Constructing the Lebesgue measure

Alex Fu

2023-01-06

To rigorously define and suitably generalize the notions of length, area, and volume is one of the basic motivations for all of measure theory. Intuitively, length, area, and volume are very similar in nature — a measurement of "how many points there are" that differs from cardinality or counting — and the only notable difference between the three is the number of dimensions in which they apply. We will show that length, area, and volume are indeed just special cases of the *n-dimensional Lebesgue measure* on $\mathbb{R}^n$.

Our construction of the Lebesgue measure follows the general blueprint of starting with the most natural definition on a simple class of sets, then extending to larger classes of sets in a natural way that preserves the desired properties of a measure. The existence of nonmeasurable sets like the Vitali set implores us to be careful in this multistage development.

1 The semialgebra of half-open half-closed intervals

Throughout, let Ω be a set assumed to be nonempty for convenience. A **family** $\mathcal{F}$ on Ω is a collection of subsets of Ω, i.e. $\mathcal{F} \subseteq 2^{\Omega}$; we will also require that families are nonempty.

Definition 1 (Semialgebra).

A family $\mathcal{S}$ is a **semialgebra** on Ω if it satisfies the following.

1. *Empty set.* $\varnothing \in \mathcal{S}$.
2. *Closure under finite intersection.* If $A_1, \ldots, A_n \in \mathcal{S}$, then $\bigcap_{i=1}^n A_i \in \mathcal{S}$ as well.
3. *Semiclosure under complement.* If $A \in \mathcal{S}$, then there exist disjoint $B_1, \ldots, B_n \in \mathcal{S}$ such that the complement $A^c = \bigsqcup_{i=1}^n B_i$.

Definition 2 (Premeasure).

A set function $\mu\colon \mathcal{S} \to [0, \infty]$ defined on a semialgebra $\mathcal{S}$ is called a **premeasure** on $\mathcal{S}$ if it satisfies the following.

1. *Empty set.* $\mu(\varnothing) = 0$.

2. *Countable additivity, or σ-additivity.* If $A_1, A_2, \ldots \in \mathcal{S}$ are disjoint, and $\bigsqcup_{i=1}^{\infty} A_i \in \mathcal{S}$, then

$$\mu\left(\bigsqcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu(A_i).$$

What motivates the definitions above? The most basic property we want in the Lebesgue measure on $\mathbb{R}$ is that it measures the *length* of *intervals*:

$$\mathrm{m}((a,b)) = \mathrm{m}((a,b]) = \mathrm{m}([a,b)) = \mathrm{m}([a,b]) := b - a$$

for all $a, b \in [-\infty, +\infty]$, $a \leq b$. We can set $\mathrm{m}(\varnothing) := 0$ by definition, or by the convention of $(a,b) = \varnothing$ for $a \geq b$. We will also abuse notation slightly by writing $(a, +\infty] = (a, +\infty)$, and likewise for $[-\infty, b)$.

However, this is not yet a very meaningful definition of m. The countable union of disjoint open intervals is never an interval, and likewise for disjoint closed intervals. However, the half-open half-closed intervals fit together like puzzle pieces: the open endpoint of one interlocks perfectly with the closed endpoint of another. We can thus check that m is indeed countably additive on these intervals; otherwise, countable additivity is not even well-defined, as sets like $\bigsqcup_{i=1}^{\infty}(a_i, b_i)$ are never in the domain of $\mathrm{m}(\cdot)$.

The choice between $(a,b]$ and $[a,b)$ is mostly arbitrary; we will take $(a,b]$, since probability theory prefers sets of the form $(-\infty, x]$ for its cumulative distribution functions $F(x) := \mu((-\infty, x])$. We define

$$\mathcal{S} := \{(a,b] : a, b \in [-\infty, +\infty]\}.$$

In fact, the properties of $(a,b]$ are what motivate the definition of a semialgebra in the first place.

Proposition 1 (The half-open half-closed intervals form a semialgebra)**.**

$\mathcal{S}$ above is a semialgebra.

Proof. We perform some regular checks.

1. $\varnothing \in \mathcal{S}$ per the convention that $(a,b] = \varnothing$ for $a > b$.

2. It suffices to show closure under pairwise intersections. Let $(a_1, b_1], (a_2, b_2] \in \mathcal{S}$. Then $(a_1, b_1] \cap (a_2, b_2] = (\max\{a_1, a_2\}, \min\{b_1, b_2\}] \in \mathcal{S}$, where the intersection may be empty.

3. Let $(a,b] \in \mathcal{S}$. Then $(a,b]^c = (-\infty, a] \sqcup (b, +\infty]$ is a disjoint union of intervals in $\mathcal{S}$.

$\square$

The significance of Proposition 1 lies in a later theorem which states that premeasures on semialgebras can be extended to full-fledged measures on σ-algebras, where this extension is unique in the special case of σ-finiteness (which holds for $\mathbb{R}^n$). To invoke this result, we need the following check:

Proposition 2 (Length is a premeasure).

Let $\mathcal{S}$ be the semialgebra of half-open half-closed intervals in $\mathbb{R}$, and let $\mathrm{m}\colon \mathcal{S} \to [0,\infty]$ be given by $\mathrm{m}((a,b]) \coloneqq b-a$. Then m is a premeasure on $\mathcal{S}$.

Proof. Let us first show that m is *finitely additive* on $\mathcal{S}$. Let $A_1,\ldots,A_n \in \mathcal{S}$ be disjoint, and suppose that $A = \bigsqcup_{i=1}^n A_i \in \mathcal{S}$ as well, which allows us to write $A_i \coloneqq (a_i,b_i]$ and $A \coloneqq (a,b]$. We may further assume that $a_1 \le b_1 = a_2 \le b_2 = a_3 \le \cdots \le b_n$ by reindexing the A_i without loss of generality. Then

$$\sum_{i=1}^n \mathrm{m}(A_i) = \sum_{i=1}^n (b_i - a_i) = b_n - a_1 = b - a = \mathrm{m}(A).$$

Now, suppose that $(a,b] = \bigsqcup_{i=1}^\infty (a_i,b_i] \in \mathcal{S}$. The infinite case is tricky because we cannot rearrange the $(a_i,b_i]$ such that $a_1 \le b_1 = a_2 \le \cdots$ willy-nilly. However, we can break up the equality

$$b - a = \sum_{i=1}^\infty (b_i - a_i)$$

into two inequalities, then prove each direction separately by reducing to the finite case.

- The easier inequality is $b - a \ge \sum_{i=1}^\infty (b_i - a_i)$: it suffices to show $b - a \ge \sum_{i=1}^n (b_i - a_i)$ and pass to the limit as $n \to \infty$. For finitely many intervals, we may assume $a_1 \le b_1 = a_2 \le \cdots = a_n \le b_n$ without loss of generality. Then

  $$\sum_{i=1}^n (b_i - a_i) = b_n - \sum_{i=2}^n (a_i - b_{i-1}) - a_1 \le b_n - a_1 \le b - a,$$

  where each $a_i - b_{i-1} \le 0$ and $a \le a_1 \le b_n \le b$.

- For the harder inequality, we give ourselves an epsilon's worth of room. It suffices to show $b - a \le \sum_{i=1}^\infty (b_i - a_i) + \varepsilon$ for any $\varepsilon > 0$, as we can take $\varepsilon \downarrow 0$. We will reduce to the finite case by cleverly using the **compactness** of closed intervals in $\mathbb{R}$. First, the epsilon allows us to approximate

  $$[a', b] \subseteq (a,b] \subseteq \bigcup_{i=1}^\infty (a_i, b_i] \subseteq \bigcup_{i=1}^\infty (a_i, b_i'),$$

  where $a' = a + \frac{\varepsilon}{2}$ and $b_i' = b_i + \frac{\varepsilon}{2^{i+1}}$. We write $\varepsilon_i \coloneqq \frac{\varepsilon}{2^{i+1}}$ for convenience, noting $\sum_{i=1}^\infty \varepsilon_i = \frac{\varepsilon}{2}$. With this approximation, we invoke the compactness of $[a', b]$ to find a finite subcover such that

  $$[a', b] \subseteq \bigcup_{k=1}^n (a_{i_k}, b_{i_k}').$$

  Now we can reindex these open intervals. Let O_1 be an interval (a_{i_1}, b_{i_1}') containing a'. If $b_{i_1}' > b$, then we are done; otherwise, let O_2 be any (a_{i_2}, b_{i_2}') containing b_{i_1}', and continue inductively. This process clearly terminates in finitely many steps, resulting in $O_1, \ldots, O_n$ with *interlaced* endpoints

  $$a_{i_k} < a_{i_{k+1}} < b_{i_k}' < b_{i_{k+1}}'.$$

From here, the rest of the proof proceeds naturally.

$$\sum_{k=1}^{n}(b'_{i_k} - a_{i_k}) = b'_{i_n} - \sum_{k=2}^{n}(a_{i_k} - b'_{i_{k-1}}) - a_{i_1} \geq b'_{i_n} - a_{i_1} = b - a - \frac{\varepsilon}{2}$$

using the fact that $[a', b] \subseteq \bigcup_{k=1}^{n} O_k$ and $a' = a + \frac{\varepsilon}{2}$. Using the other approximation $b'_i = b_i + \varepsilon_i$,

$$\sum_{k=1}^{n}(b'_{i_k} - a_{i_k}) \leq \sum_{i=1}^{\infty}(b'_i - a_i) = \sum_{i=1}^{\infty}(b_i - a_i) + \frac{\varepsilon}{2}.$$

Combining these two inequalities, we are done.

$\square$

2 The algebra of finite disjoint unions

For a premeasure μ on a semialgebra $\mathcal{S}$, the natural next step is to extend μ by defining

$$\mu\left(\bigsqcup_{i=1}^{n} A_i\right) := \sum_{i=1}^{n} \mu(A_i)$$

for disjoint $A_1, \ldots, A_n \in \mathcal{S}$. This in fact extends μ to the *algebra* generated by $\mathcal{S}$, the set of all finite unions of $\mathcal{S}$.

Definition 3 (Algebra).

A family $\mathcal{A}$ is an **algebra** on Ω if it satisfies the following.

1. *Nonempty.* $\varnothing \in \mathcal{A}$.
2. *Closure under finite union.* If $A_1, \ldots, A_n \in \mathcal{A}$, then $\bigcup_{i=1}^{n} A_i \in \mathcal{A}$ as well.
3. *Closure under complement.* If $A \in \mathcal{A}$, then $A^c \in \mathcal{A}$ as well.

It is clear that $\Omega \in \mathcal{A}$, and closure under finite intersection holds by De Morgan's laws. Algebras are also semialgebras, but not the converse, though we can find a fairly canonical algebra from any semialgebra:

Proposition 3 (The algebra generated by a semialgebra).

Let $\mathcal{S}$ be a semialgebra. Then the algebra $\mathcal{A}$ generated by $\mathcal{S}$, i.e. the minimal algebra containing $\mathcal{S}$, is precisely the set of all finite (disjoint) unions of elements in $\mathcal{S}$.

Proof. Let $\mathcal{F}$ be the set of all finite unions of elements in $\mathcal{S}$. Any algebra containing $\mathcal{S}$ must contain $\mathcal{F}$, so it suffices to show that $\mathcal{F}$ is itself an algebra, in which case it is the minimal algebra containing $\mathcal{S}$.

1. $\varnothing \in \mathcal{F}$ because $\varnothing \in \mathcal{S}$, or because $\varnothing$ is the empty union.

2. A finite union of finite unions remains a finite union, so $\mathcal{F}$ is closed under finite unions.

3. Let $A \in \mathcal{F}$, where $A = \bigcup_{i=1}^n A_i$ for some $A_1, \dots, A_n \in \mathcal{S}$. By De Morgan's laws, $A^c = \bigcap_{i=1}^n A_i^c$. Now, by the properties of a semialgebra,

$$A^c = \bigcap_{i=1}^{n} \bigsqcup_{j=1}^{m_i} B_{i,j}$$

for some $B_{i,j} \in \mathcal{S}$. We leave it as an exercise in managing indices to show that A^c is a (disjoint) union of various intersections of $B_{i,j}$, which belong to $\mathcal{S}$.

Lastly, let us check that $\mathcal{F}$ is equal to the set of all finite *disjoint* unions in $\mathcal{S}$. We observe that we can *disjointize* any finite union:

$$\bigcup_{i=1}^{n} A_i = A_1 \sqcup (A_2 \setminus A_1) \sqcup (A_3 \setminus (A_1 \cup A_2)) \sqcup \cdots .$$

If $A_1, \dots, A_n \in \mathcal{S}$, then it is clear that the disjointized $B_i = A_i \setminus \bigcup_{j=1}^{i-1} A_j = A_i \cap \bigcap_{j=1}^{i-1} A_j^c$ belong to $\mathcal{S}$ as well. In other words, every $A \in \mathcal{F}$ can be expressed as a finite *disjoint* union in $\mathcal{S}$. $\square$

With Proposition 3, the extension of m to the algebra generated by $\mathcal{S}$ is well-defined. What's more, we preserve countable additivity.

Proposition 4 (Extension of premeasure to algebra remains a premeasure).

Let μ be a premeasure on the semialgebra $\mathcal{S}$, and extend μ to the algebra $\mathcal{A}$ generated by $\mathcal{S}$ by

$$\bar{\mu}\left(\bigsqcup_{i=1}^{n} A_i\right) := \sum_{i=1}^{n} \mu(A_i).$$

Then $\bar{\mu}$ is a premeasure on $\mathcal{A}$; in particular, $\bar{\mu}$ is countably additive.

Proof. Let $A_1, A_2, \dots \in \mathcal{A}$ be disjoint, and suppose $A = \bigsqcup_{i=1}^\infty A_i \in \mathcal{A}$. By Proposition 3, for each i, there exist disjoint $B_{i,1}, \dots, B_{i,n_i} \in \mathcal{S}$ such that $A_i = \bigsqcup_{j=1}^{n_i} B_{i,j}$. Then

$$\sum_{i=1}^{\infty} \bar{\mu}(A_i) = \sum_{i=1}^{\infty} \sum_{j=1}^{n_i} \mu(B_{i,j}).$$

We also observe that $A = \bigsqcup_{i,j} B_{i,j}$, so $\bar{\mu}(A) = \sum_{i,j} \mu(B_{i,j})$. By Tonelli's theorem, since the summands are all nonnegative, the two series are equal:

$$\bar{\mu}(A) = \sum_{i=1}^{\infty} \bar{\mu}(A_i).$$

$\square$

3 The Borel σ-algebra on $\mathbb{R}$

While going from a semialgebra to an algebra admits a very simple characterization, there is, in general, no constructive way to describe the σ-algebra generated by an algebra.

Definition 4 (σ-algebra).

A family Σ is a **σ-algebra** on Ω if it satisfies the following.

1. *Nonempty.* $\varnothing, \Omega \in \Sigma$.
2. *Closure under countable union.* If $A_1, A_2, \ldots \in \Sigma$, then $\bigcup_{i=1}^{\infty} A_i \in \Sigma$ as well.
3. *Closure under complement.* If $A \in \Sigma$, then $A^c \in \Sigma$ as well.

The pair (Ω, Σ) is called a **measurable space**.

By De Morgan's laws, σ-algebras are also closed under countable intersection. The σ-algebra generated by an algebra $\mathcal{A}$ is, per usual, the minimal σ-algebra containing $\mathcal{A}$.

One might ask if m is countably additive on the algebra $\mathcal{A}$ generated by $\mathcal{S}$, why not stop there? Why do we need to extend m to a σ-algebra? For one, σ-algebras are vastly richer classes of sets than algebras.

Proposition 5 (The Borel σ-algebra).

The σ-algebra generated by the set of half-open half-closed intervals $\mathcal{S}$ is called the **Borel σ-algebra** $\mathcal{B}$ on $\mathbb{R}$, and contains all open and closed sets of $\mathbb{R}$, all countable unions and intersections of those, and all countable unions and intersections of *those*, $\ldots$, so on and so forth.

Note that "all countable unions of countable intersections of $\ldots$ in $\mathcal{A}$" still does not encompass all of $\mathcal{B}$. The introduction of countable infinity to the definition clearly introduces a greater level of complexity to the sets we can measure. For this reason, and the fact that we most often want to work with *sequences* of sets, that we require countable unions and intersections to be well-defined operations, that we require measures be *countably* additive instead of simply finitely additive.

As a side note, while σ-algebras encode the notion of *measurability* in general, in probability spaces, they capture the more specific idea of *information*. A measurable set is renamed an *event*, whose probability we should be able to measure or determine. Countability is key to describing various forms of convergence, which apply to long-term behaviors even in finite time horizons.

Definition 5 (Measure).

A set function $\mu \colon \Sigma \to [0, \infty]$ defined on a σ-algebra Σ is a **measure** if it satisfies the following.

1. *Empty set.* $\mu(\varnothing) = 0$.

2. *Countable additivity*. If $A_1, A_2, \ldots \in \Sigma$ are disjoint, then

$$\mu\left(\bigsqcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu(A_i).$$

The triple (Ω, Σ, μ) is called a **measure space**.

It is nonobvious how we might extend m on $\mathcal{A}$ to an actual measure $\bar{\mathrm{m}}$ on $\mathcal{B}$ if the sets in $\mathcal{B}$, also known as *Borel sets*, do not admit any explicit description in terms of sets in $\mathcal{A}$. The key is again *approximation*: we will approximate the measure of a Borel set arbitrarily closely from above by the measures of countable covers from the algebra $\mathcal{A}$.

Theorem 1 (Carathéodory's extension theorem)**.**

Let μ be a premeasure on an algebra $\mathcal{A}$. Then there exists a measure $\bar{\mu}$ on Σ, the σ-algebra generated by $\mathcal{A}$, such that $\bar{\mu} \restriction \mathcal{A} \equiv \mu$; that is, $\bar{\mu}$ is an extension of μ. If μ is also σ-finite, then $\bar{\mu}$ is unique.

Proof. We leave the proof to a future note. □

To be covered: outer measures, the monotone class lemma, Dynkin's π-λ theorem, the completion of a σ-algebra and the Lebesgue σ-algebra, product measures, and the Lebesgue measure in n dimensions.

■