$\mathbf{APX_1}$: A Theory for Probabilistic Polynomial-Time Reasoning

Lijie Chen Jiatu Li **Igor Oliveira** Ryan Williams

Bounded Arithmetic Day @ Warwick

March 6, 2026

(**Warning:** Some parts of this talk assume basic familiarity with bounded arithmetic!)

Overview

- **Bounded arithmetic** extends complexity theory by capturing not only the computational resources required by algorithms, but also the complexity of proving their correctness.
- $\mathbf{PV}/\mathbf{PV_1}$ (Cook, 1975) is a robust theory for **deterministic polynomial-time reasoning**.
- Many TCS proofs rely on **probabilities** (expectation, concentration, probabilistic method, tail bounds, ...).
- Existing probabilistic framework $\mathbf{APC_1}$ (Jeřábek, 2007) uses a powerful counting principle (**dWPHP**), **possibly stronger than needed**.
- We propose theory $\mathbf{APX_1}$: a weaker bounded arithmetic theory closer to **"probabilistic polynomial-time reasoning"**.

Why weaken APC_1?

- **Philosophical:** Want probabilities but dWPHP axiom might not be "feasible" **[ILW'23]**: Given $C\colon \{0,1\}^n \to \{0,1\}^{n+1}$, how to find/certify y such that, for all x, $C(x) \neq y$?
- **Unprovability of complexity-theoretic conjectures:** Want an appropriate theory, but not more: strong axioms complicate unprovability arguments.
- **Reverse mathematics of TCS:** understand minimal axioms required to prove theorems. APC_1 is an overly powerful base theory for correspondences weaker than dWPHP.
- **Proof complexity of derandomization:** Can we formulate and study the feasible provability of prBPP = prP?

Comparing PV_1, APX_1, and APC_1

- $\mathbf{PV_1}$ formalizes polynomial-time functions + induction for feasible predicates.
- $\mathbf{APC_1} = \mathbf{PV_1} + \mathbf{dWPHP}(\mathrm{PV})$, provides approx. counting and probabilistic reasoning.
- **dWPHP** plays two roles:
  - (i) enables approximate counting (via hardness & NW PRG formalization);
  - (ii) acts as a strong counting/combinatorial principle to prove probability inequalities.

Key Design Principle for $\mathbf{APX_1}$: Approximate counting as a central primitive (rather than deriving it from a stronger pigeonhole principle).

APX_1 in one slide: The approximate counting oracle P

- Extend the language of PV with a new function symbol $P(C, \Delta)$.
- Inputs: Boolean circuit $C : \{0,1\}^n \to \{0,1\}$ and a precision parameter Δ.
- Intended meaning: $P(C, \Delta)$ provides a rational approximation to

$$p(C) \triangleq \mathbb{P}_{x \leftarrow \{0,1\}^n}[C(x) = 1]$$

  within additive error $\delta \approx 1/|\Delta|$.

  (For convenience, we often write $P_\delta(C)$ for $P(C, \Delta)$, where $\delta = 1/|\Delta|$.)

- Add **"simple axioms"** that force P_δ to behave like "counting over the hypercube" **(?)**

Axioms for P_δ

Axioms are universal PV(P)-equations; $\beta^{-1} \in \mathrm{Log}$ is a "slack parameter" (think of it as $o(1)$).

1. **Basic Axiom:** $P_\delta(C)$ is a rational in $[0,1]$ (with feasibility/output-length bounds).
2. **Boundary Axiom:** if C is syntactically constant then $P_\delta(C) \in \{0,1\}$ and matches C.
3. **Precision Consistency:** for any precision parameters $\delta_1^{-1}, \delta_2^{-1}, \beta^{-1} \in \mathrm{Log}$,

$$|P_{\delta_1}(C) - P_{\delta_2}(C)| \leq \delta_1 + \delta_2 + \beta.$$

4. **Local Consistency:** if C has input variables and $\mathrm{Fix}_b(C)$ fixes the last bit to b,

$$\left|P_\delta(C) - \tfrac{1}{2}\left(P_\delta(\mathrm{Fix}_0(C)) + P_\delta(\mathrm{Fix}_1(C))\right)\right| \leq 2\delta + \beta.$$

Soundness / correctness of the axioms

- Let $\mathbb{N}$ be the standard model of PV_1. Consider a length-bounded $[0,1]$-valued interpretation $\widetilde{P}$ of $P(C,\Delta)$ as a correct approximate counting procedure

$$\widetilde{P}(C,\Delta) \ \in p(C) \pm 1/|\Delta|$$

  that is exact on constant circuits. Then $\langle \mathbb{N}, \widetilde{P} \rangle$ is a model of APX_1.

- Conversely, if $\langle \mathbb{N}, \widehat{P} \rangle$ is a model of APX_1, then

$$\widehat{P}(C,\Delta) \ \in p(C) \pm 1/|\Delta|.$$

- **Consequence:** the axioms characterize the intended notion of **approximate counting**.

Computational perspective

$$\mathbf{PV_1} \lesssim \mathbf{APX_1} \lesssim \mathbf{APC_1}$$

deterministic polytime

det. polytime + **CAPP**

" prBPP $\subseteq$ prP "

det. polytime + **Range Avoidance**

" $\exists f \notin \mathrm{SIZE}[2^{\varepsilon n}]$ "

Main Results

I: Probabilistic toolkit in APX_1

From the axioms, APX_1 derives:

- **Invariance principles:**
  - semantic invariance (equivalent circuits have close P_δ values),
  - permutation invariance (relabelling inputs barely changes P_δ).
- **Probabilistic method:** if $P_\delta(C) \geq \delta + \beta$ for some $\delta^{-1}, \beta^{-1} \in \mathrm{Log}$ then $\exists x\, C(x) = 1$.
- Useful notions of **feasible random variables** and **approximate expectation**.
- **Standard inequalities in approximate form:** union bound, Markov, Chebyshev, one-sided error reduction, restricted version of the Chernoff bound, among other results.

Example: Feasible random variables and approximate expectation

A **feasible random variable** X is specified by:

$$X = (V, n, C), \qquad C : \{0,1\}^n \to V \subseteq \mathbb{Q},$$

where V is an explicit set of size $\mathrm{poly}(n)$.

We define **indicator circuits** $C_v(x) \triangleq \mathbb{1}[C(x) = v]$, and set

$$\mathbb{E}_\delta[X] \triangleq \sum_{v \in V} v \cdot P_\delta(C_v).$$

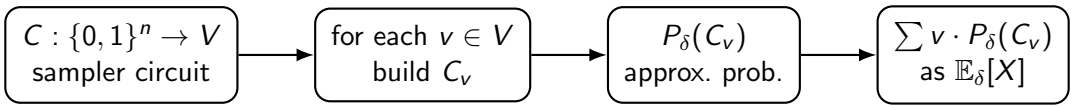

II: Formalizing TCS theorems in APX_1

APX_1 is strong enough to formalize several nontrivial results, including:

- **Yao's distinguisher-to-predictor transformation.**
- **Schwartz–Zippel lemma** (via the alternative proof technique from [AT'25]).
- **Blum–Luby–Rubinfeld (BLR) linearity testing.**
- **Circuit lower bounds:** average-case AC^0 lower bounds for parity.

Upshot: APX_1 enables probabilistic reasoning **without relying on any pigeonhole principle**.

Example: Parity vs AC^0

- **Average-case lower bound in $\mathbf{APX_1}$.** For constants $k, d \geq 1$, $\mathbf{APX_1}$ proves:
  For every $n \in \mathrm{Log}$, any depth-d AC^0 circuit C of size $\leq n^k$ agrees with $\oplus_n$ on at most

$$\tfrac{1}{2} + \tfrac{1}{n^k} + \delta + \beta$$

  fraction of inputs (measured via P_δ using circuit $T_C(x) \triangleq \mathbb{1}[C(x) = \oplus_n(x)]$).

As a byproduct of our refined treatment of AC^0 circuits in bounded arithmetic, we also show:

- **Worst-case lower bound in $\mathbf{PV_1}$.** For constants $k, d \geq 1$, $\mathbf{PV_1}$ proves:
  For every $n \in \mathrm{Log}$ and depth-d AC^0 circuit C of size $\leq n^k$,

$$\exists x \in \{0,1\}^n \text{ such that } C(x) \neq \oplus_n(x).$$

Previous formalizations used pigeonhole principles and were only known in APC_1 [K'95, MP'20].

III: Relative strength of APX_1

- Trivially: $\mathbf{PV_1} \subseteq \mathbf{APX_1}$ (i.e., $\mathbf{APX_1}$ extends $\mathbf{PV_1}$).

- $\mathbf{APX_1}$ is **interpretable** in $\mathbf{APC_1}$ (via a conservative extension where $P(C, \Delta)$ can be simulated by NW-style terms).

- Under plausible assumptions, $\mathbf{APX_1}$ is **strictly weaker** than $\mathbf{APC_1}$:

  Assume JLS-secure i$\mathcal{O}$ and coNP $\not\subseteq$ i.o.NP/poly. Then there is a $\forall\Sigma_2^b$ PV-sentence provable in APC_1 but unprovable in APX_1.

IV: $\forall\exists$-Witnessing for APX_1

We introduce the computational problem **Refuter(Yao)**:

- Input is a circuit G (**"Yao Procedure"**) that, given a "flat distribution" D (an m-tuple of n-bit strings), outputs an index i and a predictor circuit P of size s.
- A solution is a **refutation** that G is correct, i.e., a distribution D such that the predictor $(i, P) \leftarrow G(D)$ **fails** to predict the i-th bit of D with advantage $> \frac{1}{2} + \delta$.

(In a natural parameter regime, a random D is likely a solution $\Rightarrow$ the problem lies in TFZPP.)

Witnessing Theorem (Informal). If $\mathbf{APX_1}$ proves $\forall \boldsymbol{x}\, \exists \boldsymbol{y}\, \varphi(\boldsymbol{x}, \boldsymbol{y})$ (φ is an open PV-formula), then the associated search problem reduces in deterministic polytime to **Refuter(Yao)**.

Note: Refuter(Yao) reduces to **Lossy-Code** (corresponding to rWPHP).

V: Reverse mathematics of average-case lower bounds

We use $\mathbf{APX_1}$ as a **base theory** for classifying **average-case/randomized lower bounds**.

Representative result (very informal). The following statements are equivalent over $\mathbf{APX_1}$:

- Counting variants of retraction weak pigeonhole principles (**#rWPHP**):

  For any deterministic compressor-decompressor pair with encoding length $m < n$, an ε-fraction of inputs cannot be correctly decompressed.

- Randomized one-way communication lower bounds for Set Disjointness.

Interpretation: $\mathbf{APX_1}$ is expressive enough to **state** (via P_δ) and **establish** these equivalences, while still being "lightweight", as required in reverse mathematics.

From Local Consistency to Probabilistic Reasoning: A Useful Technique

The pointwise-to-global technique

Consider Boolean circuits $C_1, C_2 \colon \{0,1\}^n \to \{0,1\}$, where $C_1 \triangleq \neg C_2$.

Q. Does APX_1 prove that complementation is consistent, i.e., $P_\delta(C_1) + P_\delta(C_2) \approx 1$?

Not obvious: $P_\delta(\cdot)$ only satisfies **"local"** constraints, while this statement is **"global"**.

The **pointwise-to-global technique** allows us to connect a global statement to a local, pointwise statement (within APX_1):

For **every fixed string** $a \in \{0,1\}^n$, let $C[a]$ be circuit C fixed (hardwired) with input a.

Using the **boundary axiom**, $\mathrm{APX}_1 \vdash P_\delta(C_1[a]) + P_\delta(C_2[a]) = 1$.

Therefore, the desired statement holds **pointwise**.

The pointwise-to-global technique, cont'd

Now to an example of a **global consequence**: Suppose towards a contradiction that

$$P_\delta(C_1) + P_\delta(C_2) \;\leq\; 0.99.$$

By **local consistency** of $P_\delta(\cdot)$ and **averaging**, we can fix variable x_n to a bit a_n such that:

$$P_\delta(\mathrm{Fix}_{a_n}(C_1)) + P_\delta(\mathrm{Fix}_{a_n}(C_2)) \;\lesssim\; P_\delta(C_1) + P_\delta(C_2).$$

Repeating n times, we get a fixed string $a \in \{0,1\}^n$ such that:

$$P_\delta(C_1[a]) + P_\delta(C_2[a]) \;\lesssim\; P_\delta(C_1) + P_\delta(C_2) \;\leq\; 0.99.$$

But this **contradicts the pointwise statement** for a. □

(The formal proof proceeds by induction on n (available in APX_1), and employs **precision consistency** and the slack parameter β to handle the cumulative error from each "$\lesssim$".)

Open Problems & Concluding Remarks

Open problems

1. **Approximate counting in $\mathbf{PV_1}$?**
   Is there a PV function symbol $\tilde{P}(C, \Delta)$ that satisfies the APX_1 axioms *provably in* PV_1?

   (This implies conservativity of APX_1 over PV_1 and prBPP = prP with a feasible proof.)

2. **Conservativity:**
   Is APX_1 conservative over PV_1 for sentences not mentioning $P(C, \Delta)$?

   (A positive answer to the previous question would provide a positive answer here.)

3. **Proof derandomization vs computation derandomization:**
   If APX_1 is conservative over PV_1, does it follow that prBPP = prP?

4. **$\mathbf{APX_1}$ vs $\mathbf{APC_1}$:**
   Can we separate APC_1 and APX_1 using a $\forall\Sigma^b_1$-sentence? Does $\mathrm{APX}_1 \vdash \mathrm{rWPHP(PV)}$?

Takeaways

- $\mathbf{APX_1}$ axiomatizes approximate counting with a remarkably limited set of axioms.
- From these, it builds a workable probability toolkit (expectation, inequalities, ...).
- Strong enough to formalize several nontrivial results, yet plausibly weaker than $\mathbf{APC_1}$.
- It enables a program of reverse mathematics for average-case lower bounds.
- Finally, it motivates several research directions, including improved formalizations, unprovability results, reverse mathematics, and the feasible provability of derandomization.

Thanks!