

---

# Perturbation, Optimization, and Statistics

Editors:

**Tamir Hazan**

*Technion - Israel Institute of Technology  
Technion City, Haifa 32000, Israel*

tamir.hazan@technion.ac.il

**George Papandreou**

*Google Inc.  
340 Main St., Los Angeles, CA 90291 USA*

gpapan@google.com

**Daniel Tarlow**

*Microsoft Research  
Cambridge, CB1 2FB, United Kingdom*

dtarlow@microsoft.com

This is a draft version of the author chapter.

The MIT Press  
Cambridge, Massachusetts  
London, England



**Chris J. Maddison**

*University of Toronto*

*Toronto, Canada*

cmaddis@cs.toronto.edu

*Simulating samples from arbitrary probability distributions is a major research program of statistical computing. Recent work has shown promise in an old idea, that sampling from a discrete distribution can be accomplished by perturbing and maximizing its mass function. Yet, it has not been clearly explained how this research project relates to more traditional ideas in the Monte Carlo literature. This chapter addresses that need by identifying a Poisson process model that unifies the perturbation and accept-reject views of Monte Carlo simulation. Many existing methods can be analyzed in this framework. The chapter reviews Poisson processes and defines a Poisson process model for Monte Carlo methods. This model is used to generalize the perturbation trick to infinite spaces by constructing Gumbel processes, random functions whose maxima are located at samples over infinite spaces. The model is also used to analyze  $A^*$  sampling and  $OS^*$ , two methods from distinct Monte Carlo families.*

---

## 7.1 Introduction

The simulation of random processes on computers is an important tool in scientific research and a subroutine of many statistical algorithms. One way to formalize this task is to return samples from some distribution given access to a density or mass function and to a pseudorandom number generator that returns independent uniform random numbers. “Monte Carlo methods”, a phrase originally referring to the casinos of Monte Carlo, is a catchall for algorithms that solve this problem. Many Monte Carlo methods exist for spe-


cific distributions or classes of distributions (Walker, 1977; Devroye, 1986), but there are a few generic principles. One principle is to simulate a Markov chain whose stationary distribution is the distribution of interest. Work on these Markov chain Monte Carlo methods has exploded over the past few decades, because of their efficiency at sampling from complex distributions in high dimensions. Their downside is that convergence can be slow and detecting convergence is hard. A second principle is propose samples from a tractable distribution and accept them according to a correction factor. These accept-reject Monte Carlo methods are the workhorses of modern statistical packages, but their use is restricted to simple distributions on low dimensional spaces.

Recently, a research program has developed around another principle for sampling from discrete distributions, the so called ‘‘Gumbel-Max trick’’. The trick proceeds by simulating a random function  $G : \{1, \dots, m\} \rightarrow \mathbb{R}$  whose maximum is located at a sample. Sampling therefore reduces to finding the state that maximizes  $G$ . This trick has the same complexity as better known methods, but it has inspired research into approximate methods and extensions. Methods that abandon exactness for efficiency have considered introducing correlated  $G$  with a variety of applications (Papandreou and Yuille, 2011; Tarlow et al., 2012; Hazan et al., 2013). Chen and Ghahramani (2015) consider bandit algorithms for optimizing  $G$  over low dimensional spaces when function evaluation is expensive. Maddison et al. (2014) generalized  $G$  with Gumbel processes, random functions over infinite spaces whose maxima occur at samples of arbitrary distributions, and introduced  $A^*$  sampling, a branch and bound algorithm that executes a generalized Gumbel-Max trick. Kim et al. (2016) introduced a related branch and bound algorithm tailored to discrete distributions and successfully sampled from a large fully connected attractive Ising model. Taken together, this view of simulation as a maximization problem is a promising direction, because it connects Monte Carlo research with the literature on optimization. Yet, its relationship to more established methods has not been clearly expressed. This chapter addresses that need by identifying a model that jointly explains both the accept-reject principle and the Gumbel-Max trick.

As a brief introduction, we cover a simple example of an accept-reject algorithm and the Gumbel-Max trick shown in Figure 7.1. Suppose we are given a positive function  $f : \{1, \dots, m\} \rightarrow \mathbb{R}^+$ , which describes the unnormalized mass of a discrete random variable  $I$ ,

$$\mathbb{P}(I \in B) = \sum_{i \in B} \frac{f(i)}{\sum_{j=1}^m f(j)}, \quad B \subseteq \{1, \dots, m\}. \quad (7.1)$$

The following algorithms return an integer with the same distribution as  $I$ .



**Figure 7.1:** Two simple Monte Carlo methods for a discrete distribution described by positive function  $f$  via (7.1). The left hand plot shows the first accepted sample  $J$  in an accept-reject scheme; note that  $U < f(J)$ . The right hand plot shows a sample  $I^*$  in the Gumbel-Max trick;  $I^*$  is the state that achieves the maximum  $G^* = \max_i \log f(i) + G(i)$ .

The accept-reject algorithm is,

1. Sample  $J$  uniformly from  $\{1, \dots, m\}$ ,  $U$  uniformly from  $[0, \max_{i=1}^m f(i)]$ ,
2. If  $U < f(J)$ , return  $J$ , else go to 1.

We can intuitively justify it by noticing that accepted pair  $(J, U)$  falls uniformly under the graph of  $f(i)$ , Figure 7.1. The sample  $J$ , which is accepted or rejected, is often called a *proposal*. The Gumbel-Max trick proceeds by optimizing a random function,

1. For  $i \in \{1, \dots, m\}$  sample an independent Gumbel random variable  $G(i)$ .
2. Find and return  $I^* = \operatorname{argmax}_{i=1}^m \log f(i) + G(i)$ .

Because the random values  $\log f(i) + G(i)$  can be seen as a perturbed negative energy function, the function  $G$  is often called a *perturbation*. Uniform and Gumbel random variables are included among the standard distributions of statistical computing packages. So these algorithms, while inefficient, are simple to program.

Considering their apparent differences and the fact that they have been studied in distinct literatures, it is surprising that both algorithms can be unified under the same theoretical framework. The framework rests on the study of Poisson processes, a random object whose value is a countable set of points in space (Kingman, 1992; Daley and Vere-Jones, 2007). The central idea is to define a specific Poisson process, called an exponential race, which models a sequence of independent samples arriving from some distribution. Then we identify two operations, corresponding to accept-reject and the Gumbel-Max trick, which modify the arrival distribution of exponential


paces. In this view a Monte Carlo method is an algorithm that simulates the first arrival of an exponential race, and many existing algorithms fall into this framework.

Section 7.2 reviews Poisson processes and studies the effect of operations on their points. Section 7.3 introduces exponential races and studies the accept-reject and perturb operations. In Section 7.4 we construct Gumbel processes from exponential races and study the generalized Gumbel-Max trick. In Section 7.5 we analyze  $A^*$  sampling and  $OS^*$  (Dymetman et al., 2012) and show how they use perturb and accept-reject operations, respectively, to simulate the first arrival of an exponential race. All of our Poisson process results are either known or elementary extensions; the correctness and behaviour of the Monte Carlo methods that we study have all been established elsewhere. Our contribution is in identifying a theory that unifies two distinct literatures and in providing a toolset for analyzing and developing Monte Carlo methods.

## 7.2 Poisson processes

### 7.2.1 Definition and properties

A Poisson process is a random countable subset  $\Pi \subseteq \mathbb{R}^n$ . Many natural processes result in a random placement of points: the stars in the night sky, cities on a map, or raisins in oatmeal cookies. A good generic mental model to have is the plane  $\mathbb{R}^2$  and pinpricks of light for all points in  $\Pi$ . Unlike most natural processes, a Poisson process is distinguished by its complete randomness; the number of points in disjoint subsets are



**Figure 7.2:** The set of  $*$  is a realization of a Poisson process in the plane. Counts in sets  $A, B, C$  are marginally Poisson and are independent for disjoint sets.

independent random variables, see Figure 7.2. In this section we review a general Poisson process theory culminating in two theorems, which describe how they behave under the generic operations of removing or relocating their points. In the next section we restrict our view to a specific Poisson process and two specific operations, which correspond to accept-reject and Gumbel-Max. Our study is situated in  $\mathbb{R}^n$  for intuition, but these results generalize naturally; for more information, the ideas of this section are adapted from the general treatment in Kingman (1992). Readers familiar with that treatment can safely skip this section

To identify a realization of a random countable set  $\Pi \subseteq \mathbb{R}^n$ , we use counts of points in subsets  $B \subseteq \mathbb{R}^n$ ,

$$N(B) = \#(\Pi \cap B).$$

where  $N(B) = \infty$  if  $B$  is infinite, see Figure 7.2 again. Counts are nonnegative and additive, so for any realization of  $\Pi$   $N(B)$  satisfies

1. (*Nonnegative*)  $N(B) \geq 0$ ,
2. (*Countably additive*) For disjoint  $B_i \subseteq \mathbb{R}^n$ ,  $N(\cup_{i=1}^{\infty} B_i) = \sum_{i=1}^{\infty} N(B_i)$ .

Set functions from subsets of  $\mathbb{R}^n$  to the extended reals  $\mathbb{R} \cup \{\infty, -\infty\}$  that are nonnegative and countably additive are called measures. Measure theory is a natural backdrop for the study of Poisson processes, so we briefly mention some basic concepts. In general measures  $\mu$  assign real numbers to subsets with the same consistency that we intuitively expect from measuring lengths or volumes in space. If  $\mu(\mathbb{R}^n) = 1$ , then  $\mu$  is a probability distribution. Because it is not possible to define a measure consistently for all possible subsets, the subsets  $B \subseteq \mathbb{R}^n$  are restricted here and throughout the chapter to be from the Borel sets, a nice measurable family of subsets. The Borel sets contain almost any set of interest, so for our purposes it is practically unrestricted. Integration of some function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  with respect to some measure  $\mu$  naturally extends Riemann integration, which we can think about intuitively as the area under the graph of  $f(x)$  weighted by the instantaneous measure  $\mu(dx)$ . When a measure is equal to the integral of a nonnegative function  $f : \mathbb{R}^n \rightarrow \mathbb{R}^{\geq 0}$  with respect to  $\mu$ , we say  $f$  is the *density* with respect to  $\mu$ .

The Poisson process receives its name from the marginal distribution of counts  $N(B)$ .  $N(B)$  is Poisson distributed on the nonnegative integers parameterized by a rate, which is also its expected value.

**Definition 7.1** (Poisson random variable).  *$N$  is a Poisson distributed*

random variable on  $k \in \{0, 1, \dots\}$  with nonnegative rate  $\lambda \in \mathbb{R}^{\geq 0}$  if

$$\mathbb{P}(N = k) = \exp(-\lambda) \frac{\lambda^k}{k!}.$$

This is denoted  $N \sim \text{Poisson}(\lambda)$ .  $N \sim \text{Poisson}(0)$  and  $N \sim \text{Poisson}(\infty)$  are the random variables whose values are 0 and  $\infty$  with probability one. If  $N \sim \text{Poisson}(\lambda)$ , then  $\mathbb{E}(N) = \lambda$ .

The Poisson distribution is particularly suited to modelling random counts, because it is countably additive in the rate.

**Lemma 7.1.** *If  $N_i \sim \text{Poisson}(\lambda_i)$  independent with  $\lambda_i \in \mathbb{R}^{\geq 0}$ , then*

$$\sum_{i=1}^{\infty} N_i \sim \text{Poisson} \left( \sum_{i=1}^{\infty} \lambda_i \right).$$

*Proof.* (Kingman, 1992). Let  $S_m = \sum_{i=1}^m N_i$  and assume  $\lambda_i > 0$  without loss of generality. Then for  $S_2$ ,

$$\begin{aligned} \mathbb{P}(S_2 = k) &= \sum_{r=0}^k \mathbb{P}(N_1 = r, N_2 = k - r) \\ &= \sum_{r=0}^k \exp(-\lambda_1) \frac{\lambda_1^r}{r!} \exp(-\lambda_2) \frac{\lambda_2^{k-r}}{(k-r)!} \\ &= \frac{\exp(-\lambda_1 - \lambda_2)}{k!} \sum_{r=0}^k \binom{k}{r} \lambda_1^r \lambda_2^{k-r} \\ &= \frac{\exp(-\lambda_1 - \lambda_2)}{k!} (\lambda_1 + \lambda_2)^k. \end{aligned}$$

By induction Lemma 7.1 also holds for  $S_m$ . For infinite sums the events  $\{S_m \leq k\}$  are nonincreasing. Thus,

$$\mathbb{P}(S_{\infty} \leq k) = \lim_{m \rightarrow \infty} \mathbb{P}(S_m \leq k) = \sum_{j=1}^k \lim_{m \rightarrow \infty} \exp \left( - \sum_{i=1}^m \lambda_i \right) \frac{(\sum_{i=1}^m \lambda_i)^j}{j!}.$$

□

Because expectations distribute over infinite sums of positive random variables, the Poisson rate  $\mu(B) = \mathbb{E}(N(B))$  must also be a measure.

Instead of starting with a definition of Poisson processes, we work backwards from an algorithmic construction. Algorithm 7.1 is a procedure that realizes a Poisson process  $\Pi$  for a specified mean measure  $\mu$ . Algorithm 7.1 iterates through a partition  $\{B_i\}_{i=1}^{\infty}$  of  $\mathbb{R}^n$ . For each  $B_i$  it first decides the number of points to place in  $\Pi$  by sampling a Poisson with rate given by the measure,  $N_i \sim \text{Poisson}(\mu(B_i))$ . Then, it places  $N_i$  points by sampling



---

**Algorithm 7.1** A Poisson process  $\Pi$  with  $\sigma$ -finite nonatomic mean measure  $\mu$

---

```

Let  $\{B_i\}_{i=1}^\infty$  be a partition of  $\mathbb{R}^n$  with  $\mu(B_i) < \infty$ 
 $\Pi = \emptyset$ 
for  $i = 1$  to  $\infty$  do
   $N_i \sim \text{Poisson}(\mu(B_i))$ 
  for  $j = 1$  to  $N_i$  do
     $X_{ij} \sim \mu(\cdot \cap B_i) / \mu(B_i)$ 
     $\Pi = \Pi \cup \{X_{ij}\}$ 
  end for
end for

```

---

independently from the probability distribution proportional to  $\mu$  restricted to  $B_i$ . Normally,  $X \sim \mathcal{D}$  is just a statement about the marginal distribution of  $X$ . In the context of an Algorithm box we also implicitly assume that it implies independence from all other random variables. We should note that Algorithm 7.1 operates on volumes and samples from  $\mu$ . This is not an issue, if we think of it as a mathematical construction. It would be an issue, if we set out to simulate  $\Pi$  on a computer.

Algorithm 7.1 will occasionally have pathological behaviour, unless we restrict  $\mu$  further. First, we require that each subset  $B_i$  of the partition has finite measure; if  $\mu(B_i) = \infty$ , then Algorithm 7.1 will stall when it reaches  $B_i$  and fail to visit all of  $\mathbb{R}^n$ . If a partition  $\{B_i\}_{i=1}^\infty$  with  $\mu(B_i) < \infty$  exists for measure  $\mu$ , then  $\mu$  is called  $\sigma$ -finite. Second, we want the resulting counts  $N(B_i)$  to match the number of points placed  $N_i$ . This can be ensured if all of the points  $X_{ij}$  are distinct with probability one. It is enough to require that  $\mu(\{x\}) = 0$  for all singleton sets  $x \in \mathbb{R}^n$ . This kind of measure is known as nonatomic.

The crucial property of the sets  $\Pi$  produced by Algorithm 7.1 is that the number of points  $N(A_j)$  that fall in *any* finite collection  $\{A_j\}_{j=1}^m$  of disjoint sets are independent Poisson random variables. Clearly, the counts  $N(B_i)$  for the partitioning sets of Algorithm 7.1 are independent Poissons; it is not obvious that this is also true for other collections of disjoint sets. To show this we study the limiting behaviour of  $N(B)$  by counting the points placed in  $B_i \cap B$  and summing as Algorithm 7.1 iterates over  $\mathbb{R}^n$ .

**Theorem 7.2.** *Let  $\Pi \subseteq \mathbb{R}^n$  be the subset realized by Algorithm 7.1 with  $\sigma$ -finite nonatomic mean measure  $\mu$  and  $A_1, \dots, A_m \subseteq \mathbb{R}^n$  disjoint.  $N(B) = \#(\Pi \cap B)$  for  $B \subseteq \mathbb{R}^n$  satisfies*

1.  $N(A_j) \sim \text{Poisson}(\mu(A_j))$ ,
2.  $N(A_j)$  are independent.

*Proof.* Adapted from Kingman (1992). Let  $B_i$  be the partition of Algo-

rithm 7.1 with  $\mu(B_i) > 0$  without loss of generality. With probability one,

$$N(A_j) = N(\cup_{i=1}^{\infty} B_i \cap A_j) = \sum_{i=1}^{\infty} N(B_i \cap A_j).$$

Consider the array of  $N(B_i \cap A_j)$  for  $i \in \{1, 2, \dots\}$  and  $j \in \{1, \dots, m\}$ . The rows are clearly independent. Thus, by Lemma 7.1 it is enough to show

1.  $N(B_i \cap A_j) \sim \text{Poisson}(\mu(B_i \cap A_j))$ ,
2.  $N(B_i \cap A_j)$  for  $j \in \{1, \dots, m\}$  are independent,

Let  $A_0$  be the complement of  $\cup_{i=1}^m A_i$ . Because  $\mu$  is nonatomic, each point is distinct with probability one. Thus,

$$\begin{aligned} \mathbb{P}(N(B_i \cap A_0) = k_0, \dots, N(B_i \cap A_m) = k_m | N_i = k) = \\ \frac{k!}{k_0! \dots k_m!} \prod_{j=0}^m \frac{\mu(B_i \cap A_j)^{k_j}}{\mu(B_i)^{k_j}} \end{aligned}$$

with  $k_0 = k - \sum_{j=1}^m k_j$ . Now,

$$\begin{aligned} \mathbb{P}(N(B_i \cap A_1) = k_1, \dots, N(B_i \cap A_m) = k_m) = \\ \sum_{k=\sum_j k_j}^{\infty} \exp(-\mu(B_i)) \frac{\mu(B_i)^k}{k!} \frac{k!}{k_0! \dots k_m!} \prod_{j=0}^m \frac{\mu(B_i \cap A_j)^{k_j}}{\mu(B_i)^{k_j}} \\ \sum_{k_0=0}^{\infty} \prod_{j=0}^m \exp(-\mu(B_i \cap A_j)) \frac{\mu(B_i \cap A_j)^{k_j}}{k_j!} \\ = \prod_{j=1}^m \exp(-\mu(B_i \cap A_j)) \frac{\mu(B_i \cap A_j)^{k_j}}{k_j!}. \end{aligned}$$

finishes the proof. □

Notice that the partition in Algorithm 7.1 has an indistinguishable effect on the eventual counts  $N(B)$ . In fact there may be entirely different algorithms that realize random subsets indistinguishable from  $\Pi$ . This motivates the standard definition for deciding whether a random process is Poisson.

**Definition 7.2** (Poisson process). *Let  $\mu$  be a  $\sigma$ -finite nonatomic measure on  $\mathbb{R}^n$ . A random countable subset  $\Pi \subseteq \mathbb{R}^n$  is a Poisson process with mean measure  $\mu$  if*

1. For  $B \subseteq \mathbb{R}^n$ ,  $N(B) \sim \text{Poisson}(\mu(B))$ .
2. For  $A_1, \dots, A_m \subseteq \mathbb{R}^n$  disjoint,  $N(A_j)$  are independent.

Algorithm 7.1 together with Theorem 7.2 is an existence proof for Poisson

processes. Poisson processes are generic models for procedures that place points completely randomly in space. In later sections we specialize them to model the sequence of points considered by Monte Carlo methods.

### 7.2.2 Mapping and thinning a Poisson process

We are ultimately interested in understanding how the operations of accept-reject and the Gumbel-Max trick modify distributions. They are special cases of more generic operations on the points  $X \in \Pi$  of a Poisson process, which modify its measure. Accept-reject corresponds to the stochastic removal of points based on their location. The Gumbel-Max trick corresponds to the deterministic relocation of points. Here we study those operations in some generality.

The stochastic removal of points  $X \in \Pi$  is called thinning. To count the number of points that remain after thinning, we need their joint distribution before thinning. If we restrict our attention to one of the subsets  $B_i$  of the partition in Algorithm 7.1, then the distribution is clear: conditioned on  $N(B_i) = k$ , each point is distributed identically and independently (i.i.d.) as  $\mu$  restricted to  $B_i$ . This property turns out to be true for any subset  $B \subseteq \mathbb{R}^n$  of finite measure.

**Lemma 7.3.** *Let  $\Pi \subseteq \mathbb{R}^n$  be a Poisson Process with  $\sigma$ -finite nonatomic mean measure  $\mu$  and  $B \subseteq \mathbb{R}^n$  with  $0 < \mu(B) < \infty$ . Given  $N(B) = k$ , each  $X_i \in \Pi \cap B$  for  $i \in \{1, \dots, k\}$  is i.i.d. as,*

$$X_i | \{N(B) = k\} \sim \mu(\cdot \cap B) / \mu(B).$$

*Proof.* The proof is uninformative, so we leave it to the Appendix.  $\square$

Intuitively, this result ought to be true, because we could have realized  $\Pi$  via Algorithm 7.1 with  $B$  as one of the partitioning sets.

Now suppose we remove points  $X \in \Pi$  independently with probability  $1 - \rho(X)$ , where  $\rho : \mathbb{R}^n \rightarrow [0, 1]$  is some integrable function. For  $B$  with finite measure, given  $N(B)$  the probability of keeping  $X \in \Pi \cap B$  is

$$\mathbb{P}(\text{keep } X | N(B) = k) = \mathbb{E}(\rho(X) | N(B) = k) = \int_B \frac{\rho(x)}{\mu(B)} \mu(dx). \quad (7.2)$$

By summing over the value of  $N(B)$ , we can derive the marginal distribution over the number of remaining points. This is the basic strategy of the Thinning Theorem.

**Theorem 7.4** (Thinning). *Let  $\Pi \subseteq \mathbb{R}^n$  be a Poisson Process with  $\sigma$ -finite nonatomic mean measure  $\mu$  and  $S(x) \sim \text{Bernoulli}(\rho(x))$  an independent*

Bernoulli random variable for  $x \in \mathbb{R}^n$  with integrable  $\rho : \mathbb{R}^n \rightarrow [0, 1]$ , then

$$\text{thin}(\Pi, S) = \{X : X \in \Pi \text{ and } S(X) = 1\} \quad (7.3)$$

is a Poisson process with mean measure

$$\mu^*(B) = \int_B \rho(x) \mu(dx).$$

*Proof.* Originally from Lewis and Shedler (1979). Let  $B \subseteq \mathbb{R}^n$ . Define,

$$N^*(B) = \#(\text{thin}(\Pi, S) \cap B)$$

$N^*(B)$  clearly satisfies the independence property and the result is trivial for  $\mu(B) = 0$ . For  $0 < \mu(B) < \infty$ ,

$$\mathbb{P}(N^*(B) = k) = \sum_{j=k}^{\infty} \mathbb{P}(N(B) = j) \mathbb{P}(k \text{ of } S(X_i) = 1 | N(B) = j).$$

Let  $\bar{\mu}^*(B) = \mu(B) - \mu^*(B)$ . By (7.2),

$$\begin{aligned} &= \sum_{j=k}^{\infty} \exp(-\mu(B)) \frac{\mu(B)^j}{j!} \binom{j}{k} \frac{\mu^*(B)^k}{\mu(B)^k} \frac{\bar{\mu}^*(B)^{j-k}}{\mu(B)^{j-k}} \\ &= \exp(-\mu^*(B)) \frac{\mu^*(B)^k}{k!} \sum_{j=k}^{\infty} \exp(-\bar{\mu}^*(B)) \frac{\bar{\mu}^*(B)^{j-k}}{(j-k)!} \\ &= \exp(-\mu^*(B)) \frac{\mu^*(B)^k}{k!}. \end{aligned}$$

For  $\mu(B) = \infty$ , partition  $B$  into subsets with finite measure. The countable additivity of integrals of nonnegative functions and of Poisson random variables (Lemma 7.1) finishes the proof.  $\square$

A measurable function  $h : \mathbb{R}^n \rightarrow \mathbb{R}^n$  that relocates points  $X \in \Pi$  is easy to analyze if it is 1-1, because it will not relocate two distinct points to the same place. The key insight is that we can count the points relocated to  $B \subseteq \mathbb{R}^n$  by counting in the preimage  $h^{-1}(B)$ ; the so-called Mapping Theorem.

**Theorem 7.5** (Mapping). *Let  $\Pi \subseteq \mathbb{R}^n$  be a Poisson process with  $\sigma$ -finite nonatomic mean measure  $\mu$  and  $h : \mathbb{R}^n \rightarrow \mathbb{R}^n$  a measurable 1-1 function, then*

$$h(\Pi) = \{h(X) : X \in \Pi\}$$

is a Poisson process with mean measure

$$\mu^*(B) = \mu(h^{-1}(B))$$

*Proof.* Adapted from Kingman (1992).  $h$  is 1-1, therefore

$$\#\{h(X) : X \in \Pi \cap B\} = \#\{X \in \Pi : X \in h^{-1}(B)\} \sim \text{Poisson}(\mu(h^{-1}(B))).$$


Pre-images preserve disjointness, so the independence property is guaranteed. 1-1 functions map partitions of the domain to partitions of the range, so  $\mu^*$  is still  $\sigma$ -finite.  $\square$

## 7.3 Exponential races

### 7.3.1 Definition and first arrivals distribution

In this section we specialize the Poisson process to model the sequence of points considered by accept-reject and the Gumbel-Max trick. We call the model an exponential race as a reference to a classical example. An exponential race (occasionally race for short) is a Poisson process in  $\mathbb{R}^+ \times \mathbb{R}^n$ , which we interpret as points in  $\mathbb{R}^n$  ordered by an arrival time in the positive reals  $\mathbb{R}^+$ . The ordered points of an exponential race have a particularly simple distribution; the location in  $\mathbb{R}^n$  of each point is i.i.d. according to some arrival distribution and the rate at which points arrive in time depends stochastically on the normalization constant of that arrival distribution. The Thinning and Mapping Theorems of Poisson processes have corresponding lemmas for exponential races, which describe operations that modify the arrival distribution of an exponential race. The ultimate value of this model is that a variety of apparently disparate Monte Carlo methods can be interpreted as procedures that simulate an exponential race. In Section 7.5 we present Monte Carlo methods which produce samples from intractable distributions by operating on the simulation of an exponential race with a tractable distribution. In this section we define an exponential race for an arbitrary finite nonzero measure  $P$ , discuss strategies for simulating exponential races when  $P$  is tractable, and derive two operations that modify the arrival distribution of exponential races.

For motivation we review the traditional exponential race example (see Durrett, 2012). Imagine instantaneous flashes of light arriving in time at  $m$  distinct points  $p_j$  scattered in  $\mathbb{R}^2$ . Suppose the arrival times of the flashes at each  $p_j$  are determined by independent Poisson processes  $\Pi_j \subseteq \mathbb{R}^+$  with mean measure  $\lambda_j((0, t]) = \lambda_j t$  and  $\lambda_j > 0$ , see Figure 7.3. The question is



**Figure 7.3:** The realization of an exponential race with points arriving at  $p_j \in \mathbb{R}^2$ . The left hand plot shows the location of arrivals in the plane  $\mathbb{R}^2$  and the first arrival at time  $t$  at  $p_3$ . The right hand plot shows future arrival times at the four points.

which point will get the first flash of light and how long do we need to wait? The first arrival at  $p_j$  is after time  $t$  iff  $\Pi_j \cap (0, t]$  is empty,

$$\mathbb{P}(T_j > t) = \mathbb{P}(\#\Pi_j \cap (0, t] = 0) = \exp(-\lambda_j t). \quad (7.4)$$

(7.4) is the complementary cumulative distribution function of an exponential random variable, which we briefly review.

**Definition 7.3** (Exponential random variable).  *$E$  is an exponential random variable distributed on positive  $t \in \mathbb{R}^+$  with nonnegative rate  $\lambda \in \mathbb{R}^{\geq 0}$  if*

$$\mathbb{P}(E > t) = \exp(-\lambda t). \quad (7.5)$$

*This is denoted  $E \sim \text{Exp}(\lambda)$  and  $E \sim \text{Exp}(0)$  is the random variable whose value is  $\infty$  with probability one. If  $E \sim \text{Exp}(1)$ , then  $E/\lambda \sim \text{Exp}(\lambda)$ .*

Thus, the location and time of the first arrival is determined by the minimum of  $m$  exponential random variables. For exponential random variables this is particularly easy to analyze; the minimum is an exponential random variable with rate  $\sum_{j=1}^m \lambda_j$  and it is achieved at the  $j$ th variable with probability proportional to the rate  $\lambda_j$ . Surprisingly, these values are independent.

**Lemma 7.6.** *Let  $E_j \sim \text{Exp}(\lambda_j)$  independent with nonnegative  $\lambda_j \in \mathbb{R}^{\geq 0}$ . If*

$$E^* = \min_{1 \leq j \leq m} E_j \text{ and } J^* = \operatorname{argmin}_{1 \leq j \leq m} E_j,$$

*and at least one  $\lambda_j > 0$  then*

1. *The density of  $E_j$  with  $\lambda_j > 0$  is  $\lambda_j \exp(-\lambda_j t)$  for  $t \in \mathbb{R}^+$ ,*
2.  *$E^* \sim \text{Exp}(\sum_{j=1}^m \lambda_j)$ ,*

3.  $\mathbb{P}(J^* = k) \propto \lambda_k$ ,
4.  $E^*$  is independent of  $J^*$ .

*Proof.* 1. The derivative of  $1 - \exp(-\lambda_j t)$  is  $\lambda_j \exp(-\lambda_j t)$ .

2., 3., 4. Note that with probability 1 the  $E_j$  will be distinct, so

$$\begin{aligned}
 \mathbb{P}(J^* = k, E^* > t) &= \mathbb{P}(\cap_{j \neq k} \{E_j > E_k > t\}) \\
 &= \int_t^\infty \lambda_k \exp(-\lambda_k x) \prod_{j \neq k} \exp(-\lambda_j x) dx \\
 &= \frac{\lambda_k}{\sum_{j=1}^m \lambda_j} \int_t^\infty (\sum_{j=1}^m \lambda_j) \exp(-\sum_{j=1}^m \lambda_j x) dx \\
 &= \frac{\lambda_k}{\sum_{j=1}^m \lambda_j} \exp(-\sum_{j=1}^m \lambda_j t).
 \end{aligned}$$

This finishes the lemma.  $\square$

The extension of exponential races to arbitrary distributions on  $\mathbb{R}^n$  is straightforward. The  $m$  Poisson processes of the example are together a single Poisson process on  $\mathbb{R}^+ \times \mathbb{R}^n$  with mean measure  $(\lambda \times P)((0, t] \times B) = \sum_{j=1}^m t \lambda_j 1_B(p_j)$ .  $\lambda \times P$  is the product measure on  $\mathbb{R}^+ \times \mathbb{R}^n$ , where each is respectively equipped with  $\lambda((0, t]) = t$  and  $P(B) = \sum_j \lambda_j 1_B(p_j)$ . Extending this idea to an arbitrary finite measure  $P$  (not just the discrete measures) is the key idea behind exponential races. Notice that  $P$  in our example is atomic, which is fine, because the product measure  $\lambda \times P$  is not atomic. On the other hand, we want the points arriving in  $\mathbb{R}^n$  to correspond to the probability distribution  $P(\cdot)/P(\mathbb{R}^n)$ , so we will require that  $P$  is finite,  $P(\mathbb{R}^n) < \infty$ , and nonzero,  $0 < P(\mathbb{R}^n)$ . Also, in contrast to Poisson processes, exponential races have a natural ordering in time.

**Definition 7.4** (Exponential race). *Let  $P$  be a finite nonzero measure on  $\mathbb{R}^n$ . A random countable subset  $R \subseteq \mathbb{R}^+ \times \mathbb{R}^n$  is an exponential race with measure  $P$  if the following hold*

1.  $R$  is a Poisson process with mean measure  $\lambda \times P$ .
2.  $R$  is totally ordered by time, the first coordinate.

If  $R = \{(T_i, X_i)\}_{i=1}^\infty$ , then we assume the enumeration corresponds to the ordering so that  $i < j$  implies  $T_i < T_j$ .

We can realize an exponential race with a slight modification of Algorithm 7.1; use the partition of rectangles  $B_i = (i-1, i] \times \mathbb{R}^n$ , and sort points by their time variable. This is not the most direct characterization, so instead we derive the joint distribution of the first  $m$  ordered points in Theorem 7.7. The distribution of the countably infinite set  $R$  is completely described by

the joint distribution of the first  $m$  points for all finite  $m$ . The proof of Theorem 7.7 shows that the locations  $X_i$  are independently distributed as  $P(\cdot)/P(\mathbb{R}^n)$  and the interarrival times  $T_i - T_{i-1}$  are independent and exponentially distributed with rate  $P(\mathbb{R}^n)$ . This theorem is the cornerstone of this chapter, because it suggests a strategy for proving the correctness of Monte Carlo methods; if we can prove that the output of an algorithm  $(T, X)$  is the first arrival of an exponential race with measure  $P$ , then Theorem 7.7 guarantees that the location  $X$  is a sample from  $P(\cdot)/P(\mathbb{R}^n)$ .

**Theorem 7.7.** *Let  $P$  be a finite nonzero measure on  $\mathbb{R}^n$ ,  $X_i \sim P(\cdot)/P(\mathbb{R}^n)$  independent, and  $E_i \sim \text{Exp}(P(\mathbb{R}^n))$  independent, then first  $m$  points  $\{(T_i, X_i)\}_{i=1}^m$  of any exponential race  $R \subseteq \mathbb{R}^+ \times \mathbb{R}^n$  with measure  $P$  have the same joint distribution as*

$$\{(\sum_{j=1}^i E_j, X_i)\}_{i=1}^m.$$

*Proof.* Let  $T(t, B)$  be the time of the first arrival in  $B$  after time  $t \geq 0$ ,

$$T(t, B) = \min\{T_i : (T_i, X_i) \in R \cap (t, \infty) \times B\}. \quad (7.6)$$

$R \cap ((t, s+t] \times B)$  is finite with probability one for all  $s > 0$ , so (7.6) is well defined.  $T(t, B) - t$  is an exponential random variable, because

$$\mathbb{P}(T(t, B) - t > s) = \mathbb{P}(N((t, s+t] \times B) = 0) = \exp(-P(B)s).$$

$T(t, B)$  and  $T(t, B^c)$  are independent, by Poisson process independence.

We proceed by induction. The event  $\{T_1 > s, X_1 \in B\}$  is equivalent to  $\{T(0, B^c) > T(0, B) > s\}$ .  $P(B) > 0$  or  $P(B^c) > 0$ , so by Lemma 7.6,

$$\mathbb{P}(T_1 > s, X_1 \in B) = \mathbb{P}(T(0, B^c) > T(0, B) > s) = \exp(-sP(\mathbb{R}^n)) \frac{P(B)}{P(\mathbb{R}^n)}.$$

Now, assume Theorem 7.7 holds for  $k$ . The event

$$\{T_i = t_i, X_i = x_i\}_{i=1}^k$$

is completely described by counts in  $(0, t_k] \times \mathbb{R}^n$  and thus independent of

$$\{T(t_k, B^c) > T(t_k, B) > s + t_k\}$$



Thus

$$\begin{aligned}
& \mathbb{P}(T_{k+1} - T_k > s, X_{k+1} \in B | \{T_i = t_i, X_i = x_i\}_{i=1}^k) \\
&= \mathbb{P}(T(t_k, B^c) > T(t_k, B) > s + t_k | \{T_i = t_i, X_i = x_i\}_{i=1}^k) \\
&= \mathbb{P}(T(t_k, B^c) > T(t_k, B) > s + t_k) \\
&= \exp(-sP(\mathbb{R}^n)) \frac{P(B)}{P(\mathbb{R}^n)}
\end{aligned}$$

concludes the proof.  $\square$

### 7.3.2 Simulating an exponential race with a tractable measure

If  $Q$  is a tractable finite nonzero measure on  $\mathbb{R}^n$ , that is we have a procedure for computing  $Q(\mathbb{R}^n)$  and sampling from  $Q(\cdot)/Q(\mathbb{R}^n)$ , then Theorem 7.7 suggests Algorithm 7.2 for simulating an exponential race  $R$  with measure  $Q$ . Algorithm 7.2 simulates the points of an exponential race in order of arrival time. It does not terminate, but we can think of it as a coroutine or generator, which maintains state and returns the next arrival in  $\mathbb{R}$  each time it is invoked. As a simple example consider the uniform measure  $Q((a, b]) = b - a$  on  $[0, 1]$ . Algorithm 7.2 for this  $Q$  simulates a sequence of arrivals  $\{(T_i, X_i)\}_{i=1}^\infty$  with arrival location  $X_i \sim \text{Uniform}[0, 1]$  and interarrival time  $T_{i+1} - T_i \sim \text{Exp}(1)$ , see the left hand plot of Figure 7.4.

As with the initial discrete example, in which we constructed an exponential race from  $m$  independent Poisson processes, this is not the only approach. More generally, if  $\{B_i\}_{i=1}^m$  is any finite partition of  $\mathbb{R}^n$  such that  $Q(\cdot \cap B_i)$  is tractable, then we can simulate  $R$  by simulating  $m$  independent exponential races  $R_i$  with measure  $Q(\cdot \cap B_i)/Q(B_i)$  via Algorithm 7.2 and sorting the result  $\cup_{i=1}^m R_i$ . This can be accomplished lazily and efficiently with a priority queue data type, which prioritizes the races  $R_i$  according to which arrives next in time. It also possible to split the races  $R_i$  online by partitioning  $B_i$  and respecting the constraint imposed by the arrivals already generated in  $B_i$ . We highlight a particularly important variant, which

---

**Algorithm 7.2** An exponential race  $R$  with finite nonzero measure  $Q$


---

```

 $R = \emptyset$ 
 $T_0 = 0$ 
for  $i = 1$  to  $\infty$  do
   $E_i \sim \text{Exp}(Q(\mathbb{R}^n))$ 
   $X_i \sim Q(\cdot)/Q(\mathbb{R}^n)$ 
   $T_i = T_{i-1} + E_i$ 
   $R = R \cup \{(T_i, X_i)\}$ 
end for

```

---



**Figure 7.4:** Two methods for simulating an exponential race. The left hand plot shows the first arrivals of a uniform exponential race on  $[0, 1]$  simulated by Algorithm 7.2. The right hand plot shows the first arrivals of an exponential race simulated over a space partitioning tree. Dashed lines dominate the set in which an arrival is first.

features in  $A^*$  sampling in Section 7.5. Consider an infinitely deep tree in which each node is associated with a subset  $B \subseteq \mathbb{R}^n$ . If the root is  $\mathbb{R}^n$  and the children of each node form a partition of the parent, then we call this a space partitioning tree. We can realize an exponential race over a space partitioning tree by recursively generating arrivals  $(T, X)$  at each node  $B$ . Each location  $X$  is sampled independently from  $Q(\cdot \cap B)/Q(B)$ , and each time  $T$  is sampled by adding an independent  $\text{Exp}(Q(B))$  to the parent's arrival time. The arrivals sorted by time over the realization of the tree form an exponential race. See Figure 7.4.

### 7.3.3 Transforming an exponential race with accept-reject and perturb

Most finite nonzero measures  $P$  on  $\mathbb{R}^n$  are not tractable. Monte Carlo methods accomplish their goal of sampling from intractable distributions by transforming samples of tractable distributions. In this subsection we present accept-reject and perturb operations, which transform a realization of an exponential race with measure  $Q$  into a realization of an exponential race with a distinct measure  $P$ . In practice  $Q$  will be tractable and  $P$  intractable, so that simulating an exponential race with an intractable measure can be accomplished by simulating the points of an exponential race with a tractable measure, for example via Algorithm 7.2, and transforming it with accept-reject or perturb operations. The accept-reject and perturb operations are named after their respective literatures, accept-reject corresponds to rejection sampling and perturb corresponds to the Gumbel-Max trick. The correspondence between the perturb operation and the Gumbel-Max trick may not be obvious, so we discuss this in Section 7.4.

Let  $Q$  and  $P$  be finite nonzero measures in  $\mathbb{R}^n$ . We assume that they have

densities  $g$  and  $f$  with respect to some base measure  $\mu$ ,

$$Q(B) = \int_B g(x)\mu(dx) \quad P(B) = \int_B f(x)\mu(dx). \quad (7.7)$$

We assume that  $g$  and  $f$  have the same support and their ratio is bounded,

$$\text{supp}(f) = \text{supp}(g) \quad \frac{f(x)}{g(x)} \leq M \text{ for all } x \in \text{supp}(g) \quad (7.8)$$

where  $\text{supp}(g) = \{x \in \mathbb{R}^n : g(x) \neq 0\}$ . The assumption  $\text{supp}(f) = \text{supp}(g)$  can be softened here and throughout the chapter to  $\text{supp}(f) \subseteq \text{supp}(g)$ , but it complicates the analysis. The accept-reject strategy is to realize more points than needed from an exponential race with measure  $MQ(\cdot)$  and stochastically *reject* points with probability equal to the ratio of instantaneous rates of arrival,  $f(x)/(g(x)M)$ . The perturbation strategy is to realize just the points needed from an exponential race with measure  $Q$ , but to *perturb* the arrival times according to the transformation  $t \rightarrow tg(x)/f(x)$  for all points arriving at  $x$ .

Before we present the proofs, consider the following intuition. Imagine taking a long exposure photograph of the plane as instantaneous flashes arrive according to an exponential race with measure  $Q$ . The rate at which points arrive will determine the intensity of a heat map with regions receiving more points brighter than those receiving fewer. Over time the relative intensities will correspond to the probability distribution proportional to  $Q$ . If someone were just ahead of us in time and stochastically discarded points that arrived in  $B$  or delayed points in  $B$  relative to points in  $B^c$ , then our perception of the likelihood of  $B$  would change. Mired in time, we would not be able to distinguish whether points were discarded, reordered, or the true measure  $Q$  was in fact different.

The correctness of these operations on an exponential race can be justified as special cases of the Thinning and Mapping Theorems.

**Lemma 7.8** (Accept-Reject). *Let  $Q$  and  $P$  be finite nonzero measures on  $\mathbb{R}^n$  under assumptions (7.7) and (7.8). If  $R \subseteq \mathbb{R}^+ \times \mathbb{R}^n$  is an exponential race with measure  $MQ(\cdot)$  and  $\text{accept}(t, x) \sim \text{Bernoulli}(\rho(t, x))$  is i.i.d. for all  $(t, x)$  with probability*

$$\rho(t, x) = \frac{f(x)}{g(x)M},$$

*then  $\text{thin}(R, \text{accept})$ , from (7.3), is an exponential race with measure  $P$ .*

*Proof.* By the Thinning Theorem, the mean measure of  $\text{thin}(R, \text{accept})$  is

$$\iint_B \frac{f(x)}{g(x)M} g(x)M \mu(dx) \lambda(dt) = \iint_B f(x) \mu(dx) \lambda(dt) = (\lambda \times P)(B).$$

for  $B \subseteq \mathbb{R}^+ \times \text{supp}(g)$ . The subsampled  $(T_i, X_i)$  are in order and thus an exponential race with measure  $P$ .  $\square$

**Lemma 7.9** (Perturbation). *Let  $Q$  and  $P$  be finite nonzero measures on  $\mathbb{R}^n$  under assumptions (7.7) and (7.8). If  $R \subseteq \mathbb{R}^+ \times \mathbb{R}^n$  is an exponential race with measure  $Q$  and*

$$\text{perturb}(t, x) = \left( t \frac{g(x)}{f(x)}, x \right),$$

*then  $\text{sort}(\text{perturb}(R))$  is an exponential race with measure  $P$  where  $\text{sort}$  totally orders points by the first coordinate, time.*

*Proof.*  $\text{perturb}$  is 1-1 on  $\text{supp}(f)$ , so the Mapping Theorem applies. It is enough to check the mean measure of  $\text{perturb}(R)$  on subsets of the form  $B = (0, s] \times A$  for  $s \in \mathbb{R}^+$  and  $A \subseteq \text{supp}(g)$ ,

$$\iint_{h^{-1}(B)} g(x) \lambda(dt) \mu(dx) = \int_A g(x) s \frac{f(x)}{g(x)} \mu(dx) = (\lambda \times P)(B).$$

Thus, sorting  $\text{perturb}(T_i, X_i)$  forms an exponential race with measure  $P$ .  $\square$

## 7.4 Gumbel processes

### 7.4.1 Definition and construction

The central object of the Gumbel-Max trick is a random function over a finite set whose values are Gumbel distributed. Gumbel valued functions over a finite choice set are extensively studied in random choice theory, where there is a need for a statistical model of utility (Yellott, 1977 for example). The extension to Gumbel valued functions over continuous spaces has been explored in random choice theory (Malmberg, 2013) and in the context of Monte Carlo simulation (Maddison et al., 2014). Following Maddison et al. (2014) we will refer to this class of Gumbel valued functions on  $\mathbb{R}^n$  as Gumbel processes. Gumbel processes underpin the recent interest in perturbation based Monte Carlo methods, because their maxima are located at samples from probability distributions, see also (Papandreou and Yuille, 2011; Tarlow et al., 2012; Hazan et al., 2013; Chen and Ghahramani, 2015; Kim et al.,

2016). In this section we clarify the connection between Gumbel processes and our development of exponential races. We will show that the value of a Gumbel process at  $x \in \mathbb{R}^n$  can be seen as the log transformed time of the first arrival at  $x$  of some exponential race. This has the advantage of simplifying their construction and connecting the literature on the Gumbel-Max trick to our discussion. Related constructions have also been considered in the study of extremal processes (Resnick, 2007). In this subsection we define and construct Gumbel processes. In the next subsection we discuss their simulation and present a generalized Gumbel-Max trick derived from the Perturbation Lemma.

The Gumbel distribution dates back to the statistical study of extrema and rare events (Gumbel and Lieblein, 1954). The Gumbel is a member of a more general class of extreme value distributions. A central limit theorem exists for these distributions — after proper renormalization the maximum of an i.i.d. sample of random variables converges to one of three possible extreme value distributions (Gedenko, 1948). The Gumbel is parameterized by a location  $\mu \in \mathbb{R}$ .

**Definition 7.5** (Gumbel random variable).  *$G$  is a Gumbel distributed random variable on  $\mathbb{R}$  with location  $\mu \in \mathbb{R}$  if*

$$\mathbb{P}(G \leq g) = \exp(-\exp(-g + \mu))$$

*This is denoted  $G \sim \text{Gumbel}(\mu)$  and  $G \sim \text{Gumbel}(-\infty)$  is the random variable whose value is  $-\infty$  with probability one. If  $G \sim \text{Gumbel}(0)$ , then  $G + \mu \sim \text{Gumbel}(\mu)$ .*

The Gumbel distribution has two important properties for our purposes. The distribution of the maximum of independent Gumbels is itself a Gumbel — a property known as max-stability — and the index of the maximum follows the Gibbs distribution: if  $G(i) \sim \text{Gumbel}(\mu_i)$ , then

$$\max_{1 \leq i \leq m} G(i) \sim \text{Gumbel}(\log \sum_{i=1}^m \exp(\mu_i)) \quad \operatorname{argmax}_{1 \leq i \leq m} G(i) \sim \frac{\exp(\mu_i)}{\sum_{i=1}^m \exp(\mu_i)}.$$

The Gumbel-Max trick of the introduction for sampling from a discrete distribution with mass function  $f : \{1, \dots, m\} \rightarrow \mathbb{R}^+$  is explained by taking  $\mu_i = \log f(i)$ . It is informative to understand these properties through the Gumbel's connection to the exponential distribution.

**Lemma 7.10.** *If  $E \sim \text{Exp}(\lambda)$  with nonnegative rate  $\lambda \in \mathbb{R}^{\geq 0}$ , then*

$$-\log E \sim \text{Gumbel}(\log \lambda).$$

*Proof.*  $\mathbb{P}(-\log E \leq g) = \mathbb{P}(E \geq \exp(-g)) = \exp(-\exp(-g + \log \lambda))$   $\square$

Therefore the distribution of the maximum and argmaximum of Gumbels is explained by Lemma 7.6, because passing a maximization through  $-\log$  becomes a minimization.

A Gumbel process  $G : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{-\infty\}$  is a Gumbel valued random function. Their characterizing property is that the maximal values of a Gumbel process over the subsets  $B \subseteq \mathbb{R}^n$  are marginally Gumbel distributed with a location that scales logarithmically with the volume of  $B$  according to some finite nonzero measure  $P$ ,

$$\max_{x \in B} G(x) \sim \text{Gumbel}(\log P(B))$$

Implicit in this claim is the assertion that the maximizations  $\max_{x \in B} G(x)$  are well-defined — the maximum exists — for all  $B \subseteq \mathbb{R}^n$ .

**Definition 7.6** (Gumbel process). *Let  $P$  be a finite nonzero measure on  $\mathbb{R}^n$ ,  $G : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{-\infty\}$  a random function, and*

$$G^*(B) = \max_{x \in B} G(x). \tag{7.9}$$

*$G$  is a Gumbel process with measure  $P$  if*


1. *For  $B \subseteq \mathbb{R}^n$ ,  $G^*(B) \sim \text{Gumbel}(\log P(B))$ .*
2. *For  $A_1, \dots, A_m$  are disjoint,  $G^*(A_i)$  are independent.*

Note, the event that  $\text{argmax}_{x \in \mathbb{R}^n} G(x)$  lands in  $B \subseteq \mathbb{R}^n$  depends on which of  $G^*(B)$  or  $G^*(B^c)$  is larger. Following this reasoning one can show that the  $\text{argmax}$  over  $\mathbb{R}^n$  is distributed as  $P(\cdot)/P(\mathbb{R}^n)$ .

The study of Gumbel processes can proceed without reference to exponential races, as in Maddison et al. (2014), but our construction from exponential races is a convenient shortcut that allows us to import results from Section 7.3. Consider the function that reports the arrival time of the first arrival at  $x \in \mathbb{R}^n$  for an exponential race  $R$  with measure  $P$ ,

$$T(x) = \min\{T_i : (T_i, x) \in R\}$$

This function is almost surely infinite at all  $x$ , but for any realization of  $R$  it will take on finite value at countably many points in  $\mathbb{R}^n$ . Moreover, the minimum of  $T(x)$  over subsets  $B \subseteq \mathbb{R}^n$  is well-defined and finite for sets with positive measure  $P(B) > 0$ ; it is exponentially distributed with rate  $P(B)$ . In this way we can see that  $-\log T(x)$  is Gumbel process, Figure 7.5.



**Figure 7.5:** Constructing a uniform Gumbel process  $G : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{-\infty\}$  on  $[0, 1]$  with an exponential race. The left hand plot shows the first arrivals  $*$  of a uniform exponential race  $R$ . The right hand plot shows  $G(x)$  set to  $-\log$  the time  $T(x)$  of the first arrival at  $x$ . The graph of  $G(x)$  extends downwards to  $-\infty$  taking on finite value at all points in  $[0, 1]$  that have arrivals and  $-\infty$  for all points with no arrivals.

**Theorem 7.11.** *Let  $R \subseteq \mathbb{R}^+ \times \mathbb{R}^n$  be an exponential race with measure  $P$ .*

$$G(x) = -\log \min\{T_i : (T_i, x) \in R\} \tag{7.10}$$

*is a Gumbel process with measure  $P$ .*

*Proof.* First, for  $x \in \mathbb{R}^n$

$$\min\{T_i : (T_i, x) \in R\} = T(0, \{x\}),$$

where  $T(0, B)$  is the first arrival time in subset  $B \subseteq \mathbb{R}^n$  defined in (7.6) from Theorem 7.7. Thus  $G^*(B)$  of (7.9) is well defined, because

$$G^*(B) = \max_{x \in B} -\log \min\{T_i : (T_i, x) \in R\} = -\log T(0, B).$$

$G^*(B)$  inherits the independence properties from Poisson process independence. Finally, Lemma 7.10 gives us the marginal distribution of  $G^*(B)$ .  $\square$

#### 7.4.2 Simulating a Gumbel process and the Gumbel-Max trick

Gumbel processes are relevant to Monte Carlo simulation in the same sense that we motivated exponential races — if we can simulate the maximum value of a Gumbel process with measure  $P$ , then its location is a sample from the distribution  $P(\cdot)/P(\mathbb{R}^n)$ . Maddison et al. (2014) gave an algorithm for simulating Gumbel processes with tractable measures and a generalized Gumbel-Max trick for transforming their measure. We present those results

**Algorithm 7.3** A Gumbel process with finite measure  $Q$ 


---

```

Initialize  $G(x) = -\infty$  for all  $x \in \mathbb{R}^n$ .
 $(\Omega_1, G_0, i) = (\mathbb{R}^n, \infty, 1)$ 
while  $Q(\Omega_i) > 0$  do
   $G_i \sim \text{TruncGumbel}(\log Q(\Omega_i), G_{i-1})$ 
   $X_i \sim Q(\cdot \cap \Omega_i)/Q(\Omega_i)$ 
   $G(X_i) = G_i$  % assign  $G(x)$  at  $X_i$  to  $G_i$ 
   $\Omega_{i+1} = \Omega_i - \{X_i\}$ 
   $i = i + 1$ 
end while

```

---

derived from our results for exponential races.

The Gumbel process  $G$  from construction (7.10) has value  $-\infty$  everywhere except at the countably many arrival locations of an exponential race. Therefore, for tractable measures  $Q$  we could adapt Algorithm 7.2 for exponential races to simulate  $G(x)$ . The idea is to initialize  $G(x) = -\infty$  everywhere and iterate through the points  $(T_i, X_i)$  of an exponential race  $R$  setting  $G(X_i) = -\log T_i$ . To avoid reassigning values of  $G(x)$  we refine space as in Section 7.3.2 by removing the locations generated so far. Algorithm 7.3 implements this procedure, although it is superficially different from our description. In particular the value  $G(X_i)$  is instead set to a truncated Gumbel  $G_i \sim \text{TruncGumbel}(\log Q(\Omega_i), G_{i-1})$ , a Gumbel random variable with location  $\log Q(\Omega_i)$  whose domain is truncated to  $(-\infty, G_{i-1}]$ . The connection to Algorithm 7.2 can be derived by decomposing the arrival times  $T_i = \sum_{j=1}^i E_j$  for  $E_j \sim \text{Exp}(Q(\Omega_j))$  and then considering the joint distribution of  $G_i = -\log(\sum_{j=1}^i E_j)$ . A bit of algebraic manipulation will reveal that


$$G_i | G_{i-1} \sim \text{TruncGumbel}(\log Q(\Omega_i), G_{i-1})$$

Thus, translating between procedures for simulating Gumbel processes and procedures for simulating exponential races is as simple as replacing chains of truncated Gumbels with partial sums of exponentials.

For continuous measures removing countably many points from the sample space has no effect, and in practice the removal line of Algorithm 7.3 can be omitted. For those and many other measures Algorithm 7.3 will not terminate; instead it iterates through the infinitely many finite values of  $G(x)$  in order of their rank. For discrete measures with finite support Algorithm 7.3 will terminate once every atom has been assigned a value.

Finally, for simulating Gumbel processes with intractable measures  $P$  the Perturbation Lemma of exponential races justifies a generalized Gumbel-Max trick. The basic insight is that multiplication by the ratio of densities  $g(x)/f(x)$  becomes addition in log space.





**Figure 7.6:** A continuous Gumbel-Max trick. The left hand plot shows the maximal values of a uniform Gumbel process  $G(x)$  on  $[0, 1]$ . The right hand plot shows the result of perturbing  $\log f(x)$  with  $G(x)$ . Notice that the ordering of values changes, and  $X^*$  is now the location of the maximum  $G^* = \max_x \log f(x) + G(x)$ . Therefore,  $X^*$  is a sample from the distribution with density proportional to  $f(x)$ .

**Lemma 7.12** (Gumbel-Max trick). *Let  $Q$  and  $P$  be finite nonzero measures on  $\mathbb{R}^n$  with densities  $g$  and  $f$  under assumptions (7.7) and (7.8). If  $G : \mathbb{R}^n \rightarrow \mathbb{R} \cap \{-\infty\}$  is a Gumbel process with measure  $Q$ , then*

$$G'(x) = \begin{cases} \log f(x) - \log g(x) + G(x) & x \in \text{supp}(g) \\ -\infty & \text{otherwise} \end{cases}$$

*is a Gumbel process with measure  $P$ . In particular for  $G^* = \max_{x \in \mathbb{R}^n} G'(x)$  and  $X^* = \text{argmax}_{x \in \mathbb{R}^n} G'(x)$ ,*

$$G^* \sim \text{Gumbel}(\log P(\mathbb{R}^n)) \quad X^* \sim P(\cdot)/P(\mathbb{R}^n)$$

*Proof.* Arguing informally, this follows from the Perturbation Lemma applied to our construction (7.10) of Gumbel processes. For  $x \in \text{supp}(g)$

$$\log f(x) - \log g(x) + G(x) = -\log \min\{T_i g(x)/f(x) : (T_i, x) \in R\}.$$

See Maddison et al. (2014) for a formal proof. □

When  $Q$  is the counting measure on  $\{1, \dots, m\}$ , Lemma 7.12 exactly describes the Gumbel-Max trick of the introduction. This brings full circle the connection between accept-reject and the Gumbel-Max trick.

A Gumbel process is not profoundly different from an exponential race, but the difference of perspective — a function as opposed to a random set — can be valuable. In particular consider the following generalization of a result from Hazan and Jaakkola of this book. Let  $G : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{-\infty\}$  be a Gumbel process with measure  $P$  whose density with respect to  $\mu$  is  $f$ . If

$G^* = \max_{x \in \mathbb{R}^n} G(x)$  and  $X^* = \operatorname{argmax}_{x \in \mathbb{R}^n} G(x)$ , then

$$\mathbb{E}(G^*) = \log P(\mathbb{R}^n) + \gamma \quad \mathbb{E}(-\log f(X^*) + G^*) = H(f) + \gamma,$$

where  $H(f)$  is the entropy of a probability distribution with probability density function proportional to  $f$  and  $\gamma$  is the Euler-Mascheroni constant. Therefore the representation of probability distributions through Gumbel processes gives rise to a satisfying and compact representation of some of their important constants.


## 7.5 Monte Carlo methods that use bounds

### 7.5.1 Rejection sampling

In this section we present practical Monte Carlo methods that use bounds on the ratio of densities to produce samples from intractable distributions. We show how these methods can be interpreted as algorithms that simulate the first arrival of an exponential race. The basic strategy for proving their correctness is to argue that they perform accept-reject or perturb operations on the realization of an exponential race until they have provably produced the first arrival of the transformed race. We start by discussing the traditional rejection sampling and a related perturbation based method. Then we study OS\* (Dymetman et al., 2012), an accept-reject method, and A\* sampling (Maddison et al., 2014), a perturbation method. These algorithms have all been introduced elsewhere in the literature, so for more information we refer readers to the original papers.

Throughout this section our goal is to draw a sample from the probability distribution proportional to some measure  $P$  with density  $f$  with respect to some base measure  $\mu$ . We assume, as in the Accept-Reject and Perturbation Lemmas, access to a tractable proposal distribution proportional to a measure  $Q$  with density  $g$  with respect to  $\mu$  such that  $f$  and  $g$  have the same support and the ratio  $f(x)/g(x)$  is bounded by some constant  $M$ . For example consider the sample space  $\{0, 1\}^n$  whose elements are bit vectors of length  $n$ . A proposal distribution might be proportional to the counting measure  $Q$ , which counts the number of configurations in a subset  $B \subseteq \{0, 1\}^n$ . Sampling from  $Q(\cdot)/Q(\{0, 1\}^n)$  is as simple as sampling  $n$  independent Bernoulli(1/2).

Rejection sampling is the classic Monte Carlo method that uses bound information. It proposes  $(X, U)$  from  $Q$  and Uniform[0, 1], respectively, and accepts  $X$  if  $U \leq f(X)/(g(X)M)$ . The algorithm terminates at the first acceptance and is normally justified by noticing that it samples uniformly from the region under the graph of  $f(x)$  by rejecting points that fall between



**Figure 7.7:** Algorithms **REJ** and **PER** for measure  $P$  on  $[0, 1]$  with proposal measure  $Q$ . The densities of  $Q$  and  $P$  are shown on the left hand side as densities over  $x \in [0, 1]$ .  $\circ$  are arrivals of the race with measure  $Q$ ,  $*$  of the race with measure  $P$ . Both plots show the proposals considered until the first acceptance. For **PER** opaque solid lines represent the perturb operation.  $T_4$  is the fourth arrival from the race with measure  $Q$ .  $T_4/M$  is the lower bound on all future arrivals, and thus all  $*$  points to the left of  $T_4/M$  are in order.

$g(x)M$  and  $f(x)$ , see the left hand graph on Figure 7.7 for an intuition. The acceptance decision also corresponds exactly to the accept-reject operation on exponential races, so we can interpret it as an procedure on the points of an exponential race. We call this procedure **REJ** for short,

```

for  $(T_i, X_i) \in R$  simulated by Algorithm 7.2 with measure  $MQ(\cdot)$  do
   $U_i \sim \text{Uniform}[0, 1]$ .
  if  $U_i < f(X_i)/(g(X_i)M)$  then return  $(T_i, X_i)$ 
  end if
end for

```

The Accept-Reject Lemma guarantees that the returned values  $(T, X)$  will be the first arrival of an exponential race with measure  $P$ , and Theorem 7.7 guarantees that  $X$  is a sample from  $P(\cdot)/P(\mathbb{R}^n)$ . This is the basic flavour of the arguments of this section.

The Perturbation Lemma has a corresponding procedure, which uses the bound  $M$  to provably return the first arrival of a perturbed exponential race. It is shown on the right hand side of Figure 7.7, and we call it **PER**.

```

 $(T^*, X^*) = (\infty, \text{null})$ 
for  $(T_i, X_i) \in R$  simulated by Algorithm 7.2 with measure  $Q$  do
  if  $T^* > T_i g(X_i)/f(X_i)$  then
     $T^* = T_i g(X_i)/f(X_i)$ 
     $X^* = X_i$ 
  end if
  if  $T_{i+1}/M \geq T^*$  then return  $(T^*, X^*)$ 

```

**end if**  
**end for**

In this procedure  $(T_i, X_i)$  iterates in order through the arrivals of an exponential race with measure  $Q$ . The perturbed times  $T_i g(X_i)/f(X_i)$  will form a race with measure  $P$ , but not necessarily in order.  $(T^*, X^*)$  are variables that track the earliest perturbed arrival so far, so  $T^*$  is an upper bound on the eventual first arrival time for the race with measure  $P$ .  $T_{i+1}$  is the arrival time of the next point in the race with measure  $Q$  and  $M$  bounds the contribution of the perturbation, so  $T_{i+1}/M$  is a lower bound on the remaining perturbed arrivals. When  $T^*$  and  $T_{i+1}/M$  cross,  $(T^*, X^*)$  is guaranteed to be the first arrival of the perturbed race.

**REJ** and **PER** can be turned into generators for iterating through all of the arrivals of an exponential race with measure  $P$  as opposed to just returning the first. For **REJ** it is as simple as replacing **return** with **yield**, so that each time the generator is invoked it searches until the next acceptance and returns. For **PER** we must store every perturbed arrival until its eventual order in the race with measure  $P$  is determined. This can be accomplished with a priority queue  $\mathcal{U}$ , which prioritizes by earliest arrival time,

$\mathcal{U} = \text{minPriorityQueue}()$

**for**  $(T_i, X_i) \in R$  simulated by Algorithm 7.2 with measure  $Q$  **do**  
 $\mathcal{U}.\text{pushWithPriority}(T_i g(X_i)/f(X_i), X_i)$   
**if**  $T_{i+1}/M \geq \min \mathcal{U}$  **then yield**  $\mathcal{U}.\text{pop}()$   
**end if**  
**end for**

$\mathcal{U}$  takes the place of  $T^*$  and  $X^*$  in **PER**. The highest priority arrival on  $\mathcal{U}$  will be the earliest of the unordered perturbed arrivals and  $T_{i+1}/M$  is a lower bound on all future perturbed arrivals. When  $T_{i+1}/M \geq \min \mathcal{U}$ , the earliest arrival on  $\mathcal{U}$  is guaranteed to be the next arrival. It is informative to think of the generator version of **PER** via Figure 7.7. The lower bound  $T_{i+1}/M$  is a bound across space that advances rightward in time, every arrival to the left of  $T_{i+1}/M$  is in order and every arrival to the right is unordered.

Consider the number of iterations until the first acceptance in **REJ** and **PER**. At first it seems that both algorithms should have different runtimes. **REJ** is obviously memoryless, and it seems wasteful — no information accumulates. On the other hand **PER** accumulates the earliest arrival and its termination condition depends on a history of arrivals. Unfortunately, both algorithms have the same geometric distribution over the number of arrivals considered. Arguing informally, the lower bound  $T_{i+1}/M$  of **PER** plotted over the iterations will form a line with slope  $(MQ(\mathbb{R}^n))^{-1}$ . **PER** terminates when this line crosses the first arrival time of the perturbed race.

The first arrival of a race with measure  $P$  occurs at  $P(\mathbb{R}^n)^{-1}$  in expectation, so we expect the crossing point to occur on average at  $MQ(\mathbb{R}^n)/P(\mathbb{R}^n)$  iterations. This is the same as the expected runtime of **REJ**.

**Lemma 7.13.** *Let  $K(\mathbf{REJ})$  and  $K(\mathbf{PER})$  be the number of proposals considered by the rejection and perturbation sampling algorithms. Then*

$$\mathbb{P}(K(\mathbf{REJ}) > k) = \mathbb{P}(K(\mathbf{PER}) > k) = (1 - \rho)^k \text{ with } \rho = \frac{P(\mathbb{R}^n)}{Q(\mathbb{R}^n)M}.$$

Thus  $K(\mathbf{REJ})$  and  $K(\mathbf{PER})$  are geometric random variable with

$$\mathbb{E}(K(\mathbf{REJ})) = \mathbb{E}(K(\mathbf{PER})) = \frac{1}{\rho}$$

*Proof.* The probability of accepting a proposal at any iteration of **REJ** is

$$\mathbb{E}(f(X_i)/(g(X_i)M)) = \int \frac{f(x)}{g(x)M} \frac{g(x)}{Q(\mathbb{R}^n)} \mu(dx) = \rho.$$

Each decision is independent, so the probability of  $k$  rejections is  $(1 - \rho)^k$ .

**PER** exceeds  $k$  iterations if  $T_i g(X_i)/f(X_i) > T_{k+1}/M$  for all  $i \leq k$ . Because the  $X_i$  are i.i.d.,

$$\mathbb{P}(K(\mathbf{PER}) > k \mid \{T_i = t_i\}_{i=1}^{k+1}) = \prod_{i=1}^k \mathbb{P}(t_i/t_{k+1} > f(X)/(g(X)M)),$$

where  $X \sim Q(\cdot)/Q(\mathbb{R}^n)$ . Given  $T_{k+1} = t_{k+1}$  the  $T_i$  for  $i \leq k$  are i.i.d.  $T_i \sim \text{Uniform}(0, t_{k+1})$  by Lemma 7.3. Thus  $T_i/T_{k+1} \sim \text{Uniform}(0, 1)$  i.i.d.

$$\mathbb{P}(K(\mathbf{PER}) > k) = \prod_{i=1}^k \mathbb{P}(U > f(X)/(g(X)M)) = (1 - \rho)^k$$

finishes the proof. □

### 7.5.2 Adaptive bounds

Lemma 7.13 is disappointing, because it suggests that reasoning about perturbations is as inefficient as discarding proposals. The problem is fundamentally that information carried in the bound  $M$  about the discrepancy between  $g(x)$  and  $f(x)$  is static throughout the execution of both algorithms. Considering a contrived scenario will illustrate this point. Suppose that for every failed proposal we are given a tighter bound  $M_{i+1} < M_i$  from some oracle. Both **REJ** and **PER** can be adapted to take advantage of these adaptive bounds simply by dropping in  $M_i$  wherever  $M$  appears.

In this case **PER** is distinguished from **REJ**. **REJ** makes an irrevocable

decision at each iteration. In contrast **PER** simply pushes up the lower bound  $T_{i+1}/M_i$  without erasing its memory, bringing it closer to accepting the earliest arrival so far. Indeed, the probability of this oracle rejection sampling exceeding  $k$  proposals is

$$\mathbb{P}(K(\mathbf{OREJ}) > k) = \prod_{i=1}^k (1 - \rho_i) \text{ where } \rho_i = P(\mathbb{R}^n)/(Q(\mathbb{R}^n)M_i).$$

On the other hand, the probability of this oracle perturbation sampling exceeding  $k$  proposals is

$$\mathbb{P}(K(\mathbf{OPER}) > k) = \prod_{i=1}^k \mathbb{P}(U > f(X)/(g(X)M_k)) = (1 - \rho_k)^k,$$

or the probability of rejecting  $k$  proposals *as if* the  $M_k$ th bound was known all along. By tracking the earliest arrival so far **OPER** makes efficient use of adaptive bound information, reevaluating all points in constant time.

### 7.5.3 OS\* adaptive rejection sampling and A\* sampling

The difference between **REJ** and **PER** exposed by considering adaptive bounds motivates studying OS\* and A\* sampling, Monte Carlo methods that use realistic adaptive bounds. Both methods iteratively refine a partition  $\{B_i\}_{i=1}^m$  of  $\mathbb{R}^n$ , which allows them to use regional bounds  $M(B_i)$ , where  $f(x)/g(x) \leq M(B_i)$  for  $x \in B_i$ . As with **REJ** and **PER**, OS\* and A\* sampling are only distinguished by how they use this information. OS\* reasons about accept-reject operations, A\* sampling about perturb operations. In contrast to the relationship between **REJ** and **PER**, A\* sampling makes more efficient use of proposal samples than OS\*.

OS\* and A\* sampling must compute volumes and samples of subsets under the proposal measure  $Q$ . It will be possibly intractable to consider any possible  $B_i \subseteq \mathbb{R}^n$ , so a user must implicitly specify a nice family  $\mathcal{F}$  of subsets that is closed under a user-specified refinement function  $\text{split}(B, x)$ . Hyperrectangles are a simple example. All together, the user must provide,

1. finite nonzero measure  $P$  with a method for computing the density  $f(x)$ .
2. finite nonzero proposal measure  $Q$  with methods for sampling restricted to  $B \in \mathcal{F}$ , computing measures of  $B \in \mathcal{F}$ , and computing the density  $g(x)$ .
3. partitioning set function  $\text{split}(B, x) \subseteq \mathcal{F}$  for  $B \in \mathcal{F}$  that partitions  $B$ .
4. bounding set function  $M(B)$  for  $B \in \mathcal{F}$ ,  $f(x)/g(x) \leq M(B)$  for  $x \in B$ .

Specific examples, which correspond to experimental examples, are given in

**Algorithm 7.4 OS\*** adaptive rejection sampling for  $P$  with proposal  $Q$ 


---

```

 $\mathcal{P}_0 = \{\mathbb{R}^n\}$ 
 $T_0 = 0$ 
for  $i = 1$  to  $\infty$  do
   $B_i \sim \mathbb{P}(B) \propto Q(B)M(B)$  for  $B \in \mathcal{P}_{i-1}$ 
   $X_i \sim Q(\cdot \cap B_i)/Q(B_i)$ 
   $E \sim \text{Exp}(\sum_{B \in \mathcal{P}_{i-1}} M(B)Q(B))$ 
   $T_i = T_{i-1} + E$ 
   $U_i \sim \text{Uniform}[0, 1]$ 
  if  $U_i < f(X_i)/(g(X_i)M(B_i))$  then
    return  $(T_i, X_i)$ 
  else
     $\mathcal{C} = \text{split}(B_i, X_i)$ 
     $\mathcal{P}_i = \mathcal{P}_{i-1} - \{B_i\} + \mathcal{C}$ 
  end if
end for

```

---

the Appendix.

OS\* (**OS\*** for short) is in a family of adaptive rejection sampling algorithms, which use the history of rejected proposals to tighten the gap between the proposal density and the density of interest. The name adaptive rejection sampling (ARS) is normally reserved for a variant that assumes  $\log f(x)$  is concave (Gilks and Wild, 1992). Accept-reject decisions are independent, so any adaptive scheme is valid as long as the rejection rate is not growing too quickly (Casella et al., 2004). Our proof of the correctness appeals to exponential races, and it works for a wider range of adaptive schemes than just **OS\***.

In more detail, **OS\*** begins with the proposal density  $g(x)$  and a partition  $\mathcal{P}_0 = \{\mathbb{R}^n\}$ . At every iteration it samples from the distribution with density proportional to  $\sum_{B \in \mathcal{P}_{i-1}} g(x)M(B)1_B(x)$  in a two step procedure, sampling a subset  $B \in \mathcal{P}_{i-1}$  with probability proportional to  $Q(B)M(B)$ , and then sampling a proposal point  $X$  from the distribution with density  $g(x)$  restricted to  $B$ . If  $X$  is rejected under the current proposal, then  $\mathcal{P}_{i-1}$  is refined by splitting  $B$  with the user specified  $\text{split}(B, X)$ . There is a choice of when to refine and which subset  $B \in \mathcal{P}_{i-1}$  to refine, but for simplicity we consider just the form the splits the subset of the current proposal. **OS\*** continues until the first acceptance, see Algorithm 7.4.

**Theorem 7.14** (Correctness of OS\*). *Let  $K(\mathbf{OS}^*)$  be the number of proposal samples considered before termination. Then*

$$\mathbb{P}(K(\mathbf{OS}^*) > k) \leq (1 - \rho)^k \text{ where } \rho = \frac{P(\mathbb{R}^n)}{Q(\mathbb{R}^n)M(\mathbb{R}^n)}$$

and upon termination the return values  $(T, X)$  of  $OS^*$  are independent and

$$T \sim \text{Exp}(P(\mathbb{R}^n)) \quad X \sim \frac{P(\cdot)}{P(\mathbb{R}^n)}.$$

*Proof.* The situation is complicated, because the proposals  $\{(T_i, X_i)\}_{i=1}^\infty$  of  $OS^*$  are not an exponential race. Instead, we present an informal argument derived from a more general thinning theorem, Proposition 14.7.I. in Daley and Vere-Jones (2007). Let  $g_i(x)$  be the proposal density at iteration  $i$ ,

$$g_i(x) = \sum_{B \in \mathcal{P}_{i-1}} g(x)M(B)1_B(x).$$

Clearly,  $g_i(x)$  depends on the history of proposals so far and  $f(x) \leq g_i(x) \leq g(x)M(\mathbb{R}^n)$  for all  $i$ . Let  $R$  be an exponential race with measure  $M(\mathbb{R}^n)Q(\cdot)$  and  $U_j \text{Uniform}[0, 1]$  i.i.d. for each  $(T_j, X_j) \in R$ . Consider the following adaptive thinning procedure, subsample all points of  $R$  that satisfy  $U_j \leq g_i(X_j)/(g(X_j)M(\mathbb{R}^n))$  where  $g_i(X_j)$  is defined according to the refinement scheme in  $OS^*$ , but relative to the history of *points subsampled from  $R$  in the order of their acceptance*. It is possible to show that the sequence of accepted points  $\{(T_i, X_i, U_i)\}_{i=1}^\infty$  have the same marginal distribution as the sequence of proposals in  $OS^*$ .

Thus, we can see  $OS^*$  and  $REJ$  as two separate procedures on the same realization of  $R$ . For the termination result, notice that  $REJ$  considers at least as many points as  $OS^*$ . For partial correctness, the points  $(T_i, X_i, U_i)$  such that  $U_i < f(X_i)/g_i(X_i)$  are exactly the subsampled points that would have resulted from thinning  $R$  directly with probability  $f(x)/(g(x)M(\mathbb{R}^n))$ . Thus, by the Accept-Reject Lemma, the returned values  $(T, X)$  will be the first arrival of an exponential race with measure  $P$ .  $\square$

$A^*$  sampling ( $A^*$  for short) is a branch and bound routine that finds the first arrival of a perturbed exponential race. It follows  $PER$  in principle by maintaining a lower bound on all future perturbed arrivals. The difference is that  $A^*$  maintains a piecewise constant lower bound over a partition of space that it progressively refines. On every iteration it selects the subset with smallest lower bound, samples the next arrival in that subset, and refines the subset unless it can terminate. It continues refining until the earliest perturbed arrival is less than the minimum of the piecewise constant lower bound. The name  $A^*$  sampling is a reference to  $A^*$  search (Hart et al., 1968), which is a path finding algorithm on graphs that uses a best-first criteria for selecting from heuristically valued nodes on the fringe of a set of visited nodes.  $A^*$  sampling was originally introduced by Maddison et al. (2014) as an algorithm that maximizes a perturbed Gumbel process. We define it over an exponential race for the sake of consistency. Usually, it is better to work



**Algorithm 7.5** A\* sampling for  $P$  with proposal  $Q$ 


---

```

 $\mathcal{L}, \mathcal{U} = \text{minPriorityQueue}(), \text{minPriorityQueue}()$ 
 $T_1 \sim \text{Exp}(Q(\mathbb{R}^n))$ 
 $\mathcal{L}.\text{pushWithPriority}(T_1/M(\mathbb{R}^n), \mathbb{R}^n)$ 
for  $i = 1$  to  $\infty$  do
   $(T_i/M(B_i), B_i) = \mathcal{L}.\text{pop}()$ 
   $X_i \sim Q(\cdot \cap B_i)/Q(B_i)$ 
   $\mathcal{U}.\text{pushWithPriority}(T_i g(X_i)/f(X_i), X_i)$ 
   $E \sim \text{Exp}(Q(B_i))$ 
   $T = T_i + E$ 
  if  $\min(\min \mathcal{L}, T/M(B_i)) < \min \mathcal{U}$  then
     $\mathcal{C} = \text{split}(B_i, X_i)$ 
    while  $\mathcal{C} \neq \emptyset$  do
       $C \sim \mathbb{P}(C) \propto Q(C)$  for  $C \in \mathcal{C}$ 
       $\mathcal{L}.\text{pushWithPriority}(T/M(C), C)$ 
       $\mathcal{C} = \mathcal{C} - \{C\}$ 
       $E \sim \text{Exp}(\sum_{C \in \mathcal{C}} Q(C))$ 
       $T = T + E$ 
    end while
  else
     $\mathcal{L}.\text{pushWithPriority}(T/M(B_i), B_i)$ 
  end if
  if  $\min \mathcal{L} \geq \min \mathcal{U}$  then
    return  $\mathcal{U}.\text{pop}()$ 
  end if
end for

```

---

with a Gumbel process to avoid numerical issues.

In more detail,  $\mathbf{A}^*$  searches over a simulation of an exponential race organized into a space partitioning tree, as in the right hand plot of Figure 7.4, for the first arrival of the perturbed race. The tree is determined by the splitting function  $\text{split}(B, x)$ . Each node  $v$  of the tree is associated with a subset  $B_v \subseteq \mathbb{R}^n$  and an arrival  $(T_v, X_v)$  from an exponential race with measure  $Q$ .  $\mathbf{A}^*$  iteratively expands a subtree of internal visited nodes, taking and visiting one node from the current fringe at each iteration. The fringe  $\mathcal{L}$  of the visited subtree is always a partition of  $\mathbb{R}^n$ . Each subset  $B \in \mathcal{L}$  is associated with the arrival time  $T$  of the next arrival of the race with measure  $Q$  in  $B$ . Therefore  $T/M(B)$  is a lower bound on all future perturbed arrivals in  $B$ .  $\mathcal{L}$  is implemented with a priority queue that prioritizes the subset  $B$  with the lowest regional bound  $T/M(B)$ . As  $\mathbf{A}^*$  expands the set of visited nodes the lower bound  $\min \mathcal{L}$  increases.

$\mathcal{L}$  is initialized with the root of the tree  $\{(T_1/M(\mathbb{R}^n), \mathbb{R}^n)\}$ . At the start of an iteration  $\mathbf{A}^*$  removes and visits the subset  $(T_i/M(B_i), B_i)$  with lowest lower bound on  $\mathcal{L}$ . Visiting a subset begins by realizing a location  $X_i$  from  $Q(\cdot \cap B_i)/Q(B_i)$  and pushing the perturbed arrival  $(T_i g(X_i)/f(X_i), X_i)$  onto

another priority queue  $\mathcal{U}$ .  $\mathcal{U}$  prioritizes earlier arrivals by the perturbed arrival times  $T_i g(X_i)/f(X_i)$ . In this way  $\mathbf{A}^*$  decreases the upper bound  $\min \mathcal{U}$  at each iteration.

$\mathbf{A}^*$  attempts to terminate by simulating the next arrival time  $T > T_i$  in  $B_i$  of the race with measure  $Q$ . If  $\min \mathcal{U} \leq \min(\min \mathcal{L}, T/M(B_i))$ , then the top of  $\mathcal{U}$  will not be superseded by future perturbed arrivals and it will be the first arrival of the perturbed race. If termination fails,  $\mathbf{A}^*$  refines the partition by splitting  $B_i$  into a partition  $\text{split}(B_i, X_i)$  of children. Arrival times for each of the children are assigned respecting the constraints of the exponential race in  $B_i$ . Each child  $C$  is pushed onto  $\mathcal{L}$  prioritized by its lower bound  $T/M(C)$ . Because the lower bounds have increased there is a second opportunity to terminate before continuing.  $\mathbf{A}^*$  checks if  $\min \mathcal{U} \leq \min \mathcal{L}$ , and otherwise continues, see Algorithm 7.5. As with **PER**,  $\mathbf{A}^*$  can be turned into a generator for iterating in order through the points of the perturbed race by replacing the **return** statement with a **yield** statement in Algorithm 7.5.

**Theorem 7.15** (Correctness of  $\mathbf{A}^*$  sampling). *Let  $K(\mathbf{A}^*)$  be the number of proposal samples considered before termination. Then*

$$\mathbb{P}(K(\mathbf{A}^*) > k) \leq (1 - \rho)^k \text{ where } \rho = \frac{P(\mathbb{R}^n)}{Q(\mathbb{R}^n)M(\mathbb{R}^n)}$$

and upon termination the return values  $(T, X)$  of  $\mathbf{A}^*$  sampling are independent and

$$T \sim \text{Exp}(P(\mathbb{R}^n)) \quad X \sim \frac{P(\cdot)}{P(\mathbb{R}^n)}.$$

*Proof.* Adapted from Maddison et al. (2014). The proposals are generated lazily in a space partitioning tree. If  $\{(T_i, X_i)\}_{i=1}^\infty$  are the arrivals at every node of the infinite tree sorted by increasing  $T_i$ , then  $(T_i, X_i)$  forms an exponential race with measure  $Q$ .

For the termination result, each node  $v$  of the tree can be associated with a subset  $B_v$  and a lower bound  $T_v/M(B_v)$ . One of the nodes will contain the first arrival of the perturbed process with arrival time  $T^*$ .  $\mathbf{A}^*$  visits at least every node  $v$  with  $T_v/M(B_v) > T^*$ . If  $M(B)$  is replaced with a constant  $M(\mathbb{R}^n)$ , then this can only increase the number of visited nodes. The last step is to realize that  $\mathbf{A}^*$  searching over a tree with constant bounds  $M(\mathbb{R}^n)$  searches in order of increasing  $T_v$ , and so corresponds to a realization of **PER**. The distribution of runtimes of **PER** is given in Lemma 7.13.

For partial correctness, let  $(T, X)$  be the return values with highest priority on the upper bound priority queue  $\mathcal{U}$ . The arrival time of unrealized perturbed arrivals is bounded by the lower bound priority queue  $\mathcal{L}$ . At termination  $T$  is less than the top of the lower bound priority queue. So

$P$	$Q$	$\Omega$	$N$	$\bar{K}(\mathbf{OS}^*)$	$\bar{K}(\mathbf{A}^*)$
clutter posterior	prior	$\mathbb{R}$	6	9.34	7.56
clutter posterior	prior	$\mathbb{R}^2$	6	38.3	33.0
clutter posterior	prior	$\mathbb{R}^3$	6	130	115
robust Bayesian regression	prior	$\mathbb{R}$	10	9.36	6.77
robust Bayesian regression	prior	$\mathbb{R}$	100	40.6	32.2
robust Bayesian regression	prior	$\mathbb{R}$	1000	180	152
fully connected Ising model	uniform	$\{-1, 1\}^5$	-	4.37	3.50
fully connected Ising model	uniform	$\{-1, 1\}^{10}$	-	19.8	15.8

**Table 7.1:** Comparing  $\mathbf{A}^*$  and  $\mathbf{OS}^*$ . Clutter and robust Bayesian regression are adapted from Maddison et al. (2014) and the Ising model from Kim et al. (2016).  $\Omega$  is the support of the distribution;  $N$  is the number of data points; and  $\bar{K}(\mathbf{OS}^*)$  and  $\bar{K}(\mathbf{A}^*)$  are averaged over 1000 runs. More information in the Appendix.

no unrealized points will arrive before  $(T, X)$ . By Lemma 7.9  $(T, X)$  is the first arrival of an exponential race with measure  $P$ .  $\square$

#### 7.5.4 Runtime of $\mathbf{A}^*$ sampling and $\mathbf{OS}^*$

$\mathbf{A}^*$  and  $\mathbf{OS}^*$  are structurally similar; both search over a partition of space and refine it to increase the probability of terminating. They will give practical benefits over rejection sampling if the bounds  $M(B)$  shrink as the volume of  $B$  shrinks. In this case the bound on the probability of rejecting  $k$  proposals given in Theorems 7.14 and 7.15 can be very loose, and  $\mathbf{OS}^*$  and  $\mathbf{A}^*$  can be orders of magnitude more efficient than rejection sampling. Still, these methods scale poorly with dimension.

The cost of running  $\mathbf{A}^*$  and  $\mathbf{OS}^*$  will be dominated by computing the ratio of densities  $f(x)/g(x)$  and computing bounds  $M(B)$ . Because the number of bound computations is within a factor of 2 of the number of density computations, the number of evaluations of  $f(x)/g(x)$  (equivalently number of proposals) is a good estimate of complexity. Table 7.1 presents a summary of experimental evidence that  $\mathbf{A}^*$  makes more efficient use of density computations across three different problems. For each problem the full descriptions of  $P$ ,  $Q$ ,  $M(B)$ , and  $\text{split}(B, x)$  are found in the Appendix.

The dominance of  $\mathbf{A}^*$  in experiments is significant, because it has access to the same information as  $\mathbf{OS}^*$ . There are at least two factors that may give  $\mathbf{A}^*$  this advantage. First, if all lower bounds increase sharply after some exploration  $\mathbf{A}^*$  can retroactively take advantage of that information, as in

Section 7.5.2. Second,  $\mathbf{A}^*$  can take advantage of refined bound information on the priority queue  $\mathcal{L}$  before proposing the next sample. Still, the difference in search strategy and termination condition may counteract these advantages, so a rigorous theory is needed to confirm exactly the sense in which  $\mathbf{A}^*$  and  $\mathbf{OS}^*$  differ. We refer readers to Maddison et al. (2014) for more detailed experiments.

---

## 7.6 Conclusion

The study of Poisson processes is traditionally motivated by their application to natural phenomenon, and Monte Carlo methods are developed specifically for them (Ripley, 1977; Geyer and Møller, 1994). We considered the inverse relationship, using Poisson processes to better understand Monte Carlo methods. We suspect that this general perspective holds value for future directions in research.

Monte Carlo methods that rely on bounds are not suitable for most high dimensional distributions. Rejection sampling scales poorly with dimensionality. Even for  $\mathbf{A}^*$  sampling there are simple examples where adaptive bounds become uninformative in high dimensions, such as sampling from the uniform hypersphere when using hyperrectangular search subsets. Still, specialized algorithms for limited classes of distributions may be able to take advantage of conditional independence structure to improve their scalability.

Another direction is to abandon the idea of representing arbitrary distributions, and study the class of distributions represented by the maxima of combinations of lower order Gumbel processes. This is the approach of the perturbation models studied in Papandreou and Yuille; Gane et al.; Hazan and Jaakkola; Tarlow et al.; and Keshet et al. of this book. In these models a Gumbel process over a discrete space is replaced by sums of independent Gumbel processes over discrete subspaces. The maxima of these models form a natural class of distributions complete with their own measures of uncertainty. An open direction of inquiry is developing efficient algorithms for optimizing their continuous counterparts.

Our study of Poisson processes and Monte Carlo methods was dominated by the theme of independence; the points of an exponential race arrive as independent random variables and accept-reject or perturb do not introduce correlations between the points of the transformed race. Continuing in this direction it is natural to investigate whether other Poisson process models or other operations on an exponential race could be used to define a new class of Monte Carlo methods. In a separate direction the Markov Chain Monte Carlo (MCMC) methods produce a sequence of correlated samples

whose limiting distribution is the distribution of interest. The theory of point processes includes a variety of limit theorems, which describe the limiting distribution of random countable sets (Daley and Vere-Jones, 2007). It would be interesting to see whether a point process treatment of MCMC bears fruit, either in unifying our proof techniques or inspiring new algorithms.

---

## Acknowledgements

We would like to thank Daniel Tarlow and Tom Minka for the ideas, discussions, and support throughout this project. Thanks to the other editors Tamir Hazan and George Papandreou. Thanks to Jacob Steinhardt, Yee Whye Teh, Arnaud Doucet, Christian Robert for comments on the draft. Thanks to Sir J.F.C. Kingman for encouragement. This work was supported by the Natural Sciences and Engineering Research Council of Canada.

---

## Appendix: Proof of Lemma 7.3

*Proof.* The lemma is trivial satisfied for  $k = 0$ . For  $k > 0$  and  $B_i \subseteq B$  we will express

$$\mathbb{P}(\{X_i \in B_i\}_{i=1}^k | N(B) = k) \quad (7.11)$$

in terms of counts. The difficulty lies in the possible overlap of  $B_i$ s, so we consider  $2^k$  sets of the form

$$A_j = B_1^* \cap B_2^* \cap \dots \cap B_k^*$$

where  $*$  is blank or a complement, and  $A_1$  is interpreted as  $B \cap B_1^c \cap \dots \cap B_k^c$ . The  $A_j$  are a disjoint partition of  $B$ ,

$$B_i = \cup_{j \in I(i)} A_j, \quad B = \cup_{j=1}^{2^k} A_j,$$

where  $I(i) \subseteq \{1, \dots, 2^k\}$  is some subset of indices. Let  $\mathcal{J} = I(1) \times I(2) \times \dots \times I(k)$ , so that each  $s \in \mathcal{J}$  is a vector indices  $(s_1, s_2, \dots, s_k)$  associated with the disjoint events  $\{X_i \in A_{s_i}\}_{i=1}^k$ . Thus,

$$\mathbb{P}(\{X_i \in B_i\}_{i=1}^k | N(B) = k) = \sum_{s \in \mathcal{J}} \mathbb{P}(\{X_i \in A_{s_i}\}_{i=1}^k | N(B) = k).$$

For  $s \in \mathcal{J}$ , let  $n_j(s) = \#\{i : s_i = j\}$  be the number of indices in  $s$  equal to  $j$  and notice that  $\sum_{j=1}^{2^k} n_j(s) = k$ . To relate the probability if specific numbering  $\{X_i \in A_{s_i}\}_{i=1}^k$  with counts  $\{N(A_j) = n_j(s)\}_{j=1}^{2^k}$ , we discount by

all ways of the arranging  $k$  points that result in the same counts.

$$\begin{aligned} \mathbb{P}(\{X_i \in A_{s_i}\}_{i=1}^k | N(B) = k) &= \frac{\prod_{j=1}^{2^k} n_j(s)! \mathbb{P}(\{N(A_j) = n_j(s)\}_{j=1}^{2^k})}{k! \mathbb{P}(N(B) = k)} \\ &= \frac{\prod_{j=1}^{2^k} \mu(A_j)^{n_j(s)}}{\mu(B)^k}. \end{aligned}$$

Thus (7.11) is equal to

$$\sum_{s \in \mathcal{J}} \frac{\prod_{j=1}^{2^k} \mu(A_j)^{n_j(s)}}{\mu(B)^k} = \prod_{i=1}^k \frac{\sum_{j \in I(i)} \mu(A_j)}{\mu(B)} = \prod_{i=1}^k \frac{\mu(B_i)}{\mu(B)}$$

□

## Appendix: Experimental details

### Clutter posterior

This example is taken exactly from Maddison et al. (2014). The clutter problem (Minka, 2001) is to estimate the mean  $\theta \in \mathbb{R}^n$  of a Normal distribution under the assumption that some points are outliers. The task is to sample from the posterior  $P$  over  $w$  of some empirical sample  $\{(x_i)\}_{i=1}^N$ .

$$\begin{aligned} f_i(\theta) &= \frac{0.5 \exp(-0.5 \|\theta - x_i\|^2)}{(2\pi)^{n/2}} + \frac{0.5 \exp(-0.5 \|x_i\|^2 / 100^2)}{100^n (2\pi)^{n/2}} \\ \log g(\theta) &= -\frac{\|\theta\|^2}{8} \quad \log f(\theta) = \log g(\theta) + \sum_{i=1}^N \log f_i(\theta) \\ (a, b) &= \{y : a_d < y_d \leq b_d\} \text{ for } a, b \in \mathbb{R}^n \\ M((a, b)) &= \prod_{i=1}^N f_i(x^*(a, b, x_i)) \quad x^*(a, b, x)_d = \begin{cases} a_d & \text{if } x_d < a_d \\ b_d & \text{if } x_d > b_d \\ x_d & \text{o.w.} \end{cases} \end{aligned}$$

$$\begin{aligned} \text{split}((a, b), x) &= \{(a, b) \cap \{y : y_s \leq x_s\}, (a, b) \cap \{y : y_s > x_s\}\} \\ &\text{where } s = \underset{d}{\operatorname{argmax}} b_d - a_d \end{aligned}$$

Our dataset was 6 points  $x_i \in \mathbb{R}^n$  of the form  $x_i = (a_i, a_i, \dots, a_i)$  for  $a_i \in \{-5, -4, -3, 3, 4, 5\}$ .

### Robust Bayesian regression

This example is an adaption from Maddison et al. (2014) with looser bounds. The model is a robust linear regression  $y_i = wx_i + \epsilon_i$  where the noise  $\epsilon_i$  is distributed as a standard Cauchy and  $w$  is a standard Normal. The task is to sample from the posterior  $P$  over  $w$  of some empirical sample  $\{(x_i, y_i)\}_{i=1}^N$ .

$$\log g(w) = -\frac{w^2}{8}$$

$$\log f(w) = \log g(w) - \sum_{i=1}^N \log(1 + (wx_i - y_i)^2)$$

$$M((a, b)) = \prod_{i=1}^N M_i((a, b)) \quad M_i((a, b)) = \begin{cases} \exp(a) & \text{if } y_i/x_i < a \\ \exp(b) & \text{if } y_i/x_i > b \\ \exp(y_i/x_i) & \text{o.w.} \end{cases}$$

$$\text{split}((a, b), x) = \{(a, x], (x, b]\}$$

The dataset was generated by setting  $w^* = 2$ ;  $x_i \sim \text{Normal}(0, 1)$  and  $y_i = wx_i + \epsilon$  with  $\epsilon \sim \text{Normal}(0, 0.1^2)$  for  $i \leq N/2$ ; and  $x_i = x_{i-N/2}$  and  $y_i = -y_{i-N/2}$  for  $i > N/2$ .

### Attractive fully connected Ising model

This is an adaptation of Kim et al. (2016). The attractive fully connected Ising model is a distribution over  $x \in \{-1, 1\}^n$  described by parameters  $w_{ij} \sim \text{Uniform}[0, 0.2]$  and  $f_i \sim \text{Uniform}[-1, 1]$ .

$$\log g(x) = 0$$

$$\log f(x) = \sum_i f_i x_i + \sum_{i < j \leq n} w_{ij} x_i x_j$$

We considered subsets of the form  $B = \{x : x_i = b_i, i \in I\}$  where  $I \subseteq \{1, \dots, n\}$  and  $b_i \in \{0, 1\}$ . We split on one of the unspecified variables  $x_i$  by taking variable whose linear program relaxation was closest to 0.5.

$$\text{split}(B, x) = \{B \cap \{x : x_i = 0\}, B \cap \{x : x_i = 1\}\}$$

$\log M(B)$  is computed by solving a linear program relaxation for the following type of integer program. Let  $b_i \in \{0, 1\}$  for  $1 \leq i \leq n$  and  $b_{ijkl} \in \{0, 1\}$  for  $1 \leq i < j \leq n$  and  $k, l \in \{0, 1\}$ .

$$\min_x \sum_i -f_i b_i + f_i(1 - b_i) + \sum_{1 \leq i < j \leq n} \sum_{k, l \in \{0, 1\}} (-1)^{kl + (1-l)(1-k)} w_{ij} b_{ijkl}$$

subject to the constraints for  $1 \leq i < j \leq n$ ,

$$\begin{aligned} \sum_{l \in \{0,1\}} b_{ij0l} &= 1 - b_i & \sum_{k \in \{0,1\}} b_{ijk0} &= 1 - b_j \\ \sum_{l \in \{0,1\}} b_{ij1l} &= b_i & \sum_{k \in \{0,1\}} b_{ijk1} &= b_j \end{aligned}$$

as the subsets  $B$  narrowed we just solved new linear programs with constants for the fixed variables.

## 7.9 References

- G. Casella, C. P. Robert, and M. T. Wells. *Generalized Accept-Reject sampling schemes*, volume 45 of *Lecture Notes–Monograph Series*. Institute of Mathematical Statistics, 2004.
- Y. Chen and Z. Ghahramani. Scalable Discrete Sampling as a Multi-Armed Bandit Problem. *ArXiv e-prints*, June 2015.
- D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes, Volume II: General Theory and Structure*. Springer, 2007.
- L. Devroye. *Non-Uniform Random Variate Generation*. Springer, 1986.
- R. Durrett. *Essentials of Stochastic Processes*. Springer, 2012.
- M. Dymetman, G. Bouchard, and S. Carter. The OS\* Algorithm: a Joint Approach to Exact Optimization and Sampling. *arXiv preprint arXiv:1207.0742*, 2012.
- B. Gedenko. On a local limit theorem of the theory of probability. *Uspekhi Mat. Nauk*, 3, 1948.
- C. J. Geyer and J. Møller. Simulation procedures and likelihood inference for spatial point processes. *Scandinavian Journal of Statistics*, 1994.
- W. R. Gilks and P. Wild. Adaptive rejection sampling for Gibbs sampling. *Applied Statistics*, 1992.
- E. J. Gumbel and J. Lieblein. *Statistical Theory of Extreme Values and Some Practical Applications: a Series of Lectures*. US Govt. Print. Office, 1954.
- P. Hart, N. Nilsson, and B. Raphael. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *Systems Science and Cybernetics, IEEE Transactions on*, 4(2), 1968.
- T. Hazan, S. Maji, and T. Jaakkola. On Sampling from the Gibbs Distribution with Random Maximum A-Posteriori Perturbations. In *NIPS*. 2013.
- C. Kim, A. Sabharwal, and S. Ermon. Exact Sampling with Integer Linear Programs and Random Perturbations. In *AAAI*, 2016.
- J. Kingman. *Poisson Processes*. Oxford University Press, 1992.
- P. A. W. Lewis and G. S. Shedler. Simulation of nonhomogeneous poisson processes by thinning. *Naval Research Logistics Quarterly*, 26(3), 1979.
- C. J. Maddison, D. Tarlow, and T. Minka. A\* Sampling. In *NIPS*. 2014.
- H. Malmberg. Random Choice over a Continuous Set of Options. 2013.
- T. P. Minka. Expectation propagation for approximate Bayesian inference. In *UAI*,



- 2001.
- G. Papandreou and A. Yuille. Perturb-and-MAP Random Fields: Using Discrete Optimization to Learn and Sample from Energy Models. In *ICCV*, 2011.
- S. I. Resnick. *Extreme Values, Regular Variation and Point Processes*. Springer, 2007.
- B. D. Ripley. Modelling Spatial Patterns. *Journal of the Royal Statistical Society. Series B (Methodological)*, 1977.
- D. Tarlow, R. P. Adams, and R. S. Zemel. Randomized Optimum Models for Structured Prediction. In *AISTATS*, 2012.
- A. J. Walker. An efficient method for generating discrete random variables with general distributions. *ACM Transactions on Mathematical Software*, 3(3), 1977.
- J. I. Yellott. The relationship between Luce's choice axiom, Thurstone's theory of comparative judgment, and the double exponential distribution. *Journal of Mathematical Psychology*, 15(2), 1977.