Hardcore Measures, Dense Models and Low Complexity Approximations

Sita Gakkhar Russell Impagliazzo Valentine Kabanets

May 25, 2012

1 Introduction

Main results:

1) black-box reductions; preserving the parameters, constructiveness, and the complexity assumptions.

$$\text{HCL}_{\text{STRONG}} \xrightarrow{\text{IMPAGLIAZZO}} \text{DMT}^{\substack{\text{PSEUDODENSITY}\\ \text{PSEUDORANDOM}}} \longrightarrow \text{LCAT}$$

2) Common generalization: the constructive LP duality (constructive MinMax) DMT: For every class $\mathcal{F}$, a distribution ρ in $\mathcal{U} = (U, \sigma)$, and parameters $\epsilon, \delta \in (0, 1)$, either have a δ-dense $(\epsilon, \mathcal{F})$-model for ρ in $\mathcal{U}$ which moreover has low complexity relative to $\mathcal{F}$, or have a small-complexity function (the average of f's in $\mathcal{F}$) that is a universal distinguisher between ρ and any given δ-dense measure.

$$\begin{array}{ccccc} \text{HCL}_{\text{STRONG}} & \longleftarrow & \text{MINMAX DMT} & \longrightarrow & \text{DMT}^{\substack{\text{PSEUDODENSITY}\\ \text{PSEUDORANDOM}}} \\ & & \downarrow & & \\ & & \text{LCAT} & & \end{array}$$

2 Preliminaries

2.1 Distributions, measures, and tests

A *probability space* is the pair $\mathcal{U} = (U, \sigma)$, where U is a finite set and $\sigma : U \to \{0, 1\}$ is a probability distribution over U (so $\sum_{x \in U} \sigma(x) = 1$). Without loss of generality, we assume that $\text{SUPPORT}(\sigma) = U$; otherwise, we define U to be $\text{SUPPORT}(\sigma)$. We write $x \in \sigma$ to indicate that $x \in U$ is randomly sampled according to the probability distribution σ.

A function $\mu : U \to [0, 1]$ is called a *measure* over U. Informally, a measure μ can be thought of as a "fuzzy" set where, for each $x \in U$, $\mu(x)$ measures the likelihood of x being in the set.[1] Subsets $A \subseteq U$ correspond to the Boolean measures $A : U \to \{0, 1\}$ where $A(x) = 1$ iff $x \in A$.

We will also think of functions $f : U \to [0, 1]$ as "tests". For a function $f : U \to [0, 1]$, we define its *complement* $\bar{f} = 1 - f$. A set $\mathcal{F}$ of functions $f : U \to [0, 1]$ is called a *class* if $\mathcal{F}$ is closed under complementation, i.e., if $f \in \mathcal{F} \Leftrightarrow \bar{f} \in \mathcal{F}$ for every $f : U \to [0, 1]$. A *Boolean class* is a class of Boolean functions $f : U \to \{0, 1\}$.

Given any two probability distributions ρ and σ over U, we can express σ as a convex combination $\sigma(x) = \delta \cdot \rho(x) + (1 - \delta) \cdot \tau(x)$ of ρ and some distribution τ, for $\delta \in [0, 1]$. In this case, to sample from σ, we need to sample from ρ with probability at least δ, implying that ρ is δ-dense in σ. We have the following definition.

[1] More precisely, μ specifies a probability distribution over subsets of U, where a random subset R is picked by placing each $x \in U$ into R with probability $\mu(x)$, independently.

Definition 2.1 (density of distribution)**.** For probability distributions ρ and σ over U, we say that ρ is δ-*dense in* σ if, for every $x \in U$, $\sigma(x) \geq \delta \cdot \rho(x)$. The maximum value δ, $0 \leq \delta \leq 1$, such that ρ is δ-dense in σ is called the *density of* ρ *in* σ, and is denoted by $d_\sigma(\rho)$.

For a subset $A \subseteq U$, its density in (U, σ), denoted $d_\sigma(A)$, is the probability that a random sample from σ lands inside A, i.e., $d_\sigma(A) = \mathbb{E}_{x \in \sigma}[A(x)]$. More generally, we have the following.

Definition 2.2 (density of measure)**.** For $\mu : U \to [0, 1]$, the *density of the measure* μ *in the probability space* $\mathcal{U} = (U, \sigma)$, denoted by $d_\sigma(\mu)$, is $d_\sigma(\mu) = \mathbb{E}_\sigma[\mu]$.

Next we define probability distributions induced by measures.

Definition 2.3 (induced distribution)**.** For a nonzero measure μ over a probability space (U, σ), the *induced probability distribution*, denoted by μ_σ, is $\mu_\sigma(x) = \mu(x)\sigma(x)/d_\sigma(\mu)$ for each $x \in U$.

Note that, for a set $\emptyset \neq A \subseteq U$, $A_\sigma = \sigma|_A$, where $\sigma|_A$ is σ conditioned on sampling from A.

We relate the densities of a measure μ and the corresponding induced distribution μ_σ.

Claim 2.4. *For a probability space (U, σ) and a nonzero measure μ over U, $d_\sigma(\mu_\sigma) \geq d_\sigma(\mu)$, with equality when μ is a Boolean measure (i.e., a subset of U).*[2]

Proof. For every $x \in U$ such that $\mu(x) \neq 0$, we have $\sigma(x) = \mu_\sigma(x) \cdot (d_\sigma(\mu)/\mu(x)) \geq d_\sigma(\mu) \cdot \mu_\sigma(x)$, as $\mu(x) \leq 1$. For any $x \in U$ with $\mu(x) = 0$, we also get that $\sigma(x) \geq 0 = d_\sigma(\mu) \cdot \mu_\sigma(x)$.

For a nonzero Boolean measure μ, take any $x_0 \in U$ such that $\mu(x_0) = 1$. We have $\sigma(x_0) \geq d_\sigma(\mu_\sigma) \cdot \sigma(x_0)\mu(x_0)/d_\sigma(\mu)$, implying that $d_\sigma(\mu_\sigma) \leq d_\sigma(\mu)$. □

The following claim relates a measure μ and its scaled version $\mu' = c \cdot \mu$, for a constant $c > 0$.

Claim 2.5. *For any measures μ over (U, σ) and $\mu' = c \cdot \mu$, for a constant $c > 0$, we have (i) $d_\sigma(\mu') = c \cdot d_\sigma(\mu)$, and (ii) $\mu_\sigma = \mu'_\sigma$.*

Proof. For *(i)*, $d_\sigma(\mu') = \mathbb{E}_\sigma[c \cdot \mu] = c \cdot \mathbb{E}_\sigma[\mu] = c \cdot d_\sigma(\mu)$. For *(ii)*, for every $x \in U$, we have $\mu'_\sigma(x) = c\mu(x)\sigma(x)/d_\sigma(\mu') = \mu(x)\sigma(x)/d_\sigma(\mu) = \mu_\sigma(x)$. □

2.2 Indistinguishability

For any function $f : U \to [0, 1]$ and a probability distribution ρ over the universe U, we define $f[\rho] = \mathbb{E}_\rho[f] = \sum_{x \in U} f(x)\rho(x)$. For $f : U \to [0, 1]$ and any measure μ over a probability space (U, σ), we define

$$f_\sigma[\mu] = \mathbb{E}_{\mu_\sigma}[f] = \sum_{x \in U} f(x)\mu(x)\sigma(x)/d_\sigma(\mu) = \frac{1}{d_\sigma(\mu)} \cdot \mathbb{E}_\sigma[f \cdot \mu].$$

Note that $f[\rho] = d_\rho(f)$ is the density of f (when viewed as a measure) in the probability space (U, ρ); similarly, $f_\sigma[\mu]$ is the density of f in the probability space (U, μ_σ).

Recall that distributions ρ_1 and ρ_2 over U are ϵ-close if, for every subset $T \subseteq U$, we have $|T[\rho_1] - T[\rho_2]| \leq \epsilon$. By restricting our tests to a set $\mathcal{F}$ of $[0, 1]$-valued functions, we get the notion of closeness, or indistinguishability, relative to $\mathcal{F}$.

Definition 2.6 (indistinguishable distributions)**.** For $\epsilon \in [0, 1]$ and a set $\mathcal{F}$ of $[0, 1]$-valued functions, distributions ρ_1 and ρ_2 over U are $(\epsilon, \mathcal{F})$-*indistinguishable* if, for every $f \in \mathcal{F}$, $|f[\rho_1] - f[\rho_2]| \leq \epsilon$. A distribution ρ over U is $(\epsilon, \mathcal{F})$-*pseudorandom over* (U, σ) if it is $(\epsilon, \mathcal{F})$-indistinguishable from σ.[3]

[2] In general, $d_\sigma(\mu_\sigma)$ can be arbitrarily bigger than $d_\sigma(\mu)$. For example, let $\mu(x) = \epsilon$ for every $x \in U$, where $0 < \epsilon < 1$. Then $d_\sigma(\mu) = \epsilon$. But $\mu_\sigma = \sigma$, and so $d_\sigma(\mu_\sigma) = 1$.

[3] Note that for σ uniform over U, we get the standard notion of a pseudorandom distribution.

Definition 2.7 (approximation)**.** For two functions $g : U \to [0,1]$ and $h : U \to [0,1]$ and a set $\mathcal{F}$ of $[0,1]$-valued functions, we say that h is an $(\epsilon, \mathcal{F})$-*approximation* of g over the probability space (U, σ) if, for all $f \in \mathcal{F}$, $|\mathbb{E}_\sigma[f(g-h)]| \leq \epsilon$.

Finally, we define the notion of a model for a set and for a distribution.

Definition 2.8 (model)**.** For a set $A \subseteq U$, a measure μ over the probability space (U, σ), and a parameter $\epsilon \in [0,1]$, we say that μ is an $(\epsilon, \mathcal{F})$-*model for* A if the induced distributions A_σ and μ_σ are $(\epsilon, \mathcal{F})$-indistinguishable. Similarly, for a distribution ρ over U, a measure μ over U is an $(\epsilon, \mathcal{F})$-*model* for ρ in the probability space (U, σ) if ρ and μ_σ are $(\epsilon, \mathcal{F})$-indistinguishable.

Remark 2.9. *If μ is a model for some distribution ρ (or a set A), then, by Claim 2.5, so is $\mu' = c \cdot \mu$, for any $c \in (0,1)$. Thus, there is a model for ρ of any density smaller than that of μ.*

Using the definition of density for distributions for a restricted collection $\mathcal{F}$ of tests, we get the following definition of *pseudo-density* for distributions and sets.

Definition 2.10 (pseudo-density)**.** For distribution σ and ρ over U, for a set $\mathcal{F}$ of $[0,1]$-valued functions over U, and parameters $\epsilon, \delta \in [0,1]$, we say that the distribution ρ has $(\epsilon, \mathcal{F})$-*pseudo-density* δ *inside* σ if, for all $f \in \mathcal{F}$, $f[\sigma] \geq \delta \cdot f[\rho] - \epsilon$.

For a probability space $\mathcal{U} = (U, \sigma)$, a set $A \subseteq U$ has $(\epsilon, \mathcal{F})$-*pseudo-density* δ *inside* $\mathcal{U}$ if the induced distribution A_σ has $(\epsilon, \mathcal{F})$-pseudo-density δ inside σ.

A test $f \in \mathcal{F}$ such that $f[\sigma] < \delta \cdot f[\rho] - \epsilon$ witnesses the fact that ρ has density less than δ; such a test is called (ϵ, δ)-*distinctive*. The negative ϵ term is needed to ensure that $poly(1/\epsilon)$ random samples are enough to verify that a given $f \in \mathcal{F}$ is indeed (ϵ, δ)-distinctive. Thus, ρ has pseudo-density δ inside σ if there are no (ϵ, δ)-distinctive tests in $\mathcal{F}$.

2.3 Hardness of Boolean functions

Here we recall the definitions of hardness and hardcore measures for Boolean functions, where we allow not necessarily Boolean classes $\mathcal{F}$ of tests.

Definition 2.11 (hardness and hardcore)**.** For a Boolean function $g : U \to \{0,1\}$ over a probability space $\mathcal{U} = (U, \sigma)$, a class $\mathcal{F}$ of functions over U, and $\delta \in [0,1]$, we say that g is $(\delta, \mathcal{F})$-*hard in* $\mathcal{U}$ if, for all $f \in \mathcal{F}$, we have $\mathbb{E}_\sigma[|f-g|] \geq \delta$.

For a Boolean function $g : U \to \{0,1\}$ over a probability space $\mathcal{U}$, a class $\mathcal{F}$, and an $\epsilon \in [0,1]$, we say that a measure μ over U is an $(\epsilon, \mathcal{F})$-*hardcore for* g *in* $\mathcal{U}$ if, for all $f \in \mathcal{F}$, $|\mathbb{E}_{\mu_\sigma}[|f-g|] - 1/2| \leq \epsilon$.

Remark 2.12. *Since a class $\mathcal{F}$ is closed under complement, to show that a measure μ is an $(\epsilon, \mathcal{F})$-hardcore for a Boolean function g over $\mathcal{U} = (U, \sigma)$, it suffices to argue that $\mathbb{E}_{\mu_\sigma}[|f-g|] \leq 1/2 + \epsilon$. Indeed, applying this upper bound to the complement $\bar{f} = 1 - f$, and using the identity $|1 - f - g| = 1 - |f - g|$ (true for a Boolean g and any f), we also get the lower bound $\mathbb{E}_{\mu_\sigma}[|f-g|] \geq 1/2 - \epsilon$. Similarly, the lower bound $1/2 - \epsilon$ also implies the upper bound $1/2 + \epsilon$. Thus, it suffices to prove either an upper or a lower bound on $\mathbb{E}_{\mu_\sigma}[|f-g|]$.*

2.4 Complexity relative to $\mathcal{F}$

Given a class $\mathcal{F}$ of functions $f : U \to [0,1]$, we consider as "low complexity" those functions $h : U \to [0,1]$ that "can be easily expressed" in terms of functions $f \in \mathcal{F}$, using some "simple" operations. By "easily expressed", we will mean that h has a small-size circuit whose inputs are elements $f \in \mathcal{F}$ and real constants, and whose gates are labeled by the allowed operations. We will allow the following operations: addition, multiplication, threshold, and truncation. Here, for $f \in \mathcal{F}$ and $\theta \in [0,1]$, the *threshold* $th_\theta[f]$ is defined as $th_\theta[f](x) = 1$ if $f(x) \geq \theta$,

and $th_\theta[f](x) = 0$ otherwise. The *truncation* $trunc[f]$ is defined as $trunc[f](x) = f(x)$ if $0 \le f(x) \le 1$, $trunc[f](x) = 1$ if $f(x) > 1$, and $trunc[f](x) = 0$ if $f(x) < 0$.[4] In addition, we allow *limited exponentiation*: for a constant $\beta \in [0,1]$ and a function $f \in \mathcal{F}$, we allow the function $\beta^{f(x)}$.

Remark 2.13. *For a* Boolean *class $\mathcal{F}$, the operation of exponentiation is* not necessary *as we can write $\beta^{f(x)} = 1 - (1-\beta)\cdot f(x)$. We can also avoid exponentiation if we assume that a class $\mathcal{F}$ is closed under thresholds, i.e., if $f \in \mathcal{F}$ implies $th_\theta[f] \in \mathcal{F}$ for all $\theta \in [0,1]$.*

More formally, we have the following.

Definition 2.14 (complexity relative to $\mathcal{F}$)**.** For a function $h : U \to [0,1]$, its *complexity relative to $\mathcal{F}$*, denoted by $\textsc{comp}_{\mathcal{F}}[h]$, is the minimum size of a circuit C computing h, where C is using real constants and functions $f \in \mathcal{F}$ as inputs, and whose gates are labeled by "$+$", "$\times$", "th_θ", or "$trunc$". If, in addition, the operation of limited exponentiation is allowed, we use the notation $\textsc{comp}^*_{\mathcal{F}}[h]$. For $t \in \mathbb{N}$, define $\mathcal{F}_t = \{f \mid \textsc{comp}_{\mathcal{F}}[f] \le t\}$ and $\mathcal{F}^*_t = \{f \mid \textsc{comp}^*_{\mathcal{F}}[f] \le t\}$.

Of special importance to us will be the following set of low-complexity linear threshold functions. For a class $\mathcal{F}$ and a parameter $\lambda \in \mathbb{N}$, define the set of Boolean functions

$$\textsc{Th}_\lambda[\mathcal{F}] = \left\{ th_\theta\left[\tfrac{1}{\lambda'}\sum_{i=1}^{\lambda'} f_i\right] \mid f_i \in \mathcal{F} \setminus \{\mathbf{0},\mathbf{1}\}, \theta \in [0,1], \lambda' \le \lambda \right\},$$

where $\mathbf{0}$ and $\mathbf{1}$ are the constant 0 and 1 functions, respectively. That is, $\textsc{Th}_\lambda[\mathcal{F}]$ is a collection of linear threshold functions of particularly simple form, built from at most λ non-constant functions from $\mathcal{F}$. By definition, we have $\textsc{Th}_\lambda[\mathcal{F}] = \textsc{Th}_\lambda[\mathcal{F} \cup \{\mathbf{0},\mathbf{1}\}]$. Also, for a Boolean class $\mathcal{F}$, $\textsc{Th}_\lambda[\mathcal{F}]$ is also a class.

3 Reductions

Here we give the following *efficient black-box reductions*: Constructive Strong Hardcore Lemma $\Rightarrow$ Constructive Dense Model Theorem $\Rightarrow$ Low Complexity Approximation. We consider *generic* formulations of the respective theorems where certain parameters are left unspecified, with their actual values dependent on a particular proof of the theorem.

> **Generic Constructive Strong Hardcore Lemma:** For every $\epsilon, \delta \in (0,1)$, there is a $\lambda = \lambda(\epsilon,\delta)$ so that the following holds for any probability space $\mathcal{U} = (U,\sigma)$, any class $\mathcal{F}$, and any Boolean function g over U. If g is $(\delta,\mathcal{F}')$-hard in $\mathcal{U}$ for some $\mathcal{F}' \subseteq \mathcal{F}_\lambda$, then there is a (2δ)-dense measure $\mu \in \mathcal{F}''$ that is $(\epsilon,\mathcal{F})$-hardcore for g over $\mathcal{U}$, where $\mathcal{F}'' \subseteq [|g - \mathcal{F}|]^*_\lambda$.

An actual Strong Hardcore Lemma would specify the function λ as well as the form of sets of tests $\mathcal{F}'$ and $\mathcal{F}''$. For example, in case of a Boolean class $\mathcal{F}$, Holenstein [Hol05] achieves $\lambda(\epsilon,\delta) = O(\epsilon^{-2}\delta^{-2})$, $\mathcal{F}' = \textsc{Th}_\lambda[\mathcal{F}]$, and $\mathcal{F}''$ consisting of functions of the form $trunc[c_0 + \sum_i^t c_i \cdot |g - f_i|]$ for $f_i \in \mathcal{F}$, $c_i \in \mathbb{R}$, and $t \le \lambda$. Barak et al. [BHK09] achieve $\lambda(\epsilon,\delta) = O(\epsilon^{-2}\log\delta^{-1})$, but for $\mathcal{F}''$ consisting of more complicated functions in $[|g-\mathcal{F}|]_\lambda$; in case $\mathcal{F}$ is not Boolean, $\mathcal{F}'' \subseteq [|g-\mathcal{F}|]^*_\lambda$.

Next we state the Dense Model Theorem.

> **Generic Constructive Dense Model Theorem:** For every $\epsilon, \delta \in (0,1)$, there is a $\lambda = \lambda(\epsilon,\delta)$ such that the following holds for any probability space $\mathcal{U} = (U,\sigma)$, any class $\mathcal{F}$, and any distribution ρ over U. If ρ has $(\epsilon,\mathcal{F}')$-pseudodensity δ inside σ for some $\mathcal{F}' \subseteq \mathcal{F}_\lambda$, then there is a δ-dense $(O(\epsilon/\delta),\mathcal{F})$-model $\mu \in \mathcal{F}''$ for ρ in $\mathcal{U}$, where $\mathcal{F}'' \subseteq \mathcal{F}^*_\lambda$.

[4] For a function $g : U \to \mathbb{R}$, we have $trunc[g](x) = g(x)(1 - th_1[g](x)) + th_1[g](x)$, and so the operation $trunc$ is not strictly necessary, but is included for convenience.

A non-constructive form of the Dense Model Theorem with the uniform distribution σ and $\lambda(\epsilon,\delta) = O((\delta/\epsilon)^2 \log \epsilon^{-1})$ was given by Reingold et al. [RTTV08]. A constructive version was proved by Zhang [Zha11], achieving $\lambda(\epsilon,\delta) = O((\delta/\epsilon)^2 \log \delta^{-1})$ and $\mathcal{F}' = \mathrm{TH}_\lambda[\mathcal{F}]$.

Remark 3.1. *The parameter $O(\epsilon/\delta)$ for the quality of the δ-dense model in the formulation of the Dense Model Theorem above is tight, up to a constant factor.*

Indeed, consider the case of the uniform distribution σ. For any $0 < \epsilon \leq \delta < 1$, let $S \subseteq U$ be of density $\epsilon/2$, and let $T \subseteq S$ be of density ϵ/δ inside S. Consider the class of tests $\mathcal{F} = \{T, \bar{T}\}$, where T is the indicator function of the set T, and $\bar{T}$ is its complement. Observe that $T[\sigma] \geq 0 = \delta(\epsilon/\delta) - \epsilon = \delta \cdot T[S_\sigma] - \epsilon$, and $\bar{T}[\sigma] = 1 - T[\sigma] \geq 1 - \epsilon \geq \delta \cdot (1 - T[S_\sigma]) - \epsilon$. Hence, for every $f \in \mathrm{TH}_\lambda[\mathcal{F}]$ (for any λ), we have $f[\sigma] \geq \delta \cdot f[S_\sigma] - \epsilon$, and so S has $(\epsilon, \mathrm{TH}_\lambda[\mathcal{F}])$-pseudodensity δ inside σ (for any value λ).

On the other hand, for any measure μ of density δ, we have $T[\mu_\sigma] = (1/d_\sigma(\mu)) \cdot \mathbb{E}_\sigma[T \cdot \mu] \leq (1/\delta) \cdot \mathbb{E}_\sigma[T] \leq \epsilon/(2\delta)$. For S_σ, we have $T[S_\sigma] = \epsilon/\delta$. Hence, $|T[S_\sigma] - T[\mu_\sigma]| \geq \epsilon/(2\delta)$. Thus, the test $T \in \mathcal{F}$ can $\epsilon/(2\delta)$-distinguish between S_σ and μ_σ, for every δ-dense measure μ.

The pseudorandom formulation of the Dense Model Theorem is as follows:

> **Generic Constructive Dense Model Theorem (pseudorandom version):** For every $\epsilon, \delta \in (0,1)$, there is a $\lambda = \lambda(\epsilon,\delta)$ such that the following holds for any probability space $\mathcal{U} = (U,\sigma)$, any class $\mathcal{F}$, and any distributions ρ and τ over U. If ρ is δ-dense inside τ, where τ is $(\epsilon,\mathcal{F}')$-pseudorandom in $\mathcal{U}$, then there is a δ-dense $(O(\epsilon/\delta),\mathcal{F})$-model $\mu \in \mathcal{F}''$ for ρ in $\mathcal{U}$, where $\mathcal{F}'' \subseteq \mathcal{F}'^*_\lambda$.

Trevisan et al. [TTV09] achieve $\lambda(\epsilon,\delta) = O(\epsilon^{-2})$ and $\mathcal{F}' = \mathcal{F}_\lambda$. Zhang [Zha11] gets $\lambda = O((\delta/\epsilon)^2 \log \delta^{-1})$ and $\mathcal{F}' = \mathrm{TH}_\lambda[\mathcal{F}]$.

Remark 3.2. *As before, the relation between density and model parameters is tight for the pseudorandom formulation of the Dense Model Theorem as well. We can modify our earlier construction to get a $(O(\epsilon), \mathrm{TH}_\lambda[\mathcal{F}])$-pseudorandom R, with S δ-dense in R, and the rest of the argument is identical.*

Let σ, $T \subseteq S \subseteq U$, and $\mathcal{F}$ be as in Remark 3.1. Define $R = S \cup V$ where $V \subset \bar{S}$ is such that S has density δ inside R. Observe that $T[\sigma] \leq \epsilon/2$. On the other hand, $T[R_\sigma] = \delta \cdot T[S_\sigma] = \delta \cdot (\epsilon/\delta) = \epsilon$. Hence, $|T[R_\sigma] - T[\sigma]| \leq \epsilon$. It follows that R is $(\epsilon, \mathrm{TH}_\lambda[\mathcal{F}])$-pseudorandom (for every λ), and S is δ-dense inside R. But, as we argued in Remark 3.1, the test T will $\epsilon/(2\delta)$-distinguish between S_σ and μ_σ, for every δ-dense measure μ.

Remark 3.3. *It is easy to see that the pseudodensity formulation of the Dense Model Theorem implies the pseudorandom formulation. If a distribution ρ is δ-dense inside a $(\epsilon,\mathcal{F}')$-pseudorandom distribution τ, then ρ also has $(\epsilon,\mathcal{F}')$-pseudodensity at least δ. Indeed, by assumption, $\tau(x) \geq \delta \cdot \rho(x)$ for all $x \in U$. Hence, for every f, we have $f[\tau] \geq \delta \cdot f[\rho]$. By the definition of pseudorandomness, we have for every $f \in \mathcal{F}'$ that $|f[\tau] - f[\sigma]| \leq \epsilon$. It follows that, for every $f \in \mathcal{F}'$, $f[\sigma] \geq f[\tau] - \epsilon \geq \delta \cdot f[\rho] - \epsilon$, as required.*

Finally, we state the Low Complexity Approximation theorem.

> **Generic Low Complexity Approximation Theorem:** For every $\epsilon > 0$, there exists a $\lambda = \lambda(\epsilon)$ such that the following holds for any probability space $\mathcal{U} = (U,\sigma)$ and any class $\mathcal{F}$ over U. Every function $g : U \to [0,1]$ has an $(\epsilon,\mathcal{F})$-approximation h in $\mathcal{U}$ such that $\mathbb{E}_\sigma[h] = \mathbb{E}_\sigma[g]$.

Trevisan et al. [TTV09] get $\lambda = O(\epsilon^{-2} \log \epsilon^{-1})$ in the general case (and, in the case of a Boolean $\mathcal{F}$, $\lambda = O(\epsilon^{-2})$).

3.1 Strong Hardcore Lemma implies Dense Model Theorem

Assuming the generic Constructive Strong Hardcore Lemma, parametrized by a function $\lambda(\epsilon,\delta)$ and test sets $\mathcal{F}'$ and $\mathcal{F}''$ over U, we derive the Dense Model Theorem which inherits the type of tests of $\mathcal{F}'$ and $\mathcal{F}''$ from the Hardcore Lemma, and whose complexity parameter λ' is $\lambda(O(\epsilon),O(\delta))$. For concreteness, we prove the Dense Model Theorem assuming the Hardcore Lemma with $\lambda(\epsilon,\delta)=O(\epsilon^{-2}\log\delta^{-1})$, $\mathcal{F}'=\textsc{Th}_\lambda[\mathcal{F}]$ and $\mathcal{F}''=[|g-\mathcal{F}|]^*_\lambda$ (i.e.,, assuming Barak et al.'s formulation of the Hardcore Lemma [BHK09]). However, we stress that our reduction is "black-box", and could use any other version of the Strong Hardcore Lemma.

Theorem 3.4. *Given $\epsilon,\delta\in(0,1)$ such that $\epsilon\le\delta/3$, there exists a $\lambda=O(\epsilon^{-2}\log\delta^{-1})$ such that the following holds. Let $\mathcal{U}=(U,\sigma)$ be any finite probability space, and let $\mathcal{F}$ be any class of functions over U. If a distribution ρ over U has $(\epsilon,\textsc{Th}_\lambda[\mathcal{F}])$-pseudodensity δ in σ, then there exists a $(12\cdot\epsilon/\delta,\mathcal{F})$-model $\mu\in\mathcal{F}^*_\lambda$ for ρ of density $\delta-3\epsilon$.*[5]

Proof. Set $\bar\delta:=\delta/(1+\delta)$ (implying $\delta=\bar\delta/(1-\bar\delta)$) and $\bar\epsilon:=\epsilon/6$. Define $\hat U:=\{(0,x)\mid x\in U\}\cup\{(1,x)\mid x\in\textsc{Support}[\rho]\}$, along with the following distribution $\hat\sigma$ over $\hat U$: for $(b,x)\in\hat U$, where $b\in\{0,1\}$ and $x\in U$, define

$$\hat\sigma(b,x)=\begin{cases}(1-\bar\delta)\cdot\sigma(x) & \text{if } b=0\\ \bar\delta\cdot\rho(x) & \text{if } b=1.\end{cases}$$

Since $\textsc{Th}_\lambda[\mathcal{F}]=\textsc{Th}_\lambda[\mathcal{F}\cup\{\mathbf{0},\mathbf{1}\}]$, we can assume without loss of generality that $\mathbf{0},\mathbf{1}\in\mathcal{F}$. Associate with each $f\in\mathcal{F}$ a function $\hat f:\hat U\to\{0,1\}$ such that, for any $(b,x)\in\hat U$, $\hat f(b,x):=f(x)$. Define $\hat{\mathcal{F}}=\{\hat f\mid f\in\mathcal{F}\}$. Clearly, $\hat{\mathcal{F}}$ is a class of Boolean functions over $\hat U$.

Consider $g:\hat U\to\{0,1\}$ where $g(b,x)=b$ for every $(b,x)\in\hat U$. Our tests $\hat f\in\hat{\mathcal{F}}$, on input (b,x), ignore b and use x only. Such tests have difficulty in computing g, as we show next.

Claim 3.5. *For $\hat\delta:=\bar\delta-\epsilon+\epsilon\bar\delta$, the function g is $(\hat\delta,\textsc{Th}_\lambda[\hat{\mathcal{F}}])$-hard in $(\hat U,\hat\sigma)$.*

Proof. Suppose there is $\hat\phi\in\textsc{Th}_\lambda[\hat{\mathcal{F}}]$, corresponding to $\phi\in\textsc{Th}_\lambda[\mathcal{F}]$, such that $\mathbb{E}_{\hat\sigma}[|g-\hat\phi|]<\hat\delta$. By considering separately inputs $\{1\}\times U$ and $\{0\}\times U$, we have $\mathbb{E}_{\hat\sigma}[|g-\hat\phi|]=\bar\delta\cdot\mathbb{E}_\rho[1-\phi]+(1-\bar\delta)\cdot\mathbb{E}_\sigma[\phi]$, and so $\bar\delta\cdot(1-\phi[\rho])+(1-\bar\delta)\cdot\phi[\sigma]<\bar\delta-\epsilon(1-\bar\delta)$. Dividing both sides of this inequality by $(1-\bar\delta)$ and using $\delta=\bar\delta/(1-\bar\delta)$, we get $\phi[\sigma]<\delta\cdot\phi[\rho]-\epsilon$, contradicting the pseudodensity δ of ρ. □

By the Strong Hardcore Lemma, for $\lambda=O(\bar\epsilon^{-2}\log(\hat\delta^{-1}))=O(\epsilon^{-2}\log\delta^{-1})$, there exists an $(\bar\epsilon,\hat{\mathcal{F}})$-hardcore measure $\eta\in[|g-\hat{\mathcal{F}}|]^*_\lambda$ of density at least $2\hat\delta$ over $(\hat U,\hat\sigma)$. Define $\eta_1(x):=\eta(1,x)\in\mathcal{F}^*_\lambda$ and $\eta_0(x):=\eta(0,x)\in\mathcal{F}^*_\lambda$. We get

$$d_{\hat\sigma}(\eta)=\bar\delta\cdot d_\rho(\eta_1)+(1-\bar\delta)\cdot d_\sigma(\eta_0)\ge 2(\bar\delta-\epsilon+\epsilon\bar\delta). \tag{1}$$

We will argue that η_0 is a dense model for ρ. First we lower-bound $d_\sigma(\eta_0)$ and $d_\rho(\eta_1)$.

Claim 3.6. *$d_\sigma(\eta_0)\ge\delta-(7/3)\epsilon$, and $d_\rho(\eta_1)\ge 1-(7/3)\epsilon/\delta$.*

Proof. Since $\mathbf{0},\mathbf{1}\in\hat{\mathcal{F}}$, by the definition of hardcore we get that both $\mathbb{P}_{\eta_{\hat\sigma}}[g=1]$ and $\mathbb{P}_{\eta_{\hat\sigma}}[g=0]$ are in the interval $[\frac12-\bar\epsilon,\frac12+\bar\epsilon]$, and so are within $2\bar\epsilon$ of each other. We have $\mathbb{P}_{\eta_{\hat\sigma}}[g=1]=\sum_{x\in U}\eta_{\hat\sigma}(1,x)=(1/d_{\hat\sigma}(\eta))\cdot\sum_{x\in U}\eta_1(x)\hat\sigma(1,x)=(1/d_{\hat\sigma}(\eta))\cdot\bar\delta\cdot d_\rho(\eta_1)$, and similarly, $\mathbb{P}_{\eta_{\hat\sigma}}[g=0]=\sum_{x\in U}\eta_{\hat\sigma}(0,x)=(1/d_{\hat\sigma}(\eta))\cdot\sum_{x\in U}\eta_0(x)\hat\sigma(0,x)=(1/d_{\hat\sigma}(\eta))\cdot(1-\bar\delta)\cdot d_\sigma(\eta_0)$. It follows that $|\bar\delta\cdot d_\rho(\eta_1)-(1-\bar\delta)\cdot d_\sigma(\eta_0)|\le 2\bar\epsilon\cdot d_{\hat\sigma}(\eta)\le 2\bar\epsilon$. Together with Eq. (1), this implies $d_\sigma(\eta_0)\ge(\hat\delta-\bar\epsilon)/(1-\bar\delta)=(\bar\delta-\epsilon+\epsilon\bar\delta-\bar\epsilon)\cdot(1+\delta)\ge\delta-2(\epsilon+\bar\epsilon)=\delta-7\epsilon/3$, and $d_\rho(\eta_1)\ge(\hat\delta-\bar\epsilon)/\bar\delta=(\hat\delta-\bar\epsilon)(1+\delta)/\delta\ge(\delta-(7/3)\epsilon)/\delta=1-(7/3)\epsilon/\delta$. □

[5] The 3ϵ slack in the density can be moved into the error term by averaging the measure μ with the constant 1 measure; we skip the details.

Next we show that η_0 is a model for ρ in (U,σ). To this end, we first argue that $(\eta_0)_\sigma$ is indistinguishable from $(\eta_1)_\rho$, and that $(\eta_1)_\rho$ is indistinguishable from ρ by tests in $\mathcal{F}$; applying the triangle inequality will then conclude the proof of the theorem.

Claim 3.7. *The distributions $(\eta_0)_\sigma$ and $(\eta_1)_\rho$ are $(\epsilon,\mathcal{F})$-indistinguishable.*

Proof. Let $f \in \mathcal{F}$ be arbitrary. For the corresponding test $\hat{f} \in \hat{\mathcal{F}}$, we get by the definition of hardcore that $\mathbb{E}_{\eta_{\hat{\sigma}}}[|\hat{f}-g|] \in [\frac{1}{2}-\bar{\epsilon},\frac{1}{2}+\bar{\epsilon}]$. Conditioning on $g=0$ and $g=1$, we get

$$\mathbb{E}_{\eta_{\hat{\sigma}}}[|\hat{f}-g|] = \mathbb{E}_{\eta_{\hat{\sigma}}}[f \mid g=0]\cdot\mathbb{P}_{\eta_{\hat{\sigma}}}[g=0] + \mathbb{E}_{\eta_{\hat{\sigma}}}[1-f \mid g=1]\cdot\mathbb{P}_{\eta_{\hat{\sigma}}}[g=1]. \tag{2}$$

We have

$$\mathbb{E}_{\eta_{\hat{\sigma}}}[f \mid g=0] = \frac{\sum_{x\in U} f(x)\eta_0(x)\sigma(x)(1-\bar{\delta})}{\sum_{x\in U}\eta_0(x)\sigma(x)(1-\bar{\delta})} = \mathbb{E}_{(\eta_0)_\sigma}[f] = f[(\eta_0)_\sigma], \tag{3}$$

and, similarly,

$$\mathbb{E}_{\eta_{\hat{\sigma}}}[1-f \mid g=1] = \frac{\sum_{x\in U}(1-f(x))\eta_1(x)\rho(x)\bar{\delta}}{\sum_{x\in U}\eta_1(x)\rho(x)\bar{\delta}} = \mathbb{E}_{(\eta_1)_\rho}[1-f] = 1-f[(\eta_1)_\rho]. \tag{4}$$

Also, since $\mathbf{0},\mathbf{1}\in\hat{\mathcal{F}}$, we get by the definition of hardcore that both $\mathbb{P}_{\eta_{\hat{\sigma}}}[g=0]$ and $\mathbb{P}_{\eta_{\hat{\sigma}}}[g=1]$ are in the interval $[\frac{1}{2}-\bar{\epsilon},\frac{1}{2}+\bar{\epsilon}]$. Combining this with Eqs. (2)–(4) yields

$$1-\epsilon \le \frac{1-2\bar{\epsilon}}{1+2\bar{\epsilon}} \le 1-f[(\eta_1)_\sigma]+f[(\eta_0)_\rho] \le \frac{1+2\bar{\epsilon}}{1-2\bar{\epsilon}} \le 1+\epsilon,$$

where we used that $\bar{\epsilon}=\epsilon/6\le 1/6$. We conclude that $|f[(\eta_0)_\rho]-f[(\eta_1)_\sigma]|\le\epsilon$, as required. □

Claim 3.8. *The distributions $(\eta_1)_\rho$ and ρ are $(11\cdot\epsilon/\delta,\mathcal{F})$-indistinguishable.*

Proof. Let $f\in\mathcal{F}$ be arbitrary. We have that $f[(\eta_1)_\rho]-f[\rho]$ equals

$$\frac{\mathbb{E}_\rho[f\cdot\eta_1]}{d_\rho(\eta_1)} - \mathbb{E}_\rho[f] = \frac{1}{d_\rho(\eta_1)}\cdot\mathbb{E}_\rho[f\cdot(\eta_1-d_\rho(\eta_1))] \le \frac{1}{d_\rho(\eta_1)}\cdot\mathbb{E}_\rho[1-d_\rho(\eta_1)] = \frac{1}{d_\rho(\eta_1)}-1,$$

where for the inequality we first used $f(x)\ge 0$ for all $x\in U$ to get $f(x)\cdot(\eta_1(x)-d_\rho(\eta_1))\le f(x)\cdot(1-d_\rho(\eta_1))$, and then used $1-d_\rho(\eta_1)\ge 0$ to get $f(x)\cdot(1-d_\rho(\eta_1))\le 1\cdot(1-d_\rho(\eta_1))$. By Claim 3.6, $d_\rho(\eta_1)\ge 1-(7/3)\epsilon/\delta$, and so, $1/d_\rho(\eta_1)-1\le(\epsilon/\delta)/(3/7-\epsilon/\delta)\le(21/2)\epsilon/\delta$, where we used our assumption that $\epsilon\le\delta/3$ to get the lower bound $3/7-\epsilon/\delta\ge 3/7-1/3=2/21$.

Thus, $f[(\eta_1)_\rho]-f[\rho]\le(10.5)\epsilon/\delta$, for every $f\in\mathcal{F}$. Since $\mathcal{F}$ is closed under complement, we also get for every $f\in\mathcal{F}$ that $1-f[(\eta_1)_\rho]-(1-f[\rho])=f[\rho]-f[(\eta_1)_\rho]\le(10.5)\epsilon/\delta$. □

Finally, we argue that η_0 is a model for ρ in σ. Let $f\in\mathcal{F}$ be arbitrary. By the triangle inequality and Claims 3.7 and 3.8, we get $|f[\rho]-f[(\eta_0)_\sigma]|\le|f[\rho]-f[(\eta_1)_\rho]|+|f[(\eta_1)_\rho]-f[(\eta_0)_\sigma]|\le(10.5)\epsilon/\delta+\epsilon\le 12\epsilon/\delta$. Hence, η_0 is a $(12\cdot\epsilon/\delta,\mathcal{F})$-model for ρ in (U,σ) of density at least $\delta-3\epsilon$. □

3.2 Dense Model Theorem implies Low-Complexity Approximation Theorem

For concreteness, we assume Zhang's formulation of the pseudorandom version of the Dense Model Theorem with $\lambda(\epsilon,\delta)=O(\epsilon^{-2}\log\delta^{-1})$ and $\mathcal{F}'=\mathrm{TH}_\lambda[\mathcal{F}]$ and $\mathcal{F}''=\mathcal{F}^*_\lambda$. Using this DMT, we derive the following Low-Complexity Approximation Theorem.

Theorem 3.9. *For any $\epsilon>0$ there is a $\lambda(\epsilon)=O(\epsilon^{-2})$ such that the following holds. For any probability space $\mathcal{U}=(U,\sigma)$ and any class $\mathcal{F}$ on U, every function $g:U\to[0,1]$ has an $(\epsilon,\mathcal{F})$-approximation $h\in\mathcal{F}^*_\lambda$ in $\mathcal{U}$ such that $\mathbb{E}_\sigma[h]=\mathbb{E}_\sigma[g]$.*

Proof. Let $\delta = d_\sigma(g)$. We may assume, without loss of generality, that $\delta \geq 1/2$; otherwise, we work with $\bar{g}$. By Claim 2.4, the induced distribution g_σ has density at least α inside σ, where σ is obviously $(\epsilon', \mathcal{F}')$-pseudorandom in $\mathcal{U}$ for any $\epsilon' > 0$ and any set $\mathcal{F}'$ of tests. Set $\epsilon' = \epsilon/c$ for some large constant c. By the Dense Model Theorem, we get a δ-dense measure μ that is an $(\epsilon, \mathcal{F})$-model for g_σ. The measure μ is in $\mathcal{F}^*_{O(\epsilon^{-2})}$ since $\delta \geq 1/2$.

If $d_\sigma(\mu) > \delta$, we consider the scaled-down version $c' \cdot \mu$ for some $c' < 1$ so that $d_\sigma(c' \cdot \mu) = \delta$; recall that such scaling preserves the property of being a model. Let h be the resulting (possibly scaled) measure. Note that $\mathbb{E}_\sigma[h] = \mathbb{E}_\sigma[g]$ by construction.

By the definition of a model, we get, for every $f \in \mathcal{F}$, $|f[h_\sigma] - f[g_\sigma]| \leq \epsilon$, which is equivalent to

$$\left| \frac{\mathbb{E}_\sigma[f \cdot h]}{d_\sigma(h)} - \frac{\mathbb{E}_\sigma[f \cdot g]}{d_\sigma(g)} \right| \leq \epsilon.$$

But since $d_\sigma(h) = d_\sigma(g) = \delta$, we conclude that $|\mathbb{E}_\sigma[f \cdot (h - g)]| \leq \epsilon\delta \leq \epsilon$, as required. □

Remark 3.10. *It is possible to strengthen the conclusion on the complexity of h so that $h \in \mathcal{F}_\lambda$ (i.e., drop the need for limited exponentiation operations). By averaging, one can show that if a test f witnesses that h is not an ϵ-approximation of g, then there is a threshold $\theta \in [0, 1]$ such that the Boolean test $th_\theta[f]$ witnesses that too. Consequently, it suffices to get a model with respect to the set $\mathcal{F}^{th} = \{th_\theta[f], 1 - th_\theta[f] \mid f \in \mathcal{F}, \theta \in [0, 1]\}$ of Boolean tests. Since, in the proof of Theorem 3.9, we can assume pseudorandomness of σ with respect to an arbitrary set $\mathcal{F}'$ of tests, we can apply the Dense Model Theorem with $\mathcal{F}^{th}$ as our class $\mathcal{F}$. As a result, we end up working with the Boolean class $\mathcal{F}$, and so limited exponentiation is not needed.*

4 The MinMax formulation of the Dense Model Theorem

Here we show that the Strong Hardcore Lemma, Dense Model Theorem (both pseudodense and pseuodrandom versions), and Low-Complexity Approximation Theorem all follow from the following LP Duality, or MinMax, formulation of the Dense Model Theorem:

> **Constructive MinMax version of Dense Model Theorem:** For every $\epsilon, \delta \in (0, 1)$, there is a $\lambda = \lambda(\epsilon, \delta)$ so that the following holds for any probability space $\mathcal{U} = (U, \sigma)$, any class $\mathcal{F}$, and any distribution ρ over U:
>
> - either ρ has a δ-dense $(\epsilon, \mathcal{F})$-model $\mu \in \mathcal{F}'' \subseteq \mathcal{F}^*_\lambda$ in $\mathcal{U}$,
> - or there is a function F, the average of fewer than λ functions f from $\mathcal{F}$, and a constant $c > 0$ such that F is an ϵ/c-distinguisher between ρ and any given δ-dense measure: for any δ-dense measure γ, $F[\gamma_\sigma] - F[\rho] > \epsilon/c$.

4.1 Deriving the Dense Model Theorem

Reingold et al. [RTTV08] showed how the MinMax version of DMT implies the pseudodense formulation of DMT: assuming the existence of the universal distinguisher F, one shows that ρ does not have $(\Omega(\epsilon\delta), \textsc{Th}_\lambda[\mathcal{F}])$-pseudodensity δ. For completeness, we present their argument below.

Suppose F is such that $F[\gamma_\sigma] > F[\rho] + \epsilon/c$, for some constant $c > 0$. Then $\bar{F}[\rho] > \bar{F}[\gamma_\sigma] + \epsilon/c$. Order elements of U such that $\bar{F}(x_i) \geq \bar{F}(x_{i+1})$. Let n be the largest integer such that $d_\sigma(\{x_i : i \in [n]\}) < \delta$, define measure γ by $\gamma(x_i) = 0$ for $i \geq n + 2$, $\gamma(x_i) = 1$ for $i \in [n]$ and $\gamma(x_{n+1}) = c$ where $c \in (0, 1]$ is such that $d_\sigma(\gamma) = \delta$. By averaging, there exists $\kappa \in [0, 1]$ such that for $\Phi = th_\kappa[\bar{F}] \in \textsc{Th}_\lambda[\mathcal{F}]$,

$$\Phi[\rho] > \Phi[\gamma_\sigma] + \epsilon/c. \tag{5}$$

Since $1 \geq \mathbb{E}_\rho[\Phi] > \mathbb{E}_{\gamma_\sigma}[\Phi] + \epsilon/c$, we get $1 > \mathbb{E}_{\gamma_\sigma}[\Phi]$. Since Φ is Boolean, there must be an $x^* \in \textsc{support}[\gamma]$ such that $\Phi(x^*) = 0$, and, in particular, $\Phi(x_{n+1}) = 0$ (since $\bar{F}(x_{n+1}) =$

$min_{x \in \text{SUPPORT}[\gamma]}[\bar{F}(x)]$). So, for every $x \in \text{SUPPORT}[\bar{\gamma}] \subset \{x_i : i \geq n+1\}$, $\Phi(x) = 0$, implying $\mathbb{E}_\sigma[\Phi \cdot \bar{\gamma}] = 0$. Using the identity $d_\sigma(\gamma) \cdot \mathbb{E}_{\gamma_\sigma}[\Phi] = \mathbb{E}_\sigma[\Phi \cdot \gamma]$, we get $\mathbb{E}_\sigma[\Phi] = \mathbb{E}_\sigma[\Phi \cdot \gamma] + \mathbb{E}_\sigma[\Phi \cdot \bar{\gamma}] = \delta \cdot \mathbb{E}_{\gamma_\sigma}[\Phi]$. By Eq. (5), we conclude that $\Phi[\sigma] < \delta \cdot \Phi[\rho] - \epsilon\delta/c$.

A very similar argument also shows that the pseudorandom version of DMT is implied by the same constructive MinMax formulation. We show that the existence of the universal distinguisher F implies that any distribution τ in which ρ is δ-dense cannot be $(\Omega(\epsilon\delta), \text{TH}_\lambda[\mathcal{F}])$-pseudorandom. Indeed, as in the argument above, we get Φ such that $\delta \cdot \Phi[\rho] - \Phi[\sigma] > \epsilon\delta/c$. By the δ-density of ρ inside τ, we get $\Phi[\tau] \geq \delta \cdot \Phi[\rho]$, and hence, $\Phi[\tau] - \Phi[\sigma] > \epsilon\delta/c$.

4.2 Deriving the Low-Complexity Approximation Theorem

Let $\delta = d_\sigma(g)$ and assume that $\delta \geq \frac{1}{2}$ (otherwise we can work with $\bar{g}$). Set $\rho := g_\sigma$. Recall that ρ has density at least δ.

By the constructive MinMax version of DMT, either we have a δ-dense model $\mu \in \mathcal{F}^*_\lambda$ for ρ, or we have a universal distinguisher F. In the former case, μ is an approximation to g (by the same argument as in the proof of Theorem 3.9). The latter case is impossible. Indeed, suppose $F : U \to [0,1]$ is such that, for some constant $c > 0$, $F[\mu_\sigma] - F[\rho] > \epsilon/c$ for every measure μ over U of density $d_\sigma(\mu) = \delta$. For $\mu = g$, this yields $0 > \epsilon/c$, a contradiction.

4.3 Deriving the Strong Hardcore Lemma

Using the constructive MinMax version of DMT, we derive the Strong Hardcore Lemma. More precisely, we first show that every hardcore measure, however small, has the optimal pseudodensity, and hence, has the model of the optimal density, which is the required hardcore measure. Then we give a simple construction of hardcore measure whose support consists of just two elements. Together, these two results imply the Strong Hardcore Lemma.

4.3.1 From small to large hardcore measure

Suppose that μ_0 is a $(\epsilon/(2c), \mathcal{F})$-hardcore measure for a Boolean function $g : U \to \{0,1\}$, where $\mathcal{F}$ is an arbitrary class of tests over the probability space $\mathcal{U} = (U, \sigma)$, and $c > 0$ is a large enough constant. Define the new class of tests $\hat{\mathcal{F}} := \{|g - f| \mid f \in \mathcal{F}\}$. Define $\rho := (\mu_0)_\sigma$.

By the constructive MinMax version of DMT applied to ρ, $\hat{\mathcal{F}}$, and 2δ, we have that either ρ has a 2δ-dense $(\epsilon, \hat{\mathcal{F}})$-model $\mu^* \in \hat{\mathcal{F}}^*_\lambda$, or there is a universal distinguisher F between ρ and any 2δ-dense measure γ.

In the former case, it is easy to see that μ^* must be a hardcore measure. Indeed, by the definition of a model, we have that, for all $f \in \mathcal{F}$, $|\mathbb{E}_{(\mu^*)_\sigma}[|g - f|] - \mathbb{E}_{(\mu_0)_\sigma}[|g - f|]| \leq \epsilon$. By the definition of hardcore, we have that $|\mathbb{E}_{(\mu_0)_\sigma}[|g - f|] - \frac{1}{2}| \leq \epsilon/(2c)$. Hence, $|\mathbb{E}_{(\mu^*)_\sigma}[|g - f|] - \frac{1}{2}| \leq (1 + 1/(2c))\epsilon$, implying that μ^* is $((1 + 1/(2c))\epsilon, \mathcal{F})$-hardcore measure for g in $\mathcal{U}$.

It remains to argue that the latter case is impossible. We show that the existence of F implies that g is not δ-hard. Indeed, suppose we have a function $F = \frac{1}{T}\sum_{t=1}^{T}(|g - f_t|)$, with $f_t \in \mathcal{F}$ for all $1 \leq t \leq T < \lambda$, such that $F[\mu_\sigma] - F[(\mu_0)_\sigma] > \epsilon/c$ for every measure μ over U with $d_\sigma(\mu) = 2\delta$. Since μ_0 is an $\epsilon/(2c)$-hardcore for g, we get that $|g - f_t|[(\mu_0)_\sigma] \in [\frac{1}{2} - \epsilon/(2c), \frac{1}{2} + \epsilon/(2c)]$ for every $1 \leq t \leq T$, and hence, $F[\mu_\sigma] > \frac{1}{2} + \epsilon/(2c)$. Conditioning on $g = 0$ and $g = 1$, we get

$$F = \frac{1}{T}\sum_{t=1}^{T}|g - f_t| = \left|g - \frac{1}{T}\sum_{t=1}^{T} f_t\right|,$$

and hence, $\mathbb{E}_{\mu_\sigma}[|g - \Phi|] > \frac{1}{2}$, where $\Phi := \frac{1}{T}\sum_{t=1}^{T} f_t$. The contradiction is now achieved by the following claim, due to Holenstein [Hol05]; for completeness, we give its proof below.

Claim 4.1 ([Hol05]). *Suppose $\Phi : U \to [0,1]$ is such that, for all measures μ over U with $d_\sigma(\mu) = 2\delta$, $\mathbb{E}_{\mu_\sigma}[|g - \Phi|] > \frac{1}{2}$. Then there is a $\theta \in [0,1]$ such that $\mathbb{E}_\sigma[|g - th_\theta[\bar{\Phi}]|] < \delta$, where $\bar{\Phi} = 1 - \Phi$ is the complement of Φ.*

Note that $th_\theta[\bar{\Phi}] \in \mathrm{TH}_T[\mathcal{F}]$. Since g is assumed $(\delta, \mathrm{TH}_\lambda[\mathcal{F}])$-hard in $\mathcal{U}$, we get a contradiction.

Proof of Claim 4.1. Define $\alpha_c(x) := 2|g - \Phi| - 1$ and $\alpha_1(x) := 2\bar{\Phi}(x) - 1$. Order elements $\alpha_c(x)$ from smallest to largest, inducing the ordering on the elements $x \in U$: $x_1, x_2, x_3, \ldots, x_{|U|}$. Define $S = \{x_1, \ldots, x_n\}$ for the biggest $1 \le n \le |U|$ such that $d_\sigma(S) < 2\delta$. Define the measure μ over U as follows: $\mu(x_i) = 1$ for $1 \le i \le n$, $\mu(x_j) = 0$ for $j > n+1$, and $\mu(x_{n+1}) = c$ for $0 < c \le 1$ so that $d_\sigma(\mu) = 2\delta$. Note that $\textsc{Support}[\mu_\sigma] = \{x_1, \ldots, x_{n+1}\}$.

We can represent σ as a convex combination of the induced distributions μ_σ and $(\bar{\mu})_\sigma$, where $\bar{\mu} = 1 - \mu$ is the complement of μ. More precisely, we have, for every $x \in U$,

$$\sigma(x) = d_\sigma(\mu) \cdot \mu_\sigma(x) + (1 - d_\sigma(\mu)) \cdot (\bar{\mu})_\sigma(x). \tag{6}$$

It is easy to see that the support of $(\bar{\mu})_\sigma$ is either $\{x_{n+2}, \ldots, x_{|U|}\}$ (if $\mu(x_{n+1}) = 1$), or $\{x_{n+1}, x_{n+2}, \ldots, x_{|U|}\}$ (if $\mu(x_{n+1}) < 1$).

Define $\kappa := \max_{x \in \textsc{Support}[\mu_\sigma]}\{\alpha_c(x)\} = \alpha_c(x_{n+1})$. Note $\mathbb{E}_{\mu_\sigma}[\alpha_c(x)] > 0$ as, by assumption, $\mathbb{E}_{\mu_\sigma}[|g - \Phi|] > \frac{1}{2}$. This also means that $\kappa > 0$.

Consider the probabilistic Boolean function Ψ (with internal randomness r) which, on input $x \in U$, outputs 1 with probability $trunc[\frac{1}{2} + \frac{\alpha_1(x)}{2\kappa}]$. It is not hard to see that, on each input $x \in U$, we have $\mathbb{P}_r[\Psi(x) = g(x)] = trunc[\frac{1}{2} + \frac{\alpha_c(x)}{2\kappa}]$.

Since $\alpha_c(x_i) \ge \kappa$ for all $i \ge n+1$, we get $\mathbb{P}_r[\Psi(x_i) = g(x_i)] = 1$ for all $i > n$. Hence, in particular, $\mathbb{E}_{(\bar{\mu})_\sigma}[\mathbb{P}_r[\Psi(x) = g(x)]] = 1$. Next, for $1 \le i \le n+1$, we have $\alpha_c(x_i) \le \kappa$, and so $\frac{1}{2} + \frac{\alpha_c(x_i)}{2\kappa} \le 1$. Therefore, $\mathbb{E}_{\mu_\sigma}[\mathbb{P}_r[\Psi(x) = g(x)]]$ is at least

$$\mathbb{E}_{\mu_\sigma}\left[\frac{1}{2} + \frac{\alpha_c(x)}{2\kappa}\right] = \frac{1}{2} + \frac{\mathbb{E}_{\mu_\sigma}[\alpha_c(x)]}{2\kappa} > \frac{1}{2}.$$

Using Eq. (6), we get $\mathbb{E}_\sigma[\mathbb{P}_r[\Psi(x) = g(x)]] = 2\delta \cdot \mathbb{E}_{\mu_\sigma}[\mathbb{P}_r[\Psi(x) = g(x)]] + (1-2\delta) \cdot \mathbb{E}_{(\bar{\mu})_\sigma}[\mathbb{P}_r[\Psi(x) = g(x)]] > 2\delta/2 + (1-2\delta) = 1 - \delta$. By averaging, we can fix the internal randomness of Ψ, preserving this inequality. The resulting deterministic algorithm has the form: On input $x \in U$, compute $\frac{1}{2} + \frac{2\bar{\Phi}(x)-1}{2\kappa}$, and output 1 iff the result is at least α, for some threshold $\alpha \in [0,1]$ (the same for all inputs x). This can be equivalently written as $th_\theta[\bar{\Phi}]$ for $\theta = \frac{1}{2} + (\alpha - \frac{1}{2})\kappa$, where $\theta \in [0,1]$ as $0 < \kappa \le 1$ and $0 \le \alpha \le 1$. □

Thus, we get that any hardcore measure has a model of optimal density. In order to use this implication, we need to argue the existence of some hardcore measure μ_0. Note that μ_0 can be of any small density, which makes it easier to argue its existence. Below we show that there is hardcore measure μ_0 on a two-element set.

4.3.2 Two-element hardcore measure

Given a Boolean function g that is $(\delta, \mathrm{TH}_\lambda[\mathcal{F}])$-hard on U with respect to a distribution σ, our goal is to construct a 2δ-dense $(\epsilon, \mathcal{F})$-hardcore measure for g over $\mathcal{U} = (U, \sigma)$. To this end, we extend U with two new elements, extend g and $\mathcal{F}$ to the new universe, extend the distribution σ to the new universe, and argue that the set of two new elements is a hardcore for the extension of g and that the extension of g is still δ-hard. We then apply our argument from the previous subsection, to get a "large" hardcore measure for the extension of g. Finally, we argue that the restriction of this hardcore measure to the original universe U is hardcore for the original function g of almost optimal density 2δ, and can be easily made to be of density 2δ.

In more detail, we extend U with two new elements x_1 and x_2. Let $X = \{x_1, x_2\}$, and let $U_* = U \cup X$. Let $x_0 \in U$ be some fixed element. We define the extension g_* of g as follows:

for $x \in U$, set $g_*(x) = g(x)$, and define $g_*(x_1) = 1 - g_*(x_2) = g(x_0)$. Extend each $f \in \mathcal{F}$ to U_* by defining $f_*(x) = f(x)$ for $x \in U$, and defining $f_*(x_1) = f_*(x_2) = f(x_0)$. Let $\mathcal{F}_*$ be the corresponding class of extended tests. Note that our definition of $\mathcal{F}_*$ ensures that $\mathcal{F}_*$ is indeed closed under complement, and so is still a class. Also note that for every $h \in \mathrm{TH}_\lambda[\mathcal{F}_*]$, we have $h(x_1) = h(x_2) = h(x_0)$, while $g(x_1) \neq g(x_2)$. So, each such h agrees with g_* on exactly one of the two elements in X.

Define a distribution σ_* on U_* as σ on U with probability $(1 - \epsilon\delta)$ and as the uniform distribution on X with probability $\epsilon\delta$. Let $\mathcal{U}_* = (U_*, \sigma_*)$. Define $\mu_0(x_1) = \mu_0(x_2) = 1$, and $\mu_0(x) = 0$ for all $x \in U$.

Claim 4.2. *The Boolean function g_* is $(\delta + \epsilon\delta(\frac{1}{2} - \delta), \mathrm{TH}_\lambda[\mathcal{F}_*])$-hard over $\mathcal{U}_*$, and the measure μ_0 is $(0, \mathcal{F}_*)$-hardcore for g_* over $\mathcal{U}_*$.*

Proof. First we argue hardness of g_*. Since for all $h \in \mathrm{TH}_\lambda[\mathcal{F}_*]$, $|g_*(x_1) - h(x_1)| = 1 - |g_*(x_2) - h(x_2)|$, we get $\mathbb{E}_X[|g_* - h|] = \frac{1}{2}$ (for the uniform distribution over X). It follows that $\mathbb{E}_{\sigma_*}[|g_* - h|] = (1 - \epsilon\delta) \cdot \mathbb{E}_\sigma[|g - h|] + \epsilon\delta \cdot \mathbb{E}_X[|g_* - h|] \geq (1 - \epsilon\delta)\delta + \epsilon\delta/2$.

Next, to argue that μ_0 is a $(0, \mathcal{F}_*)$-hardcore, note that $(\mu_0)_{\sigma_*}$ is the uniform distribution over X, and that, for every $f_* \in \mathcal{F}_*$, $\mathbb{E}_X[|g_* - f_*|] = \frac{1}{2} \cdot (|g(x_0) - f(x_0)| + |1 - g(x_0) - f(x_0)|) = \frac{1}{2} \cdot (|g(x_0) - f(x_0)| + 1 - |g(x_0) - f(x_0)|) = \frac{1}{2}$, where we used the identity $|1 - g - f| = 1 - |g - f|$ (true for Boolean g and any bounded f). □

Let $\delta' = \delta + \epsilon\delta(\frac{1}{2} - \delta)$. By remark 2.12, we may assume that $\delta < 1/2 - \epsilon$ (since otherwise we get that the constant 1 measure is already an $(\epsilon, \mathcal{F})$-hardcore for g). Thus, we have that $\delta' > \delta$.

By the previous subsection, we get a $2\delta'$-dense $(\epsilon, \mathcal{F}_*)$-hardcore measure μ_* for g_* over $\mathcal{U}_*$. Define the measure μ over U as the restriction of μ_* to U, i..e, for all $x \in U$, set $\mu(x) = \mu_*(x)$. We argue next that μ is a large hardcore measure for g over $\mathcal{U}$.

Claim 4.3. *μ is $((9/8)\epsilon, \mathcal{F})$-hardcore for g over $\mathcal{U}$, and $d_\sigma(\mu) \geq (2\delta' - \epsilon\delta)/(1 - \epsilon\delta)$.*

Proof. We have $2\delta' \leq d_{\sigma_*}(\mu_*) = \mathbb{E}_{\sigma_*}[\mu_*] = (1 - \epsilon\delta) \cdot \mathbb{E}_\sigma[\mu] + \epsilon\delta \cdot \mathbb{E}_X[\mu_*] \leq (1 - \epsilon\delta) \cdot d_\sigma(\mu) + \epsilon\delta$, which yields the required density lower bound.

Next, for every $f_* \in \mathcal{F}_*$, we have $\frac{1}{2} - \epsilon \leq \mathbb{E}_{(\mu_*)_{\sigma_*}}[|g_* - f|] = \frac{1}{d_{\sigma_*}(\mu_*)} \cdot \mathbb{E}_{\sigma_*}[|g_* - f_*| \cdot \mu_*]$. The latter is at most

$$\frac{1}{2\delta'} \cdot ((1 - \epsilon\delta) \cdot \mathbb{E}_\sigma[|g - f| \cdot \mu] + \epsilon\delta \cdot \mathbb{E}_X[|g_* - f_*| \cdot \mu_*]) \leq \frac{\mathbb{E}_\sigma[|g - f| \cdot \mu]}{d_\sigma(\mu)} \cdot \frac{d_\sigma(\mu)}{2\delta'} + \frac{\epsilon\delta}{4\delta'}.$$

We may assume that $d_\sigma(\mu) \leq 2\delta'$, since we can always scale μ down. Thus, we can upperbound the right-hand side of the last equation by $\mathbb{E}_{\mu_\sigma}[|g - f|] + \epsilon\delta/(4\delta')$. Using the definition of δ', we can upperbound the last expression by $\mathbb{E}_{\mu_\sigma}[|g - f|] + \epsilon/8$. It follows that $\frac{1}{2} - \epsilon \leq \mathbb{E}_{\mu_\sigma}[|g - f|] + \epsilon/8$, implying by Remark 2.12 that μ is a $((9/8)\epsilon, \mathcal{F})$-hardcore for g over $\mathcal{U}$. □

Finally, if $d_\sigma(\mu) < 2\delta'$, we modify μ as follows: Define the measure $\mu' = (1 - \epsilon\delta)\mu + \epsilon\delta$ (it is easy to see that $0 \leq \mu'(x) \leq 1$ for all $x \in U$). We have the following.

Claim 4.4. *The measure μ' is 2δ-dense $((13/8)\epsilon, \mathcal{F})$-hardcore for g over $\mathcal{U}$.*

Proof. It is easy to see that $d_\sigma(\mu') = (1 - \epsilon\delta) \cdot d_\sigma(\mu) + \epsilon\delta \geq 2\delta' > 2\delta$. It is also easy to see that $d_\sigma(\mu') \geq d_\sigma(\mu)$.

To argue that μ' is a hardcore, we upperbound $\mathbb{E}_{(\mu')_\sigma}[|g - f|] = \frac{1}{d_\sigma(\mu')} \cdot \mathbb{E}_\sigma[|g - f| \cdot \mu']$ by

$$\frac{(1 - \epsilon\delta)\mathbb{E}_\sigma[|g - f| \cdot \mu]}{d_\sigma(\mu')} + \frac{\epsilon\delta\mathbb{E}_\sigma[|g - f|]}{d_\sigma(\mu')} \leq \frac{\mathbb{E}_\sigma[|g - f| \cdot \mu]}{d_\sigma(\mu)} + \frac{\epsilon\delta}{2\delta} = \mathbb{E}_{\mu_\sigma}[|g - f|] + \epsilon/2.$$

Since μ is a $((9/8)\epsilon, \mathcal{F})$-hardcore, we conclude (using Remark 2.12) that μ' is $((13/8)\epsilon, \mathcal{F})$-hardcore for g over $\mathcal{U}$, as required. □

5 On-Line Learning Algorithm

Here we derive an algorithmic version of the MinMax formulation of the Dense Model Theorem, using the online-learning framework of [FS99], adapted to the setting of dense models by [BHK09, Zha11]. The learning algorithm uses Bregman projections, which we discuss next.

5.1 Bregman projections

Given a measure μ over (U, σ) of density greater than δ, we can always get a measure μ' of density exactly δ by setting $\mu' = \gamma \cdot \mu$, for some constant $0 < \gamma < 1$. Such scaling doesn't change the induced distribution μ_σ, and so μ' stays a model if μ was a model. What if we have a measure μ of density *less* than δ, and we want to get from μ a "closely related" measure of density δ? A natural idea would be to define $\mu'' = c \cdot \mu$ for some constant $c > 1$. The problem is that $c \cdot \mu$ may not be $[0,1]$-bounded. So it is natural to define $\mu'' = \mathit{trunc}[c \cdot \mu]$, for the smallest constant $c \geq 1$ such that μ'' has density δ if such a c exists. It turns out that the measure μ'' defined this way is exactly the result of "projecting" the measure μ onto the set of measures of density at least δ, using the notion of *Bregman projection*. For the notion of "distance" between measures (and so also between a measure and its Bregman projection), it is convenient to use the *Kullback-Leibler divergence*. We define these notions next.

For measures μ and ν over a probability space $\mathcal{U} = (U, \sigma)$, the *Kullback-Leibler divergence between μ and ν over $\mathcal{U}$* is $\mathrm{D}(\mu||\nu) = \mathbb{E}_{x \in \sigma}[\mu(x) \cdot \log(\mu(x)/\nu(x))] - \mathbb{E}_{x \in \sigma}[\mu(x)] + \mathbb{E}_{x \in \sigma}[\nu(x)]$.

Definition 5.1 (Bregman projection)**.** Let $\emptyset \neq \Gamma \subseteq [0,1]^N$ be any closed convex set of measures, and let $\nu \in [0,1]^N$ be any measure. The *Bregman projection* of ν onto Γ, denoted by $P_\Gamma \nu$, is $P_\Gamma \nu = \arg\min_{\gamma \in \Gamma} \mathrm{D}(\gamma||\nu)$.

Note that, for $\nu \in \Gamma$, we have $P_\Gamma \nu = \nu$ (since $\mathrm{D}(\gamma||\nu) \geq 0$, and $\mathrm{D}(\gamma||\nu) = 0$ iff $\gamma = \nu$).

For $\delta \in [0,1]$, we denote by $P_\delta \nu$ the Bregman projection of a measure ν onto the set $\Gamma_\delta = \{\nu \in [0,1]^N \mid d_\sigma(\nu) \geq \delta\}$ of δ-dense measures; note that Γ_δ is a convex and closed set. As shown by [BHK09, Lemma 3.1], the projection P_δ has the following nice form. For a measure ν over $\mathcal{U}$ with $d_\sigma(\textsc{support}[\nu]) \geq \delta$, for some $\delta \in [0,1]$, we have $P_\delta \nu = \mathit{trunc}[c \cdot \nu]$ for the smallest constant $c \geq 1$ such that the measure $\mathit{trunc}[c \cdot \nu]$ has density δ in $\mathcal{U}$.

We use the divergence as a "distance" between measures. The crucial property of Bregman projections is given by the following theorem of Bregman.[6] Intuitively, it says that if a measure $\gamma \in \Gamma$, for some closed convex set Γ, is close to a measure ν (which may be outside of Γ), then γ is as close (or even closer) to the projection $P_\Gamma \nu$ of ν onto Γ.

Theorem 5.2 ([Bre67])**.** *Let Γ be any non-empty closed convex set of measures, and let γ, ν be measures, with $\gamma \in \Gamma$. Then $\mathrm{D}(\gamma||P_\Gamma \nu) + \mathrm{D}(P_\Gamma \nu||\nu) \leq \mathrm{D}(\gamma||\nu)$, and hence, $\mathrm{D}(\gamma||P_\Gamma \nu) \leq \mathrm{D}(\gamma||\nu)$.*

Bregman projections can be relaxed to *approximate* Bregman projections:

Definition 5.3 (approximate Bregman projections)**.** For a non-empty closed convex set Γ of measures and a parameter $\alpha \geq 0$, a measure $\nu^* \in \Gamma$ is called an *α-approximate Bregman projection* of a measure ν onto Γ, denoted by α-$P_\Gamma \nu$, if for all $\gamma \in \Gamma$, $\mathrm{D}(\gamma||\nu^*) \leq \mathrm{D}(\gamma||P_\Gamma \nu) + \alpha$.

Approximate Bregman projections onto Γ_δ are efficiently computable:

Lemma 5.4 ([BHK09])**.** *For $\delta \in [0,1]$, let ν be a measure over $\mathcal{U} = (U, \sigma)$ such that $P_\delta \nu = \mathit{trunc}[c \cdot \nu]$ for $c \in [1, 1+\zeta]$, where $\zeta > 0$. Suppose we have oracle access to ν, and can sample an element from $\mathcal{U}$ in time t. Then, for any $p, \epsilon \in (0,1)$, we can compute an implicitly represented approximate projection $(\epsilon\delta)$-$P_\delta \nu$ in time $O\left(t\delta^{-1}\epsilon^{-2}(\log\log\zeta\epsilon^{-1} + \log p^{-1})\right)$, with probability $1-p$. Moreover, the computed approximate projection has the form $\mathit{trunc}[\tilde{c} \cdot \nu]$, for some $\tilde{c} \in [1, 1+\zeta]$.*

[6] D is a Bregman divergence associated with a generalized entropy function $\textsc{Ent}_\sigma(\mu) = -\mathbb{E}_\sigma[\mu \cdot \log \mu]$, which can be shown to be a Bregman function; see [CZ97] for details.

Finally, we get the following generalization of the Total Loss Lemma due to[BHK09] for the general probability space $\mathcal{U} = (U, \sigma)$. The lemma captures the setting of a two-player game where one of the players is starting with some measure μ^1, and in each round t updates her current measure to the new measure μ^{t+1} by taking into account the "penalty" function m^t received from the other player; the update is through a simple multiplicative procedure, followed by an approximate Bregman projection onto the desired set Γ of measures. Intuitively, the lemma says that the total loss incurred in the game from using measures μ^t is not much worse than the total loss from using an *arbitrary single* measure μ in every round t of the game, where μ could depend on all the penalty functions m^t played in the game.

The proof of the following lemma is identical to that of Lemma 4.1 in [BHK09].[7]

Lemma 5.5 (Total Loss Lemma [BHK09])**.** *For a probability space $\mathcal{U} = (U, \sigma)$, let Γ be a non-empty closed convex set of measures over U. Let $\epsilon \in (0, \frac{1}{2})$ and let $T \in \mathbb{N}$ be arbitrary. Let $\mu^1 \in \Gamma$ be an arbitrary measure over U, and, for $1 \leq t < T$, let $m^t : U \to [0, 1]$ be an arbitrary function ("penalty"). Define, for each $1 \leq t < T$, the measures $\nu^{t+1}(x) = (1-\epsilon)^{m^t(x)} \cdot \mu^t(x)$ and $\mu^{t+1} = \alpha\text{-}P_\Gamma \nu^{t+1}$. Then, for every measure $\mu \in \Gamma$, we have*

$$\sum_{t=1}^{T} \mathbb{E}_\sigma(\mu^t \cdot m^t) - \tfrac{\alpha}{\epsilon} \cdot T \leq (1+\epsilon) \cdot \sum_{t=1}^{T} \mathbb{E}_\sigma(\mu \cdot m^t) + \tfrac{1}{\epsilon} \cdot \mathrm{D}(\mu || \mu^1). \tag{7}$$

5.2 DMT algorithm

Now we state the following online-learning algorithm **OLL** which is essentially due to Zhang [Zha11], based on [FS99, BHK09]. We generalize the algorithm to arbitrary finite probability space $\mathcal{U} = (U, \sigma)$.

The algorithm is given as input a distribution ρ over U, a class $\mathcal{F}$ of tests $f : U \to [0, 1]$, and parameters $0 < \epsilon, \delta < 1$. Starting with the constant δ-dense measure μ^1 such that $\mu^1(x) = \delta$ for all $x \in U$, the algorithm iterates doing multiplicative updates and approximate Bregman projections to get a new δ-dense measure μ^{t+1} from the current measure μ^t. The algorithm uses as "penalty" functions m^t the tests f which witness that μ^t is not a model for ρ.

Algorithm 1 OLL$[\rho, \mathcal{F}, \delta, \epsilon]$ (Generalized DMT-ALGORITHM[Zha11])

$t = 1$, define the measure μ^1 by $\mu^1(x) = \delta$, for each $x \in U$
$T = \frac{16}{\epsilon^2} \log(\frac{1}{\delta})$, $\alpha = \frac{\epsilon^2 \delta}{16}$
while $t < T$ **do**
  if $\exists f \in \mathcal{F}$ such that $f[(\mu^t)_\sigma] - f[\rho] > \epsilon$ **then**
    $m^t = f$ for some $f \in \mathcal{F}$ satisfying $f[(\mu^t)_\sigma] - f[\rho] > \epsilon$
    $\nu^{t+1}(x) = (1 - \frac{\epsilon}{4})^{m^t(x)} \mu^t(x)$
    $\mu^{t+1} = \alpha\text{-}\ P_\delta \nu^{t+1}$
    $t = t + 1$
  else
    return MODEL, μ^t
  end if
end while
return DISTINGUISHER, $F = \frac{1}{T} \sum_{t=1}^{T} m^t$

After at most $T = O(\epsilon^{-2} \log \delta^{-1})$ iterations, either we get, for some $1 \leq t \leq T$, a measure μ^t that is a δ-dense model for the distribution ρ, or we get that the average F of all penalty functions m^t is a "universal" distinguisher in the sense that the same test F witnesses that *every* δ-dense measure is not a model for ρ. More precisely, we have the following.

[7] [BHK09] had a typo: the factor $1/\epsilon$ was missing in the last term on the right-hand side of Eq. (7).

Theorem 5.6 (Analysis of **OLL**). *For a finite probability space $(\mathcal{U}, \sigma)$, a class $\mathcal{F}$, a distribution ρ over U, and parameters $\epsilon, \delta \in (0,1)$, the algorithm* **OLL** *finds*

1. *either a δ-dense $(\epsilon, \mathcal{F})$-model μ^t for ρ, for some $1 \leq t \leq T$,*
2. *or a test $F : U \to [0,1]$ such that $F[\mu_\sigma] - F[\rho] > \epsilon/4$ for every measure μ with $d_\sigma(\mu) = \delta$ (thereby witnessing that ρ has no δ-dense $(\epsilon/4, \mathcal{F})$-model).*

Proof. Observe that the algorithm **OLL** outputs "MODEL" only if, for some μ^t, it holds that $f[(\mu^t)_\sigma] - f[\rho] \leq \epsilon$ for every $f \in \mathcal{F}$. Since the class $\mathcal{F}$ is closed under negation, the same inequality holds also for $\bar{f}$, which implies that $f[\rho] - f[(\mu^t)_\sigma] \leq \epsilon$ for every $f \in \mathcal{F}$ as well. So we get that $|f[(\mu^t)_\sigma - f[\rho]| \leq \epsilon$ for all $f \in \mathcal{F}$, and hence, μ^t is a model for ρ. By construction, all measures μ^j, $1 \leq j \leq T$, are δ-dense, and so μ^t is a δ-dense $(\epsilon, \mathcal{F})$-model for ρ.

Now suppose that the algorithm **OLL** outputs "DISTINGUISHER". We will show that in this case the function F constructed by the algorithm is such that, for every measure μ over U with $d_\sigma(\mu) = \delta$,

$$F[\mu_\sigma] - F[\rho] > \tfrac{\epsilon}{4}. \tag{8}$$

Observe that F has complexity T relative to $\mathcal{F}$ as $F = \frac{1}{T}\sum_{t=1}^{T} m^t$ for $m^t \in \mathcal{F}$. Using Eq. (8) and the fact that F is an average of tests m^t from $\mathcal{F}$, we get that some m^t (with t depending upon μ) is a witness to μ not being an $(\epsilon/4, \mathcal{F})$-model for ρ. Hence, ρ has no δ-dense $(\epsilon/4, \mathcal{F})$-model.

Now we prove Eq. (8). First, by the construction of F, we have

$$\sum_T \mathbb{E}_{(\mu^t)_\sigma}[m^t] > \sum_T \mathbb{E}_\rho[m^t] + \epsilon T, \tag{9}$$

where $\sum_T$ denotes the summation $\sum_{t=1}^{T}$. By the Total Loss Lemma, Lemma 5.5, we get

$$\sum_T \mathbb{E}_\sigma[\mu^t \cdot m^t] \leq (1 + \tfrac{\epsilon}{4}) \sum_T \mathbb{E}_\sigma[\mu \cdot m^t] + \tfrac{4}{\epsilon}\mathrm{D}(\mu||\mu^1) + \tfrac{4\alpha}{\epsilon}T. \tag{10}$$

We have $\mathbb{E}_\sigma[\mu^t \cdot m^t] = d_\sigma(\mu^t) \cdot \mathbb{E}_{(\mu^t)_\sigma}[m^t]$, and hence, using the fact that $d_\sigma(\mu^t) \geq \delta$ for all $1 \leq t \leq T$ and Eq. (9), we conclude that

$$\sum_T \mathbb{E}_\sigma[\mu^t \cdot m^t] \geq \delta \cdot \sum_T \mathbb{E}_{(\mu^t)_\sigma}[m^t] > \delta \left(\sum_T \mathbb{E}_\rho[m^t] + \epsilon \cdot T \right) = \delta \cdot T \cdot (\mathbb{E}_\rho[F] + \epsilon)\,. \tag{11}$$

Similarly, we have that $\mathbb{E}_\sigma[\mu \cdot m^t] = d_\sigma(\mu) \cdot \mathbb{E}_{\mu_\sigma}[m^t] = \delta \cdot \mathbb{E}_{\mu_\sigma}[m^t]$, and so,

$$\sum_T \mathbb{E}_\sigma[\mu \cdot m^t] = \delta \cdot T \cdot \mathbb{E}_{\mu_\sigma}[F]. \tag{12}$$

We also have

$$\mathrm{D}(\mu||\mu^1) = \mathbb{E}_\sigma\left[\mu \cdot \log\left(\tfrac{\mu}{\mu^1}\right)\right] + d_\sigma(\mu_1) - d_\sigma(\mu) \leq \delta \log \delta^{-1}, \tag{13}$$

using the fact that $\mu^1(x) = \delta$ for all $x \in U$, and hence $d_\sigma(\mu^1) = \delta = d_\sigma(\mu)$, as well as the inequality $\mu(x) \cdot \log(\mu(x)/\delta) \leq \mu(x) \cdot \log(1/\delta)$, which follows from $0 \leq \mu(x) \leq 1$ and $0 \log 0 = 0$.

Using Eqs. (11)–(13) inside Eq. (10), we get

$$\delta T \cdot (\mathbb{E}_\rho[F] + \epsilon) < \left(1 + \frac{\epsilon}{4}\right) \delta T \cdot \mathbb{E}_{\mu_\sigma}[F] + \frac{4}{\epsilon}\delta \log \delta^{-1} + \frac{4\alpha}{\epsilon} \cdot T. \tag{14}$$

Dividing both sides of Eq. (14) by δT, and using the definition of $T = 16\epsilon^{-2} \log \delta^{-1}$ and $\alpha = \epsilon^2\delta/16$, we get $\mathbb{E}_{\mu_\sigma}[F] - \mathbb{E}_\rho[F] > \epsilon - (\epsilon/4) \cdot \mathbb{E}_{\mu_\sigma}[F] - (\epsilon/4) - (\epsilon/4)$. The latter is at least $\epsilon/4$, since $\mathbb{E}_{\mu_\sigma}[F] \leq 1$. □

Remark 5.7. *Consider a two-player zero-sum game where the first player chooses $f \in \mathcal{F}$ and the second player chooses a δ-dense measure μ, with the payoff for the first player given by $f[\mu_\sigma] - f[\rho]$. Then* **OLL**$[\rho, \mathcal{F}, \delta, \epsilon]$ *returns either a mixed strategy for the second player which does well against every strategy of first player, or vice versa for the first player. Thus* **OLL** *finds approximately optimal mixed strategies for this class of zero-sum games between δ-dense measures and algorithms $\mathcal{F}$. In fact, the algorithm* **OLL** *is the Multiplicative Weights Algorithm of Freund and Schapire [FS99] for approximately solving two-player zero-sum games, combined with taking Bregman projections (as in the Smooth Boosting Algorithm of Kale[Kal07] and Barak et al. [BHK09]); the projections onto the set Γ_δ of δ-dense measures are taken to ensure that the strategy of the second player at each stage is a δ-dense measure.*

5.3 Constructiveness of Algorithm OLL

The function F computed in the "DISTINGUISHER" branch of the algorithm **OLL** clearly satisfies the conditions of the Constructive MinMax version of Dense Model Theorem with $\lambda = 16\epsilon^{-2} \log \delta^{-1}$.

The complexity of the measures μ^j, for $1 < j \leq T$, produced by the algorithm is $O(j)$. Indeed, μ^1 has complexity 1. For $j > 1$, ν^j is obtained from μ^{j-1} using one exponentiation and multiplication. Finally, the approximate Bregman projection of ν^j requires two extra operations: scalar multiplication and truncation, by Lemma 5.4. Overall, μ^j is obtained from μ^{j-1} using a constant number of operations. So, the complexity of model μ^t returned by the algorithm **OLL** in its "MODEL" branch is at most $O(t) \leq O(\lambda)$. Thus, we have:

Corollary 5.8. *Algorithm* **OLL** *provides the constructive MinMax version of DMT for $\lambda(\epsilon, \delta) = O(\epsilon^{-2} \log \delta^{-1})$ and $c = 4$.*

5.4 Avoiding exponentiation

As observed earlier, if $\mathcal{F}$ is a Boolean class, the operation of limited exponentiation is not needed. In general, when $\mathcal{F}$ is a class of $[0,1]$-bounded functions, we can use thresholding to get from $f \in \mathcal{F}$ satisfying the condition in the **if**-statement of the algorithm **OLL** a *Boolean* function $f' = th_\theta(f)$, for some $\theta \in [0,1]$, such that f' satisfies the same condition.

Indeed, first it is easy to see that, for any function $g : U \to [0,1]$, we have for every $x \in U$, $g(x) = \mathbb{E}_{\theta \in [0,1]}[th_\theta[g(x)]]$. Hence, for any g and any distribution π over U, we get $\mathbb{E}_\pi[g] = \mathbb{E}_{\theta \in [0,1]}[\mathbb{E}_\pi[th_\theta[g]]]$. Finally, using this equality, the linearity of expectation, and averaging, we conclude that if there is some f such that $f[(\mu^t)_\sigma] - f[\rho] > \epsilon$, then there is also some $\theta \in [0,1]$, such that, for $f' = th_\theta[f]$, we have $f'[(\mu^t)_\sigma] - f'[\rho] > \epsilon$.

It follows that for an arbitrary class $\mathcal{F}$, we either get a model μ^t of low complexity *without using exponentiation*, or get a universal distinguisher F that is the average of few *thresholded* functions from $\mathcal{F}$. Thus, we can trade the simplicity of the model for the extra complexity of the universal distinguisher.

5.4.1 Algorithmic issues

Given access to an oracle which produces a function $f \in \mathcal{F}$ witnessing that a current measure μ^t is not an $(\epsilon, \mathcal{F})$-model for ρ, the algorithm **OLL** will *efficiently construct* either a dense model μ for ρ or a universal distinguisher F. Thus, we actually get an *algorithmic* constructive MinMax version of DMT.

In case one wishes to avoid the exponentiation operation following the "thresholding" approach outlined above, one needs to find appropriate thresholds θ efficiently. This can be done by random sampling, with some slight loss in parameters. We explain this point in more detail next.

Consider the following modification **OLL**$'$ of the algorithm **OLL**.

Algorithm 2 OLL′$[\rho, \mathcal{F}, \delta, \epsilon]$ (Modified DMT-ALGORITHM)

$t = 1$, define the measure μ^1 by $\mu^1(x) = \delta$, for each $x \in U$
$T = \frac{256}{\epsilon^2} \log(\frac{1}{\delta})$, $\alpha = \frac{\epsilon^2 \delta}{256}$
while $t < T$ **do**
  if $\exists f \in \mathcal{F}$ such that $f[(\mu^t)_\sigma] - f[\rho] > \epsilon$ **then**
    let $f \in \mathcal{F}$ be any function such that $f[(\mu^t)_\sigma] - f[\rho] > \epsilon$
    for each $0 \leq n \leq \lceil 2\epsilon^{-1} \rceil$, set $f_n = th_{n\epsilon/2}(f)$
    let $0 \leq n^* \leq \lceil 2\epsilon^{-1} \rceil$ be such that, with high probability, $f_{n^*}[(\mu^t)_\sigma] - f_{n^*}[\rho] > \epsilon/4$
    $m^t = f_{n^*}$
    $\nu^{t+1}(x) = (1 - \frac{\epsilon}{16})^{m^t(x)} \mu^t(x)$
    $\mu^{t+1} = \alpha$- $P_\delta \nu^{t+1}$
    $t = t + 1$
  else
    return MODEL, μ^t
  end if
end while
return DISTINGUISHER, $F = \frac{1}{T} \sum_{t=1}^{T} m^t$

The estimation of $f_n[\rho]$ and $f_n[(\mu^t)_\sigma]$ is done by random sampling; by Chernoff bounds, we can achieve the additive error at most $\epsilon/8$ for each, with high probability, in time $poly(\epsilon^{-1})$. Assuming that the required integer value n^* exists (and hence can be efficiently found by random sampling), the rest of the algorithm is the same as before, and so the old analysis of Theorem 5.6 applies.

It remains to argue that n^* always exists. To this end, we shall need the following basic properties of the threshold operation.

Lemma 5.9. *Let $g : U \to [0, 1]$ be any function, and let ρ and τ be any distributions over U. Suppose $\mathbb{E}_\rho[g] > \mathbb{E}_\tau[g] + \epsilon$ for some $\epsilon \in [0, 1]$. Then there exist θ, n satisfying:*

1. *[RTTV08] $\theta \in [\epsilon/2, 1]$ and $\mathbb{E}_\rho[th_\theta[g]] > \mathbb{E}_\tau[th_{\theta-\epsilon/2}[g]] + \epsilon/2$.*
2. *$n \in \mathbb{N}, n \leq \lceil 2\epsilon^{-1} \rceil$ such that $\mathbb{E}_\rho[th_{n\epsilon/2}[g]] > \mathbb{E}_\tau[th_{n\epsilon/2}[g]] + \epsilon/2$.*

Proof. For item (1), suppose it fails, i.e., for every $\theta \in [\epsilon/2, 1]$, $\mathbb{E}_\rho[th_\theta[g]] \leq \mathbb{E}_\tau[th_{\theta-\epsilon/2}[g]] + \epsilon/2$. Using the fact that $th_t[g(x)]$ as a function of t with x fixed is piece-wise constant with only discontinuity at $g(x) = t$, we get $\mathbb{E}_\rho[g] = \mathbb{E}_\theta[\mathbb{E}_\rho[th_\theta[g]]] = \int_0^{\epsilon/2} \mathbb{E}_\rho[th_\theta[g]]dt + \int_{\epsilon/2}^1 \mathbb{E}_\rho[th_\theta[g]]dt \leq \epsilon/2 + \int_0^1 (\mathbb{E}_\tau[th_\theta[g]] + \epsilon/2)\, dt$, yielding $\mathbb{E}_\rho[g] \leq \mathbb{E}_{\theta\in[0,1]}[\mathbb{E}_\tau[th_\theta[g]]] + \epsilon = \mathbb{E}_\tau[g] + \epsilon$. A contradiction.

Finally, θ in item (1) must satisfy $\theta \in [n\epsilon/2, (n+1)\epsilon/2]$ for some $n \in \mathbb{N}$ with $n \leq 2\lceil \epsilon^{-1} \rceil$. It follows that $n\epsilon/2 \in [\theta - \epsilon/2, \theta]$. By item (1), we get $\mathbb{E}_\rho[th_{n\epsilon/2}[g]] \geq \mathbb{E}_\rho[th_\theta[g]] > \mathbb{E}_\tau[th_{\theta-\epsilon/2}[g]] + \epsilon/2 \geq \mathbb{E}_\tau[th_{n\epsilon/2}[g]] + \epsilon/2$, where we used the fact that $th_\alpha[v] \geq th_\beta[v]$ for any $v, \alpha, \beta \in [0, 1]$ such that $\alpha \leq \beta$. □

The existence of n^* required by the modified algorithm **OLL′** follows from item (2) of Lemma 5.9. Using our previous analysis of algorithm **OLL**, we get the following analysis of the modified algorithm **OLL′**.

Theorem 5.10 (Analysis of **OLL′**)**.** *For a finite probability space $(\mathcal{U}, \sigma)$, a class $\mathcal{F}$, a distribution ρ over U, and parameters $\epsilon, \delta \in (0, 1)$, the algorithm **OLL′** finds*

1. *either a δ-dense $(\epsilon, \mathcal{F})$-model $\mu^t \in \mathcal{F}_{O(t)}$ for ρ, for some $1 \leq t \leq T$,*
2. *or a test $F : U \to [0, 1]$ such that $F[\mu_\sigma] - F[\rho] > \epsilon/16$ for every measure μ with $d_\sigma(\mu) = \delta$, where $F = \frac{1}{T} \sum_{t=1}^{T} m^t$, with m^t a thresholded function from $\mathcal{F}$.*

References

[BHK09] B. Barak, M. Hardt, and S. Kale. The uniform hardcore lemma via approximate Bregman projections. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1193–1200, 2009.

[Bre67] L.M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967.

[CZ97] Y. Censor and S.A. Zenios. *Parallel optimization: theory, algorithms, and applications*. Numerical mathematics and scientific computation. Oxford University Press, 1997.

[FS99] Y. Freund and R.E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29:79–103, 1999.

[Hol05] T. Holenstein. Key agreement from weak bit agreement. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, pages 664–673, 2005.

[Kal07] S. Kale. Boosting and hard-core set constructions: a simplified approach. *Electronic Colloquium on Computational Complexity*, 14(131), 2007.

[RTTV08] O. Reingold, L. Trevisan, M. Tulsiani, and S.P. Vadhan. Dense subsets of pseudorandom sets. In *Proceedings of the Forty-Ninth Annual IEEE Symposium on Foundations of Computer Science*, pages 76–85, 2008.

[TTV09] L. Trevisan, M. Tulsiani, and S.P. Vadhan. Regularity, boosting, and efficiently simulating every high-entropy distribution. In *Proceedings of the Twenty-Fourth Annual IEEE Conference on Computational Complexity*, pages 126–136, 2009.

[Zha11] J. Zhang. On the query complexity for showing dense model. *Electronic Colloquium on Computational Complexity*, 18:38, 2011.