MODELING INDIVIDUAL CONCEPTS AS GRAPH THEORETICAL NETWORKS

Sarah Solomon, John Medaglia, & Sharon Thompson-Schill

Department of Psychology, University of Pennsylvania

MOTIVATION

① Concepts (e.g. CHOCOLATE) can be instantiated in many different forms (e.g., bar, truffle), and our conceptual system must be flexible enough to capture this variation.

② Whereas traditional models[1] define concepts as static structures, we aim to model concepts in a way that can accommodate the variation of conceptual information across instances.

③ We model concepts as graph-theoretical networks, with properties represented as nodes and their associations as edges.

④ Instead of relying on how properties correlate with *concepts*, does the correlation of properties with *each other* play a role in the structure of basic-level concepts?

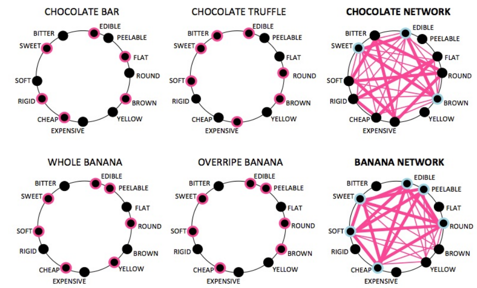

APPROACH

CHOCOLATE, BANANA, BOTTLE, TABLE, PAPER

① Construct a **set of properties** and **define various subkinds** of each of our concepts.

② Measure **property strengths for each subkind** for each concept.

③ Create our novel **network models** and **traditional models** in which each concept is represented by a property vector.

④ Collect testing data based on images of concept exemplars, and **classify these data** using both our network models (*graph alignment*) and traditional models (*correlation classifier*).

NETWORK MODELS

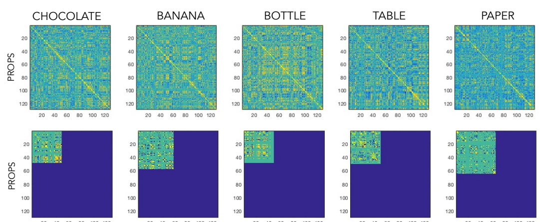

Participants (*N*=66) listed properties for each concept, and generated subkinds for each concept. We compiled these to result in a 129-property vector. Separate participants (*N*=198) selected which properties apply to each subkind of each concept.

① **Property vectors for each subkind were correlated with each other to create each concept's network model.** Properties are represented as nodes and their associations represented as edges. *Warmer colors represent high covariation across subkinds of a concept. (top)*

② For each concept, the model was restricted to only include properties that were present in at least one subkind. *(bottom)*

TRADITIONAL MODELS

For each concept, data were collapsed across subkinds to represent the presence or strength of that property for that concept.

① **Binary:** Each concept is represented as a vector designating whether each of the 129 properties applies to that concept.

② **Weighted:** Each concept is represented as a vector representing the strength (*probability across subkinds*) of each of the 129 properties.

③ **Weighted + Restricted:** Each concept is represented as a vector representing the strength of properties, restricted to only include properties above threshold.

GRAPH ALIGNMENT

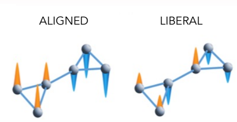

The graph alignment technique can be used to assess the degree to which a signal (*vector*) aligns with a network (*graph*). A signal is highly aligned if the magnitude of the nodes corresponds tightly to that expected by the network's organization.[2] In our case, the concept networks define what kind of signals we expect from individual exemplars.

TESTING DATA

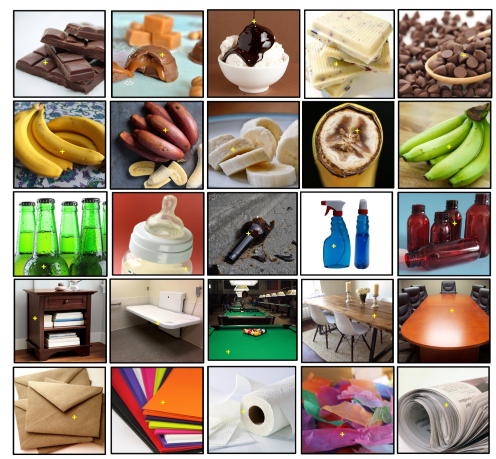

A separate group of subjects (*N*=60) saw one exemplar per concept, and selected which of the 129 properties apply to that specific object. Testing data included 300 exemplar input-vectors, 60 per concept.

CLASSIFICATION RESULTS

threshold=0

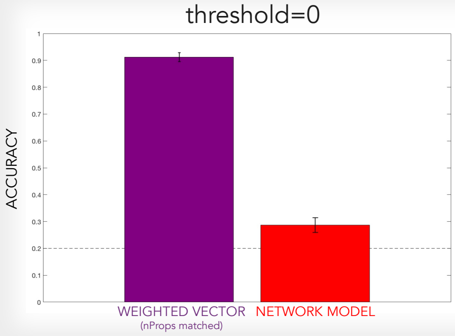

Even when property strength cannot play any role in classification, the network model is able to classify which concept was seen at above-chance levels. This suggests **that property correlations within a concept might play a role in conceptual structure.**

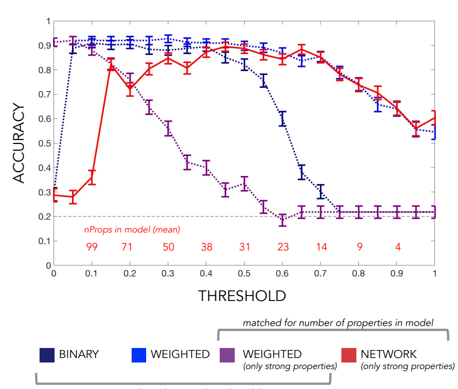

threshold=0.7

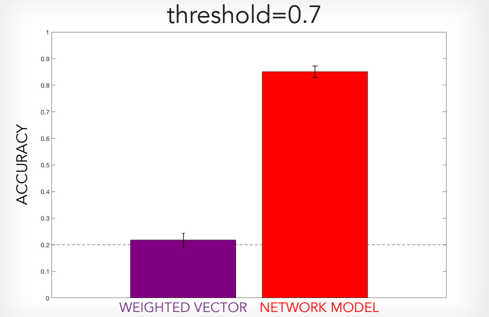

When **property correlations is combined with property strength** (by reducing networks to high-strength properties), the network model is highly successful at classifying concepts, whereas a traditional vector model matched for the same number of properties fails. Thus, a concept may be represented by **a small set of strong properties and the ways they correlate with each other** across instances.

CONCLUSION

① We created novel network-models for concepts, and show that these models are successful at classifying individual exemplars.

② This suggests that property covariations may contribute useful information to the structure of basic-level concepts.

③ Even when no property strength is captured in the model (threshold=0), the network models still capture conceptual information, using property covariance alone.

④ When property strength is incorporated into the model (threshold=0.7) and the number of properties in the model is reduced, the network model does well, whereas traditional vector-based models fail. This suggests that specific relations among the strongest properties tell us a lot about a concept.

⑤ Future extensions include extracting network-science measures of flexibility and modularity and linking these to conceptual phenomena.

ACKNOWLEDGEMENTS

1. *Tyler, Moss, Durrant-Peatfield, & Levy (2000). Brain and Language.*

2. *Medaglia, Huang, Karuza, Thompson-Schill, Ribeiro, & Bassett (2016). arXiv preprint: 1611.08751*

This research was supported by Grant R01 DC015359 awarded to STS and an NSF Graduate Research Fellowship awarded to SHS.