

Technologies for harnessing knowledge: a European research agenda

Hans-Georg Stork*

European Commission, Directorate-General Information Society

Introduction and overview

The European Commission supports efforts to create advanced networking facilities such as Grids, for highly demanding and challenging applications, laying and strengthening the “pipes, tubes and hoses” that make communication happen, within and between research and education communities, and certainly beyond.

This note addresses a complementary topic. It is about the substance conveyed by these “pipes, tubes and hoses”, a substance that fuels not only scientific and scholarly discourse but also market processes and many activities of potential benefit to society at large. We call this substance “digital content”.

It comes in countless varieties and it would be a futile exercise to try to cover the entire spectrum in a short note. Therefore, we confine ourselves to one of the presumably noblest forms of "digital content", to "knowledge content" that is, or content that represents specifically human knowledge about aspects or phenomena of what we call the "real world".

This note is about the contributions that are made possible through European funding, not only to advancing the “state-of-the-art” of technologies that are specific to this kind of digital content, but also to the actual use of these technologies in application systems that render or support services, meeting needs that arise in various sectors of our societies.

These contributions commonly take the form of projects, involving partners from all over Europe. Projects are proposed in compliance with work programmes that underly more or less regularly published Calls for Proposals. The work programmes in turn, reflect - to a certain extent - prevailing policies pertaining to research, markets and society at large.

This “hierarchy” sets the agenda of the present note: We remind ourselves of the political framework insofar as it bears upon the topic at issue. We clarify what we actually mean when we talk about "knowledge" in the context of technology. Thirdly, we report on Commission supported activities furthering "Technologies for harnessing Knowledge" under the past 5th Framework Programme. Lastly, we give an overview of what is and will be in the offing under the recently adopted 6th Framework Programme.

The political backdrop

The most salient political reference is the one to the "Lisbon Summit", the extra-ordinary European Council of Heads of State and Government, who came together in Lisbon, in March 2000. The Council agreed to set the ambitious goal for the European Union to become - by the year 2010 - *"the most competitive and dynamic knowledge-based economy in the world"*.

* UCISA Conference 2003; the views expressed in this note are those of the author and do not necessarily engage the European Commission.

For reasons that should be obvious we forgo discussing the likely rationale underlying this goal. It would take us too far afield - into the realms of global politics and economics, presumably. Rather, we focus on one of the more practical and operational aspects.

Although stated simply, reaching said goal is a rather complex undertaking. First and foremost, it requires creating the conditions that would allow to make full use of the advances in information and communication technology achieved in the last decade of the previous century.

The most dramatic advance has no doubt been made in global networking and it has found its most visible expression in the development of the Internet and the World Wide Web. Its sheer scale has sparked a flurry of new research and revived somewhat dormant research areas.

It also facilitates new types of relationships among enterprises and between enterprises and their customers. It encourages new approaches to dealing with a host of societal issues, from access to public sector information to health care, education, training and independent learning.

The action plan, known as “e-Europe 2005”, that has been designed by the European Commission to meet the expectations set by the Council, addresses precisely these four areas: e-business, e-government, e-health and e-learning. It foresees policy measures as well as best practice and demonstration projects. Benchmarks have been defined and mechanisms have been put in place for national and European policy makers to exchange pertinent information.

The IST programme, now one of the priority areas of the 6th Framework Programme, is seen as a key component of that plan. In fact, the Council Decision on FP6-IST refers explicitly to the Lisbon Council and the e-Europe Action Plan: *“The IST thematic priority will contribute directly to realising European policies for the knowledge society as agreed at the Lisbon Council of 2000, the Stockholm Council of 2001, and reflected in the e-Europe Action Plan. It will ensure European leadership in the generic and applied technologies at the heart of the knowledge economy.”*

Technology for harnessing knowledge

The new IST programme translates the message of the “Lisbon Summit” into an agenda for research and technology development. Indeed, if by the end of this decade Europe is to become the most competitive and dynamic knowledge-based economy then harnessing knowledge for improving productivity and the quality of life is an imperative sine qua non.

In a nutshell, this translation can be phrased as follows: if we want to improve the ways we apply knowledge to the development of all kinds of technology, to our organisations and to our institutions, then we must also learn to apply technology to knowledge more effectively and efficiently.

IST under FP6 acknowledges this imperative by giving technologies dealing with knowledge a prominent place in its work programmes.

This begs the question: “How can technology be applied to knowledge?” Of course, we all know more or less how to apply knowledge to creating and using technology; we do it all the time. But the other way round?

Any reasonable answer to that question presupposes at least some understanding of the term “knowledge”. Unfortunately, that term denotes a concept philosophers have ruminated

for thousands of years, producing tomes of learned treatises. Yet it has remained elusive a notion at best.

Fortunately, we can have recourse to the opinion of experts, researchers and developers, who were invited to a workshop in Luxembourg, in Spring, last year, to deliberate the ins and outs of what computer technology can do for "knowledge".

That workshop was designed to illuminate the "Knowledge Technologies" part of the then proposed text for a Council Decision on the specific programmes under FP6. It described the objectives of "Knowledge Technologies" as "... *to provide automated solutions for creating and organising virtual knowledge spaces (e.g. collective memories) so as to stimulate ... new content and media services and applications.*"

So, until further notice, we quote - more or less verbatim - from the minutes of that workshop:

Knowledge is a prerequisite for acting purposefully in a given environment or domain. This includes decision making, planning, collaborating or simply finding things. Knowledge is always about something: objects, processes, phenomena, etc..

To be amenable to "*automated* (i.e. computer-implementable) *solutions*" knowledge must be formally represented. Computer-based representations of knowledge about objects and processes (digital or not) capture to a certain extent, the *semantics* of these objects and processes.

While computer-based knowledge representation is a form of digital content it is often also about digital content which itself may or may not represent knowledge.

The remit of *Knowledge Technologies* therefore includes "*meta-content*" (about digital content of all sorts) that makes all other forms of knowledge (and occasionally nonsense) more accessible and usable.

Adding, automatically or interactively, explicit semantics to content, services and processes - and thereby producing knowledge representations - is a key function of knowledge technology tools.

A second important class of functions of Knowledge Technology tools and artefacts can be characterised broadly as "acting upon semantically-enriched content (including service descriptions)".

Let us take a slightly closer look at these functions:

(i) Adding explicit semantics to content, services and processes

Knowledge about content, services and processes is made explicit through formal *descriptions* of such entities. These descriptions are usually referred to as *metadata*. However, while formal description is necessary for agents to act upon, it is not sufficient. Descriptive terms are "understandable" (or "meaningful") only if their meaning has been defined somehow, somewhere. This is usually done through *ontologies* which provide meaning that can be operated upon. They embed terms in contexts (of other terms) and/or stipulate rules a given term (or set of terms) must obey.

Metadata and *ontologies* demarcate broad and inter-related research areas. Pertinent problems and solutions depend largely on content types and usage environments. But there are generic problem classes, for example:

- ontology construction (including ontologies for multimedia objects) and management ("knowledge lifecycle" support)

- metadata extraction / capture
- semantic annotation
- domain-, context-, user- and task-oriented indexing
- semantic indexing of multimedia content
- ontology learning

And there are numerous classes of technologies and/or approaches likely to provide generic solutions, for instance:

- data/text mining for "knowledge discovery" (such as concept detection and fact extraction) in data bases or large text repositories
- machine learning (e.g. for automatic classification)
- semantic analysis of audiovisual content (segmentation, object extraction, etc.)
- speech, face, gesture and emotion recognition (e.g. through natural language analysis and cognitive vision)

Specialised services could be offered to support both knowledge acquisition and knowledge elicitation (e.g. for semantic annotation of content). *Peer-to-peer* networks for instance, would lend themselves to implementing methods allowing semantics to "emerge" from node-to-node interaction.

Given the sheer amount of content in global distributed systems, solutions to some of the above problems (e.g. semantic annotation, multimedia content analysis, etc.) may require powerful computing resources as provided for instance by *Grid* computing technologies.

(ii) Acting upon semantic descriptions

Semantic content (service, process, ...) description (based on suitable ontologies) enables *software agents* in distributed systems to co-operate and to perform complex transactions and other operations (such as searching, filtering and integrating information), on their users' behalf and without extensive user intervention.

Semantic descriptions relieve implementers of the burden of "hard-coding" semantics in agents and thus contribute to achieving interoperability. They can also help *human agents* ("users") to make sense of and interact with content, services and processes in distributed systems.

Services provided by software are of particular relevance in distributed systems such as the World Wide Web.

Indeed, the term "*Semantic Web*" usually refers to the formal framework (in terms of models and languages) needed to provide software agents with ontology-based (i.e. semantic) descriptions of all sorts of web-addressable entities, allowing also context-sensitive *service discovery*, *mediation* and *composition*.

Grids will become powerful service providers, accessible through Semantic Web interfaces.

The "interfacing with knowledge" aspects (i.e. making content and services accessible, intelligible and actionable to people) are equally intriguing. Subjects falling into that category and allowing for largely generic R&D include:

- semantics based navigation and browsing

- semantic search engines with domain-, context-, user- and task-sensitive query construction support
- knowledge-(viz. semantics-)based dialogue management
- user profiling, personalisation, customization (e.g. through particular "views on knowledge")
- “visualizing knowledge”
- device dependent interfacing.

The technologies elucidated so far bear on many broad application domains, for example:

- business and public administrations,
- science, engineering, medicine, law and (many) other professions,
- (general and specific) information services,
- education and training,
- “memory institutions / facilities” (harbouring intellectual and cultural assets, “digital libraries”),
- entertainment (games, interactive digital TV, etc.),

providing the underpinning for

- workflow and collaboration support,
- transaction oriented systems (e.g. for e-commerce),
- proactive portals for community building,
- retrieval, filtering, profiling and recommender systems, as well as
- document change and innovation management.

Lastly, it goes without saying that *Knowledge Technologies* draw on various Computer Science sub-disciplines such as formal modelling, logics and languages, information retrieval, (multimedia) databases, image analysis, cognitive vision, etc., but also on "trans-disciplines" such as Cognitive Science. Therefore R&D addressing *Knowledge Technologies* is necessarily multi-disciplinary.

“Knowledge technologies” in the IST part of FP5

The workshop was meant to clarify issues with a view to preparing work programmes under the current (sixth) Framework Programme. It goes without saying, however, that the research topics, technologies and application needs it referred to are by no means new.

On the contrary: *natura non facit saltus* - nature does not jump, unless the quanta are very, very small - and neither do research and technology development. Key challenges - such as ‘real-world’ analysis or access to and use of digital content - tend to persist and solutions tend to evolve in a never ending “*challenge - response*” cycle.

It is therefore no surprise that many of the issues addressed have been picked up - with or without the "knowledge" or "semantics" qualifiers - under previous EC R&D programmes and of course elsewhere.

Let us then take a look at some of the contributions made possible through the IST part of the previous, the fifth, Framework Programme.

As one may recall, it was partitioned into four “key actions”:

- Systems & Services for the Citizen (KA1)
- New Methods of Work & Electronic Commerce (KA2)
- Multimedia, Content and Tools (KA3)
- Essential Information Society Technologies and Infrastructures (KA4)

In addition, there was ample room for

- Future and Emerging Technologies (FET), as well as so called
- Cross-Programme Actions (CPA)

The Key Actions as well as the FET domain and the CPAs offered so-called Action Lines to guide potential proposers. Several of these Action Lines addressed topics that can with full justification be subsumed under “Knowledge Technologies and applications”. For instance:

KA2

- Corporate knowledge management
- Knowledge Management for eCommerce and eWork
- Technology Building Blocks for Trust and Security

KA3

- Authoring and design systems
- Content management and personalisation
- Media representation and access: new models and standards
- Access to digital collections of cultural and scientific content
- Content-processing for domestic and mobile multimedia platforms
- Information visualisation
- Semantic Web Technologies

KA4

- Engineering of intelligent services
- Methods and tools for intelligence and knowledge sharing
- Information management methods

FET

- Open domain (FET OPEN)
- Universal information ecosystems
- The disappearing computer
- Global computing: co-operation of autonomous and mobile entities in dynamic environments

CPAs

- CPA9: GRID Technologies and their applications

... not to mention KA1 in whose Action Line descriptions the term “intelligence” abounds. Presumably, “intelligence” is intimately related to “knowledge”.

An assessment of all projects funded under the IST Programme 1998-2002 identified well over 750 projects emerging from all IST Calls, that address in one way or other the applications and some of the technologies at issue. They deal with knowledge and information management, agent technologies, optimisation tools and decision support systems, supply chain management, generic organisational tools and many other relevant topics¹.

It is not possible to go into great detail here. Rather, we should take a closer look at one specific area labeled '*Information access, filtering, analysis and handling*' (IAF for short) and one of the modules of Key Action III (*Multimedia Content and Tools*).

Its objectives, according to the early FP5 documents, were to support the development of “...*advanced technologies for the management of information content to empower the user to select, receive and manipulate ... only the information required when faced with an ever increasing range of heterogeneous sources.*”

And these technologies should lead to “... *improvements in the key functionalities of large-scale multimedia asset management systems (including the evolution of the World Wide Web) to support the cost-effective delivery of information services and their usage.*”²

Picking up on the “*evolution of the World Wide Web*” clue, a specific Action Line of the IST Work Programme 2001 was dedicated to '*Semantic Web Technologies*' (Action Line III.4.1), thereby underlining the importance (in terms of research challenges and expected impact) of 'semantics issues' for achieving the declared goals of the IAF module of IST.

While triggered by the now well-known “Semantic Web” vision developed at the W3C, under Tim Berners-Lee, its scope was considerably broader than the original formal and informal Semantic Web notes, issued by the W3C, might have insinuated. It encompassed four interrelated R&D tracks as an orientation for submitting project proposals. These tracks can be summarised briefly as follows:

- creating a usable formal framework in terms of formal methods, models, languages and corresponding tools for semantically sound machine-processable resource description
- fleshing out the formal skeletons by developing and applying techniques for knowledge discovery (in databases and text repositories), Ontology learning, multimedia content analysis, content-based indexing, ...
- acting in a semantically rich environment, performing resource and service discovery, complex transactions, semantic search and retrieval, filtering and profiling, supporting collaborative filtering and knowledge sharing, ...
- making it understandable to people through device-dependent information visualisation, semantics-based and context-sensitive navigation and browsing, semantics-based dialogue management, ...

Its scope was hence wide enough to include problems such as the automatic or semi-automatic creation of semantic annotation of all forms of content and resources (thus

¹ <http://www.cordis.lu/ist/cpt/ipa.htm>

² Quoted from <http://www.cordis.lu/ist/b-oj-en5.htm#KA3>

creating a link to multimedia resource description), or for instance, ontology learning in peer-to-peer systems.

It also provided some continuity with respect to previous Key Action III activities (notably on 'media representation and access' and digital libraries) and the already mentioned activities supported by other IST departments.

Last but not least it provided a sharper focus on the problems of creating and using knowledge representations, in the context of large-scale distributed systems, such as the World Wide Web.

Focus and scope were largely retained in Work Programme 2002 as part of Key Action III's 'Preparing for future research activities' Action Line (AL III.5.2). Moreover, Work Programme 2002, in one of its 'Cross Programme Activities (CPA)', took account of the new trend that has surfaced over the last couple of years: the application of Grid technologies to “*knowledge discovery in ... large distributed datasets, using cognitive techniques, data mining, machine learning, Ontology engineering, information visualisation, intelligent agents...*”³, all more or less directly pertinent to the Semantic Web vision.

Calls for submission of proposals to these Action Lines, published in July (AL III.4.1) and November (AL III.5.2 & CPA9) 2001, respectively (Calls 7 and 8), drew altogether nearly one hundred submissions involving several hundred participating organisations. They resulted in a significant growth, by 17 projects, of a portfolio of projects that are all poised to contribute in one way or other, to making the "Semantic Web" happen, with as general an interpretation of that concept as possible.

Said portfolio had been “initialised” by *On-To-Knowledge* and *Ibrow*, probably the first Semantic Web projects ever to receive public funding in Europe - if not in the world.

*On-To-Knowledge*⁴, funded under Key Action 4, has become one of the birthing grounds of OWL, the proposed new Web Ontology language, currently under discussion at the W3C. *Ibrow* (*An Intelligent Brokering Service for Knowledge-Component Reuse on the World Wide Web*)⁵ already started in 1997, under FP4, when the terms "Semantic Web" and "Web Services" had not yet been coined or widely used.

Just finished, it already has a worthy successor: *SWWS*, “*Semantic Web enabled Web Services*”, supported under the “Semantic Web Technologies” Action Line. *SWWS* develops means for “*describing, recognising, configuring, combining, comparing and negotiating Web services, supporting Web service discovery and scalable mediation*”.

In recognition of the central role ontologies are likely to play in building the 'Semantic Web', the European Commission, through its IST Programme, also supports the 'Thematic Network' *OntoWeb*⁶, a platform for fostering collaboration between industry and academia, on creating a 'semantic infrastructure' for applications in many different areas (e-business, Web services, multimedia asset management, community webs, etc.). Through *OntoWeb*, European researchers and practitioners also have an opportunity to make more targeted contributions to international standardisation activities and to the W3C process.

It is an interesting exercise to analyse the “Semantic Web Technologies” portfolio with a view to categorising projects roughly along (at least) four dimensions: (i) generic problem

³ Quoted from WP2002, CPA9

⁴ <http://www.ontoknowledge.org/>

⁵ <http://www.ibrow.org/>

⁶ <http://www.ontoweb.org/>

class (such as the ones identified earlier: 'making semantics explicit' and 'acting upon explicit semantics'), (ii) technical solutions (e.g. automatic versus semi-automatic and interactive tools), (iii) type of content (e.g. text, corporate databases, multimedia objects, web pages, man-machine interaction records, sensor generated data streams, etc.) and (iv) application domain. This has been done elsewhere⁷ and we forego presenting this analysis here.

Suffice to say that applications are of course a must in the IST programme. Technologies must not be developed for the sake of developing technologies. Projects should not benefit a limited constituency only, or solve just one isolated problem. Rather, projects submitted under a generic action line should, in a final analysis, yield more widely applicable results, to be demonstrated through several showcases.

And indeed, the applications targeted by FP5 “Semantic Web / Knowledge Technology” projects are in the areas that have been broadly delineated above. They range from 'hard science' and mathematics via engineering, education, training and infotainment, to health services, enterprise application integration and eCommerce. Each of these application areas would certainly deserve a presentation of its own.

Apart from the more 'technical' dimensions there are of course the political, social and economic ones. But these we have already, albeit very perfunctorily, addressed.

“Strategic priorities” in the IST part of FP6, relevant to “Knowledge Technologies”

“Ambient intelligence” has been one of the guiding mottos of the IST programme under FP5, coined by members of ISTAG, the IST advisory group. A similar motto may or may not have been coined to capture the gist of IST research under FP6. However, at least as far as “Technologies for harnessing knowledge” are concerned one might propose something along the lines of “ubiquitous knowledge” or “knowledge about everything, semantics for everything”. This would indeed be required if we were to make full use of the global wired and wireless networks that keep increasing their reach both in terms of capacity and physical access modes, evolving into an *“ubiquitous permeable web”*.

As insinuated before, the sheer scale of this impending “Evernet”, as some prefer to call it, gives “Knowledge technologies” an unprecedented boost. Of course, this is not the whole story, but it may well be one of its more important chapters. Other basic digital and non-digital technologies, especially those that tend to overwhelm us with never ending streams of data, from earth observation or physics experiments, for instance, are also likely to give “Knowledge technologies” more prominence than they have ever had in the past.

The IST programme under FP6 takes account of the importance of “Knowledge Technologies” by declaring them part and parcel of one of its prime pillars. But it must also be noted that the concept of “machine-processable knowledge”, as explained, does indeed pervade to a greater or lesser extent, all four groups of “research priorities” of FP6-IST:

- Applied IST research addressing major societal and economic challenges
- Communication, and computing infrastructures and software technologies
- Components and microsystems

⁷ Stork, Hans-Georg: Webs, Grids and Knowledge Spaces - Programmes, Projects and Prospects, Journal of Universal Computer Science, vol. 8, no. 9 (2002), 848-868. (preprint accessible from <http://www.semanticgrid.org/documents/>)

- Knowledge and interface technologies

The popular "Future and emerging technologies (FET)" sector has been retained, by the way. In both FP4 and FP5, that sector also hosted a number of projects that can be safely allocated to "Knowledge technologies" and it probably will continue to do so.

FP6-IST work programmes translate the overall "research priorities" into a number of "strategic objectives" which underly Calls for Proposals. The work programme for the period 2003-2004 lists 22 of them, 12 for the first Call and 10 for the second. They come in three categories that are almost orthogonal to the "research priority" classification:

- technology components;
- integrated systems and
- sectorial applications.

Probably none of these "strategic objectives" can be reached if the notion of "machine-processable knowledge" is ignored. Yet there are certainly some where that notion plays a dominant role.

We take a brief look at four and a half of them.

There is, **first** of all, the "strategic objective" to build "**Semantic-based knowledge systems**". It is, naturally, part of the "Knowledge and interface technologies" priority and included in the current first Call for Proposals, launched in December 2002. It probably addresses - in a manner of speaking - the core of "Knowledge technologies". Such systems would typically integrate and automate the functions needed to acquire, organise, use and share the knowledge embedded in all forms of content.

Research guided by this objective would focus on the generic technologies - such as data mining - and issues - such as interoperability - underlying such systems while projects would of course have to showcase appropriate applications. However, no specific application domain has been identified for this "strategic objective".

The **second** example is "**Cognitive Systems**". It will be included in the next IST Call to be published later this year. Its very title also indicates its nearness to "Knowledge technologies". Here the principal focus is on physically instantiated and embodied systems, capable of gaining and representing the knowledge needed to act sensibly in "real-world" environments.

Attaining this objective requires building on a formidable body of research that has been going on for many years. It has in fact been one of the early dreams of the computer age. Today, given the current state-of-the-art of basic digital technologies, we may indeed get much closer to realising this dream more satisfactorily than ever before. It has been, is and will be a multidisciplinary endeavour, involving experts in computer vision, natural language understanding, robotics, artificial intelligence, mathematics and cognitive neuroscience.

Example number **three** is "**GRID-based systems for solving complex problems**". The concept is well known by now. In the UK in particular, it has gained some popularity very much thanks also to the eScience programme.

We have already mentioned Grids in the context of one of the last "Cross Programme Actions" of FP5. In FP6, Grid related research will aim to expand further the concept from "computational" to "knowledge Grid".

Indeed, a major role ascribed to processes running within the "knowledge layer" of a Grid is to assist in making sense of the huge amounts of data generated by, say, scientific instruments such as particle accelerators, gene sequencers, telescopes, satellites and a gamut of sensors. Hence, while operating Grids may require much of what we identified as "Knowledge technologies" - for instance semantically grounded resource and service description, network management agents and the like - they will also provide a host of services in "problem solving and knowledge creation environments", supporting and enhancing collaborative work in virtual communities and organisations.

The **fourth** example, "**Networked businesses and governments**", is a direct and explicit response to the e-Europe challenges of e-business and e-government. Both provide prime motivation for developing "Knowledge technologies" in the first place. Successful enterprise application integration and e-commerce (B2B as well as B2C) hinge on the semantic interoperability of enterprise application systems and e-commerce platforms. The same holds, *mutatis mutandis*, for services rendered by public administrations accessible online and for organisational networking in general.

Ontologies, for product and service description for instance, are instrumental in making this happen. And the software entities known as agents enter this stage not only as users of the 'semantic infrastructure' but also as tools for its creation and maintenance.

Example **fourPOINTfive** is the second half of "strategic objective" "**Technology-enhanced learning and access to cultural heritage**": So, it is "**Access to cultural heritage**". Well, it is certainly not only the heritage from the past that would be of interest here but also what we leave to future generations. And it certainly includes the scientific and scholarly assets of the past and the present - for the future. This is the remit of what we usually call the "memory institutions": libraries, archives, museums, public or private.

And it is more: it is a formidable testing ground for most of the "Knowledge technologies" discussed in this note. As a matter of fact, this domain has also been one of the birthing grounds of these technologies. It is probably not too far-fetched to declare the thesauri, controlled vocabularies and authority files - that have been used in these institutions for ages - worthy predecessors of modern metadata schemata and ontologies.

Solid research work still seems to be required to improve the accessibility and usability of cultural and scientific resources. This holds for both of the basic classes of "Knowledge technologies": we need, for example, highly automated digitisation processes that capture simultaneously and to the largest extent possible the "semantics" of the objects in question; and we also need interoperable *digital library services*, providing high-bandwidth access to distributed and highly interactive repositories of culture, history and science.

But why have we skipped the "Technology-enhanced learning" part of this "strategic objective"?

At this juncture, we should come back to the fundamental question raised earlier: "What can technology do for knowledge?" The answers given were all - more or less - about "machines", as if the question had been: "How can we put our (or whatever) knowledge into machines?"

Yet there is something that should be considered much nobler than "machine-knowledge". The most advanced and most challenging "cognitive system" is after all ... Homo Sapiens. That system is based on knowledge represented somehow - we still do not know exactly how - in the bio-wetware of our brains. And that knowledge is the ultimate source of all innovation, of scientific and technical advance and in a final analysis, of machine

knowledge as well. It is gained through learning. And unlike machines, we all do that differently.

So, human knowledge and human learning are very special and they should be treated as such. The fundamental question here is "How does knowledge enter the human brain?" But we may of course also ask: "How can technology help people acquire knowledge and use it to their and their kin's advantage?". Or simply, to use the terminology of our "strategic objective": "How can technology enhance learning?"

That question has been on the agenda for a very long time. Many answers have been given, not always explicit or based on thorough research but often through products that were simply put at the teachers' and students' disposal, challenging their ability to organise themselves. Like the village common where people tread out paths as they go about their errands.

And clearly, modern computer and communication technology that deals with digital content representing knowledge and providing information, the very substance that also fuels the learning process, lends itself to many answers.

It stimulates our imagination and lets us find new ways of making - figuratively speaking - more intelligent slates, blackboards, textbooks, science cabinets, workbenches, classrooms and lecture theatres, catering to the needs of the individual learner or specific communities of learners.

This is quite different from, say, putting networked computers with Internet connections in a classroom simply to make our kids fit into the working environments they are likely to encounter when they grow up. One is "technology for learning", the other is "learning for technology".

In the past, European Programmes, such as DELTA, the early distance learning initiative, the Telematics for Education and Training sector under the fourth Framework Programme, and the Education and Training area under the fifth, have contributed greatly to "technology for learning".

However, applying technology to learning with a view to improving the learning process requires a thorough understanding of that process, and a thorough understanding of the "cognitive systems" that people embody, as individuals and in organisations. Just providing the "gadgets" may not be enough.

What does "improving the learning process" mean? How do we measure the improvement? This may be a somewhat naïve question likely to allow many detailed answers. But whenever new ideas are actually put forward on how to improve the learning process through computer and communication technology there should be a compelling case behind them.

So, last but by no means least, the second half of the fifth "strategic IST objective" is also about knowledge. Not as much about "harnessing knowledge" as about "growing knowledge". It is of course about technology as well but also and equally about pedagogy and organisational issues. And it is about making that case.

The aim is to apply what is currently at the leading edge in for instance, networking, multimedia, virtual and augmented reality, virtual presence and simulation, to creating new effective learning environments for people, as individuals and as members of groups or organisations. The formally representable "machine knowledge" we were dealing with up until "strategic objective" fourPOINTfive will most certainly also play a crucial role in these "technology-enhanced learning environments".

Afterword

"Knowledge technologies" in the context of science, business, public and private organisations, culture, learning and many other areas, clearly demonstrate yet again the close interdependence in general between technology and technology application.

Both are inseparably linked in an infinite loop similar to the "challenge - response" cycle mentioned before: applications pose challenges; technologies are developed in response to such challenges and may make new applications possible, desirable or necessary.

Applications in turn are driven by societal needs while society - including the markets - expects viable products and services in return. However, the crucial element in this picture is "research". It is at the heart of the matter, our engine of curiosity, creating the knowledge of what exists, what can be done, how things can be put together, and so on.

Learning lubricates and fuels that engine. And in spite of hope or hype to the contrary: the engines composed of "Knowledge technology" parts will not step in or take over if our own engine of curiosity fails. All they can and should do for us is to support our engine in getting results faster and perhaps more reliably.

And for all we know they may not be able either to realise the dream of Gottfried Wilhelm Leibniz, polymath, Sir Isaac's contemporary and great European who, in 1678, proposed the "intelligent agent" approach to solving philosophical problems:

"When controversies arise, no more a dispute will be necessary among two philosophers than among two calculators. For it will be enough to take pencils and abacuses in hands, and say to each other: let us compute!"

But perhaps the man-made knowledge engines will bring us closer to realising another vision, enunciated by H.G. Wells in his essay "World Brain", well before Vannevar Bush invented the Memex machine:

"A great new world is struggling into existence. But its struggle remains catastrophic until it can produce an adequate knowledge organisation. ... An immense, an ever-increasing wealth of knowledge is scattered about the world today, a wealth of knowledge and suggestion that - systematically ordered and generally disseminated - would probably ... suffice to solve all the mighty difficulties of our age."

That was in 1938. At the time Wells seemed to be fairly pessimistic about the situation one would have to start from. He continued:

"But the knowledge is still dispersed, unorganised, impotent in the face of adventurous violence and mass excitement."

Let's hope we can do a better job today. In Europe and the whole world. More than "Knowledge technologies" is needed. Above all, we must learn.