Semi-automatic Configuration of Temporal Abstractions of Time Series Data Using CBR

G. Leonardi (*), A. Bottrighi, S. Montani and L. Portinale

(*) Universitá di Pavia, Pavia (Italy)
Universitá del Piemonte Orientale, Alessandria (Italy)

Abstract. In the haemodialysis domain, we are implementing a case-based architecture aimed at configuring Temporal Abstractions (TA) to be applied to time series data. The advantage of a case-based approach is the one of "quickly" obtaining a suitable parameter configuration, simply by looking at the most similar already configured case, where configured cases are indexed by means of contextual information. The obtained configuration, together with the time series data, is then used as an input to a TA processing module, able to provide a set of qualitative states, trends and significant combinations of both as an output. TA processing results can finally be analysed to evaluate the whole reasoning activity, possibly leading to a revision of the case base content. The main contribution of this work focuses on the configuration of TA and on the design of the processing module to be exploited in RHENE, a system for case-based retrieval in haemodialysis, able to work both on raw time series data and on pre-processed (by means of TA) ones.

1 Introduction

Parameter configuration is a critical issue in many Artificial Intelligence (AI) processes, especially when they are applied to complex domains like medical ones. Configuration usually demands for domain knowledge, which could be unavailable, or whose elicitation, exploitation and maintenance could be too time-consuming in practice. In several applications, however, the use of knowledge about the contextual situation under examination, together with the nature of the examined data, represents an appropriate means for parameter setting; the main difficulty in these cases is the one of selecting a criterion to find the most suitable configuration from a large number of possible ones. If we define a (prototypical) case as a set of $\langle feature, value \rangle$ pairs keeping contextual information, and store the suggested parameter configuration as the corresponding solution, Case-based Reasoning (CBR) [5, 1] can be resorted to in order to fulfill this task. The main advantage of a CBR approach obviously stands in the fact that the knowledge acquisition process for configuring parameters is made easier by the use of already configured cases, retrieved because similar to the current input situation [9, 7].

The Temporal Abstractions (TA) [10, 4] methodology is an AI process which requires a non-trivial parameter configuration phase, and which could benefit

from a CBR approach like the ones sketched above. TA are resorted to in order to map large amounts of temporal information, such as the ones embedded in a time series, to a compact representation, able not only to summarize the original data themselves, but also to abstract and highlight meaningful behaviors in them.

The basic principle of TA methods is to move from a *point-based* to an *interval-based* representation of the data, where: (i) the input points (*events* henceforth) are the elements of the discretized time series; (ii) the output intervals (*episodes* henceforth) aggregate adjacent events sharing a common behaviour, persistent over time. More precisely, the method described above should be referred to as *basic* TA [4]. Basic abstractions can be further subdivided into *state* TA and *trend* TA. *State* TA are used to extract episodes associated with *qualitative levels* of the monitored feature, e.g. low, normal, high values; *trend* TA are exploited to detect specific *patterns*, such as increase, decrease or stationarity, from the time series. *Complex* TA [4], on the other hand, aggregate two series of episodes into a set of higher level episodes (i.e., they abstract output intervals over pre-calculated input intervals). In particular, they are used to search for specific *temporal* relationships between episodes that can be generated from a basic abstraction or from other complex abstractions. The relation between time intervals can be any of the temporal relations defined by Allen [2]. This kind of TA can be exploited to extract patterns that depend on the course of several features, or to detect patterns of complex shapes (e.g. a peak) in a single feature.

In this work we propose a case-based architecture for parameter configuration of TA on time series data, in the domain of End Stage Renal Disease (ESRD). ESRD is a severe chronic condition that corresponds to the final stage of kidney failure. Without medical intervention, ESRD leads to death. Haemodialysis is the most widely used treatment method for ESRD; it relies on an electromechanical device, called haemodialyzer, which, thanks to an extra-corporal blood circuit, is able to clear the patient's blood from catabolites, to re-establish acid-base equilibrium and to remove water in excess. On average, haemodialysis patients are treated for four hours three times a week. Each single treatment is called a haemodialysis session. Haemodialyzers typically allow users to collect several variables during a session, most of which are in the form of time series.

In the last several years we have developed a system, called RHENE, able to retrieve haemodialysis cases, which can then be visually analyzed by physicians, or provided as an input to an automatic reasoner (for additional details, see [8]). In particular, RHENE originally worked on raw time series retrieval (just taking advantage of mathematical dimensionality reduction techniques, such as the Discrete Fourier Transform). We are currently planning to extend it in the direction of enabling the retrieval of pre-processed time series, where pre-processing will be realized by means of TA - a step which justifies the work described in this paper, even though the methodology we present here is general enough to be resorted to by a (stand-alone) TA tool applicable to very different domains.

Since TA are used here as a dimensionality reduction technique to search for similar cases, this approach has been defined to be used in systems working

off-line. In the case of RHENE, for example, all the time series belonging to the (already collected) haemodialysis sessions will be pre-processed to obtain cases whose features are reduced by means of TA. As a future work, the retrieval of similar cases will be then performed off-line, using distance metrics based on interval similarity (please refer to [14], [13]). Please note that, even if the architecture has been defined as a "black box" to be used in off-line system, the retrieval of the configuration of the parameters defining the abstractions to be searched for in a signal can be useful also in real-time systems, for example to be used in intensive-care units to generate alerts as soon a particular state of the signal is reached, or when a pattern representing a critical situation is found. In this case, the configuration module can be used as is, while it is necessary to restructure the TA processing module to work with real-time data, rather than performing the transformation from time series data into TA in a single step.

The main contribution of this paper focuses on the description of a prototypical context and on the configuration of the associated TA, and the presentation of the design of the TA processing module, which will be integrated in RHENE. This work is built on the basis of the formal definitions of context and configuration, which can be found in [9]. Also, in [9], the reader can find the formal definition and description of the TA, of their intervals, of how the complex TA are built on the basis of the simple TA discovered by the TA process, and the relations between the intervals.

The rest of the paper is organized as follows. In section 2 we introduce our case representation details, and we sketch the TA tool architecture; in section 2.1 we then describe the functionality of the TA processing module, along with an example, while in section 2.2 we analyse TA-based case revision. Finally section 3 concludes the paper.

2 Case-based Reasoning for TA configuration and processing

In TA processing, parameter configuration is a very important phase because the TA output results heavily depend on the value assigned to specific parameters. In particular, for trend TA, the following (main) parameters need to be set [3]: Minimum/Maximum Rate (i.e. the minimum and maximum slope allowed for the trend episode); Minimum/Maximum Duration (i.e. the minimum and the maximum duration in time for the trend episode). As regards state TA, on the other hand, we need to specify [3]: Lower/Upper Bound (i.e. the lower and upper bounds of data values allowed for the state episode); Minimum/Maximum Duration (defined as above).

In our application, a case stored in the case base is defined as follows [9]: (i) problem description: the *context description*, composed by patient and haemodialysis session characteristics which tend to be stable in the long/medium run (such as patient's age and session duration); (ii) case solution: the *configuration* of the various signals (i.e. of the time series variables collected by the haemodialyzer). In turn, the configuration of each signal consists of a list of state and trend TA

symbols to be searched for in the time series to which the configuration refers, together with the corresponding parameter values, and optionally, of a list of suitable combinations of the obtained states and trends, known as *joint TA*. A joint TA is actually a special case of a complex TA; indeed if t is an instance of a trend having validity in the time interval I_t and s is an instance of a state having validity in the time interval I_s, then $j = \langle t, s \rangle$ is a complex TA based on the Allen's relation t R s in the time interval $I_j = I_t \cap I_s$, where R is any of the following: *overlaps, during, starts, finishes, equal* and their inverse relations [2]. If the list of joint TA is empty, all $\langle trend, state \rangle$ pairs will be calculated.

The initial Case Base is set up with prototypical cases derived from the medical knowledge provided by a specialist. The specialist can also suggest a maximum distance threshold to be verified when a case is retrieved.

An input case contains, together with the *context description*, a set of raw time series, instances of the signals on which TA must be extracted.

The system we are implementing is conceived as a two-module architecture, composed as follows (see the dashed boxes on a gray background in figure 1):

- a module for TA parameter *configuration*;
- a module for TA *processing*.

Given an input case, the TA parameter configuration module, which was already described in [9], retrieves the less distant (i.e. most similar) case (without loss of generality), with respect to the input case context description.

The similarity distance is evaluated using the HEOM distance function described in [12].

If the distance between input case and retrieved case exceeds the given threshold, the user can stop the consultation, or he can allow the system to process the time series using the retrieved configuration. On the basis of the results, the Case Base can be revised as explained in Section 2.2.

The retrieved *configuration* information, corresponding to the signals present in the input case, is extracted and passed to the TA processing module, together with the raw data. The TA processing module, whose implementation represents one of the original contributions of this work, provides a set of qualitative states, trends and suitable combinations of both as a result, and is extensively described in section 2.1.

The processing of a time series exploiting a given configuration may provide a significant feedback for the possible revision of the used configuration itself. In particular, the resulting TA series can be relied upon to guide case base maintenance, by suggesting the elimination of non-representative cases, and by keeping the knowledge embedded in the case library always up to date. As described in figure 1, the two modules thus give birth to a closed loop architecture, where parameter configurations suggested by the case-based retrieval module are adopted for TA processing, while the obtained TA series are evaluated to support case base maintenance. Details about case base maintenance are provided in section 2.2.

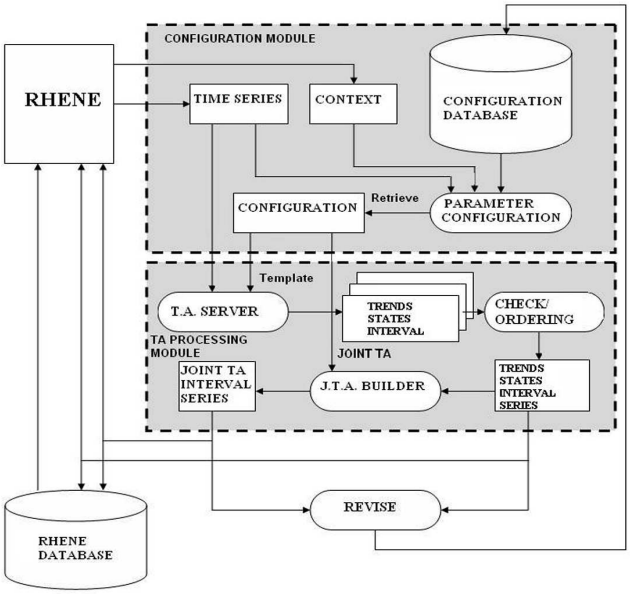

Fig. 1. Case-based configuration of TA parameters, and TA processing. The dashed boxes on a gray background represent the already implemented modules, while the boxes on a white background depict future extensions, meant to realize the integration between the TA tool and the RHENE system [8]

2.1 TA processing module

The aim of the TA processing module is to obtain a TA time series, given a raw time series with its corresponding configuration, which is provided by the configuration module (for the sake of simplicity, we are considering to deal with only one signal). The process is articulated in two steps. In the first step, the *TA Server* (figure 1) identifies the instances of the allowed trends and states. In our architecture, the search of predefined patterns (both trend and state) in the time series is implemented by using the TA web-service described in [3]. The web-service takes as input the retrieved configuration, represented in XML format, together with the raw data to be abstracted. The output is a set of XML documents, one for each searched pattern (trend or state). Each document contains the instances of the pattern found in the data (i.e. the time intervals in which the pattern has been found). In the second step, the *Check/Ordering* submodule in figure 1 collects all XML documents, which are generated by *TA Server*, and creates an ordered series of trend and state instances. This submodule manages two kind of situations: **gaps** and **overlaps**. A gap is a time interval (in the original data) where no instance of an allowed trend/state has been found. The *Check/Ordering* submodule creates an instance of a special symbol UT (Unknown Trend) or US (Unknown State) each time a gap is found for a trend or a state respectively. Overlaps occur when two or more instances of allowed trend/state cover the same time interval. In our approach, we allow the existence of partial overlaps; in particular, just situations involving the following Allen's relations are not allowed: *during, starts, finishes, equals* and their inverse relations. They represent situations of "temporal overlap" among intervals that are considered as not consistent and caused by errors in the specification of configuration parameters (for instance an *Increasing Trend* during a *Stable Trend*). On the other hand, the Allen's relation *overlaps* (and its inverse *overlapped-by*) is allowed, only if the intersection interval does not exceed a given threshold. If a configuration problem is detected, the user is asked to revise the parameters of the involved trends or states. If the analysis is considered as valid, the *Check/Ordering* submodule builds the sequence of the instances of the found patterns.

Joint TA are obtained by the *JTA Builder* submodule in figure 1. This submodule takes as input the ordered series of instances obtained so far and the joint template of the retrieved configurations; it then builds all the pairs of trend and state instances produced by the *TA Server*, which are allowed by the joint template. Each pair is associated with a time interval corresponding to the intersection of the time intervals of the two basic instances composing the joint TA.

An example: in order to illustrate TA configuration and processing, we will now show an example in which we will abstract two time series describing the haematic volume (HV) and diastolic pressure (DP) behaviour of a patient during an haemodialysis session, and we will show the different results obtained by abstracting the same time series with a slightly modified configuration.

The behaviour of the HV is extremely important, because it is strictly related to the water reduction from the patient's blood during the haemodialysis session. The correct behaviour of this parameter is composed by two phases. In the first phase the session starts, and the water is extracted from the blood at high speed, until the blood pressure equals the intra-cellular pressure of the cells in the blood. In the second phase, the water passes from the cells to the blood, and is again extracted from the blood by the haemodialyzer. During this second phase the HV decrease is stable and constant. Therefore, the correct HV behaviour is composed by a first exponential decrease followed by a linear decrease until the end of the session.

Furthermore, the reduction of the water from the blood causes a constant decrease of the blood pressure, during the haemodialysis session. This behaviour is correct and, even if it can sometimes cause minor problems to the patient (e.g. light head spinning), it is necessary to achieve a good water and metabolites reduction.

The behaviour explained before can be altered if the patient is affected by particular health problems. For example, hypotension can cause problems in both HV and DP. In particular, considering the HV, hypotension introduces two kinds of problems: the low pressure can alter the first phase, slowing down the water reduction speed and generating an inefficient linear decrease; moreover, it can cause instability and difficulties in the water extraction process. These problems can also force the health care professional to interrupt the haemodialysis session before the usual four hours duration. Looking at the DP, the patient could suffer seriously if the haemodialysis session starts with a low DP value, or the hypotension causes sudden falls of the DP. In this situation, the subsequent decreasing trend, leading to a critical low value of the pressure, or the occurrence of a sudden drop episode can expose the patient to the risk of collapse, forcing the health care professional to give him anti-hypotension drugs.

A prototypical case describing the hypotensive context and the configuration of the parameters for this kind of analysis can be defined as follows. Please consider that this description is intentionally simplified, for brevity and to keep this example clear. More complex contexts can be defined using additional features, such as the sex of the patient, her weight before and/or after the haemodialysis session, an optional pharmacological treatment, the outcome of the session, and many more.

- Context description:
  - systolic pressure under 110 mmHg, or;
  - diastolic pressure under 60 mmHg;
  - session duration: 3 hours and 30 minutes;
  - age range: 64 years and over.

- Configuration (for the HV signal):
  - Expected Trends:
    - EXD = EXponential Decrease;
    - LD = Linear Decrease;

    * FD = Fall Decrease;
    * ST = Stationary;
    * INC = Increase;
  - Expected States:
    * PS = Positive State (HV increases, while in this state, which is dangerous);
    * NS = Negative State (HV decreases, which is correct);
  - Expected Joint Symbols: no symbols specified, therefore all pairs $\langle trend, state \rangle$ are allowed.

- Configuration (for the DP signal):
  - Expected Trends:
    * SI = Strong Increase;
    * SD = Strong Decrease;
    * ST = Stationary;
  - Expected States:
    * Normal = Normal blood pressure;
    * Hypo = Hypotension (usually tolerated by the patient);
    * SH = Strong Hypotension (with the risk of collapse);
  - Expected Joint Symbols: no symbols specified, therefore all pairs $\langle trend, state \rangle$ are allowed.

As an example, we show the definition of the main parameters of the EXD trend and NS state.

- EXD trend:
  - SymbolTrend = EXD
  - MinimumRate = 0.150 pts/min
  - MaximumRate = 0.250 pts/min
  - MinimumDuration = 10 min
  - MaximumDuration = no bound

- NS state:
  - SymbolState = NS
  - LowerBound = - infinite
  - UpperBound = 0
  - MinimumDuration = 10 min
  - MaximumDuration = no bound

We can now apply the configuration defined to analyse a case representing the haemodialysis session of a patient suffering from hypotension.

This case reports the session 126 of patient 40. Here, considering the DP signal, we can see that the blood pressure is stable and the state is Normal for the entire session. This means that the patient, despite her health conditions, did not suffer for problems related to the blood pressure. Looking at the HV, we can see that the plot starts with an exponential decrease, followed by a linear decrease for the rest of the session. This is the correct behaviour of the signal, as explained

before. We can also observe a partial overlap between the first and the second trend (exponential decrease and linear decrease, respectively). This overlap is due to the fact that the signal changes its slope smoothly, therefore a part of the signal is interpreted both as exponential decrease and linear decrease. We expected this situation, therefore we specified a proper threshold for the overlap of these two trends. In this case the threshold has been respected, therefore we obtained a good quality output from the TA process, and the sequence of trends and states highlights the good quality of the haemodialysis session. The plots of the DP and HV for the patient 40, session 126 can be seen in figures 2 and 3, respectively. Figures 4 and 5 show the legends for the interpretation of trends and states.



Fig. 2. Diastolic pressure of patient 40, session 126

The analysis of the abstraction obtained for these two signals, produces this result:

- *Gaps and overlaps*: there are no gaps in the DP signal, and no gaps in the HV signal. A partial overlap is found in the trend sequence of the HV signal, but the threshold for this overlap has been respected, therefore there are no critical issues;
- *Missing abstractions*: the analysis found missing trends in either DP and HV. In particular, the SI and SD trends, together with the Hypo and SH states have not been found in the abstraction process of the DP, while the FD, ST and INC states and the NS state have not been found in the abstraction process of the HV.

Considering the behaviour of these two signals, we can say that the absence of the states listed before is not an error, because in a good haemodialysis session

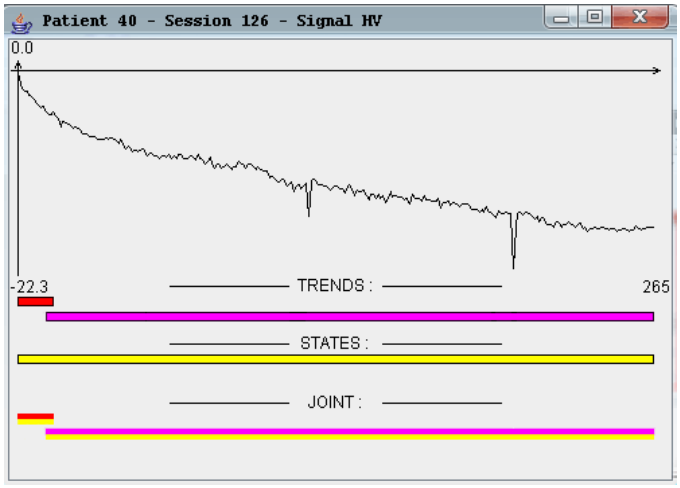

Fig. 3. Haematic volume of patient 40, session 126

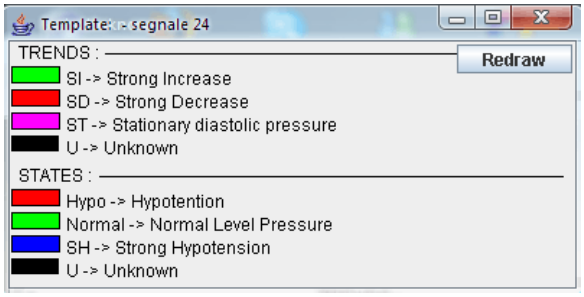

Fig. 4. Legend for diastolic pressure

we expect that the state of the DP remains Normal and the state of the HV must be Negative for the whole session. A different discussion may be carried out concerning the missing trends: looking at the plot of the DP, we can see clearly that increasing and decreasing trends are present in the signal, but they have been included in a single Stable trend. The HV signal shows at least two evident spikes, included in the trend of Linear Decrease. Finally, even if the threshold has been respected, we can find a way to avoit the overlap in the HV. This means that we made at least three mistakes in defining the parameters:

- the minimum duration for the SI and SD trends configured for the DP signal is too high. The episodes in the signal are shorter than the 10 minutes set in the configuration;
- the parameters set for the elaboration of the signal (not listed before, for simplicity, but described in [9]) force the TA process to strongly smooth the signal during the elaboration of the signal. Because of this, the spikes in the

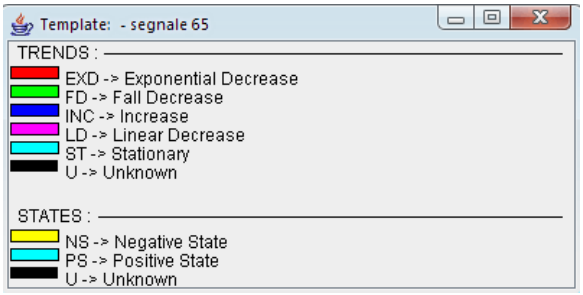

Fig. 5. Legend for haematic volume

HV and the increasing and decreasing trends in the DP are smoothed and considered as part of the Linear Decrease and Stable trends, respectively;
- the minimum rate and maximum rate set for the EXD trend must be restricted to obtain a better identification of this trend.

The configuration is then updated with these modifications:

- the minimum duration for the SI and SD trends is set in 3 minutes;
- the parameters set for the elaboration of the signal has been modified, for example the parameter LocalWindow (see [9] for details) for all the trends of the DP and HV has been reduced to avoid an excessive smoothing;
- the minimum rate and maximum rate for the EXD has been set to 0.190 and 2.70 pts/min, respectively.

The output of the TA process with the new configuration is shown in fig. 6, and 7. Here, we can see that two Decreasing trends and an Increasing trend have been correctly identified in the DP signal, while in the HV the overlap has been almost completely eliminated. Furthermore, this signal presents two intervals where no trends have been identified. These Unknown intervals correspond to the two spikes which have been previously smoothed and included in the Linear Decrease. This new outcome can be considered improved, compared to the first outcome considered in this example, therefore we are confident that the TA process carried out a finer treatment of the signal.

2.2 Revise and case base maintenance

The output of the TA processing module can be used to revise the knowledge contained in the case base. Case-based revision is an important process in the whole CBR cycle, since it may affect the actual performance of a CBR-based system (see [6] for a survey on the possible case base maintenance policies).

In particular, in our approach we identify cases which are candidate for revisions by investigating:

1. the frequency of the each suggested abstraction in the input situations;

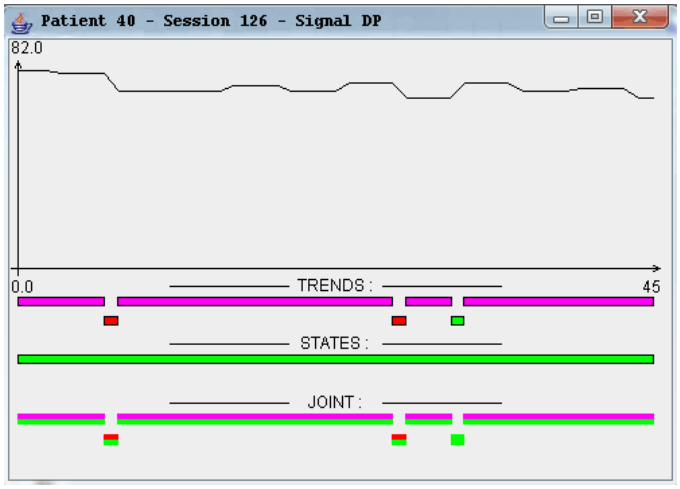

Fig. 6. Diastolic pressure of patient 40, session 126 with new parameters

2. the quality of the extracted TA series.

Frequency of occurrence is measured as follows. Given an input context, its best match is retrieved from the case base. For each input time series, the TA listed in the case solution are extracted, by applying the suggested parameter configurations. Each time a TA allowed by the retrieved case (e.g. *Increasing Trend*) is actually identified in the form of some episodes in the input data, a related counter variable is incremented. A very low value of a counter variable corresponds to an abstraction which is rarely identified in practice. Observing it suggests that the prototypical configuration might be inappropriate as regards that particular TA. We manage these situations by *shadowing* the TA at hand: whenever the case is retrieved again, the shadowed TA is not calculated anymore on the new input data.

While the frequency of occurrence of symbols (issue 1) allows to identify a possibly badly defined prototype, an insufficient quality of the extracted TA series (issue 2) may emerge due to two reasons: the input data are strongly affected by noise, or the prototypical case needs to be revised. Supposing that noisy data have been correctly filtered, we should further distinguish between two situations: (i) the retrieved case was the most similar to the input context; nevertheless, their distance was higher than a proper threshold, defined on the basis of medical knowledge; (ii) the retrieved case and the input context have indeed a very high similarity. In the first situation, we have identified a competence gap region in the case base, since the input context is not well represented by any of the items in the library. The solution is thus the one of acquiring a new case from scratch. In the second situation, we proceed as follows:

– if the output TA series suffer from the presence of the *unknown* symbol (see section 2.1) over long time intervals and some pertinent TA symbol had

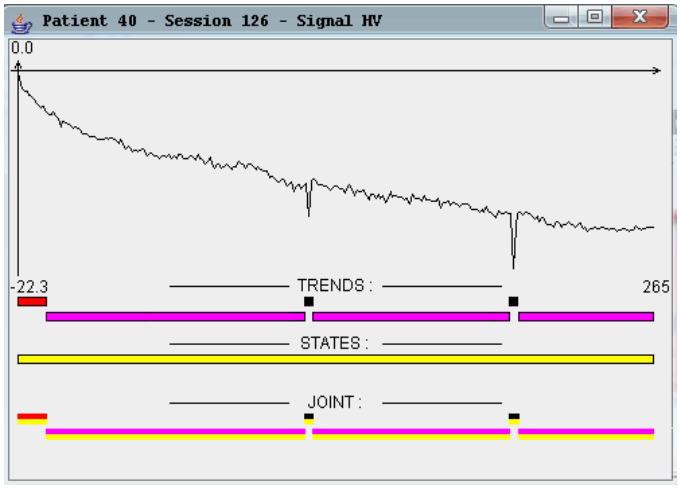

Fig. 7. Haematic volume of patient 40, session 126 with new parameters

been shadowed in the prototype, we re-admit the symbol into the list of the abstractions to be calculated, and process the input time series again;
- if the output quality does not improve, or if the suggestion above could not be implemented (because no shadowed symbol was present, or because the output quality was low due to the presence of overlaps of different TA), the retrieved case has to be carefully revised, in search of conflicting or improper parameter settings that might have been originally introduced.

In our approach, case base maintenance, whose implementation is still a work in progress, is thus conceived as a semi-automatic procedure, to be always supervised by a domain expert. However, such a partially data-driven approach is very appealing in the haemodialysis domain, where a well established knowledge about $\langle context, configuration \rangle$ pairs does not exist.

In other contexts, where the domain knowledge can be fully exploited and formalized a-priori, this approach does not need to be supervised by a domain expert, therefore an automatic system can be built on the basis of this architecture.

3 Conclusions

In this paper, we have proposed a case-based architecture tackling the problem of configuring and processing TA obtained from raw time series data, which is being implemented in the haemodialysis domain. The CBR approach does not require an explicit domain model and avoids the need of defining the right configuration for each possible contextual situation to be handled. Moreover, a CBR system can learn new knowledge by acquiring new cases or by revising cases

which are already stored, on the basis of a detailed evaluation of the problem solving activity (TA processing in our case).

Our work appears to be innovative in the literature panorama for at least two main reasons: (i) it is one of the very few contributions in which automatic knowledge acquisition techniques (such as CBR) are applied [11], with the aim of adopting the acquired knowledge for TA processing; (ii) it moves in the direction of integrating TA processing and data analysis in a more complex architecture, in which the challenges offered by the medical domain can be better afforded [7, 11].

The implementation of the TA tool described in this paper (see the dashed boxes on a gray background in figure 1) is being quickly completed. As a subsequent step, we will realize the integration of the tool itself with the system RHENE, to which we will provide the capability of retrieving cases whose features are time series pre-processed by means of TA (see the modules on a white background in figure 1). The overall work will be followed by a testing phase on real patients' data. To this end, we plan to provide the service to the physicians of the Nephrology and Dialysis Unit of the Vigevano Hospital in Italy.

References

1. A. Aamodt and E. Plaza. Case-based reasoning: foundational issues, methodological variations and systems approaches. *AI Communications*, 7:39–59, 1994.
2. J.F. Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23:123–154, 1984.
3. R. Bellazzi, C. Larizza, and G. Lanzola. An http-based server for temporal abstractions. In *Proc. IDAMAP'99*, pages 52–62, 1999.
4. R. Bellazzi, C. Larizza, and A. Riva. Temporal abstractions for interpreting diabetic patients monitoring data. *Intelligent Data Analysis*, 2:97–122, 1998.
5. J.L. Kolodner. *Case-Based Reasoning*. Morgan Kaufmann, San Mateo, CA, 1993.
6. D.B. Leake, B. Smyth, D.C. Wilson, and Q. Yang(eds). Special issue on maintaining case based reasoning systems. *Computational Intelligence*, 17(2):193–398, 2001.
7. S. Montani. Exploring new roles for case-based reasoning in heterogeneous ai systems for medical decision support. *Applied Intelligence (in press)*.
8. S. Montani and L. Portinale. Accounting for the temporal dimension in case-based retrieval: a framework for medical applications. *Computational Intelligence*, 22:208–223, 2006.
9. L. Portinale, S. Montani, A. Bottrighi, G. Leonardi, and J. Juarez. A case-based architecture for temporal abstraction configuration and processing. In *Proc. IEEE International Conference on Tools with Artificial Intelligent (ICTAI)*, page 667–674. IEEE Computer Society, Los Alamitos, 2006.
10. Y. Shahar. A framework for knowledge-based temporal abstractions. *Artificial Intelligence*, 90:79–133, 1997.
11. M. Stacey and C. McGregor. Temporal abstraction in intelligent clinical data analysis: a survey. *Artificial Intelligence in Medicine*, 39:1–24, 2007.
12. D.R. Wilson and T.R. Martinez. Improved heterogeneous distance functions. *Journal of Artificial Intelligence Research*, 6:1–34, 1997.

13. R . Schmidt , L . Gierl. A prognostic model for temporal courses that combines temporal abstraction and case-based reasoning. *International Journal of Medical Informatics*, 74(2-4):307-315, 2005.
14. P. Perner. Introduction to Case-Based Reasoning for Signals and Images. *Case-Based Reasoning on Signals and Images*, Ed. P. Perner, Springer-Verlag, Berlin, 2007.