POLITECNICO DI TORINO
Repository ISTITUZIONALE

Adaptive COVID-19 Screening of a Subpopulation

Original
Adaptive COVID-19 Screening of a Subpopulation / Di Stefano, F., Gasparini, M.. - 406:(2022), pp. 105-116. (SIS2021) [10.1007/978-3-031-16609-9_8].

Availability:
This version is available at: 11583/2976619 since: 2023-03-15T09:52:17Z

Publisher:
Springer

Published
DOI:10.1007/978-3-031-16609-9_8

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

04 November 2026

Adaptive COVID-19 screening of a subpopulation

Fulvio Di Stefano and Mauro Gasparini

Abstract Methods are sought to test adaptively whether a subpopulation proportion follows the same time evolution as the population proportion. The motivating case study is the COVID-19 screening in a university community, taking into account the time evolution of the pandemic in the whole country.

Key words: Outbreak detection, Control Charts, dynamic threshold, SARS-CoV-2

1 Introduction

During the ongoing pandemic era, the need arises among members of certain working, studying or social commmunities to undergo a reinforced screening to immediately identify outbreaks within the community, as outlined for example in [1]. In our case study, a screening process is planned within Politecnico di Torino (POLITO), a public university, to prevent clusters among students who come in person to attend classes (most classes are also given in an online mode) during the first semester of academic year 2021-2022, running from 27 September 2021 to 14 January 2022. Based on the case study, the aim of this work is to compare methods to test, repeatedly and dynamically, whether a sub-population of interest is on average similar to the general population with respect to a certain binary characteristic (e.g. infected/not infected) with a time-changing distribution.

Outbreak detection is widely treated in literature and many studies have been conducted on the topic [2, 3]. Tukey's fences [4] and other well known static meth-

Fulvio Di Stefano
Dipartimento di Scienze Matematiche "Giuseppe Luigi Lagrange", Politecnico di Torino e-mail: fulvio.distefano@polito.it

Mauro Gasparini
Dipartimento di Scienze Matematiche "Giuseppe Luigi Lagrange", Politecnico di Torino e-mail: mauro.gasparini@polito.it

ods for outlier detection [5] are aimed at identifying excessive presence of some characteristic in random samples. Attribute control charts (see for example [6]) apply the same ideas to time series data. This work extends those models by introducing forecasting techniques to obtain a general methodology which combines the forecasting of the characteristic with the detection of the excess of that characteristic in a subpopulation. Thus, the main advantage of control charts, which perform well in detecting sudden and large deviation of the characteristic of interest [3], is combined with a forecasting technique which takes into account the variability of the data. The methodology is then applied to the POLITO case study to obtain an adaptively varying threshold for the COVID-19 screening in POLITO, taking into account the time evolution of the pandemic in the whole country, i.e. an alert threshold which is not fixed but is able to adapt to the predicted future evolution of the pandemic.

2 Methods

2.1 Modeling time evolving proportions

Let $P_t \ \forall t \geq 0$ denote the proportion of individuals in a general population of approximately constant size N_P who carry a characteristic of interest. Suppose also that P_t is an unknown stochastic process over time, meaning that the number of individuals who carry the characteristic is a random variable with a time-changing distribution. Let $N_S \leq N_P$ be the size of a well-defined subpopulation of interest. Let $p_t \ \forall t \geq 0$, be proportion of individuals who carry the characteristic of interest in the subpopulation: p_t is a distinct stochastic process over time t. If the subpopulation is conformal to the general population it is a subset of, i.e. the subpopulation and the general population are homogeneous at all possible scales of observations, p_t should be approximately equal to P_t. However, the characteristic of interest may evolve differently in the subpopulation with respect to the general population. At a given time t_0, estimates of P_t for some previous time steps $t \leq t_0$ are available. These estimates, denoted as $\hat{P}_t \ \forall t \leq t_0$, are based on samples of varying sizes, but this extra source of uncertainty is ignored in this work, for reasons explained in detail later on. The following methodology aims at predicting P_{t_0+1}, having obtained estimates of $P_t, t \leq t_0$, to obtain a statistical test of whether the subpopulation proportion p_{t_0+1} is significantly larger than P_{t_0+1}. The interest in one-sided upper tests is due to the fact the we are concerned with the possibility of an excessive proportion p_t with respect to P_t, as exemplified by the POLITO case study, where evidence for an excessive proportion in the subpopulation would cause the reinforcing of restrictive measures such as confinement or distance learning.

2.2 Forecasting using ARMA models

Given the observed time series $\hat{P}_t \ \forall t \leq t_0$, one of the objectives is to forecast the future value P_{t_0+1}. The reason is two-fold: to have a reference value which moves over time, according to the progress of the underlying characteristic of interest in the general population, and to incorporate all the information collected up to this point in the next prediction.

Well-known methods to forecast the future values of a time series when the underlying process is unknown are ARMA models (a good textbook introduction can be found in [7]), very popular in the econometric literature. ARMA models have been generalized in several ways and can easily be adapted to many time series. Far from claiming ARMA models are sufficient for all predictions, we propose instead to use them as a working tool to update a population reference value.

Each ARMA model is constructed to have two order parameters, denoted (p,q), which have to be "identified" based on data with an empirical model selection procedure, and several unknown regression parameters, which have to be estimated based on data. A generic ARMA(p,q) model can be applied to our variable of interest P_t, after a preliminary logarithmic transformation, to account for the inherent positiveness of P_t (the fact that P_t is also bounded above by 1 can be neglected, since P_t is usually a small value, way closer to 0 than to 1):

$$\log(P_t) = K + a_1\log(P_{t-1}) + ... + a_p\log(P_{t-p}) + \varepsilon_t + b_1\varepsilon_{t-1} + ... + b_q\varepsilon_{t-q}$$

where K is the underlying mean of the process, the ε_t are the error terms, a_i $i = 1,...,p$ are the coefficients associated to the auto-regressive part of the model and b_j $j = 1,...,q$ are the coefficients associated to the moving average part of the model. The ε_t are assumed to be independent and identically distributed, following a normal distribution $N(0,\sigma^2)$.

The estimation of the order and the coefficients of the model is widely treated in the literature [7, 8] and goes beyond the scope of this work. However, it can briefly be mentioned that the values of p and q can be chosen to minimise the Bayesian information criterion (BIC) [8, 9].

In our setup, the estimates $\hat{P}_t \ \forall t \leq t_0$ can be used as surrogate P_t to estimate the parameters of the model. Therefore, having determined the order (p,q) of the model and estimated the parameters $\hat{a}_i$ and $\hat{b}_j$, for $i = 1,...,p$, $j = 1,...,q$, $\hat{K}$ and $\hat{\sigma}^2$ and having obtained the realizations of the errors $\hat{\varepsilon}_t \ \forall t \leq t_0$, we can proceed with the forecasting. Given all available information up to t_0, the prediction for the next time-step can be calculated as:

$$\log(\tilde{P}_{t_0+1}) = \hat{K} + \hat{a}_1\log(\hat{P}_{t_0}) + ... + \hat{a}_p\log(\hat{P}_{t_0-p}) + \hat{b}_1\hat{\varepsilon}_{t_0} + ... + \hat{b}_q\hat{\varepsilon}_{t_0-q}$$

This quantity can be thought to be approximately normally distributed with variance σ^2, estimated by $\hat{\sigma}^2$ for practical purposes.

2.3 Detecting excessive presence of the characteristic of interest in the subpopulation

Given the forecast for the next time period in the whole population, inference can be made over the presence, and in particular over possible excessive presence, of the characteristic of interest in the subpopulation. Suppose then a random sample of n_S out of N_S individuals are tested for the presence of the characteristic in the subpopulation. Let X_t be the number of individuals who carry the characteristic of interest among the tested individuals n_S at time t; then X_t has a hypergeometric distribution with (discrete) density

$$\text{Prob}_{p_t}(X_t = x) = \frac{\binom{N_S p_t}{x}\binom{N_S(1-p_t)}{n_S-x}}{\binom{N_S}{n_S}}$$

which, for N_S large compared to n_S, can be approximated by the binomial density

$$\text{Prob}_{p_t}(X_t = x) = \binom{n_S}{x} p_t^x (1-p_t)^{n_S-x}$$

where p_t is an unknown parameter evolving over time. We would like to test formally, at level $1-\alpha$, the system of hypotheses

$$\begin{cases} H_0 : p_{t_0+1} = P_{t_0+1} \\ H_A : p_{t_0+1} > P_{t_0+1} \end{cases}$$

Following the methods described in the previous section, at time t_0 we have a forecast $\tilde{P}_{t_0+1}$ for the next period, accompanied by an estimate $\hat{\sigma}$ of its uncertainty. We can proceed in different ways.

Method 1: direct thresholding. Use the normal approximation retrieved from the ARMA model to obtain an explicit threshold τ_{1,t_0+1} and use the decision rule

> "Reject H_0 if $X_{t_0+1} > \tau_{1,t_0+1} := n_S \exp(\log(\tilde{P}_{t_0+1}) + z_{1-\alpha}\hat{\sigma})$",

where $z_{1-\alpha}$ is the $1-\alpha$ quantile of the normal distribution.

Method 2: binomial testing. A more traditional approach would be to perform a standard binomial test at a significance level approximately equal to $1-\alpha$ (due to the discreteness of the binomial distribution), with decision rule

> "Reject H_0 if $X_{t_0+1} > \tau_{2,t_0+1}$",

where τ_{2,t_0+1} is the $(1-\alpha)$-quantile of the binomial distribution with parameters n_S and $\tilde{P}_{t_0+1}$. This method disregards the uncertainty about P_{t_0+1}.

Method 3: normal testing. If n_S is large enough, the binomial distribution can be approximated by a normal distribution, so that the following decision rule is obtained:

"Reject H_0 if $X_{t_0+1} > \tau_{3,t_0+1} = n_S\tilde{P}_{t_0+1} + z_{1-\alpha}\sqrt{n_S\tilde{P}_{t_0+1}(1-\tilde{P}_{t_0+1})}$".

The idea behind these techniques is to combine the main advantage of control charts, which perform well in detecting sudden and large deviation from the average [3], with a forecasting technique which is able to take into account the evolution of the characteristic of interest over time.

2.4 A naive fixed threshold

The three methods previously presented are opposed to a naive method which consists in giving an alert if, at some α level type I error, the null hypothesis $H_0 : p_t = p_N$ at time $t \geq 0$ is rejected, using binomial testing. p_N is a pre-determined level of diffusion of the charactestic of interest in the subpopulation which is considered acceptable.

Method 4: naive fixed threshold binomial testing. Binomial testing using a fixed threshold results in the following decision rule.

"Reject H_0 if $X_{t_0+1} > \tau_N$",

where τ_N , is the $(1-\alpha)$-quantile of the binomial distribution parameters n_S and p_N.

This procedure is formally equivalent to a binary attribute control chart [6] and does not take into account the evolution of the characteristic of interest over time.

3 A case study: COVID-19 testing at Politecnico di Torino.

3.1 Current testing.

In POLITO, during the first semester of the academic year 2021-2022, an estimate of $N_S = 21870$ students plan to attend classes in person, and the university screens sample of students on site to develop a system of alert to predict possible outbreaks. This cluster detection system works alongside the screening and the prevention measures every national system is implementing. In particular, $n_S = 250$ oropharyngeal swabs are carried out every Monday, Wednesday and Friday, due to procedural and technical constrains, leading to the total of 750 swabs weekly. We will discuss the previous test using $n_S = 250$, whereas in practice the same test will be repeated three times a week. The test will be conducted on a random sample of students who have booked to attend their lessons on that particular day. Students can refuse to be tested. However, up to the 3 December 2021, the number of refusals to the tests has been of a few units. Usually, students who are symptomatic and acknowledge to be positive do not come in presence, given the possibility to follow online lessons.

This fact remarks the necessity of a screening procedure which is able to capture possible unnoticed clusters, like the ones caused by asymptomatic cases, inside the university.

Following Section 2, the probability of a positive test in a particular day $t \geq 0$ is denoted by p_t. Therefore, following Method 4, an alert will be given if, at some α level type I error, the null hypothesis $H_0 : p_t = p_N$ at time $t \geq 0$ is rejected. Using $\alpha = 0.20, n_S = 250$ and $p_N = 0.015$, a rough guess of the average pandemic situation in the country in September 2021, we obtain $\tau_N = 5$ (the reason for using such a large level of type I error is explained below). This procedure is the one currently implemented in POLITO, which the authors have helped to set up. Up to the 3 December 2021, the screening has produced very few positive tests, far below the threshold, due to particularly stringent rules for accessing the POLITO site.

The current procedure does not take into account the fact that the pandemic is evolving over time. In particular, in Italy the evolution of the pandemic is monitored by Istituto Superiore di Sanità and Protezione Civile, who provide day by day data on the evolution of the pandemic [10]. These data give a clue on how the pandemic is evolving in Italy and, if the pandemic is worsening in the whole country, it is expected that it will also worsen in POLITO, provided that the above homogeneity assumption holds. The use of national data and not regional (Piemonte) data is two-fold. Notwithstanding the fact that the current regulation still allocates Italian regions in different risk areas according to the regional spread of the pandemic, all restrictions of movements (like travelling between regions) and many social restrictions do not apply to people who have obtained the so-called "green pass certification", which is mandatory also to physically access universities. Moreover, just over a third of POLITO students residing in Italy come from the Piemonte region: a huge number of students come from the south of Italy and a big component comes also from the centre and north-east of Italy. Therefore, given the heterogeneous regional composition of the students which plan to attend class in person and the freedom guaranteed to them by the current regulation, which permits to people in possession of the green pass certification to not experience any of the regional constraints in place, national data are believed to better represent the considered subpopulation.

In Figure 1, the weekly percentage of positive tests in Italy since the beginning of the pandemic is shown. Now, for reasons related to the way these data are collected in Italy, they do not properly represent the proportion of infected people at time t. As a matter of fact, molecular tests are much more precise than antigenic tests and are carried out mainly to confirm cases reported via the latter ones, increasing the estimate of positive cases since they may be reported multiple times. However, the data on antigenic tests decrease the estimate of positive cases, since a huge number of these tests are carried out by not vaccinated people to access many working activities, in compliance with the current regulation. Moreover, the number of weekly antigenic tests is approximately twice the number of molecular tests in this phase of the pandemic, but it is not consistent over time. For these reasons, the percentage of positive tests probably overestimates the true proportion of infected people at time t, However, it is difficult to quantify how large the overestimation is, but the expected percentage of positive tests in POLITO might be lower than those numbers. Rather

than arbitrarly reduce by a significant fraction those percentages, we prefer to use the original data but work with a large type I error, following the criterion that in this situation prevention is better than cure, and a false alarm does not entail dramatic consequences. In particular, the authors will help in analysing the data coming from the screening process and, if necessary, start an alarm. Given the alarm, a series of containment measures will be taken, starting from retesting the positive people with molecular tests - more precise than the rapid tests used in POLITO - then escalating to quarantineing people and possibly switch to online teaching for entire classes. We will therefore use a large $\alpha = 0.20$ and take as our $\hat{P}_t \ \forall t \geq 0$ the values shown in Figure 1. These are calculated by dividing the total number of new cases reported in a particular week over the total number of tests (both antigen and molecular) conducted in that particular week. The rationale for using percentages instead of counts is that they are known to help in facing inconsistency in reporting of cases[2], a well-known topic in outbreak detection.

Using the dataset from the Istituto Superiore di Sanità and Protezione civile and the methodology described earlier, a forecast on the percentage of positive tests in Italy next week can be obtained. If the pandemic is worsening (or improving) in the whole country, the ARMA(p,q) model will provide an adequate prediction, which can be used for a more accurate alarm system with respect to a fixed threshold. The use of the weekly time series permits to have a quantity which can be compared to the percentage of positive tests among the students at POLITO. Moreover, the data are grouped weekly for two reasons: the number of tests during the week is not constant but depends on the day of the week, and the fact that the tests at the Politecnico are not planned daily, as for those in the whole country.

Finally, due to the highly non-stationary evolution of the pandemic, which depends on many factors such as restrictive measures, vaccines, seasonality, heterogeneous intensity of screening (e.g. differing search rates for asymptomatic individuals), it is recommendable to use only the last portion of the time series; in this work we use only the last 16 weeks of observations and disregard the previous ones dinamically in case we are interested in different t_0 times.

More accurate methods to predict the number of SARS-CoV-2 positives in a general population are available in literature: the SIS model [11], the SIPRO model [12], the SIDARTHE model [13] and its extentions [14], the Covasim model [15], and many others (for example [16, 17, 18]). However, the number of weekly tests in Italy is not constant over time and social and economical measures are continuously taken by the governments to contain the pandemic, making a precise estimation of P_t very hard. In this situation, far from proposing ARMA models as a well-thought realistic model for epidemic prevalence, the proposal in this work is to use ARMA models as a working and adaptive tool which performs well enough to give one-time step ahead predictions, and no further in time.

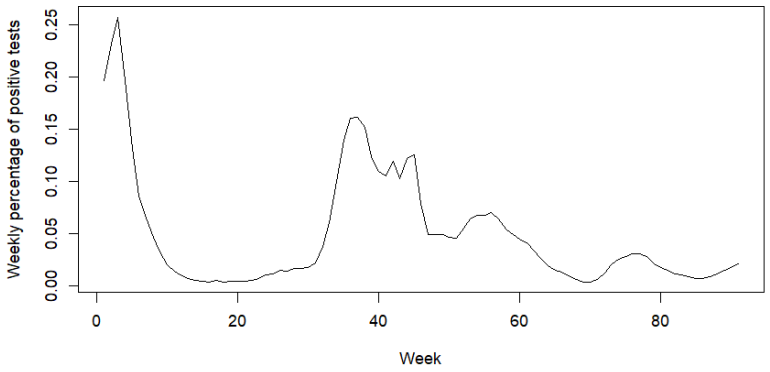


Fig. 1 Percentage of weekly positive SARS-CoV-2 tests in Italy since the beginning of the pandemic (Week 0 is 24-02-2020) up to the end of November 2021 (Week 91 is 28-11-2021).

3.2 A proposal for adaptive testing

Using R version 4.1.2 and the package forecast [19], the methodology of Section 2 can be applied to the dataset. Suppose to start the estimation at week 74 after the beginning of the pandemic, i.e. 25 July 2021. At this date, the weekly percentage of positive tests in Italy is $\hat{P}_{74} = 0.025$, and the curve is slightly increasing. Using R, it is possible to fit an ARMA(p,q) model to the dataset using the previous 16 weeks' data.

The best fitting model, according to the minimization of the BIC, is ARMA(4,0). The following parameters are estimated: $\hat{a}_1 = 2.723$, $\hat{a}_2 = -3.334$, $\hat{a}_3 = 2.235$, $\hat{a}_4 = -0.705$, $\hat{K} = -3.876$, $\hat{\sigma}^2 = 0.009$. Using this model, we can forecast $\tilde{P}_{75} = 0.026$. This number is bigger than $\hat{P}_{74}$, as expected since the pandemic is slightly worsening in this period. Also, the three thresholds described in subsection 2.3 are calculated: $\tau_{1,75} = 8$, $\tau_{2,75} = 9$,$\tau_{3,75} = 9$, where $\tau_{1,75}$ and $\tau_{3,75}$ have been rounded to the next integer. Therefore, if the number of positives on Monday, Wednesday or Friday of week 75 among the 250 swabs at the Politecnico is greater or equal than the threshold of our interest, we start an alert since we reject the null hypothesis that the proportion of positives in the university p_{75} is equal to the proportion of positives in the whole country P_{75}: there is evidence for an ongoing outbreak in POLITO.

After collecting the national data for week 75, we can proceed to estimate the threshold for the next week. And then repeat this procedure over time. Figure 2 shows the three different thresholds calculated from week 75 to week 92 since the beginning of the pandemic, using all the data available from the previous 16 weeks. Comparing this Figure with Figure 1, it is clear that the progress of the pandemic has been adequately incorporated in the model. The decrease in the weekly percentage of positive after week 77 is captured by the variable thresholds, as it is for

the uprising trend after week 86. Of the three different thresholds, τ_1 is the most conservative, resulting in the lowest number of positives to be achieved to start an alarm. On the other hand, τ_2 and τ_3 give very similar results, with τ_2 being the most conservative of the two. This is because the rationale behind the two thresholds is the same, but τ_3 is derived from a normal approximation of binomial distribution, which traces the widely used asymptotic test for proportions, while τ_2 is the exact quantile of a binomial distribution. In any case, all these methods outperform the fixed threshold τ_N, which is not able to capture any of the fluctuation of the progress of the pandemic: it can result in a lower threshold compared to the others when the pandemic is in an expansion phase, or in a higher threshold comparing to the others when the pandemic is in a regressive phase.

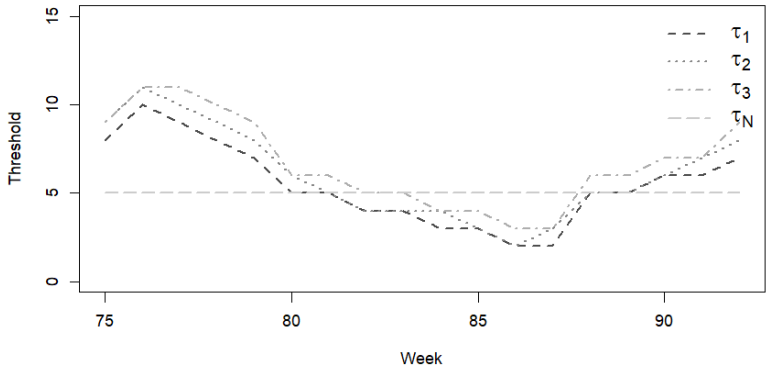


Fig. 2 The variable and fixed thresholds for the SARS-CoV-2 swabs at Politecnico di Torino.

3.3 Operating characteristics of adaptive testing

Given the available data on the progress of the pandemic $\hat{P}_t$, some operating characteristics can be calculated for the different thresholds to check their properties. As an example, it is possible to start from the previously calculated $\tau_{1,75} = 8$, $\tau_{2,75} = 9$, $\tau_{3,75} = 9$, $\tau_{N,75} = 5$ and the real percentage of positives tested in Italy in week 75 $\hat{P}_{75} = 0.028$, to retrieve power and type I error of the thresholds for this week. In particular, supposing that the positive tests at POLITO follow a binomial distribution with parameters n_S and p_{75}, the type I error can be defined as:

$$\alpha_{i,75} = 1 - \sum_{x=1}^{\tau_{i,75}} \binom{n_S}{x} (\hat{P}_{75})^x (1 - \hat{P}_{75})^{n_S - x}$$

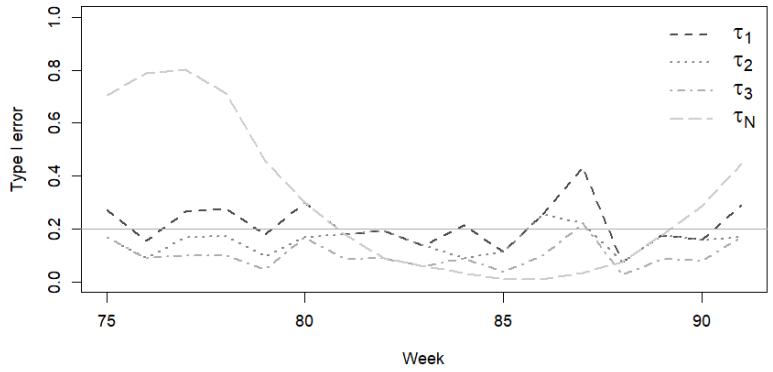


Fig. 3 Type I error for the different thresholds. The solid line is 0.2.

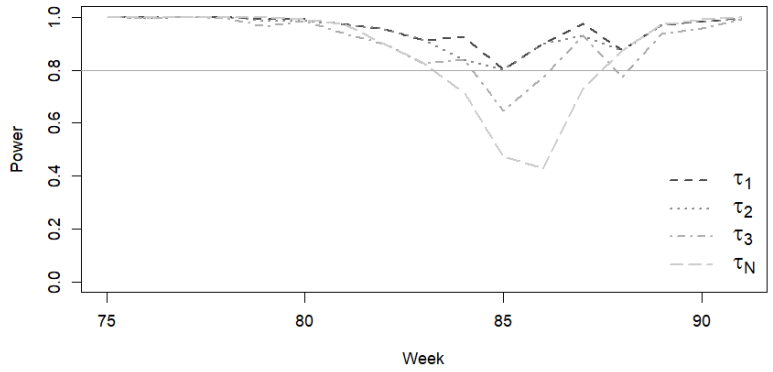


Fig. 4 Power for the different thresholds. The solid line is 0.8.

for $i = 1,2,3,N$; while the power can be defined as:

$$(1-\beta_{i,75}) = 1 - \sum_{x=1}^{\tau_{i,75}} \binom{n_S}{x} (3\cdot\hat{P}_{75})^x (1-3\cdot\hat{P}_{75})^{n_S-x}$$

for $i = 1,2,3,N$. This type I error identifies the probability to overcome the given threshold when in truth $p_{75} = \hat{P}_{75}$, resulting a false alarm. On the other hand, the power is the probability to overcome the given threshold when in truth $p_{75} = 3\cdot\hat{P}_{75}$, resulting a right alarm. For week 75 $\alpha_{1,75} = 0.271$, $\alpha_{2,75} = 0.168$, $\alpha_{3,75} = 0.168$, $\alpha_{N,75} = 0.705$, $(1-\beta_{1,75}) = 0.999$, $(1-\beta_{2,75}) = 0.998$, $(1-\beta_{3,75}) = 0.998$, $(1-\beta_{N,75}) = 1$.

In Figure 3 and Figure 4 the type I error and power for the different thresholds are presented. These are calculated using the available weekly data of $\hat{P}_t$ from week 75 to 91, as shown above. It can be seen that the fixed threshold gives an extremely high type I error but also a very high power when the pandemic is worsening, but its type I error is controlled under 0.20 and its power decreases down to 0.429 when the pandemic is in a regressive phase. This means that in a expansive phase of the pandemic there is a high risk of a false alarm, because of the more plausible high number of positives, while in a regressive phase it is realistic to not detect a possible cluster inside the university, as a consequence of the high threshold. Instead, the type I error of τ_1 is almost everywhere above 0.07 and below 0.30, with a peak of 0.431 at week 87 (when the weekly percentage of positive tests changes its convexity and starts growing again); however, its power is above 0.8 in all weeks. As regards τ_2 and τ_3, these give very similar results. Their type I error is controlled under 0.20 everywhere except around week 87, with τ_2 peaking at 0.256 at week 86 and at 0.221 at week 87 and τ_3 peaking at 0.221 at week 87. The power of τ_2 is above 0.8 all of the time, while the power of τ_3 is above 0.8 most of the times, except on week 86 and 88 where it goes down to 0.772 and 0.776, respectively.

A brief summary of the operating characteristics is shown in Table 1, where it can be seen that the variable thresholds have better operating characteristics with respect to a fixed threshold: the proposed methodology results in a lower number of false alarms and in improved detection of possible outbreaks.

Table 1 Summary of the operating characteristics of the different thresholds.

Threshold	Fixed	Type I error	Power
τ_1	No	(0.07,0.44)	(0.80,1)
τ_2	No	(0.07,0.26)	(0.80,1)
τ_3	No	(0.02,0.23)	(0.64,1)
τ_N	Yes	(0.01,0.80)	(0.42,1)

4 Discussion

In this work, a methodology is proposed to identify how conformal a subpopulation is to a general population with respect to the distribution of a binary variable. This study was motivated by a case study on the SARS-CoV-2 tests in POLITO, to identify outbreaks inside the university via the screening process organized with oropharyngeal swabs three days a week.

Making use of a very general ARMA(p,q) model, three thresholds which vary over time have been determined. These thresholds are used to test the equality of the proportion of individuals with the characteristic of interest in the subpopulation and the general population. Via the case study, it has been shown that the three presented variable thresholds are able to capture the progress of the underlying process, outperforming a fixed threshold in terms of operating characteristics. The three presented thresholds exhibit different properties: threshold τ_1 performs very well in terms of power, but type I error is not controlled at the however large $\alpha = 0.20$ level we decided to work with; thresholds τ_2 and τ_3 have controlled type I error at a $\alpha = 0.20$ level, but a little less power with respect to τ_1.

Some limitations and future extension of this work regard the possibility to use a more accurate model to predict the COVID-19 pandemic in Italy and therefore obtain more accurate thresholds for the case study. However, this work aims at being very general, in order to be adapted to various possible scenarios.

Acknowledgements The authors would like to thank Paola Lerario and Maurizio Galetto for insights and details on the organization of the screening procedure in Politecnico di Torino and Enrico Bibbona for some critical discussions. The authors would like to thank also two anonymous reviewers, who greatly helped to improve the quality of the manuscript.

References

1. European Center for Disease prevention and Control. COVID-19 clusters and outbreaks in occupational settings in the EU/EEA and the UK. (2020)
2. Buckeridge D.L., Burkom H., Campbell M., Hogan W.R., Moore A.W.: Algorithms for rapid outbreak detection: a research synthesis. Journal of Biomedical Informatics. (2005) doi: 10.1016/j.jbi.2004.11.007
3. Leclère B., Buckeridge D.L., Boëlle P.Y., Astagneau P., Lepelletier D.: Automated detection of hospital outbreaks: A systematic review of methods. PLOS ONE. (2017) doi: 10.1371/journal.pone.0176438
4. Tukey J.: Exploratory data analysis. Addison-Wesley Pub. Co. (1977)
5. Hawkins D.M.: Identification of Outliers. Springer (1980)
6. Montgomery D.C.: Introduction to Statistical Quality Control. John Wiley & Sons (2019)
7. Brockwell P.J., Davis R.A.: Time Series: Theory and Methods. Springer (2009)
8. Stoica P., Selen Y.: Model-order selection. IEEE Signal Processing Magazine (2004) doi: 10.1109/msp.2004.1311138
9. Schwarz G.: Estimating the Dimension of a Model. The Annals of Statistics (1978) doi: 10.1214/aos/1176344136
10. https://github.com/pcm-dpc/COVID-19/. Cited 20 Dec 2021.

11. Kermack W.O., McKendrick A.G.: A contribution to the mathematical theory of epidemics. Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character (1927) doi: 10.1098/rspa.1927.0118
12. Amongero M., Bibbona E., Mastrantonio G.: Analysing the Covid-19 pandemic in Italy with the SIPRO model. Book of short papers SIS 2021
13. Giordano G., Blanchini F., Bruno R., et al.: Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy. Nature Medicine (2020) doi: 10.1038/s41591-020-0883-7
14. Giordano G., Colaneri M., Filippo A.D., et al.: Modeling vaccination rollouts, SARS-CoV-2 variants and the requirement for non-pharmaceutical interventions in Italy. Nature Medicine (2021) doi: 10.1038/s41591-021-01334-5
15. Kerr C.C., Stuart R.M., Mistry D., et al.: Covasim: An agent-based model of COVID-19 dynamics and interventions. PLOS Computational Biology (2021) doi: 10.1371/journal.pcbi.1009149
16. Farcomeni A., Maruotti A., Divino F., Jona-Lasinio G., Lovison G.: An ensemble approach to short-term forecast of COVID-19 intensive care occupancy in Italian regions. Biometrical Journal (2020) doi: 10.1002/bimj.202000189
17. Fokas A.S., Dikaios N., Kastis G.A.: Mathematical models and deep learning for predicting the number of individuals reported to be infected with SARS-CoV-2. Journal of The Royal Society Interface (2020) doi: 10.1098/rsif.2020.0494
18. Kissler S.M., Tedijanto C., Goldstein E., Grad Y.H., Lipsitch M.: Projecting the transmission dynamics of SARS-CoV-2 through the postpandemic period. Science (2020) doi: 10.1126/science.abb5793
19. https://cran.r-project.org/web/packages/forecast/index.html. Cited 20 Dec 2021.