Development and Use of BOVID-3, An Expert System for Veterinarians Involved in Diagnosis, Treatment and Prevention of Diseases of Cattle

P. Brightling,[1] M.T. Larcombe,[2] B.C. Blood[2] & P.C. Kennedy[3]
[1]University of Melbourne,
Veterinary Clinical Centre, Princes Hwy.
Werribee, Victoria, Australia 3030
Email p.brightling@vet_science.unimelb.edu.au
[2]Animal Information Management, 209 Watton St,
Werribee, Victoria, Australia 30303
[3]861 Laketerrace Circle, Davis, California, USA 95616

Abstract

Case management is a complex process in which veterinarians utilise a large amount of information about diseases. This project was undertaken to develop a comprehensive, up-to-date information source on diseases of cattle, which could be easily interrogated and was presented in a form that supported common veterinary diagnostic decision-making techniques. The objective was to provide an information management tool which would be valuable for practitioners and government officers and could be used to enhance undergraduate education in cattle medicine.

BOVID-3 is compiled from the veterinary literature and the experience of a panel of senior cattle clinicians and pathologists. Over 1040 different diseases have their clinical signs, clinicopathological findings, gross and histological lesions specified. Each disease also has a summary of its cause and transmission, major risk factors, confirmatory tests, recommended treatments and control measures (in general terms), and cross-references in standard texts.

In case investigations, field veterinarians aim to establish differential lists that rank diseases in order of likelihood. BOVID-3 includes a classical probability analysis routine to model that approach

The program is currently used in 42 sites internationally. It is a primary resource for senior undergraduate veterinary students at the University of Melbourne where it has been integrated into the cattle medicine curriculum. Through use of BOVID-3, students learn the basis of the diagnostic strategies they are employing and how to access and manipulate information sources. These are facets of the deep learning approaches that innovative contemporary veterinary educators now foster.

Introduction

Making a diagnosis is a critical step in the path from case presentation to treatment, and requires a fine blend of keen observation, astute interpretation and wise selection. Our direct experience with cattle cases in the field, and in teaching undergraduate veterinary students, has highlighted for us just how complex the diagnostic decision-making process really is, and stimulated us to find methods of exploring and supporting it.

The ability to observe an animal efficiently and comprehensively and to select the key abnormalities present, is the hallmark of the successful clinician. Students must spend much of their learning time devoted to acquiring these skills.

Once the key abnormality has been identified, there are two basic methods of proceeding to selection of an appropriate short list of differential diseases: to work from "first principles" to a logical determination of the patho-anatomical lesion and causative agent(s) for the condition or to match the pattern of signs observed in the case with the signs recorded for known diseases. Both techniques are used by nearly all veterinary clinicians, (often in any one case) but our observations over more than 50 years indicate that the most common method is a pattern-matching approach. The clinician searches for the diagnosis of "best-fit".

Using this approach, the veterinarian creates a short list of diagnostic hypotheses based on the recogni-

Adapted from the Proceedings, XIX World Buiatrics Congress, Edinburgh, Scotland; July 8-12, 1996.

tion of critical clinical sign. The list is usually ranked in order of probability, based on the clinician's knowledge of the frequency of local occurrence of the diseases. Specific questions are then asked which support or refute the hypotheses. The clinician seeks supporting evidence until the condition matches the sign pattern of the hypothesized diagnosis sufficiently well— until he or she is sufficiently certain of the diagnosis.

This method has the advantage of economising on time spent with the case and on clinicopathological tests because everything that is done is relevant to diseases which are strong possibilities. It is not as successful when the clinician does not have a very good mental database of diseases to chose from, or knowledge of the relative occurrence rates of diseases in the area. It is also less successful when the disease in question is a true or rare one, or when a second disease is present as well.

We have observed that students often have particular difficulty in establishing lists of disease hypotheses and choosing appropriate questions to test them. They may fail to include the correct diagnosis on their initial lists, may be unable to recall discriminating signs, or not know how to rank the likely occurrence of diseases (the "which ones are common?" question).

It was in the area of matching case data to differential disease lists that we considered significant support could be offered through a computerised database. Our aim was to establish a comprehensive, up-to-date dataset that could be easily interrogated. It needed to provide manageable, credible differential lists in order of probability, with summary information about additional observations to be made to improve diagnostic acuity, and guides to treatment and control of each disease. It also needed to show why certain observations were required to differentiate diseases - to be a transparent expert system which helped students understand the process of diagnosis. If this could be achieved it would be a powerful case management support tool for veterinarians in the field, and enhance education by promoting deep learning and alleviating the need for undergraduates to rote learn a large amount of detail from lists of diseases.

The most appropriate approach was judged to be generation of an electronic database of cattle diseases and their signs and use of Bayes' theorem to calculate the conditional probability of the disease being present given that particular sign(s) had been observed.

Materials and Methods

A classical probability approach gives an elegant method of determining a list of best-fit diagnoses, but relies on a detailed dataset for the domain of the program. To cover the complete field of cattle medicine we required lists of all diseases recorded in cattle and all signs used to define them, an estimate of the likely occurrence of each sign in all diseases, and a prevalence estimate for all diseases.

The list of diseases

The list of diseases to be included was generated from our prime resource text Veterinary Medicine (Radostits, Blood & Gay, 1994). Additional disease names were then added from other texts on lameness and reproductive diseases.

Some diseases were included in several forms because although they had a common aetiology, their signs or epidemiology differed (e.g. cardiac and nervous forms, amongst others, of bovine viral leucosis), and they represented distinct diagnoses which were helpful to clinicians.

New diseases were added as they appeared in the literature. The contents of 84 journals were checked on a monthly basis.

The list of signs

The first list of signs was extracted from Veterinary Medicine, and additional signs were included by the authors when required. Initially the dataset only included clinical and clinicopathological signs, but in the most recent version of the program, gross and histological findings have also been added.

To avoid the confusion caused by use of synonyms, a set vocabulary was adopted, and a definition of each sign provided in an on-screen glossary.

Probability of each sign occurring in every disease

To calculate the probability of a disease given the observation of certain sign(s), it was necessary to have an estimate of: the conditional probability of observing a sign given that the disease was present and an estimate of the frequency of that sign (whether it was a common abnormality to many other diseases). This required an estimate of the probability of observing each sign in every disease.

Quantitative data for these estimates were found to be very scarce in veterinary medicine. The method adopted was for two of the authors to independently score each sign, and confer where their estimates differed by more than 30%. A third opinion was sought if the authors maintained a difference of more than 30%, and a consensus decision reached on the score to be assigned.

The prevalence of each disease

Each disease had a prevalence score assigned, based on a relative ranking—with the most common disease occurring at 100 and the very rare and exotic diseases scored as less than one. These scores were based on the authors' assessment of disease occurrence in the southern, temperate region of Australia but the program was designed to allow all prevalence scores to be easily reassigned by users, to local data, where required.

The application program

The dataset and the algorithm for its use were prepared as an application program for an IBM compatible platform, running DOS version 3.0 or above. The prototype (known as BOVID) was released in 1989, and revised (BOVID-2) in 1991. For the current version (BOVID-3, 1995), necropsy findings were added to the dataset and design features were chosen to minimise keystrokes, and allow users to save, retrieve and add their own notes for each case.

Refining and testing the database and program

The initial dataset was tested by running 240 cases from the case records of one of the authors. Changes to the database and algorithm were made and a further 65 cases described in a regular newsletter from an Australian regional veterinary laboratory were entered. These cases were independent of the clinical experience of the authors and provided sufficient detail on clinical findings to enter cases as they would have been attended in the field.

The application program was made available as a prototype and 31 sites installed the program and provided feed-back on its use.

After the re-design of the program and inclusion of the necropsy findings, BOVID-3 was demonstrated to nine experienced cattle practitioners in Australia, who were requested to enter details of cases which they had attended, where they were highly confident (based on clinicopathological, necropsy or clinical findings such as response to treatment) of the diagnosis. They were asked to record the time required to enter each case, and comment on the differential list generated by BOVID-3 (to identify diagnoses which were listed but were inappropriate, nominate diagnoses which should be included but were omitted, and comment on the order in which the diagnoses were listed). Details on 27 cases were collected.

Results

1044 separate diseases are included in the database, each with its signs and a summary of the cause and transmission of the disease, breed susceptibilities, major epidemiological risk factors, confirmatory tests or procedures, recommended treatments (in general terms), recommended control measures, and references pages in standard texts. 1191 clinical, paraclinical, clinicopathological and pathological findings are defined.

The database has a matrix of over one million estimates of the probability of signs occurring in diseases. More than 95% of these are the expression of the collective opinion of the authors of the program. The prevalence ranking of diseases is also the opinion of the authors, but is easily edited by users.

The algorithm used with the database models the hypothetico-deductive approach to diagnosis. When a small group of signs is entered, (with one marked as a key sign), a list of diagnoses is created which could explain that bracket of observations. The program then has an "Interrogation" mode in which it checks the database for the signs (or lesions) which would differentiate the possible diagnoses on the initial list. The results for each of these observations are requested and the list of probable diagnoses is rearranged. Thus the important signs which are NOT present in the case are also established. During the Interrogation mode, the disease which has prompted the current question is highlighted. In 23 of 27 cases assessed through BOVID-3 by experienced cattle clinicians, the resultant list of differentials was described as complete and in expected order. In three cases the order of the diseases listed was changed and in one case which involved two separate disease syndromes the program failed to offer an acceptable diagnosis. The time required was less than 10 minutes in all cases.

The program was considered to be sufficiently robust to make it a primary resource for senior veterinary students, and its use was incorporated into the cattle teaching curriculum at the University of Melbourne.

Discussion

BOVID-3 has been developed as a computerised problem-oriented index to the diseases of cattle, using probabilities to differentiate the likelihood of different diagnoses in cases in the field. It is a huge database, compiled from the veterinary literature but without the advantage of statistical data from a substantial collection of case material. In the current program it is possible to save case details, so there is now the potential of such a dataset being created prospectively by users. The data is regularly up-dated.

The program uses Bayes' Theorem to calculate the probabilities of differentials, given that a number (often a large number) of signs have been observed to be present or absent. When more than one sign is used the calculation sums the effects of each sign, which is reasonable provided that the signs are independent. Sometimes this is not the case, for example pallor of the mucosae and pallor of the conjunctivae are not independent observations. This makes the calculated probability values less accurate but the order of the diagnoses is not likely to be greatly altered. The program cannot reliably generate a credible differential list when two concurrent diseases are present in a patient.

In veterinary medical education the volume of information which students face is now so large that it is imperative that they learn to access required data when they need it (Whithear, Browning, Brightling & McNaught, 1994). BOVID-3 has been incorporated into the bovine medicine curriculum at the University of Melbourne in two ways: as an information source for assignments introduced in lectures, and in case investigations.

Lecture topics are organised around a series of presenting signs and students are encouraged to use the database to explore comprehensive lists of diseases which could include those signs. For example, one assignment requires students to compare all those diseases in which oral mucosal lesions may be observed, determine which key signs distinguish them, and which are common or exotic to a particular environment.

Clinical case material is provided for the students in two main ways: actual cases are obtained for teaching purposes where the students retain virtually all the decision-making responsibility, and video/slide case simulations are presented in tutorials. BOVID-3 is used to establish and narrow differential lists for both of these types of investigations, and the process of diagnosis is examined during its use.

Any diagnostic assistance program can only offer an opinion based on the clinical signs (or necropsy lesions) that the student has recorded. The more complete the examination, the better the student's observation skills, the more accurate the diagnosis will be. This program has a significant educational benefit because the Interrogation function reminds the student of body functions that should be examined, and may prompt a focused re-examination of the patient: the essential clinical skills of observation and interpretation are highlighted. A list of specific clinical procedures and clinicopathological tests which would be useful in the particular case is also be generated, leading the student to focused rather than "survey" laboratory requests.

The development of a comprehensive cattle medicine database which supports diagnostic and case management decisions, encourages careful observation of the patient, and can be used in undergraduate and field work, provides a significant new tool in advancing veterinary undergraduate and continuing education. The program is now in use in 42 sites, including 24 practices, 10 government veterinary units, and 8 veterinary schools.

References

Radostits OM, Blood DC & Gay CC (1994) Veterinary Medicine 8th edn., Bailliere Tindall, London. Whithear KG, Browning GF, Brightling P & McNaught C (1994) Veterinary education in the era of information technology. Aust Vet J 71:106-108.