The Topology of Statistical Verifiability

Konstantin Genin and Kevin T. Kelly
Department of Philosophy, Carnegie Mellon University

Abstract

In many epistemic applications of topology, open sets are interpreted as hypotheses deductively verifiable by true *propositional* information that rules out relevant possibilities. However, in statistical data analysis, one routinely receives random samples logically compatible with every statistical hypothesis. We bridge the gap between propositional and statistical data by solving for the *unique* topology on probability measures in which the open sets are exactly the *statistically verifiable* hypotheses. Furthermore, we extend that result to a topological characterization of learnability in the limit from statistical data. These results extend topological models of inquiry that have emerged in diverse fields such as domain theory [Abramsky and Jung, 1994, Vickers, 1996], formal learning theory [Yamamoto and de Brecht, 2010], epistemology and philosophy of science [Kelly, 1996, Schulte and Juhl, 1996, Genin and Kelly, 2015, forthcoming, Baltag et al., 2015], statistics [Dembo and Peres, 1994, Ermakov, 2013] and modal logic [Wáng and Ågotnes, 2013, Bjorndahl, 2013].

Main Result

$$\begin{tikzcd}
\Pi^0_2 \ \text{limiting refutable} \arrow[rr, leftrightarrow, "\neg"] & & \Sigma^0_2 \ \text{limiting verifiable} \\
& \Delta^0_2 \ \text{limiting decidable} \arrow[ul, phantom, "\subseteq"] \arrow[ur, phantom, "\subseteq"] & \\
\Pi^0_1 \ \text{deductively refutable} \arrow[uu, "\cap_{n\in\mathbb{N}}\Pi_1^{\mathsf{c}} \cup \Pi_1"] \arrow[rr, leftrightarrow, "\neg"] \arrow[ur, phantom, "\subseteq"] & & \Sigma^0_1 \ \text{deductively verifiable} \arrow[uu, "\cup_{n\in\mathbb{N}}\Sigma_1 \cap \Sigma_1^{\mathsf{c}}"'] \arrow[ul, phantom, "\subseteq"] \\
& \Delta^0_1 \ \text{deductively decidable} \arrow[ul, phantom, "\subseteq"] \arrow[ur, phantom, "\subseteq"] &
\end{tikzcd}$$

In both the statistical and propositional settings, the Borel hierarchy captures the structure of methodological possibility. In the propositional setting, Σ^0_1 hypotheses are the open sets generated by the information basis. In the statistical setting, Σ^0_1 hypotheses are the open sets of the weak topology. The Σ^0_2 hypotheses are all the hypotheses H that can be expressed in the form

$$H = \cup_{i\in\mathbb{N}} U_i \cap V_i^{\mathsf{c}},$$

where $U_i, V_i \in \Sigma^0_1$. The Π^0_n hypotheses are complements of Σ^0_n hypotheses. Finally, $\Delta^0_n = \Sigma^0_n \cap \Pi^0_n$. The diagram shows the inclusions between levels of the Borel hierarchy.

Propositional Setting

W is a set of possible worlds, and $\mathcal{I}_w$ is the set of all information states true in w. Information states logically refute incompatible possibilities. We assume that:

1. For every w, $\mathcal{I}_w$ is non-empty;
2. The set of all information states $\mathcal{I} = \cup_w \mathcal{I}_w$ is countable;
3. for every E, F in $\mathcal{I}_w$, there is G in $\mathcal{I}_w$ such that $G \subseteq E \cap F$.

The underlying idea is that a sufficiently diligent inquirer in w eventually receives information as strong as an arbitrary information state E in $\mathcal{I}_w$. Since that is true of both E and F, there must be true information as strong as $E \cap F$.

It follows from assumptions (1) and (3) that $\mathcal{I}$ is a topological basis. Therefore, the closure of $\mathcal{I}$ under union, denoted $\mathcal{T}$, is a topological space we call the *information topology*.

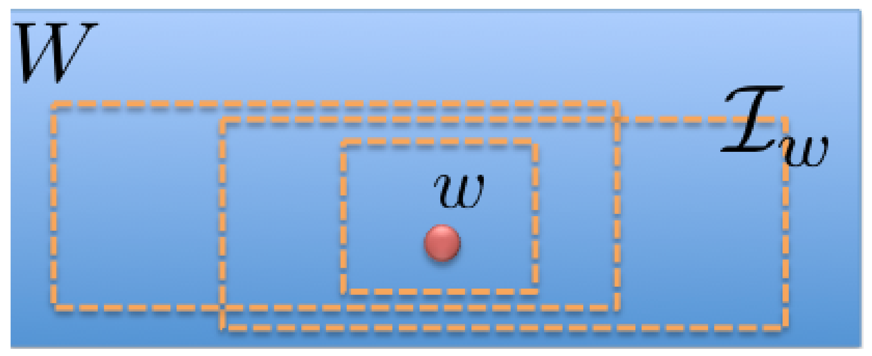

Logical Verifiability

A *method* is a function from information states to propositions. Method $L(\cdot)$ is *infallible* iff its output is always true, i.e. iff $w \in L(E)$ for all $E \in \mathcal{I}_w$. A *verifier* for H is an infallible method that converges to belief in H iff H is true. That is, $L(\cdot)$ is a verifier for H iff

1. $L(\cdot)$ is infallible and
2. $w \in H$ iff there is $E \in \mathcal{I}_w$ such that $L(F) \subseteq H$ for all $F \in \mathcal{I}_w$ entailing E.

H is *verifiable* iff there exists a verifier for H. H is *refutable* iff its complement is verifiable. H is *decidable* iff H is both verifiable and refutable.

Theorem. Hypothesis H is verifiable iff H is open in the information topology.

Logical Verifiability in the Limit

A *limiting verifier* for H is a method that converges to true belief in H iff H is true. That is, $L(\cdot)$ is a limiting verifier for $H \subseteq W$ iff

1. $w \in H$ iff there is $E \in \mathcal{I}_w$ such that $L(F) \subseteq L(E) \subseteq H$ for all $F \in \mathcal{I}_w$ entailing E.

H is *limiting verifiable* iff there exists a limiting verifier of H. H is *limiting refutable* iff there exists a limiting verifier of H^{c}. H is *limiting decidable* iff it is limiting verifiable and refutable.

Theorem. Hypothesis H is limiting verifiable iff H is Σ^0_2 in the information topology.

Statistical Setting

W is a set of Borel measures on a sample space S, equipped with a topology. The topology on S reflects what is verifiable about the sample itself. As in the propositional setting, it is *verifiable* that sample ω lands in A iff A is open, and it is *decidable* whether sample ω falls into region A iff A is clopen. Suppose that A is the closed interval $[1/2, \infty]$, and the sample ω lands right on the end-point $1/2$ of A. Given enough time and computational power, the sample ω can be specified to arbitrary, finite precision. But no finite degree of precision: $\omega \approx .50$; $\omega \approx .500$; $\omega \approx .5000$; $\ldots$ suffices to determine that ω is truly in A. But the mere possibility of a sample hitting the boundary of A does not matter statistically, if the chance of such a sample is zero. A Borel set $A \subseteq S$ for which $p_w(\mathsf{bdry}(A)) = 0$ for every world in W is said to be *almost surely decidable*. We assume that the topology on S has a basis of almost surely decidable regions.

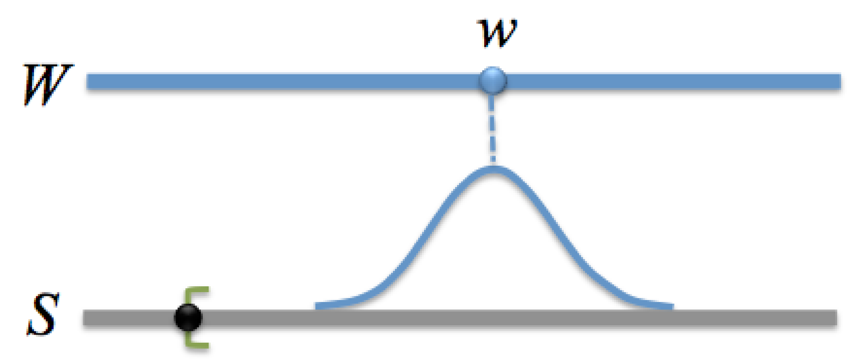

Statistical Verifiability

A family $\{\lambda_n\}_{n\in\mathbb{N}}$ of feasible tests of $H^{\mathsf{c}} \subseteq W$ is an *almost sure α-verifier* of H iff

1. $\sum_{n=1}^{\infty} p_w^n[\lambda_n^{-1}(H)] \leq \alpha$ for all $w \in H^{\mathsf{c}}$ and
2. $p_w^{\infty}[\liminf_{n\to\infty} \lambda_n^{-1}(H)] = 1$ for all $w \in H$.

Hypothesis $H \subseteq W$ is *almost surely α-verifiable* iff there is an almost sure α-verifer of H. H is *almost surely verifiable* iff H is almost surely α-verifiable, for every $\alpha > 0$. H is *almost surely refutable* iff H^{c} is almost surely verifiable. H is *almost surely decidable* iff H is almost surely verifiable and refutable.

Theorem. Hypothesis H is almost surely verifiable iff H is open in the weak topology.

Statistical Verifiability in the Limit

A family $\{\lambda_n\}_{n\in\mathbb{N}}$ of feasible methods is a *limiting almost sure verifier* of $H \subseteq W$ iff

1. $w \in H$ iff there is $H' \subseteq H$, s.t. $p_w^{\infty}[\liminf_{n\to\infty} \lambda_n^{-1}(H')] = 1$;
2. $w \notin H$ iff for all $H' \subseteq H$, $p_w^{\infty}[\limsup_{n\to\infty} \lambda_n^{-1}(H')] = 0$.

H is *limiting a.s. verifiable* iff there is a limiting a.s. verifier of H. H is *limiting a.s. refutable* iff H^{c} is limiting a.s. refutable. H is *limiting a.s. decidable* iff H^{c} is limiting a.s. verifiable and refutable.

Theorem. Hypothesis H is limiting almost surely verifiable iff H is Σ^0_2 in the weak topology.

Feasible Methods

A statistical *method* is a measurable function from samples to propositions over W. A *test* of a statistical hypothesis $H \subseteq W$ is a statistical method $\psi : \Omega \to \{W, H^{\mathsf{c}}\}$. Call $\psi^{-1}(W)$ the *acceptance region*, and $\psi^{-1}(H^{\mathsf{c}})$ the *rejection region* of the test. The *power* of test $\psi(\cdot)$ is the worst-case probability that it rejects truly, i.e. $\inf_{w\in H^{\mathsf{c}}} p_w[\psi^{-1}(H^{\mathsf{c}})]$. The *significance level* of a test is the worst-case probability that it rejects falsely, i.e. $\sup_{w\in H} p_w[\psi^{-1}(H^{\mathsf{c}})]$. A method is *feasible in μ* iff the preimage of every element of its range is almost surely decidable in μ. A method is *feasible* iff it is feasible in every world in W. Methods that are not feasible are impossible to implement.

Weak Topology

A sequence of measures $(w_n)_n$ *converges weakly* to w, written $w_n \Rightarrow w$, iff $p_{w_n}(A) \to p_w(A)$ for every A almost surely clopen in w. It is immediate that $w_n \Rightarrow w$ iff for every w-feasible test $\psi(\cdot)$, $p_{w_n}(\psi \text{ rejects}) \to p_w(\psi \text{ rejects})$. It follows that no feasible test of $H = \{w\}$ achieves power strictly greater than its significance level. Furthermore, every feasible method that correctly infers H with high chance in w, exposes itself to a high chance of error in "nearby" w_n.

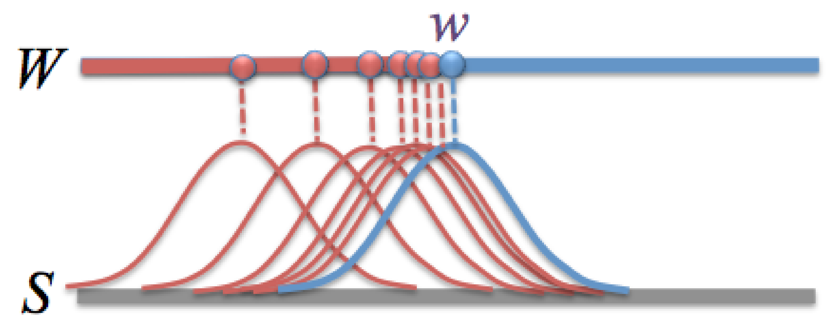

References

Samson Abramsky and Achim Jung. Domain theory. *Handbook of logic in computer science*, 3:1–168, 1994. doi: 10.1.1.50.8851.

Alexandru Baltag, Nina Gierasimczuk, and Sonja Smets. On the solvability of inductive problems: a study in epistemic topology. In *Proceedings of the Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge*, 2015. doi: 10.4204/EPTCS.215.7.

Adam Bjorndahl. Subset space public announcement logic revisited. *arXiv preprint arXiv:1302.4009*, 2013.

Amir Dembo and Yuval Peres. A topological criterion for hypothesis testing. *The Annals of Statistics*, pages 106–117, 1994. doi: 10.1214/aos/1176325360.

Mikhail Ermakov. On distinguishability of hypotheses. *arXiv preprint arXiv:1308.4295*, 2013.

Konstantin Genin and Kevin T. Kelly. Theory choice, theory change, and inductive truth-conduciveness. In *Proceedings of the fifteenth conference on Theoretical Aspects of Rationality and Knowledge*, 2015.

Konstantin Genin and Kevin T. Kelly. Learning, theory choice, and belief revision. *Studia Logica*, forthcoming.

Kevin T. Kelly. *The Logic of Reliable Inquiry*. Oxford University Press, 1996.

Oliver Schulte and Cory Juhl. Topology as epistemology. *The Monist*, 79(1): 141–147, 1996. doi: 10.5840/monist19967916.

Steven Vickers. *Topology Via Logic*. Cambridge University Press, 1996.

Yì N Wáng and Thomas Ågotnes. Subset space public announcement logic. In *Logic and Its Applications*, pages 245–257. 2013. doi: 10.1007/978-3-642-36039-8_22.

Akihiro Yamamoto and Matthew de Brecht. Topological properties of concept spaces (full version). *Information and Computation*, 208(4):327–340, 2010. doi: 10.1016/j.ic.2009.08.001.