This document summarizes essential points regarding my scientific research contribution to date and its potential for the immediate future. It consists of:

- a short statement on my **perspective** based on what **background**;
- what originally triggered my **research motivation** in order to give a better appreciation of "where this all comes from;"
- a succint overview of what has been my **lifelong pursuit** starting right after my PhD up until today;
- a digest of my **current potential** in terms of research and development to build software tools and applications exploiting the current context of opportunities;
- a **synthesis** of the points above.

The following is minimally technical and contains links provided to the reader in need of more information.

Perspective and Background

A simple high-altitude perspective on what currently attracts my interest is to identify, implement, test, adapt, and exploit, in the light of arising opportunities offered by the latest technology, the various ideas in computer-based Knowledge Representation and Automated Reasoning that I and others have conceived and elaborated over the past 40 years.

As for background, my qualifications and career have been in scientific research, with principal interests in software development for Artificial Intelligence (AI)—formalisation, application, and practice.[1] My area of expertise is in the conception and implementation of tools for knowledge processing—representation and acquisition, automated reasoning (theory and practice). My work has now become even more relevant in the emerging context of superabundant data and the need to extract, represent, and use knowledge from raw data. I believe my ideas have direct practical, marketable, money-making potential for applications in very "hot" areas where decision-making relying on knowledge-based intelligent assessment of data is critical.

The most important of my research contributions could roughly be summed up as the Order-Sorted Feature ($\mathcal{OSF}$) constraint formalism. This formalism has mostly been exploited in Declarative Programming—*viz.*, Constraint Logic Programming (CLP). It has also proven invaluable in particular for Natural Language Processing (NLP)—*viz.*, so-called "Head-driven Phrase Stucture Grammar" (HPSG) and Unification Grammar (UG) technology.[2]

The $\mathcal{OSF}$ formalism, while simple and intuitive, provides a powerful means for expressing and using data and knowledge declaratively. It is relevant both to deduction (using $\mathcal{OSF}$ unification) and induction (using $\mathcal{OSF}$ generalization).[3] It has thus provided a solid basis for expressing and processing conceptual knowledge with what could be dubbed "constraint-based *Ontological Programming*."[4] The latter designates the declarative specification of knowledge structures where semantic properties of concepts are expressed as constraints which can then be efficiently enforced through constraint-solving. In particular (but not only), it provides a clean executable means to specify and use formal distributed ontologies. Also, being modular in terms of constraints and algebraic structures, its formal basis readily accommodates fuzzy or probabilistic reasoning. All this gives it a versatile power to be exploited in applications relying on knowledge and data bases distributed over the Internet.

In what remains of this document, I shall introduce myself better in order to justify this claim, starting at the origin.

[1] `http://www.hassan-ait-kaci.net/`

[2] For an instructive historical account on this point, please refer to Section 9, Pages 82–84, of this article by Gert Smolka.

[3] The latter is invaluable in learning.

[4] For lack of a prettier name.

Research Motivation

I did my undergraduate education in Mathematics and Informatics in France, graduating in 1976. I took a double major in Operations Research and Computer Science. I pursued my graduate studies in the USA first in Operations Research, starting at Stanford University then continuing at the Wharton School of Business, at the University of Pennsylvania (Penn). There, I studied for a PhD in OR, completing all doctoral course requirements and passing doctoral prelims. I acquired advanced competence in mathematical programming, deterministic and stochastic modeling, decision theory, multiple-criteria decision making, Utility Theory, and Saaty's theory of hierarchies and eigenvalue analysis.

In 1980, while a graduate research assistant in Wharton's Department of Decision Sciences, I contributed to the integration of a semantic network knowledge base into a decision support system: the Decision Aiding Information System (DAISY). My computer science background had served me implementing several OR models only to realize that the real challenge was not in implementing any specific model, but in being able *to formalize the modeling process* itself. This had to involve being able to describe and represent formally knowledge about models and what criteria characterized the process of identifying their adequacy to specific situations. I then tried to explicate how I would approach the "design of automated consultant" in terms of the technology of the time (this was when Prolog was not yet known and still being conceived, and the term "Expert System" had not been invented yet).

This made me realize that the essential challenge in what I had in mind was not really in OR but resided in AI—more specifically in Knowledge Representation and Automated Reasoning. My interest in those issues made me then decide to transfer to Penn's departement Computer and Information Science in the Moore School of Electrical Engineering. There, I had the opportunity to become proficient in the theory and implementation of Theorem Proving and Declarative Programming. That gave me all the formal tools to contribute with the design of a Knowledge-Base Language. I obtained my MSc in 1982 and my PhD in 1984, both at Penn, as well as my research director *habilitation* at Paris 7 university. I proposed using sorted graph unification to express and reason with ontological knowledge in which I laid out the basis of what has eventually become known as the $\mathcal{OSF}$ constraint formalism, thus giving me the means to enhance Logic Programming and Functional Programming both in expressivity and performance, both programming constructs used in turn to magnify the declarativeness and efficiency of ontological programming as one pleased. This has proven invaluable in the context of large-scale knowledge and data processing and distributed architectures.

Lifelong Pursuit (1984–today)

Let me now brush a rapid account of my research life since my PhD.

- ***1984–1988: MCC***—Conducted research at the Microelectronics and Computer Technology Corpration in declarative "new generation" programming languages and architectures, functional, logic, constraint, and object-oriented programming, type theory, symbolic computation, abstract machine compilers. This thas the first prototype of what was developed next at DEC as the $\mathcal{LIFE}$ CLP programming language.
- ***1988–1994: DEC***—Research Project Director (the Paradise Project) Digital Equipment Corporation, Paris Research Laboratory, Rueil-Malmaison, France. I set up and led a research project on executable constraint specifications. This effort extends my previous work on $\mathcal{LIFE}$ to encompass theory and practice of constraint-based programming, including the design and implementation of a multi-paradigm environment for advanced applications with enhanced programming productivity. The design took form as the Wild $\mathcal{LIFE}$ CLP language and system.
- ***1995–2001: SFU***—Full professorship appointment with tenure at Simon Fraser University, Burnaby, BC, Canada; Director, Intelligent Software Group. Until December 1996, NSERC Industrial Chair

in Intelligent Software Systems. Research: Objects and Constraints for Intelligent Internet Resource Processing. Interests in programming paradigms based on logic, constraints, and type theory, and the formal interplay of computation and approximation. Teaching - CMPT117: Introduction to Internet Programming in Java and its Friends, CMPT-212: MS Windows Application Programing with C++, CMPT-383: Principles of Programming Languages, CMPT-379: Compiler Design, CMPT-384: AI/Symbolic Computing, CMPT-883: Graduate seminar on advanced Internet Application Design using Java, XML, and network object/service brokering.

- ***2000–2009: ILOG***—Distinguished Scientist at ILOG, Gentilly, France. in charge of research projects with the Product Division (R&D). Technology watch and technology transfer. Areas of expertise: Theory and practice of Artificial Intelligence: Formal logic, Business Rule Management Systems, Constraint-based computing, Semantic Web, Knowledge management and discovery, Probabilistic Graphical Models. Since 2005: Principal ILOG member of the W3C Working Group on a Rule Interchange Format (RIF), a crucial part of the Semantic Web effort.
- ***2009–2013: IBM***—Senior Technical Staff Member reconducted in all my capacities at ILOG, now an IBM company since January 2009.[5] IBM ILOG's scientific coordinator for the European Project ONTORULE. Facilitator between newly created IBM Center for Advanced Studies (CAS) in Paris and IBM Canada.
- ***2013–2105: UCBL***—Recipient of an ANR Chair of Excellence heading the research project $\mathcal{CEDAR}$ (Constraint Event-Driven Automated Reasoning),[6] which was carried out over a two-year period in Villeurbanne, France, at the *Laboratoire d'InfoRmatique en Image et Systèmes d'information* (LIRIS) of the *Université Claude Bernard Lyon 1* (UCBL). The project set up a systematic experimentation with $\mathcal{OSF}$ knowledge representation technology departing from the official, but not really tested, W3C standards based on Description Logic such as OWL, and addressing two essential challenges: scalability and distribution in knowledge base management systems.
- ***2015–2106: UCBL***—Director of the $\mathcal{L}iv\mathcal{EMUSIC}$ Project (Living-Environment Monitoring Use Scenario with Intelligent Control),[7] Programme Avenir Lyon Saint-Etienne (PALSE) at the LIRIS at the UCBL. The project developed a convincing use case in intelligent living-environment monitoring. The technology demonstrated in this use case leverages knowledge representation and automated reasoning as means to enable "smart" monitoring of living environments, focusing on urban and social milieu. It combines more semantic-based techniques for distributed knowledge and information processing from areas such as knowledge representation and reasoning, and intelligent query processing.

I can state unequivocally that my lifelong pursuit has been the specification and efficient implementation of the $CLP(\mathcal{OSF})$ paradigm incarnated in the $\mathcal{LIFE}$ programming language. I am confident that this experience is now more than ever appropriate in the context of intelligent processing of massively distributed knowledge and data such as, in particular, the so-called Semantic Web.

Current Potential

Mine has indeed been a lifelong pursuit. As we stand, I now have gathered all the required experience enabling me to build and demonstrate the power of efficient (distributed) reasoning exploiting (distributed) knowledge and data in real time through streamed incremental (distributed) processes evolving cooperatively as information is fleshing out (retrieved and/or synthesized)—locally or over a network. Incrementality is paramount and is captured formally as the invariant illustrated in Figure 1—namely, *approximation and evaluation must commute*. Why is this important? Because it is the key to distributed concurrency and it is formally characterized as an intuitive (indeed innocuous, one might even say!), yet pragmatically formidable,

[5] http://www-03.ibm.com/press/us/en/pressrelease/26403.wss

[6] http://cedar.liris.cnrs.fr/

[7] http://cedar.liris.cnrs.fr/livemusic/

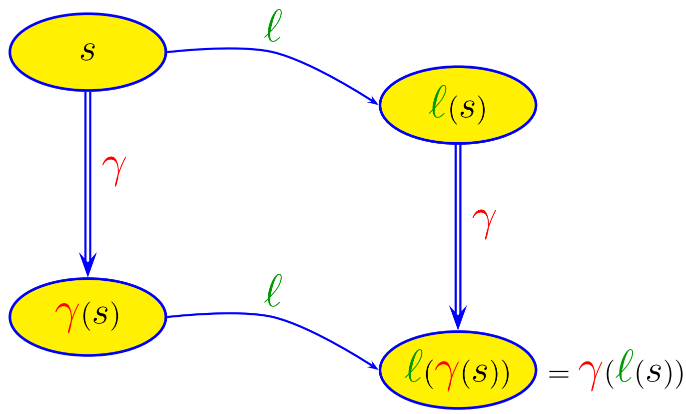

Arrow :	**Process**	**Operation**	**Denotation**
γ :	Approximation	Constraint Implication	$\mathcal{OSF}$ Endomorphism
ℓ :	Evaluation	Function Application	$\mathcal{OSF}$ Feature Projection

Figure 1: Incremental structure inheritance is commutativity of approximation and evaluation

property. The $\mathcal{OSF}$ constraint computational formalism that I have relentlessly striven to explore and develop with the help of colleagues and students over the years definitely possesses this property indeed.

All is done through logically and functionally constrained order-sorted featured graphs, where the ordering is that of approximation and feature projection is evaluation (cached features can thus be also used for lazy evaluation). This is also a remarkably clean and simple modular approach since all is abstracted through the algebraic structures of the underlying mathematical operations expressing a specific variety of approximation precision (such as various monoids or, more expressively, distributive rings and lattices). This can thus be made to tolerate partial, fuzzy, probabilistic, or even partially inconsistent reasoning as long as approximation and evaluation are allowed to commute—which is always possible in meaningful situations, and thus realizable. Distributed concurrency is the key to making all this work in harmony with the Internet (whether of sites, things, agents, services, *etc.*).

Synthesis

As an academic research statement, this short document has attempted to give a concise but clear general idea of the origin and value of my research contribution, and its potential for now and times to come.

- ***Why?***—It is my conviction that today's distributed net-oriented computing environment is more than ever appropriate for me to demonstrate the various claims I make in the previous sections.
- ***How?***—My confidence is such that with adequate means granted to me (research position, personnel, and work conditions), I can commit to specifying and implementing a convincing demonstration system in five years at most.
- ***What?***—As my experience has abundantly shown me, high-caliber graduate students, engineers, and post-doctoral fellows, can be organized into a multi-instrument orchestra to conduct on scores I can compose. As well, benefitting from the proximity of, and interaction with, colleagues that are experts in complementary yet necessary areas and whose own activities my ideas can profit, is therefore ideal.

It is my hope that this statement has achieved its goal. I am available for further questions.