

A formally verified proof of the Central Limit Theorem (preliminary report)

Jeremy Avigad, Johannes Hölzl, and Luke Serafin

May 21, 2014

Abstract

We describe a formally verified proof of the Central Limit Theorem in the Isabelle proof assistant.

1 Introduction

If you roll a fair die many times and compute the average number of spots showing, the result is likely to be close to 3.5, and the odds that the average is far from the expected value decreases roughly as the area under the familiar bell-shaped curve. Something similar happens if the measurement is continuous rather than discrete, such as when you repeatedly toss a needle on the ground and measure the angle it makes with respect to a fixed reference line. Even if the die is not a fair die or the geometry of the needle and the ground makes some angles more likely than others, the distribution of the average still approaches the area under a bell-shaped curve centered on the expected value. The width of the bell depends on both the variance of the random measurement and the number of times it is performed. Made precise, this amounts to a statement of the Central Limit Theorem.

The theorem was proved by De Moivre in the eighteenth century in the special case where the repeated experiment is a toss of a fair coin. It was generalized by Laplace in the early nineteenth century, and later by Cauchy and others. However, the first rigorous proof of the Central Limit Theorem at level of generality hinted at above was given by Lyapunov in 1901.

Even in more rudimentary formulations, the Central Limit Theorem says, in a sense, that the behavior of random or chaotic events becomes highly predictable when averaged over the long term. In 1889, Sir Francis Galton waxed poetic about this conclusion:

I know of scarcely anything so apt to impress the imagination as the wonderful form of cosmic order expressed by the “Law of Frequency

of Error.” The law would have been personified by the Greeks and deified, if they had known of it. It reigns with serenity and in complete self-effacement, amidst the wildest confusion. The huger the mob, and the greater the apparent anarchy, the more perfect is its sway. It is the supreme law of Unreason. Whenever a large sample of chaotic elements are taken in hand and marshaled in the order of their magnitude, an unsuspected and most beautiful form of regularity proves to have been latent all along. [2, page 66]

The result lies at the heart of modern probability theory. Many generalizations and variations have been studied; for example, relaxing the requirement that the repeated measurements are independent of one another and identically distributed, or providing additional information on the rate of convergence.

In this note, we report on a formalization of the Central Limit Theorem in the Isabelle proof assistant. We feel that this formalization is interesting and valuable for a number of reasons. Not only is the Central Limit Theorem fundamental to contemporary probability theory and the study of stochastic processes, but so is almost all of the machinery developed to prove it, ranging from ordinary calculus to the properties of real distributions and characteristic functions. There is a pragmatic need to have statistical claims made in engineering, risk analysis, and financial computation subject to formal verification, and the library we have built provides infrastructure for such practical efforts.

The formalization is also a good test for Isabelle’s libraries, proof language, and automated reasoning tools. As will become clear below, the proof draws on a very broad base of facts from analysis, topology, measure theory, and probability theory, providing a useful evaluation of the robustness and completeness of the supporting libraries. Moreover, the concepts build on one another. For example, a measure is a function from a class of sets to the reals, and reasoning about convergence of measures involves reasoning about sequences of such functions. The operation of forming the characteristic function is a functional taking a measure to a function from the reals to the complex numbers, and the convergence of such functionals is used to deduce convergence of measures. Thus the conceptual underpinnings are as deep as they are broad, and working with the notions exercises Isabelle’s mechanisms for handling abstract mathematical notions.

Although the Central Limit Theorem, as formulated in the next section, has been completely verified, the proofs and supporting libraries are still under development. The proofs can be found online at <https://github.com/avigad/isabelle>. This report provides only a brief overview, and we will provide more detail in a subsequent version.

2 The measure-theoretic formulation

For our formalization, we followed Billingsley’s oft-used textbook, *Probability and Measure* [1], which provides an excellent introduction to these subjects.

Here we briefly review the core concepts and give a precise statement of the Central Limit Theorem.

A *measure space* $(\Omega, \mathcal{F})$ consists of a set Ω and a σ -algebra $\mathcal{F}$ of subsets of Ω , that is, a collection of subsets of Ω containing the empty set and closed under complements and countable unions. Think of Ω as the set of possible states of affairs, or possible outcomes of an action or experiment, and each element E of $\mathcal{F}$ as representing the set of states or outcomes in which some *event* occurs — for example, that a card drawn is a face card, or that Spain wins the World Cup. A *probability measure* μ on this space is a function that assigns a value $\mu(E)$ in $[0, 1]$ to each event E , subject to the following two conditions:

1. $\mu(\emptyset) = 0$, and
2. μ is countably additive: if (E_i) is any sequence of disjoint events in $\mathcal{F}$, $\mu(\bigcup_i E_i) = \sum_i \mu(E_i)$.

Intuitively, $\mu(E)$ is the “probability” that E occurs.

The collection $\mathcal{B}$ of *Borel subsets* of the real numbers is the smallest σ -algebra containing all intervals (a, b) . A *random variable* X on the measure space $(\Omega, \mathcal{F})$ is a measurable function from $(\Omega, \mathcal{F})$ to $(\mathbb{R}, \mathcal{B})$. Saying X is measurable means that for every Borel subset B of the real numbers, the set $\{\omega \in \Omega \mid X(\omega) \in B\}$ is in $\mathcal{F}$. Think of X as some real-valued measurement that one can perform on the outcome of the experiment, in which case, the measurability of X means that if we are given any probability measure μ on $(\Omega, \mathcal{F})$, then for any Borel set B it makes sense to talk about “the probability that X is in B .” In fact, if X is a random variable, then any measure μ on $(\Omega, \mathcal{F})$ gives rise to a measure ν on $(\mathbb{R}, \mathcal{B})$, defined by $\nu(B) = \mu(\{\omega \in \Omega \mid X(\omega) \in B\})$. A measure on $(\mathbb{R}, \mathcal{B})$ is called a *real distribution*, or, more simply, a *distribution*, and the measure ν just described is called *the distribution of X* .

If X is a random variable, the *mean* or *expected value* of X with respect to a probability measure μ is $\int X d\mu$, the integral of X with respect to μ . If m is the mean, the *variance* of X is $\int (X - m)^2 d\mu$, a measure of how far, on average, we should expect X to be from its average value.

Note that passing from μ and X to its distribution ν means that instead of worrying about the probability that some abstract event occurs, we focus more concretely on the probability that some measurement on the outcome lands in some set of real numbers. In fact, many theorems of probability theory do not really depend on the abstract space $(\Omega, \mathcal{F})$ on which X is defined, but, rather, the associated distribution on the real numbers. Nonetheless, it is often more intuitive and convenient to think of the real distribution as being the distribution of a random variable (and, indeed, any real distribution can be represented that way).

One way to define a real distribution is in terms of a *density*. For example, in the case where $\Omega = \{1, 2, 3, 4, 5, 6\}$, we can specify a probability on all the subsets of Ω by specifying the probability of each of the events $\{1\}, \{2\}, \dots, \{6\}$. More generally, we can specify a distribution μ on $\mathbb{R}$ by specifying a function f such that for every interval (a, b) , $\mu((a, b)) = \int_a^b f x \, dx$. The measure μ is

then said to be the real distribution with density f . In particular, the *normal distribution* with mean m and variance σ^2 is defined to be the real distribution with density function

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-m)^2}{2\sigma^2}},$$

a “bell shaped curve” centered at m . When $m = 0$ and $\sigma = 1$, this is called the *standard normal distribution*.

Let $X_0, X_1, X_2, \dots$ be any sequence of independent random variables, each with the same distribution μ , mean m , and variance σ^2 . Here “independent” means that the random variables $X_0, X_1, \dots$ are all defined on the same measure space $(\Omega, \mathcal{F})$, but they represent independent measurements, in the sense that for any finite sequence of events $B_1, B_2, \dots, B_k$ and any sequence of distinct indices $i_1, i_2, \dots, i_k$, the probability that X_{i_j} is in B_j for each j is just the product of the individual probabilities that X_{i_j} is in B_j . For each n , let $S_n = \sum_{i < n} X_i$. Notice that each S_n is really a measurable function on $(\Omega, \mathcal{F})$, which is to say, it is a random variable; and so it is natural to ask how its values are distributed. We can shift the expected value of S_n to 0 by subtracting nm , and scale the variance to 1 by dividing by $\sqrt{n\sigma^2}$. The Central Limit Theorem says that the corresponding quantity,

$$\frac{S_n - nm}{\sqrt{n\sigma^2}},$$

approaches the standard normal distribution as n approaches infinity.

All that remains to do is to make sense of the assertion that a sequence of distributions $\mu_0, \mu_1, \mu_2, \dots$ “approaches” a distribution, μ . For distributions that are defined in terms of densities, the intuition is that over time the graph of the density should look more and more like the graph of the density of the limit. For example, if you flip a coin a number of times and graph all the possible values of the average number of ones, the discrete points plotted over the possibilities $0, 1/n, 2/n, 3/n, \dots, 1$ start to look like a bell-shaped curve centered on $1/2$. The notion of *weak convergence* makes the notion of “starts to look like” precise.

If μ is any real distribution, then the function $F_\mu(x) = \mu((-\infty, x])$ is called the *cumulative distribution function* of μ . In words, for every x , $F_\mu(x)$ returns the likelihood that a real number chosen randomly according to the distribution is at most x . Clearly $F_\mu(x)$ is nondecreasing, and it is not hard to show that F_μ is right continuous, approaches 0 as x approaches $-\infty$, and approaches 1 as x approaches ∞ . Conversely, one can show that any such function is the cumulative distribution function of a unique measure. Thus there is a one-to-one correspondence between functions F satisfying the properties above and real distributions.

The notion of weak convergence can be defined in terms of the cumulative distribution function:

Definition 2.1. Let (μ_n) be a sequence of real distributions, and let μ be a real distribution. Then μ_n converge weakly μ , written $\mu_n \Rightarrow \mu$, if $F_{\mu_n}(x)$ approaches $F_\mu(x)$ at each point x where F_μ is continuous.

To understand why we need to exclude points of continuity of F_μ , for each n , consider the probability measure μ_n that puts all its “weight” on $1/n$, which is to say, for any Borel set B , $\mu(B) = 1$ if and only if B contains $1/n$. Then F_{μ_n} is the function that jumps from 0 to 1 at $1/n$. Intuitively, it makes sense to say that μ_n approaches the real distribution μ that puts all its weight at 0. But for every n , $F_{\mu_n}(0) = 0$, while $F_\mu(0) = 1$, which explains why want to exclude the point 0 from consideration. Notice that since F_μ is a monotone function, it can have at most countably many points of discontinuity, so we are excluding only countably many points.

That the notion of weak convergence is robust is supported by the fact that there are a number of equivalent characterizations. The following theorem is sometimes known as the *Portmanteau Theorem*:

Theorem 2.2. *The following are equivalent:*

- $\mu_n \Rightarrow \mu$
- $\int f d\mu_n$ approaches $\int f d\mu$ for every bounded, continuous function f
- If A is any Borel set, ∂A denotes the topological boundary of A , and $\mu(\partial A) = 0$, then $\mu_n(A)$ approaches $\mu(A)$.

The Central Limit Theorem can now be stated precisely as follows:

Theorem 2.3. *Let $X_0, X_1, X_2, \dots$ be a sequence of independent random variables with mean 0, variance σ^2 , and common distribution μ . Let $S_n = (X_0 + X_1 + \dots + X_{n-1})$. Then the distribution of $S_n/\sqrt{n\sigma^2}$ converges weakly to the standard normal distribution.*

The restriction to mean 0 does not result in any loss of generality, since if each X_i has mean m , we can apply Theorem 2.3 to the sequence $(X_i - m)$ to obtain the more general statement above. Our formulation of Theorem 2.3 in Isabelle is depicted as Figure 1.

3 Supporting libraries

In this section, we discuss some of the libraries of theorems in Isabelle that supported our formalization, and some of the ways we found it useful or necessary to extend those libraries.

To start with, Isabelle’s libraries include a development of the notion of a sigma-additive measure on a space, Lebesgue integration with respect to any such measure, and standard properties of integration, including the monotone convergence theorem, the dominated convergence theorem, and Fubini’s theorem. For non-finite measures, such as Lebesgue measure on the real line, measures and integrals take values in the *extended reals*, which include $\pm\infty$. Probability measures, however, are instance of finite measures, which always take values in the reals. The measure theory library also includes the definition of an independent family of sets and an independent family of random variables,

```

theorem (in prob_space) central_limit_theorem:
  fixes
    X :: "nat  $\Rightarrow$  'a  $\Rightarrow$  real" and
     $\mu$  :: "real measure" and
     $\sigma$  :: real and
    S :: "nat  $\Rightarrow$  'a  $\Rightarrow$  real"
  assumes
    X_indep: "indep_vars ( $\lambda$ i. borel) X UNIV" and
    X_integrable: " $\bigwedge$ n. integrable M (X n)" and
    X_mean_0: " $\bigwedge$ n. expectation (X n) = 0" and
     $\sigma$ _pos: " $\sigma > 0$ " and
    X_square_integrable: " $\bigwedge$ n. integrable M ( $\lambda$ x. (X n x)2)" and
    X_variance: " $\bigwedge$ n. variance (X n) =  $\sigma^2$ " and
    X_distrib: " $\bigwedge$ n. distr M borel (X n) =  $\mu$ "
  defines
    "S n  $\equiv$   $\lambda$ x.  $\sum$  i<n. X i x"
  shows
    "weak_conv_m ( $\lambda$ n. distr M borel ( $\lambda$ x. S n x / sqrt (n *  $\sigma^2$ )))
      (density lborel standard_normal_density)"

```

Figure 1: The Central Limit Theorem

and some of their basic properties. These were crucial to our formalization. This library is described in [4, 5], and served our purposes well, providing, in particular, automated methods for establishing tedious measurability claims. The formalization helped us round out the library and fill some minor omissions, adding to its robustness.

When μ is a measure, we found it useful to introduce notation for the integral $\int_A f d\mu$ over a set A . Using the characteristic function of A , this is interpreted as $\int f \cdot 1_A d\mu$. We also found it useful to restate common facts from the integration library (addition formulas, etc.) “relativized” to a set. In hindsight, it would have been smoother to take integration over a set to be a basic notion, and then take $\int f d\mu$ to be notation for integration over the entire type in question.

We also drew on Isabelle’s library of topological notions, including limits and continuity. Isabelle also has powerful library for dealing with limits, described in [6], which supports reasoning about limits of sequences and limits of functions. In the latter case, the relevant concepts are general enough to allow both the arguments and the result to approach a particular real number, or infinity, or negative infinity. Moreover, the domain and range can be restricted, for example, allowing us to reason about limits and continuity from the left or right. These notions are central to our formalization.

Isabelle library includes a good development of notions from point-set topology. We were primarily concerned topological properties of the real numbers. For example, recall that one of the equivalence characterizations of the weak convergence of measures asserts that the sequence converges on any measurable set whose boundary has measure 0. Facts about cardinality also came into play. For example, we had to prove that a monotone increasing function on the reals

has only countably many points of discontinuity, and hence, since any nonempty interval is uncountable, one can find a point in any such interval at which the function is continuous.

Isabelle’s analysis library includes a theory of differentiation, and standard properties of sine, cosine, and exponentiation, and their Taylor series expansions. There is also a development of multivariate analysis, based on John Harrison’s development of the notions in HOL light. This includes a formalization of the Fréchet derivative for arbitrary real normed vector spaces, and the gauge integral on any Euclidean space. We only needed the univariate versions. Note that this means that there are two distinct versions of integration, since there is also the measure-theoretic notion, defined for functions from any measure space to the reals. There is some overlap, however; for example, the library shows that the two integrals, taken over a finite interval, are equal for Riemann integrable functions on that interval.

For our formalization, we had to make extensive use of ordinary calculus, for example, calculating $\int_0^\infty (\sin x)/x \, dx$ and $\int_0^\infty x^k e^{-x^2} \, dx$ for every k . These are tricky because the limits of integration make them “improper” integrals. Text-book presentations of these calculations rely explicitly or implicitly on various appeals to limits, and one has to take care to ensure that expressions in questions are integrable. Both calculations involve changes of variable (the “substitution theorem” from calculus) and integration by parts, and the first calculation relies on Fubini’s theorem. To support the calculations, we defined the notion of integration over an interval, $\int_a^b f(x) \, dx$, with two interesting twists: a and b are extended reals, which is to say, they can be $\pm\infty$, and also b is allowed to be less than a , in which case $\int_a^b f(x) \, dx = -\int_b^a f(x) \, dx$. We developed versions of the fundamental theorem of calculus, the substitution lemma, and so on, compatible with this notion. In particular, our library supports arguments where one calculates an integral $\int_a^b f(x) \, dx$, and then allows a and b to approach the endpoints of the interval of interest (for example, letting them go off to $\pm\infty$, respectively). (For our purposes, it was sufficient to have the substitution lemma for integrals of continuous functions, but Manuel Eberl has recently generalized the result to arbitrary measurable functions.)

We also extended the library for measure theory and probability to include many of the facts mentioned in the last section. Establishing the correspondence between real distributions and cumulative distribution functions required a compactness argument and the Carathéodory extension theorem. Proving the harder implications of the Portmanteau Theorem required proving *Skorohod’s theorem*, a generally useful fact that shows that under general conditions it is possible to represent a sequence of measures on a common space. This required a long and delicate argument.

Finally, as will become clear in the next section, we also need notions of integration and differentiation for functions from $\mathbb{R}$ to $\mathbb{C}$, the complex numbers. This is a starting point for the development of complex analysis (e.g. as carried out by Harrison [3]), but it is much simpler: the integral of a function $f : \mathbb{R} \rightarrow \mathbb{C}$ is simply defined in terms of the integral of its real and imaginary parts, and

similarly for the derivative. Since the complex numbers can be viewed as a vector space over the reals, we could show that the natural definition of the complex derivative of a function $f : \mathbb{R} \rightarrow \mathbb{C}$ coincides with derivative of f considered as function from $\mathbb{R}$ to a real-valued vector space. For the integral, however, we had to define the new notion and “import” basic properties of real-valued integrals to the complex version (in addition, of course, to establishing properties that are specific to the complex integral). In fact, the three notion of integral we have considered here (the multivariate version, the version for functions from an arbitrary measure space to the reals, and the version for functions from an arbitrary measure space to the complex numbers) admit a common generalization, which would unify the library substantially.

4 Overview of the proof

Contemporary proofs of the Central Limit Theorem rely on the use of *characteristic functions*, a powerful method that dates back to Laplace. If μ is a real-valued distribution, its characteristic function $\varphi(t)$ is defined by

$$\varphi(t) = \int_{-\infty}^{\infty} e^{itx} \mu(dx).$$

In words, $\varphi(t)$ is the integral of the function $f(x) = e^{itx}$ over the whole real line, with respect to μ . Notice that for each $t \neq 0$, the function e^{itx} is periodic with period $2\pi/t$. It might be helpful to think of e^{itx} , as a function of x , as like a sine or cosine in x whose period depends on t . (Indeed, $e^{itx} = \cos(tx) + i \sin(tx)$.) Notice that $\varphi(0) = 1$, the measure of the entire real line. The characteristic function of a real distribution μ is a Fourier transform of the measure μ , and when $t \neq 0$, $\varphi(t)$ “detects” periodicity in the way that the real distribution μ distributes its “weight” over different parts of the real line.

The *Levy Uniqueness Theorem* asserts that if μ_1 and μ_2 have the same characteristic function, then $\mu_1 = \mu_2$. In other words, a measure μ can be “reconstructed” from its characteristic function, and so the characteristic function of a measure determines the measure uniquely. To prove the uniqueness theorem, it is enough to show that $\mu_1((a, b]) = \mu_2((a, b])$ for real numbers a and b with the property that $\mu_1(\{a\}) = \mu_2(\{a\}) = \mu_1(\{b\}) = \mu_2(\{b\}) = 0$. The proof involves approximating the indicator function of $(a, b]$ in terms of the complex exponential (this is where the function $(\sin x)/x$ comes into play), unfolding the definition of the characteristic function, and doing an explicit calculation with integrals and limits.

Let (μ_n) be a sequence of distributions, where each μ_n has characteristic function φ_n , and let μ be a distribution with characteristic function φ . The *Levy Continuity Theorem* states that φ_n converges to φ weakly if and only if $\varphi_n(t)$ converges to $\varphi(t)$ for every t . Proving the “only if” direction is easy, but the other direction is a lot harder. It relies on a result known as *Helly’s Selection Theorem*, a form of compactness for the space of distributions.

Remember that the CLT asserts that if (X_n) is a sequence of random variable satisfying certain hypotheses, and, for each n , μ_n is a certain distribution defined in terms of $X_1, \dots, X_n$, then μ_n converges weakly to the standard normal distribution. The considerations of the previous paragraph provide a simple strategy for proving the theorem: writing the characteristic function of μ_n as φ_n for each n , we need only show that φ_n approaches the characteristic function of the standard normal distribution pointwise.

To implement this strategy, one needs to know that the characteristic function of the standard normal distribution is $\varphi(t) = e^{-t^2/2}$. Establishing this fact took more work than we thought it would. Many textbook proofs of this invoke facts from complex analysis that were unavailable to us. Billingsley [1, page 344] sketches an elementary proof, which required calculating the moments and absolute moments of the standard normal distribution. (This is where the calculations of $\int_{-\infty}^{\infty} x^k e^{-x^2/2} dx$, mentioned in the last section, were needed. A prior calculation by Sudeep Kanav covered the cases $k = 0, 1$, which provide the base cases for an inductive proof.) Filling in the details involved carrying out careful computations with integrals and power series approximations to e^x .

The other component to implementing the strategy involves computing the characteristic functions of the distributions μ_n , which are defined in terms of finite sums of the random variables $X_0, X_1, \dots$. This is where the utility of characteristic functions becomes apparent. If X is a random variable, define the characteristic function of X to be the characteristic function of its distribution. A straightforward calculation involving properties of independent random variables shows that if $X_1, \dots, X_n$ are independent random variables then the characteristic function of $X_1 + \dots + X_n$ is the product of the characteristic functions of each X_i . This facilitates the computation of the characteristic functions of the distributions μ_n .

Once all these components are were in place, putting the pieces together was not hard. Given the continuity theorem, the characteristic function of the standard normal distribution, the result on the characteristic functions of sums of random variables, and a power series approximation to the complex exponential function, the proof of the Central Limit Theorem is quite short. In our formalization, it was only about 120 lines long.

It is a mistake to judge the complexity of a formalization by the number of lines. For one thing, different styles of writing proofs can compress or expand the number of lines dramatically. Moreover, length is not a good proxy for difficulty: some proofs are short but hard (for example, because they involve combining heterogeneous facts in tricky ways) and others are long and easy. For what it is worth, however, the files in our repository contain about 13,000 lines.

References

- [1] Patrick Billingsley. *Probability and measure*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, third edition, 1995. A Wiley-Interscience Publication.

- [2] Francis Galton. *Natural inheritance*. Macmillan, London, 1889.
- [3] John Harrison. Formalizing basic complex analysis. *Studies in Logic, Grammar and Rhetoric*, 10(23):151–165, 2007.
- [4] Johannes Hölzl. *Construction and Stochastic Applications of Measure Spaces in Higher-Order Logic*. PhD thesis, Institut für Informatik, Technische Universität München, October 2012.
- [5] Johannes Hölzl and Armin Heller. Three chapters of measure theory in Isabelle/HOL. In Marko C. J. D. van Eekelen, Herman Geuvers, Julien Schmaltz, and Freek Wiedijk, editors, *Interactive Theorem Proving (ITP 2011)*, volume 6898 of *LNCS*, pages 135–151, 2011.
- [6] Johannes Hölzl, Fabian Immler, and Brian Huffman. Type classes and filters for mathematical analysis in Isabelle/HOL. In Sandrine Blazy, Christine Paulin-Mohring, and David Pichardie, editors, *Interactive Theorem Proving (ITP 2013)*, volume 7998 of *LNCS*, pages 279–294. Springer, 2013.