Optimizing Student Models for Causality

Benjamin Shih

shih@cmu.edu

Carnegie Mellon University

Problem

Models for student behavior, such as production rules, are often theory-driven. However, there is a wide gap between a theory and an effective student model. This gap includes overly complex models and incorrect parameter values, both of which can negatively impact experimental results.

The following diagram illustrates one path between a new hypothesis and an experimental trial. The Model Optimization step is added by the research presented here, and is an attempt to close the gap between the hypothesis and an effective student model.

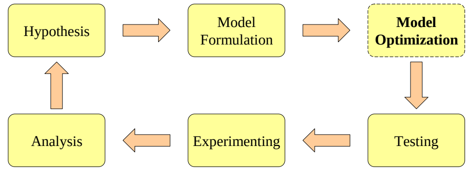

An example of how this approach can work in practice is research on the help-seeking tutor by Aleven et. al.

▪**Hypothesis**: Student help-seeking behavior is suboptimal, and intervening on this behavior can improve learning.

▪**Formulation**: Aleven et. al. created a production rule model with 15 rules and 10 thresholds (parameters) designed to capture help-seeking errors.

▪**Testing**: The model was tested on data from a prior (unrelated) experiment, resulting in a -0.60 correlation between the number of help-seeking errors and pre-post learning gain.

▪**Experimenting**: There have been several experiments with promising results, but no significant gains in pre-post learning gain.

The goal of this research is to improve experimental results through pre-experimental model optimization.

Method

The goal of an optimization method is, when given a cognitive model, to accomplish two things: tune the parameter values and reduce the model complexity.

▪**Tune parameters:** The original help-seeking model had many parameters, such as the time required to read a hint. Properly choosing values for these parameters is vital to the accuracy of the model, both in theoretical terms and in practical ones.

▪**Reduce model complexity:** Complex models with many rules are hard to generalize, hard to understand, and tend to have poor generalization to new data sets.

I use two data sources, one for optimizing (training) the model and the other for testing the model's generalization to other data sets.

▪**Training**: This data is from an experiment on self-explanation, and was used in the initial, non-experimental tests of the help-seeking model.

▪**Testing**: This data is from a pilot study for the help-seeking tutor. I only use the control group that did not receive the intervention.

I consider two general optimization approaches:

▪**Stepwise regression**: Backwards stepwise regression removes rules based on tests of statistical significance. It is a standard model reduction algorithm.

▪**Causal optimization**: The relationship between the help-seeking production rules, the pretest scores, and the posttest scores, can be represented in a graph. **A causal search algorithm, GES, can help to determine the causal relationships between the variables.** Production rules are then eliminated from the model if they are not causally linked to the post-test score.

Results

The following table shows the models generated by each optimization method. Each model is specified in terms of the steps used to generate it, e.g. "Tune + Stepwise + Tune" requires tuning the parameters, performing a stepwise regression, and then tuning the parameters for the resulting (reduced) model.

The correlations specified are between the number of classified help-seeking errors and pre-post learning gain.

Model	# of Rules	Training Correlation	Test Correlation
Original	15	-0.60	-0.07
Original + Tune	15	-0.62	0.18
Stepwise	5	-0.13	-0.22
Stepwise + Tune	5	-0.20	-0.23
Tune + Stepwise	5	-0.53	0.09
Tune + Stepwise + Tune	5	-0.54	0.03
GES	3	-0.47	-0.04
GES + Tune	**3**	**-0.48**	**-0.23**
Tune + GES	2	-0.48	-0.04
Tune + GES + Tune	**2**	**-0.49**	**-0.20**

The original model has a very high correlation with learning gain in the initial data, but it does not generalize well.

The stepwise approach *sometimes* gives smaller models and better generalization. Often, however, the generalization is almost as bad as the in the original model.

The causal search approach gives very small models, fantastic generalization, solid training set performance, accomplishes all of this reliably, and it gives a strong theoretical basis for belief that the relationships are causal.

Causal search for model optimization has the potential to improve students models, improve interventions, and decrease the likelihood of inconclusive experiments. This result, however, is itself inconclusive.