Evolutionary Dynamics

Supreet Saini

Chemical Engineering

Indian Institute of Technology Bombay

Week 11

Lecture 55

Hi everyone, welcome back to the next video. We'll continue our discussion with this surprising observation from the LTEE: that landscapes may not have peaks, and adaptation may continue forever. Let's just look at where we were. So we have this time versus fitness data. This is the point at 20k.

This is the point at 40,000 generations. And using experimental fitness data from the first 20,000 generations, I built two models. Both these models made excellent predictions of the experimental data. This was one. And the second model predicted something like this.

Both these models were made with the data collected for the first 20,000 generations. After that, these models were asked to predict what would happen in the next 20,000 generations. This is the prediction made by model 2. And this is the prediction made by model 1. When we actually computed the fitnesses of the next 20,000 generations, we saw that these fitness values in the subsequent 20,000 generations were much better predicted by model 2 compared to model 1.

Remember model 2 was of the form that y is equal to a.t plus 1. to the power b. We call these a dash and b dash depending on their values determined by the data from the first 20,000 generations. And model 1 was y is equal to 1 plus a naught t plus b naught by t. And what these results clearly show that this model is a better predictor. Now, if that is the case, then the question is that this model also says that there is no fitness peak because fitness will keep on increasing. There is no Fmax associated with fitness on a landscape.

So there is no fitness peak. In this one, if I plug t equal to infinite, infinity, y also will reach the value infinity, which means that fitness will keep on increasing forever. And that seems very counterintuitive. Eventually, when it comes to cell division, there are

laws of physics that have to be respected. And cell division is a process that does need finite amount of time.

And cells cannot divide into two infinitely fast. So how do we think about the fact that this model is doing a better job at prediction as compared to this one? And I think there are two ways, there are two reasons to not be sort of dismayed at the fact that fitness doesn't have peak and that model is a better predictor than the other one. And the first reason is that the prediction that this model is making is better in a very short term, is better in the time window of 20,000 to 40,000 only. We don't know evolutionary change takes place, large evolutionary change takes place over millions of years.

And we don't know what's going to happen if we were to ask these two models to predict what would happen in the next 1 million generations. So, to illustrate this idea, let's imagine this, and now let's change the scale so that this is 20,000 and this is 40,000. And we start with a fitness of 1, and model 1 predicts something like this. And, of course, there is an Fmax here, which model 1 will not cross. So, this is Fmax.

On the other hand, the second model tells us a slightly different story. It also starts with one, but maybe its trajectory is 40,000. Maybe its trajectory was something like this. Let me draw model one again, and I get something like this. So, maybe over a very long period of time, the model predictions look like this.

This is always capped by Fmax, which it will not cross. So, what we have seen so far with the experimental data from LTEE is the fact that both these models do a great job at prediction in 20,000, but when it comes to 40,000, model predictions 2 becomes a better predictor. This is model 2. But over a really long time scale, maybe the fitness values, as we move forward in time, will begin to look something like this.

And hence, when I look at this, let us say this is half a million generations. If I look over evolutionary time scales, really long time scales, then I would still find that the reality is between Model 1 and Model 2, and Model 1 is sort of a better predictor compared to Model 2. So, Just because the model that said there is no fitness peak was a better predictor in the window of 20,000 to 40,000 generations does not mean it will continue to be a better predictor as we move forward in time. So, a scenario like this might play out.

So, we can summarize this by saying that Model 2 is a better predictor between 20k and 40k generations. This does not imply that it's going to be a better predictor at all times.

That's the first reason I think why it shouldn't be so hard for us to believe that Model 2 is a better predictor here.

That's sort of the first reason. The second reason is a commentary on this facet that Lenski wrote on the website dedicated to LTEE. And remember that while Model 2 says there is no peak, while Model 2 implies there is no peak, the rate at which fitness is increasing—the rate of increase of fitness—is actually continuously slowing down in Model 2. In Model 2.

So the time required for fitness to increase by a certain amount is going to get increasingly longer as as we move towards higher number of generations. Fitness is increasing, but that rate is slowing down all the time. And we have this fitness related to this model, which is A dash T plus 1 to the power B dash in model 2. And what this model, what the values of A dash and B dash show that in 1 billion generations, If I were to plug in T as 1 billion generations, remember T is time measured in number of generations, the Y that I get, the fitness that I will get as a predictor from model 2 is that the E. coli after 1 generation, although who knows if it will still be E. coli after 1 billion generations, E. coli or the bacterial species that we get after 1 billion generations will divide every

I think it's of the order of every five minutes. So while five minutes is extremely rapid, we know of bacteria that divide in times which are close to that time. Is it inconceivable that in a time which is as long as 1 billion generations, the adaptation that is taking place in those flasks will not evolve these organisms to divide every 5 minutes? To my mind, this prediction that this model makes is not impossible. So while the idea that no fitness peak and hence infinite fitnesses seems a little ridiculous and counterintuitive, I think these two mitigating factors sort of take care of what Model 2 is trying to tell us.

Although I won't be around to see what happens in one billion generations, but I wouldn't be surprised if E. coli was dividing every five minutes or six minutes or something of that nature. Once one billion years or one billion generations of this experiment have been processed. All right. So that's the discussion of this paper. This paper was published in 2013 in the Journal Science.

Next, we come to perhaps what is considered the most interesting observation from the LTEE, and this was something that happened in the 2000s. So, let us go back to the experiment, and we have 12 identical flasks. These are 12 identical lines which are being evolved in identical environments, which were seeded by an identical equalized strain

and so on and so forth. So, everything about them has been precisely the same ever since the experiment was started in 1988. The culture media in which they are growing has this characteristic that we are trying to study their adaptation.

We are trying to study how bacteria cope with low amounts of glucose as the only carbon source available to them. So, low amounts of glucose. As the carbon source in the environment. So, low amounts of glucose. Now, however, at the start of the experiment in the '80s, what was also done was that this media was supplemented with citrate in the flasks.

Citrate was added to this culture media when the experiment was started. And remember that the conditions in which these bacteria and flasks are being propagated forward are 37 degrees Celsius with shaking, which means that these growing bacteria have access to oxygen. And this, while it wasn't realized at the time, led to the most surprising observations from this experiment. E. coli has this characteristic property that in anaerobic conditions, when oxygen is not around, it can utilize citrate as a carbon source. But in oxic conditions, when there is oxygen around,

E.coli cannot use citrate as a carbon source. That is a characteristic property of E.coli which has also been used historically to perform diagnostics on whether a pathogen or a bacterial isolate is E. coli or not, that it cannot use citrate in aerobic conditions. Despite that, and obviously this was known at the start of the experiment, and despite that citrate was added to the culture media when the experiment was started. The sort of the purpose behind adding citrate was that it chelates metal ions better.

So it makes historically citrate has been added to culture media to make these growth curves more uniform and more consistent and not depend so much on the quality of water and other reagents that are being added to the media. So that is the premise of adding citrate to these flasks. But what that also leads us to is the fact that there is this lots of resource available to bacteria which is not being currently used. The experiment is being done in this environment. However, if the environment was this, then E. coli can use this resource for their growth.

But imagine a bacteria in the flask where everybody is fighting for the small amounts of glucose present. If I as an individual pick up a mutation which allows me to utilize citrate for carbon source, that would be a great beneficial mutation because that will give me unfiltered access to this giant pool of reserve that is present in the culture conditions. So I could get access to it, but obviously there would need to be mutations which need to

happen before I can start to utilize citrate. But what I would like you to realize is that in this flask, there is some amount of glucose. Let us say this is glucose.

And there is lots of citrate. This is citrate. And The bacteria can only utilize glucose, and now if one of these bacteria picks up a mutation that allows it to use citrate for growth as well, then this mutation would be a great mutation because it allows unfiltered access to this great resource, which is not being used by anybody else in the population. However, it was not anticipated that such a mutation would happen, and there was no basis for us to think that such a mutation could take place and such an evolutionary change would result from this experiment.

So, what has happened since the start of the experiment is that we have these 12 lines of these experiments, one of these lines at around generation 31,000. So, in order to understand this, you have to understand the logistics of how an experiment like this is conducted. Since subculture from one flask to another flask happens every 24 hours, So what this would mean—and the same happens in my lab, for instance—is that since this subculture of one in 200 is happening every day at the same time.

Let's say this transfer happens at 10 a.m. This is full of bacteria, and we are going to transfer. So at 10 a.m., somebody needs to come to the lab. Take a one in 1000 fraction of this and transfer it to a new flask, which contains the same recipe of nutrients available to it. And because this is happening every day, you develop a familiarity of what a flask looks like when you come in after 24 hours of growth.

So, flasks which have media that do not support a lot of growth will have some bacteria in them. And these bacteria will manifest as this haziness in the culture. As more and more bacteria grow, if you have a flask which contains lots of resources, then more bacteria will be present. And those cultures, after 24 hours of growth, will look much denser. So, if you are a microbiology student who has done this kind of thing, you develop an intuition for how dense a culture is every day, and you expect that it has to be the same density today as well.

However, In the course of this experiment, what happened at some point around this number of generations is that, on one day when somebody came to transfer, 11 flasks had—the technical term is—optical density. And this is just a measure of how densely the culture had grown in the flask. So, in 11 flasks, the optical density was normal, as usual. But in one flask, the density of the culture was much higher.

So, as someone who performs evolution experiments, when you see something like this happen, your first instinct is to say that what must have happened is that some contaminant bacteria must have gotten into the flask. Either during the process of transfer or something must have gone wrong, and you had this contaminant enter this particular flask. As a result, this culture is just totally overcrowded because of this contaminant growth. Because the bacteria that I am working with do not grow to such high densities. I know that from prior experience of every single day when I come in the morning to do these transfers.

Hence, there must have been something else that happened. And the most common cause of that something else is that some contamination occurred in the growth media. So. That's sort of the first explanation and the most common and obvious explanation when something like this happens. So, this culture, this flask was checked for contaminants.

However, it was found that there was no contamination, and the flask contained only bacteria from the experiment. Bacteria from the experiment. And then further analysis revealed So, if that is the case, why has the number of bacteria increased so much? Because these bacteria are growing only on glucose as the carbon source, and there is only a small amount of glucose as carbon source.

So, glucose cannot support the carbon requirements of so many cells to sustain such a high density. And then it was found that actually this flask's density has increased because it contains cells which are growing on citrate in the presence of oxygen, which E. coli doesn't do. So, that is what we have as the most surprising observation from LTEE. This happened soon after the turn of the century, a few years after 2000.

So, what is going on here? Now, even after 80,000 generations, we only have one line that is growing on citrate. Citrate. None of these 11 lines have learned to use citrate in this time. None of these 11 lines have citrate utilizers.

So, what does that tell us? That obviously some mutation happened which facilitated the utilization of citrate in this environment. But what sort of mutation was that, that only one line in all this time—remember, this happened at around 30,000 generations? Another 50,000 generations have since passed, but citrate utilization as a capacity has not yet evolved in any of the other 11 lines. So, what does that tell us?

What could be the reason that only one out of 12 lines has done it in all this while? And again, the fitness advantage to be had—if I, as an individual in one of those flasks,

evolved the capacity to utilize and grow on citrate—the advantage that I have compared to other bacteria in the population is huge. So, why hasn't this happened? Why hasn't this ability evolved in other flasks? So, this could be because of citrate.

The ability to utilize citrate could have evolved because of two reasons. The first possible reasons, two possible reasons. The first possible reason says that a very fortuitous lucky mutation happened, a very lucky or a rare mutation happened in flask I, in flask of line I, which facilitated utilization of citrate. Now, that makes sense because maybe it required duplication of certain parts of the genome, followed by a deletion, followed by a SNP. So, all these events needed to happen and maybe this was a very precise and a very lucky mutation that needed to happen.

And because suppose this is the genome of the individual. maybe you needed just this exact piece of DNA maybe to get flipped maybe to maybe you needed this region to get flipped and only then citrate utilization would evolve any any more or any less of the region being flipped wouldn't be able to do it and hence that kind of event happens only very very rarely and this was a very rare such mutation that happened in this line it could have that rare mutation could have happened anywhere else alternatively Line I acquired mutations in the time leading up to acquired mutations leading up to 30,000 such that Such that individuals in this line were primed to use citrate.

These mutations happened which brought the cell to a point where it was just ready to utilize citrate and one more mutation and then it is ready to utilize citrate. Such that one additional acquisition occurred. one additional mutation would lead to citrate utilization. These prior mutations may have conferred benefit in terms of being able to grow better on glucose, but one of the side effects of acquisition of these prior mutations was that if one of the cells in this population acquired this one precise mutation, a regular one, not a rare mutation, if this one happened, then it is ready to utilize citrate. If this same mutation happened in any other line, but because its prior mutations are other mutations, this wouldn't be able to use citrate because this mutation only confers the ability to utilize citrate when it happens in conjunction with these precise mutations.

So, this mutation itself was not lucky. But one of these lines needed to have these prior mutations such that this mutation could happen and then subsequent utilization of citrate could occur. So these two are alternative hypotheses. And we'll discuss which one of these two hypotheses was found to be correct. And what would be the experiment that you do?

To check which of these two hypotheses was the actual one leading to citrate utilization. We will continue that in the next video. Thank you.