A STANDARDIZED FRAMEWORK FOR MODELING NON-PHARMACEUTICAL INTERVENTIONS IN INDIVIDUAL-BASED INFECTIOUS DISEASE SIMULATIONS

Johannes Ponge¹, Janik Suer², Bernd Hellingrath¹, and André Karch²

¹Dept. of Information Systems, University of Münster, Münster, NRW, GERMANY ²Inst. of Epidemiology & Social Medicine, University of Münster, Münster, NRW, GERMANY

ABSTRACT

Individual-based infectious disease simulations play a fundamental role in the evaluation of intervention strategies before implementing them as public policy during emerging epidemics. While public health offices provide a vast array of potential measures in their preparedness plans, their mode of operation in practice is conditioned by a plethora of regional legal, economic, and demographic factors. This work introduces the Trigger-Strategy-Measure (TriSM) formalization as our main contribution and its implementation in the German Epidemic Microsimulation System (GEMS). TriSM is a standardized formalization for modeling complex interventions in individual-based infectious disease simulations, focusing on granularity, extensibility, expressiveness, and usability. We demonstrate TriSM's capabilities in six simulation case studies where we apply nuanced intervention strategies to a COVID-19-like outbreak scenario and evaluate their effectiveness. Our work contributes to the ongoing efforts of increasing fidelity and expatiating implicit assumptions in individual-based infectious disease models through a standardized formalization.

1 INTRODUCTION & RELATED WORK

An infectious disease outbreak can pose a major threat to the well-being of entire populations. In the absence of vaccination programs, the immediate public health response is generally the introduction of nonpharmaceutical interventions (NPIs). Individual-based infectious disease simulations are widely used to evaluate the anticipated impact of intervention strategies before their enactment (Matrajt and Leung 2020; Prem et al. 2020). A large effort to build such a modeling framework is carried out in the OptimAgent consortium, where a team of 14 German and international research institutions collaborate to build the German Epidemic Microsimulation System (GEMS). GEMS is a flexible individual-based modeling framework written in the Julia programming language designed to work on a German population model (~84 million individuals), providing the basis for evaluating intervention strategies for contactcommunicable diseases. The range of potential intervention strategies that shall be modeled can be obtained from public health institutions such as the WHO, CDC and RKI, Germany's federal disease control and prevention agency. They provide guidelines on a wide spectrum of interventions in their pandemic response plans (Centers for Disease Control and Prevention 2022; Robert Koch-Institut 2017; World Health Organization 2023), including quarantining, testing, mask-wearing, school-, border- and non-essential business closures, and even widespread lockdowns. Often, these strategies are provided as collections of detailed specific measures, e.g., defining a proposed sequence of actions (such as self-isolation and mandated testing) upon experiencing symptoms (Centers for Disease Control and Prevention 2022).

However, as the COVID-19 pandemic has demonstrated, their organization in practice is very much dependent on the regionally applicable regulatory body, resource availability (for example testing kits), public infrastructure, demographic or social structures of the affected population, or the current state of the pandemic. In Germany, the mode of operation deviated significantly even across federal state borders with varying numbers of permitted outer-household contacts, testing mandates or closure policies. Although, as

we observe in our simulation studies, even nuanced variations in strategy design can have a significant impact on the anticipated effectiveness. Moreover, novel intervention strategies emerged throughout the course of the pandemic such as mobile phone apps for contact tracing (Jalabneh et al. 2021). Additionally, given the already ambiguous context about the definition of interventions, Jewell et al. (2020) find that infectious disease models often lack a sufficient explanation of underlying assumptions. In the context of intervention modeling, these could, for example, be the assumed sensitivity of testing kits (true-positive rate) or the assumed contact tracing success rate.

We suggest that a formalized approach to intervention modeling in individual-based simulations can (1) apply a certain structure to the multifaceted landscape of available intervention strategies improving reusability and comparability and (2) facilitate the explication of underlying assumptions to foster model transparency and eventually enhance the decision makers' trust in the model. Based on the aforementioned considerations, such a formalization shall satisfy the following requirements:

Granularity. Given the various definitions of interventions, models used to investigate their effects must be granular enough to capture their nuances. Even simple-sounding measures like "self-isolation" can vary in duration, isolation location (households vs. special locations), testing requirements, exceptions (e.g., doctor visits), and subsequent measures such as isolating contact persons.

Extensibility. The COVID-19 pandemic has demonstrated that infectious disease mitigation is carried out in a highly dynamic environment and intervention strategies are adjusted once new data (e.g., about the effectiveness of mask-wearing) or new technology (e.g., app-based contact tracing) becomes available.

Expressiveness. All models are abstractions. The complexity of agent-based models can suggest the illusion of realism (Jewell et al. 2020) where implicit assumptions are confounded with the ground truth. For models informing public health decisions potentially impacting the well-being of entire populations, it is imperative to make any assumptions about the evaluated NPI strategies as explicit as possible.

Usability. Public health decision-makers are the stakeholders of both the simulation outcomes but also the modeling tools as potential users. With them rarely being trained software developers, we reason that usability becomes a critical success factor, enabling subject matter experts without a pronounced background in programming to interact with the system.

There are several generalized individual-based NPI modeling approaches in existence, albeit only as dedicated NPI-modules of larger individual-based infectious disease modeling frameworks. The Framework for Reconstructing Epidemic Dynamics (FRED) offers a variety of intervention strategies, both on population- and individual level (Grefenstette et al. 2013; Paparian et al. 2012). There is, however, no uniform formalization or common interface for users to extend custom intervention measures. Intervention mechanisms are injected into the core simulation by means of a supplementary script in the framework's language. The Epidemiological Modeling Software (EMOD) offers one of the most comprehensive frameworks for NPI modeling. It employs so-called "Campaigns" that can be triggered by certain conditions (such as an infection), which possibly contain multiple intervention events (Eckhoff and Wenger 2016). These campaigns can be granularly customized and extended with parameters such as dates, target population, triggering conditions, and further cascading events (Bershteyn et al. 2018). The campaign files explicitly define the when, who/where, why, and what, expressing all assumptions about the interventions clearly. However, EMOD's complexity might come at the cost of usability as they state, a team of professional software developers is working alongside the disease researchers for extending the platform as well as devising disease applications (Bershteyn et al. 2018). The Generic Population Model (GePoc) is a highly versatile agent-based model of the entire Austrian population, prominently used for decision support during the Austrian COVID-19 pandemic (Bicher et al. 2018; Bicher et al. 2022). It provides a very granular intervention interface where strategies are events on a timeline and applied to either person-agentbehavior or locations (Bicher et al. 2022). However, as the code to this primarily commercial tool is not public, we find ourselves unable to conclusively assess the framework's extensibility or usability.

In this paper we introduce the TriSM (Trigger-Strategy-Measure) approach to uniformly model complex non-pharmaceutical interventions in individual-based infectious disease simulations and its concrete implementation in GEMS. We argue that formalizing NPIs with respect to triggers, strategies, and

measures provides a significant contribution to the infectious disease modeling community regarding model extensibility, expressiveness, and comparability. We demonstrate TriSM's capabilities of recreating granular intervention policies in six scenario-based applications. Its usability, as perceived by the potential user group of subject matter experts and public health decision advisors, was verified through a two-day hands-on workshop in early 2024. Chapter 2 provides a brief introduction to the GEMS framework, setting the ground for the proposition of the TriSM approach in Chapter 3, followed by a scenario-based demonstration of its capabilities in Chapter 4 before offering a discussion and conclusion in Chapter 5.

2 THE GERMAN EPIDEMIC MICRO-SIMULATION SYSTEM (GEMS)

The German Epidemic Micro-Simulation System (GEMS) is an agent-based simulation framework for modeling the spread of infectious diseases in heterogeneous populations. The terminology in the following descriptions are borrowed from the Grimm et al. (2020) ODD Protocol for reporting on agent-based models. GEMS was developed with the specific aim of creating a highly adaptable agent-based simulation framework to be applicable as a decision support tool during pandemics. Within this section, GEMS will be briefly introduced, including the relevant components and the internal procedures.

2.1 Entities

The GEMS framework includes two types of entities: individuals, i.e., individuals of the population, and settings, i.e., places where the individuals can have contacts.

Individuals represent the individual members of the population. They possess various attributes, including age, sex, health status, or education level. Individuals are associated with specific settings. The associated settings correspond to the physical or social contexts in which contact with other individuals might occur, such as their household. Each individual is only associated with one setting of a specific type.

Settings are the contexts in which individuals can encounter contacts. Each specific setting belongs to a setting type, e.g., there are multiple households that all belong to the setting type *Household*. Currently, the included setting types are *Households*, *Workplaces* and *Schools*. The framework, however, is open to adding further setting layers, e.g., to model friendship networks. Each specific setting includes a list of individuals that are associated with this setting and a flag describing if the setting currently includes at least one infectious agent. The more general setting type, e.g., *Household*, includes a contact parameter, which denotes the average number of contacts an individual has in a setting of this setting type per simulation tick.

2.2 Internal Processes

GEMS works with a discrete-time model where each tick is of the same length. The following section elaborates on the most central processes that occur once every tick as well as the stepping function.

The **infection procedure** is used to spread the disease among the individuals. The procedure is applied during the step function for each active setting, i.e., each setting containing at least one infectious agent. This is done in order to improve the performance of the simulation. Details on this *active setting approach* are discussed in Ponge et al. (2023). For the specific setting, the infection procedure iterates through all infectious individuals and draws contact persons using the setting type-specific contact distribution parameters. A predefined infection rate is used to evaluate the chance of this encounter being an infectious contact. If so, the contact person's disease progression (from being exposed to being removed) is calculated, and all settings this person is associated with are flagged as active to be handled in the next tick.

The **disease progression** advances individuals in their current disease state according to the underlying SEIR (Susceptible-Exposed-Infectious-Removed) base model as illustrated in Figure 1. Individuals are initially susceptible. Upon infection, they become exposed, i.e., infected but not yet infectious. The individual progresses to becoming infectious and being able to induce further infections. Eventually, the individual either recovers or dies. While GEMS offers the possibility of reinfections, these are not considered in this paper, i.e., individuals acquire full immunity after being infected. The times at which transitions occur are again drawn from distributions that are given by the specific model applied.

GEMS includes further distinctions in the infectious compartment, which describe the severity of the case and are displayed as subcategories of the infectious compartment in Figure 1. Individuals can have an asymptomatic case, i.e., become infectious but not develop symptoms, have a mild case with symptoms developing after some time, a severe case where previously mild symptoms eventually worsen, and individuals are admitted to a hospital with a certain likelihood or a critical case, which in addition to severe symptoms includes guaranteed hospitalization and a likelihood of admittance to intensive care and ventilation. Once individuals are infected, the severity of the case and transition times between different states are calculated based on the likelihoods and distributions provided by the model.



Figure 1: Visualization of the employed SEIR disease progression including the severity states for the infectious compartment.

The **event queue handling** is the part of the simulation procedure that allows for the processing of interventions. Event queues are commonly used in discrete event simulations to handle the flow of simulation time (Henriksen et al. 1986). Here, the event queue is included in GEMS' discrete time simulation to efficiently keep track of the interventions to be executed. The event queue itself is a priority queue that includes all events to be executed sorted by the tick at which they should be executed. Events are themselves either individual or setting measures, which are the key blocks of the TriSM approach and will be discussed in more detail in the next chapter. Processing the events in the queue can lead to cascading events that are then inserted into the sorted event queue.

The **stepping function**, visualized in Figure 2, includes all procedures that should be executed at every tick. A single application of the stepping function advances the simulation by one tick, e.g. one day. In contrast to discrete event simulations, the size of a time-step is therefore fixed. It encompasses all previously mentioned procedures, including first iterating over all individuals and updating their disease state, e.g., if, upon infection, it was calculated that the individual should develop symptoms at the current tick, he will be set as symptomatic during this step. As the second step, it iterates over all setting types and over the settings of this type in parallel to call the infection procedure for those that are active(Ponge et al. 2023). Finally, the event queue is processed, and all events to be executed at the current tick are performed. The internal tick variable is updated, and the simulation moves to the next tick, starting the stepping function anew.



Figure 2: Visualization of the stepping function that advances the simulation one tick.

3 THE TRISM APPROACH

The TriSM (Trigger-Strategy-Measure) approach is being developed to formalize the definition of nonpharmaceutical interventions in individual-based infectious disease models by providing a standardized structure. We suggest that formulating non-pharmaceutical interventions according to TriSM encourages the explication of implicit assumptions about the course of intervention strategies and drastically enhances model expressiveness as it requires the adherence to a standardized notation, comparable to the EMOD approach in Bershteyn et al. (2018). TriSM, however, follows a leaner implementation consisting of three building blocks: Triggers, Strategies, and Measures.

3.1 Triggers, Strategies, and Measures

Figure 3 provides a schematic overview of the TriSM building blocks implemented in GEMS.

Triggers are certain model state change events or criteria that, when met, provoke the execution of an intervention strategy. Such triggers can be associated with individuals, settings, or overall model states. An easy example of a trigger related to an individual is the so-called *SymptomTrigger* which is set off once an individual starts to experience symptoms. *DateTriggers* provoke the execution of a strategy at a certain point in simulation time e.g. to steer predefined lockdowns. *IncidenceTriggers* rely on the current disease dynamics and activate a strategy once a certain threshold incidence value is met. Triggers are implemented as callback hooks into the GEMS framework. For example, the main simulation routine will execute all *SymptomTriggers* once an individual progresses to the symptomatic phase of her disease history.



Figure 3: Triggers, Strategies, and Measures as implemented in GEMS.

Strategies are defined as an enclosing set of concrete intervention measures. Each measure has a property *Offset* defining the number of time steps between the triggering of the strategy and the execution of the measure. Furthermore, strategies can attach so-called *Conditions* and *Delays* to measures. *Conditions* can restrict the execution of measures, e.g. to only apply measures in a certain geographical region or to only process measures once a certain case incidence threshold was met. *Delays* can alter the tick *Offset*, e.g. to model the increasing amount of time it takes to successfully trace a contact person the longer a contact is in the past. A strategy always has a "focal object" of one of the model entities. In the context of our GEMS-based implementation, these are either an individual or a setting. Thus, strategies are divided into *IStrategy* and *SStrategy* types, respectively. An example of an *IStrategy* is the isolation of an individual upon experiencing symptoms and an antigen test after the first week to evaluate their qualification for a premature release. This example corresponds to simulation scenario 3, presented in chapter 4.2.

Measures are the building blocks of all intervention strategies. They define single actions with examples such as isolating an individual, testing an individual for a certain pathogen, or closing a particular school. Moreover, measures can be passive in nature, i.e., they are not used to directly affect the simulation but rather "detect" certain relationships in the model. Examples of such measures are backward tracing of infectious contacts or the detection of all members of a workplace. Processing these measures provides a list of new focal objects (contact persons, workplace members, etc.) to trigger subsequent strategies for each of them. This makes it easy to piece together more complex scenarios (e.g. tracing infectious contact persons and isolating their respective households).

3.2 Processing Strategies and Measures

The pseudocode below illustrates how triggered strategies are being processed. For simplicity reasons, the code blocks omit measure *Delays* and *Conditions*. The function generates one event for each of the measures associated with the strategy and its focal object (individual or setting). All events are then pushed to the event queue and labelled with the expected execution time (sum of current tick and measure offset).

```
function trigger_strategy(strategy, focal_object)
    foreach (offset, measure) in strategy.measure_list
        event_queue.enqueue(tick + offset, Event(measure, focal_object))
        end
end
```

Each measure type (e.g. *SelfIsolation* or *TraceContacts*) has a dedicated *process_measure* function. Julia's dispatcher ensures that the correct function is being executed dependent on the argument measure's type. After performing all model state changes (e.g. isolating individuals or tracing infectious contacts), the function returns either *nothing* (concluding the measure execution) or a tuple consisting of a vector of new focal objects (e.g. contact persons) and a subsequent strategy (e.g. *isolate_contact_persons*) as seen in the code below. The latter will trigger the provided strategy for all focal objects using the code above. The order in which measures are being processed is given by their scheduled execution time in the event queue.

```
function process_measure(measure, focal_object)
    ...// change model state, e.g. "isolating" individual
    return {nothing, (Vector<focal_objects>, strategy)}
end
```

This approach fosters extensibility as implementing a new type of measure (e.g. app-based contact tracing as mentioned in the introduction) can be supplemented by defining a new measure struct and providing the corresponding *process_measure* function. The unified function signatures ensure that new measures can be seamlessly combined with preexisting ones. Moreover, Julia's packaging architecture permits users to supply custom measures without requiring the adaptation of GEMS' code base.

4 APPLICATIONS

In the following, we present six simulated infection scenarios illustrating TriSM's capability of modeling fine-grained interventions within the GEMS framework. The scenarios are deliberately focused on nuanced variations of self-isolation- and testing strategies, demonstrating how TriSM supports the subtle adjustment of intervention strategies, which can, according to our simulations, result in significant variances in their effectiveness assessment.

4.1 Base Infection Model

The population being used to evaluate the example scenarios was synthetically generated and includes 100,000 individuals. While actual representations of real-world populations can be created based on census data (Ponge et al. 2021), we rely on a completely artificial homogeneous population to isolate the effect of the evaluated intervention strategies in the scenarios. All individuals are associated with one household and one workplace each, such that the average workplace size is equal to 50.15 and the average household size is equal to 2.3, comparable to the average household size in Germany (Statistisches Bundesamt 2024). This corresponds to a total of 1,994 workplaces and 50,134 households. Their size distributions are displayed in Figure 4 (a). The number of contacts per tick in the household and workplace settings follows a Poisson distribution with $\lambda = 1$ and $\lambda = 2$, respectively.



Figure 4: (a) Setting size distributions for Workplaces and Households, (b) disease progression.

The simulated disease has an infection likelihood of 10% for contact between infectious and non-infected individuals. Every infection has a symptomatic (mild) progression corresponding to developing symptoms on average three days after exposure, following a Poisson distribution ($\lambda = 3$), and becoming infectious on average one tick earlier, following a Poisson distribution ($\lambda = 1$). No deaths are included in the model such that all individuals recover on average seven days after developing symptoms. The duration of the symptomatic phase follows a Poisson distribution ($\lambda = 7$). Individuals acquire full immunity after an infection and reinfections are not considered. Figure 4 (b) displays the distribution of disease states relative to the days after exposure. The initial state of the simulation corresponds to 1% of the population being infected with the modelled disease (1,000 individuals). The simulation terminates after 120 days.

4.2 Intervention Scenarios

We introduce five intervention scenarios using the structure offered by TriSM all regarded to isolation and testing of symptomatic individuals. Additionally, we provide a baseline scenario of the unmitigated epidemic progression for the previously defined disease. Simplistic examples were chosen to illustrate the granularity of the TriSM approach and the relevance of such an approach in infectious disease modeling, as minor differences in simple interventions can have a major influence on the intervention's effectiveness.

Figure 5 provides a schematic overview of the five intervention scenarios using TriSM building blocks. Green shapes represent triggers; red boxes correspond to strategies, and blue ones to measures. As mentioned in Chapter 3, a strategy can be interpreted as a set of measures. Due to the structure of TriSM the components and parametrizations of the scenarios can be clearly (and graphically) formalized. The scenarios start for each infected individual once they develop symptoms (symptom triggers). The structure of TriSM now allows for an easy combination and granular parametrization of the individual components enabling the construction of the different scenarios. In the first two scenarios, this causes symptomatic individuals to go into household isolation immediately and stay there for 14 days (Scenario 1) or seven days (Scenario 2). This will effectively prevent all subsequent infections at workplaces but can still cause infections in households. Scenario 3 sends symptomatic individuals in isolation immediately for a maximum of 14 days but subjects all individuals to a test on day seven. A positive test causes an individual to remain isolated for the remainder of the time. A negative result will trigger the premature release strategy ("end isolation". The test sensitivity (true-positive rate) and specificity (true-negative rate) is 100%. The test will always correctly detect a current infection. Scenarios 4 and 5 assume that only a positive test conditions isolation. The scenarios suppose that test results will be available after two days. While a positive test will trigger a 14-day isolation in both scenarios, we vary the test sensitivity between 100% (Scenario 4) and 50% (Scenario 5). A test with 50% sensitivity would probably not be authorized for public use; however, in our scenario, this value also includes false negatives caused by improper application.



Figure 5: Five intervention scenarios combining isolation and testing measures.



Figure 6: Results of 100 runs for each scenario. Dashed lines correspond to the arithmetic means, low opacity areas indicate the range of values and high opacity the 95% confidence bands.

4.3 Simulation Results

We ran the baseline and five intervention scenarios 100 times each. The results are presented in Figure 6 where each row corresponds to one of the scenarios. The columns illustrate daily new infections (including the 1,000 initial infections), applied tests, and the number of currently isolated individuals per day. Dashed lines represent the arithmetic mean values of the three variables across the 100 simulation runs, colored areas indicate the value ranges across all simulation runs and the confidence bands in low and high opacity, respectively. The arithmetic mean totals of the three variables for each scenario are given in Table 1.

Table 1: Arithmetic mean of the cases, tests and days spent in quarantine for the baseline and five intervention scenarios.

	Cases	Tests	Quarantine Days	R _{eff}
[Baseline]: No interventions	74,852.57	0	0	1.87
[Scenario 1]: Isolation (14)	3,093.34	0	41,838.58	0.70
[Scenario 2]: Isolation (7)	4,572.61	0	30,707.18	0.82
[Scenario 3]: <i>Isolation</i> (14) + <i>Test</i> (7, 100%)	3,108.29	3,000.52	29,449.30	0.70
[Scenario 4]: Test (2, 100%) + Isolation (14)	15,631.53	14,775.56	198,186.10	1.11
[Scenario 5]: Test (2, 50%) + Isolation (14)	53,525.39	50,639.63	340,814.81	1.51

It becomes apparent from the baseline scenario that an unmitigated epidemic causes an attack rate of ~74,8%, i.e., infects roughly three-quarters of the population within the first 120 days in this hypothetical model. This corresponds to an effective reproduction rate (R_{eff}) of 1.87 for the unmitigated progression. Isolating individuals for 14 days after they become symptomatic pushes R_{eff} below 1, thus cutting the total attack rate to around 3,1%, containing a large outbreak (Scenario 1). However, this strategy causes mass isolations accumulating to ~42 thousand person-days (PDs) spent in quarantine. Reducing the isolation duration to seven days (Scenario 2) effectively reduces the PDs spent in isolation by roughly 26,6% to a mean total of 30,707 days but increases the total case numbers by almost 48% to 4,572 as individuals who are still infectious are being released from isolation. Given that the symptomatic period follows a Poisson distribution with $\lambda = 7$ the cumulative distribution function shows that ~40% of the individuals are infectious. In Scenario 3, we observe that it is possible to combine both, the low infection numbers from Scenario 1 with the reduced isolation days from Scenario 2 by only releasing individuals at day seven who tested negative. However, this strategy requires the availability of at least 3,000 testing kits.



Figure 7: (a) Mean effective generation time in Scenario 4 throughout the simulation, (b) percentage of individuals that are infectious and (un)tested relative to the days after exposure.

Scenarios 4 and 5 emphasize the importance of intervention timing and test availability. Isolating individuals only after a positive test with a two-day delay, e.g. due to a limited testing capacity, causes a

five-fold spike in infections (15.6% attack rate) and tests (14.775 tests) compared to Scenario 3 and an almost seven-fold increase in PDs spent in isolation (198,186 PDs) in Scenario 4. In both cases, the effective reproduction rate is above the criticality threshold (1.11 and 1.51). The reason for the drastic increase becomes apparent when looking at the average difference in exposure times between two direct successors in an infection chain, also referred to as *Generation Time*. Figure 7 (a) shows that the mean empirical generation time in Scenario 4 is around four and a half days in the beginning and slowly moving towards five days towards the end of the simulation time. Figure 7 (b) displays the percentage of infectious and (un)tested individuals relative to the days after their exposure. Given that the onset of symptoms happens at day three on average, a two-day isolation delay (triggered by entering the symptomatic phase) suggests that individuals spend the entire mean generation time non-isolated in this scenario. This is also visually displayed in Figure 7 (b). Here the onset of testing is at tick three while already more than 50% of individuals have become infectious. However, starting from day four after exposure, the number of untested infectious individuals, which in Scenario 4 are always quarantined, decreases. This causes the overall case numbers to be significantly reduced relative to the baseline but performs worse than immediate action. In the case of Scenario 5 it will require 17 times the number of testing kits and cause 12 times the number of PD in isolation as Scenario 3 while not being able to contain the outbreak and only preventing $\sim 28.5\%$ of the infections compared to the baseline scenario. This significantly less effective scenario is due to the reduced test sensitivity of 50%, resulting in only 50% of the tested individuals in Figure 7 (b) being identified as positive cases and thus sent into isolation.

5 DISCUSSION & CONCLUSION

The previous scenarios demonstrate that subtle variations of intervention strategy design can induce significant variations in their effectiveness assessment and how the TriSM approach can reproduce these nuances in detail. Furthermore, as all measures are processed with an associated "focal object", it would even be theoretically viable to have each individual follow a distinct set of intervention strategies. This would certainly denote the *most granular* approach possible to intervention modeling, albeit being highly unlikely in practice.

The interface-based architecture of TriSM allows users to supplement custom intervention measures and thus making the approach highly extensible. It enables disease modelers to react quickly to the introduction of new technologies during public health crises (such as app-based contact tracing).

We argue that formalizing interventions using the TriSM approach induces a high level of expressiveness regarding their parameterization. The reporting of implicit assumptions is not primarily dependent on the completeness of the model documentation, which might be the case for FRED (Grefenstette et al. 2013), but is inherently present in the composition of measures that constitute the modeled strategies. While one could argue that custom measures can still induce implicit model state changes in the respective implementation of the *process_measure* function, the TriSM-based formalization would indicate the explicit inclusion of such custom measures. The core codebase remains untouched.

Given that the target audience of such a modeling approach predominantly consists of decision-makers and infectious disease modelers with varying technical backgrounds, usability is paramount. In order to validate our modeling approach, we introduced the GEMS framework, including the TriSM formalization, to an expert group during the 2nd National Conference on Infectious Disease Modeling of the Modeling Network for Severe Infectious Diseases (MONID) in Halle, Germany. More than 20 peers, including but not limited to epidemiologists, infectious disease modelers, public health decision-makers, statisticians, social scientists, and computer scientists, participated. In a concluding user experience survey with 17 of the participants, 70% answered the question, "*I found it easy to model intervention scenarios with TriSM*", with *tend to agree* or *agree completely*. While all attendees voiced their interest in using TriSM and GEMS, half of the participants were eager to even implement custom NPIs using the TriSM formalization. The feedback made us confident to believe that the interfaces TriSM offers are also usable by non-technicians, e.g., decision-makers. Participants also highlighted the very expressive nature of TriSM, suggesting that GEMS and TriSM have another strong application in teaching infectious disease modeling.

Our experience in building individual-based simulations of infectious diseases taught us that modeling non-pharmaceutical interventions remains a delicate topic. In this work, we presented a standardized formalization that can facilitate the design and evaluation of nuanced intervention strategies. Moreover, as our workshop results suggest, it enhances the expressiveness of models and enables non-experts, e.g., decision-makers, to comprehend the "mechanics" of interventions. We focused on isolation- and testing strategies exclusively to exemplify the importance of granular intervention design. The GEMS framework, however, contains all necessary TriSM building blocks to simulate further measures such as contact tracing, isolation of entire households, school- or workplace closure, and pool testing. Upcoming extensions will also include pharmaceutical interventions, notably vaccination programs. TriSM was implemented using the GEMS framework, although we suggest that the approach of formalizing interventions into triggers, strategies, and measures is not exclusive to GEMS, thus providing a contribution to the broader individualbased infectious disease modeling community. Furthermore, subsequent research could theorize about the generalizability of TriSM's implementation and offer a universal plug-and-play package that might be usable in simulations going beyond the infectious disease modeling context.

So far, all simulations carried out with GEMS and the TriSM approach are based on hypothetical scenarios and a clear limitation compared to established frameworks such as EMOD, which have been utilized in multiple infectious disease research studies for over a decade (Bershteyn et al. 2018). While we argue that the TriSM formalization yields a clear benefit in terms of usability, it remains to be evaluated whether this proposition can hold up when employing it for actual policy advice. We propose that future research subject the established frameworks and GEMS to a comparing case study of a real-world infectious disease scenario, evaluating their granularity, expressiveness, extensibility, and usability from a decision-maker perspective.

ACKNOWLEDGMENTS

This work is part of the "OptimAgent" project. The project "OptimAgent" was funded by grants (no. 031L0299X) of the German Federal Ministry of Education and Research (BMBF). "OptimAgent" is part of the Modeling Network for Severe Infectious Diseases (MONID). The authors are responsible for the content of this publication.

REFERENCES

- Bershteyn, A., J. Gerardin, D. Bridenbecker, C. W. Lorton, J. Bloedow, R. S. Baker et al. 2018. "Implementation and Applications of EMOD, an Individual-Based Multi-Disease Modeling Platform". *Pathogens and Disease* 76(5).
- Bicher, M., C. Rippinger, D. Brunmeir, C. Urach, M. Zechmeister, and N. Popper. 2022. Agent-Based SARS-CoV-2 Simulation Model: Model Specification. https://www.dwh.at/projects/covid-19/Covid19_Model-20230322.pdf, accessed 11th April 2024.
- Bicher, M., C. Urach, and N. Popper. 2018. "Gepoc ABM: a generic agent-based population model for Austria". In 2018 Winter Simulation Conference (WSC), 2656–2667. IEEE Press. https://doi.org/10.1109/WSC.2018.8632170
- Centers for Disease Control and Prevention. 2022. Ending isolation and precautions for people with COVID-19: interim guidance. https://stacks.cdc.gov/view/cdc/113474, accessed 11th April 2024.
- Eckhoff, P. A., and E. A. Wenger. 2016. "The EMOD Individual-Based Model". In *Spatial Agent-Based Simulation Modeling in Public Health*, 185–208. John Wiley & Sons, Ltd.
- Grefenstette, J. J., S. T. Brown, R. Rosenfeld, J. DePasse, N. T. B. Stone, P. C. Cooley et al. 2013. "FRED (a Framework for Reconstructing Epidemic Dynamics): an open-source software system for modeling infectious diseases and control strategies using census-based populations". *BMC Public Health* 13(1):940.
- Grimm, V., S. F. Railsback, C. E. Vincenot, U. Berger, C. Gallagher, D. L. DeAngelis et al. 2020. "The ODD Protocol for Describing Agent-Based and Other Simulation Models: A Second Update to Improve Clarity, Replication, and Structural Realism". *Journal of Artificial Societies and Social Simulation* 23(2).
- Henriksen, J. O., R. M. O'Keefe, C. D. Pegden, R. G. Sargent, and B. W. Unger. 1986. "Implementations of time (panel)". In 1986 Winter Simulation Conference (WSC), 409–416. https://doi.org/10.1145/318242.318467

- Jalabneh, R., H. Z. Syed, S. Pillai, E. H. Apu, M. R. Hussein, R. Kabir et al. 2021. "Use of mobile phone apps for contact tracing to control the COVID-19 pandemic: A literature review". *Applications of artificial intelligence in COVID-19* :389–404.
- Jewell, N. P., J. A. Lewnard, and B. L. Jewell. 2020. "Predictive mathematical models of the COVID-19 pandemic: Underlying principles and value of projections". *Jama* 323(19):1893–1894.
- Matrajt, L., and T. Leung. 2020. "Evaluating the Effectiveness of Social Distancing Interventions to Delay or Flatten the Epidemic Curve of Coronavirus Disease". *Emerging infectious diseases* 26(8):1740–1748.
- Paparian, J., S. Brown, D. Burke, and J. Grefenstette. 2012. "FRED Navigator: An interactive system for visualizing results from large-scale epidemic simulations". In 2012 IEEE 8th International Conference on E-Science, 1–5.
- Ponge, J., D. Horstkemper, B. Hellingrath, L. Bayer, W. Bock, and A. Karch. 2023. "Evaluating Parallelization Strategies for Large-Scale Individual-based Infectious Disease Simulations". In 2023 Winter Simulation Conference (WSC), 1088–1099. IEEE. https://doi.org/10.1109/WSC60868.2023.10407633
- Ponge, J., M. Enbergs, M. Schungel, B. Hellingrath, A. Karch, and S. Ludwig. 2021. "Generating Synthetic Populations Based On German Census Data". In 2021 Winter Simulation Conference (WSC), 1–12. IEEE. https://doi.org/10.1109/WSC52266.2021.9715369
- Prem, K., Y. Liu, T. W. Russell, A. J. Kucharski, R. M. Eggo, N. Davies et al. 2020. "The Effect of Control Strategies to Reduce Social Mixing on Outcomes of the COVID-19 Epidemic in Wuhan, China: A Modelling Study". *The Lancet. Public health* 5(5):e261-e270.

Robert Koch-Institut. 2017. "Nationaler Pandemieplan Teil I". https://doi.org/10.25646/112, accessed 11th April 2024.

Statistisches Bundesamt. 2024. Haushalte und Familien: Haushalte nach Haushaltsgröße und Haushaltsmitgliedern. https://www.destatis.de/DE/Themen/Gesellschaft-Umwelt/Bevoelkerung/Haushalte-Familien/Tabellen/1-2-privathaushaltebundeslaender, accessed 3rd April 2024.

World Health Organization. 2023. "Infection prevention and control guideline for coronavirus disease 2019 (COVID-19): Executive summary", World Health Organization. https://iris.who.int/handle/10665/375233, accessed 11th April 2024.

AUTHOR BIOGRAPHIES

JOHANNES PONGE is a Research Assistant at the Chair for Information Systems and Supply Chain Management at the University of Münster, Germany. He holds an M.Sc. in Information Systems. His work focuses on the development of simulation-based decision support systems for infectious disease mitigation and intervention. He is a developer of the German Epidemic Micro-Simulation System (GEMS). His email address is johannes.ponge@ercis.uni-muenster.de.

JANIK SUER is a Research Assistant at the Clinical Epidemiology Unit at the University of Münster, Germany. He holds an M.Sc. in Physics. The focus of his research is on the development of large agent-based infectious disease models and the parametrization of contact. He is a developer of the German Epidemic Microsimulation System (GEMS). His email address is janik.suer@uni-muenster.de.

BERND HELLINGRATH is a Professor and Head of the Chair for Information Systems and Supply Chain Management at the University of Münster, Germany. His research interests deal with the broader area of modeling and simulating with a distinct focus on the context of crisis-management and humanitarian logistics. His email address is bernd.hellingrath@ercis.uni-muenster.de.

ANDRÉ KARCH is a Professor of Clinical Epidemiology, and the Head of the Clinical Epidemiology Unit at the University of Münster, Germany. He holds an MD and an MSc in Epidemiology. His research in the field of infectious disease epidemiology focuses on dynamic transmission models for public health decision-making. His email address is akarch@uni-muenster.de.