Mathematics, on the Back of a Napkin

Muhammed Tariq

April 2025

> "What can you really know
> without knowing anything?"

1 Introduction

At the age of twenty-three, Sir Isaac Newton upended a thousand years of Aristotelian thought with the publication of his equation for universal gravitation:

$$F = G\frac{m_1 m_2}{r^2}$$

Even reinforced by Greek astronomer Claudius Ptolemy, the orthodox belief that objects only moved if an external force drove that motion – centuries of formerly founded thinking – was rendered obsolete overnight, the foundations of classical thought torn through with the inexorability of a hurricane.

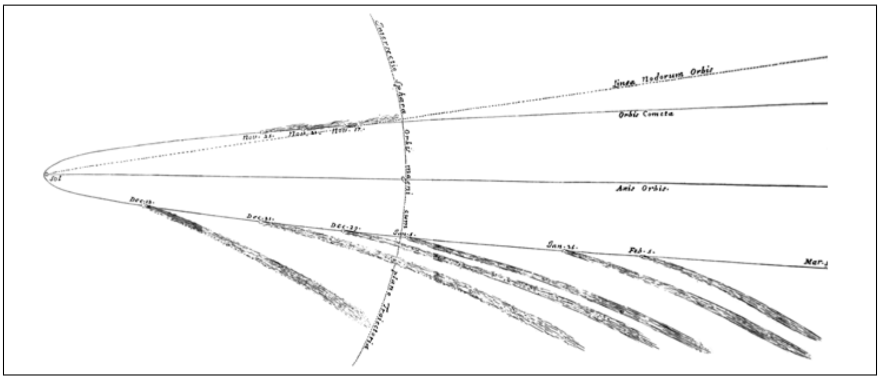

Figure 1: Isaac Newton demonstrated the Universal Law of Gravitation by showing that a comet visible during 1680 and 1681 followed the path of a parabola. [1]

At the age of sixteen, I set out to discover how many tennis balls would fit inside the Eiffel Tower. One of us made history; the other, a very strange spreadsheet. While the product of these efforts isn't out of the passion I have for mass-produced leisure-grade foam-core spheres or France's most famous lattice tower, it does reflect something else entirely.

At twenty-nine minutes past five, on a Monday morning in July of 1945, the world's first atom bomb detonated in the desert ninety-six kilometres north-west of Alamogordo, New Mexico. A stunned forty seconds later, contemplating this historic spectacle, a cluster of scientists stood at the base camp, where the shockwave burst through. The Italian-American physicist Enrico Fermi was the first to stir, witnessing the culmination of the project he had helped initiate.

Prior to the explosion, Fermi tore a sheet of notebook paper into small bits. As he felt the first quiver of the shockwave spreading outwards, he released the shreds into the air above his head, which fluttered down and away from the ever-growing mushroom cloud on the horizon, coming to a rest on the ground about two and a half yards behind him. After a brief mental calculation, Fermi proclaimed that the bomb's energy was equivalent to that produced by ten thousand tons of TNT (10 kilotons, or KT). Calculations from more sophisticated areas near the blast zone revealed the now-accepted value of 21KT [2], well within an order of magnitude of accuracy, which was really all that was required.

Enrico Fermi's extraordinary ability to make excellent approximate calculations with little to no actual data justifies the coinage of what is known today as a **Fermi problem** – an estimation problem involving dimensional analysis (an examination of the relationships between different physical quantities), or an approximation of extreme scientific calculations.

Enlightened by Fermi's paper scraps and XKCD's "*What If?*", and perhaps driven by a similar desire to understand the unknowable, I decided to tackle a few Fermi problems of my own. Some serious. Some... less so.

Let's embark on a trip to France, together.

Figure 2: The Eiffel Tower. [3]

2 The Fermi Problems

2.1 La Tour aux Balles

QUESTION: How many ITF-approved, pressurised Type 2 tennis balls fit in the Eiffel Tower?

RESPONSE: A few billion. Give or take a baguette.

Tennis balls, at the tour level, have a surprisingly non-negligible variation in their diameters. The International Tennis Federation (ITF) defines their official, standardised diameters by class:

- Type 1 (fast): 6.54 - 6.86cm
- Type 2 (medium) 6.54 - 6.86cm
- Type 3 (slow) 7.00 - 7.30cm
- Type 4 (high altitude) 6.54 - 6.86cm

The ITF maintains that Type 3 and Type 4 are only allowed for use on tennis courts 1219m above sea level (to balance out the increased surface area leading to more drag with the reduced air density at higher altitudes leading to less of it), so it seems most reasonable to average the range for the most conventionally used tennis ball, a Type 2 ball, giving a diameter of 6.7cm (0.067m) [4].

It is surprising that there are no accurate estimates, at all, regarding the volume enclosed by the shape of the Eiffel Tower. Modelling it in exact curved fashion would make this problem needlessly complex for an order-of-magnitude estimation – it can be argued that the sheer increase in the amount of work required is not worth the uncertain gain in accuracy.

Therefore, most reasonably, a spherical cow would be considered[1] by simplifying the tower to a square-based pyramid, with base width 124.9m and height 309.6m [5]. The following calculations are, as a result, very simple:

$$\begin{aligned} V &= \frac{1}{3}b^2h \\ &= \frac{1}{3}(124.9)^2(309.6) \\ &= 1609921.032\text{m}^3 \\ &\approx 1\,600\,000\text{m}^3 \end{aligned}$$

It would be dishonest to say that I'm satisfied with this approximation, however. As the tower in actuality curves inwards, the calculated volume is a comically vast overestimate (see Figure 3). This doesn't hold a candle to the order of magnitude that the final answer would have, so it would be entirely appropriate for a Fermi estimation, but I still wanted to do better.

The tower can be modelled more realistically as **a solid of revolution**[2] by using a function which describes how its cross-sectional width narrows with height, which we could then obtain a function for the variation in cross-sectional area $A(y)$ from, and then integrate for the volume:

$$V = \int_0^h A(y)\,\mathrm{d}y$$

Through a visual judgement of the shape of the Eiffel Tower, an exponential function appears to be most accurate in capturing the shape of the curve of one of its edges.[3] In any case, the error gained in incorrectly assuming the scaling law of the taper (how it changes as you go up – linear, exponential or something entirely different) is severely outweighed by the magnitude of the error in assuming that the tower is a square-based pyramid, so going down this route is still best.

The general form for an exponential function is $y = ae^{bx}$ where, in this scenario, x represents the height, y denotes the "radius" of the cross-section at that height (the half-width), and a and b represents constants to be determined from data. We can graph this exponential function by considering various

[1] Once upon a time, milk production at a dairy farm was low, so the farmer wrote to the local university, asking for help from academia. A multidisciplinary team of professors was assembled, headed by a theoretical physicist, and two weeks of intensive on-site investigation took place. The scholars then returned to the university, notebooks crammed with data, where the task of writing the report was left to the team leader. Shortly thereafter the physicist returned to the farm, saying to the farmer, 'I have the solution, but it works only in the case of spherical cows in a vacuum.' If assumptions were currency, theoretical physicists would be trillionaires by lunchtime.

[2] A 3D object created by rotating a 2D shape or curve around an axis. You could alternatively break it down into geometric primitives (simple 3D shapes) instead, but I wanted to have some more fun with this.

[3] Their use is also likely; they're generally used in skyscrapers to distribute load effectively and minimise material usage.

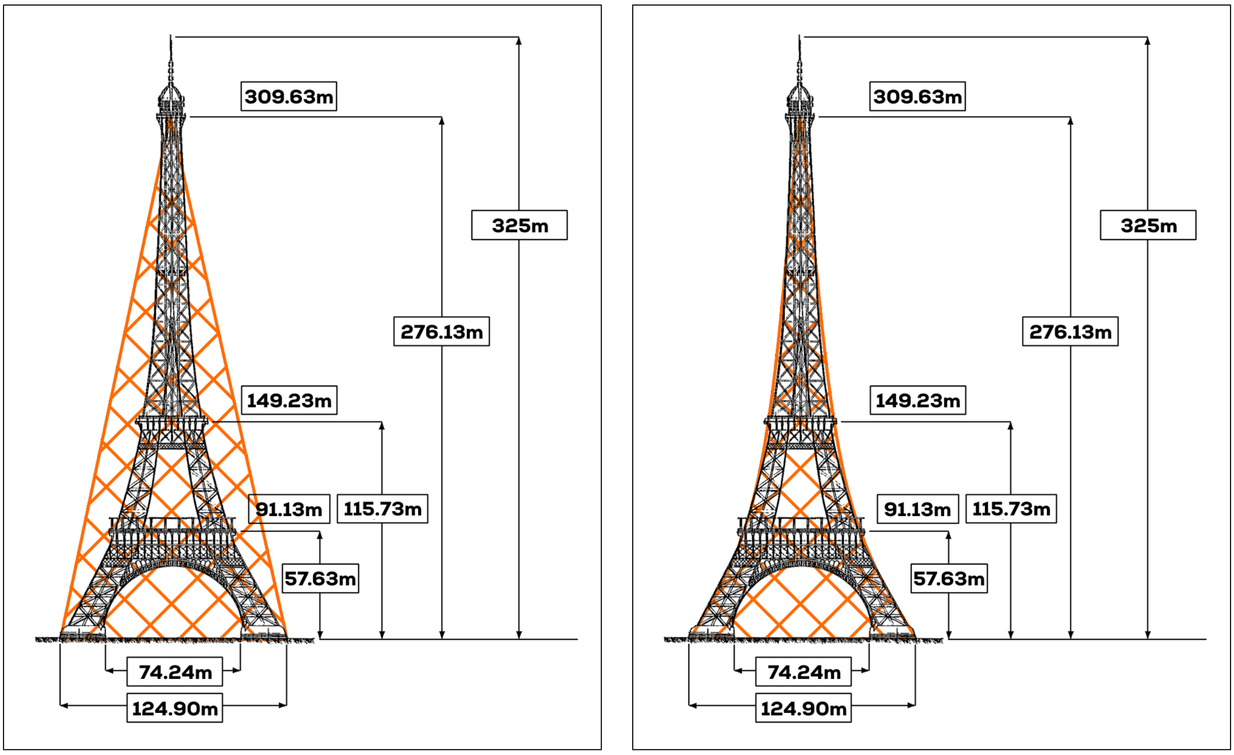

(a) Volume as a square-based pyramid: $V = \frac{1}{3}b^2h$

(b) Volume with exponential taper: $V = \int_0^h 4a^2e^{2bx}\,\mathrm{d}x$

Figure 3: Comparison of Eiffel Tower volume modelling approaches. The orange cross-hatched area represents the "volume" calculated relative to the actual volume of the tower. [6]

points on the Eiffel Tower at its different stages. Sources report slightly varying platform dimensions depending on whether they measure from edge-to-edge or include usable floor area, but after a painstaking analysis of the dimensions, I determined these co-ordinates that are likely to fit on the graph of the shape of the curve, in the form of height against half-width (x, y):

- (0, 62.45); ground level
- (57.63, 31.68); first floor
- (115.73, 18.44); second floor
- (276.13, 8.13); third floor.

When these co-ordinates are plotted in Desmos, it does certainly look somewhat like an exponential curve!

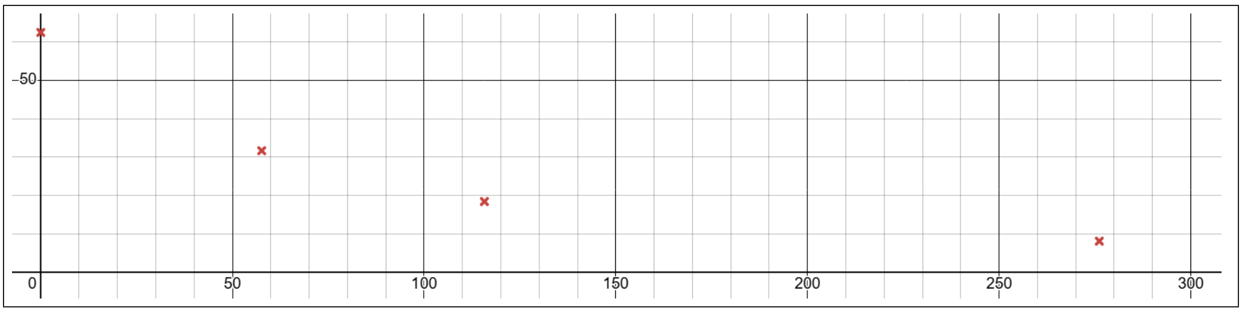

Figure 4: A graph of height against the half-width of the Eiffel Tower.

We can now take $\ln y$ against x, which allows us to confirm the curve's exponential nature as this should result in a straight line plot, for which we can draw a line of best fit. Going back to the general form for an exponential curve:

$$\begin{aligned} y &= ae^{bx} \\ \Rightarrow \ln y &= \ln\left(ae^{bx}\right) \\ \ln y &= \ln a + bx \\ \ln y &= bx + \ln a \end{aligned}$$

Evaluating $\ln y$ for all data points gives 4.134, 3.455, 2.915, and 2.095 respectively for each point, to three decimal places. Plotting $\ln y$ against x, and drawing a line of best fit[4], yields:

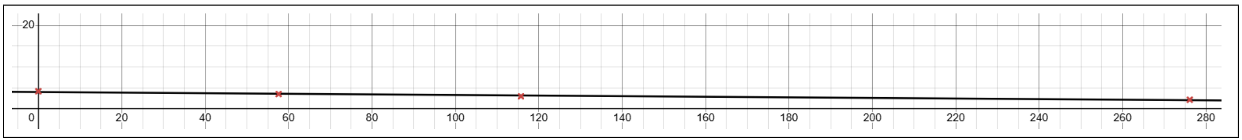

Figure 5: A graph of height against the natural log of the half-width of the Eiffel Tower.

This means that the gradient $b = -0.0070862$, and the y-intercept $\ln a = 3.94604$, which means that $a = e^{3.94604} \approx 51.73011$. This gives an exponential curve function of:

$$y = 51.73011e^{-0.0070862x}$$

As this considers the half-width of the Eiffel Tower, this function can be called $W_h(x)$. When plotted over the co-ordinates before, it lines up somewhat well, but very much imperfectly:

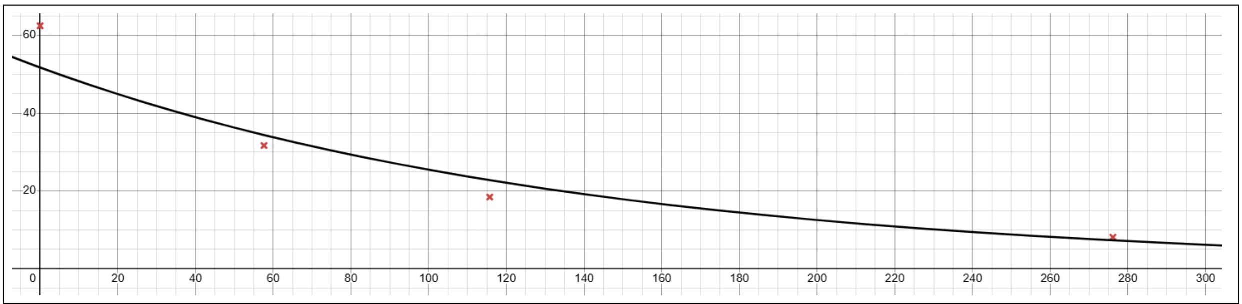

Figure 6: The exponential model $W_h(x) = 51.73011e^{-0.0070862x}$ compared to real edge data points. While the curve underestimates the width at the base ($x = 0$), it aligns more closely with the tower's actual shape in the upper regions — where most of the volume lies. This supports the model's use for a reasonable Fermi approximation despite its imperfect fit.

This is expected when the points span such a wide range, and when the data doesn't perfectly follow a single exponential decay all the way down – which most real towers don't. After all, if you look at any real image of the Eiffel Tower, the curve of the tower changes abruptly through its stages, almost like a piecewise function (which uses different line equations for different intervals). Using them could increase the accuracy of the model, albeit at an unjustifiably high level of effort even compared to the excessiveness of what we're doing now – the marginal increase in precision is greatly outweighed by the added complexity, offering no significant improvement in the context of an order-of-magnitude estimation. In any case, since most of the tower's volume lies within the range where it does fit well, the error is now certainly negligible in the context of a Fermi problem.

The beauty of Fermi problems is that if you think you have made an underestimation or overestimation in a particular value, you can adjust other values to compensate for that (at your own peril). Generally, overestimates and underestimates tend to help cancel each other out.

We can move forward by using calculus to estimate the volume of the tower, by integrating the square cross-sectional area over the tower's height. Let $W(x)$ denote the full width of the tower and $A(x)$ denote

[4] I could have also skipped this entire step and used exponential regression, a built-in function in Desmos. I chose this approach because it doesn't require any special techniques and can be done with any spreadsheet or graphing software.

the cross-sectional area at height x, with V representing the volume of the tower:

$$\begin{aligned}
W(x) &= 2W_h(x) && \text{(full width is double the half-width)} \\
&= 2\left(51.73011e^{-0.0070862x}\right) \\
&= 103.46022e^{-0.0070862x} \\
A(x) &= W(x)^2 && \text{(assuming a perfect square cross-section)} \\
&= \left(103.46022e^{-0.0070862x}\right)^2 \\
&= 10704.01712e^{-0.0141724x} \\
V &= \int_0^h A(x)\,\mathrm{d}x \\
&= \int_0^{309.63} 10704.01712e^{-0.0141724x}\,\mathrm{d}x \\
&= 10704.01712\int_0^{309.63} e^{-0.0141724x}\,\mathrm{d}x \\
&= 10704.01712\left[\frac{1}{-0.0141724}e^{-0.0141724x}\right]_0^{309.63} && \text{(reverse of chain rule)} \\
&= 10704.01712\left[\frac{1}{-0.0141724}\left(e^{-0.0141724[309.63]} - e^{-0.0141724[0]}\right)\right] \\
&= \frac{10704.01712}{-0.0141724}\left(e^{-4.38820} - e^{0}\right) \\
&\approx -755272.016\,(-0.98758) \\
&\approx 745889.2184\text{m}^3
\end{aligned}$$

Rounded to two significant figures, the volume of the Eiffel Tower, V, is $750\,000\text{m}^3$. This differs from the square-based pyramid estimate by a factor of $2.1\dot{3}$, which was surprising to me initially until I had another look at the comparison between the modelling approaches. It does seem very likely that the square-based pyramid estimate is a massive overestimate, but I also think that my model underestimates the volume (to a lesser degree), and that the real value lies somewhere slightly higher. Nevertheless, it most definitely appears more accurate than the square-based pyramid we started with.

The volume of an average Type 2 pressurised tennis ball is obtained with a very simple calculation:

$$\begin{aligned}
V_b &= \frac{4}{3}\pi\left(\frac{0.067}{2}\right)^3 \\
&\approx 1.57479 \cdot 10^{-4}\,\text{m}^3
\end{aligned}$$

A very lazy approach to figuring out the number of tennis balls n that could fit would be to simply divide the volume of the tower by the volume occupied by each tennis ball. This would go forth as such:

$$\begin{aligned}
n &= \frac{V}{V_b} \\
&= \frac{750000}{1.57479 \cdot 10^{-4}} \\
&= 4762539767 \\
&\approx 5 \cdot 10^9
\end{aligned}$$

This approach gratuitously dismisses the entire field of sphere packing in mathematics, and is nothing short of a crime against geometry (even if it would suffice for an order-of-magnitude estimate).

Packing spheres in a given volume, say, 1m^3, depends primarily on the fraction of space occupied by them once they are jammed together in a mechanically stable arrangement. While random packings of identical spheres can vary with the packing efficiency φ being 54%, dense and efficient packings have a value of 64%, and the most efficient packings are at 74%. These limits arise from the balance between geometry (how neighbouring spheres fill space) and mechanical stability (the frictional contacts that keep spheres from slipping past each other) [7].

Under normal, lower-friction circumstances, spheres (e.g. atoms or marbles) naturally fall into close-packed structures because it's energetically favourable. However, with the tennis balls we're playing with here, their high-friction fuzzy surface means that they resist doing so, meaning that there is a much higher likelihood of ending up with random loose packing ($\varphi \approx 55\%$) if simply dropped in without care. Furthermore, friction could cause jamming where balls lock into place prematurely, leading to more voids and air gaps. Ultimately, the packing efficiency we'll end up with is dependent on whether we've decided to manually stack every individual tennis ball with care within the tower, or if we've simply dumped them in. For a more detailed explanation on the different approaches to sphere packing, see Appendix 4.1.

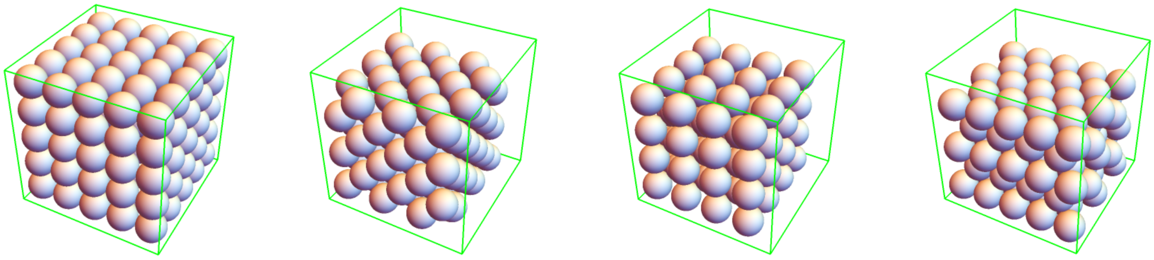

(a) SC; $\varphi \approx 52.4\%$ (b) FCC; $\varphi \approx 74\%$ (c) BCC; $\varphi \approx 68\%$ (d) HCP; $\varphi \approx 74\%$

Figure 7: A comparison of different packing approaches with spheres of diameter 1 in a container of length 5. While this may give the appearance that simple cubic (SC) packing is most efficient, this is only by virtue of the coincidence that a perfect number of spheres fit within the dimensions of the box. Even a tiny change in the box's dimensions could make SC packing vastly more inefficient, which is why its value of φ is much lower than the rest. Ultimately, the number of spheres that can fit in a container depends largely on its dimensions [8].

I will assume a more efficient packing as a means of compensating for a probable underestimate in the calculation of my volume, as well as because I want to consider a scenario in which one might genuinely attempt to maximize how many tennis balls could fit, which represents the true upper limit of the scenario. Given that $\varphi = 0.64$:

$$\begin{aligned} n &= \frac{V\varphi}{V_b} \\ &= \frac{750000 \cdot 0.64}{1.57479 \cdot 10^{-4}} \\ &= 3048025451 \\ &\approx 3 \cdot 10^9 \end{aligned} \qquad \square$$

In conclusion, the Eiffel Tower holds approximately 3 billion tennis balls. Serve accordingly.

Now that we're done – if you'll excuse me – I have a spreadsheet full of foam-core leisure-grade projectiles to close.

2.2 Interstellar Waldo

QUESTION: Suppose I have hidden an atom in a completely random location within the observable universe. How long would it take you to find it?

RESPONSE: Text me when you're 10^{63} years in. Call me if you've gone 10^{65} without any luck. Find me if entropy has maximised to no avail.

The observable universe is very big[citation needed]. Estimates place its diameter at about 93 billion light years, which increases at a rate of about $1.96 \cdot 10^9 \text{ms}^{-1}$, about 6.5 times faster than the speed of light in empty space [9].[5] Given that, it's a very good thing we're doing this now rather than later.

We will assume that, throughout this ordeal, you carry with you a magical ray gun that can identify whether any atom you point to using the gun is the atom I've hidden within 10^{-9} seconds, which has a range of 10^7m, meaning it can identify any specific atom within that radius. Foolishly idealistically, we'll also assume that you won't accidentally miss any atoms, and that you can super precisely point your gun at a specific atom.

[5] This faster-than-light motion is possible in curved space-time, or equivalently under the gravitational force's influence.

An absolute lower bound for the time taken, t, would be to randomly fire at every atom in the observable universe (ignoring the range of the gun for now) until you find the one I've placed. The mass

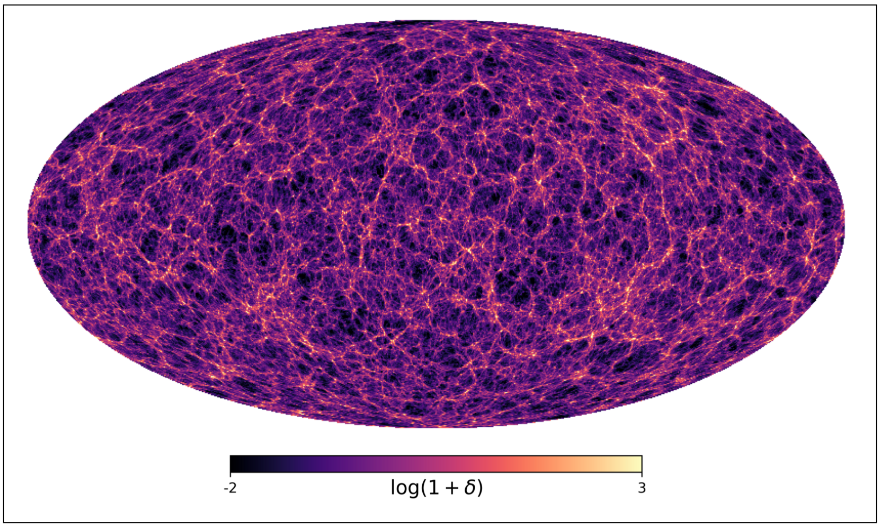

Figure 8: A map of matter density fluctuations within the observable universe [10].

of all matter, M, in the observable universe is frequently quoted as $1.45 \cdot 10^{53}$kg [11]. Assuming that all atoms are hydrogen atoms (which are about 74% of all atoms in the Milky Way by mass), you can estimate the total number of atoms N by dividing the mass of all ordinary matter by the mass of a hydrogen atom, m_h, which is $1.67 \cdot 10^{-27}$kg.

$$\begin{aligned} N &= \frac{M}{m_h} \\ &= \frac{1.45 \cdot 10^{53}}{1.67 \cdot 10^{-27}} \\ &\approx 8.68 \cdot 10^{79} \end{aligned}$$

Now, we search through each one of them every nanosecond:

$$\begin{aligned} t &= \left(8.68 \cdot 10^{79}\right)\left(10^{-9}\right) \\ &\approx 8.7 \cdot 10^{70}\text{s} \end{aligned}$$

This amounts to about $2.76 \cdot 10^{63}$ years, which is the time it takes to have a 100% chance of finding the atom. This amount of time is so unfathomably large that I've resorted to telling you to go to Appendix 4.2 if you'd like to visualise its sheer insanity.

Unfortunately, most of the matter in the universe is far out of reach for any human being. The main problem for a prospective atom-scanner would be the immense amount of time required for interstellar space travel (as our range is limited), which will certainly take up more time than individually scanning every atom.

Let's very generously assume that the maximum velocity man-made propulsion can manage is exactly the speed of light, $c \approx 3 \cdot 10^8 \text{ms}^{-1}$. Treating the observable universe as a sphere with radius 46.5 billion light years across, its volume, V_u, is calculated as follows:

$$\begin{aligned} 1\,\text{ly} &= \left(3 \cdot 10^8\right)(60)(60)(24)(365)\text{m} \\ &= 9.4608 \cdot 10^{15}\text{m} \\ \Rightarrow\; 46.5 \cdot 10^9\,\text{ly} &= \left(46.5 \cdot 10^9\right)\left(9.4608 \cdot 10^{15}\right)\text{m} \\ &= 4.399272 \cdot 10^{26}\text{m} \\ V_u &= \frac{4}{3}\pi\left(4.399272 \cdot 10^{26}\right)^3 \\ &\approx 3.56641 \cdot 10^{80}\text{m}^3 \end{aligned}$$

We will also assume that all the atoms in the universe are uniformly distributed. Although they absolutely are not, it should suffice for an average and certainly suffices for an order-of-magnitude estimate. The following, then, will be a calculation to determine how many atoms there are in an average spherical slice of the universe of radius 10^7m, $n_{(10^7)}$. Starting off with the number of atoms in 1m^3, n_1:

$$\begin{aligned} n_1 &= \frac{N}{V_u} \\ &= \frac{8.68 \cdot 10^{79}}{3.56641 \cdot 10^{80}} \\ &\approx 0.24338 \\ \Rightarrow\ n_{(10^7)} &= n_1 \cdot 10^7 \\ &= 0.24338 \cdot 10^7 \\ &= 2.4338 \cdot 10^6 \end{aligned}$$

Considering that only about 0.0000000000000000000042% of the universe contains any matter [12], I'm not too surprised that there's less than a quarter of an atom per cubic metre of space.

Surprisingly, we see a return to sphere packing! Considering that you're really smart, let's assume that you've used the most efficient face-centred cubic (FCC) arrangement to pack as many 10^7-radius spheres within the observable universe, as we're trying to check every region within the observable universe, all of which has the radius of our range. We'll also assume that travel between spheres while checking atoms constantly covers the empty regions around them, so their consideration has a negligible effect on the time taken. Given that $\varphi = 0.74$, the number of 10^7-radius spheres that can be packed within the observable universe, η, is determined as follows:

$$\begin{aligned} \eta &= \frac{V_u \varphi}{10^7} \\ &= \frac{\left(3.56641 \cdot 10^{80}\right)(0.74)}{10^7} \\ &\approx 2.63914 \cdot 10^{73} \end{aligned}$$

Visually, you can observe that the average distance between spheres in an FCC packing is $2r = 2 \cdot 10^7$:

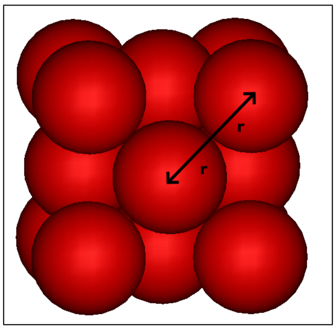

Figure 9: A very small set of spheres packed in an FCC arrangement.

In order to scan every atom, you would have to travel between every sphere, scan all the atoms within the sphere, and then move on to the next, repeating the process for every sphere. The time it takes to travel from one sphere to the next and scan all the atoms within that sphere, t_1, is determined as follows:

$$\begin{aligned} t_1 &= \frac{2 \cdot 10^7}{3 \cdot 10^8} + \left(2.4338 \cdot 10^6\right)\left(10^{-9}\right) \\ &\approx 0.06910\text{s} \end{aligned}$$

Now repeating this for every sphere we have in order to find t:

$$\begin{aligned} t &= t_1 \eta \\ &= (0.06910)\left(2.63914 \cdot 10^{73}\right) \\ &= 1.82366 \cdot 10^{72}\text{s} \end{aligned}$$

□

As you might reasonably expect, the travel time adds two magnitudes onto the answer. The difference between 10^{70} and 10^{72} may not seem like a lot, but it is extremely considerable – you may very well be able to run a 5K, but could you do a 500K?

This obviously represents the amount of time to have searched 100% of all atoms in the observable universe. It may very well be, with your amazing luck and genius, that you discover the atom I've hidden within a fraction of that time.

I have some catastrophic news, though. Due to the accelerating expansion of the universe (specifically, the expansion of the space between interstellar objects), for every megaparsec (Mpc, about $3.26 \cdot 10^6$ ly) of distance from an observer, a galaxy appears to be receding at roughly 67-73kms^{-1}. Let's say you've covered 5% of the search area, and so about 10^{71} seconds have passed. At that time scale, the universe's expansion will have massively accelerated due to dark energy, leading to exponential expansion.

In the far future, the **Hubble parameter** (which denotes the rate at which objects in the universe are moving away from each other) will asymptotically approach 57km s^{-1} Mpc^{-1}, decreasing from its current value of about 70km s^{-1} Mpc^{-1} [13]:[6]

The relationship between the scale factor of the universe $a(t)$, the Hubble parameter H, and cosmic time t, is represented as such:

$$a(t) \propto e^{Ht}$$

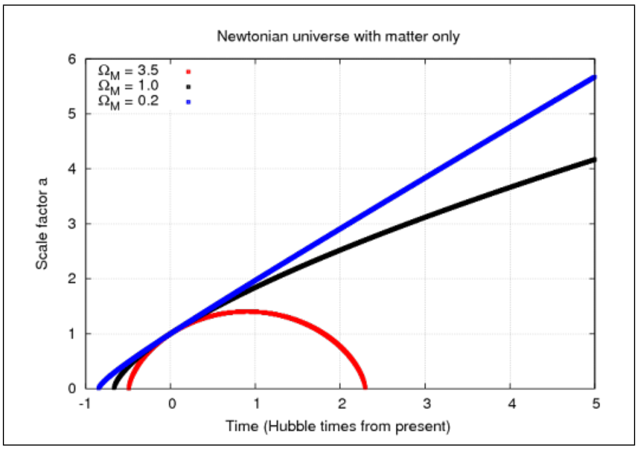

Figure 10: A graph showing how the scale factor of the universe changes with time. The different lines represent the expansion rates with different values for the matter density of the universe, which is a current topic of debate.

Assuming that you will have stayed in roughly the same position relative to a galaxy on the opposite end of the universe, which is nearly entirely true, the factor by which the galaxy will have recessed by, $D(t)$ is approximately this number:[7]

$$\begin{aligned} D(t) &\approx e^{\left(1.8 \cdot 10^{-18}\right)\left(10^{71}\right)} \\ &= e^{1.8 \cdot 10^{53}} \end{aligned}$$

This number is so incalculably large that I've resorted to telling you to go to Appendix 4.3 to demonstrate how large this is. An attempt to try to calculate how long it would actually take considering this would be fruitless; I may as well say infinity.

I have just one question, though. Did you think about scanning yourself?

[6]Yes – despite the fact that the Hubble parameter is decreasing, the rate of expansion of the universe is still accelerating. This is because while the expansion per distance (the Hubble parameter) is slowing down, the distances themselves are still increasing. Since expansion depends on both the rate per unit and the total distance, and the increasing distances outweigh the drop in the Hubble parameter, the rate of expansion is still accelerating.

[7]I am ignoring other numbers in this calculation as they're massively outweighed.

2.3 ...while writing this essay?

The following bite-sized estimation problems are all based on some aspect of the writing of this essay so far. Currently, there are 4786 words and 33794 characters in the LaTeX code, which is obviously due to increase, but I can't accurately predict how many I'll finish off with, and recalculating after I'm done isn't as fun as continuing on to the problems I've planned for next.

Let's dig into estimating now, and apologise to future-me for losing my sanity later.

2.3.1 How much weight have I lost...

> QUESTION: Through the physical activities I have conducted creating this essay, how much weight have I lost?
>
> RESPONSE: Counterfactually, the answer is negative. Realistically, it's just the weight of my expectations.

Since August of 2024, I've managed to lose 23.8kg, bringing my weight down from 101.6 to 77.8kg. The improvements to my health, mental and physical, have been astronomical. So, naturally, the best way to supplement that is to spend my time holed up in my room, endlessly typing away, trading kilos for keystrokes.

A 185-pound (83.9kg) person burns 61 calories per 30 minutes spent doing "computer work" [14]. I'm not quite at that weight, so I'll gratuitously assume I burn 56 calories per half-hour.

An estimation of the time I've spent on this essay can be obtained by looking at its history. I discovered this competition fairly late, and so the majority of it is a cram, but the time I've spent on it so far in minutes, t, is as follows:[8]

$$\begin{aligned} t &= 388 + 161 + 34 + 8 + 339 + 51 + 182 + 154 + 102 \\ &= 1419 \text{ minutes} \end{aligned}$$

The following calculation to determine c, the number of calories burned, is quite straight-forward:

$$\begin{aligned} c &= \frac{1419}{30} \cdot 56 \\ &= 2648.8\,\text{kcals} \end{aligned}$$

This is a genuinely surprising result! It turns out I have burned far more calories sitting at my desk and

(a) Me, before the half-marathon

(b) Me, after the half-marathon

[8] I really have enjoyed spending all this time writing this.

typing, compared to when I ran the Hampton Court Half Marathon a few weeks ago (where I burned 1728 kcals).

In order to lose 1kg of fat, you must burn 7700 more calories than you consume. Ignoring the fact that this considers **total daily energy expenditure** – the number of calories consumed subtracted by the number burned – the mass of fat lost is calculated as such:

$$\begin{aligned} m &= \frac{2648.8}{7700} \\ &= 0.344\,\text{kg} \end{aligned} \qquad \square$$

To my continued fat loss prospects, though, it's such a pity that Nando's exists.

2.3.2 How many seconds would elapse for a monkey...

> QUESTION: If a monkey were to type characters at random, how long would it take for it to eventually type this essay so far?
>
> RESPONSE: You'd better take this over trying to find my atom.

The infinite monkey theorem states that a monkey hitting keys independently, and at random, on a keyboard for an infinite amount of time will almost surely type any given text, including the complete works of William Shakespeare.

The "monkey" isn't actually real. It stands in for an abstract device that produces an endless random sequence of letters and symbols, with each proceeding character input independently to the previous, and with all keys having an equal probability of being pressed. So, assuming spherical monkeys in a vacuum...

To level the playing field, I will allow my monkey, who I'll call Bob, to use my exact setup. There are 110 keys on my keyboard, and yes, I counted them all. I've also permitted it to not have to bother with LaTeX file management by doing it for them, so all they will have to worry about is typing the essay using a keyboard.

I'll assume that Bob has a consistent typing speed of 600 characters per minute (CPM), which is highly above average. As mentioned above, Bob has to type every one of the 33794 characters on my essay accurately, and in the correct order. The probability of Bob typing in a single button and getting it correct, c_1, and the probability of typing the whole essay successfully, c_{33794}, are determined as such:

$$\begin{aligned} c_1 &= \frac{1}{110} \\ \Rightarrow c_{33794} &= \left(\frac{1}{110}\right)^{33794} \\ &\approx 10^{-68986.82} \end{aligned}$$

I think we're being unnecessarily harsh on Bob though. I've certainly made typos throughout this essay and had to go back and correct them – why can't he?

To give Bob a better chance, we could model this scenario **binomially**. This is because each key press is independent of the other, each key press is its own Bernoulli trial (it either matches my essay, succeeding, or doesn't, where it fails), and the probability of matching or not is constant throughout. While you may be right to question whether there are a fixed number of trials, we can model this slightly differently by considering 33794 trials (each trial being an individual key press) with a $\frac{1}{110}$ chance of success. Calculating $P(33793 \text{ successes})$ essentially gives us the probability that, for any given attempt, it will have one typo, and this can be done with any number of typos n_t where you calculate $P(33794 - n_t$ successes).

Let the random variable X denote the number of correctly typed characters in one attempt to generate the essay:

$$X \sim B\left(33794, \frac{1}{110}\right)$$

We can then begin to evaluate the probabilities of a certain number of typos in a given attempt using the binomial formula:

$$
\begin{aligned}
P(X=x) &= \binom{n}{x} p^x (1-p)^{n-x} \\
\Rightarrow P(X=33793) &= \binom{33794}{33793} \left(\frac{1}{110}\right)^{33793} \left(1-\frac{1}{110}\right)^{33794-33793} \\
&= 33794 \left(\frac{1}{110}\right)^{33793} \left(\frac{109}{110}\right) \\
&= \frac{3683546}{110} \left(\frac{1}{110}\right)^{33793} \\
&= \left(3683546 \cdot 110^{-1}\right) \left(110^{-33793}\right) \\
&\approx 5.51904 \cdot 10^{-68981} \\
\Rightarrow P(X=33792) &= \binom{33794}{33792} \left(\frac{1}{110}\right)^{33792} \left(1-\frac{1}{110}\right)^{33794-33792} \\
&= 571000321 \left(\frac{1}{110}\right)^{33792} \left(\frac{109}{110}\right)^2 \\
&= 560665687.1 \left(\frac{1}{110}\right)^{33792} \\
&\approx 1.01645 \cdot 10^{-68974}
\end{aligned}
$$

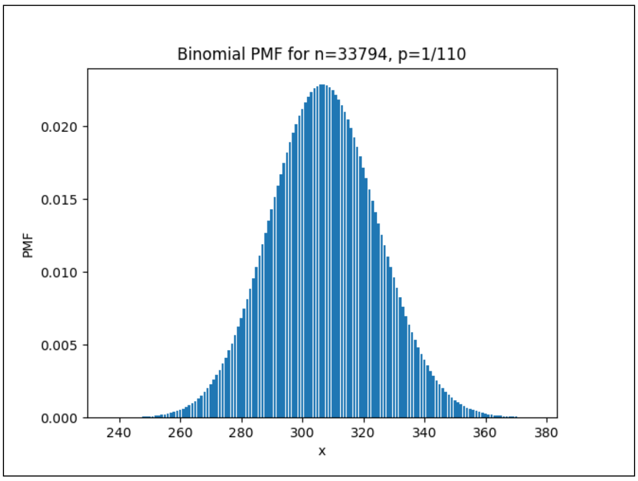

Figure 12: The binomial probability distribution for $X \sim B\left(33794, \frac{1}{110}\right)$. Note how the average number of correct characters is roughly 307, with the graph flattening out to practically 0 as it moves towards 33794. This shows just how immensely unlikely a successful attempt is.

You can fairly easily see how you could decrease the number of correct inputs needed, which increases the probability by a few orders of magnitude (yet it still remains astronomically small). Let's say we'll allow the monkey to make mistakes on 1% of the essay, meaning that Bob can make a maximum of 337 mistakes. We would like to determine $P(X \geq 33794 - 337) = P(X \geq 33457)$, which can be done as such:

$$
\begin{aligned}
P(X \geq 33457) &= \sum_{k=33457}^{33794} P(X=k) \\
&= \sum_{k=33457}^{33794} \binom{33794}{k} \left(\frac{1}{110}\right)^{33794} \left(1-\frac{1}{110}\right)^{33794-k} \\
&\approx 1.64178 \cdot 10^{-68170}
\end{aligned}
$$

This answer is hundreds of orders of magnitude larger, yet still unfathomably small. These, unfortunately, are the best odds Bob's got. The number of trials t needed to have a cumulative 100% probability of typing this essay so far is obtained by a simple calculation:

$$\begin{aligned} t &= \frac{1}{1.64178 \cdot 10^{-68170}} \\ &\approx 6.09095 \cdot 10^{68169} \end{aligned}$$

We'd now stand to benefit most from figuring out the **average duration** of an attempt. To start off with, we'd need to find the expected number of characters typed until the 337th typo occurs. This is captured perfectly with a **negative binomial distribution**. Where $\mathbb{E}[X]$ denotes the expected number of characters typed until $n_t > 337$ and p represents the probability of a typo, if we know the probability of a typo is $1-p$, on average it will take $\frac{1}{1-p}$ trials to get one failure. For r failures:

$$\begin{aligned} \mathbb{E}[X] &= r \cdot \frac{1}{1-p} \\ &= \frac{r}{1-p} \\ &= \frac{337}{\left(\frac{109}{110}\right)} \\ &\approx 340.09174 \end{aligned}$$

For every attempt, Bob will type ~ 340 characters before failure occurs. Hilariously, this does mean that for the average attempt, Bob only types 3 correct characters. Poor you...

The following calculation to find out the time needed to complete one trial, T_1 and then to complete all of them, T_t, is fairly straight-forward:

$$\begin{aligned} T_1 &= \frac{340.09174}{600} \\ &= 0.56682\,\text{minutes} \\ &= 34.00917\,\text{s} \\ \Rightarrow T_t &= T_1 t \\ &= 34.00917 \cdot \left(6.09095 \cdot 10^{68169}\right) \\ &= 2.07148 \cdot 10^{68171}\,\text{s} \end{aligned}$$

□

This number is so incomprehensibly large that not even WolframAlpha wants to evaluate it!

That being said... between this and the atom, I'd still pick this.

2.3.3 What impact have I had on global warming...

Global warming poses one of the most serious long-term civilisational risks to humanity in recorded history. Since record-keeping began in 1880, global surface temperatures have increased by 1.28°C, with the last decade being the warmest on record [15]. Although writing this essay has been a joy, I would be lying if I said my computer's power consumption wasn't solely responsible for this.

Figure 13: My power supply, the Corsair RM850.

My computer does not draw the full capacity of my 850W power supply throughout its running, however.

That being said, considering that I have a screen that runs at 240 Hz as well as a power-hungry CPU and GPU, I'll (rather unsubstantiatedly) assume that my computer constantly draws 250W when browsing – or in this case, just typing the essay up. Using the figure for the time obtained from 2.3.1, we can figure out my total energy use, E:

$$\begin{aligned} t &= 1419\,\text{minutes} \\ &= 85140\,\text{s} \\ \Rightarrow E &= 250t \\ &= 21285000\,\text{J} \end{aligned}$$

The UK's fuel mix is comprised of 6.3% coal and 35.0% natural gas [16], with the remainder being nuclear or renewables (which have negligible greenhouse gas emissions). This means that the amount of the energy used to write my essay that came from coal, E_c, and that came from natural gas, E_g, are:

$$\begin{aligned} E_c &= 21285000 \cdot 0.063 \\ &= 1340955\,\text{J} \\ E_g &= 21285000 \cdot 0.35 \\ &= 7449750\,\text{J} \end{aligned}$$

A commonly used type of coal is bituminous coal, which produces $\sim 24\,000\,000 \text{Jkg}^{-1}$ when burned. This is about 70-80% carbon by weight, so we will use 75% for this estimate. This means that 1kg of coal contains 0.75kg of carbon.

We must now, unfortunately, deep dive into your PTSD from GCSE Chemistry. Feel free to gloss your eyes over the next few lines!

The combustion of carbon is expressed as this chemical reaction:

$$C + O_2 \rightarrow CO_2$$

The A_r (relative atomic mass) of carbon is $12\,\text{gmol}^{-1}$, and the A_r of $CO_2 = 44\,\text{gmol}^{-1}$. As the units need not be converted, 1 kg of carbon produces $\frac{44}{12} \cdot 1 = 3.\dot{6}\,\text{kg}$ of CO_2. Therefore, 1 kg of coal would produce $0.75 \cdot 3.67 = 2.75\,\text{kg}$ of CO_2. A quick calculation with the energy output of 1 kg of coal shows that the CO_2 per joule $\approx 1.15 \cdot 10^{-7}\text{kgJ}^{-1}$. This leads to a carbon dioxide emission from coal of $\left(1.15 \cdot 10^{-7}\right)(1340955) = 0.15421\,\text{kg}$.

We'll use methane as a model for natural gas, with a combustion energy of $\sim 50\,000\,000\,\text{kg}^{-1}$. The combustion of methane is expressed in this chemical reaction:

$$CH_4 + 2O_2 \rightarrow CO_2 + 2H_2O$$

One mole of CH_4 (16 g) produces 1 mole of CO_2 (44 g). So $\frac{44}{16} = 2.75\,\text{g}$ of CO_2 per gram of methane, and so 1 kg of methane produces 2.75 kg of CO_2. This is the same as for coal, but when you calculate it considering the energy output of 1kg of methane, you get $5.5 \cdot 10^{-8}$ – half as much CO_2 release as coal. Anyway, this leads to a carbon dioxide emission from methane of $\left(5.5 \cdot 10^{-8}\right)(7449750) = 0.40974\,\text{kg}$.

The total CO_2 emissions from both energy sources, C, is simply calculated by addition:

$$\begin{aligned} C &= 0.15421 + 0.40974 \\ &= 0.56395\,\text{kg} \end{aligned}$$

Phew! If your eyes started rolling, you can place them back over the text now.

Over all of recorded history, humans have emitted $1.21122 \cdot 10^{16}$ kg of carbon dioxide into the atmosphere (which I obtained by summing values on the spreadsheet here [17]). This means that my percentage contribution to global warming, G, is as follows:

$$\begin{aligned} G &= \frac{0.56395}{1.21122 \cdot 10^{16}} \cdot 100 \\ &\approx 4.65602 \cdot 10^{-15} \end{aligned} \qquad \square$$

This number, expressed in long form, is 0.00000000000000465602%.

Phew, the Earth gets to live another day!

2.4 Simian Suffering for an S-Rank

> Question: If I give the monkey Bob my exact setup and open the rhythm game *osu!*, how likely is it that for a given song, he will obtain an S-rank?
>
> Response: Well, let's just say that his accuracy is absolutely bananas.

osu! is a rhythm game about clicking circles that appear on the screen to the beat of music playing in the background. The objective of the game is to be as accurate as possible in terms of aim (ensuring that your cursor is within the area of the hit circle) and timing (ensuring that you tap the hit circle at just the correct time), aided by "approach circles" that appear and close in on hit circles, where you have to tap as soon as the approach circle makes contact with the hit circle.

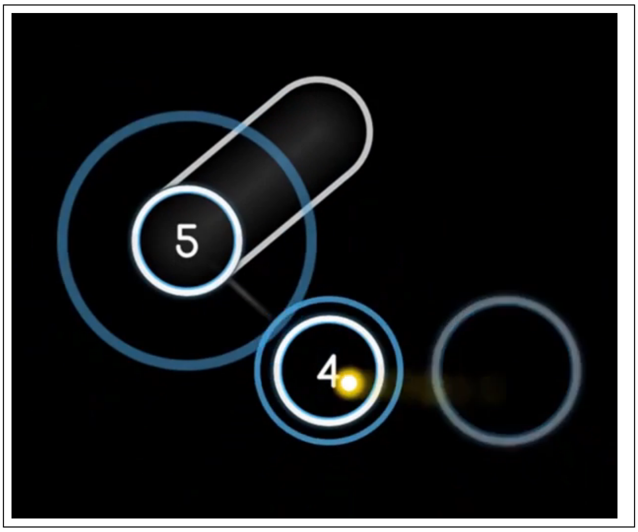

Figure 14: Two hit-objects in *osu!*. One of them is a circle, which only requires a timely tap. The other (towards the top left) is a slider, which requires a tap followed by holding down the input key and following a "slider ball" until you reach the end of the slider.

Tapping a circle at a slightly incorrect time leads to a small drop in accuracy, which is measured as a percentage with the maximum being 100%. Tapping with a more incorrect time leads to a greater drop in accuracy. If your cursor is not on the note when it's due to be hit, or you hit far too early/late, the note is "missed", which causes an even greater drop in accuracy and disqualifies you from being able to receive an S-rank.

The ranking system in *osu!* is as follows:[9]

- SS: 100% accuracy. All notes are hit with perfect accuracy with no misses.
- S: $\geq$ 95% accuracy. No misses are permitted.[10]
- A: $\geq$ 90% accuracy. Any amount of misses are permitted as long as it doesn't take the accuracy below the threshold.
- B: $\geq$ 80% accuracy.
- C: $\geq$ 70% accuracy.
- D: $<$ 70% accuracy.

We must make some assumptions about the beatmaps (the songs, which have hit circles and sliders mapped to their beats) before continuing. Most beatmaps in *osu!* are two minutes long, and so have roughly 250 sliders and 250 circles respectively.[11] We'll also let the circle size be 4, and the overall difficulty 8.5, which are both arbitrary designations to the amount of space circles take up on the screen,

[9] In *osu!lazer*.

[10] Slider breaks, where you let go in the middle of a slider, are permitted, which is why I've ignored them in this estimation.

[11] For simplicity's sake, we are assuming that this song does not have spinners.

and the size of the timing windows involved for a perfect, slightly imperfect hit, very imperfect hit, and a miss, respectively.

The playfield in *osu!* is 512 by 384 units wide. The published formula for the radius of a circle in units can be used to work out the area of it, A, followed by what fraction F of the playfield the circle covers:

$$\begin{aligned} r &= 54.4 - (4.48 \cdot \text{circle size}) \\ &= 54.4 - (4.48 \cdot 4) \\ &= 36.48 \\ A &= \pi r^2 \\ &= \pi (36.48)^2 \\ &\approx 4180.80134 \\ F &= \frac{4180.80134}{512 \cdot 384} \\ &\approx 0.02126 \end{aligned}$$

Let's go through the requirements for an S-rank one more time. Bob must successfully aim at every hit circle **and** obtain an accuracy above or equal to 95%. Well, for every instance of a note (a circle or slider) we know how likely it is that Bob will successfully aim at it, as it's just 0.02126. As for accuracy, that's a lot more complicated.

In *osu!*, the degrees to which you're accurate are determined by certain categories – a 300 is a perfect hit, a 100 is a slightly mistimed hit (either a bit too early or late) and a 50 is a very mistimed hit. Where n_{300} denotes the number of 300's hit, with the same applying for n_{100} and n_{50}, the accuracy formula is:

$$\begin{aligned} \text{Accuracy} &= \frac{300n_{300} + 100n_{100} + 50n_{50}}{300(n_{300} + n_{100} + n_{50} + n_X)} \\ \Rightarrow 0.95 &\leq \frac{300n_{300} + 100n_{100} + 50n_{50}}{300(500)} \\ 142500 &\leq 300n_{300} + 100n_{100} + 50n_{50} \\ 2850 &\leq 6n_{300} + 2n_{100} + n_{50} \end{aligned}$$

The task is to now find all values $n_{300}, n_{100}, n_{50} \in \mathbb{Z}^+$ that satisfy this system of equations:

$$\begin{cases} 2850 & \leq 6n_{300} + 2n_{100} + n_{50} \\ 500 & = n_{300} + n_{100} + n_{50} \end{cases}$$

This is obviously immensely impractical. While it is technically possible to model the probability of achieving an S-rank through multinomial distributions, doing so would not only require evaluating a hyper-dimensional inequality but also potentially losing my sanity in the process. Bob isn't worth it.

What we do know is, when input into WolframAlpha, the lowest-bound solution for n_{300} is:

$$462.5 < n_{300} \leq 500$$

We can use this value as a proxy for what we're trying to achieve; in this case, we'll simply declare that 463 300's qualify Bob for an S-rank, even though under small edge cases where he also hits many 50's, it would be slightly under 95%. It suffices for a Fermi estimation problem, at the very least, even if it isn't so mathematically rigorous. Furthermore, we'll know for a fact that our answer will be an overestimate.

So, in a given attempt, how likely is Bob to get above 463 perfect hits? First, let's work out our probability of hitting the note perfectly. The formula for the hit window (in milliseconds) can be used to work out the probability of a perfect hit, P_p, as such:

$$\begin{aligned} \text{For a 300: hit window} &= 80 - (6 \cdot \text{overall difficulty}) \\ &= 80 - (6 \cdot 8.5) \\ &= 29\,\text{ms} \\ &= 0.029\,\text{s} \end{aligned}$$

At any given time there is likely a note on the screen, and so any random hits would likely fall within at least the 50 range. The monkey is also probably randomly mashing the keyboard. As a result, we can assume that the range of hit opportunities lie within said range. Continuing on with this:

$$\begin{aligned} \text{For a 50: hit window} &= 200 - (10 \cdot \text{overall difficulty}) \\ &= 200 - (10 \cdot 8.5) \\ &= 115\,\text{ms} \\ &= 0.115\,\text{s} \end{aligned}$$

From there, the probability of hitting a 300 is very simple:

$$\begin{aligned} P_p &= \frac{0.029}{0.115} \\ &= 0.25217 \end{aligned}$$

Now, the probability of Bob aiming correctly and hitting a 300 is obtained, H_1 through easy conditional probability:

$$\begin{aligned} H_1 &= 0.25217 \cdot 0.02126 \\ &\approx 5.36122 \cdot 10^{-3} \end{aligned}$$

Let's head back to our binomial distribution – let the random variable Y denote the number of correctly timed hits in one attempt to play the song (assuming perfect aim):

$$Y \sim B\left(500,\ 5.36122 \cdot 10^{-3}\right)$$

We are trying to work out $P(Y \geq 463)$ Using the same technique as before:

$$\begin{aligned} P(Y \geq 463) &= \sum_{k\,=\,463}^{500} P(Y = k) \\ &= \sum_{k\,=\,463}^{500} \binom{500}{k} \left(\frac{1}{110}\right)^{500} \left(1 - \frac{1}{110}\right)^{500-k} \\ &\approx 2.105 \cdot 10^{-965} \end{aligned}$$

□

Hey, man! It's better than Bob writing my essay... maybe he has a future in something after all![12]

2.5 The Weight of What Was

QUESTION: How much does the average relationship weigh?

RESPONSE: When you lay awake at night, do I ever cross your mind? 'Cause you still cross mine...

A relationship is a living manifestation of its participants' combined beliefs and ideologies. Evolutionarily speaking, humans are integrally geared towards maximising their chances of finding and securing a mate. Your body does this through the targeted releasing of a neurochemical cocktail designed to spark attraction and closeness, bond you long enough to support offspring, and to make you irrationally attached to a fellow human being.

A cascade of dopamine, oxytocin, and vasopressin marries you to someone else's memory, deliberately designed to hijack your brain, even if the person in question no longer exists to you.

The average relationship lasts 17.2 months (1 year, 5 months) [18], and feelings of emotional distress largely subside after roughly 9 months after a breakup, although this varies massively.

It is hilariously impossible to estimate the precise volume to which "love chemicals" are released within our bodies, let alone along an entire relationship. But when did that ever stop us?

The key chemicals involved (for which their changes are at all discernible) are listed as follows:

- Dopamine: Normal plasma concentration of $\sim$ 10-50pg ml^{-1} (picograms per millilitre), which equates to a total mass of roughly $\sim$ 0.05-0.25μg in human blood.

[12] For further investigation: how likely is it for him to set a score comparable to the likes of top players?

- Oxytocin: Normal plasma concentration of $\sim$ 1-10pg ml^{-1}, which equates to a total mass of roughly $\sim$ 0.005-0.05μg in human blood.
- Norepinephrine: Normal plasma concentration of $\sim$ 100-400pg ml^{-1} (picograms per millilitre), which equates to a total mass of roughly $\sim$ 0.5-2μg in human blood.

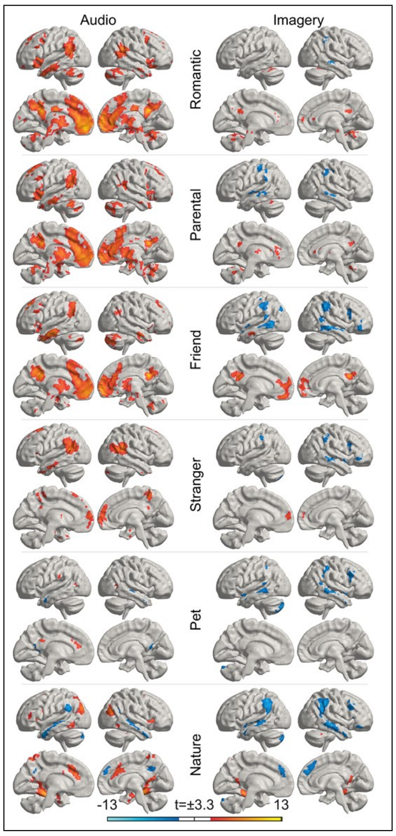

Figure 15: The brain regions activated (orange) and deactivated (blue) while listening to stories of love (left) and imagining (right) feelings of love for different objects as contrasted with neutral control stories.

During the course of a relationship, the concentrations of these chemicals within the bloodstream fluctuates wildly. We will assume that during a "love event", these chemicals increase threefold. This means that the total mass of the chemicals increases from $\sim$ 0.3025 to 22.6875μg.

A "love event" is probably one of the most vague things I have stated during this essay, and so many things here have been vague. Generally, a combination of distinct thoughts, direct interactions, emotional spikes and deliberate expressions put this in the ballpark of 20-30 events per day.

I've always thought of a breakup as a curve that asymptotically approaches 0, but never quite hits it. Funnily enough, the exponential decay curve I'm imagining matches the format for our Eiffel Tower taper!

$$f(t) = ae^{bt}$$

In this case, a represents the initial intensity (the emotional spike, $22.6875 \cdot 25 = 22.6875$, assuming the emotions start off at the same frequency and intensity as in the relationship, perhaps just negative feelings now), which likely matches the max threefold increase of chemicals, b, the decay constant, denotes how fast you're able to "move on", and t represents time, in days. Let's gratuitously assume that after 9 months (270 days), $f(t)$ reduces to 5% of its original value. As, at the start, $f(t) = 22.6875$:

$$\begin{aligned} f(270) = (22.6875 \cdot 0.05) &= 22.6875e^{270b} \\ 1.134375 &= 22.6875e^{270b} \\ 0.05 &= e^{270b} \\ \ln 0.05 &= 270b \\ \frac{\ln 0.05}{270} &= b \\ &\approx -0.0110953 \end{aligned}$$

Therefore, our breakup exponential decay equation is:

$$f(t) = 22.6875e^{-0.0110593t}$$

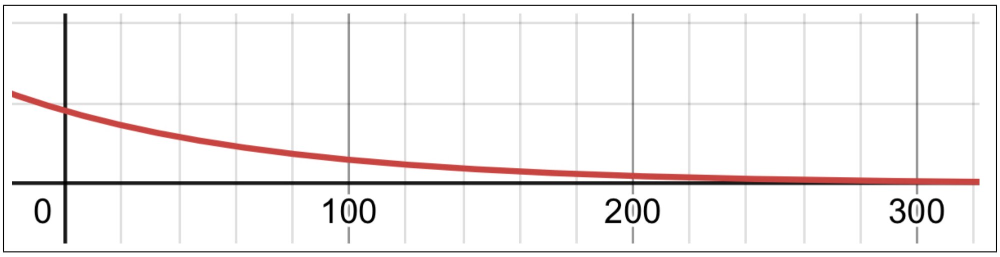

Figure 16: The graph $f(t) = 22.6875e^{-0.0110953t}$. If you somehow fit the niche of a heartbroken person and a maths nerd, you should soon observe tears falling from your eyes soon.

All we need to do is simply integrate that between $t = 0$ and 270! This will give C_b, the mass of chemicals released by the body during a breakup. We can use the exact same technique as in our Eiffel Tower estimate:

$$\begin{aligned} C_b &= \int_0^{270} f(t)\,\mathrm{d}t \\ &= \int_0^{270} 22.6875e^{-0.0110953t}\,\mathrm{d}t \\ &= 22.6875 \int_0^{270} e^{-0.0110953t}\,\mathrm{d}t \\ &= 22.6875 \left[\frac{1}{-0.0110953} e^{-0.0110953t}\right]_0^{270} \\ &= 22.6875 \left[\frac{1}{-0.0110953}\left(e^{-0.0110953[270]} - e^{-0.0110953[0]}\right)\right] \\ &= 22.6875 \left[\frac{1}{-0.0110953}\left(e^{-2.995731} - 1\right)\right] \\ &= \frac{22.6875}{-0.0110953}\left(e^{-2.995731} - 1\right) \\ &\approx \frac{22.6875}{-0.0110953} \cdot -0.95 \\ &\approx 1942.54549\mu\text{g} \end{aligned}$$

Calculating the mass of chemicals released over the duration of an average relationship, C_r and then total, C_t, is fairly simple:

$$\begin{aligned} C_r &= (22.6875 \cdot 25)(516) \\ &= 11706.75\mu\text{g} \\ \Rightarrow C_t &= C_r + C_b \\ &= 11706.75 + 1942.5459 \\ &\approx 13649.29549\mu\text{g} \\ &\approx 0.01365\text{g} \end{aligned}$$

□

The thing that struck out to me in particular was the difference in masses between the chemicals released during the course of a relationship, and the chemicals released during a breakup. This could very well be a reflection of my poor maths and rampant idealised assumptions, but I think it signifies something a lot more profound.

The total mass of all chemicals released in the body during a relationship – a proxy for the all the illuminating waves of feelings and emotions that run through you – was calculated to be roughly $12\,000\mu$g. Contrastingly, and almost symbolically, the mass of the chemicals released following a breakup was 6 times lower, at roughly $2\,000\mu$g. The variation behind these numbers is absolutely immense, and I'm very confident that there's a strong possibility that they're completely incorrect (and such are the beauty of Fermi estimations), but the fact that the sum total of all of those transformational positive experiences and feelings is half an order of magnitude higher than the ruinous pain that follows from a breakup is so, immensely telling.

It is melancholy to know that people you come to love will dip into and out of your life, but from them you will retain lovely pieces and memories. You will be forever shaped by the powerful experiences you have with them, and by far that occurs for the better, even despite the immense emotional turmoil.

Lest you otherwise turn into me, someone crying onto their maths-laden napkin at four in the morning.

3 Conclusion

I've either made Enrico Fermi very proud, or he's currently rolling in his grave.

I hope you've enjoyed reading these very unusual and perhaps funny problems as much as I have writing them! This essay is quite long but isn't necessarily intended to be read as a whole – each estimation problem is its own little piece that can be read completely on its own. I've drawn each question from something I've found interest in and I really loved coming up with them and conjuring up methods to solve the unsolvable.

Fermi problems may seem like a bit of a party trick at this point. After all, if you've read everything up to here, you'll have been watching me explain how to use a binomial distribution to model the time it takes for a monkey to type this essay, or how you can use calculus to figure out the volume of the Eiffel Tower so that I can shove tennis balls inside of it.

However, they do really have uses in real-world applications:

- They can make you more confident on the decisions you make: instead of relying on vibe-based guesses, a quick back-of-the-hand calculation could give you a very reasonable answer when you don't have enough information.
- They're used almost like sanity checks in STEM: results can quickly be checked for if they're reasonable – such as when I compared my Eiffel Tower square-based pyramid volume to the solid of revolution's volume.
- Organisations, such as effective charities, use these rough calculations to determine how cost-effective a particular intervention is, which is especially useful in global-scale issues where exact numbers are impossible to determine. This allows them to better decide where funding should be spent, which can quite literally save more lives.
- During emergency situations, Fermi logic gives a ballpark estimate fast before full data is available. For example, during COVID-19, early models estimated infection spread by using approximate reproduction numbers and assumptions about transmission, which helped governments make early policy decisions.

So perhaps, unlike Sir Isaac Newton, I haven't solved the mysteries of the universe – but at least I've managed to produce a marginally entertaining essay about monkeys, tennis balls, France's most iconic landmark, and love.

And I have, at least, proven that you can know a lot without knowing anything.

4 Appendix

4.1 But what is sphere packing?

Mathematicians love weird, niche, and oddly specific things. In this case, it is the desire to pack as many balls in a container as possible, without squashing or deforming them, in perfect Euclidean harmony. Aside from the fact that many seemingly abstract problems in maths, such as this one, have had real-world, impactful effects, these problems are also genuinely fun to work on.

The least efficient sphere packing method is obtained by increasing the spacing between spheres arbitrarily, i.e. the worst thing you could do is place one sphere in a container, and then give up. That's not a very fun answer though!

The (most realistic)[13] worst packing is typically **random loose packing**, with an average efficiency (denoted by φ) of ~ 0.55. This typically occurs when you dump many spherical objects into a container, which (especially when there are high amounts of friction between them) causes jamming and many empty gaps between spheres. In this arrangement, 55% of the container's volume is occupied by the spheres, whereas the remainder is just empty space; voids between the balls. Products containing loosely packed items are often labelled with the message, "contents may settle during shipping". This is because during shipping, the bumping of the container leads to the packing density of the items inside to increase, making it appear as though the container isn't actually full as the items inside now take up less space. The message is therefore added as an assurance to the consumer that it's full, just on a mass basis.

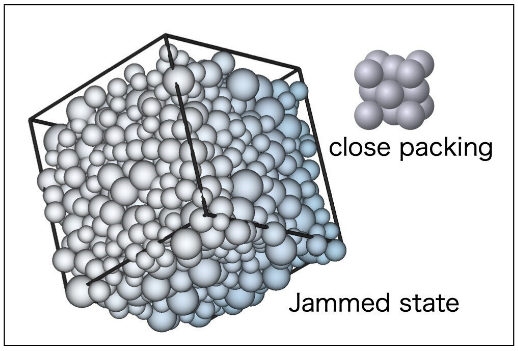

Figure 17: A visualisation of what random packing looks like compared to close packing. [19]

The shaking or bumping of a container which has items randomly and loosely packed into it leads to the next packing arrangement in the hierarchy – **random close packing** (RCP). In this arrangement, $\varphi \approx 0.64$, and the higher packing efficiency can be attributed to how spheres, when shaken, can move past others (overcoming frictional forces) to move to a more energetically stable state. Specifically, they slide into lower positions and tighter spots, reducing gaps and increasing the number of contacts per sphere.

In order to obtain higher packing efficiencies, randomness no longer suffices. The densest packing arrangements, **hexagonal close packing** and **face-centred cubic** packing, have $\varphi \approx 0.7408$, and is obtained by combining layers of compactly arranged spheres in a specific way.

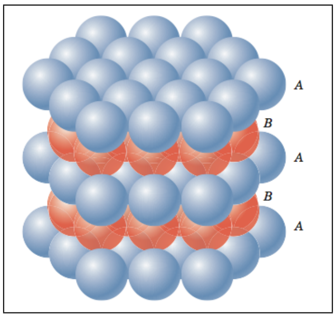

Figure 18: A visualisation of hexagonal close packing. [20]

[13] There exist even worse packings that stay strictly jammed, i.e. all spheres are constrained by their neighbours to stay in one location, with the worst having a value of $\varphi \approx 0.49365$.

4.2 Watching Entropy Win

This explanation is heavily inspired by Scott Czepiel's writings on how big 52! is, which I have referenced here [21].

Imagine setting a timer for $8.7 \cdot 10^{70}$ seconds and that, beyond all logic, you are standing on the equator of the Sun. To help pass the time, you decide to play a game.

You start by waiting one million years, before then playing the UK National Lottery. If you lose, you wait another million before playing another one. When you win, you can take one step forwards, wait another million years, and then repeat the same steps. Once you have finally walked around the entire circumference of the Sun, remove one gram from its mass, and then do the same thing again, waiting a million years before playing another lottery ticket, and if you win you take one step forwards...

Once you have removed all of the Sun's mass, draw 5 playing cards. If they are anything other than a royal flush, replace the Sun's mass, and repeat all the steps above. When you finally get a royal flush, check your timer. This is (almost certainly) how much time you have left:

$$7.9 \cdot 10^{70}\text{s}$$

You would have to repeat this entire process about 11 times to finally drain out the timer. And that's for the absolute lower bound estimate... yet before you could even win your first lottery ticket, the Sun would likely already have turned into a red giant and dissipated, ruining the game. In fact, every star in the observable universe will have died by the time you have walked around the Sun for the first time, and smaller black holes will have been extinguished due to Hawking radiation by the time you're finally done.

4.3 Watching Entropy Lose

There is no activity in the world that I could give you to pass the amount of time it takes to cover the distance $e^{1.8 \cdot 10^{53}}$. You could repeat the process detailed in Appendix 4.1 a googol times, and still not even come close to passing a billionth of the time it takes to cover that distance. Travelling at the speed of light through it would help, but with respect to the sheer magnitude of this number, you may as well be travelling at 1ms^{-1}. It makes no difference. What I will do is list out the events that proceed in the far future of the universe, until $e^{1.8 \cdot 10^{53}}$ seconds[14] elapse:

- 180 million years: Due to the gradual slowing of Earth's rotation, a day will be one hour longer than it is today.
- 2.8 billion years: Barring all human advancement, all remaining life has certainly gone extinct due to the Sun's rising temperature.
- 5.4 billion years: The Sun is no longer a main sequence star, and has become a red giant.
- 8 billion years: The Sun becomes a white dwarf with half its original mass.
- 10^{14} years: All star formation ends in galaxies, marking the transition to the Degenerate Era. At this point, the universe has expanded by a factor of 10^{2554}. A short while after this point, all stars in the universe will have exhausted their fuel, and the only stellar-mass objects remaining will be remnants (white dwarfs, neutron stars, black holes, brown dwarfs).
- 10^{38} years: If proton decay can occur, all remaining planets and stellar-mass objects will have done so.
- 10^{110} years: All black holes, including ultramassive black holes, will have decayed due to Hawking radiation.
- 10^{1500} years: Assuming that protons do not decay, all matter will have fused together to form iron-56 or decayed from a higher mass element into iron-56, forming iron stars.
- $10^{10^{26}}$ years: Around this time, all iron stars have collapsed via quantum tunnelling into black holes (assuming proton decay does not occur). The subsequent evaporation of these black holes into subatomic particles is instantaneous on this time scale.

[14] This could be expressed in nanoseconds or star lifespans – this number is so big that the timescales have become completely irrelevant.

- $10^{10^{50}}$ years: The estimated time for a Boltzmann brain to appear in a vacuum via a spontaneous entropy decrease.

Not even WolframAlpha could evaluate this, but a tedious division process shows that, on average, $10^{7.81 \cdot 10^{52}}$ Boltzmann brains could form before you covered that distance, assuming you were walking at 1ms^{-1}. Travelling at the speed of light would still lead to an incomprehensible number of them forming. By the time you've covered even a fraction of the distance, space would stretch so violently that even light trying to cross it would be like a tortoise on a treadmill made of molasses... on a glacier... during a time freeze.

I would imagine that the original mission – recovering the one atom I've hidden – would take so long due to the universe's immense expansion, that quantum tunnelling could generate new inflationary events, leading to new Big Bangs giving birth to new universes. You could watch entropy win, and then lose, before you'd recover my atom.

I think it's time to get a better ray gun.

References

[1] Isaac Newton. *Philosophiæ Naturalis Principia Mathematica*. Translated as "Mathematical Principles of Natural Philosophy". Royal Society, 1687.

[2] Richard Rhodes. *The making of the atomic bomb*. New York: Simon & Schuster, 1986.

[3] Benh Lieu Song. *Eiffel Tower, Paris France*. Wikimedia Commons. 2018. URL: https://commons.wikimedia.org/w/index.php?curid=48220385/.

[4] Miller S. "Modern tennis rackets, balls, and surfaces". In: *Br J Sports Med* 40.5 (2006). Accessed: 2025-04-02, pp. 401–405. DOI: 10.1136/bjsm.2005.023283. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC2577483/.

[5] Société d'Exploitation de la Tour Eiffel. *The Eiffel Tower: Key Figures*. https://www.toureiffel.paris/en/the-monument/key-figures. Accessed: 2025-04-02.

[6] Paris 16. *Cropped version of sizes of Eiffel Tower*. Wikimedia Commons. 2009. URL: https://commons.wikimedia.org/wiki/Category:Architectural_drawings_of_the_Eiffel_Tower#/media/File:Eiffel_sizes.png.

[7] Ping Wang et al. "Jamming II: Edwards' statistical mechanics of random packings of hard spheres". In: *Physica A: Statistical Mechanics and its Applications* 390.3 (2011), pp. 427–455.

[8] Ed Pegg Jr. *Sphere Packing*. Wolfram Demonstrations Project. 2011. URL: https://demonstrations.wolfram.com/SpherePacking/.

[9] John Terning Itzhak Bars. *Extra Dimensions in Space and Time*. Springer Science & Business Media, 2009.

[10] Dr. Volker Springel. *Simulating the observable Universe with MillenniumTNG*. Gauss Centre for Supercomputing. 2022. URL: https://www.gauss-centre.eu/results/astrophysics/simulating-the-observable-universe-with-millenniumtng.

[11] P. C. W Davies. *Cosmic jackpot: why our universe is just right for life*. Boston: Houghton Mifflin, 2007.

[12] HowStuffWorks. *If you were to move all of the matter in the universe into one corner, how much space would it take up?* Blog Post. Accessed: 2025-04-06. 2023. URL: https://www.gauss-centre.eu/results/astrophysics/simulating-the-observable-universe-with-millenniumtng.

[13] Various Editors. *Hubble's law*. Wikipedia Article. Accessed: 2025-04-06. 2025. URL: https://en.wikipedia.org/wiki/Hubble%27s_law.

[14] Unknown. *Calories burned in 30 minutes for people of three different weights*. Harvard Health Publishing. Accessed: 2025-04-06. 2004. URL: https://mymission.lamission.edu/userdata/cooperck/docs/Calories%20burned%20from%20Harvard.pdf.

[15] NASA. *Global Temperature*. Article. Accessed: 2025-04-07. 2024. URL: https://climate.nasa.gov/vital-signs/global-temperature/?intent=121#:~:text=Overall%2C%20Earth%20was%20about%202.65,change%20in%20global%20surface%20temperatures..

[16] Department for Energy Security and Net Zero. *Fuel mix disclosure data table*. gov.uk Article. Accessed: 2025-04-07. 2024. URL: https://www.gov.uk/government/publications/fuel-mix-disclosure-data-table/fuel-mix-disclosure-data-table.

[17] Max Roser Hannah Ritchie. *CO2 emissions*. Article. Accessed: 2025-04-07. 2024. URL: https://ourworldindata.org/co2-emissions.

[18] Benson NF Freeman H Simons J. "Romantic Duration, Relationship Quality, and Attachment Insecurity among Dating Couples". In: *Int J Environ Res Public Health* 3.20 (2003). Accessed: 2025-04-07, pp. 401–405. DOI: 10.3390/ijerph20010856. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC9820285/.

[19] Osaka University. *Random close packing or jamming of spheres in a container*. Blog Post. Accessed: 2025-04-07. 2004. URL: https://phys.org/news/2021-05-random-spheres.html.

[20] Callister & Rethwisch 5e. *Hexagonal Close Packed Crystal Structure (HCP)*. Educational Article. Accessed: 2025-04-07. URL: https://www.e-education.psu.edu/matse81/node/2134.

[21] Scott Czepiel. *52 Factorial*. Blog Post. Accessed: 2025-04-06. 2016. URL: https://czep.net/weblog/52cards.html.