Geophys. J. Int. (1989) **98**, 223–232

Artificial intelligence techniques in seismic signal interpretation

C. Chiaruttini
Istituto di Geodesia e Geofisica, Università di Trieste, Via Università 7, 34123 Trieste, Italy

V. Roberto and F. Saitta
Dipartimento di Matematica e Informatica, Università di Udine, Via Zanon 6, 33100 Udine, Italy

Accepted 1989 January 23. Received 1988 October 5; in original form 1988 May 31

SUMMARY
We show that system architectures based on artificial intelligence (AI) techniques can be applied to the automation of seismic signal interpretation. In particular, *blackboard systems* appear to provide flexible and efficient ways to model the strategy a human expert adopts for the analysis. We also present the design of the Seismic Network Analyser, (SNA), a blackboard system applied to the interpretation of signals from a local seismometric network. Results, examples and performances of a prototype implementation of SNA are reported and extensively discussed.

Key words: knowledge-based systems, automatic interpretation, seismic networks, signal understanding, expert systems, artificial intelligence

1 INTRODUCTION

Seismic networks are the main source of information in seismology. Much effort has been devoted so far to the automation of signal acquisition, specific steps in the analysis, and preliminary interpretation. As a result, many tasks are now successfully accomplished by means of numerical packages (Allen 1982; Blandford 1982). However, the goal of automating the whole routine seismological analysis (Keilis-Borok 1972) is far from achieved, and seismogram interpretation is still a bottleneck in the routine work of observatories, as the skill of trained analysts is required to recognize seismic waveforms. Furthermore, the interpretation is to some extent subjective: different analysts may give different interpretations, and the interpretation of the same analyst may vary with time. Therefore the use of automatic systems for routine analysis can guarantee a higher degree of homogeneity in seismic databases.

2 ARTIFICIAL INTELLIGENCE APPROACH TO SIGNAL INTERPRETATION

We address our attention to the control schemes which have been used so far in the analysis programs. All of them exhibit a rigid, pre-defined control flow, according to the technique known as procedural programming: default steps in a program are sequential, branchings are explicitly coded (as if-then-else constructs, for instance) and cannot be modified at run-time.

This poses severe limitations when the data to be processed show highly variable patterns, are incomplete and/or affected by a substantial noise level. Then, branchings are the rule, not the exception, and a program must exhibit enough flexibility to handle a quantity of unexpected situations. In these cases (i.e. when our knowledge is uncertain or incomplete) it is desirable for a program to adapt its strategy based on the actual data and to integrate useful contextual information, simulating the human expert analyst when doing his/her interpretive work. In addition, modelling the control flow in a non-procedural way helps to avoid an inappropriate use of numerical modules, postponing expensive computations to the stage in which they are actually needed.

Artificial intelligence (AI) (Barr & Feigenbaum 1982) provides a potentially useful approach to signal understanding. We shall now give an outline of the most relevant techniques to be applied in seismic network signal interpretation.

3 KNOWLEDGE-BASED SYSTEMS

In a previous paper (Chiaruttini & Roberto 1988a) we analysed the objects, strategies and reasoning techniques involved in the analyst's job. We identified two distinct aspects of the expert knowledge: a *domain knowledge*, concerning all basic notions on seismic wave generation, propagation and detection; and a *problem-solving knowledge*, including the ways to organize and use the domain knowledge when interpreting data. This twofold character of knowledge is closely reflected in a software architecture, named knowledge-based system (KBS), which has been proposed in the AI literature (see for example Frost 1986). A typical KBS architecture is sketched in Fig. 1.

Three basic modules are present:

(a) *User interface*. Two major tasks are accomplished by

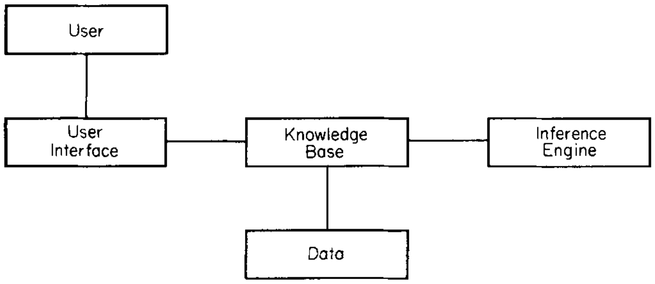

Figure 1. Basic scheme of a knowledge-based system.

this module:

- User-system dialogue;
- Justification: the path to the interpretation is traced back, and the reasoning chain is expressed as a set of strings in natural language.

(b) *Knowledge base.* A collection of data and symbol structures to include the necessary domain knowledge, and to keep track of the status of computation.

(c) *Inference engine.* The encoding of what has been called problem-solving knowledge. The purpose of this module is to select the currently relevant fragments of knowledge, i.e. those who match the current status of the knowledge base, and execute the actions suggested by the selected fragments. This, in turn, has the effect of modifying the knowledge base itself, thus triggering a new selection-execution cycle.

In order to meet our need of a flexible signal understanding system, two further capabilities are required of a KBS.

First, numerical signal processing modules are necessary and must be suitably interfaced to non-numerical ones (i.e. those implementing logical operations like inferences, associations, decisions).

Second, suitable focus-of-attention mechanisms must be provided in order to assure acceptable performances. This means, for example, that spurious or noisy tracks have to be quickly recognized as such, and excluded from the subsequent analysis tasks. Stated more generally, the system has to gain as soon as possible a rough idea of the information content of the data (if any), and consequently choose the most appropriate strategy of analysis.

A specialization of the KBS non-procedural programming scheme, known as blackboard system (Nii 1986), has been proposed for similar signal understanding goals (Nii & Feigenbaum 1978; Erman *et al.* 1980; Nii *et al.* 1982).

4 BLACKBOARD SYSTEMS

A blackboard system shows a number of distinctive features. Besides the user interface module, there are two major components:

(a) *Knowledge sources* (KS). Separate and independent programs performing specialized tasks. The aim of a KS is to provide information which contributes to the construction of an event interpretation. Each KS is not 'called' by another program, but knows its own activation conditions, and is self-activating.

(b) *Blackboard data structure.* This is the system *global database,* sometimes divided into a permanent and a volatile substructure. The latter, also named *working memory,* contains all information concerning the current status of the computation. It may be in turn divided into a *current best hypothesis* (CBH), reporting the current interpretation of the data, a *problem list,* with all the problems the system is not currently able to solve, and an *agenda,* containing all KS whose activation conditions currently fit the database. It should be observed that the KS produce changes on the blackboard, and only on it; in addition, interactions among KS take place solely through changes on the blackboard.

As far as control is concerned, no explicit control structures are coded in a pure blackboard system. The computation flows via successive KS self-activations: a KS responds to a situation on the database; its activation, in turn, modifies the blackboard, thus allowing a new KS to be activated as a consequence according to the scheme outlined in the previous section.

An interesting situation arises when more than one KS matches the blackboard configuration. Since on a sequential machine only one KS at a time can be executed, one must define suitable criteria to select a KS among the set of potentially active ones. A *Control module* may be added to perform this task, as well as any kind of knowledge application strategy which is peculiar of the problem to be solved.

A few observations are of interest to our signal understanding goals. A blackboard system, being a collection of specialized modules, naturally includes all data-processing tasks, which cooperate to construct an event interpretation in an incremental way. Furthermore, a high degree of modularity is desirable in order to guarantee readability, modifiability, expandability to the system itself. Focus-of-attention strategies can be embedded in a blackboard architecture in a number of ways: for example, scheduling the KS execution by means of a suitable control module.

On this basis we designed the Seismic Network Analyser (SNA), a KBS for the interpretation of data from a local seismometric network.

5 SNA: DESIGN AND IMPLEMENTATION

5.1 Architecture

SNA is composed of four basic units, whose functions are now briefly discussed. The structural scheme is shown in Fig. 2.

(a) *User interface.* It is intended for user-system dialogue and justification.

(b) *Permanent database.* It is made up of:

- *Facts*: encode the static knowledge about the domain (e.g. the coordinates of the stations and the velocities of seismic phases);
- *Rules*: implement the basic inferential processes in the form of condition-action pairs;
- *Knowledge sources*: are the true problem-solving units in SNA.

At the data level the KS extract features and transform numbers into symbols; at the symbolic level, they build or modify the current hypothesis. In particular, at this level the

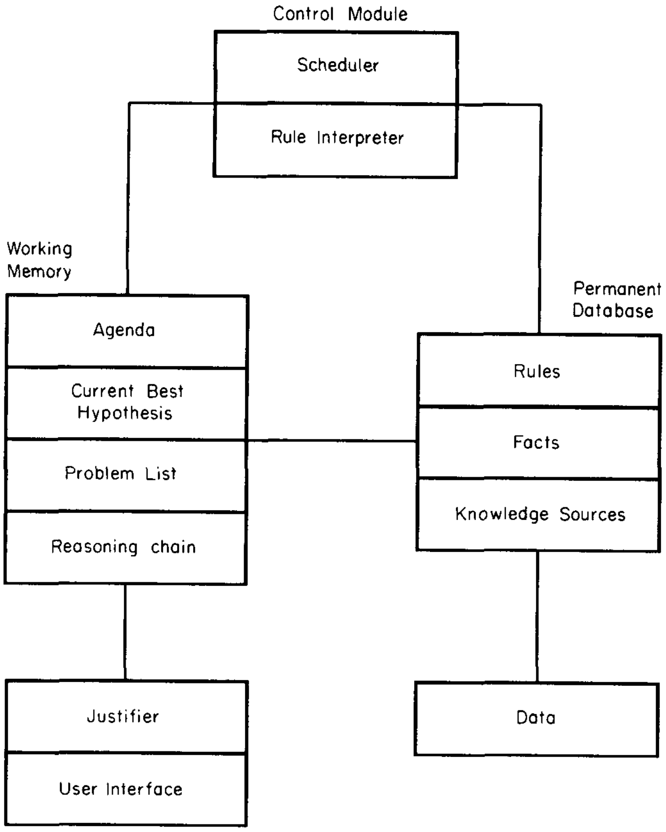

Figure 2. Architecture of SNA.

KS perform three kinds of activities: evaluation and solution of pending problems; validation of hypothesis elements; increment of the hypothesis.

From a structural point of view, a KS is made up of the following subunits: *Activator*: logical unit to check the KS activation conditions against data in the blackboard; *Executor*: logical unit to perform the KS problem-solving task; *Feature-extractor*: numerical processing unit, working at the signal level; *Translator*: interface unit between numerical and logical processing modules. The first two units are present in all the KS, the last two are present in the KS that manipulate the data.

(c) *Working memory.* A data and symbol structure composed of:

- *Agenda*: keeps track of all the KS whose pre-conditions match those of the current situation;
- *Current best hypothesis* (CBH): contains the current state of the interpretation and is structured according to an object hierarchy (Chiaruttini & Roberto 1988b), corresponding to the semantically relevant concepts in the expert analysis: the *event,* the *trace,* the *seismogram,* the *noise,* and the *phase*;
- *Problem list*: contains unresolved situations to be subsequently analysed;
- *Reasoning chain*: a list of explanations in natural language of the inferences made by the system; also this data structure reflects the aforementioned object hierarchy.

(d) *Control module.* It is made up of:

- *Rule interpreter* acting at the rule level, to select, match and execute single rule prescriptions;
- *Scheduler,* which is the priority selector for a KS to be executed.

Stated in a different way, our architecture may be seen as a hybrid problem-solving paradigm, made up of two basic kinds of activities: the numerical processing and the symbolic reasoning. We shall briefly analyse both in the following sections.

5.2 Numerical processing in SNA

Numerical processing techniques in SNA are adopted by the feature extractor parts of the KS. The purpose of the numerical processing modules is essentially to individuate the semantically relevant elements in the data, if any.

This should be done, initially, through simple computations, in order to quickly discard noisy traces, and/or generate hypotheses for candidate events; this, in turn, will drive the subsequent processing steps. In other words, numerical processing should help in the simulation of the focus-of-attention activity of the expert analyst. The techniques we adopt, in the present version of the system, are segmentation, statistical estimates and spectral estimates.

After pre-processing, in which wide-band noise is removed, the traces are segmented into intervals of fixed duration (e.g. 1 s). Each interval is described in terms of amplitude and frequency: the former is estimated by the peak absolute amplitude, the latter by half the zero-crossing rate. This segmentation results in data compression of about two orders of magnitude, still preserving the relevant features of the signals. The segmented time series is treated as a bivariate distribution and scanned by a signal detection procedure in order to estimate the onset of a candidate seismogram, identify its possible end, extract the basic features (amplitude, frequency and duration) both in the signal and in the initial noise sequence. All this processing is done by the feature extractor of the KS named *Initialize.*

The signal-to-noise ratio and the duration of the candidate signals are used to rate their clarity with the following criteria: signals of high amplitude and long duration are more likely seismograms than short and low amplitude ones. The clarity is rated in three levels: clear, probable and possible. Too short signals (1 s or less) are immediately discarded, and the same holds for signals with an envelope exceedingly skewed with respect to the baseline. A detailed account on this processing is given by Chiaruttini (1989).

The above procedure provides a first arrival time estimate that is correct in most cases. When a more sensitive algorithm is necessary for the identification of the signal onset, we use a frequency domain detector. The algorithm we implement is that of Goforth & Herrin (1981) after replacing the Walsh transform with the Fourier transform. The detector is the feature extractor part of the KS *Search-P.* The best results are obtained analysing non-overlapping intervals 0.7 s wide. This allows arrival time determinations accurate enough for the preliminary analysis, while the lowest spectral amplitude estimate (at about 1.4 Hz) is adequate to the pass-band of short period instruments.

The combination of these algorithms provide both efficiency and sensitivity: the whole record is quickly scanned and the more sophisticated processing is applied only when necessary, as will be illustrated later. Precise time picking procedures are not yet implemented: in the SNA design this task is left to a later stage in the analysis when a sound event hypothesis exists.

5.3 Reasoning techniques in SNA

SNA basically adopts an opportunistic reasoning model (Nii 1986; Erman *et al.* 1980) which means that signal data are used to build the hypothesis elements, as well as to validate some of them, with no pre-defined order: in principle, the KS are loosely coupled pieces of knowledge whose activation is controlled by the current status of the working memory.

More specific reasoning strategies can be combined to the basic scheme; as an example, we briefly outline some ways in which the focus of attention of an expert analyst is modelled in SNA.

The features of the traces, extracted by the numerical processors, are used to quickly understand whether an event is present. The number and clarity of seismograms is used to rate the clarity of the event again in three levels: clear, probable and possible. Only clear and probable events are analysed. Possible events are those made of a small number of ambiguous signals: they are simply indicated as such for later inspection by the analyst. For the events selected in this way, a true problem-solving process is started: the seismograms with the highest clarity pattern are used to formulate an hypothesis at all semantically relevant levels. All subsequent steps are driven by this hypothesis, which is gradually incremented or modified, according to the results of the KS activations.

A meta-level reasoning strategy (Davis 1980) is also implemented in SNA by assigning priority levels to the KS: as pointed out earlier, the scheduler is charged with selecting a KS according to the assigned priority. The top priority level is given to the problem evaluation, then to the validation, the KS incrementing the hypothesis are executed with the lowest priority. This reasoning strategy models that of a 'cautious' analyst, who prefers to base his reasoning on checked information. In addition, within the KS sharing the same priority, those acting at the higher semantic levels are privileged, as will be further clarified in Section 6.1.

We stress here that the strategy can be changed by just changing the priority table of the KS.

6 THE PROTOTYPE: RESULTS AND EXAMPLES

A prototype version of SNA has been designed and implemented on a DEC Microvax II machine, equipped with VMS 4.7 operating system.

The logical part of SNA is written in the production system language OPS5 (Brownston *et al.* 1985), and consists of nearly 230 rules, for a total of roughly 7000 lines of code. The numerical operations on seismic traces, the input/output and part of the justification are performed by FORTRAN 77 modules, which include roughly 5000 lines, suitably interfaced to the OPS5 code.

The prototype has been designed to interpret the digital records collected by the North Eastern Italy Seismometric Network. The network is installed for the surveillance of the local seismicity and is equipped with short period instruments. A data-acquisition system monitors the seismometers and starts recording when some trigger occurs. The output of this system in 1987 February is currently used as test sample for SNA. The automatically recording stations were 11, at that time, covering the whole Friuli area, all equipped with single component instruments. The aperture of this subnetwork is about 100 km.

The small aperture of the network allows the assignment of a distance class also to events, not only to seismograms: local events are those occurring within 200 km from the network centroid, regional events are those located at distances from 200 to 2000 km, and teleseismic events are those farther away.

The implementation of SNA reflects the characteristics of the test network: reliability and efficiency in the interpretation of local events is emphasized; teleseismic events are considered of least importance; only the vertical component is considered.

Apart from pre-processing, a typical event is fully analysed in 60–150 CPU-seconds, while false triggers are discarded in about 10 s.

6.1 The knowledge sources

At the moment, SNA is composed of 12 KS. For the most relevant of them we list the activation conditions and the actions taken. The KS are listed in the order of their priority in the strategy table.

Seis-Cons-Prob

Activation. There is some consistency problem, i.e. at least a pair of inconsistent seismograms has been found (see the KS *Seis-Cons* for the definition of inconsistency).

Execution. The problems are examined to find the seismogram(s) responsible for the inconsistencies. The basic criterion is that the signal with the largest number of inconsistencies is excluded from the analysis. In case this condition occurs for more than one signal, that with the lowest signal-to-noise ratio is discarded.

Comp-Prob

Activation. There is some compatibility problem, i.e. at least a pair of incompatible phases has been detected (see the KS *Compatibility* for the definition of incompatibility).

Execution. The problems are examined to sort out the phase(s) more likely responsible for the incompatibilities. Phases incompatible with seismograms of higher clarity are considered wrong; at the same clarity level, the phases with the highest number of incompatibilities are considered wrong; if this number is the same for two phases, they are both considered suspect. It is inferred whether the wrong phase times are early or late, and an hypothesis of arrival time window is made, based on the available information: the event location or the time of the phase in compatible stations.

Seis-Cons

Activation. There is some seismogram not checked for consistency.

Execution. For each pair of seismograms, the ratio of frequencies and duration is checked. The seismograms are inconsistent if one of the ratios is greater than (or smaller than) an upper (lower) threshold value (set, for example, at 10 and 1/10, respectively). This test guarantees that the signals examined do not show totally different patterns.

Compatibility

Activation. There is some phase not checked for compatibility.

Execution. The compatibility of the arrival time of the phase on each pair of stations is checked. The times are compatible if their difference does not exceed the travel time for the distance between the stations.

Event

Activation. There are consistent seismograms not yet considered, or the attributes (e.g. initial time) of some seismogram have been updated.

Execution. An event hypothesis is made, if there is none; consistent seismograms are assigned to the event; the clarity of the event is rated; its distance class is inferred; the focal parameters are also determined. The event is rated as clear, probable or possible, based on the number and clarity of its seismograms. The distance class is determined from the majority of seismogram distance classes. The location is presently based on the first recording station: in case of a local event, the coordinates of the station and the first arrival time are taken as preliminary estimates of the event parameters; in case of a distant event, it is assigned an azimuth range. For local events the duration magnitude is evaluated.

Initialize

Activation. No trace has been scanned.

Execution. Traces are pre-processed to eliminate wide-band noise (e.g. spikes) and correct the baseline; segmentation and scanning of traces, clarity rating of candidate signals, exclusion of short and skewed 'signals' are also performed. The hypotheses of candidate signals (also named intervals) are made here.

Seismogram

Activation. Some phase attribute was modified, or the distance class of a seismogram differs from that of the event.

Execution. If the time of the first arrival was modified, then the seismogram initial time is updated. The distance class of the seismogram is inferred, based on the frequency of the first arrival. Above 7 Hz the distance is local, between 3 and 7 Hz is regional and below 3 Hz teleseismic. The distance class of the seismogram is redefined equal to that of the event, if they are different. The first arrival phase type is determined according to the seismogram distance class: *Pg* for local, *Pn* for regional, and *P* or *PKP* for teleseismic events respectively.

Intv-to-Seism

Activation. On a trace there is a candidate signal (interval) but there is still no seismogram.

Execution. The candidate signal is shifted from the class interval to the class seismogram; a hypothesis of first arrival is formulated. The distance class is defined local if the seismogram frequency is greater than the corresponding threshold.

Intv-to-Noise

Activation. The interpretation of the event is over and some candidate signal is still not processed.

Execution. The left-over interval is shifted to the class noise. Note that we assume here that there is at most one event in a record.

Search-P

Activation. There is an hypothesis of time window for an incompatible first arrival.

Execution. The detector is initialized by analysing the noise spectrum, and is run on the expected window. The first of these actions is performed only once for a trace, as obvious.

From the above list it is apparent that besides the aforementioned hierarchy of KS types, within a type, the KS making inferences on the higher levels of the hypothesis have higher priority (e.g. *Seis-Cons,* acting on seismograms, has higher priority than *Compatibility,* acting on phases, and similarly *Event* has higher priority than *Search-P*). A third control level on the clarity of seismograms exists: the interpretation starts with the signals of higher clarity and the less clear ones are in turn included. This is also a part of the focus-of-attention strategy in SNA, as pointed out in Section 5.3.

6.2 An event and its interpretation

A typical recording from the network, a clear local event of magnitude about 2.0, is reported in Fig. 3. The same figure shows also the *Pg* times detected by SNA with an accuracy of 1.0–1.5 s, depending on the signal-to-noise ratio.

The interpretation developed by SNA is reported in Fig. 4: it is the content of the CBH when no more KS can be fired. All the semantically relevant levels are present: the event, the trace, the seismogram, the phase (only one station has been reported): the capability of SNA to develop a global interpretation right at the beginning of its analysis is one of its distinctive features. This is due to the problem-solving paradigm adopted, which integrates coarse and fine grain knowledge in a unique framework.

A further peculiarity is that such interpretation has been reached with the simple numerical techniques outlined in Section 5.2: this is because the logical operations play an

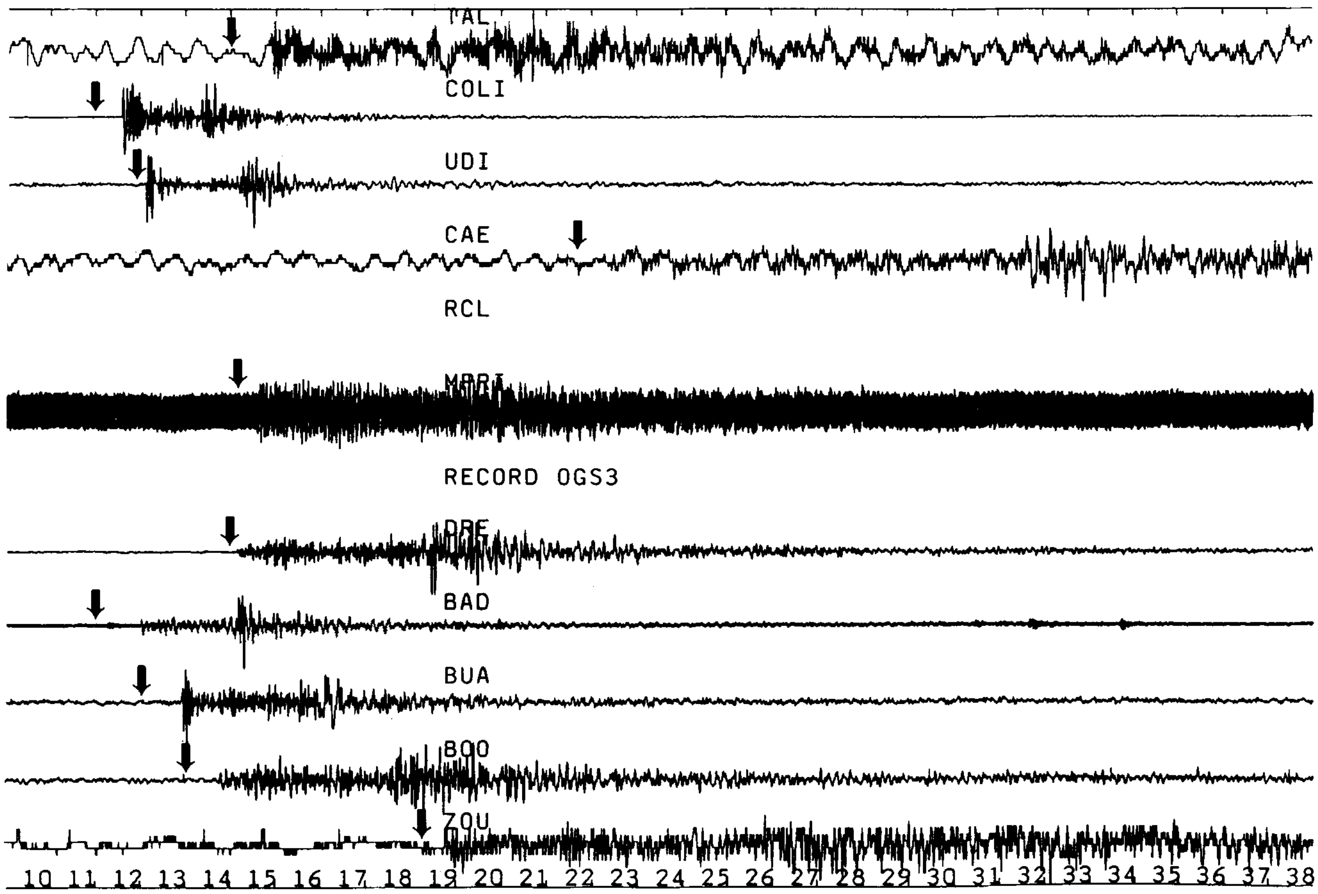

Figure 3. A clear local event recorded by the North-Eastern Italy Seismometric Network. The time unit is 1 s; the first 10 s of background noise are not shown; the black vertical arrows indicate the arrival times of *Pg* phases detected by SNA.

essential role in our system, namely inferences, associations of facts and findings, consistency analyses of features detected in different seismograms. In essence, logical processing allows a more efficient use of contextual information and its integration in the problem-solving process. The redundancy of information (i.e. typically more than one station record the same event) is fully exploited by compatibility and consistency analyses, which cures possible instabilities.

On the other hand, the limits of the present status of the interpretation are also apparent from this example: for instance, no S-wave analysis is performed, which prevents the system from locating regional events; in addition, the event location and origin time are affected by a large uncertainty, due to the crudeness of the estimates.

However, these faults are due to lack of knowledge coded in the system: its modularity guarantees that new knowledge can be added in a natural way, for example by adding further KS modules.

Let us now make a brief overview on the test performed so far. The test data were chosen to contain a few tens of particularly noisy network records. All records containing no events were quickly discarded, most of them as purely noisy records; the others, having some strong noise bursts, were classified as possible events left to the inspection of the analyst.

All the events were correctly interpreted at the event level; misinterpretations occurred at the seismogram and phase levels, at a rate of roughly one per event; this is also due to the fact that not enough knowledge is presently coded in SNA.

6.3 Decision-making and justification in SNA

Let us look more closely at the problem-solving activity in SNA, still referring to the event in Fig. 3.

If we think in terms of sequential activation of modules, we can only *a posteriori* show the list of knowledge sources which were activated to solve the quoted problem; we also give a brief account on the related processing tasks.

Initialize. The traces in the record are pre-processed. The trace of station RCL is found to be null (i.e. mostly made of zeros) and therefore discarded. Clear candidate signals are found at BAD, DRE, UDI, COLI, probable signals at BOO and BUA, and possible ones at ZOU, BAD, MPRI, CAE, COLI, TAL. Note that two signals are found at BAD and COLI: in both cases one is clear and the other possible. A noise lasting less than 1 s is also detected at BAD.

Intv-to-Seism. Seismogram hypotheses are made from clear and probable candidate signals (BOO, BUA, BAD, DRE, UDI, COLI). The distance class of all these seismograms is inferred to be local. *Pg*-phase hypothesis is made accordingly for the first arrivals.

Seis-Cons. Consistency check of clear seismograms

```
                C U R R E N T    B E S T    H Y P O T H E S I S

In Record #          3 I have identified the following EVENTS:

1 LOCAL          CLEAR

- - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -

The STATIONS are:

ZOU  BOO  BUA  BAD  DRE  MPRI RCL  CAE  UDI  COLI TAL
```

a)

```
                          E V E N T    D I S P L A Y

EVENT #            1 is LOCAL           and rated CLEAR

Magnitude                   :   2.2
Origin time                 :   12.0          s
Epicentre  latit. and  longit.:  46.2            13.3        deg
                                +/-    20.          km

Recorded at Stations:

CAE  ZOU  MPRI TAL  BOO  BUA  BAD  DRE  UDI  COLI
```

b)

```
                          T R A C E    D I S P L A Y

The TRACE of Station ZOU   has the following attributes: Normal

The Background Noise features are:

Duration         :   11.0         s
Amplitude        :   25.          mV
Frequency        :   20.7         Hz
Number of Spikes:          24

I identified the following Intervals:

         1 POSSIBLE       Seismogram
```

c)

```
                    S E I S M O G R A M    D I S P L A Y

Seismogram rated POSSIBLE         recorded at ZOU
of Event #            1    LOCAL              CLEAR

Initial time:   19.3          +/-  0.7            s
Duration    :   19.6          s
S. N. R.    :   2.5

I have recognized the following Phase(s): PG
```

d)

```
                          P H A S E    D I S P L A Y

PG    Phase of Seismogram recorded at ZOU
of Event #            1 LOCAL

Arrival time:   19.3          +/-  0.7            s
S. N. R.    :   3.5
Frequency   :   11.2          Hz
```

Figure 4. The CBH for some of the interpreted elements. All the semantically relevant levels are shown: the event (a), the trace (b), the seismogram (c), and the phase (d).

(BAD, DRE, UDI, COLI) is done; no error has been found.

Compatibility. Compatibility check of *Pg* phases of clear and probable seismograms is done; no error found.

Event. A hypothesis of clear local event is made, based on clear seismograms. The event is located close to the station COLI; its magnitude is evaluated.

Seis-Cons. The consistency of probable seismograms at BOO and BUA is checked; no error is found.

Event. The probable seismograms are added to the event; the event parameters are up-dated.

Intv-to-Seism. Hypotheses of possible seismograms are made for stations ZOU, MPRI, CAE, TAL. For the same stations a hypothesis of first arrival is made, but the phase type is left undefined as there is no frequency estimate, due to the low signal-to-noise ratio.

Seis-Cons. The consistency of possible seismograms is checked; no error is found.

Event. Possible seismograms are added to the event; the event parameters are up-dated.

Seismogram. The distance class of possible seismograms is inferred to be local according to that of the event. No such hypothesis was made before, due to lack of reliable frequency information. The first arrival is inferred to be *Pg*.

Compatibility. The *Pg* phase at CAE is incompatible with those at BOO, BUA, BAD, MPRI, UDI, COLI, TAL; the *Pg* phase at ZOU is incompatible with those at BOO, BUA, BAD, DRE, MPRI, UDI, COLI, TAL. The initial hypothesis of onset time was in fact 32 s at CAE (the S-wave) and 33 s at ZOU (the peak of the seismogram).

Comp-Prob. The *Pg* phase times at ZOU and CAE are inferred to be late; arrival time windows are hypothesized based on the event location.

Event. The event parameters are up-dated.

Search-P. The first arrival is searched and found by the detector at ZOU and CAE.

Compatibility. The time of *Pg* phases at ZOU and CAE is found compatible with all the others.

Seismogram. The initial time of seismograms of ZOU and CAE is updated.

Event. The event parameters are updated.

Intv-to-Noise. The possible candidate signals from BAD and COLI are considered noise.

The example illustrates how contextual analysis is implemented in SNA, and how the hypothesis is built up incrementally: information about seismograms and phases yield information about the event, and information about the event, in turn, yield information and expectations about

a)

```
Ev Sei Ph  Tr     --------------  R E A S O N I N G   C H A I N  --------------

1                 An EVENT HYPOTHESIS is tentatively made since there is
                  at least one consistent Seismogram.

1 BAD             This SEISMOGRAM is PART OF the EVENT,
                  as it is part of a consistent set.

1 DRE             This SEISMOGRAM is PART OF the EVENT,
                  as it is part of a consistent set.

1 UDI             This SEISMOGRAM is PART OF the EVENT,
                  as it is part of a consistent set.

1 COLI            This SEISMOGRAM is PART OF the EVENT,
                  as it is part of a consistent set.

1                 The EVENT was assigned a DISTANCE CLASS based on the
                  relative majority of Distance Classes of its Seismograms.

1                 The EPICENTRE and ORIGIN TIME of the LOCAL EVENT
                  were evaluated.

1                 The DURATION MAGNITUDE of the Local Event was EVALUATED.

1                 The EVENT is rated as CLEAR,
                  since it consists of more than NSisCl clear Seismograms.

1 BOO             This SEISMOGRAM is PART OF the EVENT,
                  as it is part of a consistent set.

1 BUA             This SEISMOGRAM is PART OF the EVENT,
                  as it is part of a consistent set.

1                 The PREVIOUS LOCATION of the event is CONFIRMED.

1                 The DURATION MAGNITUDE of the Local Event was EVALUATED.

1 ZOU             This SEISMOGRAM is PART OF the EVENT,
                  as it is part of a consistent set.

1 MPRI            This SEISMOGRAM is PART OF the EVENT,
                  as it is part of a consistent set.

1 CAE             This SEISMOGRAM is PART OF the EVENT,
                  as it is part of a consistent set.

1 TAL             This SEISMOGRAM is PART OF the EVENT,
                  as it is part of a consistent set.

1                 The PREVIOUS LOCATION of the event is CONFIRMED.

1                 The DURATION MAGNITUDE of the Local Event was EVALUATED.
```

Figure 5. The justification of reasoning about the hypothesis elements shown in Fig. 4: the event (a), the trace (b), the seismogram (c), and the phase (d).

b)

```
Ev Sei Ph  Tr   --------------  R E A S O N I N G   C H A I N  --------------

           ZOU  The number of SPIKES in the Trace, and the NULL_TRACE flag
                are added to the Trace attributes.

           ZOU  The TRACE has been scanned to define the still missing
                attributes and to identify CANDIDATE SIGNALS.

           ZOU  The CANDIDATE SIGNAL was rated as POSSIBLE.

   ZOU     ZOU  A hypothesis of SEISMOGRAM is made out of the
                candidate signal of higher clearness in the trace.
```

c)

```
Ev Sei Ph  Tr   --------------  R E A S O N I N G   C H A I N  --------------

   ZOU     ZOU  A hypothesis of SEISMOGRAM is made as there is an Event
                hypothesis and a candidate signal rated POSSIBLE

   ZOU          A PHASE is HYPOTHESIZED at the beginning of Seismogram; its
                TYPE is left UNDEFINED as the Seismogram has no distance class

1  ZOU          This SEISMOGRAM is PART OF the EVENT,
                as it is part of a consistent set.

   ZOU          The SEISMOGRAM DISTANCE CLASS was REVISED
                and set equal to that of the Event.

   ZOU          The time of the first arrival was modified,
                then the INITIAL TIME of the SEISMOGRAM is UPDATED.
```

d)

```
Ev Sei Ph  Tr   --------------  R E A S O N I N G   C H A I N  --------------

   ZOU  PG      The FIRST ARRIVAL PHASE TYPE was inferred, based on the
                Seismogram distance class.

   TAL  PG      The TIMES of these PHASES are INCOMPATIBLE.
   ZOU  PG

   MPRI PG      The TIMES of these PHASES are INCOMPATIBLE.
   ZOU  PG

   BOO  PG      The TIMES of these PHASES are INCOMPATIBLE.
   ZOU  PG

   BUA  PG      The TIMES of these PHASES are INCOMPATIBLE.
   ZOU  PG
   BAD  PG      The TIMES of these PHASES are INCOMPATIBLE.
   ZOU  PG

   DRE  PG      The TIMES of these PHASES are INCOMPATIBLE.
   ZOU  PG

   UDI  PG      The TIMES of these PHASES are INCOMPATIBLE.
   ZOU  PG

   COLI PG      The TIMES of these PHASES are INCOMPATIBLE.
   ZOU  PG

   ZOU  PG      The PHASE TIME is LATE as it is not compatible
                with the Phase of Seismogram(s) of higher Clearness.

   ZOU  PG      The FIRST ARRIVAL was searched and FOUND by the Detector.
```

seismograms and phases. As a consequence, the processing focus switches between the hypothesis elements in an opportunistic way: no systematic loop in the sequence of KS activation may be traced.

We point out that although feedback can be implemented also in procedural systems (Anderson 1978), it is a natural outcome when using AI techniques, where the design mainly concerns the *conditions* for module activation rather than the *sequence* of activation.

Figs 3 and 4 also show the ability of SNA to exploit the information available in the data, down to very low signal-to-noise ratios – note the correct identification of first arrivals at ZOU, MPRI, CAE and TAL.

Another distinctive feature of SNA is its capability to trace back the reasoning chain followed to reach an interpretation.

The justifier module in our system adopts the already mentioned object hierarchy: therefore the user can ask the system to justify its behaviour at any semantic level. Fig. 5 reports the justification for some of the hypothesis elements.

7 CONCLUSIONS

In this paper we argued that artificial intelligence techniques provide a new and promising approach to the problem of automating the interpretation of seismic network recordings.

In particular, knowledge-based systems with blackboard architectures offer a suitable problem-solving paradigm, which can model in an effective way the activity of the expert analyst. On this basis, the expert system SNA has been designed; a prototype has been developed and is being tested on real data.

The results we presented and discussed in the previous sections confirm our expectations.

We believe that a number of distinctive features have proved to be very interesting when compared with the current procedural systems.

First of all, the integration of knowledge of different kinds and from different domains guarantees at the same time high sensitivity and robustness to the analysis. In this respect, the interpretation of records with low signal-to-noise ratio can be afforded.

Secondly, the logical control in the system, not pre-defined nor explicitly coded, ensures flexibility to the problem-solving process. In this context, we have shown how focus-of-attention techniques and the ease of feedback ensure efficiency to the interpretation. As a result, simple-minded and inexpensive numerical computations are sufficient to reach a thorough hypothesis of the event; more precise computations can be postponed to later stages in the analysis.

Moreover, the modularity of the architectural scheme allows an easier maintenance and expandibility: for example, new knowledge can be added by adding new KS.

Finally, the justification allows the user to gain a deeper insight into the system's behaviour. Therefore the comparison with the analyst's reasoning is direct, which greatly simplifies testing and debugging procedures; in addition, justification may prove very useful for teaching and training purposes.

Work is in progress to further extend the capabilities of SNA.

ACKNOWLEDGMENTS

We thank Dr M. Russi of the Osservatorio Geofisico Sperimentale of Trieste for providing the test data, and Dr F. Gentile of the same observatory for helpful discussions about seismogram interpretation.

REFERENCES

Allen, R., 1982. Automatic phase pickers: their present use and future prospects, *Bull. seism. Soc. Am.*, **72,** S225–S242.

Anderson, K. R., 1978, Analysis of microearthquake network data, *Geoexploration*, **16,** 159–175.

Barr, A. & Feigenbaum, E. A., 1982. *The Handbook of Artificial Intelligence*, Pitman, London.

Blandford, R. R., 1982. Seismic event discrimination, *Bull. seism. Soc. Am.*, **72,** S69–S87.

Brownston, L., Farrell, R., Kant, E. & Martin, N., 1985. *Programming Expert Systems in OPS5,* Addison-Wesley, Reading, Massachusetts.

Chiaruttini, C., 1989. *Focus-of-attention techniques in the automatic interpretation of seismograms,* Research Report UDMI/05/89/RR, University of Udine, Italy.

Chiaruttini, C. & Roberto, V., 1988a. Automation of seismic network signal interpretation: an artificial intelligence approach, *Il Nuovo Cimento*, **11C,** 327–338.

Chiaruttini, C. & Roberto, V., 1988b. SNA: a knowledge-based system for seismic network signal interpretation, in *Signal Processing IV: Theories and Applications*, pp. 1205–1208, eds. Lacoume, J. L., Chehikian, A., Martin, N. and Malbos, J., North-Holland, Amsterdam.

Davis, R., 1980. Meta-rules: Reasoning about control, *Artificial Intelligence*, **15,** 179–221.

Erman, L. D., Hayes-Roth F., Lesser, V. D. & Reddy, R. D., 1980. The Hearsay II speech understanding system: integrating knowledge to resolve uncertainty, *ACM Comput. Surv.*, **12,** 213–253.

Frost, R. A., 1986. *Introduction to Knowledge Base Systems*, Collins, London.

Goforth, T. & Herrin, E., 1981. An automatic seismic signal detection algorithm based on the Walsh transform, *Bull. seism. Soc. Am.*, **71,** 1351–1360.

Keilis-Borok, V. I., 1972. Seismology and logic, in *Computational Seismology*, pp. 1–9, ed. Keilis-Borok, V. I., Consultant Bureau, New York.

Nii, H. P., 1986. Blackboard systems: the blackboard model of problem solving and the evolution of blackboard architectures, *AI Magazine*, Summer, pp. 38–53.

Nii, H. P. & Feigenbaum, E. A., 1978. Rule-based understanding of signals, in *Pattern-directed Inference Systems*, pp. 483–501, eds Waterman, D. A. & Hayes-Roth F., Academic Press, New York.

Nii, H. P., Feigenbaum, E. A., Anton, J. J. & Rockmore, A. J., 1982. Signal-to-symbol transformation: the HASP/SIAP case study, *AI Magazine*, Spring, pp. 23–35.