

Towards Recommender System Supported Contact Tracing for Cost-Efficient and Risk Aware Infection Suppression*

Vladimir Marbukh^{1,*}

¹ National Institute of Standards & Technology, Information Technology Laboratory, 100 Bureau Dr., Gaithersburg, Maryland, USA

Abstract

In public health, contact tracing is the process of identifying people who may have been exposed to an infected person. Contact tracing performance criteria, which include infection suppression, protection of high-risk individuals, and cost-efficiency, are not necessarily aligned with each other. Pareto optimization of the corresponding inherent trade-offs, especially at the early stages of infection, is typically unrealistic due to insufficient information on infection propagation, risk factors, prevention and treatment options, etc. We suggest that contact tracing performance can be significantly improved with the support of a specialized Recommender System (RS). Based on the combination of up-to-date contact tracing and medical data, RS can identify and test through Exposure Notification System (ENS) not only high-risk individuals but also potential superspreaders to suppress infection propagation. Due to incomplete information, the dynamic nature of the problem, and a large state and action spaces, the RS should be supported by Deep Reinforcement Learning (DRL) for solving the corresponding Partially Observable Markov Decision Process (POMDP).

Keywords

contact tracing, exposure notifications, recommender system, deep reinforcement learning, partially observable Markov decision process

1. Introduction & Motivation

In public health, contact tracing is the process of identifying people who may have been exposed to an infected person, subsequent testing them for infection, and isolating or treating the infected [1]. Contact tracing performance criteria include infection suppression, protection of high-risk individuals, and cost-efficiency. These criteria are not necessarily aligned with each other, e.g., given testing capacity, infection suppression requires high priority testing for the potential super spreaders, while protection of high-risk individuals requires testing them with higher priority. Given the testing priorities, the existing Google/Apple Exposure Notification (GAEN) technology [2] can support an Exposure Notification System (ENS) by allowing public health authorities to quickly notify people for subsequent testing. GAEN is a framework and protocol specification developed by Apple Inc. and Google to facilitate digital contact tracing during the COVID-19 pandemic to augment more traditional contact tracing techniques using Android or iOS smartphones.

Extensive research on COVID-19 has revealed that while risk-aware, multi-criteria optimization of contact tracing has significant potential, realization of this potential requires deep knowledge of the infection propagation mechanisms, medical prognoses and treatment options for infected individuals with different risk profiles [3]. Even though COVID-19 originated almost five years ago, such knowledge is still lacking [4, 5], which suggests that a contact tracing system should have the ability to collect and make sense of all up-to-date available information on infection. This can be achieved with the

HealthRecSys'24: The 6th Workshop on Health Recommender Systems co-located with ACM RecSys 2024

* Official contribution of the National Institute of Standards and Technology; not subject to copyright in the United States. Certain equipment, instruments, software, or materials are identified in this paper in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement of any product or service by NIST, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

✉ marbukh@nist.gov (V. Marbukh)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

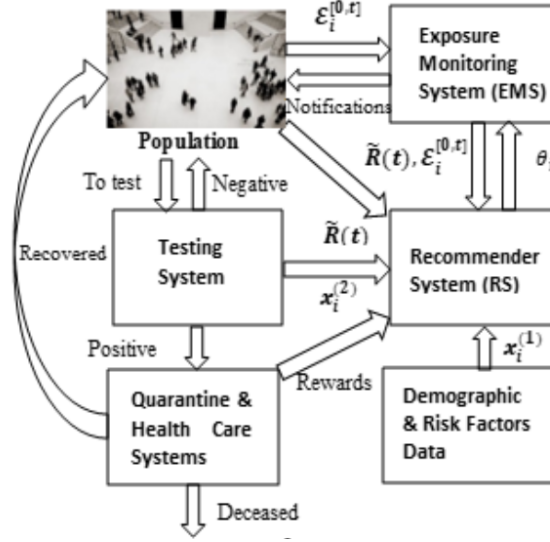


Figure 1: Recommender System supported Contact Tracing.

support of a specialized Recommender System (RS). Given testing capacity, the RS should utilize the up-to-date contact tracing, medical, and all other available relevant data to identify and through Exposure Notification System (ENS) notify individuals to be tested [2, 6]. Due to incomplete information, the dynamic nature of the problem, and a large state and action space the RS should be supported by Deep Reinforcement Learning (DRL) for solving the corresponding Partially Observable Markov Decision Process (POMDP) [7, 8]. POMDP describes the evolution of health status of each participating individual, where infectious status may not be observable and testing decisions are constrained by available testing capacity. Since a positive test result for some individual may reveal increased accumulated exposure for other individuals due to proximity to the newly discovered infection spreaders, the problem cannot be decoupled. These interdependencies significantly complicate the problem. The paper is organized as follows. Section 2 outlines operations and flow of information in the proposed RS supported contact tracing, and section 3 provides some technical details on accumulated exposure evaluation.

2. RS Supported Contact Tracing

Figure 2 presents a highly aggregated view of a Recommender System supported Contact Tracing System.

The Exposure Monitoring System (EMS) monitors “accumulated exposure to infection” for each participating individual i , $\mathcal{E}_i^{[0,t]}$ (defined in the next section) in near real time t , and feeds this information into the RS. RS also gets available information on demographic and risk factors of participating individuals $x_i^{(1)}$ as well as health status of both participating and not participating individuals who went through Health Care System $x_i^{(2)}$. Note that participating individuals are likely to consent to revealing their health information since they would benefit from accounting for their risk factors, e.g., advanced age, preexisting conditions, etc. For not participating individuals, some relevant information, which does not require revealing individual identity, can be obtained without violating their privacy.

RS is also fed the estimate of the infection reproduction number $R(t)$, i.e., the average number of new infections produced by one infected individual during his/her lifetime: $\tilde{R}(t) \approx R(t)$. Estimate $\tilde{R}(t)$ may combine information from EMS, the Health Care System, and possibly from other tracing mechanisms not shown in Figure 1, e.g., from manual tracing. Infection suppression requires keeping the infection reproduction number less than one: $R(t) < 1$. Due to numerous uncertainties in the $R(t)$ estimation: $\tilde{R}(t) \approx R(t)$, the infection suppression condition is $\tilde{R}(t) \leq 1 - \epsilon$, where “safety margin” $\epsilon < 1$ depends on the confidence level of the corresponding estimate. The reward of the RS supported Contact Tracing

is quantified by the negative loss $-L(t)$, where $L(t) = L_{ec}(t) + L_{sc}(t)$. Here economic loss due to lost productivity and cost of testing/treatment is $L_{ec}(t)$, and “social cost” quantifying suffering and, most importantly, deaths due to the infection $L_{sc}(t)$. The cost estimates are provided to RS by the Health Care System and Agencies collecting and processing economic data.

System evolution is described by POMDP $\delta(t) = (\delta_i(t))$, where component $\delta_i(t)$ describes the health status of participating individual i , i.e., “non-infected,” “infected,” “deceased.” “Non-infected” and “infected” states may not be observable which makes process $\delta(t)$ partially observable. The decision to test a participating individual reveals his/her infected or not-infected status at a certain cost due to limited testing capacity. RL employs DRL to make testing decisions on the basis of $\mathcal{E}_i^{[0,t]}$, $x_i^{(1)}$, $x_i^{(2)}$. Constraints on the infection reproduction number can be incorporated through penalty function $h(\tilde{R}(t))$ which is flat for $\tilde{R}(t) \leq 1 - \varepsilon$ and sharply increases for $\tilde{R}(t) > 1 - \varepsilon$.

Our conjecture is that (near) optimal notification strategy is threshold-based: individual i should be notified at the first moment $t = \theta_i > 0$ when this individual’s accumulated exposure reaches threshold $\hat{\mathcal{E}}_i$:

$$\theta_i = \inf_{t \geq 0} \{t : \mathcal{E}_i^{[0,t]} \geq \hat{\mathcal{E}}_i\}, \quad (1)$$

where threshold $\hat{\mathcal{E}}_i = \Delta(\mathcal{E}, x)$ depends on the history of former testing decisions/results combined with medical and demographic data of these individuals. The function $\Delta(\mathcal{E}, x)$ can be evaluated by employing a Deep Supervised Learning (DSL) algorithm. Note that in practice, notification strategy may operate on the basis of a small number of risk groups [3], which may be defined and then redefined by an on-line clustering algorithm. Assumptions of homogeneity and large number of individuals within each group, simplifies optimization of group-specific thresholds in (1).

3. Accumulated Exposure

For each participating individual i , the contact tracing system identifies “accumulated exposure” to another participating individual j during time interval $[0, t]$ as follows:

$$\mathcal{E}_{ij}^{[0,t]} = \int_0^t \phi[(\tilde{d}/d_{ij}(\tau))^\alpha] d\tau, \quad (2)$$

where the corresponding instantaneous exposure rate $\phi(z)$ is an increasing function of $z > 0$, the distance between individuals i and j at moment τ is $d_{ij}(\tau)$, and $\tilde{d} > 0$, $\alpha \geq 1$ are some parameters. Individual i accumulated exposure to infection during time interval $[0, t]$ is defined as the aggregated exposure to all known spreaders during this time interval:

$$\mathcal{E}_i^{[0,t]} = \sum_j \int_0^t \pi_j(\tau) \phi[(\tilde{d}/d_{ij}(\tau))^\alpha] d\tau, \quad (3)$$

where $\pi_j(\tau) = 1$ if individual j is infected at moment τ and $\pi_j(\tau) = 0$ otherwise.

Consider some examples. As currently defined by the CDC [1], a high-risk COVID-19 exposure is a contact with a person who tests/tested positive for SARS-CoV-2 which takes place at a distance of less than two meters for a total of 15 minutes or more over a 24-hour period. In this case, $\tilde{d} = 2$ m, $\alpha \rightarrow \infty$, $\phi(x) \equiv \min(x, 1)$, and thus an individual is assumed exposed if $\mathcal{E}^{[0,T]} = \int_0^{24} \mathbb{1}(d(t) - 2 \text{ m}) dt > 15$ min, where $\mathbb{1}(x) = 0$ if $x \leq 0$ and $\mathbb{1}(x) = 1$ otherwise. In another example [9], $\tilde{d} = 2$ m, $\phi(x) \equiv x$, and thus an individual is assumed exposed if $\mathcal{E}^{[0,T]} = \int_0^{24} (2/d(t))^\alpha dt > 15$ min.

Finally note that available information on accumulated exposure to specific individuals can be used to identify “infection superspreaders” who otherwise could be unidentified, e.g., due to being asymptomatic or for any other reason. This can be done with known algorithms [10] on undirected exposure graph G where nodes i and j are connected if $\mathcal{E}_{ij}^{[0,t]} \geq \tilde{\mathcal{E}}^{[0,t]}$, and $\tilde{\mathcal{E}}^{[0,t]} > 0$ is a properly defined threshold.

References

- [1] CDC, Contact tracing, 2024. URL: <https://www.cdc.gov/museum/pdf/cdm-pha-stem-lesson-contact-tracing-lesson.pdf>.
- [2] C. Nebeker, D. Kareem, A. Yong, R. Kunowski, M. Malekinejad, E. Aronoff-Spencer, Digital exposure notification tools: A global landscape analysis, *PLOS Digital Health* 2 (2023) 1–16. doi:10.1371/journal.pdig.0000287.
- [3] D. Acemoglu, V. Chernozhukov, I. Werning, M. D. Whinston, Optimal Targeted Lockdowns in a Multigroup SIR Model, *American Economic Review: Insights* 3 (2021) 487–502. doi:10.1257/aeri.20200590.
- [4] C. F. Manski, F. Molinari, Estimating the COVID-19 infection rate: Anatomy of an inference problem, *Journal of Econometrics* 220 (2021) 181–192. doi:10.1016/j.jeconom.2020.04.041.
- [5] W. Rauch, H. Schenk, N. Rauch, M. Harders, H. Oberacher, H. Insam, R. Markt, N. Kreuzinger, Estimating actual SARS-CoV-2 infections from secondary data, *Scientific Reports* 14 (2024) 6732. doi:10.1038/s41598-024-57238-0.
- [6] X. Chen, L. Yao, J. McAuley, G. Zhou, X. Wang, Deep reinforcement learning in recommender systems: A survey and new perspectives, *Knowledge-Based Systems* 264 (2023) 110335. doi:10.1016/j.knosys.2023.110335.
- [7] X. Xiang, S. Foo, Recent advances in deep reinforcement learning applications for solving partially observable markov decision processes (POMDP) problems: Part 1—fundamentals and applications in games, robotics and natural language processing, *Machine Learning and Knowledge Extraction* 3 (2021) 554–581. doi:10.3390/make3030029.
- [8] X. Xiang, S. Foo, H. Zang, Recent advances in deep reinforcement learning applications for solving partially observable markov decision processes (POMDP) problems part 2—applications in transportation, industries, communications and networking and more topics, *Machine Learning and Knowledge Extraction* 3 (2021) 863–878. doi:10.3390/make3040043.
- [9] K. Sayrafian, B. Cloteaux, V. Marbukh, Impact of using soft exposure thresholds in automatic contact tracing, in: *2022 IEEE International Conference on E-health Networking, Application & Services (HealthCom), 2022*, pp. 1–6. doi:10.1109/HealthCom54947.2022.9982790.
- [10] D. Zhang, Y. Wang, Z. Zhang, Identifying and quantifying potential super-spreaders in social networks, *Scientific Reports* 9 (2019) 14811. doi:10.1038/s41598-019-51153-5.