Applied Logic: a Manifesto

Lawrence S. Moss

March, 2005

Contents

1 What is applied logic?

My main purposes in this essay are to introduce applied logic as a research area, to situate it in a broader context, to make the case that it is a significant and worthwhile enterprise, and to detail some of its research areas. These are my main overt purposes. But my "covert" purpose is to write something that might open new doors for young students with interests in logic. In my younger days I remember the excitement I felt from subjects at the borders of mathematics, computer science, and linguistics. It was not until many years later that I began to think more explicitly about the "politics" of what I and many others have been doing. I think it would have helped me to ponder a manifesto or two along the way, even in my high school or college days. So my hope about this article is that somewhere, sometime, somebody will pick it up and . . ..

While many people know something about logic, I take it that the idea of applied logic will be unfamiliar to all readers. This is because it doesn't exist in the same institutional sense as other fields, and only a few books have the words " applied logic" in their title. There really is no consensus on what "applied logic" means, so in effect, I'm making a proposal here. To explain matters, I'll need to go into a fair amount of detail about logic, and also about mathematics, computer science, philosophy, and other fields. But to keep things short, most of my discussion of these matters will be offered with only minimal support.

The main points The reader can find most of the main points in the section headings and **boldface** lead phrases. Many of these are slogans that I present in a deliberately provocative way.

First things first, we should say what the subject is about. It is always difficult to define fields, but I take applied logic to be defined and characterized in the following ways:

1. It is the application of logical and mathematical methods to foundational matters that go beyond the traditional areas of mathematical logic. The central domain of application at the present time is computer science, but it also has significant applications in other fields.

2. It also is extension of the boundaries of logic to include *change*, *uncertainty*, *fallibility*, and *community*. Far from being the study of matters which are absolutely black or white and never change, and which exist in the mind of a single person, applied logic aims to study communication via brushstrokes of gray. In this way, it is a reconstitution of the study of *foundations*.

3. Its ultimate interest is a concern with human reasoning, so it will ultimately lead to a rapprochement with psychology, artificial intelligence, and cognitive science. But even before this happens, the development of *tractable logical systems* are the most conspicuous *applications* of logic in many fields.

4. Applied logic is an interdisciplinary field, and this has its own set of difficulties and opportunities.

Most of this essay is a discussion of these points. But in the spirit of a manifesto, I do want to make some grandiose claims: applied logic is the most vigorous branch of logic. And if one is interested in current research on the topics that motivate and animate logic in the first place – the concepts of formal reasoning, truth, meaning, paradox, proof, and computation – then applied logic is the place to look.

2 Mathematics and logic, but different from Mathematical Logic

> Logic is the study of reasoning;
> and mathematical logic is the study of the type of reasoning done by mathematicians.
>
> Joseph R. Shoenfield, **Mathematical Logic**, 1967.

I know that most readers will recognize the split in logic between "mathematical logic" and "philosophical logic". At the same time, they may be surprised to hear that the difference is not mainly about whether the subject is "mathematical": philosophical logic has technical sides that use and inspire mathematics. The difference has to do more with what the targets of the studies are. Mathematical logic is a much-better-defined area, and so it makes sense to discuss it first. The quote above is the first sentence from one of the standard textbooks on mathematical logic.

I did not re-read Shoenfield's entire book, but I doubt very much that the word "reasoning" appears many times after the first sentence. Going further, I do not think we can take mathematical logic seriously as a study of mathematical reasoning. There are numerous reasons, and they all echo broader philosophical claims. First of all, one would think that in studying "mathematical reasoning", one would be interested in "reasoning" in other areas. That is, one would expect some sort of engagement with psychology. Yet ever since Frege if not earlier, mathematical logic has rejected the idea of an engagement with psychology at any level. Even putting this aside, in examining the considerable body of work in mathematical logic, we find nothing of key aspects of "flesh-and-blood" mathematical reasoning, such as the use of symbols, diagrams, formulas; hunches, nothing about evidence, and mistakes; nothing about why some types of mathematical reasoning are more interesting than other types; nothing about different mathematical fields and what they have or do not have in common; nothing about how the faculty of mathematical reasoning is acquired; and again, nothing at all about how it is related to any other human faculty. One would think that a subject devoted to mathematical reasoning would in part be interested in *all* of these issues.

2.1 Mathematical logic and mathematics

If mathematical logic is not the study of mathematical reasoning, what is it? Mathematical logic has three aspects:

1. It is a *foundational* discipline which studies *an idealization of* mathematical reasoning, the reasoning done by perfect beings with no resource limitations who reason in a way captured by the axiomatic method. It is mainly concerned with idealizations of the concepts of *truth*, *proof*, *computation*, and *infinity*. It is traditionally divided into four areas that correspond to these: model theory, proof theory, recursion theory, and set theory.

2. It also deals with a host of theory-internal questions. Each area of mathematical logic is now an active branch of mathematics, and like any branch of mathematics, there are many questions whose primary interest will be to those inside.

3. Finally, it is concerned with applications to other areas of mathematics.

I think there are reasons not to be happy with the received view of mathematical logic as "the foundations" of mathematics, but I will not go into that here. I think it is fair to say that this foundational contribution of mathematical logic is what is usually meant by speaking of mathematical logic as a foundation of mathematics, and that despite my (and many others') quibbles on whether it is *the ultimate* foundation, the foundational achievements of logic are of permanent interest.

For many years, the theory-internal questions of mathematical logic were the most important parts of the subject. There are whole fields of active study, such as the theory of recursively enumerable degrees, large cardinals, the fine structure of the constructible universe, and infinitary logic, which are mathematically interesting but really quite far in motivation from any foundational matters or from the study of mathematical reasoning. These questions are the real aim of Shoenfield's book, for example. Theory-internal questions constitute most of any active field, and indeed most of the articles in the *Journal of Symbolic Logic* would fall in this category.

The dream of many logicians has been to apply logic to settle serious questions of mathematics. For many years, this dream went mostly unrealized. There were some exceptions, such as Tarski's work on real closed fields, in some of the celebrated results of set theory that had implications for the foundational questions of analysis; and indeed in the whole field of non-standard analysis. And in very recent times we see significant applications of mathematical logic in mathematics, so much so that perhaps mathematicians outside will get interested in one or another of the branches of mathematical logic. Two areas where this happens are model-theoretic algebra, with its connections to areas like arithmetic algebraic geometry; and descriptive set theory, with its connections to topology, analysis, functional analysis, and other areas.

It seems fair to say that the main thrust of mathematical logic is not the foundational contribution in point (1): this is saved for textbooks and introductory courses, or it comes up only when justifying the subject. The most valued work is in (2) answering hard technical questions about areas that have arisen because of the subject itself, and (3) contributions to more central areas of mathematics. Looking at conference programs and invited lectures, this latter thrust seems on the rise and destined to become the most important one for mathematical logic. Certainly when I talk to young people interested in the mathematica logic, I encourage them to aim for area (3).

2.2 Where applied logic differs

I think that all of the contributions in (1) – (3) above are interesting and good. But I don't think that they are the only interesting things to do in logic. My main purpose in this essay is to present an alternative agenda. It will extend the foundational impulse which motivated mathematical logic in the first place, it will have its own internal questions, and it will be a field with applications – indeed, as the name suggests, the applications are going to be at the center of the subject.

Once again, the closest direction for applied logic is the one that mathematical logic seems to be finished with, the foundational contribution clarifying some of the central aspects of reasoning in mathematical and formal contexts. But this idealization is a two-edged sword. Usually it is presented as an advantage, because *negative results* about the idealization imply negative results about the real thing. For example, it was shown early on that there are natural things that one might want to write a computer program to do which simply cannot be done by an idealized computer. (One example: write a program which looks at other programs and tells for sure whether or not the input program will go into an infinite loop.) Since this cannot be done on even an idealized computer, it follows that it cannot be done on a real computer either. This illustrates how when considering negative results, it might be fine to work with idealizations.

However, the other side of the sword is that sometimes the idealization might be so questionable as to make the inference from the idealization problematic. An example here concerns the most celebrated result of mathematical logic, Gödel's Incompleteness Theorem. There are many people who believe that this result implies that human beings are not computers. This may or may not be the case, but I think it's a mistake to think that the Incompleteness Theorem gives us conclusive evidence, or even suggestive evidence. The inference from the technical result to the philosophical point is questionable precisely because there is no reason to take an understanding of the Incompleteness Theorem, to be a good idealization of intelligence or what it means to be a person.

Returning to my topic, questions about *real* reasoning, or about aspects of it that we can model mathematically, are going to be important for applied logic. In this sense, applied logic is carrying forward the program of applying mathematics to the human world. This is the crux of the difference. For applied logic, mathematics will be a tool to use. And although I think that blends of applied logic and cognitive science will ultimately tell us more about mathematics than mathematical logic has as yet told us, I don't take this to be the one and only goal of applied logic.

2.3 Applied mathematics is good mathematics

> "Applied mathematics is bad mathematics."
>
> Paul Halmos, "Applied mathematics is bad mathematics," in Lynn Arthur Steen (ed.), **Mathematics Tomorrow**, Springer, 1981

This enterprise of applied logic builds on and uses all the results of mathematical logic, but it is not aimed back at mathematics the way mathematical logic is. My argument in this section is that applied logic should be recognized as an area of applied mathematics.

Halmos' quote above is the title of his paper on the subject of the relation of pure and applied mathematics, one of the few papers devoted exclusively to that topic. He writes, that the concept of *motion* "plays the central role in the classical conception of what applied mathematics is all about." And in a passage comparing pure and applied mathematics, he states:

> The motivation of the applied mathematician is to understand the world and perhaps to change it; the requisite attitude (or, in any event, a customary one) is sharp focus (keep your eye on the problem); the techniques are chosen for and judged by their effectiveness (the end is what's important); and the satisfaction comes from the way the answers checks against reality and can be used to make predictions. The motivation of the pure mathematician is frequently just curiosity; the attitude is more that of a wide-angle lens than a telescopic one (is there a more interesting and perhaps deeper question nearby?); the choice of technique is dictated at least in part by its harmony with the context (half the fun is getting there); and the satisfaction comes from the way the answer illuminates unsuspected connections between ideas that had once seemed to be far apart.

Halmos makes it clear that he values applied mathematics, but as his overall title indicates, he is partial to pure mathematics above all else.

I think that applied logic could well be considered as applied mathematics. It is not based on the concept of motion, but as I mentioned above regarding *change*, some evidently related concepts are at the heart of it. Applied logic is about understanding the world, and to a very limited extent, changing it. On the other hand, it is more like a wide-angle lens than a telescope, so the analogy is not perfect. But overall, based on what Halmos writes (as quoted above and elsewhere in the article), I think it's fair to say that applied logic is closer in spirit to applied mathematics than pure mathematics.

I must add that usually applied mathematics is not taken to include discrete subjects. I say "usually" here; sometimes discrete math topics are included in applied mathematics. But one need only look at departments of Applied Mathematics in the USA to see my point. (Incidentally, it seems clear that *theoretical computer science* fits Halmos' criteria for applied mathematics rather well.) And most applied mathematicians would be surprised, I think, to consider any branch of logic in the same category. Conversely, logic is rarely seen as an applied subject. So my entire discussion is intended to make a point that is controversial.

2.4 Applied logic is applied mathematics

> Throughout the centuries the great themes of pure mathematics, which were conceived without thought of usefulness, have been transformed to essential tools for scientific understanding. ... this transformation is now happening to mathematical logic, and ... a subject of applied logic is emerging akin in its range and power to classical applied mathematics.
>
> Anil Nerode, "Applied logic", in R.E. Ewing et al (eds.)
> **The Merging of Disciplines: New Directions in Pure, Applied, and Computational Mathematics**, Springer 1986.

I believe that Nerode is right: applied topics in logic are in the process of coalescing around a set of questions and research agendas that will constitute a coherent subject matter. I would like to think of this as applied mathematics in the same kind of way that other areas are applied mathematics: it certainly involves doing new mathematics, and doing interesting mathematics at that; but the choice of problems and viewpoints is driven primarily by modeling phenomena which exist out in the big world.

Nerode's article is the one of the few I can point to that makes the case for applied logic. His paper is mostly a compendium of examples and does not attempt to systematically present applied logic. He is most interested in applications to computer science, and I'll have more to say about this in Section 4 below. But applied logic is a very interdisciplinary study, with additional contributions and applications from artificial intelligence, cognitive science, economics, and linguistics, and with fundamental interactions with computer science, mathematics and philosophy. Before we get to that, it would be good to contrast applied logic with its much better-known cousin, mathematical logic.

3 Applied philosophical logic

Traditionally, logic has been divided into "mathematical logic" and "philosophical logic." At most institutions in the US which feature significant activity in logic, this division is a useful one. Few people bridge the gap.

I take philosophical logic to be the continuing foundational study that I mentioned above in connection with mathematical logic. In addition, I take it to be the home of formal, mathematical studies of all of the important concepts which somehow did not make it into the purview of mathematical logic. E. J. Lowe[1] holds that the subject's main areas are

> (1) theories of reference, (2) theories of truth, (3) problems of logical analysis (for example, the problems of analysing conditional and existential statements), (4) problems of modality (that is, problems concerned with necessity, possibility and related notions), and (5) problems of rational argument.

I contrast philosophical logic with *philosophy of logic*, and by this I have in mind more the relation of logic to more central branches of philosophy such as epistemology and metaphysics. All of my remarks in this essay are about philosophical logic rather than philosophy of logic.

My feeling is that the pure/applied continuum and the mathematical/philosophical continuum are somewhat orthogonal. Specifically, there are many applied subjects that are applications of topics originating in philosophical logic. These are mainly in Lowe's areas (4) and (5) above. I'll have more to say about one such topic from (4), epistemic logic, in Section 5.2 below. Overall, I think that philosophical logic is the source of many problems and research connections with applied logic. This is mainly because whole areas of philosophical logic have been given new life by connections to computer science. I'll return to this point after discussing the relation of computer science with applied logic.

I also have an overall feeling about the foundational problems that motivated logic in the last century. To be blunt, we're were a different world in 2000 than we were in 1900. Many of the questions that seemed so pressing back then have lost some of their appeal. Very few people today want to fight the old fights about the Axiom of Choice, or about predicativity, or a host of other issues. Instead, we have a host of new questions, and new areas that at this point are in need of mathematical insight: what would models of computation look like which are appropriate to the brain as we know it? What is information? What are the best ways to represent the fallible, uncertain, and sometimes-incorrect knowledge that we all have? What are the most efficient ways for a computer to manage large amounts of changing information? What are we to make of the failure of logic to be a "magic bullet" in artificial intelligence?

3.1 Applied philosophical logic = theoretical AI

The slogan here is perhaps a bit of an overstatement, but the point is that work on the theoretical questions in artificial intelligence often looks back at earlier discussions in philosophical logic. One area where this happens is in the study of knowledge; I'll say more on this below. Another is the study of *context*: how is it to be represented, and what role does it play in reasoning? If one wants to build a robot and make it *rational*, then the hard problem of deciding what rationality means will lead back to the parallel philosophical literature.

4 What does computer science have to do with it?

Applied logic is the most vibrant and relevant form of contemporary logic. It is primarily the study of logic that is relevant to, and in symbiosis with, computer science. So it is worthwhile at this point to go into detail about the relation of applied logic to computer science.

[1] Lowe's survey article is available at `http://www.dur.ac.uk/~dfl0www/modules/logic/PHILLOG.HTM`

4.1 Logic is the Calculus of computer science

Logic is a surprisingly prevalent tool in Computer Science. NSF's Directorate for Computer and Information Science and Engineering (CISE) had a workshop[2] two years ago called "The Unreasonable Effectiveness of Logic in Computer Science." The point is that some areas of logic get used again and again in formulating the main notions of computer science. Relational databases are close to first-order relational structures, and model theory is therefore an appropriate tool. Programming language "types" are best understood with the help of much older tools from logic like typed lambda calculi. The study of abstract data types is quite close to universal algebra, so equational logic is prominent there. Verifying that a computer program or a piece of hardware does what it was designed to do requires formalization, and this formalization inevitably uses the tools of logic. Interestingly enough, the tools in verification often come originally from areas of logic where time and change are studied, so they ultimately derive from philosophical logic. Turning to artificial intelligence (AI), there was a time when AI was taken to be one big application of logic. The celebrated P=NP problem, the problem which has been called "computer science's gift to mathematics", is often cast as a problem in logic: is there a polynomial time algorithm to determine whether a boolean formula is satisfiable or not? And all the other main problems of complexity theory have logical versions.

The widespread use of logic in computer science goes back about twenty or thirty years. Although logic is not seen as a specific area of computer science the way it is in mathematics and philosophy, there are those who believe that logic is even more important in computer science than it is in mathematics, that large parts of computer science are applications of two parent disciplines: electrical engineering and logic. The slogan here is: **Logic is the Calculus of computer science.**

4.2 Computer science motivates logic

Just as physics was a great spur to the development of applied mathematics, so computer science will be a motivating field for applied logic. It is not surprising that nearly all of the applications mentioned above were not applications of existing theory. For the most part, the applications called out new theory, new mathematics. And this new theory is developing at a rapid pace. Here are the examples again, with mention of the new work that has come up: Database theory has given us *finite model theory* and connections of logic and probability theory. The issue of types in programming languages is now of keen interest in category theory, and to follow current developments one really needs a good background in that subject. Programming language semantics has also been a motivating force in current developments in proof theory such as linear logic. The universal algebra/computer science border has given many new questions: for example, the classical questions of universal algebra usually are asked without reference to complexity. Verification has given us numerous flavors of dynamic logic, and I'll return to this in Section 5.2 below. It also has revitalized higher-order logic. AI has lead to non-monotonic logic, to blends of logic and probability, to automated theorem proving and knowledge-based programming. And complexity theory has lead to descriptive complexity theory and to learning theory.

Why has computer science been so powerful of a driving motor for logical applications

[2]There is also a nice survey article by essentially the same people as the presenters of the CISE workshop: Joseph Y. Halpern, Robert Harper, Neil Immerman, Phokion G. Kolaitis, Moshe Y. Vardi, and Victor Vianu, "On the unusual effectiveness of logic in computer science." Bulletin of Symbolic Logic 7 (2001), no. 2, 213–236.

and for developments inside of logic? Here are two related reasons: First, the whole tenor of computer science is toward applications that actually run. Traditional systems of logic are the right place to look to find the appropriate theory, at least at first glance. But at the same time, the main body of technical work in logic is based on idealizations: complexity doesn't matter, mistakes are unimportant, everything is relevant to everything else, the world may be modeled as an unchanging totality of facts, etc. Each of these has to be abandoned or at least seriously modified to make real progress in the fields I've mentioned above. The closer we get to the human world, the more we need to re-think things. And it is this reconsideration of old idealizations which has lead to many new developments in logic.

4.3 Going beyond the traditional boundaries of logic

As I mentioned above, there was a time when the logical paradigm in AI was the leading one. This is no longer the case. In AI itself, even those who do believe logic has a key role often resort to new varieties of logic, such as default logic and non-monotonic logic. These differ from standard logical systems because one can "take back" parts of arguments, or jump to conclusions (in some sense). Going further, *probabilistic methods* are now recognized as critical, not only in AI but also in fields of interest for applied logic, such as cognitive science and computational linguistics. There is a recognition that uncertainty and randomness are not flaws; they are design tools. So connectionist modeling is now widespread in cognitive science. Dealing with uncertainty is a major research area in AI. Statistical methods in natural language processing are widely believed to outperform deterministic methods.

My point here is to suggest that a coming key area for applied logic is going to be some sort of rapprochement with all of the mathematical areas that turn out to be important in modeling on the same set of phenomena. This is especially important for cognitive science, and I'm encouraged by some very recent developments that relate connectionist models to non-monotonic logics.

A postscript: one of the interesting developments in recent years is the degree to which the 'declarative' and 'stochastic' sides *do* turn out to talk to each other. Perhaps this is because both are getting something right. In any case, I take the project for applied logic *not* to be the one of defending declarative frameworks or making improvements on them, but rather the one of accepting the points of the "other" side and working towards a stronger synthesis.

5 Other case studies

I have mentioned that the primary application area of applied logic is to computer science. The applications there are so well-developed that people who work on them might not even be interested in the foundational questions that I take to be an important part of applied logic. In this section, I present case studies and research questions in applied logic whose main motivation is areas outside of computer science.

5.1 Neural networks and non-monotonic logic

People investigating learning, categorization, memory, and other areas of cognitive science often use *neural networks*. There are many different kinds of neural nets, but they share the features of processing information numerically, and of doing so in a distributed way. This contrasts with serial, symbolic processing that is more natural for the computational models like Turing machines. Those kinds of computational models seem better-suited to model activities like logical reasoning. The name "neural" comes from the view that the brain, too, is a neural network. As it happens, it is much easier to use neural network models to "learn" than to give an account of learning in general. It is easier to use the models than it is to understand what they are doing.

The need for some synthesis between the symbolic/serial and numerical/parallel models has been felt by researchers for quite a while. Two people who have done work on this include Peter Gärdenfors and Reinhard Blutner. Their overall idea is view symbolic computation as a higher-level description of what is going on in connectionist models. In other words, we would like to explain emergence of abstract symbols from subsymbolic data such as weights in a trained neural network. The logical tool employed is *non-monotonic logic*, the same subject I mentioned above in connection with logic in AI.

5.2 Dynamic epistemic logic

Modal logic is the study of logical systems which involve some qualification of the concept of truth. For examples, one studies the differences between "true", "possibly true" and "necessarily true". Epistemic logic is a branch of modal logic that deals particularly with concepts having to do with *knowledge* and *belief*. Other branches of modal logic study concepts like *before* and *after* (temporal logic), or *obligation* and *permission* (deontic logic). One can sense that the all of these areas are going to be of interest not only in philosophy but also in cognitive science and artificial intelligence.

Incidentally, modal logic in general and epistemic logic in particular are subject conspicuously missing from mathematical logic. I think this is all unfortunate, because modal logic is one of the most applicable fields of logic, and also because it has connection with many areas of mathematics. One can find to papers where modal logic is mixed with general topology, dynamical systems, universal algebra, and boolean algebra.

It probably would have surprised the early workers in the subject that their ideas would be useful in computer science, but this has indeed happened. Epistemic logic overall is one of the areas of philosophical logic that has benefited a lot from computer science, and I'll go into this below. Just the same, it might have surprised the computer scientists of a few years back that economists have become interested in the subject. Overall, it is today a richer field than ever before.

The dynamic turn I mentioned in connection with applied mathematics Halmos' point that *motion* "plays the central role in the classical conception of what applied mathematics is all about." Interestingly enough, a parallel to motion is playing a central role in many areas of applied logic. This parallel notion is *change*. Models incorporating change are prevalent in computer science, since a computational process is one in which values (of variables, or registers) change. So the logics of computer science are generally logics of change. These themselves are modal logics; temporal logics are the most prevalent kind in applied work, but others are as well. As it happens, ideas from dynamic logic have made their way back into epistemic logic. This happens primarily in connection with the issue of *common knowledge*, a matter which I won't discuss in detail.

There are now areas of epistemic logic where one is concerned with the modeling of epistemic actions and the change of knowledge that comes from them. One proposes and studies models for notions like *public announcements*, *cheating in games*, *wire-tapping*, etc. The models use ideas from both epistemic and dynamic logic, and their study involves a lot of non-trivial mathematics. The overall idea of dynamic epistemic logic is to get a good mathematical treatment of all of the notions I mentioned above. Later, one would like to see how the work plays out in areas like the modeling of conversation between people, or in models for computer security.

Connections with economics/game theory Another area where logic might shed light on matters in the social world is in economics and game theory. Here there are basic questions concerned with how agents interact, what rationality and belief come to, and how communication and action change the world. As it happens, the tools of modal logic that I mentioned above play a role in this work. A further connection to computer science comes in when one looks at *auction theory* or *mechanism design*. One interesting source to look at "Logic for Mechanism Design – A Manifesto" by Marc Pauly and Michael Woolridge.[3] Their proposal is to use a logic called *alternating-time temporal logic* as a language to formally define social procedures such as voting systems, auctions, and algorithms for fair division. The parallel is with the uses of logic in defining (specifying) computer programs; there is also a definite sense in which social algorithms are a "many-agent" version of computer programs. Perhaps the fullest expression of work in this direction is Rohit Parikh's program of *social software*. What all of this suggests to me is that applied logic is poised to dramatically enlarge the traditional scope of logic, and that in some sense it is already doing just that.

5.3 Linguistics, logic, and mathematics

> Precisely constructed models for linguistic structure can play an important role, both positive and negative, in the process of discovery itself. By pushing a precise but inadequate formulation to an unacceptable conclusion, we can often expose the exact source of this inadequacy and, consequently, gain a deeper understanding of the linguistic data. More positively, a formalized theory may automatically provide solutions for many problems other than those for which it was explicitly designed. I think that some of those linguists who have questioned the value of precise and technical development of linguistic theory have failed to recognize the productive potential in the method of rigorously stating a proposed theory and applying it strictly to linguistic material with no attempt to avoid unacceptable conclusions by *ad hoc* adjustments or loose formulation.
>
> Noam Chomsky, **Syntactic Structures**, Mouton, 1957.

For some, mathematics is a fortress that should remain on guard against contact with the world. For others, it is part of the light of the mind, the light by which we understand the world. In this section, I am concerned with the relation of linguistics and mathematics. The fields of linguistics that I have most contact with are *syntax* (the study of phrase structure and sentence structure) and *semantics*; these are also the branches of linguistics that are closest to logic and theoretical computer science.

The most important linguist in modern times is Noam Chomsky. The quote above makes the case why mathematics has something to say to linguistics. This point is immediately appealing to people like me who value the use of mathematics in the social sciences, and I remember being inspired in this direction many years ago.

[3] See `www.csc.liv.ac.uk/ mjw/pubs/gtdt2003.pdf`

Unfortunately, the story does not end here. For various reasons, Chomsky in later work has *reversed* his position. To this day, he is not in mathematical results concerning the grammars he proposes; and the majority of syntacticians follow suit. The set of people interested in the mathematical study of the syntactic formalisms is fairly small. It almost seems like they are a small remnant of candle-holders in a crowd that would rather walk in the dark. However, I am optimistic about the long-term prospects of formal work in linguistics. This is partly because I'm optimistic about the use of mathematics in all fields, and partly because I think the results we already have are interesting enough to pursue further. In any case, the history of the interaction of mathematics and linguistics is an interesting one, and so I'll discuss a few aspects of it. As we'll see, there are echoes of the pure math/applied math split here, too.

Remarks on syntax and semantics in computer science and linguistics Overall, I think that work on syntax and semantic in linguistics is harder than parallel work in computer science. The main reason for this is that computer languages are human creations, and so we know what they are supposed to be like and what things are supposed to mean. Syntactic problems concerning computer languages are mostly non-existent. Semantic problems do exist and are highly non-trivial. But looking at natural language, everything is harder. We have no direct evidence for the existence of any of the traditional syntactic categories (such as noun phrase, verb phrase, etc.) These are all theoretical constructs that differ from framework to framework. In a sense, when we look at sentences in a natural language, the structures we posit are our own; they are not self-evident, or found in a manual, or in any way obvious. With semantics, things are even harder. Psycholinguistic evidence is hard to come by, and even if we had more of it I'm not sure it would always be useful in semantic theory.

One sees many instances of concepts from theoretical computer science filtering back to linguistics. The main reason for this is that there is so much more activity in computer science than linguistics, and so much more sophisticated mathematics. I think it's fair to say that nearly many of the technical innovations in linguistics (especially syntax) in the past 25 years are borrowings from theoretical computer science. (The main exception here is that the main concepts of formal language theory were formulated first by Chomsky for linguistics (where they are today largely forgotten), and these quickly became a "classical" area of theoretical computer science.) This includes all of the interaction of applied logic and linguistics.

Logic in computational linguistics Computational linguistics is concerned with all matters related to the project of getting a computer to process language. This includes speech recognition, parsing, syntax, semantics, pragmatics. I think computational linguistics bears the same relation to mainstream linguistics that applied math bears to pure math. In fact, going back to the quote from Halmos on page 5, it seems that the description of applied math fits computational linguistics even better theoretical computer science, and for that matter even better than applied math itself!

It is no surprise to find logical aspects of computational linguistics: as we have seen, this is bound to happen with all serious work with computers that requires a theoretical approach. Here are some of the many ways logic enters in: various proposals for syntax use systems like *linear logic* (an outgrowth of proof theory) or *logic programming* (coming from Horn clause logic, and now the basis of the language Prolog). An area common to logic and linguistics is *categorial grammar*, based on work by Joachim Lambek from the 1950's and 60's. Yet another concerns precisely the latest work in the Chomskyan tradition, the minimalist program. Here, despite the fact that Chomsky didn't want to consider formalization, people have done exactly that. They've characterized the languages which are generable by "MG grammars" in terms of classical formal language theory. This interesting result came via a lot of other work, some of it involved formal language theory, and some of it involved proof-theoretic grammars related to Lambek's categorial grammars.

A potentially far-reaching application of logic in the study of the syntactic formalisms is the program of *model-theoretic syntax* initiated by James Rogers in the 1980's. Rogers applied seminal work in decidability coming from mathematical logic to the syntactic formalisms. (So we see again that applied logic builds on core areas of mathematical logic.) In more detail, the basic idea here is to translate syntactic formalisms which are used by linguists into a single language. That language happens to be a version of *monadic second-order logic*. It so happens that Michael Rabin had shown the decidability of the monadic second order logic on trees, and others had made the connection between definability in that logical language and notions coming from formal language theory. So Rogers' proposal applies one of the most traditional measures of complexity coming from logic, that of asking for a given notion and a given language whether the notion is definable in the language.

Semantics Semantics as a formal discipline owes a tremendous amount to logic. Richard Montague in the 1960's and 70's pioneered the adaptation of semantical methods from logic to fragments of natural language. This was the first, and therefore the most decisive development, in the field. Logic is at the heart of the fairly new discipline of *computational semantics*, too.

Beyond logic As it happens, even the best ideas from logic and classical linguistics are not in practice as accurate or as fast for natural language processing as some very simple heuristics coming from statistics. One gets even better results with more sophisticated models, and perhaps the best models to date come from Markov branching processes and random fields. These models work in the sense of being, say, 94.54% correct for their tasks. But they work on the basis of correlations rather than principled explanations. (So again we have Halmos' wide-angle and telescopic lenses.) Going further, it is quite an interesting matter to *mix* declarative (logical) and statistical methods in linguistics (and in vision and planning as well). This relates to my point earlier that applied logic will need to go beyond the traditional borders of logic. In this case, what is needed is a principled blend of logic and statistics.

5.4 But is it dead?

My proposal calls for the extension of logic to incorporate a number of aspects that traditionally are missing: uncertainty, social aspects, context, and so on. One reaction to this is that doing so would lead to the *end of logic*. I am reminded of Keith Devlin's book **Goodbye, Descartes: The End of Logic and the Search for a New Cosmology of Mind** (Johy Wiley & Sons, 1997). Devlin's book is a popular account of the history of logic and the related areas that I am dealing with in this essay. His book closes with the discussion pertinent to the subtitle, that what we are seeing is an end of logic – especially logic conceived of as requiring the duality, the divorce of mind and body. It also makes the case for "soft mathematics" and has proposals on what this means. Much of the points in the book are consonant with my message here. But my feeling is that everything that would lead one to declare a death of logic could just as easily be seen the opposite way: rather than dead, the subject is rather only in its infancy (strange as that sounds for one of the most classical subjects). One would like to think that in a hundred years, or a thousand, that trends in applied logic will give rise to a subject of permanent human importance, a revitalization of the classical subject of logic that takes into account the many things missing from the subject as we now know it.

6 Being as catholic as possible

Main Entry: catholic
Etymology: ... from Greek katholikos universal, general, from katholou in general,
... 2 : COMPREHENSIVE, UNIVERSAL; especially : broad in sympathies, tastes, or interests
Webster's Collegiate Dictionary

It should be clear that applied logic is multidisciplinary: since it is outward looking, it thrives on interactions with many other fields. There are institutional challenges for applied logic, as there are with any interdisciplinary endeavor. What I want to do here is to make a few comments on those challenges, and what can be done to help.

The Pigeonhole Principle is for the birds When mathematicians speak of the Pigeonhole Principle, they have in mind a fundamental fact: if you have more pigeons than pigeonholes, then when you put the birds in the holes, at least two are going to end up in the same place. This is not the principle that I object to. Instead, I feel hampered by the principle that *people* should be pigeonholed according to what they study, or by their academic departments. For applied logic as I'm thinking of it, departments are really not of great value. We should try to see beyond the boundaries of disciplines.

One of the success stories in the field has been the European Summer School in Logic, Language, and Information. This annual school is *the* showcase for many areas of work, including those that I'm calling applied logic. Its interdisciplinary spirit is inspiring, and it would seem to be important to emulate and strengthen it.

But even with modest successes, we have problematic points. The biggest paradox with an interdisciplinary field concerns exactly the phrase *interdisciplinary field*. One must value, on the one hand, contributions to outside areas, and on the other hand, solutions to problems that come internally. And the other challenge with interdisciplinary work concerns the matter of making connections between fields. In many cases, doing this is not appreciated, and at other times it even feels like a poor use of time and effort. My feeling here is that these hard "political" problems will be with us for some time.

Towards a new logic curriculum Another challenge comes in the matter of training people to do applied logic. Here I am more conservative than in other sections of this essay. I don't think there's any substitute for thorough grounding in more established fields. In fact, to be able to do serious work in applied logic one would probably have to have a good grounding in at least two fields. Where I think it makes sense to think about change is in the logic curriculum itself, specifically in the *second* (and further) courses in the subject. In the undergraduate curriculum, these second courses are frequently ones aimed at set theory or recursion theory. I see no reason why this has to be the case: after all, subjects can be presented with different emphases, and so why should we not present a serious, mathematically engaging treatment of logic that goes in the direction of applied logic? There are already texts on logic for computer science that do this. In other directions, one could easily imagine a second course in logic that emphasized both the challenging mathematics and the stimulating interest coming from modal logic. Another alternative would be to teach some of the areas of interaction of logic and linguistics, and in this way introduce model theory and proof theory. Surely doing this is not only The graduate logic curriculum, too, can be reworked. Here I think that all of the traditional areas of mathematical logic can be presented in ways that emphasize the applied side. A course in model theory, for example, could turn into the study of logical systems: students would then learn quite a bit more about completeness theorems for many logical systems (maybe even non-standard ones) than in the traditional classes. One can imagine other applied logic graduate classes, and surely those of us who are doing this should be talking to each other.

The price one would pay for a non-traditional curriculum is that students trained like this would not be trained to do research in traditional areas. But given that applied logic is likely to be interesting to a wide range of students, and given that applied logic stands ready to blossom into an important field, this price is not too steep.