Warm-up question

How can we tell if a signal is from aliens?

Hello world

FRBs. Fast radio bursts are brief (~ms), high-intensity radio wave signals (~1400 MHz), probably originating outside the galaxy. Since they are so far away, they must be extraordinarily energetic to stand out from background noise.

- Roughly, how would we put a *lower bound* on the distance, using an interferometer with multiple receivers? *Hint*: think about the shape of the radio wave's wavefront.
- The wavefront from a point source is *curved*. This means that if it is sufficiently nearby, an interferometer with multiple receivers will receive the signal at slightly different times. There is some maximum, interferometer-dependent distance at which this curvature can be distinguished, which is how we get a lower bound.
- [Advanced.] This curvature idea is not the full story. Really, we need to consider diffraction, and how close infinity looks to an antenna.
- The lower bound on distance for some FRBs is 10,000 km, coming from the UTMOST interferometer in Australia. If students are interested, we can give a rough sense of where this comes from, but it turns out to involve some antenna physics [really Fourier optics, but let us ban these words]. For an incoming signal at wavelength λ, and D the size of the antenna, the furthest point source it can resolve into a point is at distance $2*D^2/\lambda$. Anything further out looks like a source at infinity. UTMOST consists of two parabolic detectors about 800 m long, and we are looking at radio waves around $\lambda \approx 20$ cm. What is the lower bound?
- Plugging numbers into the formula, I get around 6000 km, which is in the ballpark.

There is no consensus about what causes FRBs. Less exotic hypotheses include collisions of compact objects (such as neutron stars or black holes), supernovae, "blitzars" (neutron stars collapsing to form black holes), magnetars (neutron stars with powerful magnetic fields), and stellar flares. Of various exotic hypotheses, the most whimsical is extraterrestrial life.

- Although FRBs are by definition transient, some sources are "repeating", with multiple FRBs coming from the same place in the sky. Which hypotheses does this rule out?
- Presumably any mechanism which cannot be repeated at the same point, such as collisions, supernovae, or blitzars.
- LIGO has been operating at the same time as we have observed FRBs, and has not seen any gravity signals. What do you think this rules out?
- Probably the same non-repeating mechanisms, which we would expect to generate powerful gravity waves.

- (Time scales? Not much variation in the signal, so perhaps it's shorter than the time scale of the mechanism generating it.

***Pulsars*.** FRBs are not the first astrophysical signals to be attributed to aliens. In 1967, a periodic source of radio waves was discovered, called a *pulsar*. At the time, no known astrophysical mechanism created periodic pulses, so its discoverers jokingly christened the source *LGM-1*, for "Little Green Men". Maybe it was an extraterrestrial clock!

- Think of some ways that periodic behaviour can arise in nature.
- Rotation, pendulums, orbits, and in fact any perturbation around a stable minimum.
- Does periodicity seem like a signature of intelligence?
- Um, no. It's everywhere! I doubt anyone seriously thought pulsars were extraterrestrial, but rather that a natural, periodic mechanism for the signal would soon come along. And so it did!
- But does repetition rule out periodicity? (Pedram's point about resolution and structure.)

Soon afterwards, a *second source* of periodic signals was discovered, very far from the first, effectively ruling out alien signals. By 1968, theorists realised that pulsars were magnetised, rapidly spinning *neutron stars*, shooting out charged particles along the axis of magnetic field: something like a lighthouse crossed with a railgun! As these particles brake, they emit EM radiation, including the radio waves first observed.

- Why does a second pulsar rule out aliens? What does this suggest about FRBs?
- A natural mechanism easily explains a second source: if nature can do it once, it can do it again! Intelligent, technologically advanced aliens may spread around their local star system, but they are likely to generate different-looking signals, and sources are unlikely to be far apart.
- Similarly, FRBs have multiple sources. They are probably not messages from aliens.
- Are we being too skeptical? Could a civilization have conquered large parts of the universe? How could we constrain how large this angular patch is/are there constraints? (Cosmic history, inflation, etc.)

Hydrogen line. Perhaps the most promising candidate for an alien transmission was the Wow! Signal (1977), a high-intensity, 72 second pulse of radio waves from somewhere in the Sagittarius constellation. It has never been seen again.

- Multiple sources seem to rule out aliens. Does non-repetition from a single source make it more or less likely?
- In fact, repetition is one of the *requirements* for a candidate extraterrestrial message, at least as defined by METI. This is partly due to the high standard of evidence needed when explaining something with aliens!

- But an intelligent civilisation unintentionally broadcasting will probably send multiple, single-source transmissions, while a deliberate message (e.g. a distress call) will probably be repeated to maximise its chances of success.
- But Sean's point: with a natural mechanism, you would expect to see it again! It could be super rare or anomalous. We haven't been measuring for very long on cosmic time scales… Some stuff is just rare.

The frequency of Wow! was around 1420 MHz, and this is part of the reason it excites astronomers so much. Let's see why! Atomic hydrogen, the most abundant substance in the universe, absorbs and emits light at 1420 MHz, corresponding to the difference between its ground and first excited states. This is called the *hydrogen line,* referring to the characteristic lines of emission and absorption spectra. The universe also happens to be quiet around that frequency, in a small band called the microwave window.

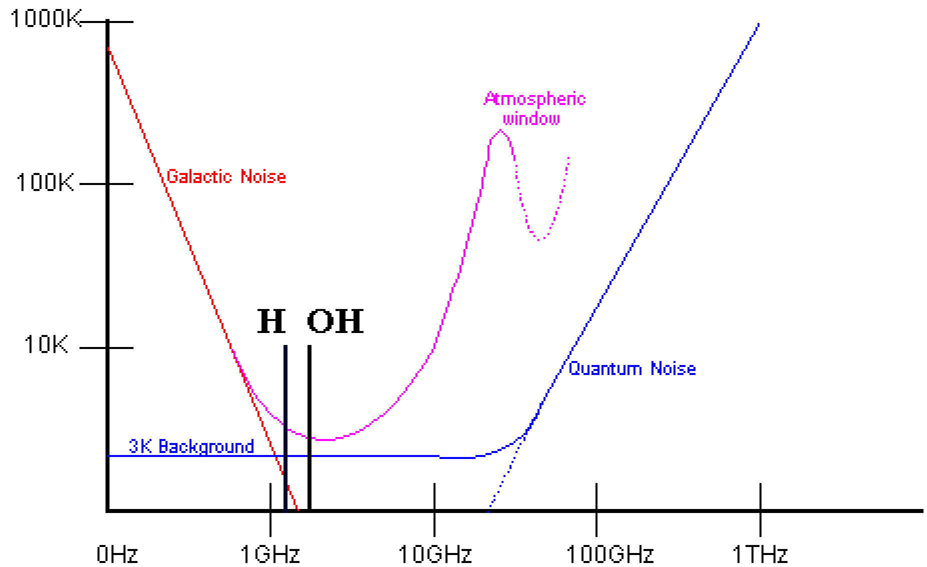

- List some possible sources of observational noise.
- There is the CMB, a low-level hum at 3K throughout the universe. There is "galactic noise" from synchrotron radiation (from particles braking in magnetic fields), and emission from the atmosphere and human-made sources. Finally, quantum noise is in the detectors themselves.
- On earth, it is illegal to transmit on frequencies in the microwave window. Why?
- Signals from earth can bounce off clouds, the ionosphere, etc, and back into our radio telescopes. We might confuse Uncle Bob's ham radio for a giant ball of hydrogen or signs of extraterrestrial intelligence!

The hydrogen line is a perfect lens for peering into deep space, since the region between stars is mostly hydrogen atoms: an interstellar chorus of voices singing at 1420 MHz. This also makes the hydrogen line an excellent candidate frequency for communication between intelligent species.

- Can you think of reasons an intelligent species might not use the hydrogen line?
- Here are some random thoughts. First of all, they might not know about the hydrogen line, either because they are less advanced or have developed in a different direction (e.g. pure maths).
- They could also be much more advanced, communicating using an entirely different method that we are ignorant of, but which is just as obvious to suitably developed civilisations.
- Finally, perhaps the most compelling reason – aliens might not want anyone to find them! Who knows what's out there? It only takes one evil intergalactic civilisation for this to be a bad idea… and in fact, SETI is primarily focused on listening for this reason!

Advanced notes. *What is Wow!?* What could the Wow! signal be? Although the central frequencies are similar to an FRB, the Wow! signal is long and narrowband (7s, Δf ~10 kHz), while FRBs are short and broadband (~ms, Δf ~100 MHz). Given the non-repetition and lack of structure, I personally incline towards the (boring) hypothesis that it is a terrestrial radio source which got reflected, or a transient natural occurrence in the Sagittarius constellation.

Why does the ISM glow? The ISM is illuminated by a continuous spectrum of starlight, which usually includes the hydrogen line. Hydrogen atoms will absorb photons at that frequency, and after a short delay, re-emit them in a random direction, explaining why the ISM glows. You might worry that the ISM will scatter its own signal. But the ISM is low density (around 1 atom per cubic centimetre), and elastic (Rayleigh) scattering is also inefficient for long wavelengths, scaling as 1/λ^4.

The Cosmic Water-Hole. The microwave window is also home to the *hydroxyl line* at 1660 MHz, associated with oscillations of hydroxyl (OH) ions, which the ISM also contains. Since H and OH are dissociation products of water, Barney Oliver dubbed the spectral gap between H and OH lines the *Cosmic Water-Hole*, since it would be a natural place for intelligent life to congregate. As Oliver put it: "Where shall we meet our neighbors? At the water-hole, where species have always gathered."

Inference and information

Alien abduction. In the examples above, you may have noticed that we blame aliens as a last resort, i.e. when we have no other plausible explanations. In the words of Sherlock Holmes, "when you have eliminated the impossible, whatever remains, however improbable, must be the truth." Let's write the form of reasoning up a bit more carefully: if some phenomenon X can only plausibly be explained by assuming Y, then (until further evidence comes along) we can regard

X as evidence of Y. It may seem a little weak, but this type of reasoning, called *abductive explanation* or *inference to the best explanation*, is ubiquitous in both science and everyday life.

- Can you think of examples of inference to the best explanation?
- Can you think of any issues with inference to the best explanation?
- I'll focus on an obvious one: our list of candidate explanations may be bad! For instance, we might regard X as evidence for a very silly hypothesis Y simply because Y was the least silly explanation on our list. (Think of the X Files. No pun intended!)
- What does "silly" mean? It could be something like "requiring many additional assumptions" or "in conflict with existing evidence". I think if we have silly options on our list, we should also have an option like "we don't know enough about the domain of X to abductively support our silly options". But in some cases, X could indeed constitute stronger evidence for a silly hypothesis than the hypothesis that our list of explanations is inadequate (once again, cf. The X Files).
- How confident are you that we can infer the existence of aliens from this sort of reasoning?
- Students will have different opinions. I personally think the option "we don't know enough about astrophysics" is more convincing than "aliens did it" unless we have specific evidence for intelligent life. We'll explore specific evidence now!

Frequency space. Can we do better than aliens as a last resort? In other words, can we find features of a signal that would constitute *direct* evidence for extraterrestrial intelligence? As a warm-up, let's think about what *we* look like.

- What signals does our civilisation unintentionally broadcast into space?
- A lot of radio waves! Radio, television, and radar all use frequencies in the radio band. There will be a smattering of other higher frequencies, but these are most lower volume and more easily scattered by the atmosphere.
- Do you think these signals would look natural or intelligent to aliens in our vicinity?
- A jumble of television and radio broadcasts probably looks like noise to a casual observer. But aliens might do better by looking in frequency space

Suppose we have some method for taking a waveform and splitting it into its component frequencies. (The mathematical tool which does this is called *Fourier analysis*, but we won't worry about the details.) We can represent signals on a Cartesian plane with frequency f on the x-axis and amplitude A on the y-axis, i.e. a point (f, A). We call this the *frequency domain*, in contrast to the *time domain* where we usually observe signals. For instance, in the time domain, a single frequency component of frequency 1 Hz and strength 1 (in some units) looks as follows:

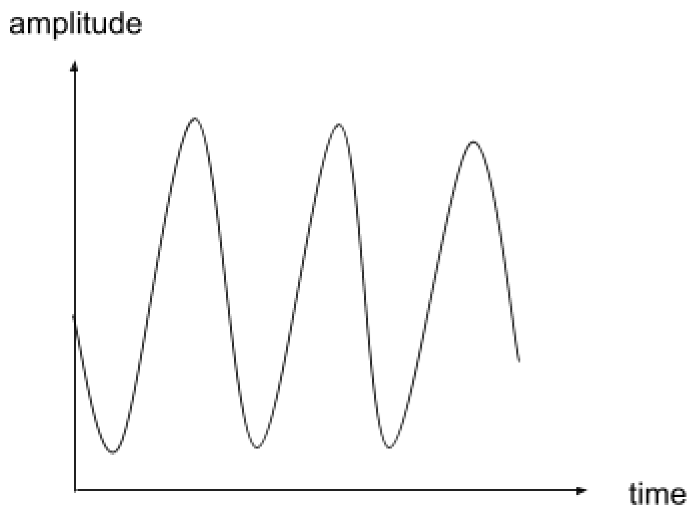

In the frequency domain, this looks as follows:

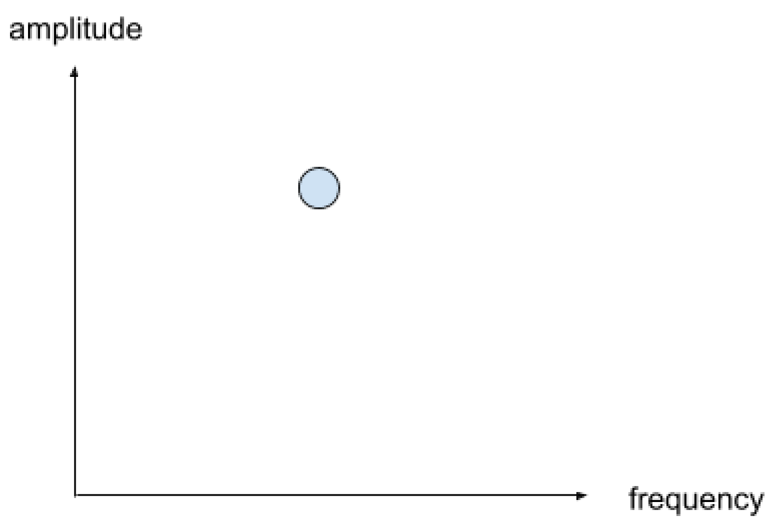

Another example is the static you see on an untuned television, called *white noise*. This is a sum of every possible frequency! We can represent the frequencies as a flat, straight line in frequency space. We say white noise has a *flat spectrum*. There are other types of coloured noise, and we can talk about that if people are interested.

- [Sum of sine waves, i.e. work out example from beat frequency. Can we do some sort of low-tech Fourier decomposition for multiple waves this way?]

Modulation. When you tune into a radio station, you are picking a specific frequency. You might wonder how music, made out of a huge range of different frequencies, can be magically encoded on a single frequency.

- Imagine a point in frequency space, (f, A). Without changing f, how could we encode information in the position of the point?
- Wiggling it up and down! That is, we can vary the amplitude A with time in order to encode a signal. This scheme is called *amplitude modulation*.
- Can you imagine an equally simple scheme which keeps the amplitude, rather than frequency, fixed? What other schemes can you think of?
- Wiggle it side to side! This is *frequency modulation*.
- As for other schemes, let them go wild!
- Now imagine that you are an alien, looking at signals from earth in frequency space. What might you expect to see? What features would suggest intelligence?
- Since we primarily use frequency and amplitude modulation on earth, you would expect to see many little points in frequency space wiggling up and down or left and right. (A point is probably more likely a very narrow band, and bands can broaden due to various effects, you will probably see wiggling lumps of frequency rather than points.)

Message structure. Using modulation schemes is certainly consistent with intelligence. Then again, there are many beautiful *evolved* mechanisms in the natural world, much more elaborate than frequency modulation, that did not require an intelligent designer. To be certain, aliens might want to analyse the actual signals we encode using modulation. This leads to a much harder and more interesting question: *how can we tell a message is intelligent, even if we don't speak the language*?

Let's think about some general features of language. For simplicity, we'll represent everything in binary digits (*bits*), i.e. as a sequence of 1s and 0s, which we imagine an alien race examining for signs of intelligence. The first thing to observe is that some things are *too simple* to count as language, for instance the signals 0000… or 01010101….

- How many bits do you need to encode a letter? A word? A paragraph?
- For n bits, you have 2^n possible sequences. There are 26 letters, so you need ≈5 bits (corresponding to 2^5 = 32 options) to encode a letter. A word is maybe 4 letters on average, corresponding to 20 bits, and a paragraph of 50 words words will take 1000 bits ≈ 1Mbit. Note that a *byte* is 8 bits, so a paragraph takes around 125 bytes.
- Why are the simple repeating messages 0000… and 0101... unlikely to contain messages? Would any repeating sequence suffer from the same problem?
- A sequence like 00000… contains only one bit of information: 0. Similarly, the message 010101… contains two bits: 01. This is not enough to encode a word! But a long repeating sequence could certainly contain enough information to encode letters.

In some sense, these simple examples don't contain enough information to say anything interesting. One simple way of quantifying this is *compressibility*: I can compress 0000… and 0101… into very short descriptions:

```
0000... = repeat 0, 0101… = repeat 01.
```

A compressible string is *described by a shorter program*, and the amount of information contained means roughly the length of the shortest program. On the other end of the spectrum is a completely random string of 1s and 0s, obtained for instance by flipping a coin. Although this has a simple probabilistic description (`flip coin`) it has no simple *deterministic* description. In fact, for a long random string, the shortest description is likely to be the string itself! There are certain statistical tests we can perform to see if a string is random.

- Can you think of a simple test to tell if a long string is likely to be random?
- Many options, but a simple one is checking that equally likely substrings occur equally often, e.g. 0 and 1, or 00, 01, 10, 11.
- Although a random string contains lots of information, does it seem likely to be a language?
- Probably not. Even though random strings carry a lot of information, language is not just a mechanism for conveying pure information. We communicate over "lossy" channels, subject to e.g. ambiguity, distraction, mishearings, etc. These can corrupt the message completely (e.g. mishearing a "not"). A language should have some built-in form of *error correction* or *redundancy* to guard against misunderstanding.

So, we expect our signals to be both information-rich (the shortest descriptions are still long enough to say something interesting, e.g. more than 100 bytes), but have built-in redundancy. This might be something as simple as repetition (e.g. repeat each bit 3 times), or more complicated (e.g. in English, certain words tend to be correlated, like in "rubber ducky", and aspects like grammar, sentiment, and so on). A fairly general and math-y way of saying this is that *different aspects of the message are correlated*. To sum up: the message should be compressible (due to redundancy) but not too compressible (so it says something interesting).

- Try taking a list of related words, getting rid of some letters, and getting your friends to guess what the words are. You have successfully compressed if they work out what the words are. What's the best you can do?
- Here's a list of related words, see if you can figure them out: apl, bnn, pnpl, pr. This is a compression of about 50%. I hope the students can do better!
- How compressible do you think English is? In other words, how many bits do you think we need per letter, once all the redundancy is taken into account?
- This is a fun factoid which I think will probably surprise people. According to Claude Elwood Shannon, the inventor of information theory, the true rate is about 1 bit per letter. Since it naively takes 5 bits to specify a letter, we can compress English by around 80%!

- We need a much more clever scheme, and possibly a computer, to achieve this compression, but on average it is doable.

– David Wakeham