Approximate Knowledge Compilation by Online Collapsed Importance Sampling

Tal Friedman and Guy Van den Broeck

Motivation

Factor Graphs:

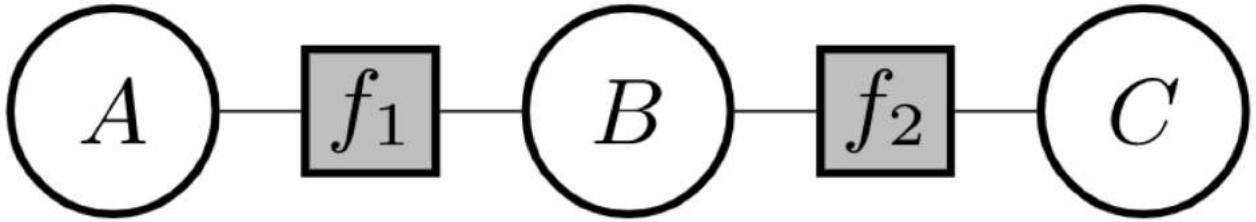

Motivation

Factor Graphs:

Great! But asking queries is hard

Motivation

Factor Graphs:

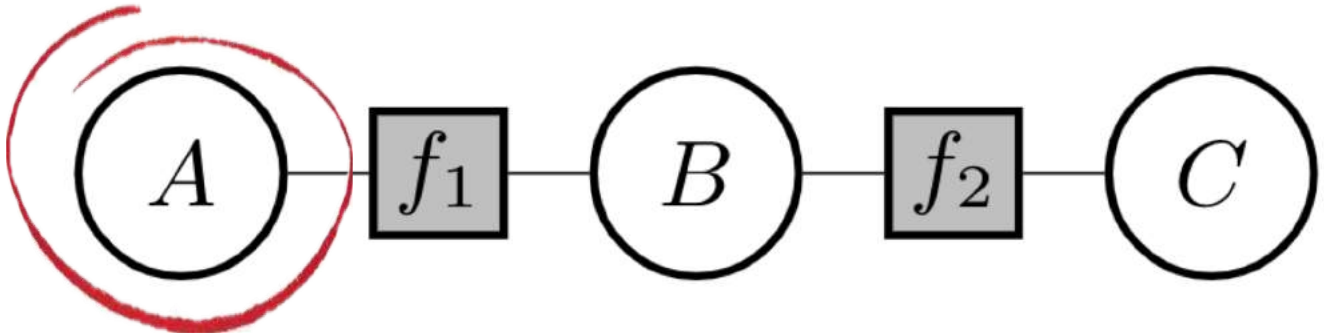

Motivation: Arithmetic Circuit

- Exact inference: Use Knowledge Compilation (e.g. BDD, SPN)

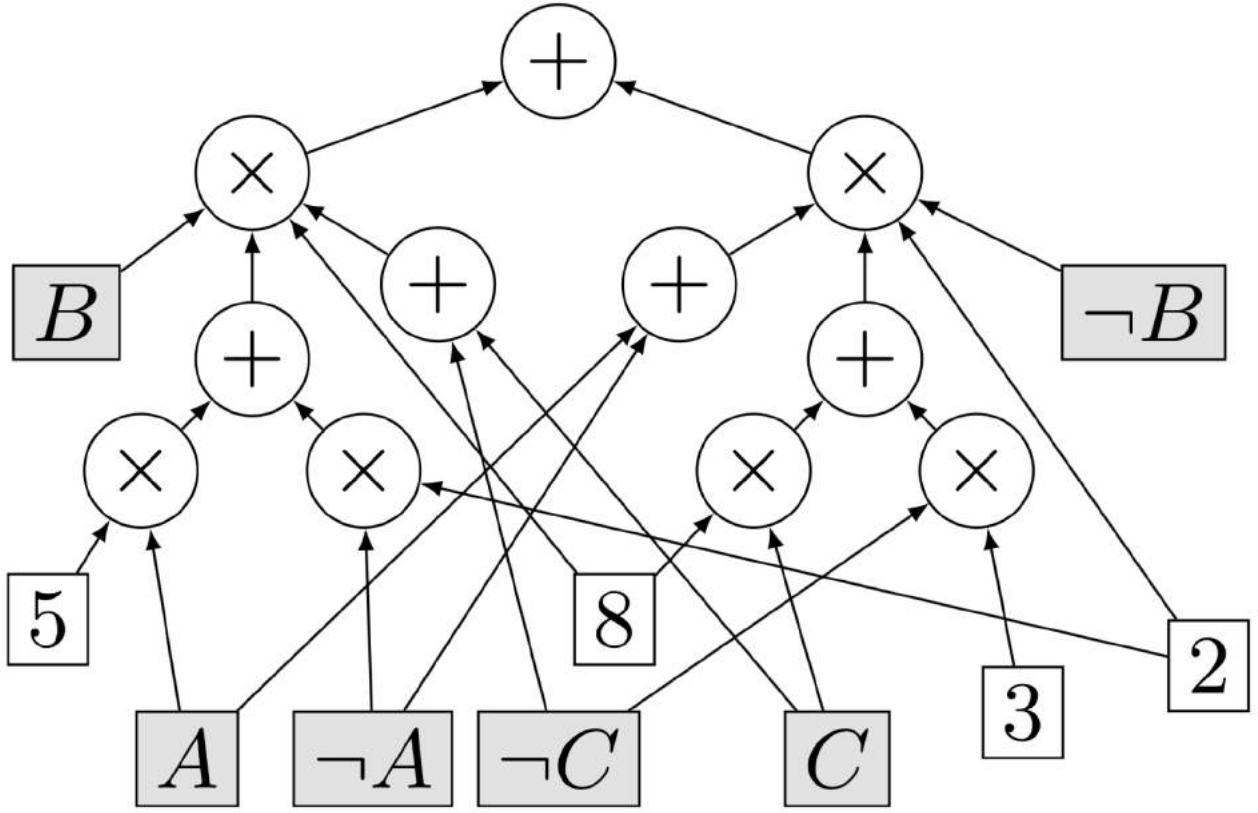

- **Tractable** form: easy queries + operations
- Take advantage of further independence properties, logical structure

*But they don't **scale**!*

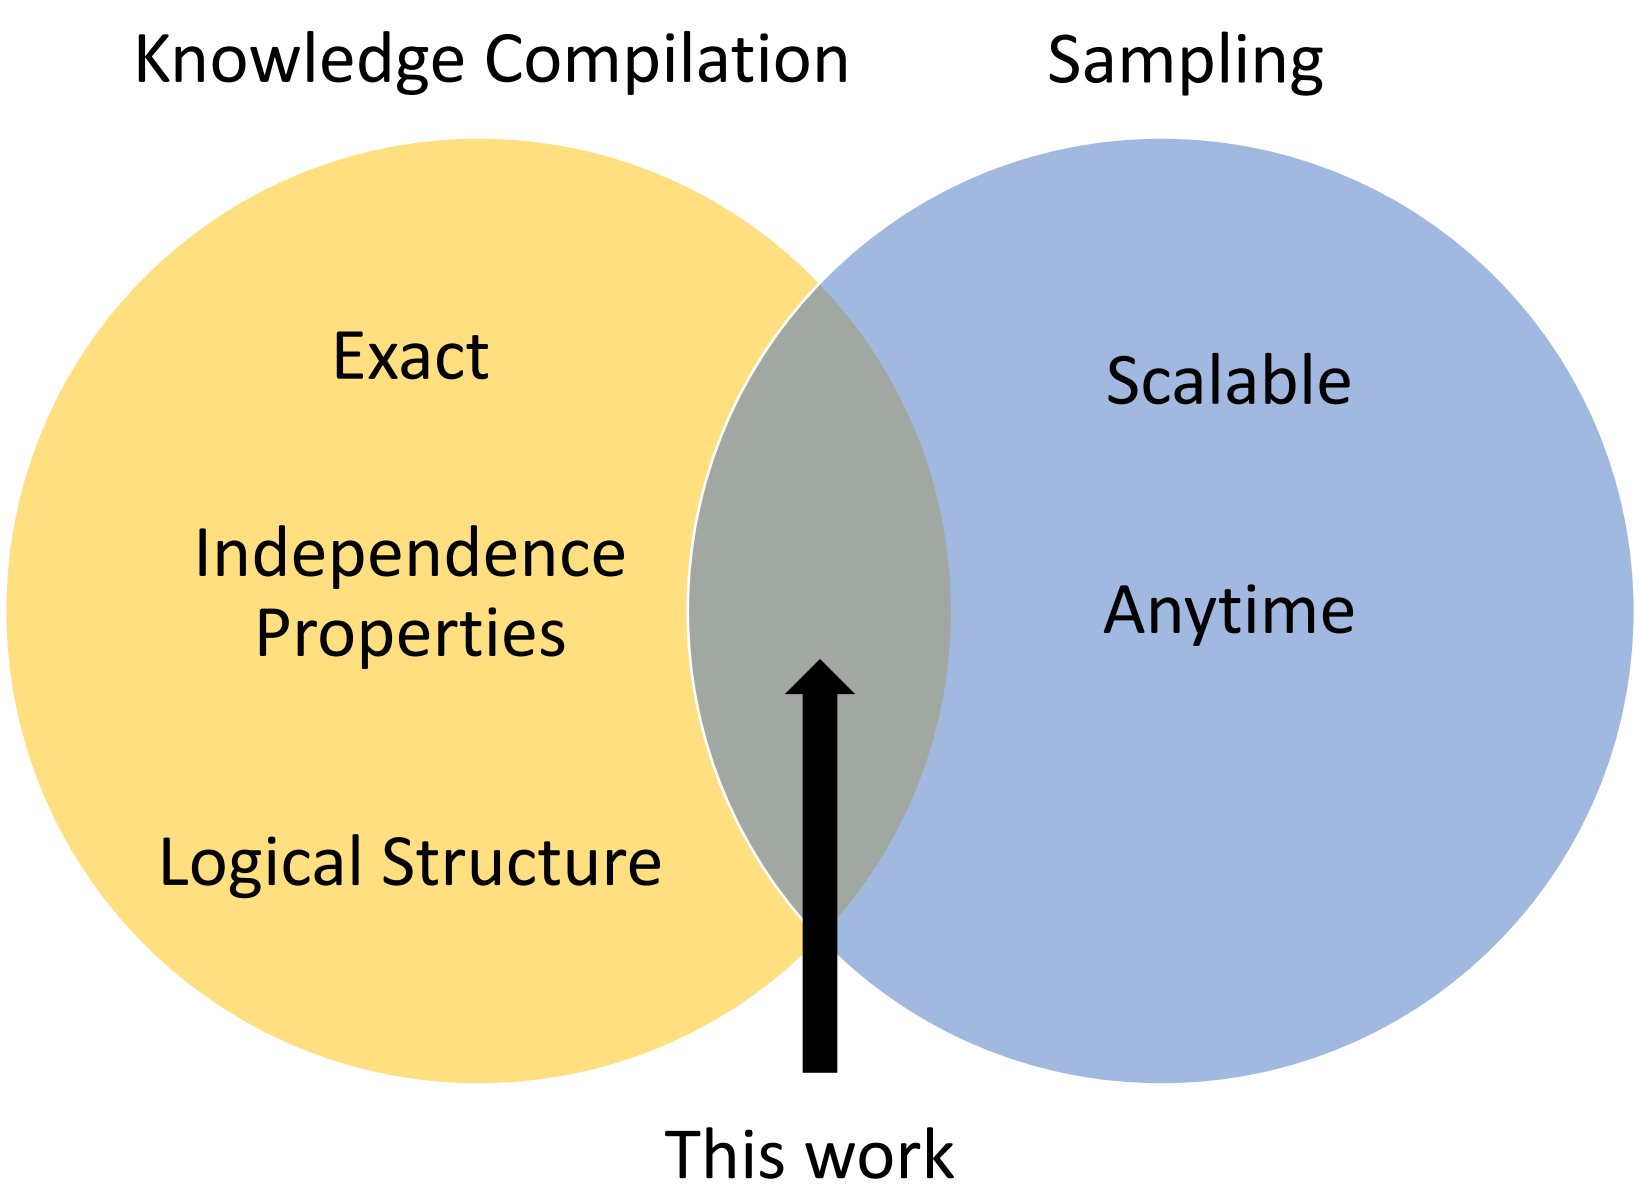
Knowledge Compilation
Sampling
Exact
Independence Properties
Logical Structure
Scalable
Anytime
This work

Collapsed Sampling (Rao-Blackwell)

Sampling on some variables, exact inference conditioned on sample

Collapsed Sampling (Rao-Blackwell)

Sampling on some variables, exact inference conditioned on sample

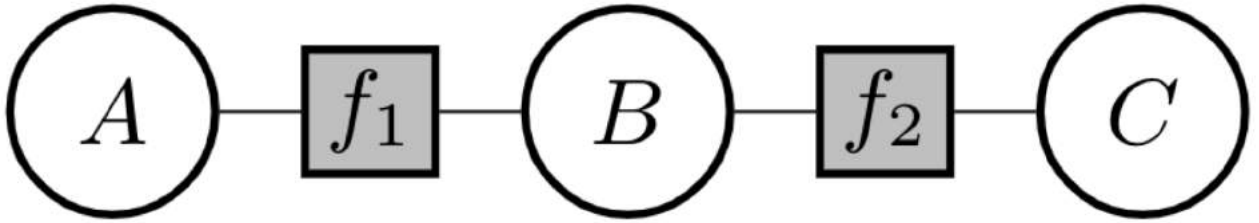

Collapsed Sampling (Rao-Blackwell)

Sampling on some variables, exact inference conditioned on sample

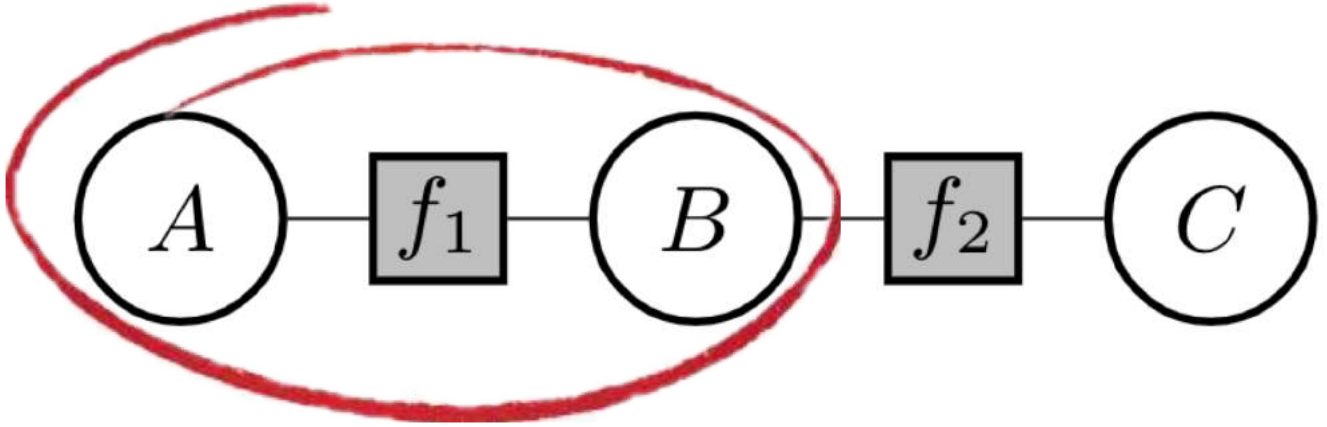

Sample A,B

Collapsed Sampling (Rao-Blackwell)

Sampling on some variables, exact inference conditioned on sample

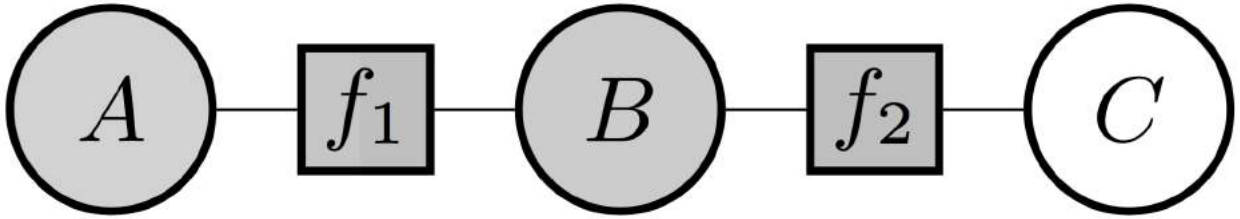

Observe sampled values

Collapsed Sampling (Rao-Blackwell)

Sampling on some variables, exact inference conditioned on sample

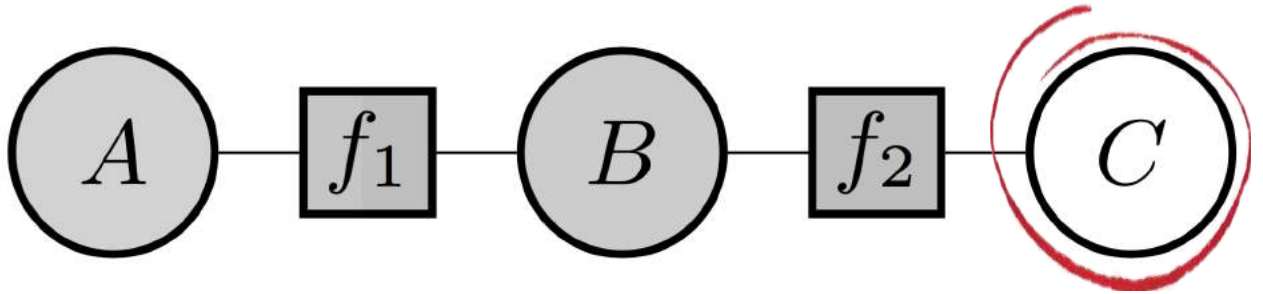

Compute exactly P(C|A,B)

What to Sample?

- Is it even possible to pick a correct set a priori?
- Consider a network of potential smokers, with friendships sampled

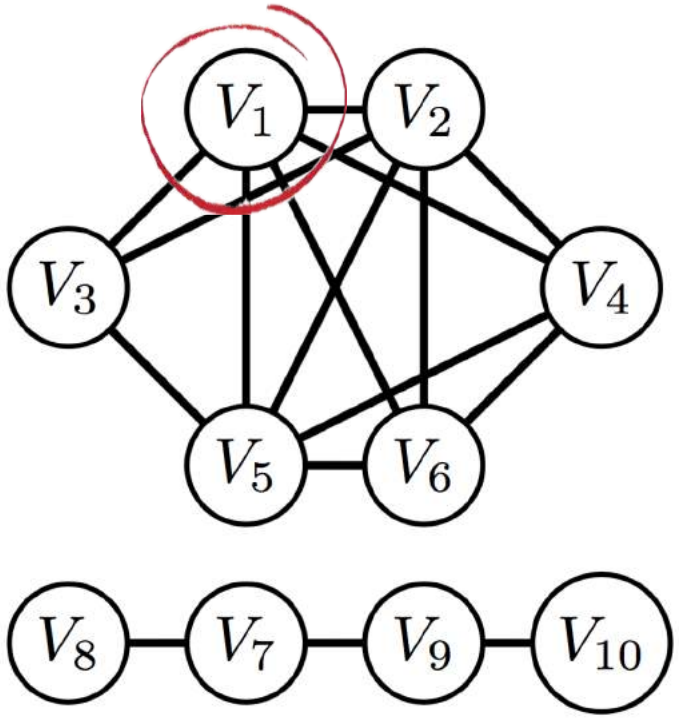

Sample 1

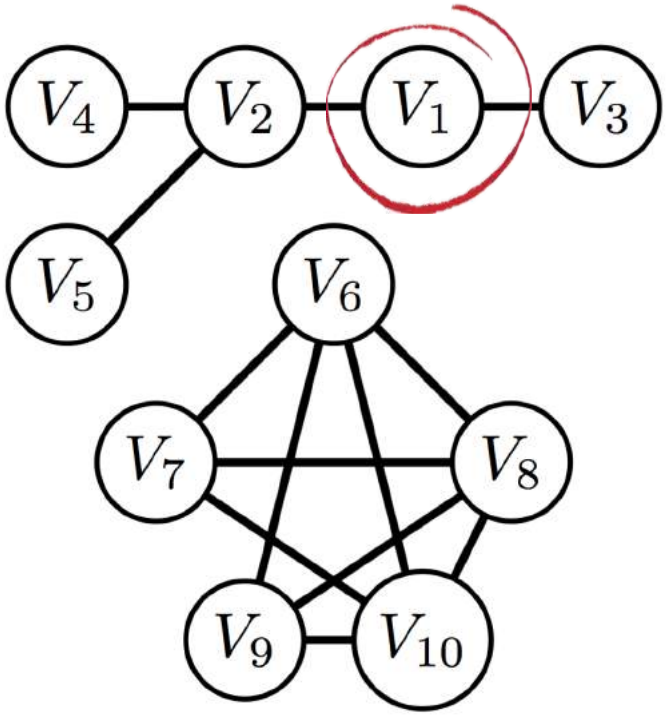

Sample 2

Online Collapsed Sampling

Choose *on-the-fly* which variable to sample next, based on result of sampling previous variables

Theorem: Still unbiased

How?

1. What/when do we sample?

How?

1. What/when do we sample?
2. How do we sample?

How do we Sample?

- Importance Sampling
- Need a proposal for **any** variable conditioned on **any other** variables

How?

1. What/when do we sample?
2. How do we sample?
3. How do we do exact inference?

Exact Inference

How do we do exact inference conditioned on different variables?

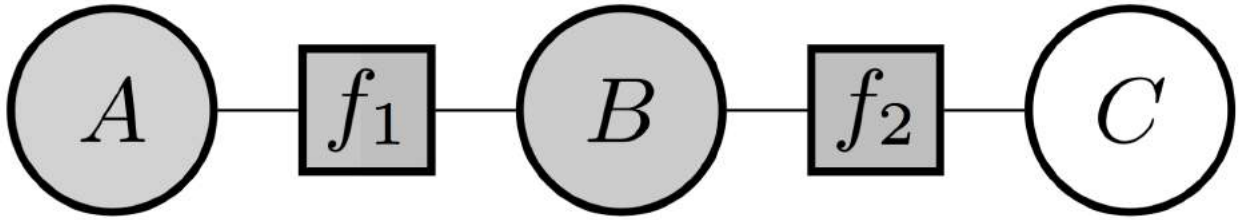

Exact Inference

- How do we do exact inference conditioned on different variables?

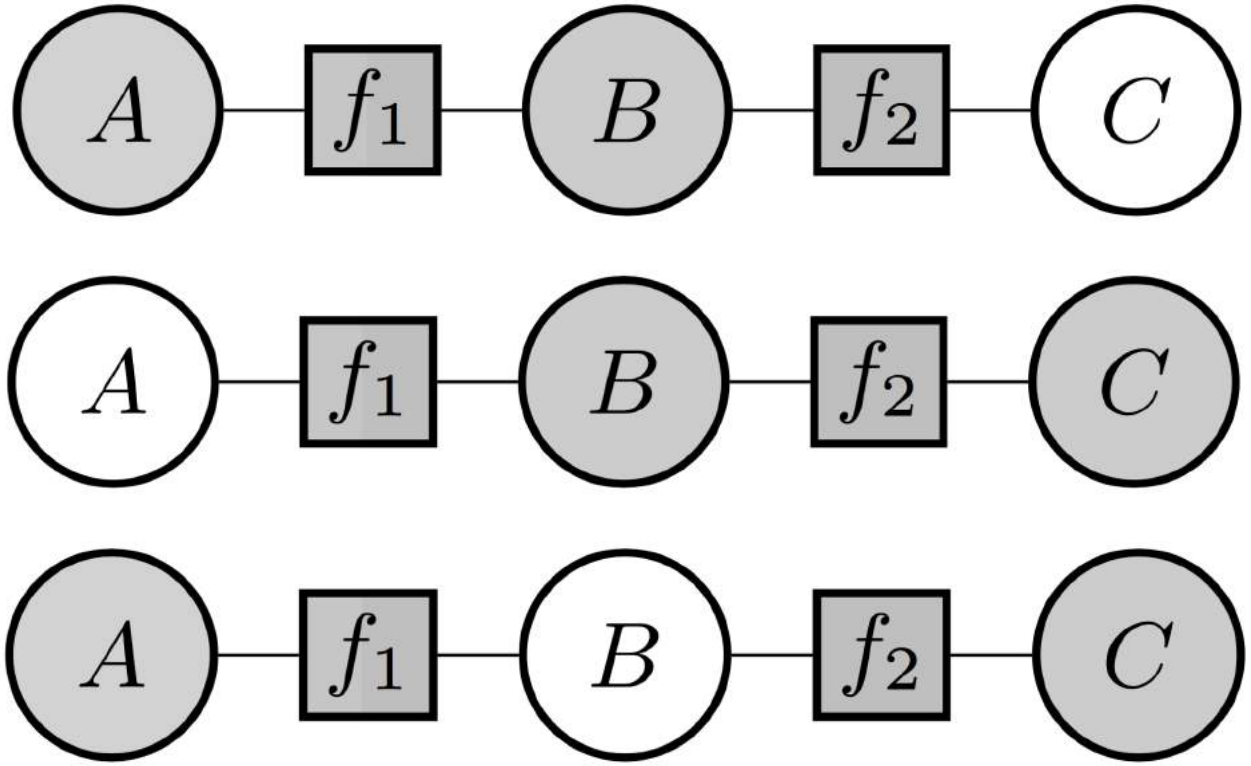

Collapsed Compilation

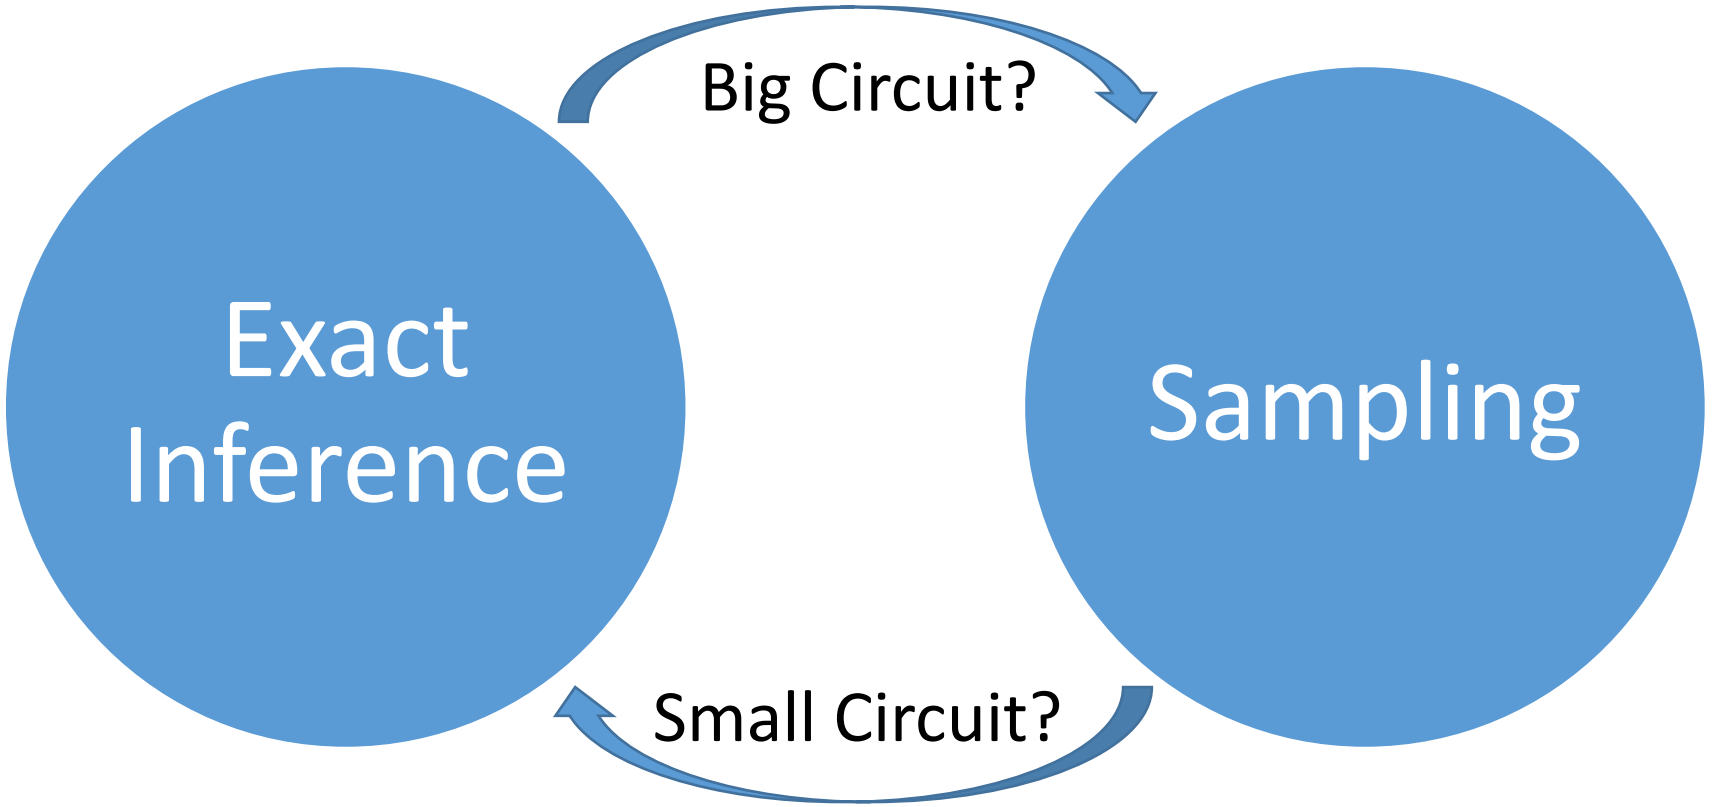

Result: A circuit for factor graph with some sampled variables

Collapsed Compilation

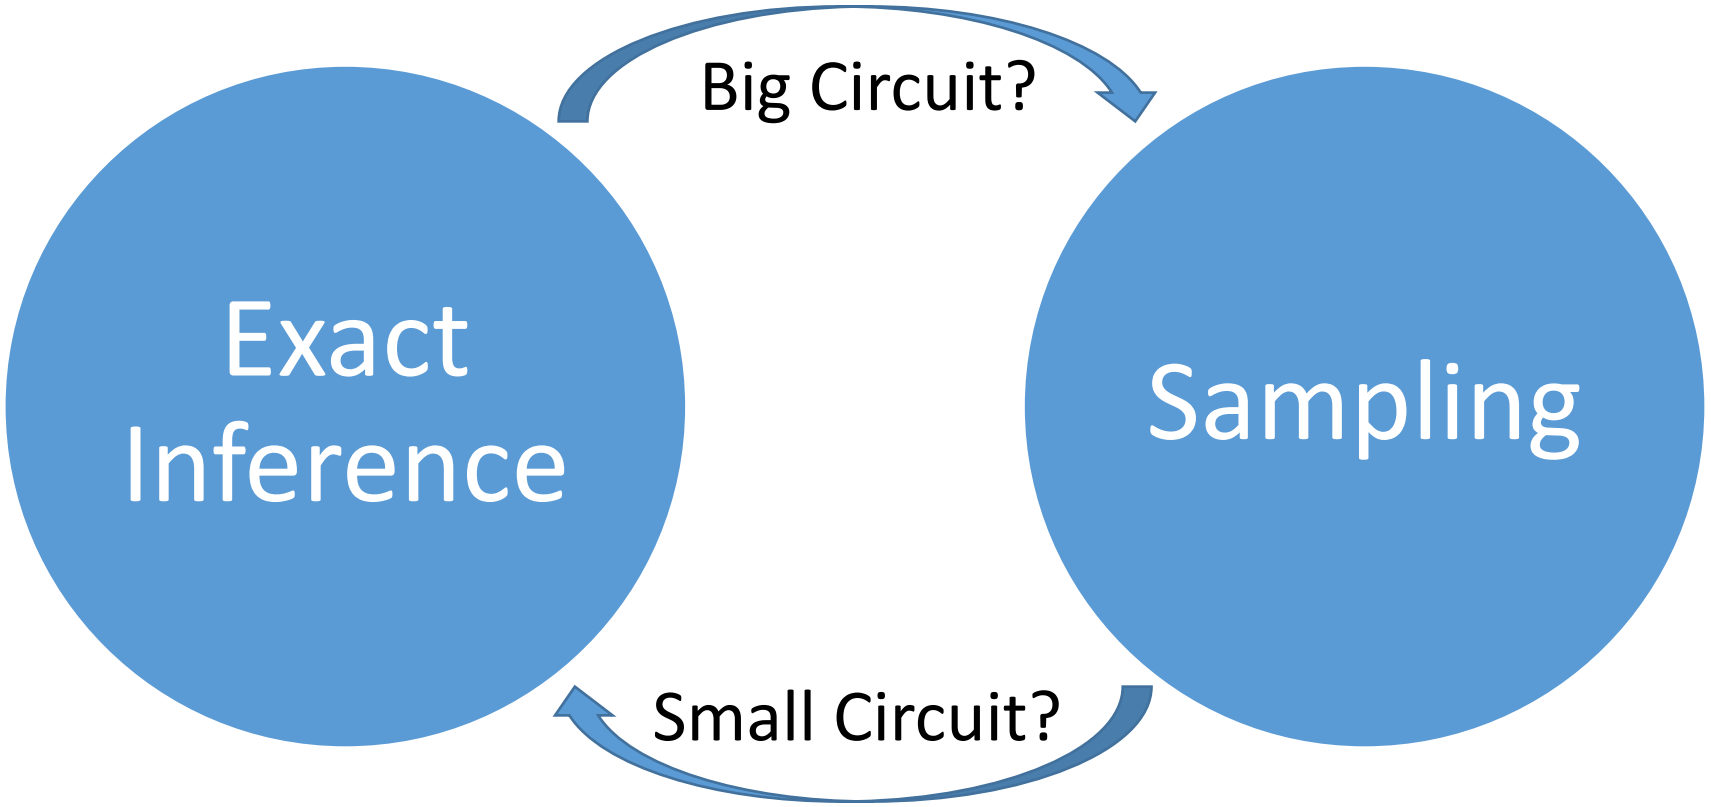

1. What/when do we sample?
2. How do we sample?
3. How do we do exact inference?

Collapsed Compilation

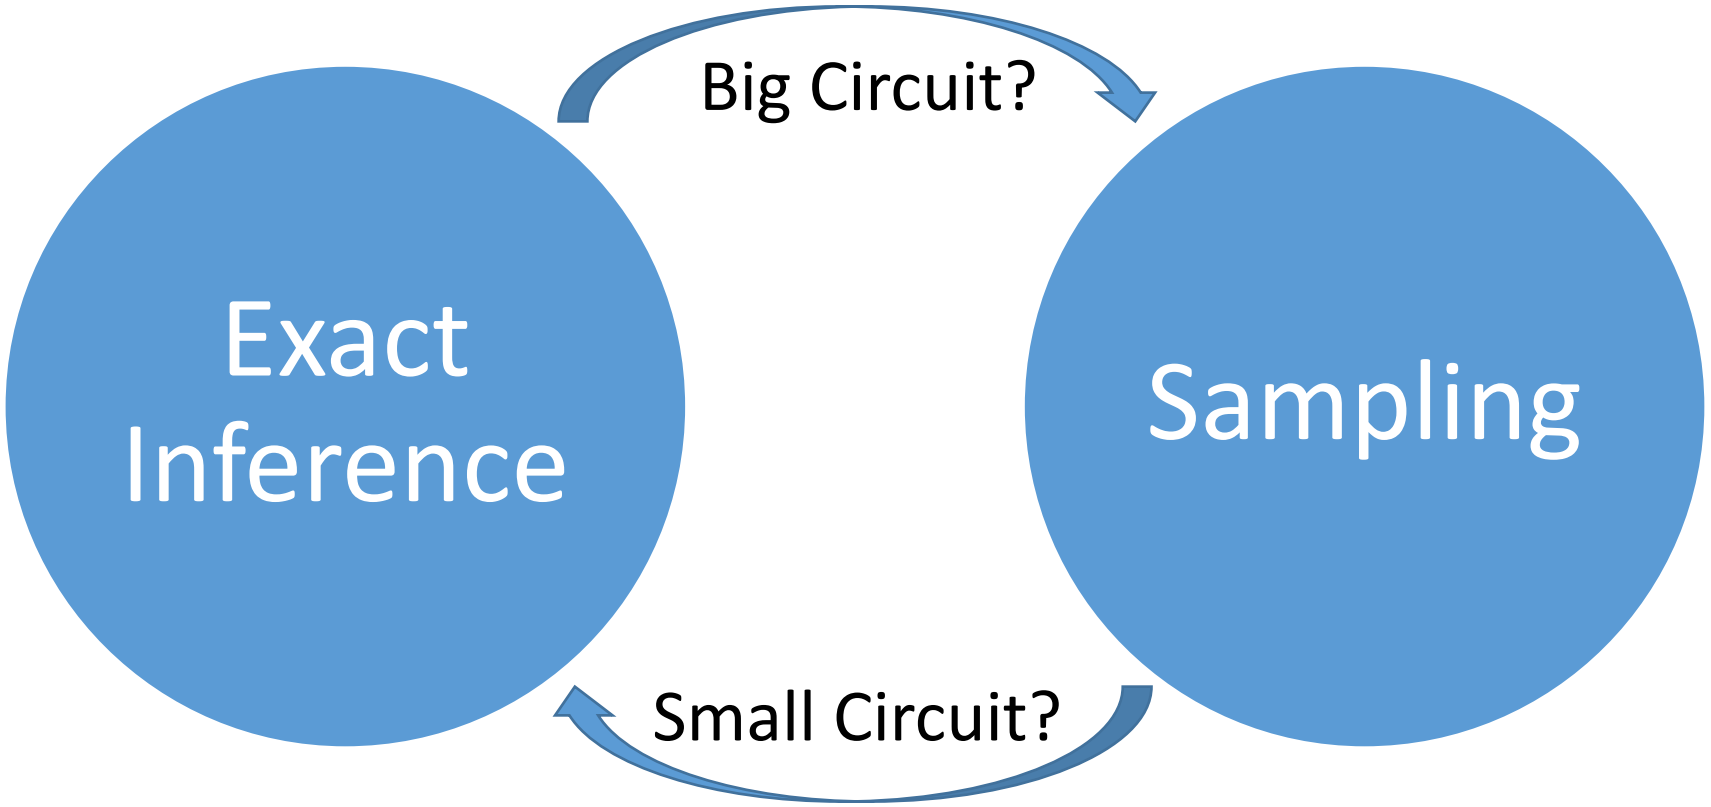

1. **What/when do we sample?**
2. How do we sample?
3. How do we do exact inference?

What/when do we sample?

When: Circuit too big

What: Heuristic on current circuit

Collapsed Compilation

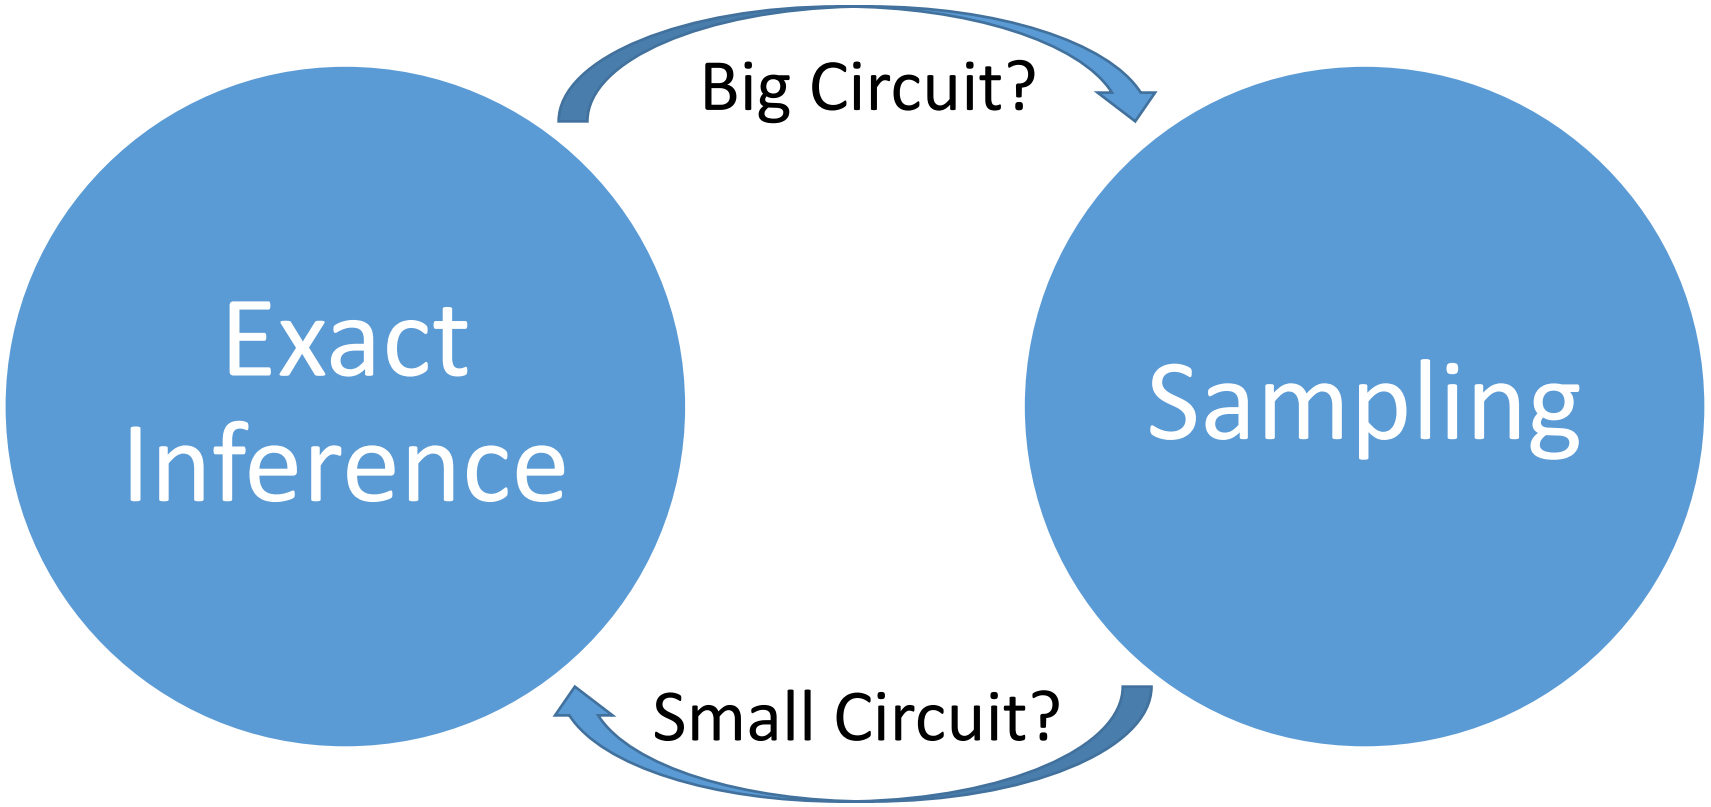

1. What/when do we sample?
2. **How do we sample?**
3. How do we do exact inference?

Motivation: Arithmetic Circuit

- Exact inference: Use Knowledge Compilation (e.g. BDD, SPN)

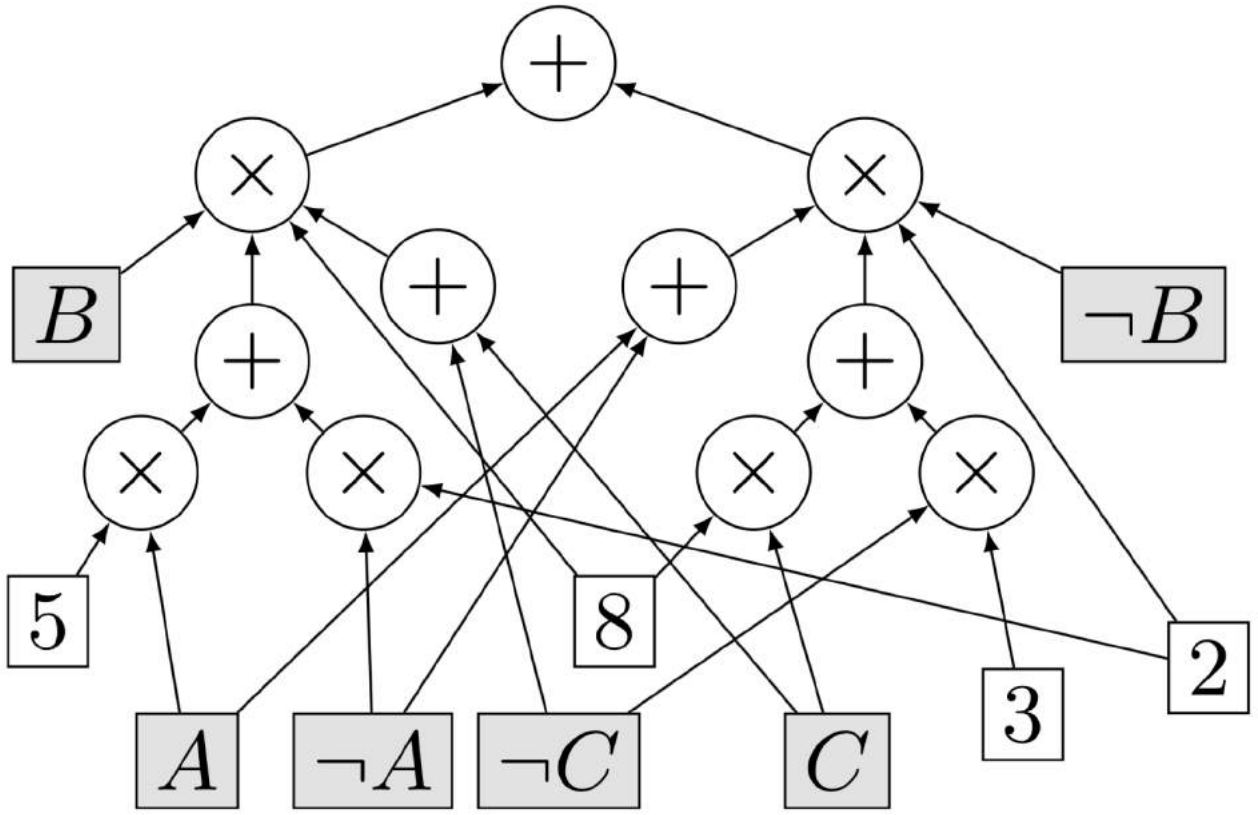

- **Tractable** form: easy queries + operations
- Take advantage of further independence properties, logical structure

How do we sample?

Compute the marginal of the variable in the current circuit!

Collapsed Compilation

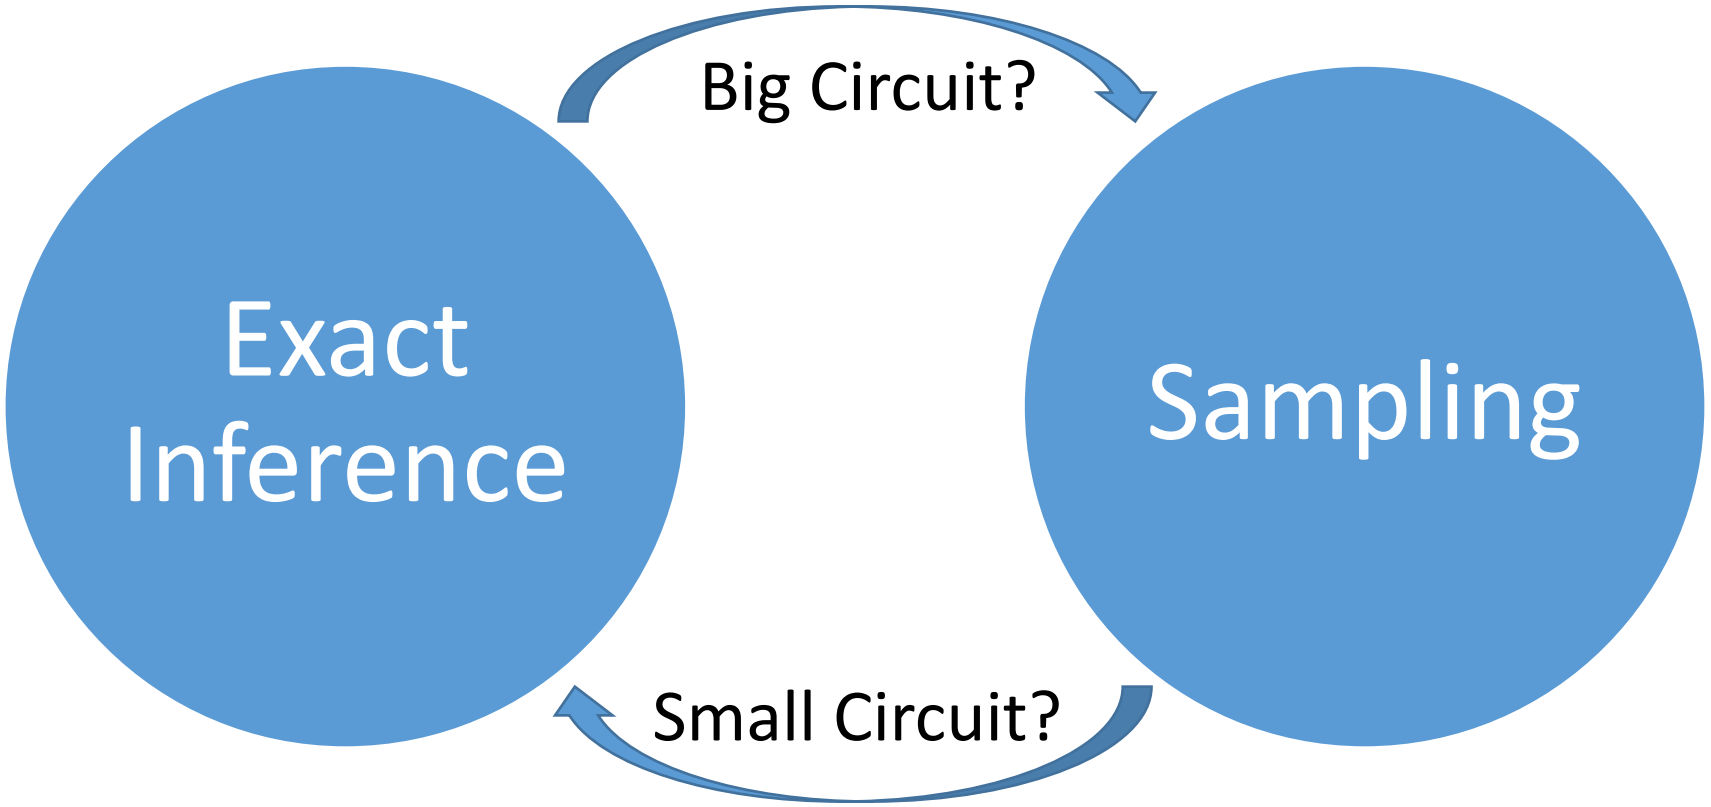

1. What/when do we sample?
2. How do we sample?
3. **How do we do exact inference?**

Conditional Exact Inference

Result is a circuit: **any** joint can be computed efficiently & exactly

Online Collapsed Importance Sampling

0.6 x

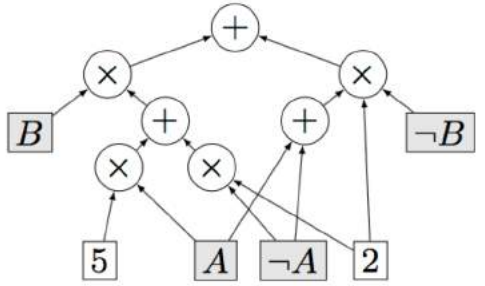

0.9 x

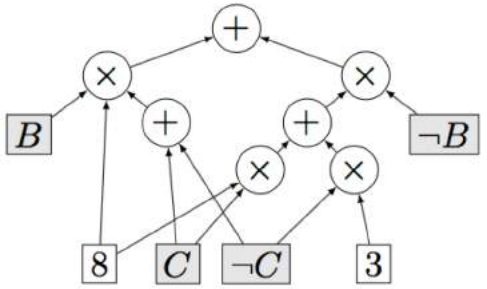

1.1 x

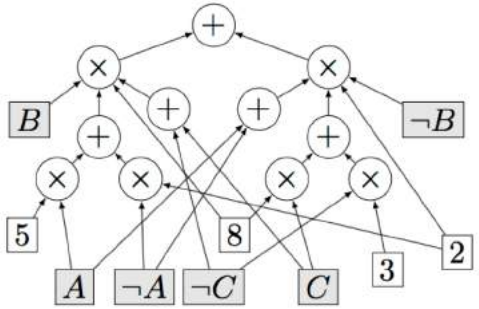

Approximate any query!

Experiments

- Approximate marginal in factor graph
- Algorithmically limit exact inference

Experiments

Table 2: Hellinger distances across methods with internal treewidth and size bounds

Method	50-20	75-26	DBN	Grids	Segment	linkage	frust
EDBP-100k	2.19e−3	3.17e−5	6.39e−1	1.24e−3	1.63e−6	6.54e−8	4.73e−3
EDBP-1m	7.40e−7	2.21e−4	6.39e−1	1.98e−7	1.93e−7	5.98e−8	4.73e−3
SS-10	2.51e−2	2.22e−3	6.37e−1	3.10e−1	3.11e−7	4.93e−2	1.05e−2
SS-12	6.96e−3	1.02e−3	6.27e−1	2.48e−1	3.11e−7	1.10e−3	5.27e−4
SS-15	9.09e−6	1.09e−4	(Exact)	8.74e−4	3.11e−7	4.06e−6	6.23e−3
FD	9.77e−6	1.87e−3	1.24e−1	1.98e−4	6.00e−8	5.99e−6	5.96e−6
MinEnt	1.50e−5	3.29e−2	1.83e−2	3.61e−3	3.40e−7	6.16e−5	3.10e−2
RBVar	2.66e−2	4.39e−1	6.27e−3	1.20e−1	3.01e−7	2.02e−2	2.30e−3

Knowledge Compilation

Collapsed Compilation

Sampling

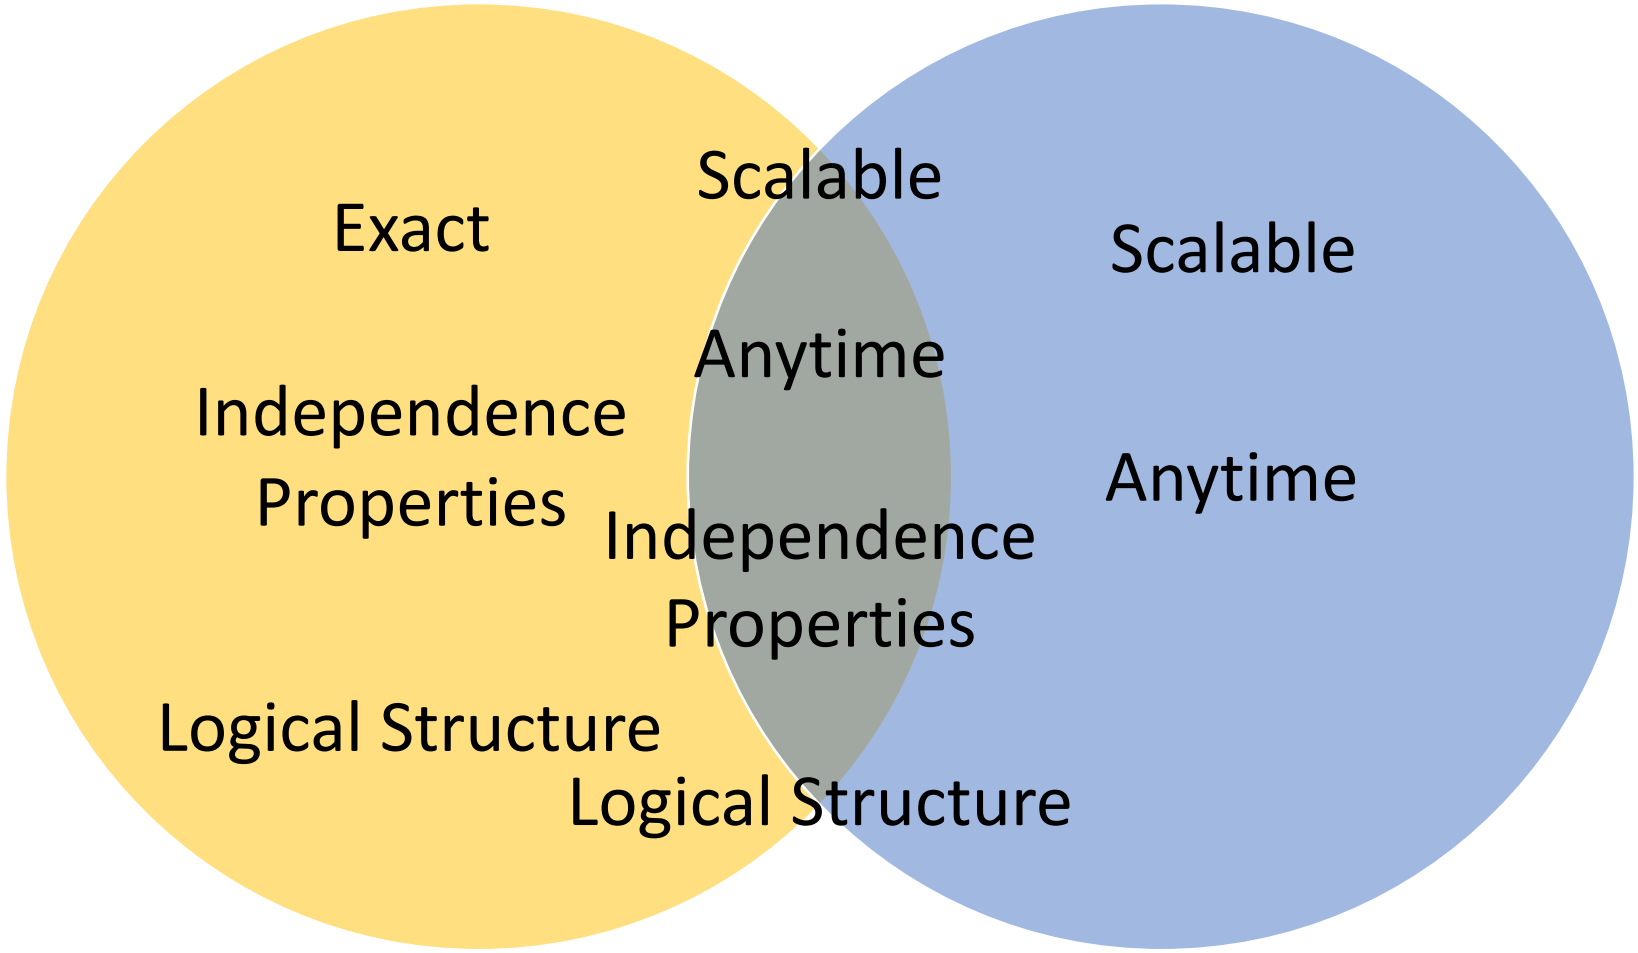

Thanks!

Poster: Room 210 #5

Code: github.com/UCLA-StarAI/Collapsed-Compilation