Investigating Fairness in Epidemic Control

Sam Vanspringel[1*][0009-0007-4194-9262], Alexandra Cimpean[1*][0000-0002-6559-2831], Ann Nowé[1][0000-0001-6346-4564], and Pieter Libin[1,3][0000-0003-3906-758X]

Vrije Universiteit Brussel, Brussel, Belgium, *: equal contribution
{sam.vanspringel,ioana.alexandra.cimpean,ann.nowe,pieter.libin}@vub.be

Abstract. Effective pandemic mitigation strategies should reduce both disease burden (e.g., hospitalizations) and societal impact. Previous work approached this challenge with a multi-objective reinforcement learning method that balances hospitalizations against the loss of social contacts (i.e., social burden). To build on this, we recognize that fairness should also be considered when designing mitigation strategies. To this end, we introduce social burden fairness, which aims for a more equitable distribution of social burden, weighted by age-specific hospitalization risk. We experimentally evaluate our fairness-based approach against a baseline that jointly optimizes hospitalizations and overall social burden, but does not treat fairness as an explicit optimization objective. We then investigate how integrating age-oriented fairness constraints influences the learned mitigation strategies.

Keywords: COVID-19 · Fair epidemic mitigation · Multi-objective Reinforcement Learning

Due to the complex and non-linear nature of epidemics, developing efficient mitigation strategies is challenging [5, 7]. Key priorities during an epidemic are to prevent excessive hospitalizations and keep social restrictions to a minimum [6, 8]. To balance these possibly conflicting objectives, multi-objective reinforcement learning (MORL) offers an approach to learn mitigation strategies in complex epidemic models [5]. These policies can assist decision-makers in real-life epidemic settings.

Previous work optimizes hospitalizations and overall social burden using Pareto Conditioned Networks (PCN) [7]. However, representing the social burden as overall lost contacts in the population is a simplification, particularly when hospitalization risks vary between age groups. To address this, we define a fairness notion that captures the lost contacts weighted by age-based hospitalization risk.

To learn policies, PCN restricts social contacts in school, work, and leisure. Additionally, we employ the Multi-Objective Belgian COVID gym environment (MOBelCov) to simulate the epidemic trajectories affected by these restrictions [7]. The environment encapsulates the age-structured compartment model introduced by Abrams et al., which was fitted to incidence data and seroprevalence data of the Belgian COVID-19 epidemic [1]. PCN iteratively adjusts the proportion of contacts allowed in each domain, aiming to balance hospitalizations ($\mathcal{R}_{ARH}$) and social burden fairness ($\mathcal{R}_{SBF}$). We then investigate how integrating age-oriented fairness constraints influences the learned mitigation strategies.

To define social burden fairness, we measure the loss of social contacts due to the restrictions on school, work, and leisure [7]. It considers both susceptible (S) and recovered (R) individuals, comparing the original contact matrix C with the reduced contact matrix $\hat{C}$. The lost contacts of these non-infected groups are weighted by the hospitalization risk (h_i) of the different age groups (i).

$$\begin{aligned}\mathcal{R}_{SBF}(s,a,s') = -\Big(&\sum_{i=1}^{K}\sum_{j=1}^{K}\big[(\hat{C}-C)_{ij}\, S_i(s)\tfrac{1}{h_i} + (\hat{C}-C)_{ij}\, S_j(s)\tfrac{1}{h_j}\big] \\ &+ \sum_{i=1}^{K}\sum_{j=1}^{K}\big[(\hat{C}-C)_{ij}\, R_i(s)\tfrac{1}{h_i} + (\hat{C}-C)_{ij}\, R_j(s)\tfrac{1}{h_j}\big]\Big)\end{aligned}$$

As the ICU capacity was a key factor in determining mitigation measures in Belgium, we compare two policies that accumulate at most 2,000 ICU patients at their peaks throughout the epidemic trajectories. [3, 2].

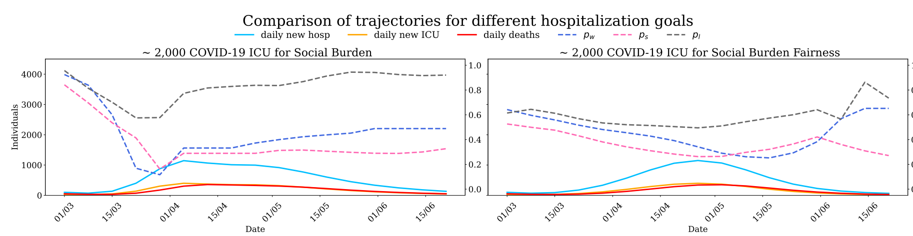

Fig. 1. Comparing two epidemiological trajectories for PCN optimized on [$\mathcal{R}_{ARH}$, $\mathcal{R}_{SB}$] (left) [7] and [$\mathcal{R}_{ARH}$, $\mathcal{R}_{SBF}$] (right) with a budget of 5 changes. The epidemiological state is illustrated by daily new hospitalizations, ICU patients, and deaths. The dashed lines depict the proportional amount of contacts allowed per domain.

Figure 1 shows the difference in learned policies. The left policy shows a loose approach considering leisure contacts, while school and work are restricted notably at the beginning. These restrictions are then reduced slowly. In a policy optimized for $\mathcal{R}_{SBF}$ (Fig. 1, right), the restrictions are less divergent and increased gradually. Towards the end of the trajectory, the restrictions on work and leisure are decreased the most. The restriction on schools remains roughly consistent.

Disease burden and societal impact are not the sole epidemic objectives. Future research could investigate integrating economic objectives alongside fairness considerations to improve and broaden the mitigation strategies. Furthermore, evaluating PCN in epidemic contexts with varying hospitalization risks, such as the 1918 influenza pandemic, would deepen this work [4].

References

1. Abrams, S., Wambua, J., Santermans, E., Willem, L., Kuylen, E., Coletti, P., Libin, P., Faes, C., Petrof, O., Herzog, S.A., Beutels, P., Hens, N.: Modelling the early phase of the belgian covid-19 epidemic using a stochastic compartmental model and studying its implied future trajectories. Epidemics **35** (6 2021)
2. Belgian Health Care Knowledge Centre: Performance of the Belgian health system – report 2024. Kce report, Belgian Health Care Knowledge Centre (KCE) (2024)
3. Faes, C., Hens, N., Gilbert, M.: On the timing of interventions to preserve hospital capacity: lessons to be learned from the belgian sars-cov-2 pandemic in 2020. Archives of Public Health **79**(1), 164 (2021)
4. Gagnon, A., Miller, M.S., Hallman, S.A., Bourbeau, R., Herring, D.A., Earn, D.J., Madrenas, J.: Age-specific mortality during the 1918 influenza pandemic: Unravelling the mystery of high young adult mortality. PLOS ONE **8**(8), 1–9 (08 2013)
5. Libin, P.J., Moonens, A., Verstraeten, T., Perez-Sanjines, F., Hens, N., Lemey, P., Nowé, A.: Deep reinforcement learning for large-scale epidemic control. In: Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track: European Conference, ECML PKDD 2020, Ghent, Belgium, September 14–18, 2020, Proceedings, Part V. pp. 155–170. Springer (2021)
6. Miranda, M.N.S., Pingarilho, M., Pimentel, V., Torneri, A., Seabra, S.G., Libin, P.J.K., Abecasis, A.B.: A tale of three recent pandemics: Influenza, hiv and sars-cov-2. Frontiers in Microbiology **Volume 13 - 2022** (2022)
7. Reymond, M., Hayes, C.F., Willem, L., Rădulescu, R., Abrams, S., Roijers, D.M., Howley, E., Mannion, P., Hens, N., Nowé, A., Libin, P.: Exploring the pareto front of multi-objective covid-19 mitigation policies using reinforcement learning. Expert Systems with Applications **249**, 123686 (2024)
8. Witteveen, A.B., Young, S.Y., Cuijpers, P., Ayuso-Mateos, J.L., Barbui, C., Bertolini, F., Cabello, M., Cadorin, C., Downes, N., Franzoi, D., et al.: Covid-19 and common mental health symptoms in the early phase of the pandemic: An umbrella review of the evidence. PLoS medicine **20**(4), e1004206 (2023)