Philosophy of Computing and Information: Five Questions

Johan van Benthem, Amsterdam & Stanford, http://staff.science.uva.nl/~johan/

September 2007

Question 1 Why were you initially drawn to computational and/or informational issues?

Frankly, the term 'information' had a softish ring in my student days. It was associated with Dutch philosophers writing shallow pamphlets about Everything with a soup of half-digested bits and pieces of Shannon's information theory. Nor did 'computation' have an immediate inspiring appeal. It suggested a drill in decision methods, or an auxiliary task of implementation like writing a computer program – while its cognate 'calculating' was even decidedly objectionable. But over time, I have come to love our Editor's two themes: partly through developments in my own field of logic, and also through the rise of the discipline of *computer science*. The latter term is still somewhat unfortunate, as it suggests a dance around machines, and an auxiliary crowd of mechanics greasing wheels and serving customers. But what I have in mind is the austere Latin term "Informatica", still used in The Netherlands for the field. That has the ring of the fundamental scientific study of information, it sounds like a classy relative of "logica", and from the start, it associates with computation, a link which the questions of this interview keep alive in English with a valiant host of slashes /. All this is why I also like the modern term *Informatics*, which suggests the right mixture of themes. Enough of terminology now, and on to ideas!

I became a logician at an early age – and our field seems information-laden from the start. We tell students that valid logical inferences 'unpack the information' in given data, and in modern dynamic logics, we show them how events of observation and communication 'update the information' of rational agents. Indeed, an embarrassment of riches threatens. There are many different notions of information in logic, ranging from more deductive to more semantic views: a diversity to which I will return below. But even so, the curious thing is this. Logic has official definitions for its central concepts of proof, computation, truth, or definability, but not of information! And somehow, many logicians feel this is significant. We do not need this popular notion in the mechanics or even the foundations of our formal systems – or, as Laplace said to Napoleon, who inquired into the absence of God in his *Mécanique Céleste*: "Sire, je n'avais pas besoin de cette hypothèse".

One important push taking information more seriously was that of Jon Barwise and John Perry around 1983, who created 'situation semantics' as a radical alternative to the ancien régime in philosophical and mathematical logic. On their view, triggered by developments in cognitive psychology and philosophical epistemology, logic should study the information available in rich distributed environments (with both physical and human components), and the resulting information flow. At the same time, with a band of allies, Barwise and Perry started the 'Center for the Study of Language and Information' at Stanford, which quickly became a hot-bed of lively interdisciplinary encounters between philosophers, computer scientists, linguists, and psychologists. *CSLI* still exists today. After all, academic paradigms are like religions: the faithful do not just want rousing sermons and moving ceremonies, but also imposing architecture. Many still worship at the Stanford temple, though few have become out-and-out situation theorists.

At the same time, in Europe the study of natural language semantics underwent an informational turn. Jeroen Groenendijk and Martin Stokhof introduced information of language users in defining meanings of key linguistic constructions, including dynamic speech acts like questions. With Peter van Emde Boas, a pioneer in the study of parallels between natural and programming languages, and Frank Veltman, who had developed an update semantics for conditional expressions, they redefined meaning as 'potential for information update' based on abstract computation in appropriate state spaces. Similar ideas were found in the influential discourse representation theory of Irene Heim and Hans Kamp. By 1986, all this had become so natural that we started an 'Institute for Language, Logic and Information' *ITLI* in Amsterdam, which is still in full swing today. Incidentally, terminological 'capitalism' mattered even then. In 1991, *ITLI* was renamed to *ILLC*, the *Institute for Logic, Language, and Computation*, as colleagues felt the *I* of 'information' was soft, while a *C* of 'computation' suggested depth and real labour. While we were at it, around 1990, with like-minded colleagues across Europe, we also set up the European Association for Logic, Language and Information *FoLLI*: no *C* there, though it does incorporate interfaces with computer science. *FoLLI's* annual ESSLLI Summer Schools have become a tradition traveling all over the continent. I hope that one day, just as the 'Olympic Games' transcended their Greek roots, they will travel all over the world.

Just to be sure, a serious interest in logical theories of information does not force a break with the tradition: one can also use classical tools. Around 1990, I became interested in uses of *modal logic*, my first and maybe still my truest love, as a general theory of process structure. This had to do with interests in process equivalence, expressive power, and computational complexity – but also: information! Modal logic seems well-suited as a calculus of information – and that at two levels, which reflect the tandem with computation in this Volume. First, possible worlds can represent information that agents have: witness the 'information stages' of intuitionistic logic or the 'information ranges' of epistemic logic. But second, dynamic processes of inference, observation, or communication continually *change* these static representations. And modal logic can describe those, too. Statics and dynamics come together in modern logics of what may be called *intelligent interaction* – and this is no coincidence. Logics of information should take the systematic Tandem View that information cannot be understood in isolation from the processes which convey and transform it. *No information without transformation*! And this is not just a slogan. The *Handbook of Modal Logic* (Elsevier, Amsterdam 2006) provides some powerful machinery, and illustrations of how this methodology works in many settings.

The Tandem View shows particularly well in modern *epistemic logic*, one major strand in logical studies of information. Epistemic logic was proposed by Hintikka in philosophy in the 1960s, and independently by Aumann in economics in the 1970s. Since the 1980s, when Joe Halpern and his colleagues at *IBM* San Jose started the *TARK* conferences on Reasoning about Knowledge and Rationality, the field has flowered at the interface of computer science, philosophy, and economics. *TARK* has been one more major influence putting information on the map in logic. But one must be down to earth. Modern epistemic logic is not an account of the philosopher's Holy Grail of 'True Knowledge', whose demands are so strict that no mortal can ever aspire to it – but rather one of the more mundane, but also much more useful, notion 'to the best of my current information'.

In the 1990s, a further notable new force was the rise of 'Informatics': a new academic conglomerate of disciplines sharing a natural, not funding-driven, interest in information and computation as themes cutting through old boundaries between humanities, social, and natural sciences. C.P. Snow deplored, but did not heal, the divide between the 'Two Cultures': Informatics is a seismic force which can redraw academic territories. By now, there are Informatics faculties in Bloomington and Edinburgh, to name a few. We still lack one in Amsterdam, though The Dream is still alive in many hearts and minds.

Thus, information and computation have been major forces in shaping my own intellectual development, my interactions with others, and even my organisational activities. Here you might wonder: why mention the latter in an interview like this? Well, I find it hard to separate individual research from interactions with colleagues and students. My current work on logics of games even intensifies an awareness of these social aspects and the intellectual power of interaction. But I even find it hard to separate this research interest from community-building activities. Call the latter the road of easy money and power if you like (before you have tried to run an institute), or the thankless life of public service (after you have). Either way, information and computation are powerful concepts that, to me and others, call irresistibly, not just for reflection, but also for broader academic action.

Question 2 What example(s) from your work (or the work of others) best illustrates the fruitful use of a computational and/or informational approach for foundational researches and/or applications?

I am not sure what 'foundational' means in this setting, but let me mention some examples that seem important to me. First, a focus on information and computation goes far beyond the immediate necessities of signal engineering or computer programming. And I am not primarily interested in the penetration of computing technology and *ICT* into our society. But as it happens, this technology also comes with a genuine flow of deep new ideas. Informatics offers new ways of conceptualizing scientific questions – and in doing so, it redraws boundaries in Academia in beneficial ways. We saw already how casting language users as information-processing agents has reshaped linguistics and philosophical epistemology. Here, computation starts as a pretty metaphor, but it then runs deep. Techniques for modeling and calculizing which have been developed for the narrower purposes of programming digital computers turn out to work just as well, when understood at an appropriate level of abstraction, for tasks such as grasping meanings, engaging in successful communication, or planning intelligent interaction.

Many colleagues have contributed decisively here, not just those mentioned so far. Peter Gärdenfors' pioneering work in the 1980s on belief revision showed how a central process in scientific methodology but also domestic human cognition, viz. mechanisms of *self-correction* on the basis of new information, involves precise structures of a computational nature. More broadly, 'logical AI' in the tradition of John McCarthy merged narrower issues in computer science with essential questions of understanding the world of *common sense* which have exercised people in the humanities and social sciences. In the same spirit, Dov Gabbay has taken informational–computational viewpoints across a wide range of topics, including argumentation, temporal reasoning, or the abductive formation of new hypotheses. But one can equally well cite Samson Abramsky's recent computational analysis of quantum mechanics using compositional modeling of programs by linear *logic games* in a category-theoretic setting. And these computational models are crossing over to cognitive science. E.g., around 2000, various people (including, in different ways, Reinhard Blutner, Michiel van Lambalgen, and Hannes Leitgeb) have shown that computation-based default reasoning is close to the working of *neural networks*, perhaps even that of the human brain – thereby laying to rest sterile polemics between logic-based and neuroscience based approaches to reasoning. My final example concerns a different part of Academia once more, viz. Rohit Parikh's program of *social software*. This is an ambitious attempt at applying computational-informational thinking to the analysis and design of actual social procedures, the glue that holds society together. This links up with game theory, social choice theory, and other parts of the social sciences, as well as with cognitive science. Thus, information and computation have a fundamental impact across our universities. In all the samples I mentioned, it seems fair to say that this stance transforms existing fields, giving them a richer set of tools, new friends, and even more importantly, a much richer agenda of significant questions to address.

My own work may also help illustrate how informatics emerges. At the 1987 Logic Colloquium in Granada, I presented a paper 'Semantic Parallels in Natural Language and Computation'. It shows how then new computational ideas like 'circumscription' make sense in linguistics, while ‘abstract data types’ revitalize old studies of empirical theories in the philosophy of science. And in line with the earlier Tandem View, it presents a modal logic of information stages together with informational processes of update, contraction, and revision over these. While this was still largely programmatic, my favourite vehicle for pursuing this concretely has become *epistemic logic*, suitably understood.

Let's first consider statics. Epistemic logic encodes a natural intuition of *information as range*, viz. all those worlds which one's current candidates for the actual state of affairs. Knowledge as 'to the best of my current information' then quantifies universally: an agent knows what is true in all her current candidates. To a logician, this at once raises an issue. What about the other natural, existentially quantified, idea of knowledge as ‘having some piece of *evidence* for a proposition’? In a lecture at *TARK* 1993, I proposed merging range and evidence views into one calculus. Exciting combinations have appeared since, such as Dov Gabbay's 'labeled deductive systems' and Sergei Artemov's ‘logic of proofs'. Thus, information has many natural aspects – and one wants to know if it is just a loose family of concepts, or whether it supports deeper links and combinations between these.

Next, my book *Exploring Logical Dynamics* (1996) tried to unify achievements by many people putting cognitive actions at centre stage as first-class citizens in logical theory. This 'Dynamic Turn' included belief revision theory, dynamic semantics, discourse representation, and other research lines in computer science, linguistics, and philosophy. The book showed that modal logic, with its nice balance between expressive power and computational complexity of languages, can unify theories of processes and information. In doing so, new themes arose. In particular, modal semantics becomes a probe to analyze existing logics into 'core' versus 'wrappings'. And then, interpreting first-order predicate logic over modal state spaces, a surprise occurred. One discovers a decidable core calculus of sequential procedures starting from assignments of objects to variables and testing of atomic facts. Standard Tarski semantics wraps this in a special structure theory of 'full assignment spaces' – which leads to the usual undecidability. What the modal models add here is a deeper understanding of the phenomenon of *dependence* between variables. This suggests an alternative view of *information as correlation*: but more on that below.

Around 2000, Willem Groeneveld, Jelle Gerbrandy and Hans van Ditmarsch finished their dissertations on information update in epistemic logic, while Alexandru Baltag visited ILLC as a post-doc. Since then, there has been a development of *dynamic epistemic logics* providing a super-structure to static logics of information by also describing information-carrying actions and events explicitly, with concrete procedures updating current epistemic models. Thus, the Tandem View lives inside one logical system. Recently, with students and colleagues, I have extended this approach to deal with belief revision by changing doxastic plausibility relations over models, and even to changes in agents' preferences.

Dynamic epistemic logic is a framework for modeling information in observed events, but also a calculus for such scenarios, with crucial ‘dynamic system equations’ telling us which information agents have after which events. A natural next step here also takes in agents' goals in acquiring and processing information, 'making sense' of the flow. With that, we get into *games* and *strategic interaction* over time. 'Logic and Games' is a lively information and computation-related current research area which is shaping up these days, but it is the topic of my interview in another *Five Questions* volume. One may even claim that information makes most sense in longer-term interaction – witness Keith Devlin's tongue-in-cheek definition of information as 'the ping-pong ball of communication'.

But there is much more to the relation of logic and information. With Pieter Adriaans, professor of learning systems in Amsterdam, I am editing a *Handbook of the Philosophy of Information*. By accident, it fell upon me to write the chapter on 'Logical Theories of Information', together with Maricarmen Martinez. Here is a summary of its conclusion. "Logic as theory of information is a legitimate perspective which puts many things in an attractive new light. One now pursues statics and dynamics, with intertwined accounts of information structure and dynamic processes that manipulate it. Thus, epistemic logic is an information theory of *range, knowledge*, and *observation-based update*. In doing so, we encounter the essential role of agents, and how they take information: often in interaction with others. But there is another basic aspect of information, its ‘aboutness’ and its links to the reality that we are interested in. Situation theory focuses on *correlation* and *dependence* between different parts of distributed systems. This is a complementary view of how agents access information in a structured world, and why it is there for them to pick up and communicate in the first place. The situation-theoretic perspective is not in conflict with the epistemic one, but rather a natural complement. But logic even offers a third major view, now more syntactic, of *information as code,* in the form of proof systems and other syntax-based calculi that drive inferential processes of *elucidation*. At this third stage, we link up with information processing as computation, and quantitative views of information. There are significant issues of compatibility and co-operation between all these different views , and we have merely indicated some possible merges, leaving the question if a Grand Unification is desirable, or even possible, to others."

Question 3 What is the proper role of computer science and/or information science in relation to other disciplines?

It seems hard to say what computer or information science is, as practitioners do not agree among themselves. In The Netherlands, the field has a history of ideological quarrels and lack of a united front to the outside world. Even base curricula change dramatically over time, as if there is no common treasury of consolidated insights. With this caveat out of the way, and sticking to my favourite term of ‘Informatics’, here is what I would say.

Information and computation are a new focus of research with a flavour like physics since the 17th century: scientific thought for professionals and technological innovation in society go hand in hand. While computers are changing our social world, computer scientists are changing academia. I will focus on the latter *cultural role*, which is a social one, too, provided fundamental insights make their way outside in an appropriate manner.

Starting from the 1930s, informatics has produced a steady stream of new notions and insights which affect the way we conceptualize problems across the university. The start was spectacular with Turing machines, which made computation a mathematical notion, with deep results such as the undecidability of natural questions like the Halting Problem. Over time, the study of computation generated further powerful ideas, which – de-coupled from their initial practical setting – turned out to have wide applicability. Examples are Automata Theory, Complexity Theory, Semantics of Programs, Type Theory and Linear Logic, Process Theory, Data Base Theory, Artificial Intelligence, and the list is growing still. These topics show a natural unity, in that they track the growing sophistication and sweep of a joint study of representation of data plus methods of computation over these. Just consider a few examples. The work by Dijkstra and Hoare in the 1960s on the effects of structured sequential programs run on single computers led to the *dynamic logic* of Salwicki, Pratt, and others, a general paradigm for describing actions and events which has found its way as far as linguistics and philosophy. Around 1980, distributed systems and parallel computation on many computers posed a new challenge. Here, a major innovation was the work of Milner and others on *process algebra,* which is now a general theory of communicating processes reaching out as far as physics and biology. And computation continues to inspire new fundamental theory. Just look at the *co-algebra* of Aczel and others, the study of never-ending computation over infinite data streams, which cannot be constructed, only observed. Co-algebra already has repercussions for mathematical proof methods in analysis and set theory. Finally, consider the realities of the Internet. This is best described as a mixed society of human and virtual agents engaging in interactions like those of ordinary life – cooperating at times, but also competing for scarce resources. Modern agent-based theories for internet computation meet with philosophical logics of knowledge, belief, and intentions, social choice and theories of organization, and economic *game theory*. In 2005, I wrote a survey paper for the conference *Computing in Europe* entitled 'Computation as Conversation'. It shows how this mixture of computational, philosophical, game-theoretic, and social themes generates interesting parallels across disciplines, and new concrete questions for a joint theory of rational action. Theorems about computation give insights on what can be achieved through conversation, while conversely, conversational models are a powerful metaphor for computational design. My point is that all this is already happening: we just need to see things for what they are.

In fact, the story is potentially endless. One could easily generate other lists with similar points, looking at knowledge representation, general artificial intelligence, constraint satisfaction, algorithmics, complexity, security, learning theory, diagrammatic reasoning, or image processing. In each case, an initial practical setting induces fundamental theory whose themes, properly understood, far transcend computer science in a narrower sense.

What is missing? Well, I often wonder why so little of the above is known to a general audience. Most academics view informatics as just an auxiliary discipline, the handmaiden of implementation – and some computer scientists even pander to this subservient stance. As a final indignity: 'information science' means courses on libraries and how to order your file systems. To change this, in a fashion term, what is lacking is a good *Narrative*, a story line which makes it clear what has really been happening in the sciences of computation and information, and what cultural impact this has had, and still should have.

Question 4 What do you consider the most neglected topics and /or contributions in late 20th century studies of computation and/or information?

So much is happening today that ‘neglect’ is not the first term that comes to mind! But the more exciting the landscape, the more new ridges to be climbed – so this question can always be answered. Here are a few themes on which I would like to see progress, if only for my own edification, going back to some issues touched upon in the preceding text.

One basic issue concerns the very notion of information. Is it really a coherent notion with consistent intuitions? Let us start with semantics. Current work on modal and epistemic logic concentrates on *information as range*. But there was also the situation-theoretic view of *information as correlation*. My Handbook chapter with Maricarmen Martinez tries to see the grand pattern behind both, and merges are possible. But we still lack a consensus on the appropriate level of generality here – though I suspect it might lie with the general logics of *dependence* studied, e.g., by van Lambalgen, Hodges, and Väänänen.

But logic also has a combinatorial syntactic perspective on information, as the structure 'unpacked' by inference. This alternative view of information-as-elucidation is closer to logical proof and code-based computation, which come with their own elaborate theory. Again, there is no necessary conflict here: our Handbook chapter gives joint models of update and elucidation in specific settings. But I am not aware of one widely accepted paradigm combing the inferential sense of information with the semantic one. Over the years, there have been interesting attempts in the work of Carnap, Dunn, Hintikka, Parikh, or Scott. But in my view, the price of abstract unification is still sometimes lack of exciting content. Even so, there is reason for optimism. Modern logic does have deep connections between proof and semantics. Just recall the celebrated result inaugurating the field, viz. Gödel's Completeness Theorem from 1929 showing that semantic validity and syntactic provability coincide for first-order predicate logic, the paradigm of modern logic.

Outside of logic, there is the more general interface between algorithmics and semantics, the two major prongs of the study of information and computation. There is interesting recent work of Abramsky, Adriaans, and others (cf. various chapters in the *Handbook of the Philosophy of Information*) on information flow in computation. Now, there need not be any unification here – maybe just a choice of 'life-style'. Occasionally, I get drawn into discussions (like with the Halting Problem, there is no decision method predicting if a conversation will turn out frustrating) with informatics colleagues telling me semantics has had its day, and algorithmics should take over. Perhaps, but I am not yet ready to give up.

Further desiderata! I already mentioned the lack of Grand Narrative explaining the cultural role of informatics. One cause may be the lack of an agenda of *Grand Challenges*. In a pamphlet written with my Amsterdam colleagues Arnold Smeulders and Martin Stokhof in 1999, we tried the following seven, beyond the logic-oriented story I have told:

1. Why is cognition so efficient? (E.g., what goes on as you are grasping this interview?)
2. How do observation and internal reasoning combine in successful planning and acting?
3. How to integrate content from different carriers: e.g., symbolic and graphic?
4. What is most efficient social organisation for intelligent persons and machines?
5. Which conservation laws of complexity govern information processing?
6. How to integrate the Two Traditions: algorithmics, code, physical signal transmission, and semantic information from logic, linguistics, philosophy, and social sciences?
7. What are the computational mechanisms driving major cognitive activities like learning, and can we use them for practical social impact through our educational system?

I would definitely not plead for this particular list, but rather for an agenda where issues like these are stated and discussed – much in the spirit of Luciano Floridi's work.

Finally, here is another desideratum. I said that in principle, informatics has the same 'success formula' as modern physics, where fundamental research happens simultaneously with technological innovation. In the short term, this is also a problem, as the fast delights of engineering, applications that matter, and recognition may overwhelm the slow delights of insight in the light of eternity. Nevertheless, there is no necessary conflict – and there is no blame to be assigned. To the contrary, what we really have is an exciting Triangle with three vertices. There is empirical phenomena around us, fundamental theory behind this, but also a third activist stance, viz. the *design of new practices* inspired by theoretical ideas concerning information and computation. We already live in mixed societies populated partly by our biological off-spring and partly by virtual citizens put there by the information and computer sciences. These mixed societies raise fundamental issues again, such as understanding successful interactions between agents of very different abilities. You may want to recall, if you know your informatics culture, that the original *Turing Test* in AI was already of this kind. It did not ask if a computer can fully emulate a human being, but rather, whether in a group consisting of a computer and a human being, the latter could detect the type of the former. The third activist stance is itself worth pursuing – and an earlier-mentioned program like Parikh's 'social software' shows how surprising that can be. Informatics can change the world, and the distance between theory and practice is so small sometimes that this itself calls for reflection, if not participation.

Question 5 What are the most important open problems concerning computation and/or information and what are the prospects of progress?

I answered this question under the preceding one. Let me just throw in two more things.

One concerns the *Handbook of the Philosophy of Information*, mentioned several times already. That book, to appear in 2008, brings together all major approaches to information in academia, and it draws basic philosophical lines – aiming for peaceful co-existence rather than grand unification. The authors represent both semantic and algorithmic traditions – with the latter all the way to quantitative Shannon information theory, Kolmogorov complexity, and probability theory. And here I will admit to a private perplexity when viewing all these chapters. To borrow a phrase, I would like to understand the 'unreasonable effectiveness' of quantitative information theories. To me, the idea that one can measure information flow one-dimensionally in terms of one number of bits, or some other measure, seems patently absurd. But in reality, it is spectacularly successful, often much more so than anything produced in my world of logic and semantics. Why?

Finally, further merges may be on the horizon. Is information and computation truly a natural frontier in Academia? Many topics on its agenda blend into empirical issues in *cognitive science*. Editing a recent 2007 issue of "*Topoi*" on 'Logic and Psychology' with Helen and Wilfrid Hodges, we found, often to our own surprise, how information, reasoning, computation, neural networks, learning, vision, and brain function form one natural conglomerate, where biological function meets computational design. In particular, 'competing paradigms' from the last century are now meeting in surprising ways. Who would have thought that neural nets compute much as default logics and event calculus in computer science, or that game theory would team up with linguistics and experimental psychology? Taking biological and psychological facts seriously is not uncontroversial in logical circles, but 'Information, Computation, and Cognition' may be the way to go.