

<https://doi.org/10.1038/s41746-025-01529-x>

# Identifying and forecasting importation and asymptomatic spreaders of multi-drug resistant organisms in hospital settings

Check for updates

Jiaming Cui<sup>1,2</sup>, Jack Heavey<sup>3</sup>, Eili Klein<sup>4,5,6</sup>, Gregory R. Madden<sup>7,8</sup>, Costi D. Sifri<sup>7,8</sup>, Anil Vullikanti<sup>3,9</sup> & B. Aditya Prakash<sup>1</sup> ✉

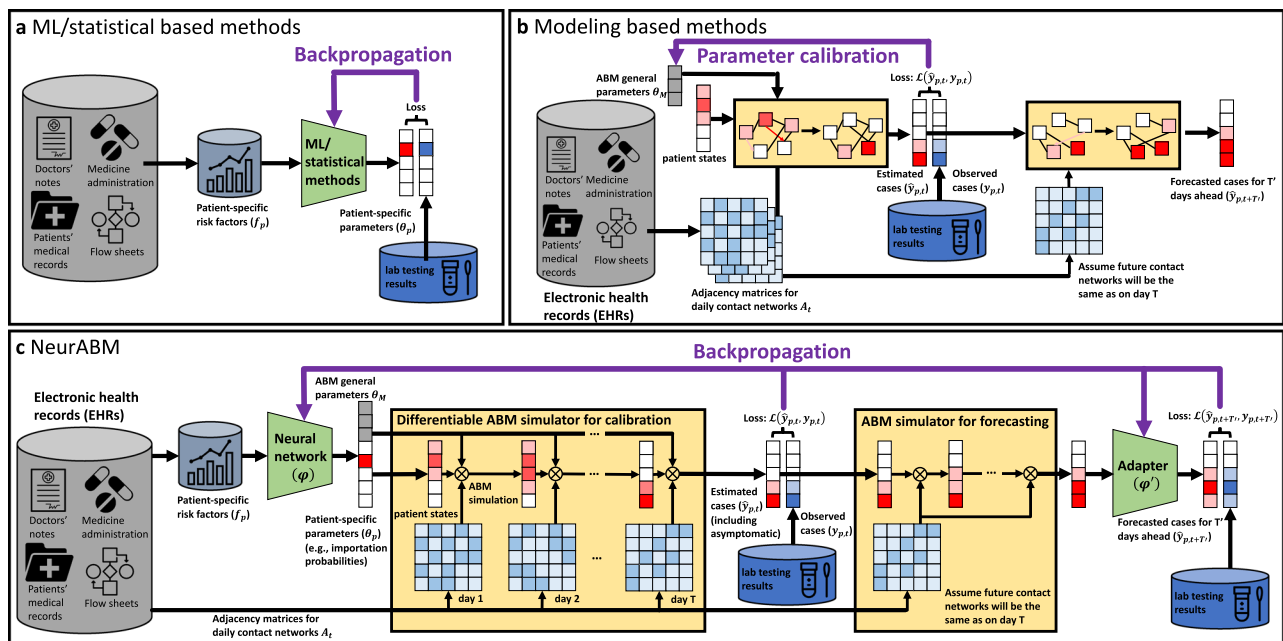
Healthcare-associated infections (HAIs) from multi-drug resistant organisms (MDROs) pose a significant challenge for healthcare systems. Patients can arrive at hospitals already infected ("importation") or acquire infections during their stay ("nosocomial infection"). Many cases, often asymptomatic, complicate rapid identification due to testing limitations and delays. Although recent advancements in mathematical modeling and machine learning have aimed to identify at-risk patients, these methods face challenges: transmission models often overlook valuable electronic health record (EHR) data, while machine learning approaches typically lack mechanistic insights into underlying processes. To address these issues, we propose NeurABM, a novel framework that integrates neural networks and agent-based models (ABM) to leverage the strengths of both methods. NeurABM simultaneously learns a neural network for patient-level importation predictions and an ABM for infection identification. Our findings show that NeurABM significantly outperforms existing methods, marking a breakthrough in accurately identifying importation cases and forecasting future nosocomial infections in clinical practice.

Healthcare-associated infections (HAIs), especially those caused by multi-drug resistant organisms (MDROs), pose a significant threat to patient safety and burden the healthcare system with increased costs due to longer hospital stays and more expensive therapies<sup>1–6</sup>. Approximately 3% of hospitalized patients in the United States acquire an HAI during their stay, resulting in more than 35,000 deaths annually<sup>7–10</sup>. In many occasions, patients may have already been colonized (pathogens present on patients without causing disease<sup>11</sup>) or infected but may be asymptomatic at admission (i.e., importation cases)<sup>12</sup>. For instance, the European Center for Disease Prevention and Control (ECDC) estimates that importation cases contribute to 13% of HAI cases in Germany and 18.9% in Spain<sup>13</sup>. These importation cases can spread HAI-causing pathogens either directly via patient-patient contacts or indirectly via healthcare workers (HCWs) and contaminated physical surfaces<sup>14–16</sup> and lead to nosocomial infection cases,

which can also be asymptomatic but further spread pathogens to additional healthy patients<sup>17</sup>.

Despite the critical concerns associated with importation and nosocomial infection cases, identifying these cases rapidly and accurately still remains a challenging problem. Current methods to identify them include surveillance tests<sup>18,19</sup>, machine learning-based methods<sup>20,21</sup>, and transmission modeling-based methods<sup>22–28</sup>. However, each of these methods suffers from unavoidable drawbacks. For example, surveillance tests, such as culture or PCR tests, are common in hospitals; however, they are costly, require time to process, and are not 100% accurate<sup>29</sup>. Additionally, they are not applicable for all MDROs; typically, they can be used for only a subset of MDROs<sup>30</sup>. Machine learning and statistical techniques use patients' electronic health record (EHR) data to predict the probability of importation and nosocomial infection cases<sup>20,21</sup> (see Fig. 1a). However, the performance

<sup>1</sup>College of Computing, Georgia Institute of Technology, Atlanta, GA, 30332, USA. <sup>2</sup>Department of Computer Science, Virginia Tech, Blacksburg, VA, 24060, USA. <sup>3</sup>Department of Computer Science, University of Virginia, Charlottesville, VA, 22904, USA. <sup>4</sup>Center for Disease Dynamics, Economics & Policy, Washington, DC, 20015, USA. <sup>5</sup>Department of Emergency Medicine, Johns Hopkins School of Medicine, Baltimore, MD, 21205, USA. <sup>6</sup>Department of Epidemiology, Johns Hopkins Bloomberg School of Public Health, Baltimore, MD, 21205, USA. <sup>7</sup>Division of Infectious Diseases & International Health, Department of Medicine, University of Virginia School of Medicine, Charlottesville, VA, 22903, USA. <sup>8</sup>Office of Hospital Epidemiology/Infection Prevention & Control, UVA Health, Charlottesville, VA, 22904, USA. <sup>9</sup>Biocomplexity Institute, University of Virginia, Charlottesville, VA, 22904, USA. ✉e-mail: [badiyap@cc.gatech.edu](mailto:badiyap@cc.gatech.edu)



**Fig. 1 | NeurABM framework.** **a** Traditional machine learning and statistical methods use patients' risk factors to predict the probability of being importation and/or future nosocomial infection cases. It is hard to explicitly incorporate epidemiological mechanisms. **b** Modeling based methods are built on mechanistic models to capture the HAI spread in healthcare facilities. Here, observed cases are used to train an ABM to (1) learn underlying ABM parameters and (2) pin the ABM states to the spread dynamics until today, which allows people to run ABM simulations for  $T'$  more days for forecasting. However, they cannot leverage the risk factor of each patient from the EHR and have to make assumptions for future contact networks. **c** To address this issue, we propose NeurABM framework to couple both neural networks and ABM simultaneously. Specifically, it is composed of two main components: the neural network component and the agent-based model (ABM) simulator component. The neural network component takes patient-specific risk factors data collected from EHR as input and outputs both the ABM

parameters (applicable to every patient) and patient-specific parameters (importation probabilities in this paper). The ABM simulator component takes daily contact networks collected from EHR and the parameters as input, runs simulations for  $T$  days, and outputs the probability that each patient is in the *Carriage* state for each day. We then compare them with the ground-truth observations (via lab testing) and compute the first loss. To forecast future nosocomial infections, we start from predictions on day  $T$  and run simulations for  $T'$  more days. Note that these  $T'$  more days are for future, we assume that the contact networks will be the same as on day  $T$ , and use another adapter network to correct the potential bias caused by this assumption and get the final forecast. We also compute the loss of this forecast with ground-truth observations. With both losses, we backpropagate the total loss to adjust neural network parameters. This design allows us to train both the neural network and the ABM simultaneously in an end-to-end manner, mitigating the issues encountered when using either component individually.

of machine learning methods has not proven to be sufficiently robust for clinical practice due to many reasons, including imbalance (since HAI cases form a very small fraction of the entire patient population) and bias in the data (since testing is generally not done in a systematic manner). Moreover, explicitly incorporating physical epidemiological mechanisms into machine learning frameworks remains an open problem. ML frameworks tend to focus on data-driven correlations rather than understanding and using complex causal relationships and dynamics of disease transmission<sup>31</sup>. Finally, modeling-based methods are based on detailed mechanistic models (e.g., compartmental mixing models<sup>26,27</sup> and agent-based models (ABMs)<sup>22–24</sup>) to capture the transmission dynamics of HAIs within a healthcare facility. They are calibrated to infections in the hospital and use projections from such models for prediction. Although ABMs have used information about contact networks between patients and providers within healthcare facilities to model the infection status of an individual patient, they still cannot directly incorporate the risk factor of each patient from the EHR data, such as medications, lab results, vital signs, and device use history into modeling. Besides, they rely on contact networks to run simulations, which makes their forecasts for future nosocomial infections less reliable since people have to make assumptions on future contact networks to run the model (see Fig. 1b). As we will also show in the later results section, these limitations lead to suboptimal performance in identifying importation and forecasting future nosocomial infection cases.

From a broader perspective, improved identification of importation cases and forecasts of future nosocomial infections could also improve hospital management. Practices to prevent the spread of MDROs within hospitals, such as testing, quarantine, and isolation, consume limited

healthcare resources and could only be implemented on a small group of patients. Meanwhile, decisions about prioritizing limited resources are critical to ensure effective and efficient delivery of healthcare services<sup>32</sup>, where agent-based models<sup>33–36</sup> and machine learning<sup>37</sup> have proven to be invaluable tools for optimizing these decisions.

In this work, we propose a new framework, NeurABM, to identify HAI importation and forecast future nosocomial infection cases by *coupling* a neural network and an ABM and training simultaneously. We use methicillin-resistant *Staphylococcus aureus* (MRSA) as an example HAI in later sections. Figure 1 shows an overview of our framework. As shown in Fig. 1c, the neural network  $\phi$  estimates the importation probability for each patient using EHR data, while the ABM incorporates MRSA dynamics and is used to estimate the MRSA infection probability. After training, NeurABM runs as a discrete time process; at each time step  $t$ , the following two steps are performed: (1) the neural network estimates the importation probability (i.e., identifies importation cases) for each new patient who enters the hospital, (2) the ABM (which keeps track of the disease states of all patients in hospital till time  $t - 1$ ) runs the next step of the disease simulation to estimate all disease states (i.e., identify nosocomial infection cases, including those which are asymptomatic) at time  $t$ . When forecasting for future, we assume that the contact networks will be the same as on day  $T$ , and use another adapter network  $\phi'$  to learn and mitigate the bias caused by this assumption. The parameters of NeurABM consist of two parts: those of the neural network and those of the ABM (e.g., general parameters such as transmission and recovery rates, and patient-specific parameters such as importation probabilities in this paper); these parameters are learned by minimizing a loss function that considers the errors in the ABM projections

and ground truth incidence data from EHR. Since the dynamics of MRSA transmission depend on the importation model (and conversely, through this kind of training process), this approach couples the neural network and ABM and is trained end-to-end, which mitigates the issues in using either of them individually. NeurABM significantly extends the work of Chopra and Rodriguez et al.<sup>38</sup>, which was the first method to consider such a joint deep learning and ABM approach, by introducing approximation techniques to scale the disease model, incorporating rich patient-level EHR data, and new techniques to train the pipeline robustly.

We demonstrate the performance of NeurABM using EHR data for patients at the University of Virginia (UVA) hospital intensive care units (ICUs). Our results show that NeurABM not only identifies importation cases but also forecasts future nosocomial infection cases better than other machine learning or modeling-based baselines. Note that the NeurABM is a general framework that integrates both neural networks and mechanistic models in an end-to-end way, one which can be easily extended to other ABMs or EHR data and study other clinical problems.

## Results

EHR data was used from the UVA hospital to construct patient contact networks (used by the ABM) and collect patient risk factors (used by the neural network). We use the SIS-ABM model<sup>39,40</sup> as the ABM for disease transmission in NeurABM. Ground-truth MRSA infections are identified from lab test results for each patient in the EHR. For each week  $k$ , contact networks, patient risk factors, and lab test results until week  $k - 1$  were used to train the NeurABM and identify importation cases before week  $k - 1$ . We then ran the SIS-ABM model for 7 more days to infer the infection states of patients for week  $k$ . Two setups were assessed: (1) When identifying current nosocomial infection cases in hospital, contact networks in week  $k$  are used in this process, similar to the approach discussed in previous work<sup>22</sup> for detection of asymptomatic cases, although their ABM is different. They also did not consider the importation problem. (2) When forecasting future nosocomial infection cases, no data in week  $k$  is used, which exactly follows a real-world forecasting setup that no future information is known. For evaluation, we compare NeurABM with a broad range of baselines in machine learning categories (feedforward neural network<sup>41</sup>, decision tree<sup>42</sup>, naive bayes<sup>43</sup>, XGBoost<sup>44</sup>, and Autoencoder+KNN<sup>45,46</sup>), modeling categories (the SIS-ABM model<sup>39,40</sup>, SILI-ABM model<sup>22</sup>), and clinical heuristic categories (length of stay<sup>22</sup>). Note that since the SILI-ABM model<sup>22</sup> is not designed to identify importation cases, we only compare it for identifying nosocomial infection cases.

### Identifying importation cases

In Fig. 2a, we show the precision-recall curves<sup>42</sup> for NeurABM and other methods. Intuitively, precision measures the accuracy of positive predictions, while recall measures the ability to identify all actual positives. Note that in clinical practice, very low precision is not very useful, since this means too many tests and treatments do not help to identify and treat MRSA cases. Therefore, we always expect high recall with not-too-low precision. Following previous work<sup>47</sup>, we consider precision smaller than 0.25 as clinically inapplicable and focus on three important precision levels: 0.25, 0.5, and 0.75 (dashed gray lines). NeurABM consistently achieves the highest recall at precision levels of 0.25, 0.5, and 0.75, demonstrating the effectiveness of the NeurABM framework. For example, a precision of 0.25 means that, on average, one out of four patients identified as importation cases are actually importation cases. Additionally, when precision is 0.25, NeurABM achieves 0.74 recall, while the best baseline (feedforward neural network) only achieves a recall of 0.40. This means that NeurABM could identify 55 out of 74 importation cases in UVA ICU, while the feedforward neural network could only identify 29 out of 74, missing 26 more patients than NeurABM. Finally, the area under the precision-recall curve (AUPRC) for NeurABM is the largest (0.60) among all methods. In Fig. 2b, we show how the negative predictive value (NPV) changes with threshold changes (when the estimated probability is higher than the threshold, we classify this patient as a MRSA importation case, and vice versa). The NPV is the fraction of the

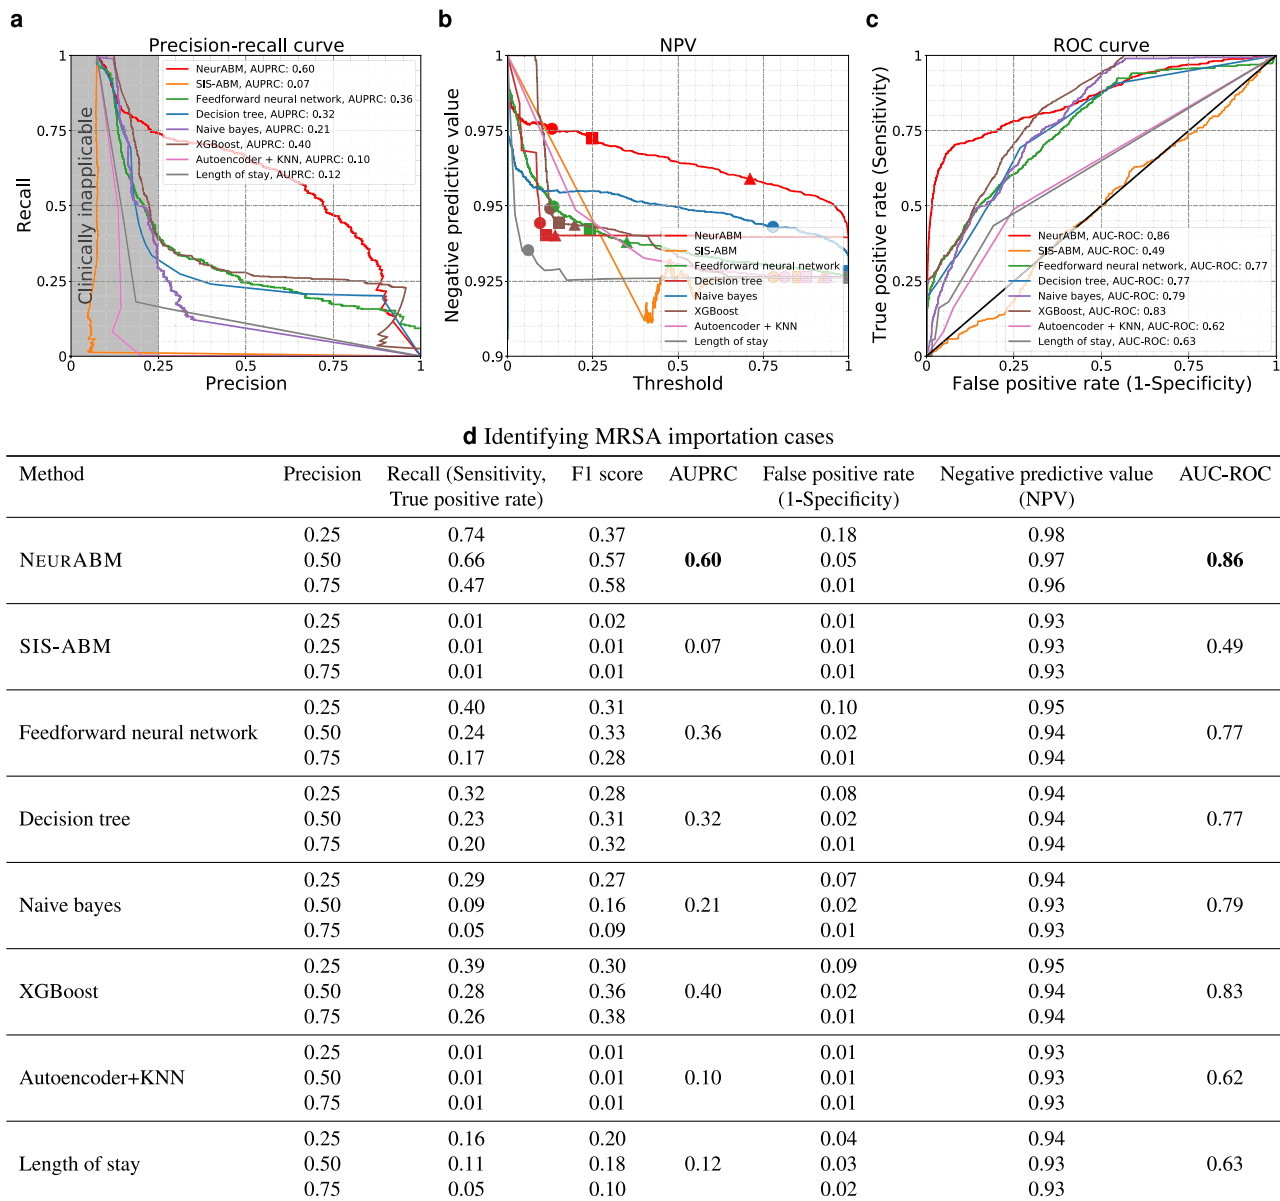
number of true negative cases over predicted negative cases. Intuitively, a higher NPV means that when we predict a patient as negative, he or she is less likely to be a false negative patient that we fail to identify. We can see that NeurABM's NPV rate is always higher than 0.9 and other baselines, indicating that NeurABM can identify the importation cases well with fewer missing/undetected patients. Figure 2c presents the receiver operating characteristic (ROC) curve<sup>42</sup>, which illustrates the trade-off between the true positive rate (the proportion of actual positives correctly identified, which is recall) and the false positive rate (the proportion of negatives incorrectly classified as positives, which is 1-specificity) across different classification thresholds. Here, the area under the curve (AUC-ROC) for our framework is still the largest (0.86) compared to other baselines. Intuitively, an AUC-ROC of 0.86 means that, on average, there is an 86% probability that NeurABM will predict higher importation probabilities for an importation case compared to a non-importation case. In the table in Fig. 2d, we also list the recall, F1 score, and false positive rate corresponding to different precision values. NeurABM consistently achieves the highest recall and F1 score with a given precision, indicating the effectiveness of our framework in identifying MRSA importation cases.

### Identifying current nosocomial infection cases

Next, we assessed the effectiveness of NeurABM to identify current nosocomial infection cases of MRSA. As shown in Fig. 3a, the x-axis and y-axis represent precision and recall, respectively. The red curve represents the results for NeurABM. As shown in the figure, the area under the precision-recall curve for our framework is the largest (0.69) compared to other baselines. The dashed gray lines correspond to the precision of 0.25, 0.5, and 0.75. Again, NeurABM consistently achieves the highest recall with precision equal to 0.25, 0.5, and 0.75, indicating that our framework is effective. For example, a precision of 0.5 means that, on average, one out of two patient days we identified are actual nosocomial infection patient days. Additionally, when precision is 0.5, NeurABM achieves 0.81 recall, while the best baseline (XGBoost) only achieves a recall of 0.60. This means that NeurABM was able to identify 331 out of 408 nosocomial infection patient days in UVA ICU, while XGBoost could only identify 245 out of 408, missing 86 more patient days than NeurABM. In Fig. 3b, we show how the NPV changes with the threshold for classification. We can see that the NPV rate is always higher than 0.9 and other baselines, indicating that NeurABM can identify current nosocomial infection cases well with fewer missing/undetected patients. The ROC curve in Fig. 3c demonstrates that the AUC-ROC for the NeurABM framework is the highest (0.88) compared to other baselines. Intuitively, this means there is an 88% probability that NeurABM will assign higher probabilities to actual current nosocomial infection cases compared to non-nosocomial infection cases. In the table in Fig. 3d, NeurABM always achieves the highest recall and F1 score with a given precision, demonstrating the effectiveness of our framework in identifying current nosocomial MRSA infection cases.

### Forecasting future nosocomial infection cases

We then directed our attention to examining the ability of NeurABM to identify current nosocomial infection cases of MRSA. As shown in Fig. 4a, the x-axis and y-axis represent precision and recall, respectively. The red curve represents the results for NeurABM. As shown in the figure, the area under the precision-recall curve for our framework is the largest (0.85) compared to other baselines. The dashed gray lines correspond to the precision of 0.25, 0.5, and 0.75. Again, NeurABM always achieves the highest recall with precision equal to 0.25, 0.5, and 0.75, indicating that our framework is effective. For example, when precision is 0.75, NeurABM achieves 0.88 recall, while the best baseline (Autoencoder+KNN) only achieves a recall of 0.5. This means that we could forecast 275 out of 312 future nosocomial infection patient days in UVA ICU, while the Autoencoder+KNN could only forecast 156 out of 312, missing 119 more patient days than NeurABM. In Fig. 4b, we show how the NPV changes with the threshold for classification. We can see that the NPV



**Fig. 2 | The performance in identifying importation cases.** **a** The precision-recall curves (PRC). The x-axis represents precision, and the y-axis represents recall. The red and other color curves represent NeurABM and other baselines. A larger area under the precision-recall curve (AUPRC) indicates better performance. AUPRC values are listed in the legends, and NeurABM has the highest AUPRC value. **b** The negative predictive value (NPV) with different thresholds. The x-axis is the threshold for classification, and the y-axis is the NPV value. Circles, squares, and triangles correspond to the thresholds and NPV values where precision is 0.25, 0.5, and 0.75, respectively. A higher NPV value indicates fewer missing importation cases that are

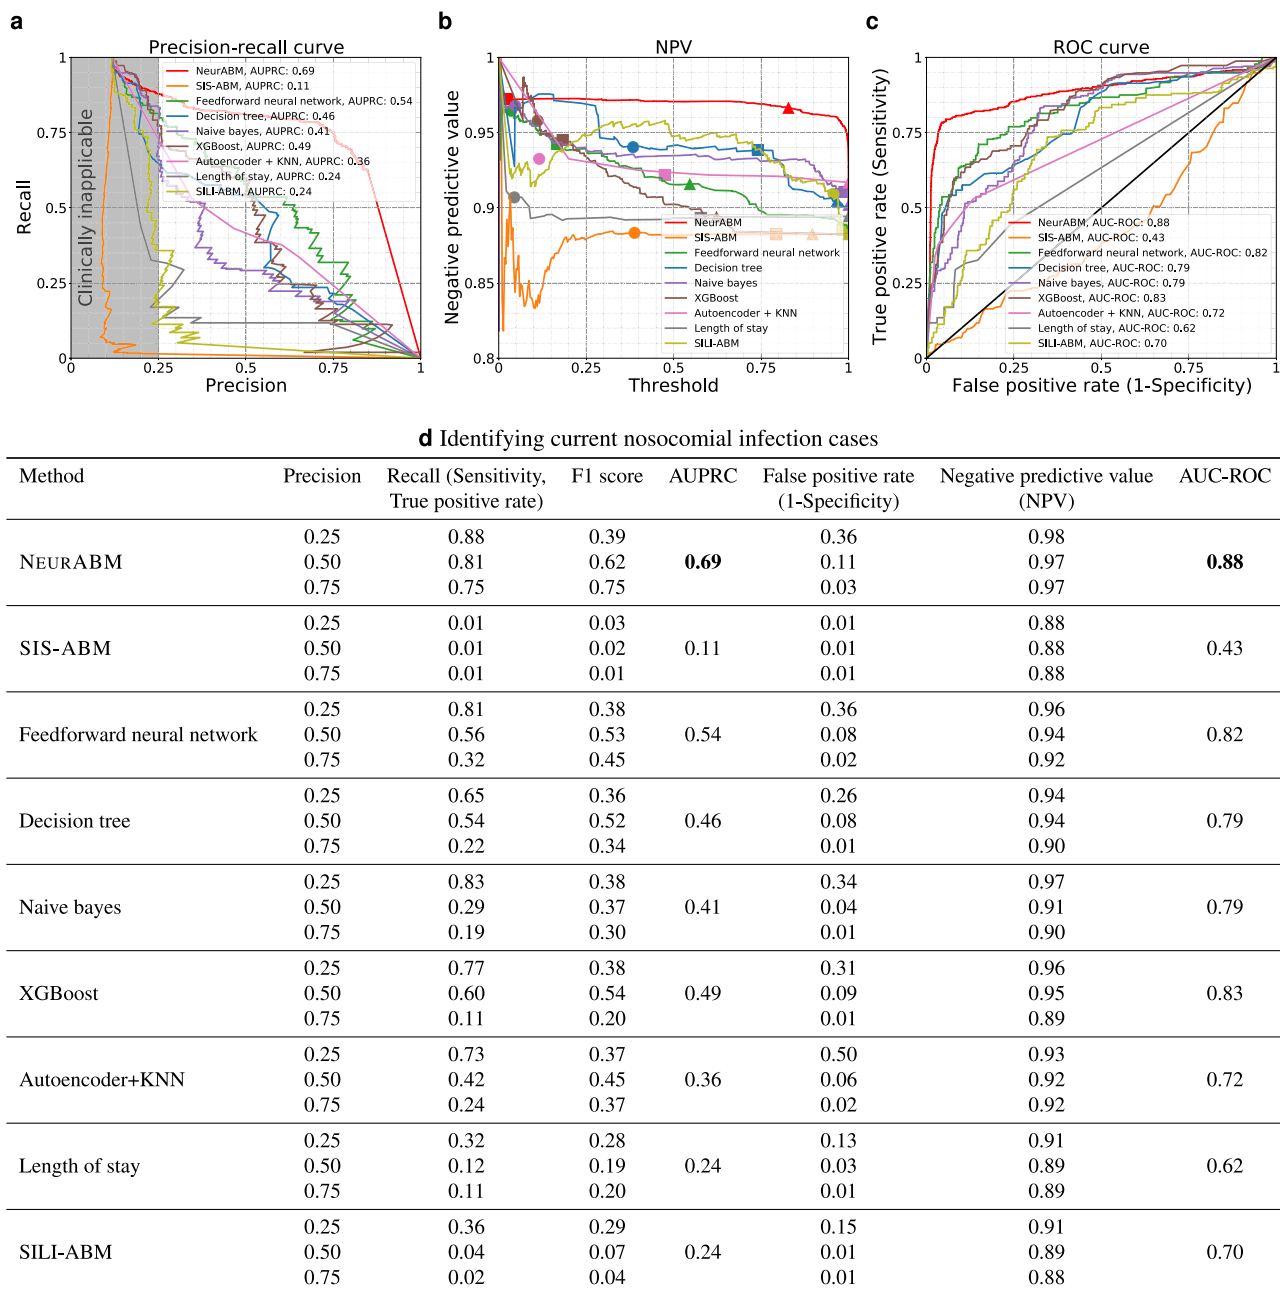
not identified and therefore better performance, and NeurABM has the highest NPV values. **c** The receiver operating characteristic (ROC) curves in identifying MRSA importation cases. The x-axis is the false positive rate, and the y-axis is the true positive rate. A larger area under the ROC (AUC-ROC) indicates better performance. AUC-ROC values are listed in the legends, and NeurABM has the highest AUC-ROC value. **d** The recall, F1 score, AUPRC, false positive rate, NPV, and AUC-ROC under different precisions (0.25, 0.5, 0.75). The best AUPRC and AUC-ROC are in bold.

rate is always higher than 0.9 and other baselines, indicating that NeurABM can forecast future nosocomial infection cases well with fewer missing/undetected patients. In Fig. 4c, the AUC-ROC for NeurABM framework is the largest (0.92) compared to other baselines. Intuitively, this means there is a 92% probability that NeurABM will assign higher probabilities to future nosocomial infection cases compared to non-nosocomial infection cases. In the table in Fig. 4d, NeurABM always achieves the highest recall and F1 score with a given precision, demonstrating the effectiveness of our framework in forecasting future nosocomial MRSA infection cases. We also ran experiments to forecast for 2 weeks ahead (see Supplementary Information), and the results show that NeurABM still performs better than other baselines.

### Forecasting high-risk MRSA cases

Identifying MRSA cases in high-risk patients can help target infection control interventions that reduce transmission. Following the approach of Pei et al.<sup>22</sup>, we consider a strategy of testing patients based on the ranked carriage probability forecasted by NeurABM and other baselines, and determined the proportion of MRSA cases that can be identified. This is shown in Figure 5, where we rank all patients according to the forecasted infected probability of each method from high to low, and test patients in this order. For example, if we only test 20% of patients based on the ranked risk estimated by NeurABM, we could identify 88% (86 out of 98) of carriage patients. Meanwhile, the best baseline (XGBoost) could only identify 65% (64 out of 98) of them (i.e., 22 less patients than NeurABM). Therefore,





**Fig. 3 | The performance in identifying current nosocomial infection cases.** **a** The precision-recall curves. The red and other color curves represent NeurABM and other baselines. Higher AUPRC is better, and NeurABM has the highest AUPRC value. **b** The negative predictive value with different thresholds. Circles, squares, and triangles correspond to the thresholds and NPV values where precision is 0.25, 0.5, and 0.75, respectively. Higher NPV value is better, and NeurABM has the highest

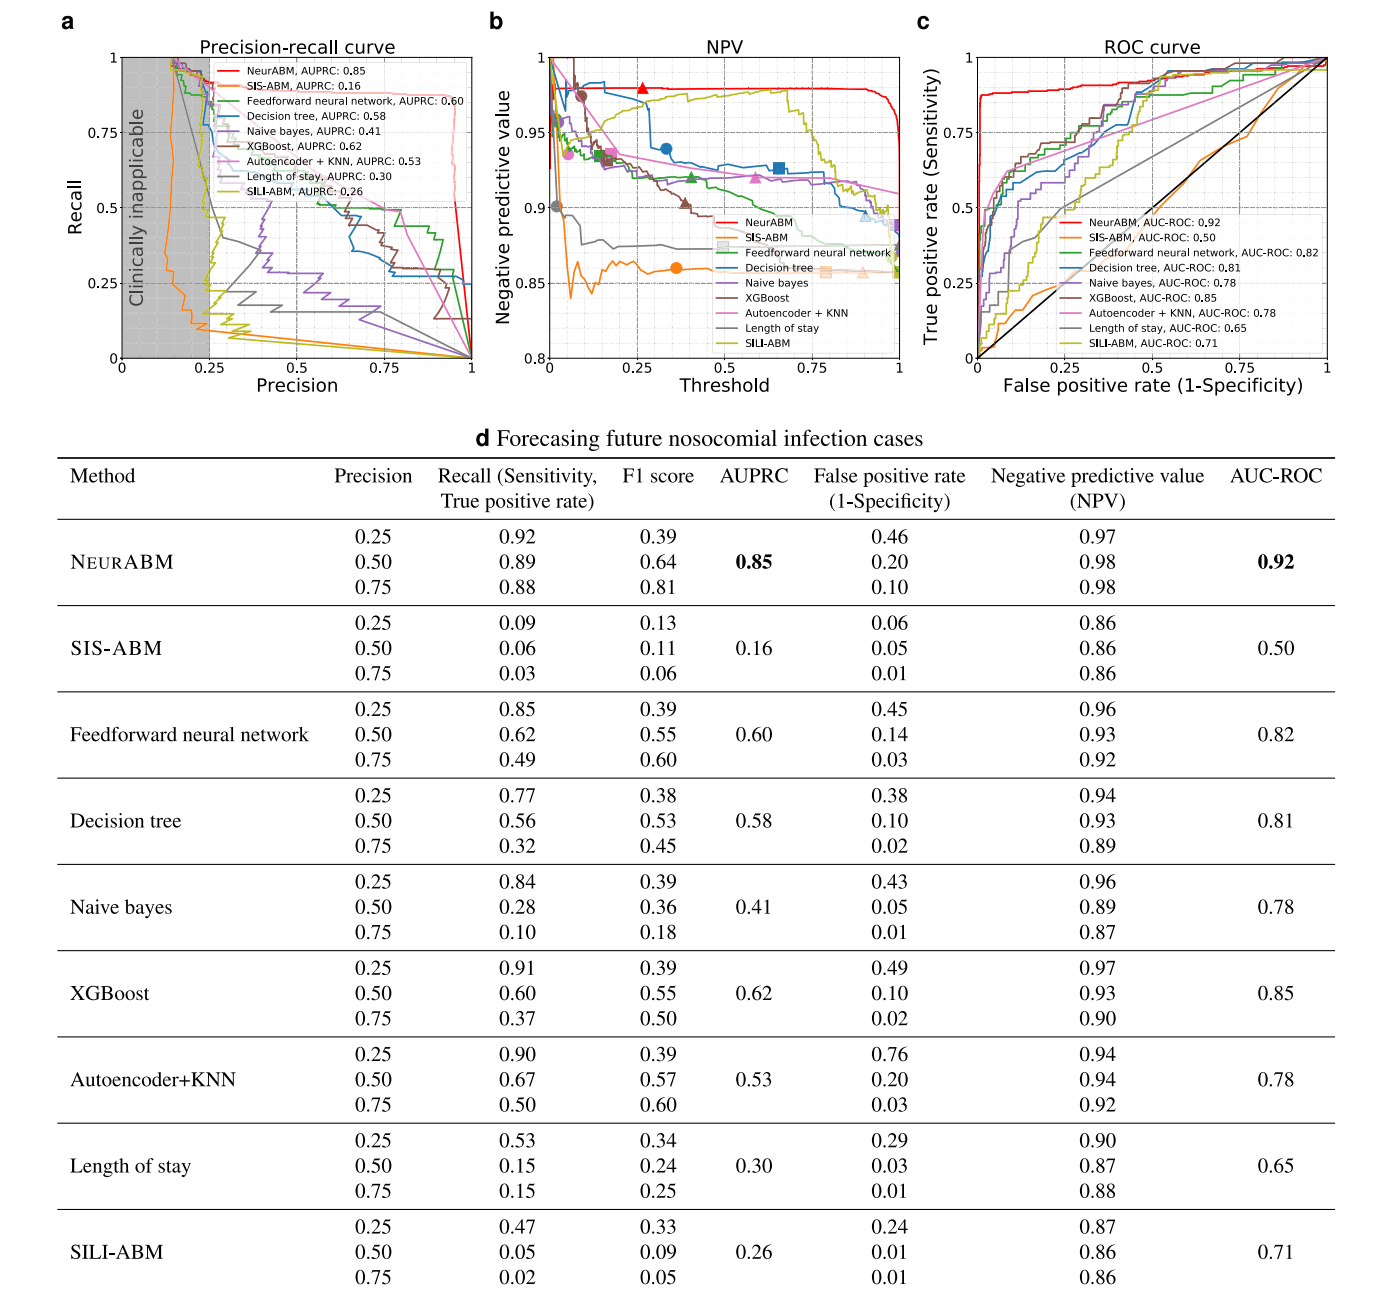
NPV values. **c** The receiver operating characteristic curves in identifying MRSA nosocomial infection cases. Higher AUC-ROC is better, and NeurABM has the highest AUC-ROC value. **d** The recall, F1 score, AUPRC, false positive rate, NPV, and AUC-ROC under different precisions. The best AUPRC and AUC-ROC are in bold.

NeurABM can always forecast more future nosocomial MRSA infection cases (y-axis) given the same test budget (x-axis), which suggests that our framework is effective and practical in real-world clinical settings.

### Case study: explanation of neural network

To better demonstrate the interpretability of the NeurABM framework, we further investigated which patient risk factors are considered to have a high risk of being importation cases in the previous results shown in Fig. 2 by NeurABM. As shown in Fig. 6, each dot represents a patient, and the color of the dot represents the risk factor value (red means higher, and blue means lower). Dots with a higher impact value (to the right of the control line) mean that NeurABM tends to classify this

patient as having a higher probability of being an importation case (and vice versa). Here, we can see that patients who have had contact with more MRSA patients in the past 7/14 days prior to ICU admission are considered as having high risk. In addition, we found that patients with a device usage history, and who come from or were discharged to other healthcare facilities last time, are more likely to be considered as importation cases. This also aligns with real-world observations from clinicians<sup>48</sup>. It is important to note that the SHAP values are primarily used to explain the outputs of a neural network, and the results presented in Fig. 6 are intended to highlight the interpretability of the NeurABM framework rather than providing direct clinical recommendations or supporting causal conclusions.

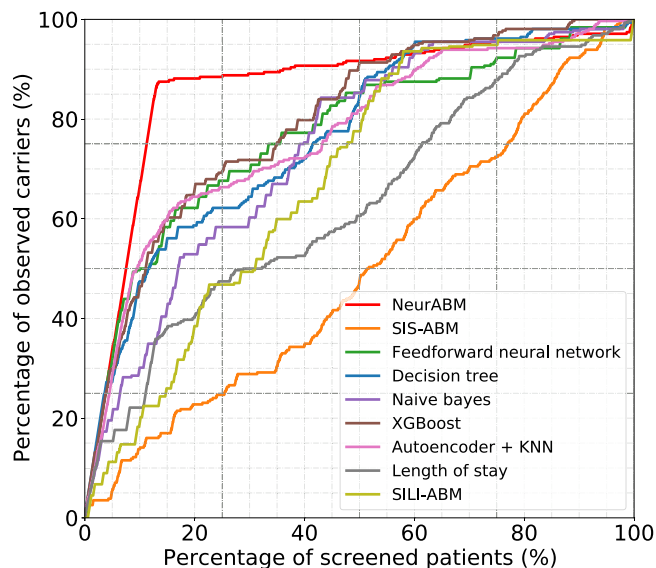


**Fig. 4 | The performance in forecasting future nosocomial infection cases. a** The precision-recall curves. The red and other color curves represent NeurABM and other baselines. Higher AUPRC is better, and NeurABM has the highest AUPRC value. **b** The negative predictive value with different thresholds. Circles, squares, and triangles correspond to the thresholds and NPV values where precision is 0.25, 0.5, and 0.75, respectively. Higher NPV value is better, and NeurABM has the highest NPV values. **c** The receiver operating characteristic curves in forecasting future MRSA nosocomial infection cases. Higher AUC-ROC is better, and NeurABM has the highest AUC-ROC value. **d** The recall, F1 score, AUPRC, false positive rate, NPV, and AUC-ROC under different precisions. The best AUPRC and AUC-ROC are in bold.

Discussion

Identifying healthcare-associated infections importation and forecasting future nosocomial infection cases have long been a challenging problem. Previous research works relied either solely on machine learning and statistical techniques, which directly use patients’ risk factors collected from EHR data for prediction while overlooking the epidemiological mechanisms involved in HAI spread<sup>20,21</sup>, or solely on mechanistic models that capture the transmission dynamics of HAI within healthcare facilities while ignoring individual patients’ specific risk factors<sup>22</sup>. In contrast, NeurABM integrates both deep learning and ABMs simultaneously in a lock-step to learn importation probabilities. This framework allows the use of diverse data sources from EHRs by deep learning methods and ABMs. Furthermore, the


contact network utilized in the SIS-ABM model within NeurABM plays a crucial role and underscores the importance of contact representations in identifying nosocomial infections. Overall, NeurABM is a novel framework for supporting a diverse class of HAI surveillance and control questions. Our experiments show that NeurABM not only identifies MRSA importation cases and current nosocomial cases but also forecasts future nosocomial infection cases in UVA ICUs, with performance that can be considered clinically useful (e.g., high recall with not-too-low precision as shown in Fig. 2a, Fig. 3a, and Fig. 4a); in contrast, as shown in experiments, prior methods using EHR data haven’t achieved this level of performance. For example, when ranking the patients based on their probability in the carriage state and test the top 20%, we could identify 88% of carriage patients



**Fig. 5 | Percentage of identified MRSA cases by screening “high-risk” patients.** For each patient in the UVA ICUs, we use each method to estimate their MRSA infection probability and rank them according to this probability from high to low. We then screen different percentages of patients (x-axis) and see how many actual MRSA cases can be forecasted (y-axis). As seen in the figure, NeurABM can always identify more MRSA cases than other baselines.

while the best baseline could only identify 65% of them (86 out of 98 vs. 64 out of 98). Moreover, the NPV for our method is always higher than 0.9, indicating that NeurABM can identify and forecast these cases well with fewer missing/undetected patients. Even though importation cases are identified retrospectively, the results remain practical and valuable. Typically, MRSA culture tests require a 3–5 day turnaround time, while the NeurABM framework can provide a rapid estimation during this period, potentially offering actionable insights before culture results are available. Our case study also reveals the risk factors that are highly related to MRSA importations, which confirms clinical observations and allows clinicians to better respond to potential MRSA importation cases in their hospital.

However, our framework is not without limitations. One limitation is that the SIS-ABM model may oversimplify MRSA carriage by defining just two states: *susceptible* and *carriage*. There are different forms of MRSA carriage, including various clinical infection types such as skin abscess or bloodstream infection, which may confer different levels of risk for transmission. Similarly, the SIS-ABM model does not have an explicit “colonization” state. More advanced epidemiological models like multi-strain SIR models could more closely reflect realistic scenarios and further improve performance. Another limitation is that we are currently only using the co-location data to construct the contact networks and assuming that the entities co-located in the same room will have interactions with each other. Recent works have shown that this well-mixed assumption, even for short periods of time or small spatial locations, can be sub-optimal in terms of prediction, while better-designed, norm-based spatial formalization can better capture the interactions between entities in continuous space<sup>36,37</sup>. Despite the above disadvantages, NeurABM is quite general and can be extended to use more complex ABMs that incorporate these properties separately, depending on the needs in the application and data availability. Besides, MRSA nares testing is widely used in clinical practice to guide decision-making on whether contact precautions and MRSA-active antibiotics should be implemented, and agent-based models (such as the SIS-ABM model used in this paper) are well established for studying MRSA transmission<sup>39,40</sup>. Therefore, extending agent-based methods to predict MRSA testing outcomes for potential future clinical application is a natural progression of this work, aligning with their



**Fig. 6 | The patient risk factors that are considered as having a high risk of being importation cases by the trained neural network.** Color of the dots represent risk factor values, with red indicating higher values. A higher impact means that NeurABM is more likely to consider the patient as an importation case.

established utility in modeling transmission dynamics. Additionally, we are currently evaluating NeurABM only in the UVA ICU due to data availability; adding more patient risk factors (e.g., chronic diseases history or Apache II scores) or conducting more experiments in other settings and healthcare facilities could provide a more comprehensive assessment of NeurABM.

Nevertheless, NeurABM opens up a new direction of research for using both rich patient risk factors from EHR data as well as agent-based models designed with epidemiological knowledge and can be used for many other questions about HAI spread in hospitals. The specific architecture of our method, which runs the deep learning method and ABM step at each time allows us to add predictors at different stages within this framework to address other kinds of questions. For instance, people can also forecast future contact networks instead of naively assuming future contact networks may be the same as previous ones to better forecast future nosocomial infections. The ABM can also be enriched with representations of interventions implemented in the hospital (e.g., how patients are selected and put under contact precautions with identified with MRSA). For patient-specific parameters, we focus on importation probabilities in this work. However, many other parameters such as recovery rate can also be patient-specific and our NeurABM framework can be easily extended to it. The performance of these variations depends on data availability, and is a promising topic for future research.

In terms of real-world applications, the NeurABM framework is highly flexible, as it is not tied to any specific agent-based model. Instead, users can easily replace the SIS-ABM model in NeurABM with their own ABM designs into NeurABM, allowing people to adapt to other healthcare-associated infections, such as *Clostridioides difficile*. Moreover, the datasets required for NeurABM are straightforward to extract from electronic health records<sup>22</sup> (see the Methods section for more details). Our experiments also adhere to a real-world setup: when forecasting future nosocomial infections in week  $k$ , the model only requires data up to week  $k - 1$ . This ensures that

NeurABM operates without data leakage and is suitable for real-time implementation. The fact that NeurABM is amenable to such adaptations quite easily is an indication of its generality.

## Methods

### Dataset

We extract three different types of patient data based on the electronic health records (EHR) from the University of Virginia hospital: patient demographic information and risk factors (e.g., comorbidities, medical history), lab testing, and contact network data.

- **Patient risk factor data:** This dataset consists of risk factors for all patients in ICUs. From the EHR dataset, we collected 19 different risk factors for each patient, all of which are available before ICU admission. From July 1, 2019, to December 31, 2019, there were 1117 patients in UVA ICUs, and 74 of them were MRSA importation cases (all patients received an MRSA test within  $(t - 3, t + 3)$  days of being admitted into one of the ICUs, and patients who tested positive for MRSA within this range are considered as importation cases). A list and description of each risk factor are provided in the Supplementary Information.
- **Lab testing data:** This dataset consists of infection data for each patient. There were two different types of tests to diagnose MRSA: culture tests and polymerase chain reaction (PCR) tests. However, since a negative culture test cannot disqualify an individual from MRSA infection, we focused on only positive culture tests and both positive and negative PCR tests. For a given patient  $p$  on a given day  $t$ ,  $y_{p,t} = 1$  represents that the patient was tested positive on day  $t$  or if their most recent test in the past was positive. Likewise,  $y_{p,t} = 0$  if this patient was tested negative on day  $t$  or if their most recent test was negative.
- **Contact network data:** This dataset consists of a series of ward-level co-location contact networks  $A_t$  comprising three different entities: patients, healthcare workers (HCWs), and locations, and each network is for one specific day. From the EHR dataset, we can collect the movement information of patients and HCWs (e.g., the ward that a patient stayed, and when the doctors and nurses visited a specific ward). Note that these movement information also includes start and end times; we can infer whether these patients and HCWs were co-located (i.e., time overlapped) at any specific location. Specifically, if two patients or HCWs  $v_1, v_2$  colocated at location  $l$  on day  $t$ , we would create edges between  $v_1$  and  $v_2$ ,  $v_1$  and  $l$ ,  $v_2$  and  $l$  on day  $t$  in  $A_t$ . However, because of the nature of this data, individuals such as support staff or patient guests are not tracked, and thus are not included in the network. Additionally, HCW-HCW colocations are not tracked in rooms where care is not administered, such as break rooms. From July 1, 2019 to December 31, 2019, there were 1117 patients (and 6445 patient days), 2385 healthcare workers, and 77 locations in the UVA ICU.

### Problem setup

We use the trained NeurABM model to identify importation and nosocomial infection cases. For importation cases, the patient-specific parameter  $\theta_p$  in this work is importation probability for each patient. Specifically, for each week  $k$ , we used the contact networks, patient risk factors, and lab testing results until week  $k$  to train the NeurABM, and our task is to identify the importation cases from them until week  $k$ . Note that NeurABM framework do not access to the ground-truth importation cases data. Instead, we only use the ground-truth importation cases data for evaluation.

For identifying current nosocomial infection cases in hospital, we follow the setup of a previous work<sup>22</sup>. In this setup, since we want to identify current nosocomial infections in week  $k$ , we should be exactly in week  $k$  (e.g., at the end of week  $k$ ). Therefore, for each week  $k$ , we used the contact networks, patient risk factors, and lab testing results until week  $k - 1$  to train the NeurABM, and then ran the SIS-ABM model for 7 more days to infer the infection states of all patients for week  $k$ . Since we are identifying current nosocomial infection cases in week  $k$ , the contact network in week  $k$  is known, while the lab test results are not fully available (because of the delays in getting the lab test results). For example, if we were at the end of week 40

(beginning of October), we would train NeurABM on all the data from week 28 (beginning of July) to week 39. As for week 40, we would only use the contact network in that week and use the ABM simulator to identify the nosocomial infection cases in week 40. Then, at the end of week 41, we would train on the data from week 29 to 40 and identify the nosocomial infection cases in week 41. We repeated this procedure until we were at the end of week 52.

For forecasting future nosocomial infection cases, we followed a real-world step-forward scenario that made weekly predictions. In this setup, since we want to forecast future nosocomial infection cases in week  $k$ , we should be before week  $k$  (e.g., at the end of week  $k - 1$ ). Therefore, for each week  $k$ , we used the contact networks, patient risk factors, and lab testing results until week  $k - 1$  to train the NeurABM, and then ran the SIS-ABM model for 7 more days to forecast the infection states of all patients for week  $k$  without knowing any information for week  $k$ . For example, if we were at the end of week 39, we would train on the data from week 28 to week 39 to forecast for week 40 (i.e., no information from week 40 is used). Then, at the end of week 40, we would train on the data from week 29 to 40 and forecast for week 41. We repeated this procedure until we were at the end of week 52. Intuitively, the key difference between the setup of identifying current nosocomial infection cases and forecasting future nosocomial infection cases is whether the contact network in the targeted week  $k$  is known or not.

### Transmission model

In this work, we use the SIS-ABM model in NeurABM to capture the MRSA spread dynamics in UVA ICUs<sup>39</sup>. SIS-ABM is a pathogen load-based model that keeps track of pathogen load on all people and locations using a load vector  $I_t$ . For each patient  $i$ , they can either be in the *Susceptible* ( $S$ ) or *Carriage* ( $C$ ) state. Specifically, the probability of transitioning from  $S$  to  $C$  is proportional to the amount of pathogen on this patient  $I_t(i)$ , which can be formulated as a linear dose-response function  $\beta I_t(i)$  ( $\beta$  is the disease infectivity parameter). Once in the carriage state, the patient keeps shedding more pathogen loads at each step, which can later be transferred to its neighbors (including both people and locations). Such a shedding process continues until the patient recovers with a recovery probability  $\delta$ .

For the pathogen load transfer, as described in the previous text, the SIS-ABM model uses daily contact networks  $A_t$  to capture the exchange of pathogens among patients, HCWs, and locations. Specifically, we construct a transfer matrix  $R_t$  for each day  $t$ , where  $R_{ijt} = \tau_{ijt} A_{ijt}$ . Here,  $A_t$  is the adjacency matrix of contact networks on day  $t$ , and  $\tau_{ijt}$  is the transfer ratio parameter (the ratio of pathogen being transferred (or remaining if  $i = j$ ) from patient/HCW/location  $j$  to patient/HCW/location  $i$  on day  $t$ ). Specifically, based on the kinds of nodes of  $i$  and  $j$ , we have 8 kinds of transfer ratios:  $\tau_{PP}, \tau_{PH}, \tau_{PL}, \tau_{HP}, \tau_{HH}, \tau_{HL}, \tau_{LP}, \tau_{LH}$ . It also uses  $\gamma_P, \gamma_H, \gamma_L$  to denote the natural pathogen reduction rate on patient, HCW, and location nodes. Using this  $R_t$  and  $I_t(i)$ , the SIS-ABM model updates the pathogen loads every day as a linear operation. We also restrict the column-sums of  $R_t$  to be less than or equal to 1, which implies that the total amount of pathogen cannot increase after transfer (i.e.,  $|R_t I_t| \leq |I_t|$ ). Note that susceptible patients may still carry a small amount of pathogen loads and spread them to others, and HCWs and locations are always in the *susceptible* state, which means that they can spread the MRSA pathogen loads but are non-infectable. We provide more details in the Supplementary Information.

### NeurABM framework

As shown in Fig. 1, the NeurABM framework is composed of two parts: the neural network part (green block, parameterized by  $\phi$ ) and the agent-based model simulator part (yellow blocks).

For the neural network component, we take the risk factors  $f$  (where  $f_p$  is for patient  $p$ ) as input and then estimate both the patient-specific parameters  $\theta_p$  (which is a vector and each element  $\theta_p$  is for patient  $p$  and only influenced by the patient itself's risk factors  $f_p$ , in this work it is importation probability for each patient) and ABM parameters  $\theta_M$  (i.e., the general parameters that apply to every patient in the SIS-ABM model, including disease infectivity parameter  $\beta$ , recovery probability  $\delta$ , pathogen shedding



rate  $\alpha$ , natural pathogen reduction rate for patients, HCWs, and locations  $\gamma_P$ ,  $\gamma_H$ ,  $\gamma_L$ , transfer ratios from different kinds of nodes  $\tau_{PP}$ ,  $\tau_{PH}$ ,  $\tau_{PL}$ ,  $\tau_{HP}$ ,  $\tau_{HH}$ ,  $\tau_{HL}$ ,  $\tau_{LP}$ ,  $\tau_{LH}$  together. The neural network is parameterized by  $\phi$  and we use  $\theta_p$ ,  $\theta_M = \text{NN}(f; \phi)$  to represent it. We list all ABM parameters in Supplementary Information.

For the ABM simulator, we implement the simulation process of the SIS-ABM model using matrix operations in a differentiable way. ABM simulator takes the adjacency matrices of contact networks  $A_t$  and the parameters learned by neural networks ( $\theta_p$ ,  $\theta_M$ ) as input, and simulates MRSA spread in  $T$  days to estimate patient states on each day  $\hat{y}$  (where  $\hat{y}_{p,t}$  is for patient  $p$  on day  $t$ ). Specifically, the simulation process of the SIS-ABM model can be decomposed into three substeps: (1) pathogen load transmission where the load transfers via contact edges, (2) updating the states for each patient based on their pathogen loads and recovery probability, and (3) updating the timestep from day  $t$  to day  $t + 1$ . This process can be repeated for arbitrary steps to simulate the MRSA spread over  $T$  days. We use  $\hat{y} = \text{ABM}(A; \theta_p, \theta_M)$  to represent it.

We also have another ABM simulator for forecasting. Since we cannot exactly know future contact networks, we assume they will be the same as the contact network on day  $T$  (i.e.,  $A_T$ ). Therefore, it takes the patient states on day  $T$  as the input and forecast for  $T'$  days ahead. It can be represented by  $\hat{y}'_F = \text{ABM}_F(\hat{y}; A_T, \theta_p, \theta_M)$ . However, the assumption on the future contact networks is inaccurate and lead to bias in  $\hat{y}'_F$  we forecast. Therefore, we use another adapter neural network  $\phi'$  to revise  $\hat{y}'_F$  and give our final forecast  $\hat{y}_F$ . We formulate it as  $\hat{y}_F = \text{Adapter}(\hat{y}'_F, f; \phi')$ .

With the aforementioned neural network and the ABM simulator, we then integrate them together to train simultaneously. Specifically, one training epoch comprises the following steps.

- Step 1: We feed the risk factor data  $f$  into the neural network as the input to estimate the patient-specific parameters  $\theta_p$ . In this work,  $\theta_p$  is the probability of being importation cases for each patient  $p$ . It is a vector of size  $N$ , where  $N$  is the number of patients in the contact network. Meanwhile, the neural network will also give the general ABM parameters  $\theta_M$  that are applied to all patients (e.g.,  $\beta$ ,  $\alpha$ , ...).
- Step 2: We then feed  $\theta_p$ ,  $\theta_M$ , and the contact networks  $A_t$  into the ABM simulator and simulate for  $T$  steps. The output will be the vector  $\hat{y}$  of size  $N \times T$ , in which  $\hat{y}_{p,t}$  represents the probability of being in the state carriage for patient  $p$  on day  $t$ .
- Step 3: We compare the estimated carriage probability  $\hat{y}$  with the corresponding ground-truth observations (i.e., known carriage patients based on lab testing)  $y$ . We use the weighted binary cross entropy loss (BCE loss)<sup>42</sup>  $\mathcal{L}(\hat{y}, y) = \sum_p \sum_t w_{pos} y_{p,t} \log(\hat{y}_{p,t}) + w_{neg} (1 - y_{p,t}) \log(1 - \hat{y}_{p,t})$  as the loss function. Here  $w_{pos}$  and  $w_{neg}$  are the weights for positive and negative observations. We set  $w_{pos} : w_{neg} = \sum_p \sum_t \mathbb{1}[y_{p,t} = 0] : \sum_p \sum_t \mathbb{1}[y_{p,t} = 1]$ , where  $\mathbb{1}[\cdot]$  is the indicator function, which is 1 if the condition is true, and 0 otherwise.
- Step 4: Meanwhile, we also feed the  $\hat{y}_{p,t}$  for day  $T$  to another ABM simulator to forecast for  $T'$  days ahead. Note that these  $T'$  days are for future and we cannot access to the real contact networks when forecasting, we assume that future contact networks will be the same as what we have for day  $T$ . Let  $\text{ABM}_F$  be this ABM simulator for forecasting, and  $\hat{y}'_F$  of size  $N \times T'$  as rough forecast output (in which  $\hat{y}'_{p,t+T'}$  represent the probability of being in the state carriage for patient  $p$  on day  $t + T'$ ), we formulate it as  $\hat{y}'_F = \text{ABM}_F(\hat{y}; A_T, \theta_p, \theta_M)$ .
- Step 5: However, the assumption that future contact networks will be the same as on day  $T$  is inaccurate and leads to bias in  $\hat{y}'_{p,t+T'}$ . To tackle this, we use another adapter neural network  $\phi'$  to revise this rough output to get the revised forecast output  $\hat{y}_F$  of size  $N \times T'$ , in which  $\hat{y}_{p,t+T'}$  represents the probability of being in the state carriage for patient  $p$  on day  $t + T'$ . We formulate it as  $\hat{y}_F = \text{Adapter}(\hat{y}'_F, f; \phi')$  and compute the BCE loss  $\mathcal{L}(\hat{y}_F, y_F) = \sum_p \sum_t w_{pos} y_{p,t+T'} \log(\hat{y}_{p,t+T'}) + w_{neg} (1 - y_{p,t+T'}) \log(1 - \hat{y}_{p,t+T'})$ .
- Step 6: With the total loss  $\mathcal{L}(\hat{y}, y) + \mathcal{L}(\hat{y}_F, y_F)$ , and the differentiable ABM simulator, we can calculate the gradient of the loss with respect to

the neural network parameters  $\phi$  and adapter parameters  $\phi'$  via backpropagation. This allows us to better tune the neural network and learn more reasonable parameters as the input for the ABM simulator.

The above steps are repeated until the total loss converges. Sometimes, directing the whole NeurABM framework can be extremely hard, since the framework can be really deep. Intuitively, if we use  $T$  days for ABM calibration and forecast for  $T'$  days ahead, the depth of the framework becomes at least  $T + T'$ . To address this problem, an iterative training approach following previous work can be helpful: We can iteratively freeze different parts of the frameworks and train the other parts, allowing us to simplify the training NeurABM by breaking it down into two substeps: (1) First, we focus on training the neural network and the ABM simulator while keeping the adapter and ABM simulator for forecasting frozen, which allows us to first learn reasonable ABM parameters. (2) We then focus on the adapter and the ABM simulator for forecasting while keeping others frozen, which allows us to learn reasonable parameter for adapter. These two substeps above can be repeated until convergence. More details are provided in the Supplementary Information.

## Baselines

To compare NeurABM with current modeling or machine learning-based methods, we also compare with other baselines including machine learning-based methods (neural network<sup>41</sup>, decision tree<sup>42</sup>, naive bayes<sup>43</sup>, XGBoost<sup>44</sup>, Autoencoder+KNN<sup>45,46</sup>), mechanistic modeling-based methods (SIS-ABM model<sup>39,40</sup>, SILI-ABM model<sup>22</sup>), and clinical heuristic methods (length of stay<sup>22</sup>).

For machine learning based methods, we train two models: one for identifying importation cases and another for identifying nosocomial infection cases. For importation cases, we train on the ground-truth importation cases from January 2019 to June 2019 and test on July to December (i.e., the same time period for NeurABM). Note that NeurABM identifies importation cases without access to ground-truth data from January 2019 to June 2019. In contrast, the machine learning-based baselines require ground-truth labels as training data. To enable these baselines to function, we provided this additional information, yet they still could not outperform NeurABM. For identifying current and forecasting future nosocomial infection cases in each week  $k$ , we train on data until week  $k - 1$  and test on week  $k$ . For modeling-based methods, we run their models following their original papers<sup>22,39</sup> and take the average infected probability of 100 simulations as the probabilities of being importation cases and nosocomial infection cases. For clinical heuristic baselines, the length of stay will consider patients staying longer in the hospital to have higher probabilities.

Since the outputs of NeurABM framework and baseline models are probabilities of being classified as importation or nosocomial infection cases, we applied varying thresholds (ranging from 0 to 1) to convert these probabilities into binary classification outcomes. This allows us to predict whether a patient is an importation or nosocomial infection case when the probability exceeds the threshold, and vice versa. By adjusting the thresholds, we aim to compare different methods more comprehensively. We plotted precision-recall curves (e.g., Fig. 2a), ROC curves (e.g., Fig. 2c), and the changes in negative predictive value (e.g., Fig. 2b) across different thresholds. We also provide a detailed description of the metrics used to evaluate performance in the Supplementary Information.

## Inclusion & Ethics statement

This work is a retrospective observational study and follows the STROBE statement checklist<sup>49</sup>. This study was approved by the University of Virginia Institutional Review Board for Health Sciences Research (IRB-HSR-22410), and the written consent was waived due to the retrospective nature of the study.

## Data availability

The outputs of our model are available on GitHub via link: <https://github.com/AdityaLab/NeurABM>. The electronic health record (EHR) data used in developing the models is not available since it is highly sensitive, and we do not have permission to release it. However, we provide the code, a demo, and a synthetic dataset on GitHub.

## Code availability

The code of our model are also available on GitHub via link: <https://github.com/AdityaLab/NeurABM>. It also contains a detailed readme file to help run the code.

Received: 26 November 2024; Accepted: 19 February 2025;

Published online: 07 March 2025

## References

- Stone, P. W. Economic burden of healthcare-associated infections: An american perspective. *Expert Rev. Pharmacoecon. outcomes Res.* **9**, 417–422 (2009).
- Sharma, A. et al. The cost of contact precautions: A systematic analysis. *Can. J. Infect. Control* **35**, 163–169 (2020).
- Roth, J. A. et al. Direct costs of a contact isolation day: a prospective cost analysis at a swiss university hospital. *Infect. control Hosp. Epidemiol.* **39**, 101–103 (2018).
- Valiquette, L., Chakra, C. N. A. & Laupland, K. B. Financial impact of health care-associated infections: When money talks. *Can. J. Infect. Dis. Med. Microbiol.* **25**, 71–74 (2014).
- Howden, B. P. et al. Staphylococcus aureus host interactions and adaptation. *Nat. Rev. Microbiol.* **21**, 380–395 (2023).
- Lowy, F. D. Staphylococcus aureus infections. *N. Engl. J. Med.* **339**, 520–532 (1998).
- Centers for Disease Control and Prevention. Antibiotic Resistance Threats in the United States (2019).
- Kallen, A. J. et al. Health care-associated invasive mrsa infections, 2005–2008. *JAMA* **304**, 641–647 (2010).
- Fernandez-Gracia, J., Onnela, J.-P., Barnett, M. L., Eguluz, V. M. & Christakis, N. A. Influence of a patient transfer network of us inpatient facilities on the incidence of nosocomial infections. *Sci. Rep.* **7**, 2930 (2017).
- Rocha, L. E. et al. Dynamic contact networks of patients and mrsa spread in hospitals. *Sci. Rep.* **10**, 1–10 (2020).
- Milstone, A. M. et al. Methicillin-resistant staphylococcus aureus colonization and risk of subsequent infection in critically ill children: importance of preventing nosocomial methicillin-resistant staphylococcus aureus transmission. *Clin. Infect. Dis.* **53**, 853–859 (2011).
- European Centre for Disease Prevention and Control. Point prevalence survey of healthcare-associated infections and antimicrobial use in european acute care hospitals: 2011–2012 (2013).
- Suetens, C. et al. Prevalence of healthcare-associated infections, estimated incidence and composite antimicrobial resistance index in acute care hospitals and long-term care facilities: results from two european point prevalence surveys, 2016 to 2017. *Eurosurveillance* **23**, 1800516 (2018).
- Bhalla, A. et al. Acquisition of nosocomial pathogens on hands after contact with environmental surfaces near hospitalized patients. *Infect. Control Hosp. Epidemiol.* **25**, 164–167 (2004).
- Albrich, W. C. & Harbarth, S. Health-care workers: source, vector, or victim of mrsa? *Lancet Infect. Dis.* **8**, 289–301 (2008).
- Snyder, G. M. et al. Detection of methicillin-resistant staphylococcus aureus and vancomycin-resistant enterococci on the gowns and gloves of healthcare workers. *Infect. Control Hosp. Epidemiol.* **29**, 583–589 (2008).
- Smith, D. L., Dushoff, J., Perencevich, E. N., Harris, A. D. & Levin, S. A. Persistent colonization and the spread of antibiotic resistance in nosocomial pathogens: resistance is a regional problem. *Proc. Natl. Acad. Sci.* **101**, 3709–3714 (2004).
- Perreault, S. K., Binks, B., McManus, D. S. & Topal, J. E. Evaluation of the negative predictive value of methicillin-resistant staphylococcus aureus nasal swab screening in patients with acute myeloid leukemia. *Infect. Control Hosp. Epidemiol.* **42**, 853–856 (2021).
- Luteijn, J., Hubben, G., Pechlivanoglou, P., Bonten, M. & Postma, M. Diagnostic accuracy of culture-based and pcr-based detection tests for methicillin-resistant staphylococcus aureus: a meta-analysis. *Clin. Microbiol. Infect.* **17**, 146–154 (2011).
- Raschpichler, G. et al. Development and external validation of a clinical prediction model for mrsa carriage at hospital admission in southeast lower saxony, germany. *Sci. Rep.* **10**, 17998 (2020).
- Jang, H., Pai, S., Adhikari, B. & Pemmaraju, S. V. Risk-aware temporal cascade reconstruction to detect asymptomatic cases: For the cdc mind healthcare network. In *2021 IEEE International Conference on Data Mining (ICDM)*, 240–249 (IEEE, 2021).
- Pei, S., Liljeros, F. & Shaman, J. Identifying asymptomatic spreaders of antimicrobial-resistant pathogens in hospital settings. *Proc. Natl. Acad. Sci.* **118**, e2111190118 (2021).
- Jang, H. et al. Evaluating architectural changes to alter pathogen dynamics in a dialysis unit. In *2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 961–968 (IEEE, 2019).
- Adhikari, B., Lewis, B., Vullikanti, A., Jimenez, J. M. & Prakash, B. A. Fast and near-optimal monitoring for healthcare acquired infection outbreaks. *PLoS Comput. Biol.* **15**, e1007284 (2019).
- Montella, E. et al. Healthcare associated infections in the neonatal intensive care unit of the -aefederico ii- university hospital: Statistical analysis and study of risk factors. In *2021 International Symposium on Biomedical Engineering and Computational Biology, BECB 2021* <https://doi.org/10.1145/3502060.3503650> (Association for Computing Machinery, New York, NY, USA, 2022).
- Mietchen, M. S., Short, C. T., Samore, M., Lofgren, E. T. & in Healthcare Program (MInD-Healthcare), C. M. I. D. Examining the impact of icu population interaction structure on modeled colonization dynamics of staphylococcus aureus. *PLoS Comput. Biol.* **18**, e1010352 (2022).
- van Kleef, E., Robotham, J. V., Jit, M., Deeny, S. R. & Edmunds, W. J. Modelling the transmission of healthcare associated infections: a systematic review. *BMC Infect. Dis.* **13**, 1–13 (2013).
- Haghpahan, F., Lin, G., Klein, E. & Program, C. M.-H. Deconstructing the effects of stochasticity on transmission of hospital-acquired infections in icus. *R. Soc. Open Sci.* **10**, 230277 (2023).
- Dangerfield, B., Chung, A., Webb, B. & Seville, M. T. Predictive value of methicillin-resistant staphylococcus aureus (mrsa) nasal swab pcr assay for mrsa pneumonia. *Antimicrobial agents Chemother.* **58**, 859–864 (2014).
- Centers for Disease Control and Prevention, MDRO Prevention Strategies, <https://www.cdc.gov/healthcare-associated-infections/php/preventing-mdros/mdro-prevention-strategies.html>.
- Rodríguez, A. et al. Machine learning for data-centric epidemic forecasting. *Nat. Mach. Intell.* **6**, 1122–1131 (2024).
- Savchenko, E. & Bunimovich-Mendrazitsky, S. Investigation toward the economic feasibility of personalized medicine for healthcare service providers: the case of bladder cancer. *Front. Med.* **11**, 1388685 (2024).
- Cui, J. et al. Modeling relaxed policies for discontinuation of methicillin-resistant staphylococcus aureus contact precautions. *Infect. Control Hosp. Epidemiol.* **45**, 833–838 (2024).
- Hunter, E. et al. Assessing the impact of contact tracing with an agent-based model for simulating the spread of covid-19: The irish experience. *Healthc. Analytics* **4**, 100229 (2023).
- Kerr, C. C. et al. Covasim: an agent-based model of covid-19 dynamics and interventions. *PLOS Comput. Biol.* **17**, e1009149 (2021).

36. Lazebnik, T. & Alexi, A. Comparison of pandemic intervention policies in several building types using heterogeneous population model. *Commun. Nonlinear Sci. Numer. Simul.* **107**, 106176 (2022).
37. Lazebnik, T. Data-driven hospitals staff and resources allocation using agent-based simulation and deep reinforcement learning. *Eng. Appl. Artif. Intell.* **126**, 106783 (2023).
38. Chopra, A. et al. Differentiable agent-based epidemiology. In *ICAAMS*, 1848–1857 (2023).
39. Jang, H. et al. Evaluating architectural changes to alter pathogen dynamics in a dialysis unit: For the cdc mind-healthcare group. In *2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, ASONAM '19, 961–968 <https://doi.org/10.1145/3341161.3343515> (Association for Computing Machinery, New York, NY, USA, 2020).
40. Cui, J. et al. Using spectral characterization to identify healthcare-associated infection (hai) patients for clinical contact precaution. *Sci. Rep.* **13**, 16197 (2023).
41. Pedregosa, F. et al. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
42. Hastie, T., Tibshirani, R., Friedman, J. H. & Friedman, J. H. *The elements of statistical learning*, vol. 2 (Springer, 2009).
43. Bishop, C. Pattern recognition and machine learning. *Springer Google Sch.* **2**, 5–43 (2006).
44. Chen, T. & Guestrin, C. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794 (2016).
45. Pulgar, F. J., Charte, F., Rivera, A. J. & Del Jesus, M. J. Aeknn: An autoencoder knn-based classifier with built-in dimensionality reduction. *Int. J. Comput. Intell. Syst.* **12**, 436–452 (2018).
46. Garouani, M., Ahmad, A., Bouneffa, M. & Hamlich, M. Autoencoder-knn meta-model based data characterization approach for an automated selection of ai algorithms. *J. Big Data* **10**, 14 (2023).
47. Martinez, D. A. et al. Early prediction of acute kidney injury in the emergency department with machine-learning methods applied to electronic health record data. *Ann. Emerg. Med.* **76**, 501–514 (2020).
48. Shorr, A. F. et al. A risk score for identifying methicillin-resistant staphylococcus aureus in patients presenting to the hospital with pneumonia. *BMC Infect. Dis.* **13**, 1–7 (2013).
49. STROBE, Checklist of items that should be included in reports of cohort studies, <https://www.strobe-statement.org/download/strobe-checklist-cohort-studies-pdf>.

## Acknowledgements

This work was supported in part by the NSF (Expeditions CCF-1918770, CAREER IIS-2028586, RAPID IIS-2027862, Medium IIS-2403240, Medium IIS-1955883, Medium IIS-2106961, PIPP CCF-2200269), Centers for Disease Control and Prevention Modeling Infectious Diseases In Healthcare

program (5U01CK000589), National Institutes of Health (K23AI163368 to G.R.M.) and National Center for Advancing Translational Science (UL1TR003015, KL2TR003016 to G.R.M.), faculty research award from Facebook, Dolby faculty research award, and funds/computing resources from Georgia Tech.

## Author contributions

J.C., J.H., E.K., G.M., A.V. and B.A.P. conceived the experiments, J.C., and J.H. conducted the experiments, J.C., J.H., E.K., G.M., A.V. and B.A.P. analysed the results. J.C., J.H., E.K., G.M., C.S., A.V. and B.A.P. reviewed the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41746-025-01529-x>.

**Correspondence** and requests for materials should be addressed to B. Aditya Prakash.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025