

# Typed Abstractions for Causal Probabilistic Programming

**Théo Wang\***, Dario Stein\*, Eli Bingham, Jack Feser, Ohad Kammar,  
Michael Lee, Jeremy Yallop


# Our work in context

- **Causal Inference (Pearl)**
  - “Does X cause Y”
  - “Given that we saw Y, had X not been the case, would Y still happen?”
  - Wide applications in empirical sciences (e.g. climate modelling, clinical trials etc.) + Machine learning
- **Causal Probabilistic Programming**
  - Fully-fledged PPLs + First class causal primitives
  - E.g. MultiVerse (Perov et al 2020), OmegaC (Tavares 2021), ChiRho (Basis)
- **Contribution:** ChiRho as a Haskell library
  - Disentangled its dependence on effect handlers and vectorisation
  - Typed abstractions
  - Clarified its semantics

# Overview

- Crash course in causal inference à la Pearl with an example.
- Using typed abstractions à la ChiRho to solve this example.


# A simple example: Dario's Car



```
1 model :: (MonadDistribution m) => m Bool
2 model =
3     fuel <- bernoulli(0.8)
4     battery <- bernoulli(0.9)
5     let carStarts = fuel && battery
6     return carStarts
```

**Dario's car didn't start today. Had he refuelled, would the car have started?**

# A simple example: Dario's Car



```
1 model :: (MonadDistribution m) => m Bool
2 model =
3     fuel <- bernoulli(0.8)
4     battery <- bernoulli(0.9)
5     let carStarts = fuel && battery
6     return carStarts
```

Factual World

**Dario's car didn't start today. Had he refuelled, would the car have started?**


Observing the Factual World

Counterfactual World

Intervening on the Counterfactual World

Result of the query

**Solving the problem *à la Pearl***



Factual World


```

1 model :: (MonadDistribution m) => m Bool
2 model =
3     fuel <- bernoulli(0.8)
4     battery <- bernoulli(0.9)
5     let carStarts = fuel && battery
6     return carStarts
  
```

Counterfactual World


**Dario's car didn't start today. Had he refuelled, would the car have started?**

**(1) Twinning** (e.g. Jacobs '19, Pearl '09) to split the model into two correlated copies/worlds (factual & counterfactual)



```

1 fuel <- bernoulli(0.8)
2 battery <- bernoulli(0.9)
3 let fuel' = fuel
4 let battery' = battery
5 let carStarts = fuel && battery
6 let carStarts' = fuel' && battery'
7 return carStarts'
  
```



```

1 fuel <- bernoulli(0.8)
2 battery <- bernoulli(0.9)
3 let fuel' = fuel
4 let battery' = battery
5 let carStarts = fuel && battery
6 let carStarts' = fuel' && battery'
7 return carStarts'

```


Dario's car didn't start today. Had he refuelled, would the car have started?

Observing the Factual World

Intervening on the Counterfactual World

Result of the query

## (2) Intervening with `do(fuel'=True)` and observing `carStarts`



```

fuel <- bernoulli(0.8)
battery <- bernoulli(0.9)
let fuel' = True -- do-intervention
let battery' = battery
let carStarts = fuel && battery
let carStarts' = fuel' && battery'
condition(carStarts == False) -- conditioning
return carStarts'


```

→ About 64%

# So far...



## Causal operations as program transformations



## First class causal primitives

**χp**

+ Types

Typed and automatic causal  
reasoning à la ChiRho

# Implementing Interventions: the Caus monad

- **Intervention points**: typed identifiers of program locations where interventions are allowed

$$\text{Caus}(X) = \text{InterventionEnv} \rightarrow \text{Prob}(X)$$

- A **model** is a suspended computation executed under **intervention instructions** which specify intervention point behaviour.

```
1 model :: (MonadDistribution m) =>
2     InterventionPoint m Bool -> Caus m (...)
3 model fuelPt =
4     fuel <- bernoulli(0.8)
5     fuelInt <- new_ fuelPt fuel    -- New intervention point for fuel
6     ...
```

# Implementing Interventions ct'd

- **Do** intervention as an environment transformation; apply on model by precomposition

$\text{do} : \text{InterventionPoint}(X) \rightarrow \text{InterventionInstr}(X) \rightarrow (\text{InterventionEnv} \rightarrow \text{InterventionEnv})$

$\text{model} : \text{Caus}(X) = \text{InterventionEnv} \rightarrow \text{Prob}(X)$

$\text{InterventionEnv} \xrightarrow{\text{do } \text{ipt } \text{instr}} \text{InterventionEnv} \xrightarrow{\text{model}} \text{Prob}(X)$

# Implementing Twinning (briefly)

- **Strategy:** eagerly compute all worlds of interest at the same time.
- **Which worlds?:** determined by interventions
- **How?** Multivalue applicative
- Check out my poster!

# Putting it together

Dario's car didn't start today. Had he refuelled, would the car have started?

```
1 model :: MonadDistribution m =>
2     InterventionPoint m Bool -> Caus m (MVal Bool)
3 model fuelPt = do
4     fuel <- sample (pure (bernoulli 0.8));
5     fuelInt <- new_ fuelPt fuel;
6     battery <- sample (pure (bernoulli 0.9));
7     let carStarts = ((&&) <$> fuelInt) <*> battery;
8     return carStarts
```

**Model code**

**Inference code**

```
1 fuelPt :: InterventionPoint Bool <- createKey -- (impl detail)
2 let intervenedConditionedModel :: Caus m Bool = do
3     -- attach intervention instructions for 'fuelPt'
4     carStarts <- do_ fuelPt (Value True)
5             "fuelTrue" (model fuelPt);
6     -- condition on the factual value of carStarts
7     condition (getFactual carStarts == False);
8     -- return its counterfactual value
9     return (getCounterfactual carStarts);
10    -- Run inference
11    avg <- infer (run intervenedConditionedModel)
12    print avg -- approx 0.64
```

# Outlook

- Naively, Pearl's intervention and twinning are program transformations.
- But we can do them automatically and safely using Typed ChiRho (<https://github.com/causal-ppl/chirho-haskell>)
  - Interventions = reader monad
  - Counterfactuals à la ChiRho = applicative
  - More detail on the poster 😊
- Highly expressive (implementation has more examples!)
- Towards semantics for Causal PPLs?

# Backup slides

# Haskell ChiRho vs Python ChiRho

Haskell ChiRho	Python ChiRho
Intervention points	Pyro's string names of sample statements
Do interventions	The do handler in ChiRho (with some difference in the order in which intervention instructions are run)
MVal type	PyTorch tensors
Intervention names	Named tensor dimensions generated and managed by ChiRho and its handlers, e.g. MultiWorldCounterfactual
Applicative structure	Tensor broadcasting
sample: MVal Caus a -> Caus MVal a	Vectorised sampling

# Implementing Twinning (Counterfactuals)

**Strategy:** eagerly compute all worlds of interest at the same time.

$$\text{MVal}(X) \triangleq \coprod_{\substack{N \subseteq_{\text{fin}} \text{Names} \\ \text{Branching points}}} (2^N \rightarrow X)$$

For each generated world (choice of factual (0) or counterfactual (1) per branching point), an X value.

**E.g.**

$$M_1 = ([], \{() \mapsto a\})$$

$$M_2 = ([b], \{(0) \mapsto a_0, (1) \mapsto a_1\})$$

**Which Worlds are of interest?**

**2 Interventions**

1. **do(A=1)**
2. **do(B=2)**

**4 Worlds of interest**

- Fully Factual
- 1 is applied
- 2 is applied
- Both are applied

**Sharing information across worlds via broadcasting**

$$\text{pure} : X \rightarrow \text{MVal}(X)$$

$$\langle *\rangle : \text{MVal}(X \rightarrow Y) \rightarrow \text{MVal}(X) \rightarrow \text{MVal}(Y)$$

$$\text{pure}(f) \langle *\rangle M_1 \langle *\rangle M_2 = ([b], (0) : f(a, a_0), (1) : f(a, a_1))$$

# Graded monad

$$\text{Caus } [A_1 \dots A_n] \ X \triangleq \prod_i \text{List}(\text{InterventionInstr}(A_i)) \rightarrow P(X)$$