

What are the hyperreals?

Saúl Rodríguez

Sets small and large

Theorem

There exists a collection $\mathcal{F} \subseteq \mathcal{P}(\mathbb{N})$ of subsets of \mathbb{N} such that, for every $A, B \subseteq \mathbb{N}$:

Intuition: $A \in \mathcal{F}$ means 'A contains almost all natural numbers'.

1. $A \in \mathcal{F}$ iff $\mathbb{N} \setminus A \notin \mathcal{F}$.
2. $\mathbb{N} \in \mathcal{F}$, finite sets are not in \mathcal{F} .
3. If $A \in \mathcal{F}$ and $A \subseteq B$, then $B \in \mathcal{F}$. Equivalently, if $A \notin \mathcal{F}$ and $B \subseteq A$, then $B \notin \mathcal{F}$.
4. If $A, B \in \mathcal{F}$ then $A \cap B \in \mathcal{F}$. Equivalently, if $A, B \notin \mathcal{F}$ then $A \cup B \notin \mathcal{F}$.

Such a collection \mathcal{F} is called a 'non-principal ultrafilter in \mathbb{N} '.

There are many ($2^{2^{\aleph_0}}$) non-principal ultrafilters in \mathbb{N} , but we cannot construct any of them explicitly. We need the axiom of choice (AC) to do it (but ultrafilter lemma is weaker than AC).

Big sets - they are basically full measure sets

A non-principal ultrafilter \mathcal{F} defines a **finitely** additive probability measure $\mu : \mathcal{P}(\mathbb{N}) \rightarrow \{0, 1\}$, by $\mu(A) = 1$ if $A \in \mathcal{F}$ and $\mu(A) = 0$ otherwise.

We say that ' \mathcal{F} -almost all $n \in \mathbb{N}$ satisfy some property P ' when

$$\{n \in \mathbb{N}; n \text{ satisfies } P\} \in \mathcal{F}.$$

Ultrafilters have some not so nice properties: e.g. set of even numbers is in \mathcal{F} iff set of odd numbers is not in \mathcal{F} (even though they are translates of each other).

Ultrafilters define limits

Let $(x_n)_n$ be a sequence of real numbers, let $L \in \mathbb{R}$.

Definition

We say that $L = \lim_{n \rightarrow \infty} x_n$ when for any neighborhood $U = (L - \varepsilon, L + \varepsilon)$ of L we have $x_n \in U$ for all n except finitely many.

A bounded sequence of real numbers $(x_n)_n$ always has accumulation points, but does not necessarily have a limit. For example, $x_n = (-1)^n$. We can use ultrafilters to 'force' any bounded sequence to have a limit:

Definition

We say that $L = \lim_{n \rightarrow \mathcal{F}} x_n$ when for any neighborhood $U = (L - \varepsilon, L + \varepsilon)$ of L we have $x_n \in U$ for \mathcal{F} -almost all n .

In particular, $L = \lim_{n \rightarrow \infty} x_n$ implies $L = \lim_{n \rightarrow \mathcal{F}} x_n$.

Ultrafilters define limits

Theorem

Every bounded sequence $(x_n)_n$ of real numbers has an \mathcal{F} -limit $L = \lim_{n \rightarrow \mathcal{F}} x_n \in \mathbb{R}$.

Proof.

For each $x \in \mathbb{R}$ let $A_x = \{n \in \mathbb{N}; x_n < x\}$. Note that for some big constant M , we have $A_{-M} = \emptyset$, $A_M = \mathbb{N}$. Thus, $L = \inf\{x \in \mathbb{R}; A_x \in \mathcal{F}\}$ is defined.

Then, for each $\varepsilon > 0$, $A_{L+\varepsilon} \in \mathcal{F}$. Similarly, $A_{L-\frac{\varepsilon}{2}} \notin \mathcal{F}$. So $x_n < L + \varepsilon$ for \mathcal{F} -almost all n and $x_n > L - \varepsilon$ for \mathcal{F} -almost all n . So $x_n \in (L - \varepsilon, L + \varepsilon)$ for \mathcal{F} -almost all n . □

In general, sequences in any compact space have limits along ultrafilters.

In the case of $x_n = (-1)^n$, we will have $\lim_{n \rightarrow \mathcal{F}} (-1)^n = 1$ if the set of even numbers is in \mathcal{F} , and $\lim_{n \rightarrow \mathcal{F}} (-1)^n = -1$ if not.

Infinitesimals

An ordered field is a field $(K, +, \cdot)$ with a total order \leq such that, for all $a, b, c \in K$:

1. If $a \leq b$ then $a + c \leq b + c$.
2. If $0 \leq a, b$ then $0 \leq ab$.

For $x \in K$ we denote $|x| = x$ if $x \geq 0$ or $|x| = -x$ if $x \leq 0$.

We say that an element $\varepsilon \in K$ is infinitesimal if for all $n \in \mathbb{N}$ we have $|n\varepsilon| < 1$. An ordered field is said to be *Archimedean* if it contains no nonzero infinitesimals. For example, \mathbb{R} is Archimedean.

There are many ways to construct a non-Archimedean extension of \mathbb{R} . The construction we will use was introduced by Edwin Hewitt in 1948.

The hyperreals

Let $\prod_{n \in \mathbb{Z}} \mathbb{R}$ be the set of sequences $(x_n)_n$ of real numbers.

The ordered field of hyperreal numbers, ${}^*\mathbb{R}$, is defined as the quotient $\frac{\prod_{n \in \mathbb{Z}} \mathbb{R}}{\sim}$, where $(x_n)_n \sim (y_n)_n$ if $x_n = y_n$ for \mathcal{F} -almost all n .

We denote by $[x_n]_n \in {}^*\mathbb{R}$ the class of a sequence $(x_n)_n$.

The operations are given by

$$[x_n]_n + [y_n]_n = [x_n + y_n]_n$$

$$[x_n]_n \cdot [y_n]_n = [x_n \cdot y_n]_n.$$

The order is given by $[x_n]_n \leq [y_n]_n$ if $x_n \leq y_n$ for \mathcal{F} -almost all n .

The sum identity is $[0]_n$ and the product identity is $[1]_n$.

The fact that the operations/order are well defined and ${}^*\mathbb{R}$ is an ordered field is a good list of exercises to practice the properties of ultrafilters.

Infinitesimals in the hyperreals

Note that the map $\mathbb{R} \rightarrow {}^*\mathbb{R}; x \mapsto [x]_n$ is a homomorphism of fields, so ${}^*\mathbb{R}$ is a field extension of \mathbb{R} . For any $x \in \mathbb{R}$ we denote $x = [x]_n$. A number $x \in {}^*\mathbb{R}$ is an infinitesimal iff for all real $\varepsilon > 0$ we have $-\varepsilon < x < \varepsilon$.

For example, the number $x = \left[\frac{1}{n}\right]_n$ is an infinitesimal, as for all $\varepsilon > 0$ we have $-\varepsilon < \frac{1}{n} < \varepsilon$ for \mathcal{F} -almost all n , so

$$-\varepsilon = [-\varepsilon]_n \leq \left[\frac{1}{n}\right]_n \leq [\varepsilon]_n = \varepsilon.$$

Limited numbers

Definition

Say ${}^*\mathbb{R}$ is *limited* if $-N \leq x \leq N$ for some $N \in \mathbb{N}$. If not, we say that x is *unlimited*.

For example, $[n]_n$ is an unlimited number.

Proposition

Every limited hyper-real number is a sum of a real number and an infinitesimal.

Proof.

Let $[x_n]_n$ be limited, so that $-N \leq x \leq N$ for some $N \in \mathbb{N}$. We may assume $|x_n| \leq N$ for all n , so that $(x_n)_n$ is bounded. Let $L = \lim_{n \rightarrow \mathcal{F}} x_n$.

Then by definition of L , for all ε we have $x_n \in (L - \varepsilon, L + \varepsilon)$ for \mathcal{F} -almost all n . So $x_n - L \in (\varepsilon, \varepsilon)$ for all n .

Meaning that the number $[x_n]_n - L = [x_n - L]_n$ is infinitesimal, and $[x_n]_n = L + ([x_n]_n - L)$, concluding the proof. \square

Limits of probability measure spaces

We can use ultrafilters to construct a natural notion of 'limit' of a sequence of probability measure spaces.

Let $(X_n, \mathcal{B}_n, \mu_n)$ be a sequence of measure preserving systems. We can define a 'limit' measure preserving system $(X_\infty, \mathcal{B}_\infty, \mu_\infty)$ in the following way.

The set X_∞ will be defined by

$$X_\infty = \frac{\prod_{n \in \mathbb{N}} X_n}{\sim},$$

where $(x_n)_n \sim (y_n)_n$ if $x_n = y_n$ for \mathcal{F} -almost all n .

We denote by $[x_n]_n \in X_\infty$ the class of $(x_n)_n$.

In general, if $(X_n)_{n \in \mathbb{N}}$ is a sequence of sets, the set X_∞ defined above is known as the *ultraproduct* of $(X_n)_n$ along \mathcal{F} .

Internal subsets of X_∞

Definition

An *internal* subset of X_∞ is a subset of the form

$$\lim_{n \rightarrow \mathcal{F}} A_n := \{[x_n]_n \in X_\infty; x_n \in A_n \forall n\},$$

for some sequence of subsets $A_n \subseteq X_n$. Note that

$$\lim_{n \rightarrow \mathcal{F}} (A_n \cap B_n) = \left(\lim_{n \rightarrow \mathcal{F}} A_n \right) \cap \left(\lim_{n \rightarrow \mathcal{F}} B_n \right)$$

$$\lim_{n \rightarrow \mathcal{F}} (A_n \cup B_n) = \left(\lim_{n \rightarrow \mathcal{F}} A_n \right) \cup \left(\lim_{n \rightarrow \mathcal{F}} B_n \right)$$

$$\lim_{n \rightarrow \mathcal{F}} (X_n \setminus A_n) = X_\infty \setminus \left(\lim_{n \rightarrow \mathcal{F}} A_n \right)$$

Thus, internal subsets form an algebra of subsets of X_∞

A limit probability measure

We will let $\mathcal{B}_\infty \subseteq \mathcal{P}(X_\infty)$ be the σ -algebra generated by internal subsets of the form $\prod_{n \rightarrow \mathcal{F}} A_n$, where $A_n \in \mathcal{B}_n$. Then,

Proposition

There is a probability measure $\mu_\infty : \mathcal{B}_\infty \rightarrow [0, 1]$ such that for all sets $A_n \in \mathcal{B}_n$ ($n \in \mathbb{N}$) we have

$$\mu_\infty \left(\lim_{n \rightarrow \mathcal{F}} A_n \right) = \lim_{n \rightarrow \mathcal{F}} \mu_n(A_n).$$

We can take limits of uniformly bounded sequences of functions:

Proposition

Let $f_n : X_n \rightarrow [0, 1]$ be measurable for all n . Then, the function $f_\infty : X_\infty \rightarrow [0, 1]$ defined by $f_\infty([x_n]_n) = \lim_{n \rightarrow \mathcal{F}} f_n(x_n)$ is measurable, and satisfies

$$\int_{X_\infty} f_\infty d\mu_\infty = \lim_{n \rightarrow \mathcal{F}} \int_{X_n} f_n d\mu_n.$$

Measure preserving systems

If $T_n : X_n \rightarrow Y_n$ is measure preserving for all n , then we have a measure preserving map $T_\infty : X_\infty \rightarrow Y_\infty$ given by

$$T_\infty([x_n]_n) = [T_n x_n]_n.$$

Thus, we can take the limit of a sequence of measure preserving systems.

An application to recurrence

Definition

We say that $A \subseteq \mathbb{N}$ is a *set of recurrence* if, for any probability m.p.s. (X, \mathcal{B}, μ, T) and any $B \in \mathcal{B}$ with $\mu(B) > 0$, we have $\mu(B \cap T^a B) > 0$ for some $a \in A$.

Using limits of probability measures, it is not hard to prove that:

Proposition

If $A \subseteq \mathbb{N}$ is a set of recurrence, then for any $\varepsilon > 0$ there exists $A_0 \subseteq A$ finite and $\delta > 0$ such that for any probability m.p.s. (X, \mathcal{B}, μ, T) and any $B \in \mathcal{B}$ with $\mu(B) > \varepsilon$ we have $\mu(B \cap T^a B) > \delta$ for some $a \in A_0$.

Bibliography

Robert Goldblatt. *Lectures on the Hyperreals: An Introduction to Nonstandard Analysis*. Springer, 1998.