Proceedings of the 2025 Winter Simulation Conference
E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.

EVALUATING EPIDEMIC SCENARIOS WITH AGENT-BASED SIMULATION: A CASE STUDY FROM UK PUBLIC HEALTH WORKSHOP

Masiar Ghorbani[1], Anastasia Anagnostou[1], Arindam Saha[1], Tasin Islam[1], Nura T. Abubakar[1], Kate Mintram[1], Simon J. E. Taylor[1], and Derek Groen[1]

[1]Dept. of Computer Science, Modelling & Simulation Group, Brunel University of London, UK

ABSTRACT

The growing complexity of public health emergencies requires modeling tools that are both scientifically robust and operationally scalable. As part of the EU Horizon 2020 STAMINA project, we deployed the Flu and Coronavirus Simulator (FACS), a geospatial agent-based model designed to simulate the spread of infectious diseases at local and regional levels. This paper presents a case study from a UK Public Health Workshop, where FACS supported the evaluation of epidemic response scenarios. We describe how FACS integrates demographic, spatial, and epidemiological data, and outline key enhancements, such as location-based parallelization and FabSim3-enabled automation, which enable large-scale simulation. We detail the scenario designs and outcomes, highlighting the intersection of simulation projections and intervention planning. Finally, we reflect on communicating results to stakeholders and bridging the gap between modeling and policy. This work demonstrates how geospatially grounded, scalable agent-based simulations can provide meaningful insights into regional intervention planning within operational timeframes.

1 FACS: THE FLU AND CORONAVIRUS SIMULATOR

The COVID-19 pandemic, triggered by the novel Coronavirus, posed unprecedented challenges globally, requiring timely and informed decisions across multiple levels of governance (Rockett et al. 2020). Understanding the dynamics of infectious disease transmission, particularly at sub-national scales, proved critical for effectively managing such crises. Traditional models, such as the Susceptible, Infected, Recovered (SIR) framework, have primarily focused on national-scale analyses, often overlooking regional heterogeneity and the impact of localised intervention measures (Thomas et al. 2020).

Several agent-based and compartmental models were developed during the COVID-19 pandemic to address these limitations and support public health decision-making at national and local levels. Notable examples include CovidSim (Ferguson et al. 2020), an early agent-based model developed at Imperial College London; Covasim (Kerr et al. 2021), an open-source agent-based model focused on intervention analysis; models from the JUNIPER Consortium, including OpenABM-Covid19; CityCOVID (Liu et al. 2021), a granular urban-scale agent-based simulator, CitySEIRCast (Bilal et al. 2023), which combined SIR modeling with city-specific forecasts, and CALMS (Mintram et al. 2022), to predict the lifelong impacts of COVID-19 on the health and economy. Compared to these models, FACS emphasizes high-resolution geospatial integration, operational scalability through parallelization, and scenario customization explicitly tailored for regional and local epidemic response.

FACS builds on the Simulation Development Approach (SDA), a structured framework that integrates geospatial layouts, demographic profiles, and epidemiological parameters to support detailed, region-specific simulations and forecasting (Imran et al. 2022). SDA outlines the steps needed to build simulation models by constructing realistic environments that incorporate data on residential distributions, building types, mobility patterns, and local healthcare facilities, while also finding potential bottlenecks in the development process. This structured methodology ensures that the resulting models are transparent, reproducible, and adaptable to new regions or evolving public health challenges.

As FACS has evolved, the SDA has remained central to its architecture, supporting adaptability to larger-scale and more complex scenarios while preserving its core strength: providing insights into the localised impacts of disease spread and public health interventions. Designed to reflect real-world complexity, FACS integrates demographic, epidemiological, and spatial data to simulate how agents interact, move, and respond to evolving health policies.

Figure 1 presents the overall FACS modeling workflow, from scenario selection and data ingestion to model construction, refinement, and execution. The diagram also highlights the integration of public health measures (e.g., travel restrictions, facility closures) and the deployment of FabSim3 (Groen et al. 2016; Groen et al. 2023) for high-performance computing resources to support large-scale simulations. FabSim3 supports a plugin-based architecture tailored to domain-specific applications, including pandemic modeling through the FabCovid19 plugin. This plugin streamlines the process of running FACS at scale by packaging simulation inputs into region-specific formats, enabling users to efficiently customise and launch many large-scale epidemiological simulations across diverse HPC environments.

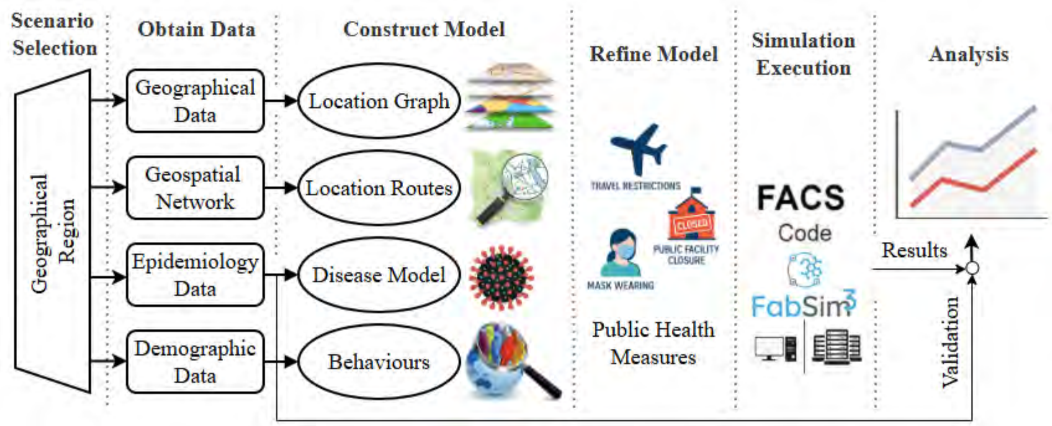

Figure 1: This diagram shows the complete FACS pipeline from scenario selection to analysis. Data from geographical, epidemiological, and demographic sources is used to construct the simulation model, define agent behaviours, and embed public health interventions. Simulations are executed using the FACS engine and FabSim3, with results validated and analysed to support policy decision-making.

FACS has been previously validated and applied in multiple contexts, including national and international collaborations, including STAMINA H2020 (Bakalos et al. 2022), as a flexible and scalable simulation tool. While earlier publications have focused on FACS's technical design and initial applications, this paper highlights its operational deployment in a public health workshop setting, emphasising its role in real-time scenario analysis and policy evaluation.

2 FACS: A TOOL FOR LOCAL EPIDEMIC DECISION-MAKING SUPPORT

Agents in FACS represent individuals within the population, each assigned attributes such as age, health status, and mobility patterns. These agents interact within an Ecosystem comprising different location types, including homes, schools, hospitals, and public spaces, each defined by parameters such as size, density, and function. Agent interactions and transitions between health states (e.g., susceptible, exposed, infectious, recovered) are governed by agents' behavioural rules, their spatial proximity, and probabilistic transmission dynamics, allowing interventions such as vaccination campaigns for virus mutations and lockdown measures.

As part of the EU Horizon 2020 STAMINA project, FACS facilitated communication between modellers, policymakers, and public health stakeholders by enabling scenario-driven exploration of intervention strategies. The simulator was evaluated and deployed in trials across a wide range of partner countries, including the UK, Turkey, Spain, Slovenia, Romania, Lithuania, Greece, the Netherlands, the Czech Republic, and Austria. These trials varied in scope, ranging from full-scale regional modeling to focused intervention scenario testing, demonstrating FACS's adaptability to diverse public health contexts. This was further demonstrated through a case study with the London North West University Healthcare NHS Trust, where FACS was used to forecast hospital and ICU admissions across the boroughs of Brent, Ealing, and Harrow. Model outputs were validated against anonymised NHS data and refined through expert consultation, highlighting FACS's capacity to support rapid, data-driven decision-making under real-world constraints (Imran et al. 2022).

In the Workshop Dry-run, we simulated eighteen regions, including Cumbria, Lancashire, Cheshire, Sussex, Surrey, and Berkshire, organised into two distinct geographical clusters: the North West and the South East, as shown in Figure 2. This subdivision was designed to provide us with greater control over computational parallelization, while enabling a comparative analysis between two demographically and geographically distinct areas.

Figure 2: This map highlights the North West and South East regions of England, where FACS was deployed to simulate COVID-19 transmission dynamics as part of the UK Public Health Agency Workshop activities. The subdivision into two distinct areas enabled the controlled testing of parallelization strategies and supported a comparative analysis of intervention effectiveness across regions with different demographic and mobility characteristics.

Combining fine-grained agent-based simulation with real-world geospatial and demographic data enabled stakeholders to explore targeted intervention strategies and anticipate region-specific outcomes. In particular, the subdivision into North West and South East regions provided a valuable testbed for assessing how computational scalability could be achieved while preserving epidemiological relevance. Building on these foundations, the next stage of development focused on enhancing FACS's ability to operate at scale, automate workflows, and produce insights across large and complex health system scenarios.

3 MAKING FACS SCALABLE AND CAPABLE OF SUPPORTING ACTIONABLE RESULTS

Developing an accurate and flexible simulator is an essential first step, but it is not, on its own, sufficient for real-world application. Throughout this work, we use "actionable" to refer to the simulation capacity to generate robust, reproducible results that can meaningfully support decision-making processes in public health planning. To be operationally sound, a simulator must be capable of representing complex regional dynamics, scaling to large populations, and accommodating modifications in response to evolving public health requirements. Without these capabilities, even the most sophisticated models risk remaining academic exercises, disconnected from the practical demands of epidemic response.

A key architectural decision to make FACS actionable involved distributing the workload across more computational resources. To achieve this, we adopted a location-based parallelization strategy, similar to the approach used in the FLEE model (Ghorbani et al. 2024; Groen et al. 2024), in which physical environments, such as schools, offices, and hospitals, are assigned to separate processes. This approach enables parallelization and concurrent computation across the simulation space while minimizing communication overhead between agents.

Figure 3 presents a high-level view of the FACS ecosystem and its parallelization strategy. Agents are assigned demographic attributes and initialised with needs that drive them to visit different location types, such as shopping centres or markets, where exposure risk depends on proximity-based interactions. The simulation distributes location types across multiple processes, with each process managing a group of locations and executing agent interactions independently. At the end of each timestep, processes synchronise and update the global state before advancing the simulation clock. This loop continues until completion, after which results are aggregated and exported for analysis.

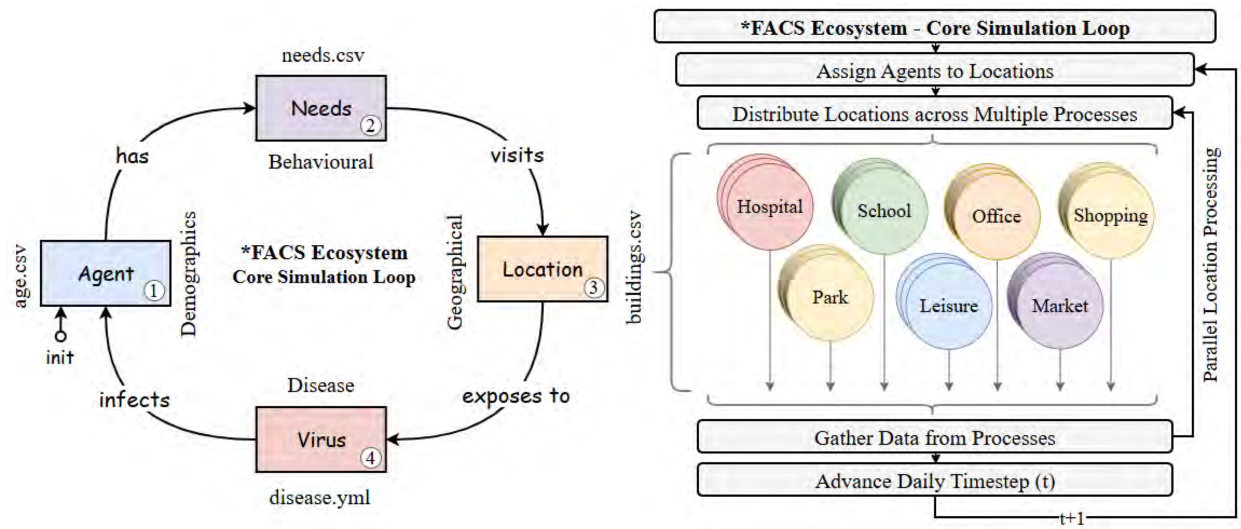

Figure 3: Overview of the FACS simulation architecture. (Left) The core simulation loop, where an agent has needs that prompt them to visit a location. By visiting locations, the agent may become exposed to or spread infections. (Right) The parallel processing strategy, in which location types are distributed across multiple processes and run concurrently to advance the simulation's timestep.

Additionally, we integrated FACS with FabSim3, a job management toolkit developed by the Modelling & Simulation Group at Brunel University of London. FabSim3 automates the pre-processing, execution, and post-processing of simulation tasks across high-performance computing environments such as the ARCHER2 Supercomputer (Beckett et al. 2024) through the FabCovid19 plugin. Additional supercomputing facilities can be integrated into this plugin with minimal adaptation.

While FabSim3 enabled efficient deployment on supercomputing platforms, the UK Public Health Workshop simulations demanded more computational logistics. Generating results required exploring numerous scenarios, each replicated hundreds of times to ensure statistical robustness. Due to the high cost and restricted allocation windows on supercomputers, we sought more scalable and cost-efficient high-performance computing resources to support high-throughput large-scale simulations.

Drawing on prior experience with Grid Computing from the European Organisation for Nuclear Research, and its lxplus computing systems, as well as the UK Tier-2 GridPP network, we introduced additional resources available at Brunel University of London to support FACS large-scale simulations (Britton et al. 2009). Although full FabSim3 integration was not yet possible due to differences in authentication workflows, such as certificate authority management and proxy generation, we successfully used Grid computing manually to run thousands of simulation instances in parallel, substantially accelerating scenario evaluation for the workshop.

In the next section, we explain how FACS has been applied to modeling healthcare operations and epidemiological forecasting. We also discuss how it has been integrated with discrete-event simulation frameworks, such as CHARM, to help decision-making in complex healthcare environments.

4 SIMULATION AS A DECISION SUPPORT TOOL IN PUBLIC HEALTH

In the face of uncertainties surrounding disease transmissibility, public compliance, healthcare capacity, and emerging variants, simulation tools provide a controlled environment for testing assumptions, anticipating outcomes, and informing evidence-based decisions (Howerton et al. 2023).

Agent-based simulations (Taylor 2014), in particular, enable decision-makers to explore "what-if" scenarios that reflect real-world heterogeneity, such as how disease dynamics differ between communities, how interventions such as school closures or travel restrictions affect transmission, or how hospital systems might cope with surge conditions (Imran et al. 2022). These insights are crucial not only for national-level strategy but also for regionally tailored responses, where mobility patterns, demographic structures, and resource availability vary widely (Saha et al. 2023).

FACS is designed with interoperability in mind, enabling integration within broader simulation toolchains to address complex, multi-layered challenges in public health planning. One such integration is with the dynamiC Hospital wARd Management (CHARM) model (Anagnostou et al. 2022), a discrete-event simulation tool focused on hospital operations planning. CHARM simulates dynamic patient flows across Intensive Care Units (ICU), emergency departments, and elective care units, accounting for fluctuating bed availability and ward reconfiguration strategies. When linked with FACS in a sequential hybrid architecture, CHARM receives localised hospitalisation forecasts generated from FACS epidemiological outputs, using them as input triggers for dynamic resource allocation simulations (Anagnostou et al. 2013; Anagnostou and Taylor 2024).

As illustrated in Figure 4, multiple CHARM instances can be initialised concurrently, each corresponding to a different healthcare facility, using parallelized FACS outputs. This design allows for scalable analysis of healthcare system pressures under various outbreak conditions while maintaining regional granularity. The FACS–CHARM integration thus enables decision-makers to not only anticipate disease burden at the population level but also to simulate operational responses at the hospital level, supporting more precise, system-aware intervention planning across various parts of the healthcare network.

In this context, simulation tools such as FACS are not merely an academic exercise but a practical decision-support tool, especially when embedded into public health workflows and aligned with stakeholder priorities. Its value lies in its ability to uncover dynamics that might otherwise go unnoticed and support informed, timely, and locally relevant public health action.

5 SCENARIO DESIGN AND SIMULATION RESULTS

Designing meaningful simulation scenarios requires close collaboration between developers and stakeholders to ensure that both epidemiological contexts and public health interventions are accurately represented. In our case studies, scenario selection was guided by verifying the completeness and relevance of input data, including high-resolution regional maps, demographic profiles, and geolocated infrastructure such as residential buildings, offices, schools, and public facilities. This systematic process also emphasized the importance of transparent documentation to enable future validation, comparison, and reuse of scenarios.

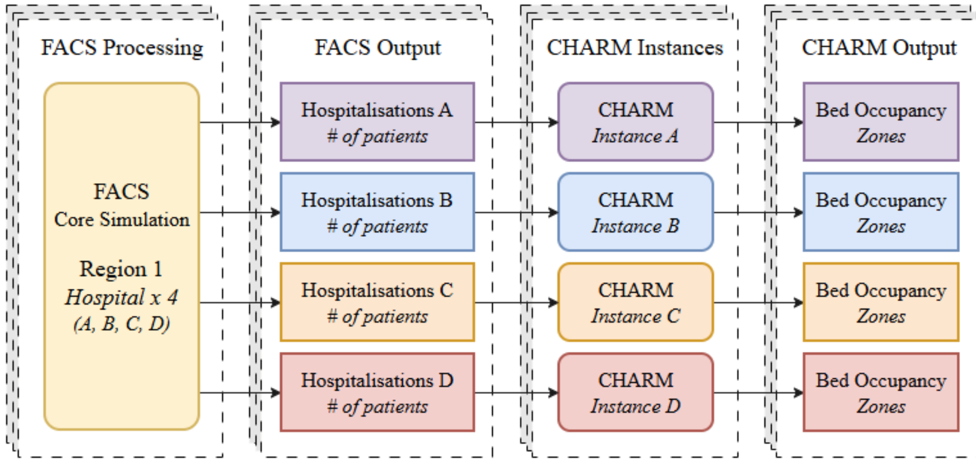

Figure 4: FACS–CHARM integration for parallelized healthcare simulation. This architecture supports scalable analysis of healthcare system pressures based on epidemic forecasts. Each layer presents a FACS simulation for one geographical region (e.g., Cumbria, Sussex, Oxfordshire, etc.). The hospitalisation data from each facility is used to initialise an instance of the CHARM simulation, which models bed occupancy and patient flow. These instances of CHARM can be executed independently and concurrently.

FACS supports the simulation of complex disease dynamics through flexible input configurations. The system is capable of modeling diseases with different transmission rates and intervention responses, including mutation handling and vaccination rollouts. Once the model is configured, we run extensive simulations comprising hundreds of replications across various parameter variations. This approach ensures statistical robustness and increases confidence in the results by accounting for stochastic variability in agent behaviours and interactions.

The policy requirements and population behaviours were carefully encoded into the simulation setup, ensuring alignment with real-world decision-making contexts. Epidemiological parameters, including infection rates, incubation periods, vaccination effects, and variant profiles, were validated through expert consultation as necessary. Building on this foundation, we designed and executed a series of FACS simulation scenarios tailored to the Public Health Workshop, focusing on evaluating intervention strategies across distinct regional clusters.

Extensive pilot simulations, involving hundreds of replications across various parameter variations, were conducted to capture stochastic variability in agent behaviours within each simulation scenario. In this context, "pilot" refers to preliminary simulation runs designed to verify model setup, parameter sensitivity, and computational scalability before full-scale execution. Each region, such as Cumbria, Sussex, or Berkshire, was modelled as an independent simulation, minimizing data dependencies and enabling parallel execution across high-performance computing resources.

Simulation outputs provided temporal and spatial trends, including hospital admissions and ICU demand, which were validated against available empirical data. Results were presented through intuitive visualisations and scenario comparisons to support transparent, evidence-based discussions.

To explore the impact of different school reopening strategies on COVID-19 transmission dynamics, eight scenarios were developed in collaboration with the UK Public Health Agency, as summarized in Table 1. These scenarios varied the timing, extent, and structure of school resumptions across the simulated regions, and were defined as follows:

A key metric for monitoring epidemic trends is the time-varying reproduction number, R(t), defined as the average number of secondary infections generated by a single infected individual (Pouw et al. 2021). While extracting R(t) from agent-based models such as FACS is non-trivial due to population heterogeneity,

Table 1: Scenario timeline and school reopening strategies evaluated during the UK Public Health Workshop.

<table>
<tr><th>Date</th><th colspan="8">Strategy</th></tr>
<tr><td>01/03/2020 – 15/06/2020</td><td colspan="8">Restrictions and measures following the actual timeline</td></tr>
<tr><td rowspan="2">15/06/2020 – 24/07/2020</td><th>Scenario 1</th><th>Scenario 2</th><th>Scenario 3</th><th>Scenario 4</th><th>Scenario 5</th><th>Scenario 6</th><th>Scenario 7</th><th>Scenario 8</th></tr>
<tr><td>All schools stay shut</td><td>All schools resume</td><td>Early years resume</td><td>Primary schools resume</td><td>Secondary schools resume</td><td>Transition years resume</td><td>Two weeks ON/OFF</td><td>Classes split AM/PM</td></tr>
<tr><td>24/07/2020 – 15/08/2020</td><td colspan="8">Restrictions and measures following the actual timeline</td></tr>
</table>

we approximate it using infection counts with a 10-day lag:

$$R(t) = \frac{I(t)}{I(t-10)} \tag{1}$$

Where I(t) represents new infections on day t, and the denominator approximates the infectious period. This method provides a practical estimate of R, although it relies on population-level averaging and does not account for transmission paths. More granular alternatives, such as transmission trees, can offer insights into which individual transmitted the illness to others, but are less practical for scenario-level comparisons. In Table 2, we calculated an average peak R of 2.15 in the North West and 2.03 in the South East across all eight scenarios.

Table 2: R-values for North West (NW) and South East (SE). For full R-values, see Figure 7.

Date	R-value NW	R-value SE
15/06/2020	0.93	0.92
04/07/2020	1.14	1.11
15/07/2020	2.15	2.03
24/07/2020	1.66	1.61

In addition to estimating the reproduction number, we approximated hospitalisation trends based on ICU admissions. On average, patients admitted to the ICU spent approximately two days in the hospital beforehand, and roughly one in ten hospitalised patients required ICU care. Using this relationship, we estimated the number of hospitalisations H(t), from ICU admissions C(t), as follows:

$$H(t) = 10 \times C(t+2) \tag{2}$$

Here, H(t) estimates the number of hospitalizations on day t, and C(t) represents the number of ICU admissions, assuming 10% of hospitalized patients require ICU care approximately two days after admission. This approximation helps align simulation outputs with historical data for Validation.

To evaluate the impact of different school reopening strategies on epidemic trajectories, we conducted simulation experiments across all eight scenarios outlined previously. Figure 5 presents the predicted hospitalisation trends per 100,000 population for both the North West and South East regions. Model predictions represent each scenario in comparison to available historical data. Differences in peak magnitude, timing, and resource demand across scenarios provide insights into the relative epidemiological and healthcare impacts of various reopening strategies.

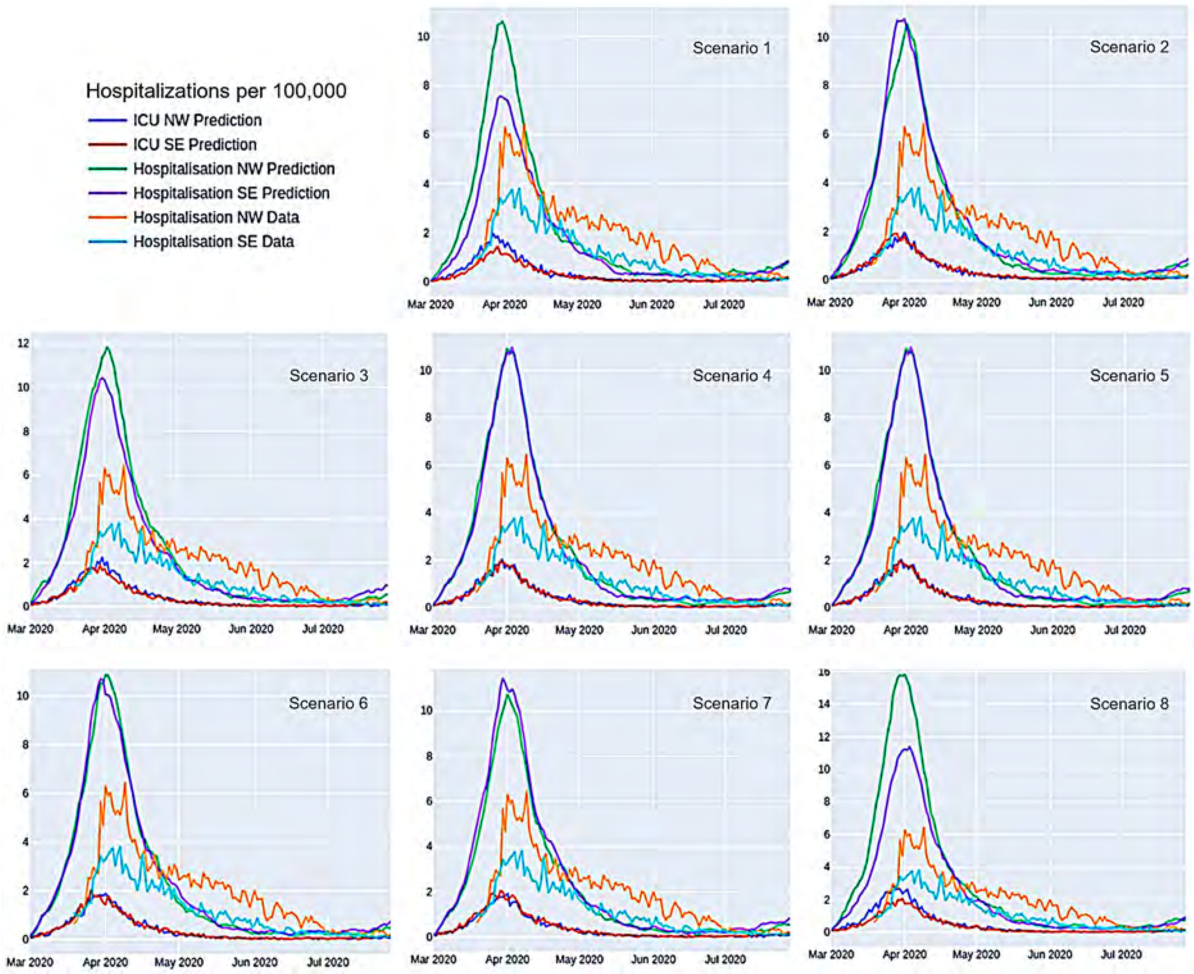

Figure 5: Predicted hospitalisations per 100,000 population across Scenarios 1–8. Each panel displays model projections for the North West (NW) and South East (SE) regions, alongside historical data. Scenarios vary by school reopening strategy, with impacts reflected in the magnitude and timing of epidemic peaks and hospital resource demand.

Simulations were first conducted independently for each of the 18 regions, covering eight distinct school reopening scenarios per region. Results were then aggregated into North West and South East groupings for comparative analysis. Figure 6 shows validation results, with simulated hospitalisation trends plotted against observed data and shaded areas indicating 95% confidence intervals. Although both infection and hospitalisation trends were compared to historical records, more weight was given to hospitalisation data due to its greater reliability and accuracy. The simulations accurately captured the timing and magnitude of the epidemic peak, supporting their use in scenario evaluation.

To further evaluate the impact of school reopening strategies, we analysed the progression of the reproduction number R(t) across the North West and South East regions. Figure 7 presents the simulated R-value trajectories under all eight scenarios, with solid coloured lines representing different intervention strategies and vertical dashed lines indicating key policy dates. Variations in the magnitude and timing of R-value increases highlight how different reopening approaches influenced transmission dynamics regionally. Notably, scenarios involving full school reopening exhibited sharper rises in R-values compared to staggered or partial reopening strategies.

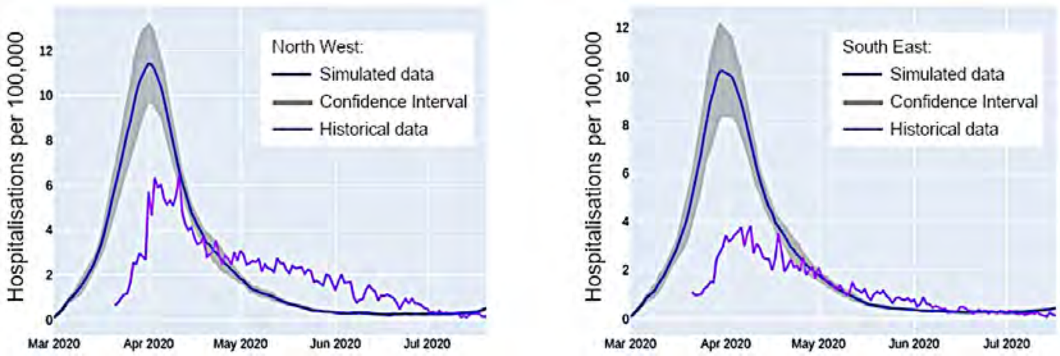

Figure 6: Validation of simulated hospitalisation trends against historical data for the North West (left) and South East (right) regions. Solid lines represent the FACS simulation mean; shaded areas denote 95% confidence intervals across replication, with the shaded areas containing approximately 95% of the simulation ensemble results, indicating aleatoric uncertainty in the code.

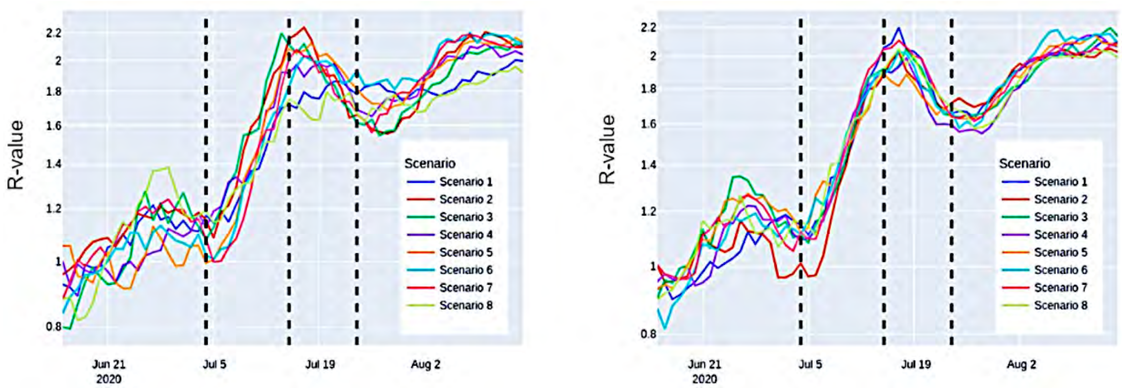

Figure 7: Simulated reproduction number (*R*-value) trajectories for the North West (left) and South East (right) regions across all eight school reopening scenarios. Solid coloured lines represent different scenarios, and vertical dashed lines indicate key policy dates. These results show the dynamic response of *R*-values to different intervention strategies and highlight regional differences in epidemic resurgence risk. The R-value curves begin slightly later than the scenario timelines, as their computation relies on preceding infection data and stabilises only once sufficient case history has been generated within the simulation.

R-values below 1.0 indicate that the epidemic is shrinking, while values above 1.0 reflect continued transmission and potential growth of the outbreak. R-values greater than 2.0 are expected during early phases of an epidemic, or in scenarios where interventions are limited or delayed.

To support the robustness of our results, we refer to prior work where we conducted a comprehensive sensitivity analysis of the FACS model using sensitivity analysis and Sobol indices (Saha, Ghorbani, Suleimenova, Anagnostou, and Groen 2023). This analysis examined how key epidemiological parameters, including infection rates, incubation period, and mild recovery period, influenced health outcomes across different geographic configurations. The findings indicated that the sensitivity of model outputs is context-dependent. For instance, the infection rate was more dominant in spatially segregated regions, while recovery dynamics played a larger role in mixed-population settings. These insights validated the model's responsiveness to key assumptions and provided a broader context for interpreting variability in the scenario outcomes presented here.

6 CONNECTING DATA TO POLICY: IMPLICATIONS AND INSIGHTS

Turning complex simulations into policy-relevant insights requires more than technical accuracy; it demands clarity, contextualisation, and trust. FACS supports this process through visual outputs designed to convey both the direction and confidence of predicted trends. These include spatial heatmaps of infection clusters, scatter plots for key sites (e.g., schools, supermarkets), and time-series overlays comparing simulated and historical hospitalisations. Together, these outputs help identify local hotspots, assess intervention effects, and validate patterns against observed outcomes.

To support time-sensitive decision-making, results are summarised in structured reports, slide packs, and dashboards. These include scenario assumptions, ICU forecasts, and cumulative infections, with visuals prioritised for interpretability. Outputs are also shared with analysts and domain experts to support model refinement and internal alignment.

Direct engagement with stakeholders and policymakers, via workshops, dry runs, and briefings, was central to aligning simulations with decision-making contexts. We worked closely with the UK Health Security Agency and other partners to ensure outputs were correctly interpreted. Special attention was given to results that challenged expectations, such as delayed peaks or limited effects from anticipated interventions. In these cases, open discussion of model assumptions and uncertainties helped foster trust and turn scepticism into productive dialogue.

Rather than offering definitive forecasts, decision-support tools such as FACS act as boundary objects between technical and policy communities. By revealing scenario sensitivities and trade-offs, they help structure conversations around risk, capacity, and timing. Ultimately, this approach strengthens the connection between simulation data and real-world public health planning, ensuring that insights are not just generated but also actionable.

7 CONCLUSION AND FUTURE WORK

The COVID-19 pandemic led to the development of many modeling frameworks, each offering distinct strengths and perspectives. FACS represents one such approach, developed with a particular emphasis on localised decision support, high-performance scalability, and flexible data integration. It incorporates modeling, simulation, and decision support into a modular framework. Its agent-based architecture captures behavioural, spatial, and temporal complexity while remaining adaptable across scales. The simulation engine enables high-resolution scenario analysis, with outputs validated against empirical data and translated into accessible visualisations that support policy engagement.

Deploying FACS in a live public health workshop involved technical and collaborative challenges. A core team gathered and cleaned spatial, demographic, and epidemiological data, ran thousands of simulations, and worked closely with stakeholders to align outputs with expectations. Tailored visualisations supported interpretation and trust, enabling the timely delivery of validated insights.

As part of this effort, we simulated the spread of COVID-19 in eighteen regions across North West and South East of England, and hospital demand in the boroughs of Brent, Ealing, and Harrow. The workshop brought together NHS analysts, public health officials, modellers, and developers to compare tools, refine outputs, and improve interpretability. A roundtable addressed simplification for non-technical users and clarification of behavioural assumptions. Forecasts were validated against anonymised NHS data and refined with clinical input. Stakeholders used FACS to explore reopening scenarios and capacity thresholds, with feedback emphasising interpretability and alignment with internal analytics.

Outputs, including time-series plots, heatmaps, and ICU projections, were presented to support structured discussions around timing, demand, and risk thresholds. Overlaying simulations with historical data improved clarity, especially for non-modeling participants. Stakeholders valued collaborative interpretation, transparent communication, GDPR compliance, real-time visualisation, and access to manuals and training. In response, we enhanced stakeholder-facing tools, expanded disease coverage to measles, and supported broader deployment alongside other modeling efforts.

While we cannot verify a direct influence on policy decisions, the workshop provided a venue for introducing simulation tools to public health stakeholders. As computational modeling becomes increasingly embedded in policy, such engagements help clarify assumptions, foster trust, and build familiarity for future integration into decision-making processes.

Looking ahead, we are exploring machine learning for post-simulation analysis, with initial steps led by students. We also plan to refactor and document the codebase to support long-term maintainability and facilitate adoption by external users. The FACS framework remains open-source and available at https://github.com/djgroen/FACS, including the FabCovid19 plugin: https://github.com/djgroen/FabCovid19.

ACKNOWLEDGMENTS

This work has been supported by the SEAVEA ExCALIBUR project, which has received funding from EPSRC under grant agreement EP/W007711/1, as well as by the STAMINA project, which has received funding from the European Union Horizon 2020 research and innovation programme under grant agreement no 883441.

REFERENCES

Anagnostou, A., D. Groen, S. J. E. Taylor, D. Suleimenova, N. Abubakar, A. Saha, *et al*. 2022. "FACS-CHARM: A Hybrid Agent-Based and Discrete-Event Simulation Approach for COVID-19 Management at Regional Level". In *2022 Winter Simulation Conference (WSC)*, 1223–1234 https://doi.org/10.1109/WSC57314.2022.10015462.

Anagnostou, A., A. Nouman, and S. J. E. Taylor. 2013. "Distributed Hybrid Agent-Based Discrete Event Emergency Medical Services Simulation". In *2013 Winter Simulation Conference (WSC)*, 1625–1636 https://doi.org/10.1109/WSC.2013.6721545.

Anagnostou, A., and S. J. E. Taylor. 2024. "Hybrid Simulation in Healthcare Applications". In *Hybrid Modeling and Simulation: Conceptualizations, Methods and Applications*, edited by M. Fakhimi and N. Mustafee, 271–294. Cham: Springer Nature Switzerland https://doi.org/10.1007/978-3-031-59999-6_11.

Bakalos, N., M. Kaselimi, N. Doulamis, A. Doulamis, D. Kalogeras, M. Bimpas, *et al*. 2022. "STAMINA: Bioinformatics Platform for Monitoring and Mitigating Pandemic Outbreaks". *Technologies* 10(3):63 https://doi.org/10.3390/technologies10030063.

Beckett, G., J. Beech-Brandt, K. Leach, Z. Payne, A. Simpson, L. Smith *et al*. 2024. "ARCHER2 Service Description". Technical Report, EPCC, University of Edinburgh https://doi.org/10.5281/zenodo.14507040.

Bilal, S., W. Zaatour, Y. A. Otano, A. Saha, K. Newcomb, K. Soo, *et al*. 2023. "CitySEIRCast: An Agent-Based City Digital Twin for Pandemic Analysis and Simulation". *medRxiv* https://doi.org/10.1101/2023.12.22.23300481.

Britton, D., A. J. Cass, P. E. L. Clarke, J. Coles, D. J. Colling, A. T. Doyle *et al*. 2009. "GridPP: The UK Grid for Particle Physics". *Philosophical Transactions of the Royal Society A* 367:2–18 https://doi.org/10.1098/rsta.2009.0036.

Ferguson, N. M., D. Laydon, G. Nedjati-Gilani, N. Imai, K. Ainslie, and M. Baguelin. 2020. "Report 9: Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand". Technical Report, Imperial College London https://doi.org/10.25561/77482.

Ghorbani, M., D. Suleimenova, A. Jahani, A. Saha, Y. Xue, K. Mintram, *et al*. 2024. "Flee 3: Flexible Agent-Based Simulation for Forced Migration". *Journal of Computational Science* 81:102371 https://doi.org/10.1016/j.jocs.2024.102371.

Groen, D., H. Arabnejad, D. Suleimenova, W. Edeling, E. Raffin, Y. Xue, *et al*. 2023. "FabSim3: An Automation Toolkit for Verified Simulations Using High Performance Computing". *Computer Physics Communications* 283:108596 https://doi.org/10.1016/j.cpc.2022.108596.

Groen, D., A. P. Bhati, J. Suter, J. Hetherington, S. J. Zasada, and P. V. Coveney. 2016. "FabSim: Facilitating Computational Research Through Automation on Large-Scale and Distributed E-Infrastructures". *Computer Physics Communications* 207:375–385 https://doi.org/10.1016/j.cpc.2016.05.020.

Groen, D., N. Papadopoulou, P. Anastasiadis, M. Lawenda, L. Szustak, S. Gogolenko, *et al*. 2024. "Large-Scale Parallelization of Human Migration Simulation". *IEEE Transactions on Computational Social Systems* 11(2):2135–2146 https://doi.org/10.1109/TCSS.2023.3292932.

Howerton, E., P. Simen, and A. Bleakley. 2023. "Simulation Models to Inform Pandemic Response: Lessons from COVID-19". *American Journal of Public Health* 113(9):974–980 https://doi.org/10.2105/AJPH.2023.307374.

Imran, M., H. Arabnejad, D. Suleimenova, I. Sassoon, A. Marshan, A. Serrano-Rico, *et al*. 2022. "FACS: A Geospatial Agent-Based Simulator for Analysing COVID-19 Spread and Public Health Measures on Local Regions". *Journal of Simulation* 16(4):355–373 https://doi.org/10.1080/17477778.2020.1800422.

Kerr, C. C., R. M. Stuart, D. Mistry, R. G. Abeysuriya, K. Rosenfeld, and G. R. Hart. 2021. "Covasim: An Agent-Based Model of COVID-19 Dynamics and Interventions". *Nature Communications* 12(1):1181 https://doi.org/10.1038/s41467-021-21465-9.

Liu, Y., S. Ying, M. Modjeski, L. Ang, A. Ghaffarizadeh, and P. Macklin. 2021. "CityCOVID: A High-Fidelity Agent-

Based Model of COVID-19 Transmission to Support Decision-Making". *ACM Transactions on Modeling and Computer Simulation* 31(4):1–27 https://doi.org/10.1145/3468496.

Mintram, K., A. Anagnostou, N. Anokye, E. Okine, D. Groen, A. Saha, *et al*. 2022. "CALMS: Modelling the Long-Term Health and Economic Impact of COVID-19 Using Agent-Based Simulation". *PLOS ONE* 17(8):1–19 https://doi.org/10.1371/journal.pone.0272664.

Pouw, N., J. van de Maat, K. Veerman, J. ten Oever, N. Janssen, E. Abbink, *et al*. 2021. "Clinical Characteristics and Outcomes of 952 Hospitalized COVID-19 Patients in the Netherlands: A Retrospective Cohort Study". *PLOS ONE* 16(3):1–15 https://doi.org/10.1371/journal.pone.0248713.

Rockett, R. J., A. Arnott, C. Lam, R. Sadsad, V. Timms, K. Gray, *et al*. 2020. "Revealing COVID-19 Transmission in Australia by SARS-CoV-2 Genome Sequencing and Agent-Based Modeling". *Nature Medicine* 26(9):1–10 https://doi.org/10.1038/s41591-020-1000-7.

Saha, A., M. Ghorbani, D. Suleimenova, A. Anagnostou, and D. Groen. 2023. "Impact of Geography on the Importance of Parameters in Infectious Disease Models". *arXiv:2310.02449*.

Taylor, S. J. E. 2014. *Agent-Based Modeling and Simulation*. OR Essentials. London: Palgrave Macmillan UK https://doi.org/10.1057/9781137453648.

Thomas, L. J., P. Huang, F. Yin, X. I. Luo, Z. W. Almquist, J. R. Hipp *et al*. 2020. "Spatial Heterogeneity Can Lead to Substantial Local Variations in COVID-19 Timing and Severity". *Proceedings of the National Academy of Sciences* 117(39):24180–24187 https://doi.org/10.1073/pnas.2011656117.

AUTHOR BIOGRAPHIES

MAZIAR GHORBANI is an Associate Lecturer in Computer Science at Brunel University of London and a member of its Modelling & Simulation Group. He co-developed the FACS and contributed to the H2020 STAMINA project. His email address is maziar.ghorbani@brunel.ac.uk.

ANASTASIA ANAGNOSTOU is a Senior Lecturer and Co-Lead of the Modelling & Simulation Group at Brunel University of London. Anastasia led contributions to the STAMINA project, and she is the Program Chair for WSC 2025. Her email address is anastasia.anagnostou@brunel.ac.uk.

ARINDAM SAHA is a Senior Research Data Scientist at UCL and an associate of Brunel University of London. His work focuses on agent-based modeling and emulation. His email address is arindam.saha@ucl.ac.uk.

TASIN ISLAM is a Research Fellow in Computer Science at Brunel University of London. He contributed to the STAMINA project and works on agent-based modeling and Artificial Intelligence. His email address is tasin.islam@brunel.ac.uk.

NURA TAJJANI ABUBAKAR is an Associate Lecturer at Brunel University of London with research interests in distributed simulations. He contributed to the STAMINA project and other EU projects. His email address is nura.abubakar2@brunel.ac.uk.

KATE MINTRAM is an Associate Lecturer at Brunel University of London. She works on agent-based simulation in healthcare. She contributed to the STAMINA project. Her email address is kate.mintram@brunel.ac.uk.

SIMON J. E. TAYLOR is a Professor of Computer Science at Brunel University of London, specialising in modeling & simulation, and digital infrastructures. He co-founded the Journal of Simulation and the UK Operational Research Society Simulation Workshop Series. His email address is simon.taylor@brunel.ac.uk.

DEREK GROEN is a Reader in Computer Science at Brunel University of London and developer of the FACS and FLEE models. He leads technical work on several simulation projects, including STAMINA and SEAVEA. His email address is derek.groen@brunel.ac.uk.