“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

Brad Power
August 3, 2022

“I don't think we know how to build these (mechanistic simulation) models (that would predict drug response or resistance for individual patients). There's too much uncertainty in the models themselves to be clinically informative at this point. That's both from a technical perspective, because the models need to be big, and we just don't have a lot of the formalisms and computational tools to do it.” – Marc Birtwistle

Meeting Summary

Marc Birtwistle, PhD, Associate Professor of Chemical and Biomolecular Engineering and Bioengineering, and Alex Feltus, PhD, Professor, Department of Genetics and Biochemistry, Clemson, led a discussion on "Basic Research into Simulation Models that Could Eventually Guide Clinical Decisions".

What are the complex decisions faced by advanced cancer patients that simulation models might help?

> There is much room for improvement in making treatment decisions for advanced cancer patients. For example, although genomically-targeted therapies work for some people that have a mutation, it doesn't always work for everybody that has the mutation. A treatment can also eventually fail due to development of resistance. Personalized drug combinations can offer better outcomes, but there is no evidence for most of the many potential combinations from randomized clinical trials. If we had a good tumor simulation model, we could prioritize what types of drugs might be useful for a given patient, or we could even start talking about what types of dosing or scheduling might be better than others.

What are the challenges in developing simulation models to describe cancer dynamics?

- **Dynamic**: Drugs in pharmacology are dynamic. The tumors adapt on multiple time scales. The time of day when drugs are administered matters. Dosing matters. Probably the simplest dynamic we can think about is when you treat a single cell with a drug – it is usually going to have some sort of a stress response. It's going to change the genes that it's expressing to try to help it deal with the fact that now you're trying to kill it. It does things like upregulate pumps that help it to pump the drug out of the cell. These are very well established mechanisms, and those are things that can really affect drug response, so are important to capture. The aspect of the dynamics that may be arguably the most important one to try to get a handle on is when we're thinking about what drugs we start with. Then as the subclonal makeup of that tumor changes, then what do we do? Do we attack the dominant subclone first? And then the ones that are initially in a lower proportion and maybe more aggressive, or do we take out those other low proportion ones first?

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

- **Heterogeneity**: Across every axis that you look at in cancer, there's heterogeneity. If you look across patients, it’s not just that a patient has prostate cancer, each patient’s tumor is unique. If you look within one patient's tumor, all of the cells within that tumor can be different. You can have different genetic subclones within that tumor that might respond differently to drugs. And even within the same genetic subclone, there is heterogeneity due to other random processes that happen in the cells. If you look at that tumor in a spatial sense, there are different microenvironmental factors, different oxygen concentrations, different immune local environments that can control drug responses.
- **Multiple Pathways**: Gene expression isn’t linear, it’s more a network. Multiple pathways intersect to explain how the cancer evolves and behaves.

How do you build simulation models to describe cancer dynamics?

Simulation models can be empirical (based on observations of experience, per the scientific method) or mechanistic (based on a theory of how the system is structured and works). Mechanistic models are preferable because they can predict, fill in blanks, and are interpretable. Empirical models depend on large, clean datasets to infer patterns. Biochemistry provides biochemical models which can be built upon.

When will simulation models be ready for clinical use?

Simulation models are in the world of research and basic science. There's too much uncertainty in the models to be clinically informative at this point. That's both from a technical perspective, because the models need to be big, and we just don't have a lot of the formalisms and computational tools to do it.

Alex Feltus: *“Marc’s stuff is probably years away from being truly translational. But I think Marc’s stuff is the stuff that's going to change everything.”*

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

Meeting Notes

Brad Power: This is going to be a discussion about simulations and how they might in the future guide cancer treatment decisions. I'm pleased we have Marc Bertwhistle and Alex Feldus to lead this discussion. I got to know Alex through Pete Kane and through Bill Paseman, for whom Alex was running a hackathon. Bill Paseman has a rare kidney cancer. Alex introduced us to Marc.
We're interested in learning about how simulations might help guide treatment decisions.
We have identified that personalized solutions for patients are often off label, particularly for drug combinations. There are not enough randomized clinical trials that say, “this drug combination is indicated for this use, or this drug combination is inside the standard of care.” How could you have confidence in making such a prescription if you were the treating physician? How would you feel good about something that's not obvious because there's good clinical evidence to support it? At least some confidence might come from understanding simulations or models of cancer and its dynamics, and how it progresses, and how it's impacted, and putting everything together into a simulation model.

Marc Birtwistle: I'm excited and nervous. I've never talked to a forum like this before. I was struggling a little bit on how to talk about the work that we do in my lab and how to present it. Hopefully it's useful and informative for everybody here.

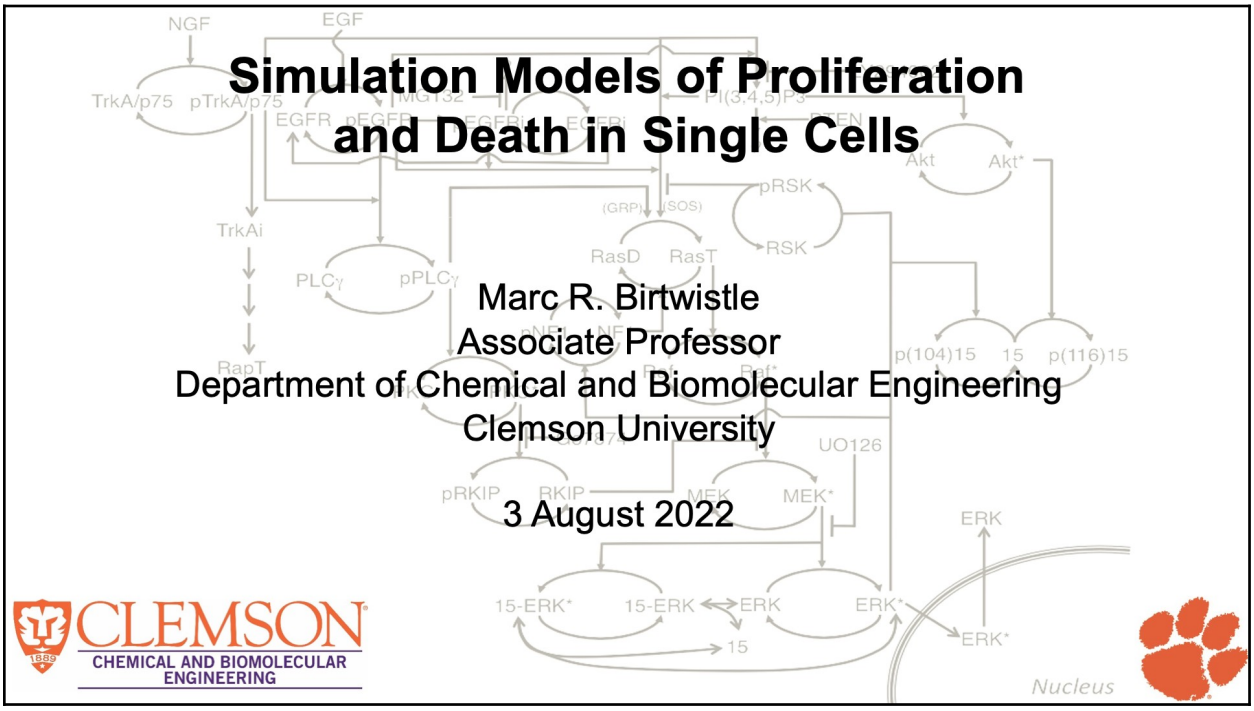

I am an associate professor here in the chemical and biomolecular engineering department at Clemson. I'm a chemical engineer by training, but I've been working in systems biology and signal transduction and cancer systems biology pretty much for my entire research career. An

engineer with that sort of thinking, but also blending that with cancer biology and trying to do research.

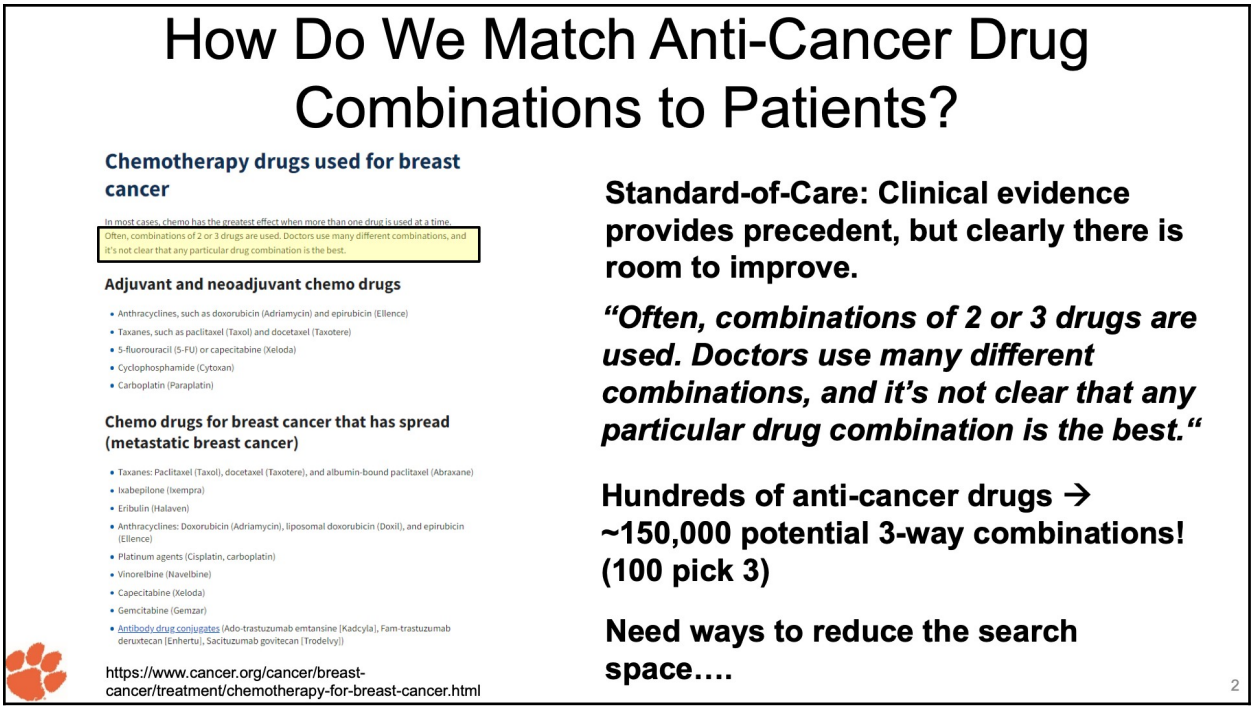

One of the big questions that the whole field is interested in is how we match anticancer drug combinations to patients. I heard Brad talk about that in a way that really resonated with me – it's really not a very well established or known thing. Some quick Google searching will give you a lot of information about what drugs could be used, or sometimes are used, for particular types of cancer, for pretty much any cancer type. I put some information here on the left for breast cancer. You can see dozens of them. Some are traditional chemotherapies, some are more targeted chemotherapies. The reason why these are used is because there is some clinical evidence that they are effective. But there's obviously still room to improve because cancer is still quite a deadly disease for many, many people. I wanted to highlight a pretty common knowledge, but combinations of two or three or sometimes even more drugs are used. Doctors use many different combinations, and it's not clear that any particular drug combination is the best. How do we actually make any traction on that problem?

One of the big roadblocks in the way of that, is to get solid numbers on how many anticancer drugs there are that are FDA approved right now. It was a hard number to pinpoint. There are at least hundreds of them that are approved and maybe sometimes drugs used off label that aren't for anticancer indications, but could be used by physicians.

Let's just say that you had a hundred anticancer drugs, and you wanted to figure out different three-way combinations. A little bit of math will tell you there's over 150,000 different three-way combinations. We need some way to reduce that search space to answer this question.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

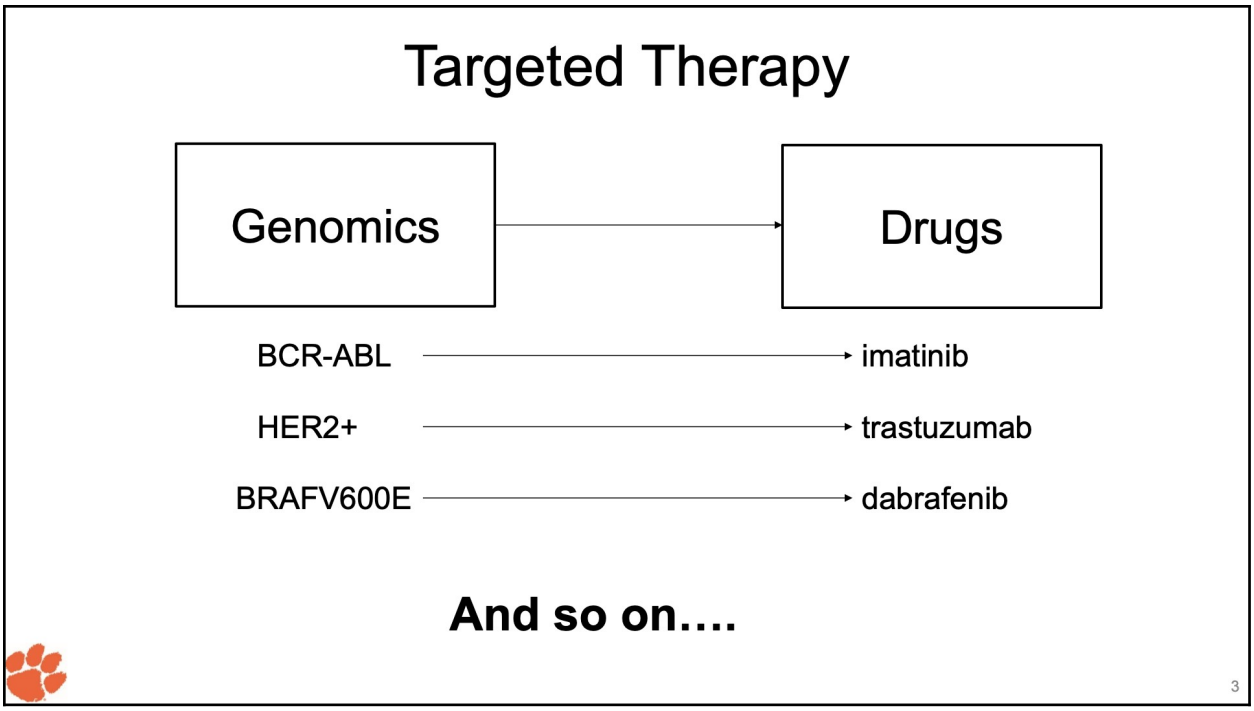

One of the ways to do that currently is what we call targeted therapy. The idea is that if we know something about the genomics of somebody's tumor – the mutations that are driving the tumor's behavior – maybe we can directly match a drug to that mutation. There are lots and lots of examples of this that have been developed over the past several decades. One of the original ones was imatinib for a BCR-ABL fusion protein mutation. If you have a HER2 positive breast cancer, there are several drugs available, one of which is a monoclonal antibody called trastuzumab or Herceptin. If you have a melanoma and you have a particular point mutation in BRAF called BRAF V600E, there are multiple small molecule kinase inhibitors that are available, one of which is called dabrafenib. You can go on and on with the examples of targeted therapies that have been developed. They have done a good job, but there's still room for improvement. I'd like to highlight a couple of reasons why there's room for improvement, and what some of that improvement might look like.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

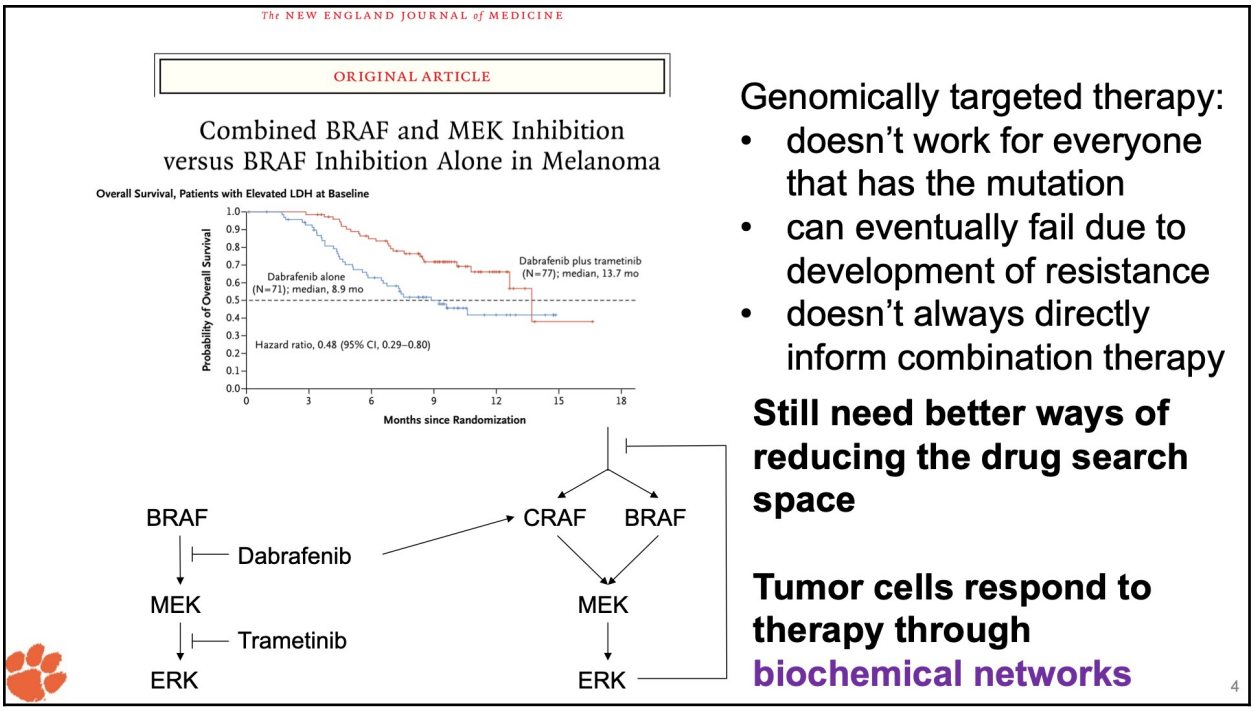

This paper shows this Kaplan-Meier survival plot (depicting survival time) if you combine two different inhibitors to treat melanoma, one is inhibiting BRAF in patients who have the BRAF mutation with the drug dabrafenib that I just showed, and then combining it with a drug called trametinib, a MEK inhibitor. These proteins essentially act in a signaling pathway.
A signaling pathway is a system of proteins that work together in cells to send signals that eventually control what a cell does. Like, is it going to divide? Is it going to die? Is it going to move? Is it going to grow? Things like that. BRAF activates MEK, and MEK activates this protein ERK, which ends up doing a lot of that signal transmission in the cell to tell a cell to grow or not grow. By inhibiting these pathways, you can turn off a signal, but why would you turn off a signal and then turn off a signal that's directly downstream of it? It really doesn't make much sense genetically that a drug combination like this would be effective, but clinically it was shown to be quite effective because patients survived longer.

What's really going on there?

What's going on in a cell is not so simple as a linear pathway. What's been shown is that some of these BRAF inhibitors can activate a different isoform of the protein called CRAF, and then CRAF can activate the pathway in a parallel manner. There are also things like feedback loops inside signal transduction systems like engineered systems. We have negative feedback loops to help us suppress the effective noise.

This RAF-MEK-ERK pathway has such negative feedback loops that help to fight the effects of something like a drug to keep that pathway on.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

Although genomically targeted therapies work for some people that have a mutation, it doesn't always work for everybody that has the mutation. It can also eventually fail due to development of resistance. Both of these points can be seen from any Kaplan-Meier survival plot, like this one here, or other ones that highlight the effectiveness of targeted therapies.

My other point is that it doesn't always directly inform combination therapy because of this linear pathway thinking. Why would a drug combination targeting the same pathway actually be effective?

There's a lot more that we need to consider to understand those sorts of things and to be able to predict them. I think targeted therapy helps us to reduce that search space, but we still need better ways. And one of the ways that we should be focusing on is really understanding more how biochemical networks drive response to therapy.

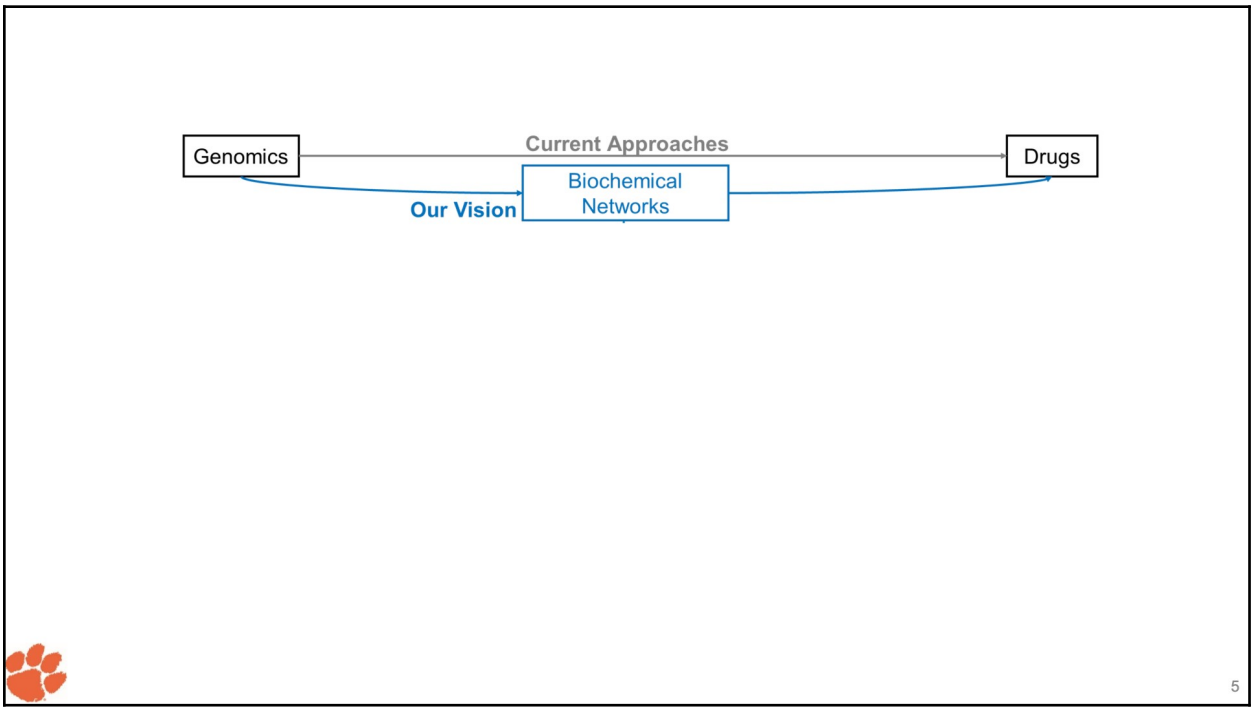

One of the ways I think about this problem in my lab is to use genomics as a foundation because we know genetics really drives the behavior of most tumors. But instead run that information through some understanding of a biochemical network, and use that to better inform what kind of drugs might be useful or effective for a particular patient. In doing so, I like to talk a little bit about models, and the way that I think about models.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

Let’s Talk About Models

“A **model** is an informative representation of an object, person or system”

Mouse Models of Cancer

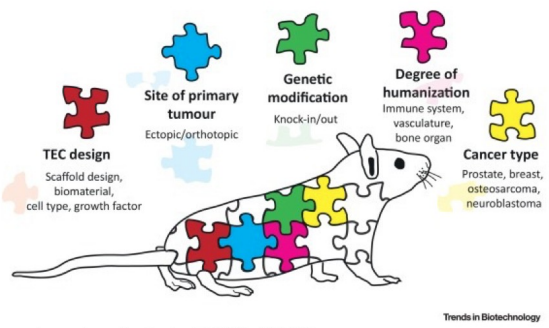

Landgraf et al., TIBS, 2018

Cell Line Models of Cancer

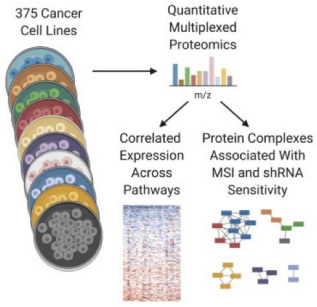

Nusinow et al., Cell, 2020

The large combinatorial anti-cancer drug space presents practical roadblocks for experimental models

6

Usually I'm talking to an audience that doesn't work with mathematical models or simulation models very much. I like to start out by talking about experimental models, two of which that are very widely used in cancer research are mouse models and cell line models.
We can learn a lot about the basic biology of cancer or how particular types of cancer could be sensitive or resistant to different drugs or drug combinations by using experimental models that have some recapitulation of the eventual human cancer that we care about. One of the problems though with experimental models is that the large combinatorial drug space, if we're talking about, as I mentioned before, two- and three-, maybe even more, drug combinations, and the sheer number of anticancer drugs, we just can't do that many experiments to take a brute force approach of understanding that mapping from genetics to drug combinations.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

Let’s Talk About Simulation Models

https://www.frasca.com/wp-content/uploads/2018/10/Cessna-C172-Simulator.jpg

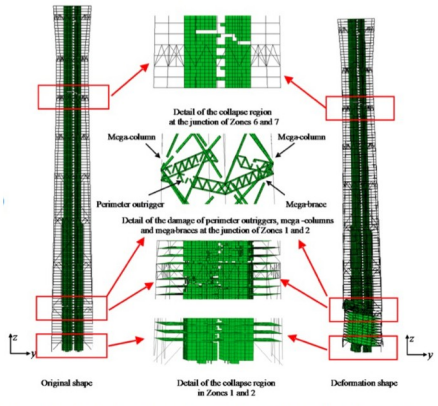

The earthquake-induced collapse of a super-tall mega-braced frame-core tube building (H=550m) is simulated with proposed finite element model based on the fiber-beam and multi-layer shell models.

https://www.researchgate.net/figure/The-earthquake-induced-collapse-of-a-super-tall-mega-braced-frame-core-tube-building_fig1_237047370

https://mars.nasa.gov/imgs/mars2020/mars2020-sky-crane.jpg

Simulation models facilitate design choices across engineering.

7

That's where this idea of simulation models comes in. It’s an experimental model that is just a more convenient way of looking at or analyzing a system that we care about.
Those simulation models are in principle the same, and they're very much used across engineering. I highlighted a couple of examples here: a flight simulator – if people want to understand how changes to an aircraft might show up; if we want to try to understand how to design earthquake-proof buildings. We use simulation models to try to help us with that. Or if you want to land something on Mars, it really helps to have good simulation models, because that's really hard to test before you get there.

My point is that simulation models facilitate the design choices across engineering. Maybe this is something that we can start to try to develop in the context of oncology.

It wouldn't be easy, but we can think about it, and that's where my lab tries to sit much more in a basic science sense at this point. How do we actually build those models, and what should we

build into those models to have it be “good”? What if we had it though? What could we do?

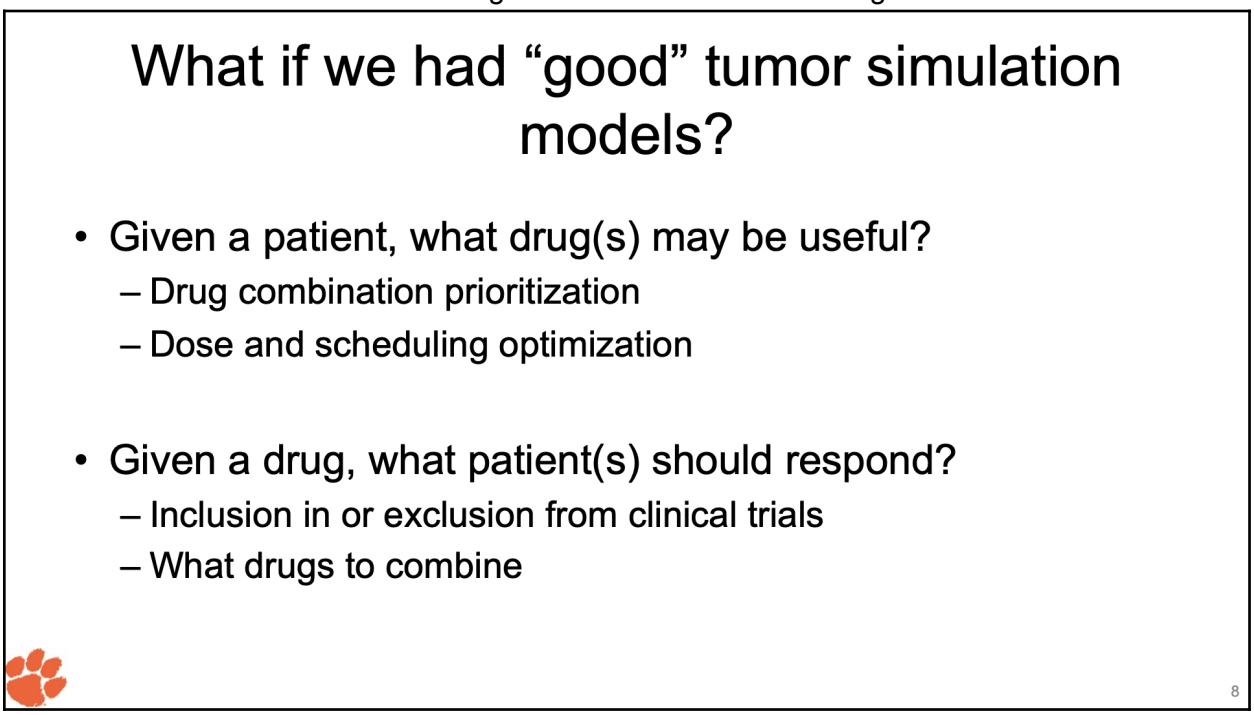

If we had a good tumor simulation model, we could start to answer two kinds of questions in my vision: (1) given a patient we could prioritize what types of drugs might be useful, or we could even start talking about what types of dosing or scheduling might be better than others, or (2) if we're taking more from the perspective of drug development, if we're given a drug, what sort of patients should or should not respond? We could start to make more informed decisions about inclusion or exclusion from clinical trials, or given a drug we're trying to develop, what are some other drugs that might be good or not good to combine with that drug?

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

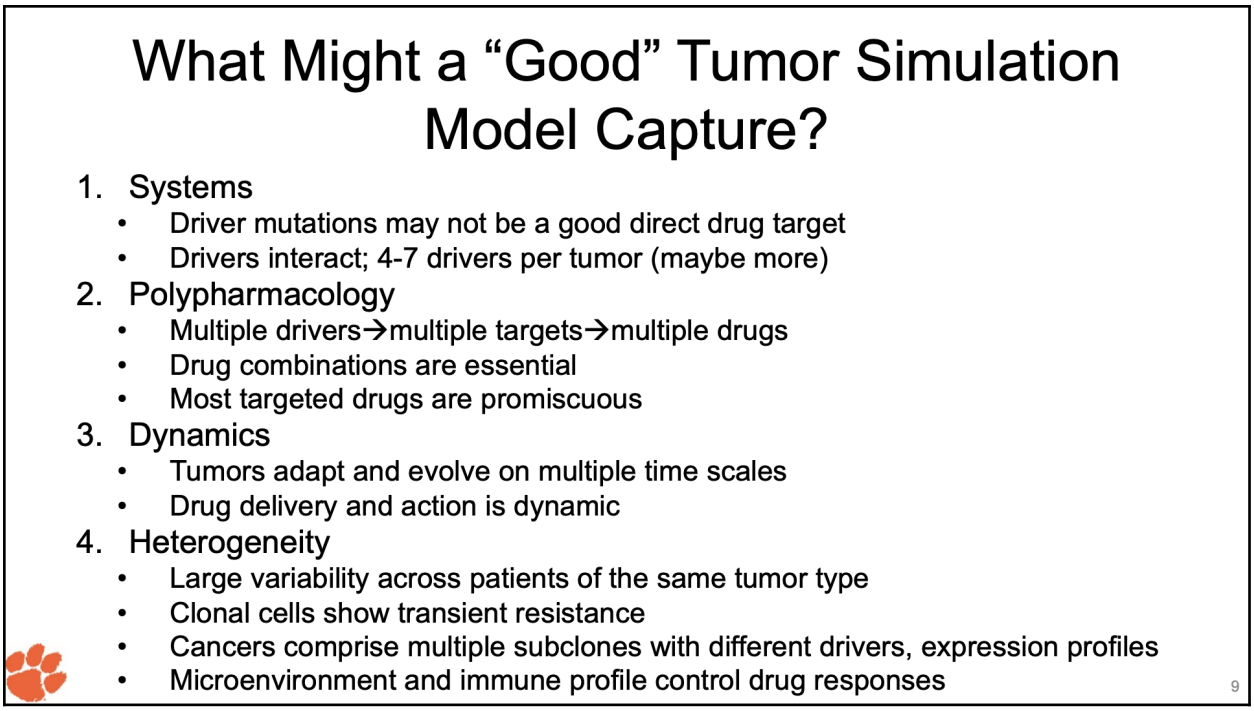

What do we need to capture in such a model? There are lots of aspects of cancer biology that we know are important, but are just hard to get a handle on in terms of building a simulation model. There is also lots of just incomplete knowledge, and it makes it really challenging. But we know there are several important aspects that tend to drive response to therapy, and I listed four of them here.

Some of the examples I highlighted before highlight the fact that we need to think about the systems that are actually driving drug responses inside of cells, not just as genetic mutations, but what is the system where this genetic mutation lives, and how does that drive that cellular decision to grow or divide or die?

Most tumors aren't driven by one mutation; there are multiple mutations. I've seen estimates of around four to seven different mutations, and the fact that we have different driver mutations probably means that we need multiple drugs to hit them.

We also need to consider that these are dynamic systems. Drugs in pharmacology in general are dynamic. The tumors adapt. They evolve on multiple time scales.

One of the most challenging aspects of the disease that has become much more highlighted and understood in the past 10 years or so is the heterogeneity of the disease in every direction. Across every axis that you look at in cancer, there's heterogeneity. If you look across patients, not just every patient that has prostate cancer, their tumor is unique. It's not quite the same, even though they might have some shared features. If you look within one patient's tumor, all of the cells within that tumor can be different. You can have different genetic subclones within that tumor that might respond differently to drugs. And even within the same genetic subclone, there

is heterogeneity due to other random processes that happen in the cells. And if you then look at that tumor more in a spatial sense, there are different microenvironmental factors, different oxygen concentrations, different immune local environments that can control drug responses. There are a lot of things that a good tumor simulation model should capture. When you're building a simulation model, there are lots of different options of formalisms and ways to go about it.

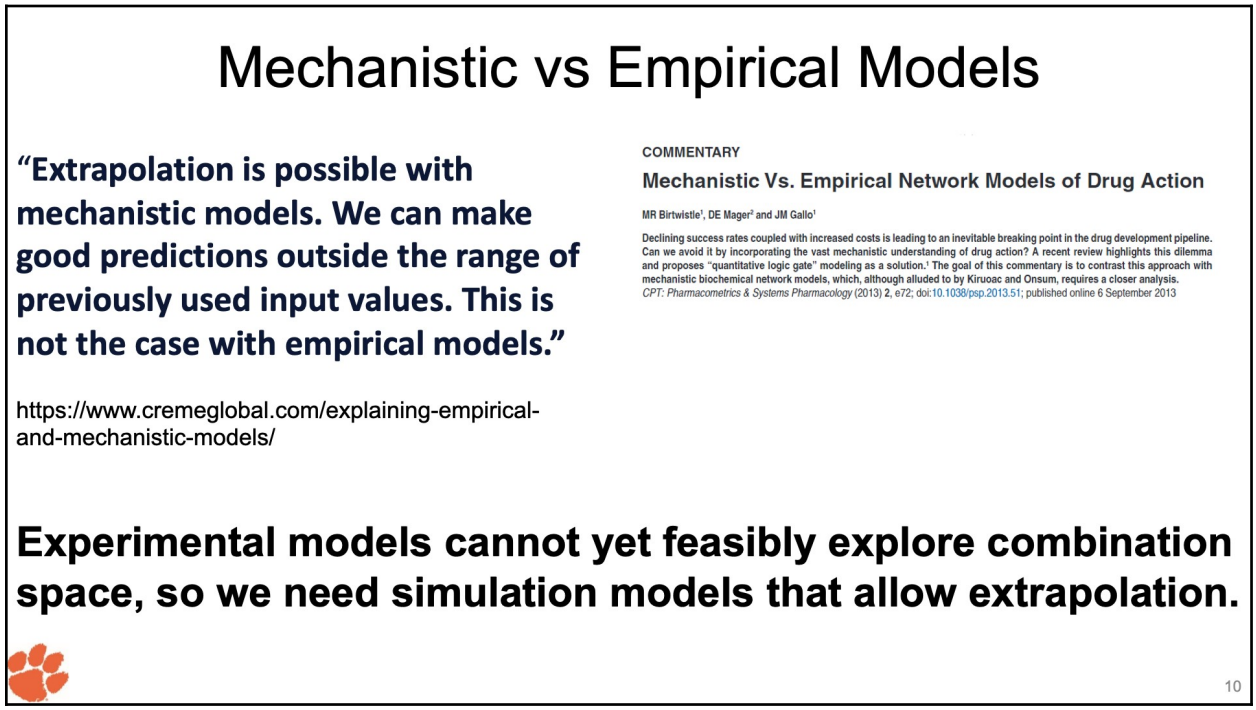

Two very broad classes of models are mechanistic or empirical models. When you formulate a simulation model in a mechanistic way, inherently you have some ability to do what's called extrapolation. That is making some prediction outside of the data you use to build that model. That is a really important feature of any good tumor simulation model because we can't get experimental data on all the drug combinations that we want to predict.
We need some model that inherently has extrapolability. We published a commentary on this 10 years ago, but it's a pretty widely accepted idea in modeling that there is this difference between mechanistic and empirical models. That's not to bash empirical models that are much more statistical or based on data crunching. They definitely are useful, but for this particular application, I think mechanistic models are important to consider.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

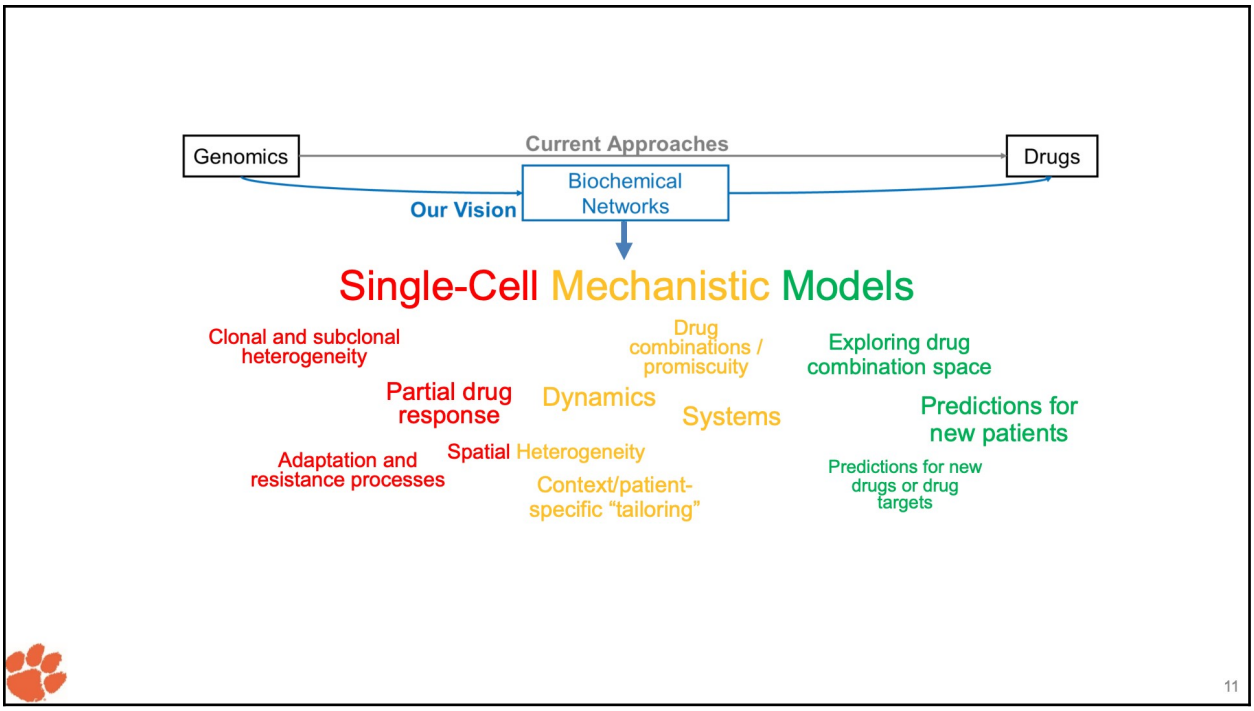

In my lab we try to build mechanistic models of how single cancer cells would respond to drugs. We try to do that by building in what we understand about these biochemical networks that control how those cells respond to drugs. I tried to color code what I believe to be the features of what I thought should be included eventually in a good tumor simulation model, and why these different ways that we build models could help contribute to that.

Anonymous Caregiver: The thing that really intrigued me most about everything that you said was the idea of dynamics and what's changing. How should we be thinking differently about the changes that are happening in real time? We talk about a lot of static decisions or static tests and diagnostics we can do. I would like to hear your highest level thinking about how we should think moving forward about the fact that we have a lot of things moving and changing every minute.

Marc Birtwistle: There are multiple mechanisms in multiple time scales. Let me try to parse those out a little bit, and then go from there. Probably the simplest dynamic we can think about is when you treat a single cell with a drug. It is going to usually have some sort of a stress response. It's going to change the genes that it's expressing to try to help it deal with the fact that now you're trying to kill it. It does things like upregulate pumps that help it to pump the drug out of the cell. These are very well established mechanisms, and those are things that can really affect drug response, so are important to capture.

Another really important thing is that a tumor is composed of heterogeneous subclones. You can imagine by the time somebody gets a tumor that is clinically presented, it's been evolving in that person for a very long time, and not every cell in there has the same evolutionary history and the same set of mutations, because it had to change and survive just like a species.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

People refer to these subclones as different types of cells within that tumor that have different mutations because of that historical process, the way that it developed. If we were hitting the tumor with all the same drugs, some of those subclones are going to respond and some of them are not, and the ones that respond are not going to grow and/or die, and the ones that are resistant to that drug are eventually going to take over that tumor.

That aspect of the dynamics may be arguably the most important one to try to get a handle on when we're thinking about what drugs we start with. Then as the subclonal makeup of that tumor changes, then what do we do? Do we attack the dominant subclone first? And then the ones that are initially in a lower proportion and maybe more aggressive, or do we take out those other low proportion ones first?

Brad Power: You're describing in very similar words\ to what Bob Gatenby talked about in terms of evolutionary and game theory.

Saed Sayad: I have many years of experience in mechanistic and data-driven modeling.
I see many targets here, but I strongly suggest focusing on two questions, instead of many targets, because it’s very complex.

1. The bottleneck in the treatment of cancer is the way one patient responds to treatments and the another patient doesn't. Why?
2. What's the source of resistance? Why do we get resistance? Can we simulate and even answer this question?

Marc Birtwistle: That's still a mystery to me. Why do targeted therapies fail? Sometimes two patients come in that are both HER2 positive. You give them the targeted drug, and one responds, and one doesn't. I don't know what the most important factors are driving that, as you said. I don't think anybody really does. But it's interesting to focus on. Can we develop a simulation model that at least provides some explanations for why that might happen? I like the idea of focusing in that way.

Alex Feltus: Isn't variation in metabolism a key component of that? Are pharmaco aspects, for example, how your liver processes the drug, an explanation for some of that?

Saed Sayad: We have some information, but really we cannot say if I give this drug X to this patient, this is good, but don't give it to other patients right now. The only option we have is through an experimental approach: get the group of patients, get their genomics or mixed data, and see if we can find any model based on those data, which we can predict the probability of response to this drug based on the genetics map.
We don't know how to predict this based on a mechanistic model. We have some sort of data-derived model using -omics data, micro RNA, or protein array. We can predict some level of resistance. This guy is going to be more suitable for this one, but we don't have any mechanistic model.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

Marc Birtwistle: I have been reading some stuff lately, I think it was in breast cancer, talking about whether it's night versus day, it affects what those cells are doing. We know metabolically that's the case. I've also heard in the literature about chronotherapy. When you give the drugs in the day relative to the circadian clock can matter as well.

Saed Sayad: The biggest problem is when you have the cancer cell or any disease cell compared to the normal set, the changes are not just in a few genes but in the thousands of genes. Building a model to mimic this behavior almost needs a quantum machine.

Brad Power: It’s a very complex system with many complex factors in a network, like Marc was saying.

Rick Stanton: I'm an electrical engineer by training that shifted to Amgen. I did a lot of simulations trying to help us develop drugs and understand pathways. When I moved over to Human Longevity, we tried to assess the dominant mutations that were probably starting the cancer. You have a pyramid of accumulated mutations. We were assessing dominance by allele fraction. You may have more advanced thinking. That was back in 2017.
I am also curious about your platform. Engineers love this stuff. This is how we think, such as feedback loops. Could you please comment on any of the details? I used to simulate via MATLAB Simulink because it was easy and supported encapsulation.
I'm a stage four prostate cancer patient, and Brian and I and others in our Prostate Cancer Lab are facing therapeutic decisions. We have DNA and RNA seq data. That's the easiest to get. Spatial has been elusive for us so far.
Before we dive into all these details, at the end of the day, we're going to be wondering, is there any way we could work together? If our data is open on the web, is there anything you could do to help us?

Marc Birtwistle: That's one of the things where I was really nervous about talking in this forum. The work that I do is much more in the realm of basic science. I don't think we know how to build these models. There's too much uncertainty in the models themselves to be clinically informative at this point. That's both from a technical perspective, because the models need to be big, and we just don't have a lot of the formalisms and computational tools to do it.
We just don't understand all of the relevant mechanisms going on within cancer cells and how some of those mutations actually change some of the systems. That's my view now. I would love it if this kind of stuff could eventually do that. With work, it could be the case, but not right now.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

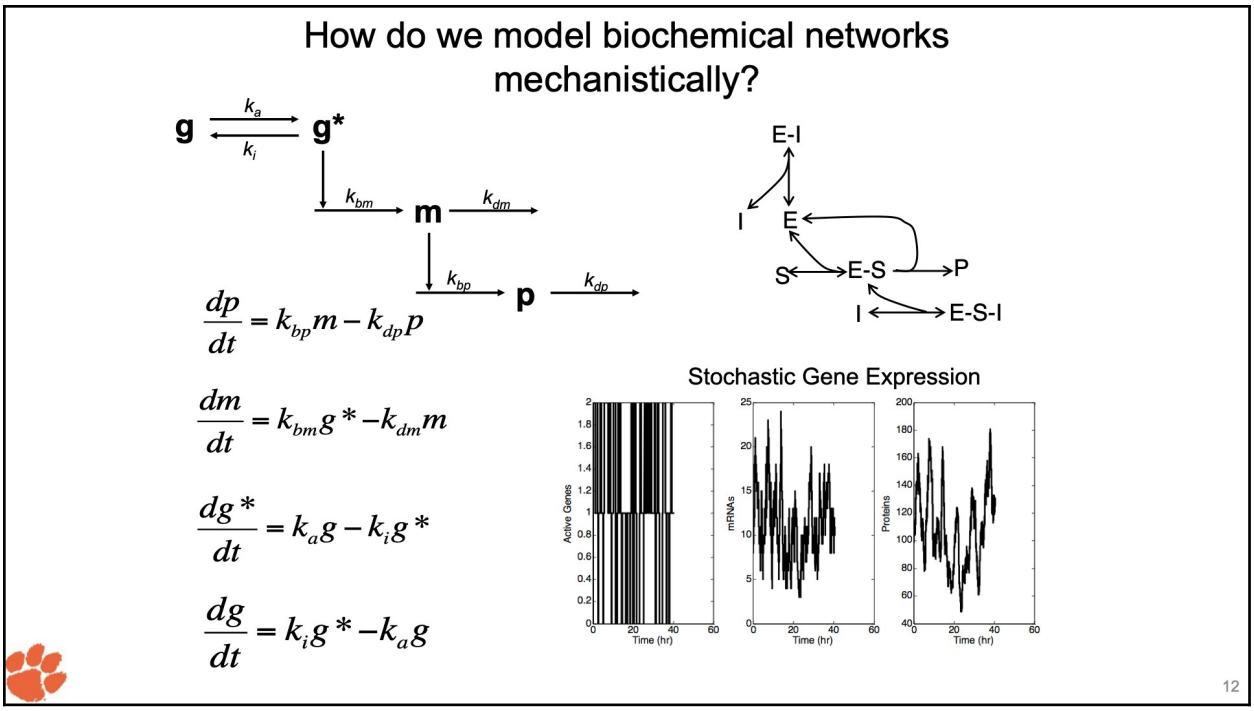

How we go about it: we essentially formulate things in terms of chemical reaction kinetics. We view what's going on inside of cells or within the extracellular space, within the tumor, as reaction diffusion processes. There are pretty standard ways of describing those from physics and chemical engineering. In one example here, I'm showing a gene expression cascade. We have a gene, g, that can be active or inactive. In its active state, it can be transcribed to make some mRNA, m, and the mRNA can be degraded and turned over. And when the mRNA is around, it can be translated to make some protein, p, and, and then the protein can also be turned over and integrated.

If we want to try to understand how much mRNA or how much protein is in a cell at any given moment of time, we make some assumptions on what those reactions are, how fast they go, called a rate law assumption. Then we add up all the ones that are making something and subtract all the ones that are taking it away. It’s like molecular accounting, like you balance your checkbook, we can balance some molecules in the cell. Then we know at least from the simulation, how many molecules are there at any given moment at the time.

The formalism we use to do that is called a differential equation. In this case, they're ordinary differential equations that capture what I just said in words. We can do that for gene expression. This is a standard way of thinking about an enzyme in the cell and how it might interact with something called a substrate, s, or a drug or inhibitor. The inhibitor combines that protein in different types of ways and inhibits its action, and these are modeled in the same way. Another thing we can do to model one important aspect of the heterogeneity is that some of the reactions and cells are very stochastic in nature because they have low molecule numbers. For example, for most of the genes, you have two copies of the gene because you inherited one from your mom and one from your dad. The probability that it is going to be on or off at any given moment of time is random in time. It's more burst-like, as opposed to smooth and continuous. By representing the systems in this way, it gives us kind of a clean handle to model

what we call stochastic gene expression, and the idea that mRNA levels and protein levels actually fluctuate quite a bit over time in a single cell. If you take a snapshot across the cell population, you have quite a broad distribution of expression. That's the formalisms that we try to use to describe what we think we know about cells and signaling pathways.

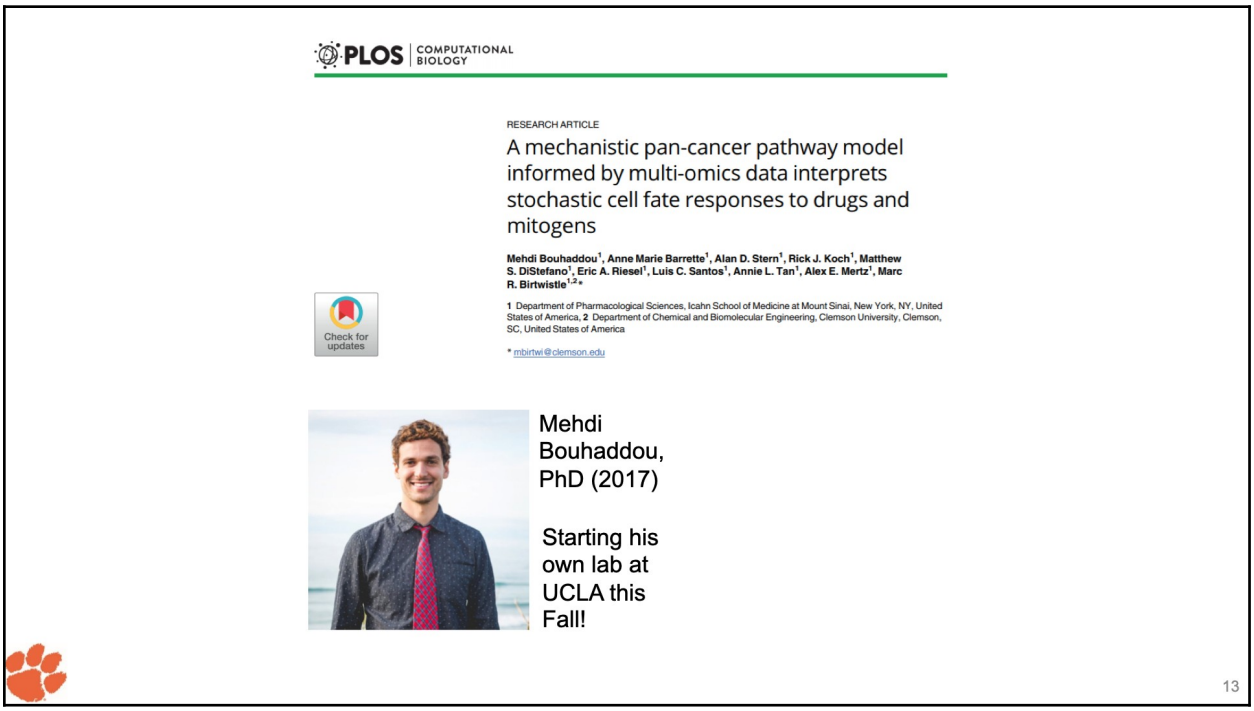

Our first paper on the topic was about five years ago now in PLOS Computational Biology. First we had to decide what we were going to model if we're going to try to model what's going on in a single tumor cell.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

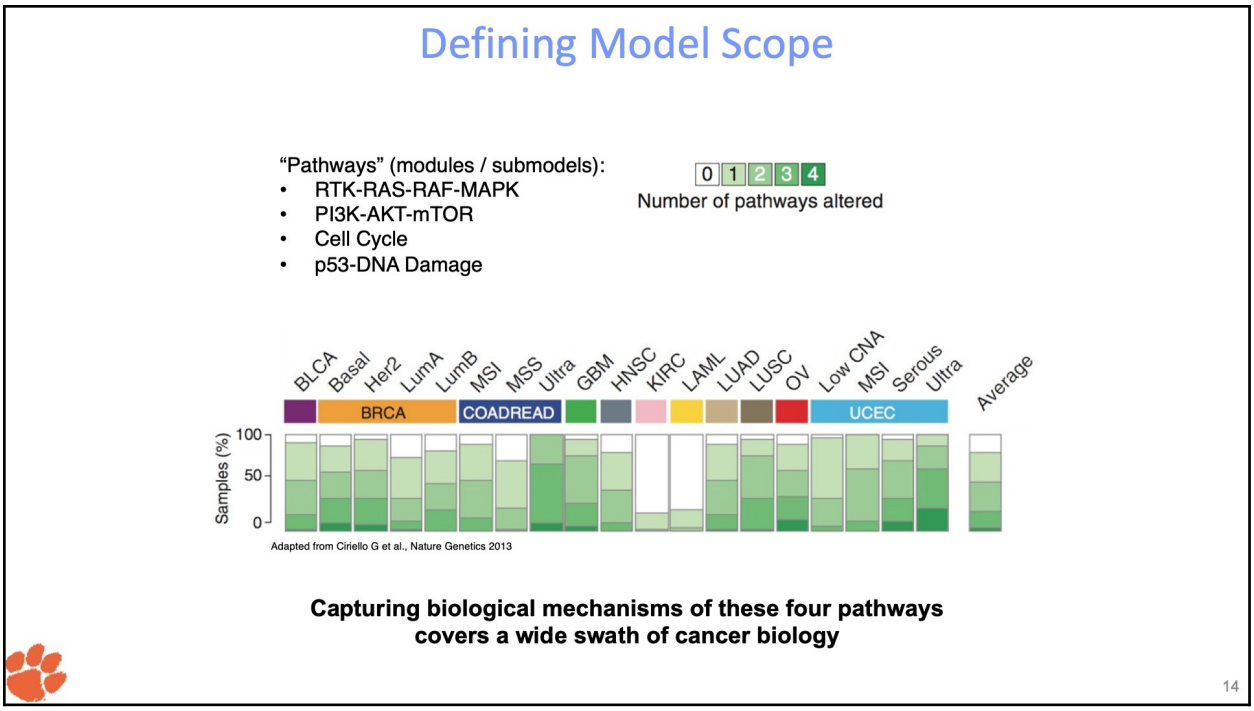

We can't model everything yet, but we want to model more than people have been. We looked at data from The Cancer Genome Atlas at the time and tried to understand: where do we need to draw our boundaries for a good first effort?

This is a figure showing that if we at least built a model that captured the known biology of these four different pathways: (1) receptor tyrosine kinases, (2) PI3 kinase-AKT-mTOR, (3) cell cycle, and (4) DNA damage, we would capture quite a large amount of the known alterations across different cancer types. This figure shows across cancer types as assayed by The Cancer Genome Atlas, how many of these pathways were altered across the different patients that were assayed. You see a lot of shading in there saying that this seems like a good place to start. Where do we start?

I was involved in building a lot of models of these individual pathways. At the time there were other people that were modeling and publishing their models of these different types of pathways. We went to find models and asked how we actually put them together in a way that is biologically consistent with what is known.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

15

Model is Composed of Pathway-Specific Models from the Literature

Submodel	Origin
Receptor Tyrosine Kinase	• Birtwistle MR et al., MSB 2007 • Bouhaddou M & Birtwistle MR, Mol Biosys 2014 • Many others
Proliferation & Growth	• Birtwistle MR et al., MSB 2007 • Nakakuki T et al., Cell 2010 • Kriegsheim A et al., Nat Cell Bio 2009
DNA Damage	• Batchelor E et al., MSB, 2011
Cell Cycle	• Gerard C & Goldbeter A, PNAS 2009
Apoptosis	• Albeck JG et al., Plos Biology 2008
Expression	Central Dogma (first order)

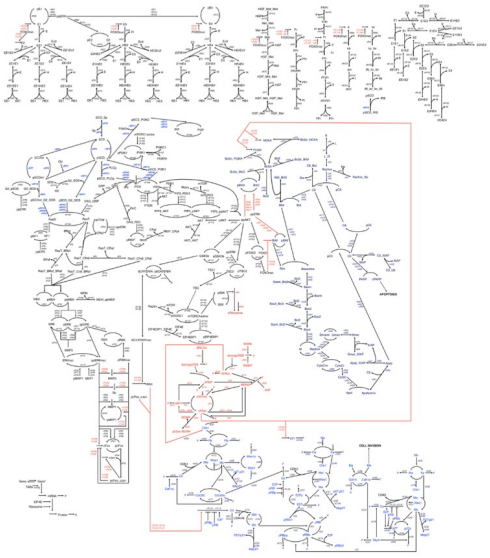

System of ordinary differential equations (ODEs)
~800 species and ~2800 reactions

15

We took all these models, and we put them together, and we came up with this kind of monstrous schematic for how we think at least these core pathways are working together to control how a cancer cell proliferates, grows, dies – processes like that. We ended up with a very large system of ordinary differential equations with a lot of different biochemical reactions. I joke that when I don't want a student to join my lab, I just start out and show them this picture, and it scares them away. That's an incredibly oversimplified view of the problem though.

That’s crazy! You’re insane!

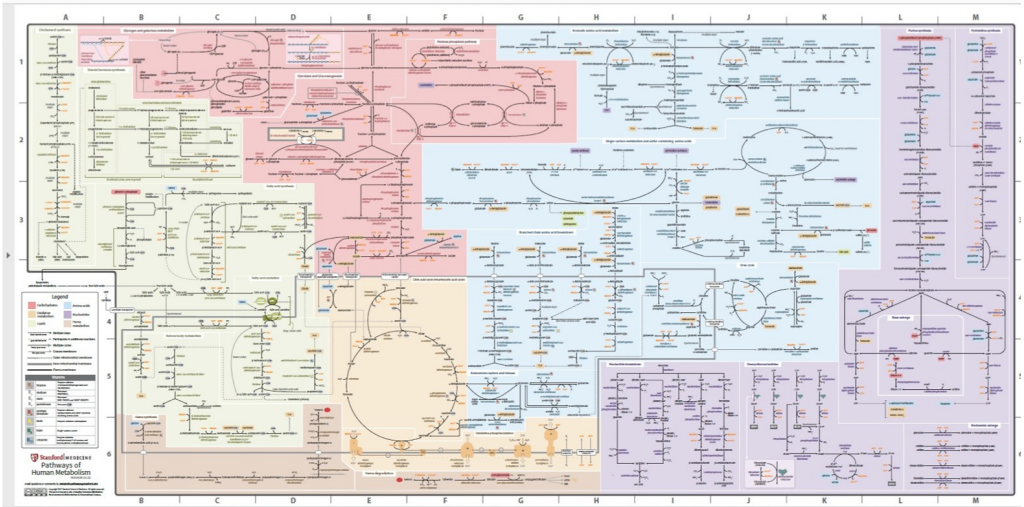

https://mededucation.stanford.edu/pathways-download/

Things like this have existed for metabolism for quite a while

16

It actually is much more complicated, which is scary on its own. Then people say, “man, you're crazy”. How do you even model all that stuff? I try to connect it back to metabolic pathways, which we've had maps like this one for since the sixties or seventies. For a long time, we've known that human metabolism is complicated like that.

We've mapped out a lot of these pathways because there have been a lot of good scientists from a lot of labs focusing on one particular reaction or one gene and where it fits in the overall map. We have a lot of that for signal transduction. We have to start putting it together.
What do we actually use to simulate?

Yes, we used MATLAB first. Our first version of the model was in MATLAB, and it was coded by a PhD student that did that work in my lab. He was not a computer scientist. We got it to work. But when I started trying to hand that model off to the next people in the lab, nobody could figure out how to change the model. It was a nightmare because it wasn't structured the way a software engineer would. MATLAB isn’t open source. That makes it even harder for people to use.

The next project in the lab, which I collaborated with Alex on, was to make this model formulated in a way that was easy to change, easy to use, and open source. That just got published a couple of months ago where Jamal is a postdoc in my lab. We reformulated this model so that it was essentially scalable open source. Now it's easy to use. We can add pathways, for example, pretty easily. If there's an existing model, you can explore alternative mechanisms pretty easily just by changing a couple text files. It's all compiled. It runs really easily and quickly. Thanks to Alex, we can run this on cloud computing resources pretty easily in reliable ways.
That answers how we code it and how we do it? That's where we're at right now. As a testament to that, we had a summer student who was interested in coming to my lab from France. She's a pharmacy student, super smart. It would be interesting to have her. She was interested in working with these models. I said, okay, let's try to add some targeted drugs into the model. We had a list of 15 or 20 we were interested to add, and she was able to within a week kick up, use the model, install it, and start changing it and adding these different drug mechanisms right away. That's exactly what we wanted out of this work.

Alex Feltus: What's the approximate cost to buy your model?

Marc Birtwistle: It's free. Zero.

Alex Feltus: It’s open source.

Marc Birtwistle: I'm not a company for the model. I just published a model and hope people can find it useful and build on it.

Rick Stanton: What language is your model in?

Marc Birtwistle: Python. It's mostly runnable through Jupyter Notebooks, which is a type of Python that's very easy to interact with. We made it compatible with a language called SBML Systems Biology Markup Language, which is a more standard way of describing reactions and ways that software can understand them. We have those two things working together, and it works with some other packages that then takes that SBML representation and compiles it into C code, so that it's really, really fast. We can use high powered ODE solvers to do this, like before, in some reasonable time.

Brian McCloskey: You are working on the right problem, which is: how do we intelligently integrate combinatorials into patients’ care? We've had a number of these conversations, including with Ally Perlina of CureMatch. I know that you're very early in the translational pipeline. You're doing basic science.

As you know, further ahead in that pipeline to me, the patient, you have a couple of constituents that can create roadblocks to actually using this technology. One is big pharma. When you're talking about combinatorials, you're talking about drugs, likely from different pharmaceutical companies. How do we get them to adopt this?

The second thing is: I have had a number of conversations with my medical oncologist about how to integrate, for example, functional precision medicine, testing, organoids, et cetera, into my care and not just my oncologist, but a number of other oncologists. I consistently run into this challenge of, “we don't really believe in ex vivo models.” I'd love to get your thoughts in terms of how to address those two constituencies, and the barriers that they may create, because obviously you are doing this because you care about helping cancer patients like myself, and a number of other people on this call. I'm trying to ground this a little bit and maybe figure out how to help you to bring this science to fruition, to make a difference in patients' lives.

Marc Birtwistle: Those are really interesting, challenging questions. I have to admit, I haven't thought much about how pharma companies would perceive this, and if they perceive it as a threat. That could be a problem. I don't understand enough about the drug development pipeline for anticancer drugs to predict how they might respond. Or if all these things are already FDA approved, I couldn't see why they would be opposed to selling more of their drug. Would a Roche drug play nice with a Pfizer drug, if that turns out to be the best combo? I don't see why not, but I feel that I'm too naive about that whole dynamic to really say anything there. I should think more about that because if there are real roadblocks there, and I identify them now, I could do something about it a lot more easily in the development.

Brian McCloskey: There are challenges from not buying into ex vivo models because it doesn't accurately simulate what's happening in vivo. I've heard that from a number of doctors, to the challenges of actually even enabling some of these ex vivo models. I'm going to get biopsied soon. For some of these clinicians, the whole process of getting tissue just seems daunting.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

Marc Birtwistle: It's a model. Clinicians would say a mouse model is terrible because it doesn't get everything right. It does get some things right. I guess the ex vivo models are kind of newish. There's not a good understanding of where they're right and where they're not, which is probably part of the trepidation on their part. The one aspect of ex vivo models I like a lot, especially when it's directly from that patient's tumor biopsy, is that you're getting some real data from cells that directly then came from that patient. It's not in the true microenvironment, and it's not having all those other factors from your body, but that to me is better than a mouse model or a cell line model that just is the disease you have, but isn't quite like it.
The other thing about it is that you can do high throughput screening with those approaches. You could grow out your cells a little bit, and put them in multiwell plates, and then do screening with single drugs, and find out at least as single agents, which of these single agents seems to have some kind of efficacy against the cells that actually came from your tumor. I feel like that would be super useful information to be able to inform something. In the absence of nothing, that would be pretty good, and we could use that kind of data eventually if we have a good simulation model that could be really strong constraints for such a model to help it do a better job of predicting. I could see it as very useful in that sense too.

Brad Power: You're saying the same thing that I was going to raise: I was talking with Matt Vander Heiden, who's the head of the MIT Koch Center for Integrative Cancer Institute. We were trying to come up with a problem that MIT computer scientists would get into because it's a difficult problem. If it's an easy problem, they're not interested. One of the problems we were brainstorming was: We know cell lines aren't a real situation that occurs in patients. We know that mouse models aren't real, organoids aren't real, and functional testing isn't real, but could we put a simulation wrapper around whatever the signal is coming out of those models and compensate for whatever is not real about them? In other words, to make them more real, to simulate something that's more real, more personalized, more human. That to me is an interesting problem. And it sounds like you were almost saying the same thing.

Marc Birtwistle: I would totally agree with that. That sounds right. Nail on head.

Brad Power: Alex: would you please speak to your role, your connection with Marc, and comment on some of the hackathons you've run, and some of the work you do with students?

Alex Feltus: I've worked with a few patient hackathons. I worked with a patient a couple years ago, a pancreatic cancer patient, who I got hooked up probably through Bill Paseman. Bill's the nexus of patient-centered research. We were trying to get biopsies. It took us two months. You guys are sick, and we couldn't get the biopsy for a long time. I talked to the oncologist. I talked to the cancer center, all that kind of stuff. They lost the sample. By the time they went in and got another biopsy done, he was gone a couple weeks later. We didn't have a chance to do any fancy genomics. It taught me a lot about the realities. He was a really fantastic guy.
I’m trying to democratize research that's locked up in my lab. I'm trying to open it up to the planet, and we're doing it through hackathons. We have done hackathons with Bill Paseman's RNA seq data.

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

Rick, this goes to your question. We have to deal with his RNA seq data. We had to have the raw sequence data to go back and process it the same way that was processed to unify The Cancer Genome Atlas, and another project called GTEx (Genotype Expression Portal), which has normal samples. To be able to process that we have an algorithm that goes through and finds RNA level biomarkers that are different in a single tumor. It works with any equals one it's trained on, on the tumors in normal tissues that are in the databases. It's not perfect, but it seems to be pretty good. But the point is that we do these hackathons. If the data is truly open source, like Bill Paseman’s open source data, we can just download it from the internet and analyze it. Then we can do hackathons with the data. A problem I have with forums like this quite honestly, is that there's an hour to talk, and there's a lot of heads on here. There's not a chance to really get into discussions that sometimes save years. To figure out Marc’s stuff is probably years away from being truly translational. But I think Marc’s stuff is the stuff that's going to change everything.

If we can get the data, and Rick this is a specific question, can you have the raw sequence data? If you can go back to the raw data then you can process it in a way that can be used to do analytics in a hackathon.

Rick Stanton: I have access to my primary tumor, which was sequenced twice, once by Tempus, and once by Ashion, which was via City of Hope via TGen. I can get that.
I can process that or anyone can process that via the Xena portal at UC Santa Cruz. I've downloaded that data, which has been processed with the toil recompute. That has been a consistent processing of TCGA. I'm working with that data now. I can get that data, but unfortunately only myself because the other guys on the call have been sequenced via Tempus, which we only get counts and TPM.
I'll definitely follow up with what kind of data and see what we can do. It'd be interesting for your team, and it'd be super helpful for us.

Brian McCloskey: I can probably get it through Tempus. At one juncture we had like a BAM file for both my primary and met tumors. I just have to reach back out to them.

Alex Feltus: Let's have a conversation if you want about having a hackathon where we can get a group of people to process that. We need to go back and process it to reanalyze it, to put into an algorithm, do it a certain way with the raw data. This stuff takes weeks to process sometimes and get the analysis done. That's what the hackathon is for and why we don't do 24-hour hackathons. We do perpetual hackathons. We started a hackathon with Bill Paseman back in April for kidney cancer, and it's still going on.

—

Followup Email from Marc Birtwistle Responding to Questions (Mostly from Rick Stanton) in the Zoom Chat

“Simulations for Predicting Treatment Response” (Marc Birtwistle) [#20]

I wanted to address chat items that might not have been fully discussed in the meeting. Please see below. I enjoyed the chance to talk about our work and why we think it is useful to embark on, and also to discuss with you all about all such things.

First off - fantastic work!!! immune signaling?

> *We have added an IFNg pathway to the model recently, but don’t yet branch outside of models of tumor cells. That will be an important scale up step to consider immune (and other) cell types and where they are spatially.*

measurement assays?

> *We built the model purposely to take as input gene copy number, RNAseq, and proteomic (if available) data to “initialize” the model to be as consistent as possible with a particular biological sample (cell line / patient) . As a side note, we've recently developed a high throughput western blot technique (called mesowestern) and are trying to commercialize it (got word an NIH grant for small business is going to be funding our company Blotting Innovations). Western blots are great for functional protein measurements that are more trusted. We think that they could find more clinical use if they were more high throughput and robust which is one thing we are trying to do.*

spatial?

> *We don’t do spatial yet but agree it is really important. We are collaborating with a company named Simbiosys who builds spatial models based on MRI data to put our models in theirs….you should check out simbiosys! I can put you in touch with people I work with there if you're interested.*

RNASeq?

> *Yes we use transcriptome data to set mRNA levels in the model.*

We tried to assess the dominant software/platform to model? simulink? Any AI? machine/deep?

> *As mentioned in the discussion we now use python / open-source solvers (fast). I believe machine / deep learning could be integrated with mechanistic models to help change model structure when we are uncertain about certain connections in modeled networks (we have a project ongoing about it…)*