Qualitative Abduction and Prediction:

Regularities over Various Expert Domains

E. Burattini*° M. De Gregorio
*Istituto di Cibernetica, CNR,Via Toiano 6 - I-80072 Arco Felice, Napoli, Italy
tel. +39-(0)81-8534204(8), fax +39-(0)81-5267654,
°LAFORIA Institut Blaise Pascal - Paris VI - 4,Place Jussieu, Paris, France
e-mail: ernb@sole.cib.na.cnr.it; massimo@sole.cib.na.cnr.it

(to be published by *Information and Decision Technology*,
Elsevier Sci. Pub., North Holland)

Abstract

We describe work on designing and implementing actual Diagnostic Expert Systems in the areas of medicine and civil engineering.

A classification is introduced on the basis of the nature and structure of individual domains, and case studies are comparatively examined in the light of this classification. The significance of qualitative reasoning for these expert systems is discussed and a rough distinction between three commonly encountered situations is introduced as a basis for choosing between appropriate types of models for knowledge representation and processing.

A Non-Diagnostic Expert System, in the area of architecture, is also introduced showing that the same methodologies used in diagnostic domains can be applied.

Résumé

Nous présentons des travaux sur le projet et réalisation de systeme expert diagnostique dans le domains de la médecine et du bâtiment. Nous introduisons une classification fondée sur la nature et sur la structure de chaque domaine, et nous comparons différents exemples suivant cette classification. Nous discutons de l'importance du raisonnement qualitatif pour ces systèmes experts comme base du choix des types de modèles appropriés pour la représentation et l'élaboration de la connaissance.

Nous examinons aussi un système expert non-diagnostique, dans le domaine de l'architecture, et nous montrons qu'on peut utiliser la même méthodologie que dans le cas du diagnostic.

1 Introduction

The task of a diagnostic expert system can be roughly described as that of isolating a set of explanatory hypotheses for the insurgence of certain types of anomalies observed in objects belonging to its intended domain of application. For example, an expert system concerned with the diagnosis of skin diseases must advance explanations for abnormal skin features observed in particular human beings, and one concerned with the diagnosis of statical wearout in buildings must determine the causes of statical damages detected in particular buildings (such as, e.g., certain types of cracks in the walls). A diagnostic expert system can thus be viewed as a particular type of problem solver.[1] The problem's statement is a description of an abnormal state observed in an object belonging to its domain of application, and a solution is given by a set of possible explanations for the occurrence of this abnormal state.[2]

In order to produce a diagnosis, or to explain a situation from observed facts, an expert system makes use of a knowledge base extracted from human experts. Knowledge bases usually provide a very impoverished representation of the knowledge available to human experts: the latter are generally required to state the information directly involved into their decisions, which represents only an end product of their theoretical and experimental knowledge [8]. This information is ultimately derived from the overall theories, models and inductive generalizations available to the experts, but even when refined theoretical knowledge about the behaviour of the objects in the intended domain of application is available, this is rarely represented in the knowledge base of the expert system. Background theoretical knowledge plays a crucial role, as it were, only from the outside of the expert system – in formulating and justifying the correctness of the information that human experts communicate to the knowledge engineer. The accuracy and intersubjective acceptability of the obtainable knowledge base as well as the most appropriate inference strategies depend crucially on the state of our knowledge in the given domain of application.

This variability is reflected in the expert systems in three different domains (medicine, civil engineering and architecture) examined in sections 3 to 5. But let us first focus on common features of diagnostic processes.

2 Commonalities between diagnostic systems

The central inferential processes involved in diagnostic processes are <u>abduction</u> (for selecting possible explanatory hypotheses of observed data) and <u>prediction</u> (for searching new supporting information for the selected explanatory hypotheses).

---

[1]Cf. the informative discussion in [23].

[2]Some explicata of the imprecise ideas of diagnostic problem and solution have been discussed in the literature; representative cases are in [18] and [20].

The term 'abduction' indicates a mode of inference whose conclusion is merely plausible. In limited domains, where the direct relations between facts and corresponding explanatory hypotheses are known[3], abductive inferences can be represented by the following schema:

$$\frac{\begin{array}{l}\mathtt{A} \rightarrow \mathtt{B}\\ \mathtt{B}\end{array}}{\mathbf{A}} \tag{1}$$

in which, being the rule `A→B` known, and supposed true the observation `B`, one can suppose true (*abduce*) the left hand **A** of the rule.

Abduction is a key process for generating hypotheses which can justify observed facts; it can be viewed as the *step* from observed facts to possible explanations. The conclusion of a deduction is a logical consequence of its antecedents, whereas abduction generates hypotheses which need an external confirmation.

In order to evaluate the credibility of the inferred hypothesis, instead of adopting a numerical (e.g. bayesian) inference schema as, for example, in Charniak and Mc Dermott [10], where the most probable inferred hypothesis is assumed to be the best, we adopt a "qualitative" evaluation schema. We notice that such an approach, already present in Peirce [4] has been adopted in Hintikka [15], Reggia *et al.* [18], Reiter [20], Eshghi and Kowalski [14] and in Burattini *et al.* [5][6] who reject probabilistic or numerical constructs.

With the term 'prediction' we indicate a mode of inference in which the antecedent of the rule is only supposed to be true and, in this case, also the conclusion of the rule is supposed to be true. This inference is formally represented by the usual schema for *modus ponens*, with the crucial proviso that **A** is obtained through an hypothesis formation process:

$$\frac{\begin{array}{l}\mathtt{A} \rightarrow \mathtt{C}\\ \mathbf{A}\end{array}}{\mathbf{C}} \tag{2}$$

This inference is useful when one looks for more confirming (or falsifying) information for explanatory hypotheses. In fact, suppose one has both rules (1) and (2) in a knowledge base. After observing `B`, one can assume **A** as a possible explanation for the presence of `B`. Which additional facts could be present (true) if **A** were the right hypothesis? Using predictive inferences one produces the set of possible consequences of **A** (in our case just {**C**}), and hypothesis **A** is assessed on the basis of this set.

---

[3] Writing '`A→B`' we mean that '`A` is a possible explanation of `B`'. Thus '→' defines a relation between the set of hypotheses and the set of facts.

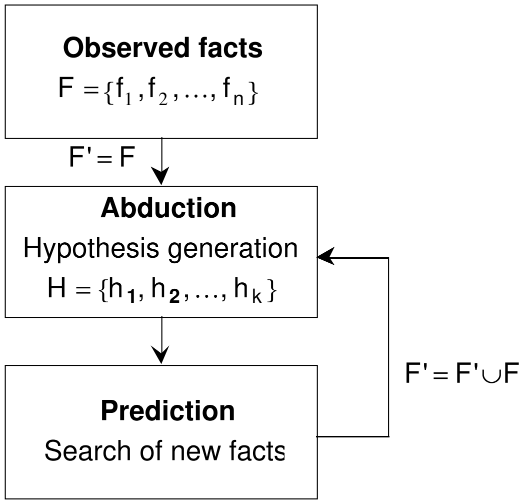

Figure 1: *Abductive-predictive process*

An overall abductive-predictive cycle commonly encountered in diagnostic problem solving is shown in fig. 1. More fine-grained, domain-dependent differences between various diagnostic processes are brought out with the description of the expert systems given in the ensuing sections.

3 A medical domain

The expert domain is pediatric gastroenterology [13]. As in many other medical domains, the diagnosis is formulated on the basis of information which is not always certain and complete. Furthermore, many medical domains are characterized by a *poorly structured knowledge* [9], i.e., there is no comprehensive and highly predictive theory for the behaviour of the objects in these domains. Relationships between observable anomalies and possible explanations are mostly derived from statistical correlations taken on a sample of a given population, with experts often disagreeing on their predictive value and generalizability to other populations. Sometimes even significant statistical correlations are unavailable, and human experts can only communicate examples of their decisions, by associating sets of anomalies and diagnoses, with the indication that these associations are generalizable to "similar" cases. Under these circumstances, it is problematic to obtain an informative protocol of intermediate reasoning processes of the experts, and this leads to the implementation of a system using only rudimentary, mostly one-step inferential processes.

Many medical diagnostic problems fall into this category, but the situation is typical of several areas of interest for expert system research. For example, meteorology, continental drift theory, vulcanology and seismology are characterized by highly qualitative and weakly predictive theories or models, widely supplemented by inductive generalizations on past events.

Lacking an adequate theoretical development, diagnostic activities can be heavily influenced by subjective factors, and expert systems developed from poorly structured knowledge often provide only a very approximate model of the diagnostic activity of individual experts.

Many instances of medical expert reasoning follow the basic abductive-predictive cycle described in section 2: from the knowledge of some observed signs an initial set of diagnostic hypotheses is elicited; subsequently, a query process takes place in order to isolate other signs enabling one to confirming, disconfirming or differentiating between the selected hypotheses.

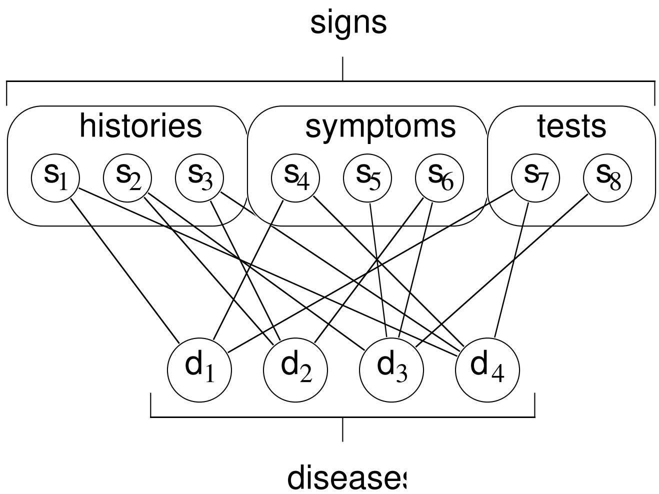

Figure 2: *knowledge base structure*

The knowledge base In this system, the knowledge base is represented as a bipartite graph. The two sets of nodes represent signs and diseases, respectively. Edges link signs with their possible explanations. The set of sign nodes is split into three classes: *histories* (signs arising from the patient anamnesis); *symptoms* (signs arising from medical examination); *tests* (signs arising from laboratory tests and/or therapies). Splitting the set of sign nodes into classes allows one to introduce a qualitative evaluation. For example, one can assert that a *test* is more relevant than a *symptom*, that is, the abductive support for a disease of a *test* is stronger than that provided by a *symptom* or by a *history*. For the signs present in fig. 2 one can assert, for example, that: s_2 p s_4 p s_7, where 'p' means 'less important than'. Furthermore, these classes denote different sorts of knowledge in terms of the roles they play in eliciting diagnostic hypotheses.

The inference engine Organizing the knowledge base as a bipartite graph facilitates the implementation of the abductive-predictive process showed in fig. 1. Each edge of the graph is considered by the system as a relation between one sign and one disease. For example, in fig. 2 the relations expressed by the edges of the graph are: $d_1 \rightarrow s_1$, $d_1 \rightarrow s_4$, …, $d_4 \rightarrow s_8$.

If one denotes by $O_s = \{s_2, s_3, s_8\}$ the set of observed signs, to justify observations O_s, the system goes from sign nodes to disease nodes following the arcs of the graph to obtain a set of possible diseases (abduction); the result of this operation is the set of diseases $D_s = \{d_2, d_3, d_4\}$. To confirm these diseases the system, using predictive inferences, asks the following question: given D_s, which other signs could be present in addition to the observed signs (O_s)? The abductive-predictive process allows one to generate a set of diagnostic hypotheses on the current situation and to locate the possible questions that may help confirming the elicited hypotheses.

Query and hypothesis evaluation The search for further information is based on the assumption that the degree of support for a disease depends on the diversity of the observed signs (i.e., a diagnostic hypothesis which explains signs belonging to different classes is much more credible than an hypotheses which explains only signs belonging to the same class). This assumption is adopted also in one of the inference schemata proposed by Polya [17]:

$$\begin{array}{l} A \rightarrow B_1, B_2, \ldots, B_n, B_{n+1} \\ B_1, B_2, \ldots, B_n \\ B_{n+1} \text{ differs to a great extent from } B_1, B_2, \ldots, B_n \\ \hline \mathbf{A} \text{ much more credible .} \end{array}$$

Supposing one knows the rules $A \rightarrow B_1$, ..., $A \rightarrow B_n$ and $A \rightarrow B_{n+1}$, and supposing that the facts (in our case, signs) B_1, ..., B_n have been observed and accepted as true, one advances $\mathbf{A}$ as possible cause for B_1, ..., B_n. If, in addition, B_{n+1} is observed and accepted as true, since B_{n+1} differs to a great extent from B_1, ..., B_n, then cause $\mathbf{A}$ becomes more credible.

Since we have divided information into classes, the search for further information must detect signs belonging to classes about which no information is available. This is quite close to the query strategy followed by physicians. For example, after collecting information about the patient's anamnesis they look for symptoms, and possibly prescribe tests and/or therapies.

Since the signs are ordered according to their importance in eliciting diseases, the inferential process, based on the interleaving of abductive and predictive subprocesses, produces a list of diseases with a preference ordering. This ordering is induced by the organization into classes of signs involved in the elicitation process [6].

In the example of fig. 2, if the observed set of signs is $O_s = \{s_2, s_3, s_8\}$ then the list of ordered diseases, providing a possible explanation for O_s, is $d_2 \; \underline{p} \; d_4 \; \underline{p} \; d_3$. The meaning of '$\underline{p}$' is 'less likely than'. We notice that the total ordering of the set of signs induces a partial order on the set of diseases. Since d_3 is a possible explanation for both s_2 and s_8, and these signs belong to different classes of information, the expert system ranks d_3 as the best diagnosis.

A theoretical approach based on multiset theory [6] was developed to evaluate both signs and diagnoses.

4 An expert system in civil engineering domain

Now, we review work on the diagnosis of the statical wearout of buildings [7]. Here the domain has a *partially structured knowledge* [9], i.e., there are fairly precise theories and models relative to the objects in the domain of the diagnostic expert system, though not for every aspect of their behaviour: one may exploit a wealth of information derived from the physical laws of mechanics, the data relative to construction techniques, and to characteristic features of the materials used. However, the structure of, e.g., the construction materials may not be fully known, and a complete theory of their behaviour under stress may not be available. Thus, past experience on "similar" situations may be needed to integrate theoretical knowledge in designing an adequate knowledge base. Such information might be needed for the diagnosis of malfunctioning in many types of man-made systems, whose complexity is a practical obstacle to the design of a usable knowledge base derived solely from models and theories.

The effects, that is, structural damages, and the relations between causes and effects (damages), are mainly derived from theories.

The main diagnostic task here is that of partitioning causes into two sets: primary causes, i.e. those to which the observed damages can be mainly ascribed, and secondary causes, i.e. the causes which may have contributed to the insurgence of the observed damages only in the presence of primary causes.

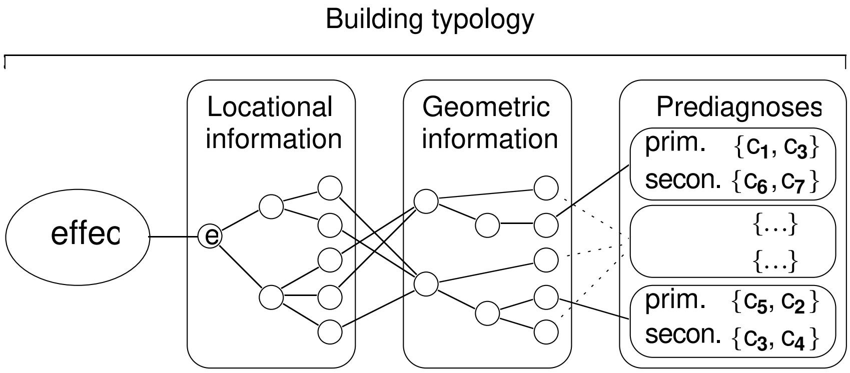

Figure 3: *Knowledge base*

The knowledge base The knowledge structure is represented by a set of trees. The root of each tree is a damage, the leafs represent the qualified damage and the terminal nodes are the sets of possible causes (fig. 3). Each tree is labeled by one or more building typology. The cause-effect links (trees) are taken into account in solving a diagnostic problem only if they are labelled with the selected building typology (contextual information - analogous to sex, age, and other background information in medical domains). For each observed damage, the experts are usually capable of determining the associate sets of possible causes and to distinguish between primary and secondary causes.

A set of possible questions is associated to each cause. The system can ask these questions in order to strengthen, weaken or discard causes. These questions, asked by the system before giving the final diagnosis, concern chronological and environmental information.

The inference engine In this case, expert reasoning develops into three stages. First, an augmenting (or qualifying effect) and refining process of the observed damages takes place, which involves, e.g., a more precise description of the damages (i.e. width, size, etc.); then a prediagnosis is obtained on the basis of this information (abduction); the final diagnosis is reached via a query process driven by the causes present in the prediagnosis (prediction). In this final step, a cause refinement process takes place in order to discriminate between primary and secondary causes and to rule out some elements of the prediagnosis. This requires interactions with the user, enabling the system to determine the contextual and environmental situation in which the damages appeared.

The inferential process traces a path along the tree, driven by the information acquired step by step from the user (in fig. 4 the thick lines represent one of the possible paths). All the explored paths are stored by the system to be used later when providing explanations for the proposed diagnosis.

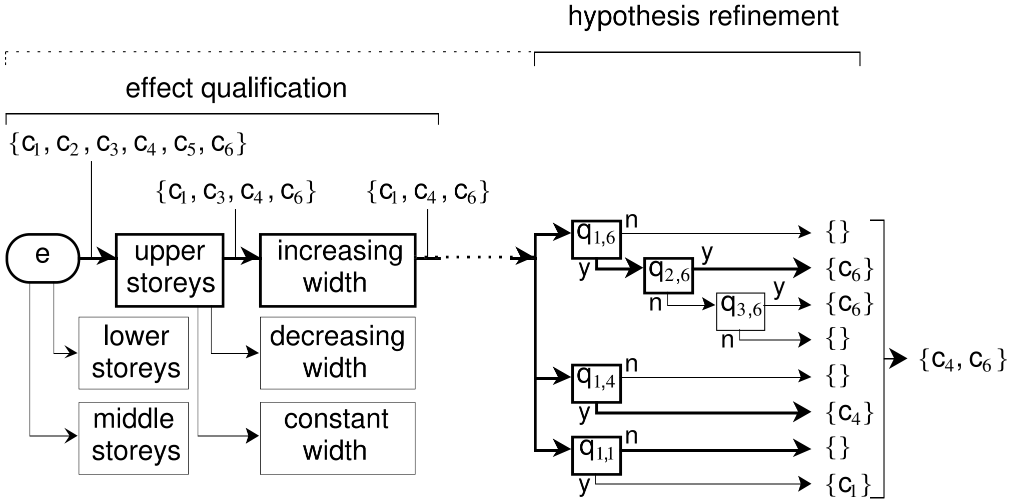

e	vertical fracture in the wall	$q_{1,6}$	are there relevant temperature differences between the inside and the outside of the building?
c_1	undulatory earthquake shock	$q_{2,6}$	did the fracture appear after the insurgence of these temperature differences?
c_2	foundation sinking	$q_{3,6}$	did the fracture extend after the insurgence of these temperature differences?
c_3	vertical load increase	$q_{1,4}$	is there a thrust of the roof?
c_4	thrust of the roof	$q_{1,1}$	did an earthquake occur?
c_5	thrust of vaults and arches		
c_6	thermal actions		

Fig. 4: *Qualifying and refinement processes*

At least three relevant differences between the medical domain and the engineering domains can be brought out:

- in the medical system the interaction with the user takes place by means a natural-language-like interface; in the civil engineering system, in order to avoid mistakes due to imprecise descriptions of the damages, a prevalently iconic language was choosen, thereby eliminating sources of uncertainty and ambiguity of the input data;
- in the statical system, for each step of the effect qualification process, the experts can isolate the set of possible causes with their value (as shown in the left side of fig. 4). This is made possible by the knowledge structure of the civil engineering

domain, in which the cause-effect links are mainly derived by means of the statical laws of physics;

- in the final phase of the hypothesis refinement process (see the right side of fig. 4) one can ask directly a question about the cause under examination. For example, if one of the hypotheses present in the prediagnosis is *earthquake*, the system can ask the user whether an earthquake occurred (at this stage, the system has no information about the context in which it is solving the diagnostic problem). Thus, the expert system can evaluate again the cause *earthquake* and eliminate it from the final diagnosis in case of a negative answer. This sort of interaction is impossible in medical domains: physicians cannot directly ask the patient this kind of questions.

5 Classification of architectural styles

The third system, currently under development, is supposed to provide an historical and architectonic description of old non-monumental buildings in order to identify the different periods in which the building was made, modified, partially destroyed or enlarged. The expert's reasoning is hypothetical, that is, hypotheses are elicited from the analysis of historical, architectural, and structural characteristic signs.

Here, causal reasoning is only marginal involved, no real cause-effect relations hold between signs and hypotheses. The problem is mainly that of recognizing which historical period the observed signs belong to.

Clearly, this system is a non-diagnostic expert systems. It seems likely that the methods and principles described in section 2 can be extended and applied to non-diagnostic tasks (see also Reggia and Peng [19]).

Uncertainty arises both in the stage of collecting observable signs and when gathering contextual and historical information. These data should drive the inferential engine to reconstruct the history of the building.

The main signs which provide relevant historical information are:

- building techniques,
- building materials,
- arrangement of architectonic elements (portal, windows, etc.),
- stylistic character of architectonic elements.

Unfortunately, most of these signs are not easily recognizable looking at the building, and for some of them there is no bibliographic information. In these circumstances, expert's knowledge can only be suggestive, and experts can only "qualitatively" assess the importance of that sign in the final decision.

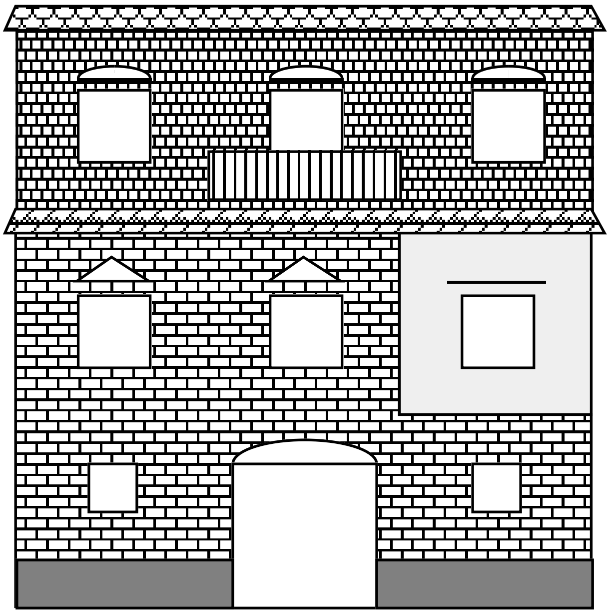

Fig. 5: *Building map*

Expert reasoning In the historical analysis, we distinguish among three main processes: coarse examination, fine examination, and synthesis.

Coarse examination

In this part of the process, the system reproduces a map of the materials (mainly stones and mortars) present in the building (see fig. 5) and advances hypotheses on the historical periods to which those materials belong to. The steps followed by the system are:

1. building typology and building techniques analysis;
2. building materials and their location analysis;
3. location and shape analysis of hewn stones;
4. coherence control of the data acquired in steps 1 to 3 (for instance, verify whether there is a chronological coherence between the hewn stones, made of a particular material, and the mortar used to bind them).

As in the previous systems, predictive inference is adopted to ask the user information for weakly supported hypotheses.

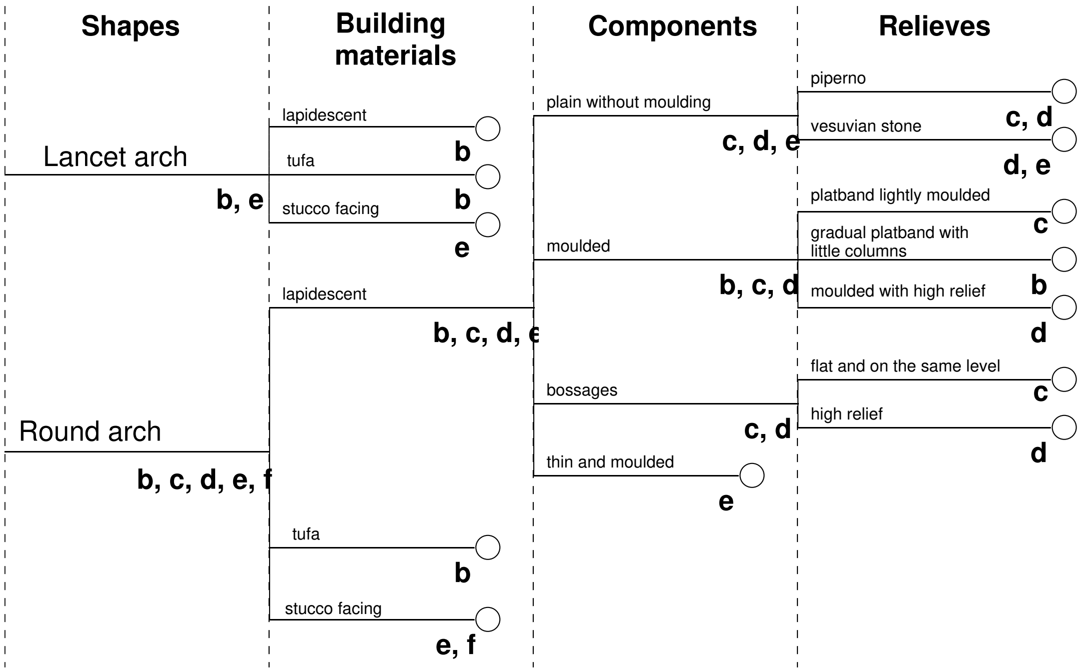

a - Before VII cen., **b** - Medieval VII-XIV cen., **c** - Renaissance XV-XVI cen.,
d - Baroque XVII-XVIII cen., **e** - Nineteenth-century, **f** - Twentieth century.

Fig. 6: *Inferential tree for two types of portals*

Fine examination

Here, the historical analysis concerns the investigation of the relevant architectonic elements: portal, window, balcony, footing, eaves. The investigation process is the same for each of these elements. Fig. 6 shows a sketch of the inferential tree for two types of portals.

The portal analysis is segmented into four parts.

Firstly, one determines the shape (i.e. lancet, round, etc.); secondly, the building materials of the portal are inspected; thirdly, one isolates the architectural components of the portal; fourthly, the relieves are analyzed.

At each step, a subset of the hypotheses concerning the classification of the portal is isolated, refined, and reduced as one progresses from shape to relief. This is quite close to the inferential process adopted in the civil engineering domain (one can notice the similarity between the trees of fig. 4 and of fig. 6).

The *fine examination* allows the system to date each element with an error range of 50 years.

Synthesis

After the *coarse* and *fine examination* the system produces a map of the building allowing one to distinguish between areas belonging to different periods. The period set for each area can be fairly wide. The *synthesis* operation is supposed to reduce this range by intersecting different information on the same area. For example, if the material around the portal belongs to the period 1650-1740 and the portal belongs to the period 1700-1750, the system will set the period for that part of the building to 1700-1740.

A further check of chronological coherence is made by the system in order to verify whether the parts of the building dated with different periods could have been assembled together. For example, it is unlikely that a lower part of building is older than its upper parts; the rare occurrences of this situation have to be justified on the basis of healing operations made possible by the construction technologies available at the period the system is focusing on. To deal with this kind of situation, the system is provided with a set of production rules given by the experts; these rules express knowledge on building techniques, building materials, etc., used in the different periods one is concerned with.

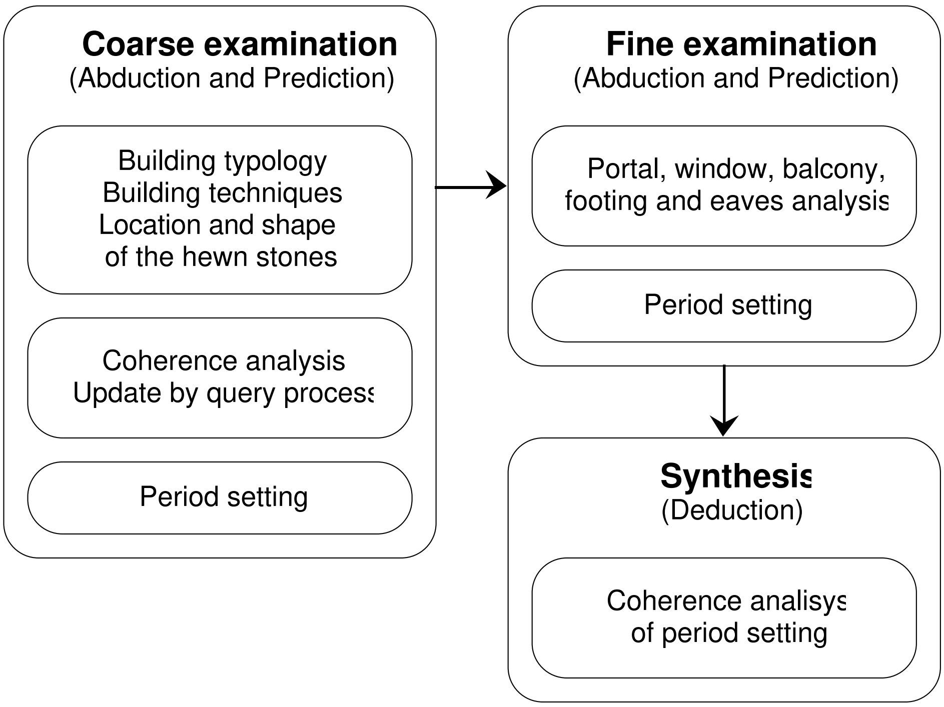

Fig. 7: *Historical and architectonic analysis*

To sum up, in the *coarse examination*, the experts mainly work on the basis of knowledge given by examples or by analysis of similar cases (poorly structured knowledge). In the *fine examination*, predictive reasoning is adopted starting from the hypotheses generated by the abductive inference process. In the last phase (*synthesis*), the system checks possible "incoherences" in the set of selected hypotheses, eliminates and/or ranks differently elements of this set; this phase is driven by a deductive inference engine on a set of production rules. The overall historical and architectonic analysis of old non-monumental buildings is showed in fig. 7.

6 Concluding remarks

Many diagnostic expert systems adopt a numerical paradigm to infer hypotheses, starting from facts and rules, and testing them on the basis of available data. The expert's knowledge, concerning data, hypotheses and rules of inference, is expressed in terms of statistical weights, certainty coefficients, fuzzy values, etc. Examples of this approach are: MYCIN [3] based on a quantitative theory of change of belief; MILORD [22] and CADIAG-2 [1] based on fuzzy logic; OASES [2] based on the Dempster-Shafer possibility theory; POSCON [21] based on subjective bayesian probability.

To simulate expert reasoning, these approaches need laborious mathematical computations and deductive inferences. In contrast, our approach is based on qualitative considerations (mainly about fact-fact and fact-hypothesis relationships) and follows an abductive inferential strategy (in order to advance explanatory hypotheses).

The case studies reported here suggest more general observations on the treatment of uncertainty.

Poorly structured knowledge domains can hardly be treated on the basis of frequentistic, probabilistic or fuzzy approaches. The medical and architectural experts we have consulted were unable to give a numerical evaluation of the uncertainty affecting the elements of the knowledge base. For example, in the architectural domain the problem is that of reconstructing the history of the building on the basis of observed architectonic characters, and to verify the plausibility of this reconstruction investigating its compatibility with typological, historical and architectonic knowledge. Such knowledge can only be expressed in qualitative terms. In the medical domain, the qualitative evaluation of diagnostic hypotheses is based on a total ordering, given by the experts, of the set of signs. This ordering represents the relevance of a sign with respect to a diagnostic hypothesis. For example, the expert generally regards an anamnestic sign as less relevant than a symptom or a laboratory test. The total ordering of the set of signs induces an ordering onto the set of diseases. The final diagnosis, proposed by the expert system, is represented by an ordered list of diseases. The first element of the list represents the highest ranked hypothesis.

In the case of *partially structured knowledge*, our experience on the diagnosis of building wearout suggests that a numerical evaluation of the importance of some damages (for example, the width of a crack in the wall and so on) may be useful for the elaboration of a pre-diagnosis, whereas the refinement process of the causes belonging to the pre-diagnosis is made again by means of purely qualitative reasoning. In this domain, the expert, for each damage and for every step of the augmentation process, is capable of isolating the possible causes and their (primary or secondary) role in the pre-diagnosis. The step from pre-diagnosis to final diagnosis is performed by a query process in which the expert system can reassess some diagnostic hypotheses through a refinement process and qualitative ranking.

The architectural system indicates the usefulness of abductive and predictive reasoning also in non-diagnostic settings. In fact, this classification system describes a situation rather than looking for an explanatory cause of observed anomalies. In this domain, we have not settled yet on an adequate knowledge representation, for a novel and complicated phenomenology has emerged: in the coherence analysis, for example, technical knowledge and historical data merge, and reasoning based both on deductive logic and common sense are involved. Only deep analyses of expert behaviour in critical cases may help us overcoming conflicts between classification hypotheses.

We have discussed here only two of the three different levels of knowledge structure in diagnostic expert system domains [9]. The third domain is characterized by *fully structured knowledge*: well-developed, stable theories and models for every aspect of the behaviour of the objects are available there. This condition applies, for example, to many human artifacts which, moreover, have a low degree of complexity (such as some mechanical devices or the circuitry of some electronic machines). An appropriate knowledge base for the diagnosis of their malfunctioning may be designed solely on the basis of the theories and models available. Detailed protocols of the stepwise inferential processes leading to diagnoses can be obtained from the experts, resulting in the implementation of a system capable of planning a sophisticated diagnostic strategy, which includes the achievement of subgoals and interactions with the environment.

For better understanding the differences between *fully structured knowledge* and *poorly or partially structured knowledge* let us remind that for domain in which the knowledge is well structured in general no uncertainty arises on the causal relations while some uncertainties may happen on the presence or absence of specific signs. In such a case statistics or large collection of data are available and then bayesian or certainty factors models may be applied. For these reasons it seems evident that qualitative reasoning models are inappropriate.

In conclusion we wish to remark that the main applications of Expert Systems fall into the *poorly or partially structured knowledge* domains. Due to the formlessness of those domains there are neither mathematical nor logical models to describe them. Only the expertise of few human beings may be helpful. This expertise does not accomplish quantitative measure of uncertainty but simply a qualitative description. Starting from this remark we proposed a qualitative paradigm which seems to match well enough the natural way of expert's thinking.

Acknowledgement

The authors are indebted to Guglielmo Tamburrini for helpful comments and criticisms.

References

[1] Adlassnig, K.P., Grabner, G., The Vienna Medical Expert System in the Hospital and on a Personal Computer, in [11], (1987), 95-107.

[2] Biswas, G., & Abramczyk, R., & Oliff, M. (1987). OASES: An Expert System for Operations Analysis - The System for Cause Analysis, IEEE Transactions on Systems, Man, and Cybernetics, vol. 17, no. 2, pp. 133-1125.

[3] Buchanan, B.G., & Shortliffe, E.H., *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, (Addison-Wesley Publishing Company, 1984).

[4] Buchler J., *Philosophical writings of Peirce* (Dover Publication Inc., New York, 1987).

[5] Burattini E., Criscuolo G., Filosa A., Giani U., Mele F., SIGN: a question answer non-numerical abductive expert system, in: *Proc. of the 9th European Meeting on Cybernetics and Systems Research*, Vienna, Austria, (1988) 467–473.

[6] Burattini E., De Glas M., De Gregorio M., A symbolic model for the representation of uncertain knowledge. Theory and applications, in: *Proc. of the IEEE Conference on Systems, Man and Cybernetics*, Chicago, Illinois, USA, (October 18–21, 1992) 272–278.

[7] Burattini E., De Gregorio M., Pagano M., An Expert System for the Diagnosis of Building Statical Wearout, in: *Proc. of the World Congress on Expert Systems*, Orlando, Florida, USA, (December 16–19, 1991) 1941–1949.

[8] Burattini E., Tamburrini G., A pseudo-neural system for hypothesis selection, *Int. J. of Intelligent Systems* **7** (1992) 521–545.

[9] Burattini E., Tamburrini G., Neural parallel processing and diagnostic expert systems, in: Marinaro M., Scarpetta G., eds., *Structure: from Physics to General Systems*, Festschrift volume in Honour of E.R. Caianiello on his seventieth birthday, (World Scientific Publishers, Singapore, 1992) 263–274.

[10] Charniak E., McDermott D., *Introduction to Artificial Intelligence* (Addison-Wesley, Reading, Mass., 1987) chap. 8.

[11] Chytil, M. K., Engelbrecht, R., *Medical Expert Systems using personal computer* (Sigma Press, Wilmslow, Cheshire, U. K., 1987).

[12] Console L., Torasso P., Integrating Models of the Correct Behaviour into Abductive Diagnosis, in: *Proc. of the 9th ECAI*, Stockholm, Sweden, (1990) 160–166.

[13] De Gregorio M., *Un Sistema Esperto per la diagnosi in gastroenterologia pediatrica,* Tesi di Laurea in Fisica, Università degli Studi di Napoli (1990).

[14] Eshghi K., Kowalski R.A., Abduction compared with negation as failure, *Proc. 6th International Conference on Logic Programming*, Lisbon, Portugal, (1989) 234-255.

[15] Hintikka J., Hintikka M. B., Sherlock Holmes confronts Modern Logic, in: Barth E. M., Martens J. L., eds., *Argumentation: Approaches to Theory Formation Containing the Contributions to the Groningen Conference on the Theory of Argumentation* (Benjamins, Amsterdam, 1978).

[16] Konolige K., Abduction Versus Closure in Causal Theories, *Artificial Intelligence* **53** (1992) 225–272.

[17] Polya G., *Mathematics and Plausible Reasoning*, vol. II, Patterns of Plausible Reasoning (Princeton Univ. Press, 1987).

[18] Reggia J. A., Nau D. S., Wang P. Y., Diagnostic Expert Systems Based on a Set Covering Model, *Int. J. Man-Machine Studies* **19** (1983) 437–460.

[19] Reggia J. A., Peng Y., Parsimonious Covering Theory and Non-Diagnostic Tasks, in *Automated Abduction* (Working Notes), 1990 Spring Simposium Series, Stanford University, March 27-29, 120-124.

[20] Reiter R., A Theory of Diagnosis from First Principles, *Artificial Intelligence* **32** (1987) 57–96.

[21] Schaafsma, W., & Ambergen, A.W., POSCON - a Probabilistic Approach to Medical Expert Systems, in [12], (1987), 175-186.

[22] Sierra G. C. A., *MILORD: arquitectura multi-nivell per a sistemes experts en classificacio*, Tesi Doctoral, Universitat Politecnica de Catalunya, España (1989).

[23] Torasso P., Console L., *Diagnostic problem solving* (North Oxford Academic, London 1989).

ERROR: undefined
OFFENDING COMMAND:

STACK: