

Optimizing PPL model evaluation with graphs and applicative functors

Viktor Palmkvist

Museum of Natural History

Inference in Probabilistic Programming

- Core idea: probabilistic model = probabilistic program
- Inference runs and modifies programs

Markov Chain Monte Carlo (MCMC)

1. Produce a sample (s_0).
2. Produce another sample (s').
3. Compare the two likelihoods.
 1. If proposal is likelier, accept (i.e., $s_1 := s'$).
 2. Otherwise, accept with probability $\frac{L(s')}{L(s_0)}$ (reject means $s_1 := s_0$).

Observation: most models are “smooth” \Rightarrow “close” samples have similar likelihood.

Problem: how do we efficiently produce a proposal that is
“close” to the previous sample?

Defining “Close”

- General, must handle *all* programs.

(Definition: “close” means one `assume` changed)


Example:

```
1 model function example() => Real {  
2   assume a ~ Gaussian(0.0, 1.0);  
3   observe expensive(a) ~ Gaussian(0.0, 1.0);  
4   assume b ~ Gaussian(0.0, 1.0);  
5   observe expensive(b) ~ Gaussian(0.0, 1.0);  
6   assume c ~ Gaussian(0.0, 1.0);  
7   observe expensive(c) ~ Gaussian(0.0, 1.0);  
8   return a + b + c;  
9 }
```

} Unchanged


} Unchanged

Data Dependencies



It's an applicative functor!


A Probabilistic Applicative Functor



```
1 type PVal a
2 let pure  : a -> PVal a = ...
3 let map   : (a -> b) -> PVal a -> PVal b = ...
4 let apply : PVal (a -> b) -> PVal a -> PVal b = ...
5
6 let assume : PVal (Dist a) -> PVal a = ...
7 let weight : PVal Float -> () = ...
```


mcore

Generating and Optimizing a Graph



```
1 map f (pure a) =  
2   pure (f a)  
3 apply (pure f) a =  
4   map f a  
5 map f (map g a) =  
6   map (f . g) a  
7 map f (apply a b) =  
8   apply (map (lam a. f . a))  
   a) b
```

Graphical Models



- This looks like a graphical model
 - ...with deterministic nodes
 - ...and observations as hexagons
- What about universality?

TreePPL is a Universal Probabilistic Programming Language

- Statically unbounded number of random variables
- How do we express this?

```
1 model geometric(p: Real) => Int {  
2   assume c ~ Bernoulli(p);  
3   if c {  
4     return 1 + geometric(p);  
5   }  
6   return 1;  
7 }
```

tppl

- Probabilistically guarded recursion
- **Problem:** this cannot be expressed with a (finite) applicative functor


Solution: Monad

```
1 type PVal a
2 let pure : a -> PVal a = ...
3 let map : (a -> b) -> PVal a -> PVal b = ...
4 let apply : PVal (a -> b) -> PVal a -> PVal b = ...
5 let bind : (a -> PVal b) -> PVal a -> PVal b = ...
6
7 let assume : PVal (Dist a) -> PVal a = ...
8 let weight : PVal Float -> () = ...
```

mcore

- Intuition:
 - ▶ `apply` can run things “in parallel”.
 - ▶ `bind` must run things “in sequence”.

Manual Geometric Distribution as a Graph



```
1 model geometric(p: Real) => Int {  
2   assume c ~ Bernoulli(p);  
3   if c {  
4     return 1 + geometric(p);  
5   }  
6   return 1;  
7 }
```

tppl

- `bind` builds and discards sub-graphs as needed
- This can be expensive

Conclusion

Not mentioned here and future work

- Implementation strategies (e.g., mutability or not)
- Integrating pruning, delayed sampling, etc.

Take-aways

- Applicative functor gives us speed
 - ▶ ...by modelling data-flow and avoiding recomputation
- Monad gives us expressivity
 - ▶ ...by nesting graphs
- Together, they generalize graphical models to handle universality

Thank you!