

**Article:**

Rebecca D Ganetzky and Stephen R Master.
Machine Learning for the Biochemical Genetics Laboratory.
Clin Chem 2020 <https://doi.org/10.1093/clinchem/hvaa168>

Guest: Dr. Stephen Master is the Chief of the Division of Laboratory Medicine and Medical Director of the Michael Palmieri Laboratory for Metabolic and Advanced Diagnostics at the Children's Hospital of Philadelphia, and an Associate Professor of Pathology and Laboratory Medicine at the Perelman School of Medicine of the University of Pennsylvania.

Bob Barrett: This is a podcast from *Clinical Chemistry*, sponsored by the Department of Laboratory Medicine at Boston Children's Hospital. I am Bob Barrett.

Machine learning and artificial intelligence are two terms that were coined in the 1950s, but only now are beginning to be put into solving practical problems. Only within the past few years have machine learning algorithms been demonstrated to automate the interpretation and analysis of clinical chemistry data in a variety of situations.

In the September 2020 issue of *Clinical Chemistry*, Rachel Carling and her colleagues in the U.K. published a paper on a machine learning approach for the automated interpretation of amino acid profiles in human plasma. The same issue contains an accompanying editorial titled, "Machine Learning for the Biochemical Genetics Laboratory." It was authored by Dr. Rebecca Ganetzky and Dr. Stephen Master, and we are pleased to have Dr. Master here as our guest in today's podcast.

He is the Chief of the Division of Laboratory Medicine and Medical Director of the Michael Palmieri Laboratory for Metabolic and Advanced Diagnostics at the Children's Hospital of Philadelphia, and an Associate Professor of Pathology and Laboratory Medicine at the Perelman School of Medicine of the University of Pennsylvania.

So, Dr. Master, first of all, what exactly is machine learning, and why would it be significant for the clinical laboratory?

Stephen Master: Well, machine learning is a branch of artificial intelligence that uses computer algorithms that can learn from data. So, for example, if I want to predict a model of whether someone is going to buy a product, I might want to create a model based on their age, their income, their internet-browsing history, or a number of other factors that might reflect how attractive that product would be to them. To create that kind of a model using machine learning, I would start with a set of training data that had all the relevant information about a group of people, along with whether

they had purchased the product, and the algorithm would then tune the parameters of the model in a way that tried to maximize its ability to make correct predictions, and I could then test that machine learning model by taking data from a new group of people and seeing how well the model predicted whether they would purchase the product.

So, I've given you a common commercial example, but the same type of approach has been applied in a number of healthcare settings. Anatomic pathologists, just to pick one example, have really aggressively explored the use of machine learning to automatically identify concerning areas of tissue slides. In the clinical laboratory, we generate a tremendous amount of quantitative patient data on a daily basis, and so it's natural to think that machine learning provides a powerful tool to add interpretative value to the work we do, that is, to integrate patterns of lab values and other medical data into models that predict disease or prognosis in patients.

Bob Barrett: Well, I mentioned earlier that the September 2020 issue of *Clinical Chemistry* contains a report describing the use of machine learning to analyze amino acid profiles in the biochemical genetics laboratory. What is new with this approach and why is it important?

Stephen Master: Well, when you think about applications of machine learning to laboratory medicine, one obvious area would be the biochemical genetics laboratories, and there are two specific reasons that I say this.

The first is that these labs diagnose and manage inborn errors of metabolism using assays such as plasma and urine amino acid profiles, urine organic acids, or acylcarnitine profiles, and each of these assays that I just mentioned shares the characteristic that one test measures a significant number of different compounds in the patient, and that means not only that there's a lot of data generated for each patient, but also just as importantly, that the patients who get these tests can be easily compared.

So, unlike some other laboratory results, you can imagine where one patient might have gotten one set of tests and another patient might have gotten a different set of tests, every patient who gets these particular tests that we're talking about should have a large set of comparable data. And the second reason why I think the biochemical genetics lab is a particularly interesting area in which to apply machine learning is that interpreting complex assays, such as amino acid profiles, requires a fair degree of training and skill, and the amino acid profiles alone, there are over 50 known disorders that can be reflected in these results and some of them can be difficult to spot. So, traditionally, this

means that an expert interpretation is always performed in reporting these results.

When you think about it though, finding subtle patterns and complex datasets is something that machine learning can sometimes do quite well. So, I think that one of the significant aspects of the work in this paper is that it has begun to address the question, can we use machine learning to aid, augment, or even replace the work of the traditional human laboratorian in interpreting amino acid profiles?

Bob Barrett: So, what are the particular challenges when using machine learning in these types of applications?

Stephen Master: Well, I've discussed the fact that successful machine learning relies on having a very good training dataset and the more data, the better. Beyond that, it's not enough just to have data, known diagnoses for training. It's also important to have a second, or even a third, independent dataset that can be used to tune the performance of machine learning algorithm and to independently validate how well it performs as a predictor. So, without sufficient data and the right data, even the best machine learning algorithm won't perform well, and I would say in my experience, getting the data that you need is probably the main challenge.

Bob Barrett: What about the situation with relatively rare diseases? How can you gather sufficient data for the machines to adequately learn?

Stephen Master: That's a great question. As you can imagine, if it's difficult to collect enough data for common diseases, it's even harder to collect enough data for rare diseases. And you can see that the authors of this current paper really wrestle with that issue, so even though they were able to amass a fairly sizeable cohort, over 2,000 cases actually, they were still forced to group a number of very rare, diverse conditions into what they called "rare, inherited metabolic disease" as a category, simply because of the very small number of positives. And then, in order to prove that the classifier was robust, they used the smallest group of external-quality assurance specimens that represented a subset of the diseases that they wanted to classify, and I don't say this to take anything away from their work, actually. I think it just shows the challenges of doing this with rare conditions. I actually think there are several approaches that might be helpful here, but I think that the best strategy for the field, honestly, is likely to involve getting biochemical genetics labs to create larger consortia, to share patient data and, therefore, get a better representation of some of these rare conditions.

Bob Barrett: Well, what about the risks of using machine learning algorithms? How can a lab ensure that their results are accurate?

Stephen Master: Well, any time that we talk about using a diagnostic algorithm, it raises very important issues that we have to think about with respect to patient safety, and even liability. The current discussion goes back at least a decade, and it led to some very specific recommendations by the U.S. Institute of Medicine in 2012, regarding how to validate the performance of these types of machine learning models. So, I think it's worth stating clearly that any lab that's thinking of implementing a machine learning model should make sure that they really, really understand the best practices, pitfalls, and issues surrounding correct validation. But when thinking specifically about risk, though, I have to say that I think that one really interesting idea is the paradigm of using these types of tools to augment rather than replace human diagnosticians. So, in the case of amino acids, a model that could reliably tell you that a profile looks normal would still be useful in helping to triage the cases that a human should examine more carefully, even if it couldn't necessarily identify every single rare disorder on its own.

Bob Barrett: Well, let's look ahead. What do you foresee as the future for this type of approach in the next five, even ten years? How should clinical laboratorians think about the significance of machine learning for their future practice?

Stephen Master: Yeah. I think that it is inevitable that these approaches will become more and more important throughout healthcare and particularly, in laboratory medicine context. The reality is that it isn't a question of whether they'll be developed and used, but rather who's going to develop them and use them, and that's where I actually feel very strongly that we in the laboratory medicine community, not just in biochemical genetics as we've been discussing, but more broadly across lab medicine, need to make sure that we're involved.

There are important characteristics of our data that can affect machine learning, and those characteristics may not always be apparent from outside the clinical lab. So, I really think that we want to drive this as a field to make sure that the models that are created using laboratory data are well-designed and appropriate for patient care, and frankly, that's why I'm delighted that this study we've been discussing today has been published in *Clinical Chemistry*, because I think it reflects the exciting kinds of work that can result when recent developments in machine learning are combined with solid laboratory medicine.

Bob Barrett:

That was Dr. Stephen Master from the Children's Hospital of Philadelphia and the University of Pennsylvania. He is co-author of an editorial on machine learning and the biochemical genetics laboratory that appears in the September 2020 issue of *Clinical Chemistry* along with an original scientific paper applying machine learning to the interpretation of amino acid profiles. I'm Bob Barrett. Thanks for listening.