On the Concept of Class and Its Role in the Future of Machine Learning

Lev Goldfarb

Faculty of Computer Science
University of New Brunswick
Fredericton, NB
goldfarb@unb.ca
http://www.cs.unb.ca/~goldfarb

Abstract. My objective is to explain why a completely inadequate *focus on the two central and inseparable concepts—the concepts of class (of objects) and class representation*—is responsible for the lack of *adequate* progress in machine learning, and AI in general. I suggest that the main reason for this lack of progress is reliance on conventional formalisms, mainly the vector space and logical, which *cannot in principle* support a satisfactory concept of class. On the other hand, the orientation towards *new, class-oriented, representational formalisms*—if the underlying informational hypothesis about the nature of classes in the universe is vindicated—would establish machine learning as a new kind of natural science.

1 Introduction

Machine learning is a burgeoning area of research and one of, if not the central field of AI. My contention, however, is that this state of affairs should not be confused with the maturity of the field, as is usually done. Rather, it is due to the field's *strategic* importance. In this paper, I explain some of the reasons for the former claim about the maturity of the field and suggest that a change in the orientation of the field towards class-oriented representational formalisms is due. My paper [1] is recommended as clarifying related issues.

As mentioned in the abstract, I would like to discuss the inherent, or structural, limitations of the two main conventional formalisms, vector space and logical, which are responsible for the present orientation of machine learning and related areas. Although these two formalisms are quite different, both structurally and "semantically", I would like to explain why none of them can accommodate the (central) concepts of interest, i.e. that of *class and class representation.* I also claim that without the capability of fully addressing these central concepts, no formalism for machine learning can be considered adequate. In other words, I propose to judge the maturity of an information processing field not by the amount of research work performed in the field, as is usually done in relatively new fields, but by the progress towards clarifying its central concepts. Moreover, the critical spirit of science demands that we should be prepared to abandon any information processing formalism once it falls far short in meeting the central informational goals of the field.

I hope that my claim about the centrality of the above two concepts in machine learning is not very controversial: if the central goal is to inductively learn *classes*, the corresponding formalism must, above all, be capable of supporting a satisfactory concept of class (including class representation), otherwise, from scientific point of view, we are engaged in quite dubious activity of *"learning" something we don't know what.* Of course, the usual (and non-scientific) justification of the status quo, which I have heard quite often, is that "it works". Hence, in this paper, among other things, I want to consider the reasons why it *only appears to "work"*.

So, which features of the present state of affairs are responsible for this confusion? I am convinced that the confusion is a consequence of studying classification without the benefit of relying on a satisfactory concept of class. The question that we must ask is this: Can one claim that the information-processing *science* of classification is in a satisfactory state, when the *basic underlying concept of class remains as obscure as it has ever been*? Moreover, is it an appropriate justification of the present situation to claim that the present framework "works" when there is nothing really *scientifically* interesting/revealing in the fact that, for example, in the case of a vector space formalism—having carefully preselected zillion features and a very large training set for the English handwritten character recognition problem—the system begins to classify the incoming data with a very good error rate? After all, such character data was, in the first place, *designed (by humans) to be easily distinguishable.* Is it appropriate, in a key information processing field, to accept such *non-informative* [1] frameworks for classification as satisfactory?

I am also convinced that the above misunderstanding is a serious sign of immaturity of our and, in fact, all information processing fields. We must always remember that

> today, in the age of computing, almost any model, including a very poor one, could be "made [and hence appear] to work" in a number of applications, especially when given enough human resources. That does not at all mean that the model has a scientific merit. What makes a new model scientifically attractive is its important explanatory value, i.e. it must hypothesize a qualitatively new, non-trivial feature of reality that one should be able to verify experimentally (otherwise, at best, it is not a "new" model). If no interesting/novel feature of reality is being hypothesized, the model, accordingly, does not have much scientific value. In AI, one must insist on this criterion to an even greater extent, since there is a strong intuitive expectation that any natural environment is "meaningful", or "full of meaning". Thus, in pattern recognition (and machine learning), from a scientific point of view, a useful formalism is supposed to advance a "useful" hypothesis about the structure of pattern classes in the Universe. [Of course,] . . . there is a longstanding scientific practice concerning the verification of such hypotheses. But, above all, one must have such a hypothesis. [1]

[1] i.e. not clarifying the basic underlying concept of class.

In other words, we must look for formalisms—or, more accurately, *representational formalisms*—that would, above all, clarify the concept and the informational nature of the class, including class representation. Such formalisms, for the first time, would allow for *experimental confirmation* of the class structure hypothesis proposed by them, by verifying it against the structure of classes in nature, *as has been the practice in all natural sciences*. Moreover, I strongly believe that, in developing machine learning and related areas, we should be guided by the *hypothesis that the informational structure of universe is based on (evolving) class representations*. Such orientation would establish machine learning as a new kind of natural science, the first information processing *science*. I am also convinced that without such general orientation machine learning will not succeed. For example, the main consequence of the present, incredibly "promiscuous", classification frameworks—which allows for an arbitrary concept of class, e.g. the "class" of bachelors or the "class" of large stones—is that we still don't know what a class is and hence cannot deal adequately with a wide variety of problems in pattern recognition, classification, data mining, and information organization in general.

2 What is a Class?

In this section, I would like only to hint at a fundamentally new way of looking at objects and their classes that has partly inspired the present paper. [2] In the following sections, however, I will not rely on any *special* considerations of this section.

To address the idea of class, first, it is important to realize that *focusing on the languages humans use cannot help us in this endeavor* (see Section 4). On the other hand, we can agree that elements in a class are bound together via some *similarity* related to the *"nature" of the class*. The main question then is this: How do we understand/interpret "similarity" and "nature of the class"? The reason I used the singular "question" is to suggest that the two ideas should be considered together to clarify each other, i.e. we should view the "similarity" as induced by the "structure" of the class. I am convinced that in order to come to grips with these issues we must look to nature for their clarification, thus orienting our information processing science towards natural sciences. The latter does not mean, of course, that we should follow the structure of existing natural sciences, but rather that our *inspiration* must come from the structure of classes in nature.

So, what is the nature of classes in the universe? Although we have addressed the issue in Part I of [3], there are several general points that would be important to state here. Above all, learning from the history of science in the 20^{th} century, it would be wise—and even necessary— to view classes, as everything else in nature, in the light of their evolution, including their formation, modification, and so on. Moreover, what evolves are not just the classes but also their elements

[2] A much fuller elaboration developed over many years can be found in [2].

[3]. In other words, quite naturally, elements and their classes co-evolve *together*, and the chosen *representational formalism must be capable of supporting such reality*, which is outside the capabilities of the current mathematical formalisms due to their reliance on the representation of an object by a "point" instead of by a structured entity!

Thus, the next step is to choose some natural domain, in which classes play a key role and *which would most explicitly suggest* to us how to think about classes in general. It should not be surprising that I propose to look towards biology and biological categories as serving this purpose. Although biological entities are very complex, their general evolution, if interpreted appropriately, must exhibit the same general pattern as the rest of objects in nature.

> In the 17th-century thought a gulf was fixed between matter and mind, the nature of each being conceived as totally exclusive of and different from the other. The modern idea of nature involves the direct opposite of this position. As the final products of the evolutionary process are life and mind, and as these are its higher phases, we must presume that *what is evolving theroughout the process is more fully and adequately manifested in these forms than in forms which are prior to them* and lower in the scale. The further the process advances the more dominant in . . . [its] product are the characteristics of the mind. We can but conclude , therefore, that [these universal characteristics] . . . *must be immanent in all its phases.* [4, my emphasis]

So what do biological classes suggest? One critical feature of biological "objects" stands out: *any* organism is not built from scratch but rather its instantiation requires following some kind of stored "formative history" [4]. It appears quite reasonable to extend this form of instantiation *to all objects in the universe*, including man-made objects, where the "formative" history should be interpreted reasonably broadly. Indeed, stones, pencils, web pages all have their formative histories, albeit of different "kinds". For example a web page has a quite complex formative history related to its conception and execution.

Most important to us, in order to address properly the above question about the nature of the class, I want to propose the following *fundamental ontological postulate* (which is fully consistent with the above hypothesis about the informational organization of the universe [5]): as is the case with a biological organism, *any object in nature* exists along with its formative/generative history. Hence the "similarity" of objects should now be understood as the "similarity" of their formative histories. Next, since it is the object's formative history that reveals its similarity or dissimilarity with other objects, this formative history must be captured in the *object's representation.* As to the nature of the physical storage of an object's generative history, I can think of several hypotheses—the weak,

[3] In this connection, we must remember that even the *meaning* of words used in a human language continuously change with time, which is not reflected in the words themselves.

[4] From the theoretical point of view, the questions of how and where it is stored should not concern us at this stage.

[5] See the last paragraph of section 1.

which suggests that it is stored locally, and the strong, which suggests that it may (also) be stored globally—but this paper is hardly a place for the discussion of this topic.

It is understood that, when constructing such object representation, *a particular agent must rely on its own, internal, representational resources to capture the corresponding formative history*, which is thus agent-dependent, as opposed to the "actual" formative history. And so an agent's sensory resources and experience affect such representations.

Now it becomes clear that classes evolve not only through the appearance of new class elements but also via evolution/modification of their elements' formative histories. Moreover, the concept of *class representation* emerges, then, as directly related to the generative scheme for constructing the class elements, i.e. of their formative histories (or, which from formal point should be the same thing, their representations).

In the following sections, I will rely on one general and I hope not very controversial point: the central concept of machine learning is that of *class*, and it is inseparable from *the concept of class representation*. As a consequence, in order to be considered as a serious candidate for use in machine learning, *the formalism should be able to support a meaningful and informative—both structurally and intuitively—concept of class representation*. By the "structural meaningfulness" of the class representation I mean a *constructive nature of the (formal) class definition*: it means that this definition must explicitly specify—*via the basic (representational) operations* of the underlying formalism—how the class elements are to be constructed [6]. Without such a constructive class specification, it is practically impossible (but by itself it is not sufficient) to insure meaningfulness of the class concept.

3 The Inadequacy of Numeric Formalisms

In this section, I address the situation with the concept of class as it relates to by far the most popular and actually ubiquitous applied framework, the vector space formalism. Because the concept of measurement has evolved with this, numeric, formalism, it is no exaggeration to say that, practically, this is the only universal formalism relied upon in all of natural sciences. Hence, one should not be surprised that the latter is the main reasons for adopting it in machine learning, while the other reason has to do with the attendant fact of availability within that formalism of highly developed (albeit irrelevant) formal machinery, including statistical machinery.

Within this formalism, an object is represented as a point in a multi-dimensional positive inner-product vector (Euclidean) space, hence *the only conceivable way* to delineate a class of objects, if at all, is to construct some "boundaries", linear or non-linear, enclosing the relevant points, separating them from the rest of the space. Moreover, since the inductive learning in this case, by definition, is based

[6] The latter is also a standard requirement for any basic definition in modern algebra.

on a finite training set of vectors, it is expected that the enclosed parts of the vector space that were delineated—the *decision regions*—would contain almost all (in a probabilistic sense) points from the corresponding class and almost no points from other classes under the consideration. This expectation is based on the so-called “compactness hypothesis” [5].

Thus, in the ubiquitous Euclidean vector space, a class is specified via the class “boundaries” that are selected from some fixed family of *decision surfaces.* [7] Because these surfaces play such a critical role in the “representation” of the corresponding class, let us, first, consider the relation, if any, between the surfaces and the class being learned. Indeed, does this family of surfaces have *any* relation to the corresponding class of objects? It is easy to see that *no meaningful relations can exist*, since the family itself simply has to be chosen in an ad hoc manner (see the next paragraph). Furthermore, the final specification of class via such surfaces adds practically nothing to our knowledge about the class in question. Consequently, since the “structure” of the decision surfaces tell us nothing about the class, the obtained in this case “representation” of the class is not “meaningful” at all, i.e. it reveals nothing about the nature/structure of the class. On the other hand, it is very hard to believe that *biological class representations* are not meaningful; and there are also serious reasons to believe that the latter, in contrast to the vector space “class representations”, are generative (dreams, paintings, etc., see e.g. [13]).

Second, in view of the intrinsic (algebraic) structure of the vector space, there is no, and there cannot be any, criteria as to which family of decision functions/surfaces to choose based on a particular family of actual classes. Hence, this choice must remain, both formally and intuitively, ad hoc, i.e. absolutely extraneous to the underlying structure of the class and, in case of non-linear surfaces, also structurally extraneous to the underlying (linear), or *representational*, structure of the vector space itself.

Third, since the choice of the family of decision functions must remain ad hoc, from a pragmatic point of view, one should assume that *any* small (especially non-linear) perturbations of the decision surfaces in resulting family of decision surfaces is as good as the original. Therefore, *we have no reliable/stable formal structure to associate with the class representation*, which from formal point of view is a *very* unsatisfactory situation.

Putting all of the above considerations together and looking afresh and as impartially as possible at the class decision surfaces in a vector space, it is hard not to admit that they cannot serve the role, either structurally [8] or intuitively, of the corresponding class representation. Thus, *the best that the vector space formalism can offer for machine learning is “classification” without the concept of class, if such a thing make sense at all.* Of course, it is always possible to defend the status quo by denying the importance of the capability of a representational

[7] Actually, there are uncountably many of such families.

[8] Probably, the most important point here is that they cannot play the role of a constructive specification of the class.

formalism to accommodate meaningful concepts of class and class representation, but I find such arguments just that—a simple justification of the status quo.

Finally, two interesting (peripheral) observations are in order: one about the metric restrictions of the Euclidean vector space and the other about the *representational conservatism* associated with its use. The first observation concerns the lessons that should have been, but were not, learned from one of the best known physical theories, special relativity theory, which emerged about a hundred years ago, including the introduction of the Minkowski space. This theory clearly implies that some variables—for example, time and space—are "non-commensurable" and this fact should influence the (ubiquitous Euclidean) inner product that specifies the geometry on the underlying vector space, i.e. the inner product must now become indefinite instead of positive-definite. Surprisingly, these major developments in physics and mathematics had little effect on the areas of applied mathematics (including statistics) that had a direct impact on machine learning. Despite my major work [6], written over twenty years ago, several dozen papers, e.g. [7–10], and at least two recent doctoral theses [11, 12] discussing the utility of the pseudo-Euclidean spaces for machine learning, pattern recognition, and data mining, the latter work had relatively little effect on the mainstream research work in these areas. I mentioned this mainly to emphasize a *strong representational conservatism of researchers* and not just in machine learning but practically in all other areas, even when it comes to relatively minor modifications in the *underlying structure* of a conventional data representation formalism as opposed to the technical growth *within a fixed representational formalism*. Thus, my second observation has to do with this quite understandable, but often underestimated, sociological fact that any non-trivial modifications in a conventional "representational" formalism (which are quite rare)—or, its replacement by a completely different one (which are historically extremely rare)—*require a considerable retraining effort on the part of the researchers involved*.

4 The Inadequacy of Logical Formalisms

Logical formalisms emerged as the result of historically sustained attempts to develop a *formal* "restriction" of the ordinary human languages, mainly for the needs of mathematics. Two key points are worth keeping in mind, both being the consequences of this observation. First, in many ways the structure of all logical formalisms was inspired by the structure of written, mainly Indo-European, languages (subject, predicate, etc.). And second, implicitly following the use of traditional languages, in developing logical formalisms, no serious effort has been invested to bypass *the human mind as the only "consumer", or interpreter, of such formalisms*. In other words, such formalisms are not strictly "representational" in the sense that they are not *directly interfacable* with the actual "physical", or external to the mind, events in nature, since it had been implicitly presumed that these formalisms will be used by humans *for the purposes of more precise written communication between humans*. In this connection, I like to quote an important observation by one of the prominent logicians of the first half of the 20^{th} century, Bertrand Russell, who in the second half of his life changed his views on the primacy of logical formalisms: "Nature herself cannot err, because *she makes no statements*. It is men who may fall into error, *when they formulate propositions*." [14, emphasis is added]. This "obvious" observation is very important to keep in mind, simply because it reminds us that *our present obsession with varius kinds of propositional (and non-propositional) languages* is not bringing us any closer to mother nature.

To give away the plot of my main argument, historically, human (and other, artificial,) *languages have evolved on top of more sophisticated/powerful sensory mechanism*, where the latter is directly interfacable with the "physical", or external, events and is responsible for producing class representations. In particular, the language mechanism does not have and does not need the capability to construct a wide variety of class representations: it just needs *to label and manipulate the outputs of the sensory mechanism*. Therefore, emulating language structure does not lead us to the elucidation of the nature and representation of the class.

To clarify the role of speech and language mechanism, the main points I want to address have to do with the evolutionary place of those as well as sensory mechanism. Again, I hope that *today* my view of it as the central information processing mechanism is not as controversial in AI as it might had been one to two decades ago [9]. In fact, the biological evidence supporting my position is really overwhelming. To realize this, it is enough to recall that except humans no other species have language, which, moreover, was acquired *only* within the last million years. It goes without saying that by that time all our sensory mechanisms (for various sensory modalities) were almost fully developed: the relatively minor modifications in the sensory machinery that have occurred since then were due to the evolution of some sensory functions to support the emerging speech and language.

So what is the role of sensory mechanisms? There is plenty of evidence to suggest that the main role of sensory mechanisms is the *classification* of incoming stimuli to support the main functions of life, i.e. orientation in the environment [16, 17, 18], and of course, to be able to classify one must have the ability to *represent* the relevant classes of objects. This answer to the question would become much more apparent when the above hypothesis about the informational organization of the universe (see section 2) is gradually verified. There is, however, some "evidence" towards this end. Indeed, sensory classification and class representation mechanisms have been evolving over several billions of years and are ubiquitous throughout all organisms, including bacteria. The sheer variety of such mechanisms—*which have evolved equally easy in extremely varied environments, including those of the early Earth*—suggests that their efficient development must have been guided by a *single* abstract form of class representation,

[9] See Nilsson's paper [15]: "This paper presents the view that artificial intelligence (AI) is primarily concerned with propositional languages for representing knowledge and the techniques for manipulating these representations."

which in turn could have come about only because it preexisted the biological evolution (according to the above informational hypothesis).

Returning to the nature and structure of the natural "language mechanism", as I have proposed in [19] (and explained above), such a mechanism, of necessity, relies on the main *outputs* of the "sensory mechanism", i.e. on class representations. And it is this organization that allows the language machinery to deal effectively—via nouns, verbs, adjectives, adverbs, and syntax—with various attributes of the learned classes, relationships between the classes, etc. Moreover, if this hypothesis about the nature of the sensory mechanism is true, which is very likely, then the primary biological storage mechanism must have evolved to support *exclusively*, or at least *mainly*, the storage of classes, i.e. of class representations. Hence not only the inputs but also the (stored) "outputs" of the language mechanism should conform to the "class format". For example, it is quite possible that the abstract verb "to love" is stored by means of the association with the potentially infinite class that in the conventional set-theoretic language can be described as follows: the class of pairs, in each of which the second component is either an agent or an agent's state standing in the particular relation to the first component which is an agent (e.g. "nature loves to hide"). [10]

I would like to draw attention to the central role of generativity in the Chomsky's formal grammar model as the syntactic model of language. What is interesting about this, syntactic, side of language modeling is that it is fully consistent with the proposed above generativity of the class representation. So, according to the above hypothesis about the primacy of sensory, or class based, mechanism, I suggest that the former generativity is something that was "inherited" from the latter one. [11]

Thus, in complete contrast to what Nilsson was proposing in [15] (and what many AI researchers have often implicitly assumed)—that "it seems reasonable to distinguish between peripheral and central processes, in which the *peripheral* ones are those that are quite close to the boundary between the environment and the animal or machine that inhabits it"—I am suggesting that it is the language mechanism that is "peripheral" and built on top of the central, sensory, mechanism, where the latter is responsible for class representation and classification. Accordingly, since the language mechanism was not supposed to deal with external, or sensory, events, to develop the science of machine learning we need to look for formalisms fundamentally different from propositional formalisms in-

[10] This class can be represented in the above generative manner.

[11] I believe the main reason why the concept of generativity, *so far*, turned out to be not very useful in machine learning and many related areas has to do with the much deeper, representational, fact that, under the conventional string representation, generativity cannot be integrated into this representation, i.e. it is not captured in the string representation, and must be stored *separately*, in the form of the grammar itself. As a consequence, given a small training set of strings, it is imossible to recover reliably the original grammar. In our representational formalism [3], this inadequacy has been remedied.

spired by the surface structure of Indo-European languages. As to the argument similar to the above (that "it works"), my above considerations also apply.

Finally, I want to mention just three (almost "randomly" selected out many) relatively recent references discussing inadequacies logical formalisms. Marvin Minsky in [20] discusses the inadequacy of logic-based approaches to AI and the need for approaches not based on logic. Mike Oaksford and Nick Chater in [21] discuss the "inadequacy of logic as an account of everyday human reasoning". I also recommend section nine of my paper [1], which discusses some "additional" inadequacies of the logical formalisms.

5 Conclusion

My objective was to draw attention to a completely inadequate focus that the two central and inseparable concepts—those of *class of objects and its representation* [12]—have received in machine learning and AI in general. I suggested that the main reason for this situation has to do with the non-trivial fact that conventional formalisms, e.g. the vector space and logical, *cannot in principle* support a satisfactory concept of class, and hence *within them* these concepts cannot be attained. These concepts can be reached only within radically different, class-oriented, representational formalisms, which were not considered here and which should clarify the nature of *structural* object representation (see [3]).

On the other hand, a more thoughtful research into the above two concepts also suggests the corresponding informational hypothesis about the nature of classes in universe, which, if vindicated, would establish machine learning (and in fact AI and computer science) as a new kind of *natural science* studying the nature of information processing in universe.

Acknowledgments. I thank Dmitry Korkin for his comments and Ian Scrimger for his help with formatting.

References

1. Goldfarb, L.: Representational formalisms: What they are and why we haven't had any. In: Goldfarb, L. (ed.): What Is a Structural Representation (in preparation) http://www.cs.unb.ca/~goldfarb/ETSbook/ReprFormalisms.pdf
2. Goldfarb, L. (ed.) : What Is a Structural Representation (in preparation) http://www.cs.unb.ca/~goldfarb/ETSbook/ETS5.pdf
3. Goldfarb, L., Gay, D., Golubitsky, O., Korkin, D.: What is a structural representation? A proposal for an event-based representational formalism. In: Goldfarb, L. (ed.): What Is a Structural Representation (in preparation) http://www.cs.unb.ca/~goldfarb/ETSbook/ETS5.pdf
4. Harris, E.E.: Nature, Mind and Modern Science. Gorge Allen & Unwin, London (1954) 205–206

[12] They are related to the class extension and class intention.

5. Arkadev, A.G., Braverman, E.M.: Computers and Pattern Recognition, Thompson Book Co. (1967)
6. Goldfarb, L.: A new approach to pattern recognition. In: Kanal, L.N., Rosenfeld, A. (eds.): Progress in Pattern Recognition, Vol. 2. Elsevier (1985) 241–402
7. Liang, J.J., Clarson, V.: A New approach to classification of brainwaves. Pattern Recognition, Vol. 22, No. 6 (1989) 767–774
8. Gudivada, V., Raghavan, V.: Content-based image retrieval systems. IEEE Computer, Vol. 28, No. 9 (1995) 18–22
9. Graepel, T., Herbrich R., Bollmann-Sdorra, P., Obermayer, K.: Classification on pairwise proximity data. Advances in Neural Information Processing Systems **11** (1999) 438–444
10. Pekalska, E., Paclik, P., Duin R.P.W.: A Generalized kernel approach to dissimilarity-based classification. Journal of Machine Learning Research, Special Issue on Kernel Methods, Vol. 2, No. 2 (2002) 175–211
11. Gutkin, A.: Towards Formal Structural Representation of Spoken Language: An Evolving Transformation System (ETS) Approach. Ph.D. thesis, School of Informatics, University of Edinburgh (2005)
12. Pekalska, E.: Dissimilarity representations in pattern recognition. Concepts, theory and applications. ASCI Dissertation Series **109**, Delft University of Technology (2005)
13. Jackendoff, R.: Patterns in the Mind: Language and Human Nature. Basic Books (1994)
14. Russell, B.: Wisdom of the West, Crescent Books (1989) 311
15. Nilsson, N.: Artificial intelligence: Engineering, science, or slogan? AI Magazine 3(1) (Winter 1981-82) 2–9
16. Cohen, H., Lefebvre, C. (eds.): Handbook of Categorization in Cognitive Science.Elsevier, Amsterdam (2005)
17. Matlin, M.W., Foley, H.J.: Sensation and Perception. Allyn & Bacon, Boston (1997)
18. Goldstein, E. B.: Sensation and Perception. Wadsworth, Pacific Grove, CA (2007)
19. Goldfarb, L.: On the foundations of intelligent processes I: An evolving model for pattern recognition. Pattern Recognition 23(6) (1990) 595–616
20. Minsky, M.: A Framework for Representing Knowledge. In: Winston, P. (ed.): The Psychology of Computer Vision. McGraw-Hill, New York (1975) 211–277
21. Oaksford, M. and Chater, N.: The probabilistic approach to human reasoning. Trends in Cognitive Sciences, Vol.5, No.8 (2001) 349-357. http://www.dectech.org/publications/LinksNick/ReasoningAndDecisionMaking/The probabilistic approach to human reasoning.pdf