

Modular Probabilistic Models

via Algebraic Effects


Minh Nguyen, Roly Perera, Meng Wang, Nicolas Wu

University of Bristol




Interacting with a probabilistic model

A *probabilistic model* is a set of relationships between certain random variables:



$\lambda x.$	}	input
$\mu \sim \text{Normal}(0, 3)$	}	parameters
$c \sim \text{Normal}(0, 2)$	}	
$\sigma \sim \text{Uniform}(1, 3)$	}	output
$y \sim \text{Normal}(\mu * x + c, \sigma)$	}	



What might this look like in a PPL (probabilistic programming language)?

A possible simulation

```
simulateLinRegr x μ c σ = do
  y ← sample (normal (μ * x + c) σ)
  return y
```

A possible inference

```
inferLinRegr x y = do
  μ ← sample(normal 0 3)
  c ← sample(normal 0 2)
  σ ← sample(uniform 0 3)
  observe (normal (μ * x + c) σ) y
  return (μ, c, σ)
```

*new model
interactions
=
new model
implementations*



Interacting with a probabilistic model

Simulation in WebPPL

```
var linRegr = function(x, mu, c,  $\sigma$ ) {  
  y = sample(Normal(mu * x + c,  $\sigma$ ), y)  
  return y  
}
```

Simulation in Anglican

```
(defquery linRegr [x, mu,  $\sigma$ , c]  
(let [y (sample(normal (c + (* mu) x))  $\sigma$ ))] (let [mu (sample mu-prior)  
  c (sample c-prior)  
   $\sigma$  (sample sigma-prior)  
  predictive (fn [x] (normal (c + (map (* mu) x))  $\sigma$ ))] (observe (predictive x) y)  
  { :mu mu : $\sigma$   $\sigma$  :predictor (predictive x) })))
```

Inference in WebPPL

```
var linRegr = function(x, y) {  
  mu = sample(Normal(0, 3))  
  c = sample(Normal(0, 2))  
   $\sigma$  = sample(Uniform(0, 2))  
  observe(Normal(mu * x + c,  $\sigma$ ), y)  
  return (mu, c,  $\sigma$ )  
}
```

Inference in Anglican

```
(defquery linRegr [x mu-prior  $\sigma$ -prior c-prior y]  
(let [mu (sample mu-prior)  
  c (sample c-prior)  
   $\sigma$  (sample sigma-prior)  
  predictive (fn [x] (normal (c + (map (* mu) x))  $\sigma$ ))] (observe (predictive x) y)  
  { :mu mu : $\sigma$   $\sigma$  :predictor (predictive x) })))
```



Motivation 1: Multimodal models

Multimodal model: A model whose random variables can be specialised to sample or observe modes

Wasabaye: A Haskell PPL for multimodal models

Example: Linear regression

Inputs = x
Parameters = μ, c, σ
Outputs = y

$\mu \sim \text{Normal}(0, 3)$
 $c \sim \text{Normal}(0, 2)$
 $\sigma \sim \text{Uniform}(1, 3)$
 $y \sim \text{Normal}(\mu * x + c, \sigma)$

What might this look like in Wasabaye?

"model environment"

```
linRegr :: Observables env ["μ", "c", "σ", "y"] Double  
        => [Double] -> Model env es [Double]
```

```
linRegr xs = do
```

```
  μ <- normal 0 3 #μ
```

```
  c <- normal 0 2 #c
```

```
  σ <- uniform 1 3 #σ
```

```
  ys <- mapM (λx -> normal (m * x + c) σ #y) xs
```

```
  return ys
```

"observable variable"



Motivation 1: Multimodal models

```
linRegr :: Observables env ["μ", "c", "σ", "y"] Double
        => [Double] -> Model env es [Double]

linRegr xs = do
  μ <- normal 0 3 #μ
  c <- normal 0 2 #c
  σ <- uniform 1 3 #σ
  ys <- mapM (λx -> normal (m * x + c) σ #y) xs
  return ys
```

Interacting with a multimodal model

do -- simulation


```
let xs      = [0 .. 100]
    envin   = (#μ := [3]) • (#c := [0]) • (#σ := [1]) • (#y := [])
    ys      <- simulate (linRegr xs) envin
```

← We observe μ, c, σ
We sample y

-- inference

```
let envin   = (#μ := []) • (#c := []) • (#σ := []) • (#y := ys)
    (envouts, weights) <- lw 1000 (linRegr xs) envin
    let μs = map (get #μ) envouts
    return (μs, weights)
```

← We sample μ, c, σ
We observe y





Motivation 2: First-class models

So, PPLs with multimodal models do already exist.

But models generally aren't first-class citizens.


Supported model features	Wasabaye	Gen	Turing	Stan	Pyro
Multimodal	●	●	●	●	◐
Modular	●	●	●	○	●
Higher-order	●	◐	○	○	●
Type-safe	●	○	○	●	○

- Full support
- ◐ Partial support
- No support



Compositional multimodal models

Hidden Markov Model (HMM)



We can decompose a HMM into two sub-models:

```
type TransModel env es x = x -> Model env es x
```

```
type ObsModel env es x y = x -> Model env es y
```

And then define a HMM as a higher-order model:

```
hmm :: TransModel env es x -> ObsModel env es x y -> Int -> x -> Model env es x
```

```
hmm transModel obsModel n x0 = do
```

```
  let hmmNode x = do x' <- transModel x
                    y' <- obsModel x'
                    return x'
```

```
  foldl (>=>) return (replicate n hmmNode) x0
```

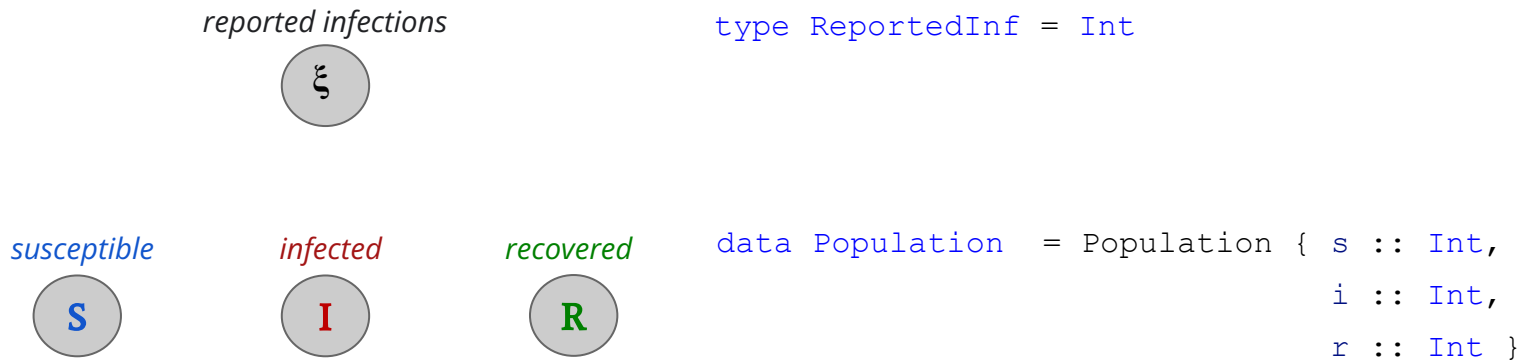
```
(>=>) :: (a -> Model env es
b)
-> (b -> Model env es
c)
```

```
-> (a -> Model env es
```




Modelling an epidemic: The SIR model

Setting: We assume a fixed population of **susceptible**, **infected**, and **recovered** individuals.



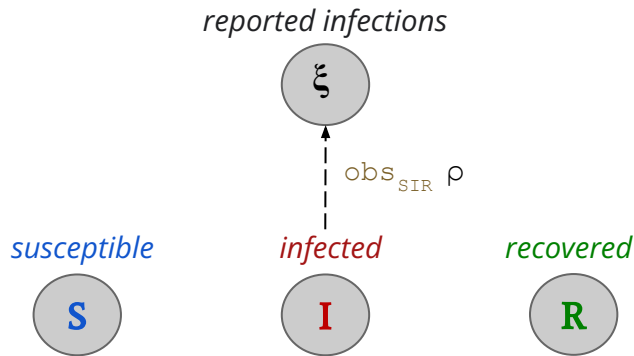
SIR Model: During an epidemic, how do these populations vary over time (days)?



Modelling an epidemic: The SIR model

SIR observation model

```
type ObsModel env es x y = x -> Model env es y
```



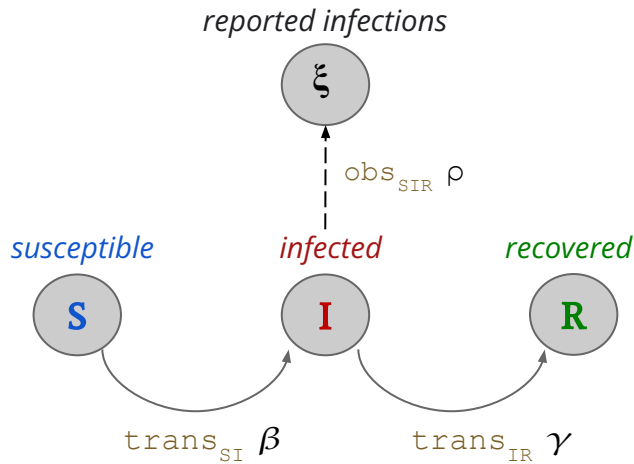
```
obsSIR :: Observable env "ξ" Int  
=> Double -> ObsModel env es Population ReportedInf  
obsSIR ρ (Population _ i _) = poisson (ρ * i) #ξ
```



Modelling an epidemic: The SIR model

SIR transition model

```
type TransModel env es x = x -> Model env es x
```



```
transSI :: Double -> TransModel env es Population
transSI β (Population s i r) = do
  δsi <- binomial' s (1 - exp (-β * i / (s + i + r)))
  return $ Population (s - δsi) (i + δsi) r



transIR :: Double -> TransModel env es Population
transIR γ (Population s i r) = do
  δir <- binomial' i (1 - exp (-γ))
  return $ Population s (i - δir) (r + δir)
```

```
transSIR :: Double -> Double
          -> TransModel env es Population
transSIR β γ = transSI β ==> transIR γ
```



Modelling an epidemic: The SIR model

SIR as a Hidden Markov Model



```
obsSIR ρ (Population _ i _)
  = poisson (ρ * i) #ξ
```

```
transSIR β γ ω φ
  = transSI β ==> transIR γ
    ==> transRS ω
    ==> transSV φ
```

```
let sir0 = Population { s = 760, i = 1, r = 0 }
  n_days = 100
hmm (obssir ρ) (transSIR β γ ω φ) n_days sir0
```



Models as Algebraic Effects

```
newtype Model env es a =  
  Model { runModel :: (Member Dist es, Member (ObsReader env) es) => Prog es a }
```

Infrastructure:

A program containing *syntactic operations* `op` belonging some *effect* `E` in the *signature* `es`

```
data Prog es a = Val a | Op op k   where op : E ∈ es
```

Effect 1: Primitive distributions

```
data Dist a where  
  Normal    :: Double -> Double -> Maybe Double -> Dist Double  
  Bernoulli :: Double -> Maybe Bool -> Dist Bool  
                                     optional observed value
```

Effect 2: Reading observable variables from a model environment `env`

```
data ObsReader env a where  
  Ask :: Observable env x a => ObsVar x -> ObsReader env (Maybe a)  
                                     env assigns x a list of type [a]   type-level string  
                                     #foo :: ObsVar "foo"
```



Models as Algebraic Effects

Example: Desugaring models

```
coinFlip :: (Observable env "p" Double,  
            Observable env "y" Bool)  
          => Model env es Bool
```

```
coinFlip = do  
  p <- uniform 0 1 #p  
  y <- bernoulli p #y  
  return y
```

desugars to



```
coinFlip :: (Observable env "p" Double,  
            Observable env "y" Bool)  
          => Model env es Bool
```

```
coinFlip = do  
  maybe_p <- send (Ask #p)  
  p        <- send (Uniform 0 1 maybe_p)  
  maybe_y <- send (Ask #y)  
  y        <- send (Bernoulli p maybe_y)  
  return y
```



Model Environments

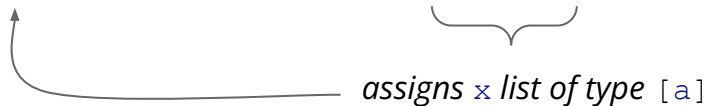
Model environments:

Extensible records from observable variables to lists of values

```
data Env env where
```

```
ENil  :: Env '[]
```

```
ECons :: [a] -> Env env -> Env ((x := a) : env)
```

 *assigns x list of type [a]*

```
env :: Env ["μ" := Double, "c" := Double, "σ" := Double, "y" := Double]
```

```
env = (#μ := [3.0]) • (#c := [0.0]) • (#σ := [1.0]) • (#y := []) • nil
```

```
loop n = do
```

```
  y <- normal 0 1 #y
```

```
  if n <= 0 then return () else loop (n-1)
```



Executing Models with Effect Handlers

An effect handler interprets an effect in es

```
handler :: Prog es a -> Prog es' b
```

Handler 1: Handling reading of observable variables:

```
handleObsReader :: Env env -> Prog (ObsReader env : es) a -> Prog es a
handleObsReader env (Op (Ask x) k) = do
  let maybe_v = getHead x env
      env'     = setTail x env
  in handleObsReader env' (k maybe_v)
```

Handler 2: Handling primitive distributions:

```
handleDist :: Prog (Dist : es) a -> Prog (Observe : Sample : es) a
handleDist (Op (Normal  $\mu$   $\sigma$  maybe_v) k) =
  case maybe_v of Just v -> (handleDist . k) (send $ Observe (Normal ..) v)
                 Nothing -> (handleDist . k) (send $ Sample (Normal ..) )
```

```
data Observe a where
  Observe :: Dist a -> a -> Observe a
```

```
data Sample a where
  Sample :: Dist a -> Sample a
```



Executing Models Compositionally

Interpreting multimodal models to samples and observes

```
handleCORE :: Model env es a -> Prog (Observe : Sample : es') a
```

```
handleCORE env = handleObsReader env ◦ handleDist ◦ runModel
```

```
coinFlip = do  
  p <- uniform 0 1 #p  
  y <- bernoulli p #y  
  return y
```

handle_{CORE}
→

(#p := [0.7]) • (#y := [])

```
coinFlip = do  
  p <- send (Op (Observe (Uniform ..) 0.7  
  y <- send (Op (Sample (Bernoulli ..))  
  return y
```




Executing Models Compositionally

Interpreting multimodal models to samples and observes

```
handleCORE :: Model env es a -> Prog (Observe : Sample : es') a
```

```
handleCORE env = handleObsReader env ◦ handleDist ◦ runModel
```

Simulation



```
traceSamples :: Member Sample es => Prog es a -> Prog (a, SampleTrace)
```

```
handleObserve :: Prog (Observe : es) a -> Prog es a
```

```
handleObserve (Op (Observe d y) k) = handleObserve (k y)
```

```
handleSample :: Prog (Sample : []) a -> IO a
```


```
handleSample (Op (Sample d) k) = IO.sample d >>= (handleSample ◦ k)
```



Executing Models Compositionally

Likelihood weighting

`handleSample` ◦ `handleObserveLW 0` ◦ `traceSamples` ◦ `handleCORE`




Accumulate the log probabilities of observed values

```
handleObserveLW :: Double -> Prog (Observe : es) a -> Prog es (a, Double)
```

```
handleObserveLW p (Op (Observe d y) k) = handleObserveLW (p + logProb d y) (k y)
```

Metropolis-Hastings

`handleSampleMH α` ◦ `handleObserve` ◦ `traceLogProbs` ◦ `traceSamples` ◦ `handleCORE`



selectively sample from chosen address

for each Observe + Sample operation: install and handle a `State` effect



Executing Models Compositionally



What I'm up to

Formalising the language metatheory

So far
A lambda calculus based on algebraic effects,
extended with: random variables, primitive distributions, multimodal models.

Aims
Understanding what metatheory and properties we would like to show for this language

- (Denotational) semantics of probabilistic models under model environments:
what underlying distributions do they denote?

Using row polymorphism to elegantly express effects and model environments

Exploring effect handlers for compositional inference

You can play with Wasabaye!

<https://github.com/min-nguyen/wasabaye/>



Example program

$\Omega = (\mu : \text{Double}) \cdot (\sigma : \text{Double}) \cdot (c : \text{Double}) \cdot (y : \text{Double})$

$\rho : \Omega$
 $\rho = (\mu, [3]) \cdot (\sigma, [1]) \cdot (c, [3]) \cdot (y, [1])$

$M : \text{Double} ! \text{Dist}_{\Omega}$

```
let linearRegression : Double → Double ! DistΩ
  linearRegression = model(x : Double).
    (let μ ~ normal (1, 2)      in
     let σ ~ uniform (1, 3)    in
     let c ~ normal (0, 5)     in
     let y ~ normal (μ * x + c, σ) in
     return y)
in linearRegression 7
```

$N : \text{Double} ! \text{Observe} \cdot \text{Sample}$

```
with { return x          → return x
      , distφ (A, Just y) k → let y' ← observeφ (A, y)
                               in k y'
      , distφ (A, Nothing) k → let y' ← sampleφ A
                               in k y' }
handle (linearRegression 7)
```

Note: $\text{let } x = V \text{ in } M \rightsquigarrow (\lambda x \rightarrow M) V$

Note: the reduction shown above will only partially happen, as evaluation will get stuck on the first unhandled operation $\mathcal{E}[\text{op } V]$.

$\rho' : \Omega$
 $\rho' = (\mu, [1]) \cdot (\sigma, [1]) \cdot (c, [1]) \cdot (y, [1])$

$M : \text{Double} ! \text{Dist}_{\Omega}$

```
let μ ← distnormal ((1, 2), Just 3)      in
let σ ← distuniform ((1, 3), Just 1)    in
let c ← distnormal ((0, 5), Just 3)     in
let y ← distnormal ((μ * 7 + c, σ), Nothing) in
return y
```

→ *

$N : \text{Double} ! \text{Observe} \cdot \text{Sample}$

```
let μ ← observenormal ((1, 2), 3)      in
let σ ← observeuniform ((1, 3), 1)    in
let c ← observenormal ((0, 5), 3)     in
let y ← samplenormal (μ * 7 + c, σ) in
return y
```

→ *

