Optimistic Optimization of a Brownian

JEAN-BASTIEN GRILL, MICHAL VALKO, RÉMI MUNOS
jbgrill@google.com and valkom@google.com and munos@google.com

Brownian

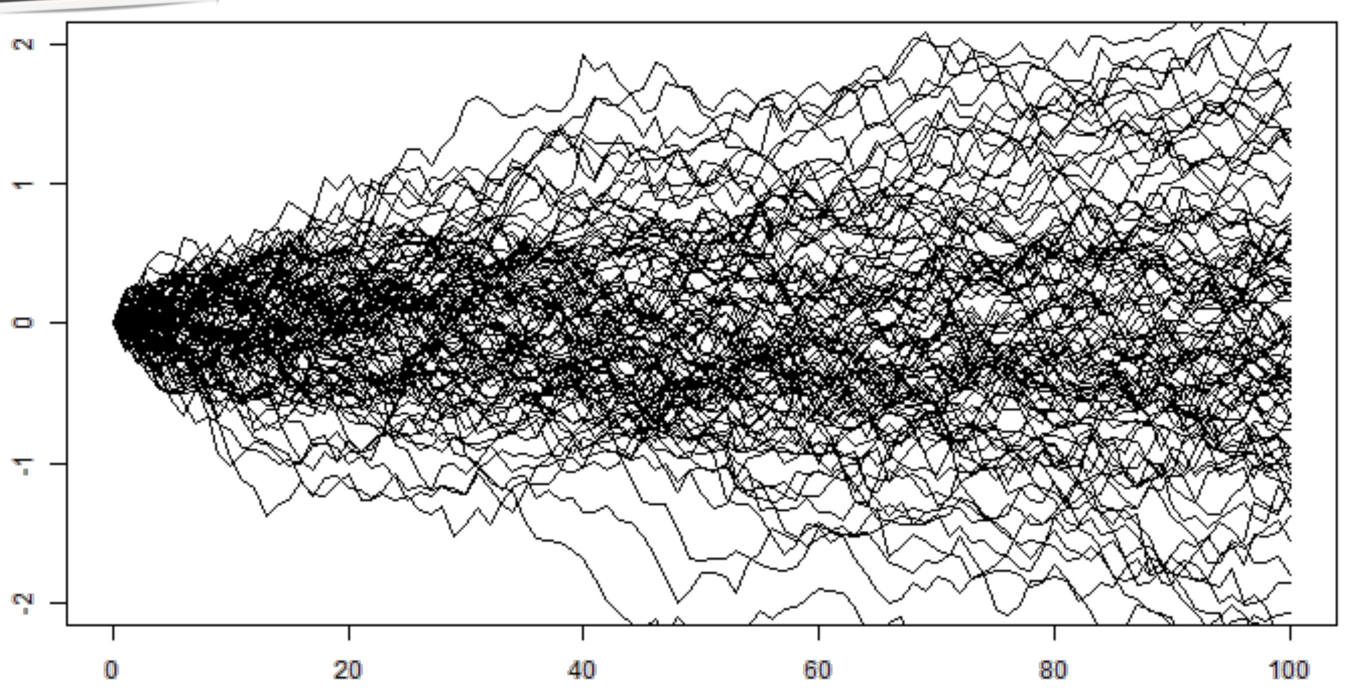

À la une

- GOAL: Finding maximum of a Brownian motion
- PRIOR WORK: Polynomial sample complexity
- CONTRIBUTION: Exponentially fast algorithm
- SOLVES OPEN PROBLEM: dimension of a Brownian

Setting

- **Every round t**: based on history $W(t_1), \dots, W(t_{n-1})$
  - choose t_n and observe $W(t_n)$
- **Objective**: return $\widehat{t}$ s.t. w.p. 1-ε $M - W(\widehat{t}) \le \varepsilon$
  - where $M \triangleq \sup_{t\in[0,1]} W(t)$

Who cares

- Type 1: **Exists prior to optimization**
  - simulation of financial stocks
- Type 2: **Observing the function creates it**
  - Gaussian process/Thomson sampling
- Type 3: **As a tool beyond learning**
  - computationally sample a solution of a stochastic differential equation in the work of Hefter and Herzwurm (2017)

Prior work

- Al-Mharmah and Calvin (1996)
  - a **non-adaptive** method
  - sample complexity: 1/√ε
- Calvin (2017)
  - **adaptive** method
  - better than any polynomial
  - *does not guarantee an exponential rate*

Algorithm

Algorithm 1 OOB algorithm

1: **Input**: ε
2: **Init**: $\mathcal{I} \leftarrow \{[0,1]\}$, $t_1 = W(1)$
3: **for** $i = 2, 3, 4, \dots$ **do**
4: $[a,b] \in \arg\max_{I\in\mathcal{I}} B_I$ {*break ties arbitrarily*}
5: **if** $\eta_\varepsilon(b-a) \le \varepsilon$ **then**
6: **break**
7: **end if**
8: $t_i \leftarrow W\left(\frac{a+b}{2}\right)$
9: $\mathcal{I} \leftarrow \{\mathcal{I} \cup [a, \frac{a+b}{2}] \cup [\frac{a+b}{2}, b]\} \setminus \{[a,b]\}$
10: **end for**
11: **Output**: location $\widehat{t}_\varepsilon \leftarrow \arg\max_{t_i} W(t_i)$ and its value $W(\widehat{t}_\varepsilon)$

$$B_{[a,b]} \triangleq \max(W(a), W(b)) + \eta_\varepsilon(b-a), \quad \text{where} \quad \forall \delta > 0 \text{ s.t. } \varepsilon\delta \le \frac{1}{2}, \quad \eta_\varepsilon(\delta) \triangleq \sqrt{\frac{5\delta}{2}\ln\left(\frac{2}{\varepsilon\delta}\right)}$$

Guarantees

THEOREM: For any ε < 1/2

$$\mathbb{P}\left[M - W\left(\widehat{t}_\varepsilon\right) > \varepsilon\right] \le \varepsilon$$ (correctness)

$$\mathbb{E}[N_\varepsilon] \le c \log^2(1/\varepsilon)$$ (sample complexity)

Corollary: Also (δ,ε) PAC

Proof

- 1. **Correctness:** algorithm definition + the law of Brownian bridge
- 2. at OOB evaluates pretty much only near-optimal points
- Denisov (1984): rewrite the motion as two Brownian meanders
- By Durett et al. (1977) the expected number of near-optimal points is bounded as $\mathbb{E}[\mathcal{N}_h(\eta)] \le 6\eta^2 2^h$ which is $\mathcal{O}(\log(1/\varepsilon))$

Open problem

- Munos (2011) classifies functions according to (d, C) to:
  - easy, d = 0, exponentially fast optimization
  - difficult, d ≥ 0, polynomially fast optimization
- Open questions for a Brownian process:
  - what is its dimension d
  - how fast can we optimize it

Challenge: Brownian motion is a stochastic process!

Our answers that solve the open problem:

- ∀ε, w.p. 1-ε, W(t) is ℓ_ε-Lipschitz + ∃ C(ε) s.t. Brownian ∈ (d, C(ε))
- there is no (d, C) with C < ∞ such that Brownian ∈ (d, C)
- we can optimize it with sample complexity of $O(\log^2(1/\varepsilon))$

Experiments

the last paper 100%

the first paper 100%