

`tsoutliers` R Package for Automatic Detection of Outliers in Time Series

Javier López-de-Lacalle*

Draft version: May 30, 2014

Abstract

Time series data often undergo sudden changes that alter the dynamics of the data transitory or permanently. These changes are typically non-systematic and cannot be captured by standard time series models. That's why they are known as exogenous or outlier effects. Detecting outliers is important because they have an impact on the selection of the model, the estimation of parameters and, consequently, on forecasts. An automatic procedure described in the literature to detect outliers in time series is implemented in the `tsoutliers` R package.

Key words: outliers, time series, R.

1 Introduction

Time series data often undergo sudden changes that affect the dynamics of the data transitory or permanently. Sometimes, the practitioner may have some *a priori* information about the existence of these effects. A new regulation, strike periods, exceptional natural phenomena or a change in the methodology used to collect the data are some examples of events that may alter the overall dynamics of the data. In other cases, there is no such clear cut information available. The information may be unknown by the practitioner or it may not be feasible to carry out a review of all the historical events related to the context of the data. Regardless of their known or unknown source, these changes are typically non-systematic and cannot be captured by standard time series models (autoregressive integrated moving average models or structural time series models). A preliminary analysis of the data may reveal the existence of exceptional events that do not respond to patterns that can be captured by time series models, i.e. effects that are exogenous to the model. This analysis can be cumbersome or unfeasible if the data set is large. Moreover, sometimes these effects

*E-mail: javlacalle@yahoo.es.

do not stand out in graphical representations or *ad hoc* auxiliary regressions and remain masked within the overall dynamics of the data. When some *a priori* knowledge about these effects exists, then a regressor variable can be explicitly defined and included in the model. These regressors are known as *intervention* variables. When this exogenous effect is not related to any known event then we say that there is an *outlier* in the time series. We consider five types of outliers: innovational outlier, additive outlier, level shift, temporary change and seasonal level shift. Sometimes when the time and shape of an outlier is detected an *ex post* explanation can be found to it. In that case, the general structure of the outlier regressor should be replaced by the corresponding intervention variable.

Detecting and correcting the effect of outliers is important because they have an impact on the selection of the model, the estimation of parameters and, consequently, on forecasts and other results pursued by the analysis such as seasonal adjustment. In practice, we observe that many time series are affected by outliers and they are more common than intervention variables, especially when a large and heterogeneous set of time series is analysed, as it is the case of the analyses carried out by National Institutes of Statistics and Central Banks.

An automatic procedure described in the literature to detect and adjust the series for outliers is implemented in the package `tsoutliers` of R.

The remaining of the paper is organised as follows. ARIMA and structural time series models are briefly introduced in Section 2. The types of outliers considered in `tsoutliers` are described in Section 3. Section 4 describes the procedure to detect outliers. The performance of the implementation and a comparison with other specialised software is discussed by means of simulation exercises in Section 5. The usage of the package is illustrated with real examples in Section 6. Section 7 concludes.

2 Time series models

2.1 ARIMA time series models

A stationary ARMA model of orders (p, q) for a series y_t observed at time points $t = 1, 2, \dots, n$ can be defined as follows:

$$y_t = \sum_{i=1}^p \alpha_i y_{t-i} + a_t + \sum_{i=1}^q \theta_i a_{t-i}, \quad \text{where } a_t \sim NID(0, \sigma_a^2),$$

or more concisely

$$\phi(B)y_t = \theta(B)a_t, \quad \text{where } B \text{ is the lag operator.}$$

All the roots of the autoregressive polynomial, $\phi(B)$, and the moving average polynomial, $\theta(B)$, lie outside the unit circle. The parameters of the model $\alpha_{1,\dots,p}$, $\theta_{1,\dots,q}$ and σ_a^2 can be estimated by maximum likelihood (see for instance [1], [2, Chapter 5], and [10, Chapter 22]).

For non-stationary data, the observed series y_t is transformed by means of a differencing filter and then an ARMA model is defined for the differenced data. An ARIMA model for the observed series y_t is defined as follows:

$$\phi(B)\alpha(B)y_t = \theta(B)a_t. \quad (1)$$

The differencing filter that renders the data stationary is given by $\alpha(B)$, which is an autoregressive polynomial that contains all its roots on the unit circle.

2.1.1 Automatic ARIMA model selection

The detection of outliers can be conducted along with the model selection procedure implemented in the `forecast` package. For details see [6, 11].

2.2 Structural time series models

The basic structural model (BSM) is a pure variance structural model commonly used in applications. This model plays a central role in the approach advocated in [5] for time series analysis. A detailed view of the features and theoretical properties of this model can be found, for instance, in [5, Chapter 2], [2, Chapter 8] and [4, § 3.2]. The model is defined as follows:

$$\begin{aligned} \text{observed series: } y_t &= \mu_t + \gamma_t + \epsilon_t, & \epsilon_t &\sim \text{NID}(0, \sigma_\epsilon^2); \\ \text{latent level: } \mu_t &= \mu_{t-1} + \beta_{t-1} + \xi_t, & \xi_t &\sim \text{NID}(0, \sigma_\xi^2); \\ \text{latent drift: } \beta_t &= \beta_{t-1} + \zeta_t, & \zeta_t &\sim \text{NID}(0, \sigma_\zeta^2); \\ \text{latent seasonal: } \gamma_t &= \sum_{j=1}^{s-1} -\gamma_{t-j} + \omega_t, & \omega_t &\sim \text{NID}(0, \sigma_\omega^2), \end{aligned}$$

for $t = s, \dots, n$, where s is the periodicity of the data.

The BSM encompasses models that are common in applications: the local level model, that consists of a random walk with a deterministic drift, β_0 , plus a noise component, ϵ_t ; the local trend model, where the drift follows a random walk. Setting $\sigma_\omega^2 = 0$ yields a model with deterministic seasonality. Setting also $\gamma_1 = \dots = \gamma_{s-1} = 0$ removes the seasonal component and gives the local trend model. Adding the restriction $\sigma_\zeta^2 = 0$ yields the local level model.

The state space form of the BSM is given by the following representation:

$$\begin{aligned} y_t &= Z\alpha_t + \epsilon_t, & \epsilon_t &\sim \text{NID}(0, \sigma_\epsilon^2), \\ \alpha_t &= T\alpha_{t-1} + R\eta_t, & \eta_t &\sim \text{NID}(0, V), \\ \alpha_0 &\sim \text{N}(a_0, P_0), \end{aligned} \quad \text{with } V = \begin{bmatrix} \sigma_\xi^2 & 0 & 0 \\ 0 & \sigma_\zeta^2 & 0 \\ 0 & 0 & \sigma_\omega^2 \end{bmatrix},$$

for $t = 1, \dots, n$. For $s = 4$, the matrices of this representation are defined as follows:

$$y_t = \begin{bmatrix} 1 & 0 & 1 & 0 & 0 \end{bmatrix} \alpha_t + \epsilon_t,$$

$$\alpha_t \equiv \begin{bmatrix} \mu_t \\ \beta_t \\ \gamma_t \\ \gamma_{t-1} \\ \gamma_{t-2} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & -1 & -1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix} \times \begin{bmatrix} \mu_{t-1} \\ \beta_{t-1} \\ \gamma_{t-1} \\ \gamma_{t-2} \\ \gamma_{t-3} \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \times \begin{bmatrix} \xi_t \\ \zeta_t \\ \omega_t \end{bmatrix}.$$

The initial state vector, a_0 , and its covariance, P_0 , can be chosen beforehand or included in the set of parameters to be estimated. Given the variance parameters, the Kalman filter and smoother can be applied to extract an estimate of the the latent components (level, trend and seasonal).

3 Types of outliers

A general representation of an outlier can be given by the following expression: $L(B)I(t_j)$, where $L(B)$ is a polynomial of lag operators and $I(t_j)$ is an indicator variable that takes on the value 1 at time $t = j$ when the outlier springs and the value 0 elsewhere. We consider five types of outliers: innovational outlier (IO), additive outlier (AO), level shift (LS), temporary change (TC) and seasonal level shift (SLS). The polynomial $L(B)$ for each type of outlier is defined as follows:

$$\begin{aligned} \text{IO: } L(B) &= \frac{\theta(B)}{\alpha(B)\phi(B)}; & \text{LS: } L(B) &= \frac{1}{(1-B)}; & \text{SLS: } L(B) &= \frac{1}{(1-B^s)}; \\ \text{AO: } L(B) &= 1; & \text{TC: } L(B) &= \frac{1}{(1-\delta B)}. \end{aligned} \tag{2}$$

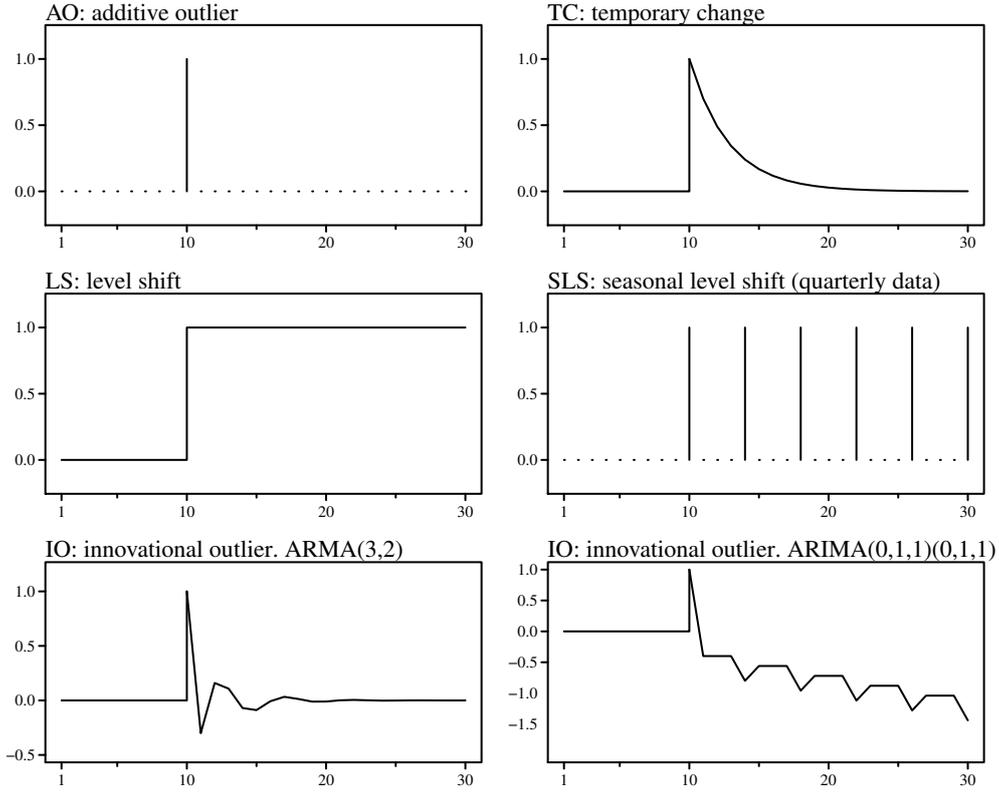
The value of δ is usually set equal to 0.7 and s is the periodicity of the data (e.g. 4 in quarterly data and 12 in monthly data). Each of them has a specific structure and influence on the data. Figure 1 shows a unit impulse for each type of outlier occurring at time point $t = 10$.

4 Automatic detection procedure

The procedure described in [3] for the automatic detection of outliers is implemented in the package `tsoutliers`. Next, we introduce the procedure following the notation of the original reference. At the same time, we give an insight into the implementation in package `tsoutliers` by showing and discussing the main parts of the code.

Let us define y_t^* as the observed series subject to m outliers with weights w . The ARIMA model

Figure 1: Unit impulse for different types of outliers



for the observed series is written as:

$$y_t^* = \sum_{j=1}^m \omega_j L_j(B) I_t(t_j) + \frac{\theta(B)}{\phi(B)\alpha(B)} a_t. \quad (3)$$

The estimated residuals, which are contaminated with the outliers, are given by:

$$\pi(B)y_t^* \equiv \hat{e}_t = \sum_{j=1}^m \omega_j \pi(B) L_j(B) I_t(t_j) + a_t, \quad (4)$$

where the coefficients of the power series expansion $\pi(B) = \sum_{i=0}^{\infty} \pi_i B^i$ can be determined from the relation (see for instance [2, Chapter 3]):

$$\pi(B) = \frac{\phi(B)\alpha(B)}{\theta(B)}.$$

Since all the roots of $\theta(B)$ lie outside the unit circle, the coefficients π_i satisfy $\sum_{i=0}^{\infty} |\pi_i| < \infty$ and vanish to zero.

Equations (3) and (4) constitute the grounds of the procedure for the detection of outliers, which consists of the following stages:

stage I: Locate outliers. Given an ARIMA model fitted to the data, outliers are detected and located by checking the significance of all types of outliers at all possible time points. This is done by means of the corresponding t -statistics in the regression equation (4).

stage II: Remove outliers. Given a set of potential outliers, an ARIMA model is chosen and fitted according to the regression equation (3). The significance of the outliers is reassessed in the new fitted model. If an ARIMA selection procedure is used a new model may be selected at this stage. Those outliers that are not significant are removed from the set of potential outliers.

stage III: Iterate stages I and II, first for the original series and then for the adjusted series.

4.1 Stage I: locate outliers

Replacing in equation (4) $L(B)$ by the definition of each type of outlier defined above and multiplying by $\pi(B)$, we can see that the equation becomes, respectively for each type of outlier:

$$\begin{aligned} \text{IO: } \hat{e}_t &= \omega I_t(t_j) + a_t, & \text{AO: } \hat{e}_t &= \omega \pi(B) I_t(t_j) + a_t, \\ \text{LS: } \hat{e}_t &= \omega \frac{\pi(B)}{1-B} I_t(t_j) a_t, & \text{TC: } \hat{e}_t &= \omega \frac{\pi(B)}{1-\delta B} I_t(t_j) a_t. \end{aligned} \tag{5}$$

[3] adopt the following t -statistics to identify and locate outliers:

$$\hat{\tau}_k = \frac{\hat{\omega}_k(t_j)}{\hat{\sigma}_a \left(\sum_{i=t_j}^n x_{k,i}^2 \right)^{1/2}}, \quad \text{with } \hat{\omega}_k = \frac{\sum_{i=t_j}^n \hat{e}_t x_{k,i}}{\sum_{i=t_j}^n x_{k,i}^2},$$

where k stands for each type of outlier, $k = \text{IO, AO, LS, TC}$. [3] derive the structure of each regressor $x_k = \pi(B)L_j(B)I_t(t_j)$ employed in equation (4). The residual standard deviation $\hat{\sigma}_a$ is estimated as the median absolute deviation (MAD) defined as:

$$\hat{\sigma}_a = 1.483 \times \text{median}(|\hat{e} - \tilde{e}|), \text{ with } \tilde{e} = \text{median}(\hat{e}).$$

At first glance, the computation of all the τ statistics may seem more cumbersome than it actually is. In principle, the above regression equations should be run for all types of outliers and for all the potential locations of outliers $t_j = 1, 2, \dots, n$. However, we can take advantage of the fact that, given a type of outlier, the regressors x_k for different time points t_j are shifted versions of x_k defined for $t_j = 1$. As we shall see, we do not need to explicitly build all the regressors and run a regression for each time point t_j and for each type of outlier.

The code below loads the required packages and generates a series from an ARIMA(1,0,1) model with three exogenous shocks (two additive outliers and one level shift). In the sequel, we will use this simulated series to illustrate the procedure.

```
R> library("tsoutliers")
R> set.seed(123)
```

```

R> y <- arima.sim(model = list(ar = 0.7, ma = -0.4), n = 120)
R> y[15] <- -4
R> y[45] <- 5
R> y[80:120] <- y[80:120] + 5
R> y <- round(y, 2)

```

Next, we start the procedure. First, an ARIMA model is chosen and fitted by means of the function `auto.arima` available in the package `forecast`. No drift is included in the model and the Bayesian information criterion (BIC) is used to select the model. The autoregressive and moving average coefficients as well as the residuals and its length are extracted from the fitted model. The standard deviation of residuals, $\hat{\sigma}_a$, is computed by means of the MAD method defined above.

```

R> fit <- forecast::auto.arima(x = y, allowdrift = FALSE, ic = "bic")
R> pars <- coefs2poly(fit)
R> resid <- residuals(fit)
R> n <- length(resid)
R> sigma <- 1.483 * quantile(abs(resid - quantile(resid, probs = 0.5)), probs = 0.5)

```

Upon the information obtained from the fitted model we can compute the τ statistics as follows. The coefficients of $\pi(B)$ are obtained by means of the function `ARMAtoMA` from the `stats` package. Notice that this function computes the coefficients in the relation $\psi(B) = \frac{\theta(B)}{\phi(B)} = \frac{1}{\pi(B)}$. After changing the sign of the parameters and passing the AR coefficients as the MA coefficients and vice versa, this function returns the coefficients of the polynomial $\pi(B)$. Looking at equations (5), we can see that the τ statistics related to the innovational outlier can be easily obtained as the residuals divided by $\hat{\sigma}_a$.

```

R> picoeefs <- c(1, ARMAtoMA(ar = -pars$macoefs, ma = -pars$arcoef, lag.max = n-1))
R> tauI0 <- resid / sigma

```

The τ statistics for the AO outliers are obtained taking advantage of the fact that the regressors involved in equation (5) for each time point $t = t_j$ are shifted versions of the case $I(t_j = 1)$.

```

R> padded.resid <- c(resid, rep(0, n-1))
R> xy <- as.vector(na.omit(filter(x = padded.resid,
+   filter = rev(picoeefs), method = "conv", sides = 1)))
R> xx <- rev(cumsum(picoeefs^2))
R> tauA0 <- xy / (sigma * sqrt(xx))

```

Similarly, we obtain the τ statistics related to the LS and TC after filtering $\pi(B)$ through $1/(1-B)$ and $1/(1-\delta B)$, respectively.

```

R> di.picoefs <- diffinv(picoefs)[-1]
R> xy <- as.vector(na.omit(filter(x = padded.resid,
+   filter = rev(di.picoefs), method = "conv", sides = 1)))
R> xx <- rev(cumsum(di.picoefs^2))
R> tauLS <- xy / (sigma * sqrt(xx))
R> delta <- 0.7
R> didelta.picoefs <- filter(x = picoefs, filter = delta, method = "rec")
R> xy <- as.vector(na.omit(filter(x = padded.resid,
+   filter = rev(didelta.picoefs), method = "conv", sides = 1)))
R> xx <- rev(cumsum(didelta.picoefs^2))
R> tauTC <- xy / (sigma * sqrt(xx))

```

We have seen that the computation of all the τ statistics involves fewer computations than may seem necessary at first glance and it can be done without explicit loops. Now, let's look at some of the values that we obtained for these statistics.

```

R> tau.stats <- data.frame(IO = tauIO, AO = tauAO,
+   LS = tauLS, TC = tauTC, row.names = NULL)
R> round(tau.stats[c(14:16,44:46,78:82),], 3)

```

	IO	AO	LS	TC
14	1.119	1.386	0.105	-0.406
15	-4.103	-4.797	-0.930	-2.397
16	2.322	1.613	2.655	2.865
44	-0.535	-1.096	0.786	1.245
45	4.934	5.517	1.605	3.216
46	-2.883	-2.405	-2.518	-2.640
78	1.755	-0.028	4.411	1.595
79	1.215	-0.734	4.432	2.316
80	4.325	2.984	4.981	4.271
81	1.958	1.093	2.751	2.189
82	1.231	0.582	1.934	1.695

Those statistics that are in absolute value higher than a threshold identify the time point of a potential outlier. The code below identifies those statistics that are higher than the threshold in absolute value. Here, we use the threshold or critical value `cval = 3.5`. The potential outliers are stored in a `data.frame` named `mo` containing the type of outlier, the observation index and the t -statistic.

```

R> cval <- 3.5
R> indIO <- which(abs(tauIO) > cval)
R> indAO <- which(abs(tauAO) > cval)

```

```

R> indLS <- which(abs(tauLS) > cval)
R> indTC <- which(abs(tauTC) > cval)
R> mo <- data.frame(type = factor(c(rep("IO", length(indIO)),
+   rep("AO", length(indAO)), rep("LS", length(indLS)), rep("TC", length(indTC))),
+   levels = c("IO", "AO", "LS", "TC")),
+   ind = c(indIO, indAO, indLS, indTC),
+   tstat = c(tauIO[indIO], tauAO[indAO], tauLS[indLS], tauTC[indTC]))
R> cbind(mo[order(mo[, "ind"])[1:5],], " " = "   ", mo[order(mo[, "ind"])[6:10],])

```

	type	ind	tstat	type	ind	tstat
1	IO	15	-4.102540	LS	79	4.432065
4	AO	15	-4.797319	IO	80	4.325142
2	IO	45	4.933745	LS	80	4.980832
5	AO	45	5.517405	TC	80	4.271211
6	LS	78	4.410770	<NA>	NA	NA

We can see that, for a given time point, there are more than one type outlier exceeding the threshold (e.g. at $t = 15$, $\tau_{IO} = -4.103$ and $\tau_{AO} = -4.797$). Duplicates are removed keeping the type of outlier with the highest t -statistic in absolute value.

```

R> for (i in mo[, "ind"][duplicated(mo[, "ind"])])
+ {
+   ind <- which(mo[, "ind"] == i)
+   moind <- mo[ind,]
+   if ("IO" %in% moind[, "type"]) {
+     tmp <- moind[which(moind[, "type"] != "IO"),]
+   } else
+     tmp <- moind[which.max(abs(moind[, "tstat"]))],]
+   mo <- mo[-ind,]
+   mo <- rbind(mo, tmp)
+ }
R> mo[order(mo[, "ind"]),]

```

	type	ind	tstat
4	AO	15	-4.797319
5	AO	45	5.517405
6	LS	78	4.410770
7	LS	79	4.432065
8	LS	80	4.980832

The function `locate.outliers` performs the location of outliers in the lines shown above except that it adds one additional rule: when a LS is identified at consecutive time points only the point

related to the highest τ statistic is kept. Thus, if we run the function, we obtain the result stored above in object `mo` except that the LS at times 78 and 79 are not included since $|4.411| < |4.980|$ and $|4.432| < |4.980|$.

```
R> locate.outliers(resid, pars, cval, types = c("IO", "AO", "LS", "TC"))
```

```

  type ind  coefhat    tstat
4  AO  15 -4.450352 -4.797319
5  AO  45  5.118357  5.517405
8  LS  80  3.452909  4.980832

```

The function `locate.outliers` is the baseline of stage I of the procedure. The complete stage consists of an outer loop and an inner loop. The inner loop is implemented in the function `locate.outliers.iloop` and can be sketched as follows:

```
R> mo.iloop <- NULL
R> maxit.iloop <- 4
R> iter <- 0
R> resid.iloop <- resid
R> while (iter < maxit.iloop)
+ {
+   mo <- locate.outliers(resid.iloop, pars, cval, types = c("IO", "AO", "LS", "TC"))
+   if (nrow(mo) == 0)
+     break
+   mo.iloop <- rbind(mo.iloop, mo)
+   oxreg <- outliers.regressors(pars, mo, n, weights = TRUE)
+   resid.iloop <- resid.iloop - rowSums(oxreg)
+   iter <- iter + 1
+ }
R> mo.iloop
R> cat(sprintf("number of iterations: %s\n", iter))
```

```

  type ind  coefhat    tstat
4  AO  15 -4.450352 -4.797319
5  AO  45  5.118357  5.517405
8  LS  80  3.452909  4.980832
number of iterations: 1

```

First, the following objects are defined: `mo.iloop` stores the potential set of outliers; `maxit.iloop` is the maximum number of allowed iterations; `iter` is the counter; and `resid.iloop` stores the residuals adjusted for the outlier effects. Then the inner loop begins. At each iteration the effect of the outliers on the residuals is removed and a new check for the presence of outliers is carried out.

The function `outliers.regressors` builds the regressors x_k from equation (5) weighted for the parameter estimates (column "coefhat" in object `mo`). The inner loop stops when no additional outliers are detected or when a maximum number of allowed iterations is reached. The output shows that only one iteration was required to identify outlying observations.

The outer loop is implemented in the function `locate.outliers.oloop`. The main parts of the code are shown below. First, some variables are created: `y.oloop` and `fit.oloop` to avoid overwriting the objects `y` and `fit`; `moall`, the object that will store information about the outliers; `maxit.oloop` maximum number of allowed iterations and the counter `iter`.

```
R> y.oloop <- y
R> fit.oloop <- fit
R> moall <- NULL
R> maxit.oloop <- 4
R> iter <- 0
R> while (iter < maxit.oloop)
+ {
+   pars <- coefs2poly(fit)
+   resid.oloop <- residuals(fit.oloop)
+   resid.med <- quantile(resid.oloop, probs = 0.5)
+   sigma <- 1.483 * quantile(abs(resid.oloop - resid.med), probs = 0.5)
+   mo.iloop <- locate.outliers.iloop(resid.oloop, pars, cval,
+     types = c("IO", "AO", "LS", "TC"), maxit = maxit.iloop)
+   if (nrow(mo.iloop) == 0)
+     break
+   moall <- rbind(moall, mo.iloop)
+   oeff <- outliers.effects(mo.iloop, n, weights = TRUE, pars = pars)
+   y.oloop <- y.oloop - rowSums(oeff)
+   fit.oloop <- arima(y.oloop, order = fit$arma[c(1,6,2)],
+     seasonal = list(order = fit$arma[c(3,7,4)]))
+   iter <- iter + 1
+ }
R> moall
R> cat(sprintf("number of iterations: %s\n", iter))

  type ind  coefhat   tstat
4  AO  15 -4.450352 -4.797319
5  AO  45  5.118357  5.517405
8  LS  80  3.452909  4.980832
number of iterations: 1
```

At each iteration, the required required is extracted from the last fitted model (parameter estimates,

residuals and MAD of residuals). The inner loop is run to identify outliers and then the model is fitted again for the series adjusted for the outliers identified at the current iteration. The function `outliers.effects` creates a matrix of outliers effects as shown in Figure 1; `weights` is set to `TRUE` in order to weight each outlier by the estimated coefficient instead of using a unit impulse. The loop stops when no more outliers are detected by the inner loop or when the maximum number of allowed iterations is reached. The final set of outliers is printed. We can see that it is in agreement with the relatively prominent shocks added to the simulated series.

For completeness, we show how a call to the function `locate.outliers.oloop` would be defined to run the inner and outer loops and obtain the results shown above.

```
R> stage1 <- locate.outliers.oloop(y, fit, types = c("IO", "AO", "LS", "TC"),
+   cval = cval, maxit.iloop = maxit.iloop)
R> stage1$outliers

   type ind  coefhat    tstat
4   AO  15 -4.450352 -4.797319
5   AO  45  5.118357  5.517405
8   LS  80  3.452909  4.980832
```

4.2 Stage II: remove outliers

Although the detection of outliers is already made in stage I, we must be aware that the initial selection of the ARIMA model may be affected by the presence of outliers (especially the order of integration) and vice versa, that is, the detection of outliers may be affected by the choice of the model and the parameter estimates based on the observed series.

To avoid finding spurious outliers caused by an improper choice or estimation of a model, the equation (3) is fitted for the observed series including the outlier regressors identified in stage I. If any of the outliers turn to be non-significant then they are removed from the set of potential outliers. The function `remove.outliers` performs this operation as follows.

```
R> moall <- stage1$outliers
R> xreg <- outliers.effects(moall, n, weights = FALSE, pars = pars)
R> iter <- 0
R> while (TRUE)
+ {
+   fit <- forecast::auto.arima(x = y, xreg = xreg, allowdrift = FALSE, ic = "bic")
+   xregcoefs <- coef(fit)
+   ind <- match(colnames(xreg), names(xregcoefs))
+   xregcoefs <- xregcoefs[ind]
```

```

+   tstats <- xregcoefs / sqrt(diag(fit$var.coef)[ind])
+   ref <- which(abs(tstats) < cval)
+   if (length(ref) > 0)
+   {
+     moall <- moall[-ref,]
+     xreg <- data.matrix(xreg[,-ref])
+   } else
+     break
+   if (nrow(moall) == 0)
+     break
+   iter <- iter + 1
+ }

```

The `data.frame` `moall` contains information about the set of outliers detected in stage I. The function `outliers.effects` builds the regressors and stores them in the object `xreg`. Notice that here the regressors are not the variables x_k used in stage I but a variable capturing the exogenous effect according to each type of outlier as was defined in equation (2) and depicted in Figure 1. An ARIMA model is selected again for the observed data by means of `forecast::auto.arima` including the object `xreg` as regressor variables. Then the t -statistics are computed and stored in `tstats`. If any of them is below the threshold (in absolute value) then the corresponding outlier is removed from `moall`. The process is iterated until all the outliers are significant in the last model selected for the data. This process to remove outliers corresponds to the option `method = "en-masse"` in function `remove.outliers`. Alternatively, the significance of the outliers can be tested in a bottom-up procedure (see the documentation of the function).

A summary output for this stage is printed below.

```

R> print(fit)
R> moall[, "coefhat"] <- xregcoefs
R> moall[, "tstat"] <- tstats
R> print(moall)

Series: y
ARIMA(1,0,0) with zero mean

Coefficients:
      ar1      A015      A045      LS80
0.3023 -4.6067  5.4875  4.6667
s.e.  0.0909  0.8736  0.8689  0.1986

sigma^2 estimated as 0.8077:  log likelihood=-157.51

```

```
AIC=325.02  AICc=325.55  BIC=338.96
```

```
  type ind  coefhat    tstat
4  AO  15 -4.606657 -5.273256
5  AO  45  5.487542  6.315486
8  LS  80  4.666688 23.492144
```

Despite no outliers were removed, a different model is now selected. Although we did not print it, the model chosen at the beginning of stage I in the previous subsection was an ARIMA(0, 1, 1), now an ARIMA(1, 0, 0) is chosen. When the level shift is not considered, the data cannot be modelled by a stationary ARMA model. Despite there is a deterministic level shift a stochastic trend is considered and the first order differencing filter is applied. After detecting and including the effect of the level shift, its effect is estimated and a stationary ARMA model is chosen by the automatic model selection procedure.

4.3 Stage III: iterate stages I and II for the adjusted series

Here, the stages I and II are iterated for the adjusted series. If new outliers are found they are added to the set of outliers found in the previous iteration.

The original series can be adjusted for outliers defined in `moall` as follows. The function `outliers.effects` builds the outlier variables according to each type. The total outliers effect is measured and stored in `oeff`. The series adjusted for outliers is stored in `yadj` (the linearized series).

```
R> tmp <- coefs2poly(fit)
R> xreg <- outliers.effects(moall, n, weights = FALSE, pars = tmp)
R> xregcoefs <- coef(fit)
R> ind <- match(colnames(xreg), names(xregcoefs))
R> xregcoefs <- xregcoefs[ind]
R> oeff <- xreg %*% cbind(xregcoefs)
R> yadj <- y - oeff
R> plot(y, col = "gray80")
R> lines(yadj)
```

After choosing a model for the adjusted series and calling directly the function `locate.outliers`, we see that no additional outliers are detected and thus the procedure is finished. Figure 2 shows the original data (gray line), the adjusted series (black line), the location of the detected outliers (red points) and their estimated effects (red line).

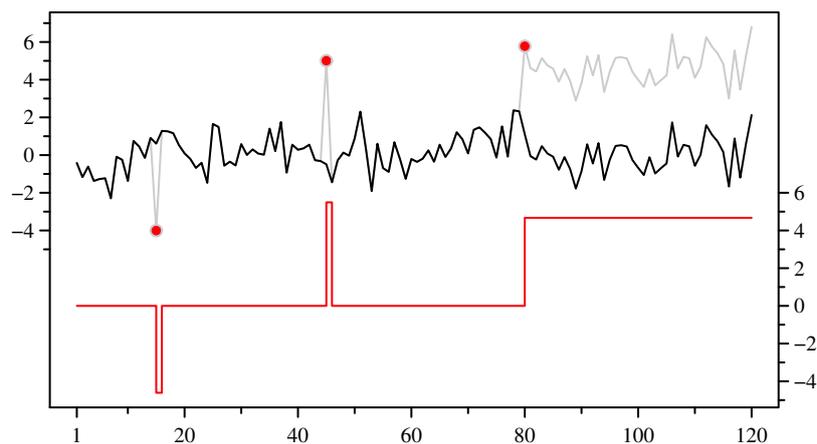
```
R> fit <- forecast::auto.arima(x = yadj, allowdrift = FALSE, ic = "bic")
R> tmp <- coefs2poly(fit)
```

```
R> resid <- residuals(fit)
R> sigma <- 1.483 * quantile(abs(resid - quantile(resid, probs = 0.5)), probs = 0.5)
R> locate.outliers(resid, pars, cval, types = c("IO", "AO", "LS", "TC"))

[1] type    ind    coefhat tstat
<0 rows> (or 0-length row.names)
```

In most cases, the user will be interested in running the procedure in one call rather than by stages as we did here for illustration. That's precisely the point of an automatic procedure. The function `tsoutliers0` runs stages I and II and the function `tsoutliers` iterates the procedure running `tsoutliers0` first for the original series and then for the adjusted series. These functions are relatively flexible and allows the user to control several parameters of the procedure such as the maximum number of iterations in the loops, the threshold and the arguments passed to `forecast::auto.arima` to select the ARIMA model, among other parameters.

Figure 2: Original series, adjusted series and estimated outlier effects



4.4 Polishing rules

In practice, we may encounter some series for which the procedure as described above would not return a reliable result. If the initial model and parameter estimates do not fit well the overall dynamics of the data, the procedure may find too many outliers or the estimate of the outlier effects may not be accurate enough. In this case, the correction for outliers may induce some artificial effects in the adjusted series that may in turn be identified as outliers in the next run of the procedure. Next, we discuss some polishing rules that are included in the code of package `tsoutliers`. Some of them have already been mentioned above.

- By default, the threshold against which the t -statistics are compared is determined based on the sample size, n . `cval` = 3 if $n \leq 50$; `cval` = 4 if $n \geq 450$; `cval` = $3 + 0.0025 \times (n - 50)$ otherwise.
- In the adjusted series, which is expected to be cleaner from outlying observations, the threshold or critical value, `cval`, can be reduced by a factor `cval.reduce` so that the threshold `cval` \times $(1 - \text{PC})$ is used.
- In a previous version of the package this rule was applied: The outliers AO, LS, and TC have precedence over IO when the t -statistic related to more than one type of outlier exceed the threshold `cval` at a given time point, i.e., the IO is ignored and the other type of outlier is kept. In the current version this rule is not applied but the IO type of outlier is not considered by default.
- If at a given iteration of the inner or outer loop an outlier is found at a time point where an outlier was already detected in a previous iteration, then type of outlier that was first detected is kept.
- If the t -statistics related to the LS type for outlier are beyond the critical value for consecutive time points, only the point with the highest t -statistic is kept; i.e., we do not consider two level shifts at a given observation and the very next time point.
- The inner and outer loops of stage I are stopped if a predefined maximum number of iterations is reached. This does not seem necessary for the outer loop but a limit is set for safety.
- When regular or seasonal differences are taken, the residuals related to the first observations may be overly erratic. If the maximum value of the first residuals are beyond a threshold then they are set to zero in function `locate.outliers.oloop`. In this case, the first observations are not checked for the presence of outliers. The number of observations that may be omitted is $d + D \times s$, where d is the number of regular differences, D is the number of seasonal differences and s is the periodicity of the data. The threshold is defined as 3.5 times the standard deviation of the remaining residuals.

4.5 Adaptation of the procedure to structural time series models

The procedure to detect outliers can be adapted to the framework of the basic structural model introduced in § 2.2.

The t -statistics related to IO are computed as in the ARIMA model since the regressor is a unit impulse variable.

```
R> tauIO.stsm <- resid / sigma
```

The object `resid` is the residuals from structural model and `sigma` is the MAD of the residuals that was defined before.

The statistics for AO are computed similarly as in the ARIMA model. Here, the Kalman filter is run taking as input the indicator variable $I(1)$. To avoid the problems inherent to the initialization of the Kalman filter, `n.start` additional zeros are included in the indicator variable. They act as warming observations and are removed afterwards from the final vector of coefficients `f`. The computations for the remaining time points take advantage of the fact that the regressors are shifted versions of the regressor obtained for $I(1)$, as we did in the ARIMA case.

```
R> n.start <- 50
R> I <- ts(rep(0, length(resid) + n.start - 1),
+   start = start(resid), frequency = frequency(resid))
R> I[n.start] <- 1
R> tmp <- KFKSDS::KF(I, ss)
R> f <- tmp$v[-seq.int(n.start-1)]
R> ao.xy <- as.vector(na.omit(
+   filter(x = c(resid, rep(0, n-1)), filter = rev(f), method = "conv", sides = 1)))
R> ao.xx <- rev(cumsum(f^2))
R> tauAO.stsm <- ao.xy / (sigma * sqrt(ao.xx))
```

The object `ss` is a list containing the matrices of the state space representation of the structural model. If the package `stsm` [9] was used to fit the model, the function `stsm.class::char2numeric` can be used to easily define this object.

The statistics for the LS reuse the values obtained above and stored in `ao.xy`. The inverse of the difference filter is applied on it and then the statistics are computed as usual.

```
R> ls.xy <- rev(diffinv(rev(ao.xy))[-1])
R> ls.xx <- rev(cumsum(diffinv(f)[-1]^2))
R> tauLS.stsm <- ls.xy / (sigma * sqrt(ls.xx))
```

The operations for the TC are the same as for LS except that instead of the inverse of the differencing filter $(1 - L)$, the inverse of $(1 - \delta L)$ is applied.

```
R> tc.xy <- rev(filter(x = rev(ao.xy), filter = delta, method = "rec"))
R> dinvf <- filter(x = f, filter = delta, method = "rec")
R> tc.xx <- rev(cumsum(dinvf^2))
R> tauTC.stsm <- tc.xy / (sigma * sqrt(tc.xx))
```

Given the t -statistics, the automatic detection procedure can be applied as described in the previous section, except that the structural model is fitted instead of an ARIMA model. Notice also that since the components of the structural model are defined beforehand, there is no need to reassess the choice of the model but only to refit the parameters of the model.

5 Simulation exercises

Under development, sorry.

6 Applications to real data

6.1 Two popular time series

Time series methods described in the literature are often applied to two popular time series: 1) the Nile time series, measurements of the annual flow of the river Nile at Ashwan in the period 1871-1970 and 2) the classic Box & Jenkins airline data, monthly totals of international airline passengers in the period 1949 to 1960. These data are available in the `datasets` package that comes with the base R system.

The Nile time series is a common example of the local level model, the simplest structural time series model. Next, we run `tsoutliers` for the local level model; `"linear-regression"` is chosen as the `remove.method`; the procedure is applied only on the original series, `maxit = 1`.

```
R> resNile1 <- tsoutliers(y = Nile, types = c("AO", "LS", "TC"),
+   maxit = 1, remove.method = "linear-regression",
+   tsmethod = "stsm", args.tsmode = list(model = "local-level"))
R> resNile1$fit$call$xreg<-NULL
R> resNile1
```

Warning message:

```
In maxlik.fd.scoring(m = <S4 object of class "stsm">, xreg = c(0, :
  Possible convergence problem. Maximum number of iterations reached.
```

Call:

```
maxlik.fd.scoring(m = <S4 object of class structure("stsm", package = "stsm.class">,
  step = NULL, information = "expected")
```

Variance parameters:

	var1	var2
Estimate	11600	819.3
Std. error	1564	380.7

Coefficients of regressor variables:

	A07	LS29	A043	TC46
Estimate	-148.4	2.545	-135.6	131.3
Std. error	85.5	120.917	85.5	111.5

Log-likelihood: -625.8488

Convergence: 0

Number of iterations: 101

Outliers:

	type	ind	time	coefhat	tstat
1	AO	7	1877	-148.441	-1.73613
2	LS	29	1899	2.545	0.02105
3	AO	43	1913	-135.570	-1.58559
4	TC	46	1916	131.312	1.17789

We can see that some outliers were detected. However, the algorithm to fit the model (by default, maximum likelihood in the frequency domain by means of the scoring algorithm) did not converged and was stopped after 100 iterations. The standard errors related to the coefficients of the outlier regressors bear this problem out as well. Although the detection of outliers can be extended to structural models, we may expect more difficulties to fit the model compared to the ARIMA model. In the ARIMA framework, the following results are obtained:

```
R> resNile2 <- tsoutliers(y = Nile, types = c("AO", "LS", "TC"),
+   maxit = 1, remove.method = "bottom-up", tsmethod = "auto.arima",
+   args.tsmethod = list(allowdrift = FALSE, ic = "bic"))
R> resNile2
```

Series: Nile

ARIMA(0,0,0) with non-zero mean

Coefficients:

	intercept	LS29	A043
	1097.7500	-242.2289	-399.5211
s.e.	22.6783	26.7793	120.8446

sigma² estimated as 14401: log likelihood=-620.65

AIC=1249.29 AICc=1249.71 BIC=1259.71

Outliers:

```

      type ind time coefhat tstat
1  LS  29 1899  -242.2 -9.045
2  AO  43 1913  -399.5 -3.306

```

A strong level shift is detected at observation 29 (year 1899), which is also suggested by the graphical representation of the series (not shown here). Curiously enough, the model selection procedure did not find any ARIMA structure when the level shift is detected and included in the model.

The model ARIMA(0,1,1)(0,1,1) has been found to fit well to many economic time series. This model is often called the airlines model since its application to the airlines passengers data is a popular example of the Box & Jenkins methodology. The outliers detection procedure can be run specifying beforehand the use of this or any other ARIMA model instead of running the model selection procedure. To do so, we set `tsmethod = "arima"` to indicate that an ARIMA model will be used. The ARIMA model is defined in the argument `args.tsmethod`, which is a list that may contain any of the arguments accepted by the familiar function `stats::arima`.

```

R> resAirP <- tsoutliers(y = log(AirPassengers), types = c("AO", "LS", "TC"),
+   maxit = 1, remove.method = "bottom-up", tsmethod = "arima",
+   args.tsmethod = list(order = c(0, 1, 1), seasonal = list(order = c(0, 1, 1))))
R> resAirP

```

```

Series: log(AirPassengers)
ARIMA(0,1,1)(0,1,1)[12]

```

Coefficients:

	ma1	sma1	A029	LS54	A062	A0135
	-0.3320	-0.4965	0.0959	-0.0967	-0.0803	-0.1032
s.e.	0.0958	0.0800	0.0230	0.0262	0.0227	0.0261

```

sigma^2 estimated as 0.001025: log likelihood=243.23
AIC=-472.47  AICc=-471.56  BIC=-452.34

```

Outliers:

```

      type ind   time coefhat tstat
1  AO  29 1951:05  0.09590  4.169
2  LS  54 1953:06 -0.09673 -3.689
3  AO  62 1954:02 -0.08032 -3.531
4  AO 135 1960:03 -0.10322 -3.948

```

6.2 Consumer price indices

Table 1 reports the best ARIMA model chosen by TRAMO both when the presence of potential outliers is considered and when it is omitted. The value of the Jarque-Bera test for normality [7] is also reported. Rejection of the null hypothesis of normality of residuals at the 5% significance level is indicated with an asterisk (the reference critical value is $\chi_{2,0.05}^2 = 5.99$).

[Table 1 about here.]

Including outlier regressors may lead to different ARIMA models and may especially affect the order of integration necessary to render stationary residuals. However, in these series we do not observe changes in the order of integration of the regular and seasonal part of the model. As regards the normality of residuals, we observe a significant change when outliers are taken into account. When outliers are omitted the null of normality of residuals is rejected in all cases except in the series ‘011300’. When the series is checked for the presence of outliers normality of residuals is achieved in 8 out of the 11 series where the null was rejected.

The same exercise is conducted in R. The automatic procedures for the selection of the ARIMA model and the detection of outliers implemented respectively in the packages `forecast` and `tsoutliers` are employed. The function `forecast::auto.arima` is run with the options `allowdrift = FALSE` and `ic = "bic"`, i.e., a drift is not included and the Bayesian information criterion is used; the arguments `maxit = 1` (the process is applied once on the original series) and `remove.method = "bottom-up"` are used in `tsoutliers`. Results are reported in Table 2. It is observed that including outliers increases the number of cases where normality is achieved, although in fewer cases compare to TRAMO. It seems that TRAMO tends to favour the choice of a seasonal differencing filter compared to `forecast`. Although not reported, when model selection is not performed and the model chosen by TRAMO is passed to the function `tsouliers`, normality is achieved in the same cases as those reported for TRAMO in Table 1.

[Table 2 about here.]

The outliers detected by each implementation are reported in Table 3.

[Table 3 about here.]

6.3 Industrial production indices

We may be interested in including other external regressors in the model. For example, calendar effects are often relevant when analysing industrial production indices. Regressors other than

outliers can be specified by means of the argument `xreg`. These regressors are included in the process of detection of outliers.

Calendar effects (by default trading day and Easter) can be defined by means of the function `calendar.effects`. See the documentation of this function for details.

Below, we show some of the options available in `tsoutliers`. We apply the procedure on the Italian industrial production index, `gipi`. This series is used by [8] to illustrate the relevance of seasonal level shifts. We work with the logarithms of the data, although in the reference paper it is not clear whether they use logarithms or levels. We specify the ARIMA(0,1,1)(0,1,1) model (the model chosen in the reference paper) and set the same critical value used in the reference paper, `cval = 3.5`.

```
R> data("bde9915")
R> gipi <- log(bde9915$gipi)
R> ce <- calendar.effects(gipi)
R> resGIPI1 <- tsoutliers(y = gipi, xreg = ce, cval = 3.5,
+   types = c("AO", "LS", "TC", "SLS"), maxit = 1, remove.method = "bottom-up",
+   tsmethod = "arima", args.tsmethod = list(order = c(0, 1, 1),
+     seasonal = list(order = c(0, 1, 1))))
R> resGIPI1
```

Series: `gipi`
ARIMA(0,1,1)(0,1,1)[12]

Coefficients:

	<code>ma1</code>	<code>sma1</code>	<code>trading-day</code>	<code>Easter</code>	<code>SLS28</code>	<code>A044</code>	<code>SLS49</code>	<code>SLS92</code>	<code>SLS164</code>
	-0.5756	-0.9265	0.0085	-0.0273	-0.0620	0.1068	-0.0606	0.0594	0.0704
s.e.	0.0603	0.1176	0.0005	0.0080	0.0156	0.0213	0.0121	0.0118	0.0141

`sigma^2` estimated as 0.0005178: `log likelihood=388.13`
AIC=-756.26 AICc=-754.95 BIC=-724.38

Outliers:

	<code>type</code>	<code>ind</code>	<code>time</code>	<code>coefhat</code>	<code>tstat</code>
1	SLS	28	1983:04	-0.06198	-3.975
2	AO	44	1984:08	0.10679	5.022
3	SLS	49	1985:01	-0.06062	-4.994
4	SLS	92	1988:08	0.05945	5.024
5	SLS	164	1994:08	0.07043	4.995

We see that calendar effects are significant and several SLS are found. The outliers detected in the reference paper are AO 1984:08, AO 1987:01, SLS 1994:08 and SLS 1988:08. Two more SLS

and one less AO are found with the options chosen above. It must be noticed that here we specified the airlines model beforehand while in the reference paper a model selection procedure is used. Hence, despite the airlines model is eventually chosen, different models may have been used in the procedure.

Running the procedure along with model selection the ARIMA(1,0,1)(2,1,1) is chosen and a SLS at 1985 : 01 is found.

```
R> resGIPI2 <- tsoutliers(y = gipi, xreg = ce, types = c("AO", "LS", "TC", "SLS"),
+   maxit = 1, remove.method = "bottom-up", tsmethod = "auto.arima",
+   args.tsmethod = list(allowdrift = FALSE, ic = "bic"))
R> resGIPI2
```

Warning message:

```
In auto.arima(x = c(4.45782959808938, 4.47278099794235, 4.56746831880408, :
```

```
Unable to fit final model using maximum likelihood. AIC value approximated
```

Series: gipi

```
ARIMA(1,0,1)(2,1,1)[12]
```

Coefficients:

	ar1	ma1	sar1	sar2	sma1	trading-day	Easter	SLS49
	0.9875	-0.6187	0.2762	-0.3598	-0.7191	0.0087	-0.0182	-0.0652
s.e.	0.0114	0.0573	0.0911	0.0831	0.0722	0.0007	0.0070	0.0178

```
sigma^2 estimated as 0.0006939: log likelihood=372.5
```

```
AIC=-726.99 AICc=-725.93 BIC=-698.26
```

Outliers:

	type	ind	time	coefhat	tstat
1	SLS	49	1985:01	-0.06517	-3.664

Using `remove.method = "en-masse"`, the model ARIMA(0,1,1)(0,0,2) was chosen and the a SLS at 1988 : 08 is found. Changing the method to discard outliers leads also to a different set of outliers with the ARIMA(0,1,1)(0,1,1) model.

```
R> resGIPI3 <- tsoutliers(y = gipi, xreg = ce, cval = 3.5,
+   types = c("AO", "LS", "TC", "SLS"), maxit = 1, remove.method = "en-masse",
+   tsmethod = "arima", args.tsmethod = list(order = c(0, 1, 1),
+     seasonal = list(order = c(0, 1, 1))))
R> resGIPI3
```

Series: gipi

```
ARIMA(0,1,1)(0,1,1)[12]
```

Coefficients:

	ma1	sma1	trading-day	Easter	A044	SLS49	A0176
	-0.5972	-0.6995	0.0089	-0.0200	0.0898	-0.0638	0.0963
s.e.	0.0600	0.0699	0.0005	0.0085	0.0224	0.0177	0.0236

sigma² estimated as 0.0007009: log likelihood=369.61

AIC=-723.23 AICc=-722.38 BIC=-697.73

Outliers:

	type	ind	time	coefhat	tstat
1	AO	44	1984:08	0.08984	4.010
2	SLS	49	1985:01	-0.06376	-3.592
3	AO	176	1995:08	0.09631	4.080

In the next example the number of seasonal differences is fixed to one, $D = 1$. Despite the same model as in `resGIPI2` shown above is chosen and the same method to remove outliers is used, here two additional SLS are found. The reason for that is that although the final model is the same in both cases, different models may probably have been considered during the process, in particular a different order of integration may have been used, leading to a different set of outliers obtained at each step.

```
R> resGIPI4 <- tsoutliers(y = gipi, xreg = ce, types = c("AO", "LS", "TC", "SLS"),
+   maxit = 1, remove.method = "bottom-up", tsmethod = "auto.arima",
+   args.tsmethod = list(allowdrift = FALSE, D = 1, ic = "bic"))
R> resGIPI4
```

Series: gipi

ARIMA(1,0,1)(2,1,1)[12]

Coefficients:

	ar1	ma1	sar1	sar2	sma1	trading-day	Easter	SLS49	SLS92	SLS176
	0.9957	-0.6041	0.2186	-0.3719	-0.8479	0.0084	-0.0188	-0.0600	0.0435	0.0773
s.e.	0.0066	0.0572	0.0866	0.0816	0.0799	0.0007	0.0068	0.0133	0.0119	0.0189

sigma² estimated as 0.0005625: log likelihood=386.29

AIC=-750.58 AICc=-749.01 BIC=-715.46

Outliers:

	type	ind	time	coefhat	tstat
1	SLS	49	1985:01	-0.06002	-4.500

2 SLS 92 1988:08 0.04349 3.668

3 SLS 176 1995:08 0.07729 4.092

Despite using different options we found some differences in the results, a SLS at 1985 : 01 was found in the four cases discussed above. This outlier seems to be therefore the most relevant or robust result. In the reference paper this outlier was not identified. Our examples may differ from the application shown in the reference paper in some issues such as the inclusion of a possible holidays variables or other issues not mentioned in the reference paper.

7 Conclusion

We have introduced the `tsoutliers` R package. The package implements an automatic procedure for the detection of outliers in time series. The original procedure proposed in the literature relies on the framework of ARIMA models. The package `tsoutliers` adapts the procedure to the context of structural time series models. Currently this extension is an experimental version that requires further work since the results were not always reliable with real data. The main difficulty is the estimation of the structural model, which should provide feasible parameter estimates and standard errors for the regressor variables.

The functions available in the package `tsoutliers` allows the user to do a manual run of each step of the procedure, i.e. not in automatic mode. Thus, the package is also useful to track and inspect the behaviour of the procedure and come up with ideas for possible improvements or enhancements.

The application of the `gipi` time series showed the flexibility of the `tsoutliers` interface to set and choose different options that may be relevant in practice. Those examples suggested also that pursuing a fully automatic procedure may become too cumbersome. Therefore, alternative manual runs of the automatic procedure are advisable when possible.

References

- [1] George E. P. Box and Gwilym M. Jenkins. *Time Series Analysis: Forecasting and Control*. San Francisco: Holden-Day, 1970.
- [2] Peter J. Brockwell and Richard A. Davis. *Introduction to Time Series and Forecasting*. Springer Texts in Statistics. Springer, 2002.

- [3] Chung Chen and Lon-Mu Liu. Joint Estimation of Model Parameters and Outlier Effects in Time Series. *Journal of the American Statistical Association*, 88(421):284–297, 1993.
- [4] James Durbin and Siem Jan Koopman. *Time Series Analysis by State Space Methods*. Oxford Statistical Science Series. Oxford University Press, 2001.
- [5] Andrew C. Harvey. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press, 1989.
- [6] Rob J. Hyndman and Yeasmin Khandakar. Automatic Time Series Forecasting: The `forecast` Package for R. *Journal of Statistical Software*, 27(3):1–22, 2008.
- [7] Carlos M. Jarque and Anil K. Bera. Efficient Test for Normality, Homoscedasticity and Serial Independence of Residuals. *Economic Letters*, 6(3):255–259, 1980.
- [8] Regina Kaiser and Agustín Maravall. *Measuring Business Cycles in Economic Time Series*. Springer Verlag, 2001.
- [9] Javier López-de-Lacalle. *stsm: Structural Time Series Models*, 2014. R package version 1.3.
- [10] D. S. G. Pollock. *A Handbook of Time-Series Analysis Signal Processing and Dynamics*. Academic Press, 1999.
- [11] Rob J Hyndman with contributions from George Athanasopoulos, Slava Razbash, Drew Schmidt, Zhenyu Zhou, Yousaf Khan, Christoph Bergmeir, and Earo Wang. *forecast: Forecasting Functions for Time Series and Linear Models*, 2014. R package version 5.3.

Table 1: Test for normality of residuals in best ARIMA models chosen by software TRAMO

Series	Omitting outliers		Including outliers	
	Model	JB	Model	JB
foodpr	(2, 1, 0)(1, 0, 1)	637.60 *	(1, 1, 1)(1, 0, 1)	33.99 *
igxe00	(0, 1, 2)(0, 1, 0)	354.30 *	(3, 1, 0)(0, 1, 1)	173.60 *
011600	(0, 1, 0)(0, 1, 1)	8.16 *	(3, 1, 1)(0, 1, 1)	0.39
011000	(1, 1, 1)(0, 1, 1)	10.65 *	(1, 1, 1)(0, 1, 1)	3.18
010000	(1, 1, 0)(1, 0, 0)	8.54 *	(2, 1, 0)(0, 1, 1)	3.35
011200	(1, 1, 0)(0, 1, 1)	28.12 *	(1, 1, 0)(0, 1, 1)	8.96 *
nrgy00	(0, 1, 1)(0, 0, 0)	19.09 *	(0, 1, 1)(0, 0, 0)	19.23 *
011700	(0, 1, 1)(0, 1, 1)	15.23 *	(0, 1, 0)(0, 1, 1)	0.67
foodun	(0, 1, 1)(1, 0, 0)	8.29 *	(0, 1, 1)(0, 1, 1)	0.21
011300	(0, 1, 1)(0, 1, 1)	0.05	(0, 1, 0)(0, 1, 1)	2.58
000000	(0, 1, 1)(0, 1, 1)	10.54 *	(0, 1, 1)(0, 1, 1)	1.76
serv00	(1, 1, 0)(0, 1, 1)	290.10 *	(3, 1, 1)(0, 1, 1)	2.66

Table 2: Test for normality of residuals in best ARIMA models chosen in R

Series	Omitting outliers ^a		Including outliers ^b	
	Model	JB	Model	JB
foodpr	(2, 1, 0)(0, 0, 1)	558.47 *	(1, 1, 1)(2, 0, 2)	38.64 *
igxe00	(2, 1, 3)(0, 1, 0)	127.02 *	(3, 1, 5)(0, 0, 2)	16.34 *
011600	(0, 1, 0)(2, 0, 1)	11.59 *	(1, 1, 0)(2, 0, 2)	0.16
011000	(1, 1, 0)(1, 0, 1)	18.65 *	(1, 1, 0)(1, 0, 1)	7.47 *
010000	(1, 1, 1)(2, 0, 2)	14.05 *	(2, 1, 0)(2, 0, 2)	3.46
011200	(1, 1, 0)(2, 0, 1)	60.15 *	(1, 1, 0)(2, 0, 1)	8.3 *
nrgy00	(1, 1, 0)(1, 0, 0)	20.39 *	(1, 1, 0)(1, 0, 0)	19.72 *
011700	(1, 1, 0)(1, 0, 1)	36.32 *	(1, 1, 0)(1, 0, 1)	1.59
foodun	(1, 1, 0)(2, 0, 2)	13.06 *	(1, 1, 0)(2, 0, 2)	0.52
011300	(1, 1, 0)(2, 0, 2)	4.15	(0, 1, 1)(0, 1, 1) ^c	3.88
000000	(0, 1, 1)(2, 1, 0)	10.27 *	(1, 1, 2)(1, 0, 1)	3.66
serv00	(1, 1, 2)(2, 1, 0)	359.06 *	(1, 1, 2)(2, 0, 0)	230.04 *

^aAutomatic model selection is carried out by means of the function `forecast::auto.arima` for the Bayesian information criterion.

^bOutliers are detected by means of function `tsoutliers` which in turn calls the function `forecast::auto.arima`.

^cThe Airlines model is used. No model selection is performed for this series since the default options failed when outliers were considered.

Table 3: Outliers detected in the HICP data set

Series	TRAMO ^a				tsoutliers ^b			
	Outliers				Outliers			
foodpr	AO145	LS157	LS171	LS180	TC145	LS157	LS171	TC180
	TC189	TC202	LS214		TC189	TC202	LS214	
igxe00	AO134	TC145	TC259		TC259			
011600	AO79	LS85	AO138	TC163	AO79	LS85	AO210	LS225
	AO210	LS225						
011000	TC85				TC85			
010000	AO23	TC85	TC145		AO23	AO85	TC145	
011200	AO28				LS29			
nrgy00	TC129				None			
011700	AO25	TC85	TC105		AO25	TC85	TC105	
foodun	TC145				AO23	AO85	TC145	
011300	AO120	TC157	TC162	LS171	AO120	TC162	LS171	
000000	TC19	LS215	LS231	AO235	TC19	LS215	LS231	AO236
serv00	TC7	AO16	TC132	LS144	AO16	TC19	AO219	AO280
	AO219	AO280						

^aContinuation of results reported in the fourth and fifth columns of Table 1.

^bContinuation of results reported in the fourth and fifth columns of Table 2.