1-3 present processes that satisfy the canonical assumption.

Lemma 1 Let W_t be a standard Brownian Motion. Assume that T is a stopping time such that $T < \infty$ a.s. Then for all $-\infty < \lambda < \infty$

$$E \exp\{\lambda W_T - \lambda^2 T/2\} \leq 1.$$

Lemma 2 Let M_t be a continuous, square-integrable martingale, with $M_0 = 0$. Then, for all $-\infty < \lambda < \infty$,

$$\exp\{\lambda M_t - \lambda^2 < M >_t /2\} \leq 1.$$

If M_t is ony assumed to be a continuous local martingale, the inequality is also valid (by application of Fatou's lemma).

Lemma 3 Let $\{d_i\}$ be a sequence of variables adapted to an increasing sequence of σ-fields $\{\mathcal{F}_i\}$. Assume that the d_i's are conditionally symmetric (i.e., $\mathcal{L}(d_i|\mathcal{F}_{i-1}) = \mathcal{L}(-d_i|\mathcal{F}_{i-1})$). Then,

$$E \exp\{\lambda \sum_{i=1}^{n} d_i - \lambda^2 \sum_{i=1}^{n} d_i^2/2\} \leq 1,$$

for all $-\infty < \lambda < \infty$.

Theorem 2. (de la Peña, Klass and Lai (AOP)) *Let A, B with $B > 0$ be random variables satisfying the canonical assumption for all $\lambda \in \mathbf{R}$. Then*

$$P(\frac{|A|}{\sqrt{B^2 + (EB)^2}} > x) \leq \sqrt{2}\exp(-x^2/4)$$

for all $x > 0$.

The proof of this result is based on the following lemma.

Lemma *Let A, B with $B > 0$ be two random variables satisfying the canonical condition for all $\lambda \in \mathbf{R}$. Then for all $y > 0$,*

$$E\frac{y}{\sqrt{B^2 + y^2}} \exp\{\frac{A^2}{2(B^2 + y^2)}\} \leq 1.$$

Proof: Multiplying both sides of the canonical condition by $(2\pi)^{-1/2} y \exp(-\lambda^2 y^2/2)$ (with $y > 0$) and integrating over λ, we obtain by using Fubini's theorem that

$$1 \geq \int_{-\infty}^{\infty} E\frac{y}{\sqrt{2\pi}} \exp\Big(\lambda A - \frac{\lambda^2}{2}B^2\Big) \exp\Big(-\frac{\lambda^2 y^2}{2}\Big) d\lambda$$

$$= E\Big[\frac{y}{\sqrt{B^2 + y^2}} \exp\Big\{\frac{A^2}{2(B^2 + y^2)}\Big\} \times$$

$$\int_{-\infty}^{\infty} \frac{\sqrt{B^2 + y^2}}{\sqrt{2\pi}} \exp\{-\frac{B^2 + y^2}{2}\Big(\lambda^2 - 2\frac{A}{B^2 + y^2}\lambda + \frac{A^2}{(B^2 + y^2)^2}\Big)\} d\lambda\Big]$$

$$= E\Big[\frac{y}{\sqrt{B^2 + y^2}} \exp\Big(\frac{A^2}{2(B^2 + y^2)}\Big)\Big]. \qquad \square$$

By Schwarz's inequality,

$$E\exp\left\{\frac{A^2}{4(B^2+y^2)}\right\} \le \left\{\left(E\frac{y\exp\{\frac{A^2}{2(B^2+y^2)}\}}{B^2+y^2}\right)\left(E\sqrt{\frac{B^2+y^2}{y^2}}\right)\right\}^{1/2}$$
$$\le \left(E\sqrt{\frac{B^2}{y^2}+1}\right)^{1/2}.$$

Since $E\sqrt{\frac{B^2}{y^2}+1} \le E(\frac{B}{y}+1)$, the special case $y = EB$ above gives

$$E\exp(A^2/[4(B^2+(EB)^2]) \le \sqrt{2}.$$

Then, using Markov's inequality and this we get

$$P\{\frac{|A|}{\sqrt{B^2+(EB)^2}} \ge x\} = P\{\frac{A^2}{4(B^2+(EB)^2)} \ge \frac{x^2}{4}\} \le \sqrt{2}\exp(-x^2/4).$$

Corollary (Law of the Iterated Logarithm)

Let $d_i, \mathcal{F}_i$ a sequence of conditionally symmetric random variables. Then, using Lemma 3 and a mixture introduced by Robbins and Sigmund we obtain the following,

$$\limsup_{n\to\infty} \frac{\sum_{i=1}^n d_i}{\sqrt{2\sum_{i=1}^n d_i^2 \log\log \sum_{i=1}^n d_i^2}} \le 1,$$

on the set $\{\lim_{n\to\infty} \sum_{i=1}^n d_i^2 = \infty$.

Conditionally Independent (Tangent) Decoupling

Please refer to de la Pena and Giné (1999) and de la Pena (2019) for a more comprehensive account of the theory as well as additional pertinent references.

The theory of martingale inequalities has been central in the development of modern probability theory. This theory has been expanded widely through the introduction of the conditionally independent (tangent) decoupling principle which provides a general approach for handling problems involving dependent variables.

Let $\{d_i\}$ and $\{e_i\}$ be two sequences of random variables adapted to the σ-fields $\{\mathcal{F}_i\}$. Then $\{d_i\}$ and $\{e_i\}$ are said to be tangent with respect to $\{\mathcal{F}_i\}$ if, for all i,

$$\mathcal{L}(d_i|\mathcal{F}_{i-1}) = \mathcal{L}(e_i|\mathcal{F}_{i-1}),$$

where $\mathcal{L}(d_i|\mathcal{F}_{i-1})$ denotes the conditional probability law of d_i given $\mathcal{F}_{i-1}$.

Let $d_1, ..., d_n$ be an arbitrary sequence of dependent random variables adapted to an increasing sequence of $\sigma-$fields $\{\mathcal{F}_i\}$. Then, one can construct a sequence $e_1, ..., e_n$ of random variables which is conditionally independent given $\mathcal{G} = \mathcal{F}_n$. The construction proceeds as follows: First we take e_1 and d_1 to be two independent copies of the same random mechanism. Having constructed $d_1, ..., d_{i-1}; e_1, ..., e_{i-1}$, the ith pair of variables d_i and e_i comes from i.i.d. copies of the same random mechanism, given $\mathcal{F}_{i-1}$. It is easy to see that using this construction and taking

$$\mathcal{F}'_i = \mathcal{F}_i \vee \sigma(e_1, ..., e_i),$$

the sequences $\{d_i\}$, $\{e_i\}$ satisfy

$$\mathcal{L}(d_i|\mathcal{F}'_{i-1}) = \mathcal{L}(e_i|\mathcal{F}'_{i-1}) = \mathcal{L}(e_i|\mathcal{G}),$$

and the sequence $\{e_1, ..., e_n\}$ is conditionally independent given $\mathcal{G} = \mathcal{F}_n$

A sequence $\{e_i\}$ of random variables satisfying the above conditions is said to be a *decoupled* tangent version of $\{d_i\}$. Moreover, the sequences $\{d_i - e_i\}$, is conditionally symmetric.

Corollary (A Universal LIL)

Let $d_i, \mathcal{F}_i$ a sequence of arbitrarily dependent variables without moment assumptions. Let e_i be its associated decoupled sequence. Then,

$$\limsup_{n\to\infty} \frac{\sum_{i=1}^n (d_i - e_i)}{\sqrt{2\sum_{i=1}^n (d_i - e_i)^2 \log\log \sum_{i=1}^n (d_i - e_i)^2}} \leq 1,$$

on the set $\{\lim_{n\to\infty} \sum_{i=1}^n (d_i - e_i)^2 = \infty\}$,

A symmetrised extension of Kolmogorov's LIL.

NIPS (2011)

Improved Algorithms for Linear Stochastic Bandits

Yasin Abbasi-Yadkori
abbasiya@ualberta.ca
Dept. of Computing Science
University of Alberta

Dávid Pál
dpal@google.com
Dept. of Computing Science
University of Alberta

Csaba Szepesvári
szepesva@ualberta.ca
Dept. of Computing Science
University of Alberta

Abstract

We improve the theoretical analysis and empirical performance of algorithms for the stochastic multi-armed bandit problem and the linear stochastic multi-armed bandit problem. In particular, we show that a simple modification of Auer's UCB algorithm (Auer, 2002) achieves with high probability constant regret. More importantly, we modify and, consequently, improve the analysis of the algorithm for the for linear stochastic bandit problem studied by Auer (2002), Dani et al. (2008), Rusmevichientong and Tsitsiklis (2010), Li et al. (2010). Our modification improves the regret bound by a logarithmic factor, though experiments show a vast improvement. In both cases, the improvement stems from the construction of smaller confidence sets. For their construction we use a novel tail inequality for vector-valued martingales.

1 Introduction

Linear stochastic bandit problem is a sequential decision-making problem where in each time step we have to choose an action, and as a response we receive a stochastic reward, expected value of which is an unknown linear function of the action. The goal is to collect as much reward as possible over the course of n time steps. The precise model is described in Section 1.2.

Several variants and special cases of the problem exist differing on what the set of available actions is in each round. For example, the standard stochastic d-armed bandit problem, introduced by Robbins (1952) and then studied by Lai and Robbins (1985), is a special case of linear stochastic bandit problem where the set of available actions in each round is the standard orthonormal basis of $\mathbb{R}^d$. Another variant, studied by Auer (2002) under the name "linear reinforcement learning", and later in the context of web advertisement by Li et al. (2010), Chu et al. (2011), is a variant when the set of available actions changes from time step to time step, but has the same finite cardinality in each step. Another variant dubbed "sleeping bandits", studied by Kleinberg et al. (2008), is the case when the set of available actions changes from time step to time step, but it is always a subset of the standard orthonormal basis of $\mathbb{R}^d$. Another variant, studied by Dani et al. (2008), Abbasi-Yadkori et al. (2009), Rusmevichientong and Tsitsiklis (2010), is the case when the set of available actions does not change between time steps but the set can be an almost arbitrary, even infinite, bounded subset of a finite-dimensional vector space. Related problems were also studied by Abe et al. (2003), Walsh et al. (2009), Dekel et al. (2010).

In all these works, the algorithms are based on the same underlying idea—the *optimism-in-the-face-of-uncertainty* (OFU) principle. This is not surprising since they are solving almost the same problem. The OFU principle elegantly solves the exploration-exploitation dilemma inherent in the problem. The basic idea of the principle is to maintain a confidence set for the vector of coefficients of the linear function. In every round, the algorithm chooses an estimate from the confidence set and an action so that the predicted reward is maximized, i.e., estimate-action pair is chosen optimistically. We give details of the algorithm in Section 2.

A Proof of Theorem 1

For the proof of Theorem 1 we will need the following two lemmas. Both lemmas use the same assumptions and notation as the theorem. The first lemma is a standard supermartingale argument adapted to randomly stopped, vector valued processes.

Lemma 8. *Let $\lambda \in \mathbb{R}^d$ be arbitrary and consider for any $t \geq 0$*

$$M_t^\lambda = \exp\left(\sum_{s=1}^{t}\left[\frac{\eta_s \langle \lambda, X_s\rangle}{R} - \frac{1}{2}\langle \lambda, X_s\rangle^2\right]\right) .$$

Let τ be a stopping time with respect to the filtration $\{F_t\}_{t=0}^\infty$. Then M_τ^λ is almost surely well-defined and

$$\mathbf{E}[M_\tau^\lambda] \leq 1 .$$

Proof of Lemma 8. We claim that $\{M_t^\lambda\}_{t=0}^\infty$ is a supermartingale. Let

$$D_t^\lambda = \exp\left(\frac{\eta_t \langle \lambda, X_t\rangle}{R} - \frac{1}{2}\langle \lambda, X_t\rangle^2\right) .$$

Observe that by conditional R-sub-Gaussianity of η_t we have $\mathbf{E}[D_t^\lambda \mid F_{t-1}] \leq 1$. Clearly, D_t^λ is F_t-measurable, as is M_t^λ. Further,

$$\mathbf{E}[M_t^\lambda \mid F_{t-1}] = \mathbf{E}[M_1^\lambda \cdots D_{t-1}^\lambda D_t^\lambda \mid F_{t-1}] = D_1^\lambda \cdots D_{t-1}^\lambda\, \mathbf{E}[D_t^\lambda \mid F_{t-1}] \leq M_{t-1}^\lambda ,$$

showing that $\{M_t^\lambda\}_{t=0}^\infty$ is indeed a supermartingale and in fact $\mathbf{E}[M_t^\lambda] \leq 1$.

Now, we argue that M_τ^λ is well-defined. By the convergence theorem for nonnegative supermartingales, $M_\infty^\lambda = \lim_{t\to\infty} M_t^\lambda$ is almost surely well-defined. Hence, M_τ^λ is indeed well-defined independently of whether $\tau < \infty$ holds or not. Next, we show that $\mathbf{E}[M_\tau^\lambda] \leq 1$. For this let $Q_t^\lambda = M_{\min\{\tau,t\}}^\lambda$ be a stopped version of $(M_t^\lambda)_t$. By Fatou's Lemma, $\mathbf{E}[M_\tau^\lambda] = \mathbf{E}[\liminf_{t\to\infty} Q_t^\lambda] \leq \liminf_{t\to\infty} \mathbf{E}[Q_t^\lambda] \leq 1$, showing that $\mathbf{E}[M_\tau^\lambda] \leq 1$ indeed holds. □

The next lemma uses the "method of mixtures" technique (cf. Chapter 11, de la Peña et al. 2009). In fact, the lemma could also be derived from Theorem 14.7 of de la Peña et al. (2009).

Lemma 9 (Self-normalized bound for vector-valued martingales). *Let τ be a stopping time with respect to the filtration $\{F_t\}_{t=0}^\infty$. Then, for $\delta > 0$, with probability $1-\delta$,*

$$\|S_\tau\|^2_{\overline{V}_\tau^{-1}} \leq 2R^2 \log\left(\frac{\det(\overline{V}_\tau)^{1/2}\det(V)^{-1/2}}{\delta}\right) .$$

Proof of Lemma 9. Without loss of generality, assume that $R = 1$ (by appropriately scaling S_t, this can always be achieved). Let

$$V_t = \sum_{s=1}^{t} X_s X_s^\top \qquad\qquad M_t^\lambda = \exp\left(\langle \lambda, S_t\rangle - \frac{1}{2}\|\lambda\|^2_{V_t}\right) .$$

Notice that by Lemma 8, the mean of M_τ^λ is not larger than one.

Let Λ be a Gaussian random variable which is independent of all the other random variables and whose covariance is V^{-1}. Define

$$M_t = \mathbf{E}[M_t^\Lambda \mid F_\infty] ,$$

where F_∞ is the tail σ-algebra of the filtration i.e. the σ-algebra generated by the union of the all events in the filtration. Clearly, we still have $\mathbf{E}[M_\tau] = \mathbf{E}[\mathbf{E}[M_\tau^\Lambda \mid \Lambda]] \leq 1$.

14.2 Moment and Exponential Inequalities via Pseudo-Maximization

Consider the canonical assumption (14.4). If the random function $\exp\{\theta'A - \theta'C\theta/2\}$ could be maximized over θ inside the expectation, taking the maximizing value $\theta = C^{-1}A$ in (14.4) would yield $E\exp\{A'C^{-1}A/2\} \leq 1$. This in turn would give the exponential bound $P(\|C^{-1/2}A\| > x) \leq \exp(-x^2/2)$. Although we cannot interchange the order of $\max_\lambda$ and E that is needed in the above argument, we can integrate both sides of (14.4) with respect to a probability measure F on θ and use Fubini's theorem to interchange the order of integration with respect to P and F. To achieve an effect similar to maximizing the random function $\exp\{\theta'A - \theta'C\theta/2\}$, F would need to assign positive mass to and near $\theta = C^{-1}A$ that maximizes $\exp\{\theta'A - \theta'C\theta/2\}$, for all possible realizations of (A, C). This leads us to choose probability measures of the form $dF(\theta) = f(\theta)d\theta$, with f positive and continuous. Note that

$$\int_{\mathbb{R}^d} e^{\theta'A-\theta'C\theta/2} f(\theta)d\theta = e^{A'C^{-1}A/2} \int_{\mathbb{R}^d} e^{-(\theta-C^{-1}A)'C(\theta-C^{-1}A)/2} f(\theta)d\theta. \qquad (14.7)$$

Let $\lambda_{\max}(C)$ and $\lambda_{\min}(C)$ denote the maximum and minimum eigenvalues of C, respectively. Since $(\theta - C^{-1}A)'C(\theta - C^{-1}A) \geq \lambda_{\min}(C)\|\theta - C^{-1}A\|^2$, it follows that as $\lambda_{\min}(C) \to \infty$,

$$\int_{\mathbb{R}^d} e^{-(\theta-C^{-1}A)'C(\theta-C^{-1}A)/2} f(\theta)d\theta \sim \frac{(2\pi)^{m/2}}{\sqrt{\det C}} f(C^{-1}A). \qquad (14.8)$$

Combining (14.7) with (14.8) yields Laplace's asymptotic formula that relates the integral on the left-hand side of (14.7) to the maximum value $\exp(A'C^{-1}A/2)$ of $\exp\{\theta'A - \theta'C\lambda/2\}$. Thus integration of $\exp(\theta'A - \theta'C\theta/2)$ with respect to the measure F provides "pseudo-maximization" of the integrand over θ when $\lambda_{\min}(C) \to \infty$. By choosing f appropriately to reflect the growth rate of $C^{-1/2}A$, we can extend the moment and exponential inequalities in Sect. 12.2 to the multivariate case. In particular, we shall prove the following two theorems and a related lemma.

Theorem 14.7. *Let A and C satisfy the canonical assumption* (14.4). *Let V be a positive definite nonrandom matrix. Then*

$$E\left[\sqrt{\frac{\det(V)}{\det(C+V)}}\exp\left\{\frac{1}{2}A'(C+V)^{-1}A\right\}\right] \leq 1, \qquad (14.9)$$

$$E\exp\{A'(C+V)^{-1}A/4\} \leq \left\{E\sqrt{\det(I+V^{-1}C)}\right\}^{\frac{1}{2}}. \qquad (14.10)$$

Proof. Put $f(\theta) = (2\pi)^{-d/2}\sqrt{\det V}\exp(-\theta'V\theta/2)$, $\theta \in \mathbb{R}^d$, in (14.7) after multiplying both sides of (14.4) by $f(\theta)$ and integrating over θ. By Fubini's theorem,

JMLR: Workshop and Conference Proceedings vol 35:1–17, 2014

lil' UCB : An Optimal Exploration Algorithm for Multi-Armed Bandits *

Kevin Jamieson KGJAMIESON@WISC.EDU
Matthew Malloy MMALLOY@WISC.EDU
Robert Nowak NOWAK@ECE.WISC.EDU
University of Wisconsin

Sébastien Bubeck SBUBECK@PRINCETON.EDU
Princeton University

Abstract

The paper proposes a novel upper confidence bound (UCB) procedure for identifying the arm with the largest mean in a multi-armed bandit game in the fixed confidence setting using a small number of total samples. The procedure cannot be improved in the sense that the number of samples required to identify the best arm is within a constant factor of a lower bound based on the law of the iterated logarithm (LIL). Inspired by the LIL, we construct our confidence bounds to explicitly account for the infinite time horizon of the algorithm. In addition, by using a novel stopping time for the algorithm we avoid a union bound over the arms that has been observed in other UCB-type algorithms. We prove that the algorithm is optimal up to constants and also show through simulations that it provides superior performance with respect to the state-of-the-art.

Keywords: Multi-armed bandit, upper confidence bound (UCB), iterated logarithm

1. Introduction

This paper introduces a new algorithm for the *best arm* problem in the stochastic multi-armed bandit (MAB) setting. Consider a MAB with n arms, each with unknown mean payoff $\mu_1, \dots, \mu_n$ in $[0, 1]$. A sample of the ith arm is an independent realization of a sub-Gaussian random variable with mean μ_i. In the *fixed confidence setting*, the goal of the best arm problem is to devise a sampling procedure with a single input δ that, regardless of the values of $\mu_1, \dots, \mu_n$, finds the arm with the largest mean with probability at least $1 - \delta$. More precisely, best arm procedures must satisfy $\sup_{\mu_1,\dots,\mu_n} \mathbb{P}(\widehat{i} \neq i^*) \leq \delta$, where i^* is the best arm, $\widehat{i}$ an estimate of the best arm, and the supremum is taken over all set of means such that there exists a unique best arm. In this sense, best arm procedures must automatically adjust sampling to ensure success when the mean of the best and second best arms are arbitrarily close. Contrast this with the *fixed budget setting* where the total number of samples remains a constant and the confidence in which the best arm is identified within the given budget varies with the setting of the means. While the fixed budget and fixed confidence settings are related (see Gabillon et al. (2012) for a discussion) this paper focuses on the fixed confidence setting only.

The best arm problem has a long history dating back to the '50s with the work of Paulson (1964); Bechhofer (1958). In the fixed confidence setting, the last decade has seen a flurry of

* Part of the research described here was carried out at the Simons Institute for the Theory of Computing. We are grateful to the Simons Institute for providing a wonderful research environment.

Note that the algorithm obtains the optimal query complexity of $\mathbf{H}_1 \log(1/\delta) + \mathbf{H}_3$ up to constant factors. We remark that the theorem holds with any value of λ satisfying (7). Inspection of (7) shows that as $\delta \to 0$ we can let λ tend to $\left(\frac{2+\beta}{\beta}\right)^2$. We point out that the sample complexity bound in the theorem can be optimized by choosing ε and β. For a setting of these parameters in a way that is more or less faithful to the theory, we recommend taking $\varepsilon = 0.01$, $\beta = 1$, and $\lambda = \left(\frac{2+\beta}{\beta}\right)^2$. For improved performance in practice, we recommend applying footnote 2 and setting $\varepsilon = 0$, $\beta = 0.5$, $\lambda = 1 + 10/n$ and $\delta \in (0, 1)$, which do not meet the requirements of the theorem, but work very well in our experiments presented later. We prove the theorem via two lemmas, one for the total number of samples taken from the suboptimal arms and one for the correctness of the algorithm. In the lemmas we give precise constants.

4. Proof of Theorem 2

Before stating the two main lemmas that imply the result, we first present a finite form of the law of iterated logarithm. This finite LIL bound is necessary for our analysis and may also prove useful for other applications.

Lemma 3 *Let $X_1, X_2, \ldots$ be i.i.d. centered sub-Gaussian random variables with scale parameter σ. For any $\varepsilon \in (0,1)$ and $\delta \in (0, \log(1+\varepsilon)/e)^2$ one has with probability at least $1 - \frac{2+\varepsilon}{\varepsilon}\left(\frac{\delta}{\log(1+\varepsilon)}\right)^{1+\varepsilon}$ for all $t \geq 1$,*

$$\sum_{s=1}^{t} X_s \leq (1+\sqrt{\varepsilon})\sqrt{2\sigma^2(1+\varepsilon)t\log\left(\frac{\log((1+\varepsilon)t)}{\delta}\right)}.$$

Proof We denote $S_t = \sum_{s=1}^{t} X_s$, and $\psi(x) = \sqrt{2\sigma^2 x \log\left(\frac{\log(x)}{\delta}\right)}$. We also define by induction the sequence of integers (u_k) as follows: $u_0 = 1$, $u_{k+1} = \lceil (1+\varepsilon)u_k \rceil$.
Step 1: Control of $S_{u_k}, k \geq 1$. The following inequalities hold true thanks to an union bound together with Chernoff's bound, the fact that $u_k \geq (1+\varepsilon)^k$, and a simple sum-integral comparison:

$$\begin{aligned}
\mathbb{P}\left(\exists k \geq 1 : S_{u_k} \geq \sqrt{1+\varepsilon}\,\psi(u_k)\right) &\leq \sum_{k=1}^{\infty} \exp\left(-(1+\varepsilon)\log\left(\tfrac{\log(u_k)}{\delta}\right)\right) \\
&\leq \sum_{k=1}^{\infty}\left(\tfrac{\delta}{k\log(1+\varepsilon)}\right)^{1+\varepsilon} \leq \left(1+\tfrac{1}{\varepsilon}\right)\left(\tfrac{\delta}{\log(1+\varepsilon)}\right)^{1+\varepsilon}.
\end{aligned}$$

Step 2: Control of $S_t, t \in (u_k, u_{k+1})$. Adopting the notation $[n] = \{1, \ldots, n\}$, recall that Hoeffding's maximal inequality[3] states that for any $m \geq 1$ and $x > 0$ one has

$$\mathbb{P}(\exists\, t \in [m] \text{ s.t. } S_t \geq x) \leq \exp\left(-\tfrac{x^2}{2\sigma^2 m}\right).$$

2. Note δ is restricted to guarantee that $\log(\frac{\log((1+\varepsilon)t)}{\delta})$ is well defined. This makes the analysis cleaner but in practice one can allow the full range of δ by using $\log(\frac{\log((1+\varepsilon)t+2)}{\delta})$ instead and obtain the same theoretical guarantees.
3. It is an easy exercise to verify that Azuma-Hoeffding holds for martingale differences with sub-Gaussian increments, which implies Hoeffding's maximal inequality for sub-Gaussian distributions.

Victor H. de la Pena
Department of Statistics
Columbia University

TAYLOR'S LAW

It has been used as a tool in studies in Demography, Biology and Physics, among others. Thousands of papers have been dedicated to the study of Taylor's Law. Important surveys on the topic are Kendal (2004), Eisler et al (2008) and Meng (2015). This talk is partly based partly in joint work with Mark Brown and Joel Cohen (2017).

Taylor's laws *Nature 1961*

In multiple sets of samples, the variance of population density is proportional to a power of the mean population density.

$$\text{variance} = a(\text{mean})^b, \ a > 0.$$
$$\log(\text{variance}) = \log(a) + b \cdot \log(\text{mean}).$$
$$\text{variance}/(\text{mean})^b = a, \ a > 0.$$

Taylor stated no model of error (deviations from exact equality).

Aphids, $1.29 \leq b \leq 2.95$

Each dot is (mean, var) in 1 year from sites in Great Britain, one species per line:

spatial TL.

Taylor, Woiwod & Perry, J. Anim. Ecol. 1980

3/4/2016

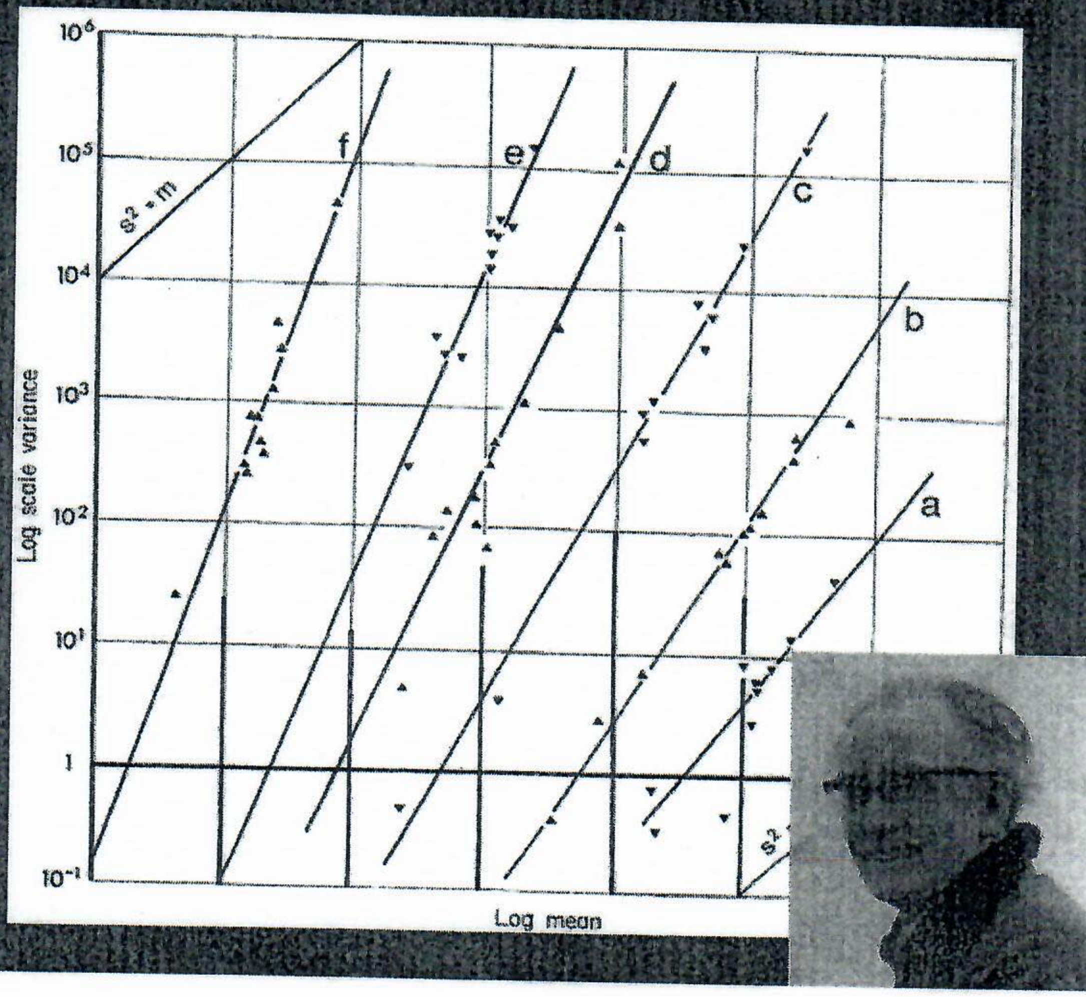

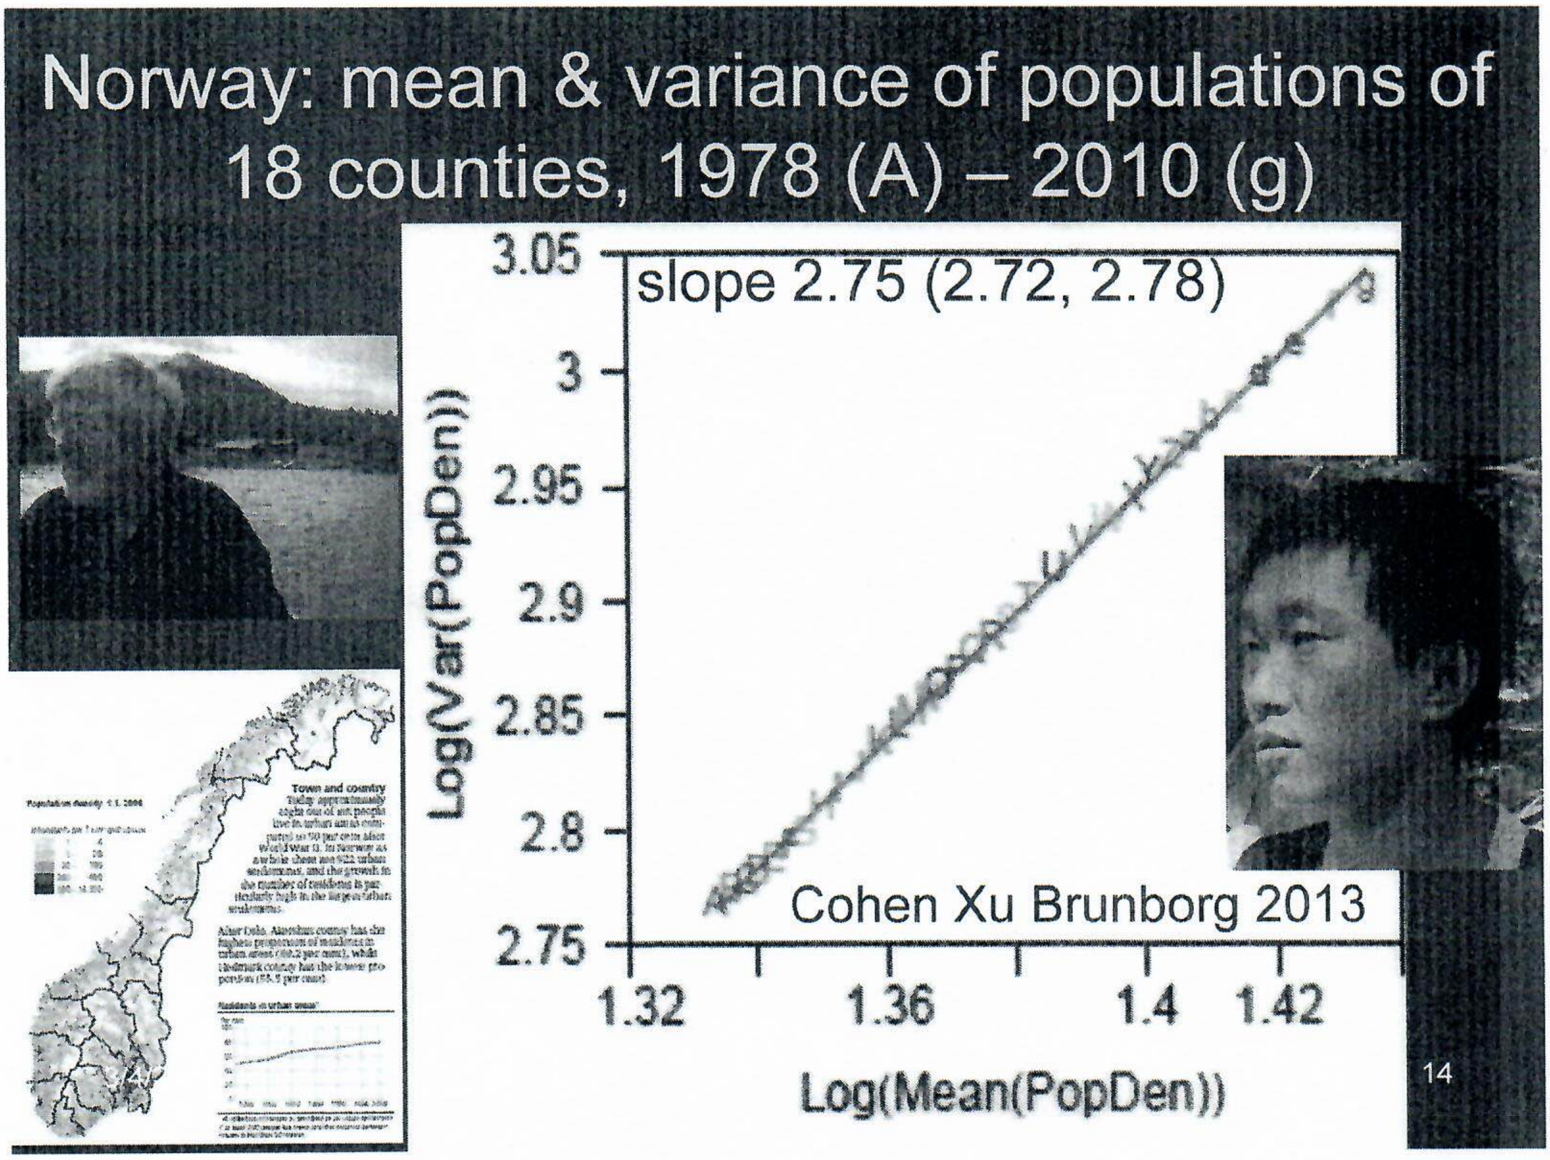

Norway: mean & variance of populations of 18 counties, 1978 (A) – 2010 (g)
slope 2.75 (2.72, 2.78)
Cohen Xu Brunborg 2013
Log(Var(PopDen))
Log(Mean(PopDen))
3.05
3
2.95
2.9
2.85
2.8
2.75
1.32
1.36
1.4
1.42
Town and country
14