Learning-Theoretic Epistemology

Epistemic justification consists in adopting the most reliable means for attaining the ends of inquiry. Since our ends are determined by the kind of problem we are trying to solve, they vary with the epistemic context.

A learning-theoretic analysis of a learning problem consists in: (1) specifying the **problem context**: what *question* are we trying to answer? what relevant *information* can we hope to acquire as inquiry continues? (2) analysing the **inherent difficulty** of the problem: given the kind of information we can hope to acquire, what is the *best sense* in which our problem can be solved? and (3) identifying the **methods** that solve the problem in the best sense that its inherent difficulty allows.

An epistemically justified method must converge to the best answer to the question as information accumulates. But that is only a minimal requirement. More stringently, the method must converge *as efficiently as possible* given the inherent difficulty of the problem. Moreover, scientific preferences for theoretical virtues like simplicity, unity, or testability are justified by proving that they are *necessary* for efficient convergence. For example, Ockham's razor is justified by demonstrating that it is a necessary condition of *progressive* convergence to the best answer.

Topology as Epistemic Geometry

Topology is usually thought of as a kind of abstract geometry. That makes it sound remote from the concerns of epistemology. I claim that this is mistaken: topology is an even more natural fit for epistemology than logic or probability theory. It allows us to solve for the inherent difficult of a learning problem by characterizing exactly how well information resolves answers to the question.

A point w of an abstract space is a *limit point* of a region A if there are points lying in A that get closer and closer to w. These points may never arrive at w, but they are said to approximate it *arbitrarily well.* Two regions of space are *well-separated* if no point of the one is a limit point of the other.

Interpret the points of the space as the live epistemic possibilities in our context. A possibility w is *a limit point* of an answer A if all information that would ever be available in w is compatible with some possibility in which A is true. Then, the separation relation precisely captures how well, or poorly, information resolves answers. The less information separates answers, the greater the inherent difficult of our question. Furthermore, theoretical simplicity tracks relations of underdetermination. If A and B are competing answers, say that A is *simpler than* B if all information compatible with A is also compatible with B. That definition validates our intuitive judgements in a wide range of problems.

The Topology of Statistical Inquiry

A common criticism of the topological framework is that it appeals to an idealization: inquirers are afforded information that logically refutes incompatible possibilities. Sober (2015) and others doubt whether our arguments can be made to apply in probabilistic settings where no possibility is ever logically refuted. I resolve that worry in my dissertation by generalizing the learning-theoretic framework to apply literally to statistical inference as practiced in the data-driven sciences. This was done by solving for the unique topology on probability measure in which topological separation tracks *statistical* underdetermination. That advance allows for epistemological insight, developed by philosophers dealing with traditional epistemic concerns, to flow directly into scientific practice. That yields concrete methodological recommendations e.g., it suggests modifications to hypothesis testing methods so that true findings are more likely to be replicated; it also hones and justifies statistical procedures for inferring causal structure from observational data. Topological learning theory promises a unified framework in which to study the entire variety of problems arising in Machine Learning, Data Science and Statistics.