Prediction of prophylactic replacement of voice prosthesis in laryngectomized patients - A retrospective cohort study
Suppl. Monte-Carlo simulations

Anne N. Heirman MD, Vincent van der Noort PhD, Rob van Son PhD, Japke F. Petersen MD PhD, Lisette van der Molen PhD, Gyorgy B. Halmos MD PhD, Richard Dirven MD PhD, Michiel W.M. van den Brekel MD PhD

Simulations: R. van Son

Tue May 10 13:52:11 2022

R version 3.6.1 (2019-07-05) [1]

Prediction of prophylactic replacement of voice prosthesis in laryngectomized patients – A retrospective cohort study

Suppl. Monte-Carlo simulations of optimal time-to-change for voice prostheses.

Introduction

The aim of the project is to investigate the optimal time-to-change for voice prostheses that balances the number of leakages (device fails) events and the number of device changes required. The standard policy is *Wait-to-Leakage*, where the user will return to the care-center for a device change when the device starts leaking (100% leakage). In profylactic changing, the devices are changed at a scheduled time to reduce the number of leakage events. In this study, the target number of leakage events is set to 30% of devices leaking before change, DL70, meaining that 70% of device are changed before they fail. As devices are changed before they fail, the mean time of use of the devices is shorter than under the WtL policy. Therefor, the number of leakage events per year for DL70 profylactic changes will not be 30% of these events under the WtL policy, but higher. How much higher, or even higher than under the WtL policy, depends on the details of the device life-times. The data available from real patients is sparse and little is known about the statistics of the device life-times in vivo. Therefore, it is not possible to evaluate the possible device change policies and the possibilities of minimizing the number of leakage events.

Another option is to model the device life-times and change policies in a Monte-Carlo simulation where it is possible to model various device life-time distributions and the effect of different policies on the number of leakage events. To make the simulations scalable and numerically comparable between different probability distibutions, the simulations are organized according to the mean and standard deviation (SD), i.e., the Coefficient of Variation (CV) which equals $SD/Mean$, instead of the median and range. As device life-times are always positive, none of the distributions tested are Gaussian (Normal).

The average number of device fail events is the product of the probability that a device fails before it is scheduled to be changed, times the average number of devices used. These two numbers depend on the mean device time-to-leakage in the patient and the variation in the device life-times. The simulations are run for many average life-times and standard deviations, and four types of probability distributions, Truncated Normal, Poisson, Uniform, and Exponential. The average life-time of the device before leakage, averaged over all simulated recipients, is 66 (± 59) days.

The device life-times are simulated with the parameters given in Table 1. Five device change policies are investigated (see Table 2 and the *Details* section) where *Wait-to-Leakage* (WtL) is the baseline on which the other policies are evaluated. The *Ideal* policy assumes that the distribution of the device life times is known for each patient and an optimal (ideal) time can be selected for the interval between changes. This policy simulates the best possible case if the life-times of individual devices are not predictable.

Ideal device change policy

Using this ideal case it can be investigated how outcomes depend on the variation in the device life-times. This is evaluated first on the different statistical distributions themselves, with 10^6 simulated devices and varying Coefficient of Variation, CV = $SD/Mean$ (Figure 1). The important parameters prove to be the Time-between-Leakages and the size of the CV) in the WtL policy. If f is the probability of a device leaking before the scheduled change and $\#Devices$ the number of devices used during a period, the time between leakages is the $\#Devices_{WtL}/(f \cdot \#Devices_{Ideal})$ times the average device life-time in the WtL policy. For a wide range of probability distributions it is found that the Time-to-Leakage using the *Ideal* policy, relative to the WtL times, is close to a linear function of the CV of the distribution, see Figure 1.

– In Figure 1, a dotted blue line illustrates the relation between time-between-leakages and $SD/Mean$ (= CV) for a probability of pre-schedulled failure, $f = 30\%$ (using index values with WtL = 100). This fail probability was chosen as it illustrates some important points. First, for a truncated Normal device-life distribution with

CV = 0.50, the Time-between-Leakages = 2.04 · average device life time. This means that the number of leakages is 49% of the leakages under the WtL policy. On the right hand side axis of Figure 1, the relative number of device changes is given (index values, top is WtL=100). Looking at the CV = 0.50 point, it can be seen that the number of device changes is 2.5 · #WtL (right-hand axis). This mean, a 51% reduction in number of leakage episodes is possible with 2.5 times as many device changes.

In Figure 1, the horizontal blue dotted lines illustrate the relation between time-between-leakages and $SD/Mean$ (= CV) for a probability of pre-schedulled failure, f = 30% (DL70). This fail probability was chosen as it illustrates some important points. The intercept of the horizontal blue dotted line with the left Y-axes gives the average increase in time-to-leakage (= 204) with respect to Wait-to-Leakage (= 100) for SD/Mean = 0.50 (= CV) of a truncated Normal distribution [2]. For this time-to-leakage, the average number of leakage events per year is 51% of that under the WtL policy. The DL70 reduction requires 63% extra device changes for a CV = 0.50 (right-hand axis).

In our patient data, the median CV = 0.8 (range [0.02- 3.01]). The dotted red line and red circle in Figure 1 indicate the achievable relative Time-to-Leakage, i.e., 125, and the increase in the number of device changes required, i.e., 167%, for a CV = 0.8. This Time-to-Leakage corresponds to an average reduction in the number of leakage events per year of only 20%.

For different leakage probabilities at change, f, the Y-axis intercepts (= $100/f$) and slopes change. This indicates that it might become difficult to reach a 51% reduction in leakage episodes if the variation in device life times increases, i.e., if $SD/Mean > 0.50$. The fail probability, f, must become even smaller. The cost in increased number of device changes also becomes more extreme. From Figure 1, it is clear that with very large variance, i.e., close to $SD/Mean = 1$, it is not possible to increase the Time-to-Leakage much even when using a large number of extra device changes.

Adaptive change policy

The results in Figure 1 are based on an unrealistic *Ideal* policy where the distribution of the device life-times is known and the number of device changes is very large. An *Adaptive* policy was designed to approximate the *Ideal* policy. In the *Adaptive* policy, an optimal time for device changes is calculated from previous changes to reduce the number of leakages to f=30%. As this method is "adaptive", there is some lead time before the system reaches the desired goal. This lead time has to be minimized if the policy is useful.

The *Adaptive* policy starts with three devices which are changed only on leakage. From these three device life-times, an estimation of the optimal device change time to reach the desired leakage probability per device, f, e.g., f=30% in Figure 1. This estimation is made using a truncated Normal distribution fit on the first three device changes. This value determines the next scheduled time-to-change for the 4^{th} device. From this change on, the look-back-time is incrementally increased to 10 changes. If a device fails before the scheduled time, it's life-time is known. If the device is changed before it leaks, the life time is not known, only its scheduled time. To improve the estimation of the optimal scheduled time in the light of random fluctuations in the device life-time, each time a device is changed before it leaks, it's life time is estimated as the scheduled time plus the average of the time by which all the devices that *did* leak were changed early, with noise added to get a well ordered sequence for the quantile function (i.e., few ties).

From the 4^{th} change on, the f quantile is calculated as the f quantile of the last k device changes, with k running from 4-10. This is the next time-to-change given the history. When the next device is changed, the time since the last device change is stored. If the device had to be changed before the scheduled time because of a leakage, the time to the scheduled change is stored too. If the device did not leak before the change, the average of the time-to-schedule of the devices that failed before the scheduled time is added to the next time-to-change with a small noise term.

The *Adaptive* policy was simulated with 5000 "patients" with random average per "patient" device life-times (mean=68 days, SD=61.0), and f=30%. Each simulated patient recieved 32 devices according to the *Adaptive* policy. With an average life-time of 68 days, 32 devices would last on average for approximately 6 years of use under the WtL policy. The results are displayed in Figure 2, the upper (gray) line. It can be seen that the average device failure approaches the f=30% line asymptotically, but rather slowly. The number of leaking devices stays well over 30% for the entire period. The approach to the target frequency f=30% is much faster if an undershoot is implemented. If the target fequency is set at $f^{1.5}$ during the first 10 devices, the approach to the target value f=30% is essentially achieved at the 10^{th} device (see Figure 2, black line). This overshoot-approach showed to work well for targets from f = 5% - 30% (DL95-DL70) and a range of values for $SD/Mean$ (0.125-1) and all statistical distributions. The average value of the leakage frequency for the last 17 device changes is ~27%, which close to the target f=30%.

The results of the simulations show that an *Adaptive* policy can approximate the *Ideal* policy in ~10 device changes for a wide range of parameters.

General policies

Table 3 gives the results for using the parameters in Table 1 for the five device change policies, i.e., a truncated Normal distribution with $SD/Mean$=0.50, and a leakage frequency at change f = 30% (DL70). The *Ideal* policy with f = 30% reaches a reduction of the number of leakage events to 47% of the WtL policy, at a cost of 49% more device changes. The Adaptaion strategy also reaches this goal after the inital adaptation period.

Table 1. Device life-time Monte Carlo parameters

Parameter	Value	
Mean time to leakage*	60	(± 60.0) days
Leak target	30	%
Number of changes	32	(~5 yr)
Number of simulated patients	5000	
Length of estimation window	3	(changes)
Distribution of time-to-leakage	trunc. Normal	SD/Mean = 0.50

* Average life time for each "patient" must be above 7 days. Lower average life times were replaced with a new sample. This results in effective average life times (and SD's) that are different from the original parameters, i.e., 66 (± 59) days.

Table 2. Policy options

Policy	Description
Wait-to-Leakage	Change prothesis after it starts leaking
Fixed time	Change after a fixed time (16 days), or when leaking
Ideal	Distribution of life-times is known per patient
Initial estimate	An optimal change time is estimated per patient after an initial 3 changes
Adaptive	The optimal change time is adapted based on the previous 3-10 changes with correction factors

Table 3. Results

Policy	Mean life-time (SD)	Mean % leaking (SD)	# leaks (%)	# changes (%)
Wait-to-Leakage	68 (61)	100 (0.0)	100	100
Fixed time	15 (2)	19 (26.8)	89	464
Ideal	46 (41)	31 (0.0)	47	149
Initial estimate	9 (6)	17 (15.6)	131	780
– after first 3	3 (5)	8 (17.2)	213	2589
Adaptive (^1.5)	43 (40)	34 (8.8)	55	160
– last 17	40 (38)	27 (12.0)	46	170
Adaptive (^1.0)	47 (43)	40 (8.5)	59	146
– last 17	43 (40)	32 (11.6)	50	158

Mean time between changes and percentage of changes due to leaking prosthesis (averaged per patient). Relative number of leakage events and device changes are given in the last two columns (% of Wait-to-Leakage). Simulation parameters from Table 1.*
Note the differeces in variance (SD) of the average number of changes due to leaks.

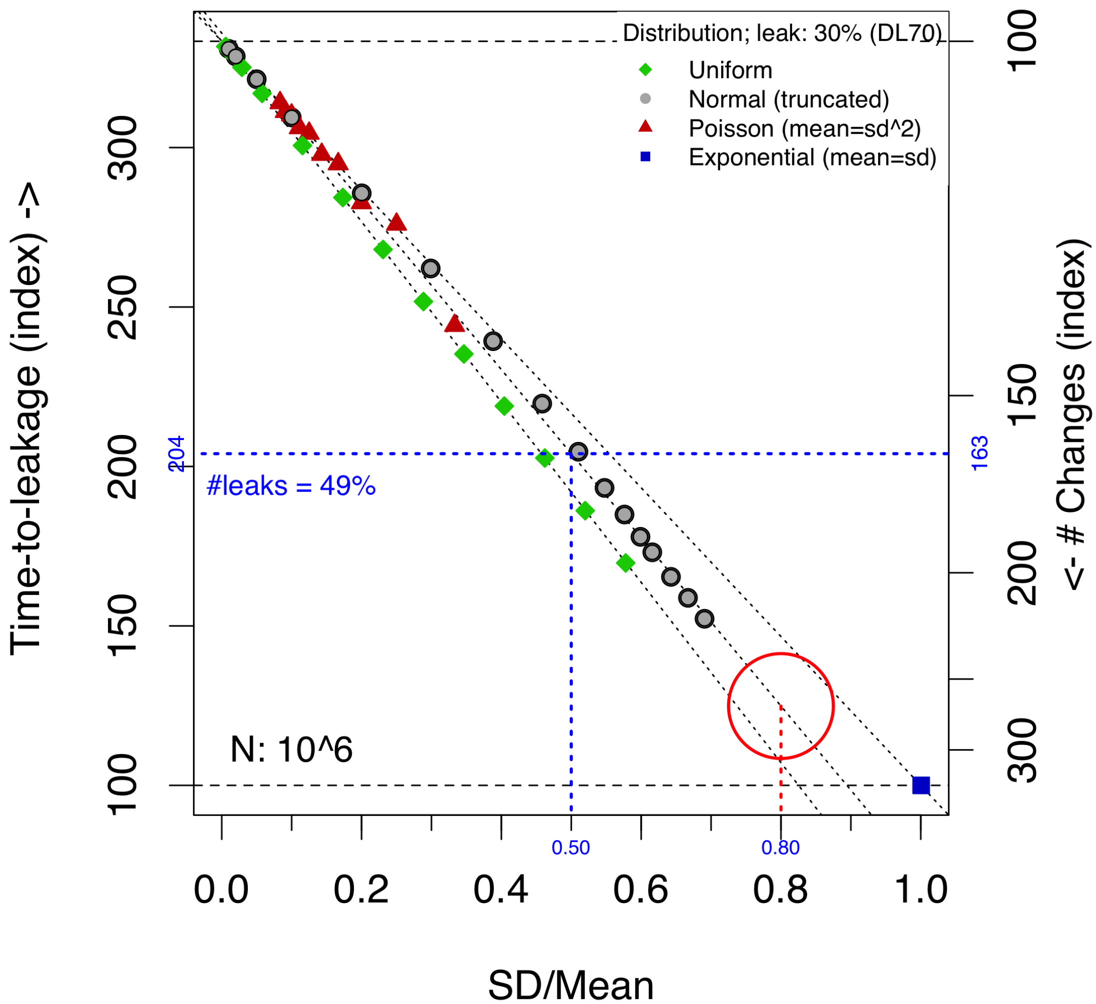

Figure 1. Mean time-to-leakage (left axis) and number of changes (right) versus the Standard Deviation as a fraction of the mean Time-to-Leakage (Coefficient of Variation) under the Ideal changing policy. Plotted are index values (Wait-to-Leakage = 100, mean = 66 days) for a target leakage fraction (f) of 30% (DL70). The simulation uses 10^6 life-times for each data point. The device leakage fraction at device-change is 30% (DL70) and the Y-axis intercept is $1/f = 333\%$. For different leakage fractions, f, the Y-axis intercepts differ. Results for four distributions of random life-times are plotted: Uniform (green diamonds), Poisson (red triangles), Normal (grey circles, truncated at 0), and Exponential (blue square). The Poisson distribution always has $SD = \sqrt{Mean}$, so the mean time-to-leakage is varied along the X-axis for the Poisson distribution. The Exponential distribution always has Mean=SD. Therefore, only a single point is plotted. The Exponential distribution always results in a mean time-to-leakage close to the mean time-to-leakage (100 in the plot). Blue and red dotted lines are example values (see text).
To prevent negative time-to-leakage for the Normal distribution, the distribution of life-times is truncated at 0. This affects both the actual mean time-to-leakage and SD. The CV = SD/Mean is calculated on the actual realized distribution. For large standard deviations, the distribution will deviate strongly from the Normal

distribution which will especially affect the resutls for small values of the leakage fraction.

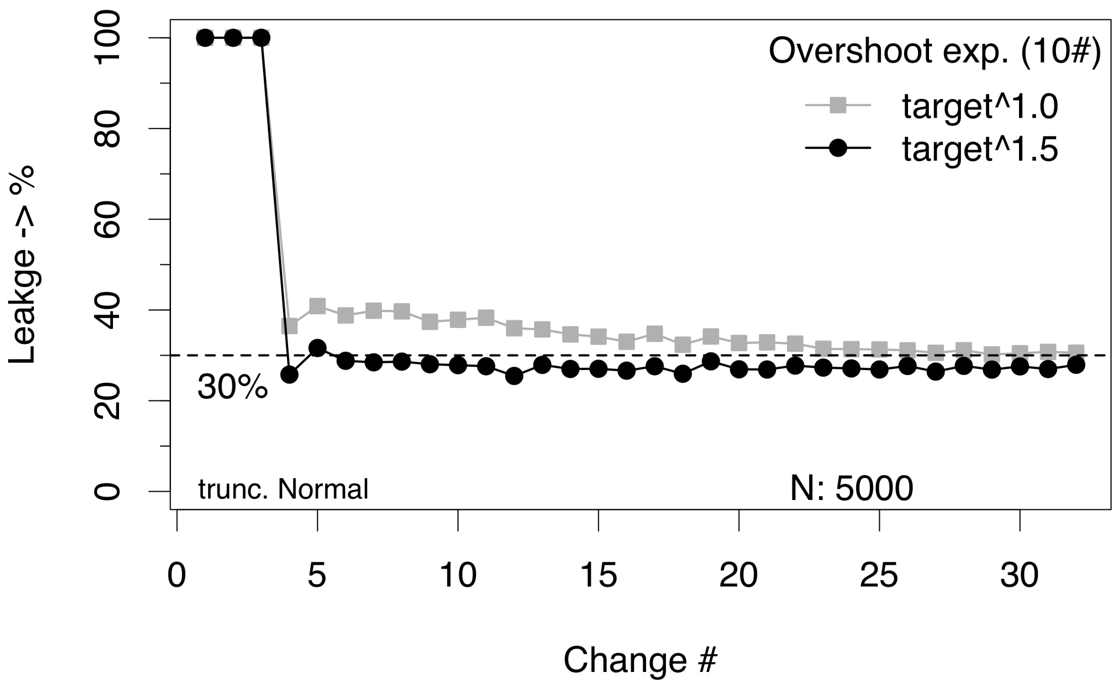

Figure 2. Average percentage leaking versus number of change for two settings of the adaptive change policy. Target long term leakage frequency: 30% (DL70). X-axis: Change number. Y-axis: The percentage of leaks at change time averaged over 5000 simulations. The simulation parameters were as above. The initial 3 changes were done when the prosthesis started to leak.

Details

Generating life-times of prosthesis

The Monte-Carlo algorithm generates 5000 sequences of 32 life times, each sequence with a specific average life time*. The individual average life-times are generated from an *Exponential* distribution with *mean=60, sd=60 (λ=0.017)* ⋆. The individual sequences are generated using a *trunc. Normal* distribution with *mean (±mean/2)* the specific average life-time (±SD).

All policies are applied to each sequence in turn.

* Average life time for each "patient" must be above 7 days. Lower average life times were replaced with a new sample. This results in effective average life times (and SD's) that are different from the original parameters, i.e., 66 (± 59) days.

Wait-to-Leakage

Prosthesis are changed when they start leaking. This is the "No Policy" baseline.

Fixed time

Prostheses are changed after 16 days for all speakers, unless they start leaking and are changed before the determined time. The number of days is estimated from the average results of the Ideal policy.

Ideal

Using the known sequence of life-times, the optimal time to change prostheses is determine for each "patient" to get as close to the desired 30% target percentage of failures. This is done by using the `quantile()` function with the target percentage.

In the case of an *Exponential* distribution (SD/Mean=1), the mean device time-to-change in percentage of the average time-to-leakage goes asymtoptically to the target leakage percentage (for high numbers of prostheses, here 10^6). This means that for an exponential distribution, the number of device leakages per year will at best be the same as the WtL policy for every profylactic change policy.

Initial estimate

The optimal time-to-change is estimated from the life-times of the first 3 prostheses. This is done by using the `qtruncnorm()` function [2] on the initial window with the target percentage. The first 3 prostheses will be used until they fail, after that, prostheses are changed after a fixed number of days or when they start leaking.

Two different statistics are given, one for for all changes, and statistics based on only the changes after the first 3.

Adaptive

After the first 3 prostheses, the following prostheses are changed after a pre-determined number of days. That number is determined after every change using the times until change of the last 3-10 prostheses. The change time is determined using the `qtruncnorm()` function [2] with the target percentage for the 4^{th} device, and the `quantile()` function later. The look-back window increases incremenatlly from 3-10. The scheduled device times are adaptated such that the probability of device leakage approaches the target value f (see above).

Two different settings for the adaptive policy are given with adaptive overshoot exponent 2 ; 1. For each setting, two different statistics are given, one for all changes, and statistics based on only the changes duting the last 17 changes.

Figure 2 give the average percentage of leakages per change number for 5000 simulated patients. It can be seen that the leakage frequency reaches the target frequency just after 10 changes.

References

[1] R Core Team (2019). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.

[2] Olaf Mersmann, Heike Trautmann, Detlef Steuer and Björn Bornkamp (2018). truncnorm: Truncated Normal Distribution. R package version 1.0-8. https://CRAN.R-project.org/package=truncnorm