



Article:

Daniel Walke, Daniel Steinbach, Thorsten Kaiser, Alexander Schönhuth, Gunter Saake, David Broneske, Robert Heyer.

SBC-SHAP: Increasing the Accessibility and Interpretability of Machine Learning Algorithms for Sepsis Prediction.

J Appl Lab Med 2025; 10(5): 1226–40. <https://doi.org/10.1093/jalm/jfaf091>

Guest: Daniel Walke studied biosystems engineering and is currently a PhD student at the Otto-von-Guericke University Magdeburg in Germany.

Randye Kaye:

Hello and welcome to this edition of *JALM* Talk from *The Journal of Applied Laboratory Medicine*, a publication of the Association for Diagnostics & Laboratory Medicine. I'm your host, Randye Kaye.

Sepsis is a life-threatening condition that occurs when the body's immune system has an extreme response to infection, leading to organ dysfunction. Sepsis is one of the leading causes of death worldwide and early detection is critical. Prompt treatment with antibiotics can significantly reduce both morbidity and mortality. Because the inflammatory response in sepsis is driven by cytokines released from neutrophils and macrophages, data from a routine complete blood count test, known as a CBC, may offer valuable clues. However, no single CBC component has proven sensitive or specific enough on its own.

Recently, machine learning models using CBC data have shown promise in enhancing the diagnostic value of this common test. Still, these models often require programming skills and can be difficult to interpret, especially when trying to understand how individual components influence the prediction.

The September 2025 issue of *JALM* features an article that proposes SBC-SHAP, an openly accessible web application for the classification of sepsis based on complete blood count data. The model may increase the interpretability and accessibility of machine learning classifiers for predicting sepsis and may enable faster detection of sepsis without the addition of new diagnostics outside of standard clinical practice.

Today, we're joined by the article's corresponding author, Daniel Walke. Daniel studied biosystems engineering and is currently a PhD student at the Otto von Guericke University Magdeburg in Germany.

Welcome, Daniel. Firstly, what is SBC-SHAP and what primary challenges in sepsis prediction does it address?

Daniel Walke:

So, SBC-SHAP is a free and open source web application. It is designed to predict a sepsis risk using a patient's complete blood count data. And it was basically created to solve two main challenges. So first of all, we have the problem with interpretability, because state-of-the-art's deep learning models like graph neural networks, for example, can already effectively analyze a time series information really, really good like complete blood count information, and they achieved really high scores in terms of accuracy.

However, the problem is that they often function as so-called black boxes. So that means that the interpretability is really low and what we try to achieve is that we can directly see for specific feature values how they contribute to the resulting sepsis prediction.

The other problem we are trying to tackle is the accessibility, because many powerful machine learning models are published on platforms like GitHub and Hugging Face, for example, and setting up these complete machine learning pipelines can be extremely complicated. It can be time-consuming, and this is something that most clinicians don't have, especially during the daily work.

And so therefore, we developed SBC-SHAP to solve exactly those two challenges, regarding the interpretability and the accessibility, by providing a really and intuitive graphical user interface where the clinicians can just enter their complete blood count information and then they receive the predicted sepsis risk. And nevertheless, we also wanted to achieve state-of-the-art performance, which should be completely interpretable so that clinicians can directly see how those decisions were made by the machine learning models.

Randye Kaye:

All right, thank you. So how exactly does SBC-SHAP work? What does the clinician need to input and what type of information will they receive?

Daniel Walke:

So, clinicians need to input patients' information inside the user interface. They need the information about the age, the sex, and an anonymous patient identifier, and then need five key blood parameters from the complete blood count information, which are specifically hemoglobin, platelets, red blood cells, white blood cells, and the mean corpuscular volume.

Additionally, we require to track the time information so that the machine learning models can incorporate previous time measurements and analyze the trends within the data. And what clinicians will receive at the end is a predicted sepsis risk that is between 0% and 100%, where 0% indicates that the sepsis risk is really low and most probably is the patient

doesn't have sepsis; however, a really high risk of like nearly 100% means that the patient most probably has a sepsis. And additionally, we receive a visualization of SHAP values that show how this prediction was made by the machine learning model.

Randye Kaye: All right, thank you. And you may have already answered this, but maybe you can expand on the SHAP values. Like what are they? How do they help users understand a specific sepsis risk prediction?

Daniel Walke: Yes. So basically, SHAP is a framework which is designed as a game theoretical approach that is used to identify which specific feature values contributed how to the final prediction. So thereby we can identify why and how a specific prediction was made. And we can differentiate between two kind of SHAP values.

So, first of all, we have like positive SHAP values. Those are in the web application visualized in red to indicate that the sepsis risk is increased by this specific feature value. And then we also have negative SHAP values which are visualized as blue bar charts to indicate that the sepsis risk is decreased.

And what is also important is the magnitude of the specific SHAP value. So basically, the size of the bar chart, whether it's higher or lower, because it shows how impactful this specific feature value is on the overall prediction. So, for example, a larger positive value indicates that this specific feature value highly increases the predicted sepsis risk, whereas for example, a lower negative value highly decreases the predicted sepsis risk.

And thereby we can directly investigate, for example, that a physiological, so normal, white blood cell count value would decrease the predicted sepsis risk, whereas for example, a really low number of platelets or a high number of white blood cells would highly increase the predicted sepsis risk.

Randye Kaye: So how does this model compare to previously proposed models?

Daniel Walke: So under the hood, compared to other models, we are first constructing a graph model based on the time series information of a specific patient and this network basically just connects all the time series information from the start measurement, from complete blood count data to all the ones that we have currently measured. And those time series information are then used to create new features that are then fed to machine learning models, like for example random forest.

One problem, however, was that those patients have a different number of time series information during the admission. So there are some that might have dozens or even hundreds of different measurements during the admission and then we can get really accurate predictions. However, there are also some patients that might have only one or two measurements and this is a problem to be accurate enough.

And what we did is that therefore we introduced a reference node that functions as a baseline, and this reference node integrates all the physiological values for all the control samples within the complete data set, and thereby the machine learning models can check whether the current values are a bit off from this baseline or not, or whether they are in a physiological range. This helped us to increase the model sensitivity from approximately 78% from previous machine learning models to over 82% while maintaining the specificity of 80%.

Randye Kaye: All right, thank you. One final question. What about the future? What future directions or potential enhancements do you foresee for this tool?

Daniel Walke: We see several directions for the future. So, first of all, what I can think about is the enhancement of the model itself, because we can further integrate some further parameters, such as, for example, the body temperature or some specific markers like procalcitonin or the C-reactive protein, and they would even further boost the classification performance.

But we could also integrate something like unstructured text data from analysis or other clinical notes. Another possible direction might be to expand the scope so that we don't only investigate or predict the sepsis risk, but probably also some other conditions and diseases that could be analyzed with blood count information, such as different types of infections.

Finally, what we could also do is that we do a clinical survey to check how useful this tool and especially specific features for the day-to-day clinical practice.

Randye Kaye: It's all very exciting. Thank you so much for joining us today.

Daniel Walke: Thank you for having me here.

Randye Kaye: That was Daniel Walke from Otto von Guericke University Magdeburg, describing the *JALM* article "SBC-SHAP: Increasing the Accessibility and Interpretability of Machine Learning Algorithms for Sepsis Prediction." Thanks for tuning in to this episode of *JALM* Talk. See you next time and don't forget to submit something for us to talk about.