Strategic NPI Optimisation: Solving a MOOP to Mitigate Healthcare Burden and Educational Loss

Leonardo Angeli[1,*], Alejandro Morales-Hernandez[2], Inneke Van Nieuwenhuyse[1], Lander Willem[3] and Niel Hens[1,3]

[1] Hasselt University, Belgium [2]; Université Libre de Bruxelles, Belgium [3]; University of Antwerp, Belgium

[*] Contact details: leonardo.angeli@uhasselt.be

Introduction

Background: Pandemic control involves measures that reduce transmission and protect healthcare capacity, but that can also disrupt education and broader societal well-being. Therefore, we frame this challenge as a multi-objective decision problem to identify intervention strategies that balance this trade-off.

Aim: We develop a simulation-based Digital Twin framework for Belgium to explore school- and workplace-related contact-scaling schedules from November 2020 to February 2021. Candidate strategies are evaluated by two objectives: minimising hospitalisation burden and preserving in-person education.

Framework

- **Decision space:** 64 contact-scaling variables $\mathbf{P}_T = [\mathbf{p}_1, \ldots, \mathbf{p}_8]$– 4 controlled school settings, 4 workplace settings over 8 two-week intervals (T).
- **Digital twin:** each candidate schedule reconstructs age-specific contact matrices from a baseline contact structure, and is propagated through a stochastic, age-structured SARS-CoV-2 transmission model.
- **Evaluation:** contact-scaling schedules are compared through the corresponding educational loss (EL) and hospitalisation burden (HB), both minimized.

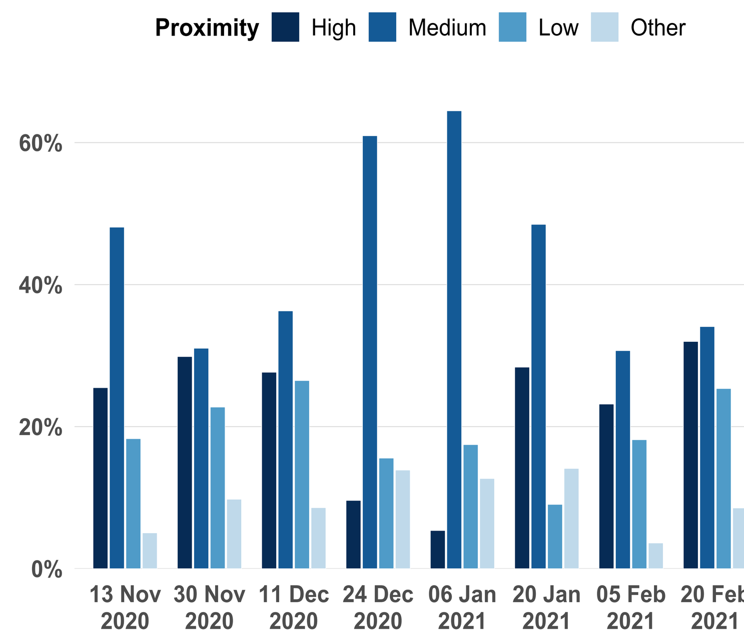

Figure 1: Contact share in workplace settings (%), by proximity level and survey wave.

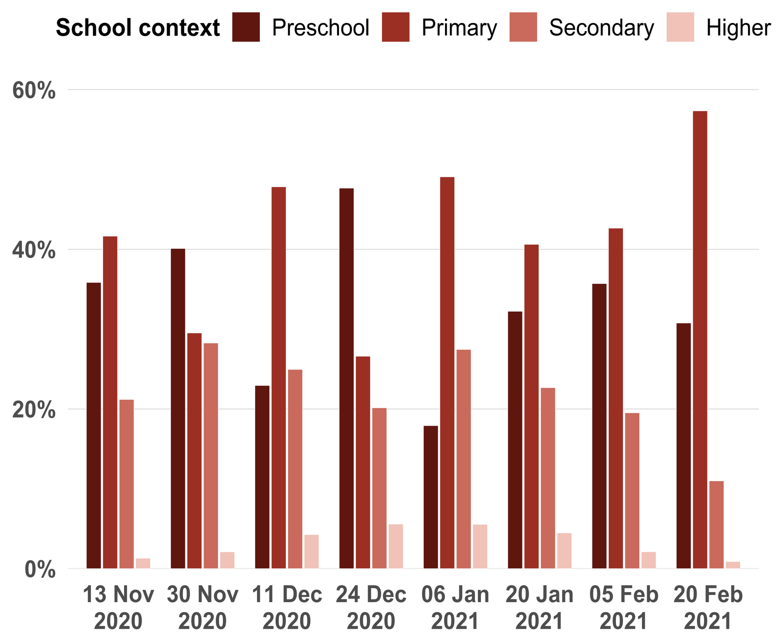

Figure 2: Contact share in school settings (%), by educational context and survey wave.

Optimisation problem

Each candidate's strategy is a contact-scaling schedule $\mathbf{P}_T^* = \{p_{l,\tau}\}$, with $p_{l,\tau} \in \left\{\frac{n}{100} \,\middle|\, n = 0, 1, \ldots, 100\right\}$.

$$\min_{\mathbf{P}_T^*} \left(EL(\mathbf{P}_T^*),\ HB(\mathbf{P}_T^*)\right)$$

$$EL = \sum_{\tau \in T} \sum_{l \in L_{\text{school}}} C_l \left(1 - \frac{p_{l,\tau}}{100}\right) w_{l,\tau}, \qquad HB = \sqrt{\frac{1}{|T|} \sum_{\tau \in T} \left[\sum_{\gamma} w_\gamma HB_\gamma(\bar{h}_{\gamma,\tau})\right]^2}.$$

- **Educational loss:** proxy for reduced in-person learning opportunities; higher penalties for primary and secondary education.
- **Hospitalisation burden:** simulated ICU and non-ICU occupancy relative to available COVID-19 bed capacity; higher penalties for ICU occupancy.
- **Data inputs:** social contact surveys, hospital-capacity data, and time-varying epidemiological parameters in an age-structured stochastic model.

Algorithmic optimisation approach

- **NSGA-II:** evolutionary multi-objective algorithm used as a traditional benchmark for approximating the Pareto front in high-dimensional decision spaces.
- **MOTPE:** Bayesian optimisation approach based on tree-structured Parzen estimators, used to guide the search towards promising regions of the objective space.
- **BO-GPR:** surrogate-assisted Bayesian optimisation using Gaussian-process regression and hypervolume-based acquisition to prioritise candidate strategies likely to improve the Pareto front.

Results

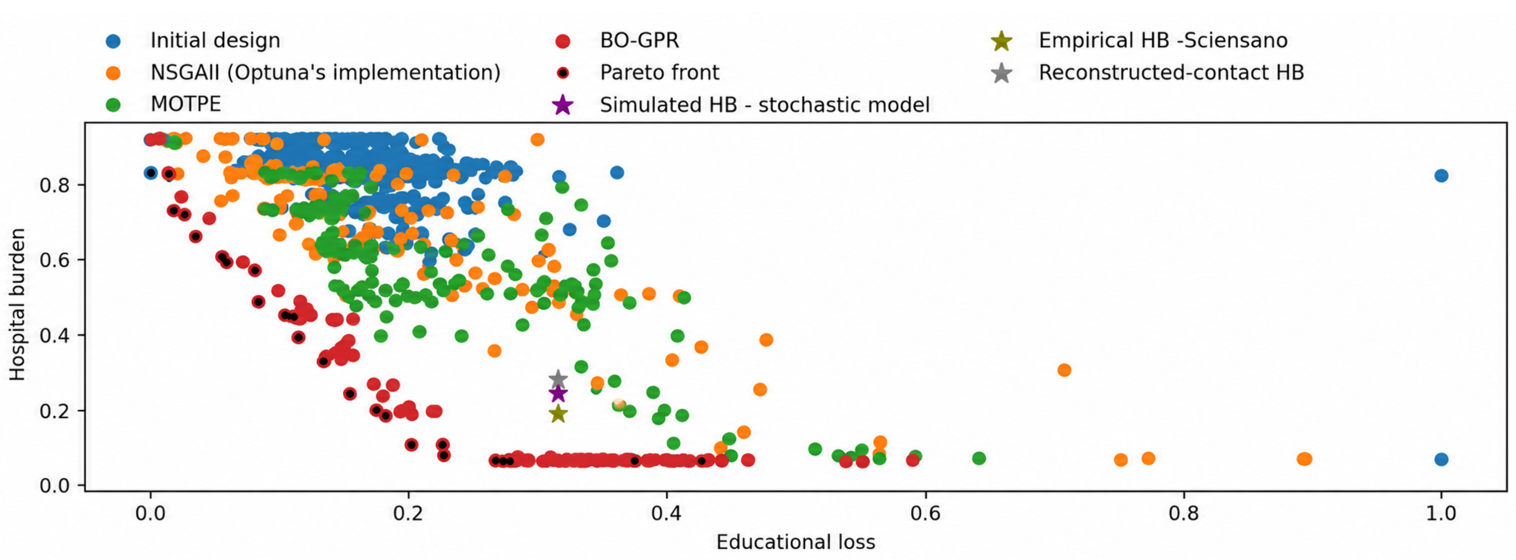

Figure 3: Pareto front exploration. Points show evaluated contact-scaling schedules $\mathbf{P}_T^*$; Stars denote benchmark objective values during the study period in Belgium.

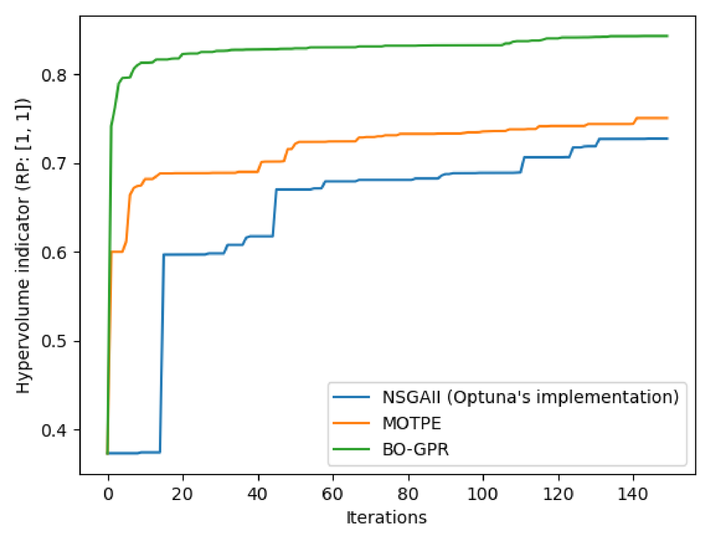

Figure 4: Hypervolume evolution across optimisation iterations. Larger values indicate better Pareto-front coverage; BO-GPR gives the strongest approximation.

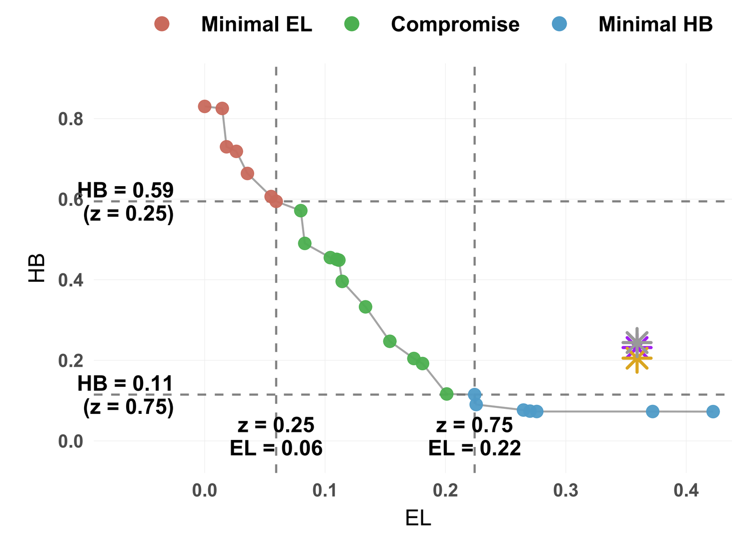

Figure 5: Pareto-front regions for decision-variable interpretation. Points on the front are ranked by increasing EL and grouped into minimal-EL, compromise, and minimal-HB regions.

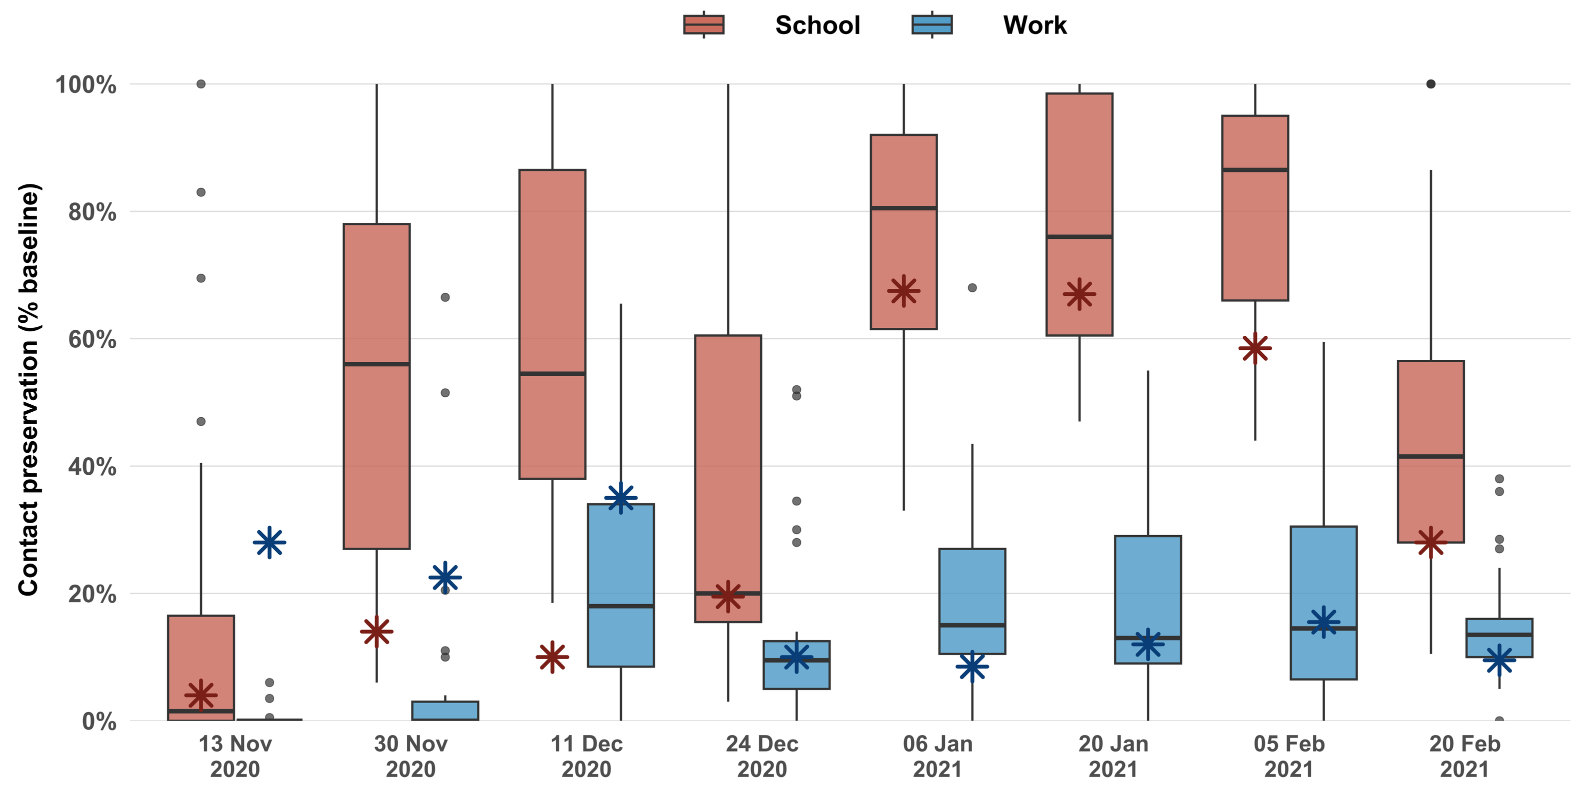

Figure 6: Non-dominated contact-scaling values for School and Work across the study period. Boxes show the interquartile range of non-dominated schedules; stars denote the average reconstructed contact levels during the study period in Belgium.

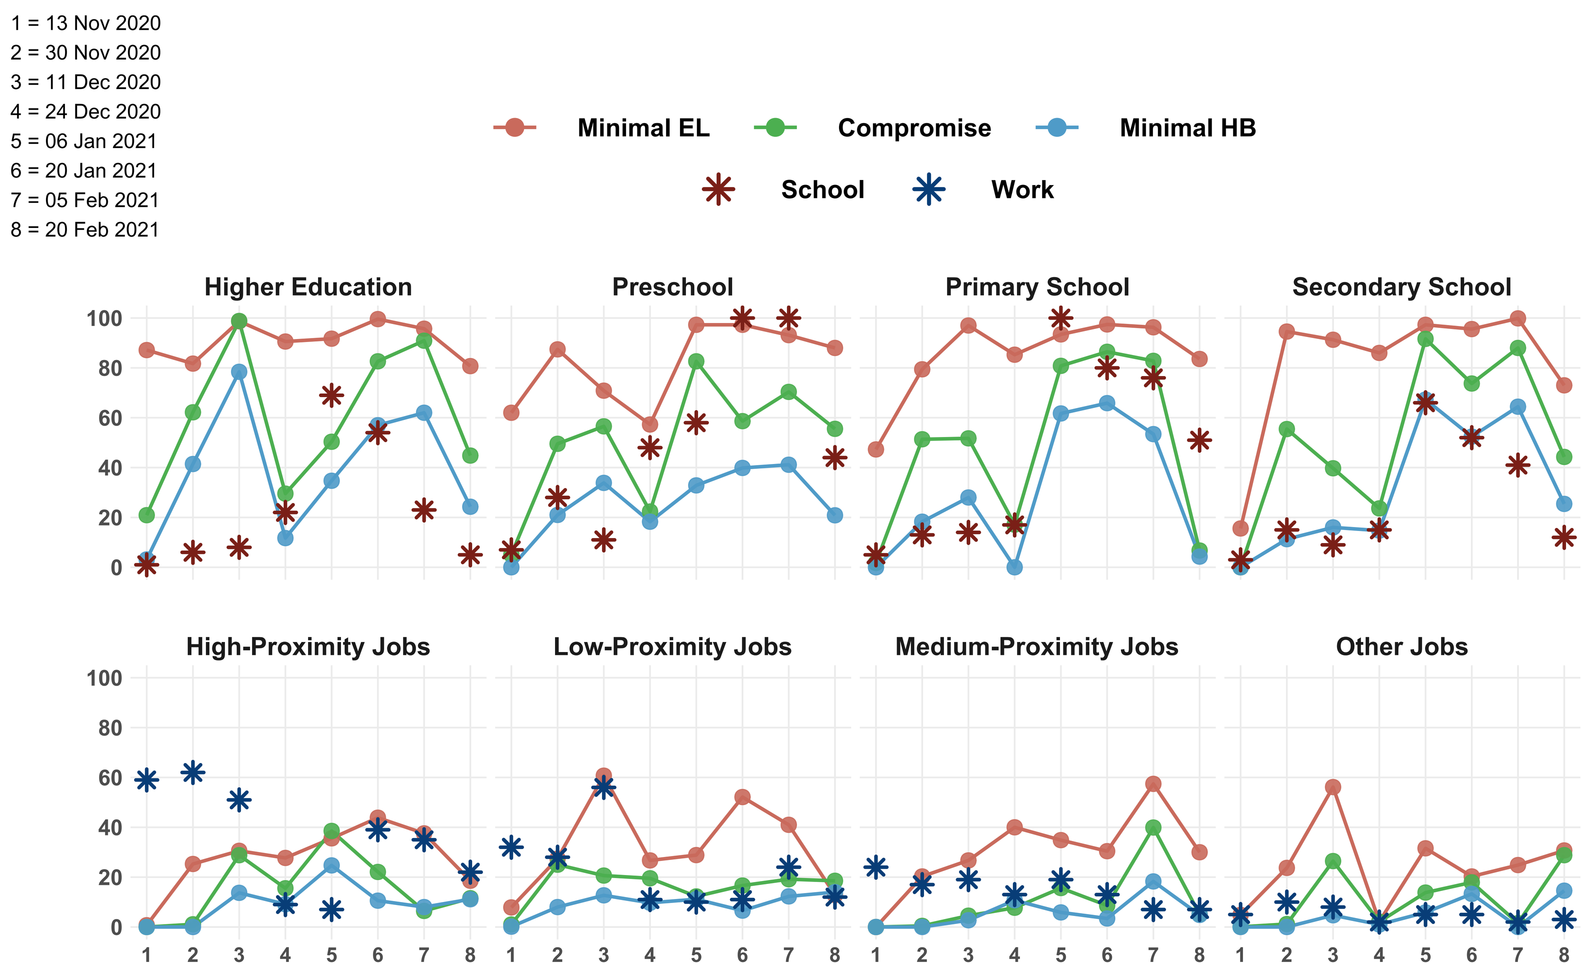

Figure 7: Setting-specific contact-scaling profiles by Pareto-front region. Points show average non-dominated schedules across the study period; stars indicate average reconstructed setting-specific contact levels during the study period in Belgium.

Take-home messages

- The optimisation identifies a clear Pareto front of non-dominated contact-scaling schedules $\mathbf{P}_T^*$, including solutions improving upon the benchmark objective values.
- BO-GPR provides the strongest Pareto-front approximation, although at a higher computational cost.
- Pareto-set schedules preserve more school contacts and reduce workplace contacts more strongly, especially from mid-November to mid-December 2020.
- Preserving primary and higher-education contacts early on, while reducing high-proximity workplace contacts more strongly, can improve upon the benchmark objectives under the model assumptions.
- Results are conditional on the objective definitions, contact-matrix reconstruction, and simulation model; stochastic variability is averaged across runs but not propagated into uncertainty-aware dominance statements.

Acknowledgments

This work has been supported by the Data Science Institute (DSI), CENSTAT, and SIMID research groups. A. Morales-Hernández received support from the Interreg project MULTIROOFS NWE0400462.

This work has been presented as a poster at the event "Digital Twins and AI in Personalized Healthcare (DT4PH): from Innovation to Implementation and Impact"

DSI DATA SCIENCE INSTITUTE UHASSELT · European Commission · UHASSELT KNOWLEDGE IN ACTION · Universiteit Antwerpen · Health Economics & Modelling Infectious Diseases, Vaccine & Infectious Disease Institute, University of Antwerp · ULB · MLG MACHINE LEARNING GROUP