Support for Expert Decision-Making: Background and Prospects

A. M. Khavtorin

Scientific Research Institute for System Analysis of the National Research Centre Kurchatov Institute, Russian Federation; Ahavtorin@mail.ru

Abstract. We reviewed research and methodologies related to the formalization of expert knowledge and the support of expert decision-making. We examined an approach developed in the 1970s and 1980s by a research group led by I. M. Gelfand, a member of the Academy of Sciences, with the participation of researchers from SRISA. We analyzed the difficulties and limitations associated with applying this methodology. We also discussed prospects for its further development and application to support expert decision-making in healthcare and other sectors of the economy.

Keywords: artificial intelligence, knowledge formalization, expert knowledge, medicine, medical informatics, diagnostic simulation

1. Introduction

Expert decision-making in various areas has been intensively researched since the 1960s. Back then, the first models and procedures simulating human thinking were developed. One of the most significant results is the combinatorial recognition algorithm (CORAL) proposed by M. Bongard in the 1960s. It simulates the human brain's pattern recognition process [5].

The weak point of these models was that they needed a highly formalized definition of the subject area and well-developed decision-making procedures. This significantly limited (though did not exclude) the application of expert systems in poorly formalized fields such as medicine.

In the 1970s and 1980s, I. Gelfand and his team analyzed the structure and mechanisms of expert knowledge. It resulted in the so-called ‘expert games’ method developed by Gelfand for formalizing expert knowledge. Diagnostic simulations have been tested in several areas of medicine, mostly for predicting the development of serious conditions in patients, and have proven to be highly effective. Further development and commercialization were hampered by its cumbersome nature and the unavailability of sufficient computing power.

Nowadays, with the rise of artificial intelligence and virtually unlimited computing resources, Gelfand's diagnostic simulation may get a new boost. The current AI limitations can be systematically circumvented using expert knowledge analysis and formalization.

Medicine is still the most promising field for diagnostic simulations.

2. Background

Since the 1960s, the progress of information technologies and cybernetics has given rise to several waves of human expectations, hopes, and disappointments. One of them is the so-called "artificial intelligence" (AI), allegedly omnipotent. Besides sci-fi scenarios, the society continues to believe that AI will:

– begin to think like a human being
– be able to make engineering and managerial decisions, substituting for humans
– take on both routine tasks and the management of complex processes
– radically change every job profile.

These expectations have overshadowed the concept of decision-making support technology, which most accurately describes the real place and role of AI. Another important issue is training AI with correct information to infer conclusions. This aspect is particularly acute for rapidly growing generative transformers.

In this study, we will use the following assumptions:

1. AI systems are considered decision-making support tools
2. The application of AI, like any other technology, is only effective in appropriately selected areas
3. The key aspect of AI efficiency is the possibility and level of formalization of the subject area.

2.1. Knowledge Formalization

Some of the first AI tools (expert systems) were in medicine. The hypothesis was that medical diagnostics could be easily represented as a tree of options, with a minimum degree of uncertainty.

In some cases, this approach indeed demonstrated greater diagnostic accuracy (or predicting the disease progression) compared to human physicians. However, such systems turned out to be highly context-dependent.

The result varies depending on the diagnostic and examination methods used, the approach of the specific doctor, etc. The root cause is an incorrect attempt at formalization of the patient definition and the decision-making procedures. The researchers failed to develop an appropriate definition language that would be understandable and unambiguously interpreted by all stakeholders.

Developers of expert systems, mainly highly skilled mathematicians and engineers, approached medicine in the same way as engineering: create a simulation model and find the optimal solution. As a result, a formalized definition of the problem and its object was initially developed independently of the problem originators (doctors), and in a (meta) language incomprehensible to them.

As noted above, this approach demonstrated good results in specific cases, but was fundamentally flawed.

Subsequently, a team led by I. Gelfand, I. investigated knowledge formalization in greater detail.

2.2. Problem Definitions and the Need for Appropriate Language

Key points of a correct formalization are the problem definition and the selection (creation) of an adequate problem/research object definition language.

Problem definition is beyond the scope of this study. We only note that definitions of identical problems can be radically different depending on the researcher's approach, current context, higher-level objectives, etc. Actually, this is called a "point of view" in systems analysis.

Let us outline only the key requirements for a correctly defined problem:

1. A clear, unambiguous list of research objects (a cohort of patients in medicine)
2. For each research object (patient), the system should give a clear answer to the question asked. In some cases, it is acceptable not to give any answer at all.
3. The solution must be verifiable and reproducible.

An appropriate problem/solution definition language is a key aspect of formalization. Yu. Vasilyev et al. [7] discussed such a language for the definition of complex systems. In the field of medicine, a language is appropriate if:

1. Problems and solutions are sufficiently easy to define in the language
2. The language is unambiguously understood by both the problem originators (doctors) and non-doctors.
3. The language allows the user to easily edit the problem/solution definition if necessary.

In general, we adhere to the hypothesis that doctors represent a patient profile using a small number of metrics. As in statistical models, we can claim that 10% of the metrics available to the doctor account for 90% of the information.

Indeed, given the large amount of information obtained during routine examinations and interviews with patients, not to mention special tests, doctors can successfully and quickly cope with the abundance of such information.

An example of an appropriate language is the one used to convey information between doctors. Doctors can express all the essential information about a patient in just a few sentences. In other words, they have a concise, succinct professional language suitable for easy definitions.

The development of an appropriate formalization language will be covered in detail in the forthcoming doctoral dissertation. It is worth noting that special attention will be paid to language generalization, making it applicable to various fields.

3. Diagnostic Simulation as a Knowledge Formalization Tool

In the 1980s, a team of mathematicians and doctors led by I. Gelfand developed an approach to expert knowledge formalization called "Gelfand's diagnostic simulation." The first and most effective applications of the approach were in medicine (see below). Diagnostic simulations were also used to predict the outcomes of snow avalanches and assess oil reservoir productivity [1].

In most cases, the approach gave excellent results, but its cumbersome nature was insurmountable in the 1980s and 1990s, preventing its widespread use.

3.1. Approach Essentials

The diagnostic simulation was based on the following hypotheses (subsequently confirmed empirically):

1. An expert’s knowledge is implicit. It manifests itself as intuitive decisions.
2. The expert makes a decision from a small number of significant features (usually no more than 7).
3. The expert is unable to identify and describe these features.
4. Any attempt to build a decision-making procedure on "let the expert explain in detail how they think" is ineffective.

The diagnostic simulation reproduces real-life medical situations. On the one hand, with this approach, the simulation is close to real medical diagnostic and treatment practice; on the other hand, it makes formal analysis of the decision-making process accessible.

During a diagnostic simulation, the facilitator (mathematician) provides the doctor with bits of information about the patient. From the doctor's responses, we can iteratively extract the features that the expert prioritizes. To illustrate the process, below is an excerpt from a transcription of such a simulation [1][2].

Example 1.

MATHEMATICIAN (M). A patient with myocardial infarction (MI) has just been admitted to your department.

Comment. It was agreed in advance that the patients in the simulation had transmural MIs, were admitted no later than 48 hours after the heart attack, and did not require transfer to the ICU. The purpose was to narrow down the questions about treatment, prognosis, etc.

DOCTOR (D). How much time has passed since the onset of the heart attack?

M. 3 hours.

D. Is the pain relieved?

M. The pain is persisting.

D. Is there cardiac insufficiency?

M. No.

D. Tell me the heart rate and blood pressure.

M. Why do you need this information if you already know there is no cardiac insufficiency?

D. First, to verify that there is no cardiac insufficiency; second, the treatment for bradycardia and tachycardia is different, as does treatment for hypertension and hypotension.

Comments. 1. Distrust of the cardiac insufficiency assessment by the attending physician is associated with its high subjectivity, especially in mild cases. 2. Additional questions asked by the mathematician are intended to gather information about the relationships between the features in a real-life situation, when a doctor analyzes the condition of a specific patient. In this case, the doctor deviates from the routine and may note relationships that they would miss or consider insignificant when answering a general question.

M. Heart rate: 80–40; blood pressure 120/80–100/70.

D. What is the respiratory rate?

M. Upon admission, 18; then increased to 24.

D. Any cyanosis?

M. Severe.

D. There is a discrepancy. Apparently, the patient did not have cardiac insufficiency upon admission, but shortly thereafter developed it (severe cyanosis, respiratory rate: 24) and arrhythmia, causing his heart rate to drop to 40.

Comment. At this point, we encountered some imperfection of the questionnaire: it did not fully track the rapid changes in the patient's condition.

M. You are right. On admission, frequent isolated and couplet premature ventricular contractions (PVCs) were observed. Then it was followed by complete heart block. What is your assessment of the severity of the patient's condition?

D. Extremely severe.

M. What is your estimation of the immediate outcome?

D. How old is the patient?

Comment. The age was needed to make a prognosis. In the following example, where the patient's prognosis was relatively favorable, the doctor did not ask this question at all.

M. 47 years old.

D. The prognosis is poor but not entirely hopeless due to the patient's relatively young age.

Comment. The real patient survived.

We completed diagnostic simulations in the following areas [1, 2, 6]:

- Prediction of the outcome of cerebral hemorrhages after different treatments
- MI prognosis
- Prediction of complications after different types of myocardial infarctions
- Prediction of the duration of sinus rhythm following atrial fibrillation ablation/elimination
- Prediction of the healing of duodenal ulcers
- Prediction of the recurrence of stomach and duodenal ulcer bleeding
- Differential diagnosis of purulent meningitis of various etiologies in infants up to one year old.

3.2. Diagnostic Simulation Procedure

Diagnostic simulations, like any sufficiently complex tool, are only effective when applied appropriately.

The authors of the approach did not explicitly address the selection of its applications, as the research was initially conducted in close collaboration with physicians, who intuitively chose suitable cases (*ironically, the expert knowledge guiding the choice of applications was not formalized*). Today, we investigated this field further and proposed the following criteria for selecting the applications:

1. The problem is recurrent.
2. The history of successful cases is available
3. The solution requires the involvement of an expert
4. There is a shortage of experts.

A diagnostic simulation is divided into stages as follows:

1. Selection of the specific area

2. Problem definition, preliminary data structuring
3. Compilation of a multi-purpose questionnaire
4. Creation of an appropriate language for the research object (patient) definition
5. Development and verification of the decision-making rules.

4. Prospects for Diagnostic Simulations

As noted above, diagnostic simulations were not further developed, largely due to their intrinsically cumbersome nature. With currently available computing power and artificial intelligence, there is reason for optimism about the revival and widespread adoption of this approach, extending beyond medicine.

Why, given the maturity of AI technologies, especially generative transformers, can diagnostic simulations find their applications today? The real effectiveness of decision-making support systems depends on the availability, size, and quality of the training dataset. For instance, AI is successfully applied in medicine if there is a steady flow of training content. In this case, the quality of diagnostics is no worse than that of a human doctor. A striking example is the analysis of X-ray images. AI successfully solves relatively simple, routine tasks. Unconventional tasks and rare cases require human expert intervention: there is just no training dataset for the AI system.

It is similar for generative transformers. Their outputs are probabilistic by design and depend primarily on the training datasets. There are generative AI systems trained on specialized texts (e.g., digitized medical records) exists, but they are not publicly available.

Current AI solutions do not cover expert knowledge.

Diagnostic simulations can reduce expert workload in some areas (see above about the selection of applications). Local generative transformers augmented with specialized information can be widely used in such simulations to generate scenarios and create specific cases.

The approach will be further developed with a focus on its medical applications. In this area, it can mature fast, since most of the Gelfand team members are still alive and ready to collaborate. We will partner with Burdenko's Research Institute of Neurosurgery. The preliminary list of applications is:

1. Indications for surgical/radiological treatment/observation for benign tumors of the skull base:

- meningiomas
- schwannomas
- neurofibromas
- osteomas
- multiple benign neoplasms.

2. Indications for various types of neurosurgical treatment

- neurovascular conflicts
- Chiari malformations
- hydrocephalus.

The preliminary analysis will identify the 3-4 most prominent applications from the above list.

The approach will be made more universal, not limited to medicine. A situation with expert knowledge similar to that in medicine also exists in many other areas of science and technology. Based on experience with diagnostic simulations outside medicine, the approach can be used in geology (exploration, reservoir productivity assessment), the operation of complex engineering systems, and agriculture.

The approach is also well aligned with the Data-Driven and Digital Economy national initiative and artificial intelligence technologies.

5. Conclusion

This study presents an approach to the formalization of expert knowledge and the analysis of expert decision-making developed under the guidance of I. Gelfand in the 1970s and 1980s. We analyzed the options for further development and various applications of this approach.

This paper forms part of a forthcoming doctoral dissertation in systems analysis, information management and processing, and statistics.

References

1. I.M. Gelfand, B.I. Rozenfeld, M.A. Shifrin Essays on the Collaboration between Mathematicians and Physicians. Moscow: Editorial URSS 2005 (in Russ.)

2. I.M. Gelfand, B.I. Rozenfeld, M.A. Shifrin "Diagnostic Simulations" for Medical Diagnostics and Prognosis. Journal of Psychology, 1985 (in Russ.)

3. M. Yakubovich Games Played by Mathematicians and Physicians, https://zdrav.fom.ru/post/igry-matematikov-i-vrachey (in Russ.)

4. Yu.B. Kotov Formalization of a Physician's Professional Knowledge for Medical Diagnostics: Doctorate dissertation. Moscow, 2002 (in Russ.).

5. Bongard M.M. The problem of Recognition Moscow.: Fizmatgiz, 1967 (in Russ.).

6. Gelfand I.M., Guberman Sh.A., Syrkin A.L., Golovnia L.D., Izvekova M.L., Alekseevskaia M.A. Kora-3 Software for Predicting the Outcome of Myocardial Infarctions // Cardiology. 1977. Vol. 17, No. 6, pp.19-23

7. Yu. M. Vasilyev, I. M. Gelfand, Sh. A. Guberman, M. L. Shik. Interaction in Biological Systems. Moscow: Priroda, 1969, No. 06, 07 (in Russ.)