

# Generative DAGs as an Interface Into Probabilistic Programming with the R Package `causact`

Adam J. Fleischhacker<sup>1</sup> and Thi Hong Nhung Nguyen<sup>2</sup>

<sup>1</sup> Adam Fleischhacker, JP Morgan Chase Faculty Fellow, University of Delaware, Newark, DE 19716 <sup>2</sup> Institute for Financial Services Analytics, University of Delaware, Newark, DE 19716

DOI: [10.21105/joss.04415](https://doi.org/10.21105/joss.04415)

## Software

- [Review](#) ↗
- [Repository](#) ↗
- [Archive](#) ↗

---

Editor: [Arfon Smith](#) ↗ 

## Reviewers:

- [@joethorley](#)
- [@ChristopherLucas](#)

Submitted: 27 April 2022

Published: 12 August 2022

## License

Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License ([CC BY 4.0](#)).

## Summary

The `causact` package provides R functions for visualizing and running inference on generative directed acyclic graphs (DAGs). Once a generative DAG is created, the package automates Bayesian inference via the `greta` package ([Golding, 2019](#)) and TensorFlow ([Dillon et al., 2017](#)). The package eliminates the need for three separate versions of a model: 1) the narrative describing the problem, 2) the statistical model representing the problem, and 3) the code enabling inference written in a probabilistic programming language. Instead, `causact` users create one unified model, a generative DAG, using a visual representation.

## Statement of Need

Bayesian data analysis mixes data with domain knowledge to quantify uncertainty in unknown outcomes. Its beautifully-simple theoretical underpinnings are deployed in three main steps ([Gelman et al., 2013](#)):

- **Modelling:** Joint probability distributions are specified to encode domain knowledge about potential data generating processes.
- **Conditioning:** Bayes rule is used to reallocate plausibility among the potential data generating processes to be consistent with both the encoded domain knowledge and the observed data. The conditioned model is known as the posterior distribution.
- **Validation:** Evidence is collected to see whether the specified model as well as the computational implementation of the model and conditioning process are to be trusted or not.

Algorithmic advances in the *conditioning* step of Bayesian data analysis have given rise to a new class of programming languages called probabilistic programming languages (PPLs). Practical and complex statistical models which are analytically intractable can now be solved computationally using inference algorithms. In particular, Markov Chain Monte Carlo (MCMC) algorithms ([Congdon, 2010](#); [Gelfand & Smith, 1990](#); [Gilks & Roberts, 1996](#)) handle arbitrarily large and complex models via highly effective sampling processes that quickly detect high-probability areas of the underlying distribution [Kruschke \(2014\)](#).

The `causact` package, presented in this paper, focuses on solving a three-language problem that occurs during Bayesian data analysis. First, there is the language of the domain expert which we refer to as the *narrative* of how data is generated. Second, there is the language of *math* where a statistical model, amenable to inference, is written. Lastly, there is the language of *code*, where a PPL language supports computational inference from a well-defined statistical model. The existence of these three languages creates friction as diverse stakeholders collaborate to yield insight from data; often mistakes get made in both communicating and translating between the three languages. Prior to `causact`, any agreed upon narrative of a

data-generating process must ultimately be modelled in code using an error-prone process where model misspecification, variable indexing errors, prior distribution omissions, and other mismatches between desired model and coded model go easily unnoticed.

To unify inference-problem narratives, the statistical models representing those narratives, and the code implementing the statistical models, `causact` introduces a modified visualization of *directed acyclic graphs* (DAGs), called the *generative DAG*, to serve as a more intuitive and collaborative interface into probabilistic programming languages and to ensure faithful abstractions of real-world data generating processes.

## Modelling with Generative DAGs

Generative DAGs pursue two simultaneous goals. One goal is to capture the narrative by building a conceptual understanding of the data generating process that lends itself to statistical modelling. And the second goal is to gather all the mathematical elements needed for specifying a complete Bayesian model of the data generating process. Both of these goals will be satisfied by iteratively assessing the narrative and the narrative's translation into rigorous mathematics using `causact` functions.

Capturing the narrative in code uses the core `causact` functions `dag_create()`, `dag_node()`, `dag_edge()`, and `dag_plate()` which are connected via the chaining operator `%>%` to build a DAG. `dag_render()` or `dag_greta()` are then used to visualize the DAG or run inference on the DAG, respectively. The simplicity with which generative DAGs are constructed belies the complexity of models which can be supported. For example, multi-level or hierarchical models are easily constructed as shown below in code for constructing and visualizing an experiment about chimpanzees (McElreath, 2020b); related data is included in `causact::chimpanzeesDF`. Each chimpanzee “actor” is given a choice to pull one of two levers - right or left. Depending on the lever pulled, the chimpanzee is either acting prosocially, pulling the lever which feeds both himself and a partner, or not prosocially where only the lever-pulling chimpanzee receives food. Figure 1 depicts a varying intercepts model where a baseline proclivity to pull the left lever varies with each of the seven chimpanzee actors being observed.

```
## load packages for example
library(greta)
library(causact)
library(tidyverse)

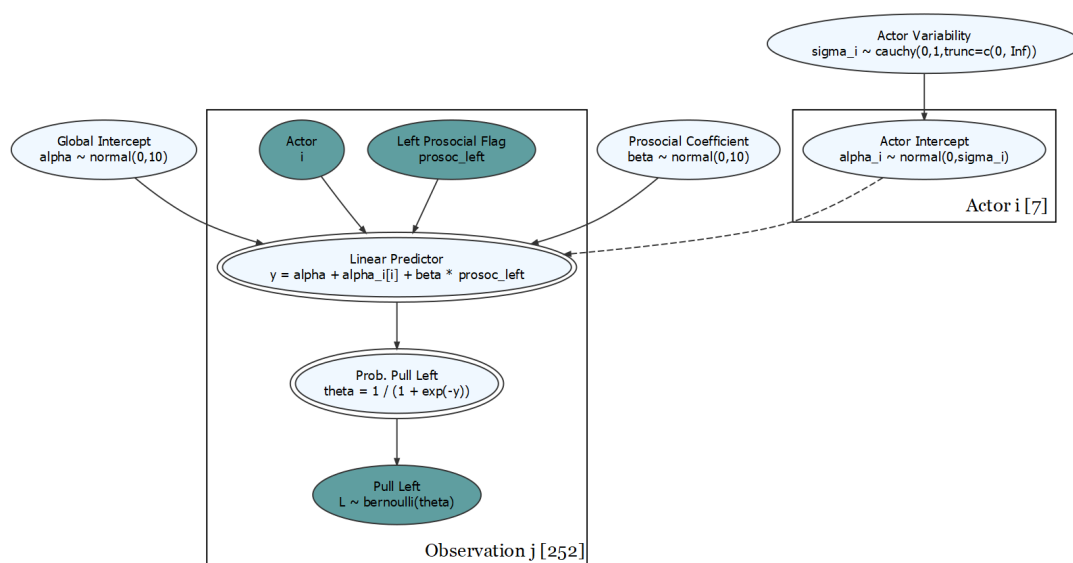
## get only the experiments with partner present
chimpDF = causact::chimpanzeesDF %>% filter(condition == 1)

## create model
graph = dag_create() %>%
  dag_node("Pull Left", "L",
    rhs = bernoulli(theta),
    data = chimpDF$pulled_left) %>%
  dag_node("Prob. Pull Left", "theta",
    rhs = 1 / (1 + exp(-y)),
    child = "L") %>%
  dag_node("Linear Predictor", "y",
    rhs = alpha + alpha_i + beta * prosoc_left,
    child = "theta") %>%
  dag_node("Global Intercept", "alpha",
    rhs = normal(0,10)) %>%
  dag_node("Actor Intercept", "alpha_i",
    rhs = normal(0,sigma_i)) %>%
  dag_node("Prosocial Coefficient", "beta",
    rhs = normal(0,10)) %>%
```

```

dag_node("Left Prosocial Flag","prosoc_left",
  data = chimpDF$prosoc_left) %>%
dag_edge(c("alpha","alpha_i","beta","prosoc_left"),"y") %>%
dag_node("Actor Variability","sigma_i",
  rhs = cauchy(0,1,trunc = c(0,Inf)),
  child = "alpha_i") %>%
dag_plate("Actor","i",
  nodeLabels = c("alpha_i"),
  data = chimpDF$actor,
  addDataNode = TRUE) %>%
dag_plate("Observation","j",
  nodeLabels = c("prosoc_left","i","y","theta","L"))
graph %>% dag_render()

```



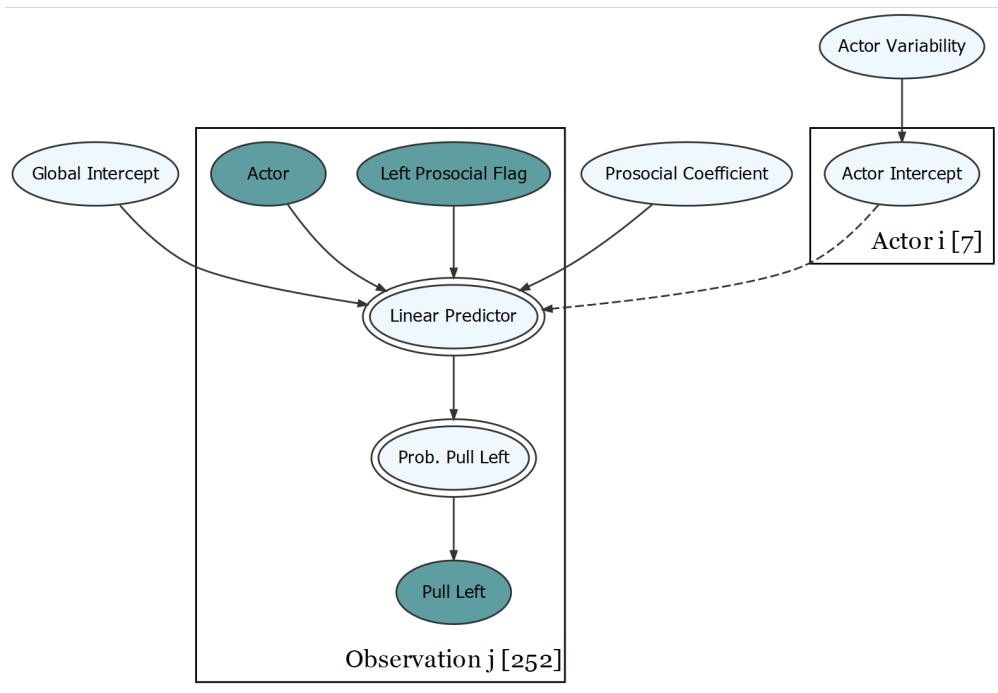
**Figure 1:** A generative DAG of modelling chimpanzee behavior.

Figure 2 replicates Figure 1 without math for less intimidating discussions with domain experts about the model using the `shortLabel = TRUE` argument (shown below). `causact` does not require a complete model specification prior to rendering the DAG, hence, `causact` facilitates qualitative collaboration on the model design between less technical domain experts and the model builder.

```

graph %>% dag_render(shortLabel = TRUE)

```



**Figure 2:** Hiding mathematical details to facilitate collaborations with domain experts.

All visualizations, including Figure 1 and Figure 2, are created via `causact`'s calls to the `DiagrammeR` package (Iannone, 2020). The `dag_diagrammer()` function can convert a `causact_graph` to a `dgr_graph` (the main object when using `DiagrammeR`) for further customizing of a visualization using the `DiagrammeR` package.

Sampling from the posterior of the chimpanzee model (Figure 1) does not require a user to write PPL code, but rather a user will simply pass the generative DAG object to `dag_greta()` and then inspect the data frame of posterior draws:

```
library(greta) ## greta uses TensorFlow to get sample
drawsDF = graph %>% dag_greta()
drawsDF

## # A tibble: 4,000 x 10
##   alpha alpha_i_1 alpha_i_2 alpha_i_3 alpha_i_4 alpha_i_5 alpha_i_6
##   <dbl> <dbl> <dbl> <dbl> <dbl> <dbl> <dbl>
## 1 0.227 -1.01 3.58 -1.40 -1.50 -0.963 0.133
## 2 0.227 -1.01 3.58 -1.40 -1.50 -0.963 0.133
## 3 -0.204 -0.489 3.84 -1.43 -1.14 -0.253 0.863
## 4 0.221 -0.503 3.83 -1.30 -0.835 -0.212 0.0866
## 5 0.189 -1.54 3.99 -1.27 -1.85 -1.60 0.0299
## 6 1.14 -1.58 2.20 -3.06 -2.58 -2.00 -1.41
## 7 -0.195 -0.405 2.75 -1.53 -0.849 -0.750 -0.481
## 8 0.832 -2.04 2.96 -3.12 -2.65 -1.67 -1.71
## 9 1.91 -2.78 1.78 -3.98 -3.15 -2.97 -2.46
## 10 1.01 -2.52 6.37 -2.74 -3.09 -2.37 -0.957
## # ... with 3,990 more rows, and 3 more variables: alpha_i_7 <dbl>,
##   beta <dbl>, sigma_i <dbl>
```

Behind the scenes, `causact` creates the model's code equivalent using the `greta` PPL, but this is typically hidden from the user. However, for debugging or further customizing a model, the

greta code can be printed to the screen without executing it by setting the `mcmc` argument to `FALSE`:

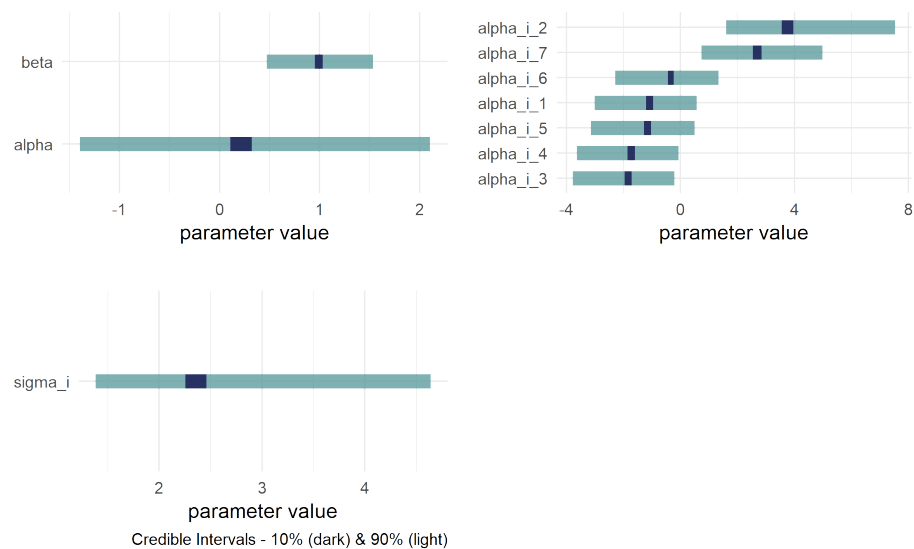
```
graph %>% dag_greta(mcmc=FALSE)

## prosoc_left <- as_data(chimpDF$prosoc_left) #DATA
## L <- as_data(chimpDF$pulled_left) #DATA
## i <- as.factor(chimpDF$factor) #DIM
## i_dim <- length(unique(i)) #DIM
## alpha <- normal(mean = 0, sd = 10) #PRIOR
## beta <- normal(mean = 0, sd = 10) #PRIOR
## sigma_i <- cauchy(location = 0, scale = 1, trunc = c(0, Inf)) #PRIOR
## alpha_i <- normal(mean = 0, sd = sigma_i, dim = i_dim) #PRIOR
## y <- alpha + alpha_i[i] + beta * prosoc_left #OPERATION
## theta <- 1 / (1 + exp(-y)) #OPERATION
## distribution(L) <- bernoulli(prob = theta) #LIKELIHOOD
## gretaModel <- model(alpha,alpha_i,beta,sigma_i) #MODEL
## meaningfulLabels(graph)
## draws <- mcmc(gretaModel) #POSTERIOR
## drawsDF <- replaceLabels(draws) %>% as.matrix() %>%
## dplyr::as_tibble() #POSTERIOR
## tidyDrawsDF <- drawsDF %>% addPriorGroups() #POSTERIOR
```

The produced greta code is shown in the above code snippet. The code can be difficult to digest for some and exemplifies the advantages of working visually using `causact`. The above code is also challenging to write without error or misinterpretation. Indexing is particularly tricky in PPL's and `causact` abbreviates posterior parameters using indexes as determined by provided data (e.g. `as.factor(chimpDF$factor)`) and not by arbitrary numbering of factors in order of appearance.

The output of `dag_greta()` is in the form of a data frame of draws from the joint posterior. To facilitate a quick look into posterior estimates, the `dagp_plot()` function creates a simple visual of 90% credible intervals. It is the only core function that does not take a graph as its first argument. By grouping all parameters that share the same prior distribution and leveraging the meaningful parameter names constructed using `dag_greta()`, it allows for quick comparisons of parameter values.

```
drawsDF %>% dagp_plot()
```



**Figure 3:** Visualizing 10% (darker fill) and 90% (lighter fill) credible intervals for latent parameters.

The code above makes the plot in [Figure 3](#) showing credible intervals for unobserved latent parameters. For example, `alpha_i_2` represents a chimp with strong preference for pulling the left lever; this chimp is affectionately referred to as “Lefty”. For further posterior plotting, users would make their own plots using `ggplot2` ([Wickham, 2016](#)), `ggdist` ([Kay, 2020](#)), or similar. For further model validation, including MCMC diagnostics, the user would use a package like `bayesplot` ([Gabry et al., 2019](#)) or `shinystan` ([Gabry, 2018](#)). For users who prefer to work with an `mcmc` object, they can extract the draws object after running the generated `greta` code from `dag_greta(mcmc=FALSE)` or find the object in the `cacheEnv` environment after running `dag_greta(mcmc=FALSE)` using `get("draws", envir = causact:::cacheEnv)`.

## Comparison to Other Packages

By focusing on generative DAG creation as opposed to PPL code, `causact` liberates users from the need to learn complicated probabilistic programming languages. As such, it is similar in spirit to any package whose goal is to make Bayesian inference accessible without learning a PPL. Perhaps the first such software was `DoodleBUGS` which provided a DAG-based graphical interface into `WinBUGS` ([Lunn et al., 2000](#)). In terms of leveraging more modern PPLs, `causact` is similar to `brms` ([Bürkner, 2017](#)), `rstanarm` ([Goodrich et al., 2020](#)), and `rethinking` ([McElreath, 2020a](#)) - three R packages which leverage `Stan` ([Stan Development Team, 2021](#)) for Bayesian statistical inference with MCMC sampling. Like the `rethinking` package which is tightly integrated with a textbook ([McElreath, 2020b](#)), a large motivation for developing `causact` was to make learning Bayesian inference easier. The package serves a central role in a textbook titled *A Business Analyst’s Introduction to Business Analytics: Intro to Bayesian Business Analytics in the R Ecosystem*. ([Fleischhacker, 2020](#)). As a point of contrast, the `DAGitty` package ([Textor et al., 2017](#)) also focuses on DAG creation/visualization, but `DAGitty`’s intent is to help ensure consistency between the causal assumptions of a researcher and the dataset to which those assumptions should apply; `DAGitty` does not create PPL code for automating inference.

## Conclusion

The `causact` modelling syntax is flexible and encourages modellers to make bespoke models. The long-term plan for the `causact` package is to promote a Bayesian workflow that philosophically

mimics the Principled Bayesian Workflow outlined by Betancourt (2020). The structure of a generative DAG is sure to be much more transparent and interpretable than most other modern machine learning workflows; this is especially true when models are made accessible to those without statistical or coding expertise. For this reason, generative DAGs can help facilitate effective communication between modelers and domain users both during the designing process of the models and when explaining the results returned by the models.

## Acknowledgements

The Stan Development team has been inspirational for this work and has formed a wonderful Bayesian inference community around their powerful language. Additionally, the books of Kruschke (2014) and McElreath (2020b) are tremendous resources for learning Bayesian data analysis and their pedagogy is aspirational. This work would not be possible without the greta dev team and special thanks to Nick Golding and Nick Tierney. Lastly, thanks to the University of Delaware students, MBAs and PhDs, who have contributed time, code, testing, and enthusiasm for this project from its beginning.

## References

- Betancourt, M. (2020). *Towards a principled Bayesian workflow*. [https://betanalpha.github.io/assets/case\\_studies/principled\\_bayesian\\_workflow.html](https://betanalpha.github.io/assets/case_studies/principled_bayesian_workflow.html).
- Bürkner, P.-C. (2017). Brms: An r package for Bayesian multilevel models using stan. *Journal of Statistical Software*, 80(1), 1–28. <https://doi.org/10.18637/jss.v080.i01>
- Congdon, P. D. (2010). *Applied Bayesian hierarchical methods*. CRC Press. <https://doi.org/10.1201/9781584887218>
- Dillon, J. V., Langmore, I., Tran, D., Brevdo, E., Vasudevan, S., Moore, D., Patton, B., Alemi, A., Hoffman, M., & Saurous, R. A. (2017). Tensorflow distributions. *arXiv Preprint arXiv:1711.10604*. <https://arxiv.org/abs/1711.10604>
- Fleischhacker, A. (2020). *A business analyst's introduction to business analytics: Intro to Bayesian business analytics in the r ecosystem*. Independently Published. ISBN: 9798667128175
- Gabry, J. (2018). *Shinystan: Interactive visual and numerical diagnostics and posterior analysis for Bayesian models*. <https://mc-stan.org/users/interfaces/shinystan>
- Gabry, J., Simpson, D., Vehtari, A., Betancourt, M., & Gelman, A. (2019). Visualization in Bayesian workflow. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 182(2), 389–402. <https://doi.org/10.1111/rssa.12378>
- Gelfand, A. E., & Smith, A. F. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410), 398–409. <https://doi.org/10.2307/2289776>
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*. CRC press. <http://www.stat.columbia.edu/~gelman/book/>
- Gilks, W. R., & Roberts, G. O. (1996). Strategies for improving MCMC. *Markov Chain Monte Carlo in Practice*, 6, 89–114.
- Golding, N. (2019). Greta: Simple and scalable statistical modelling in R. *Journal of Open Source Software*, 4(40), 1601. <https://doi.org/10.21105/joss.01601>
- Goodrich, B., Gabry, J., Ali, I., & Brilleman, S. (2020). *Rstanarm: Bayesian applied regression modeling via Stan*. <https://mc-stan.org/rstanarm>

- Iannone, R. (2020). *DiagrammeR: Graph/network visualization*. <https://github.com/rich-iannone/DiagrammeR>
- Kay, M. (2020). *ggdist: Visualizations of distributions and uncertainty*. <https://doi.org/10.5281/zenodo.3879620>
- Kruschke, J. (2014). *Doing Bayesian data analysis: A tutorial with r, JAGS, and stan*. <https://sites.google.com/site/doingbayesiandataanalysis/>
- Lunn, D. J., Thomas, A., Best, N., & Spiegelhalter, D. (2000). WinBUGS - a Bayesian modelling framework: Concepts, structure, and extensibility. *Statistics and Computing*, 10, 325–337. <https://doi.org/10.1023/A:1008929526011>
- McElreath, R. (2020a). *Rethinking: Statistical rethinking book package*. <https://github.com/rmcelreath/rethinking>
- McElreath, R. (2020b). *Statistical rethinking: A Bayesian course with examples in R and Stan*. CRC press. <https://doi.org/10.1201/9780429029608>
- Neal, R. M. (1993). *Probabilistic inference using Markov chain Monte Carlo methods*. Department of Computer Science, University of Toronto Toronto, Ontario, Canada.
- Pfeffer, A. (2016). *Practical probabilistic programming*. Manning Publ. ISBN: 9781617292330
- Stan Development Team. (2021). *Stan Modeling Language User's Guide and Reference Manual, Version 2.26*. <http://mc-stan.org/>
- Textor, J., Zander, B. van der, Gilthorpe, M. S., Liškiewicz, M., & Ellison, G. T. (2017). Robust causal inference using directed acyclic graphs: the R package “dagitty”. *International Journal of Epidemiology*, 45(6), 1887–1894. <https://doi.org/10.1093/ije/dyw341>
- Wickham, H. (2016). *ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York. ISBN: 978-3-319-24277-4