

Future directions in quantitative pathology: digital knowledge in diagnostic pathology

Peter H Bartels

The ancient Pythagorean philosophers believed in numbers. Numbers were the key to an understanding of the world and—because they were eminently practical people—to prosperity. As we enter the information age we hold similarly strong beliefs in the benefits of digital technology: numerical representation offers objectivity, discreteness, permanence, analytic evaluation. There has always been the explicit hope that quantification of knowledge in histopathology will eventually lead to better medical practice. It is not an unfounded hope. The introduction of quantitative micromorphometric procedures to diagnostic histopathology has shown the potential of information technology based methods. There is the value of objective mensuration with the prospect of standardisation of diagnostic assessment. There is the ability to establish the significance of small and early alterations. There is the proven ability to extract diagnostic information that is not reliably recognised by visual microscopic assessment. There is the ability to measure rapidly, in numerical terms, the efficacy of chemopreventive intervention and thus the ability to test a large number of potentially effective compounds. Over all, there has been outstanding progress.

Yet it may be time to take stock and to ask where we stand and where future efforts might be productive. In fact, we must ask whether we are using to a maximum the information that digital imagery offers. Questions arise as to whether the methodology underlying the quantitative analysis provides the information that we need, and even whether it is appropriate for some of the problems encountered in diagnostic and prognostic histopathology.

Image processing, diagnostic information extraction, and quantitative assessment have for years relied on methods borrowed from the material sciences and from statistical analysis, often following the lead of research in “pattern recognition.” In fact, mathematical/statistical procedures form the bulk of analytic tools in quantitative histopathology and they are indeed appropriate and immensely useful in many instances.

However, one must not forget that statistical procedures are designed to assess groups of observations and not to characterise individuals. What the clinician wants to know from a pathologist, though, is not only the specific diagnostic assessment but what the prognosis and best management plan is for a particular patient. There is an inherent mismatch between the design objective of the analytic procedures and the desired utility.

As one examines our analytic arsenal of mathematical/statistical procedures and the needs of diagnostic and prognostic histopathology it becomes evident that additional analytic approaches will have to be developed. Are we asking all the questions that we should?

Reductionism and diversity

We may begin by considering the principal rationale underlying current mathematical/statistical analysis. It is based on reductionism. A set of observations is characterised, summarised, and reduced to a “mean value.” It defines the sample, in an efficient data reduction, by a single number. The mean value of a sample allows a formal hypothesis test on whether it deviates from a standard, or from the mean value of another set of observations. The mean value allows us to abstract from the variability of values among the observations taken on the members of the sample. The mean value evolves as representative for the data set as the central limit theorem takes effect.

Variability is considered to be due to randomness. True, mathematical/statistical techniques allow us to explain the various influences contributing to the total variance—such as nucleus to nucleus variation, section to section variation, patient to patient variation. Thus one can reduce the amount of unexplained random variation by skilful experimental design and analysis of variance. Observed value differences within the sample, however, are essentially considered to be due to randomness and randomness only.

However, already when we examine two datasets representing nuclei from lesions of different grade the model is not really appropriate any more. “Low grade” and “high grade” are linguistic terms, designating instances of a given lesion at different stages of progression. Differences in the values of observed features are due, primarily, to the effects of progression—that is, of a continuous process, in addition to randomness.

The inadequacy of a simplistic analytic approach becomes even more serious when we consider the nature of the entities which we observe and measure.

In the material sciences observations generally are taken on entities that one may denote as “simple objects.” Such objects typically comprise few components, or might even, for purposes of the analysis, be considered as single component objects. Their properties are well defined and may be characterised by a moderate number of features. The values of these features may be correlated but, as a rule, they are not affected by the object’s internal feedback

Optical Sciences
Center, University of
Arizona, Tucson,
Arizona 85721, USA
P H Bartels

Correspondence to:
Dr Bartels
email: peter@
catalina.opt-sci.arizona.edu

and control mechanisms. Therefore, the behaviour of objects is fairly predictable. They are considered “linear systems” and it is for objects of this nature that the common mathematical/statistical methods were developed.

Systems in which feedback and control occurs are very difficult to predict, analytically, in their behaviour. In the physical sciences multibody systems with feedback involving only very few components are no longer tractable analytically. Such systems are “non-linear” and complex.

Now, metabolic subsystems, cells, lesions, patients, all of these entities are subject to a very large number of feedback and control mechanisms to maintain a labile equilibrium. It raises the question as to what extent members of a collective of such entities—such as patients with a lesion assigned the same diagnostic label—can be considered comparable. To what extent does the complexity of an observed entity set it sufficiently apart, for the purpose of predicting biological behaviour, to make it a unique, individual entity? One should suggest that, in addition to randomness and trendal progression, there is the role of genuine diversity to be considered in diagnostic/prognostic decision making.

Reductionism and variability owing to randomness alone are inadequate models. But they provide the model for diagnostic and prognostic classification in quantitative histopathology.

Classification and identification procedures

Visual microscopic diagnostic assessment, in practice, does not really result in a classification. Rather, it leads to an identification of a lesion. This is an important distinction, and it should by all means be carried over to procedures used in quantitative histopathology.

It is worth while enumerating the differences between a classification and an identification. What is implied, to begin with, by the concept of a diagnostic “class”?

Traditionally, entities that are similar in their properties and in the values of their properties, and that are different in these respects from other entities, have been defined as forming a “class.” The underlying assumption is that the properties considered for this assignment are not only necessary, but sufficient for the purposes of the classification process. This implies that the members of the class form a homogeneous set—that is, that they are interchangeable and that each can be expected to behave in a manner very similar to the other members of the class.

Mathematical/statistical classification procedures are based on a modest number of features, no more than are necessary and sufficient to discriminate members of one class from those of another. The classification process occurs in a closed feature space. The same set of features is used to determine class membership of unknown entities. A class assignment is made on the basis of the observed values of those features. Assignment to a class, as a matter of fact, is made on the basis of simi-

larity to the typical values seen in members of that class. In practice, for an unknown entity the values of only the features required for the classification function are computed. They are weighted, a classification score is computed, and a classification decision is made based on whether that score is lower than or exceeds a critical threshold value.

In contrast, in an identification procedure information is collected and an assessment is made as to how much confidence a decision maker can have in declaring an identity. More and more information is added until the level of confidence in a decision satisfies the diagnostician or a preset criterion value. An identification procedure operates in an open feature space. It is possible that assignment to the same identity is reached, for different entities, on the basis of different sets of features.

The difference between a classification procedure and the highly specific characterisation involved in an identification is best illustrated by a simplified example.

Consider the following numerical patterns:

5277	436
88322	493179
61616	633813857
343	2485486
4678826	5513677

These can be classified into two disjointed classes by a classification rule that if the number of odd digits is equal to or greater than the number of even digits, assign the pattern to class O; if not, assign the pattern to class E.

An identification may begin with a preclassification: for example, this pattern belongs to class E. However, the identification process then continues and attempts to establish: Is it pattern “88322”?

Furthermore, the process will determine whether this assertion can be made with high confidence. The identification thus targets a very specific neighbourhood in a pattern space. Besides, if one considered these numerical sequences as possible keys to a combination lock, then no classification would be of any use. It is the identification of the correct combination that is needed.

For many objectives in diagnostic histopathology the concept of a class or of diagnostic categories makes sense and is useful. However, the characterisation of an individual lesion by the label of a class is not sufficient for more specific assessments. The criteria and observed values used in a classification function are in fact no more than the minimum information requirements for assignment to a diagnostic category.

It is a well established fact that lesions of a given diagnostic category may, and do, show drastically different biological behaviours. Clearly, the classification is based on information sufficient to set the lesions apart from other diagnostic categories. However, evidently the information does not include properties giving rise to such different biological behaviour and clinical outcome.

A more specific “within diagnostic category” characterisation is needed: now the problems

encountered when dealing with complex entities become evident.

Reductionism and consideration of diversity are, of course, not exactly compatible. One is searching for commonality, the other for highly specific characterisation and differences. Also different is the required analytical methodology: mathematical/statistical procedures on the one hand, coding methodology and information theory on the other.

If we postulate the need for a highly specific characterisation of nuclei, a lesion, and a patient's condition, where can we obtain the required additional information?

Efficient use of diagnostic information

Such information is in fact always available and is in most cases even computed routinely and then discarded. The customary mathematical/statistical analytic models either do not allow its use or simply are not suitable to use it. In karyometry, most image analytic software packages routinely compute the values of about 100 different variables. Yet only a few, typically less than eight, are used in a discriminant function. The reason for this restrictive and wasteful use of the available information is that most statistical classification algorithms rely on an estimate of a multivariate probability density distribution. As statistical pattern recognition research has shown, the performance of a classifier based on such a principle deteriorates considerably with an increase in the number of features or "dimensions." To escape this "curse of dimensionality" only the most discriminating, or strong, features are selected and used. These are features which for two classes have widely separated mean values and low dispersions—that is, they have well separated distributions. As a result, for almost any observed object the relative likelihood for being a member of one or the other class is very high, of the order of 1000:1. Features with widely overlapping distributions for different classes typically result in relative likelihood ratios of the order of 20:1 at most, and often less. They are "weak features."

In visual diagnostic decision making, even very subtle diagnostic clues are carefully considered. Indeed, even clues that are not always seen are noted when observed. Validity is verified by observing additional, possibly subtle clues: a diagnostician makes use of redundancy to establish confidence. This is of course exactly what coding methodology does. To avoid errors in message transmission, redundancy is built into a code to provide automated error correction.

There is a rich source of diagnostic clues offered by the multitude of weak features, and it can be used for identification purposes, provided their correlations are properly taken into account. An accumulation of the diagnostic evidence offered by such features is possible—for example, by the use of inference networks. One must ask therefore whether, because diagnosticians make identification, the quantitative approaches should not also be based on identification rather than on mere classification, and thus allow a very specific characterisation. This

would allow proper consideration of diversity—something that is natural for coding methodology, but almost precluded by mathematical/statistical analytic procedures.

The potential of identification procedures has been demonstrated by the "nuclear signatures" and the "lesion signatures" derived from them, in prostatic intraepithelial neoplastic and preinfiltrating breast lesions. A lesion signature is provided by the distribution of nuclei of varying degrees of deviation from normality. Diversity of lesions of the same diagnostic label is shown by the drastically different lesion signatures.

Another and different source of diagnostic clues is offered by "sequence information." It has not been explored at all in quantitative histopathology. This is surprising since sequence information is so commonly the basis for human perception. The letter sequence in a word or the sequence of musical notes in a chord lead not to classification, but to instant and highly specific identification. The first four notes G G G E_{flat}—"fate knocking at the door"—instantly identify Beethoven's fifth symphony.

None of this is used in the usual mathematical/statistical classification algorithms. The typical "feature vector" from multivariate analysis provides an example. For a given entity the recorded values of the variables in the feature vector define a point in the p variate feature space. The order in which the features are considered is immaterial; one ends up at the defined point and that is all that matters. Thus order information is not considered. Nevertheless, one could see the feature vector as a profile with a characteristic contour.

The amount of information offered by the ordering of symbols alone is surprisingly high. Let us suppose, in analogy to a feature vector, that one had a string of six elements, say the word COMEDO. Here we could have six different symbols that would allow $6! = 720$ different arrangements. There are two identical symbols, so there are only $6!/2! = 360$ different arrangements possible.

The information contents is given by $p = 1/360 = 0.00277$, and $I = -n \log_2 p$ or $I = 6 * \log_2 360 = 6 * \log_{10} 360 / \log_{10} 2 = I = 6 * 2.556 / 0.301 = 50.9$ bits.

To provide a perspective, one should consider that an assessment involving 8 bits allows the correct choice of 1 of 256 entities, or a precision of about 0.4%.

Targeting of an analytic procedure

The problems created by the reliance on established mathematical/statistical procedures with their reductionist restrictions are encountered everywhere. The need to develop analytic methodology specifically for the needs of diagnostic and prognostic histopathology is particularly pressing where patient specific identification is required. Survival estimation provides an example. There exists a rich and elegant armamentarium of procedures to estimate survival of patients diagnosed with a given lesion. The problem is that these prognostic assessments are made for the

diagnosed population as a whole. Only rarely do these procedures allow a prognosis that is targeted specifically for the situation seen in a particular patient.

Information technology may allow the creation of massive databases needed to address this problem. Case based reasoning systems are, in principle, designed to allow highly specific individual patient targeted prognosis. The problem is that they are only in the earliest stages of development, with several major challenges still to be met. Historically, case based reasoning for prognosis is very old. Herodotus writes, around 440 BC, that “the Egyptians, when something ominous happens, they write it down and the outcome, and when the same thing happens again, they assume it will have the same outcome.”

The problem is not the principle, but the numeric and statistical validation of the basis for prediction.

The first problem is the need for a very large database of patients of a given diagnostic condition for whom the eventual clinical outcome is on record. The second challenge is that we do not know which patient observable data will reliably predict survival time. If there were such strong variables identifiable, one would not have to go for a case based reasoning approach. The best we can hope for is that sufficient numbers of data items are partially correlated to survival.

In a case based reasoning system, an exhaustive list of data items is established for a given patient. Then a very large database is searched for those cases that are the most similar in their profile of these data items. The outcome for these reference cases is used to predict outcome for the case under study.

The required sample sizes in the reference database are very large, but given the capabilities of information technology they are by no means unattainable. To give a grossly simplified example, let it be assumed that only 10 data items are considered. In practice one might have to consider from 20 to 100 such items. Furthermore let it be assumed that each of these items may take on only four different values, for example quartiles of their value range. This would lead to 4^{10} value combinations, or 1 048 576 possible different profiles. This would indeed allow a very specific signature, even though each variable range is divided only into four intervals. However, one really should have at least 10–30 reference cases for each signature. Also one must make the assumption that those come from a prognostically homogeneous population. This is by no means certain and needs to be verified.

One could allow a looser definition of similarity and reduce the requirements for the very large database, but this risks the occurrence of heterogeneity among the reference cases. In any case, clearly 10 observed variables would not possibly cover the range of data items that might affect survival time. In practice, 10 variables in a case based reasoning system would have to be the result of a preceding data reduction process.

It is most likely that the predictive information, if indeed it is represented in the large set of recorded data items, is to be found in a pattern of values for a goodly number of “weak variables.” The search for such a pattern is a problem that may be addressed by data mining techniques, for which there are at present not yet any really satisfactory procedures. One might set up a very large neural network, or develop procedures that trace patterns of mutual correlation among the variables in datasets of known outcome. Unfortunately, before data mining techniques can be applied, the original clinical database has to be available. Even though the standardised collection of such databases is technically feasible, the magnitude of the task exceeds the logistic resources available to individual researchers and even institutions. There is a need to have a professional organisation take the lead on such an endeavour, on an international basis.

While one certainly would not dispute the value of statistical procedures for survival analysis, the clinical needs call for individual patient targeted prognosis.

Analytic models inappropriate for diagnostic decision making

There is an inherent risk in transferring analytic procedures from one discipline to another. Almost as a rule procedures rest on underlying assumptions which may only be valid for the original application. A typical example is the application of decision procedures which are well proven and effective in engineering situations to clinical diagnostic decision making.

In signal detection engineering, a decision has to be made between “signal” and “noise.” For both events the detected energy shows a distribution, and the two distributions usually have some overlap. A threshold is defined to attain an unequivocal decision; however, owing to the overlap there will now be some signals that are, erroneously, considered noise. Also, there will be some noise that will erroneously be accepted as signal. This trade off is characterised by the receiver operating characteristic, or ROC curve, and it is modelled by a well developed signal detection theory.

For many years the slide diagnostic decision in designs for automated primary screening in cervical cancer has been based on the signal detection paradigm. Some criterion reflecting the chance that a cervical sample is abnormal is computed for each slide. The criterion is chosen such that it assumes high values for truly abnormal slides, and low values in slides with normal cytology. Unfortunately, the two distributions overlap. Let us assume the criterion is the number of cells detected by the cell classification module that are declared not to be unequivocally normal. Most clinical specimens with perfectly normal cytology may offer a few objects like that—they just happen to fit the values of the few features that are included in the discrimination function. Thus a threshold is set, allowing some such objects in a specimen and not triggering the slide classifi-

cation module to raise an alarm and to flag the slide for human review.

The problem is that there are truly abnormal slides in which only very few abnormal cells are present or detected, owing to clinical sampling errors or other reasons, for example the system design characteristics.

Thus, at least theoretically, it could be that the cell classification module unequivocally recognises a cell as a tumour cell, but the slide classification module, with its threshold, releases the slide as normal: there were too few alarms. The system has a built in false negative rate.

It is a decision that no human diagnostician would ever consciously make. But, as designed, an automated decision maker is not endowed with consciousness. The trend in information technology is in the direction of developing some level of machine consciousness in reasoning procedures. However, as it is now, the signal detection paradigm is not appropriate for clinical decision making as a matter of principle. In signal detection engineering an error is tolerable. Its inevitability is known and codes are designed with redundancy so that errors are self correcting. No such redundancies may be provided in the primary screening decision module.

As decisions are relegated to automated processes—be this in an algorithm, in an analytic strategy, in a hardware module—it becomes increasingly apparent that the particular needs of diagnostic decision making will require the development of its own decision procedures. The reductionist strategy of mathematical/statistical methods is not appropriate when an individual patient's situation needs to be assessed. Paradigms such as the signal detection theory, when applied without modification, may lead to wrong decisions in principle. But there is another risk, that of an engineering model not being sufficiently detailed and thus leading to wrong diagnostic decisions occasionally.

Many, if not most, engineering models assume a well defined or at least highly controlled environment. They rarely have provisions for handling exceptions. One very useful strategy for the accumulation of diagnostic evidence, such as when a number of diagnostic clues are sequentially evaluated, is the use of an inference network. Inference networks provide a numerical measure of the confidence that a set of diagnostic clues offers in support of a given diagnostic outcome. They inherently operate in a range of low relative likelihood ratios and thus lend themselves for the utilisation of weak features. Uncertainty is managed at two levels. Diagnostic outcome and observed diagnostic clue expression are related by a conditional probability matrix. Given a certain diagnostic clue is expressed at a particular level, the probability that a specific diagnostic outcome applies is provided.

These conditional probabilities are derived from a large set of observations, using well established statistical methods. The second level considers the specific case. How strongly does the observed clue expression suggest that

it represents a particular level? This is expressed by a relative likelihood ratio that the observation belongs to a given membership function for clue levels.

The practising diagnostician does not have to worry about the numerical calculations. Such diagnostic decision support systems communicate with the medical personnel through ordered reference images. The diagnostician merely moves a cursor to imagery that corresponds to the case under assessment. Inference networks thus are eminently practical and useful for diagnostic decision making, for teaching, for studies of the diagnostic recognition process, and for understanding the relative values of different diagnostic clues.

However, the common inference network is not designed to handle exceptions. If a diagnostic clue is not expressed close to the level expected for the diagnosis, or is even missing altogether, the confidence level for the appropriate diagnosis may suffer and even allow an erroneous decision.

One can modify an inference network to respond correctly to exceptions in diagnostic clue expression. At the simplest level such a modification involves provisions for clues that are known to be erratic in their expression. Then the procedure may rely on the correlation of clue expression at other evidence nodes in the network, recognise that an exception has occurred, and take appropriate corrective action.

At a somewhat more elegant level a decision support system may have the capacity to monitor the development of confidence in the alternate diagnostic outcomes. Based on information gathered from the course of previous evidence accumulation, the system may alter its strategy, even backtrack and revise decision steps already taken, including revising the level of confidence.

With such capabilities a diagnostic decision support system comes close to an automated reasoning process. The value of a reasoning capability is that it overcomes the rigidity and blindness of an algorithmic, automated procedure.

Information discovery from digital knowledge

What we have learned is that a field of knowledge as detailed, varied, and refined as diagnostic histopathology can indeed be greatly further enriched by options offered by information technology. It also has become quite clear that to do justice to the exacting intricacies of diagnostic evaluation great care has to be exercised in applying analytic procedures from other disciplines. Analytic methodology may in fact be better developed *de novo* for diagnostic decision making in many instances.

Where, then, can we expect useful future development and what is the potential for gains in diagnostic knowledge?

To date, digital representation of histopathological data has primarily led to an analytic evaluation of image information. The link to the established diagnostic knowledge, which is

represented in linguistic form, has been three-fold. Digital analytic results were used, side by side, with traditional diagnostic terms and concepts; digital results were related to pictorial reference imagery; and digital representations were related to linguistic concepts by so called interpretive transforms. These are a complete prescription for the computation of a set of entities which in their values define and delimit in numerical form exactly what a given linguistic term denotes.

Principally novel diagnostic information has been gained from the extraction of digital diagnostic clues from histologically normal appearing tissue of organs harbouring a lesion. These clues point to the presence of preneoplastic or malignant lesions. They are offered in patterns in the numerical data which are statistically significant in their diagnostic value, but which do not leave a visually perceived trace in the histopathologic image—almost in the manner of a digital watermark.

Such diagnostic information is only accessible by computation. Computation thus truly expands our ability to perceive and the gained information may properly be denoted as “digital knowledge.”

However, practically all of our knowledge in diagnostic histopathology and most newly reported knowledge presently exists and is communicated in linguistic or pictorial form. Information technology offers the capability to manipulate and organise both linguistic and pictorial knowledge in ways that could eventually lead to a systematically collected body of digital knowledge of histopathology.

Modest beginnings have been made with diagnostic expert systems or with interactive clinical diagnostic decision support systems, but one could envision much more comprehensive developments.

Information technology offers the capability to manage, research, and evaluate vast amounts of data in ways that have previously simply not been feasible. It allows searches in otherwise seemingly impenetrable data masses. The above mentioned detection of visually unperceived diagnostic clues in histological imagery provides an example of just this kind of information discovery.

The data masses to be organised and searched could be in symbolic form. They do not have to be numerically derived knowledge but could represent any other data type, from patient case histories or response to treatment at one extreme, to molecular genetic data, or immunohistochemical data, and established steps in cellular metabolism and cell transformation at the other.

Large databases can surely be expected to reveal subtle patterns in the linguistic/numerical characterisation of lesion progression. Such patterns may well be formed from a substantial number of “weak features.” They may therefore not have been readily detected and correlated with patient outcome in the past. Thus the search for multivariate weak patterns and the development of search and detection methodology for a specific charac-

terisation of cases may well be a fruitful endeavour.

A further, unaddressed problem is the rising tide of research findings and data. For even a narrowly restricted diagnostic application it has become very difficult to absorb, evaluate, and correlate the new knowledge into an individual’s knowledge store, even with the help of survey papers, reviews, and abstract services. It would be eminently useful if such data could be entered in a knowledge base structured in such a manner that automated reasoning procedures could be applied.

Research in the cognitive sciences has produced some hypothetical models of how knowledge is stored in the human brain. For a strictly limited domain such as diagnostic histopathology, some of these concepts might be practically implemented.

Information, according to such models, is stored in clusters of concepts which are closely associated and form a small network, or context. Each concept in such a network may have associations with concepts in other networks as well, and there may be a strengthening of these associations as they are invoked. Incoming information, or a “message,” is processed along a pathway that establishes itself as the conceptual elements in the message follow their associations, throughout the collection of networks.

The issue here is not the manipulation of information per se, but rather how the results of such message processing can be related to meaning.

“Meaning” is assigned to information when information is related to knowledge already existing in the information processing system. This may involve confirmation that an identical pathway through the collection of networks has been experienced before, and has been consistently associated with a given consequence, or that a specific difference in pathways is now associated with a different consequence.

It may involve testing for analogies; it may involve checking the frequency with which pathways or portions thereof have been traversed before; it may require the collection of additional information so that a conclusion can be reached that has a consequence. A very simplistic example may illustrate the establishing of an analogy and the resulting machine learning. Let it be assumed a pathway links the concepts: lesion, progression, nucleus, ploidy, genomic instability, high grade aneuploidy, mitosis, asymmetric mitosis, near diploid clone, microinvasion. It is a sequence of links that has been derived from observations in high grade PIN lesions. One possible automated reasoning sequence for data from a different organ site might be:

nuclei, ploidy, genomic instability, aneuploidy \Rightarrow microinfiltration \Rightarrow near diploid stemline? outgrowing clone?

A processing system of this nature would automatically accumulate numerical records on the frequency of concept connections, sequences of concept connections, and similarity of reasoning pathways.

At this point automated reasoning procedures approximate some capabilities which otherwise are only attained by an intelligence. In addition, though, the system provides a very important novel source of information: the statistics of pathway use throughout the knowledge base—that is, strictly digital knowledge.

One should be sparing with the term “machine intelligence.” Practically all software said to provide intelligent functions, such as diagnostic expert systems, implement canned human intelligence rather than having an intrinsic capability to find and recognise an intelligent solution to a problem. It is the automated recognition of a possible solution as such that is the difficult part. It is here that the research in cognitive sciences on the connection between information, knowledge, and meaning is relevant.

A comprehensive evaluation of findings and reasoning sequences, taken from published reports, of numerical analytical data and clinical findings would result in a numerically defined record of all explored pathways. This in turn offers a broad overview of current knowledge but also suggests areas of sparsely explored extensions.

Digital knowledge thus may indeed lead to significant information discovery, and processing systems might be designed to allow a true evolution of capabilities.

Conclusions

Immanuel Kant defined intrinsic limits to our ability to perceive reality. But he wrote a lifetime before Darwin and thus made no allowance for the power of evolution. Ironically, it was Konrad Lorenz, the renowned ethologist and his successor to the chair of philosophy at the University of Königsberg 100 years later, who pointed to biological evolution and the ensuing ability to form novel concepts as the key to gains in perception and understanding of the real world. Digital technology represents a development of similar evolutionary potential. It will lead to the gain of knowledge of a distinctly different nature—digital knowledge. This will not only increase our ability to perceive diagnostic and prognostic information, but may well help us understand the kind of complex processes involved in the development and progression of malignant lesions.