

Modular Exact Inference for Discrete Probabilistic Programs

Steven Holtzen and Guy Van den Broeck and Todd Millstein

University of California, Los Angeles
 {sholtzen, guyvdb, todd}@cs.ucla.edu



1. Motivation

- PPLs are extremely powerful and flexible, but this makes inference hard
 - PPLs have a *scaling problem*: focus primarily on small programs
- Our goal:** Focus on **discrete programs**, make a high-performance *exact* inference algorithm for this specialized setting
 - Discreteness is *very common*, many programs have discrete parts (text, graphs, computer networks, ...)
 - Discreteness is *challenging for many methods*
 - Many methods rely on differentiability
 - Low-probability observations
 - Exact inference preferable to approximate
 - Does not propagate errors
 - Suitable for high-consequence decisionmaking

2. Method

- Key idea: *factorize the inference computation* (see Figure 1a)

$$\underbrace{0.1}_{x=T} \cdot \underbrace{0.2}_{y=T} \cdot \underbrace{0.4}_{z=T} + \underbrace{0.1}_{x=T} \cdot \underbrace{0.8}_{y=F} \cdot \underbrace{0.5}_{z=T} + \underbrace{0.9}_{x=F} \cdot \underbrace{0.3}_{y=T} \cdot \underbrace{0.4}_{z=T} + \underbrace{0.9}_{x=F} \cdot \underbrace{0.7}_{y=F} \cdot \underbrace{0.5}_{z=T}$$

Versus...


$$\underbrace{0.1}_{x=T} \cdot \left(\underbrace{0.2}_{y=T} \cdot \underbrace{0.4}_{z=T} + \underbrace{0.8}_{y=F} \cdot \underbrace{0.5}_{z=T} \right) + \underbrace{0.9}_{x=F} \cdot \left(\underbrace{0.3}_{y=T} \cdot \underbrace{0.4}_{z=T} + \underbrace{0.7}_{y=F} \cdot \underbrace{0.5}_{z=T} \right)$$

- Finding and exploiting these factorization opportunities can be hard!
- We do it with binary decision diagrams (BDDs)

```

1 let x = flip1 0.1 in
2 let y = if x then flip2 0.2 else
3   flip3 0.3 in
4 let z = if y then flip4 0.4 else
5   flip5 0.5 in z
    
```

(a) Example Dice program.



(b) Compiled BDD with weighted model counts.

Figure 1: Illustration of compiling a Dice program that exploits factorization.

3. Experiments

- Show Dice can perform exact inference on *extremely large programs*
 - For instance, a 1.9 megabyte program with over 100k random variables
- Compared Dice against **Psi** and **Ace** (specialized Bayesian network solver)
- Three main experiments:
 - Common Baselines
 - Single-marginal Bayesian network inference
 - All-marginal Bayesian network inference

Benchmark	Psi (ms)	DP (ms)	Dice (ms)	# Paths	BDD Size
Grass	167	57	14	95	15
Burglar Alarm	98	10	13	250	11
Coin Bias	94	23	13	4	13
Noisy Or	81	152	13	1640	35
Evidence1	48	32	13	9	5
Evidence2	59	28	13	9	6
Murder Mystery	193	75	10	16	6

Table 1: *Baselines*. Comparison of inference algorithms (times are milliseconds). The total time for Dice is reported under the "Dice" column, and the total size of the final compiled BDD is reported in the "BDD Size" column.

Benchmark	Psi (ms)	DP (ms)	Dice (ms)	# Parameters	# Paths
Cancer	772	46	13.0	10	1.1×10 ³
Survey	2477	152	13.0	21	1.3×10 ⁴
Alarm	X	X	25.0	509	1.0×10 ³⁶
Insurance	X	X	75.0	984	1.2×10 ⁴⁰
Hepar2	X	X	54.0	48	2.9×10 ⁴⁹
Hailfinder	X	X	618	2656	2.0×10 ⁷⁶
Pigs	X	X	72	5618	7.3×10 ⁹²
Water	X	X	2876	1.0×10 ⁴	3.2×10 ⁵⁴
Munin	X	X	1998	8.1×10 ⁵	2.1×10 ⁶²²

Table 1: *Single Marginal Inference*. Comparison of inference algorithms (times are milliseconds). A "X" denotes a timeout at 2 hours of running. The total time for Dice is reported under the "Dice" column, and the total size of the final compiled BDD is reported in the "BDD Size" column.

Benchmark	Dice (ms)	Ace (ms)	BDD Size
Alarm	159	422	4.3×10 ⁵
Hailfinder	1280	522	2.1×10 ⁵
Insurance	222	492	2.3×10 ⁵
Hepar2	163	495	5.4×10 ⁵
Pigs	11243	985	2.6×10 ⁵
Water	3320	605	6.8×10 ⁴
Munin	4021194	3500	2.2×10 ⁷

Table 1: *All marginals*. A comparison between Dice and Ace on the all-marginal discrete Bayesian network inference task.

4. Conclusion

- Github: <https://github.com/SHoltzen/dice>
- Webpage: <http://dicelang.cs.ucla.edu/>
- Steven Holtzen, Guy Van den Broeck, and Todd Millstein. 2020. Scaling Exact Inference for Discrete Probabilistic Programs. Proc. ACM Program. Lang. 4, OOPSLA, Article 140 (November 2020), 37 pages. <https://doi.org/10.1145/3428208>
- Paper link: <https://arxiv.org/abs/2005.09089>

5. Acknowledgments

This work is partially supported by NSF grants #IIS-1943641, #IIS-1956441, #CCF-1837129, DARPA grant #N66001-17-2-4032, a Sloan Fellowship, and gifts by Intel and Facebook research. The authors would like to thank Jon Aytac and Philip Johnson-Freyd for feedback on paper drafts.



The screenshot shows the Dice website interface. At the top, it says "Dice The dice probabilistic programming language" and "About GitHub". Below that, it explains that Dice is a probabilistic programming language for fast exact inference. There is a code editor with the following code:

```

1 let a = flip 0.3 in
2 let b = flip 0.8 in
3 let tmp = observe a || b in
4
    
```

Below the code editor is a "Run" button. Underneath, there is a "Joint Distribution" table:

Value	Probability
true	0.348837209302
false	0.651162790698