Evolutionary Dynamics

Supreet Saini

Chemical Engineering

Indian Institute of Technology Bombay

Week 10

Lecture 48

Hi everyone, welcome to the next video in the course. So, we are starting to discuss the third phase of the course, the third and final phase of the course, where we will be discussing experimental evolution, how to perform these experiments in the lab, and what are some of the lessons that we learned. As we started out last time, there are three ways of doing evolution experiments in the lab with microbial organisms. We discussed one, which is the most widely used technique called serial subcultures, and we discussed some of the aspects and properties related to it. If you study microbial evolution experiments in the lab, well over 90 percent will be conducted in the way that we discussed in the previous video.

However, there are two more ways of doing evolution experiments in the lab. In this video, we will discuss these two other ways, and then we will move to experiments—some experiments that I think are interesting, valuable, and some of them throw surprises at us. They open a window as to how we can see evolution taking place in real time right before our eyes. We will start with that in the next video onwards. So, just to recap, the first method of doing evolution experiments was serial subculture. The advantage—why it is so widely used—is, first of all, the logistical ease of doing experiments this way is very easy. You can have hundreds of parallel test tubes going on. So, if you want to study questions such as the repeatability of evolution, you can have 100 populations evolving in the same identical environment and see what happens.

That's sort of the major advantage that we saw. However, the main disadvantage with serial subculture was that during the phase of growth, the environment is constantly changing. We sketched this graph where the x-axis is time, the y-axis is the number of bacteria, and the second axis is, let us say, the amount of glucose left in the flask. And we saw that the problem here is that if we represent the number of bacteria with green, we start with N naught, and this number increases to some NT.

And then we subculture again and bring the number back to N naught. But if we look at glucose resources—the amount of resource left in the environment—we start with some concentration of glucose resource, which is given by this number. And as growth takes place, this number goes down to zero. All the glucose is used up. So what the population is responding to during its growth phase is very high amounts of glucose because there are not many consumers there.

The demand for glucose is very low because the number of individuals in the culture is very small, but the availability of the resource is huge. Then we enter a window where demand and supply are about the right size. There are a decent number of individuals. There is still lots of glucose left, and every individual is about to grow exponentially. But then we enter into this phase where the demand for glucose is huge, but the supply is now barely there.

So as a result of that, in every cycle of this growth period that we saw—if we look at the number of bacteria over time— In every cycle, we saw that this was going on. This was the growth phase. Then we do maybe a 1:100 dilution, and then we do this again, and then we do this again. And we have these boom-and-bust type of experiments where the populations go bust.

We apply a huge dilution factor, and then they increase in number through the 24-hour period of growth. It could be 24, it could be 12, and so on and so forth. But the point is that during this window of growth, selection changes. So while we say in this experiment that we are growing, we are studying the adaptation of glucose in a low environment, we are studying the adaptation of the E. coli population in an environment where the amount of glucose is low—while we may be saying that, In reality, this selection pressure is changing with time.

Early on, glucose isn't really limiting, and very late on, it's highly limiting, and so on and so forth. So, selection is changing with time, and it's to this changing selection that we are studying the population's response in an experiment such as this. And to counter this problem with serial subculture, we have the second experimental technique, which is using chemostats. A chemostat is something that we have discussed a lot during the course. If you think about it, no environment is truly batch.

If you think of any ecological niche—whether it's a pond, whether it's our intestines, whether it's soil—no environment, no ecological niche out there is a truly batch culture, which is just closed from the rest of the world in terms of exchanging material with it. If

it's our intestines we are talking about, we continuously drink liquids and consume food. And these act as nutrients. These are sterile things that we are eating, and these act as nutrients to propel further growth. But we are also shedding bacteria out every time we go to the restroom or something.

And then, so there is an in and an out stream. If you think of bacteria in a soil environment, there is a constant exchange of matter with the environment. Rain brings in liquid. Trees have their leaves falling; as a result, organic matter is brought into the soil, bacteria die, newer species are brought in by migration, wind blows, and so on and so forth. So, every environment exchanges energy and matter with its surroundings; every niche does that, and hence, because of this exchange process, every niche can be thought of as a chemostat in some sense.

That exchange, that inlet and outlet stream we have been drawing in a chemostat, is constant V0 entry and constant V0 exit. Those entities might not be very strict; they might not be constant with time, but these inlet and outlet streams are part of every ecological niche. So, a chemostat is a good approximation of any niche we might consider. So, this is the inlet stream. This is the outlet stream.

And we do this experiment. So, this is V naught. This is rate. So, this is volume per time. What is the rate at which fresh liquid comes in and spent liquid goes out?

And this is capital V, which is the volume of the chemostat or the reactor. So let us look at this for a little bit. Chemostat. Because the environment is not changing with time, this takes care of the major objection or the major limitation of serial subculture studies. Now, this is constant with time because we have this very slow inflow of V naught and a very slow rate V naught at which spent media is taken out.

And this is well stirred. The environment inside is not changing with time. Because this is operating at a steady state. Hence, this sort of takes care of the critique of a changing environment that was there in the serial subculture technique. It is just a constant environment.

Hence, this takes care of that problem. However, the problem with the chemostat is more logistical. Because this is a continuously operating system, there needs to be a tank with fresh media from where this liquid is coming. This, of course, is waste. It can go anywhere.

We will sample some and perhaps discard the other bit. This might be for phenotypic characterization. For sequencing, all this will happen, but the problem is—and this has obviously many moving parts—there is a motor involved here which is doing this rotation throughout this process. As we will see when we discuss these evolution experiments, these experiments can go on for years. We saw that with a 1:1000 dilution, we process about 3.5 thousand generations in one year. That's about the length of a standard evolution experiment reported these days.

So these moving parts—this fresh media tank—has to be constantly replenished. It has to be kept sterile. The duration for which this has to be done is typically years in the case of microbial experiments. This often proves logistically very challenging—how to maintain all of this in sterile conditions for such an extended period of time. Secondly, serial subculturing offered us this easy extension of the experiment: instead of just doing transfers in one tube, I can set up 100 such tubes, and each one of these 100 lines is processed in parallel.

So, I can have 100 such lines, and at the end of the experiment—let us say this goes on for 10,000 generations— At the end of 10,000 generations, I can sequence every one of these 100 lines, each one of which evolved under identical environments. And I can ask questions like: Are the mutations identified in each one in the same gene, or are they the same mutations? To what extent is evolution convergent? Which means that, given an environment, these are the mutations that will take place.

And to what extent is it divergent that one line could pick up a mutation in one pathway and another line could pick up a mutation in a completely different pathway, and so on and so forth. So, I can ask questions like this when I'm doing experiments via serial subculture. It's hard here because, again, this is a logistical limitation of a chemostat because maintaining this sterile environment for a few years for one chemostat is hard enough, but doing it for 100 is much harder still. And hence, this is not a very popularly used experimental technique. One thing about the chemostat operation that we should be mindful of is considering the relative magnitudes of the two quantities we have: capital V and small v0.

Capital V is the volume of the reactor or chemostat. And small v0 is the volumetric flow rate in or out. So, this is volume per time, and this is volume. So, if we ask, what is the residence time? So, suppose

This piece of this element of liquid that is entering the chemostat at this instant—how much time does it spend here? Before, suppose I passed a small molecule which is green-colored; it's not going to be used by bacteria. It's just going to enter and go out of the chemostat. It's obviously going to spend some time. It will go around.

It will get shaken around and eventually it will go out. what is the time that it spends in this reactor and the average of that time can be estimated as so average time in chemostat is just equal to total volume divided by volumetric flow rate if this is 100 liters And I am adding 10 liters per minute. That means on an average, every liquid element spends 10 minutes before it is washed out. If this was 20 liters, then average time that it spends will be 100 upon 20, which is 5 minutes.

So, the average time that is spent in a chemostat is this. What that means is that what happens to a chemostat operation if I increase V naught? If I increase V naught too much, then what's going to happen? Imagine this chemostat, which is 100 liters, but my V naught is 20 liters per minute, which is quite substantial. That means liquid is just rushing in and rushing out.

That means every element only spends 5 minutes in the chemostat. That means I am washing away bacteria at this high rate that in 5 minutes this chemostat will be emptied if this inlet stream was stopped for some reason. In 5 minutes it will all be gone. Also with V0 I am removing certain bacteria. And in a steady state operation, those number of bacteria that are being removed have to be replaced by growth taking place inside the chemostat.

You have this exit V0 stream removing bacteria, but you have growth taking place inside the chemostat, which is adding back the numbers that have been flushed out from the system, such that the total number of bacteria in the chemostat or the population size remains constant. Suppose in this V0 flow rate at the exit, I am removing N0 number of bacteria per unit time along with V0. then N0 cell divisions must take place inside the chemostat per unit time so that the total number remains the same. So, N0 divisions per time must take place if N0 is the number of bacteria that are being flushed out of the system every single unit of time and only then will the numbers stay constant. Now, if that is the case, imagine a scenario where I am increasing V naught.

If V naught is increased beyond a certain limit, then what you are asking is that you are increasing the load. If V naught is increasing, N naught is increasing. And if N naught is increasing, that means you are saying that in per unit amount of time, these N naught

divisions must happen in the chemostat to keep the numbers same. But cells cannot divide infinitely fast. They have a finite division rate.

And as a result of this, this number of divisions actually taking place in the chemostat will be less than N naught. Because as we have seen that E. coli, even when everything is provided to E. coli, it still needs about 20 minutes to divide. But if my removal rate is so high, so that the numbers can only be kept constant if E. coli can divide in 10 minutes, then that is not going to be possible because E. coli cannot divide in 10 minutes. As a result of this, I am removing more, but cell division is not adding the same amount.

As a result, the number inside the chemostat drops. The population size in the chemostat keeps dropping. And if this exit flow stream is very high, obviously this V naught is the same as the exit flow stream. But if the V naught number is very high, then eventually all the bacteria will be flushed away from the system, and I will just have sterile liquid left in it. So, this ratio V by V naught has to be kept into account whenever we are operating a chemostat.

And that—that should time permit—we will discuss one interesting study from a chemostat evolution experiment in the lab. But for the most part, I don't think we will be revisiting this. The third way of doing evolution experiments is somewhat less frequently used, but it's a technique that is used to ask very specific types of questions. And this technique is called mutation accumulation. The main idea behind mutation accumulation is the following: when I do an experiment by serial dilution, what is going to happen?

I have a test tube, and in this test tube, growth is going to take place. When that growth happens, some variations are going to arise. And let us say the green variation is good—it increases growth rate in the environment in which I am doing this experiment—and the red variation is bad in that it reduces growth rate. So, when I subculture this in a serial subculture, when I do this subculture in the next test tube, the green guys—the progeny of the green individual—are going to outcompete. And I will be left with something like this.

And if I subculture one more time, it will perhaps be all green individuals, not red or green, as the black has been completely eliminated. In fact, what might happen is that a new great mutation, indicated by blue hair, has arisen in the population significantly. And as we go forward in time, the blue individuals will outcompete the green individuals. And this process will keep repeating itself. Newer mutations will keep coming.

But what is to note here is that in an experiment like this, mutations happening in the test tube are of three kinds. Neutral, beneficial, and deleterious. Beneficial mutations are fine because they are being selected, and those are the mutations we see as an output of the experiment. If I sequence the bacteria present in this tube, then what I will see is the green mutation. If I let this experiment go on a little longer, the blue individuals will start to outcompete the green individuals.

And if I sequence the bacteria at that point, I will identify the blue mutations. But I will never identify the deleterious mutation that happened early on. Maybe there was a neutral mutation that also happened here. Let us say this was neutral and this was a deleterious mutation. I never get a sense of how many deleterious mutations happened, how many neutral mutations happened, or how many beneficial mutations happened.

These are not the kind of answers that I get from either serial subculture or a chemostat experiment. Suppose the type of information that I wanted from these experiments was as follows: I do not want to only find out about beneficial mutations. What I want to find out through this experiment is the total spectrum of mutations that are happening. I want to find out—no. So let me write this down.

So my interest is not in adaptive mutations. Those that increase fitness, whether it is serial subculture or a chemostat. They only give me information about these adaptive mutations, which are actually able to beat drift, get established, and reach fixation. Suppose my interest is not in those mutations. Instead, my interest is in figuring out

the complete spectrum of mutations. I want to find out How many—what percentage of mutations are neutral mutations? What percentage of mutations are beneficial, and what percentage are deleterious? What is the distribution of fitness effects associated with each of these types of mutations?

Suppose that's what I wanted to find out. Then these approaches will not work, and I need a new way to do an experiment. And the key requirement of doing that experiment is that, in order to answer this question, I must design an experimental technique where selection does not act. Because if I let selection act, then selection is going to get rid of deleterious mutations. Selection is going to make it really hard for a neutral mutation to reach fixation.

Remember, the fixation probability of a neutral mutation was 1 upon n. In a large population size, that is really hard. So, I want to get rid of selection in an experiment.

And a mutation accumulation experiment helps us accomplish that. In such an experiment, evolution takes place only by drift. And I'll quickly explain how that takes place. Imagine you have a population of bacteria.

Let's say you have different types of individuals here. Let's do one more. I don't want to give selection a chance so that these five will compete between themselves and the more fit among them will win. So, what I will do is take this and spread them on a solid plate. So, this is a petri plate.

And I will spread them on this plate, not knowing where they went because they are too small to be seen, but they are somewhere. Then I will let this plate grow at 37 degrees Celsius. Let's say this is E. coli overnight. What will happen is that these five individuals, wherever they landed up, because growth conditions are great, sufficient time is there, temperature is there, nutrients are there on the plate, they will start dividing. And when I come in the next morning, I will see these five little mounds.

Each is full of bacteria. So, I will see this when I come in the following morning. Wherever those five bacteria were deposited, enough cell divisions would have taken place that starting from that one individual, this number has reached of the order of 10 to the power 8. And obviously, in these cell divisions, mutations have already taken place because cell division occurs with a certain error rate in the DNA that is given to the progeny. Then, because I don't want selection to take place, I pick one of these colonies completely randomly.

I randomize the process and maybe I will pick this one. This colony has 10 to the power 8 cells. I will pick any 5 out of them randomly and plate them on a fresh plate. And I repeat, this plate will then again be allowed to grow at 37 degrees Celsius overnight. And these individuals will grow up to become these mounds, which are called colonies.

And then I will again pick another colony. So every aspect of this experiment is randomized, and there is no point at which selection is acting. The choice of this individual, the choice of these five individuals, was random. Then, when I come in the following morning and select one of these five colonies, that choice of which one of these five colonies is random. This colony has 10 to the power of 8 cells.

Out of these 10 to the power of 8 cells, I am picking 5 randomly and spreading them on a plate, not knowing where they are going. And the following morning when I come, these 5 will have again grown into 5 colonies, and I will pick one of them again randomly. So,

every aspect of this growth experiment is random. As I keep doing a sufficient number of these transfers from one plate to another, what will happen is that sooner or later, I will start selecting cells that are now carrying mutations. And these are mutations that are being selected only on the basis of chance events.

Selection does not have anything to do with it. And if I do this experiment for a sufficiently long time, this individual that I'm selecting will have accumulated a lot of mutations in its genome. And I will sequence those mutations and then individually characterize each one of them, which will help me answer questions like: How many mutations that happened were neutral? How many were beneficial? How many were deleterious?

Again, the point is that I allowed these mutations to accumulate in a genome without selection acting, because if selection had acted, it would have gotten rid of bad mutations and only selected good ones. I don't want that. I want to let all sorts of mutations take place. And that's what happens in this experiment. And typically, this is done for maybe a few hundred transfers.

After a transfer of a few hundred times on different plates, I will get many mutations, and that gives me an answer regarding not only what the nature of mutations is, but also what the mutation rates are associated with each species in a particular environment. So, again, the key for us to note here is that in a mutation accumulation experiment, I let mutations accumulate, but without the action of selection. Mutations accumulate only under drift, and that allows me to answer questions related to the nature of mutations. That concludes our discussion of how three types of evolution experiments can be done in the lab.

Starting with the next video, we'll begin our first discussion of the first evolution experiment that we'll cover. We will continue with that in the next video. Thank you.