Proof mining in ergodic theory and additive combinatorics

Jeremy Avigad

Department of Philosophy and Department of Mathematical Sciences
Carnegie Mellon University

(joint work with Philipp Gerhardy, Ksenija Simic, and Henry Towsner)

October, 2009

Ergodic theory

A *discrete dynamical system* consists of a structure, $\mathcal{X}$, and a map T from $\mathcal{X}$ to $\mathcal{X}$:

- Think of the underlying set of $\mathcal{X}$ as the set of states of a system.
- If x is a state, Tx gives the state after one unit of time.

In ergodic theory, $\mathcal{X}$ is assumed to be a finite measure space $(X, \mathcal{B}, \mu)$, and T is assumed to be a *measure-preserving transformation*, i.e. $\mu(T^{-1}A) = \mu(A)$ for every $A \in \mathcal{B}$.

Call $(X, \mathcal{B}, \mu, T)$ a *measure-preserving system*.

The metamathematics of ergodic theory

Ergodic theory emerged from seventeenth century dynamics and nineteenth century statistical mechanics.

Since Poincaré, the emphasis has been on characterizing structural properties of dynamical systems, especially with respect to long term behavior (stability, recurrence).

Today, the field uses structural, infinitary, and nonconstructive methods that are characteristic of modern mathematics.

These are often at odds with computational concerns.

The metamathematics of ergodic theory

Central questions:

- To what extent can the methods and objects of ergodic theory be given a direct computational interpretation?
- How can we locate the "constructive content" of the nonconstructive methods?
- Can we extract additional qualitative and quantitative information from nonconstructive proofs?

I will focus on two case studies:

- the von Neumann and Birkhoff ergodic theorems
- the Furstenberg-Zimmer structure theorem, and Furstenberg's ergodic-theoretic proof of Szemerédi's theorem

The ergodic theorems

Consider the orbit x, Tx, $T^2x, \ldots$, and let $f : \mathcal{X} \to \mathbb{R}$ be some measurement. Consider the averages

$$\frac{1}{n}(f(x) + f(Tx) + \ldots + f(T^{n-1}x)).$$

For each $n \geq 1$, define $A_n f$ to be the function $\frac{1}{n}\sum_{i<n} f \circ T^i$.

Theorem (Birkhoff). For every f in $L^1(\mathcal{X})$, $(A_n f)$ converges pointwise almost everywhere, and in the L^1 norm.

A space is *ergodic* if for every A, $T^{-1}(A) = A$ implies $\mu(A) = 0$ or $\mu(A) = 1$.

If $\mathcal{X}$ is *ergodic*, then $(A_n f)$ converges to the constant function $\int f \; d\mu$.

The ergodic theorems

Recall that $L^2(\mathcal{X})$ is the Hilbert space of square-integrable functions on $\mathcal{X}$ modulo a.e. equivalence, with inner product

$$(f, g) = \int fg \; d\mu$$

Theorem (von Neumann). For every f in $L^2(\mathcal{X})$, $(A_n f)$ converges in the L^2 norm.

A measure-preserving transformation T gives rise to an isometry T on $L^2(\mathcal{X})$,

$$Tf = f \circ T.$$

Riesz showed that the von Neumann ergodic theorem holds, more generally, for any nonexpansive operator T on a Hilbert space (i.e. satisfying $\|Tf\| \leq \|f\|$ for every f in $\mathcal{H}$.)

Bounding the rate of convergence

Can we compute a bound on the rate of convergence of $(A_n f)$ from the inital data (T and f)?

In other words: can we compute a function $r : \mathbb{Q} \to \mathbb{N}$ such that for every rational $\varepsilon > 0$,

$$\|A_m f - A_{r(\varepsilon)} f\| < \varepsilon$$

whenever $m \geq r(\varepsilon)$?

Krengel (et al.): convergence can be arbitrarily slow. But computability is a different question.

Note that the question depends on suitable notions of computability in analysis.

Observations

If $(a_n)_{n \in \mathbb{N}}$ is a sequence of reals that decreases to 0, no matter how slowly, one can compute a bound on the rate of convergence from (a_n).

But there are bounded, computable, decreasing sequences (b_n) of rationals that do not have a computable limit.

There are also computable sequences (c_n) of rationals that converge to 0, with no computable bound on the rate of convergence.

Conclusion: at issue is not the *rate* of convergence, but its *predictability*.

A negative result

Theorem (A-Simic). There are a computable measure-preserving transformation of $[0,1]$ under Lebesgue measure and a computable characteristic function $f = \chi_A$, such that if $f^* = \lim_n A_n f$, then $\|f^*\|_2$ is not a computable real number.

In particular, f^* is not a computable element of $L^2(\mathcal{X})$, and there is no computable bound on the rate of convergence of $(A_n f)$ in either the L^2 or L^1 norm.

V'yugin (1998) gave a similar noncomputability result.

A positive result

Theorem (A-Gerhardy-Towsner). Let T be a nonexpansive operator on a separable Hilbert space and let f be an element of that space. Let $f^* = \lim_n A_n f$. Then f^*, and a bound on the rate of convergence of $(A_n f)$ in the Hilbert space norm, can be computed from f, T, and $\|f^*\|$.

In particular, if T arises from an ergodic transformation T, then f^* is computable from T and f.

A constructive mean ergodic theorem

It turns out that we can say more, even in situations where there is no computable bound on the rate of convergence.

The assertion that the sequence $(A_n f)$ converges can be represented as follows:

$$\forall \varepsilon > 0 \, \exists n \, \forall m \geq n \, (\|A_m f - A_n f\| < \varepsilon).$$

This is classically equivalent to the assertion that for any function K,

$$\forall \varepsilon > 0 \, \exists n \, \forall m \in [n, K(n)] \, (\|A_m f - A_n f\| < \varepsilon).$$

A constructive mean ergodic theorem

Theorem (A-G-T). Let T be any nonexpansive operator on a Hilbert space, let f be any element of that space, and let $\varepsilon > 0$, and let K be any function. Then there is an $n \geq 1$ such that for every m in $[n, K(n)]$, $\|A_m f - A_n f\| < \varepsilon$.

In fact, we provide a bound on n expressed solely in terms of K and $\rho = \|f\|/\varepsilon$. Notably, the bound is independent of X and T.

As special cases, we have the following:

- If $K = n^{O(1)}$, then $n(f, \varepsilon) = 2^{2^{O(\rho^2 \log\log \rho)}}$.
- If $K = 2^{O(n)}$, then $n(f, \varepsilon) = 2^1_{O(\rho^2)}$.
- If $K = O(n)$ and T is an isometry, then $n(f, \varepsilon) = 2^{O(\rho^2 \log \rho)}$.

A constructive pointwise ergodic theorem

The following is classically equivalent to the pointwise ergodic theorem:

Theorem (A-G-T). For every f in $L^2(\mathcal{X})$, $\lambda_1 > 0$, $\lambda_2 > 0$, and K there is an $n \geq 1$ satisfying

$$\mu(\{x \mid \max_{n \leq m \leq K(n)} |A_n f(x) - A_m f(x)| > \lambda_1\}) \leq \lambda_2.$$

We provide explicit bounds on n in terms of f, λ_1, λ_2, and K.

Bishop's *upcrossing inequalities* provides another constructive interpretation of the pointwise ergodic theorem.

Proof mining

Our constructive proof was obtained using "proof mining" methods, developed chiefly by Ulrich Kohlenbach and his students.

Terence Tao has used similar ideas to obtain constructive / quantitative versions of nonconstructive statements. He referred to such phenomena as "metastability."

Such ideas played a role in his "Norm convergence of multiple ergodic averages for commuting transformations," and in his work with Ben Green on arithmetic progressions in the primes.

Our proofs are a form of "energy incrementation" argument. The relationships between the infinitary and quantitative methods needs to be better understood.

Ergodic Ramsey theory

Let us consider an applications of ergodic theory to combinatorics.

Theorem (van der Waerden). If one colors the natural numbers with finitely many colors, then there are arbitrarily long monochromatic arithmetic progressions.

The theorem has a finitary (Π_2) statement:

Theorem. For every k and r there is an n large enough such that if one colors elements of the set $\{1, \ldots, n\}$ with r colors, there is a monochromatic arithmetic progression of length k.

van der Waerden proved this in 1927. Furstenberg and Weiss presented an elegant proof using topological dynamics in 1978.

Szemerédi's theorem

Szemerédi's theorem is a "density" version of van der Waerden's theorem.

Szemerédi's Theorem. Every set S of natural numbers with positive upper Banach density has arbitrarily long arithmetic progressions.

Equivalently:

Theorem. For every k and $\delta > 0$, there is an n large enough, such that if S is any subset of $\{1, \ldots, n\}$ with density at least δ, then S has an arithmetic progression of length k.

History

- 1936: Conjectured by Erdös and Turán
- 1952: Roth proved it for $k = 3$.
- 1969: Szemerédi proved it for $k = 4$.
- 1974: Szemerédi proved the full theorem.
- 1977: Furstenberg
  - gave an equivalent ergodic-theoretic statement,
  - provided a structural analysis of ergodic measure-preserving systems, and
  - used the latter to give a proof.
- 1979: Furstenberg and Katznelson used the structure theorem to give a streamlined proof of an even stronger result.
- 2001: Gowers gave a new proof of Szemerédi's theorem, with elementary bounds.
- 2004: Tao and Green used quantitative ergodic-theoretic methods to prove that there are arbitrarily long arithmetic progressions in the primes.

Logical analysis

The fact that powerful infinitary methods can yield explicit combinatorial results deserves logical analysis.

Recall the central questions:

- To what extent can the methods and objects of ergodic theory be given a direct computational interpretation?
- How can we locate the "constructive content" of the nonconstructive methods?
- Can we extract additional qualitative and quantitative information from nonconstructive proofs?

Furstenberg correspondence

Suppose there were a sequence of subsets S_m of $\{0, \ldots, m-1\}$ of density $\delta > 0$, with no arithmetic progression of length k.

Consider the spaces $X_m = \{0, \ldots, 2m-1\}$ with uniform distribution and shift map $Tx = x + 1 \bmod 2m$. Then for every m and $n < m$,

$$S_m \cap T^{-n} S_m \cap T^{-2n} S_m \cap \ldots \cap T^{-(k-1)n} S_m = \emptyset.$$

A compactness argument yields a space $(X, \mathcal{B}, \mu)$ and set S that gives a counterexample to the following:

Theorem. For any measure-preserving system $(X, \mathcal{B}, \mu, T)$, any set S of positive measure, and any k, there is an n such that

$$\mu(S \cap T^{-n} S \cap T^{-2n} S \cap \ldots \cap T^{-(k-1)n} S) > 0.$$

In fact, this theorem is equivalent to Szemerédi's theorem.

The Furstenberg-Katznelson-Ornstein proof

Most presentations of Furstenberg's proof use the Furstenberg-Zimmer structure theorem.

This shows that every measure preserving system has a factor, $\mathcal{Y}$, such that

- $\mathcal{Y}$ is obtained by a transfinite sequence of compact extensions, starting with the trivial factor.
- $\mathcal{X}$ is weak mixing relative to $\mathcal{Y}$.

Foreman and Beleznay have shown that, for separable systems, the construction of $\mathcal{Y}$ can extend arbitrarily far into the countable ordinals.

Towsner and I have shown that a weaker form of the second is sufficient for the FKO proof, and this happens before level ω^{ω^ω} of the hierarchy.

The Furstenberg-Katznelson-Ornstein proof

Say a set A is *UMR* if for every k, $\mu(A) > 0$ implies

$$\liminf_{n\to\infty} \frac{1}{n} \sum_{i<n} \mu(A \cap T^{-n}A \cap T^{-2n}A \cap \cdots \cap T^{-(k-1)n}A) > 0.$$

Say a factor is *UMR* if every element is *UMR*.

This is a strengthening of the desired conclusion for S.

The property of being *UMR*:

- holds the trivial factor;
- is maintained under compact extensions;
- is maintained under limits; and
- is maintained under weak mixing extensions.

Furstenberg's original proof

Say that a factor $\mathcal{Y}$ is k-characteristic if the following holds:

$$\lim_{n\to\infty} \left\| \frac{1}{n} \sum_{i<n} \prod_{l=1}^{k} T^{li} f_l - \frac{1}{n} \sum_{i<n} \prod_{l=1}^{k} T^{li} E(f_l \mid \mathcal{Y}) \right\|_{L^2(\mathcal{X})} = 0.$$

This means that the projections of the f_l's on $\mathcal{Y}$ are enough to determine the limiting behavior of the first average.

A factor $\mathcal{Y}$ is UMR if the following holds: for any nonnegative g in $L^\infty(\mathcal{X})$ and any t, if $\int g \, d\mu > 0$ then

$$\liminf_{u\to\infty} \frac{1}{u} \sum_{i<u} \int \prod_{l<t} T^{li} E(g \mid \mathcal{Y}) \, d\mu > 0.$$

This means that elements of $\mathcal{Y}$ satisfy the conclusion of Szemerédi's theorem.

Furstenberg's original proof

Define a sequence of factors $\mathcal{Y}_k$ as follows:

- $\mathcal{Y}_0$ consists of the T-invariant functions.
- For each k, $\mathcal{Y}_{k+1} = \mathcal{Z}(\mathcal{Y}_k)$, the maximal compact extension of $\mathcal{Y}_k$.

Each of these can be proved by induction on k:

- $\mathcal{Y}_k$ is $(k+1)$-characteristic.
- $\mathcal{Y}_k$ is UMR.

These yield the conclusion of Szemerédi's theorem for progressions of length $k+2$.

Towards a quantitative version

The statement that $\mathcal{Y}$ is k-characteristic is equivalent to the following:

$$\forall f_1, \ldots, f_k, \varepsilon > 0, M \, \exists n \, \forall m \in [n, M(n)]$$
$$\left\| \frac{1}{n} \sum_{i<n} \prod_{l=1}^{k} T^{li} f_l - \frac{1}{n} \sum_{i<n} \prod_{l=1}^{k} T^{li} E(f_l \mid \mathcal{Y}) \right\|_{L^2(\mathcal{X})} < \varepsilon.$$

The statement that $\mathcal{Y}$ is UMR is equivalent to the following:

$$\forall g > 0, t, V \, \exists \delta > 0, u \, \forall v \in [u, V(\delta, u)]$$
$$\frac{1}{v} \sum_{i<v} \int \prod_{l<t} T^{li} E(g \mid \mathcal{Y}) \, d\mu > \delta.$$

Towards a quantitative version

A first attempt would be to try to prove, by induction on k, a quantitative version of the statement that for every $k \geq 1$, there is an invariant factor $\mathcal{Y}$ such that

- $\mathcal{Y}$ is k-characteristic, and
- $\mathcal{Y}$ is UMR.

The problem is that the relevant $\mathcal{Y}$'s are far from computable.

Towards a quantitative version

An alternative approach: fix $\hat{k}$, and a set A with $\mu(A) > 0$. Prove by induction on k, from $\hat{k}$ down to 1, a quantitative version of the statement that for every invariant factor $\mathcal{Y}$, if

- $\mathcal{Y}$ is k-characteristic, and
- $\mathcal{Y}$ is UMR,

then the conclusion holds for A and $\hat{k}$.

When k is $\hat{k}$, the proof of the original statement is easy.

On the other hand, when $k = 1$, the quantitative MET makes it possible to satisfy the antedecent, and so the conclusion follows.

Towards a quantitative version

In symbols: for every k from $\hat{k}$ down to 1, and any invariant $\mathcal{Y}$, if

$$\forall f_1, \ldots, f_k, \varepsilon > 0, M \; \exists n \; \forall m \in [n, M(n)]$$
$$\left\| \frac{1}{n} \sum_{i<n} \prod_{l=1}^{k} T^{li} f_l - \frac{1}{n} \sum_{i<n} \prod_{l=1}^{k} T^{li} E(f_l \mid \mathcal{Y}) \right\|_{L^2(\mathcal{X})} < \varepsilon.$$

and

$$\forall g > 0, t, V \; \exists \delta > 0, u \; \forall v \in [u, V(\delta, u)]$$
$$\frac{1}{v} \sum_{i<v} \int \prod_{l<t} T^{li} E(g \mid \mathcal{Y}) \, d\mu > \delta.$$

then the conclusion holds.

Towards a quantitative version

In fact, one only has to show how the conclusion can be obtained from witnessing functions for n, δ, and μ, which in a sense determine how well $\mathcal{Y}$ is "locally k-characteristic" and "locally UMR."

It is possible that there is a bound on the conclusion that doesn't depend on $\mathcal{Y}$.

If not, one would hope to isolate the quantitative features of $\mathcal{Y}$ on which the result depends.

Tao's quantitative version

Recall that the measure space coming out of the Furstenberg construction can be viewed as a "limit" of finite spaces. Tao's quantitative proof simply uses a sufficiently large finite space.

One difficulty: constructions in the limit do not correspond to constructions in the finite spaces. For example, a factor in the limit is not a limit of factors.

Tao considers complexity-bounded approximations to the "true" ergodic-theoretic factors, for example, finite factors where the number of atoms is bounded independent of n.

It would be helpful to have a cleaner connection to the infinitary argument.

Conclusions

Goals:

- A better understanding of the relationship between the infinitary ("soft") and finitary, quantitative ("hard") methods.
- Infinitary methods that are better suited to finitary problems.
- Additional information from proofs using the infinitary methods.
- An understanding as to how and where logical strength can be avoided, and where it is necessary.

There is a lot to do:

- Dynamical systems represents represent an uneasy tension between structural and computational concerns.
- Applications to combinatorics, in particular, require both structural ideas and quantitative information.

Some references

Associated papers and talks can be found on my web page:

- "Fundamental notions of analysis in subsystems of second-order arithmetic" (with Ksenija Simic)
- "Local stability of ergodic averages" (with Philipp Gerhardy and Henry Towsner)
- "Functional interpretation and inductive definitions" (with Henry Towsner)
- "The metamathematics of ergodic theory"
- "Metastability in the Furstenberg-Zimmer tower" (with Henry Towsner)