Complexity-Theoretic Implications of Multicalibration*

(Abstract for FORC 2024)

Sílvia Casacuberta† Cynthia Dwork‡ Salil Vadhan§

Harvard University

We present connections between the recent literature on multigroup fairness for prediction algorithms and classical results in computational complexity. Hébert-Johnson, Kim, Reingold, and Rothblum [Héb+18] introduced a new measure of algorithmic fairness for predictors called *multicalibration (MC)*, which guarantees calibrated predictions across every subpopulation from a prespecified family $\mathcal{F}$ of potentially intersecting subsets of $\mathcal{X}$. *Multiaccuracy (MA)* is a relaxation of multicalibration in which the predictor is merely required to be accurate in expectation on each $f \in \mathcal{F}$. Multiaccuracy is equivalent to a regularity notion for functions defined by Trevisan, Tulsiani, and Vadhan [TTV09]. They showed that, given a class $\mathcal{F}$ of (possibly simple) functions, an arbitrarily complex function g can be approximated by a low-complexity function h that makes a small number of oracle calls to members of $\mathcal{F}$, where the notion of approximation requires that h cannot be distinguished from g by members of $\mathcal{F}$. This complexity-theoretic Regularity Lemma is known to have implications in different areas, including in complexity theory, additive number theory, information theory, graph theory, and cryptography. We focus on the following three implications: Impagliazzo's Hardcore Lemma (IHCL) [Imp95], characterizations of pseudo-average min-entropy (PAME) [VZ12; Zhe14], and the Dense Model Theorem (DMT) [GT08; TZ08; Rei+08]. Starting from the stronger notion of multicalibration, we obtain stronger and more general versions of all of IHCL, PAME, and the DMT. Here we focus on describing our generalization of IHCL (which we call IHCL++).

IHCL is a fundamental result in complexity theory stating that if a function g is $(\mathcal{F}, \delta)$-hard, meaning that $\Pr[f(x) \neq g(x)] \geq \delta$ for all $f \in \mathcal{F}$, then there is a large-enough subset H of the inputs $\mathcal{X}$ (the "hardcore set") for which the function g is $(\mathcal{F}', 1/2 - \epsilon)$-hard on H; i.e., g is indistinguishable from a random function by a family $\mathcal{F}'$ of distinguishers of complexity similar to that of $\mathcal{F}$. In [TTV09], it was shown that the Regularity Lemma implies IHCL with a hardcore size of at least $\delta|\mathcal{X}|$. Using the MC Theorem from [Héb+18] instead, we show:

Theorem (IHCL++): For every boolean function g, family of boolean functions $\mathcal{F}$, and $\epsilon > 0$, there exists a "low-complexity" partition $\mathcal{P}$ of the domain of g into $O(1/\epsilon)$ pieces such that on every large enough piece $P \in \mathcal{P}$, there exists a hardcore set $H_P \subseteq P$ of density at least $2b_P|P|$ such that g is $(\mathcal{F}', 1/2 - \epsilon/b_P)$-hard on H_P, where $b_P = \min\{\mathbb{E}_{x\sim P}[g(x)], 1 - \mathbb{E}_{x\sim P}[g(x)]\}$.

That is, instead of finding a single, globally dense hardcore set H, we find many "local" hardcore sets H_P, each of which is dense within its piece P of the partition. Importantly, IHCL++ holds for every Boolean function g, regardless of its hardness. Moreover, the density of each hardcore set H_P is such that when we bring back the original assumption stating that g is δ-hard, we recover IHCL by "gluing" together all of the hardcore sets. This global hardcore set achieves density at least 2δ, which corresponds to the optimal density parameters in IHCL [Hol05].

---

*Full version of the paper: https://arxiv.org/abs/2312.17223. To appear in STOC 2024.

†scasacubertapuig@college.harvard.edu. Supported by the Herchel Smith Undergraduate Science Research Program.

‡dwork@seas.harvard.edu. Supported by grant G-2020-13941 of the Alfred P. Sloan Foundation and the Simons Foundation collaboration project 733782.

§salil_vadhan@harvard.edu. Supported by a Simons Investigator Award.