Learning and using relational theories: supporting material

Charles Kemp, Noah D. Goodman & Joshua B. Tenenbaum
Department of Brain and Cognitive Sciences, MIT, Cambridge, MA 02139
{ckemp,ndg,jbt}@mit.edu

1 Baseline models for Experiment 1

We previously considered a Bayesian model based on our complexity measure, but Equations 2 and 3 can also be used to define a Bayesian model based on Goodman's measure. The left columns show predictions when this model is applied to the inductive tasks in our first experiment. In order to combine Goodman's model with Equations 2 and 3, each T in these equations must now refer to a candidate extension rather than a string of clauses. To keep matters simple, we did not sum over the space of all possible extensions, but considered only extensions with the lowest possible complexity score that was compatible with the training data. Goodman's model and our model make similar predictions in four of the six conditions, but the Goodman model does not make accurate predictions about the transitive and exception predicates. Since Goodman's complexity measure is not sensitive to the structural properties of these predicates, it cannot explain how these properties support inductive projections.

The right columns of Figure 5 show predictions of an exemplar model which assumes that the new letter will be just like one of the old letters. If letter 9 is like letter 1, for instance, we can make predictions about 9 by collecting all pairs involving 1 and replacing 1 with 9 wherever it occurs. Suppose that we now learn that 96 belongs to the set. The model now assumes that letter 9 is like one of the letters that has previously been paired with 6, and generates predictions according to this assumption. Figure 5 shows that the model accounts well for four of the six conditions—in particular, it predicts some of the patterns in the data for the random condition. The model, however, does not make accurate predictions about the symmetry and transitive conditions. The novel pairs for both of these conditions indicate that the new letter is unlike any of the previous letters. In the transitive condition, for instance, no letter has ever previously appeared to the left of letter 1. Two strategies are now possible: either the model can refuse to make any prediction, or it can ignore the observed pair, which is the strategy we used when generating Figure 5. Neither strategy accounts well for the responses given by subjects in the symmetry and transitive conditions.

There is an important sense in which exemplar models can eventually be subsumed by approaches like ours. The assumptions that motivate an exemplar model can be written down as a relatively simple logical theory: for instance, a theory can include clauses like $\mathtt{R}(\mathtt{X},9) \leftarrow \mathtt{R}(\mathtt{X},2)$ and $\mathtt{R}(9,\mathtt{X}) \leftarrow \mathtt{R}(2,\mathtt{X})$ which indicate that letter 9 is expected to behave just like letter 2. It is therefore realistic to aim for a theory-learning model that captures all of the effects consistent with an exemplar model as well as effects (e.g. predictions based on symmetry) that seem more consistent with a rule-based approach. Our model, however, achieves this synthesis only partially.[1]

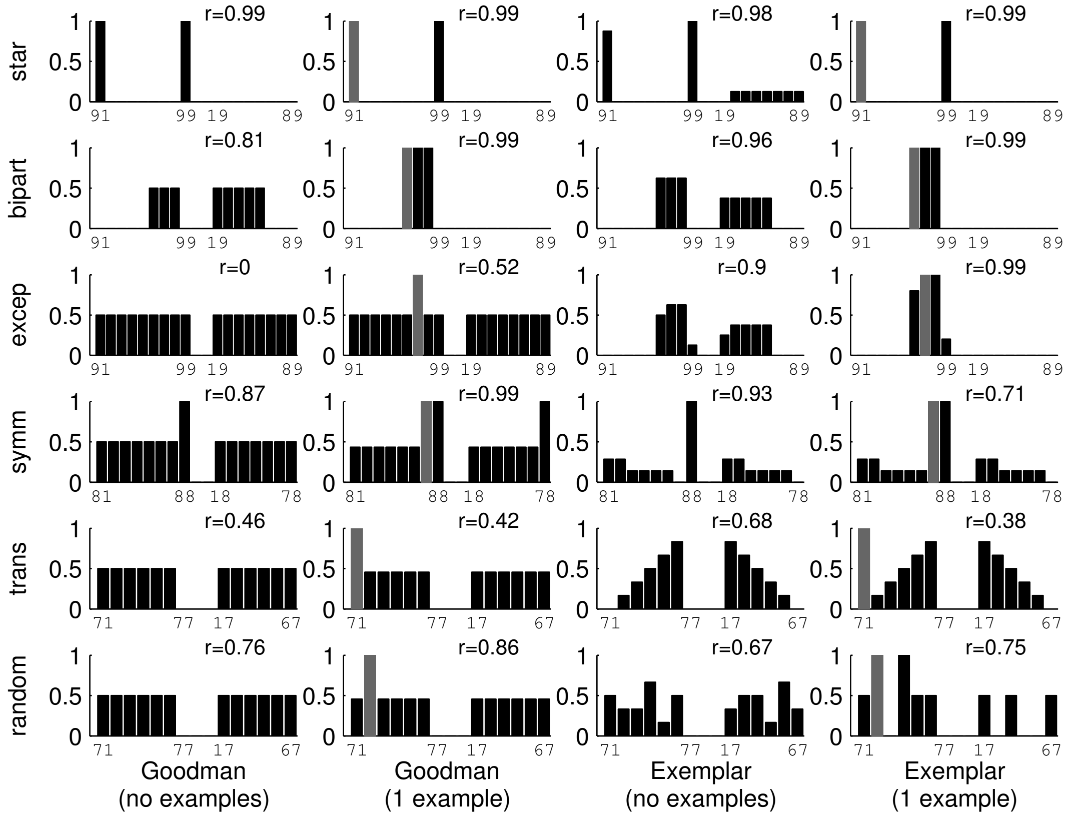

Figure 5: Model predictions for the induction task in Experiment 1. Columns 1 and 3 show predictions before any pairs involving the new item are observed. Columns 2 and 4 show predictions after a single new pair (marked with a gray bar) is observed to belong to the set. Plots for each condition include correlations with the human data.

2 Krueger data: memory and perception

Krueger [1] argues that Goodman's model can serve as an account of subjective complexity, and provides several experiments to support her claim. We suggest in this section that our model can also account for her results.

The first three columns of Figures 6 show results from three studies that helped to inspire our first experiment. In each study, subjects learned sets of triples built from letters of the alphabet. The triples in each set were presented one at a time, and subjects made multiple passes through each set. Krueger recorded the number of triples that subjects could recall at the conclusion of each pass, and used these data to construct a learning curve for each set.

Qualitative representations of the results are shown in the first row of Figure 6. In the first task, for instance, subjects were more successful at learning the set labeled 2 than the other three sets in the experiment (6a, 6b and 11). Complexity values according to our RL approach and to Goodman's model are shown in the second and third rows. Both models account well for the data, and the main difference between the two is that our model predicts a complexity difference between two sets (6a and 6b) that are considered equally complex by Goodman's model. The empirical data showed no significant difference between performance on these sets, but the learning curves for these sets

[1] Our model fails to capture some of the exemplar-based effects for two main reasons. First, clauses like $\mathtt{R}(\mathtt{X},9) \leftarrow \mathtt{R}(\mathtt{X},2)$ and $\mathtt{R}(\mathtt{X},9) \leftarrow \mathtt{R}(2,\mathtt{X})$ are equally simple according to our model, although the first appears to have a lower subjective complexity. Second, rules like $\mathtt{R}(\mathtt{X},9) \leftarrow \mathtt{R}(\mathtt{X},2)$ are considered substantially more complex than shorter rules like $\mathtt{R}(\mathtt{X},9)$.

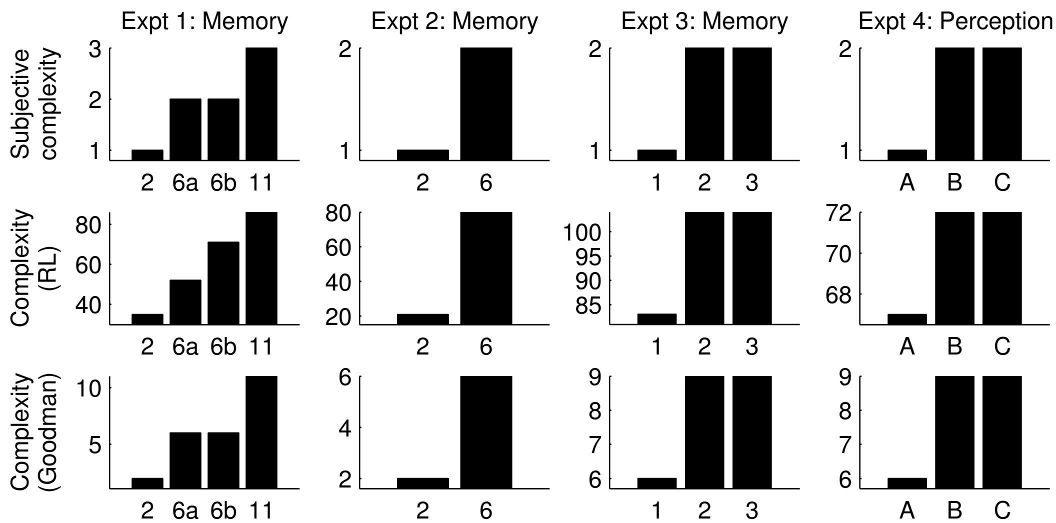

Figure 6: Data (top row) and model predictions for four experiments described by Krueger [1]. Each bar in each plot represents the complexity of a predicate, and the predicate labels (e.g. 2, 6a, A) are taken from Krueger [1]

reached an asymptote of near-perfect recall very quickly, which means that any difference between these sets would have been unlikely to emerge.

The final column of Figure 6 shows model predictions for a perceptual task where sets of triples were visually represented using graphs. Nodes in these graphs corresponded to elements in the domain, and each triple was shown by connecting three nodes to form a triangle. For each set, a single graph was created by superimposing the triangles for all triples in the set, and subjects were given pairs of graphs and asked to decide which graph was "simpler and more systematic looking." As for the previous tasks in Figure 6, both models account for the results of this experiment.

Krueger [1] ran one additional perceptual task that does not appear in Figure 6, but her thesis does not contain enough information for us to reconstruct the stimuli for this experiment. Based on the experiments we have been able to model, we conclude that our RL approach accounts for the majority of the findings that have previously been used to support the Goodman approach.

References

[1] J. T. Krueger. *A theory of structural simplicity and its relevance to aspects of memory, perception, and conceptual naturalness*. PhD thesis, University of Pennsylvania, 1979.