Department of Mathematical and Computational Sciences

Teaching Portfolio

University of Toronto

Nigel Petersen

Fall 2020 - Spring 2024

Contents

1 Probability and Statistics I

1.1 Probability and Solving Problems with R

Many of the problems that we'll encounter in our study of probability lend themselves well to *simulation*. Namely, the context of the problem is one that we can easily model by using R, and we can run a large number of trials to compute the approximate probability of success. This idea works when we think about the probability of an event as a long term average of the successful outcomes of "infinitely-many" trials. For now, we can start with something simple

Example 1. Calculate the probability of a flipping a fair coin 5 times and seeing exactly 4 heads.

A single trial of this experiment consists of flipping a fair coin 5 times, we will then record the number of heads we observe and in our code, we can check to see if the number of observed heads is exactly 4. If so, we will consider this particular trial a success. In our program, we can be a bit more general, using parameters for the probability of obtain heads upon flipping a coin, the number of flips in each trial and the number of heads we're looking for in each trial. For now, let's use a default of $10,000$ trials of our experiment.

```
CoinFlippingSimulation <- function(p, num_flips, num_heads, num_trials){
  run = 1
  prop= 0
  while(run <= num_trials){
    coin_flips = sample(c("H", "T"), size = num_flips, replace = T, prob = c(p, 1-p))
    if(length(which(coin_flips == "H")) == num_heads){
      prop = prop+ 1
    }
    run = run + 1
  }
  return(prop/num_trials)
}
```

```
> CoinFlippingSimulation(p = 0.5, num_flips = 5, num_heads = 4, num_trials = 10000)
[1] 0.1563
```

Of course, we can compute the exact value of this probability fairly easily. Suppose that p is the probability of flipping our coin and obtaining heads, n is the number of coin flips in a trial and h is the number of observed heads necessary for a successful trial. Let X denote the number of heads observed after flipping our coin n many times. Then

$$\mathbb{P}(X = h) = \binom{n}{h} p^h (1-p)^{n-h}$$

Namely $X \sim \text{Binomial}(n, p)$ and so in our original case where $p = \frac{1}{2}, n = 5, h = 4$, the probability of observing exactly 4 heads among 5 tosses of a fair coin is exactly

$$\binom{5}{4} \frac{1}{2^5} = 0.15625$$

Let's have a look at a slightly less trivial experiment, one that requires us to use a little bit more of our probability knowledge.

Example 1.1.1

Suppose we have a jar with 4 red balls and 6 black balls. 3 balls are removed from the jar at random, and their colours are unknown. What is the probability that the next ball drawn is red?

Solution. We can take cases on the colours of the 3 balls that we removed initially, and use the total law of probability. Let A_i denote the event that exactly i of the 3 balls removed initially are red, and let B denote the event of the next ball being red. Then

$$\mathbb{P}(B) = \sum_{i=0}^{3} \mathbb{P}(B|A_i)\mathbb{P}(A_i)$$

If i red balls removed initially, then there are $4-i$ red balls left in the jar, and so $\mathbb{P}(B|A_i) = \frac{4-i}{7}$. The number of ways to remove i red balls and $3-i$ black balls from the jar is simply $\binom{4}{i}\binom{6}{3-i}$, and as there are $\binom{10}{3}$ total ways to remove 3 balls from the jar, we have

$$\mathbb{P}(B) = \sum_{i=0}^{3} \mathbb{P}(B|A_i)\mathbb{P}(A_i) = \sum_{i=0}^{3} \left[\frac{4-i}{7}\right] \frac{\binom{4}{i}\binom{6}{3-i}}{\binom{10}{3}} = 0.4 \qquad \square$$

To test that our answer is correct, we can conduct this experiment several times, and determine the proportion of trials where the last ball drawn is Red. We should expect that our estimated proportion will be close to the true proportion we found above.

```
JarSamplingSimulation <- function(red, black, initial_draws, trials){
  prop = 0
  trial = 1
  while(trial <= trials){
    jar  = sample(c(rep("R", red), rep("B", black)))
    remaining_jar = sample(jar, size = length(jar) - initial_draws)
    draw = sample(remaining_jar, size = 1)
    if(draw == "R"){
      prop = prop + 1
    }
    trial = trial + 1
  }
  return(prop/trials)
}
```

```
> JarSamplingSimulation(red = 4,black = 6, initial_draws = 3, num_trials = 10000)
[1] 0.3933
```

And so with $10,000$ trials of the experiment, we see that our success ratio is approximately our exact probability. Those interested may wish to run the above code with a differing number of trials to see how it influences the accuracy of our estimate (of course, more trials will give us more accuracy). We'll be able to make the idea behind this approximation more precise at the end of the course when we discuss sequences of random variables and the weak law of large numbers. For now, let's work through one final example, this one a little bit more complicated.

Example 1.1.2: Plane Ticket Problem

Suppose you've booked a flight on a plane with 10 total seats, and each seat has exactly one passenger who has booked it. When everyone lines up to board the plane, you are the last person in line and the first person in line loses their ticket! If the first person boarding chooses a seat at random, and everyone boarding after them either sits in their assigned seat if it is available or chooses a new seat at random if it is not available, what is the probability that you sit in your assigned seat?

Before we dive into the solution, we can phrase the problem a bit more generally so that we can build up to our desired solution by considering smaller cases. We'll attempt to solve the problem in the case that there are $n \in \mathbb{N}$ seats on the plane and exactly one passenger per seat.

Solution. Certainly if there are only $n = 2$ seats on the plane, then with probability $\frac{1}{2}$ the first person boarding will not pick your seat. When $n = 3$, you'll be able to sit in your seat if the first person chooses their original seat, or the two people in front of you choose each others seats. Namely, the first person chooses their seat correctly with probability $\frac{1}{3}$, and they choose the second persons seat with probability $\frac{1}{3}$. From here, the person boarding before you must choose the remaining seat that is not yours, and they do so with probability $\frac{1}{2}$. If we define θ_n as the probability you sit in your assigned seat when there is $n \in \mathbb{N}$ seats on the plane, then

$$\theta_3 = \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{2} = \frac{1}{2}$$

When $n = 4$, things become more complicated, but a pattern starts to become apparent. The first person can sit in their assigned seat with probability $\frac{1}{4}$, but with probability $\frac{1}{4}$ they sit in the seat of the next person boarding. Notice that in this case, we are in the setting of the problem when $n = 3$, and the second person in line now has to make a choice, where their "correct" choice of seat is not their own, but the seat of the person who took theirs. So we will be successful with probability θ_3. Similarly, if the first person boarding sits in the seat assigned to the third person boarding, then the second person will sit in their assigned seat with probability 1 and the third person is now in the setting of the problem when $n = 2$, where their "correct" choice of seat is the seat corresponding to the first person boarding, so we will be successful with probability $\theta_2 = \frac{1}{2}$. Combining these cases together, we have

$$\theta_4 = \frac{1}{4} + \frac{1}{4}\theta_3 + \frac{1}{4}\theta_2 = \frac{1}{4} + \frac{1}{4} \cdot \frac{1}{2} + \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{2}$$

Aha! A pattern has started to emerge! If we enumerate the passengers boarding as $p_1, p_2, \ldots, p_n$ where you are p_n, the last person boarding, then if p_1 chooses the seat assigned to p_k, everyone in between them (namely $p_2, \ldots, p_{k-1}$) will sit in their assigned seat with probability 1, and p_k is now in the setting of the original problem with $n - k + 1$ seats, and passengers $p_k, p_{k+1}, \ldots, p_n$. It then follows that we would be successful in this case with probability $\frac{1}{n}\theta_{n-k+1}$. Thus, if there are n seats on the plane, by the total law of conditional probability

$$\begin{aligned}\theta_n &= \sum_{k=1}^{n-1} \mathbb{P}(p_n \text{ sits in their assigned seat} | p_1 \text{ sits in the seat for } p_k)\mathbb{P}(p_1 \text{ sits in the seat for } p_k) \\ &= \frac{1}{n} + \frac{1}{n}\theta_2 + \frac{1}{2}\theta_3 + \cdots + \frac{1}{n}\theta_{n-2} + \frac{1}{n}\theta_{n-1}\end{aligned}$$

From what we calculated earlier, it seems that $\theta_n = \frac{1}{2}$ for all $n \in \mathbb{N}$, where we saw this explicitly for $n = 2, 3$ and 4, so let's try to prove this using strong induction. Our base case is taken care of, so let's assume that $\theta_k = \frac{1}{2}$ for $k = 2, 3 \ldots, n-1$, and consider the case of n total seats. By the formula above, it follows that

$$\theta_n = \frac{1}{n} + \frac{1}{n}\theta_2 + \frac{1}{2}\theta_3 + \cdots + \frac{1}{n}\theta_{n-2} + \frac{1}{n}\theta_{n-1} = \frac{1}{n}\left(1 + \frac{n-2}{2}\right) = \frac{1}{2} \qquad \square$$

We've figured it out! To our surprise (or at the very least, to my surprise), the probability you end up sitting in your correctly assigned seat is $\frac{1}{2}$, no matter how many seats are on the plane! Let's see if we can verify this by using simulation. We'll make things a bit more interesting in our simulation study, and I encourage you to try this generalization for yourself. We'll assume that there is a differing number of passengers and seats on the plane (of course, with more seats than passengers). Your take home task will be to determine a closed form solution to the problem in the case of n passengers and $m \geq n$ seats on the plane. For those very interested in the problem, try to write your own simulation code to account for a selection scheme used by passengers when they have to choose a seat at random. For example, suppose they can either choose their assigned seat now with probability p and make a selection of a random available seat with probability $1 - p$. For now, let's focus on verifying our solution to the original problem.

```
PlaneSeatingSimulation <- function(num_passengers, num_seats, num_trials){
  if(num_seats < num_passengers){
    return("There are not enough seats")
  }
  run = 1
  prop= 0
  while(run <= num_trials){
    passengers = c(seq(1, num_passengers))
    seats = c(rep(0, num_seats))
    seats[sample(c(seq(1, num_seats)), 1)] = 1
    passengers = passengers[2:length(passengers)]
    while(length(passengers) > 1){
      if(seats[passengers[1]] == 1){
        choice = sample(which(seats==0), 1)
        seats[choice] = 1
      }
      else{
        seats[passengers[1]] = 1
      }
      passengers = passengers[2:length(passengers)]
    }
    if(seats[num_passengers] == 0){
      prop = prop+ 1
    }
    run = run + 1
  }
  return(prop/num_trials)
}
```

```
> PlaneSeatingSimulation(num_passengers = 10, num_seats = 10, num_trials = 10000)
[1] 0.4925
```

1.2 Sequences of Random Variables and Convergence

Traditionally, this is a topic that always confuses students the most, and there's a few reasons for this. First and foremost, it probably shouldn't be taught in this course, especially a last minute topic to squeeze in before the end of the semester (though this is just a tinfoil hat theory of mine). Secondly and more importantly, why exactly it's useful will be completely lost on you as a student until you learn more about estimation and inference, which usually comes in the continuation of this course. Nevertheless, I'll try to shed some light on how things work, why this could potentially be useful, and we'll get our hands dirty with some examples and counterexamples.

1.2.1 The Basics

A natural place to start is with the necessary definitions of convergence. A simple place to start would be what exactly a sequence of random variables is. Like with any sequence, we can think of it as either a countable collection of objects, or as a function from $\mathbb{N}$ to another set. For the sake of simplicity, we will say that a *sequence of random variables* is a collection $\{X_n\}_{n\in\mathbb{N}}$ of random variables indexed by $\mathbb{N}$, namely for every $n \in \mathbb{N}$, there is an associated random variable X_n.

Definition 1.2.1: Types of Convergence

Let (X_n) be a sequence of random variables, and X another random variable. We write F_n and F to denote the cumulative distribution functions of each X_n, and F, respectively. We say the sequence (X_n) of random variables converges to X:

1. *in probability*, written $X_n \xrightarrow{p} X$, if $\lim_{n\to\infty} \mathbb{P}(|X_n - X| > \varepsilon) = 0$ for every $\varepsilon > 0$.
2. *in distribution*, written $X_n \xrightarrow{d} X$, if $\lim_{n\to\infty} F_n(x) = F(x)$ for every point of continuity x of F.
3. *in quadratic mean*, written $X_n \xrightarrow{qm} X$, if $\lim_{n\to\infty} \mathbb{E}[(X_n - X)^2] = 0$.

For those more analysis inclined, you may recognize convergence in distribution as a particular type of pointwise convergence. Namely if $U \subseteq \mathbb{R}$ is the collection of points for which F is continuous, then $X_n \xrightarrow{d} X$ if and only $F_n \to F$ pointwise on U. Alternatively, we will adopt the following notational convention for convergence to a "constant", we write $X_n \to c$ to mean that (X_n) converges to a random variable X for which $\mathbb{P}(X = c) = 1$ (such a random variable is called degenerate)

Example 1.2.2

Suppose (X_n) is a sequence of Uniform $(0, 1)$ random variables, and set

$$X_{(n)} = \max\{X_1, \dots, X_n\} \quad \text{and} \quad Y_n = n(1 - X_{(n)})$$

Show that $X_{(n)} \xrightarrow{p} 1$ and $Y_n \xrightarrow{d} Y \sim \text{Exponential}(1)$

Proof. Let's start by showing $X_{(n)} \xrightarrow{p} 1$. Note that as each $X_n \sim$ Uniform $(0, 1)$, we have $|X_{(n)} - 1| \leq 1$, and so if $\varepsilon \geq 1$, it follows immediately that $\mathbb{P}(|X_{(n)} - 1| > \varepsilon) = 0$. Let $\varepsilon \in (0, 1)$, once again as

$X_{(n)} > 1+\varepsilon$ is not possible, we have that

$$\begin{aligned}\mathbb{P}(|X_{(n)}-1| > \varepsilon) &= \mathbb{P}(X_{(n)}-1 > \varepsilon) + \mathbb{P}(1-X_{(n)} > \varepsilon)\\ &= \mathbb{P}(X_{(n)} < 1-\varepsilon)\\ &= \mathbb{P}(X_1 < 1-\varepsilon, \ldots, X_n < 1-\varepsilon)\\ &\overset{\text{ind.}}{=} (1-\varepsilon)^n\end{aligned}$$

Since $\varepsilon \in (0,1)$, we have $|\varepsilon - 1| < 1$ and so $(1-\varepsilon)^n \to 0$. Thus, by the squeeze theorem, we conclude that $X_{(n)} \xrightarrow{p} 1$. Now let's show that $Y_n \xrightarrow{d} Y$ where $Y \sim$ Exponential(1). We know that $F_Y(y) = 1 - e^{-y}$ for $y > 0$. It suffices to find the CDF for each Y_n. Let $y > 0$, we can write

$$\begin{aligned}\mathbb{P}(Y_n \leq y) = \mathbb{P}(n(1-X_{(n)}) \leq y) = \mathbb{P}(X_{(n)} \geq 1 - \tfrac{y}{n}) &\\ = 1 - \mathbb{P}(X_{(n)} \leq 1 - \tfrac{y}{n}) &\\ = 1 - \mathbb{P}(x_1 \leq 1 - \tfrac{y}{n}, \ldots, X_n \leq 1 - \tfrac{y}{n}) &\\ \overset{\text{ind}}{=} 1 - (1 - \tfrac{y}{n})^n &\end{aligned}$$

and so as $n \to \infty$, using properties of e, we have $F_n(y) = (1-\frac{y}{n})^n \to 1 - e^{-y} = F(y)$. Of course, when $y \leq 0$, the result follows immediately as each $X_n \in (0,1)$. Thus, $Y_n \xrightarrow{d} Y$ as needed. □

Lemma 1.2.3

Let X and Y be random variables, and $y \in \Omega_Y$, the sample space of Y. Then for any $\varepsilon > 0$

$$\mathbb{P}(Y \leq y) \leq \mathbb{P}(|X - Y| > \varepsilon) + \mathbb{P}(X \leq y + \varepsilon)$$

Proof. The idea behind the proof is to partition the event $Y \leq y$ by considering $Y \leq y, |X-Y| > \varepsilon$ and $Y \leq y, |X - Y| \leq \varepsilon$. Indeed, let $\varepsilon > 0, y \in \Omega_Y$, by the total law of probability

$$\begin{aligned}\mathbb{P}(Y \leq y) &= \mathbb{P}(Y \leq y, |X-Y| > \varepsilon) + \mathbb{P}(Y \leq y, |X-Y| \leq \varepsilon)\\ &\leq \mathbb{P}(|X-Y| > \varepsilon) + \mathbb{P}(Y \leq y, X - Y \geq \varepsilon, X - Y \leq \varepsilon)\\ &\leq \mathbb{P}(|X-Y| > \varepsilon) + \mathbb{P}(Y \leq y, X \leq Y + \varepsilon)\\ &\leq \mathbb{P}(|X-Y| > \varepsilon) + \mathbb{P}(X \leq y + \varepsilon)\end{aligned}$$

□

Our goal for now will be to establish a relationship between these types of convergence. There is a special case we'll take a look at, namely when our limiting random variable is degenerate, but the strength of our convergence types can be summarized by the following theorem:

Theorem 1.2.4

Let (X_n) be a sequence of random variables, and X a random variable. Then the following hold.

1. If $X_n \xrightarrow{qm} X$, then $X \xrightarrow{p} X$.
2. If $X_n \xrightarrow{p} X$, then $X \xrightarrow{d} X$.
3. If $c \in \mathbb{R}$ and $X_n \xrightarrow{d} c$, then $X_n \xrightarrow{p} c$.

Proof. (1) This follows as a direct consequence of the squeeze theorem and the relevant definitions of convergence. Let $\varepsilon > 0$, by Markov's inequality

$$0 \leq \mathbb{P}(|X_n - X| > \varepsilon) = \mathbb{P}((X_n - X)^2 > \varepsilon^2) \leq \frac{\mathbb{E}[(X_n - X)^2]}{\varepsilon^2}$$

and the result follows as $X_n \xrightarrow{qm} X$, and hence the right side tends to 0.

(2) Let x be a point of continuity of $F = F_X$, the CDF of X, and $\varepsilon > 0$. Our goal will be to trap $F_n(x)$ between $F(x - \varepsilon)$ and $F(x + \varepsilon)$. Let $n \in \mathbb{N}$, by first applying Lemma 1.2.3 to X_n and X we have

$$\mathbb{P}(X_n \leq x) \leq \mathbb{P}(X \leq x + \varepsilon) + \mathbb{P}(|X_n - X| > \varepsilon)$$

A similar application of the lemma using $x - \varepsilon$ in place of x yields

$$\mathbb{P}(X \leq x - \varepsilon) \leq \mathbb{P}(X_n \leq x) + \mathbb{P}(|X_n - X| > \varepsilon)$$

Rewriting leaves us with $\mathbb{P}(X_n \leq x) \geq \mathbb{P}(X \leq x - \varepsilon) - \mathbb{P}(|X_n - X| > \varepsilon)$. Finally, writing everything in terms of F_n and F, we have

$$F(x - \varepsilon) - \mathbb{P}(|X_n - X| > \varepsilon) \leq F_n(x) \leq F(x + \varepsilon) + \mathbb{P}(|X_n - X| > \varepsilon)$$

As $X_n \xrightarrow{p} X$, we have $\mathbb{P}(|X_n - X| > \varepsilon) = 0$ as $n \to \infty$, and so by the Squeeze Theorem

$$F(x - \varepsilon) \leq \lim_{n\to\infty} F_n(x) \leq F(x + \varepsilon)$$

As $\varepsilon > 0$ was arbitrary, it follows that $F_n(x) \to F(x)$ and hence $X_n \xrightarrow{d} X$.

(3) Let $\varepsilon > 0$, and suppose $X \equiv c$, namely $\mathbb{P}(X = c) = 1$. As the cdf of X is $F(x) = \mathbb{I}(x \geq c)$ and $X_n \xrightarrow{d} X$, it follows that

$$\begin{aligned} \lim_{n\to\infty} \mathbb{P}(|X_n - c| > \varepsilon) &= \lim_{n\to\infty} \mathbb{P}(X_n > c + \varepsilon) + \lim_{n\to\infty} \mathbb{P}(X_n < c - \varepsilon) \\ &= 1 - F(c + \varepsilon) + F(c - \varepsilon) \\ &= 0 \end{aligned}$$

□

Conditions 2 and 3 together say that $X_n \xrightarrow{p} c$ if and only if $X_n \xrightarrow{d} c$. It's worth noting that the reverse implication does *not* generally hold in the case the limit distribution is non-degenerate. In fact, the converses of statements 1 and 2 from Theorem 1.2.4 are both not true! We can construct explicit counterexamples to both to show that each convergence type is strictly stronger than the previous.

Example 1.2.5: Probability but not Quadratic Mean

Convergence in probability is strictly weaker than convergence in quadratic mean.

Proof. Let $U \sim \text{Uniform}\,(0,1)$ and for $n \in \mathbb{N}$ define $X_n = \sqrt{n}\mathbb{I}(0 < U < \frac{1}{n})$. We claim $X_n \xrightarrow{p} 0$ yet X_n does not converge to 0 in quadratic mean. First let's show $X_n \xrightarrow{p} 0$. As the indicator function takes on values in $\{0, 1\}$ only, given $\varepsilon > 0$, for any $n \in \mathbb{N}$ with $n > \varepsilon^2$ (so that $\frac{\varepsilon}{\sqrt{n}} < 1$), we have

$$\mathbb{P}(|X_n| > \varepsilon) = \mathbb{P}(\mathbb{I}(0 < U < \tfrac{1}{n}) > \tfrac{\varepsilon}{\sqrt{n}}) = \mathbb{P}(0 < U < \tfrac{1}{n}) = \tfrac{1}{n}$$

Thus, $\mathbb{P}(|X_n| > \varepsilon) = \frac{1}{n} \to 0$ and so $X_n \xrightarrow{p} 0$. Conversely,

$$\begin{aligned}
\mathbb{E}(|X_n|^2) = \mathbb{E}(n\mathbb{I}^2(0 < U < \tfrac{1}{n})) &= n\mathbb{E}(\mathbb{I}(0 < U < \tfrac{1}{n})) \\
&= n\int_0^1 \mathbb{I}(0 < U < \tfrac{1}{n}) f_U(u)\,\mathrm{d}u \\
&= n\int_0^{\frac{1}{n}} \mathrm{d}u \\
&= 1
\end{aligned}$$

and so X_n fails to converge to 0 in quadratic mean. □

Example 1.2.6: Distribution but not Probability

Convergence in distribution is strictly weaker than convergence in probability.

Proof. Let $X \sim \text{Uniform}\,[-1, 1]$, given $k \in \mathbb{N}$ we set $X_{2k-1} = X$ and $X_{2k} = -X$. We first claim that $X_n \xrightarrow{d} X$. Indeed, as $X \sim \text{Uniform}[-1, 1]$, the CDF of X is

$$F(x) = \begin{cases} 0 & x < -1 \\ \frac{x+1}{2} & -1 \le x < 0 \\ 1 & x \ge 0 \end{cases}$$

We can start first with $x \in [-1, 0)$. Given $k \in \mathbb{N}$, we have $\mathbb{P}(X_{2k-1} \le x) = \mathbb{P}(X \le x) = F(x)$ by construction. Moreover

$$\mathbb{P}(X_{2k} \le x) = \mathbb{P}(-X \le x) = 1 - \mathbb{P}(X \le -x) = 1 - \frac{(-x+1)}{2} = \frac{x+1}{2} = F(x)$$

As $\mathbb{P}(X_{2k} \le x) = 1 - \mathbb{P}(X \le -x)$, it follows that

$$x < -1 \implies \mathbb{P}(X_{2k} \le x) = 0 \quad \text{and} \quad x > 1 \implies \mathbb{P}(X_{2k} \le x) = 1$$

Thus, $\mathbb{P}(X_n \le x) = \mathbb{P}(X \le x)$ for all $n \in \mathbb{N}$ and $x \in \mathbb{R}$, namely $X_n \xrightarrow{d} X$. To show that X_n fails to converge in probability to X, we can consider the even-indexed terms. Let $\varepsilon = \frac{1}{2} > 0$. For any $k \in \mathbb{N}$, by definition we have

$$\begin{aligned}
\mathbb{P}(|X_{2k} - X| > \tfrac{1}{2}) = \mathbb{P}(|X| > \tfrac{1}{4}) &\ge \mathbb{P}(X > \tfrac{1}{4}) \\
&= 1 - F(\frac{1}{4}) \\
&= \frac{3}{8}
\end{aligned}$$

and so X_n cannot converge in probability to X. □

1.2.2 A Simple Application

We'll end things off by recalling a major result in the course and putting it to good use in a familiar setting. It's worth noting that the world of Computational Statistics is far more vast than I've introduced here, and the problem of accurately estimating the value of a definite integral, among other interesting and useful problems, is one that is explored in further depth in our computational stats course STA312. I highly recommend! For now, let's recall the Weak Law of Large Numbers and use it to estimate various types of definite integrals.

Theorem 1.2.7: Weak Law of Large Numbers

If (X_n) is sequence of independent and identically distributed random variables with mean μ and variance $\sigma^2 < \infty$, then $\overline{X}_n \xrightarrow{p} \mu$.

Proof. This follows as a pretty quick consequence of Markov's inequality. Indeed, let $\varepsilon > 0$, Applying Markov's inequality to $|\overline{X}_n - \mu|$, it follows that

$$0 \leq \mathbb{P}(|\overline{X}_n - \mu| > \varepsilon) \leq \frac{\mathbb{E}(\overline{X}_n - \mu)^2}{\varepsilon^2} = \frac{\sigma^2}{n\varepsilon^2}$$

and the result follows by the squeeze theorem as $\frac{\sigma^2}{n\varepsilon^2} \to 0$ as $n \to \infty$. $\square$

Finally, we'll recall the continuous mapping theorem and supply a proof for those interested. Let $U \subseteq \mathbb{R}$ containing a point $c \in \mathbb{R}$, formally speaking, a function $f : U \to \mathbb{R}$ is continuous at c if for every $\varepsilon > 0$, there is $\delta > 0$ such that for any $x \in U$; if $|x - c| < \delta$, then $|f(x) - f(c)| < \varepsilon$.

Theorem 1.2.8: Continuous Mapping

Let (X_n) be a sequence of random variables and $c \in \mathbb{R}$ for which $X_n \xrightarrow{p} c$. If U is an open neighbourhood of c and $g : U \to \mathbb{R}$ is continuous at c, then $g(X_n) \xrightarrow{p} g(c)$

Proof. We'll assume that $X_n(\Omega_{X_n}) \subseteq U$ so that each $g(X_n)$ is well defined. Fix $n \in \mathbb{N}$ and $\varepsilon > 0$, as g is continuous at c, there is $\delta > 0$ such that for all $x \in U$ we have $|x - c| < \delta \Rightarrow |g(x) - g(c)| < \varepsilon$. It then follows that for any $x \in \Omega_n$, we have $|X_n(x) - c| < \delta \Rightarrow |g(X_n(x)) - g(c)| < \varepsilon$, namely

$$\{x \in \Omega_{X_n} : |X_n(x) - c| < \delta\} \subseteq \{x \in \Omega_{X_n} : |g(X_n(x)) - g(c)| < \varepsilon\}$$

Thus, by monotonicity of probability it follows that $\mathbb{P}(|X_n - c| < \delta) \leq \mathbb{P}(|g(X_n) - g(c)| < \varepsilon)$. Finally, as $X_n \xrightarrow{p} X$, together with the squeeze theorem, we conclude

$$1 = \lim_{n\to\infty} \mathbb{P}(|X_n - c| < \delta) \leq \lim_{n\to\infty} \mathbb{P}(|g(X_n) - g(c)| < \varepsilon) \leq 1 \quad \implies \quad g(X_n) \xrightarrow{p} g(c)$$ $\square$

Let's start with a motivating problem. Suppose we want to compute

$$I = \int_0^1 e^{x^2}\,\mathrm{d}x$$

If we naively try to use the Fundamental Theorem of Calculus, we'll run into trouble right away as there is no elementary anti-derivative for e^{x^2}. Let's look at this from perhaps a less obvious perspective. Let $X \sim \text{Uniform}\,[0,1]$, and set $g(x) = e^{x^2}$. We can then write

$$I = \int_0^1 g(x)\,\mathrm{d}x = \int_0^1 g(x)\frac{1}{1-0}\,\mathrm{d}x = \int_{\mathbb{R}} g(x)p_X(x)\,\mathrm{d}x = \mathbb{E}g(X)$$

Namely we can express our integral as an expected value. Why might this be useful? Well, accordingly to Theorem 1.2.7, we know that a sequence of averages will converge in probability to an expected value (population average). Namely by Theorem 1.2.8, as g is continuous, if (X_n) is a sequence of Uniform $[0,1]$ random variables, then $\overline{g(X_n)} \to I$ as $I = \mathbb{E}g(X)$ is the population mean of $\{g(X_n) :$

$n \in \mathbb{N}\}$. Our strategy will then be as follows: Generate a large IID sample (note that lower case letters indicate observed values) $x_1, \dots, x_N \sim \text{Uniform}\,[0,1]$ and then compute

$$\hat{I} = \frac{1}{N}\sum_{i=1}^{N} g(x_i) \approx \mathbb{E}g(X)$$

As we've been doing, we can implement this solution in `R` to get an approximation for I.

```
g <- function(x){
return(exp(x^2))
}
I_hat = mean(g(runif(10000, 0, 1)))
```

```
> I_hat
[1] 1.458717
```

A common integral calculator gives us an approximate value of 1.462651, so we're not extremely accurate, but we're decently close with this approach. We can generalize this to work for an arbitrary continuous function $f : [a,b] \to \mathbb{R}$. If $X \sim \text{Uniform}\,[a,b]$ then

$$I = \int_a^b f(x)\,\mathrm{d}x = (b-a)\int_a^b f(x)\frac{1}{b-a}\,\mathrm{d}x = (b-a)\int_{\mathbb{R}} f(x)p_X(x)\,\mathrm{d}x = (b-a)\mathbb{E}f(X)$$

Our approximation would involve obtaining a large IID sample $x_1, \dots, x_N \sim \text{Uniform}\,[a,b]$ and setting

$$\hat{I} = \frac{b-a}{N}\sum_{i=1}^{N} f(x_i) \approx (b-a)\mathbb{E}f(X)$$

The last thing we'll look at is when our domain of integration is unbounded, which will force us away from making use of the uniform distribution (recall that we cannot place a uniform distribution on an unbounded set) We'll generally have two approaches: Either we can make a substitution so that our domain of integration is bounded so that we can use the estimation method above, or we can compare our integral with a distribution whose support is an unbounded set.

Example 1.2.9

Approximate $I = \int_0^\infty e^{-x^2}\,\mathrm{d}x$

Solution. We'll do this in three different ways, with the final approach leading to an extremely important result in statistical inference. We'll start easy for now by making use of a friendly distribution.

Solution 1. We can write $e^{-x^2} = e^{x-x^2}e^{-x}$ and make use of the exponential distribution by setting $h(x) = e^{x-x^2}$. Namely if $X \sim \text{Exponential}(1)$, then

$$I = \int_0^\infty e^{x^2}\,\mathrm{d}x = \int_0^\infty e^{x-x^2}e^{-x}\,\mathrm{d}x = \int_{\mathbb{R}} h(x)p_X(x)\,\mathrm{d}x = \mathbb{E}h(X)$$

By sampling $x_1, \dots, x_N \overset{\text{IID}}{\sim} \text{Exponential}(1)$, we can set $\hat{I}_1 = \frac{1}{N}\sum_{i=1}^{N} h(x_i)$.

Solution 2. Recall that $\sigma : \mathbb{R} \to (0,1)$ defined by $\sigma(x) = \frac{e^x}{1+e^x}$ (sometimes called the sigmoid function) is a bijection with inverse $\sigma^{-1}(x) = \log(\frac{x}{1-x})$. We can use symmetry to help us out first and write

$$I = \int_0^\infty e^{-x^2}\,\mathrm{d}x = \frac{1}{2}\int_{\mathbb{R}} e^{-x^2}\,\mathrm{d}x$$

By setting $u = \sigma(x)$ and writing $x = \sigma^{-1}(u)$, we have

$$\mathrm{d}x = \frac{1}{u(1-u)}\,\mathrm{d}u \qquad \Longrightarrow \qquad \frac{1}{2}I = \int_0^1 \exp\left\{-\left(\log\left(\frac{u}{1-u}\right)\right)^2\right\}\frac{1}{u(1-u)}\,\mathrm{d}u = \frac{1}{2}\int_0^1 \frac{e^{-(\sigma^{-1}(u))^2}}{u(1-u)}\,\mathrm{d}u$$

Going forward we can generate a sample from Uniform $[0,1]$ as we did before. Namely, we can generate $x_1, \dots, x_N \sim$ Uniform $[0,1]$ and set

$$\hat{I}_2 = \frac{1}{2N}\sum_{i=1}^N g(x_i) \qquad \text{where} \qquad g(x) = \frac{e^{-(\sigma^{-1}(x))^2}}{x(1-x)}$$

Solution 3. The last of our solutions makes use of the normal distribution and a fundamental result in statistical inference. Let's recall that if $X \sim N(0,\sigma^2)$ then

$$\mathbb{P}(X > 0) = \int_0^\infty \frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{x^2}{2\sigma^2}}\,\mathrm{d}x$$

We wanna pick σ so that e^{-x^2} is proportional to the density of $X \sim N(0,\sigma^2)$, namely we want $2\sigma^2 = 1$, and so $\sigma = \frac{1}{\sqrt{2}}$. Thus, if $X \sim N(0,\frac{1}{2})$, then

$$I = \int_0^\infty e^{-x^2}\,\mathrm{d}x = \sqrt{\pi}\int_0^\infty \frac{1}{\sqrt{\pi}}e^{-x^2}\,\mathrm{d}x = \sqrt{\pi}\mathbb{P}(X > 0)$$

From here, we'd like to be able to write $\mathbb{P}(X > 0)$ in terms of some sort of expected value so that we can use our same approach with the Weak Law of Large Numbers. A clever observation we can make is the following

$$\mathbb{P}(X > 0) = \int_0^\infty p_X(x)\,\mathrm{d}x = \int_{\mathbb{R}} p_X(x)\mathbb{I}(x > 0)\,\mathrm{d}x = \mathbb{E}(\mathbb{I}(X > 0))$$

and so we can make use of the same technique as before! Namely if we generate $x_1, \dots, x_N \sim N(0,\frac{1}{2})$, we can set

$$\hat{I}_3 = \frac{\sqrt{\pi}}{N}\sum_{i=1}^N \mathbb{I}(x_i > 0)$$

It's worth noting that if we have an IID sample $X_1, \dots, X_N$ from a distribution Q with cumulative distribution function F, then the *empirical cumulative distribution function* $\hat{F}_N$, defined by

$$\hat{F}_N(x) = \frac{1}{N}\sum_{i=1}^N \mathbb{I}(X_i < x)$$

is an important estimator used often in inference. The quality of this estimator is established by the famed Glivenko-Cantelli Theorem which states that $\sup_x |\hat{F}_N(x) - F(x)| \xrightarrow{p} 0$, namely that $\hat{F}_N$ is a *uniformly consistent estimator* of the CDF F. This is getting a bit too far away from the material in our course, so let's finish things up by using properties of density functions to compute the exact value of I. Suppose that $X \sim N(0,\frac{1}{2})$, then using the same symmetry idea as in solution 2, we can write

$$1 = \int_{\mathbb{R}} p_X(x)\,\mathrm{d}x = \int_{\mathbb{R}} \frac{1}{\sqrt{\pi}}e^{-x^2}\,\mathrm{d}x = \frac{2}{\sqrt{\pi}}\int_0^\infty e^{-x^2}\,\mathrm{d}x \qquad \Longrightarrow \qquad I = \frac{\sqrt{\pi}}{2}$$

Lastly, we can compute each of the values of our estimators and compare their errors based on our exact value (though we will propagate error when using an approximation of π)

```r
h <- function(x){
  return(exp(x-x^2))
}
inv_sigmoid <- function(x){
  return(log(x/(1-x)))
}
g <- function(x){
  return(exp(-(inv_sigmoid(x))^2)/(x*(1-x)))
}
I_1 <- mean(h(rexp(10000, 1)))
I_2 <- mean(g(runif(10000, 0, 1)))/2
I_3 <- sqrt(pi)*mean((rnorm(10000, 0, 1/sqrt(2)) > 0))
estimators <- c(I_1, I_2, I_3)
errors <- sqrt(pi)/2 - estimators
```

```
> data.frame("Estimator" = estimators, "Error" = errors)
  Estimator        Error
1 0.8874769 -0.001249978
2 0.8811834  0.005043504
3 0.8904808 -0.004253889
```

□

2 Analysis I – II

2.1 Exploring the Intermediate Value Theorem

The Intermediate Value Theorem is one of the more fundamental theorems that we'll encounter on our journey in the course, so it's worth taking a bit of a deeper dive to explore. Like with many theorems, there's more than one way to prove it, though the three ways that we'll see in this section each have their own importance to us as budding mathematicians. As a typical "existence" based theorem, it may not be the most useful in practice when we want to figure out where our intermediate value is actually achieved within our domain. Fortunately, the first two proofs we'll see are construction based, meaning we'll explicitly construct the desired point, though the second one will likely be more useful, providing an algorithm we can employ in practice. Our third and final proof will require to take a bit of a detour using topology, and tease a particular topological you'll explore further in Analysis II.

2.1.1 The Classical Treatment

Theorem 2.1.1: Weak Intermediate Value Theorem

If $f : [a, b] \to \mathbb{R}$ is continuous and $f(a)f(b) < 0$, there is $c \in (a, b)$ such that $f(c) = 0$.

Proof. The condition $f(a)f(b) < 0$ says that one of the endpoints is positive and the other negative, let's assume $f(a) < 0 < f(b)$ (otherwise apply the following to $g := -f$). Define the set

$$U = \{c \in [a, b] : f(x) < 0 \quad \forall x \in [a, c]\}$$

As $f(a) < 0$ by assumption and $[a, a] = \{a\}$, it follows that $a \in U$, moreover $U \subseteq [a, b]$ by construction, and so $\alpha := \sup(U)$ exists by completeness. We show $\alpha \in (a, b)$ such that $f(\alpha) = 0$. Suppose $f(\alpha) > 0$, the left continuity of f at α (we can't guarantee two-sided continuity at α since $\alpha = b$ is still technically possible), there is $\delta_1 > 0$ such that $f(x) > 0$ for all $x \in (\alpha - \delta_1, \alpha]$. By the criterion for suprema, there is $c_1 \in U$ such that $\alpha - \delta_1 < c_1 \leq \alpha$, and as $c_1 \in U, f(x) < 0$ for all $x \in [a, c_1]$, and hence $f(c_1) < 0$. But $c_1 \in (\alpha - \delta_1, \alpha]$, and so $f(c_1) > 0$, a contradiction. Thus $f(\alpha) \leq 0$ and hence $\alpha < b$. Suppose now that $f(\alpha) < 0$, by the continuity of f at $\alpha \in (a, b)$, there is $\delta_2 > 0$ such that $f(x) < 0$ for all $x \in (\alpha - \delta_2, \alpha + \delta_2)$, and hence for all $x \in \left[\alpha - \frac{\delta_2}{2}, \alpha + \frac{\delta_2}{2}\right]$. Again by there criterion for suprema, there is $c_2 \in U$ such that $\alpha - \delta_2 < c_2 \leq \alpha$. As $c_2 \in U, f(x) < 0$ for $x \in [a, c_2]$ and hence $f(x) < 0$ for $x \in \left[a, \alpha + \frac{\delta_2}{2}\right]$. Thus $f(x) < 0$ for all $x \in \left[a, \alpha + \frac{\delta_2}{2}\right]$, and hence $a + \frac{\delta_2}{2} \in U$, contradicting $\alpha = \sup(U)$. Thus $\alpha \in (a, b)$ and $f(\alpha) = 0$. □

Corollary 2.1.2: Intermediate Value Theorem

If $f : [a, b] \to \mathbb{R}$ is continuous, and d is any value between $f(a)$ and $f(b)$, there is $c \in [a, b]$ such that $f(c) = d$.

Proof. Fix some d between $f(a)$ and $f(b)$. If $d = f(a)$ or $d = f(b)$, we're done, so suppose d lies strictly between them. Define $g := f - d$, so that g is continuous on $[a, b]$ and $g(a)g(b) < 0$. By Theorem 2.1.1, there is $c \in (a, b)$ such that $g(c) = 0$, namely $f(c) = d$. □

While the above proof certainly gets the job done, it doesn't exactly help us out if we wanted to actually compute the value of c in practice. Though we explicitly construct it in the proof, it can be difficult in practice to determine the set $\{c \in [a, b] : f(x) < 0 \quad \forall x \in [a, c]\}$ and hence determine c by finding the supremum.

2.1.2 Thinking Computationally

Here we can take a different approach to proving our theorem that draws from a fairly simple recursive algorithm, and make things formal with the nested interval property. Effectively, we will be performing a "binary search" on the domain interval $[a, b]$. Let's start with an example to get a feel for things

Example 2.1.3: Motivation

Approximate a root of the function $f : [0, \frac{\pi}{2}] \to \mathbb{R}$ defined by $f(x) = x - \cos(x)$.

Solution. Note that $f(0) = -1 < 0 < \frac{\pi}{2} = f(\frac{\pi}{2})$, so assuming that we already have the Intermediate Value Theorem under our belt, we know for sure that there is a root. Let's actually try to find it! Or at the very least, find a way to approximate it. Set $c = \frac{\pi}{4}$, the midpoint of the domain. The trick will be to consider the sign of $f(c)$ and use it to shrink our domain of interest down to either $[0, c]$ if $f(c) > 0$ or $[c, \frac{\pi}{2}]$ if $f(c) < 0$. Note that $f(c) \approx -0.07829 < 0$ and so we can consider $[c, \frac{\pi}{2}]$ so that once again we have the images of the endpoints having opposite signs. Consider further the midpoint $c_2 = \frac{3\pi}{8}$ of $[c, \frac{\pi}{2}]$. As $f(c_2) \approx -0.79541 < 0$, we would consider $[c_2, \frac{\pi}{2}]$. Suppose we set

$$[a_0, b_0] = [0, \tfrac{\pi}{2}] \quad c_1 = \tfrac{a_0+b_0}{2} \quad [a_1, b_1] = [c_1, b_0] \quad c_2 = \tfrac{a_1+b_1}{2} \quad [a_2, b_2] = [c_2, b_1]$$

We can continuously repeat this process of bisecting the current interval and choosing the half whose endpoints map to values with different signs. Note that by doing so we'll create a sequence of nested intervals $\{[a_n, b_n]\}_{n\in\mathbb{N}}$ where hopefully the root of f that we expect to have will live in their intersection. $\square$

Theorem 2.1.4: Intermediate Value Theorem

If $f : [a, b] \to \mathbb{R}$ is a continuous function and $f(a)f(b) < 0$, then there is $c \in (a, b)$ such that $f(c) < 0$.

Proof. As we saw in the previous example, our approach will be to make use of the nested interval property. Define $a_0 = a, b_0 = b$, as $f(a_0)f(b_0) < 0$ by our assumption, let's assume without loss of generality that $f(a) < 0 < f(b)$ and consider $c_1 = \frac{a_0+b_0}{2}$. If $f(c_1) = 0$, then we're done. If not we set

$$a_1 = \begin{cases} c_1 & \text{if} f(c_1) < 0 \\ a_0 & \text{if} f(c_1) > 0 \end{cases} \qquad b_1 = \begin{cases} b_0 & \text{if} f(c_1) < 0 \\ c_1 & \text{if} f(c_1) > 0 \end{cases}$$

Now starting from $[a_1, b_1]$, we still have $f(a_1) < 0 < f(b_1)$, and so consider $c_2 = \frac{a_1+b_1}{2}$. Once again, if $f(c_2) = 0$, we're done, but if not, we set

$$a_2 = \begin{cases} c_2 & \text{if} f(c_2) < 0 \\ a_1 & \text{if} f(c_2) > 0 \end{cases} \qquad b_2 = \begin{cases} b_1 & \text{if} f(c_2) < 0 \\ c_2 & \text{if} f(c_2) > 0 \end{cases}$$

So that $[a_2, b_2] \subseteq [a_1, b_1] \subseteq [a_0, b_0]$ and $f(a_2) < 0 < f(b_2)$. Continuous this process inductively leaves us with a nested collection of intervals $\{[a_n, b_n]\}_{n=0}^{\infty}$ such that $f(a_n) < 0 < f(b_n)$ and $b_n - a_n < \frac{b-a}{2^n}$. Thus, by the nested interval property, there is $c \in \mathbb{R}$ such that

$$\{c\} = \bigcap_{n=0}^{\infty} [a_n, b_n]$$

moreover $(a_n) \to c$ and $(b_n) \to c$. Thus, by the continuity of f , we have

$$f(a_n) \to f(c) \leq 0 \quad \text{and} \quad f(b_n) \to f(c) \geq 0$$

Thus, as $f(a), f(b) \neq 0$, we conclude $c \in (a,b)$ is such that $f(c) = 0$. $\square$

In practice, we certainly won't be able to iterate through this process indefinitely, so we need some sort of stopping condition. Typically this is done in one of two ways, we either set a threshold on the size of $f(c_n)$, or on the width of $[a_n, b_n]$, namely we stop either when $f(c_n)$ is sufficiently small, or our interval of interest is sufficiently narrow. One final thing that I'll mention is that in practice, often times we will opt to write $c_{n+1} = \frac{a_n+b_n}{2}$ as $c_n = a_n + \frac{b_n-a_n}{2}$ as the latter is more computationally stable in a floating point number system. For those more interested in computational methods like this one, I encourage to take our Numerical Methods class!

2.1.3 Fun with Topology: What really is an Interval?

Definition 2.1.5: Disconnections and Connected Sets

We say that a set $S \subseteq \mathbb{R}$ is *disconnected* if there exists disjoint open sets $U, V \subseteq \mathbb{R}$ such that $S \subseteq U \cup V$ and $U \cap S, V \cap S \neq \varnothing$, such a pair (U, V) is called a *disconnection*. We say that a set is *connected* if it is not disconnected

Though this will be our standard definition (those who are more familiar with topology already may have noticed that we explicitly avoided mentioning a notion of the "subspace topology" in our definition of a disconnection), we can rephrase it using a simple property of sets. Recall that

$$(S \cap U) \cup (S \cap V) = S \cap (U \cup V)$$

and so $S \subseteq U \cup V$ if and only if $S = S \cap (U \cup V)$, namely $S \subseteq \mathbb{R}$ is disconnected if and only if there exists disjoint open sets $U, V \subseteq \mathbb{R}$ such that $S \cap U, S \cap V \neq \varnothing$ and $(S \cap U) \cup (S \cap V) = S$.

Definition 2.1.6: Formal Intervals

A set $I \subseteq \mathbb{R}$ is called an *interval* if for every $a, b \in I$ and every $x \in \mathbb{R}$, if $a < x < b$ then $x \in I$.

Though it may seem a bit strange and unnecessary at first, with a bit of thought hopefully this definition should make sense. We say that a set I an interval if for any two distinct points $x, y \in I$, all real numbers between x and y are also in I. Let's see a simple example that will lay some groundwork for the next result. Consider the set $A = [1,2] \cup [4,5]$. This is a set that we know is not an interval, and if we were to draw it on a number line, it would make sense that it shouldn't be considered a connected set (it's the union of two separated pieces). Let's verify that A is not connected by explicitly constructing a disconnection. Note that $2, 4 \in A$, yet $3 \notin A$, showing that A is in fact not an interval. Additionally, by considering the open sets $U = (-\infty, 3)$ and $V = (3, \infty)$, it follows that (U, V) forms a disconnection of A. The exact details are generalized to form the following Proposition.

Proposition 2.1.7

If $I \subseteq \mathbb{R}$ is not an interval, then I is disconnected.

Proof. As I is not an interval, negating Definition 2.1.6 says there is $a < b \in I$ and $x \in \mathbb{R}$ such that $a < x < b$ and $x \notin I$. Define $U = (-\infty, x), V = (x, \infty)$. Then clearly U and V are open an disjoint, moreover $a \in U \cap I, b \in V \cap I$ and as $x \notin I$ we have $I \subseteq \mathbb{R} \setminus \{x\} = U \cup V$. Thus (U, V) form a disconnection and hence I is disconnected. $\square$

Proposition 2.1.8

If $I \subseteq \mathbb{R}$ is disconnected, then I is not an interval.

Proof. Suppose $I \subseteq \mathbb{R}$ admits a disconnection (U, V). Pick $x \in U \cap I, y \in V \cap I$, as $U \cap V = \varnothing$ we have $x \neq y$ and so assume without loss of generality $x < y$. Define

$$A = \{a \geq x : [x, a] \cap I \subseteq U \cap I\}$$

Note that $x \in A$ and as $y \notin I \cap U$ we have y is an upper bound for A, thus by completeness $c = \sup(A)$ exists. We show $c \notin I$. Suppose $c \in U \cap I$, then as $c \in U$ and U is open, there is $r > 0$ such that $(c - r, c + r) \subseteq U$ and hence $(c - r, c + r) \cap I \subseteq U \cap I$. As $c - r$ is not an upper bound, there is $a_0 \in A$ such that $c - \frac{r}{2} < a_0$, and hence as $[x, a_0] \cap I \subseteq U \cap I$ we have

$$\left[x, c + \frac{r}{2}\right] \cap I = \left(\left[x, c - \frac{r}{2}\right] \cap I\right) \cup \left(\left[c - \frac{r}{2}, c + \frac{r}{2}\right] \cap I\right) \subseteq U \cap I$$

Namely, $c + \frac{r}{2} \in A$, contradicting that $c = \sup(A)$ is an upper bound, and so $c \notin U \cap I$. Conversely, suppose $c \in V \cap I$. Again, as $c \in V$ and V is open, there is $r > 0$ such that $(c - r, c + r) \cap I \subseteq V \cap I$. Finally, as $(c - r, c - \frac{r}{2}] \cap I \subseteq (c - r, c + r) \cap I \subseteq V \cap I$, we have

$$\left[x, c - \frac{r}{2}\right] \cap I \cap (V \cap I) \neq \varnothing \quad \Longrightarrow \quad \left[x, c - \frac{r}{2}\right] \cap I \not\subseteq U \cap I$$

Namely $c - \frac{r}{2} \notin A$ and it follows that $c - \frac{r}{2}$ is an upper bound for A, contradicting that $c = \sup(A)$ is the least upper bound. Thus, $c \notin V \cap I$ and hence $c \notin I$ as $I = (U \cap I) \cup (V \cap I)$. Thus, as $x \in A$ and y is an upper bound for A we have $x \leq c \leq y$ and moreover as $x \in U \cap I, y \in V \cap I$, we have $x < c < y$. Thus, I is not an interval. $\square$

By combining the previous two propositions, we have that $I \subseteq \mathbb{R}$ is disconnected if and only if I is not an interval, where the contrapositive says that a set is connected in $\mathbb{R}$ if and only if it is an interval.

Proposition 2.1.9

Let $S \subseteq \mathbb{R}$ and $f : S \to \mathbb{R}$ a continuous function. If $f(S)$ is disconnected then S is disconnected.

Proof. Let (U, V) be a disconnection of $f(S)$. By the continuity of f, as $U, V \subseteq \mathbb{R}$ are open we have that $f^{-1}(U), f^{-1}(V)$ are open. We claim $(f^{-1}(U), f^{-1}(V))$ forms a disconnection of S. Indeed, note that $f^{-1}(U) \cap f^{-1}(V) = f^{-1}(U \cap V) = \varnothing$ as $U \cap V = \varnothing$, moreover

$$f^{-1}(U) \cup f^{-1}(V) = f^{-1}(U \cup V) \supseteq f^{-1}(f(S)) \supseteq S$$

as $f(S) \subseteq U \cup V$. Finally, Let $y \in f(S) \cap U$, then writing $y = f(x)$ for $x \in S$ we have $x \in S \cap f^{-1}(U)$ and so $f^{-1}(U) \cap S \neq \varnothing$ and similarly $f^{-1}(V) \cap S \neq \varnothing$. Thus, $(f^{-1}(U), f^{-1}(V))$ forms a disconnection of S and hence S is disconnected. $\square$

Corollary 2.1.10: Intermediate Value Theorem

If $f : [a, b] \to \mathbb{R}$ is a continuous function such that $f(a)f(b) < 0$, then there is $c \in (a, b)$ such that $f(c) = 0$.

Proof. As $[a, b] \subseteq \mathbb{R}$ is connected and f is continuous, by Proposition 2.1.9 $f([a, b])$ is connected. Assume without loss of generality $f(a) < 0 < f(b)$. Then as $f([a, b])$ is an interval we have $0 \in f([a, b])$ by Definition 2.1.6 and the result follows. $\square$

2.2 The Mean Value Theorem Integrals

Recall that the first Fundamental Theorem of Calculus says that if $f : [a, b] \to \mathbb{R}$ is integrable, then the function $F : [a, b] \to \mathbb{R}$ defined by

$$F(x) = \int_a^x f(t)\,\mathrm{d}t$$

is Lipschitz and morerover, F is differentiable at any point of continuity of f. In particular, if f is continuous (and hence integrable), then F is differentiable, and so we're able to use the Mean Value Theorem on F to establish some elementary results. Often times, reducing the regularity of f from continuity to plain integrability will fail preserve a result that follows from such an application of the Mean Value Theorem. In this section, we will explore a theorem that holds regardless of the continuity of f, and we'll work our way up to proving by first assuming that we have some strong regularity conditions than we may need.

2.2.1 Benefiting From Stronger Assumptions

Proposition 2.2.1

Let $f, g : [a, b] \to \mathbb{R}$ be integrable function such that g is non-negative and f is continuous. Then there is $\xi \in [a, b]$ such that

$$\int_a^b f(x)g(x)\,\mathrm{d}x = f(\xi) \int_a^b g(x)\,\mathrm{d}x$$

Proof. As f is continuous, there is $s, t \in [a, b]$ such that $f(t) \leq f(x) \leq f(s)$ for all $x \in [a, b]$, and so as g is non-negative, by monontonicity we have

$$f(t)g(x) \leq f(x)g(x) \leq f(s)g(x) \quad \Longrightarrow \quad f(t)\int_a^b g(x)\,\mathrm{d}x \leq \int_a^b f(x)g(x)\,\mathrm{d}x \leq f(s)\int_a^b g(x)\,\mathrm{d}x$$

Define $F : [a, b] \to \mathbb{R}$ by $F(x) = f(x)\int_a^b g(x)\,\mathrm{d}x$. Note that F is continuous as a constant multiple of a continuous function, and moreover as $F(t) \leq \int_a^b f(x)g(x)\,\mathrm{d}x \leq F(s)$ from our construction, by the Intermediate Value Theorem there is $\xi \in [a, b]$ such that

$$\int_a^b f(x)g(x)\,\mathrm{d}x = F(\xi) = f(\xi)\int_a^b g(x)\,\mathrm{d}x \qquad \square$$

Theorem 2.2.2

Let $f, g : [a, b] \to \mathbb{R}$ be continuous functions such that g is continuously differentiable and non-decreasing. Then there is $\xi \in [a, b]$ such that

$$\int_a^b f(x)g(x)\,\mathrm{d}x = g(a)\int_a^\xi f(x)\,\mathrm{d}x + g(b)\int_\xi^b f(x)\,\mathrm{d}x$$

Proof. Using the Fundamental Theorem of Calculus, the map $F : [a, b] \to \mathbb{R}$ defined by $F(x) = \int_a^x f(t)\,\mathrm{d}t$ is an antiderivative of f as f is continuous. Moreover, as g is continuously differentiable, we may apply integration by parts to $\int_a^b f(x)g(x)\,\mathrm{d}x$. Doing so gives

$$\int_a^b f(x)g(x)\,\mathrm{d}x = [g(x)F(x)]_a^b - \int_a^b g'(x)F(x)\,\mathrm{d}x$$

Note that F and g' are continuous, moreover $g' \geq 0$ as g is non-decreasing. From here we can apply Proposition 2.2.1 to $\int_a^b F(x)g'(x)\,\mathrm{d}x$ to obtain some $\xi \in [a,b]$ such that

$$\begin{aligned}\int_a^b f(x)g(x)\,\mathrm{d}x &= [g(x)F(x)]_a^b - \int_a^b g'(x)F(x)\,\mathrm{d}x \\ &= g(b)F(b) - g(a)F(a) - F(\xi)\int_a^b g'(x)\,\mathrm{d}x \\ &= g(b)\int_a^b f(x)\,\mathrm{d}x - \int_a^\xi f(x)\,\mathrm{d}x(g(b)-g(a)) \\ &= g(a)\int_a^\xi f(x)\,\mathrm{d}x + g(b)\int_\xi^b f(x)\,\mathrm{d}x\end{aligned}$$

□

2.2.2 Back to Basics

In order to prove the Mean Value Theorem for integrals without making any assumptions about continuity, we'll have to first establish several Lemmas to help us out.

Lemma 2.2.3: Abel's Summation by Parts

Let $\{a_i\}_{i=1}^n, \{b_i\}_{i=1}^n$ be two collections of real numbers where $A_k := \sum_{k=1}^n a_i$. Then

$$\sum_{i=1}^n a_i b_i = A_n b_n + \sum_{i=1}^{n-1} A_i(b_i - b_{i+1})$$

Proof. We proceed by induction on $n \in \mathbb{N}_{\geq 2}$. When $n = 2$, we have $a_1b_1 + a_2b_2 = (a_1+a_2)b_2 + a_1(b_1 - b_2)$, establishing our base case. Now suppose the result holds for $n \in \mathbb{N}_{\geq 2}$ and consider a sum of $n+1$ terms. Using the induction hypothesis, we can write

$$\begin{aligned}\sum_{i=1}^{n+1} a_i b_i = a_{n+1}b_{n+1} + \sum_{i=1}^n a_i b_i &= a_{n+1}b_{n+1} + A_n b_n + \sum_{i=1}^{n-1} A_i(b_i - b_{i+1}) \\ &= (A_{n+1} - A_n)b_{n+1} + A_n b_n + \sum_{i=1}^{n-1} A_i(b_i - b_{i+1}) \\ &= A_{n+1}b_{n+1} + \sum_{i=1}^n A_i(b_i - b_{i+1})\end{aligned}$$

□

We can quickly establish another result as a relatively immediate consequence of Abel's Lemma, one that will help us out in proving our final Lemma before digging into the proof of the main theorem.

Corollary 2.2.4

Let $\{a_i\}_{i=1}^n, \{b_i\}_{i=1}^n$ be two collections of real numbers where $A_k := \sum_{k=1}^n a_i$. If $b_k \geq b_{k+1} \geq 0$ for all $k = 1, \ldots, n-1$ and there are $m, M \in \mathbb{R}$ for which $m \leq A_k \leq M$ for all $k = 1, \ldots, n$, then

$$mb_1 \leq \sum_{k=1}^n a_k b_k \leq Mb_1$$

Proof. We can make use of Abel's Lemma, namely as $m \le A_k \le M$

$$\begin{aligned}\sum_{k=1}^{n} a_k b_k = A_n b_n + \sum_{i=1}^{n-1} A_i(b_i - b_{i+1}) &\le M b_n + \sum_{i=1}^{n-1} M(b_i - b_{i+1}) \\ &= M b_n + M(b_1 - b_n) \\ &= M b_1\end{aligned}$$

$$\begin{aligned}\sum_{k=1}^{n} a_k b_k = A_n b_n + \sum_{i=1}^{n-1} A_i(b_i - b_{i+1}) &\ge m b_n + \sum_{i=1}^{n-1} m(b_i - b_{i+1}) \\ &= m b_n + m(b_1 - b_n) \\ &= m b_1\end{aligned}$$

Thus $mb_1 \le \sum_{k=1}^{n} a_k b_k \le M b_1$ for all k. □

Lemma 2.2.5

If $f : [a, b] \to \mathbb{R}$ is integrable, Then

$$\lim_{\|P\|\to 0} \sup_{x\in(a,b]} (U(f|_{[a,x]}, P \cap [a, x]) - L(f|_{[a,x]}, P \cap [a, x])) = 0$$

Proof. Let $\varepsilon > 0$, as f is integrable, there is $\delta > 0$ such that for all partitions P of $[a, b]$

$$\|P\| < \delta \implies U(f, P) - L(f, P) < \frac{\varepsilon}{2}$$

Fix a partition P of $[a, b]$ such that $\|P\| < \delta$. By properties of suprema, there is $c_0 \in (a, b]$ such that

$$\begin{aligned}\sup_{x\in(a,b]} (U(f|_{[a,x]}, P \cap [a, x]) - L(f|_{[a,x]}, P \cap [a, x])) &< U(f|_{[a,c_0]}, P \cap [a, c_0]) - L(f|_{[a,c_0]}, P \cap [a, c_0]) + \frac{\varepsilon}{2} \\ &\le U(f, P \cup \{c_0\}) - L(f, P \cup \{c_0\}) + \frac{\varepsilon}{2} \\ &\le U(f, P) - L(f, P) + \frac{\varepsilon}{2} \\ &< \varepsilon\end{aligned}$$

□

From here, we are ready to state and prove stronger versions of Proposition 2.2.1 and Theorem 2.2.2

Proposition 2.2.6

Suppose $f, g : [a, b] \to \mathbb{R}$ are integrable functions where g is non-negative and non-increasing. Then there is $\xi \in [a, b]$ such that

$$\int_a^b f(x)g(x)\,\mathrm{d}x = g(a)\int_a^{\xi} f(x)\,\mathrm{d}x$$

Proof. Let $P = \{x_0, \dots, x_n\}$ be a partition of $[a, b]$ with $T = \{a, t_2, \dots, t_n\}$. For $i = 1, \dots, n$, define $a_i = f(t_i)\Delta x_i$ and $b_i = g(t_i)$. As g is non-increasing and non-negative, we have $b_i \ge b_{i+1} \ge 0$. Moreover, note that by construction, for each $k = 1, \dots, n$

$$A_k := \sum_{i=1}^{k} a_i = \sum_{i=1}^{k} f(t_i)\Delta x_i = \sigma(f|_{[a,x_k]}, P \cap [a, x_k], T \cap [a, x_k])$$

Define $F : [a, b] \to \mathbb{R}$ by $F(x) = \int_a^x f(t)\,\mathrm{d}t$, by the Fundamental Theorem of Calculus together with the Extreme Value Theorem, there is $x_m, x_M \in [a, b]$ such that $F(x_m) \leq F(x) \leq F(x_M)$ for all $x \in [a, b]$. For each $k = 1, \ldots, n$, write $P_{x_k} = P \cap [a, x_k]$ and $T_{x_k} = T \cap [a, x_k]$. As each $f|_{[a,x_k]}$ is integrable, we have that

$$\left|\sigma(f_{[a,x_k]}, P_{x_k}, T_{x_k}) - \int_a^{x_k} f(x)\,\mathrm{d}x\right| \leq U(f_{[a,x_k]}, P_{x_k}) - L(f|_{[a,x_k]}, P_{x_k})$$

By expanding the absolute values and noting that $\int_a^{x_k} f(x)\,\mathrm{d}x = F(x_k)$, we can write

$$\begin{aligned}
\underbrace{-\sup_{x\in(a,b]} (U(f|_{[a,x]}, P_x) - L(f|_{[a,x]}, P_x)) + F(x_m)}_{=:m} &\leq -(U(f_{[a,x_k]}, P_{x_k}) - L(f|_{[a,x_k]}, P_{x_k})) + F(x_k) \\
&\leq \underbrace{\sigma(f_{[a,x_k]}, P_{x_k}, T_{x_k})}_{A_k} \\
&\leq U(f_{[a,x_k]}, P_{x_k}) - L(f|_{[a,x_k]}, P_{x_k}) + F(x_k) \\
&= \underbrace{\sup_{x\in(a,b]} (U(f|_{[a,x]}, P_x) - L(f|_{[a,x]}, P_x)) + F(x_M)}_{=:M}
\end{aligned}$$

Thus, we can apply the result of Corollary 2.2.4 to bound $\sum_{i=1}^n a_i b_i = \sigma(fg, P, T)$, namely

$$mg(a) \leq \sum_{i=1}^n a_i b_i = \sigma(fg, P, T) \leq Mg(a)$$

By Lemma 2.2.5 and the Squeeze Theorem, we have $m \to F(x_m)$ and $M \to F(x_M)$ as $\|P\| \to 0$, and so $g(a)F(x_m) \leq \int_a^b f(x)g(x)\,\mathrm{d}x \leq g(a)F(x_M)$. Finally, as $H : [a, b] \to \mathbb{R}$ defined by $H(x) = g(a)F(x)$ is continuous, it follows by the Intermediate Value Theorem that there is $\xi \in [a, b]$ such that

$$\int_a^b f(x)g(x)\,\mathrm{d}x = H(\xi) = g(a)\int_a^\xi f(x)\,\mathrm{d}x$$

$\square$

Theorem 2.2.7: Integral Mean Value

Suppose that $f, g : [a, b] \to \mathbb{R}$ are integrable functions where g is non-decreasing. Then there is $\xi \in [a, b]$ such that

$$\int_a^b f(x)g(x)\,\mathrm{d}x = g(a)\int_a^\xi f(x)\,\mathrm{d}x + g(b)\int_\xi^b f(x)\,\mathrm{d}x$$

Proof. Define $h : [a, b] \to \mathbb{R}$ by $h(x) = g(b) - g(x)$. As g is continuous and non-decreasing, it follows that h is continuous and non-increasing and so applying Proposition 2.2.6 to f and h, we have that there is $\xi \in [a, b]$ such that

$$\int_a^b f(x)h(x)\,\mathrm{d}x = h(a)\int_a^\eta f(x)\,\mathrm{d}x \iff \int_a^b f(x)(g(b) - g(x))\,\mathrm{d}x = (g(b) - g(a))\int_a^\xi f(x)\,\mathrm{d}x$$

Rewriting everything in terms of f and g using the linearity of the integral, we can conclude that

$$\begin{aligned}
\int_a^b f(x)g(x)\,\mathrm{d}x &= g(b)\int_a^b f(x)\,\mathrm{d}x - g(b)\int_a^\xi f(x)\,\mathrm{d}x + g(a)\int_a^\xi f(x)\,\mathrm{d}x \\
&= g(a)\int_a^\xi f(x)\,\mathrm{d}x + g(b)\int_\xi^b f(x)\,\mathrm{d}x
\end{aligned}$$

$\square$

2.3 Uniform Convergence and Differentiation

To start, recall that a sequence of functions (f_n) on $X \subseteq \mathbb{R}$ is said to *converge uniformly* to a function $f : X \to \mathbb{R}$ if for every $\varepsilon > 0$, there is $N \in \mathbb{N}$ such that for all $n \in \mathbb{N}$ and $x \in X$

$$n > N \implies |f_n(x) - f(x)| < \varepsilon$$

Recall further that last week in tutorial we showed that boundedness, integrability, and continuity were all preserved by uniform convergence, while none of which were preserved by pointwise convergence. Here, we'll work to establish a relationship between uniform convergence and differentiability.

2.3.1 When Uniform Convergence is not Enough

As we've seen throughout tutorial, most properties that we've been working are preserved under uniform convergence, notably things like boundedness, (uniform) continuity and integrability. However, differentiability is a bit less nicely behaved. In particular, we can construct a sequence of differentiable functions converging uniformly to a non-differentiable function.

Example 2.3.1: Counterexample I

Define $f_n, f : (-1,1) \to \mathbb{R}$ by $f_n(x) = \sqrt{x^2 + \frac{1}{n}}$ and $f(x) = |x|$ for each $n \in \mathbb{N}$. Then (f_n) is a sequence of differentiable functions, $(f_n) \to f$ uniformly but f is not differentiable.

Proof. For any $n \in \mathbb{N}$, as $x^2 + \frac{1}{n}$ maps into $(0, \infty)$, we have that f_n is the composition of differentiable functions and is hence differentiable by the chain rule. Moreover, $f(x) = |x|$ is not differentiable at 0 and hence not on $(-1,1)$. We show that $(f_n) \to f$ uniformly. Indeed, let $\varepsilon > 0$, pick $N \in \mathbb{N}$ such that $\frac{1}{\sqrt{N}} < \varepsilon$. For any $n \in \mathbb{N}$ and $x \in (-1,1)$, we have

$$0 \leq |x| \leq \sqrt{x^2 + \frac{1}{n}} \leq |x| + \frac{1}{\sqrt{n}}$$

Thus $n > N \Rightarrow |f_n(x) - f(x)| \leq \frac{1}{\sqrt{n}} < \frac{1}{\sqrt{N}} < \varepsilon$, and so $(f_n) \to f$ uniformly. $\square$

This may come as a surprise, and a natural first thought might be to add additional assumptions in order to ensure the result we want. While this will eventually push us in the right direction, we can first take a look at when this won't work.

Example 2.3.2: Counterexample II

Define $f_n, f : \mathbb{R} \to \mathbb{R}$ by $f_n(x) = \frac{\sin(nx)}{n}$ and $f \equiv 0$ for each $n \in \mathbb{N}$. Then (f_n) is a sequence of differentiable functions, $(f_n) \to f$ uniformly, f is differentiable but $(f_n') \not\to f'$

Proof. Clearly both f and each f_n are differentiable. Moreover given $\varepsilon > 0$, choose $N \in \mathbb{N}$ such that $\frac{1}{N} < \varepsilon$, then for all $n \in \mathbb{N}$ and $x \in \mathbb{R}$

$$n > N \implies |f_n(x)| = \frac{|\sin(nx)|}{n} \leq \frac{1}{n} < \frac{1}{N} < \varepsilon$$

and so $(f_n) \to f$ uniformly. However, $f_n'(x) = \cos(nx)$ with does not converge pointwise on $\mathbb{R}$ as, for example, $f_n'(1) = \cos(n)$ fails to converge. $\square$

2.3.2 Filling in the Cracks

Above we saw that additionally assuming that the uniform limit was differentiable wasn't able to guarantee for us that the sequence of derivatives even converged pointwise! Now, we'll completely overcompensate and add significantly stronger conditions to guarantee our result, then work on relaxing as many as we can to ultimately strengthen our result.

Proposition 2.3.3

If (f_n) is a sequence of continuously differentiable functions on $[a,b]$, $(f_n') \to g$ uniformly and $(f_n) \to f$ uniformly, then f is differentiable and $f' = g$.

Proof. As each f_n is continuously differentiable, (f_n') is a sequence of continuous (and hence integrable) functions, and so by last week's tutorial, g is continuous (and hence integrable). From last week once again, we have that

$$\lim_{n\to\infty} \int_a^x f_n'(t)\,\mathrm{d}t = \int_a^x g(t)\,\mathrm{d}t$$

Thus, by the fundamental theorem of calculus together with the limit laws, we conclude

$$\begin{aligned} f(x) = \lim_{n\to\infty}(f_n(x) - f_n(a)) + f(a) &= \lim_{n\to\infty} \int_a^x f_n(t)\,\mathrm{d}t + f(a) \\ &= \int_a^x g(t)\,\mathrm{d}t + f(a) \end{aligned}$$

By the fundamental theorem of calculus once again, as g is continuous, $f(x) = \int_a^x g(t)\,\mathrm{d}t + f(a)$ is differentiable with derivative $g(x)$ and the result follows. □

Now that we've managed to buff our hypotheses enough to ensure the conclusion, we can strip away as many of the unnecessary assumptions we made as possible. For example, Nowhere in our proof did we use the uniform convergence of (f_n) to f. To use the limit laws as we did, we only required pointwise convergence. Additionally, though the continuously differentiable assumption was necessary for our above method of proof, it seems a bit too restrictive, so let's reduce it down to regular differentiability.

Proposition 2.3.4

If (f_n) is a sequence of differentiable functions on $[a,b]$ for which $(f_n) \to f$ and $(f_n') \to g$ uniformly, then f is differentiable and $f' = g$.

Proof. Let $c \in [a,b]$ and for simplicity let's assume $c \in (a,b)$. We aim to show that f is differentiable at c with $f'(c) = g(c)$, which we can do by analyzing the difference quotient of f, comparing it to the sequence of difference quotients of f_n. Define $\phi_n, \phi : [a,b] \setminus \{c\} \to \mathbb{R}$ by

$$\phi_n(x) = \frac{f_n(x) - f_n(c)}{x - c} \quad \text{and} \quad \phi(x) = \frac{f(x) - f(c)}{x - c}$$

We claim that $(\phi_n) \to \phi$ uniformly. Indeed, we will make use of the Cauchy criterion and show first that (ϕ_n) is uniformly Cauchy. Let $\varepsilon > 0$, as $(f_n') \to g$ uniformly and is hence uniformly Cauchy, there is $N \in \mathbb{N}$ such that $n, m > N \Rightarrow |f_n'(x) - f_m'(x)| < \varepsilon$ for all $x \in [a,b]$. Given $n \in \mathbb{N}$ and $x \in [a,b] \setminus \{c\}$, assume $n, m > N$ and that without loss of generality $x > c$. By applying the mean value theorem to $f_n - f_m$ on $[c, x]$, there is $\theta \in (c, x)$ such that

$$\frac{(f_n - f_m)(x) - (f_n - f_m)(c)}{x - c} = f_n'(\theta) - f_m'(\theta)$$

Finally, by regrouping the necessary terms, we have that

$$\begin{aligned}|\phi_n(x) - \phi_m(x)| &= \left|\frac{f_n(x) - f_n(c)}{x - c} - \frac{f_m(x) - f_n(c)}{x - c}\right| \\ &= \left|\frac{(f_n - f_m)(x) - (f_n - f_m)(c)}{x - c}\right| \\ &= |f_n'(\theta) - f_m'(\theta)| \\ &< \varepsilon\end{aligned}$$

Thus (ϕ_n) is uniformly Cauchy. To show that $(\phi_n) \to \phi$ uniformly, we show the convergence is pointwise and conclude by the uniqueness of limits and the Cauchy criterion. Given $x \in [a, b] \setminus \{c\}$, by the limit laws, as $(f_n) \to f$ uniformly (and hence pointwise) we have

$$\begin{aligned}\lim_{n\to\infty} \phi_n(x) = \lim_{n\to\infty} \left(\frac{f_n(x) - f_n(c)}{x - c}\right) &= \frac{\lim_{n\to\infty} f_n(x) - \lim_{n\to\infty} f_n(c)}{x - c} \\ &= \frac{f(x) - f(c)}{x - c} \\ &= \phi(x)\end{aligned}$$

and so $(\phi_n) \to \phi$ uniformly. To conclude that f is differentiable at c with $f'(c) = g(c)$, as $(f_n') \to g$ and $(\phi_n) \to \phi$ uniformly, there is $N_1, N_2 \in \mathbb{N}$ such that for all $n \in \mathbb{N}$ and $x \in [a, b], y \in [a, b] \setminus \{c\}$

$$n, m > N_1 \implies |f_n'(x) - f_m'(x)| < \frac{\varepsilon}{4} \quad \text{and} \quad n, m > N_2 \implies |\phi_n(y) - \phi(y)| < \frac{\varepsilon}{4}$$

Fix $n = \max\{N_1, N_2\} + 1$, as f_n is differentiable at c, there is $\delta > 0$ such that for all $x \in [a, b]$

$$0 < |x - c| < \delta \implies |\phi_n(x) - f_n'(c)| = \left|\frac{f_n(x) - f_n(x)}{x - c}\right| < \frac{\varepsilon}{4}$$

Let $x \in [a, b]$ and assume that $0 < |x - c| < \delta$. By the triangle inequality, we have

$$\begin{aligned}\left|\frac{f(x) - f(c)}{x - c} - g(c)\right| &= |\phi(x) - g(c)| \\ &\leq |\phi(x) - \phi_n(x)| + |\phi_n(x) - f_n'(c)| + |f_n'(c) - g(c)| \\ &< \frac{\varepsilon}{4} + \frac{\varepsilon}{4} + \frac{\varepsilon}{4} \\ &< \varepsilon\end{aligned}$$

□

It's worth noting that we can strengthen this result even further by weakening our assumptions on the convergence of (f_n) and further leveraging the assumption that $(f_n') \to g$ uniformly. The strongest form of our result appears as follows

Theorem 2.3.5: Uniform Differentiation

Let (f_n) be a sequence of differentiable functions on $[a, b]$, $g : [a, b] \to \mathbb{R}$ be a function such that $(f_n') \to g$ uniformly and suppose there is $x_0 \in (a, b)$ such that $(f_n(x_0))$ converges. Then there is a differentiable function $f : [a, b] \to \mathbb{R}$ such that $(f_n) \to f$ uniformly and $f' = g$.

Proof. Using Proposition 2.3.4, it suffices to prove that (f_n) is uniformly Cauchy. Let $\varepsilon > 0$, as $(f_n') \to g$ uniformly, and is hence uniformly Cauchy, there is $N_1 \in \mathbb{N}$ such that

$$n, m > N_1 \quad \implies \quad |f_n(x) - f_m(x)| < \frac{\varepsilon}{2(b - a)} \qquad \forall x \in [a, b]$$

Similarly, as $(f_n(x_0))$ is convergent and hence Cauchy, there is $N_2 \in \mathbb{N}$ such that $n, m > N_2 \Rightarrow |f_n(x_0) - f_m(x_0)| < \frac{\varepsilon}{2}$. Define $N = \max\{N_1, N_2\} \in \mathbb{N}$, let $x \in [a, b]$ and assume $n, m > N$. If $x = x_0$, we have $|f_n(x_0) - f_m(x_0)| < \frac{\varepsilon}{2}$. For $x \neq x_0$, assume without loss of generality $x > x_0$, by applying the mean value theorem to $f_n - f_m$ on $[x_0, x]$, there is $\theta \in (x_0, x)$ such that

$$|(f_n - f_m)(x) - (f_n - f_m)(x_0)| = |f_n'(\theta) - f_m'(\theta)|(x - x_0) \leq |f_n'(\theta) - f_m'(\theta)|(b - a) \tag{$\star$}$$

Finally, by the triangle inequality together with $(\star)$, we have that

$$\begin{aligned} |f_n(x) - f_m(x)| &\leq |(f_n - f_m)(x) - (f_n - f_m)(x_0)| + |f_n(x_0) - f_m(x_0)| \\ &< |f_n'(\theta) - f_m'(\theta)|(b - a) + \frac{\varepsilon}{2} \\ &< \frac{\varepsilon}{2(b - a)}(b - a) + \frac{\varepsilon}{2} \\ &= \varepsilon \end{aligned}$$

Thus, (f_n) is uniformly Cauchy and so there is $f : [a, b] \to \mathbb{R}$ for which $(f_n) \to f$ uniformly. By Proposition 2.3.4, it then follows that f is differentiable with $f' = g$. □

3 Calculus of Several Variables

3.1 Limits of Multivariable Functions

In this course, we will almost exclusively deal with functions on $\mathbb{R}^2$. To show that such a limit exists, we will either use direct evaluation (plugging in the limit point into the function), the squeeze theorem, or polar coordinates, we will study the latter two in more detail in this handout. Let's start with the squeeze theorem and some related topics.

3.1.1 Squeeze Theorem and Friends

The squeeze theorem is a powerful tool for showing that a function approaches a proposed limit, and we can use it in the same way as we did for functions of one variable. Let's see the statement of the theorem.

Theorem 3.1.1: Squeeze Theorem

Suppose $f, g, h : U \subseteq \mathbb{R}^2 \to \mathbb{R}$ are functions defined on an open neighbourhood U of (a, b). If there is an open neighbourhood $D \subseteq U$ of (a, b) such that $f(x, y) \leq g(x, y) \leq h(x, y)$ for all $(x, y) \in D \setminus \{(a, b)\}$, and

$$\lim_{(x,y)\to(a,b)} f(x, y) = L = \lim_{(x,y)\to(a,b)} h(x, y)$$

Then $\lim_{(x,y)\to(a,b)} g(x, y) = L$.

Unwrapping some of the potentially technical jargin, the squeeze theorem is saying that if we want to find the limit of a function $g(x, y)$ at (a, b), and we know that g lives between two functions f and h near (a, b), except possibly at (a, b) (namely $f(x, y) \leq g(x, y) \leq h(x, y)$ near (a, b), but not necessarily at (a, b)) and the limits of f and h are the same at (a, b), then the limit of g must be their common value. Before we get into some examples, note that the squeeze theorem involves inequalities, so let's have a look at some inequalities that will be helpful.

Proposition 3.1.2: Helpful Inequalities

Note that $3, 4, 5$ and 6 are also true with y in place of x.

1. $|x + y| \leq |x| + |y|$
2. $||x| - |y|| \leq |x - y|$
3. $0 \leq x^2 \leq x^2 + y^2$
4. $0 \leq \frac{1}{x^2+y^2} \leq \frac{1}{x^2}$
5. $0 \leq |x| \leq \sqrt{x^2 + y^2}$
6. $0 \leq \frac{1}{\sqrt{x^2+y^2}} \leq \frac{1}{|x|}$
7. $0 \leq \left|\frac{xy}{x^2+y^2}\right| \leq \frac{1}{2}$
8. $|\sin(\theta)| \leq 1$
9. $|\cos(\theta)| \leq 1$

Now that we have some inequalities to help us out, let's see an example of how we can use the squeeze theorem to show that a limit exists.

Example 3.1.3

Show that $\displaystyle\lim_{(x,y)\to(0,0)} \frac{x^2}{\sqrt{x^2 + y^2}} = 0.$

Solution. Let's use Proposition 3.1.2 to set up a helpful inequality. It looks like using inequality 3 will be helpful here. Note that all of the terms in the in the function are non-negative, and so

$$0 \leq \frac{x^2}{\sqrt{x^2+y^2}} \leq \frac{x^2+y^2}{\sqrt{x^2+y^2}} = \sqrt{x^2+y^2}$$

Using this inequality together with the Squeeze Theorem, we can see that

$$0 \leq \lim_{(x,y)\to(0,0)} \frac{y^2}{\sqrt{x^2+y^2}} \leq \lim_{(x,y)\to(0,0)} \sqrt{x^2+y^2} = 0$$

Thus, by the squeeze theorem, the limit is 0.

Note that there is an alternative way we could have used the squeeze theorem. We can write x^2 as $|x||x|$, using inequality 6 from Proposition 3.1.2, if we multiply all sides by x^2, we get

$$0 \leq \frac{x^2}{\sqrt{x^2+y^2}} \leq \frac{x^2}{|x|} = \frac{|x||x|}{|x|} \quad \Longrightarrow \quad 0 \leq \frac{x^2}{\sqrt{x^2+y^2}} \leq |x|$$

Now we're in a good position to use the Squeeze Theorem, much in the same way we did above. Doing so gives

$$0 \leq \lim_{(x,y)\to(0,0)} \frac{x^2}{\sqrt{x^2+y^2}} \leq \lim_{(x,y)\to(0,0)} |x| = 0$$

Once again by the Squeeze Theorem, we can conclude that the limit is 0. □

Often when we try to bound our function to use the squeeze theorem, our bounds are in terms of the absolute value of our function. In the case where we wish to show our function limits to 0, the following proposition is helpful.

Proposition 3.1.4

If f is a function such that $\lim_{(x,y)\to(a,b)} |f(x,y)| = 0$, then $\lim_{(x,y)\to(a,b)} f(x,y) = 0$.

Proof. I will leave this as an exercise for you to try on your own, we will go through it in my week 6 tutorials. *Hint*: $-|f(x,y)| \leq f(x,y) \leq |f(x,y)|$. □

Let's see another example, one that uses a bit of everything we've seen so far.

Example 3.1.5: A More Challenging Example

Show that $\displaystyle\lim_{(x,y)\to(0,0)} \frac{5x^\alpha y}{x^2+y^2} = 0$ for any value of $\alpha > 1$.

Solution. Let $\alpha > 1$ be given. We can make use of 7 from Proposition 3.1.2, which says that

$$\left|\frac{xy}{x^2+y^2}\right| \leq \frac{1}{2} \tag{$\star$}$$

We can multiply both sides of ($\star$) by $|5x^{\alpha-1}|$, giving us the inequality

$$0 \leq \left|\frac{5x^\alpha y}{x^2+y^2}\right| \leq \frac{5|x^{\alpha-1}|}{2}$$

Since $\alpha > 1$ we have $\alpha - 1 > 0$, meaning $\lim_{x\to 0} x^{\alpha-1} = 0$. By the Squeeze Theorem, we can see that

$$0 \leq \lim_{(x,y)\to(0,0)} \left|\frac{5x^\alpha y}{x^2+y^2}\right| \leq \lim_{(x,y)\to(0,0)} \frac{5\left|x^{\alpha-1}\right|}{2} = 0$$

Thus, $\lim_{(x,y)\to(0,0)} \left|\frac{5x^\alpha y}{x^2+y^2}\right| = 0 \Rightarrow \lim_{(x,y)\to(0,0)} \frac{5x^\alpha y}{x^2+y^2} = 0$ by Proposition 3.1.4. □

3.1.2 Showing a Limit Does Not Exist

Recall that in single variable calculus we said that a limit existed if and only if each of the one-sided limits existed and agreed, namely

$$\lim_{x\to a} f(x) = L \iff \lim_{x\to a^-} f(x) = L = \lim_{x\to a^+} f(x) = L$$

Unfortunately, this notion does not carry over nicely to functions of several variables, as we can approach our limit point (a, b) from an infinite number of directions, or paths. The generalization we are interested in is that the limit of a multivariable function must be equal along every possible path to the limit point in order to exist. Of course, it's impossible to check every single path, so this is not a practiacal way to determine if a limit exists. However, it does inspire the following criterion for multivariable limits, which will provide us with a useful means of showing that a limit does not exist.

Proposition 3.1.6: Limit Criterion

Let $f : U \subseteq \mathbb{R}^2 \to \mathbb{R}$ be a function on an open set U containing a point (a, b). If $\lim_{(x,y)\to(a,b)} f(x, y) = L$, then for any continuous path $\gamma : (-\varepsilon_0, \varepsilon_0) \to U$ such that $\gamma(0) = (a, b)$ and $\gamma(t) \neq (a, b)$ for $t \neq 0$ we have $\lim_{t\to 0}(f \circ \gamma)(t) = L$.

Let's take a step back and see what this is saying. If we know that the limit exists, then we know that the limit exists along any continuous path to the limit point. This gives us a nice way of showing that a limit does not exist, namely the contrapositive of the Limit Criterion says that if there is two continuous paths, along which the value of the limit differs, then the limit of the function does not exists. More concretely, if $f(x, y)$ is a function, and γ_1, γ_2 are two continuous paths such that $\gamma_1(0) = (a, b) = \gamma_2(0)$, and

$$\lim_{t\to 0}(f \circ \gamma_1)(t) \neq \lim_{t\to 0}(f \circ \gamma_2)(t)$$

Then $\lim_{(x,y)\to(a,b)} f(x, y)$ does not exist. This will be our primary technique to show that a limit does not exist. Let's see an example of this technique in action.

Example 3.1.7

Show that the limit $\lim_{(x,y)\to(0,0)} \frac{2xy}{x^2+2y^2}$ does not exist.

Solution. A good approach with limits in general is try to some simple paths to determine the behaviour of the function. Why don't we try along the line $y = x$, which we can parameterize as the path $\gamma_1(t) = (t, t)$ (here's a place your parameterization skills will be useful!) To compute the limit along the path γ_1, we can compose $f(x, y) = \frac{2xy}{x^2+2y^2}$ with $\gamma_1(t)$ and evaluate the limit as $t \to 0$. Doing so gives

$$\lim_{t\to 0}(f \circ \gamma_1)(t) = \lim_{t\to 0} \frac{2t^2}{t^2 + 2t^2} = \frac{2}{3}$$

This doesn't tell us that the limit doesn't exist, all it says is that if the limit does exist, then the value must be $\frac{2}{3}$ (the limit must agree along all paths to the limit point). If we can find another path to $(0,0)$ that gives a different limit, we can conclude that the limit doesn't exist. Let's try along the x-axis, namely along the path $\gamma_2(t) = (t,0)$. Doing so gives

$$\lim_{t\to 0}(f\circ\gamma_2)(t) = \lim_{t\to 0}\frac{0}{t^2} = 0$$

We've found two paths that give different limits, so we can conclude that the limit does not exist!

Note that we can be a bit more strategic than trying paths at random. One thing we can do is try a collection of paths. For example, what if we consider all linear paths to the origin of the form $y = mx$ for $m \in \mathbb{R}$. We can parameterize these as $\gamma_m(t) = (t, mt)$. If we compute the limit along the path $\gamma_m(t)$ and the value of the limit depends on m, then we can conclude that the limit does not exist. (To see why, we can simply pick different values of $m \in \mathbb{R}$) Let's try this out, the limit along $\gamma_m(t)$ is

$$\lim_{t\to 0}(f\circ\gamma_m)(t) = \lim_{t\to 0}\frac{2t(mt)}{t^2+2(mt)^2} = \lim_{t\to 0}\frac{2mt^2}{t^2(1+2m^2)} = \frac{2m}{1+2m^2}$$

This depends on $m \in \mathbb{R}$, and so we can conclude that the limit does not exist. Note that in our initial answer, we used the paths that correspond to $m = 1$ and $m = 0$. □

The collection of paths method is an effective one, it allows us to try many paths at the same time. We're not limited to just linear paths, we can try quadratic paths as well, namely $y = mx^2$, which we can write as $\alpha_m(t) = (t, mt^2)$, or even $\beta_m(t) = (mt^2, t)$, the latter of which corresponds to $x = my^2$, or any other family of paths. Let's see another example that's a bit more stubborn.

Example 3.1.8: A Stubborn One

Let $f(x,y) = \frac{x^2y}{x^4+y^2}$. Show that $\lim\limits_{t\to 0}(f\circ\alpha)(t) = 0$ along any linear path $\alpha(t)$ to the origin, but $\lim\limits_{(x,y)\to(0,0)} f(x,y)$ does not exist.

Solution. Linear paths to the origin are of the form $y = mx$ or $x = ny$ for $m, n \in \mathbb{R}$. We can consider each case. Let $\alpha_m(t) = (t, mt)$, the limit along $\alpha_m(t)$ is

$$\lim_{t\to 0}(f\circ\alpha_m)(t) = \frac{t^2(mt)}{t^4+m^2t^2} = \frac{t^2(mt)}{t^2(t^2+m^2)} = \lim_{t\to 0}\frac{mt}{t^2+m^2} = 0$$

Similarly, along the path $\alpha_n(t) = (nt, t)$, the limit is

$$\lim_{t\to 0}(f\circ\alpha_n)(t) = \frac{(n^2t^2)t}{n^4t^4+t^2} = \frac{t^2(n^2t)}{t^2(n^4t^2+1)} = \lim_{t\to 0}\frac{n^2t}{n^4t^2+1} = 0$$

This shows that the limit is 0 along any linear path to $(0,0)$, so we should try another type of path. Note that the powers of x are always twice as high as the powers of y inside $f(x,y)$. One way we can even things out is to take quadratic paths where $y = mx^2$, namely along $\gamma_m(t) = (t, mt^2)$. Doing so gives

$$\lim_{t\to 0}(f\circ\gamma_m)(t) = \lim_{t\to 0}\frac{t^2(mt^2)}{t^4+m^2t^4} = \lim_{t\to 0}\frac{m^2t^4}{t^4(1+m^2)} = \frac{m^2}{1+m^2}$$

As the limit depends on the choice of $m \in \mathbb{R}$, we conclude $\lim_{(x,y)\to(0,0)} f(x,y)$ does not exist. □

3.1.3 Polar Coordinates

Polar coordinates is our second main way of showing that a limit exists. The idea is as follows, polar coordinates are defined by

$$(x, y) = \mathbf{G}(r, \theta) = (r\cos(\theta), r\sin(\theta))$$

so if we have a limit of the form $\lim\limits_{(x,y)\to(0,0)} f(x, y)$, ideally where $f(x, y)$ contains terms like $x^2 + y^2$, we can compute the limit of $f(\mathbf{G}(r, \theta)) = f(r\cos(\theta), r\sin(\theta))$. Why this works is because our polar coordinate system satisfies $x^2 + y^2 = r^2$, and so as $(x, y) \to (0, 0), r^2 \to 0$ and hence $r \to 0$. The difficulty of using polar coordinates lies with the fact that we don't know anything about the behaviour of θ. In general, it's not true that $\theta \to 0$, because of this, we take a limit as $r \to 0$ with the understanding that θ is still a variable and hence not a constant. In a general case, we'll have something like this

$$\lim_{(x,y)\to(0,0)} f(x, y) \xrightarrow{\text{Polar Coordinates}} \lim_{r\to 0}(f \circ \mathbf{G})(r, \theta) = \lim_{r\to 0} f(r\cos(\theta), r\sin(\theta))$$

Polar coordinates can be a useful means of calculating limits of functions of two variables, but we must be careful when doing so. **Remember that θ does <u>not</u> need to approach 0, we do not know the behaviour of θ in general, only that $r \to 0$.** That being said, let's see some examples of how we can use polar coordinates to evaluate limits, and how things can go wrong if we are not careful.

Example 3.1.9

Evaluate the limit $\lim\limits_{(x,y)\to(0,0)} \dfrac{\sin(x^2 + y^2)}{x^2 + y^2}$

Solution. We can see some $x^2 + y^2$ terms inside of the function, which is a good indicator that we should try polar coordinates. When we change coordinates, we get the limit

$$\lim_{(x,y)\to(0,0)} \frac{\sin(x^2 + y^2)}{x^2 + y^2} \xrightarrow{\text{Polar Coordinates}} \lim_{r\to 0} \frac{\sin(r^2)}{r^2}$$

As we have terms that only involve r, we can treat this a limit of one variable, and use what we know from single variable calculus to evaluate it. Recall that $\lim\limits_{x\to 0} \frac{\sin(x)}{x} = 1$, so if we make the substitution $x = r^2$, as $r \to 0, x = r^2 \to 0$, and so

$$\lim_{r\to 0} \frac{\sin(r^2)}{r^2} = \lim_{x\to 0} \frac{\sin(x)}{x} = 1$$

Thus, we can conclude that $\lim_{(x,y)\to(0,0)} \frac{\sin(x^2+y^2)}{x^2+y^2} = \lim_{r\to 0} \frac{\sin(r^2)}{r^2} = 1$. □

In this example, when we converted from Cartesian to polar, the resulting function was in terms of only r, which worked out nicely for us. Let's see what happens in a more complicated case, when the resulting limit depends on θ as well.

Example 3.1.10

Determine $\lim\limits_{(x,y)\to(0,0)} \dfrac{x^3 + y^3}{x^2 + y^2}$.

Solution. Right away, we can see an x^2+y^2 term in the denominator, which is a sign to try polar coordinates. Changing coordinates gives us a new limit

$$\lim_{(x,y)\to(0,0)} \frac{x^3+y^3}{x^2+y^2} \xrightarrow{\text{Polar Coordinates}} \lim_{r\to 0} \frac{r^3\cos^3(\theta)+r^3\sin^3(\theta)}{r^2}$$

We can factor the r^3 term on the top, and cancel out with r^2 on the bottom, leaving us with

$$\lim_{r\to 0} r(\cos^3(\theta)+\sin^3(\theta))$$

If we naively treat $\cos^3(\theta)+\sin^3(\theta)$ as a constant since the limit is in terms of r, we get that the limit is 0, since $r\to 0$. Note that this is the correct answer, but the reasoning is **wrong!** Why this is 0 is as follows, we can bound $r(\cos^3(\theta)+\sin^3(\theta))$ by something that does not involve θ, namely

$$\begin{aligned} 0 \le |r(\cos^3(\theta)+\sin^3(\theta))| &= |r\cos^3(\theta)+r\sin^3(\theta)| \\ &\le |r|\underbrace{|\cos^3(\theta)|}_{\le 1}+|r|\underbrace{|\sin^3(\theta)|}_{\le 1} \\ &\le 2|r| \end{aligned}$$

Note that we used 1,8 and 9 from Proposition 3.1.2 to bound our function. Now that our upper bound is free of θ, and as $2|r|\to 0$ as $r\to 0$, the squeeze theorem tells us that

$$0 \le \lim_{r\to 0}|r(\cos^3(\theta)+\sin^3(\theta))| \le \lim_{r\to 0} 2|r| = 0 \xrightarrow{\text{Prop. 3.1.4}} \lim_{r\to 0} r(\cos^3(\theta)+\sin^3(\theta)) = 0$$

and so we can conclude that $\lim\limits_{(x,y)\to(0,0)} \frac{x^3+y^3}{x^2+y^2} = \lim\limits_{r\to 0} r(\cos^3(\theta)+\sin^3(\theta)) = 0.$ □

I'm sure you're wondering why we had to do all of this extra work when it seems pretty obvious that the limit should be 0. The idea is that whenever we have a limit with $r\to 0$ of a function of the form $g(r)f(\theta)$, where $\lim_{r\to} g(r)=0$, we need $f(\theta)$ to be bounded (by something free of θ) to conclude that $\lim_{r\to 0} g(r)f(\theta)=0$ (by the squeeze theorem), because θ is not a constant, and so $f(\theta)$ is changing. Let's see an example of where ignoring this can lead to problems.

Example 3.1.11: Where Polar Coordinates Can Fail

Show that $\lim\limits_{(x,y)\to(0,0)} \dfrac{x^2+y^2}{x}$ does not exist.

Solution. Let $f(x,y)=\frac{x^2+y^2}{x}$, and consider paths of the form $\alpha(t)=(mt^2,t)$ for $m\neq 0$, these correspond to $x=my^2$. Taking the limit along the path $\alpha(t)$ gives

$$\lim_{(x,y)\to(0,0)} \frac{x^2+y^2}{x} = \lim_{t\to 0}(f\circ\alpha)(t) = \lim_{t\to 0}\frac{m^2t^4+t^2}{mt^2} = \lim_{t\to 0}\frac{m^2t^2+1}{m} = \frac{1}{m}$$

The limit clearly depends on the choice of m, and so the limit does not exist.

However, if we use polar coordinates, the function becomes $\frac{r^2}{r\cos(\theta)}=\frac{r}{\cos(\theta)}$, if we naively treat θ as a constant, we get that

$$\lim_{(x,y)\to(0,0)} \frac{x^2+y^2}{x} = \lim_{r\to 0}\frac{r}{\cos(\theta)} \stackrel{?}{=} 0$$

since $r\to 0$. As we saw above, this limit does not exist, which means our approach goes wrong at some point. The reason why this fails is because θ is **not** a constant, and so the function $\frac{r}{\cos(\theta)}$ is undefined for $\theta=\frac{\pi}{2}$, and hence unbounded near $\theta=\frac{\pi}{2}$. □

To summarize, when you are using polar coordinates to evaluate a limit, and you are left with terms involving both r and θ, you need to bound the entire function by something that does not depend on θ, and use the squeeze theorem, as in Example 3.1.10. We can formalize this into a Criterion that we can use to help us out.

Proposition 3.1.12: Boundedness Criterion

Suppose $f(x, y)$ is a function such that $f(r\cos(\theta), r\sin(\theta)) = g(r)h(\theta)$ for single variable functions g and h. If h is globally bounded, namely there is $D > 0$ such that $|h(\theta)| \leq D$ for all θ and $\lim_{r\to 0} g(r) = 0$, then $\lim_{(x,y)\to(0,0)} f(x, y) = 0$.

Proof. Using polar coordinates, by assumption we have that $f(r\cos(\theta), r\sin(\theta)) = g(r)h(\theta)$, so

$$\lim_{(x,y)\to(0,0)} f(x, y) = \lim_{r\to 0} f(r\cos(\theta), r\sin(\theta)) = \lim_{r\to 0} g(r)h(\theta)$$

By assumption, we know that $|h(\theta)| \leq D$ for all θ, so using the squeeze theorem will be a good idea. Note that $|g(r)h(\theta)| = |g(r)||h(\theta)| \leq D|g(r)|$, and by the squeeze theorem, it follows that

$$\lim_{r\to 0} -Dg(r) \leq \lim_{r\to 0} g(r)h(\theta) \leq \lim_{r\to 0} Dg(r)$$

As $\lim_{r\to 0} g(r) = 0$, by the limits laws, $\lim_{r\to 0} \pm Dg(r) = 0$, and so by the squeeze theorem

$$\lim_{(x,y)\to(0,0)} f(x, y) = \lim_{r\to 0} g(r)h(\theta) = 0$$

□

Example 3.1.13

Use the boundedness criterion to show that $\displaystyle\lim_{(x,y)\to(0,0)} \frac{x^3 + y^3}{x^2 + y^2} = 0$.

Proof. As we saw in Example 3.1.10 we can write $\frac{x^3+y^3}{x^2+y^2}$ in polar coordinates as $r(\cos^3(\theta) + \sin^3(\theta))$. Here we can see $g(r) = r$ and $h(\theta) = \cos^3(\theta) + \sin^3(\theta)$. Certainly $\lim\limits_{r\to 0} g(r) = 0$, and by the triangle inequality

$$|h(\theta)| = |\cos^3(\theta) + \sin^3(\theta)| \leq |\cos^3(\theta)| + |\sin(\theta)| \leq 2$$

Thus, by Proposition 3.1.12, we can conclude that $\lim_{(x,y)\to(0,0)} \frac{x^3+y^3}{x^2+y^2} = 0$. □

3.1.4 Proofs of Theorems and Propositions (Optional)

As noted, this section is completely optional, and you are by no means required to understand any of this. I include the proofs for those interested, and for the sake of completeness. Both of the theorems we'll prove here can be stated for functions on $\mathbb{R}^n$, and so we will state and prove the more general results, with the understanding that the proofs will suffice for theorems above when $n = 2$. We can begin by introducing the formal definition of a limit.

Definition 3.1.14: Formal Limits

Let $\mathbf{a} \in \mathbb{R}^n$ and $f : U \to \mathbb{R}$ a function defined on an open neighbourhood of $\mathbf{a}$, except possibly at $\mathbf{a}$. We say that *the limit of $f(x)$ as $\mathbf{x}$ approaches $\mathbf{a}$ is L* if

$$(\forall \varepsilon > 0)(\exists \delta > 0)(\forall \mathbf{x} \in U)[0 < \|\mathbf{x} - \mathbf{a}\| < \delta \implies |f(\mathbf{x}) - f(\mathbf{a})| < \varepsilon]$$

In which case we write $\lim_{\mathbf{x}\to\mathbf{a}} f(\mathbf{x}) = L$.

Theorem 3.1.15: Squeeze Theorem

Suppose $f, g, h : U \subseteq \mathbb{R}^n \to \mathbb{R}$ are functions defined on an open neighbourhood U of $\mathbf{a} \in \mathbb{R}^n$. If there is an open neighbourhood $D \subseteq U$ of $\mathbf{a}$ such that $f(\mathbf{x}) \leq g(\mathbf{x}) \leq h(\mathbf{x})$ for all $\mathbf{x} \in D \setminus \{\mathbf{a}\}$, and

$$\lim_{\mathbf{x}\to\mathbf{a}} f(\mathbf{x}) = L = \lim_{\mathbf{x}\to\mathbf{a}} h(\mathbf{x})$$

Then $\lim_{\mathbf{x}\to\mathbf{a}} g(\mathbf{x}) = L$.

Proof. Let $\varepsilon > 0$ be given, as $\lim\limits_{\mathbf{x}\to\mathbf{a}} f(\mathbf{x}) = L$ and $\lim\limits_{\mathbf{x}\to\mathbf{a}} h(\mathbf{x}) = L$, $\exists \delta_1, \delta_2 > 0$ such that

$$0 < \|\mathbf{x} - \mathbf{a}\| < \delta_1 \Rightarrow |f(\mathbf{x}) - L| < \varepsilon \quad \text{and} \quad 0 < \|\mathbf{x} - \mathbf{a}\| < \delta_2 \Rightarrow |h(\mathbf{x}) - L| < \varepsilon \tag{1}$$

Take $\delta = \min\{\delta_1, \delta_2\} > 0$ and suppose $0 < \|\mathbf{x} - \mathbf{a}\| < \delta$. Then from (1) we have

$$-\varepsilon < f(\mathbf{x}) - L \quad \text{and} \quad h(\mathbf{x}) - L < \varepsilon$$

Using the inequality $f(\mathbf{x}) \leq g(\mathbf{x}) \leq h(\mathbf{x})$, we can subtract L from each term and combine with the conditions from (1) to get

$$-\varepsilon < f(\mathbf{x}) - L \leq g(\mathbf{x}) - L \leq h(\mathbf{x}) - L < \varepsilon \iff |h(\mathbf{x}) - L| < \varepsilon$$

□

Proposition 3.1.16: Limit Criterion

Let $f : U \subseteq \mathbb{R}^n \to \mathbb{R}$ be a function on an open set U containing a point $\mathbf{a} \in \mathbb{R}^n$. If $\lim\limits_{\mathbf{x}\to\mathbf{a}} f(\mathbf{x}) = L$, then for any path $\gamma : (-\varepsilon_0, \varepsilon_0) \to U$ such that $\gamma(0) = \mathbf{a}$ and $\gamma(t) \neq \mathbf{a}$ for $t \neq 0$ we have $\lim_{t\to 0}(f \circ \gamma)(t) = L$.

Proof. Let $\varepsilon > 0$ be given, as $\lim\limits_{\mathbf{x}\to\mathbf{a}} f(\mathbf{x}) \to L$ and γ is continuous (at $t = 0$), $\exists \delta, \delta_1 > 0$ such that

$$0 < \|\mathbf{x} - \mathbf{a}\| < \delta_1 \Rightarrow |f(\mathbf{x}) - L| < \varepsilon \quad \text{and} \quad 0 < |t| < \delta \Rightarrow 0 < \|\gamma(t) - \mathbf{a}\| < \delta_1 \tag{2}$$

Note that $\|\gamma(t) - \mathbf{a}\| > 0$ as $\gamma(t) \neq \mathbf{a}$ for all $t \neq 0$ by assumption. Putting everything together from (2) gives

$$0 < |t| < \delta \Rightarrow 0 < \|\gamma(t) - \mathbf{a}\| < \delta_1 \Rightarrow |f(\gamma(t)) - L| < \varepsilon$$

□

4 Student Reviews

Calculus of Several Variables

1. *Great job! This tutorial class has honestly been so enjoyable. I was able to clarify my questions and was really able to grasp new concepts.* [Fall 2020]

2. *I really liked the way of teaching as we went along during the tutorial. Outside of the tutorial, the solutions posted were very in-depth and helpful. Thank you Nigel, you really helped me out this semester! All the best with everything!* [Fall 2020]

3. *I liked how Nigel answered everyone's questions and would stay longer to just talk about random university courses and stuff. It was a fun learning environment. Also the posted notes were very detailed and clear.* [Summer 2021]

4. *I think the structure of the tutorial worked well, with a few minutes for the students to work on the problems themselves, then later taking it up as a group. I found it very helpful in improving my understanding of the content.* [Summer 2021]

5. *Please continue explaining the theorems and going through the questions in detail. Also, I really like how you choose specific questions that covered a lot of the important techniques we needed to know.* [Fall 2021]

6. *Probably one of the most helpful math tutorials I've had in university.* [Fall 2021]

7. *All of the questions we did during the tutorial were extremely helpful! thank you making math more enjoyable! thank you for a great semester!* [Winter 2022]

8. *Very helpful, any questions or doubts were cleared and thoroughly explained well. Types of questions we tackled were extremely beneficial to our understanding and helped with midterms. Even when a question was difficult, we walked through it step by step and I never felt anything was explained poorly.* [Winter 2022]

Analysis I & Analysis II

1. *Overall, the TA was really helpful and did a great job at teaching the students.* [Fall 2022 - Winter 2023]

2. *I found the tutorials very helpful, and the clear solutions helped when reviewing for tests. They also helped when working on the homework. Overall, tutorials were fun and informative.* [Fall 2023 - Winter 2024]

3. *Nigel's tutorial always keeps up to date, he put things that we didn't learn in class but were super useful in practice in his tutorial. He's always super willing to answer our questions, whether they're dumb or not, and he explains them clearly during all of his office hours. When grading our homework, he always give us useful feedback that can help us improve in proof strategies as well as our understanding of the concepts. I really appreciate Nigel's help through this school year!* [Fall 2023 - Winter 2024]

4. *Tutorials were extremely useful, helping me get adjusted to the course coming from MAT137 and giving me very useful problems related to the homework.* [Winter 2024]