Hooke in Isolation

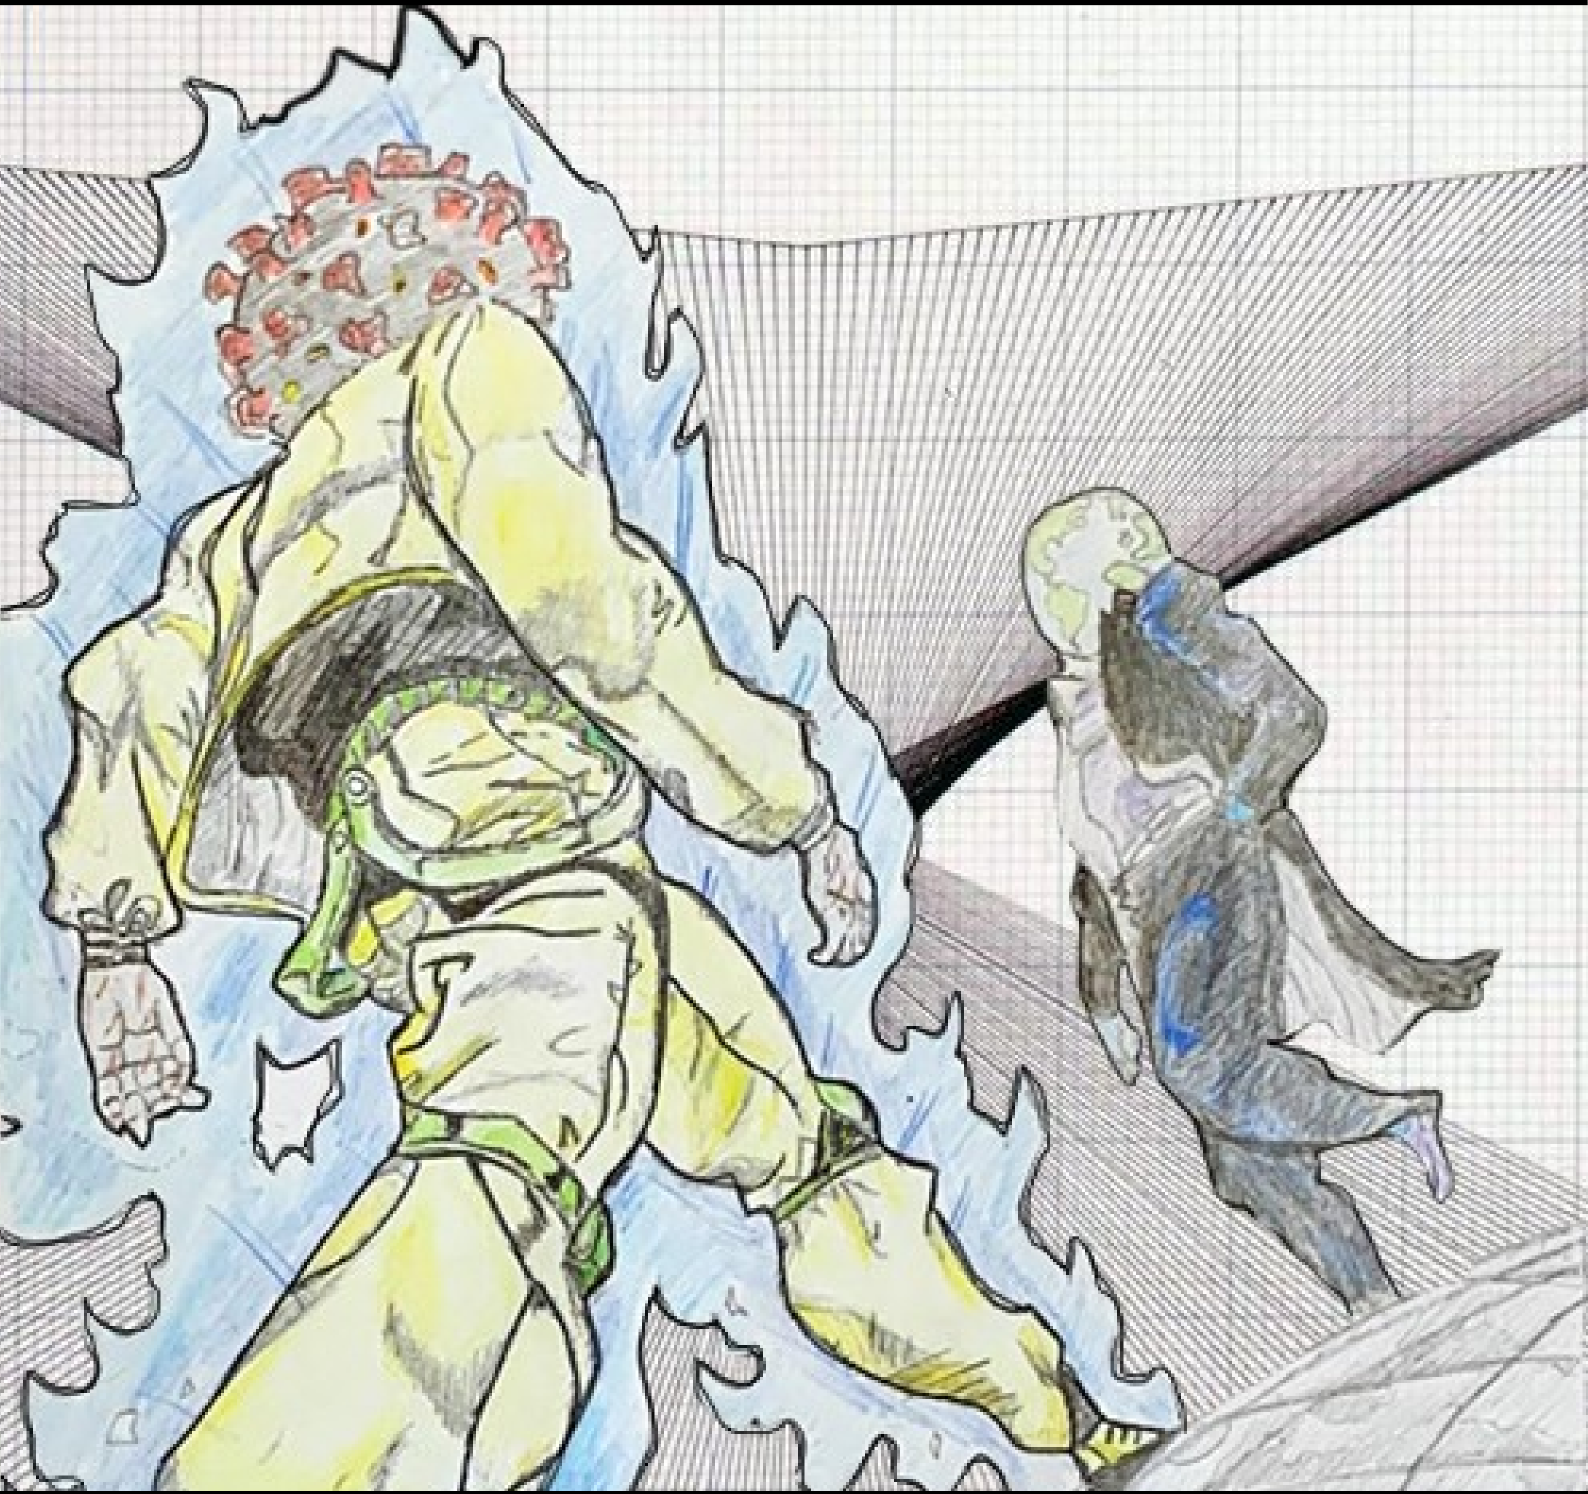

Volume 1

From the editors

Dear Reader,

Welcome to the inaugural edition of an ephemeral publication, which we sincerely hope will be short-lived…

In the situation we find ourselves in, we – a like-minded group of Upper Shell students – wanted to continue in the tradition of Hooke Magazine and to provide a sample of our scientific musings. This is not meant to replace the printed edition but is deigned as a *lite eZine* of sorts – somewhere between pamphlet and comic – to quench scientific thirsts that prevail despite our physical isolation.

The intention is not to mimic *Nature* or *JACS*, but to interest, engage and occupy thoughts. There is even a scientific cryptic crossword at the back with a £5 prize!

We will aim to produce yet more editions for as long as we're in isolation (there's not much else going on), but for now enjoy what comes next…

Best wishes,

Hooke in Isolation

Contents

The origin of the moon

George Weston

For centuries, astronomers have looked up to the sky to observe the celestial bodies. Ever since the advent of the telescope in the 1600s, astronomers have mapped out our solar system as they continue to discover its mechanics. From Kepler to Newton, scientists, with their grand ideas, brought revolutionary advances in the field of astronomy. Meanwhile, other astronomers turned their attention to the most striking object in the night sky, Earth's natural satellite: The Moon.

The telescope's huge impact on the field of astronomy cannot be exaggerated. It allowed Galileo to observe first and foremost that the Sun was the centre of the solar system; however, Galileo's telescope also allowed him to observe the Moon in more detail. Astronomers started drawing maps of the Moon noting its salient features. People started to see that the Moon was not spherical and that its terrain was similar to the Earth's. The Moon consisted of craters, Maria, and highlands that all got appropriately named. Based on their observations, astronomers started to speculate about the origins of the Moon coming up with many theories. The difficulty though was proving these theories based up a limited amount of evidence. Until the Apollo space missions, the Moon's surface and internal structure were thought to be impossible to properly analyse. Covered below are four competing theories including those whose validity has been cast into doubt by the lunar missions.

The Giant Impact Hypothesis

A celestial body, Theia, about the size of Mars, collided with a young Earth approximately 4.5 billion years ago. The impact vaporised Theia and part of the Earth resulting in a disk of debris, called the protolunar disk. The protolunar disk orbited the Earth. The debris later condensed and came together under the influence of gravity to form the Moon.

The Fission Theory

The Earth was spinning so rapidly that part of it spun off and formed the Moon. That part of the Earth is now filled by The Pacific Ocean.

Condensation (Co-accretion) Theory

The Earth and the Moon formed at the same time out of material from the solar nebula.

The Capture Theory

The Earth and the Moon were formed at different places in the Solar System. The Moon later became captured by the Earth's gravity.

Which Theory Does the Evidence Support?

The most widely accepted theory for the Moon's formation is now The Giant Impact Hypothesis. However, that was not always the case as more traditional theories (the other three) were thought to be correct. The most compelling pieces of evidence for The Giant Impact Hypothesis are lunar rocks and experiments done on the Moon. The Apollo space missions were ground-breaking in this respect since they provided actual samples from the Moon and accurate data. By comparing these results with the predicted ones, scientists could finally validate certain theories and refute others.

References:

https://www.bl.uk/learning/timeline/item104335.html

https://www.nhm.ac.uk/discover/how-did-the-moon-form.html

https://www.newscientist.com/article/2236797-oxygen-in-lunar-rocks-suggests-the-moon-formed-in-huge-collision/

http://www.psrd.hawaii.edu/Dec98/reminisces.html

By experimenting on the lunar rock samples, it became apparent that there were inconsistencies between the actual results and the ones predicted by certain theories. The Capture Theory predicts that the composition of the Earth and the Moon would be vastly different. On the other hand, the Fission Theory predicts that the composition would be the same. The evidence, however, is inconsistent with what both theories predict. In fact, the composition of the lunar rocks does bear similarity to that of rocks from the Earth while not being exactly the same. Erik Cano and others at the University of New Mexico discovered that the relative abundance of oxygen isotopes in the rocks varied and depended on where the rocks were located. The difference between the lunar rocks and the rocks from Earth was most dramatic for rocks from the deepest parts of the Moon. This supports The Giant Impact Hypothesis since it predicts that the rocks will look like a mixture of the rocks from Earth and Theia. The same line of thinking furthermore calls the Co-Accretion Theory into question since one would expect from the theory that the Earth and the Moon would be more similar and share more characteristics.

Also supporting The Giant Impact Hypothesis is the Moon's lack of volatiles in comparison with the Earth, and its small iron core. Volatiles are compounds with low boiling points, so they evaporate easily. Water is among these volatile compounds and is found in short supply on the Moon compared to the Earth. These volatile compounds are thought to be in short supply because of a collision-like event which generated a lot of heat.

The Giant Impact Hypothesis Becomes Widely Accepted

For a long time, the other theories were considered promising answers while the giant impact hypothesis was not held in the same light. As explained in previous paragraphs, the lunar missions were really the turning point and paved the way for scientists to increasingly accept the Giant Impact Hypothesis. The Conference on the Origin of the Moon was the second turning point. The conference was held in Hawaii in 1984 and saw many well-respected scientists come together to talk about the origin of the Moon. The conference was revolutionary since the general consensus now shifted away from the other traditional theories and in favour of The Giant Impact Hypothesis. Scientists attended the next major conference in 1998 in Monterey with a fresh perspective.

The Current Situation

The Giant Impact Hypothesis is now accepted, by the majority of the scientific community, as the most promising solution. More evidence about the Moon's origins is still being collected and each bit is yet another piece of the puzzle. Some evidence supports the theory. Some does not: For example, scientists recently discovered that almost all of the Moon's surface is emitting. This would suggest the presence of volatile carbon on the Moon hardly in line with the predicted lack of volatile compounds. A lack of volatile carbon would seem to indicate therefore that temperatures were not as high in the Moon's past as earlier expected. This new evidence may throw the Giant Impact Hypothesis into question yet perhaps this just shows that more work needs to be done. It is important also not to forget that the other main theories have even bigger flaws. The Giant Impact Hypothesis still has a way to go but perhaps it only needs to be refined instead of scrapped completely. With no other viable theories out there, scientists continue to place their trust in The Giant Impact Hypothesis hoping that in the end it will reward them with a satisfying answer.

Neural networks

Adam Moec

the what, the how and the why

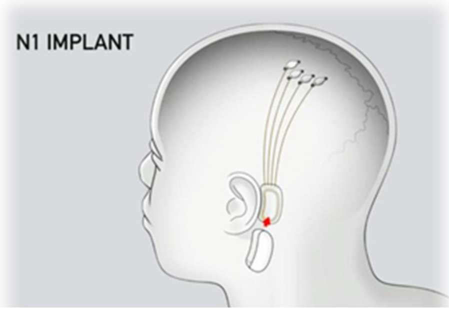

The N1 implant by Neuralink is set to soon fuse man and machine

So straight off the bat it's important to state that neural networks are a form of artificial intelligence and more specifically, a type of machine-learning algorithm. AI comes in many different shapes and sizes, all with the aim of allowing machines to "think", that is, make decisions given a set of circumstances. Probably most commonly talked about is AI's application in autonomous vehicles where the aim there is to drive from point A to point B whilst taking in many different inputs from cameras and other sensors in order to not hit anything or anyone in the process (and also respecting traffic laws). AI is set to change the way we live drastically in the next few years with companies such as Elon Musk's Neuralink eventually aiming to prevent the rapid overtaking of humans by AI through fusing man and machine and allowing a sort of symbiosis to flourish. AI is set to bring about the next seismic change to employment structures as its strengths over humans become ever more apparent, which is why I believe it is important to ride the wave and understand exactly how it works.

But where do neural networks come in; you might ask. Well they are in my opinion the most interesting of machine-learning algorithms. As their name suggests they are modelled off of neural pathways in animals. It turns out that we humans have our crucial strengths over machines in our incredible versatility. We have a long and well-recorded history of adapting and even thriving in new and initially inhospitable environments. We are able to change the way we behave based off of a change in surrounding and environment, something that is the holy grail of computing. Neural Networks allow different outputs to be reassessed as much as necessary, meaning that connections and patterns between an almost limitless amounts of inputs and outputs can be discovered and rediscovered where humans would fail. The particular interest in Neural Networks over more standard machine learning algorithms is their ability to filter out 'noisy' data and constantly reassess the importance of certain inputs over others.

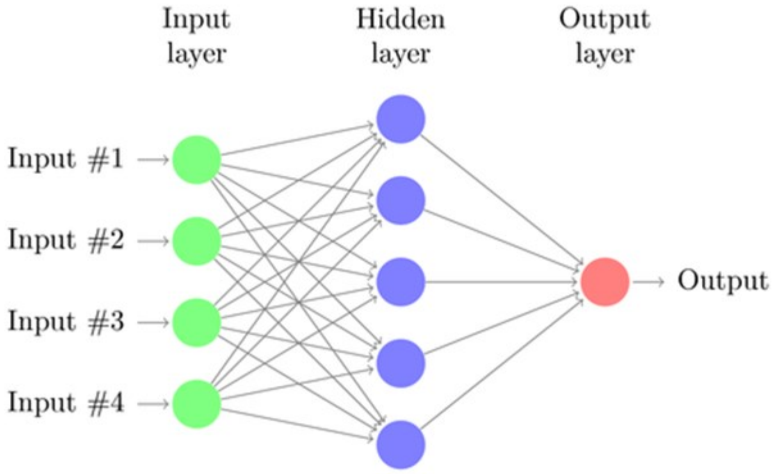

The input layer (with as many nodes as necessary) receives, yes, the input. The output layer which can either be made up of only one node or several, and most importantly and yet most mystical, the hidden layer. The hidden layer is never observed, but something does have to happen to the data as it passes from input to output layers otherwise it would all be pretty useless. It contains the “activation function” which determines how a certain node or neuron will fire depending on its inputs. The neurons of the hidden layers have outputs of 0 or 1 (off or on) depending on its inputs using these activation functions (of which there are many).

The output of the output layer will total to 1, allowing for the creation of a probability distribution where the strength of the algorithm’s “conviction” for any one answer is a probability of 0 to 1. The lines connecting input, hidden and output layers are part of what is called the network architecture, and each connection will have a particular “weight” (the importance) in relation to a neuron in the hidden layer, assessed during the training process where data is inputted with a known output and the model adjusts the different weights to best match the inputs to the outputs.

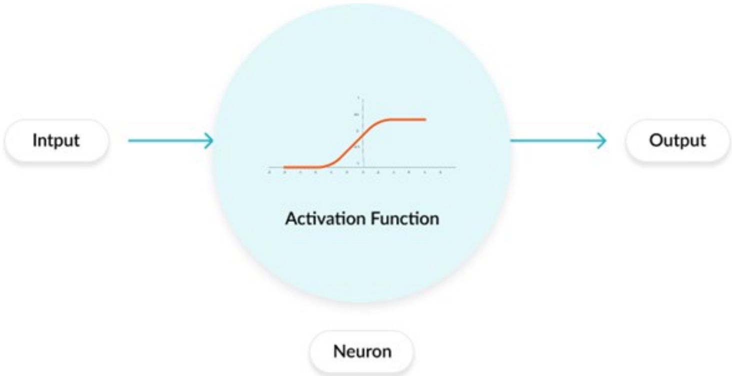

This particular neuron uses the Sigmoid Activation function:

$$f(x) = \frac{1}{1+e^{-x}}$$

Take, for example, a neural network used to tell photos of dogs and cats apart. The training data would be lots and lots of different photos of dogs and cats, all with the known label “dog” or “cat”. This training data would be put through the model and the weights (importance) of each connection between the inputs (which could be specific areas of the photo like the ears and the nose) are adjusted to create the best probability of a successful match, where every time an input is miscategorised, it would re-assess the weights of each connection to reduce the chance of error.

So as a summary, a neural network takes in a whole bunch of inputs, determines how important each one of them is in relation to each other, uses this importance to control whether a particular neuron in the hidden layer needs to be fired and then the pattern of fired neurons is used by the output layer to give a rating of confidence for a particular output. Now that I’ve (hopefully) made a little clearer how neural networks actually work; we can start to see where these algorithms would be most useful.

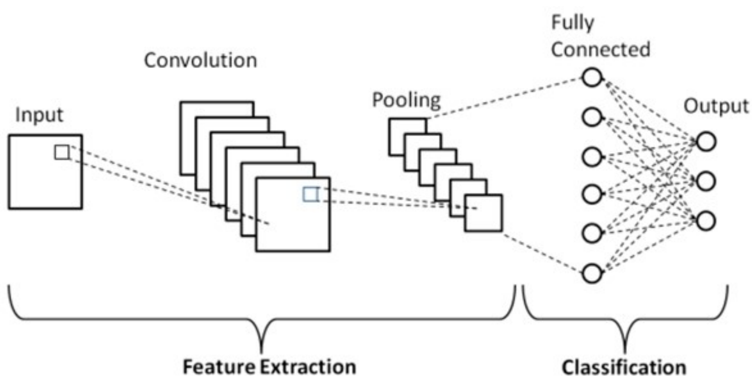

The basic architecture of an image classification neural network

Most neural networks are used to recognise specific elements amongst many inputs. This idea feeds into a lot of image recognition, speech recognition, character recognition and many more. All of these examples rely on what is the most **likely** category that a particular pattern of pixels or sounds will be in, which is exactly what a neural network is designed to do, since it generates a probability distribution across all the nodes in the output layer, each representing a category. So, using the example neural network above, if a picture of a dog is inputted then the model should output a high probability for the node representing the dog, and 1 minus that probability for the one representing the cat (therefore a lower one). With appropriate development and training, neural networks can be used to decipher hand-written text, transcribe human speech, draw connections between events that we may never have seen any reason to compare. Not only can neural networks classify, but they can also make predictions on outcomes based off of inputs, which can be particularly interesting in applications such as weather forecast, stock-predictions for bot-trading and improving pollster's voting intention prediction by including aspects such as the general sentiment in the news about particular stocks or politicians.

The fact of the matter is that neural networks are becoming ever more complex in what they can deduce, predict and categorise. Although humans have a natural drive to find patterns in nature, these algorithms can do so much more efficiently across several more dimensions than we can. Whilst the more dimensions we add, the harder it becomes to spot patterns, neural networks take all this data in their stride and can still give us increasingly reliable answers with less computational power than traditional machine-learning algorithms, and definitely less than humans. The future of technology is an ongoing move towards biomimetics. Millions of years of evolution have resulted in an incredibly efficient answer to questions such as "how do I get from point A to point B without being eaten or hitting anything?", and it thus stands to reason that we should look to nature when trying to improve the way our machines are beginning to "think" for themselves.

References

Watch Elon Musk talking about Neuralink on the Joe Rogan experience podcast at: https://www.youtube.com/watch?v=x5pgAM26wuM

https://towardsdatascience.com/understanding-neural-networks-what-how-and-why-18ec703ebd31

https://medium.com/@datamonsters/artificial-neural-networks-in-natural-language-processing-bcf62aa9151a

'HOOKE IN ISOLATION' TEAM

This first edition of *Hooke in Isolation* contains articles written exclusively by the editorial team. Offerings from the wider community are welcome in the second, and future editions.

The team worked collaboratively to design and build the structure and content, although Dr Evans and Mr Coward took the lead for this first effort.

For future editions, roles will be reassigned to others in the team—we're generous like that—so the following is merely a suggestion of responsibilities.

Editors

RJE & ETAC

Sub-editors

Ethan Kang

Ben Weiss

Contributors

Sinan Aramaz

Adam Moec

Daniel O'Keefe

George Weston

Cover art was designed by Sinan Aramaz

Moore's Law and why it's in

I remember learning about transistors, small electronic switches capable of controlling large currents and voltages, for the first time in FF Electronics. Now having gone through three years of lessons and a practical project, I know that they really are the building-blocks of all things electronic.

Simply put, Moore's Law dictates that the number of transistors in an IC double every two years while the cost halves. This exponential growth is the reason that we can all carry a mobile computer in our pockets while keeping prices under control. Nowadays, a single IC contains billions of transistors. Gordon Moore, co-founder of Intel, made a prediction in 1965 that described a doubling of transistors every year for at least another decade. In 1975, he revised the projection to a doubling every second year. This rate has held relatively steady over the years but growth was definitely faster in the first decade. However, experts say that semi-conductor advancement has been slowing since 2010 below the rate expected by Moore's Law. In fact, in 2015, Moore himself said that Moore's law would cease to apply within the next decade. Despite this it is important to consider that leading companies such as Samsung were planning to mass produce 5nm chips by 2020.

The interesting aspect of the law is that many online explanations fail to mention that it is purely an empirical relationship. It is only a correlation that is supported by observation and not on theories. Another relatively unknown fact is that there is further evidence that Moore's law will stop being relevant within a couple of decades. Dennard Scaling, which states that as transistors get smaller, their power density stays constant, so that power use stays in proportion with area. According to experts, Dennard scaling stopped in 2000s which means power usage has been increasing compared to a constant area. Eventually this will mean that Moore's Law will fail.

The law is often misreported as 18 months rather than 24 on account of a different forecast that computer chip performance would double every 18 months. This was drawn from both Dennard Scaling and Moore's Law. It is no longer that accurate as Dennard Scaling is becoming less and less accurate with every passing year. A much less known law is sometimes known as Rock's Law. As the cost of processing power to the consumer decreases, the cost for companies to fulfil the law increases on account of more R&D and testing. This has led some to say that Moore's law is a self-fulfilling prophecy. Some conspiracy theorists believe that they have intentionally slowed down the advancement of technology to protect this gold standard in the industry. With this seems far-fetched, it is a truth that when fulfilling this prophecy is not economical for the company, it will no longer be met. Whether that be because we cannot advance the miniaturization of transistors beyond atomic levels, provided that the mother of all inventions is necessity, or because research costs are too high.

It follows nicely that thought of future innovation. After Moore's law gives out,

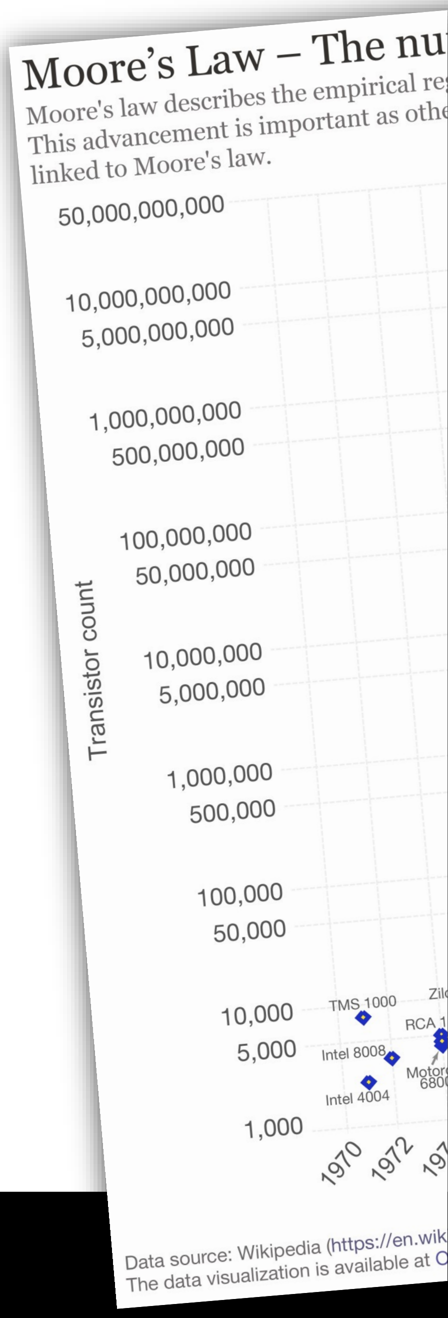

they are a couple of options for the subsequent computers. Quantum computing appears to be a frontrunner currently. The reason we cannot make transistors any smaller is because in the quantum realm, our current understanding of physics breaks down and our computers would not function in the manner in which we designed them to. While delving into qubits and superposition would provide knowledge, I couldn't really manage to do it justice in this text. In essence, we would take advantage of the fact that qubit, the equivalent of a bit, can be in multiple states at once. This means that with qubits we can have exponentially more configurations that bits. The other plausible technology is neuromorphic technology. This is essentially mimicking the brains structure to create a chip that is not built on the Von Neumann Architecture. This would mean that they would be capable of remembering and calculating just like our current technology but also make decisions and 'learn'. Just like the simulations!

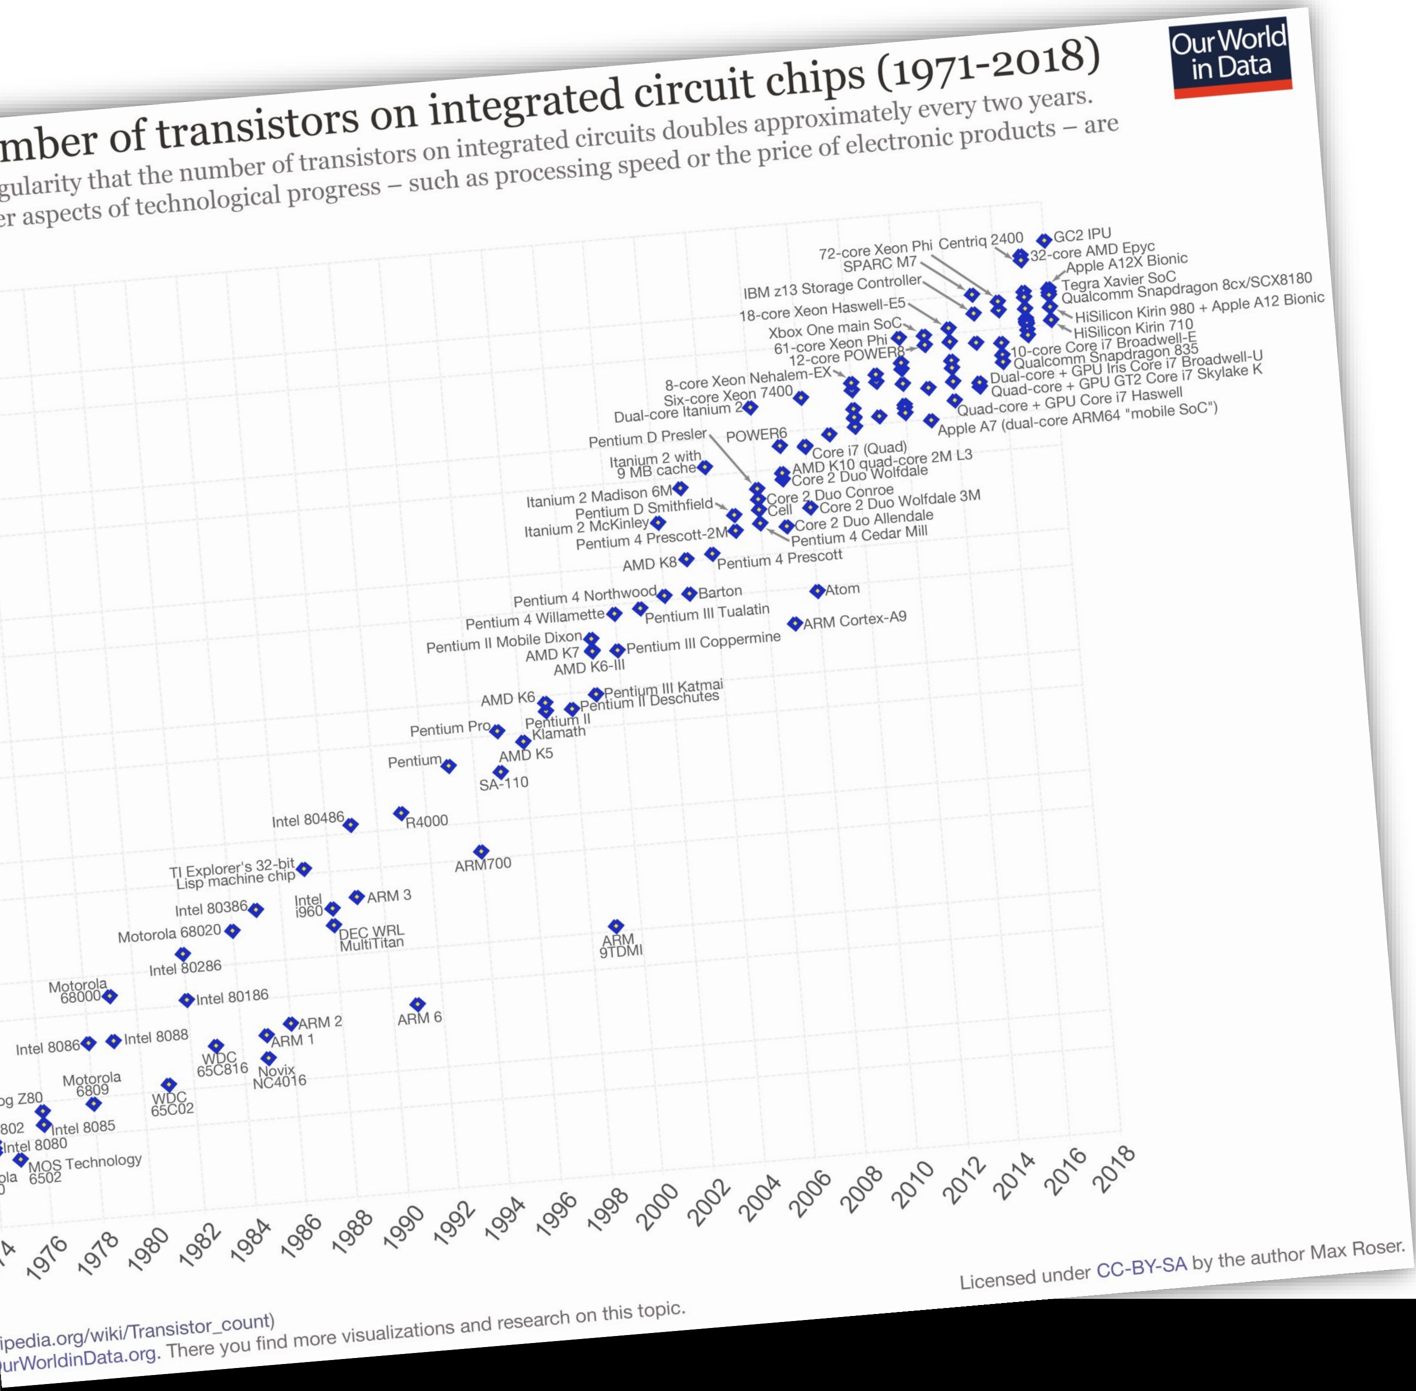

Horvath's Clocks

... and the quantification of life

Tales span back for millennia depicting the temporality and ephemerality of life. Be it through the Grim Reaper and his hourglass, life's 'hour upon the stage' as likened in Macbeth or the mysterious death-prophesising app from the 2019 horror flop 'Countdown', the trope of our days being (literally) numbered has made itself comfortable within the popular psyche. This prospect, however, has largely been explored in the abstract, pertaining to notions of fate and destiny. A curiosity for the concrete has led me to the more recent discoveries of Steve Horvath and others in the genetics field concerning the epigenetic clocks: methods of measuring age within the human body while also developing further uses such as lifespan and morbidity prediction. In this article I will discuss some of Horvath's epigenetic clocks in greater detail as well as related research in different areas of study.

Epigenetics is the field of study concerning heritable changes in gene activity and the phenotype by means other than genotype alteration, encompassing various processes and consequences of aging and development. Several epigenetic biomarkers can be indicative of age, such as histone modification in the nucleus, levels of gene expression[1] or telomere degradation[2] in the chromosome. However, the biomarker that has shown the greatest efficacy and accuracy in age prediction is that of DNA methylation, mostly to the base cytosine. This is simply the process of covalent addition of a methyl group to the DNA as catalysed by the DNA methyltransferase family of enzymes. This process affects gene function and expression as well as transcription[3], and levels of DNA methylation have been shown to strongly correlate with the aging process[4]. Horvath has likened it to a sort of epigenetic 'rust'.

Horvath first encountered this property accidentally, having noticed it during the analysis of a saliva sample. The ensuing process of age estimation involved measurements of methylation levels at distinct cytosine-guanine pair (CpG) locations across the genome which could be analysed and consequently used as a predictor of epigenetic age also known as 'DNAmAge'or DNA methylation age. This would be done by the use of a leave-one-out model in which a multivariate regression model is used to relate methylation levels and age in all the test subjects excluding one. Remaining subjects could then have their age predicted via the gathered data of the other subjects and compare it to their confirmed ages, forming a simple yet remarkable plot as shown. Strong correlation is displayed via a Pearson correlation coefficient of 0.83.[5]

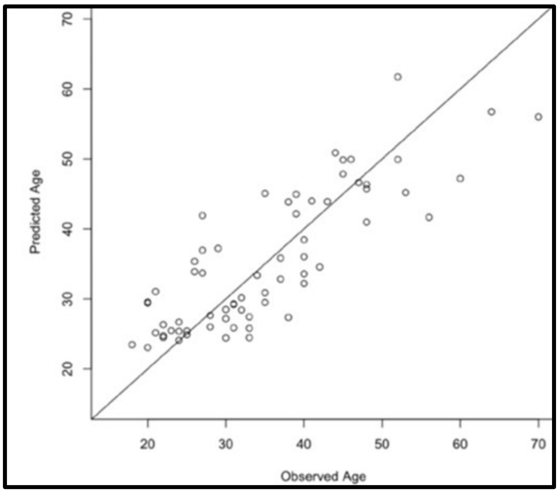

[1] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5821249/
[2] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3370421/#S6title
[3] https://pubmed.ncbi.nlm.nih.gov/1943996/
[4] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3482848/
[5] https://www.researchgate.net/publication/51469639_Epigenetic_Predictor_of_Age

From there onwards other specialised clocks continued to be developed. Clocks were developed using data from blood samples then post mortem brain tissue samples leading up to the creation of Horvath's most well-known epigenetic clock often referred to as the 'Horvath clock'. This was the first developed 'pan-tissue' clock: a much broader spectrum of samples from different parts of the human body could be used. In its creation, 7844 non-cancerous samples comprising over 21,000 CpGs were used from 82 different datasets comprised of 51 different tissues and cell types. DNAmAge was predicted using the calibrated weighted average of 353 CpGs selected automatically by an elastic net regression model, which could then compare predicted to confirmed ages as shown here in the test data. A stronger Pearson correlation coefficient of 0.96 was recorded across all the test data.[6]

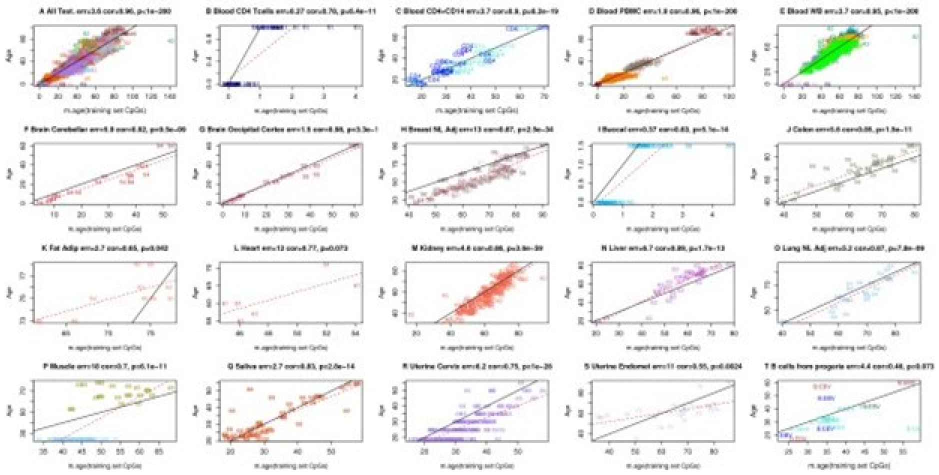

As epigenetic clocks continue to evolve, specificity and utility are invigorated. The 'skin and blood' clock tends to deviations in accuracy concerning age of fibroblasts and generates stronger predictions from skin and blood samples[7] while the GrimAge clock acts as a strong predictor of lifespan, healthspan (the proportion of life spent ailment-free) and morbidity[8].

Thus far we have only discussed DNA methylation and the epigenetic clocks in their relation to aging and age prediction, though there exists an interesting connection in the process' relation to lifespan (and so mortality). Epigenetic 'DNAmAge' does in fact relate to all-cause mortality later on in life[9], allowing for lifespan prediction as mentioned. The largest meta-analysis of this claim involved 13 cohorts comprising 13,089 people[10]. Blood samples had historically been taken from these individuals and frozen, which could be compared to the status of each individual as recorded more recently. A Cox regression model was used to assess lifespan and mortality predictions made from the initial samples, comparing survival time of subjects and measures of epigenetic aging. The data accurately predicted lifespan and time to death even after factoring in numerous confounders such as BMI, smoking history, hypertension and prior disease[11].

[6] https://genomebiology.biomedcentral.com/articles/10.1186/gb-2013-14-10-r115
[7] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6075434/
[8] https://pubmed.ncbi.nlm.nih.gov/30669119/
[9] https://genomebiology.biomedcentral.com/articles/10.1186/s13059-015-0584-6
[10] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5076441/
[11] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5076441/

continued...

This runs in a similar vein to some of Horvath's research done on mammals and the differences in longevity between mammalian species. Lifespan generally correlates to size[12] but there are some marked exceptions, such as the Brandt's bat's 40-year lifespan as compared to its 4-8g adult body mass. Research has unearthed the significance of DNA methylation in such cases, being responsible for at least 50% of longevity differences between species. Furthermore, lifespan of mammalian species has been shown to strongly negatively correlate with rate of methylation[13].

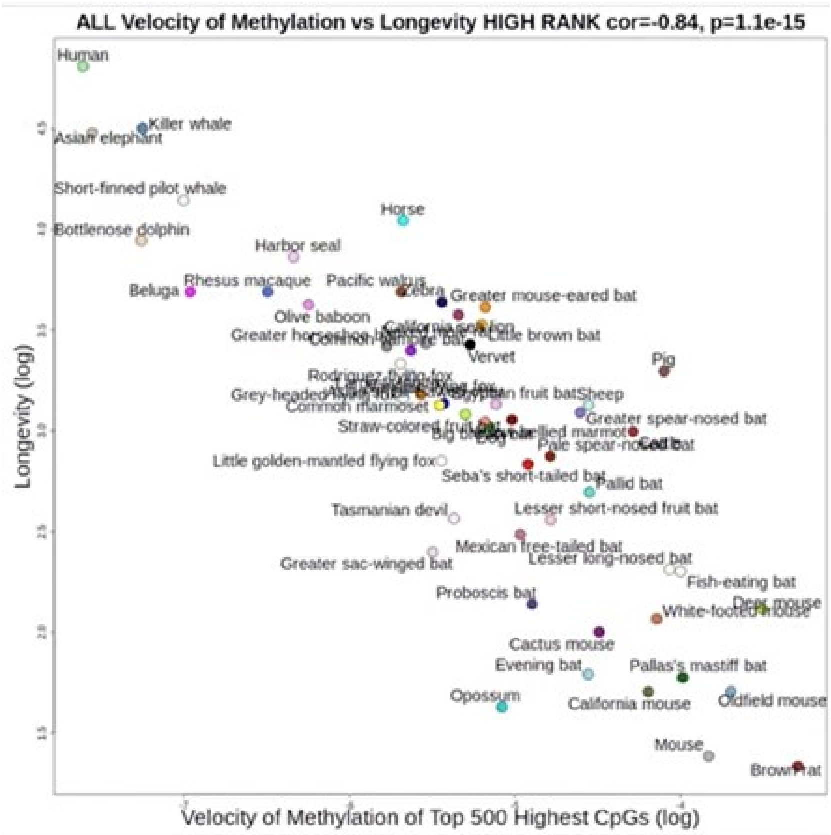

The epigenetic clocks and the insight that they provide could be valuable in several fields and studies such as forensics, developmental biology, gerontology and diagnosis as well as personal use. Moreover, the clock itself is a fairly recent progression in genetics and so still is at the cutting edge of research into new concepts and treatments moving into the future.

Overall, therefore, while science-fiction hour-and-minute precision isn't quite yet realistic, the epigenetic clock has and will continue to challenge our perception of science fact.

[12] https://pubmed.ncbi.nlm.nih.gov/15855403/
[13] https://www.youtube.com/watch?v=LuQKXux8UlE (13:02)

FROM THE ARCHIVES

This cartoon, first published in 1994, features superheroes based on two ex-Westminster teachers: Martin Robinson and Rod Beavon.

SARS-CoV-2

Adam Moec

The graph used by the government in their daily Coronavirus press briefings reveals something strange about the data on COVID-19 fatalities. The Sundays and Mondays report far fewer deaths on average than weekdays, an observation that begs the question "**Why?**".

The term for this disparity between weekday and weekend figures has been coined as the "weekend effect" by journalists, a term that already existed to, ironically, describe the usually higher mortality rate of patients admitted to hospital on weekends. Needless to say, there isn't actually a reduced mortality rate for COVID-19 patients on the weekend but rather a serious underreporting of figures due to the way that hospitals process fatalities on the weekend.

Similar instances of the "weekend effect" have been noticed throughout other hard-hit countries such as Spain for example where the term "Tuesday effect" came to describe the sudden increase of deaths seen following the weekend lows.

Although there is a noticeable increase, it is by no means as stark as that of the UK which seems to be particularly affected by weekend underreporting. No official explanation of these sudden dips has been provided by the UK government other than a lag in processing, reaffirming the need to use the 7-day rolling average instead of absolute figures to decide when we can move on with the loosening of restrictions.

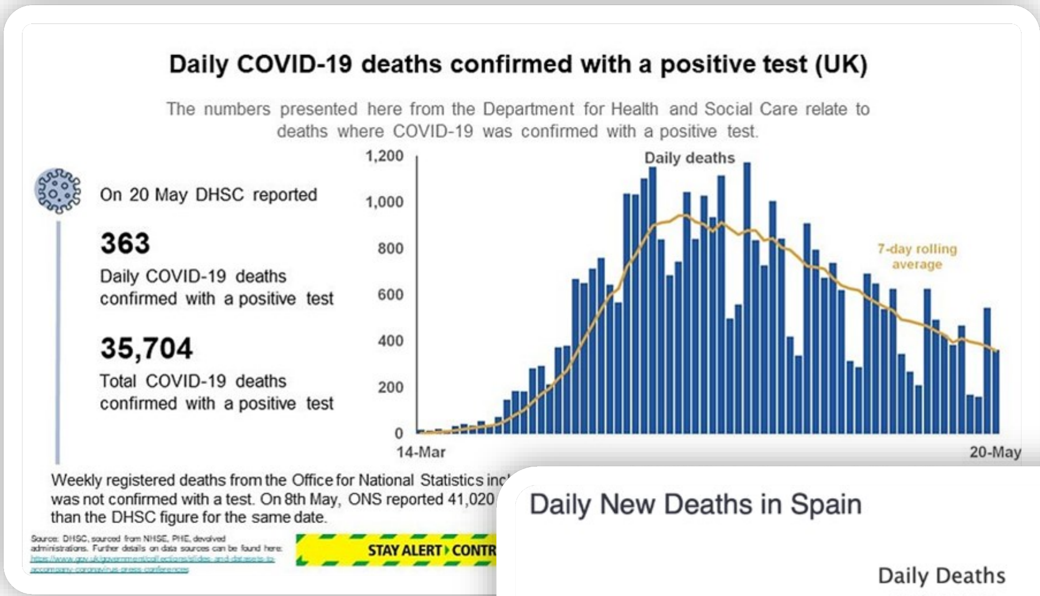

Graph taken from the Worldometer's Coronavirus data on Spain

PUZZLES

Cryptic Crossword

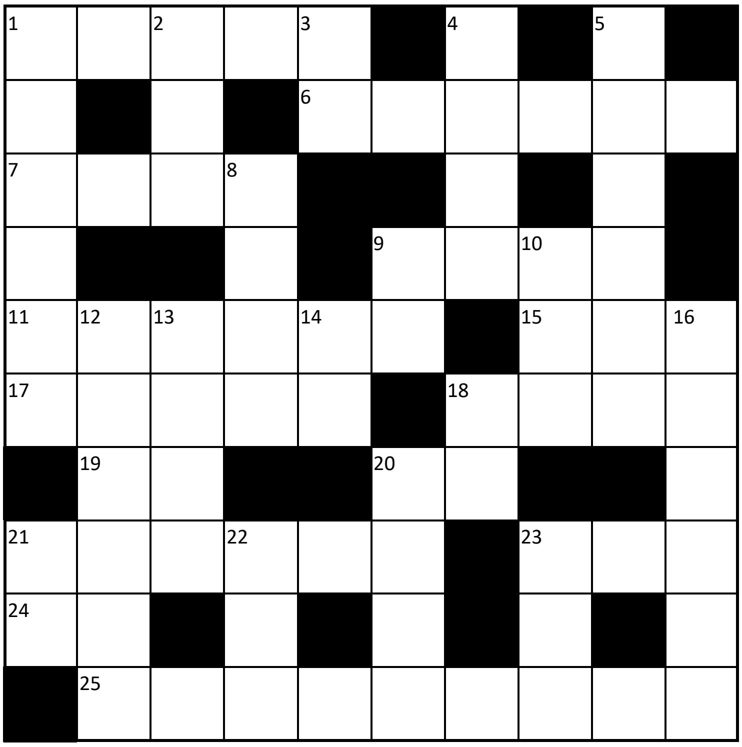

Across

1 Rugby player not quite complete scientist (5)

6 Five permitted to include moon shade (6)

7 Platinum found around each deposit (4)

9 While Astatine battles Io, Perhaps? (4)

11 Country with hesitation for a metal that's named for it (6)

15 Slimey mess involved in dripping ooze? (3)

17 Doctor Curie loses, in other words, before hesitation (5)

18 Over in space, second is a long time (4)

19 Round-about element? (2)

20 Sun god backs noble gas (2)

21 Aluminium street, for example, second ever! (6)

23 Worker starts to build Edison's engine (3)

24 Element found in cent and pence, but not in franc (2)

25 Nothing mixes with oil stain in our situation (9)

Down

1 Nervous reaction after truncating shape? Touchy! (6)

2 Sounds like it could be maiden eggs… (3)

3 Einstein's first volt unit (2)

4 Charges nothing in half-hearted taverns (4)

5 Modern heavy-weight physicist? (6)

8 Digital tree exchanging energy for current? (4)

9 Parent backs synthetic element (2)

10 Before silver duck… (3)

12 Clue in messy atoms (6)

13 Stalemate after doctor follows a woeful opening (4)

14 City's resistance following radioactive element (2)

16 Noises formed from collagen (6)

18 Hesitation at one of 4 discovered in Swedish town (2)

20 Is back in rehab, meeting on this continent (4)

21 Current charge for cooling unit (2)

22 Article on unknown oxygen compound (3)

23 Car loses weight for current measure of health (3)

The first correct solution to be submitted to the 'Hooke in isolation' team wins a £5 amazon voucher. Clues are scientific in either their construction or in their answer.

Opinions expressed in *Hooke in isolation* do not necessarily reflect official school sentiment.

Edition 1 was first released on Tuesday 9th June 2020.