CS 101. Special Topics in Computer Science: Networks of Relations. 9 units (3-0-6): first term. No formal prerequisites, but enjoyment of mathematical models of computation will be a plus. Circuits and Turing machines have provided a great foundation for understanding machines that evaluate functions, but for understanding brains they have not been particularly helpful. Relational networks are a computational model that is better suited to providing intuition for machines such as brains that are based on grounded representation and inference. The first part of this course will address the theory of relational networks.

Relations (such as the OR clauses in a satisfiability problem) are defined by the set of tuples that they accept. Networks of relations can automate inference, analogously to how networks of gates (i.e., circuits) automate function evaluation. Unlike the input/output computational paradigm, the inference process in a relational network proceeds in all directions, as known quantities are used to infer other quantities. In circuits we are used to the fact that AND/OR/NOT can implement any function, while AND/OR can only implement monotonic functions, and other sets of generators can yield other classes of functions. Similarly, for relations OR/NOT can implement any relation (as in 3-SAT for example), while other sets of generators yield other classes of relations. The complete possibilities for Boolean function classes were enumerated by Post 100 years ago, and then 50 years ago the complete possibilities for Boolean relation classes were found to have the same lattice structure as for functions, but upside down. The second part of this course will look at how networks of relations provide a foundation for thinking about how the brain processes information. Population codes encode values of variables, while synapses encode knowledge in a relational form. Standard Hebbian learning allows relations to be learned from exposure to data, and latent state variables can be inferred similarly to the digital case.

Relations can be implemented with classifiers (e.g. k-means or winner-take-all classifiers), and networks of such classifiers can both learn the statistics of their world and make inferences based on those statistics. Of course, brains do more than just infer; they need to make choices and act. The course will conclude by discussing distributed hierarchical action selection, its connections to relational networks, and its correspondence to cortical neuroanatomy, allowing simultaneous reinforcement learning in multiple areas while avoiding the credit assignment problem. Graduate students as well as interested undergraduates are welcome. Instructor: Matthew Cook.