CSE 8803 EPI: Data Science for Epidemiology, Fall 2023

Lecturer: B. Aditya Prakash
Scribe: Shawn Wahi, Jasmyn Pellebon

Oct 03, 2023
Lecture 12 : Outbreak Detection (III)

1 Lecture Summary

In this lecture, we examine various facets of the problem of outbreak detection. Outbreak detection is an important problem in epidemiology because the earlier one can identify outbreaks, the quicker one can move to prevent further spread of infections. Recent technological advancements in outbreak detection involve many interlocking parts, which include (but are not limited to) individual participation in epidemic surveillance via smart phone usage, statistical algorithms used to identify potential outbreaks, and machine learning models that predict ways of monitoring for outbreaks. The summary of this lecture will examine these recent advancements in more detail.

The following is a brief overview of the sections. In Section 2, we present an algorithm for identifying regions within a geographic boundary that are likely to contain a subset of a population that has recently experienced an outbreak. Section 3 presents a reinforcement learning algorithm that was deployed by a country during the Covid-19 pandemic to determine how best to allocate scarce testing resources. Importantly, this algorithm was able to utilize sub-population statistics to outperform approaches based on country-level statistics. Finally, in Section 4, we examine the efficacy of using smartphones to monitor epidemic trends, including the pros and cons of such approaches and the current research being done in the area.

2 Using RL for Outbreak Detection

Reinforcement learning is a machine learning technique that trains an agent based on rewarding desired behaviors and punishing undesired behaviors with different scores. Importantly, reinforcement learning algorithms can be deployed in situations that require lifelong learning, meaning that the agent's model must adapt as it is presented with new environments and information. In this section, we will observe the deployment of a reinforcement learning algorithm that was applied to the problem of allocating limited Covid-19 testing resources over the course of the pandemic.

2.1 Outbreak Detection on Border

The spread of Covid-19 was, of course, not limited to national boundaries. Each country had an interest in monitoring the state of Covid-19 infections at the borders in order to minimize the impact of the disease spreading to within the borders. These countries applied border control to prevent the importation of the disease.

Border control workers had access to real-time outbreak statistics globally and corresponding demographic features of the outbreaks, specifically different countries' data. As a first pass, one could imagine combining all of this data to construct an algorithm that determines the ideal behavior of each countries' border control to mitigate the spread of Covid-19. However, this approach comes with two key drawbacks. First, the reported Covid-19 data from each country can be unreliable for a number of reasons such as incomplete testing and deliberate misreporting. Second, there is a difference between the Covid-19 related statistics of an entire population and the people within the population who are likely to travel to a new one.

2.2 Eva in Greece

In 2019, Greece decided to deploy a reinforcement learning system called "Eva" to determine how to test populations of travelers [1]. It was needed because of the impossibility of testing every traveler given the scarcity of testing resources, which furthermore would have caused harm to the tourism industries. The general idea was to test only the travelers that were deemed the most high-risk of having an infection in order to maximize the amount of detected infected travelers. This also necessitated data collection at a scale to determine "high-risk" travelers. The Eva system was deployed successfully and communicated in a subsequent Nature article.

2.3 Eva Method

Broadly, the Eva method was composed of the following steps. The first was to classify travelers using a Passenger Locator Form where travelers input information such as flight details and companions. The second was to use Bayesian approaches to infer the risk of different travelers given the types determined by the form details. Each step's accuracy is dependent on the need to collect data about enough of the different kinds of travelers in order to classify them (and subsequently their risk) accurately. This is where the RL comes in!

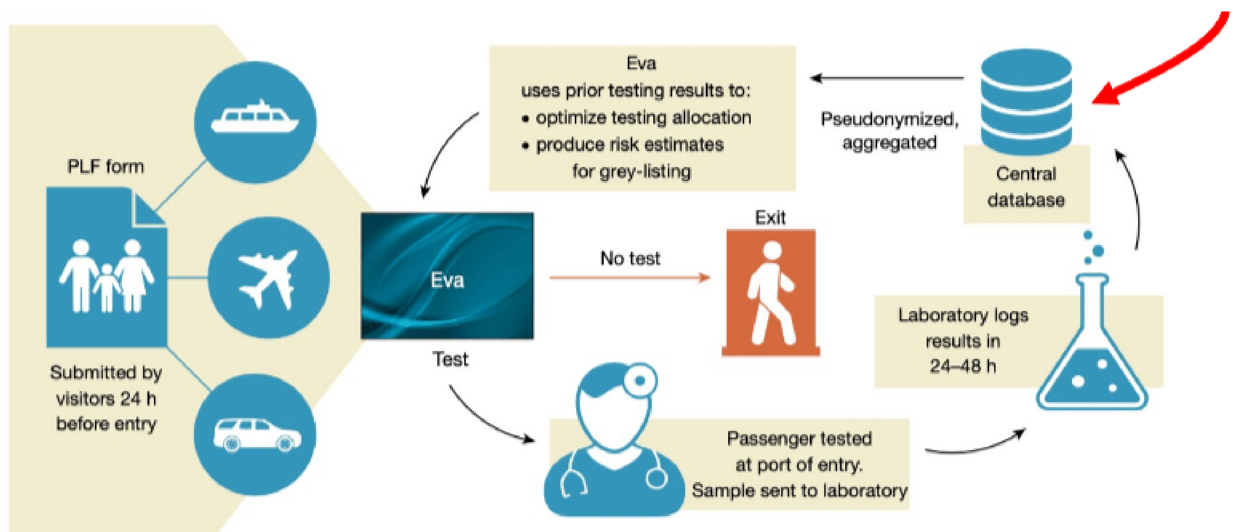

Figure 1: The Eva Method

2.4 RL in the Data Collection Loop

In order to prevent the spread of Covid-19, the highest-risk travelers needed to be detected. In order to detect the highest-risk travelers, the right kind of data needed to be collected about different traveler types. The Eva RL agent was designed to find a trade-off between "exploitation" and "exploration." The number of infectious passengers detected had to be maximized ("exploitation") in order to detect enough cases to identify risk levels of different traveler types. At the same time, multiple types of travelers had to be tested in order to guarantee that estimates of the risk-level of a particular type was accurate ("exploration"). The second target allowed the agent to detect new outbreaks in traveler types that were originally deemed low-risk.

2.5 Multi-Armed Bandit

A multi-armed bandit is a reinforcement learning agent that is trained to maximize the total reward that it receives in a set number of time steps. It can be trained to both "explore" and "exploit." The figure below is an example of a multi-armed bandit problem.

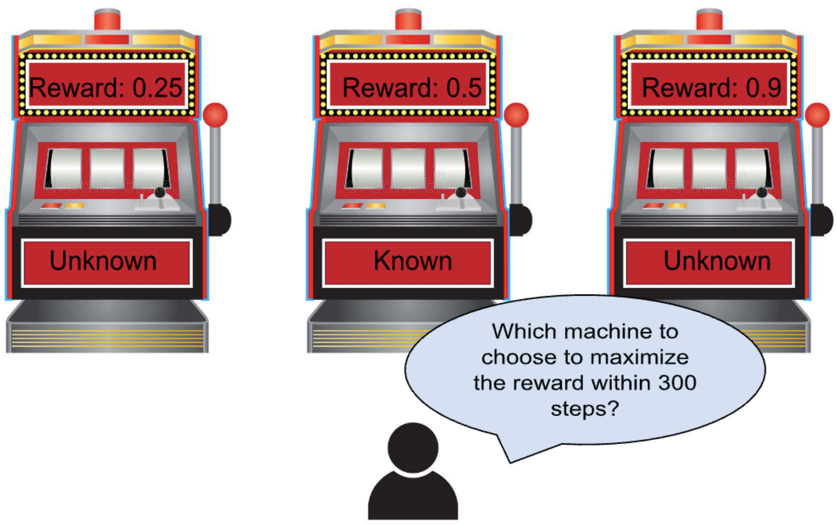

Figure 2: Multi-Armed Bandit Problem Setup

In the problem above, the agent faces the dilemma of exploring or exploiting. An strategy for exploiting would entail picking the machine the highest probability of a known reward many times over (Machine 2). An exploration-based strategy would entail picking each machine an equal amount of times and iterating over many trials to get an estimate of an "average" reward, which could then be compared against exploitation-based strategies.

In the stochastic multi-armed bandit model, a learner has several options, or "arms", each with its own independent reward distribution. The learner can choose one arm at a time and receive a reward, which may vary due to noise. Over a set number of rounds, the learner must balance between exploring all arms to gather reward information and exploiting the arm that currently seems to offer the highest reward based on estimates. In essence, the goal of learning in this model is to navigate this exploration-exploitation trade-off and maximize the total reward over the given rounds.

2.6 Gittins Index

We need to decide which "arm" (aka machine state from the related figure) to select at a time step. Gittins Index is a formula for assigning an expected reward value to each potential selected state. This can be represented by the following:

Let:

- $r(\cdot)$: reward function.
- $\pi(\cdot)$: policy or choice.
- $T(\tau|x_i)$: expected time from now to τ given x_i.

Then the simplified expression is:

$$G_i(x_i) = \sup_{\tau} \frac{E[r(\pi(x_i \text{ until } \tau))|x_i]}{E[T(\tau|x_i)]} \tag{1}$$

The numerator represents the expected reward of choosing an arm from the current time step to a future one τ. The denominator indicates the the expected amount of time from the current state to τ. Therefore $G_i(x_i)$ is therefore the reward associated with choosing x_i at the current time step, and we use maximum of these values to select the arm.

2.7 Applying Multi-Armed Bandit to Eva

Given what we know about setting up an MAB problem, we can define the Eva goal function as follows. First, we need to represent the number of travelers arriving at a certain entry point e from a specific country c on a specific day t. Let's use $A_{ce}(t)$ to represent this quantity. We also need to represent the number of travelers in this quantity who have been tested. Call this value $T_{ce}(t)$ and its maximum (the test capacity at e) $B_e(t)$. Given an unknown infected rate $R_c(t)$, we have the resulting goal of the problem:

$$\min E\left[\sum_{\tau}\sum_{c}\sum_{e} T_{ce}R_c(t)\right] \tag{2}$$

By the problem setup, this is subject to the constraints $T_{ce}(t) < A_{ce}(t)$ and $\sum_c T_{ce} < B_e(t)$.

Now we have the necessary quantities to define the high-risk assessment and data collection problem in terms of a reinforcement learning algorithm. For Eva, the machines are the travelers that arrive at an entry point on day t from a certain country, $A_{ce}(t)$. Our "reward" probability becomes the infected rate $R_c(t)$. Our Gittins Index therefore represents the risk score of different traveler populations. Now we can optimize the expected number of infected travelers tested by Eva.

2.8 High-Level Algorithm

Putting all of the above together, the first step is to test the type of travelers with the highest Gittins Index. Then, after observing the test result, the infected rate is updated for that population and a new Gittins Index is calculated for that type. This is repeated across time and populations.

2.9 Results of Eva

Recall that the overall problem was allocating limited testing resources to the riskiest populations. Eva was shown to detect more Covid-19 infections than random testing, which is indicated by the chart below.

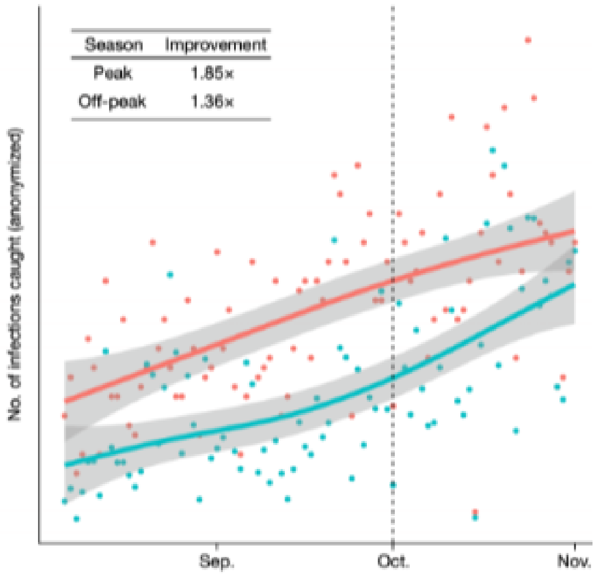

Figure 3: Eva vs Random Testing. Eva shows improvement of 1.85% during peak season, 1.36% during off-peak season.

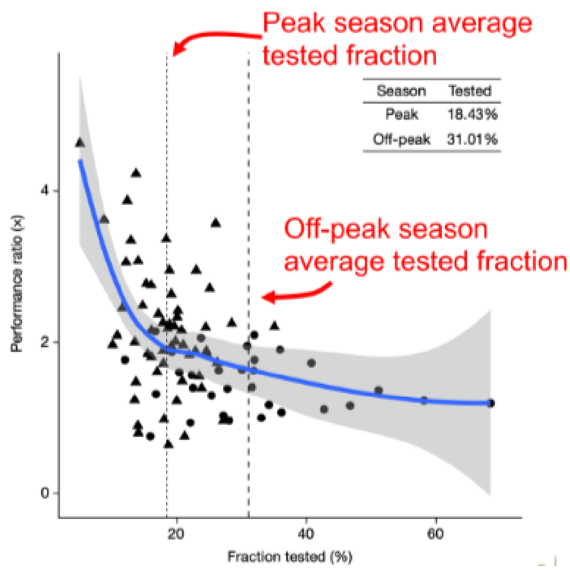

Figure 4: Eva's Efficiency. Eva averages tested fraction of 18.43% during peak season, which is 12.58% lower than off-peak season's tested fraction of 31.01%.

Eva is compared against random testing because random testing was the original method used by Greece to monitor Covid-19. Note that Eva performed better during the travel season, when there were more incoming travelers and fewer testing resources!

Ideally, Eva performs well even when testing relatively few subjects. The below figure indicates this: Eva is more efficient when the fraction tested is low. Another key metric for quantifying Eva's performance is comparing it against methods that use country-level Covid-19 metrics. The figure below indicates that Eva detects more Covid-19 infections than selecting populations based on the country-level statistics of incoming travelers such as cases and deaths. Eva outperforms these methods because it accounts for the unreliability of raw data from incoming countries and utilizes finer-grained populations (travelers vs. entire populations).

It might be counter-intuitive that country-level statistics do not result in better forecasting of high-risk groups. However, the group of practitioners behind Eva empirically demonstrate that this is in fact the case. The figure below shows that using country-level metrics as selection criteria is not more effective than simple random testing. The black line indicates randomly choosing high-risk countries. The different country-level selection criteria are indicated by the other lines. Note there is not a large difference in performance between the two!

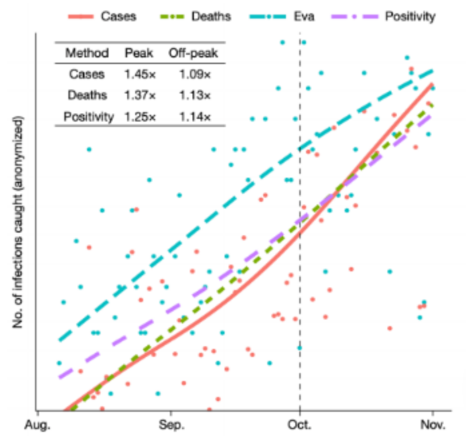

Figure 5: Eva vs. Country-Level Selection Statistics. Eva's methodology during peak season shows 1.45x increase in cases, 1.37x rise in deaths, 1.25x amplification in positivity rate compared to off-peak values.

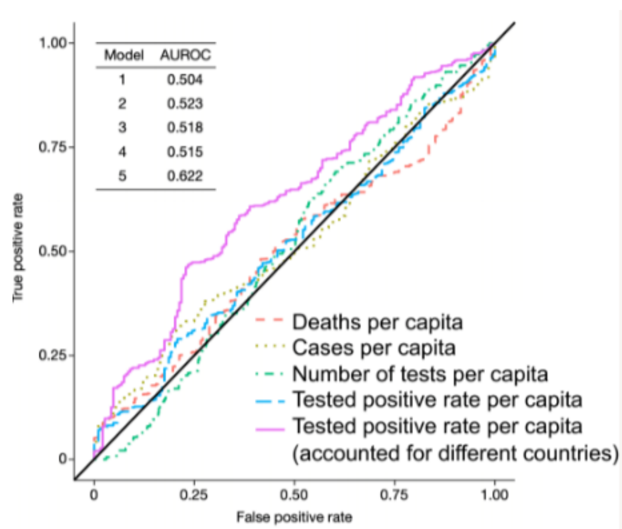

Figure 6: Comparison of Selection Methods. AUROC is an assessment of five models against metrics including deaths, cases, and tests per capita, alongside variations of tested positive rates. Model 5 demonstrates the highest AUROC of 0.622.

2.10 Lessons Learned from Eva

Measuring quantities related to health data and inter-country travel comes with many challenges, and utilizing Eva was no exception. One challenge is minimizing the request for data from individuals in order to effectively implement a policy that benefits a population as a whole. It can be ethically perilous to gather data without fully informing subjects, and everyone has a right to the privacy of their health data. Note that Eva's efficiency takes steps towards addressing these concerns. In short, it can do more with less.

One difficulty that arises with using algorithms to inform policy is the interpretability of their outputs. Any result and its metrics should be explainable in order to convince policymakers of their utility and fairness. Eva's clearly-defined reward function over many

states and repeatable performance addresses some of these concerns.
Another difficulty of deploying models of population behavior is adapting models to account for new developments in the population. Note that by modeling the situation as a multi-armed bandit problem, Eva could utilize the "explore" feature of the reward to account for new developments in the transmission rate of different populations.

2.11 Summary

Recall that the topic of the lecture is monitor outbreaks. Frequently, as in the case of Greece and Eva, we have limited resources to detect outbreaks. If we focus solely on high-risk populations, we may miss developments in other populations that make them high-risk. However, we must also avoid being too general in the surveillance of outbreaks because testing resources will not be effectively allocated where they are needed most. So we must design algorithms to account for these challenges, and reinforcement learning techniques, which are flexible in incorporating new knowledge of environments, are useful in such situations.

3 Usage of Smartphones for Covid-19 Outbreaks

Traditional surveillance of outbreaks comes with many drawbacks when compared with the surveillance capabilities offered by smartphones. In this section, we will learn about the drawbacks and benefits associated with using smartphones to survey outbreaks.

3.1 Traditional Surveillance vs Smartphone Apps

Traditionally, surveillance is conducted via clinical diagnoses of the disease in question. This is a very slow process (on the order of weeks to get data and deploy a strategy to different populations), and the slowness puts more people at risk as infected people mix with the rest of the population. Reporting results via smartphone, on the other hand, is almost real-time [2]. Furthermore, the location data they provide helps individuals make decisions according the immediate environment around them. This is of particular benefit for surveillance in regions that do not have adequate health care facilities to measure and prescribe different actions to take.

3.2 Participatory Surveillance

Participatory surveillance involves populations of individuals self-reporting their health data relevant to the epidemic of interest. This is facilitated by prevalence of smartphones, a method of surveillance that can match traditional methods but actually detect outbreaks earlier given the lower latency between positive case, reporting, data gathering, and response. There are several research groups associated with different universities developing surveillance techniques like smartphone apps that allow clinical populations to report things like vaccine efficacy and symptoms. However, one must keep in mind that this approach likely leads to biased data as sick patients tend to participate more.

3.3 Population Level Tracking

One can utilize data from the public internet to monitor outbreaks. For example, social media websites often display trends that can be utilized to predict outbreak behavior. Official websites are also useful in real-time monitoring because they can be updated quickly. However, when utilizing these indirect methods of outbreak surveillance, one runs the risk not understanding the outbreak as quickly. Because inferring outbreak statistics is harder in this regime, the disease may have spread for a while before an effective model can be created.

Beyond tracking populations, an important component of containing outbreaks is communication with at-risk populations. Social media and official websites further help with this due to the real-time communication that they allow and the localized information that can be conveyed.

3.4 Individual Risk Assessment

Individuals can utilize data gathered across populations characterized by more detail to make decisions about how to deal with an outbreak. Surveillance techniques that smartphones allow can help individuals identify risks among populations that resemble themselves demographically. Using this information, there can be a greater use of non-pharmaceutical preventative behaviors such as mask-wearing, volunteering for more testing, and social distancing.

3.5 Future Prospects

The ability to survey outbreaks via smartphones offers exciting prospects for individuals in the future to make decisions about managing their health during epidemics. Smartphones open up the potential to inform individuals about their specific infection risk during an epidemic. Individuals can also monitor infections in real-time in their communities. Models can be developed at the level of these communities in order forecast future risk for an individual.

References

[1] H. Bastani et al. Efficient and targeted covid-19 border testing via reinforcement learning. *Nature*, 599:108–113, 2021.

[2] J. Pandit et al. Smartphone apps in the covid-19 pandemic. *Nature Biotechnology*, 40:1013–1022, 2022.