

Whittemore: An embedded domain specific language for causal programming

Joshua Brulé

Department of Computer Science
University of Maryland, College Park
jbrule@cs.umd.edu

Abstract

This paper introduces Whittemore, a language for causal programming. Causal programming is based on the theory of structural causal models and consists of two primary operations: *identification*, which finds formulas that compute causal queries, and *estimation*, which applies formulas to transform probability distributions to other probability distribution. Causal programming provides abstractions to declare models, queries, and distributions with syntax similar to standard mathematical notation, and conducts rigorous causal inference, without requiring detailed knowledge of the underlying algorithms. Examples of causal inference with real data are provided, along with discussion of the implementation and possibilities for future extension.

Introduction

In computer science, programming has come to refer to two related concepts. Programmers generally think of programming as describing a process to be executed, usually by writing a sequence of instructions, i.e. code. Mathematicians often think of programming in the sense of linear programming: declaring a problem in terms of satisfiability and/or optimization criteria, to be solved by particular algorithms.

The ‘mathematical’ and ‘code’ senses of programming overlap in practice. Some notable families of programming languages that realize the more mathematical aspect of programming include logic/relational programming (Byrd 2009), and probabilistic programming (Mansinghka 2009). Logic programming is largely based on defining formulas in first-order logic. Probabilistic programming can be seen a language for defining probability distributions, with an operator that implements conditional sampling (Goodman et al. 2008).

This paper introduces “causal programming”, based on the theory of Pearl’s structural causal models (SCM) and related inference rules and algorithms (Pearl 1995) (Shpitser and Pearl 2008). The key theoretical contribution is to define causal programming as an abstraction over two primary operations: identification, which finds formulas that compute a causal query of interest, and estimation, which ap-

plies formulas to transform probability distributions to other probability distributions.

This paper describes the causal programming language “Whittemore”, which is implemented as an embedded domain specific language. The syntax and semantics of Whittemore are designed represent the underlying mathematical concepts as closely as possible. Readers familiar with both structural causal models and Lisp may wish to go directly to the “Examples and interactive computing” section.

Syntax and semantics

Whittemore is defined as a total (i.e. always terminating), purely functional subset of its host language. The reference implementation of Whittemore is in Clojure, a dialect of Lisp (Hickey 2008). This paper describes Whittemore using the same notation for expressions and data types as Clojure and generally follows the same conventions as idiomatic Clojure code.

An expression in Whittemore is a constant, symbol, or form, $(op\ expr^*)$, where op is a causal programming operator, and $expr$ is an expression. Operators are described using regular expression syntax: $?$ (optional), $*$ (0 or more), $+$ (1 or more), with non-terminals denoted by *italics*.

$$\langle expr \rangle ::= \langle constant \rangle \mid \langle symbol \rangle \mid ((op) \langle expr \rangle^*)$$
$$\langle op \rangle ::= \text{define} \mid \text{model} \mid \text{data} \mid \text{q} \mid \text{identify}$$
$$\quad \mid \text{estimate} \mid \text{measure} \mid \text{signature} \mid \langle distribution \rangle$$

Figure 1: Whittemore grammar

Constants

Constants include standard atomic data types (e.g. integer and floating point numbers, strings, booleans) as well as keywords, which are symbolic identifiers that evaluate to themselves. Keywords begin with a colon and can contain alphanumeric characters and special characters that are not reserved by the host language, e.g. `:x`, `:x'`, `:treatment`, `:z_1` are all valid keywords.

In addition to the atomic data types, constants include the following collection types, with literal syntax:

- Vectors are ordered collections of values, e.g. `[:x :y]`
- Maps are unordered collections that maps unique keys to values, e.g. `{:x 0, :y 1}`
- Sets are unordered collection of unique values, e.g. `#{:x :y}`.

Keywords and sets are optional data types, in that strings can generally be used in place of keywords, and vectors can be used in place of sets, without significantly affecting the semantics of the program. However, keywords and set notation are preferred in some cases where it is useful to have a visual distinction.

Symbols

```
(define symbol docstring? value)
```

Symbols are identifiers that normally refer to another value. The `define` operator binds a symbol to a value, and returns the value. Note that `define` cannot rebind symbols, which is necessary for Whittmore to be a purely functional language.

Identification operators

The identification operators correspond to the task of identification from population distributions, i.e. the limit of infinite samples (Heckman 2005).

Model (`model dag confounding*`)

A Model corresponds to the concept of a semi-Markovian causal diagram (Shpitser and Pearl 2008), representing a class of structural causal models. The `model` operator returns a new Model where `dag` is a map of variables to their parents, and `confounding` is a set of endogenous variables whose background variables are not independent.

```
(define front-door
  (model
    {:x []
     :z [:x]
     :y [:z]}
    #{:x :y}))
  X = fX(εX)
  Z = fZ(x, εZ)
  Y = fY(z, εY)
  εX ⊥ εY
```

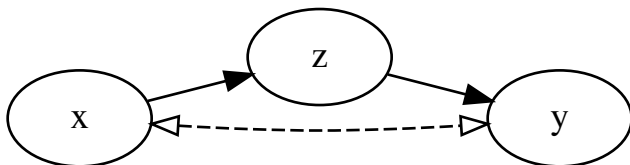


Figure 2: An expression defining a model, equivalent mathematical notation, and the resulting causal diagram.

Data (`data joint`)

Data represents the ‘signature’ of a probability function, i.e. symbolic information about a population probability distribution that is required by the underlying causal inference algorithms. Whittmore currently only supports representing knowledge of joint probability functions, for example, knowledge of the joint probability function, $P(x, y, z)$, is represented as `(data [:x :y :z])`.

Query (`q effect :do do? :given given?`)

A Query is a statistical or causal query, such that the resulting value is a probability distribution. For example, `(q [:y] :given {:x 1})` corresponds to $P(y | X = 1)$, a statistical query. `(q [:y_1 :y_2] :do {:x 0})` corresponds to $P(y_1, y_2 | do(\bar{X} = 0))$, a causal query. Whittmore does not currently support counterfactual queries, although support is planned for a future release.

Note that since `do` and `given` are both optional, they are implemented as keyword arguments in the host language. Their default values are the empty map.

Formula (`identify model data? query`)

The `identify` operator returns a Formula that computes `query`, as a function of `data`, in every SCM entailed by `model`, or a Fail, if such a Formula does not exist. If unspecified, `data` defaults to the joint observational probability function over all endogenous variables in `model`. For example, `(identify front-door (q [:y] :do {:x 0}))`, with implicit `(data [:x :y :z])`, returns the Formula:

$$\sum_z \left[\sum_x P(y | x, z) P(x) \right] P(z | x)$$

where: $x = 0$

Note that Formulas follow lexical scoping rules, e.g. only the ‘outer’ x is bound to 0. The implementation of Formulas is discussed in the ‘Implementation’ section.

The same `identify` expression, but with `(data [:x :y])` would return a Fail describing the hedge (Shpitser and Pearl 2008) that renders identification impossible. Since `identify` is based on the ID algorithm (Shpitser and Pearl 2006), it is complete; a Fail will be returned if and only if no appropriate Formula exists.

Distributions

The causal programming concept of a distribution corresponds to the mathematical concept of a probability distribution. However, an ‘impedance mismatch’ between formulas and distributions exists. Formulas represent the transformation of probability distributions to probability distributions — usually, an observational to an interventional distribution.

However, the structural causal model approach to causal inference is entirely nonparametric, whereas the evaluation of expressions of random variables requires specific knowledge of their underlying distributions.

The core problem is to provide an abstraction over the general concept of a probability distribution, while implementing distribution-specific computations. Whittemore’s solution is to rely on dynamic polymorphism, dispatching on the type of the distribution. A `Distribution` implementation respects the following protocol:

- `(distribution expr*)`
Returns an instance of the probability distribution.
- `(estimate this formula)`
Returns the result of applying a *formula* to *this* distribution, yielding a new distribution. Note that a `Query` acts as a special case of a `Formula`.
- `(measure this event)`
Returns the probability of *event*, i.e. `measure` implements the mathematical concept of a probability measure. An *event* is expected to be a map of keywords to values.
- `(signature this)`
Returns the Data ‘signature’ of the distribution.

Whittemore includes an implementation of a categorical distribution which accepts a single argument of a vector of samples (events) and infers the support of the support of the distribution. For example:

```
(define example-distribution
  (categorical
    [{:x 0, :y 0}
     {:x 0, :y 1}
     {:x 1, :y 0}
     {:x 1, :y 1}
     {:x 1, :y 1}]))
```

binds a representation of a probability distribution where $P(x, y) = 2/5$ and $P(x', y) = P(x, y') = P(x', y') = 1/5$ to the symbol `example-distribution`.

The `Distribution` protocol is user extensible; other probability distributions can be implemented in the host language without modification to Whittemore’s implementation.

Infer and ‘syntactic sugar’

The reference implementation of Whittemore provides some ‘syntactic sugar’ to make causal programming easier. In particular, the `q` operator has three versions that mimic common usage of $P()$ in probability theory:

- ‘Unbound’ query, e.g. `(q [:y] :do [:x])`, a query where *do* and *given* are vectors. An unbound query can still be provided as an argument to `identify`, but the resulting formula cannot be used as an argument to `estimate` without first providing the necessary variable bindings.

- ‘Bound’ query, e.g. `(q [:y] :do {:x 0})`, corresponding to a conditional or interventional distribution (this is considered the canonical version of a `Query`).
- ‘Event’ query, e.g. `(q {:y 1} :do {:x 0})`, corresponding to a specific probability, i.e. *effect* is an event.

Providing an event query to `estimate` implies `measure`. For example, assuming that an appropriate probability distribution is bound to the symbol `smoking`¹:

```
(estimate smoking
  (q {:y 1} :given {:x 1}))
```

Returns the probability 0.8525.

In addition, Whittemore provides the `infer` operator, which combines the functionality of `identify`, `estimate` and `measure`. For example:

```
(infer front-door smoking
  (q {:y 1} :do {:x 1}))
```

Returns the probability 0.4975.

Examples and interactive computing

Whittemore is designed to be used interactively in a notebook interface. The reference implementation of Whittemore has built-in support for Jupyter (Kluyver 2016), an open-source, interactive computational environment. Jupyter integration supports rich output; models are automatically displayed as causal diagrams, and formulas are displayed as \LaTeX math. The code examples in this section are shown along with their Jupyter notebook outputs.

These examples use some additional functions that are not part of ‘core’ Whittemore. To load data, `read-csv` parses and processes a comma-separated values file; `head` returns the first n samples for inspection. To visualize a probability distribution, `plot-univariate` returns a plot of a marginal distribution.

Resolving Simpson’s paradox

The `kidney-distribution` is the empirical probability distribution associated with a study of treatment of renal calculi, i.e. kidney stones (Charig et al. 1986). There are three categorical random variables: `treatment` \in {“surgery”, “nephrolithotomy”}, `size` \in {“small”, “large”}, and `success` \in {“no”, “yes”}. This distribution exhibits Simpson’s paradox, which despite being well known, still continues to “trap the unwary” (Dawid 1979) (Pearl 2014).

¹These examples assume that `smoking` follows the probability distribution in (Pearl 2009, Table 3.1)

```
(define kidney-dataset
  (read-csv "data/renal-calculi.csv"))

(head kidney-dataset 5)
```

:size	:success	:treatment
"small"	"yes"	"surgery"
"large"	"yes"	"nephrolithotomy"
"small"	"yes"	"surgery"
"small"	"yes"	"surgery"
"large"	"yes"	"nephrolithotomy"

Figure 3: The first 5 rows of “renal-calculi.csv”

```
(define kidney-distribution
  (categorical kidney-dataset))

(plot-univariate
  kidney-distribution :success)
```

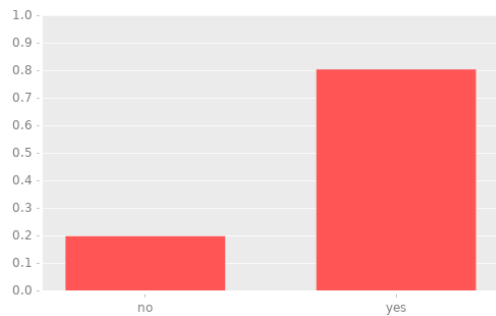


Figure 4: A plot of the marginal distribution of :success

In this distribution, the probability of success, given surgery, is less than the probability of success, given nephrolithotomy:

```
(estimate kidney-distribution
  (q {:success "yes"}
    :given {:treatment "surgery"}))
0.78
```

```
(estimate kidney-distribution
  (q {:success "yes"}
    :given {:treatment
            "nephrolithotomy"}))
0.8257142857142857
```

However, a reversal appears when conditioning on subgroups. When restricted to observing patients with small kidney stones, surgery appears to be the superior treatment:

```
(estimate kidney-distribution
  (q {:success "yes"}
    :given {:size "small"
            :treatment "surgery"}))
0.9310344827586207
```

```
(estimate kidney-distribution
  (q {:success "yes"}
    :given {:size "small"
            :treatment
            "nephrolithotomy"}))
0.8666666666666667
```

And when restricted to observing patients with large kidney stones, surgery again appears to be the superior treatment:

```
(estimate kidney-distribution
  (q {:success "yes"}
    :given {:size "small"
            :treatment "surgery"}))
0.7300380228136882
```

```
(estimate kidney-distribution
  (q {:success "yes"}
    :given {:size "small"
            :treatment
            "nephrolithotomy"}))
0.6875
```

Researchers familiar with the do-calculus will immediately recognize that resolving the question of which treatment is superior is a causal, not statistical query. The aim of causal programming is to answer such queries without the need for the user to understand the details of causal inference. The only requirement is to specify the causal diagram corresponding to the model that produced the original distribution.

```
(define charig1986
  (model
    {:size []
     :treatment [:size]
     :success [:treatment :size]}))
```

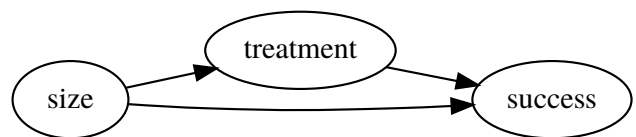


Figure 5: A model where kidney stone size affects the success of treatment and which treatment a patient receives

With respect to the casual assumptions in charig1986 (Charig et al. 1986), an estimate of the casual effect of treatment on success is easily inferred:

```
(infer charig1986 kidney-distribution
  (q {:success "yes"}
    :do {:treatment "surgery"}))
0.8325462173856037

(infer charig1986 kidney-distribution
  (q {:success "yes"}
    :do {:treatment "nephrolithotomy"}))
0.778875
```

Nonstandard adjustments

Note that Whittemore is by no means limited to the special cases of back door and front door adjustment (Pearl 1995). Causal programming easily identifies formulas for computing causal effect that involve non-standard adjustments:

```
(define concomitant-example
  "Figure 1 (f) from (Shpitser 2008)"
  (model
    {:y [:x :z_1 :z_2]
     :z_2 [:z_1]
     :z_1 [:x]
     :x []}
    #{:y :z_1}
    #{:x :z_2}))
```

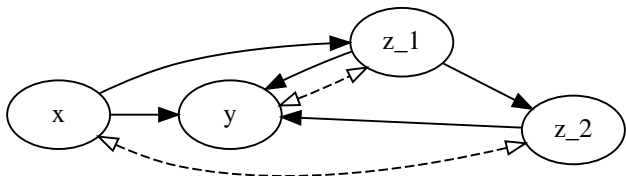


Figure 6: A model where the identification of $P(y | do(x))$ requires summing/integrating over z_1 and z_2

```
(identify concomitant-example
  (q [:y] :do [:x]))
```

$$\sum_{z_1, z_2} \left[\sum_x P(x) P(z_2 | x, z_1) \right] P(z_1 | x) P(y | x, z_1, z_2)$$

Implementation

The ID algorithm and several related algorithms have been previously implemented in the R programming language (Tikka and Karvanen 2017). Whittemore’s `identify` is a purely functional implementation of ID, designed to be easily extensible.

The Model, Data, Query and Formula types are all implemented as persistent (immutable) hash array mapped tries (Bagwell 2001) which support lookup and ‘modification’ (associating a key and value creates a new data structure)

in $\log_{32} N$ time. This provides good performance while remaining free of side effects. A considerable advantage is that the data structures can be freely shared with any other part of a program — it is impossible to corrupt a data structure since none of them can be changed.

Formulas are defined as a map of bindings of variables to values, and a *form*, which is defined recursively:

- `{:p #{vars} :given #{vars}}`
- `{:sum form :sub #{vars}}`
- `{:prod #{forms}}`
- `{:numer form :denom form}`

These forms correspond to a probability expression, summation, product, and fraction, respectively. Formulas follow lexical scoping rules, which obviates the need to rename variables — variable bindings are determined the first surrounding `:sum` that contains the variable as a subscript.

Formulas can be manipulated with standard Clojure functions and are designed to support simplification in a manner similar to a nanopass compiler (Keep 2012) where each ‘pass’ takes a valid form and applies a simple rule to reduce the form to a simpler, still valid form. For example,

```
{:numer {:p #{:y :z :x}},
 :denom {:sub #{:y},
         :sum {:p #{:y :z :x}}}}
```

Can be reduced to `{:p #{:y} :given #{:x :z}}` by applying a marginalization rule on the `:denom` form and then a conditioning rule on the resulting expression.

Additional keys can be added to the Model, Data, Query, and Formula types, without changing the semantics of a program, permitting considerable future extensibility.

Discussion

Whittemore demonstrates that the full causal inference ‘pipeline’ — from raw data to estimates of causal effect — can be effectively abstracted over. Using Clojure as host language blurs the line between a programming language and library: Whittemore syntax is similar to standard mathematical notation, but remains a subset of legal Clojure expressions.

Whittemore currently only supports causal effect identification from observational probability distributions. There are several opportunities for future extension to ‘core’ Whittemore:

- Counterfactuals. A counterfactual query operator (`cf gamma delta?`) can be introduced without affecting the existing syntax and semantics. For example, $P(Y_x | x')$, the effect of treatment on the treated (Shpitser and Pearl 2009), can be represented as `(cf [:y {:x 0}] [{:x 1} {}])`. Supporting counterfactual queries would require extending

identify to implement the IDC* algorithm (Shpitser and Pearl 2007).

- Data fusion. “Data fusion” refers to the problem of causal inference given heterogeneous sources of data (Bareinboim and Pearl 2016). This includes the problems of identification from surrogate experiments (Bareinboim and Pearl 2012a), recovery from selection bias (Bareinboim and Pearl 2012b), and causal transportability (Bareinboim and Pearl 2013). These problems could be represented as causal programs by adding additional Data signatures. For example, the availability of limited experimental data could be represented by (data joint :do surrogate?) and supported by extending identify to implement the zID algorithm (Bareinboim and Pearl 2012a).
- Causal discovery. Causal discovery algorithms generally rely on information about conditional independences between variables (Malinsky and Danks 2017). This can be introduced to causal programming by extending the Data type to include such information.

The Distribution protocol is open to extension, without modifying the implementation of Whittemore. This provides an interesting opportunity to integrate casual programming with probabilistic programming. Probabilistic programming’s primary operator is to efficiently sample from marginal and conditional probability distributions. Causal programming’s primary operator is to identify causal effects in terms of such probabilities.

Whittemore is under active development. The reference implementation is cross-platform and open source: <https://github.com/jtcbrule/whittemore>.

Acknowledgement

This research was supported in part by ONR Award N00014-19-1-2044.

References

- [Bagwell 2001] Bagwell, P. 2001. Ideal hash trees. Technical report, Es Grands Champs.
- [Bareinboim and Pearl 2012a] Bareinboim, E., and Pearl, J. 2012a. Causal inference by surrogate experiments. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence*.
- [Bareinboim and Pearl 2012b] Bareinboim, E., and Pearl, J. 2012b. Controlling selection bias in causal inference. In *Artificial Intelligence and Statistics*, 100–108.
- [Bareinboim and Pearl 2013] Bareinboim, E., and Pearl, J. 2013. Meta-transportability of causal effects: A formal approach. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics*.
- [Bareinboim and Pearl 2016] Bareinboim, E., and Pearl, J. 2016. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences* 113(27):7345–7352.
- [Byrd 2009] Byrd, W. E. 2009. *Relational programming in minikanren: Techniques, applications, and implementations*. Ph.D. Dissertation, Indiana University.
- [Charig et al. 1986] Charig, C. R.; Webb, D. R.; Payne, S. R.; and Wickham, J. E. 1986. Comparison of treatment of renal calculi by open surgery, percutaneous nephrolithotomy, and extracorporeal shockwave lithotripsy. *Br Med J (Clin Res Ed)* 292(6524):879–882.
- [Dawid 1979] Dawid, A. P. 1979. Conditional independence in statistical theory. *Journal of the Royal Statistical Society*.
- [Goodman et al. 2008] Goodman, N. D.; Mansinghka, V. K.; Roy, D.; Bonawitz, K.; and Tenenbaum, J. B. 2008. Church: a language for generative models. In *Uncertainty in Artificial Intelligence*.
- [Heckman 2005] Heckman, J. J. 2005. The scientific model of causality. *Sociological Methodology* 35(1):1–97.
- [Hickey 2008] Hickey, R. 2008. The clojure programming language. In *Dynamic Languages Symposium*.
- [Keep 2012] Keep, A. W. 2012. *A Nanopass framework for Commercial Compiler Development*. Ph.D. Dissertation, School of Informatics and Computing, Indiana University.
- [Kluyver 2016] Kluyver, Thomas, e. a. 2016. *Jupyter Notebooks-a publishing format for reproducible computational workflows*. IOS Press. 87–90.
- [Malinsky and Danks 2017] Malinsky, D., and Danks, D. 2017. Causal discovery algorithms: A practical guide. *Philosophy Compass* 13(1):e12470.
- [Mansinghka 2009] Mansinghka, V. K. 2009. *Natively probabilistic computation*. Ph.D. Dissertation, Massachusetts Institute of Technology.
- [Pearl 1995] Pearl, J. 1995. Causal diagrams for empirical research. *Biometrika* 82(4):669–688.
- [Pearl 2009] Pearl, J. 2009. *Causality*. Cambridge University Press.
- [Pearl 2014] Pearl, J. 2014. Comment: Understanding Simpson’s paradox. *The American Statistician* 68(1):8–13.
- [Shpitser and Pearl 2006] Shpitser, I., and Pearl, J. 2006. Identification of joint interventional distributions in recursive semi-markovian causal models. In *21st National Conference on Artificial Intelligence and the 18th Innovative Applications of Artificial Intelligence Conference, AAAI-06/IAAI-06*, 1219–1226.
- [Shpitser and Pearl 2007] Shpitser, I., and Pearl, J. 2007. What counterfactuals can be tested. In *Uncertainty in Artificial Intelligence*.

- [Shpitser and Pearl 2008] Shpitser, I., and Pearl, J. 2008. Complete identification methods for the causal hierarchy. *Journal of Machine Learning Research* 9(Sep):1941–1979.
- [Shpitser and Pearl 2009] Shpitser, I., and Pearl, J. 2009. Effects of treatment on the treated: Identification and generalization. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*, 514–521. AUAI Press.
- [Tikka and Karvanen 2017] Tikka, S., and Karvanen, J. 2017. Identifying causal effects with the r package causal-effect. *Journal of Statistical Software*.