

Medical diagnosis as pattern recognition in a framework of information compression by multiple alignment, unification and search

J. Gerard Wolff *

CognitionResearch.org.uk, Menai Bridge, LL59 5LR, UK

Available online 17 May 2005

Abstract

This paper describes a novel approach to medical diagnosis based on the SP theory of computing and cognition. The main attractions of this approach are: a format for representing diseases that is simple and intuitive; an ability to cope with errors and uncertainties in diagnostic information; the simplicity of storing statistical information as frequencies of occurrence of diseases; a method for evaluating alternative diagnostic hypotheses that yields true probabilities; and a framework that should facilitate unsupervised learning of medical knowledge and the integration of medical diagnosis with other AI applications.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Medical diagnosis; Information compression; Multiple alignment; SP theory; Pattern recognition; Causal reasoning

1. Introduction

The problem of providing computational support for medical diagnosis has been approached from many directions including logical reasoning, fuzzy logic, set theory, rough set theory, if-then rules, Bayesian networks, classical parametric and non-parametric statistics, artificial neural networks, case-based reasoning, support vector machines, perceptrons, possibility theory, and more, as well as various aggregations or combinations of methods [14].

This paper describes a novel approach to diagnosis based on the SP theory of computing and cognition (described below). The main attractions of this approach are:

- A format for representing diseases that is simple and intuitive.
- An ability to cope with errors and uncertainties in diagnostic information.
- The simplicity of storing statistical information as frequencies rather than conditional probabilities.
- A method for evaluating alternative diagnostic hypotheses that yields true probabilities.
- A framework that should facilitate unsupervised learning of medical knowledge and the integration of medical diagnosis with other AI applications.

* Tel.: +44 1248 712962.

E-mail address: jgw@cognitionresearch.org.uk.

URL: <http://www.cognitionresearch.org.uk/sp.htm>.

It must be stressed that the primary purpose of this paper is conceptual: to describe an approach to medical diagnosis that is significantly different from the main alternatives and with potential advantages compared with those alternatives. Although a prototype of the proposed new system exists, it is not yet a shrink-wrapped software system that is ready for immediate application. Like any other diagnostic system, the SP system must be equipped with a body of relevant knowledge and the creation of such a body of knowledge (by automatic learning or by ‘knowledge elicitation’ from experts) is a major undertaking in its own right.

Key elements of the SP theory are first described, just sufficient for present purposes. Section 3 describes how the theory may be applied to medical diagnosis, viewed as a process of pattern recognition. This section also discusses how the SP system relates to several aspects of the diagnostic process, including causal reasoning and the process of acquiring the knowledge that is needed for accurate diagnosis. Section 4 compares this new approach to medical diagnosis with some of the alternatives. The paper concludes with an outline of what still needs to be done in this programme of research and with a review of the main points that have been made.

2. The SP theory

The SP theory grew out of a long tradition in psychology that many aspects of brain function may be understood as information compression (see, for example, Refs. [2,3,24,7]). It is based on principles of *minimum length encoding*¹ pioneered by Solomonoff [21], Wallace and Boulton [23], Rissanen [17] and others (see also Ref. [13]). An overview of the theory is presented in Ref. [31] and more detail may be found in other papers cited there (see also Ref. [28]).

The SP theory has been developed as an abstract model of *any* system for processing information, either natural or artificial. In broad terms, the system receives ‘New’ information from its environment and transfers it to a repository of ‘Old’ information. At the same

time, it tries to compress the information as much as possible by finding patterns that match each other and merging or ‘unifying’ patterns that are the same.² An important part of this process is the building of ‘multiple alignments’ as described below.

The SP framework is Turing-equivalent in the sense that it can model a universal Turing machine [25] but it has much more to say about the nature of ‘intelligence’ than the Turing model of computing (or equivalent models such as lambda calculus [9] or the Post canonical system [16]).

To date, the main areas in which the SP framework has been applied are probabilistic reasoning, pattern recognition and information retrieval [26], parsing and production of natural language [27], modelling concepts in logic and mathematics [29], and unsupervised learning [32,30].

2.1. Computer models

Two computer models of the SP system have been developed:

- SP62 is a partial realisation of the theory that does not transfer any information from New to Old. This model tries to compress the New information in terms of the Old information by building multiple alignments of the kind that will be seen below. SP62 also contains procedures for calculating the probabilities of inferences that may be drawn from alignments. A slightly earlier version of this model (SP61) is described quite fully in Ref. [27]. Both versions are relatively robust and mature.
- SP70 realises all the main elements of the theory, including the transfer of information from New to Old. In addition to building multiple alignments like SP62, the model compiles one or more alternative ‘grammars’ for the information in New, using principles of minimum length encoding. This model, and its application to unsupervised learning, is described quite fully in Refs. [32,30]. More work is required to realise the full potential of this model.

¹ An umbrella term for ‘minimum message length encoding’ and ‘minimum description length encoding’.

² The term ‘unification’ in the SP theory means a simple merging of two or more identical patterns to make one. This meaning is different from but related to the meaning of the term in logic.

2.2. Representation of knowledge

In the SP system, all kinds of knowledge are stored as arrays of atomic *symbols* in one or two dimensions called *patterns*. In work to date, the main focus has been on one-dimensional patterns (i.e., sequences of symbols) but it is envisaged that, at some stage, the concepts will be generalised to patterns in two dimensions.

For present purposes, we may define patterns and symbols as follows:

- A *pattern* is a sequence of symbols bounded by end-of-pattern characters such as ‘(’ and ‘)’, not shown in the examples in this paper.
- A *symbol* is a string of non-space characters bounded by white space (space characters, line-feed characters and the like).
- Any symbol can be matched with any other symbol and, for any one pair of symbols, the two symbols are either ‘the same’ or ‘different’. No other result is permitted.
- Symbols have no intrinsic meaning such as ‘add’ for the symbol ‘+’ in arithmetic or ‘multiply’ for the symbol ‘×’. Any meaning attaching to an SP symbol takes the form of one or more other symbols with which it is associated in a given set of patterns.
- Each pattern has an associated integer value representing the absolute or relative frequency of occurrence of that pattern in some domain.

Despite the extraordinary simplicity of this format for representing knowledge, the way in which SP patterns are processed within the system means that they can model a wide variety of established representational schemes, including context-free and context-sensitive grammars, class-inclusion hierarchies,

part-whole hierarchies, discrimination networks and trees, if-then rules, and others.

2.3. Processing knowledge

A key part of the process of matching patterns is the building of ‘multiple alignments’, described and illustrated here. The process of building multiple alignments in the SP system provides a unified model for a variety of computational effects including fuzzy pattern recognition, best-match information retrieval, probabilistic and exact styles of reasoning, unsupervised learning, planning, problem solving and others, as described in Ref. [31].

2.3.1. Multiple alignments

This subsection and the ones that follow describe the main elements of the multiple alignment concept as it has been developed in the SP theory and explains how multiple alignments are created and evaluated in the SP system.

In bioinformatics, a multiple alignment is an arrangement of two or more DNA sequences or sequences of amino acid residues so that matching symbols are aligned. Fig. 1 shows a typical example. The general idea is that, by judicious ‘stretching’ of sequences, as many symbols as possible are aligned with each other. A variety of measures of the ‘goodness’ of alignments are used but they all tend to favour alignments where the number of aligned symbols is high and the gaps between them are relatively few and relatively small.

In the SP framework, the concept of multiple alignment has been modified as follows:

- One or more of the sequences (termed *patterns*, as described in Section 2.2) are classified as ‘New’ and the rest are ‘Old’.



Fig. 1. A ‘good’ alignment amongst five DNA sequences.

```

0      e p e r i m a n t a t x p i u n      0
      | | | | | | | | | |
1 < E3 e x p e r i m e n t a t      i o n > 1
    
```

Fig. 2. The best alignment formed by SP62 when it is supplied with one New pattern ('e p e r i m a n t a t x p i u n') and a dictionary of Old patterns, each one of which represents a correctly-spelled word.

- A 'good' alignment is one where the New patterns can be encoded economically in terms of the Old patterns in the alignment, as will be explained.
- Any one pattern may appear more than once in one alignment, not as two or more *copies* but as two or more *appearances* of a *single* pattern. This has implications for the way alignments are formed and for the representation of recursive structures. These aspects of the multiple alignment concept are not relevant to the main proposals here and will not be considered further in this paper. Readers who wish to know more may consult the sources cited earlier.

Normally, the SP62 model is run with relatively few New patterns and a relatively large 'database' or 'dictionary' of Old patterns. The system typically forms a set of *alternative* alignments, each one of which represents a possible encoding of the New pattern or patterns in terms of one or more of the Old patterns.

Fig. 2 shows a simple example, with one New pattern (in row 0, representing a badly-spelled version of the word 'experimentation') and one Old pattern (in row 1, representing the correctly-spelled version of the word).³ This is the best alignment produced by SP62 with the New pattern as shown and a dictionary of Old patterns, each one of which represents one word in its correctly-spelled form.

Fig. 3 is a slightly more complicated example, the best alignment produced by SP62 when it is supplied with one New pattern ('j o h n r u n s') and a set of Old patterns, each one of which represents a gram-

matical rule. This alignment shows how the sentence 'j o h n r u n s' (in row 0) may be analysed ('parsed') into its parts. The Old patterns in rows 1 to 3 represent grammatical rules: '< S < N > < V > >' in row 3 means that a (simple) sentence is composed of a noun ('N') followed by a verb ('V'), '< N 0 j o h n >' in row 2 means that 'j o h n' is a noun, and '< V 1 r u n s >' in row 1 means that 'r u n s' is a verb.

2.3.2. Evaluation of alignments

As previously mentioned, a 'good' multiple alignment is one where the New pattern or patterns can be encoded economically in terms of the Old patterns in the alignment. How is this evaluation done?

Every symbol has an associated 'weight' which is the number of bits needed to encode that symbol. And the weight is derived from the frequency of occurrence of the symbol using the Shannon–Fano–Elias method (see Ref. [10]) (a method that is similar to the well-known Huffman method). The frequency value for any symbol is derived from the frequency value of the pattern (or patterns) in which that symbol appears (as described in Section 2.2). In the context of medical diagnoses, the frequency associated with any given pattern is the frequency of occurrence of the disease that is represented by that pattern (see Section 3.1, below).

A few of the symbols (normally one or two) near the beginning of each Old pattern are classified as *identification* symbols or 'ID-symbols'. For example, the ID-symbol in the pattern '< E3 e x p e r i m e n t a t i o n >' in Fig. 2 is 'E3', and the ID-symbols in the pattern '< N 0 j o h n >' in Fig. 3 are 'N' and '0'.

A 'code' for any alignment may be derived quite simply by scanning the alignment from left to right looking for columns in the alignment that contain one code symbol, not matched with any other identical

```

0          j o h n          r u n s      0
      | | | | | | | | | |
1          | | | | | | | | | |
      | | | | | | | | | |
2 < N 0 j o h n > | | | | | | | | | | 2
      | | | | | | | | | |
3 < S < N          > < V          > > 3
    
```

Fig. 3. The best alignment formed by SP62 when it is supplied with one New pattern ('j o h n r u n s') and a set of Old patterns, each one of which represents a grammatical rule.

³ By convention, the New pattern or patterns are always shown in row 0 of alignments like those shown in Figs. 2 and 3, and the Old patterns are shown in the other rows, one pattern per row and in an order that is entirely arbitrary, without special significance. As we shall see, alignments can sometimes fit better on the page if they are rotated by 90° and in this case the New pattern or patterns are shown in column 0 with the Old patterns in the other columns, one pattern per column.

symbol. The code for the alignment is the sequence of symbols that have been found, in the given order. For example, the code derived in this way from the alignment in Fig. 2 is ‘E3’ and the code derived from the alignment in Fig. 3 is ‘S 0 1’.

The next step in the evaluation of a given alignment is the calculation of a ‘compression score’ or ‘compression difference’ as:

$$CD = B_n - B_c, \quad (1)$$

where B_n is the total size (in bits) of those symbols within the New pattern that have been matched to Old symbols within the alignment, and B_c is the total size (in bits) of the symbols in the code that has been derived as just explained. Adjustments to this score are made to take account of gaps in the alignment like those that may be seen in Fig. 2. The details of these calculations and adjustments are explained in Ref. [27].

2.3.3. The building of multiple alignments

In bioinformatics, it is generally understood that the abstract ‘space’ of alternative possible alignments between two or more sequences is, with few exceptions, astronomically large—which means that it cannot be searched exhaustively. All practical methods for finding ‘good’ alignments amongst two or more sequences use heuristic methods such as ‘hill climbing’ (or ‘descent’), ‘beam search’, ‘genetic algorithms’ or the like that search selectively and exclude large parts of the search space. With methods like these, one can find good approximate solutions in a reasonable time but one can never be sure of finding the best possible solution (unless the sequences are very short and very few). Finding good multiple alignments in the SP system is no different.

At the heart of the SP system for building multiple alignments is an improved version of ‘dynamic programming’ for finding full matches and good partial matches between two patterns (see, for example, Ref. [18]). Unlike standard versions of dynamic programming, the procedure used in the SP models:

- Can find good matches between patterns without restrictions on the lengths of the patterns.
- Can normally find several *alternative* alignments between two patterns, not just one.

- It allows the ‘depth’ or thoroughness of searching to be varied according to need.

Given one New pattern and a database of Old patterns, SP62 first builds a set of alignments, each one of which is between the New pattern and one of the Old patterns. From this set, it selects the best few alignments, using the measure described in Section 2.3.2. Each of these alignments can itself be treated as if it was a single pattern. So, in the next stage, SP62 builds larger alignments, each one of which is between one of the selected alignments and one of the Old patterns or between one of the selected alignments and another of those alignments. As before, the program selects the best of the alignments that have been formed.

The process is repeated in this way until no more alignments can be found. The process for building alignments containing two or more New patterns is a generalisation of what has been described here.

2.3.4. Unsupervised learning

At its most abstract level (Section 2), the SP model is conceived as a system that learns by transferring New information to its repository of Old information and compressing it at the same time. This abstract conception has now been realised more concretely in the form of the SP70 computer model [32,30] that is capable of learning simple grammars from raw data. However, further development of the model is needed to realise its full potential.

The current model has two stages:

- (1) From partial alignments between patterns, the model creates new patterns that are added to the repository of Old patterns as explained below.
- (2) Amongst the patterns that are generated in this way, some are ‘good’ in terms of the principles of minimum length encoding and others are ‘bad’. In the second stage of processing, the model measures the frequency with which each pattern may be recognised in the raw data and then it uses this information in a hill-climbing search amongst subsets of the Old patterns to find one or more sets of patterns that are good in terms of the principles of minimum length encoding. The remaining patterns may be discarded.

It is envisaged that, when the model is more fully developed, these two stages will be repeated so that the system can progressively bootstrap a set of patterns that are good in terms of the principles of minimum length encoding, and thus represent a distillation of the patterns of redundancy in the original data.

SP70 is currently targeted at the learning of syntax in natural languages. Given a partial alignment like this:

```

0      t h e g i r l r u n s   0
      | | | | | | | | |
1 < %1 t h e b o y   r u n s > 1

```

the program creates patterns like these:

```

< %2 t h e >
< %3 r u n s >
< %4 b o y >
< %4 g i r l >
< %5 < %2 > < %4 > < %3 > >

```

The first four are derived from coherent sequences of matched symbols and coherent sequences of unmatched symbols in the alignment and they correspond to what we would normally recognise as words. Each one has a grammatical category represented by ID-symbols such as '%2', '%3' and '%4'. Notice that 'b o y' and 'g i r l' belong to the same disjunctive category {'b o y', 'g i r l'} because they are alternatives at the same point in the original alignment and they both share the ID-symbol '%4'. The ID-symbols '0' and '1' serve to distinguish the two alternatives in that category.

The last pattern in this example ties everything together by listing the sequence of categories in the original alignment. It is an 'abstract' pattern describing the overall structure of the two original sentences.

For an application like medical diagnosis, the style of learning just described is probably not entirely appropriate. Section 3.11.2 describes how similar principles may be applied to medical data. If the potential of these ideas can be realised, the SP system should facilitate the automatic or semi-automatic construction of knowledge bases from raw medical data.

2.3.5. Computational complexity

The time complexity of the SP62 model in a serial processing environment is approximately $O(\log_2$

$n \times nm)$, where n is the size of the New pattern or patterns (in bits) and m is the total size of the patterns in Old (in bits). In a parallel processing environment, the time complexity may approach $O(\log_2 n \times n)$, depending on how well the parallel processing is applied. The space complexity in serial or parallel environments is $O(m)$. Further details may be found in Ref. [27].

In medical diagnosis, it seems reasonable to suppose that there will normally be a fairly small maximum for the number of signs and symptoms (abbreviated hereinafter as 'symptoms') exhibited by any one patient. Correspondingly, there should be a maximum size for the size of the set of New patterns that are used to represent the patient's symptoms. If we take this to be a constant value for n , then in a serial processing environment the time complexity is approximately $O(m)$ and in a parallel processing environment it may approach $O(1)$.

3. Application of the SP system to medical diagnosis

To a large extent, medical diagnosis may be viewed as a problem of (fuzzy) pattern recognition: finding the best fit between a given set of symptoms for an individual patient and the symptoms associated with one or more diseases. However, causal reasoning also has a part to play when, for example, it is understood that a given disease is caused by a bacterial or virus infection.

This section presents an example showing how the SP system may be applied to medical diagnosis, viewed as a process of pattern recognition. The system may also support causal reasoning about medical problems and this is discussed briefly in Section 3.9, below.

Other aspects of the proposals are discussed in other subsections.

3.1. Describing diseases using SP patterns

In the SP scheme, knowledge about diseases may be stored as patterns in a repository of Old information, and the symptoms for an individual patient may be represented as a set of one or more New patterns.

A pattern in the store of Old information may represent one disease and its associated symptoms, or a combination of diseases (see Section 3.7.2, below), or

it may represent a cluster of symptoms that tend to occur together in two or more different diseases (see Section 3.5, below). In addition, Old may include patterns that play supporting rôles (see Section 3.6, below).

The frequency value of each pattern may be used to represent the absolute or relative frequency with which a given disease or cluster of symptoms is found in a given population. These figures may be derived from epidemiological surveys or they may be estimated by medical experts.

By way of illustration, Fig. 4 shows five examples of such patterns, one describing the symptoms of chicken pox, another describing the symptoms of smallpox, the third describing the cluster of symptoms which is described as ‘fever’ and two more describing the class of ‘high’ temperatures (38–39 °C and 40+ °C). The number in brackets after each pattern is a very rough estimate of the relative frequencies of occurrence of the corresponding disease or condition.⁴

Each of the first three patterns begins and ends with a pair of symbols ‘<disease>...</disease>’ which indicate that the pattern describes a disease or a cluster of disease symptoms. Within each pattern, there are similar pairs of symbols, each one marking the beginning and end of a ‘field’ which describes some aspect of the disease or cluster. For example, ‘<dname> Chicken_Pox </dname>’ provides the name of the chicken pox disease, ‘<skin> rash </skin>’ describes one of its symptoms, ‘<causative_agent> chicken_pox_virus </causative_agent>’ describes what causes the disease, and ‘<treatment> chicken_pox_treatment </treatment>’ is a remarkably unhelpful description of how to treat the disease which would, of course, be much more detailed in a fully developed knowledge base.

Within the pattern for chicken pox, the field ‘<R2> fever </R2>’ indicates that ‘fever’ is one of the symptoms of the disease. However, by contrast with other fields like those just mentioned, the symbol ‘fever’ is, in effect, a reference or pointer to a cluster of symptoms such as rapid breathing, flushed face and high temperature described in the third pattern in the same figure. In a similar way, ‘<R1> flu_symptoms

</R1>’ in the pattern for smallpox is a reference or pointer to another pattern, not shown in the figure, that describes a cluster of symptoms associated with influenza and flu-like diseases. The way in which pointers like these are dereferenced in the SP system will be seen in the next section.

Readers who are familiar with XML [6] will notice that pairs of symbols like ‘<disease> ... </disease>’ or ‘<dname> ... </dname>’ are rather like the start and end tags used to mark the elements of an XML document. However, by contrast with XML and related languages such as HTML, symbols of that kind have no formal status in the SP system and the styles of symbols are not defined within the system. Any convenient style may be used such as ‘disease ... #disease’ or ‘disease ... %disease’ and in some applications it is not necessary to provide any distinctive markers for the beginnings and ends of patterns or fields. The concept of ‘field’ has no formal status in the SP system either.

3.2. Multiple alignment and medical diagnosis

The process of diagnosis may be modelled by the building of one or more multiple alignments. Fig. 6 shows the best alignment created by SP62 with a set of New patterns shown in Fig. 5 that describe ‘John Smith’ and his symptoms and a set of Old patterns like those shown in Fig. 4 that represent diseases or aspects of diseases.⁵

An alignment like this may be interpreted as the result of a process of recognition. In this case, the symptoms that have been recognised are those of influenza, as shown in column 2. The following subsections discuss aspects of the alignment and of this interpretation.

3.3. A framework pattern

In an application like this, it is convenient but not essential to include amongst the Old patterns a

⁴ The figure for smallpox is clearly too high in the world today but it will serve for the purpose of illustration.

⁵ Compared with the alignments shown in Figs. 2 and 3, the alignment in Fig. 6 has been rotated by 90° to allow the alignment to fit better on the page. As previously noted, the New patterns are shown in column 0 and the Old patterns are shown in the other columns, one pattern per column in an order that is arbitrary and without special significance.

```

<disease> chicken_pox :
  <lname> Chicken_Pox </lname>
  <R2> fever </R2>
  <appetite> normal </appetite>
  <chest> normal </chest>
  <chills> no </chills>
  <cough> no </cough>
  <diarrhoea> no </diarrhoea>
  <fatigue> no </fatigue>
  <lymph_nodes> normal </lymph_nodes>
  <malaise> yes </malaise>
  <muscles> normal </muscles>
  <nose> normal </nose>
  <skin> rash </skin>
  <throat> normal </throat>
  <weight_change> no </weight_change>
  <causative_agent> chicken_pox_virus </causative_agent>
  <treatment> chicken_pox_treatment </treatment>
</disease> (2500)

<disease> smpx :
  <lname> Smallpox </lname>
  <R1> flu_symptoms </R1>
  <appetite> normal </appetite>
  <chest> normal </chest>
  <diarrhoea> no </diarrhoea>
  <fatigue> no </fatigue>
  <lymph_nodes> normal </lymph_nodes>
  <malaise> no </malaise>
  <skin> rash with blisters </skin>
  <weight_change> no </weight_change>
  <causative_agent> smallpox_virus </causative_agent>
  <treatment> smallpox_treatment </treatment>
</disease> (5)

<disease> fever
  <breathing> rapid </breathing>
  <face> flushed </face>
  <temperature> <t1> </t1> </temperature>
</disease> (15000)

<t1> 38-39 </t1> (14705)
<t1> 40+ </t1> (147)

```

Fig. 4. Five SP patterns, one describing the symptoms of chicken pox, another describing the symptoms of smallpox, the third describing the symptoms of fever and two more describing the class of 'high' temperatures.

'framework' pattern like the one shown in column 1. This is a generalised pattern for diseases of all kinds that lists the main categories associated with diseases such as '<lname> </lname>' (the name

of the disease), '<breathing> </breathing>' (the state of the patient's breathing) and '<temperature> </temperature>' (the patient's temperature), but it does not specify specific values for any category.


```

<patient> John_Smith </patient>
<face> flushed </face>
<appetite> poor </appetite>
<breathing> rapid </breathing>
<muscles> aching </muscles>
<chills> yes </chills>
<fatigue> yes </fatigue>
<lymph_nodes> normal </lymph_nodes>
<malaise> no </malaise>
<nose> runny </nose>
<temperature> 38-39 </temperature>
<throat> sore </throat>

```

Fig. 5. The set of New patterns supplied to SP62 for the example discussed in the text. These patterns represent the patient ‘John Smith’ and his symptoms.

This framework pattern serves as an anchor point for symbols in other patterns and facilitates the formation of multiple alignments in accordance with the rules described in Ref. [31] and earlier publications.

3.4. The ordering of descriptors

In an application like medical diagnosis, it is not obvious that there is any intrinsic order to the symptoms of a disease or associated descriptors such as the name of the patient. In describing a patient’s symptoms, it should make no difference whether ‘high temperature’ is mentioned before ‘runny nose’ or the other way round.

In the SP framework, each pattern that describes a disease or a cluster of symptoms necessarily imposes an order in which categories of descriptors are specified. However, users of the system may specify the patient’s symptoms in any order that is convenient. This is because symptoms are described using a set of New patterns and there is no intrinsic order amongst the New patterns supplied to the system. In our example, New patterns were supplied to the SP62 model in the order shown in Fig. 5 but in the alignment shown in Fig. 6 they appear in a completely different order.

Notice that this freedom in the ordering of descriptors only applies to whole patterns. When two or more symbols in one pattern are matched to two or more symbols in another, the order of the symbols in one pattern must be the same as the order of the matching symbols in the other pattern.

3.5. Dereferencing of pointers

As already noted, a symbol like ‘fever’ or ‘flu_symptoms’ in one pattern may serve as a reference or pointer to another pattern that describes a cluster of symptoms that may be found in two or more different diseases.

In Fig. 6, we can see how such pointers are ‘dereferenced’ in the SP system. The symbol ‘flu_symptoms’ in column 2 is matched to the same symbol in column 3 where flu-like symptoms are listed. Likewise, the symbol ‘fever’ in column 3 is matched to the same symbol in column 4 where the symptoms of fever are listed. Fever is itself part of the cluster of flu-like symptoms.

The provision of named clusters like these saves the need to specify the corresponding symptoms redundantly in each of the diseases where such clusters appear.

3.6. Uncertainties in diagnosis

Diagnosis is not an exact process:

- Most diseases are ‘family resemblance’ or ‘polythetic’ concepts because the majority of symptoms associated with any given disease are neither necessary nor sufficient for the diagnosis of the disease: they are ‘characteristic’ of the disease in the sense that any one such symptom need not be present in every case and any of them may be associated with other diseases.
- There may be and frequently are errors in the observation or recording of symptoms.

SP62 can accommodate these kinds of uncertainty in diagnosis in two distinct ways:

- Because it looks for a global best match amongst patterns, it does not depend on the presence or absence of any particular symptom. Notice how SP62 has succeeded in constructing the alignment shown in Fig. 6 despite there being no match for ‘poor’ in the New pattern ‘<appetite> poor </appetite>’ and ‘yes’ in the New pattern ‘<fatigue> yes </fatigue>’ and no match for many of the symbols in the Old patterns.

```

0          1          2          3          4          5
          <disease> ----- <disease> ----- <disease> ---- <disease>
          flu
          : ----- :
<patient> ----- <patient>
John_Smith
</patient> ----- </patient>
          <dname> ----- <dname>
          Influenza
          </dname> ----- </dname>
          <R1> ----- <R1>
          flu_symptoms ----- flu_symptoms
          </R1> ----- </R1>
          <R2> ----- <R2>
          fever ----- fever
          </R2> ----- </R2>
<appetite> ----- <appetite>
poor ----- normal
</appetite> ----- </appetite>
<breathing> ----- <breathing> ----- <breathing>
rapid ----- rapid
</breathing> ----- </breathing>
          <chest> ----- <chest>
          normal
          </chest> ----- </chest>
<chills> ----- <chills> ----- <chills>
yes ----- yes
</chills> ----- </chills>
          <cough> ----- <cough>
          yes
          </cough> ----- </cough>
          <diarrhoea> ----- <diarrhoea>
          no
          </diarrhoea> ----- </diarrhoea>
<face> ----- <face> ----- <face>
flushed ----- flushed
</face> ----- </face>
<fatigue> ----- <fatigue> ----- <fatigue>
yes ----- no
</fatigue> ----- </fatigue>
          <headache> ----- <headache>
          yes
          </headache> ----- </headache>
<lymph_nodes> -- <lymph_nodes> ----- <lymph_nodes>
normal ----- normal
</lymph_nodes> - </lymph_nodes> ----- </lymph_nodes>
<malaise> ----- <malaise> ----- <malaise>
no ----- no
</malaise> ----- </malaise>
<muscles> ----- <muscles> ----- <muscles>
aching ----- aching
</muscles> ----- </muscles>
<nose> ----- <nose> ----- <nose>
runny ----- runny
</nose> ----- </nose>
          <skin> ----- <skin>
          normal
          </skin> ----- </skin>
<temperature> -- <temperature> ----- <temperature>
38-39 ----- <t1> ----- <t1>
          ----- 38-39
          </t1> ----- </t1>
</temperature> - </temperature> ----- </temperature>
<throat> ----- <throat> ----- <throat>
sore ----- sore
</throat> ----- </throat>
          <weight_change> ----- <weight_change>
          no
          </weight_change> ----- </weight_change>
          <causative_agent> ----- <causative_agent>
          flu_virus
          </causative_agent> ----- </causative_agent>
          <treatment> ----- <treatment>
          flu_treatment
          </treatment> ----- </treatment>
          </disease> ----- </disease> ----- </disease> --- </disease>
0          1          2          3          4          5

```

Fig. 6. The best alignment found by SP62 with the set of patterns from Fig. 5 in New (describing the symptoms of the patient ‘John Smith’) and a set of patterns in Old describing a range of different diseases and named clusters of symptoms, together with the ‘framework’ pattern shown in column 1.

Although the system does not depend on the presence or absence of any one symptom, particular symptoms can have a major impact on diagnosis, as described in Section 3.7.1, below.

- Within the SP framework, it is not necessary for every symptom of a disease to be recorded as a specific value. For example, in column 4 of the alignment in Fig. 6, the pair of symbols '<t1> </t1>' represents a set of alternative values for the temperature associated with fever. In this case, there are just two values, represented in Old by the patterns '<t1> 38–39 </t1>' (high temperature) and '<t1> 40+ </t1>' (very high temperature), as shown at the bottom of Fig. 4. The first of these patterns is shown in column 5 of the alignment, matched to the temperature of the patient shown in column 0. Being able to specify symptoms as sets of alternative values allows the system to accommodate the kind of variability which is so prominent in many diseases.

3.7. Weighing alternative hypotheses and the calculation of probabilities

In medical diagnosis, it is quite usual for the physician to consider alternative hypotheses about what disease or diseases the patient may be suffering from. The SP framework provides a model for this process in the way the system builds alternative alignments for any given pattern or set of patterns in New. Alignments—and the corresponding diagnoses—may be evaluated as follows.

As previously noted (Section 2.3.2), a 'compression difference' is calculated for each alignment as shown in Eq. (1). The value B_e that is used in that equation may be translated into an *absolute* probability for the given alignment:

$$P = 2^{-B_e}. \quad (2)$$

For any one alignment (the j th alignment) in a set of alternative alignments, $a_1 \dots a_n$, that encode the same symbols from New, a *relative* probability may be calculated as:

$$p_j = P_j / \sum_{i=1}^{i=n} P_i. \quad (3)$$

A fuller account of the way probabilities are calculated may be found in Ref. [26].

Given that the New patterns represent the symptoms of one patient at a particular time and given that each pattern in Old describes a single disease or a single cluster of symptoms that may form part of the description of one or more diseases, then each alignment formed by SP62 represents an hypothesis about any *one* disease that the patient may have.

Where alternative alignments encode different subsets of the symbols in New, it is possible that the patient may be suffering from two or more diseases at the same time. This possibility is discussed in Section 3.7.2, below. However, where two or more of the best alignments encode exactly the same symbols from New, then they represent *alternative* diagnostic hypotheses and they may be compared using values for relative probability (p).

When SP62 formed the alignment shown in Fig. 6, it also formed a similar alignment, matching exactly the same symbols in New, in which column 2 contained a pattern representing the symptoms of smallpox, instead of the pattern for influenza. The relative probability values calculated in this case were 0.99950 for influenza and 0.00049 for smallpox, reflecting the prevalence of those two diseases in the world today.⁶

3.7.1. Explaining away

The symptoms of influenza and smallpox are quite similar, except for the very distinctive rash and blisters that occur in smallpox. The example shown in Fig. 6 is silent about whether John Smith had a rash and blisters or not. If a rash and blisters had been seen to be absent, this would have been represented as '<skin> normal </skin>'. Given this lack of information about the state of the patient's skin, he may have either influenza or smallpox but he is very much more likely to have the former than the latter, as indicated by the calculated probabilities.

If '<skin> rash_with_blisters </skin>' is added to the symptoms recorded in New, and if SP62 is run again with the augmented set of symptoms, the best

⁶ There are, of course, other factors that may be relevant—such as the possibility that someone might release the smallpox virus deliberately—but in this example, knowledge of such other factors has been excluded.

alignment found by the system is similar to that shown in Fig. 6 but with the pattern for smallpox (the second pattern in Fig. 4) instead of the pattern for influenza in column 2 and with a match shown between ‘<skin> rash_with_blisters </skin>’ in the set of New patterns and the same symbols in the pattern that describes smallpox. However, in this case *there is no other alignment that matches the same symbols in New*. Consequently, the relative probability of the best alignment is 1.0. In short, the addition of one distinctive symptom to the list of symptoms has a dramatic effect on the relative probabilities calculated by the system. Instead of a vanishingly small probability for smallpox (0.00049), the system now assigns it a probability of 1.0, in accordance with our intuitions.

From this result, we may conclude that the patient certainly has smallpox and that his aching muscles and runny nose are due to smallpox, not influenza. This is the phenomenon of ‘explaining away’: “If A implies B, C implies B, and B is true, then finding that C is true makes A *less* credible. In other words, finding a second explanation for an item of data makes the first explanation less credible.” ([15], p. 7), with the emphasis as in the original).

3.7.2. A patient may suffer from two or more diseases at the same time

As noted above, it is possible for a patient to suffer from two or more diseases at the same time. Given that the Old patterns in the system describe single diseases, then the system would create two or more ‘good’ alignments, each one corresponding to one of the diseases that the patient is suffering from.

If we want the system to calculate probabilities for combinations of diseases, then the repository of Old patterns must contain patterns that represent combinations of that kind. Each such pattern may be constructed economically using references to the component diseases, in much the same way that clusters of symptoms may be referenced, as described in Section 3.5.

As with single diseases, frequency values for a combinations of diseases may be obtained from population surveys or by the judgement of medical experts. In the absence of any direct evidence of a statistical association between two or more diseases, it seems reasonable to assume that they are statistically independent. In such cases, frequency values may be

derived straightforwardly via normalised values for the frequencies of occurrence of individual diseases. Whether the frequency values for combinations of diseases are measured, estimated or derived, they can be used for the calculation of *CD* values and probabilities in exactly the same way as for single diseases.

Of course, there are so many possible combinations of diseases that it would be impossible to store information about them all. A more practical option may be to store information in Old about individual diseases and combinations of diseases that are known to have a statistical association with each other. One may assume that all other combinations of diseases are statistically independent.

3.8. Inferences and the diagnostic cycle

In a multiple alignment like the one shown in Fig. 6, any symbol within an Old pattern that is *not* matched to a symbol in New represents an inference that may be drawn from the alignment. In this example, we may infer from the alignment *inter alia* that the patient is likely to have a cough and a headache and that the standard treatment for influenza is required. Probabilities of these inferences can be calculated as described in Ref. [26].

If a ‘good’ alignment makes a prediction about some marker that may be found in the patient’s blood or something that may be observed in an X-ray, this may be interpreted as a suggestion to the physician that he or she should order an appropriate blood test or X-ray. If tests of that kind or other kinds of investigation are instigated as a result of the inferences drawn from preliminary alignments, the results of those investigations, together with the original symptoms, may be fed back into the system as New information. The system may then be run again and the alignments that are created may suggest a final diagnosis or the need for further investigation—and so on.

3.9. Causal reasoning

Apart from the kinds of inference just described, medical diagnosis often seems to involve a ‘deeper’ kind of reasoning about the causes of symptoms and diseases, using knowledge of bacteria, viruses, anatomy, physiology and so on.

The SP framework supports a variety of styles of reasoning, including probabilistic ‘deductive’ reasoning, abductive reasoning, nonmonotonic reasoning and (as we saw in Section 3.7.1) ‘explaining away’ (see Ref. [26]). So there are reasons to believe that, within the SP framework, it may be possible to extend the pattern recognition analysis described above to incorporate causal styles of reasoning.

Recent investigation has confirmed this expectation. The input–output relations of each subsystem within a larger system can be modelled in the SP framework as a set of patterns, and causal connections can be established by matching outputs to inputs. As with the analysis described above, a ‘frame-work’ pattern is also needed to ensure that alignments can be formed in an appropriate manner. These potential applications of the system need further exploration and development.

3.10. Classes and subclasses of diseases

One of the attractions of the SP system is that it allows concepts to be represented at multiple levels of abstraction (e.g., ‘cat’, ‘mammal’, ‘vertebrate’, ‘animal’) in the manner of object-oriented design and, via the building of multiple alignments, it allows a specific entity (such as “my cat Tibs”) to be recognised at several different levels of abstraction [31,26].

To some extent, this idea is already illustrated by the example shown in Fig. 6. The concept of ‘fever’, represented by the pattern in column 4 of the figure, may be seen as a superclass comprising all the diseases where the patient may be feverish. Likewise, the pattern for flu symptoms (column 3 in the figure) may be seen as a superclass of the diseases in which such symptoms may be seen.

By contrast with the classification of animals and plants, the hierarchy of diseases tends to be relatively flat. However, there is scope for the recognition of classes and subclasses in the variants of diseases such as influenza and diabetes. With the SP system, each variant of a given disease may be recorded as a pattern that specifies the symptoms that are characteristic of the variant. Provided that pattern contains a symbolic link to another pattern describing the main symptoms of the disease, there is no need to repeat those symptoms redundantly in each of the variants.

3.11. Acquisition of knowledge

Broadly speaking, the knowledge that is required in any artificial system for medical diagnosis can be obtained ‘manually’ from experts or written sources, or it may be obtained by the automatic or semi-automatic abstraction of knowledge from raw medical data, or some combination of the two. The SP system has potential to facilitate any or all of these processes.

3.11.1. Elicitation of expert knowledge

It should be apparent from the example described above that the SP system provides a means of representing medical knowledge in a form that is simple and intuitive. The simplicity of representing all knowledge as patterns is, perhaps, less important than the fact that this system allows computer-based knowledge to be expressed in a form that apparently reflects the natural structure of the original concepts.

This feature of the system should facilitate traditional kinds of knowledge elicitation from experts or written sources. Medical experts should have little difficulty in expressing their knowledge directly in the form of SP patterns. Given that such experts are often busy and their time is, in any case, expensive, there are advantages if at least some of the process of building computer-based knowledge bases can be undertaken by knowledge engineers without specialised medical training. It should be possible for such people to derive a good deal of the necessary knowledge from medical text books and other written sources.

3.11.2. Automatic or semi-automatic learning

Section 2.3.4 presented an outline description of how the SP70 model learns the kinds of structures found in the syntax of natural languages. As was indicated in that section, that style of learning is probably not entirely appropriate for medical data but the same general principles should apply. This subsection describes in outline how the SP system may be applied to the learning of medical knowledge.

The simplest kind of ‘learning’ is simply to keep a record of the symptoms of each specific patient and the corresponding diagnosis. This is the principle of

‘case-based learning’ and, in conjunction with a system that kind of find good partial matches between the description of a new patient and stored knowledge of old patients, it can be quite practical and useful in making new diagnoses. Since the SP system is capable of finding good partial matches between patterns, it could well be used in this way, as suggested in Section 4.5, below.

However, medical practitioners normally recognise a degree of abstraction in medical knowledge as described in Section 3.10: a concept like ‘fever’ describes a cluster of symptoms that is found in two or more different diseases; diseases like influenza often come in several variants or subclasses; and symptoms may be described in terms of ranges rather than specific values as in our example of temperature (Section 3.6).

As an illustration of the way in which the SP system may create these kinds of abstraction, consider two imaginary patients with two different diseases and symptoms represented by the patterns ‘X A B Y C D E Z F’ and ‘A P Q B C R D S E F T’. With these two patterns, the SP system would create an alignment like this:

```

O X A      B Y C      D      E Z F      0
   |       |       |       |       |
1  A P Q B  C R D S E  F T 1,
    
```

and from the matched symbols in this alignment it may derive the pattern ‘A B C D E F’. With the addition of some appropriate ID-symbols, this pattern may serve like the pattern for ‘fever’ in Figs. 4 and 6: it represents a cluster of symptoms that appears in two or more different diseases. Alternatively, this pattern may represent the symptoms of a general class of diseases with two variants represented by the two original patterns, modified so that the shared cluster of symptoms are replaced by a pointer to the general class.

With regard to categories like the set of two alternative temperatures represented by the patterns ‘<t1> 38–39 </t1>’ and ‘<t1> 40+ </t1>’ in Fig. 4, the SP system may derive this kind of disjunctive category in much the same manner as the disjunctive class {‘b o y’, ‘g i r l’} in the example in Section 2.3.4.

3.12. Integration

The very simple format for representing knowledge described in Section 2.2 is intended to be as nearly ‘universal’ as possible in the sense that it is designed to represent a wide range of different kinds of knowledge. This should be much more nearly true when the concept of *pattern* has been generalised to two dimensions.

In a similar way, the concept of multiple alignment described in Section 2.3 is intended to be a ‘universal’ model for a wide range of different kinds of processing: pattern recognition, information retrieval, various kinds of reasoning, and so on.

To the extent that these two objectives can be realised, they should facilitate the seamless integration of different kinds of application, including medical diagnosis. It should, for example, be relatively easy to apply a natural language interface to an SP system for medical diagnosis or to integrate the kind of visual pattern recognition needed in the diagnosis of different kinds of skin cancer with other kinds of medical expertise.

4. Comparison with alternatives

As mentioned in the introduction, a wide variety of philosophies and systems have been applied to the problem of medical diagnosis. In this section, I briefly review some of the more prominent of these approaches and compare them with the SP approach, as described in this paper.

4.1. Rule-based systems

Rule-based systems (like the well-known MYCIN system [20]) contain if–then rules where the ‘if’ side of any rule is a collection of one or more conditions for the rule to fire connected by logical operators such as ‘AND’, ‘OR’ (which may be inclusive or exclusive) and ‘NOT’. By contrast, the SP system expresses all knowledge in the form of patterns.

At first sight, SP patterns lack the expressive power of if–then rules. But the effect of such rules can be modelled within the SP system if that is required [25,29]. And if medical diagnosis is viewed as a process of pattern recognition (as in this paper), then SP patterns and the SP framework are, arguably, a more

natural and flexible medium for the representation and processing of knowledge than are if–then rules.

```
IF chills AND cough AND headache AND aching muscles AND runny nose AND sore throat
   THEN influenza (probability = 0.9)
```

Although there may be a probability associated with the rule (as shown), the rule has an intrinsic logic which, if strictly applied, means that the rule will *only* fire if all the conditions are satisfied. By contrast, a pattern like the one shown in column 3 of Fig. 6 may appear in the best alignment when any reasonably large subset of its symbols have been matched.

If one attempted to achieve this kind of flexibility with an if–then rule using combinations of AND, OR and NOT, the rule would become very complex. Alternatively, one might split up the rule into a number of smaller rules, one for each symptom or combination of two or three symptoms—but again the result would be relatively complex.

4.1.1. Probabilities

In systems like MYCIN and some of its successors, the ‘probabilities’ that the system calculates are really measures of confidence without the theoretical underpinnings of probability theory. In other systems, “. . . formal approaches based on probability theory are precise but can be awkward and non-intuitive to use.” ([11], p. 272). By contrast, the SP framework allows true probabilities to be calculated quite simply (see Section 3.7 and Ref. [26]) and strictly in accordance with established theory (as described in sources such as Ref. [10]).

4.2. Neural networks

One of the attractions of artificial neural networks for the support of medical diagnosis is that they can be trained with appropriate data, thus by-passing the need for the manual compilation of knowledge by medical experts or knowledge engineers. However, “A major drawback is that ‘knowledge’ embedded [in the neural network] is cryptically coded as a large number of weights and activation values. As a consequence, the lack of neural network validation tools is often one of the reasons limiting their use in

To illustrate this last point, we can express the distinctive features of influenza by a rule such as:

practice, especially in the context of medical diagnosis where physicians cannot trust a system without explanation of its decisions.” ([5], pp. 141–142).

While there may be scope for extracting rules from a trained neural network ([5], pp. 141–142), this adds complexity and uncertainty to the technology and defeats the other main attraction of a neural network: as a classifier of specific cases in terms of the learned knowledge.

As a system for unsupervised learning of knowledge structures from raw data, the SP system is not yet a rival to existing neural network systems. However, the system has clear potential for unsupervised learning and, if that potential can be realised, the system has the advantage that its knowledge is stored in a form that can be read and understood by people. Meanwhile, if it is supplied with knowledge about diseases derived from experts or text books, it can be used for diagnostic classification of individual patients.

4.3. Fuzzy logic and fuzzy set theory

Given the variability of diseases and other uncertainties associated with medical diagnosis (Section 3.6), the field of fuzzy logic (and fuzzy set theory) has the obvious attraction that it has been designed with the explicit intention of providing a model for ‘fuzzy’ concepts and ‘fuzzy’ operations on them (see, for example, Refs. [12,8]).

In purely theoretical terms, the field of fuzzy logic may be criticised because it introduces a fairly elaborate conceptual framework to accommodate the undoubtedly fuzzy nature of many human concepts but this conceptual framework is poorly integrated with other ideas about the nature of human cognition. By contrast, the SP theory grew out of research in psychology and it provides a unified model for several aspects of human perception and cognition [31].

Considerations of that kind may be discounted as not relevant to the practicalities of medical diagnosis.

But in that connection fuzzy logic has the drawback that it introduces another layer of complexity to the already difficult process of eliciting knowledge from medical experts [4]. There seems to be some scope for ameliorating this problem by the provision of appropriate tools [4] but the basic problem remains. By contrast, the SP system allows concepts to be expressed in a simple, intuitive manner and, at the same time, it accommodates much, perhaps all, of the fuzziness of medical diagnosis.

4.4. Bayesian networks

Two of the main differences between Bayesian networks (see, for example, Ref. [15]) and the SP system are:

- Bayesian networks focus on the binary relationship between any given node in the network and each of its parent nodes (if any). In this respect, they inherit some of the thinking behind if–then rules. By contrast, the SP system is oriented towards the representation and processing of associations (expressed as patterns) that may contain arbitrarily many elements.
- Correspondingly, any given Bayesian network stores its statistical knowledge in the form of tables of conditional probabilities, one for each node in the network. By contrast, the SP system stores its statistical knowledge in the form of integers, one for each pattern, representing the absolute or relative frequency of that pattern in some domain.

These and related differences seem to underlie some of the apparent drawbacks of Bayesian networks compared with the SP framework:

- The directional nature of Bayesian networks does not sit easily with the non-directional nature of medical syndromes.
- The process of calculating probabilities of inferences in a Bayesian network is substantially more complicated than the calculation of probabilities for alignments and inferences in the SP framework.
- The tables of conditional probabilities required for Bayesian networks are significantly more complex than simple measures of frequency that are used in the SP system. Notwithstanding the development of

special methods for eliciting conditional probabilities from experts [22], the process of building up the necessary tables of conditional probabilities is likely to be much harder than measuring or estimating an integer value for each disease, reflecting its absolute or relative frequency in a given domain.

4.5. Case-based reasoning

A major attraction of case-based reasoning in medical diagnosis (see, for example, Refs. [19,1]) is that, compared with many of the alternatives, it can considerably simplify the process of acquiring the necessary knowledge. In its simplest form, a case-based system merely requires a description of one or more specific examples of each disease and a search algorithm that can find exact matches or good partial matches between the symptoms of a given patient and one or more of the stored records.

In some respects, the SP system is like a case-based system and it could indeed be used like a case-based system. To use it in this way, each of the Old patterns should represent a specific case (including its diagnosis) and the New pattern or patterns should represent the symptoms of a patient for whom a diagnosis is required. The capabilities of the system for finding exact matches and good partial matches between patterns will allow it to retrieve patterns for previously-diagnosed cases that are similar to any given current case.

The main advantages of the SP system compared with the case-based approach to diagnosis are:

- It facilitates the description of diseases in generalised terms without the need to specify exact values for every category of symptom. In our main example, we saw how the temperature of a patient with a disease like influenza may be specified as a range of alternative values (Section 3.6). Any other category of symptom may be treated in the same way.
- It allows one to specify clusters of symptoms that are found in two or more different diseases and it allows one to describe diseases at two or more levels of abstraction (Section 3.10). Both of these things facilitate the description of diseases without the need to repeat information unnecessarily where similar patterns are found in different diseases or varieties of disease.

5. Conclusion

As we have seen, the SP system accommodates the main elements of medical diagnosis, viewed as a problem of pattern recognition, and there are reasons to believe that it may also provide support for causal reasoning in medical diagnosis. However, the SP62 model is only a prototype that serves for research and demonstration. It is not yet a system with ‘industrial strength’. The main developments that are needed to reach that goal are:

- The provision of a well-designed graphical user interface.
- There is probably scope for improvements in the search methods that are used within the system.
- There is scope for the application of parallel processing both to improve the scaling properties of the system (Section 2.3.5) and to increase absolute speeds of processing.
- Naturally, the system needs to be provided with appropriate knowledge. For each area of application, a set of patterns needs to be developed that describes the diseases and symptom clusters in that domain.
- At some stage after the development of a realistic knowledge base, the performance of the system must be validated against the judgement of human medical experts.

The potential payoff from these developments is a system that allows knowledge about diseases to be expressed in a simple, intuitive manner, that can cope with errors and uncertainties in knowledge about diseases and knowledge about individual patients, that simplifies the acquisition and storage of statistical information, that calculates true probabilities of diagnoses, that smooths the path to the automatic or semi-automatic abstraction of medical knowledge in the future, and should facilitate the integration of medical diagnosis with other kinds of application.

References

- [1] K.-D. Althoff, R. Bergmann, S. Wess, M. Manago, E. Auriol, O.I. Larichev, A. Bolotov, Y.I. Zhuravlev, S.I. Gurov, Case-based reasoning for medical decision support tasks: the Inreca approach, *Artificial Intelligence in Medicine* 12 (1998) 25–41.
- [2] F. Attneave, Some informational aspects of visual perception, *Psychological Review* 61 (1954) 183–193.
- [3] H.B. Barlow, Trigger features, adaptation and economy of impulses, in: K.N. Leibovic (Ed.), *Information Processes in the Nervous System*, Springer, New York, 1969, pp. 209–230.
- [4] K. Boegl, K.-P. Adlassnig, Y. Hayashi, T.E. Rothenfluh, H. Leitich, Knowledge acquisition in the fuzzy knowledge representation framework of a medical consultation system, *Artificial Intelligence in Medicine* 30 (2004) 1–26.
- [5] G. Bologna, A model for single and multiple knowledge based networks, *Artificial Intelligence in Medicine* 28 (2003) 141–163.
- [6] T. Bray, J. Paoli, C.M. Sperberg-McQueen, E. Maler. Extensible markup language (XML) 1.0 (second edition), Technical Report, World Wide Web Consortium, W3C Recommendation (6 October 2000). Copy: www.w3.org/TR/REC-xml.
- [7] N. Chater, Reconciling simplicity and likelihood principles in perceptual organisation, *Psychological Review* 103 (3) (1996) 566–581.
- [8] S.-M. Chen, A weighted fuzzy reasoning algorithm for medical diagnosis, *Decision Support Systems* 11 (1) (1994) 37–43.
- [9] A. Church, *The Calculi of Lambda-Conversion*, Annals of Mathematical Studies, vol. 6, Princeton University Press, Princeton, 1941.
- [10] T.M. Cover, J.A. Thomas, *Elements of Information Theory*, John Wiley, New York, 1991.
- [11] J. Fox, D. Glasspool, J. Bury, Quantitative and qualitative approaches to reasoning under uncertainty in medical decision making, in: S. Quaglini, P. Barahona, S. Andreassen (Eds.), *Proceedings of the 8th Conference on Artificial Intelligence in Medicine in Europe, AIME, Lecture Notes in Computer Science*, vol. 2101, Springer, Heidelberg, 2001, pp. 272–282.
- [12] G.J. Klir, B. Yuan, *Fuzzy Sets and Fuzzy Logic: Theory and Applications*, Prentice Hall PTR, Upper Saddle River, NJ, 1995.
- [13] M. Li, P. Vitányi, *An Introduction to Kolmogorov Complexity and Its Applications*, Springer-Verlag, New York, 1997.
- [14] P. Mangiameli, D. West, R. Rampal, Model selection for medical diagnosis, *Decision Support Systems* 36 (3) (2004) 247–259.
- [15] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*, Revised Second Printing Edition, Morgan Kaufmann, San Francisco, 1997.
- [16] E.L. Post, Formal reductions of the general combinatorial decision problem, *American Journal of Mathematics* 65 (1943) 197–268.
- [17] J. Rissanen, Modelling by the shortest data description, *Automatica-J IFAC* 14 (1978) 465–471.
- [18] D. Sankoff, J.B. Kruskall, *Time Warps, String Edits, and Macromolecules: the Theory and Practice of Sequence Comparisons*, Addison-Wesley, Reading, MA, 1983.
- [19] R. Schmidt, L. Gierl, Case-based reasoning for antibiotics therapy advice: an investigation of retrieval algorithms and

- prototypes, *Artificial Intelligence in Medicine* 23 (2001) 171–185.
- [20] E.H. Shortliffe, *Computer-Based Medical Consultations: MYCIN*, Elsevier/North Holland, New York, 1976.
- [21] R.J. Solomonoff, A Formal theory of inductive inference: Parts I and II. *Information and Control* 7 1964; 1–22 and 224–254.
- [22] L.C. van der Gaag, S. Renooij, C. Witteman, B. Aleman, B. Taal, Probabilities for a probabilistic network: a case study in oesophageal cancer, *Artificial Intelligence in Medicine* 25 (2002) 123–148.
- [23] C.S. Wallace, D.M. Boulton, An information measure for classification, *Computer Journal* 11 (2) (1968) 185–195.
- [24] J.G. Wolff, Learning syntax and meanings through optimization and distributional analysis, in: Y. Levy, I.M. Schlesinger, M.D.S. Braine (Eds.), *Categories and Processes in Language Acquisition*, Lawrence Erlbaum, Hillsdale, NJ, 1988, pp. 179–215 (Copy: www.cognitionresearch.org.uk/lang_learn.html#wolff_1988).
- [25] J.G. Wolff, ‘Computing’ as information compression by multiple alignment, unification and search, *Journal of Universal Computer Science* 5 (11) (1999) 777–815 (Copy: <http://arxiv.org/abs/cs.AI/0307013>).
- [26] J.G. Wolff, Probabilistic reasoning as information compression by multiple alignment, unification and search: an introduction and overview, *Journal of Universal Computer Science* 5 (7) (1999) 418–462 (Copy: <http://arxiv.org/abs/cs.AI/0307010>).
- [27] J.G. Wolff, Syntax, parsing and production of natural language in a framework of information compression by multiple alignment, unification and search, *Journal of Universal Computer Science* 6 (8) (2000) 781–829 (Copy: <http://arxiv.org/abs/cs.AI/0307014>).
- [28] J.G. Wolff, Information compression by multiple alignment, unification and search as a framework for human-like reasoning, *Logic Journal of the IGPL* 9 (1) (2001) 205–222 (first published in the Proceedings of the International Conference on Formal and Applied Practical Reasoning (FAPR 2000), September 2000. Copy: www.cognitionresearch.org.uk/papers/pr/pr.htm).
- [29] J.G. Wolff, Mathematics and logic as information compression by multiple alignment, unification and search, Technical Report, CognitionResearch.org.uk 2002; Copy: <http://arxiv.org/abs/MathML:/math.GM/0308153>.
- [30] J.G. Wolff, Unsupervised learning in a framework of information compression by multiple alignment, unification and search, Technical Report, CognitionResearch.org.uk 2002; Copy: <http://arxiv.org/abs/cs.AI/0302015>.
- [31] J.G. Wolff, Information compression by multiple alignment, unification and search as a unifying principle in computing and cognition, *Artificial Intelligence Review* 19 (3) (2003) 193–230 (Copy: <http://arxiv.org/abs/cs.AI/0307025>).
- [32] J.G. Wolff, Unsupervised grammar induction in a framework of information compression by multiple alignment, unification and search, in: C. de la Higuera, P. Adriaans, M. van Zaanen, J. Oncina (Eds.), *Proceedings of the Workshop and Tutorial on Learning Context-Free Grammars*, 2003, pp. 113–124 (This workshop was held in association with the 14th European Conference on Machine Learning and the 7th European Conference on Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD 2003), September 2003, Cavtat-Dubrovnik, Croatia. Copy: <http://arxiv.org/abs/cs.AI/0311045>).



Gerry Wolff is Director of CognitionResearch.org.uk. Previously, he has held academic posts in the School of Informatics, University of Wales, Bangor, and the Department of Psychology, University of Dundee, a Research Fellowship with IBM in Winchester, UK, and he has worked for 4 years as a Software Engineer with Praxis Systems in Bath, UK. His first degree is in Natural Sciences (Psychology) from Cambridge University and his PhD is in Cognitive Science from the University of Wales. He is a Chartered Engineer and a Member of the British Computer Society (Chartered IT Professional). Since 1987 his research has focussed on the development of the SP theory. Previously his main research interests were in developing computer models of language learning. He has numerous publications in a wide range of journals, collected papers and conference proceedings.