Density-one Points of Π^0_1 Classes

Mushfeq Khan

University of Wisconsin—Madison

Midwest Computability Seminar XIII,
University of Chicago,
October 1st, 2013

Outline

1. Definitions and observations
2. Dyadic density-one vs full density-one
3. What can density-one points compute?
4. Is there a minimal density-one point?

Dyadic density-one points

We use the symbol μ to refer to the uniform measure, both on Cantor space and on the unit interval.

Given $\sigma \in 2^{<\omega}$ and a measurable set $C \subseteq 2^\omega$, the shorthand $\mu_\sigma(C)$ denotes the relative measure of C in the cone above σ, i.e.,

$$\mu_\sigma(C) = \frac{\mu([\sigma] \cap C)}{\mu([\sigma])}.$$

Definition

Let C be a measurable set and X a real. The lower dyadic density of C at X, written $\rho_2(C \mid X)$, is

$$\liminf_n \mu_{X \upharpoonright n}(C).$$

Definition

A real X is a dyadic positive density point if for every Π^0_1 class C containing X, $\rho_2(C \mid X) > 0$. It is a dyadic density-one point if for every Π^0_1 class C containing X, $\rho_2(C \mid X) = 1$.

Full density-one points

Even though dyadic density seems like the natural notion of density in Cantor space, it is a simplification of the version of density that appears in the classical Lebesgue Density Theorem:

Definition

Let C be a measurable subset of $\mathbb{R}$ and $x \in \mathbb{R}$. The lower (full) density of C at x, written $\rho(C \mid x)$, is

$$\liminf_{\gamma,\delta \to 0^+} \frac{\mu((x-\gamma, x+\delta) \cap C)}{\gamma + \delta}.$$

Definition

We say $x \in [0,1]$ is a positive density point if for every Π^0_1 class $C \subseteq [0,1]$ containing x, $\rho(C \mid x) > 0$. It is a (full) density-one point if for every Π^0_1 class $C \subseteq [0,1]$ containing x, $\rho(C \mid x) = 1$.

Density on the Martin-Löf random reals

Theorem (Bienvenu, Hölzl, Miller, Nies)

If X is Martin-Löf random, then X is a positive density point if and only if it is incomplete.

Theorem (Day, Miller)

There is a Martin-Löf random real that is a positive density point (hence incomplete) but not a density-one point.

- Dyadic positive density points (and hence full positive density points) are Kurtz random.
- 1-generics are full density-one points.
- Not being a full density-one point is a Π_2^0 property. Therefore, all weak 2-random reals are full density-one points. Note that any hyperimmune-free Kurtz random is weak 2-random (Yu).
- The two halves of a dyadic density-one point are dyadic density-one. In fact, any computable sampling of a dyadic density-one point is a dyadic density-one point. Likewise for full density-one points.
- There is a Kurtz random real that is not Martin-Löf random and not a density-one point. Consider $\Omega \oplus G$ where G is weakly 2-generic.

The resulting picture

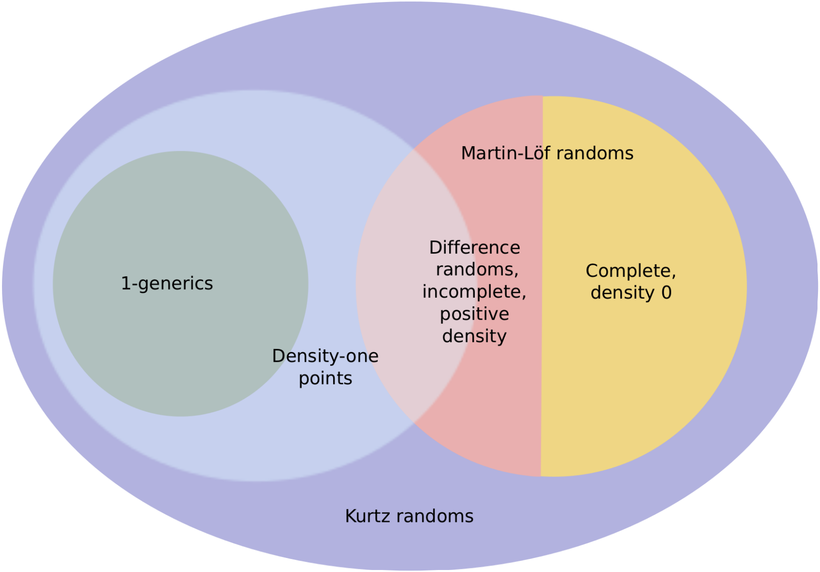

It's easy to exhibit a specific C and an X such that $\rho_2(C \mid X) \neq \rho(C \mid X)$. But is this discrepancy eliminated if we require that for every Π^0_1 class C containing X, $\rho_2(C \mid X) = 1$? In other words, are dyadic density-one points the same as full density-one points? On the Martin-Löf randoms, yes:

Theorem (K., Miller)

Let X be Martin-Löf random. Then X is a dyadic density-one point if and only if it is a full density-one point.

Some amount of randomness is necessary:

Proposition (K.)

There is a dyadic density-one point that is not full density-one.

We build a dyadic density-one point Y by computable approximation, while building a Σ^0_1 class B such that $\rho(\bar{B} \mid Y) < 1$.

The basic idea:

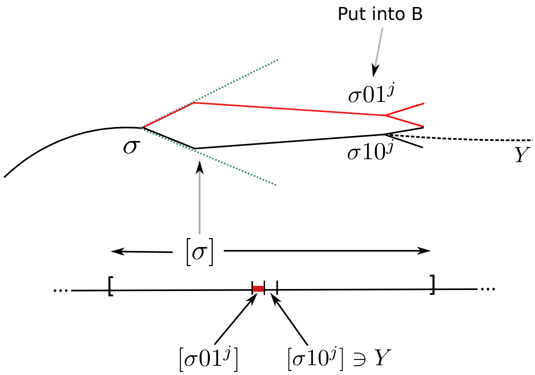

We shall be free to choose j as large as we want. Note that $[\sigma]$ is the smallest dyadic cone containing Y that can see $[\sigma 01^j]$, the "hole" that we create in $\bar{B}$, and relative to σ, this hole appears small. However, on the real line, at a certain scale around Y, the hole is quite large.

Dyadic density-one vs full density-one (contd.)

We want to place these holes infinitely often along Y, and this constitutes one type of requirement. Making Y a dyadic density-one point amounts to ensuring that for each Σ^0_1 class $[W_e]$, either

1. $Y \in [W_e]$, or
2. the relative measure of $[W_e]$ along Y goes to 0.

The basic strategy for meeting a density requirement is to reroute Y to enter $[W_e]$ if its measure becomes too big above some initial segment of Y_s. To make this play well with our hole-placing strategy, we keep the measure of B above initial segments of Y_s very small. Then if $[W_e]$ becomes big enough, we can enter it while keeping B very small along Y_s.

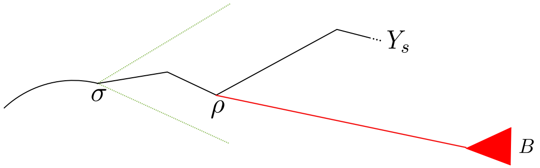

Dyadic density-one vs full density-one (contd.)

We want to place these holes infinitely often along Y, and this constitutes one type of requirement. Making Y a dyadic density-one point amounts to ensuring that for each Σ^0_1 class $[W_e]$, either

1. $Y \in [W_e]$, or
2. the relative measure of $[W_e]$ along Y goes to 0.

The basic strategy for meeting a density requirement is to reroute Y to enter $[W_e]$ if its measure becomes too big above some initial segment of Y_s. To make this play well with our hole-placing strategy, we keep the measure of B above initial segments of Y_s very small. Then if $[W_e]$ becomes big enough, we can enter it while keeping B very small along Y_s.

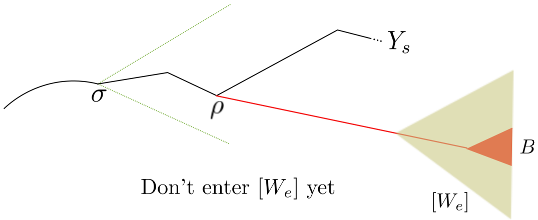

Dyadic density-one vs full density-one (contd.)

We want to place these holes infinitely often along Y, and this constitutes one type of requirement. Making Y a dyadic density-one point amounts to ensuring that for each Σ_1^0 class $[W_e]$, either

1. $Y \in [W_e]$, or
2. the relative measure of $[W_e]$ along Y goes to 0.

The basic strategy for meeting a density requirement is to reroute Y to enter $[W_e]$ if its measure becomes too big above some initial segment of Y_s. To make this play well with our hole-placing strategy, we keep the measure of B above initial segments of Y_s very small. Then if $[W_e]$ becomes big enough, we can enter it while keeping B very small along Y_s.

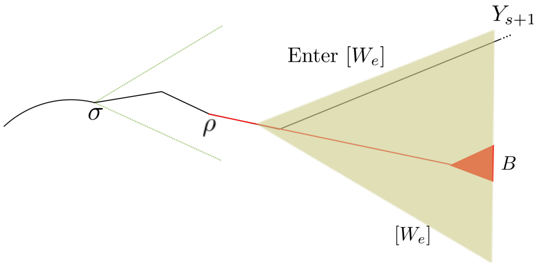

The following lemma makes this intuition precise:

Covering Lemma, dyadic version

Suppose $B \subseteq 2^\omega$ is open. Then for any ε such that $\mu(B) \leq \varepsilon \leq 1$, let $U_\varepsilon(B)$ denote the set

$$\{X \in 2^\omega : \mu_\rho(B) \geq \varepsilon \text{ for some } \rho \prec X\}.$$

Then $U_\varepsilon(B)$ is open and $\mu(U_\varepsilon(B)) \leq \mu(B)/\varepsilon$.

The lemma tells us exactly how small we have to keep B along Y_s to make it possible to act for multiple density requirements. Each time we reroute Y_s to enter a Σ^0_1 class we get a little "closer" to B, but still remain far enough away so that we can act on behalf on another, higher priority density requirement if the need arises.

Interleave hole-placing requirements with density requirements by progressively building a better and better approximation to a dyadic density-one point.

Formally, to meet the requirement $D_{e,k}$ between σ and σ' where $\sigma \preceq \sigma' \prec Y$ is to ensure that either $\sigma' \in [W_e]$ or the measure of $[W_e]$ between σ and σ' is bounded by 2^{-k} (i.e., for every τ between σ and σ', $\mu_\tau([W_e]) \leq 2^{-k}$).
We organize the construction as follows:

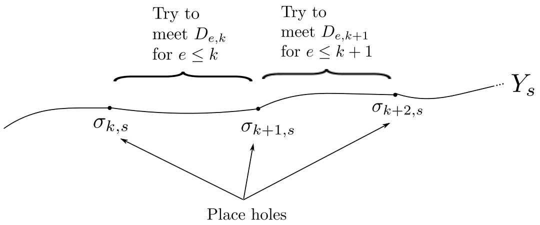

$D_{e,k}$ has higher priority than $D_{e',k}$, for $e' > e$. Above $\sigma_{k,s}$, we only act for the sake of $D_{e,k}$ if we haven't acted for the sake of a higher priority density requirement above $\sigma_{k,s}$. In sum, we have a finite-injury priority construction, where for each e, cofinitely many of the $D_{e,k}$ requirements will be satisfied. There are some details to work out, but they're routine. $\square$

1-generics are GL_1, therefore incomplete. By the theorem of Bienvenu et al., Martin-Löf random density-one points are also incomplete. But in general, density-one points can be complete. In fact, every real is computable from a full density-one point:

Theorem (K.)

For every $X \in 2^\omega$, there is a full density-one point Y such that $X \leq_T Y \leq_T X \oplus 0'$.

Because dyadic density is so much easier to work with, I'll first sketch the proof of the result for dyadic density. Even though the statement of the theorem bears a superficial resemblance to the Kučera-Gács Theorem, the method is different. For one thing, there is no Π^0_1 class consisting exclusively of density-one points. Also note that we don't get a `wtt` reduction as in the Kučera-Gács Theorem.

Basic idea: Combine the density strategy of the previous proof with coding, on a tree.
By computable approximation, we build a Δ^0_2 function tree $F : 2^{<\omega} \to 2^{<\omega}$ and a functional Γ such that for every $\sigma \in 2^{<\omega}$, $\Gamma^{F(\sigma)} = \sigma$.

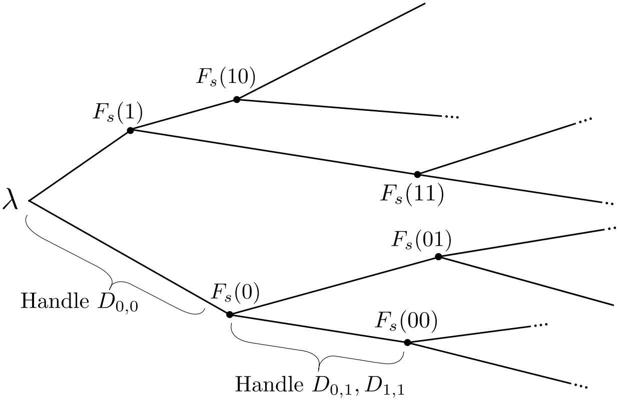

Computational strength (contd.)

To set $F_s(\sigma) = \tau$ at stage s is to code σ at the string τ. We need to ensure that we can always do this in a consistent manner. There are two ways this could go wrong:

- τ codes incorrectly (i.e., $\Gamma^\tau \mid \sigma$), or
- τ codes too much (i.e. Γ^τ properly extends σ).

For example:

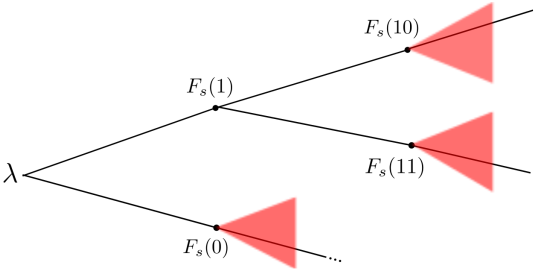

We cannot route $F_s(1)$ through the current or previous values of $F_s(10)$, $F_s(11)$ and $F_s(0)$.

In general, for every nonempty string σ, there is a Σ^0_1 class $B_{\sigma,s}$ that $F_s(\sigma)$ must avoid, and a threshold $\beta_{\sigma,s}$ below which we must keep the measure of $B_{\sigma,s}$ between $F_s(\sigma^-)$ and $F_s(\sigma)$, where σ^- is the immediate predecessor of σ.

The strategies must cooperate to maintain this condition. For example, if $\sigma = \alpha 0$, then the strategies controlling $F_s(\sigma 0)$, $F_s(\sigma 1)$ and $F_s(\alpha 1)$, all of which contribute measure to $B_{\sigma,s}$, must maintain the fact that $\mu(B_{\sigma,s})$ remains strictly below $\beta_{\sigma,s}$ between $F_s(\alpha)$ and $F_s(\sigma)$. All of this is completely within our control, since we can code on arbitrarily long strings.

For each $X \in 2^\omega$, the construction of $\bigcup_{\sigma \prec X} F(\sigma)$ is again a finite-injury priority construction. The details are easy to work out. $\square$

Real line issues

We briefly outline some of the difficulties in transferring the coding theorem for dyadic density-one points to full density-one points.

Strategies can no longer restrict their attention to dyadic cones:

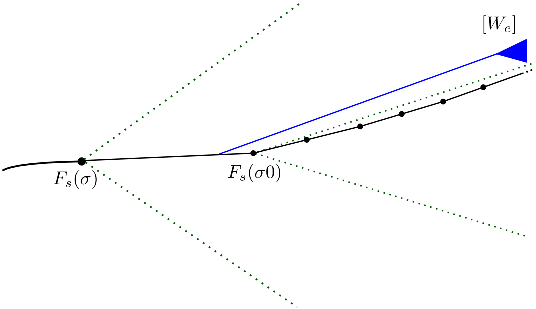

$[W_e]$ is very small relative to $F_s(\sigma)$, but it poses a threat to the path we're building.

We build a tree $\{I_\sigma : \sigma \in 2^{<\omega}\}$ of intervals with dyadic rational endpoints.

Suppose $I \supseteq I'$ are intervals in $[0, 1]$ and C is a measurable set. We say that $\mu(C)$ *is below* ε *between* I *and* I' if for every interval L such that $I \supseteq L \supseteq I'$, $\mu_L(C) < \varepsilon$.

In previous proofs, it was easy to chop up density requirements into smaller pieces such that the individual wins added up nicely. This is a little messier on the real line:

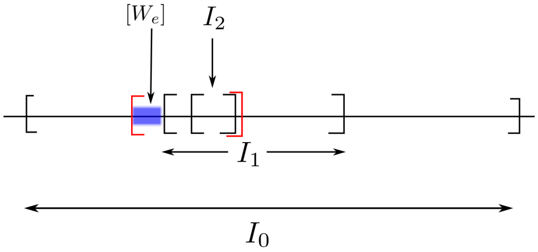

Here $\mu([W_e]) < 1/8$ between I_0 and I_1 and also between I_1 and I_2, but not between I_0 and I_2.

There is a version of the Density Drop Covering Lemma for the real line:

Lemma (Bienvenu, Hölzl, Miller, Nies)

Suppose $B \subseteq [0,1]$ is open. Then for any ε such that $\mu(B) \leq \varepsilon \leq 1$, let $U_\varepsilon(B)$ denote the set

$$\{X \in [0,1] : \exists \text{ an interval } I, X \in I, \text{ and } \mu_I(B) \geq \varepsilon\}.$$

Then $\mu(U_\varepsilon(B)) \leq 2\mu(B)/\varepsilon$.

We have to be slightly careful when applying this lemma for our construction. When we relativize this lemma to an interval L, we obtain a bound for the measure of $U_\varepsilon(B \cap L)$ within L, but in general, we are also concerned about the part of B that lies outside L. Fortunately, under the assumptions of the construction, we can obtain a bound for the measure threatened by all of B.

We skip the details. On to the next topic...

Question

Is there a density-one point of minimal degree?

Of course, 1-generics and 1-randoms cannot be minimal, since for any real $A \oplus B$ with either property, A and B are Turing incomparable. This is not true of density-one points:

Fact

There is a density-one point $A \oplus B$ with $A \equiv_T B$.

Theorem

There is a high degree that computes no density-one point.

Main idea: Combine the Sacks minimal degree construction with coding bits of $0''$ in an "almost" $0'$-computable construction. The resulting minimal real A is such that $A \oplus 0'$ can rerun the construction and recover $0''$. Instead of the usual splitting trees, use "thin" splitting trees. Code the bits of $0''$ by choosing the left or right subtree of the splitting tree.

Thanks!