
The Monte Carlo Method and New Device and Architectural Techniques for Accelerating It

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Computing systems interacting with real-world processes must safely and reliably
2 process uncertain data. The Monte Carlo method is a popular approach for comput-
3 ing with such uncertain values. This article introduces a framework for describing
4 the Monte Carlo method and highlights two advances in the domain of physics-
5 based non-uniform random variate generators (PPRVGs) to overcome common
6 limitations of traditional Monte Carlo sampling. This article also highlights recent
7 advances in architectural techniques that eliminate the need to use the Monte Carlo
8 method by leveraging distributional microarchitectural state to natively compute
9 on probability distributions. Unlike Monte Carlo methods, uncertainty-tracking
10 processor architectures can be said to be *convergence-oblivious*.

11 1 Introduction

12 Uncertainty arises when systems carry out computations based on *measurements* (aleatoric uncer-
13 tainty) or *limited knowledge* (epistemic uncertainty). Uncertainty introduces risk to actions taken
14 based on measurements or limited knowledge. Studying and quantifying how uncertainty propagates
15 through computations is a requirement when making principled decisions about the suitability of an
16 uncertain system for an application.

17 Despite the importance of quantifying and understanding uncertainty, computer architectures and
18 circuit implementations lack numerically-robust and computationally-efficient methods to program-
19 matically process and reason about uncertainty. State-of-the-art techniques often employ the Monte
20 Carlo method [1, 2, 3, 4] to estimate the effect of long sequences of arithmetic operations on inputs
21 that are uncertain when closed-form propagation of uncertainty is not possible. Monte-Carlo-based
22 methods can be sample-inefficient: the variance in the result of Monte Carlo integration using n
23 samples scales as $\frac{1}{\sqrt{n}}$ [3]. This means that if we wanted to halve the variance, we would need to
24 *quadruple* the number of samples.

25 This article presents a framework for describing Monte-Carlo-based methods (Section 2). The
26 framework poses them as the application of three steps: sampling, evaluation, and post-processing. In
27 Section 3 we describe recent advances in physics-based programmable non-uniform random variate
28 generators (PPRVGs) which can improve the sampling phase of Monte Carlo methods. Section 4
29 shows how a novel uncertainty-tracking microarchitecture, Laplace [5, 6], can provide a more efficient
30 way to represent and compute on uncertain variables. Section 6 compares the performance of Laplace
31 to the traditional Monte Carlo method.

32 2 The Monte Carlo Method

33 The phrase *Monte Carlo method* refers to a wide class of computational methods that sample from
34 random variables to calculate solutions to computational problems. The earliest example of the use
35 of a Monte Carlo method is attributed to Georges-Louis LeClerc [4], Comte de Buffon, who, in the
36 eighteenth century, simulated a value of π by dropping needles onto a lined background. He showed
37 that when the needle has the same length as the distance between parallel lines, the probability that a
38 randomly-thrown needle will overlap with a line is $\frac{2}{\pi}$. Therefore, π can be estimated by throwing a
39 large number of needles and averaging the number of times they overlap with a line.

40 2.1 The Monte Carlo Method: Sampling, Evaluation, and Post-Processing

41 The Monte Carlo method approximates a desired numerical property of the outcome of transforma-
42 tions of random variables. Practitioners use the Monte Carlo method when the desired property is not
43 available analytically or because the analytical solution is computationally expensive. The desired
44 property could be the expectation of the resulting random variable (*Monte Carlo integration*), a
45 sample from it (*Monte Carlo sampling*), or its probability density function (*Monte Carlo simulation*).

46 Suppose that we want to obtain a property from the random variable Y that is defined by transforming
47 the random variable X using the transformation $f : X \rightarrow Y$ (i.e., $Y = f(X)$). We summarize the
48 steps of the Monte Carlo method to approximate the desired numerical properties as follows:

49 1. **Sampling:** The Monte Carlo method first generates i.i.d. samples from X . Let n denote the
50 number of samples of the random variable X in the set $\{x_i\}_{i=1}^n$. This step typically uses a
51 random number generator program running on a computer that can generate pseudo-random
52 numbers from a uniform distribution. Samples from more complex random variables are
53 generated using *Monte Carlo sampling*, where the Monte Carlo method itself is used to
54 generate samples by transforming the uniform random variates. Examples of Monte Carlo
55 sampling include the Box-Muller method [7] for generating standard Gaussian samples,
56 inverse transform sampling for sampling from random variables for which an inverse
57 cumulative distribution function (ICDF) exists¹, and Markov Chain Monte Carlo (MCMC)
58 for more complex random variables [2].

59 As an alternative to Monte Carlo sampling, we can use physical hardware to efficiently
60 sample from a non-uniform random variable. Section 3 presents several such methods from
61 the research literature which can provide large-batch single-shot convergence-oblivious
62 random variate generation by exploiting physical processes that generate *non-uniform*
63 entropy and can be sampled in parallel.

64 2. **Evaluation:** The second step of the Monte Carlo method then evaluates the transformation
65 f on the set of samples $\{x_i\}_{i=1}^n$ to obtain a set $\{y_i\}_{i=1}^n$ of n samples of Y , where each
66 $y_i = f(x_i)$. This step is called *Monte Carlo evaluation*.

67 In the Monte Carlo method, the evaluation step is carried out on each sample x_i , one at a
68 time. Section 4 presents recent research on computer architectures that can process compact
69 representations of entire distributions at once, rather than one sample at a time as is the case
70 for the traditional Monte Carlo method.

71 3. **Post-processing:** In the third and final step, the Monte Carlo method approximates the
72 desired numerical property from the samples $\{y_i\}_{i=1}^n$ by applying an operation on their
73 set. For example, taking their average (as in the case of Monte Carlo integration), applying
74 the identity function (as in Monte Carlo sampling), or generating a representation of the
75 probability density function, such as a histogram (as in Monte Carlo simulation).

76 3 Physics-Based Programmable Non-Uniform Random Variate Generation

77 Section 2 described the sampling of (possibly non-uniform) random variables as the first step of the
78 Monte Carlo method. Most computing systems use pseudo-random number generators to generate
79 uniform random variates. Computers generate samples from non-uniform random variables by using
80 Monte Carlo sampling (Section 2). Since Monte Carlo methods could require large numbers of
81 samples, these methods can be computationally-expensive and can lead to a significant overhead.

¹Leemis *et al* [8] provides a good source of relationships between univariate random variables.

82 Two recent methods of generating non-uniform random variates from physical processes, Spot [9]
 83 and Grappa [10] have the following key features:

- 84 • They can *efficiently* generate *non-uniform* random variates: Spot, for example, can generate
 85 Gaussian random variables $260\times$ faster than the Box-Muller transformation running on an
 86 ARM Cortex-M0+ microcontroller, while dissipating less power than such a microcontroller.
- 87 • They are *physics-based*: Spot generates random variates using electron tunneling noise,
 88 while Grappa exploits the transfer characteristics of Graphene field-effect transistors.
- 89 • They are *programmable*: The distributions from which they can sample from are not fixed;
 90 their host systems can dynamically and digitally configure them to produce samples from a
 91 required probability distribution.

92 Due to these features, we call methods such as Spot and Grappa physics-based programmable
 93 non-uniform random variate generators (PPRVGs).

94 **Spot:** Spot is a method for generating random numbers by sampling a one-dimensional distribution
 95 associated with a Gaussian voltage noise source [11]. Using an analog-to-digital converter (ADC),
 96 Spot takes measurements of a physical process that generates Gaussian noise. Spot then maps this
 97 physically-generated univariate Gaussian to any other univariate Gaussian using only two operations:
 98 a multiplication and an addition [9]. Samples from any other non-uniform random variable are
 99 generated by creating a mixture of Gaussians.

100 **Grappa:** Grappa is a Graphene Field-Effect Transistor (GFET)-based programmable analog func-
 101 tion approximation architecture [12]. Grappa relies on the non-linear transfer characteristics of
 102 GFETs to transform a uniform random sample into a non-uniform random sample [12].

103 Grappa implements a linear least-squares Galerkin approximation [13] to approximate the ICDF
 104 of a target distribution and carry out inverse transform sampling. The required orthonormal basis
 105 functions are obtained from the GFET transfer characteristics using the Gram-Schmidt process [14].

106 Tye *et al.* showed that Monte Carlo integration using samples generated by Grappa is at least $1.26\times$
 107 faster than using a C++ lognormal random number generator. Subsequent work [12] demonstrated an
 108 average speedup of up to $2\times$ compared to MATLAB for lognormal, exponential, generalized Pareto,
 109 and Gaussian mixture distributions, with the execution time independent of the target distribution.

110 4 Beyond the Monte Carlo Method

111 Let X be a random variable and $f : X \rightarrow Y$ be a transformation of X . Denoting the resulting
 112 random variable as $Y = f(X)$, from the change of variable formula for random variables (Theorem
 113 1 in Appendix A), we obtain probability density function p_Y of Y . If p_X is the probability density
 114 function of X , then the probability density function of p_Y of Y is,

$$p_Y(y) = p_X \circ f^{-1}(y) |\det \nabla f^{-1}(y)|, \quad (1)$$

115 where $y \in Y$, and $\nabla_y f^{-1}$ is the Jacobian matrix. Using the change of variables technique of integrals,
 116 we obtain

$$\begin{aligned} \mathbb{E}_{p_X}[f(X)] &= \int_X f(x)p_X(x) dx && \text{by Equation 7 in Appendix A} \\ &= \int_Y yp_Y \circ f^{-1}(y) |\det \nabla f^{-1}(y)| dy && \text{by change of variables (integration)} \\ &= \int_Y yp_Y(y) dy && \text{by Theorem 1 in Appendix A} \\ &= \mathbb{E}_{p_Y}[Y]. \end{aligned}$$

117 Thus, if we had access to p_Y of Y , we can evaluate $\mathbb{E}_{p_X}[f(X)]$ by taking the expectation of the
 118 random variable Y with respect to the p_Y . When p_Y isn't directly accessible, we usually obtain the
 119 expectation of Y by using Monte Carlo integration. However, having access to p_Y would eliminate
 120 the need to use the Monte Carlo method completely.

121 Laplace [5, 6] is a computer microarchitecture that is capable of directly computing p_Y by representing
122 distributional information in its microarchitectural state and tracking how these distributions evolve
123 under arithmetic operations, transparently to the applications running on it. Laplace provides a
124 representation for the distribution (see Definition 3 in Appendix A) of random variables, and carries
125 out *deterministic computations on this distribution*.

126 Laplace’s in-processor distribution representation has an associated *representation size* that describes
127 the *precision* at which the probability distribution is represented. Higher values of the representa-
128 tion size result in a more accurate representation. A useful analogy is the IEEE-754 standard for
129 representing the uncountable infinite set of real numbers as floating-point numbers [15, 16] on a
130 finite-precision computer.

131 Computer architectures such as Laplace eliminate the need for using the Monte Carlo method and
132 can therefore have far-reaching consequences in areas where the Monte Carlo method is used. For
133 example, to approximate the predictive Gaussian Process posterior distribution with an uncertain
134 input, Deisenroth *et al* [17, 18] used moment-matching; Laplace could compute the posterior exactly,
135 up to the precision of the representation.

136 5 Methods

137 The remaining text compares and evaluates *Monte Carlo methods* and *Laplace-based methods*. Both
138 methods were evaluated on single-threaded applications written in the C programming language.

139 **Monte Carlo method:** We use the standard Monte Carlo method that we described in Section 2.
140 We use the pseudo-random number generator `rand` from the Standard C Library [19] to sample
141 from uniform distributions and use the modified Box-Muller method [20] as implemented by the
142 `gsl_ran_gaussian_ziggurat` function in the GNU Scientific Library [21]. We compile our code using `clang`,
143 the C family front-end to LLVM [22], with optimization set to `-O3`².

144 **Laplace:** We use Laplace as a replacement for the Monte Carlo method, as described in Section 4.
145 In our experiments, we exclusively use Laplace’s Telescopic Torques Representation (TTR) [5] as
146 provided by a commercial implementation of Laplace [23], release 2.6.

147 We compare these methods by empirically measuring and reporting the average *run time* and the
148 average *Wasserstein distance* [24] of the output to a ground truth in two different applications of
149 Monte Carlo simulation. We change the number of samples (for Monte-Carlo-based methods), or the
150 representation size (for Laplace-based methods) to observe the trade-offs between accuracy and run
151 time. For each configuration of number of samples or representation size, we repeat the experiments
152 30 times to account for variation in the process of sampling³. See Appendix C for more detail on our
153 methods. Figure 1 summarizes our results.

154 5.1 Applications

155 We carry out the experiments described above on two applications of Monte Carlo simulation.

156 **Monte Carlo Convergence Challenge Example:** Let X^{con} be the initial random variable that we
157 sample from, with its probability density function $p_{X^{\text{con}}}$ being a Gaussian mixture, given by:

$$p_{X^{\text{con}}}(x) = 0.6 \left(\frac{1}{0.5\sqrt{2\pi}} e^{-2(x-2)^2} \right) + 0.4 \left(\frac{1}{1.0\sqrt{2\pi}} e^{-\frac{(x+1)^2}{2}} \right). \quad (2)$$

158 For the Monte Carlo evaluation step of Section 2, we define a function f^{con} as a sigmoidal function:

$$f^{\text{con}}(x) = \frac{1}{1 + e^{-(x-1)}}. \quad (3)$$

²We will make all code necessary for exact replication of our experiments available through Github.

³We calculate the Wasserstein distance for Laplace’s representation of the output distributions by generating 1,000,000 samples from the representation. Therefore, we also repeat the Laplace experiments 30 times for each representation size even though Laplace’s uncertainty-tracking methods are deterministic and convergence-oblivious. The variation we see in the results in Figure 1 is therefore *only* due to sampling variance.

159 For the Traditional Monte Carlo method, we evaluated on $n \in \{4, 256, 1152, 2048, 4096, 8192,$
 160 $16000, 32000, 128000, 256000\}$. For Laplace, we evaluated on $r \in \{16, 32, 64, 256, 2048\}$.

161 **Poiseuille’s Law for Blood Transfusion** As a real-world application, we use Poiseuille’s Law, a
 162 mathematical model from fluid dynamics used to calculate the rate of laminar flow, Q , of a viscous
 163 fluid through a pipe of constant radius [25, 26]. This model is used in medicine as a simple method
 164 for approximating the rate of flow of fluids, such as blood, during transfusion [27, 28]. We look at
 165 Poiseuille’s Law applied to the case of blood transfusion using a pump with the following parameters:

- 166 ΔP : Pressure difference created by the pump, where $\Delta P \sim \mathcal{N}(5500000 \text{ mPa}, 36000^2)$.
- 167 μ : Viscosity of the fluid, where $\mu \sim \mathcal{U}(3.88 \text{ mPas}, 4.12 \text{ mPas})$.
- 168 l : Length of the tube from the cannula to the pump, where $l \sim \mathcal{U}(6.95 \text{ cm}, 7.05 \text{ cm})$.
- 169 r : Radius of the cannula, where $r \sim \mathcal{U}(0.0845 \text{ cm}, 0.0855 \text{ cm})$.

170 We assume the cannula to have a gauge of 14 (a radius of 0.85 mm) and the viscosity of blood to
 171 be 4 mPas [28]. [29] reported that for porcine blood, the uncertainty of using a ventricular assist
 172 device to measure blood viscosity in real time was ± 0.12 mPas; we use this as the uncertainty of the
 173 viscosity.

174 The flow rate Q is therefore measured in cm^3/s . Using these parameters, we can calculate the flow
 175 rate using Poiseuille’s Law:

$$Q = \frac{\pi r^4 \Delta P}{8 \mu l}. \quad (4)$$

176 For the Traditional Monte Carlo method, we evaluated on $n \in \{4, 256, 1152, 4096, 8192, 32000,$
 177 $128000, 256000, 512000, 640000\}$. For Laplace, we evaluated on $r \in \{16, 32, 64, 128, 256, 2048\}$.

178 6 Results


179 Figure 1 shows Pareto plots of the mean run time against the Wasserstein distance from the ground
 180 truth for both applications. A key observation is that the variance of the Laplace-based methods
 181 is more or less constant as we increase the representation size. Laplace carries out *deterministic*
 182 *computations on probability distributions*; this variance is caused by using a finite number of samples
 183 from Laplace’s representation of the output distribution to calculate the Wasserstein distance. It is
 184 possible to calculate the Wasserstein distance directly from the Laplace processor representation but
 185 we did not do so at the time of writing. This calculation would be deterministic since it only depends
 186 on the representation of the distribution. In contrast, each run of the Monte Carlo method results in a
 187 different output distribution; to reduce this variance we need to increase the number of samples. In
 188 this way, Laplace is *convergence-oblivious* to the number of samples.

189 Increasing the representation size larger than $r = 32$ provides a worse trade-off with the run time for
 190 both applications. Table 1 shows that for the accuracy obtained by Laplace, the equivalent Monte
 191 Carlo simulation is $113.85\times$ (for the Monte Carlo Convergence Challenge example) and $51.53\times$ (for
 192 the Poiseuille’s Law for Blood Transfusion application) slower. If much better accuracy is required,
 193 then the Monte Carlo method will need to be used. However, if the accuracy provided by Laplace
 194 is sufficient, it provides a potentially-orders-of-magnitude-faster alternative that is also *consistent*
 195 *outputs across repetitions*.

196 Tables with the numerical results are in Appendix F. Appendix F also compares histograms of the
 197 resulting distributions and provides additional discussion.

198 7 Conclusions

199 The Monte Carlo method is a powerful and historically-significant tool for solving complex problems
 200 that might otherwise be intractable. It involves three simple steps: sampling, evaluating and post-
 201 processing. Despite its versatility, the Monte Carlo method can suffer from inefficiencies. One of
 202 these is that generating samples for the first step of the Monte Carlo method is inefficient when
 203 samples are required from non-uniform probability distributions. Recent advances in physics-based
 204 random number generators, namely Spot [9, 11] and Grappa [12] address these challenges.



(a) Pareto plot for the Monte Carlo Convergence Challenge application from Section 5.1. We omitted $n = 4, 256$ for clarity; see Table 2 in Appendix F. (b) Pareto plot for the Poiseuille’s Law for Blood Transfusion application from Section 5.1. We have omitted $n = 4, 256$ for clarity.; see Table 3 in Appendix F.

Figure 1: Pareto plots between the mean run time, and the mean Wasserstein distance from the ground truth output distribution. The error bars show ± 1 standard deviation. For the Monte Carlo Convergence Challenge example (a), Traditional Monte Carlo obtains similar accuracy to Laplace with $r = 32$ at 32,000 samples. For the Poiseuille’s Law for Blood Transfusion application (b), Traditional Monte Carlo obtains similar accuracy than Laplace with $r = 32$ at 128,000 samples. In the legends, MC stands for *Traditional Monte Carlo, implemented in C*. We use the log scale on the vertical axis.

Problem	Core	Representation Size / Number of samples	Wasserstein Distance (mean \pm std. dev.)	Run time (ms) (mean \pm std. dev.)
Monte Carlo Convergence Challenge	Laplace	32	0.00167 \pm 0.00007	0.020 \pm 0.004
Monte Carlo Convergence Challenge	Traditional Monte Carlo	32000	0.00158 \pm 0.00068	2.277 \pm 0.346
Poiseuille’s Law for Blood Transfusion	Laplace	32	0.00033 \pm 0.00003	0.173 \pm 0.006
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	128000	0.00033 \pm 0.00009	8.914 \pm 1.566

Table 1: Results show the mean Wasserstein distance and the run time required the best overall configuration for Laplace and the close-to-equivalent results Monte Carlo configurations. For the Monte Carlo Convergence Challenge example, Traditional Monte Carlo takes approximately $113.85 \times$ longer than Laplace with $r = 32$. For the Poiseuille’s Law for Blood Transfusion application, Traditional Monte Carlo takes approximately $51.53 \times$ longer than Laplace with $r = 32$.

205 Techniques such as Laplace [5, 6] represent probability distributions in a computing system using
 206 an approximate fixed-bit-width representation in a manner analogous to how traditional computer
 207 architectures approximately represent real-valued numbers using fixed-bit-width representations
 208 such as the IEEE-754 floating-point [15, 16] representation. The computations of Laplace are
 209 approximations of explicit Monte Carlo methods in much the same way that computations on floating-
 210 point are approximations of arithmetic on real numbers. Laplace does not require iterative and
 211 repeated processing of samples until convergence to a target distribution is achieved, nor does it
 212 suffer from the high variance observed across Monte Carlo runs. Methods like Laplace are therefore
 213 *convergence-oblivious*.

214 References

215 [1] S. Rogers and M. Girolami, *A first course in machine learning*. Chapman and Hall/CRC, 2016.
 216 [2] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning*, vol. 4. Springer,
 217 2006.
 218 [3] S. Weinzierl, “Introduction to monte carlo methods,” *arXiv preprint hep-ph/0006269*, 2000.
 219 [4] R. L. Harrison, “Introduction to monte carlo simulation,” in *AIP conference proceedings*,
 220 vol. 1204, pp. 17–21, American Institute of Physics, 2010.

- 221 [5] V. Tsoutsouras, O. Kaparounakis, B. Bilgin, C. Samarakoon, J. Meech, J. Heck, and P. Stanley-
 222 Marbell, “The laplace microarchitecture for tracking data uncertainty and its implementation
 223 in a risc-v processor,” in *MICRO-54: 54th Annual IEEE/ACM International Symposium on*
 224 *Microarchitecture*, pp. 1254–1269, 2021.
- 225 [6] V. Tsoutsouras, O. Kaparounakis, C. Samarakoon, B. Bilgin, J. Meech, J. Heck, and P. Stanley-
 226 Marbell, “The laplace microarchitecture for tracking data uncertainty,” *IEEE Micro*, vol. 42,
 227 no. 4, pp. 78–86, 2022.
- 228 [7] G. E. Box, “A note on the generation of random normal deviates,” *Ann. Math. Statist.*, vol. 29,
 229 pp. 610–611, 1958.
- 230 [8] L. M. Leemis and J. T. McQueston, “Univariate distribution relationships,” *The American*
 231 *Statistician*, vol. 62, no. 1, pp. 45–53, 2008.
- 232 [9] J. T. Meech and P. Stanley-Marbell, “Efficient programmable random variate generation ac-
 233 celerator from sensor noise,” *IEEE Embedded Systems Letters*, vol. 13, no. 3, pp. 73–76,
 234 2020.
- 235 [10] N. J. Tye, J. T. Meech, B. A. Bilgin, and P. Stanley-Marbell, “A system for generating non-
 236 uniform random variates using graphene field-effect transistors,” in *2020 IEEE 31st International*
 237 *Conference on Application-specific Systems, Architectures and Processors (ASAP)*, pp. 101–108,
 238 IEEE, 2020.
- 239 [11] J. T. Meech, *Electron Tunneling Noise Based Programmable Random Variate Accelerator*. PhD
 240 thesis, Apollo - University of Cambridge Repository, 2023.
- 241 [12] N. Tye, *Analog-Domain Machine Learning Computations Exploiting 2D Materials Properties*.
 242 PhD thesis, Apollo - University of Cambridge Repository, 2022.
- 243 [13] A. Ern and J.-L. Guermond, *Theory and practice of finite elements*, vol. 159. Springer Science
 244 & Business Media, 2013.
- 245 [14] M. P. Deisenroth, A. A. Faisal, and C. S. Ong, *Mathematics for machine learning*. Cambridge
 246 University Press, 2020.
- 247 [15] IEEE, “IEEE standard for floating-point arithmetic,” *IEEE Std 754-2019 (Revision of IEEE*
 248 *754-2008)*, pp. 1–84, 2019.
- 249 [16] D. Goldberg, “What every computer scientist should know about floating-point arithmetic,”
 250 *ACM computing surveys (CSUR)*, vol. 23, no. 1, pp. 5–48, 1991.
- 251 [17] M. P. Deisenroth, *Efficient reinforcement learning using Gaussian processes*, vol. 9. KIT
 252 Scientific Publishing, 2010.
- 253 [18] M. Deisenroth and C. E. Rasmussen, “Pilco: A model-based and data-efficient approach to
 254 policy search,” in *Proceedings of the 28th International Conference on machine learning*
 255 *(ICML-11)*, pp. 465–472, 2011.
- 256 [19] S. Loosemore, R. Stallman, R. McGrath, A. Oram, and U. Drepper, “The gnu c library reference
 257 manual: for version 2.38,” *Free Software Foundation*, 2022.
- 258 [20] M. F. Schollmeyer and W. H. Tranter, “Noise generators for the simulation of digital communi-
 259 cation systems,” *ACM SIGSIM Simulation Digest*, vol. 21, no. 3, pp. 264–275, 1991.
- 260 [21] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken, M. Booth, F. Rossi, and
 261 R. Ulerich, *GNU scientific library*. Network Theory Limited Godalming, third ed., 2009.
- 262 [22] C. Lattner and V. Adve, “Llvm: a compilation framework for lifelong program analysis &
 263 transformation,” in *International Symposium on Code Generation and Optimization, 2004.*
 264 *CGO 2004.*, pp. 75–86, 2004.
- 265 [23] “Signaloid cloud developer platform.” <https://signaloid.com>, 2024.
- 266 [24] L. V. Kantorovich, “Mathematical methods of organizing and planning production,” *Manage-*
 267 *ment science*, vol. 6, no. 4, pp. 366–422, 1960.
- 268 [25] J. L. M. Poiseuille, *Recherches sur les causes du mouvement du sang dans les vaisseaux*
 269 *capillaires*, vol. 7. Impr. royale, 1839.
- 270 [26] B. S. Massey and J. Ward-Smith, *Mechanics of fluids*, vol. 1. Crc Press, 1998.

- 271 [27] A. Srivastava, A. Sood, S. P. Joy, and J. Woodcock, "Principles of physics in surgery: the laws of
272 flow dynamics physics for surgeons—part 1," *Indian Journal of Surgery*, vol. 71, pp. 182–187,
273 2009.
- 274 [28] E. Nader, S. Skinner, M. Romana, R. Fort, N. Lemonne, N. Guillot, A. Gauthier, S. Antoine-
275 Jonville, C. Renoux, M.-D. Hardy-Dessources, *et al.*, "Blood rheology: key parameters, impact
276 on blood flow, role in sickle cell disease and effects of exercise," *Frontiers in physiology*, vol. 10,
277 p. 1329, 2019.
- 278 [29] W. Hijikata, T. Maruyama, Y. Suzumori, and T. Shinshi, "Measuring real-time blood viscosity
279 with a ventricular assist device," *Proceedings of the Institution of Mechanical Engineers, Part*
280 *H: Journal of Engineering in Medicine*, vol. 233, no. 5, pp. 562–569, 2019.
- 281 [30] B. Presnell, "A geometric derivation of the cantor distribution," *The American Statistician*,
282 vol. 76, no. 1, pp. 73–77, 2022.
- 283 [31] V. I. Bogachev and M. A. S. Ruas, *Measure theory*, vol. 1. Springer, 2007.
- 284 [32] G. Casella and R. L. Berger, *Statistical inference*. Cengage Learning, 2021.
- 285 [33] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski,
286 P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman,
287 N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng,
288 E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero,
289 C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0
290 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature*
291 *Methods*, vol. 17, pp. 261–272, 2020.

292 Supplementary Material

293 A Mathematical Preliminaries

294 To establish a consistent framework for discussing the Monte Carlo method, we first introduce key
295 definitions and theorems that we will use throughout this article. Let \mathbb{R}^+ be the set of positive real
296 numbers (i.e., $[0, \infty)$).

297 **Definition 1 (Probability density function)** Let X be a set and p_X be a map from X to \mathbb{R}^+ ,

$$p_X : X \rightarrow \mathbb{R}^+,$$

298 that satisfies:

$$\int_X p_X(x) dx = 1.$$

299 We define p_X to be the probability density function on X .

300 **Definition 2 (Random Variable)** Let X be a set and p_X a probability density function on X . We
301 define the tuple (X, p_X) as a random variable⁴.

302 Often, the set X will be a real space \mathbb{R}^d . A random variable is a variable that can take on different
303 values from its defining set X . The probability density function p_X is a function that calculates how
304 likely X is to take on a *particular value* $x \in X$. Generating an instance value from a random variable
305 is called *sampling from the random variable*.

306 For brevity, we use the following notation: an uppercase letter, such as X , denotes a random variable
307 and the matching lowercase letter x denotes an instance value of the random variable X . We denote a
308 set of n independently and identically distributed (i.i.d.) samples or variates of a random variable X
309 as $\{x_i\}_{i=1}^n$, where i indexes this set and each x_i is an instance value of X .

310 Let $f : X \rightarrow Y$ denote a *transformation* from a random variable X to a random variable Y (i.e.,
311 $Y = f(X)$)⁵. We can also apply f to an instance value x of X to obtain an instance value $y = f(x)$
312 of Y .

313 The probability distribution of a random variable is a set function $\mathbb{P}_X : \Omega_X \rightarrow \mathbb{R}^+$ that tells us about
314 the probability of the random variable taking on the values inside the given set, where Ω_X is a set of
315 subsets of X . Ω_X must technically be a σ -algebra, but for our purposes, it can be thought of as a set
316 of sets each of which can be expressed as a countable union of disjoint sets that also belong to Ω_X ⁶.
317 Since we are only considering random variables that contain probability density functions, we define
318 the distribution of a random variable as:

319 **Definition 3 (Distribution of a Random Variable)** Given a random variable X with probability
320 density function p_X , the distribution of X is given by the set function

$$\begin{aligned} \mathbb{P}_X : \Omega_X &\rightarrow \mathbb{R}^+ \\ \omega = \bigcup_i U_i &\mapsto \mathbb{P}_X(\omega) = \sum_i \int_{U_i} p_X(x) dx, \end{aligned} \quad (5)$$

321 where U_i is a collection of disjoint sets whose union equals the input set ω .

322 If f is invertible and once-differentiable, then Theorem 1 derives the probability density function of
323 Y , denoted as p_Y [31, Chapter 3.7].

324 **Theorem 1 (Change of variables)** Given a random variable X with a probability density function
325 p_X and an invertible and once-differentiable transformation $f : X \rightarrow Y$, the probability density

⁴There can be measure theoretic random variables, such as the random variable that has the Cantor distribution, that do not admit probability density functions [30]. In this article, we do not consider such exotic random variables.

⁵ f transforms the set X such that there exists a valid probability density function p_Y over the set Y .

⁶The sets in a σ -algebra must also satisfy that countable unions and intersections of arbitrary sets also belong to Ω_X , along with X and \emptyset .

326 function p_Y of the random variable $Y = f(X)$ is given by:

$$\begin{aligned} p_Y : Y &\rightarrow \mathbb{R}^+, \\ y &\mapsto p_Y(y) = p_X \circ f^{-1}(y) |\det \nabla f(f^{-1}(y))|^{-1} \\ &= p_X \circ f^{-1}(y) |\det \nabla f^{-1}(y)|, \end{aligned}$$

327 where f^{-1} is the inverse of f and $\nabla f(\cdot)$ and $\nabla f^{-1}(\cdot)$ denote the Jacobian matrices of f and f^{-1}
328 respectively.

329 A key statistic that is often computed of a random variable is its expectation.

330 **Definition 4 (Expectation of a random variable)** Given a random variable X with probability
331 density function p_X , we define the expectation $\mathbb{E}_{p_X}[X]$ of X as

$$\mathbb{E}_{p_X}[X] = \int_X xp_X(x) dx. \quad (6)$$

332 Expectations can be calculated of transformations of random variables as well.

333 **Definition 5 (Expectation of a transformation of a random variable)** Given a random variable
334 X with probability density function p_X and a transformation $f : X \rightarrow Y$ from X to a random
335 variable Y^6 , we define the expectation $\mathbb{E}_{p_X}[f(X)]$ of the random variable $f(X)$ as

$$\mathbb{E}_{p_X}[f(X)] = \int_X f(x)p_X(x) dx. \quad (7)$$

336 This is called the Law of the Unconscious Statistician [32].

337 B Buffon's Needle

338 In this section, we describe Buffon's Needle in more detail. Let X be the random variable that
339 denotes the location of a thrown needle and f be a transformation on X defined as:

$$\begin{aligned} f : X &\rightarrow \{0, 1\}, \\ x &\mapsto f(x) = \begin{cases} 1 & \text{if needle lands on a line} \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

340 f therefore identifies whether a dropped needle lands on a line. The resulting random variable $f(X)$
341 is a Bernoulli random variable $\text{Bern}(p)$, where the probability of success p is the probability of a
342 needle landing on line. Since $\mathbb{E}_{f(X) \sim \text{Bern}(p)}[f(X)] = p$, the expectation of $f(X)$ is precisely the
343 probability of a needle landing on a line. The expectation of the resulting random variable $f(X)$ is $\frac{2}{\pi}$,
344 as shown by LeClerc [4]. One can approximate this expectation by sampling from X by dropping
345 needles, evaluating f by checking whether each needle landed on a line, and taking the average of
346 the resulting samples of $f(X)$.

347 C Methods: Additional detail

348 C.1 Measuring the run time

349 We measure the run time as the sum of the time taken to generate samples incurred during the
350 sampling step of the Monte Carlo method or the initializing step of Laplace, and the time taken
351 for the evaluation step. For both methods, we measure time using the `gettimeofday` function from the
352 Standard C library [19]. We measure the time from the start of the main entry point until the end of
353 key computations. The reported times omit any time spent by the programs on saving and reporting
354 the results. We took further measures to ensure that our results were meaningful; these are detailed in
355 the supplementary material.

356 In order to explicitly quantify the *post-processing* step of the Monte Carlo method, we compute the
357 mean and the variance of the samples obtained from the Monte-Carlo-based experiments. Such a step
358 is not necessary with Laplace, because Laplace already provides a usable representation of the output
359 distribution. We note that we are being generous to the Monte Carlo method, since the mean and the
360 variance alone does not fully capture the shape of a non-Gaussian distribution. In contrast, Laplace
361 captures the full distribution in its representation.

362 C.2 Measuring the Wasserstein Distance

363 The Wasserstein distance [24] is a metric that measures the distance between probability distributions.
364 We quantify the distance of the outputs to the ground truth using the Wasserstein distance between the
365 output distribution calculated by each approach and the ground truth output distribution. We compute
366 the ground truth output distribution by running the Monte Carlo method with 1,000,000 samples.
367 In our experiments, we calculate the Wasserstein distance using the `scipy.stats.wasserstein_distance`
368 function from the `scipy` Python package [33].

369 C.3 Experimental setup

370 Let n be the number of samples used in the sampling step of a Monte Carlo simulation. We perform
371 experiments with various values of n on an Apple M1 Pro with 16GB LPDDR5 RAM, running
372 macOS 13.5.1. This provides a baseline for the performance of the Monte Carlo method that can be
373 expected in the real-world.

374 Similarly, for Laplace, we varied the representation size r . Since the Laplace cores generate in-
375 processor representations of the output distribution, we take samples from this distribution to compute
376 the Wasserstein distance. We take 1,000,000 samples, similar to the ground truth. We do not include
377 the time taken for this sampling in the wall-clock time because this sampling is done solely to
378 calculate the Wasserstein distance and is not part of a typical use case of Laplace.

379 D Ensuring Meaningful Timing Results

380 When running the experiments on the Monte Carlo method, each repetition of an experiment was
381 run after a 5s delay. This delay ensures that we avoid buffer cache optimizations carried out by the
382 operating system.

383 We also note that we did not exploit parallelization when running Traditional Monte Carlo since the
384 available implementation of Laplace did not exploit parallelization either. We felt that this provided
385 an apples-to-apples comparison.

386 E Applications: Additional Detail

387 E.1 Monte Carlo Convergence Challenge Example

388 Here, we present a more complete description of the Monte Carlo Convergence Challenge example.
389 For ease of reading, we repeat the key equations.

390 Let X^{con} be the initial random variable that we sample from, with its PDF $p_{X^{\text{con}}}$ being a Gaussian
391 mixture. The underlying set of X^{con} is \mathbb{R} , and $p_{X^{\text{con}}}$ is given by:

$$p_{X^{\text{con}}}(x) = 0.6 \left(\frac{1}{0.5\sqrt{2\pi}} \exp(-2(x-2)^2) \right) + 0.4 \left(\frac{1}{1.0\sqrt{2\pi}} \exp\left(\frac{-(x+1)^2}{2}\right) \right). \quad (8)$$

392 For the Monte Carlo evaluation step of Section 2, we define a function f^{con} as a sigmoidal function:

$$f^{\text{con}} : X \rightarrow (0, 1), \\ x \mapsto f^{\text{con}}(x) = \frac{1}{1 + e^{-(x-1)}}. \quad (9)$$




Figure 2: Left: graphs of the analytical input and output distributions, and the evaluation function for the Monte Carlo Convergence Challenge example. Right: histograms of the output of Monte Carlo simulation with 8 (b), 256 (b), 1024 (c), and 128,000 (d) samples. The input PDF has two modes, and the output distribution is heavily influenced by them (see black curve). If only a few samples are used during Monte Carlo simulation, such as in (b), then the resulting histogram will be biased toward a single mode.

393 Let $Y^{\text{con}} = f^{\text{con}}(X^{\text{con}})$ denote the output random variable. The underlying set of Y^{con} is $(0, 1)$. Its
 394 PDF $p_{Y^{\text{con}}}(y)$ can be analytically calculated using Theorem 1:

$$\begin{aligned}
 p_{Y^{\text{con}}}(y) = & \left(0.6 \left(\frac{1}{0.5\sqrt{2\pi}} \exp(-(\text{logit}(y) - 2)^2) \right) \right. \\
 & \left. + 0.4 \left(\frac{1}{1.0\sqrt{2\pi}} \exp\left(\frac{-(\text{logit}(y) + 1)^2}{2}\right) \right) \right) \\
 & \times \left| \frac{\exp(1 - \text{logit}(y))}{(\exp(1 - \text{logit}(y)) + 1)^2} \right|^{-1}, \tag{10}
 \end{aligned}$$

395 where $\text{logit}(y)$ is:


$$\text{logit}(y) = 1 + \log \frac{y}{1 - y}. \tag{11}$$

396 . Figure 2 plots the functions of Equations 8, 9, and 10, and highlights the problem of uncertainty
 397 propagation. The function f^{con} (red curve) transforms the random variable X^{con} (PDF shown as
 398 the blue curve) into the random variable Y^{con} (PDF shown as the black curve). In particular, f^{con}
 399 transforms the two modes of X^{con} into the two modes of Y^{con} .

400 We chose this application to showcase issues with convergence in traditional Monte Carlo simulation
 401 (see Figure 2). Due to the multi-modal distribution $p_{X^{\text{con}}}$, using too few samples could bias the
 402 resulting histogram of the Monte Carlo simulation toward the largest mode and not represent the
 403 other mode well, as in Figure 2b. The fidelity of the output of Monte Carlo methods is therefore
 404 sensitive to the number of samples taken from X^{con} , and to the shape of the function f^{con} . By
 405 contrast, uncertainty-tracking processors such as Laplace are not sample-based and can be said to be
 406 *convergence-oblivious*.

407 F Additional Results and Discussion

408 We provide an abridged set of results and observations in this section. We have broken the discus-
 409 sion into three sections for comparing the relationship between the Wasserstein distance and run
 410 time, comparing the number of required dynamic instructions, and comparing histograms of output
 411 distributions.



(a) Pareto plot for the Monte Carlo Convergence Challenge application from Section 5.1. We have omitted $n = 4, 256$ for clarity. (b) Pareto plot for the Poiseuille's Law for Blood Transfusion application from Section 5.1. We have omitted $n = 4, 256$ for clarity.

Figure 3: Pareto plots between the mean run time, and the mean Wasserstein distance from the ground truth output distribution. The error bars show ± 1 standard deviation, as in Tables 2-3. For the Monte Carlo Convergence Challenge application, (a) shows that Traditional Monte Carlo obtains better accuracy than Laplace with $r = 32$ (up to 1-standard deviation) after 128,000 samples. For the Poiseuille's Law for Blood Transfusion application, (b) shows that Traditional Monte Carlo obtains better accuracy than Laplace with $r = 32$ (up to 1-standard deviation) after 2048 samples. In the legends, MC stands for *Traditional Monte Carlo, implemented in C*. We use the log scale on the vertical axis.

412 F.1 Comparing Wasserstein Distance and Run Time

413 Tables 2 and 3 show the means and the standard deviations of the run time and the Wasserstein
 414 distance⁷ for the Monte Carlo Convergence Challenge and the Poiseuille's Law for Blood Trans-
 415 fusion examples, respectively. Figure 3 plots the run time and the Wasserstein distance across all
 416 experimental variations to show the Pareto boundary for each application. In general, these results
 417 match intuition, where an increase in n improves the accuracy of the output distributions compared to
 418 the ground-truth distribution, at the expense of run time. We analyze the results for each application
 419 in turn.

420 F.1.1 Monte Carlo Convergence Challenge

421 Table 2 and Figure 3a shows the key results for this application. A Laplace configuration with $r > 32$
 422 improves the Wasserstein distance less than it worsens the run time. Therefore, we chose $r = 32$
 423 to be the *best overall* configuration of Laplace for this application. We see the Monte Carlo method
 424 requires approximately 32,000 samples to obtain a similar Wasserstein distance to Laplace with
 425 $r = 32$. For this configuration, Laplace takes $113.85\times$ less time. To obtain an accuracy that is better
 426 than Laplace with $r = 32$ up to 1-standard deviation, the Monte Carlo method requires approximately
 427 128,000 samples, for which it takes $411.25\times$ more time. Similarly, if we required the Monte Carlo
 428 method to obtain a Wasserstein distance better than Laplace with $r = 32$ up to 2-standard deviations,
 429 it would require approximately 256,000 samples, for which it takes $732.35\times$ more time.

430 F.1.2 Poiseuille's Law for Blood Transfusion

431 Table 3 and Figure 3b show that the best trade-off between accuracy and run time is made by Laplace
 432 with $r = 32$. To match the mean accuracy of this configuration of Laplace, Traditional Monte Carlo
 433 requires 256,000 samples. This takes $51.53\times$ more time than Laplace. To obtain an accuracy better
 434 than Laplace with $r = 32$ up to 1-standard deviation and 2-standard deviations, Traditional Monte
 435 Carlo requires approximately 512,000 samples. This takes $160.06\times$ more time than Laplace.

⁷The Wasserstein distances are of very different scales. The scale of Wasserstein distances will depend on the distributions being compared. However, the important insights from Table 2-3 and Figure 3 are the trends.

Problem	Core	Representation Size / Number of samples	Wasserstein Distance (mean \pm std. dev.)	Run time (ms) (mean \pm std. dev.)
Monte Carlo Convergence Challenge	Laplace	16	0.00457 \pm 0.00004	0.011 \pm 0.001
Monte Carlo Convergence Challenge	Laplace	32	0.00167 \pm 0.00007	0.020 \pm 0.004
Monte Carlo Convergence Challenge	Laplace	64	0.00097 \pm 0.00008	0.034 \pm 0.004
Monte Carlo Convergence Challenge	Laplace	128	0.00065 \pm 0.00018	0.044 \pm 0.003
Monte Carlo Convergence Challenge	Laplace	256	0.00054 \pm 0.00021	0.073 \pm 0.002
Monte Carlo Convergence Challenge	Laplace	2048	0.00042 \pm 0.00020	0.531 \pm 0.008
Monte Carlo Convergence Challenge	Traditional Monte Carlo	4	0.13781 \pm 0.06187	0.077 \pm 0.049
Monte Carlo Convergence Challenge	Traditional Monte Carlo	256	0.02136 \pm 0.01130	0.086 \pm 0.011
Monte Carlo Convergence Challenge	Traditional Monte Carlo	1152	0.00844 \pm 0.00363	0.155 \pm 0.035
Monte Carlo Convergence Challenge	Traditional Monte Carlo	2048	0.00659 \pm 0.00277	0.243 \pm 0.109
Monte Carlo Convergence Challenge	Traditional Monte Carlo	4096	0.00437 \pm 0.00205	0.381 \pm 0.125
Monte Carlo Convergence Challenge	Traditional Monte Carlo	8192	0.00307 \pm 0.00123	0.727 \pm 0.199
Monte Carlo Convergence Challenge	Traditional Monte Carlo	16000	0.00270 \pm 0.00139	1.260 \pm 0.234
Monte Carlo Convergence Challenge	Traditional Monte Carlo	32000	0.00158 \pm 0.00068	2.277 \pm 0.346
Monte Carlo Convergence Challenge	Traditional Monte Carlo	128000	0.00086 \pm 0.00035	8.225 \pm 0.849
Monte Carlo Convergence Challenge	Traditional Monte Carlo	256000	0.00067 \pm 0.00038	14.645 \pm 1.330

Table 2: Results show the mean Wasserstein distance, the run time and the factor increase in dynamic instructions required, with their 1-standard deviation errors for the Monte Carlo Convergence Challenge example. We have highlighted in bold the best overall configuration for Laplace and the close-to-equivalent results Monte Carlo configurations. Traditional Monte Carlo takes approximately $113.85\times$ more time than Laplace. To have better accuracy than the Laplace result up to 1-standard deviation and 2-standard deviations, Traditional Monte Carlo requires approximately 128,000 samples and 256,000 samples respectively. These take $411.25\times$ and $732.35\times$ more time than Laplace, respectively.

Problem	Core	Representation Size / Number of samples	Wasserstein Distance (mean \pm std. dev.)	Run time (ms) (mean \pm std. dev.)
Poiseuille’s Law for Blood Transfusion	Laplace	16	0.00085 \pm 0.00001	0.051 \pm 0.003
Poiseuille’s Law for Blood Transfusion	Laplace	32	0.00033 \pm 0.00003	0.173 \pm 0.006
Poiseuille’s Law for Blood Transfusion	Laplace	64	0.00017 \pm 0.00003	0.412 \pm 0.006
Poiseuille’s Law for Blood Transfusion	Laplace	128	0.00015 \pm 0.00004	1.406 \pm 0.017
Poiseuille’s Law for Blood Transfusion	Laplace	256	0.00015 \pm 0.00004	5.800 \pm 0.055
Poiseuille’s Law for Blood Transfusion	Laplace	2048	0.00014 \pm 0.00005	480.637 \pm 2.663
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	4	0.05379 \pm 0.01818	0.066 \pm 0.019
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	256	0.00699 \pm 0.00245	0.089 \pm 0.031
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	1152	0.00313 \pm 0.00112	0.212 \pm 0.299
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	4096	0.00183 \pm 0.00052	0.345 \pm 0.049
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	8192	0.00118 \pm 0.00038	0.676 \pm 0.204
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	32000	0.00067 \pm 0.00025	2.744 \pm 0.935
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	128000	0.00033 \pm 0.00009	8.914 \pm 1.566
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	256000	0.00023 \pm 0.00008	15.303 \pm 2.611
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	512000	0.00017 \pm 0.00004	27.690 \pm 3.361
Poiseuille’s Law for Blood Transfusion	Traditional Monte Carlo	640000	0.00015 \pm 0.00005	33.853 \pm 1.270

Table 3: Results show the mean Wasserstein distance, the run time and the factor increase in dynamic instructions required, with their 1-standard deviation errors for the Poiseuille’s Law for Blood Transfusion example. We have highlighted in bold the best overall configuration for Laplace and the close-to-equivalent results Monte Carlo configurations. Traditional Monte Carlo takes approximately $51.53\times$ more time than Laplace with $r = 32$. To have better accuracy than the Laplace result up to 1-standard deviation and 2-standard deviations, Traditional Monte Carlo requires approximately 512,000 samples. This takes $160.06\times$ more time than Laplace.

436 F.2 Comparing Histograms of Output Distributions

437 Figure 4 shows histograms of output distributions for each of the applications. A key observation
438 from these is that the outcome of Laplace, even with high representation sizes is slightly different
439 from the ground truth distributions. We can note that the distribution produced by Laplace puts higher
440 probability density at the mode, and less on the tails. This is contrasted by the histograms of the
441 Monte Carlo method, where the output distribution eventually approaches the ground truth, as the
442 number of samples is increased.

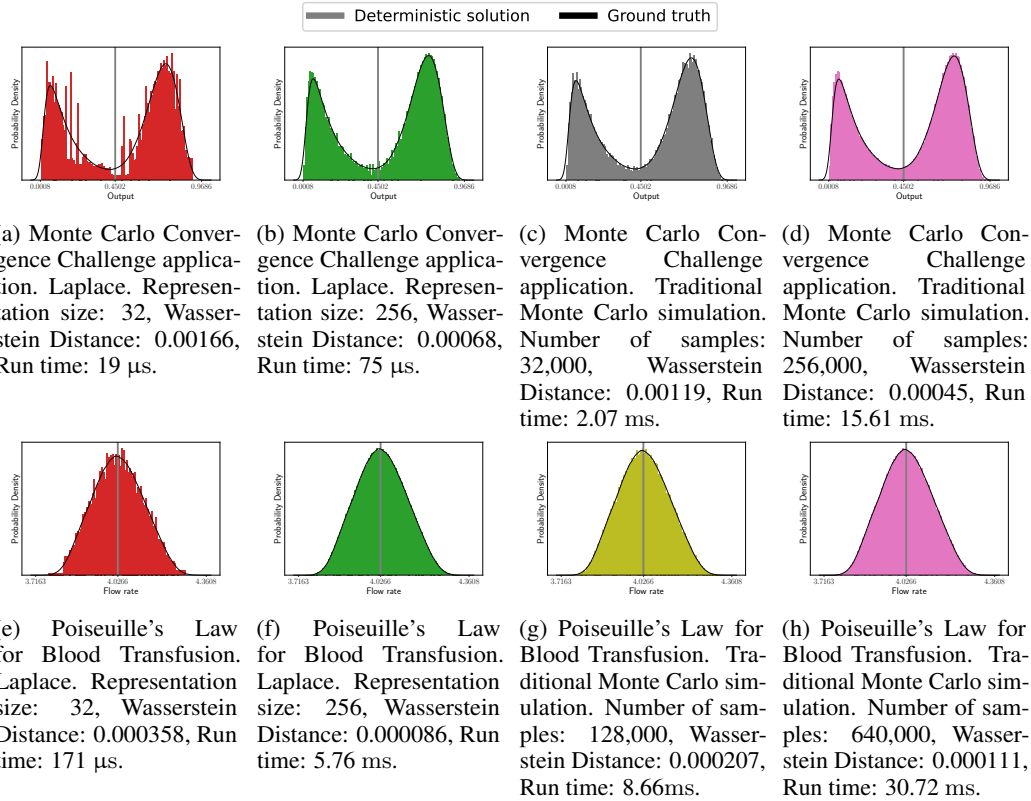


Figure 4: Histograms from example experiments on all applications, showing outputs of Laplace ((a), (b), (e), and (f)) and the Monte Carlo method ((c), (d), (g), and (h)). For the Laplace plots, we have taken 1,000,000 samples from Laplace's internal distribution representation. We set the number of histogram bins to 100 for all cases. The black outline shows a kernel density estimation of the ground truth obtained by Monte Carlo simulations with 1,000,000 samples. The gray vertical lines show the deterministic evaluation, where all uncertain input values and parameters are assumed to have taken their mean value. We also show the minimum and maximum sample values that were for the ground truth.