

Furthermore, we also updated the left middle panel in Figure 2 with a more precise estimate of the codelength for the methods. We suspect that split-on-sample GPRS's slight divergence from the predicted average is due to a numerical issue in our implementation, which we did not manage to fix before the supplementary material submission deadline.

Greedy Poisson Rejection Sampling

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 One-shot channel simulation is a fundamental data compression problem concerned
2 with encoding a single sample from a target distribution Q using a coding distribu-
3 tion P using as few bits as possible on average. Algorithms that solve this problem
4 find applications in neural data compression and differential privacy and can serve
5 as a more efficient and natural alternative to quantization-based methods. Unfor-
6 tunately, existing solutions are too slow or have limited applicability, preventing
7 their widespread adoption. In this paper, we conclusively solve one-shot channel
8 simulation for one-dimensional problems where the target-proposal density ratio
9 is unimodal by describing an algorithm with optimal runtime. We achieve this by
10 constructing a rejection sampling procedure equivalent to greedily searching over
11 the points of a Poisson process. Hence, we call our algorithm greedy Poisson rejec-
12 tion sampling (GPRS) and analyze the correctness and time complexity of several
13 of its variants. Finally, we empirically verify our theorems, demonstrating that
14 GPRS significantly outperforms the current state-of-the-art method, A* coding.

15 1 Introduction

16 It is a common misconception that quantization is essential to lossy data compression; it is merely
17 a way to discard information deterministically. In this paper, we consider the alternative, that is, to
18 discard information stochastically using *one-shot channel simulation*. To illustrate the main idea,
19 take lossy image compression as an example. Assume we have a generative model given by a joint
20 distribution $P_{\mathbf{x},\mathbf{y}}$ over images \mathbf{y} and latent variables \mathbf{x} , e.g. we might have trained a variational
21 autoencoder (VAE; Kingma & Welling, 2014) on a dataset of images. To compress a new image
22 \mathbf{y} , we encode a single sample from its posterior $\mathbf{x} \sim P_{\mathbf{x}|\mathbf{y}}$ as its stochastic lossy representation.
23 The decoder can obtain a lossy reconstruction of \mathbf{y} by decoding \mathbf{x} and drawing a sample $\hat{\mathbf{y}} \sim P_{\mathbf{y}|\mathbf{x}}$
24 (though in practice, for a VAE we normally just take the mean predicted by the generative network).

25 Abstracting away from our example, in this paper we will be entirely focused on *channel simulation*
26 for a pair of correlated random variables $\mathbf{x}, \mathbf{y} \sim P_{\mathbf{x},\mathbf{y}}$: given a source symbol $\mathbf{y} \sim P_{\mathbf{y}}$ we wish to
27 encode **a single sample** $\mathbf{x} \sim P_{\mathbf{x}|\mathbf{y}}$. A simple way to achieve this is to encode \mathbf{x} with entropy coding
28 using the marginal $P_{\mathbf{x}}$, whose average coding cost is approximately the entropy $\mathbb{H}[\mathbf{x}]$. Surprisingly,
29 however, we can do much better by using a *channel simulation protocol*, whose average coding cost
30 is approximately the mutual information $I[\mathbf{x}; \mathbf{y}]$ (Li & El Gamal, 2018). This is remarkable, since not
31 only $I[\mathbf{x}; \mathbf{y}] \leq \mathbb{H}[\mathbf{x}]$, but in many cases $I[\mathbf{x}; \mathbf{y}]$ might be finite even though $\mathbb{H}[\mathbf{x}]$ is infinite, such as
32 when \mathbf{x} is continuous. Sadly, most existing protocols place heavy restrictions on $P_{\mathbf{x},\mathbf{y}}$ or their runtime
33 scales much worse than $\mathcal{O}(I[\mathbf{x}; \mathbf{y}])$, limiting their practical applicability (Agustsson & Theis, 2020).

34 In this paper, we propose a family of channel simulation protocols based on a new rejection sampling
35 algorithm, which we can apply to simulate samples from a target distribution Q using a proposal
36 distribution P over an arbitrary probability space. The inspiration for our construction comes from an
37 exciting recent line of work which recasts random variate simulation as a search problem over a set
38 of randomly placed points, specifically a Poisson process (Maddison, 2016). The most well-known

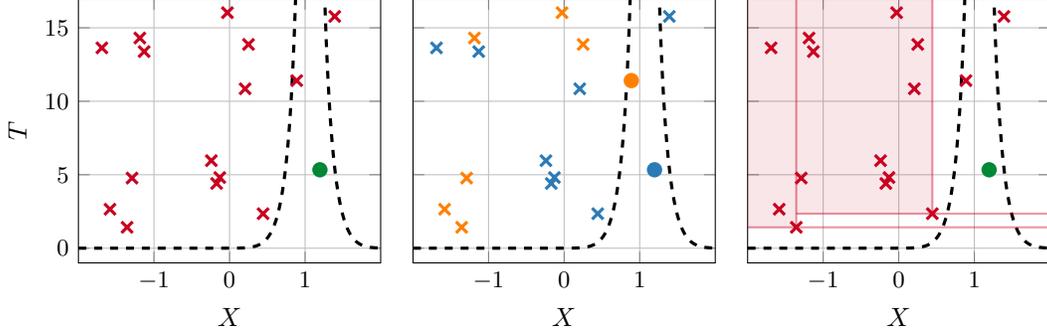


Figure 1: Illustration of three GPRS procedures for a Gaussian target $Q = \mathcal{N}(1, 0.25^2)$ and Gaussian proposal distribution $P = \mathcal{N}(0, 1)$, with the time axis truncated to the first 17 units. All three variants find the first arrival of a $(1, P)$ -Poisson process Π under the graph of $\varphi = \sigma \circ r$ indicated by the **thick dashed black line** in each plot. Here, $r = dQ/dP$ is the target-proposal density ratio, and σ is given by Equation (3). **Left:** Algorithm 3 sequentially searching through the points of Π . The green circle (●) shows the first point of Π that falls under φ , and is accepted. All other points are rejected, as indicated by red crosses (X). In practice, Algorithm 3 does not simulate points of Π that arrive after the accepted arrival. **Middle:** Parallelized GPRS (Algorithm 4) searching through two independent $(1/2, P)$ -Poisson processes Π_1 and Π_2 in parallel. Blue points are arrivals in Π_1 and orange points are arrivals in Π_2 . Crosses (X) indicate rejected, and circles (●) indicate accepted points by each thread. In the end, the algorithm accepts the earliest arrival across all processes, which in this case is marked by the blue circle (●). **Right:** GPRS with binary search (Algorithm 5), when φ is unimodal. The shaded red areas are never searched or simulated by the algorithm since, given the first two rejections, we know points in those regions cannot fall under φ .

39 examples are the Gumbel-Max trick and A* sampling (Maddison et al., 2014). Our algorithm, which
 40 we call *greedy Poisson rejection sampling* (GPRS), differs significantly from all previous approaches
 41 in terms of what it is searching for, which we can succinctly summarise as: “GPRS searches for the
 42 first arrival of a Poisson process Π under the graph of an appropriately defined function φ ”. The
 43 first and simplest variant of GPRS is equivalent to an exhaustive search over all points of Π in time
 44 order. Next, we show that the linear search is embarrassingly parallelizable, leading to a parallelized
 45 variant of GPRS. Finally, when the underlying probability space has more structure, we develop
 46 branch-and-bound variants of GPRS that perform a binary search over the points of Π . See Figure 1
 47 for an illustration of these three variants.

48 While GPRS is an interesting sampling algorithm on its own, we also show that each of its variants
 49 induces a new one-shot channel simulation protocol. That is, after we receive $\mathbf{y} \sim P_{\mathbf{y}}$, we can set
 50 $Q \leftarrow P_{\mathbf{x}|\mathbf{y}}$ and $P \leftarrow P_{\mathbf{x}}$ and use GPRS to encode a sample $\mathbf{x} \sim P_{\mathbf{x}|\mathbf{y}}$ at an average bitrate of a
 51 little more than the mutual information $I[\mathbf{x}; \mathbf{y}]$. In particular, on one-dimensional channel simulation
 52 problems where the density ratio $dP_{\mathbf{x}|\mathbf{y}}/dP_{\mathbf{x}}$ is unimodal for all \mathbf{y} , the binary search variant of GPRS
 53 leads to a protocol with an average runtime of $\mathcal{O}(I[\mathbf{x}; \mathbf{y}])$, which is optimal. This is a considerable
 54 improvement over A* coding (Flamich et al., 2022), the current state-of-the-art method.

55 In summary, our contributions are as follows:

- 56 • We construct a new rejection sampling algorithm called *greedy Poisson rejection sampling*,
 57 which we can construe as a greedy search over the points of a Poisson process (Algorithm 3).
 58 We propose a parallelized (Algorithm 4) and a branch-and-bound variant (Algorithms 5
 59 and 6) of GPRS. We analyze the correctness and runtime of these algorithms.
- 60 • We show that each variant of GPRS induces a one-shot channel simulation protocol for
 61 correlated random variables $\mathbf{x}, \mathbf{y} \sim P_{\mathbf{x}, \mathbf{y}}$, achieving the optimal average codelength of
 62 $I[\mathbf{x}; \mathbf{y}] + \log_2(I[\mathbf{x}; \mathbf{y}] + 1) + \mathcal{O}(1)$ bits.
- 63 • We prove that when \mathbf{x} is a \mathbb{R} -valued random variable and the density ratio $dP_{\mathbf{x}|\mathbf{y}}/dP_{\mathbf{x}}$
 64 is always unimodal, the channel simulation protocol based on the binary search variant of
 65 GPRS achieves $\mathcal{O}(I[\mathbf{x}; \mathbf{y}])$ runtime, which is optimal.
- 66 • We conduct toy experiments on one-dimensional problems and show that GPRS compares
 67 favourably against A* coding, the current state-of-the-art channel simulation protocol.

68 2 Background

69 The sampling algorithms we construct in this paper are search procedures on randomly placed points
70 in space whose distribution is given by a Poisson process Π . Thus, in this section, we first review
71 the necessary theoretical background for Poisson processes and how we can simulate them on a
72 computer. Then, we formulate standard rejection sampling as a search procedure over the points of a
73 Poisson process to serve as a prototype for our algorithm in the next section. Up to this point, our
74 exposition loosely follows Sections 2, 3 and 5.1 of the excellent work of Maddison (2016), peppered
75 with a few additional results that will be useful for analyzing our algorithm later. Finally, we describe
76 the channel simulation problem, using rejection sampling as a rudimentary solution and describe its
77 shortcomings. This motivates the development of greedy Poisson rejection sampling in Section 3.

78 2.1 Poisson Processes

79 A Poisson process Π is a countable collection of random points in some mathematical space Ω . In
80 the main text, we will always assume that Π is defined over $\Omega = \mathbb{R}^+ \times \mathbb{R}^d$ and that all objects
81 involved are measure-theoretically well-behaved for simplicity. For this choice of Ω , the positive
82 reals represent *time*, and the unit interval represents *space*. However, most results generalize to
83 settings when the spatial domain \mathbb{R}^d is replaced with some more general space, and we give a general
84 measure-theoretic construction in Appendix A, where the spatial domain is an arbitrary Polish space.

85 **Basic properties of Π :** For a set $A \subseteq \Omega$, let $\mathbf{N}(A) \stackrel{\text{def}}{=} |\Pi \cap A|$ denote the number of points of Π
86 falling in the set A , where $|\cdot|$ denotes the cardinality of a set. Then, Π is characterized by the following
87 two fundamental properties (Kingman, 1992). First, for two disjoint sets $A, B \subseteq \Omega$, $A \cap B = \emptyset$,
88 the number of points of Π that fall in either set are independent random variables: $\mathbf{N}(A) \perp \mathbf{N}(B)$.
89 Second, $\mathbf{N}(A)$ is Poisson distributed with *mean measure* $\mu(A) \stackrel{\text{def}}{=} \mathbb{E}[\mathbf{N}(A)]$.

90 **Time-ordering the points of Π :** Since we assume that $\Omega = \mathbb{R}^+ \times \mathbb{R}^d$ has a product space structure,
91 we may write the points of Π as a pair of time-space coordinates: $\Pi = \{(T_n, X_n)\}_{n=1}^\infty$. Furthermore,
92 we can order the points in Π with respect to their time coordinates and *index* them accordingly, i.e.
93 for $i < j$ we have $T_i < T_j$. Hence, we refer to (T_n, X_n) as the *nth arrival* of the process.

94 As a slight abuse of notation, we define $\mathbf{N}(t) \stackrel{\text{def}}{=} \mathbf{N}([0, t) \times \mathbb{R}^d)$ and $\mu(t) \stackrel{\text{def}}{=} \mathbb{E}[\mathbf{N}(t)]$,
95 i.e. these quantities measure the number and average number of points of Π that arrive
96 before time t , respectively. In this paper, we assume $\mu(t)$ has derivative $\mu'(t)$, and as-
97 sume for each $t \geq 0$ there is a conditional probability distribution $P_{X|T=t}$ with density
98 $p(x | t)$, such that we can write the mean measure as $\mu(A) = \int_A p(x | t) \mu'(t) dx dt$.

100 **Simulating Π :** A simple method to simulate Π on a com-
101 puter is to realize it in time order, i.e. at step n , simulate
102 T_n and then use it to simulate $X_n \sim P_{X|T=T_n}$. We can
103 find out the distribution of Π 's first arrival by noting that
104 by its definition, no point of Π can come before it, hence
105 $\mathbb{P}[T_1 \geq t] = \mathbb{P}[\mathbf{N}(t) = 0] = \exp(-\mu(t))$, where the sec-
106 ond equality follows from the fact that $\mathbf{N}(t)$ is Poisson dis-
107 tributed. A particularly important case is when Π is *time-*
108 *homogeneous*, i.e. $\mu(t) = \lambda t$ for some $\lambda > 0$, in which case
109 $T_1 \sim \text{Exp}(\lambda)$ is an exponential random variable with *rate* λ .
110 In fact, all of Π 's inter-arrival times $\Delta_n = T_n - T_{n-1}$ for
111 $n \geq 1$ share this simple distribution, where we set $T_0 = 0$.
112 To see this, note that

$$\mathbb{P}[\Delta_n \geq t | T_{n-1}] = \mathbb{P}[\mathbf{N}([T_{n-1}, T_{n-1} + t) \times \mathbb{R}^d) = 0 | T_{n-1}] = \exp(-\lambda t),$$

113 i.e. all $\Delta_n | T_{n-1} \sim \text{Exp}(\lambda)$. Therefore, we can use the above procedure to simulate time-
114 homogeneous Poisson processes, described in Algorithm 1. We will refer to a time-homogeneous
115 Poisson process with time rate λ and spatial distribution $P_{X|T}$ as a $(\lambda, P_{X|T})$ -Poisson process.

116 **Rejection sampling using Π :** Rejection sampling is a technique to simulate samples from a *target*
117 *distribution* Q using a *proposal distribution* P , assuming we can find an upper bound $M > 0$ for
118 their density ratio $r = dQ/dP$ (technically, the Radon-Nikodym derivative). We can formulate this

Algorithm 1: Generating a $(\lambda, P_{X|T})$ -Poisson process.

Input : Time rate λ ,
Spatial distribution $P_{X|T}$

$T_0 \leftarrow 0$
for $n = 1, 2, \dots$ **do**
 $\Delta_n \sim \text{Exp}(\lambda)$
 $T_n \leftarrow T_{n-1} + \Delta_n$
 $X_n \sim P_{X|T=T_n}$
 yield (T_n, X_n)

end

Algorithm 2: Standard rejection sampling.

Input : Proposal distribution P ,
Density ratio $r = dQ/dP$,
Upper bound M for r .
Output : Sample $X \sim Q$ and its index N .
// Generator for a $(1, P)$ -Poisson
process using Algorithm 1.
 $\Pi \leftarrow \text{SimulatePP}(1, P)$
for $n = 1, 2, \dots$ **do**
| $(T_n, X_n) \leftarrow \text{next}(\Pi)$
| $U_n \sim \text{Unif}(0, 1)$
| **if** $U_n < r(X_n)/M$ **then**
| | **return** X_n, n
| **end**
end

Algorithm 3: Greedy Poisson rejection sampling.

Input : Proposal distribution P ,
Density ratio $r = dQ/dP$,
Stretch function σ .
Output : Sample $X \sim Q$ and its index N .
// Generator for a $(1, P)$ -Poisson
process using Algorithm 1.
 $\Pi \leftarrow \text{SimulatePP}(1, P)$
for $n = 1, 2, \dots$ **do**
| $(T_n, X_n) \leftarrow \text{next}(\Pi)$
| **if** $T_n < \sigma(r(X_n))$ **then**
| | **return** X_n, n
| **end**
end

119 procedure using a Poisson process: we simulate the arrivals (T_n, X_n) of a $(1, P)$ -Poisson process
120 Π , but we only keep them with probability $r(X_n)/M$, otherwise, we delete them. This algorithm is
121 described in Algorithm 2; its correctness is guaranteed by the *thinning theorem* (Maddison, 2016).

122 **Rejection sampling is suboptimal:** Using Poisson processes to formulate rejection sampling
123 highlights a subtle but crucial inefficiency: it does not make use of Π 's temporal structure and only
124 uses the spatial coordinates. GPRS fixes this by using a rejection criterion that does depend on the
125 time variable. As we show, this significantly speeds up sampling for certain classes of distributions.

126 2.2 Channel Simulation

127 The main motivation for our work is to develop a *one-shot channel simulation protocol* using the
128 sampling algorithm we derive in Section 3. Channel simulation is of significant theoretical and
129 practical interest. Recent works used it to compress neural network weights, achieving state-of-the-art
130 performance (Havasi et al., 2018); to perform image compression using variational autoencoders
131 (Flamich et al., 2020) and diffusion models with perfect realism (Theis et al., 2022); and to perform
132 differentially private federated learning by compressing noisy gradients (Shah et al., 2022).

133 One-shot channel simulation is a communication problem between two parties, Alice and Bob,
134 sharing a joint distribution $P_{\mathbf{x}, \mathbf{y}}$ over two correlated random variables \mathbf{x} and \mathbf{y} , where we assume
135 that Alice and Bob can simulate samples from the marginal $P_{\mathbf{x}}$. In a single round of communication,
136 Alice receives a sample $\mathbf{y} \sim P_{\mathbf{y}}$ from the marginal distribution over \mathbf{y} . Then, she needs to send the
137 minimum number of bits to Bob such that he can **simulate a single sample** from the conditional
138 distribution $\mathbf{x} \sim P_{\mathbf{x}|\mathbf{y}}$. Note that Bob **does not want to learn** $P_{\mathbf{x}|\mathbf{y}}$; he just wants to simulate a single
139 sample from it. Surprisingly, when Alice and Bob have access to *shared randomness*, e.g. by sharing
140 the seed of their random number generator before communication, they can solve channel simulation
141 very efficiently. Mathematically, in this paper we will always model this shared randomness by some
142 time-homogeneous Poisson process (or processes) Π , since given a shared random seed, Alice and
143 Bob can always simulate the same process e.g. using Algorithm 1. Then, the average coding cost
144 of \mathbf{x} given Π is its conditional entropy $\mathbb{H}[\mathbf{x} | \Pi]$ and, surprisingly, it is always upper bounded by
145 $\mathbb{H}[\mathbf{x} | \Pi] \leq I[\mathbf{x}; \mathbf{y}] + \log_2 I[\mathbf{x}; \mathbf{y}] + \mathcal{O}(1)$, where $I[\mathbf{x}; \mathbf{y}]$ is the mutual information between \mathbf{x} and
146 \mathbf{y} (Li & El Gamal, 2018). This is an especially curious result, given that in many cases $\mathbb{H}[\mathbf{x}]$ is
147 infinite while $I[\mathbf{x}; \mathbf{y}]$ is finite, e.g. when \mathbf{x} is a continuous variable. In essence, this result means that
148 given the additional structure $P_{\mathbf{x}, \mathbf{y}}$, channel simulation protocols can “offload” an infinite amount of
149 information into the shared randomness, and only communicate the finitely many “necessary” bits.

150 **An example channel simulation protocol with rejection sampling:** Given Π and $\mathbf{y} \sim P_{\mathbf{y}}$, Alice
151 sets $Q \leftarrow P_{\mathbf{x}|\mathbf{y}}$ as the target and $P \leftarrow P_{\mathbf{x}}$ as the proposal distribution with density ratio $r = dQ/dP$,
152 and run the rejection sampler in Algorithm 2 to find the first point of Π that was not deleted. She
153 counts the number of samples N she had to simulate before acceptance and sends this number to Bob.
154 He decodes a sample $\mathbf{x} \sim P_{\mathbf{x}|\mathbf{y}}$ by selecting the spatial coordinate of the N th arrival of Π .

155 Unfortunately, this simple protocol is suboptimal. To see this, let $D_{\infty}[Q||P] \stackrel{\text{def}}{=} \sup_{\mathbf{x} \in \Omega} \{\log_2 r(\mathbf{x})\}$
156 denote Rényi ∞ -divergence from Q to P , and recall two standard facts: (1) the best possible upper

157 bound Alice can use for rejection sampling is $M_{opt} = \exp_2(D_\infty[Q\|P])$, where $\exp_2(x) = 2^x$, and
 158 (2), the number of samples N drawn until acceptance is a geometric random variable with mean
 159 M_{opt} (Maddison, 2016). We now state the two issues with rejection sampling that GPRS solves.

160 **Problem 1: Codelength:** However, by using the formula for the entropy of a geometric random
 161 variable and assuming Alice uses the best possible bound M_{opt} in the protocol, we find that

$$\mathbb{H}[\mathbf{x} | \Pi] = \mathbb{E}_{\mathbf{y} \sim P_{\mathbf{y}}}[\mathbb{H}[N | \mathbf{y}]] \geq \mathbb{E}_{\mathbf{y} \sim P_{\mathbf{y}}}[D_\infty[P_{\mathbf{x}|\mathbf{y}}\|P_{\mathbf{x}}]] \stackrel{(a)}{\geq} I[\mathbf{x}; \mathbf{y}],$$

162 see Appendix I for the derivation. Unfortunately, inequality (a) can be *arbitrarily loose*, hence the
 163 average codelength scales with the expected ∞ -divergence instead of $I[\mathbf{x}; \mathbf{y}]$, as would be optimal.

164 **Problem 2: Slow runtime:** We are interested in classifying the time complexity of our protocol. As
 165 we saw, for a target Q and proposal P , Algorithm 2 draws $M_{opt} = \exp_2(D_\infty[Q\|P])$ samples on
 166 average. Unfortunately, under the computational hardness assumption $\text{RP} \neq \text{NP}$, Agustsson & Theis
 167 (2020) showed that without any further assumptions, there is no sampler that scales polynomially in
 168 $D_{\text{KL}}[Q\|P]$. However, with further assumptions, we can do much better, as we show in Section 3.1.

169 3 Greedy Poisson Rejection Sampling

170 We now describe GPRS; its pseudo-code is shown in Algorithm 3. This section assumes that Q and
 171 P are the target and proposal distributions, respectively, and $r = dQ/dP$ is their density ratio. Let Π
 172 be a $(1, P)$ -Poisson process. Our proposed rejection criterion is now embarrassingly simple: for an
 173 appropriately defined invertible function $\sigma: \mathbb{R}^+ \rightarrow \mathbb{R}^+$, accept the first arrival of Π that falls under
 174 the graph of the composite function $\varphi = \sigma \circ r$, as illustrated in the left plot in Figure 1. We refer to σ
 175 as the *stretch function* for r , as its purpose is to stretch the density ratio along the time-axis of Π .

176 **Deriving the stretch function (sketch):** Let $\varphi = \sigma \circ r$, where for now σ is an arbitrary invertible
 177 function on \mathbb{R}^+ , let $U = \{(t, x) \in \Omega \mid t \leq \varphi(x)\}$ be the set of points under the graph of φ and let
 178 $\tilde{\Pi} = \Pi \cap U$. By the restriction theorem (Kingman, 1992), $\tilde{\Pi}$ is also a Poisson process with mean
 179 measure $\tilde{\mu}(A) = \mu(A \cap U)$. Let (\tilde{T}, \tilde{X}) be the first arrival of $\tilde{\Pi}$, i.e. the first arrival of Π under φ and
 180 let Q_σ be the distribution of \tilde{X} . Then, the density ratio dQ_σ/dP obeys the identity (see Appendix A):

$$\frac{dQ_\sigma}{dP}(x) = \int_0^{\varphi(x)} \mathbb{P}[\tilde{T} \geq t] dt. \quad (1)$$

181 Now we pick σ such that $Q_\sigma = Q$, for which we need to ensure that $dQ_\sigma/dP = r$. Substituting
 182 $\tau = \varphi(x)$ into Equation (1), and differentiating, we get $(\sigma^{-1})'(\tau) = \mathbb{P}[\tilde{T} \geq t]$. Since (\tilde{T}, \tilde{X}) falls
 183 under the graph of φ by definition, we find that $\mathbb{P}[\tilde{T} \geq t] = \mathbb{P}[\tilde{T} \geq t, r(\tilde{X}) \geq \sigma^{-1}(t)]$. By expanding
 184 the definition of the right-hand side, in Appendix A we obtain a time-invariant ODE for σ^{-1} :

$$(\sigma^{-1})'(\tau) = w_Q(\sigma^{-1}(t)) - \sigma^{-1}(t) \cdot w_P(\sigma^{-1}(t)), \quad \text{with } \sigma^{-1}(0) = 0 \quad (2)$$

185 where we define $w_P(h) \stackrel{\text{def}}{=} \mathbb{P}_{Z \sim P}[r(Z) \geq h]$ and define w_Q analogously. In Appendix G, we
 186 provide the analytic form of w_P and w_Q for discrete, uniform, triangular, Gaussian, and Laplace
 187 distributions. Finally, we use the inverse function theorem and integrate both sides to obtain

$$\sigma(h) = \int_0^h \frac{1}{w_Q(\eta) - \eta \cdot w_P(\eta)} d\eta. \quad (3)$$

188 Remember that picking σ according to Equation (3) ensures that GPRS is **correct by construction**. To
 189 complete the picture, in Appendix A we prove that (\tilde{T}, \tilde{X}) always exists and Algorithm 3 terminates
 190 with probability 1. Unfortunately, computing the integral in Equation (3) analytically is usually not
 191 possible. Moreover, we can show that solving it numerically is unstable as σ is unbounded. Instead,
 192 we numerically solve for σ^{-1} using Equation (2) in practice, which fortunately turns out to be stable.

193 We now turn our attention to analyzing the runtime of Algorithm 3 and find the following surprising
 194 result: the expected runtime of GPRS matches that of standard rejection sampling.

195 **Theorem 3.1** (Expected Runtime). *Let Q and P be the target and proposal distributions for Al-*
 196 *gorithm 3, respectively, and $r = dQ/dP$ their density ratio. Let N denote the number of samples*
 197 *simulated by the algorithm before it terminates. Then,*

$$\mathbb{E}[N] = \exp_2(D_\infty[Q\|P]) \quad \text{and} \quad \mathbb{V}[N] \geq \exp_2(D_\infty[Q\|P]). \quad (4)$$

Algorithm 4: Parallel GPRS with J available threads.

Input : Proposal distribution P ,
Density ratio $r = dQ/dP$,
Stretch function σ ,
Number of parallel threads J .

Output : Sample $X \sim Q$ and its code (j^*, N_{j^*}) .
 $T^*, X^*, j^*, N_{j^*} \leftarrow \infty, \text{nil}, \text{nil}, \text{nil}$

in parallel for $j = 1, \dots, J$ **do**
 $\Pi_j \leftarrow \text{SimulatePP}(1/J, P)$
 for $n_j = 1, 2, \dots$ **do**
 $(T_{n_j}^{(j)}, X_{n_j}^{(j)}) \leftarrow \text{next}(\Pi_j)$
 if $T^* < T_{n_j}^{(j)}$ **then**
 terminate thread j .
 end
 if $T_{n_j}^{(j)} < \sigma(r(X_{n_j}^{(j)}))$ **then**
 $T^*, X^*, j^*, N_{j^*} \leftarrow T_{n_j}^{(j)}, X_{n_j}^{(j)}, j, n_j$
 terminate thread j .
 end
 end
end
return $X^*, (j^*, N_{j^*})$

Algorithm 5: Branch-and-bound GPRS on \mathbb{R} with unimodal r

Input : Proposal distribution P ,
Density ratio $r = dQ/dP$,
Stretch function σ ,
Location x^* of the mode of r .

Output : Sample $X \sim Q$ and its heap index H .
 $T_0, H, B \leftarrow (0, 1, \mathbb{R})$

for $d = 1, 2, \dots$ **do**
 $X_d \sim P|_B$
 $\Delta_d \sim \text{Exp}(P(B))$
 $T_d \leftarrow T_{d-1} + \Delta_d$
 if $T_d < \sigma(r(X_d))$ **then**
 return X_d, H
 end
 if $X_d \geq x^*$ **then**
 $B \leftarrow B \cap (-\infty, X_d)$
 $H \leftarrow 2H$
 else
 $B \leftarrow B \cap (X_d, \infty)$
 $H \leftarrow 2H + 1$
 end
end

198 GPRS induces a channel simulation protocol similar to the standard rejection sampling-based one in
199 Section 2.2: The encoder simulates the arrivals of Π using shared randomness and encodes the index
200 of their accepted sample. The next theorem shows that this protocol is optimally efficient.

201 **Theorem 3.2** (Expected Codelength). *Let $P_{\mathbf{x}, \mathbf{y}}$ be a joint distribution over correlated random*
202 *variables \mathbf{x} and \mathbf{y} , and let Π be the $(1, P)$ -Poisson process used by Algorithm 3. Then, the algorithm*
203 *induces a channel simulation protocol, such that*

$$\mathbb{H}[\mathbf{x} | \Pi] \leq I[\mathbf{x}; \mathbf{y}] + \log_2(I[\mathbf{x}; \mathbf{y}] + 1) + 6. \quad (5)$$

204 See Appendix B.2 and Appendix B.3 for proofs. This result is analogous to the Poisson functional
205 representation (Li & El Gamal, 2018), and thus it can be used to provide an alternative proof of the
206 strong functional representation lemma (Theorem 1; Li & El Gamal, 2018).

207 **GPRS as greedy search:** We contrast GPRS with the well-known A* sampling algorithm (Maddison
208 et al., 2014), which also searches through the points of Π but uses a different criterion. The defining
209 feature of GPRS is that its acceptance criterion at each step is *local*, since if at step n the arrival
210 (T_n, X_n) in Π falls under φ , we will immediately accept it. Thus it is *greedy* search procedure.
211 On the other hand, the A* acceptance criterion is *global*, as the acceptance of a particular arrival
212 (T_n, X_n) depends on all other points of Π in the general case. Surprisingly, this difference between
213 the search criteria does not make a difference in the average runtimes and codelengths in the general
214 case. However, as shown in the next section, GPRS can be much faster in special cases.

215 3.1 Speeding up the greedy search

216 This section discusses two ways to improve the runtime of GPRS. First, we show how we can utilize
217 available parallel computing power to speed up Algorithm 3. Second, we propose an advanced
218 search strategy when the spatial domain Ω has more structure and show that we can obtain a **super-**
219 **exponential improvement** in the runtime from $\exp_2(D_\infty[Q||P])$ to $\mathcal{O}(D_{\text{KL}}[Q||P])$ in certain cases.

220 **Parallel GPRS:** The basis for parallelizing GPRS is the *superposition theorem* (Kingman, 1992):
221 Let Π_1, \dots, Π_J all be $(1/J, P)$ -Poisson processes; then, $\Pi = \bigcup_{j=1}^J \Pi_j$ is a $(1, P)$ -Poisson process.
222 This result makes parallelizing GPRS very simple, as shown in Algorithm 4, assuming we have J
223 parallel threads: First, we independently look for the first arrivals of Π_1, \dots, Π_J under the graph
224 of φ , yielding $(T_{N_1}^{(1)}, X_{N_1}^{(1)}), \dots, (T_{N_J}^{(J)}, X_{N_J}^{(J)})$, respectively, where the N_j corresponds to the

225 index of Π_j 's first arrival under φ . Then, we select the candidate with the earliest arrival time, i.e.
 226 $j^* = \arg \min_{j \in \{1, \dots, J\}} T_{N_j}^{(j)}$. Now, by the superposition theorem, $(T_{N_{j^*}}^{(j^*)}, X_{N_{j^*}}^{(j^*)})$ is the first arrival
 227 of Π under φ , and hence $X_{N_{j^*}}^{(j^*)} \sim Q$. Finally, Alice encodes the sample via the tuple (j^*, N_{j^*}) , i.e.
 228 which of the J processes the first arrival occurred in, and the index of the arrival in Π_{j^*} . See the
 229 middle plot in Figure 1 for an example case with $J = 2$.

230 Our next result shows that parallelizing GPRS results in a linear reduction in both the expectation
 231 and variance of its runtime and a more favourable codelength guarantee for channel simulation.

232 **Theorem 3.3** (Expected runtime of parallelized GPRS). *Let Q, P and r be defined as above, and*
 233 *let ν_j denote the random variable corresponding to the number of samples simulated by thread j in*
 234 *Algorithm 4 using J threads. Then, for all j ,*

$$\mathbb{E}[\nu_j] = (\exp_2(D_\infty[Q\|P]) - 1) / J + 1 \quad \text{and} \quad \mathbb{V}[\nu_j] \geq (\exp_2(D_\infty[Q\|P]) - 1) / J + 1. \quad (6)$$

235 **Theorem 3.4** (Expected codelength of parallelized GPRS). *Let $P_{\mathbf{x}, \mathbf{y}}$ be a joint distribution over*
 236 *correlated random variables \mathbf{x} and \mathbf{y} , and let Π_1, \dots, Π_J be $(1/J, P_{\mathbf{x}})$ -Poisson processes. Then,*
 237 *assuming $\log_2 J \leq I[\mathbf{x}; \mathbf{y}]$, parallelized GPRS induces a channel simulation protocol such that*

$$\mathbb{H}[\mathbf{x} \mid \Pi_1, \dots, \Pi_J] \leq I[\mathbf{x}; \mathbf{y}] + \log_2(I[\mathbf{x}; \mathbf{y}] - \log_2 J + 1) + 8. \quad (7)$$

238 See Appendix C for the proofs. Note that we can use the same parallelisation argument with the
 239 appropriate modifications to speed up A^* sampling / coding too.

240 **Branch-and-bound GPRS on \mathbb{R} :** We briefly restrict our attention to problems on \mathbb{R} when the density
 241 ratio r is unimodal, as we can exploit this additional structure to more efficiently search for Π 's first
 242 arrival under φ . Consider the example in the right plot in Figure 1: we simulate the first arrival
 243 (T_1, X_1) , and reject it, since $T_1 > \varphi(X_1)$. Since φ is unimodal by assumption, for $x < X_1$ we
 244 must have $\varphi(x) < \varphi(X_1)$, while the arrival time of any of the later arrivals will be larger than T_1 .
 245 Therefore, none of the arrivals to the left of X_1 will fall under φ either! Hence, it is enough to simulate
 246 Π to the right of X_1 . Repeating this argument for later arrivals, we obtain an efficient binary search
 247 procedure to find Π 's first arrival under φ , described in Algorithm 5: We simulate the next arrival
 248 (T_B, X_B) of Π within some bounds B , starting with the first arrival in $B = \Omega$. If $T_B \leq \varphi(X_B)$ we
 249 accept; otherwise, we truncate the bound to $B \leftarrow B \cap (-\infty, X_B)$, or $B \leftarrow B \cap (X_B, \infty)$ based on
 250 where X_B falls relative to φ 's mode. We repeat these two steps until we find the first arrival.

251 Since Algorithm 5 does not simulate every point of Π , we cannot use the index N of the accepted
 252 arrival to obtain a channel simulation protocol as before. Instead, we encode the *search path*, i.e.
 253 whether we chose the left or the right side of our current sample at each step. Similarly to A^* coding,
 254 we encode the path using its *heap index* (Flamich et al., 2022): the root has index $H_{root} = 1$, and for
 255 a node with index H , its left child is assigned index $2H$ and its right child $2H + 1$. As the following
 256 theorems show, this version of GPRS is, in fact, optimally efficient.

257 **Theorem 3.5** (Expected Runtime of GPRS with binary search). *Let Q, P and r be defined as above,*
 258 *and let D denote the number of samples simulated by Algorithm 5. Then,*

$$\mathbb{E}[D] = D_{\text{KL}}[Q\|P] + 3. \quad (8)$$

259 **Theorem 3.6** (Expected Codelength of GPRS with binary search). *Let $P_{\mathbf{x}, \mathbf{y}}$ be a joint distribution*
 260 *over correlated random variables \mathbf{x} and \mathbf{y} , and let Π be a $(1, P_{\mathbf{x}})$ -Poisson process. Then, GPRS with*
 261 *binary search induces a channel simulation protocol such that*

$$\mathbb{H}[\mathbf{x} \mid \Pi] \leq I[\mathbf{x}; \mathbf{y}] + \log_2(I[\mathbf{x}; \mathbf{y}] + 1) + 8 \quad (9)$$

262 See Appendix E for details and proofs of the theorems.

263 **Branch-and-bound GPRS with splitting functions:** With some additional machinery, Algorithm 5
 264 can be extended to much more general settings, such as \mathbb{R}^D and cases where r is not unimodal, by
 265 introducing the notion of a *splitting function*. For a region of space, $B \subseteq \Omega$, a splitting function
 266 `split` simply returns a binary partition of the set, i.e. $\{L, R\} = \text{split}(B)$, such that $L \cap R = \emptyset$
 267 and $L \cup R = B$. In this case, we can perform a similar tree search to Algorithm 5, captured in
 268 Algorithm 6. Starting with the whole space $B = \Omega$, we simulate the next arrival (T, X) of Π in B
 269 at each step. If we reject it, we partition B into two parts, L and R , using `split`. Then, with probability
 270 $\mathbb{P}[\tilde{X} \in R \mid \tilde{X} \in B, \tilde{T} \geq T]$, we continue searching through the arrivals of Π in R only, and in L only

Algorithm 6: Branch-and-bound GPRS with splitting function.

Input : Proposal distribution P ,
Density ratio $r = dQ/dP$,
Stretch function σ ,
Splitting function split .

Output : Sample $X \sim Q$ and its
heap index H

$T_0, H, B \leftarrow (0, 1, \mathbb{R})$

for $d = 1, 2, \dots$ **do**

$X_d \sim P|_B$

$\Delta_d \sim \text{Exp}(P(B))$

$T_d \leftarrow T_{d-1} + \Delta_d$

if $T_d < \sigma(r(X_d))$ **then**

return X_d, H

else

$B_0, B_1 \leftarrow \text{split}(B)$

$\rho \leftarrow$

$\mathbb{P}[\tilde{X} \in B_1 \mid \tilde{X} \in B, \tilde{T} \geq T_d]$

$\beta \leftarrow \text{Bernoulli}(\rho)$

$H \leftarrow 2H + \beta$

$B \leftarrow B_\beta$

end

end

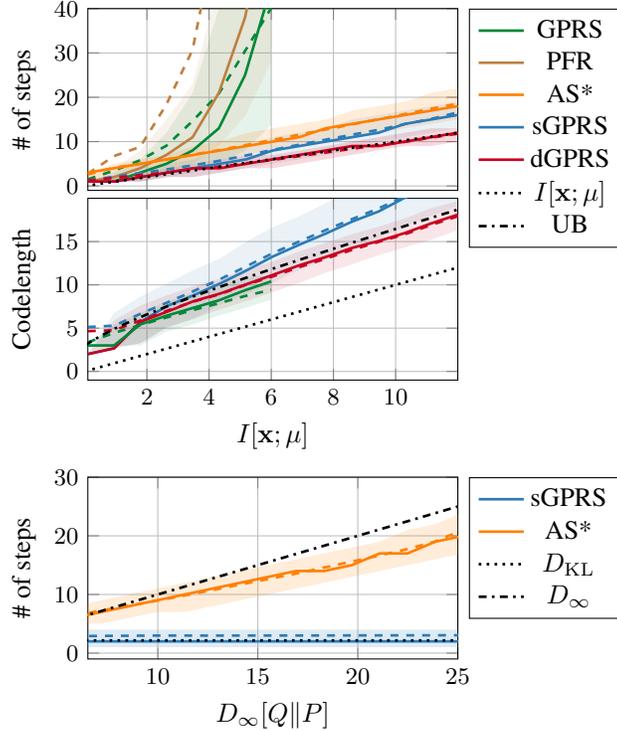


Figure 2: **Left:** Binary search GPRS with arbitrary splitting function. **Right:** Performance comparison of different channel simulation protocols. In each plot, *dashed lines* indicate the mean, *solid lines* the median and the *shaded areas* the 25 - 75 percentile region of the relevant performance metric. We computed the statistics over 1000 runs for each setting. **Top right:** Runtime comparison on a 1D Gaussian channel simulation problem $P_{\mathbf{x}, \mu}$, plotted against increasing mutual information $I[\mathbf{x}; \mu]$. Alice receives $\mu \sim \mathcal{N}(0, \sigma^2)$ and encodes a sample $\mathbf{x} \mid \mu \sim \mathcal{N}(\mu, 1)$ to Bob. The abbreviations in the legend are: *GPRS* – Algorithm 3; *PFR* – Poisson functional representation / Global-bound A* coding (Li & El Gamal, 2018; Flamich et al., 2022), *AS** – Split-on-sample A* coding (Flamich et al., 2022); *sGPRS* – split-on-sample GPRS (Algorithm 5); *dGPRS* – GPRS with dyadic *split* (Algorithm 6). **Middle right:** Average code length comparison of our proposed algorithms on the same channel simulation problem as above. *UB* in the legend corresponds to an *upper bound* of $I[\mathbf{x}; \mu] + \log_2(I[\mathbf{x}; \mu] + 1) + 2$ bits. We estimate the algorithms’ expected code lengths by encoding the indices returned by GPRS using a Zeta distribution $\zeta(n \mid \lambda) \propto n^{-\lambda}$ with $\lambda = 1 + 1/I[\mathbf{x}; \mathbf{y}]$ in each case, which is the optimal maximum entropy distribution for this problem setting (Li & El Gamal, 2018). **Bottom right:** One-shot runtime comparison of *sGPRS* with *AS** coding. Alice encodes samples from a target $Q = \mathcal{N}(m, s^2)$ using $P = \mathcal{N}(0, 1)$ as the proposal. We computed m and s^2 such that $D_{\text{KL}}[Q||P] = 2$ bits for each problem, but $D_\infty[Q||P]$ increases. *GPRS*’ runtime stays fixed as it scales with $D_{\text{KL}}[Q||P]$, while the runtime of *A** keeps increasing.

271 otherwise. We show the correctness of this procedure in Appendix F, where we also derive how the
272 quantities $\mathbb{P}[\tilde{X} \in R \mid \tilde{X} \in B, \tilde{T} \geq T]$ and σ can be computed in practice. This splitting procedure is
273 analogous to the general version of *A** sampling/coding (Maddison et al., 2014; Flamich et al., 2022),
274 which is parameterized by the same splitting functions. Note that Algorithm 5 is a special case of
275 Algorithm 6, where the underlying space is $\Omega = \mathbb{R}$, r is unimodal, and at each step for an interval
276 bound $B = (a, b)$ and sample $X \in (a, b)$ we split B into $\{(a, X), (X, b)\}$.

277 4 Experiments

278 We compare the average and one-shot case efficiency of our proposed variants of *GPRS* and a couple
279 of other channel simulation protocols in Figure 2. See the figure’s caption and Appendix H for details
280 of our experimental setup. The top two plots in Figure 2 demonstrate that our methods’ expected
281 runtimes and code lengths align very well with our theorems’ predictions and compare favourably to
282 other methods. Furthermore, we find that the mean performance is a *robust statistic* for the binary

283 search-based variants of GPRS in that it lines up closely with the median, and the interquartile range
 284 is quite narrow. On the other hand, the bottom plot in Figure 2 demonstrates the most salient property
 285 of the binary search variant of GPRS: unlike all previous methods, its runtime scales with $D_{\text{KL}}[Q\|P]$
 286 and not $D_{\infty}[Q\|P]$. Thus, we can apply it to a larger family of channel simulation problems both
 287 in theory and practice where $D_{\text{KL}}[Q\|P]$ is finite, but $D_{\infty}[Q\|P]$ is very large or even infinite in
 288 expectation, and other methods would not terminate.

289 5 Related Work

290 **Poisson processes for channel simulation:** Poisson processes were introduced to the channel
 291 simulation literature by Li & El Gamal (2018) via their construction of the *Poisson functional*
 292 *representation* (PFR) in their proof of the Strong Functional Representation Lemma. Flamich et al.
 293 (2022) observed that the PFR construction is equivalent to a certain variant of A* sampling (Maddison
 294 et al., 2014; Maddison, 2016). Thus, they proposed an optimized version of the PFR called A* coding,
 295 which achieves $\mathcal{O}(D_{\infty}[Q\|P])$ runtime for one-dimensional unimodal distributions. GPRS was
 296 mainly inspired by A* coding, and they are *dual* constructions to each other in the following sense:
 297 *Depth-limited A* coding* can be thought of as an *importance sampler*, i.e. a Monte Carlo algorithm
 298 that returns an approximate sample in fixed time. On the other hand, GPRS is a *rejection sampler*, i.e.
 299 a Las Vegas algorithm that returns an exact sample in random time.

300 **Channel simulation with dithered quantization:** *Dithered quantization* (DQ; Ziv, 1985) is an
 301 alternative to rejection and importance sampling-based approaches to channel simulation. DQ exploits
 302 that for any $c \in \mathbb{R}$ and $U, U' \sim \text{Unif}(-1/2, 1/2)$, the quantities $\lfloor c - U \rfloor + U$ and $c + U'$ are equal in
 303 distribution. While DQ has been around for decades as a tool to model and analyze quantization error,
 304 Agustsson & Theis (2020) reinterpreted it as a channel simulation protocol and used it to develop a
 305 VAE-based neural image compression algorithm. Unfortunately, basic DQ only allows uniform target
 306 distributions, limiting its applicability. As a partial remedy, Theis & Yosri (2022) showed DQ could
 307 be combined with other channel simulation protocols to speed them up and thus called their approach
 308 hybrid coding (HQ). Originally, HQ required that the target distribution be compactly supported,
 309 which was lifted by Flamich & Theis (2023), who developed an adaptive rejection sampler using HQ.
 310 In a different vein, Hegazy & Li (2022) generalize and analyze a method proposed in the appendix of
 311 Agustsson & Theis (2020) and show that DQ can be used to realize channel simulation protocols for
 312 one-dimensional symmetric, unimodal distributions.

313 **Greedy Rejection Coding:** Concurrently, Anonymous (2023) generalize Harsha et al. (2007)’s
 314 rejection sampling algorithm to arbitrary probability spaces, which they call greedy rejection coding
 315 (GRC). Furthermore, they also introduce a space-partitioning procedure to speed up the convergence
 316 of their sampler and prove that a variant of their sampler also achieves optimal runtime for one-
 317 dimensional problems where dQ/dP is unimodal. However, the construction of their method differs
 318 significantly from ours. GRC is a direct generalization of Harsha et al. (2007)’s algorithm and, thus,
 319 a more “conventional” rejection sampler, while we base our construction on Poisson processes. Thus,
 320 our proof techniques are also significantly different. It is an interesting research question whether
 321 GRC could be formulated using Poisson processes, akin to the formulation of standard rejection
 322 sampling in Algorithm 2, as this could be used to connect the two algorithms and improve both.

323 6 Discussion and Future Work

324 Using Poisson processes, we constructed greedy Poisson rejection sampling. We proved the correct-
 325 ness of the algorithm and analyzed its runtime, and showed that it could be used to obtain a channel
 326 simulation protocol. We then developed several variations on it, analyzed their runtimes, and showed
 327 that they could all be used to obtain channel simulation protocols. As the most significant result of
 328 the paper, we showed that using the binary search variant of GPRS we can achieve $\mathcal{O}(D_{\text{KL}}[Q\|P])$
 329 runtime for arbitrary one-dimensional, unimodal density ratios, significantly improving upon the
 330 previous best $\mathcal{O}(D_{\infty}[Q\|P])$ bound by A* coding.

331 There are several interesting directions for future work. From a practical perspective, the most
 332 pressing question is whether efficient channel simulation algorithms exist for multivariate problems;
 333 finding an efficient channel simulation protocol for multivariate Gaussians would already have
 334 far-reaching practical consequences.

335 References

- 336 Eirikur Agustsson and Lucas Theis. Universally quantized neural compression. *Advances in Neural*
337 *Information Processing Systems*, 33, 2020.
- 338 Anonymous. Faster relative entropy coding with greedy rejection coding. *See the Supplementary*
339 *Materials*, 2023.
- 340 Robert M Corless, Gaston H Gonnet, David EG Hare, David J Jeffrey, and Donald E Knuth. On the
341 lambert w function. *Advances in Computational mathematics*, 5:329–359, 1996.
- 342 Gergely Flamich and Lucas Theis. Adaptive greedy rejection sampling. *arXiv preprint*
343 *arXiv:2304.10407*, 2023.
- 344 Gergely Flamich, Marton Havasi, and José Miguel Hernández-Lobato. Compressing images by
345 encoding their latent representations with relative entropy coding. *Advances in Neural Information*
346 *Processing Systems*, 33, 2020.
- 347 Gergely Flamich, Stratis Markou, and José Miguel Hernández-Lobato. Fast relative entropy coding
348 with A* coding. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu,
349 and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*,
350 volume 162 of *Proceedings of Machine Learning Research*, pp. 6548–6577. PMLR, 17–23 Jul
351 2022. URL <https://proceedings.mlr.press/v162/flamich22a.html>.
- 352 Prahladh Harsha, Rahul Jain, David McAllester, and Jaikumar Radhakrishnan. The communica-
353 tion complexity of correlation. In *Twenty-Second Annual IEEE Conference on Computational*
354 *Complexity (CCC'07)*, pp. 10–23. IEEE, 2007.
- 355 Marton Havasi, Robert Peharz, and José Miguel Hernández-Lobato. Minimal random code learning:
356 Getting bits back from compressed model parameters. In *International Conference on Learning*
357 *Representations*, 2018.
- 358 Mahmoud Hegazy and Cheuk Ting Li. Randomized quantization with exact error distribution. In
359 *2022 IEEE Information Theory Workshop (ITW)*, pp. 350–355. IEEE, 2022.
- 360 Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *International Conference*
361 *on Learning Representations*, 2014. URL <https://arxiv.org/abs/1312.6114>.
- 362 J.F.C. Kingman. *Poisson Processes*. Oxford Studies in Probability. Clarendon Press, 1992. ISBN
363 9780191591242. URL <https://books.google.co.uk/books?id=VEiM-0twDHkC>.
- 364 Cheuk Ting Li and Abbas El Gamal. Strong functional representation lemma and applications to
365 coding theorems. *IEEE Transactions on Information Theory*, 64(11):6967–6978, 2018.
- 366 CA Maddison. Poisson process model for Monte Carlo. *Perturbation, Optimization, and Statistics*,
367 pp. 193–232, 2016.
- 368 Chris J Maddison, Daniel Tarlow, and Tom Minka. A* sampling. *Advances in Neural Information*
369 *Processing Systems*, 27:3086–3094, 2014.
- 370 Pat Muldowney, Krzysztof Ostaszewski, and Wojciech Wojdowski. The Darth Vader rule. *Tatra*
371 *Mountains Mathematical Publications*, 52(1):53–63, 2012.
- 372 A. Shah, W.-N. Chen, J. Balle, P. Kairouz, and L. Theis. Optimal compression of locally differentially
373 private mechanisms. In *Artificial Intelligence and Statistics*, 2022. URL <https://arxiv.org/abs/2111.00092>.
- 374
- 375 L. Theis, T. Salimans, M. D. Hoffman, and F. Mentzer. Lossy compression with Gaussian diffusion.
376 *arXiv:2206.08889*, 2022. URL <https://arxiv.org/abs/2206.08889>.
- 377 Lucas Theis and Noureldin Yosri. Algorithms for the communication of samples. In *International*
378 *Conference on Machine Learning*, 2022.
- 379 Jacob Ziv. On universal quantization. *IEEE Transactions on Information Theory*, 31(3):344–347,
380 1985.

381 A Measure-theoretic Construction of Greedy Poisson Rejection Sampling

382 In this section, we provide a construction of greedy Poisson rejection sampling (GPRS) in its greatest
383 generality. However, we first establish the notation for the rest of the appendix.

384 **Notation:** In what follows, we will always denote GPRS’s target distribution as Q and its proposal
385 distribution as P . We assume that Q and P are Borel probability measures over some Polish space Ω
386 with Radon-Nikodym derivative $r \stackrel{\text{def}}{=} dQ/dP$. We denote the standard Lebesgue measure on \mathbb{R}^n
387 by λ . All logarithms are assumed to be to the base 2 denoted by \log_2 . Similarly, we use the less
388 common notation of \exp_2 to denote exponentiation with base 2, i.e. $\exp_2(x) = 2^x$. The relative
389 entropy from Q to P is defined as $D_{\text{KL}}[Q\|P] \stackrel{\text{def}}{=} \int_{\Omega} \log_2 r(x) dP(x)$ and the Rényi ∞ -divergence
390 as $D_{\infty}[Q\|P] \stackrel{\text{def}}{=} \text{ess sup}_{x \in \Omega} \{\log r(x)\}$, where the essential supremum is taken with respect to P .

391 **Restricting a Poisson process “under the graph of a function”:** We first consider the class of
392 functions “under which” we will be restricting our Poisson processes. Let $\varphi : \Omega \rightarrow \mathbb{R}^+$ be a
393 measurable function, such that $|\varphi(\Omega)| = |r(\Omega)|$, i.e. the image of φ has the same cardinality as
394 the image of $r = dQ/dP$. This implies that there exist bijections between $r(\Omega)$ and $\varphi(\Omega)$. Since
395 both images are subsets of \mathbb{R}^+ , both images have a linear order. Hence, a monotonically increasing
396 bijection $\tilde{\sigma}$ exists such that $\varphi = \tilde{\sigma} \circ r$. Since $\tilde{\sigma}$ is monotonically increasing, we can extend it to a
397 monotonically increasing continuous function $\sigma : \mathbb{R}^+ \rightarrow \mathbb{R}^+$. This fact is significant, as it allows us
398 to reduce questions about functions of interest from Ω to \mathbb{R}^+ to invertible functions on the positive
399 reals. Since σ is continuous and monotonically increasing on the positive reals, it is invertible on the
400 extended domain. We now consider restricting a Poisson process under the graph of φ , and work out
401 the distribution of the spatial coordinate of the first arrival.

402 Let $\Pi = \{(T_n, X_n)\}_{n=1}^{\infty}$ be a Poisson process over $\mathbb{R}^+ \times \Omega$ with mean measure $\mu = \lambda \times P$. Let
403 $U \stackrel{\text{def}}{=} \{(t, x) \in \Omega \mid t \leq \varphi(x)\}$ be the set of points “under the graph” of φ . By the restriction theorem
404 (Kingman, 1992), $\tilde{\Pi} \stackrel{\text{def}}{=} \Pi \cap U$ is a Poisson process with mean measure $\tilde{\mu}(A) = \mu(A \cap U)$ for an
405 arbitrary Borel set A . As a slight abuse of notation, we define the shorthand $\tilde{\mu}(t) = \tilde{\mu}([0, t) \times \Omega)$
406 for the average number of points in $\tilde{\Pi}$ up to time t . Let $\mathbf{N}_{\tilde{\Pi}}$ denote the counting measure of $\tilde{\Pi}$ and
407 let (\tilde{T}, \tilde{X}) be the first arrival of $\tilde{\Pi}$, i.e. the first point of the original process Π that falls under the
408 graph of φ , assuming that it exists. To develop GPRS, we first work out the distribution of \tilde{X} for an
409 arbitrary φ .

410 **Remark:** Note, that the construction below only holds if the first arrival is guaranteed to exist,
411 otherwise the quantities below do not make sense. From a formal point of view, we would have
412 to always condition on the event $\mathbf{N}_{\tilde{\Pi}}(\infty) \stackrel{\text{def}}{=} \lim_{t \rightarrow \infty} \mathbf{N}_{\tilde{\Pi}}(t) > 0$, e.g. we would need to write
413 $\mathbb{P}[\tilde{T} \geq t \mid \mathbf{N}_{\tilde{\Pi}}(\infty) > 0]$. However, we make a “leap of faith” instead and assume that our
414 construction is valid to obtain our “guesses” for the functions σ and φ . Then, we show that for our
415 specific guesses the derivation below is well-defined, i.e. namely that $\mathbb{P}[\mathbf{N}_{\tilde{\Pi}}(\infty) > 0] = 1$, which
416 then implies $\mathbb{P}[\tilde{T} \geq t \mid \mathbf{N}_{\tilde{\Pi}}(\infty) > 0] = \mathbb{P}[\tilde{T} \geq t]$. To do this, note that

$$\mathbb{P}[\mathbf{N}_{\tilde{\Pi}}(\infty) > 0] = 1 - \mathbb{P}[\mathbf{N}_{\tilde{\Pi}}(\infty) = 0] \tag{10}$$

$$= 1 - \lim_{t \rightarrow \infty} e^{-\tilde{\mu}(t)}. \tag{11}$$

417 Hence, for the well-definition of our construction it is enough to show that for our choice of σ ,
418 $\tilde{\mu}(t) \rightarrow \infty$ as $t \rightarrow \infty$.

419 **Constructing σ :** We wish to derive

$$\frac{d\mathbb{P}[\tilde{X} = x]}{dP} = \int_0^{\infty} \frac{d\mathbb{P}[\tilde{T} = t, \tilde{X} = x]}{d(\lambda \times P)} dt. \tag{12}$$

420 To do this, recall that we defined $\mathbf{N}_{\tilde{\Pi}}(t) \stackrel{def}{=} \mathbf{N}_{\tilde{\Pi}}([0, t] \times \Omega)$ and define $\mathbf{N}_{\tilde{\Pi}}(s, t) \stackrel{def}{=} \mathbf{N}_{\tilde{\Pi}}([s, t] \times \Omega)$.

421 Similarly, let $\tilde{\mu}(t) \stackrel{def}{=} \mathbb{E}[\mathbf{N}_{\tilde{\Pi}}(t)]$ and $\tilde{\mu}(s, t) \stackrel{def}{=} \mathbb{E}[\mathbf{N}_{\tilde{\Pi}}(s, t)] = \tilde{\mu}(t) - \tilde{\mu}(s)$. Then,

$$\mathbb{P}[\tilde{X} \in A, \tilde{T} \in dt] = \lim_{s \rightarrow t} \mathbb{P}[\tilde{X} \in A, \tilde{T} \in [s, t]] \quad (13)$$

$$= \lim_{s \rightarrow t} \mathbb{P}[\tilde{X} \in A, \mathbf{N}_{\tilde{\Pi}}(s, t) = 1, \mathbf{N}_{\tilde{\Pi}}(s) = 0] \quad (14)$$

$$= \lim_{s \rightarrow t} \mathbb{P}[\tilde{X} \in A \mid \mathbf{N}_{\tilde{\Pi}}(s, t) = 1] \mathbb{P}[\mathbf{N}_{\tilde{\Pi}}(s, t) = 1] \mathbb{P}[\mathbf{N}_{\tilde{\Pi}}(s) = 0], \quad (15)$$

422 where the last equality holds, because $\mathbf{N}_{\tilde{\Pi}}(s) \perp \mathbf{N}_{\tilde{\Pi}}(s, t)$, since $[0, s] \times \Omega$ and $[s, t] \times \Omega$ are disjoint.

423 By Lemma 3 from Maddison (2016),

$$\mathbb{P}[\tilde{X} \in A \mid \mathbf{N}_{\tilde{\Pi}}(s, t) = 1] = \mathbb{P}[(\tilde{T}, \tilde{X}) \in [s, t] \times A \mid \mathbf{N}_{\tilde{\Pi}}(s, t) = 1] \quad (16)$$

$$= \frac{\tilde{\mu}([s, t] \times A)}{\tilde{\mu}(s, t)}. \quad (17)$$

424 Substituting this back into Equation (15) and using the fact that $\mathbf{N}_{\tilde{\Pi}}$ is Poisson, we find

$$\mathbb{P}[\tilde{X} \in A, \tilde{T} \in dt] = \lim_{s \rightarrow t} \left(\frac{\tilde{\mu}([s, t] \times A)}{\tilde{\mu}(s, t)} \cdot \tilde{\mu}(s, t) e^{-\tilde{\mu}(s, t)} \cdot e^{-\tilde{\mu}(s)} \right) / (t - s) \quad (18)$$

$$= \lim_{s \rightarrow t} \left(\tilde{\mu}([s, t] \times A) \cdot e^{-(\tilde{\mu}(t) - \tilde{\mu}(s))} \cdot e^{-\tilde{\mu}(s)} \right) / (t - s) \quad (19)$$

$$= e^{-\tilde{\mu}(t)} \cdot \lim_{s \rightarrow t} \frac{\tilde{\mu}([s, t] \times A)}{t - s} \quad (20)$$

$$= e^{-\tilde{\mu}(t)} \cdot \int_A \mathbf{1}[t \leq \varphi(x)] dP(x), \quad (21)$$

425 where the last equation holds by the definition of the derivative and the fundamental theorem of
426 calculus. From this, by inspection we find

$$\frac{d\mathbb{P}[\tilde{T} = t, \tilde{X} = x]}{d(\lambda \times P)} = \mathbf{1}[t \leq \varphi(x)] e^{-\tilde{\mu}(t)} = \mathbf{1}[t \leq \varphi(x)] \mathbb{P}[\tilde{T} \geq t]. \quad (22)$$

427 Finally, by marginalizing out the first arrival time, we find that the spatial distribution is

$$\frac{d\mathbb{P}[\tilde{X} = x]}{dP} = \int_0^\infty \frac{d\mathbb{P}[\tilde{T} = t, \tilde{X} = x]}{d(\lambda \times P)} dt \quad (23)$$

$$= \int_0^{\varphi(x)} \mathbb{P}[\tilde{T} \geq t] dt. \quad (24)$$

428 **Deriving φ by finding σ :** Note that Equation (24) holds for any φ . However, to get a correct
429 algorithm, we need to set it such that $\frac{d\mathbb{P}[\tilde{X}=x]}{dP} = \frac{dQ}{dP} = r$. Since we can write $\varphi = \sigma \circ r$ by our
430 earlier argument, this problem is reduced to finding an appropriate invertible continuous function
431 $\sigma : \mathbb{R}^+ \rightarrow \mathbb{R}^+$. Thus, we wish to find σ such that

$$\forall x \in \Omega, \quad r(x) = \int_0^{\sigma(r(x))} \mathbb{P}[\tilde{T} \geq t] dt. \quad (25)$$

432 Now, introduce $\tau = \sigma(r(x))$, so that $\sigma^{-1}(\tau) = r(x)$. Note that this substitution only makes sense for
433 $\tau \in r(\Omega)$. However, since we are free to extend σ^{-1} to \mathbb{R}^+ in any we like so long as it is monotone,
434 we may require that this equation hold for all $\tau \in \mathbb{R}^+$. Then, we find

$$\sigma^{-1}(\tau) = \int_0^\tau \mathbb{P}[\tilde{T} \geq t] dt \quad (26)$$

$$\Rightarrow (\sigma^{-1})'(\tau) = \mathbb{P}[\tilde{T} \geq \tau] \quad (27)$$

435 with $\sigma^{-1}(0) = 0$. Before we solve for σ , we define

$$w_P(h) \stackrel{def}{=} \mathbb{P}_{Y \sim P}[r(Y) \geq h] \quad (28)$$

$$w_Q(h) \stackrel{def}{=} \mathbb{P}_{Y \sim Q}[r(Y) \geq h]. \quad (29)$$

436 Note that w_P and w_Q are supported on $[0, r^*]$, where $r^* \stackrel{\text{def}}{=} \text{ess sup}_{x \in \Omega} \{r(x)\} = \exp_2(D_\infty[Q||P])$.
 437 Now, we can rewrite Equation (27) as

$$(\sigma^{-1})'(\tau) = \mathbb{P}[\tilde{T} \geq t] \quad (30)$$

$$= \mathbb{P}[\tilde{T} \geq t, \varphi(\tilde{X}) \geq t] + \underbrace{\mathbb{P}[\tilde{T} \geq t, \varphi(\tilde{X}) < t]}_{=0 \text{ due to mutual exclusivity}} \quad (31)$$

$$= \underbrace{\mathbb{P}[\tilde{T} \geq 0, \varphi(\tilde{X}) \geq t]}_{=\mathbb{P}[\varphi(\tilde{X}) \geq t], \text{ since } \tilde{T} \geq 0 \text{ always}} - \mathbb{P}[\varphi(\tilde{X}) \geq t \geq \tilde{T}] \quad (32)$$

$$= \mathbb{P}[r(\tilde{X}) \geq \sigma^{-1}(t)] - \mathbb{P}[r(\tilde{X}) \geq \sigma^{-1}(t) \geq \sigma^{-1}(\tilde{T})] \quad (33)$$

$$= w_Q(\sigma^{-1}(t)) - \sigma^{-1}(t)w_P(\sigma^{-1}(t)) \quad (34)$$

438 where the second term in the last equation follows by noting that

$$\mathbb{P}[r(\tilde{X}) \geq \sigma^{-1}(t) \geq \sigma^{-1}(\tilde{T})] = \int_{\Omega} \int_0^t \underbrace{\mathbf{1}[r(x) \geq \sigma^{-1}(t)] \mathbf{1}[r(x) \geq \sigma^{-1}(\tau)]}_{=\mathbf{1}[r(x) \geq \sigma^{-1}(t)], \text{ since } \tau < t} \mathbb{P}[\tilde{T} \geq \tau] d\tau dP(x) \quad (35)$$

$$= \int_{\Omega} \mathbf{1}[r(x) \geq \sigma^{-1}(t)] \underbrace{\int_0^t \mathbb{P}[\tilde{T} \geq \tau] d\tau}_{=\sigma^{-1}(t), \text{ by eq. (26)}} dP(x) \quad (36)$$

$$= \sigma^{-1}(t)w_P(\sigma^{-1}(t)). \quad (37)$$

439 Now, by the inverse function theorem, we get

$$\sigma'(h) = \frac{1}{w_Q(h) - h \cdot w_P(h)}, \quad (38)$$

440 thus we finally find

$$\sigma(h) = \int_0^h \frac{1}{w_Q(\eta) - \eta \cdot w_P(\eta)} d\eta. \quad (39)$$

441 Thus, to recapitulate, setting $\varphi = \sigma \circ r$, where σ is given by Equation (39) will ensure that the spatial
 442 distribution of the first arrival of Π under φ is the target distribution Q .

443 **Well-definition of GPRS and Algorithm 3 terminates with probability 1:** To ensure that our
 444 construction is useful, we need to show that the first arrival under the graph \tilde{T} exists and is almost
 445 surely finite, so that Algorithm 3 terminates with probability 1. First, note that for any $t > 0$, we
 446 have $\tilde{\mu}(t) \leq \mu(t) = t < \infty$. Since $\mathbb{P}[\mathbf{N}_{\tilde{\Pi}}(t) = 0] = e^{-\tilde{\mu}(t)}$, this implies that for any finite time t ,
 447 $\mathbb{P}[\mathbf{N}_{\tilde{\Pi}}(t) = 0] > 0$. Second, note that for $h \in [0, r^*]$,

$$w_Q(h) - h \cdot w_P(h) = \int_{\Omega} \mathbf{1}[r(x) \geq h](r(x) - h) dP(x), \quad (40)$$

448 from which

$$w_Q(r^*) - r^*w_P(r^*) = \int_{\Omega} \mathbf{1}[r(x) \geq r^*](r(x) - r^*) dP(x) \quad (41)$$

$$= \int_{\Omega} \mathbf{1}[r(x) = r^*](r(x) - r^*) dP(x) \quad (42)$$

$$= 0. \quad (43)$$

449 Thus, in particular, we find that

$$\lim_{h \rightarrow r^*} e^{-\tilde{\mu}(\sigma(h))} = \lim_{h \rightarrow r^*} \mathbb{P}[\mathbf{N}_{\tilde{\Pi}}(\sigma(h)) = 0] \stackrel{\text{eq. (34)}}{=} \lim_{h \rightarrow r^*} (w_Q(h) - h \cdot w_P(h)) = 0. \quad (44)$$

450 By continuity, this can only hold if $\tilde{\mu}(\sigma(h)) \rightarrow \infty$ as $h \rightarrow r^*$. Since $\tilde{\mu}(t)$ is finite for all $t > 0$ and σ
 451 is monotonically increasing, this also implies that $\sigma(h) \rightarrow \infty$ as $h \rightarrow r^*$. Thus, we have shown a
 452 couple of facts:

453 • $\sigma(h) \rightarrow \infty$ as $h \rightarrow r^*$, hence φ is always unbounded at points in Ω that achieve the
 454 supremum r^* . Furthermore, this implies that

$$\sigma^{-1}(t) \rightarrow r^* \quad \text{as } t \rightarrow \infty. \quad (45)$$

455 Since σ^{-1} is increasing, this implies that it is bounded from above by r^* .

- 456 • $\tilde{\mu}(t) < \infty$ for all $t > 0$, but $\tilde{\mu}(t) \rightarrow \infty$ as $t \rightarrow \infty$. Hence, by Equation (11) we have
 457 $\mathbb{P}[\mathbf{N}_{\tilde{\Pi}}(\infty) > 0] = 1$. Thus, the first arrival of $\tilde{\Pi}$ exists almost surely, and our construction
 458 is well-defined. In particular, it is meaningful to write $\mathbb{P}[\tilde{T} \geq t]$.
- 459 • $\mathbb{P}[\tilde{T} \geq t] = \mathbb{P}[\mathbf{N}_{\tilde{\Pi}}(t) = 0] \rightarrow 0$ as $t \rightarrow \infty$, which shows that the first arrival time is finite
 460 with probability 1. In turn, this implies that Algorithm 3 will terminate with probability one,
 461 as desired.

462 Note, that w_P and w_Q can be computed in many practically relevant cases, see Appendix G.

463 B Analysis of Greedy Poisson Rejection Sampling

464 Now that we constructed a correct sampling algorithm in Appendix A, we turn our attention to
 465 deriving the expected first arrival time \tilde{T} in the restricted process $\tilde{\Pi}$, the expected number of samples
 466 N before Algorithm 3 terminates and bound on N 's variance, and an upper bound on its coding
 467 cost. The proofs below showcase the true advantage of formulating the sampling algorithm using the
 468 language of Poisson processes: the proofs are all quite short and elegant.

469 B.1 The Expected First Arrival Time

470 At the end of the previous section, we showed that \tilde{T} is finite with probability one, which implies
 471 that $\mathbb{E}[\tilde{T}] < \infty$. Now, we derive the value of this expectation exactly. Since \tilde{T} is a positive random
 472 variable, by the Darth Vader rule (Muldowney et al., 2012), we may write its expectation as

$$\mathbb{E}[\tilde{T}] = \int_0^\infty \mathbb{P}[\tilde{T} \geq t] dt \stackrel{\text{eq. (26)}}{=} \lim_{t \rightarrow \infty} \sigma^{-1}(t) \stackrel{\text{eq. (45)}}{=} r^*. \quad (46)$$

473 B.2 The Expectation and Variance of the Runtime

474 **Expectation of N :** First, let $\tilde{\Pi}^C \stackrel{\text{def}}{=} \Pi \setminus \tilde{\Pi}$ be the set of points in Π above the graph of φ . Since $\tilde{\Pi}$
 475 and $\tilde{\Pi}^C$ are defined on complementary sets, they are independent. Since by definition \tilde{T} is the first
 476 arrival of $\tilde{\Pi}$ and the N th arrival of Π , it must mean that the first $N - 1$ arrivals of Π occurred in $\tilde{\Pi}^C$.
 477 Thus, conditioned on \tilde{T} , we have $N - 1 = |\{(T, X) \in \tilde{\Pi} \mid T < \tilde{T}\}|$ is Poisson distributed with mean

$$\mathbb{E}[N - 1 \mid \tilde{T}] = \int_0^{\tilde{T}} \int_{\Omega} \mathbf{1}[t \geq \varphi(x)] dP(x) dt \quad (47)$$

$$= \int_0^{\tilde{T}} \int_{\Omega} 1 - \mathbf{1}[t < \varphi(x)] dP(x) dt \quad (48)$$

$$= \tilde{T} - \tilde{\mu}(\tilde{T}). \quad (49)$$

478 By the law of iterated expectations,

$$\mathbb{E}[N - 1] = \mathbb{E}_{\tilde{T}} \left[\mathbb{E}[N - 1 \mid \tilde{T}] \right] \quad (50)$$

$$= \mathbb{E}[\tilde{T}] - \mathbb{E}[\tilde{\mu}(\tilde{T})]. \quad (51)$$

479 Focusing on the second term, we find

$$\mathbb{E}[\tilde{\mu}(\tilde{T})] = \int_0^\infty \tilde{\mu}(t) \cdot \mathbb{P}[\tilde{T} \in dt] dt \quad (52)$$

$$= \int_0^\infty \tilde{\mu}(t) \cdot \tilde{\mu}'(t) e^{-\tilde{\mu}(t)} dt \quad (53)$$

$$= \int_0^\infty u e^{-u} du = 1, \quad (54)$$

480 where the third equality follows by substituting $u = \tilde{\mu}(t)$. Finally, plugging the above and Equa-
 481 tion (46) into Equation (51), we find

$$\mathbb{E}[N] = 1 + \mathbb{E}[N - 1] = 1 + r^* - 1 = r^*. \quad (55)$$

482 **Variance of N :** We now show that the distribution of N is *super-Poissonian*, i.e. $\mathbb{E}[N] \leq \mathbb{V}[N]$.
 483 Similarly to the above, we begin with the law of iterated variances to find

$$\mathbb{V}[N] = \mathbb{V}[N - 1] = \mathbb{E}_{\tilde{T}}[\mathbb{V}[N - 1 | \tilde{T}]] + \mathbb{V}_{\tilde{T}}[\mathbb{E}[N - 1 | \tilde{T}]] \quad (56)$$

$$= \mathbb{E}[\tilde{T}] - \mathbb{E}[\tilde{\mu}(\tilde{T})] + \mathbb{V}[\tilde{T}] + \mathbb{V}[\tilde{\mu}(\tilde{T})], \quad (57)$$

484 where the second equality follows from the fact that the variance of $N - 1$ matches its mean
 485 conditioned on \tilde{T} , since it is a Poisson random variable. Focussing on the last term, we find

$$\mathbb{V}[\tilde{\mu}(\tilde{T})] = \mathbb{E}[\tilde{\mu}(\tilde{T})^2] - \mathbb{E}[\tilde{\mu}(\tilde{T})]^2 \quad (58)$$

$$= \int_0^\infty \tilde{\mu}(t)^2 \cdot \mathbb{P}[\tilde{T} \in dt] dt - 1 \quad (59)$$

$$= \int_0^\infty u^2 e^{-u} dt - 1 = 1, \quad (60)$$

486 where the third equality follows from a similar u -substitution as in Equation (54). Thus, putting
 487 everything together, we find

$$\mathbb{V}[N] = r^* - 1 + \mathbb{V}[\tilde{T}] + 1 = r^* + \mathbb{V}[\tilde{T}] \geq \mathbb{E}[N]. \quad (61)$$

488 B.3 The Codelength of the Index

489 As discussed in Section 3, Alice and Bob can realize a one-shot channel simulation protocol using
 490 GPRS for a pair of correlated random variables $\mathbf{x}, \mathbf{y} \sim P_{\mathbf{x}, \mathbf{y}}$ if they have access to shared randomness
 491 and can simulate samples from $P_{\mathbf{x}}$. In particular, after receiving $\mathbf{y} \sim P_{\mathbf{y}}$, Alice runs GPRS with
 492 proposal distribution $P_{\mathbf{x}}$ and target $P_{\mathbf{x}|\mathbf{y}}$, and the shared randomness to simulate the Poisson process
 493 Π . She then encodes the index N of the first arrival of Π under the graph of φ . Bob can decode Alice's
 494 sample $\mathbf{x} \sim P_{\mathbf{x}|\mathbf{y}}$ by simulating the same N samples from Π using the shared randomness. The
 495 question is, how efficiently can Alice encode N ? We answer this question by following the approach
 496 of Li & El Gamal (2018). Namely, we first bound the conditional expectation $\mathbb{E}[\log_2 N | \mathbf{y} = y]$,
 497 after which we average over \mathbf{y} to bound $\mathbb{E}[\log_2 N]$. Then, we use the maximum entropy distribution
 498 subject to the constraint of fixed $\mathbb{E}[\log_2 N]$ to bound $\mathbb{H}[N]$. Finally, noting that \mathbf{x} is a function of Π
 499 and N , we get

$$\mathbb{H}[\mathbf{x}, N | \Pi] = \underbrace{\mathbb{H}[\mathbf{x} | N, \Pi]}_{=0} + \mathbb{H}[N | \Pi] \quad (62)$$

$$= \underbrace{\mathbb{H}[N | \mathbf{x}, \Pi]}_{=0} + \mathbb{H}[\mathbf{x} | \Pi], \quad (63)$$

500 from which

$$\mathbb{H}[N] \geq \mathbb{H}[N | \Pi] = \mathbb{H}[\mathbf{x} | \Pi] \geq \mathbb{H}[\mathbf{x} | \Pi], \quad (64)$$

501 which will finish the proof.

502 **Bound on the conditional expectation:** Fix \mathbf{y} and set $Q = P_{\mathbf{x}|\mathbf{y}}$ and $P = P_{\mathbf{x}}$ as GPRS's target
 503 and proposal distribution, respectively. Let $(t, x) \in \mathbb{R}^+ \times \Omega$ be a point under the graph of φ , i.e.
 504 $\sigma^{-1}(t) \leq r(x)$. Then,

$$r(x) \geq \sigma^{-1}(t) \quad (65)$$

$$\stackrel{\text{eq. (27)}}{=} \int_0^t \mathbb{P}[\tilde{T} \geq \tau] d\tau \quad (66)$$

$$\geq t \cdot \inf_{\tau \in [0, t]} \left\{ \mathbb{P}[\tilde{T} \geq \tau] \right\} \quad (67)$$

$$= t \cdot \mathbb{P}[\tilde{T} \geq t]. \quad (68)$$

505 From this, we get

$$t + 1 \leq \frac{r(x)}{\mathbb{P}[\tilde{T} \geq t]} + 1 \leq \frac{r(x) + 1}{\mathbb{P}[\tilde{T} \geq t]}. \quad (69)$$

506 Next, conditioning on the first arrival (\tilde{T}, \tilde{X}) we get

$$\mathbb{E}[\log_2 N \mid \tilde{T} = t, \tilde{X} = x] \leq \log_2 \left(\mathbb{E}[N - 1 \mid \tilde{T} = t, \tilde{X} = x] + 1 \right) \quad (70)$$

$$\stackrel{\text{eq. (49)}}{=} \log_2 (t - \tilde{\mu}(t) + 1) \quad (71)$$

$$\leq \log_2 (t + 1) \quad (72)$$

$$\stackrel{\text{eq. (69)}}{\leq} \log_2 (r(x) + 1) - \log_2 \mathbb{P}[\tilde{T} \geq t] \quad (73)$$

$$= \log_2 (r(x) + 1) + \tilde{\mu}(t) \cdot \log_2 e, \quad (74)$$

507 where the first inequality follows by Jensen's inequality.

508 Now, by the law of iterated expectations, we find

$$\mathbb{E}[\log_2 N] = \mathbb{E}_{\tilde{X}} \left[\mathbb{E}_{\tilde{T} \mid \tilde{X}} \left[\mathbb{E}[\log_2 N \mid \tilde{T}, \tilde{X}] \right] \right] \quad (75)$$

$$\stackrel{\text{eq. (74)}}{\leq} \mathbb{E}_{\tilde{X}} \left[\log_2 (r(\tilde{X}) + 1) \right] + \underbrace{\mathbb{E}_{\tilde{T}} [\tilde{\mu}(\tilde{T})]}_{=1 \text{ by eq. (54)}} \cdot \log_2 e \quad (76)$$

$$= D_{\text{KL}}[Q \parallel P] + \mathbb{E}_{\tilde{X}} \left[\log_2 \left(1 + \frac{1}{r(\tilde{X})} \right) \right] + \log_2 e \quad (77)$$

$$\leq D_{\text{KL}}[Q \parallel P] + \mathbb{E}_{\tilde{X}} \left[\log_e (e^{-r(\tilde{X})}) \right] \cdot \log_2 e + \log_2 e \quad (78)$$

$$= D_{\text{KL}}[Q \parallel P] + 2 \cdot \log_2 e, \quad (79)$$

509 where the second inequality follows from switching to the natural base and using a first-order Taylor
510 approximation for e^{-x} .

511 **Bound on the marginal expectation and entropy:** Equation (79) is a one-shot bound, which yields

$$\mathbb{E}[\log_2 N \mid \mathbf{y}] \leq D_{\text{KL}}[P_{\mathbf{x} \mid \mathbf{y}} \parallel P_{\mathbf{x}}] + 2 \cdot \log_2 e. \quad (80)$$

512 Taking expectation over $\mathbf{y} \sim P_{\mathbf{y}}$, we get

$$\mathbb{E}[\log_2 N] \leq I[\mathbf{x}; \mathbf{y}] + 2 \cdot \log_2 e. \quad (81)$$

513 Finally, following Proposition 4 in Li & El Gamal (2018), the maximum entropy distribution for N
514 subject to a constraint on $\mathbb{E}[\log_2 N]$ obeys

$$\mathbb{H}[N] \leq \mathbb{E}[\log_2 N] + \log (\mathbb{E}[\log_2 N] + 1) + 1. \quad (82)$$

515 Plugging in Equation (81) into the above, we get

$$\mathbb{H}[N] \leq I[\mathbf{x}; \mathbf{y}] + \log_2 (I[\mathbf{x}; \mathbf{y}] + 2 \log_2 e + 1) + 2 \log_2 e + 1 \quad (83)$$

$$\leq I[\mathbf{x}; \mathbf{y}] + \log_2 (I[\mathbf{x}; \mathbf{y}] + 1) + 2 \log_2 e + 1 + \log_2 (2 \log_2 e + 1) \quad (84)$$

$$< I[\mathbf{x}; \mathbf{y}] + \log_2 (I[\mathbf{x}; \mathbf{y}] + 1) + 6. \quad (85)$$

516 Similarly to Li & El Gamal (2018), we can encode N using a Zeta distribution $\zeta(n \mid s) \propto n^{-s}$ with

$$s \stackrel{\text{def}}{=} \frac{1}{I[\mathbf{x}; \mathbf{y}] + 2 \log_2 e}. \quad (86)$$

517 With this choice of the coding distribution, the expected codelength of N is upper bounded by
518 $I[\mathbf{x}; \mathbf{y}] + \log_2 (I[\mathbf{x}; \mathbf{y}] + 1) + 7$ bits.

519 **C Analysis of Parallel GPRS**

520 Parallel GPRS (PGPRS) is a general technique to accelerate GPRS when parallel computing power
 521 is available and is presented in Section 3.1. Assuming that J parallel threads are available, for
 522 target distribution Q and proposal P , PGPRS simulates J Poisson processes Π_1, \dots, Π_J in parallel,
 523 all of them with mean measure $\mu_i \stackrel{\text{def}}{=} \frac{\lambda}{J} \times P$. Assuming $\left\{ \left(T_{N_1}^{(1)}, X_{N_1}^{(1)} \right), \dots, \left(T_{N_J}^{(J)}, X_{N_J}^{(J)} \right) \right\}$ are
 524 the first arrivals of Π_1, \dots, Π_J under φ , respectively, PGPRS selects the arrival with the overall
 525 smallest arrival time $J^* = \arg \min_{j \in \{1, \dots, J\}} \left\{ T_{N_j}^{(j)} \right\}$. By the superposition theorem (Kingman,
 526 1992), $\cup_{j=1}^J \Pi_j = \Pi$ is a Poisson process with mean measure $\mu = \lambda \times P$, hence $\left(T_{N_{J^*}}^{(J^*)}, X_{N_{J^*}}^{(J^*)} \right)$ will
 527 be the first arrival of Π under the graph of φ and the measure-theoretic construction in Appendix A
 528 therefore guarantees the correctness of PGPRS.

529 **Expected runtime:** We will proceed similarly to the analysis in Appendix B.2. Concretely, let
 530 $\tilde{T} = T_{N_{J^*}}^{(J^*)}$ be the first arrival time of Π under the graph of φ . Let

$$\nu_j \stackrel{\text{def}}{=} \left| \left\{ (T, X) \in \Pi_j \mid T < \tilde{T}, T > \varphi(X) \right\} \right| \quad (87)$$

531 be the number of points in Π_j that are rejected before the global first arrival occurs. By an independen-
 532 ce argument analogous to the one given in Appendix B.2, the ν_1, \dots, ν_J are independent Poisson
 533 random variables, each distributed with mean

$$\mathbb{E} \left[\nu_j \mid \tilde{T} \right] = \frac{1}{J} \left(\tilde{T} - \tilde{\mu}(\tilde{T}) \right). \quad (88)$$

534 Hence, by the law of iterated expectations, we find that we get

$$\mathbb{E} [\nu_j] \stackrel{\text{eq. (51)}}{=} \frac{r^* - 1}{J} \quad (89)$$

535 rejections in each of the J threads on average before the global first arrival occurs. Now, a thread j
 536 terminates either when the global first arrival occurs in it or when the thread's current arrival time
 537 $T_{n_j}^{(j)}$ provably exceeds the global first arrival time. Thus, on average, each thread will have one more
 538 rejection compared to Equation (89). Hence the average runtime of a thread is

$$\mathbb{E}[\nu_j + 1] = \frac{r^* - 1}{J} + 1, \quad (90)$$

539 and the average number of samples simulated by PGPRS across all its threads is

$$J \cdot \left(\frac{r^* - 1}{J} + 1 \right) = r^* + J - 1. \quad (91)$$

540 **Variance of Runtime:** Once again, similarly to Appendix B.2, we find by the law of iterated variances
 541 that the variance of the runtime in each of the j threads is

$$\mathbb{V}[\nu_j + 1] = \mathbb{V}[\nu_j] = \frac{r^* - 1}{J} + \frac{1}{J^2} \left(\mathbb{V}[\tilde{T}] + 1 \right), \quad (92)$$

542 meaning that we make an $\mathcal{O}(1/J)$ reduction in the variance of the runtime compared to regular
 543 GPRS.

544 **Codelength:** PGPRS can also realize a one-shot channel simulation protocol for a pair of correlated
 545 random variables $\mathbf{x}, \mathbf{y} \sim P_{\mathbf{x}, \mathbf{y}}$. For a fixed $\mathbf{y} \sim P_{\mathbf{y}}$, Alice applies PGPRS to the target $Q = P_{\mathbf{x}|\mathbf{y}}$ and
 546 proposal $P = P_{\mathbf{x}}$, and encodes the two-part code (J^*, N_{J^*}) . Bob can then simulate N_{J^*} samples
 547 from Π_{J^*} and recover Alice's sample.

548 **Encoding J^* :** By the symmetry of the setup, the global first arrival will occur with equal probability
 549 in each subprocess Π_j . Hence J^* follows a uniform distribution on $\{1, \dots, J\}$. Therefore, Alice can
 550 encode J^* optimally using $\lceil \log_2 J \rceil$ bits.

551 **Encoding N_{J^*} :** We can develop a bound using an almost identical argument to the one in Ap-
 552 pendix B.3. In particular, by adapting the conditional bound in Equation (74) appropriately using
 553 Equation (90), we get

$$\mathbb{E} \left[\log_2 N_{J^*} \mid \tilde{T} = t, \tilde{X} = x, J^* \right] \leq \log_2(r(x) + 1) + \tilde{\mu}(t) \cdot \log_2 e - \log_2 J. \quad (93)$$

554 Then, using this conditional bound and adapting Equation (79), we find

$$\mathbb{E}[\log_2 N_{J^*} \mid \mathbf{y}, J^*] \leq D_{\text{KL}}[P_{\mathbf{x}|\mathbf{y}} \| P_{\mathbf{x}}] + 2 \cdot \log_2 e - \log_2 J \quad (94)$$

555 to obtain a one-shot bound. Taking expectation over $\mathbf{y} \sim P_{\mathbf{y}}$, we get

$$\mathbb{E}[\log_2 N_{J^*} \mid J^*] \leq I[\mathbf{x}; \mathbf{y}] + 2 \cdot \log_2 e - \log_2 J, \quad (95)$$

556 hence, by adapting the maximum entropy bound in Equation (85), we find

$$\mathbb{H}[N_{J^*} \mid J^*] < I[\mathbf{x}; \mathbf{y}] - \log_2 J + \log(I[\mathbf{x}; \mathbf{y}] - \log_2 J + 1) + 6. \quad (96)$$

557 Thus, we finally find that the entropy of the two-part code (J^*, N_{J^*}) is upper bounded by

$$\mathbb{H}[J^*, N_{J^*}] < I[\mathbf{x}; \mathbf{y}] + \log(I[\mathbf{x}; \mathbf{y}] - \log_2 J + 1) + 7. \quad (97)$$

558 Using a Zeta distribution $\zeta(n \mid s) \propto n^{-s}$ to encode $N_{J^*} \mid J^*$ with

$$s \stackrel{\text{def}}{=} \frac{1}{I[\mathbf{x}; \mathbf{y}] + 2 \cdot \log_2 e - \log_2 J}, \quad (98)$$

559 we find that the expected codelength of the two-part code is upper bounded by

$$\mathbb{H}[J^*, N_{J^*}] < I[\mathbf{x}; \mathbf{y}] + \log(I[\mathbf{x}; \mathbf{y}] - \log_2 J + 1) + 8 \text{ bits}. \quad (99)$$

560 D Simulating Poisson Processes Using Tree-Structured Partitions

561 In this section, we examine an advanced simulation technique for Poisson processes, which is
 562 required to formulate the binary search-based variants of GPRS. We first recapitulate the tree-based
 563 simulation technique from (Flamich et al., 2022) and some important results from their work. Then,
 564 we present Algorithm 7, using which we can finally formulate the optimally efficient variant of
 565 GPRS in Appendix E. **Note:** For simplicity, we present the ideas for Poisson processes whose
 566 spatial measure P is non-atomic. These ideas can also be extended to atomic spatial measures with
 567 appropriate modifications.

568 **Splitting functions:** Let Π be a spatiotemporal Poisson process on some space Ω with non-atomic
 569 spatial measure P . In this paragraph, we will not deal with Π itself yet, but the space on which it is
 570 defined and the measure P . Now, assume that there is a function `split`, which for any given Borel
 571 set $B \subseteq \Omega$, produces a (possibly random) P -essential partition of B consisting of two Borel sets, i.e.

$$\text{split}(B) = \{L, R\}, \text{ such that } L \cap R = \emptyset \text{ and } P(L \cup R) = P(B). \quad (100)$$

572 The last condition simply allows `split` to discard some points from B that will not be included in
 573 either the left or right split. For example, if we wish to design a splitting function for a subset of
 574 $B \subseteq \mathbb{R}^d$ to split B along some hyperplane, we do not need to include the points on the hyperplane in
 575 either set. The splitting function that always exists is the *trivial* splitting function, which just returns
 576 the original set and the empty set:

$$\text{split}_{\text{trivial}}(B) = \{B, \emptyset\}. \quad (101)$$

577 A more interesting example when $\Omega = \mathbb{R}$ is the *on-sample* splitting function, where for a \mathbb{R} -valued
 578 random variable $X \sim P|_B$ it splits B as

$$\text{split}_{\text{on-samp}}(B) = \{(-\infty, X) \cap B, (X, \infty) \cap B\}. \quad (102)$$

579 This function is used by Algorithm 5 and AS* coding (Flamich et al., 2022). Another example is the
 580 *dyadic* splitting function operating on a bounded interval (a, b) , splitting as

$$\text{split}_{\text{dyad}}((a, b)) = \left\{ \left(a, \frac{a+b}{2} \right), \left(\frac{a+b}{2}, b \right) \right\}. \quad (103)$$

581 We call this dyadic because when we apply this splitting function to $(0, 1)$ and subsets pro-
 582 duced by applying it, we get the set of all intervals with dyadic endpoints on $(0, 1)$, i.e.
 583 $\{(0, 1), (0, 1/2), (1/2, 1), (0, 1/4), \dots\}$. This splitting function is used by Algorithm 6 and AD*
 584 coding. As one might imagine, there might be many more possible splitting functions than the two
 585 examples we give above, all of which might be more or less useful in practice.

Algorithm 7: Simulating a tree construction with another

Input : Proposal distribution P
Target splitting function $\text{split}_{\text{target}}$,
Simulating splitting function $\text{split}_{\text{sim}}$,
Target heap index H_{target}

Output : Arrival (T, X) with heap index H_{target} in $\mathcal{T}_{\text{target}}$, and heap index H_{sim} of the arrival in \mathcal{T}_{sim} .

```
 $\mathcal{P} \leftarrow \text{PriorityQueue}$ 
 $B \leftarrow \Omega$ 
 $K \leftarrow \lfloor \log_2 H_{\text{target}} \rfloor$ 
 $X \sim P$ 
 $T \sim \text{Exp}(1)$ 
 $\mathcal{P}.\text{push}(T, X, \Omega, 1)$ 
for  $k = 0, \dots, K$  do
  repeat
     $T, X, C, H \leftarrow \mathcal{P}.\text{pop}()$ 
     $L, R \leftarrow \text{split}_{\text{sim}}(C)$ 
    if  $B \cap L \neq \emptyset$  then
       $\Delta_L \sim \text{Exp}(P(L))$ 
       $T_L \leftarrow T + \Delta_L$ 
       $X_L \sim P|_L$ 
       $\mathcal{P}.\text{push}(T_L, X_L, L, 2H)$ 
    end
    if  $B \cap R \neq \emptyset$  then
       $\Delta_R \sim \text{Exp}(P(R))$ 
       $T_R \leftarrow T + \Delta_R$ 
       $X_R \sim P|_R$ 
       $\mathcal{P}.\text{push}(T_R, X_R, R, 2H + 1)$ 
    end
  until  $X \in B$  /* Exit loop when we find first arrival in  $B$ . */
  if  $k < K$  then
     $\{B_0, B_1\} \leftarrow \text{split}_{\text{target}}(B)$ 
     $d \leftarrow \left\lceil \frac{H_{\text{target}}}{2^{K-k}} \right\rceil \bmod 2$ 
     $B \leftarrow B_d$ 
  end
end
return  $T, X, H$ 
```

586 **Split-induced binary space partitioning tree:** Every splitting function on a space Ω induces an
587 infinite set of subsets by repeatedly applying the splitting function to the splits it produces. These sets
588 can be organised into an infinite *binary space partitioning tree (BSP-tree)*, where each node in the
589 tree is represented by a set produced by split and an unique index. Concretely, let the root of the
590 tree be represented by the whole space Ω and the index $H_{\text{root}} = 1$. Now we recursively construct the
591 rest of the tree as follows: Let (B, H) be a node in the tree, with B a Borel set and H its index, and
592 let $\{L, R\} = \text{split}(B)$ the left end right splits of B . Then, we set $(L, 2H)$ as the node's left child
593 and $(R, 2H + 1)$ as its right child. We refer to the index associated with each node as its *heap index*.

594 **Heap-indexing the points in Π and a strict heap invariant:** As we saw in Section 2.1, each point in
595 Π can be uniquely identified by their *time index* N , i.e. if we time-order the points of Π , N represents
596 the N th arrival in the ordered list. However, we can also uniquely index each point in Π using a
597 splitting function-induced BSP-tree as follows.

598 We extend each node in the BSP-tree with a point from Π , such that the extended tree satisfies a strict
599 *heap invariant*: First, we extend the root node $(\Omega, 1)$ by adjoining Π 's first arrival (T_1, X_1) and the
600 first arrival index 1 to get $(T_1, X_1, \Omega, 1, 1)$. Then, for every other extended node (T, X, B, H, N)
601 with parent node $(T', X', B', \lfloor H/2 \rfloor, M)$ we require that (T, X) is the first arrival in the restricted
602 process $\Pi \cap ((T', \infty) \times B)$. This restriction enforces that we always have that $T' < T$ and, therefore,

603 $M < N$. Furthermore, it is strict because there are no other points of Π in B between those two
604 arrival times.

605 **Notation for the tree structure:** Let us denote the extended BSP tree \mathcal{T} on Π induced by `split`.
606 Each node $\nu \in \mathcal{T}$ is a tuple $\nu = (T, X, B, H, N)$ consisting of the *arrival time* T , *spatial coordinate*
607 X , its *bounds* B , *split-induced heap index* H and *time index* N . As a slight overload of notation,
608 for a node ν , we use it as subscript to refer to its elements, e.g. T_ν is the arrival time associated
609 with node ν . We now prove the following important result, which ensures that we do not lose any
610 information by using \mathcal{T} to represent Π . An essentially equivalent statement was first proven by for
611 Gumbel processes (Appendix, “Equivalence under *partition*” subsection; Maddison et al., 2014).

612 **Lemma D.1.** *Let Π , P , `split` and \mathcal{T} be as defined above. Then, P -almost surely, \mathcal{T} contains every*
613 *point of Π .*

614 *Proof.* We give an inductive argument.

615 *Base case:* the first arrival of Π is by definition contained in the root node of \mathcal{T} .

616 *Inductive hypothesis:* assume that the first N arrivals of Π are all contained in \mathcal{T} .

617 *Case for $N + 1$:* We begin by observing that if the first N nodes are contained in \mathcal{T} , they must form a
618 connected subtree \mathcal{T}_N . To see this, assume the contrary, i.e. that the first N arrivals form a subgraph
619 $\Gamma_N \subseteq \mathcal{T}$ with multiple disconnected components. Let $\nu \in \Gamma_N$ be a node contained in a component
620 of Γ_N that is *disconnected* from the component of Γ_N containing the root node. Since \mathcal{T} is a tree,
621 there is a unique path π from the root node to ν in \mathcal{T} , and since ν and the root are disconnected in
622 Γ_N , π must contain a node $c \notin \Gamma_N$. However, since c is an ancestor of ν , by the heap invariant of \mathcal{T}
623 we must have that the time index of c is $N_c < N_\nu \leq N$ hence $c \in \Gamma_N$, a contradiction.

624 Thus, let \mathcal{T}_N represent the subtree of \mathcal{T} containing the first N arrivals of Π . Now, let \mathcal{F}_N represent
625 the *frontier* of \mathcal{T}_N , i.e. the leaf nodes’ children:

$$\mathcal{F}_N \stackrel{\text{def}}{=} \{\nu \in \mathcal{T} \mid \nu \notin \mathcal{T}_N, \text{parent}(\nu) \in \mathcal{T}_N\}, \quad (104)$$

626 where `parent` retrieves the parent of a node in \mathcal{T} . Let

$$\bar{\Omega}_N \stackrel{\text{def}}{=} \bigcup_{\nu \in \mathcal{F}_N} B_\nu \quad (105)$$

627 be the union of all the bounds of the nodes in the frontier. A simple inductive argument shows that for
628 all N the nodes in \mathcal{F}_N provide a P -essential partition of Ω , from which $P(\bar{\Omega}_N) = 1$. Let T_N be the
629 N th arrival time of Π . Now, by definition, the arrival time associated with every node in the frontier
630 \mathcal{F}_N must be later than T_N . Finally, consider the first arrival time across the nodes in the frontier:

$$T_N^* \stackrel{\text{def}}{=} \min_{\nu \in \mathcal{F}_N} T_\nu. \quad (106)$$

631 Then, conditioned on \mathcal{T}_N , T_N^* is the first arrival of Π restricted to $(T_N, \infty) \times \bar{\Omega}$, thus it is P -almost
632 surely the $N + 1$ st arrival in Π , as desired. \square

633 **A connection between the time index an heap index of a node:** Now that we have two ways of
634 uniquely identifying each point in Π it is natural to ask whether there is any relation between them.
635 In the next paragraph we adapt an argument from Flamich et al. (2022) to show that under certain
636 circumstances, the answer is yes.

637 First, we need to define two concepts, the *depth* of a node in an extended BSP-tree and a *contractive*
638 *splitting function*. Thus, let \mathcal{T} be an extended BSP-tree over Π induced by `split`. Let $\nu \in \mathcal{T}$. Then,
639 the depth D of ν in \mathcal{T} is defined as the distance from ν to the root. A simple argument shows that the
640 heap index H_ν of every node ν at depth D in \mathcal{T} is between $2^D \leq H_\nu < 2^{D+1}$. Thus, we can easily
641 obtain the depth of a node ν from the heap index via the formula $D_\nu = \lfloor \log_2 H_\nu \rfloor$. Next, we say that
642 `split` is *contractive* if **all** the bounds it produces shrink on average. More formally, let $\epsilon < 1$, and
643 for a node $\nu \in \mathcal{T}$, let \mathcal{A}_ν denote the set of ancestor nodes of ν . Then, `split` is contractive if for
644 every non-root node $\nu \in \mathcal{T}$ we have

$$\mathbb{E}_{\mathcal{T}_{\mathcal{A}_\nu} | D_\nu} [P(B_\nu)] \leq \epsilon^{D_\nu}, \quad (107)$$

645 where $\mathcal{T}_{\mathcal{A}_\nu}$ and denotes the subtree of \mathcal{T} containing the arrivals of the ancestors of ν .

646 Note that $\epsilon \geq 1/2$ for any `split` function. This is because if $\{L, R\} = \text{split}(B)$ for some set
 647 B , then $P(R) = P(B) - P(L)$. Thus, if $P(R) = \alpha P(B)$, then $P(L) = (1 - \alpha)P(B)$, and by
 648 definition $\epsilon = \max\{\alpha, 1 - \alpha\}$, which is minimized when $\alpha = 1/2$, from which $\epsilon = 1/2$.

649 For example, `splitdyad` is contractive with $\epsilon = 1/2$, while `splittrivial` is not contractive. By Lemma
 650 1 from Flamich et al. (2022), `spliton-samp` is also contractive with $\epsilon = 3/4$. We now strengthen this
 651 result using a simple argument, and show that `spliton-samp` is contractive with $\epsilon = 1/2$.

652 **Lemma D.2.** *Let ν, D be defined as above, let P be a non-atomic probability measure over \mathbb{R} with
 653 CDF F_P . Let Π a $(1, P)$ -Poisson process and \mathcal{T} be the BSP-tree over Π induced by `spliton-samp`.
 654 Then,*

$$\mathbb{E}_{\mathcal{T}_{\mathcal{A}_\nu} | D_\nu} [P(B_\nu)] = 2^{-D_\nu}. \quad (108)$$

655 *Proof.* Fix $D_\nu = d$, and let $\mathcal{N}_d = \{n \in \mathcal{T} \mid D_n = d\}$ be the set of nodes in \mathcal{T} whose depth is d .
 656 Note, that $|\mathcal{N}_d| = 2^d$. We will show that

$$\forall n, m \in \mathcal{N}_d : P(B_n) \stackrel{d}{=} P(B_m), \quad (109)$$

657 i.e. that the distributions of the bound sizes are all the same. From this, we will immediately have
 658 $\mathbb{E}[P(B_n)] = \mathbb{E}[P(B_\nu)]$ for every $n \in \mathcal{N}_d$. Then, using the fact, that the nodes in \mathcal{N}_d for a P -almost
 659 partition of Ω , we get:

$$1 = \mathbb{E} \left[P \left(\bigcup_{n \in \mathcal{N}_d} B_n \right) \right] = \mathbb{E} \left[\sum_{n \in \mathcal{N}_d} P(B_n) \right] = \sum_{n \in \mathcal{N}_d} \mathbb{E} [P(B_n)] = |\mathcal{N}_d| \cdot \mathbb{E} [P(B_\nu)]. \quad (110)$$

660 Dividing the very left and very right by $|\mathcal{N}_d| = 2^d$ yields the desired result.

661 To complete the proof, we now show that by symmetry, Equation (109) holds. We begin by exposing
 662 the fundamental symmetry of `spliton-samp`: for a node ν with left child L and right child R , the left
 663 and right bound sizes are equal in distribution:

$$P(B_L) \stackrel{d}{=} P(B_R) \mid P(B_\nu). \quad (111)$$

664 First, note that by definition, all involved bounds will be intervals. Namely, assume that $B_\nu = (a, b)$
 665 for some $a < b$ and X_ν is the sample associated with ν . Then, $B_L = (a, X_\nu)$ and $B_R = (X_\nu, b)$ and
 666 hence $P(B_L) = F_P(X_\nu) - F_P(a)$. Since $X_\nu \sim P|_{B_\nu}$, by the probability integral transform, we have
 667 $F(X_\nu) \sim \text{Unif}(F_P(a), F_P(b))$, from which $P(B_L) \sim \text{Unif}(0, F_P(b) - F_P(a)) = \text{Unif}(0, P(B_\nu))$.
 668 Since $P(B_R) = P(B_\nu) - P(B_L)$, we similarly have $P(B_R) \sim \text{Unif}(0, P(B_\nu))$, which establishes
 669 our claim.

670 Now, to show Equation (109), fix d and fix $n \in \mathcal{N}_d$. Let \mathcal{A}_n denote the ancestor nodes of n . As we
 671 saw in the paragraph above,

$$P(B_n) \mid P(B_{\text{parent}(n)}) \stackrel{d}{=} P(B_{\text{parent}(n)}) \cdot U, \quad U \sim \text{Unif}(0, 1), \quad (112)$$

672 regardless of whether n is a left or a right child of its parent. We can apply this d times to the ancestors
 673 of n find the marginal:

$$P(B_n) \stackrel{d}{=} \prod_{i=1}^d U_i, \quad U_i \sim \text{Unif}(0, 1). \quad (113)$$

674 Since the choice of n was arbitrary, all nodes in \mathcal{N}_d have this distribution, which is what we wanted
 675 to show. \square

676 Now, we have the following result.

677 **Lemma D.3.** *Let `split` be a contractive splitting function for some $\epsilon \in [1/2, 1)$. Then, for every
 678 node ν in \mathcal{T} with time index N_ν and depth D_ν , we have*

$$\mathbb{E}_{D_\nu | N_\nu} [D_\nu] \leq -\log_\epsilon N_\nu. \quad (114)$$

679 *Proof.* Let us examine the case $N_\nu = 1$ first. In this case, ν is the root of \mathcal{T} and has depth $D_\nu = 0$
680 by definition. Thus, $0 \leq -\log_\epsilon 1$ holds trivially.

681 Now, fix $\nu \in \mathcal{T}$ with time index $N_\nu = N > 1$. Let \mathcal{T}_{N-1} be the subtree of \mathcal{T} containing the first
682 $N - 1$ arrivals, \mathcal{F}_{N-1} be the frontier of \mathcal{T}_{N-1} and T_{N-1} the $(N - 1)$ st arrival time. Then, as we saw
683 in Lemma D.1, the N th arrival in Π occurs in one of the nodes in the frontier \mathcal{F}_{N-1} , after T_{N-1} .
684 In particular, conditioned on \mathcal{T}_{N-1} , the arrival times associated with each node $f \in \mathcal{F}_{N-1}$ will be
685 shifted exponentials $T_f = T_{N-1} + \text{Exp}(P(B_f))$, and the N th arrival time in Π is the minimum of
686 these: $T_\nu = T_N = \min_{f \in \mathcal{F}_{N-1}} T_f$. It is a standard fact (see e.g. Lemma 6 in Maddison (2016)) that
687 the index of the minimum

$$F_N = \arg \min_{f \in \mathcal{F}_{N-1}} T_f \quad (115)$$

688 is independent of $T_\nu = T_{F_N}$ and $\mathbb{P}[F_N = f \mid \mathcal{T}_{N-1}] = P(B_f)$. A simple inductive argument shows
689 that the number of nodes on the frontier $|\mathcal{F}_{N-1}| = N$. Thus, we have a simple upper bound on the
690 entropy of F :

$$\mathbb{E}_{F_N \mid \mathcal{T}_{N-1}, N_\nu = N}[-\log_2 P(B_\nu)] = \mathbb{H}[F_N \mid \mathcal{T}_{N-1}, N_\nu = N] \leq \log_2 N. \quad (116)$$

691 Thus, taking expectation over \mathcal{T}_{N-1} , we find

$$\log_2 N \stackrel{\text{eq. (116)}}{\geq} \mathbb{E}_{F_N, \mathcal{T}_{N-1} \mid N_\nu = N}[-\log_2 P(B_\nu)] \quad (117)$$

$$= \mathbb{E}_{D_\nu \mid N_\nu = N} [\mathbb{E}_{F_N, \mathcal{T}_{N-1} \mid D_\nu, N_\nu = N}[-\log_2 P(B_\nu)]] \quad (118)$$

$$\geq \mathbb{E}_{D_\nu \mid N_\nu = N} [-\log_2 \mathbb{E}_{F_N, \mathcal{T}_{N-1} \mid D_\nu, N_\nu = N}[P(B_\nu)]] \quad (119)$$

$$\stackrel{\text{eq. (107)}}{\geq} \mathbb{E}_{D_\nu \mid N_\nu = N} [-\log_2 \epsilon^{D_\nu}] \quad (120)$$

$$= (-\log_2 \epsilon) \cdot \mathbb{E}_{D_\nu \mid N_\nu = N}[D_\nu]. \quad (121)$$

692 The second inequality holds by Jensen's inequality. In the third inequality, we apply Equation (107),
693 and one might worry about conditioning on N_ν here. However, this is not an issue because

$$\mathbb{E}_{F_N, \mathcal{T}_{N-1} \mid D_\nu = d, N_\nu = N}[P(B_\nu)] \quad (122)$$

$$= \mathbf{1}[d \leq N - 1] \cdot \sum_{f \in \mathcal{F}_{N-1}} \mathbb{E}_{\mathcal{T}_{A_f} \mid F_N = f, D_\nu = d}[P(B_f)] \cdot \mathbb{P}[F_N = f \mid D_f = d] \quad (123)$$

$$\stackrel{\text{eq. (107)}}{\leq} \mathbf{1}[d \leq N - 1] \cdot \sum_{f \in \mathcal{F}_{N-1}} \epsilon^d \cdot \mathbb{P}[F_N = f \mid D_f = d] \quad (124)$$

$$= \mathbf{1}[d \leq N - 1] \cdot \epsilon^d. \quad (125)$$

694 Thus, we finally get

$$(-\log_2 \epsilon) \cdot \mathbb{E}_{D_\nu \mid N_\nu = N}[D_\nu] \leq \log_2 N \quad \Rightarrow \quad \mathbb{E}_{D_\nu \mid N_\nu} [D_\nu] \leq -\log_\epsilon N_\nu \quad (126)$$

695 by dividing both sides by $-\log_2 \epsilon$ and we obtain the desired result. \square

696 **Converting between different heap indices:** Assume now that we have two splitting functions,
697 $\text{split}_{\text{target}}$ and $\text{split}_{\text{sim}}$, which induce their own BSP-ordering on Π , $\mathcal{T}_{\text{target}}$ and \mathcal{T}_{sim} . Now, given
698 a $\text{split}_{\text{target}}$ -induced heap index H , Algorithm 7 presents a method for simulating the appropriate
699 node $\nu \in \mathcal{T}_{\text{target}}$ by simulating nodes from \mathcal{T}_{sim} . In other words, given a node with some heap index
700 induced by a splitting function, Algorithm 7 lets us find the heap index of the same arrival induced by
701 a different splitting function. The significance of Algorithm 7 is that it lets us develop convenient
702 search methods using a given splitting function, but it might be more efficient to encode the heap
703 index induced by another splitting function.

704 **Theorem D.4.** Let Π , $\text{split}_{\text{target}}$, $\text{split}_{\text{sim}}$, $\mathcal{T}_{\text{target}}$ and \mathcal{T}_{sim} be as above. Let $\nu \in \mathcal{T}_{\text{target}}$ with and let
705 $(T_{\text{sim}}, X_{\text{sim}}, H_{\text{sim}})$ be the arrival and its heap index output by Algorithm 7 given the above as input as
706 well as H_ν as the target index. Then, Algorithm 7 is correct, in the sense that

$$T_\nu \stackrel{d}{=} T_{\text{sim}} \quad \text{and} \quad X_\nu \stackrel{d}{=} X_{\text{sim}}, \quad (127)$$

707 and H_{sim} is the heap index of $(T_{\text{sim}}, X_{\text{sim}})$ in \mathcal{T}_{sim} .

708 *Proof.* First, observe that when $\text{split}_{\text{target}} = \text{split}_{\text{sim}}$, Algorithm 7 collapses onto just the extended
709 BSP tree construction for Π and simply returns the arrival with the given heap index H_{target} in $\mathcal{T}_{\text{target}}$.
710 In particular, the inner loop will always exit after one iteration, and every time one and only one of the
711 conditional blocks will be executed. In other words, in this case, the algorithm becomes equivalent to
712 Algorithm 2 in Flamich et al. (2022).

713 Let us now consider the case when $\text{split}_{\text{target}} \neq \text{split}_{\text{sim}}$. Denote the depth of the required node by
714 $D_\nu = \lfloor \log_2 H_\nu \rfloor$. Now, we give an inductive argument for correctness.

715 *Base case:* Consider $D_\nu = 0$. In this case, the target bounds $B = \Omega$, and the first sample we draw
716 $X \sim P$ is guaranteed to fall in Ω . Hence, for $D_\nu = 0$ the outer loop only runs for one iteration.
717 Furthermore, during that iteration, the inner loop will also exit after one iteration, and Algorithm 7
718 returns the sample (T, X) sampled before the outer loop with heap index 1. Since $T \sim \text{Exp}(1)$ and
719 $X \sim P$, this will be a correctly distributed output with the appropriate heap index.

720 *Inductive hypothesis:* Assume Algorithm 7 is correct heap indices with depths up to $D_\nu = d$.

721 *Case $D_\nu = d + 1$:* Let $\rho \in \mathcal{T}_{\text{target}}$ be the parent node of ν with arrival (T_ρ, X_ρ) . Then, $D_\rho = d$, hence
722 by the inductive hypothesis, Algorithm 7 will correctly simulate a branch $\mathcal{T}_{\text{target}}$ up to node ρ . At the
723 end of the d th iteration of the outer loop Algorithm 7 sets the target bounds $B \leftarrow B_\nu$. Then, in the
724 final, $d + 1$ st iteration, the inner loop simply realizes \mathcal{T}_{sim} and accepts the first sample after X that
725 falls inside B_ν whose time index $T > T_\rho$. Due to the priority queue, the loop simulates the nodes
726 of \mathcal{T}_{sim} in time order; hence the accepted sample will also be the one with the earliest arrival time.
727 Furthermore, Algorithm 7 only ever considers nodes of \mathcal{T}_{sim} whose bounds intersect the target bounds
728 B_ν , hence the inner loop is guaranteed to terminate, which finishes the proof. \square

729 E GPRS with Binary Search

730 We now utilise the machinery we developed in Appendix D to analyze Algorithm 5.

731 **Correctness of Algorithm 5:** Observe that Algorithm 5 constructs the extended BSP tree for the
732 on-sample splitting function. Thus, we will now focus on performing a binary tree search using
733 the extended BSP tree induced by the on-sample splitting function, which we denote by \mathcal{T} . The
734 first step of the algorithm matches GPRS's first step (Algorithm 3). Hence it is correct for the first
735 step. Now consider the algorithm in an arbitrary step k before termination, where the candidate
736 sample (T, X) is rejected, i.e. $T > \varphi(X)$. By assumption, the density ratio r is unimodal, and
737 since σ is monotonically increasing, $\varphi = \sigma \circ r$ is unimodal too. Thus, let $x^* \in \Omega$ be such that
738 $r(x^*) = r^*$, where $r^* = \exp_2(D_\infty[Q||P])$. Assume for now that $X < x^*$, the case $X > x^*$ follows
739 by a symmetric argument. By the unimodality assumption, since $X < x^*$, it must hold that for all
740 $y < X$, we have $\varphi(y) < \varphi(X)$. Consider now the arrival (T_L, X_L) of Π in the current node's left
741 child. Then, we will have $T < T_L$ and $X_L < X$ by construction. Thus, finally, we get

$$\varphi(X_L) < \varphi(X) < T < T_L, \quad (128)$$

742 meaning that the current node's left child is also guaranteed to be rejected. This argument can
743 be easily extended to show that any left-descendant of the current node will be rejected, and it is
744 sufficient to search its right-descendants only. By a similar argument, when $X > x^*$, we find that it
745 is sufficient only to check the current node's left-descendants. Finally, since both algorithms simulate
746 Π and search for its first arrival under φ , by the construction in Appendix A, the sample returned by
747 both algorithms will follow the desired target distribution.

748 **Expected runtime and codelength:** Since Algorithm 5 simulates a single branch of $\mathcal{T}_{\text{on-sample}}$, its
749 runtime is equal to the runtime of simulating that single branch. Assume that the accepted sample's
750 time index is N and its depth is D . Then, since Algorithm 5 draws one sample per depth, its runtime
751 will be exactly D steps. Then, by putting together lemmas D.2 and D.3, we get

$$\mathbb{E}_{D|N}[D] \leq \log_2 N. \quad (129)$$

752 Then, by taking expectation over N , by Equation (79) we get

$$\mathbb{E}[D] \leq D_{\text{KL}}[Q||P] + 2 \cdot \log_2 e < D_{\text{KL}}[Q||P] + 3, \quad (130)$$

753 thus in the one-shot case, the runtime of Algorithm 5 is linear in the KL divergence, which establishes
754 Theorem 3.5.

755 For the codelength result, let $\mathbf{x}, \mathbf{y} \sim P_{\mathbf{x}, \mathbf{y}}$, fix \mathbf{y} and set $Q \leftarrow P_{\mathbf{x}|\mathbf{y}}$ and $P \leftarrow P_{\mathbf{x}}$. Let H denote
756 the heap index of the sample returned by Algorithm 5, and D its depth. Recall, that $\lfloor \log_2 H \rfloor = D$,
757 hence by Equation (130), we get

$$\mathbb{E}_H[\log_2 H] < \mathbb{E}_H[\lfloor \log_2 H \rfloor] + 1 = \mathbb{E}_D[D] + 1 \leq D_{\text{KL}}[P_{\mathbf{x}|\mathbf{y}} \| P_{\mathbf{x}}] + 1 + 2 \log_2 e. \quad (131)$$

758 Taking expectation over \mathbf{y} , we get

$$\mathbb{E}[\log_2 H] \leq I[\mathbf{x}; \mathbf{y}] + 1 + 2 \log_2 e. \quad (132)$$

759 Therefore, using similar arguments to the ones in Appendix B.3, we use a Zeta distribution to encode
760 H with

$$\lambda = 1 + \frac{1}{I[\mathbf{x}; \mathbf{y}] + 1 + 2 \log_2 e}. \quad (133)$$

761 This gives an upper bound of

$$\mathbb{H}[\mathbf{x} | \Pi] \leq \mathbb{H}[H] \quad (134)$$

$$\leq I[\mathbf{x}; \mathbf{y}] + \log_2(I[\mathbf{x}; \mathbf{y}] + 2 \log_2 e + 2) + 2 \log_2 e + 2 \quad (135)$$

$$\leq I[\mathbf{x}; \mathbf{y}] + \log_2(I[\mathbf{x}; \mathbf{y}] + 1) + 2 \log_2 e + 2 + \log_2(2 \log_2 e + 2) \quad (136)$$

$$\leq I[\mathbf{x}; \mathbf{y}] + \log_2(I[\mathbf{x}; \mathbf{y}] + 1) + 8 \quad (137)$$

762 F General GPRS with Binary Search

763 We finally present a generalized version of branch-and-bound GPRS (Algorithm 6) for more general
764 spaces and remove the requirement that r be unimodal.

765 **Decomposing a Poisson process into a mixture of processes:** Let Π be a process over $\mathbb{R}^+ \times \Omega$
766 as before, with spatial measure P , and let Q be a target measure, $\varphi = \sigma \circ r$ and $U = \{(t, x) \in$
767 $\mathbb{R}^+ \times \Omega \mid t \leq \varphi(x)\}$ as before. Let $\tilde{\Pi}$ be $\Pi = \Pi \cap U$ restricted under φ and $(\tilde{T}_1, \tilde{X}_1)$ its first
768 arrival. Let $\{L, R\}$ form an P -essential partition of Ω , i.e. $L \cap R = \emptyset$ and $P(L \cup R) = 1$, and let
769 $\Pi_L = \Pi \cap \mathbb{R}^+ \times L$ and $\Pi_R = \Pi \cap \mathbb{R}^+ \times R$ be the restriction of Π to L and R , respectively. Let
770 $\tilde{\Pi}_L = \Pi_L \cap U$ and $\tilde{\Pi}_R = \Pi_R \cap U$ be the restrictions of the two processes under φ as well. Let $\tilde{\mu}_L(t)$
771 and $\tilde{\mu}_R(t)$ be the mean measures of these processes. Thus, the first arrival times in these processes
772 have survival functions

$$\mathbb{P}[\tilde{T}_1^L \geq t] = e^{-\tilde{\mu}_L(t)} \quad (138)$$

$$\mathbb{P}[\tilde{T}_1^R \geq t] = e^{-\tilde{\mu}_R(t)}. \quad (139)$$

773 Note, that by the superposition theorem, $\Pi_L \cup \Pi_R = \Pi$, hence the first arrival $(\tilde{T}_1, \tilde{X}_1)$ occurs in
774 either Π_L or Π_R . Assume now that we have already searched through the points of Π up to time τ
775 without finding the first arrival. At this point, we can ask: will the first arrival occur in Π_L , given that
776 $\tilde{T}_1 \geq \tau$? Using Bayes' rule, we find

$$\mathbb{P}[\tilde{X}_1 \in L \mid \tilde{T}_1 \geq \tau] = \frac{\mathbb{P}[\tilde{X}_1 \in L, \tilde{T}_1 \geq \tau]}{\mathbb{P}[\tilde{T}_1 \geq \tau]}. \quad (140)$$

777 More generally, assume that the first arrival of Π occurs in some set $A \subseteq \Omega$, and we know that the
778 first arrival time is larger than τ . Then, what is the probability that the first arrival occurs in some set
779 $B \subseteq A$? Similarly to the above, we find

$$\mathbb{P}[\tilde{X}_1 \in B \mid \tilde{T}_1 \geq \tau, \tilde{X}_1 \in A] = \frac{\mathbb{P}[\tilde{X}_1 \in B, \tilde{T}_1 \geq \tau, \tilde{X}_1 \in A]}{\mathbb{P}[\tilde{T}_1 \geq \tau, \tilde{X}_1 \in A]} \quad (141)$$

$$= \frac{\mathbb{P}[\tilde{T}_1 \geq \tau, \tilde{X}_1 \in B]}{\mathbb{P}[\tilde{T}_1 \geq \tau, \tilde{X}_1 \in A]}. \quad (142)$$

780 Let $(\tilde{T}_L, \tilde{X}_L)$ be the first arrival of Π_L . Then, the crucial observation is that

$$(\tilde{T}_L, \tilde{X}_L) \stackrel{d}{=} \tilde{T}_1, \tilde{X}_1 \mid \tilde{X}_1 \in L. \quad (143)$$

781 This enables us to search for the first arrival of Π under the graph of φ using an extended BSP tree
782 construction. At each node, if we reject, we draw a Bernoulli random variable b with mean equal
783 to the probability that the first arrival occurs within the bounds associated with the right child node.
784 Then, if $b = 1$, we continue the search along the right branch. Otherwise, we search along the left
785 branch.

786 Note, however, that in a restricted process Π_A , the spatial measure no longer integrates to 1. Further-
787 more, our target Radon-Nikodym derivative is $r(x) \cdot \mathbf{1}[x \in A]/Q(A)$. This means we need to change
788 the graph φ to some new graph φ_A to ensure that the spatial distribution of the returned sample is
789 still correct. Therefore, for a set A we define the restricted versions of previous quantities:

$$\tilde{\mu}_A(t) \stackrel{\text{def}}{=} \int_0^t \int_A \mathbf{1}[\tau \leq \varphi_A(x)] dP(x) dt \quad (144)$$

$$w_P(h | A) \stackrel{\text{def}}{=} \int_A \mathbf{1}[h \leq r(x)] dP(x) \quad (145)$$

790 Then, via analogous arguments to the ones in Appendix A, we find

$$\frac{d\mathbb{P}[\tilde{T}_A = t, \tilde{X}_A = x]}{d(\lambda \times P)} = \mathbf{1}[x \in A] \mathbf{1}[t \leq \varphi_A(x)] \mathbb{P}[\tilde{T}_A \geq t] \quad (146)$$

$$\frac{d\mathbb{P}[\tilde{X}_A = x]}{dP} = \mathbf{1}[x \in A] \int_0^{\varphi_A(x)} \mathbb{P}[\tilde{T}_A \geq t] dt \quad (147)$$

$$\varphi_A = \sigma_A \circ r. \quad (148)$$

791 Similarly, setting $\frac{d\mathbb{P}[\tilde{X}_A = x]}{dP} = \mathbf{1}[x \in A] \cdot r(x)/Q(A)$, and setting $\tau = \sigma_A(r(x))$, we get

$$\sigma_A^{-1}(\tau) = Q(A) \int_0^\tau \mathbb{P}[\tilde{T}_A \geq t] dt \quad (149)$$

$$\Rightarrow (\sigma_A^{-1})'(\tau) = \mathbb{P}[\tilde{T}_A \geq \tau] \quad (150)$$

$$= \mathbb{P}[\tilde{T} \geq \tau, \tilde{X} \in A]. \quad (151)$$

792 From this, again using similar arguments to the ones in Appendix A, we find

$$(\sigma_A^{-1})'(\tau) = w_Q(\sigma_A^{-1}(t) | A) - \sigma_A^{-1}(t) \cdot w_P(\sigma_A^{-1}(t) | A). \quad (152)$$

793 G Necessary Quantities for Implementing GPRS in Practice

794 Ultimately, given a target-proposal pair (Q, P) with density ratio r , we would want an easy-to-
795 evaluate expression for the appropriate stretch function σ or σ^{-1} to plug directly into our algorithms.
796 Computing σ requires computing the integral in Equation (3) and finding σ^{-1} requires solving the
797 non-linear ODE in Equation (2), neither of which is usually possible in practice. Hence, we usually
798 resort to computing σ^{-1} numerically by using an ODE solver for Equation (2).

799 In any case, we need to compute w_P and w_Q , which are analytically tractable in the practically
800 relevant cases. Hence, in this section, we give closed-form expressions for w_P and w_Q for all the
801 examples we consider and give closed-form expressions for σ and σ^{-1} for some of them. If we do
802 not give a closed-form expression of σ , we use numerical integration to compute σ^{-1} instead.

803 G.1 Uniform-Uniform Case

804 Let P be the uniform distribution over an arbitrary space Ω and Q a uniform distribution supported
805 on some subset $\mathcal{X} \subset \Omega$, with $P(\mathcal{X}) = C$ for some $C \leq 1$. Then,

$$r(x) = \frac{1}{C} \cdot \mathbf{1}[x \in \mathcal{X}] \quad (153)$$

$$w_P(h) = C \quad (154)$$

$$w_Q(h) = 1 \quad (155)$$

$$\sigma(h) = -\frac{1}{C} \log(1 - C \cdot h) \quad (156)$$

$$\sigma^{-1}(t) = \frac{1}{C} (1 - \exp(-C \cdot h)). \quad (157)$$

806 Note that using GPRS in the uniform-uniform case is somewhat overkill, as in this case, it is simply
 807 equivalent to standard rejection sampling.

808 G.2 Triangular-Uniform Case

809 Let $P = \text{Unif}(0, 1)$ and for some numbers $0 < a < c < b < 1$, let Q be the triangular distribution,
 810 defined by the PDF

$$q(x) = \begin{cases} 0 & \text{if } x < a \\ \frac{2(x-a)}{(b-a)(c-a)} & \text{if } a \leq x < c \\ \frac{2}{b-a} & \text{if } x = c \\ \frac{2(b-x)}{(b-a)(b-c)} & \text{if } c < x \leq b \\ 0 & \text{if } b < x. \end{cases} \quad (158)$$

811 For convenience, let $\ell = b - a$. Then,

$$r(x) = q(x) \quad (159)$$

$$w_P(x) = \ell - \frac{\ell^2}{2} \cdot h \quad (160)$$

$$w_Q(x) = 1 - \frac{\ell^2}{4} \cdot h^2 \quad (161)$$

$$\sigma(h) = \frac{2h}{2 - \ell \cdot h} \quad (162)$$

$$\sigma^{-1}(t) = \frac{2t}{2 + \ell \cdot t} \quad (163)$$

812 G.3 Finite Discrete-Discrete / Piecewise Constant Case

813 Without loss of generality, let Q be a discrete distribution over a finite set with K elements defined
 814 by the probability vector $(r_1 < r_2 < \dots < r_K)$ and $P = \text{Unif}([1 : K])$ the uniform distribution on
 815 K elements. Note that any discrete distribution pair can be rearranged into this setup. Now, we can
 816 embed this distribution into $[0, 1]$ by mapping P to $\hat{P} = \text{Unif}(0, 1)$ the uniform distribution over
 817 $[0, 1]$ and mapping Q to \hat{Q} with density

$$\hat{q}(x) = \sum_{k=1}^K \mathbf{1}[\lfloor K \cdot x \rfloor = k - 1] \cdot K \cdot r_k. \quad (164)$$

818 Then, we can draw a sample x from Q by simulating a sample $\hat{x} \sim \hat{Q}$ and computing
 819 $x = K \cdot \lfloor \hat{x}/K \rfloor + 1$. Hence, we can instead reduce the problem to sampling from a piecewise
 820 constant distribution. Thus, let us now instead present the more general case of arbitrary piecewise
 821 constant distributions over $[0, 1]$, with \hat{Q} defined by probabilities $(q_1 < q_2 < \dots < q_K)$ and corre-
 822 sponding piece widths $(w_1 < w_2 < \dots < w_K)$. We require, that $\sum_{k=1}^K q_k = \sum_{k=1}^K w_k = 1$. Then,
 823 the density is

$$\tilde{q}(x) = \sum_{k=1}^K \mathbf{1} \left[\sum_{j=1}^{k-1} w_j \leq x \leq \sum_{j=1}^k w_j \right] \cdot \frac{q_k}{w_k} \quad (165)$$

824 Define $r_k = q_k/w_k$. Then,

$$w_P(h) = \sum_{k=1}^K \mathbf{1}[h \leq r_k] \cdot w_k \quad (166)$$

$$w_Q(h) = \sum_{k=1}^K \mathbf{1}[h \leq r_k] \cdot w_k \cdot r_k \quad (167)$$

$$\sigma(h) = \sum_{k=1}^K \mathbf{1}[h \geq r_{k-1}] \cdot \frac{1}{B_k} \cdot \log \left(\frac{A_k - B_k r_{k-1}}{A_k - B_k \min\{r_k, h\}} \right) \quad (168)$$

825 where we defined

$$A_k = \sum_{j=k}^K w_j r_j = \sum_{j=k}^K q_j \quad (169)$$

$$B_k = \sum_{j=k}^K w_j \quad (170)$$

826 G.4 Diagonal Gaussian-Gaussian Case

827 Without loss of generality, let $Q = \mathcal{N}(\mu, \sigma^2 I)$ and $P = \mathcal{N}(0, I)$ be d -dimensional Gaussian
 828 distributions with diagonal covariance. As a slight abuse of notation, let $\mathcal{N}(x \mid \mu, \sigma^2 I)$ denote the
 829 probability density function of a Gaussian random variable with mean μ and covariance $\sigma^2 I$ evaluated
 830 at x . Then, when $\sigma^2 < 1$, we have

$$r(x) = Z \cdot \mathcal{N}(x \mid m, s^2 I) \quad (171)$$

$$m = \frac{\mu}{1 - \sigma^2} \quad (172)$$

$$s^2 = \frac{\sigma^2}{1 - \sigma^2} \quad (173)$$

$$Z = \frac{(1 - \sigma^2)^{-d}}{\mathcal{N}(\mu \mid 0, (1 - \sigma^2)I)} \quad (174)$$

$$w_P(h) = \mathbb{P}[\chi^2(d, \|m\|^2) \leq -2s^2 \ln h + C] \quad (175)$$

$$w_Q(h) = \mathbb{P}\left[\chi^2\left(d, \left\|\frac{m - \mu}{s}\right\|^2\right) \leq -2s^2 \ln h + C\right] \quad (176)$$

$$C = s^2 (2 \ln Z - d \ln(2\pi s^2)). \quad (177)$$

831 Unfortunately, in this case, it is unlikely that we can solve for the stretch function analytically, so in
 832 our experiments, we solved for it numerically using Equation (2).

833 H Experimental Details

834 **Comparing Algorithm 3 versus Global-bound A* coding:** We use a setup similar to the one used
 835 by Theis & Yosri (2022). Concretely, we assume the following model for correlated random variables
 836 \mathbf{x}, μ :

$$P_\mu = \mathcal{N}(0, 1) \quad (178)$$

$$P_{\mathbf{x}|\mu} = \mathcal{N}(\mu, \sigma^2). \quad (179)$$

837 From this, we find that the marginal on \mathbf{x} must be $P_{\mathbf{x}} = \mathcal{N}(0, \sigma^2 + 1)$. The mutual information is
 838 $I[\mathbf{x}; \mu] = \frac{1}{2} \log_2(1 + \sigma^2)$ bits, which is what we plot as $I[\mathbf{x}; \mu]$ in the top two panels in Figure 2.

839 For the bottom panel in Figure 2, we fixed a standard Gaussian prior $P = \mathcal{N}(0, 1)$, fixed $K =$
 840 $D_{\text{KL}}[Q\|P]$ and linearly increased $R = D_\infty[Q\|P]$. To find a target that achieves the desired values
 841 for these given divergences, we set its mean and variances as

$$\sigma^2 = \exp(W(A \cdot \exp(B))) - B \quad (180)$$

$$\mu = 2K - \sigma^2 + \ln \sigma^2 + 1 \quad (181)$$

$$A = 2 \ln R - 2K - 1 \quad (182)$$

$$B = 2 \ln R - 1, \quad (183)$$

842 where W is the principal branch of the Lambert W -function (Corless et al., 1996).

843 We can derive this formula by assuming we wish to find μ and σ^2 such that for fixed numbers K and
 844 R , and $q(x) = \mathcal{N}(x \mid \mu, \sigma^2)$, $p(x) = \mathcal{N}(x \mid 0, 1)$. Then, we have that

$$D_{\text{KL}}[q\|p] = K \quad \text{and} \quad \sup_{x \in \mathbb{R}} \left\{ \frac{q(x)}{p(x)} \right\} = R. \quad (184)$$

845 We know that

$$K = D_{\text{KL}}[q\|p] = \frac{1}{2} [\mu^2 + \sigma^2 - \log \sigma^2 - 1]$$

$$\log R = \log \sup_{x \in \mathbb{R}} \left\{ \frac{q(x)}{p(x)} \right\} = \frac{\mu^2}{2(1 - \sigma^2)} - \log \sigma. \quad (185)$$

846 From these, we get that

$$\begin{aligned} \mu^2 &= 2K - \sigma^2 + \log \sigma^2 + 1 \\ \mu^2 &= 2(1 - \sigma^2)(\log R + \log \sigma) \end{aligned} \quad (186)$$

847 Setting these equal to each other

$$\begin{aligned} 2K - \sigma^2 + \log \sigma^2 + 1 &= 2 \log R + \log \sigma^2 - 2\sigma^2 \log R - \sigma^2 \log \sigma^2 \\ \sigma^2 \log \sigma^2 - \sigma^2 + 2\sigma^2 \log R &= 2 \log R - 2K - 1 \\ \sigma^2 \log \sigma^2 + \sigma^2(2 \log R - 1) &= A \\ \sigma^2(\log \sigma^2 + B) &= A \\ \sigma^2 \log(\sigma^2 e^B) &= A \\ e^B \sigma^2 \log(\sigma^2 e^B) &= A e^B \\ e^{\log(\sigma^2 e^B)} \log(\sigma^2 e^B) &= A e^B \\ \log(\sigma^2 e^B) &= W(A e^B) \\ \sigma^2 &= e^{W(A e^B) - B}, \end{aligned} \quad (187)$$

848 where we made the substitutions $A = 2 \log R - 2K - 1$ and $B = 2 \log R - 1$.

849 I Rejection sampling index entropy lower bound

850 Assume that we have a pair of correlated random variables $\mathbf{x}, \mathbf{y} \sim P_{\mathbf{x}, \mathbf{y}}$ and Alice and Bob wish to
 851 realize a channel simulation protocol using standard rejection sampling as given by, e.g. Algorithm 2.
 852 Thus, when Alice receives a source symbol $\mathbf{y} \sim P_{\mathbf{y}}$, she sets $Q = P_{\mathbf{x}|\mathbf{y}}$ as the target and $P = P_{\mathbf{x}}$ as
 853 the proposal for the rejection sampler. Let N denote the index of Alice's accepted sample, which is
 854 also the number of samples she needs to draw before her algorithm terminates. Since each acceptance
 855 decision is an independent Bernoulli trial in standard rejection sampling, N follows a geometric
 856 distribution whose mean equals the upper bound M used for the density ratio (Maddison, 2016). The
 857 lower bound on the optimal coding cost for N is given by its entropy

$$\mathbb{H}[N] = -(M - 1) \log_2 \left(1 - \frac{1}{M} \right) + \log_2 M \geq \log_2 M, \quad (188)$$

858 where the inequality follows by omitting the first term, which is guaranteed to be positive since
 859 $x \mapsto -x \log_2 x$ is positive on $(0, 1)$. Hence, by using the optimal upper bound on the density ratio
 860 $M^* = \exp_2(D_\infty[Q\|P])$ and plugging it into the formula above, we find that

$$D_\infty[Q\|P] \leq \mathbb{H}[N]. \quad (189)$$

861 Now, taking expectation over \mathbf{y} , we find

$$\mathbb{E}_{\mathbf{y} \sim P_{\mathbf{y}}} [D_\infty[P_{\mathbf{x}|\mathbf{y}}\|P_{\mathbf{x}}]] \leq \mathbb{H}[N]. \quad (190)$$