Microsoft Research - Inria
JOINT CENTRE

Scientific Report 2011-2013

Microsoft Research-Inria Joint Centre

www.msr-inria.inria.fr

Introduction

Before the Joint Centre's creation in October 2005, informal collaborations already existed between Microsoft Research and Inria researchers, and this most prominently in the domain of so-called *formal methods*. These methods build on mathematical logic and programming language theory to enable automatic verification, and to some extent synthesis of software with pre-specified properties.

One of the stated goals of the Joint Centre was then to extend the applicability of formal methods to three challenging areas: distributed systems, security and cryptography, and finally pure mathematics. In the latter domain the goal is to provide computer tools to assist mathematicians in the verification of difficult mathematical proofs.

Over the years, the ambition and scope of the Joint Centre has broadened to encompass new objectives in other areas. These include combinatorial optimization based on machine learning techniques, mining of images and videos, and more recently analysis of brain and cardiac fMRI images. Behind the variety of these new domains lies a methodological unity as all rely fundamentally on machine learning.

The Centre's activity for the 2011-2013 period was structured in 8 distinct projects. As testified by the project highlights to follow, the Centre has now attained a sufficient level of maturity to produce world-class research results promising significant societal impact.

Highlights

Mathematical Components

This project aimed at developing libraries of software components formalizing mathematical theories together with a formal language for expressing mathematical proofs. In turn the validity of a proof expressed in this language could be automatically checked by computer, using the "Coq" proof assistant. From its onset in 2006, the project focused on performing a computer-checked proof of the so-called Feit-Thompson theorem, a landmark result in group theory and a key step in the classification of finite groups.

This announced goal was achieved in September 2012, thus demonstrating that the proposed approach could indeed be applied to checking "serious" (daunting, that is) mathematical proofs. In the process, the Ssreflect language matured, together with a whole library of readily re-usable components encoding many branches of Algebra such as group theory, character theory, linear algebra… These developments have already been used in independent efforts by other teams to check other challenging mathematical proofs. Our achievements in this project have been greeted in the popular press ("La Recherche", "Tangentes", "Sciences et Vie"). The fundamental work done in the process has led our researchers to contribute to a monograph published by the Princeton Institute of Advanced Studies on rethinking the logical basis of mathematics, discussed in "the New Scientist".

Secure Distributed Computations and their Proofs

This project develops tools to certify the security of transactions over the Internet. This consists on the one hand in certifying properties of embedded code written in JavaScript and on the other hand in certifying the correct execution of cryptographic encryption primitives underlying for instance the HTTPS protocol systematically used for e-commerce. On the first topic our researchers have developed a compiler synthesizing JavaScript code from source code written in F*, a language more amenable to verification. This approach bypasses the need to directly certify JavaScript code. On the second topic our researchers have identified serious bugs in widespread implementations of TLS (which underlies HTTPS among other protocols). They have developed and certified miTLS, the first TLS implementation offering such a level of security guarantees. Finally, they have developed ZQL, a tool transforming an SQL-like query into a cryptographic protocol computing the result of the query based on zero-knowledge proofs, thereby guaranteeing that the only information leaked is precisely the result of the query. ZQL has already been used to develop privacy-preserving smart-meter billing protocols. Many more applications can be envisioned, e.g. to voting. These results have been published in the best conferences of theirs domains, notably ACM POPL and Usenix Security.

Tools for Proofs

This project develops a "proof assistant" for programs written in the TLA ("temporal logic of actions") language. TLA was conceived by Leslie Lamport, one of the project leaders, to specify algorithms for distributed systems such as distributed databases, where there is no control over when executions on specific machines will complete. A first proof assistant was developed, thereby allowing to certify the validity of algorithms for solving the so-called consensus problem in the presence of corrupted machines. Consensus is a fundamental primitive in distributed systems. A subsequent goal of the project is to certify with its proof assistant the conception of schedulers of specific real-time operating systems.
In addition to publications in the leading venues in distributed computing such as the DISC conference, this work has given rise to an article in "Wired" and many keynotes. Leslie Lamport has received the 2013 Turing award in recognition of his fundamental contributions to distributed computing.

Dynamic Dictionary of Mathematical Functions

This project aimed at developing an interactive web site which would dynamically synthesize formulas, expansions and numerical evaluations of a large catalogue of so-called special functions (examples of which are Bessel and Airy functions e.g.), frequently encountered in mathematical physics among other domains. Underlying the web site, a formal calculus engine produces numerical evaluations together with error bounds at an optimized algorithmic cost. This web portal has found a large audience and is being used notably as a tool for teaching. The research performed to enable this development has been published in top venues (such as the IEEE FOCS conference) and has earned a best student paper award at the ISSAC'13 conference.

Adaptive Combinatorial Search for e-Science

This project develops efficient algorithms for solving generic combinatorial problems of the so-called "constraint satisfaction" class, whose vast range of applications goes from computer Go to workflow optimization in logistic operations. Our researchers developed solutions based on constraint programming which they applied successfully to these two application domains. They strived to parallelize their approaches leveraging so-called "Multi-Armed Bandits" techniques of adaptive learning. They also developed several generic adaptive search methods, and most notably the so-called "Covariance Matrix Adaption Evolution Strategy" (CMA-ES). The latter is the state-of-the-art in generic practical tools for black-box optimization over continuous parameters.

Scientific Image and Video Mining

This project develops methods in computer vision to process large collections of images or videos. It focuses on two main applications. Concerning image processing, it aims to build 3D models from large collections of images of a given scene. This can in turn provide tools for heritage preservation. For example, an outcome of the project has been to collect from a drone pictures of part of the Pompei archeological site; a 3D model of the site has then been synthesized and further checked against 19th century drawings of the same site, thereby revealing among other things the site's degradation since then. This part of the project led to the creation of the Iconem startup, based in Paris, in 2013. Beyond heritage preservation, the overall approach has other applications, to urban planning for instance.
Concerning video processing, the project aims to index huge catalogues of videos to allow efficient search supporting rich queries such as finding videos shot in a particular geographical area, featuring a particular actor, eg. The researchers have made steady progress towards these objectives. The corresponding results have been published not only in the best research venues of the domain (Siggraph in particular), but also in the popular press (Wall Street Journal, Wired, New Scientist).

A-Brain

This project aims to correlate patterns of brain activation as observed through functional MRI with the patient's genetic information, in order to determine which aspects of brain physiology are "inherited" rather than "acquired". The size of both genetic and fMRI data is such that cloud computing becomes an appealing approach to identify such correlations reasonably quickly. Our researchers developed statistical tests for determining such correlations, together with software solutions (specifically a middleware layer) facilitating the execution of these tests in the cloud. More precisely this middleware provides a file system abstraction hiding the location of data replicas, which it automatically handles.
The proposed approach has been tested on the Microsoft Azure cloud platform. As a result it established that the activation of a specific part of the brain (the sub-cortical nuclei) when a patient fails on a test is an inheritable trait, conditioned by the ARVCF gene. These works have been well published both in the biomedical imaging community (MICCAI conference) and in the distributed systems community (IEEE ISPA and IEEE IPDPS).

4D Cardiac MR Images

This project develops fine-grained, physiologically realistic models of human heart behaviour, to be later used to assist diagnosis of heart conditions on the basis of fMRI observation. The models are based on a precise description of mechanical and physiological processes as well as clinical descriptors used by physicians. Advanced machine learning techniques are used to de-noise the observations and determine a heart model personalized to match the patient's characteristics. A collaborative platform of cardiac images annotated with clinical information has been developed. This recent project, launched in 2011, has already been distinguished by a "Young Scientist Award" at the MICCAI 2012 conference.

Perspectives

With the renewal of the Joint Centre for the period October 2013-October 2017, we have reviewed and expanded our project portfolio.
While building on the acquired momentum, notably in **formal methods** on the one hand and in **computer vision and biomedical imaging** on the other hand, we have now explicitly identified two additional research areas in which we will invest. The first one is **Machine Learning and Big Data**. While this theme was already present in a number of projects, we want to emphasize further its importance, and have dedicated a specific project to generic Machine Learning, focusing in particular on the issues of adaptive learning and sparse representations. We have also launched a new project on efficient processing of workflows in cloud environments. Finally our fourth theme is about **Social Information Networks and Privacy**. Our objectives in this space will be (i) design of efficient algorithms for facilitating information access through Online Social Networks, e.g. through community detection for contact recommendation and incentive schemes; (ii) design of privacy notions and methods for ensuring such privacy when processing private information such as geo-localization and biometric signals. The challenge being to manage the tension between privacy and application usefulness.

Mathematical components

www.msr-inria.fr/projects/mathematical-components-2/

This project started in summer 2006.

Summary

Formalized mathematical theories can, like modern software, be built out of components. By components we mean modules that comprise both the static (objects and facts) and dynamic (proof and computation methods) contents of theories. We develop a general platform for mathematical components, based on the Coq "ssreflect" extension that was used to carry out the formalization of the Four Color Theorem. Although we used the formalization of a seminal result in Group Theory, the Odd Order Theorem, to drive our development, the components that we are developing — for Logic, Combinatorics, Set Theory, Algebra, Linear Algebra, Group Theory, Graph Theory, and Finite Field Theory — are widely usable.

Personnel

Project	Status	Last name	First name	Affiliation
Mathematical Components	Team leader	GONTHIER	Georges	Microsoft Research Cambridge
	Site leader	BERTOT	Yves	Inria Sophia Antipolis-Méditerrannée
	Researcher	MAHBOUBI	Assia	Inria Saclay-Île-de-France
	Researcher	RIDEAU	Laurence	Inria Sophia Antipolis-Méditerrannée
	Researcher	TASSI	Enrico	Inria Saclay-Île-de-France
	Researcher	THÉRY	Laurent	Inria Sophia Antipolis-Méditerrannée
	Post doc	O'CONNOR	Russell	Microsoft Research-Inria Joint Centre
	PhD student	CANO	Guillaume	Microsoft Research-Inria Joint Centre
	PhD student	COHEN	Cyril	École Polytechnique
	PhD student	DÉNÈS	Maxime	Inria Sophia Antipolis-Méditerrannée
	PhD student	GARILLOT	François	École Normale Supérieure de Paris
	PhD intern	SOJAKOVA	Kristina	Carnegie Mellon University
	PhD intern	SOLOVYEV	Alexey	University of Pittsburgh

The permanent researchers of the Mathematical Components project are Georges Gonthier (Principal Researcher, Microsoft Research), Yves Bertot (Senior Researcher, Inria), and Assia Mahboubi, Laurence Rideau, Enrico Tassi, and Laurent Théry (Researcher, Inria). Russel O'Connor was a joint postdoc with us and McMaster University (Canada) from April to July in both 2011 and 2012. We had four PhD students with the project at the start of 2011, François Garillot, Cyril Cohen, Maxime Denès and Guillaume Cano; the first two defended their thesis in 2011 and 2012, respectively, and Maxime Dénès will be defending in November 2013. Finally, Alexey Solovyev was an intern with us in 2011, as was Kristina Sojakova in 2013.

Careers

François Garillot went on to be a development engineer at Criteo (Paris), and then Typesafe (Geneva). Cyril Cohen is a postdoc at Chalmers University (Göteborg), with Thierry Coquand; he also joined him at the Institute of Advanced Studies in Princeton during the Special Year on Homotopy Type Theory. Maxime Dénès started a postdoc in Princeton with Andrew Appel in September 2013.

Highlight

On September 20, 2012 the project team reached what had been its main scientific objective since its

inception: a fully machine-checked formalization of the Feit-Thompson proof of the Odd Order Theorem. This is the first time that such an important and technically difficult proof has been formalized, and thus marks an important milestone for computer tools for mathematics.

The Odd Order Theorem states that all finite groups with an odd number of elements are *solvable*, that is can be decomposed into a series of factor groups of prime order. It is widely used in Group Theory and its applications. In particular, it is equivalent to the fact that the only *simple* (non-factorable) groups of odd order are the (cyclic) groups of integers modulo some prime *p*; as such it was the starting point of the Classification of all finite simple groups, one of the most monumental and important achievements in modern Algebra, whose proof runs over 8,000 pages.

Though smaller, at 250 pages, the Odd Order proof still represented a major technical challenge pushing the boundaries of machine theorem proving. The textbook proof required months of study to be understood by specialists, and indeed it took some 20 years of revision to weed out all the technical errors from the original text, even though its general argument was believed correct. Moreover the proof draws on about 4-500 textbook pages of background material, covering most topics in undergraduate Algebra and graduate Group Theory, from polynomials and matrices to eigenspaces, Galois theory and algebraic numbers, and from Sylow's theorems to Glauberman's ZJ-theorem and Feit's exceptional characters.

It is this range of topics that made the Feit-Thompson proof an attractive testing ground for the Mathematical Components project: it would both ensure that the tools and methods we developed would work for a variety of theories, and that these theories could be combined effectively. Perhaps the fact that the Odd Order proof was completed *as planned*, on time, is the best evidence of the success of our approach: it shows that all formalization issues had been adequately identified and resolved while working on the background theories. Indeed, this "preliminary" work was about twice that of the "main" proof (4 years / 90,000 LoC vs. 2 years / 45,000 LoC), roughly in proportion to the informal texts, even though the Feit-Thompson proof material was far more difficult.

One of the issues we had to resolve was finding effective ways of interpreting mathematical notation correctly. Working mathematicians rely on large, often implicit sets of conventions to make sense of deliberately ambiguous notation and proof arguments. Expliciting all of these to formal logical precision would make both statements and proofs impractically large. The alternative we pursued was to formalize the conventions so that they are understood by the formal proof system, using a variety of programming language and software engineering techniques, such as type inference and components. For instance we used a dependent class function type to disambiguate dot product in character theory, and a form of refinement type to assimilate subgroups to subsets. These combined techniques turned out to be more effective than we had hoped: not only were we able to match even the most elaborate notation and proof, but in some cases, such as for Thompson's critical subgroup theorem, the computer's interpretation of the notation helped us understand the proof.

While we did not uncover any significant error in the extensively revised Feit-Thompson proof, our formal proof did pick up a few dozen minor mistakes, ambiguities, and simplifications, showing that our work was not purely academic exercise. In one quaint instance, a misprint in the concluding summary of the first of the two edited volumes of the revised proof was picked up as an assumption for the second volume, though fortunately an unused one. Many of the simplifications came from refactoring the formal proof, using the precision of the formal proof to explore its dependencies. In one combinatorial proof about arrays of near-orthogonal exceptional characters, a bespoke decision procedure uncovered a new symmetry that superseded one third of the textbook proof.

Tools and software

Version 1.3 of Ssreflect, released in March 2011, featured a major expansion of the core Ssreflect 1.2 theory to cover primes, summations (and more generally "big operators"), finite sets, functions, and groups, and general algebra with matrices and polynomials, including a complete algebraic hierarchy. It also features comprehensive internal documentation for each module in the library.

Version 1.4 of Ssreflect was released in September 2012 and includes numerous improvements to the Ssreflect proof language and its implementation, including contextual patterns. This release extends further the theory library to cover general linear algebra, and advanced group theory including a comprehensive treatment of non-linear and linear group representation theory, all with extended documentation. We have also published a comprehensive Ssreflect tutorial.

Version 1.5 of Ssreflect, released in March 2014, adds more support for symmetry arguments ("without loss" and "generally have"), and support for embedding proofs in definitions. We also separated the Ssreflect core support library from the more elaborate MathComp mathematical libraries.

The CoqEAL library provides efficient proven implementations for much of the linear algebra operations in the MathComp libraries, as well as support for modules (linear algebra over Euclidean rings).

Research

The long-term scientific objective of the Mathematical Components project is to demonstrate that state-of-the art formal method tools can be successfully applied to research-level mathematics.

More precisely, our thesis is that the explanation for the limitations of the existing libraries of formal mathematics can be traced to those of the modularity tools used to compose elements of those libraries. We strive to combine modern software engineering concepts such as components with the advanced type system and type inference procedures of the Coq proof system to overcome those limitations.

As described in the Highlight section, we validated our approach by completing the formalization of the Feit-Thompson proof of the Odd Order Theorem. By early 2011 we had completed most of the required background theories, as well as the first part – Local Analysis – of the proof. In the final 18 months we completed both the Character Theory part of the proof, as well as the new background results needed, in Euclidean geometry, character theory, and Galois theory, including two novel constructive definitions of the algebraic numbers.

Along with the refactoring and final polishing of the Odd Order proof, we also worked on the formalization of Apéry's proof of the irrationality of $\zeta(3)$, with Frédéric Chyzak, who is joining the project. This project required integrating large computer algebra (CAS) computations into Coq/Ssreflect proofs.

Library organisation

Canonical Structures are a very simple mechanism: they are simply hints that let the type inference algorithm of Coq extrapolate the value of certain records from a single one of their fields. In combination with other features of the Coq system — dependent types, higher-kinded records, convertibility, coercions and user notations with inherited attributes — this apparently crude mechanism turns out to be extremely expressive, basically because it allows a library designer to reprogram the behavior of value matching, unification and type inference to suit his specific needs.

Function and predicate hierarchies In mathematics, structures are a means of associating certain operations and properties to sets; mathematical structures are organized by a complex web of inheritance relations; for example fields are a refinement of rings, while algebras are both vector spaces and rings. We had worked out previously how to use canonical structures for this purpose in formalized mathematics. While working out the formalization of character theory, we needed to extend this hierarchy to both predicates and functions. For predicates we needed to support the automatic inference of an intricate lattice of 13 different closure properties, as well as supporting such properties for local predicates, using the new Prop-irrelevance support in Ssreflect. For functions we needed a generic notion of linearity encompassing scalar and vector linearity as well as sesquilinearity in order to deal with the Hermitian geometry of characters. We could handle both reliably using Coq's Canonical Structure declaration.

Number interfaces The character theory part of the Feit-Thompson proof relies extensively on numerical inequalities involving norms and values of (complex) characters. This involves the casual use of complex quantities in real-number inequalities; naïve formalization in typed logic would have required ubiquitous and unwieldy casts to real. We discovered that by instead extending linearly real inequality to a partial order on complex numbers we could avoid a real number type altogether, as this order satisfies most properties of the real order unconditionally. We packaged these properties in a *number structure hierarchy* that conveniently supports the casual mixing of order and arithmetic.

Lemma Overloading We had previously determined that canonical structures could implement quotation – making the syntactic shape of an expression available during type inference – and had used this to formalize direct sums of matrix row spaces and vector subspaces. Along with an extensive complementary set of techniques, we have applied this to implement *lemma overloading*: a single theorem that adapts its statement to a variety of syntactic situations, using Coq's canonical structure-enhanced higher-order unification. We applied these techniques to Hoare logic proofs of heap-manipulating code, and to the generalized linearity mentioned above.

Quotients and effective algebra Cyril Cohen developed an extended framework for both constructive quotients, supporting generic construction from equivalence on types with choice, specific constructions in other cases, and providing generic means of relating values and properties of the representation and quotient types. Initially created to support the construction of rational and real algebraic numbers, this work was used extensively in the formalization of field and Galois theories. It led to the development of a similar framework for relating the simple but inefficient definition of linear operations in the MathComp library to more effective ones, and the development of the CoqEAL library.

Constructive algebra

Because the Odd Order theorem is a result in finite group theory, we expect its proof to be constructive; this should be expressible in Coq, since the core logic of Coq is also constructive, and we have set ourselves this as an additional goal. To do so we had to devise replacements for some of the non-constructive general algebra results used in the classic proof. This has led to some new and intriguing math.

Algebraic numbers. One of the reasons the proof of the Odd Order theorem proof is not obviously constructive, even though is statement is manifestly decidable, is that its character theory part uses (non-constructively defined) real and complex numbers. We resolved this by using algebraic numbers in our formal definition of characters, and providing two novel constructive definitions of algebraic numbers. The first, due to C. Cohen, obtains the algebraic reals as a discrete subset of the constructive reals, then uses a matrix-based proof of the Fundamental Theorem of Algebra (FTA) to show the closure of the complex algebraics. The second obtains the algebraics from a general closure construction on countable fields, and then uses the Artin proof of the FTA to construct the conjugation automorphism.

Galois theory Russell O'Connor completed the formalization of constructive Galois theory, including the construction of (finite) Galois groups, the Fundamental Theorem, the independence of linear characters, and Hilbert's Theorem 90. The structure supporting these constructions provides only constructive evidence that the extension is normal (in the form of a split generating polynomial); separability is asserted independently.

The Feit Thompson proof

The Character Theory part of the Feit Thompson proof relies on an entirely new set of concepts and techniques that did not appear in the Local Analysis part we had formalized in 2010. This first part relied mostly on group identities, finite set combinatorics, elementary (prime) number theory, and concrete (matrix-based) linear algebra over finite fields (for group representations). The second part considers characters, which live in abstract vector spaces over the complex algebraic numbers, equipped with a Hermitian geometry, and relies on norm inequalities with integrality constraints.

Character Theory Group characters are the traces of the matrices of a linear representation of a group, and, as is often the case, in the Feit-Thompson proof only complex characters are considered. Characters of a group *G* are naturally an integral subset of the Hermitian space of class functions over *G*; this (dependent) domain type captures remarkably well common notational conventions. Character theory is a large subfield of group theory; the Feit-Thompson proof uses almost all of its basic results.

The Dade isometry and the Suzuki inequalities. The central idea of the second part of the proof is to use characters to derive information about a hypothetical counter-example *G* to the theorem from what Local Analysis has uncovered about its maximal (proper) subgroups. Specifically, the structure of the character tables of each of these subgroups *M* is largely known, and due to way they are embedded in *G* there is an integral isometry mapping certain subspaces of class functions of *M* to class functions of *G*; in the precursor proof by Suzuki of a special case of the theorem, this isometry was just character induction, but the revised Feit-Thompson proof uses a more general construction due to Dade. If the isometry domain is coherent, that is, is the isometry extends to a subspace generated by characters, then a certain norm inequality must hold, which Suzuki showed to be impossible. This central part of the proof was formalized during a summer internship by Alexey Solovyev, who had no prior knowledge of Coq, Ssreflect, or the MathComp library.

Cyclic TI-subsets The Suzuki inequality forces the existence of a special form of maximal group, containing a cyclic TI-subgroup *W* that is disjoint from all its conjugates in *G*. Because W is a non-trivial product of odd order, this implies that induction extends to an isometry on all class functions on *W*, i.e., the entire space is coherent. Remarkably, this is proved by a purely combinatorial argument, which can be checked by off-the-shelf SMT and SAT solvers… with sufficient preprocessing in Coq. Ultimately we realized we could just as easily define a certified bespoke solver inside Coq, and the flexibility and insight afforded by this approach let us uncover a hidden symmetry that subsumed a third of the textbook proof.

Symmetry arguments The final part of the proof largely exploits the symmetry between the two "exceptional" subgroups with cyclic TI-subgroups, freely using properties established for one subgroup for the other. Formalizing faithfully this kind of loose reasoning rubbed against limitations of the Coq vernacular, but was an excellent showcase for the symmetry support in the Ssreflect proof language ("`without loss`"), and inspired further extensions of these facilities in Ssreflect 1.5 ("`generally have`").

Applications

In addition to our main effort on the Feit-Thompson proof, we have worked on other applications of our library of mathematical components.

The FORMATH consortium In 2009 we joined with the Universities of Chalmers (Sweden), Nijmegen (the Netherlands), and La Rioja (Spain) in a project to extend the ssreflect mathematical libraries to cover advanced algebra and linear algebra, some analysis, and algebraic topology. The European Community ICT Formath (Formalization of Mathematics) concluded in 2012, in particular with the successful formalization of algebraic topology computations in ssreflect.

MathComp usage In 2012 we organized a MAP Spring School on mathematical theorem proving with ssreflec, which attracted over 60 participants. Both Ssreflect and the MathComp libraries are now in active use by external groups, and features in many research publications, such as Reynaldt et al. on information theory, Smolka et al. on modal logic proof theory, Benton et al. on x86 models, and Strub et al. on elliptic curve cryptography.

Journal papers and book chapters

[1] Cyril Cohen, Assia Mahboubi. Formal proofs in real algebraic geometry: from ordered fields to quantifier elimination. Logical Methods in Computer Science 8(1) (2012).

Conference and workshop papers

[2] Jónathan Heras, María Poza, Maxime Dénès, Laurence Rideau. Incidence Simplicial Matrices Formalized in Coq/SSReflect. Calculemus/MKM 2011: 30-44, 2011.
[3] Georges Gonthier. Point-Free, Set-Free Concrete Linear Algebra. ITP 2011: 103-118, 2011.
[4] Georges Gonthier, Beta Ziliani, Aleksandar Nanevski, Derek Dreyer. How to make ad hoc proof automation less ad hoc. ICFP 2011: 163-175, 2011.
[5] Michaël Armand, Germain Faure, Benjamin Grégoire, Chantal Keller, Laurent Théry, Benjamin Werner: A Modular Integration of SAT/SMT Solvers to Coq through Proof Witnesses. CPP 2011: 135-150, 2013.
[6] Mathieu Boespflug, Maxime Dénès, Benjamin Grégoire. Full Reduction at Full Throttle. CPP 2011: 362-377, 2011.
[7] Nicolas Brisebarre, Mioara Joldes, Érik Martin-Dorel, Micaela Mayero, Jean-Michel Muller, Ioana Pasca, Laurence Rideau, Laurent Théry. Rigorous Polynomial Approximation Using Taylor Models in Coq. NASA Formal Methods 2012: 85-99, 2012.
[8] Cyril Cohen. Construction of Real Algebraic Numbers in Coq. ITP 2012: 67-82, 2012.
[9] Maxime Dénès, Anders Mörtberg, Vincent Siles. A Refinement-Based Approach to Computational Algebra in Coq. ITP 2012: 83-98, 2012.
[10] Georges Gonthier, Enrico Tassi. A Language of Patterns for Subterm Selection. ITP 2012: 361-376, 2012.
[11] Georges Gonthier. Engineering mathematics: the odd order theorem proof. POPL 2013: 1-2, 2013.
[12] Assia Mahboubi. The Rooster and the Butterflies. MKM/Calculemus/DML 2013: 1-18, 2013.
[13] Cyril Cohen. Pragmatic Quotient Types in Coq. ITP 2013: 213-228, 2013.
[14] Georges Gonthier, Andrea Asperti, Jeremy Avigad, Yves Bertot, Cyril Cohen, François Garillot, Stéphane Le Roux, Assia Mahboubi, Russell O'Connor, Sidi Ould Biha, Ioana Pasca, Laurence Rideau, Alexey Solovyev, Enrico Tassi, Laurent Théry: A Machine-Checked Proof of the Odd Order Theorem. ITP 2013: 163-179, 2013.
[15] Assia Mahboubi, Enrico Tassi. Canonical Structures for the Working Coq User. ITP 2013: 19-34, 2013.
[16] Cyril Cohen, Maxime Dénès, Anders Mörtberg. Refinements for free! CPP 2013.

Theses

[17] François Garillot. Generic Proof Tools and Finite Group Theory. PhD thesis, École Polytechnique, December 2011.
[18] Cyril Cohen. Formalized algebraic numbers: construction and first-order theory. PhD thesis, École Polytechnique, November 2012.
[19] Maxime Dénès. Étude formelle d'algorithmes efficaces en algèbre linéaire. PhD thesis, Université de Nice - Sophia Antipolis, November 2013.

Talks

[20] Georges Gonthier. Advances in the Formalization of the Odd Order Theorem (keynote). ITP 2011: 2, 2011.
[21] Georges Gonthier. Proof Engineering, from the Four Color to the Odd Order Theorem (keynote). CPP 2011, 2011.
[22] Georges Gonthier. Types in Mathematical Proofs. *Milner Symposium*, 2012.
[23] Georges Gonthier. Typing Patterns for Mathematics. *FOMCAF* 13, 2013.
[24] Georges Gonthier, Assia Mahboubi. A Computer-Checked Proof of the Odd Order Theorem. *Netherlands Mathematical Congress*, 2013.
[25] Georges Gonthier. Software engineering for mathematics (keynote). *ESEC/SIGSOFT FSE* 2013: 13, 2013.

Secure Distributed Computations and their Proofs

www.msr-inria.fr/projects/secure-distributed-computations-and-their-proofs/

Summary

We develop formal tools for programming distributed systems with constructive security guarantees. Our goal is to enable programmers to express and prove high-level security and privacy properties with a reasonable amount of effort—sometimes automatically, sometimes with mechanical assistance—as part of the development process. These properties should hold in a hostile environment, under realistic (partial) trust assumptions on the principals and machines involved in the computation, and without the need of a centralized trusted third party.

Our research project involves three complementary lines of research on cryptographic protocol verification, distributed programming languages, and decentralized data privacy. We outline below selected results in each of these lines of research for 2011–2013.

See also http://www.msr-inria.fr/projects/secure-distributed-computations-and-their-proofs/

miTLS: a verified implementation of TLS 1.2

As regards cryptographic protocol verification, our main result is a large case study that builds on the tools previously-developed at Microsoft Research-Inria: we program and verify miTLS [1]: a *reference implementation* of the TLS Internet Standard, relying on the modular, type-based cryptographic proof techniques of Fournet et al. [4]. To our knowledge, this is the first verified cryptographic protocol implementation at this scale.
TLS is possibly the most used protocol for secure communications, and also the most studied, with a 20-year history of flaws and fixes, ranging from its protocol logic to its cryptographic design, and from the Internet standard to its diverse implementations. Surprisingly, practical security for TLS remains controversial, as illustrated by many serious attacks against its mainstream implementations over the last few years.

- miTLS is a new, full-fledged implementation of TLS 1.2. Our code supports its wire formats, ciphersuites, sessions and connections, re-handshakes and resumptions, alerts and errors, and data fragmentation, as prescribed in the RFCs; it interoperates with mainstream web browsers and servers.
- At the same time, our code is carefully structured to enable its modular, automated verification, from its main API down to computational (probabilistic, polynomial-time) assumptions on its sub-protocols and underlying cryptographic algorithms such as AES and RSA.
- Finally, our code provides a flexible research tool for trying out attacks against TLS and for prototyping variants; hence, miTLS supports four experimental protocol extensions.

Our implementation consists of 7,000 lines of F#, specified using 3,000 lines of F7. We present security specifications for its main components, such as authenticated stream encryption for the record layer and key establishment for the handshake. We describe their verification using the F7 type checker. To this end, we equip each cryptographic primitive and construction of TLS with a new typed interface that captures its security properties, and we gradually replace concrete implementations with ideal functionalities. We finally typecheck the protocol state machine, and obtain precise security theorems for TLS, as it is implemented and deployed. We show how to build simple secure applications on top of our new API. We also revisit classic attacks and report a few new ones.
For additional details, see https://www.mitls.org, running our own HTTPS server on top of miTLS.

F*: a language for secure distributed systems

As regards secure distributed programming, our main result is a new language that builds upon our experience of functional programming with advanced dependently-typed systems and their applications to distributed security. The resulting language, F*, provides strong formal guarantees, and is now the basis for most of our research on programming and verifying distributed systems. Its latest implementation is available at https://github.com/FStarLang/FStar.

A new dependently-typed research language [7, 8] Distributed applications are difficult to program reliably and securely. Dependently typed functional languages promise to prevent broad classes of errors and vulnerabilities, and to enable program verification to proceed side-by-side with development. However, as recursion, effects, and rich libraries are added, using types to reason about programs, specifications, and proofs becomes challenging.

We present F*, a full-fledged design and implementation of a new dependently typed language for secure distributed programming. Our language provides arbitrary recursion while maintaining a logically consistent core; it enables modular reasoning about state and other effects using affine types; and it supports proofs of refinement properties using a mixture of cryptographic evidence and logical proof terms. The key mechanism is a new kind system that tracks several sub-languages within F* and controls their interaction. F* subsumes two previous languages, F7 and Fine. We prove type soundness (with proofs mechanized in Coq) and logical consistency for F*.

We have implemented a compiler that translates F*to .NET bytecode. F* provides access to libraries for concurrency, networking, cryptography, and interoperability with C#, F#, and the other .NET languages. The compiler produces verifiable binaries with 60% code size overhead for proofs and types, as much as a 45x improvement over the Fine compiler, while still enabling efficient bytecode verification.

We have programmed and verified nearly 50,000 lines of F* including new schemes for multi-party sessions; a zero-knowledge privacy-preserving payment protocol; a provenance-aware curated database; a suite of web-browser extensions verified for authorization properties; a cloud-hosted multi-tier web application with a verified reference monitor; the core F* typechecker itself; and programs translated to F* from other languages such as F7 and JavaScript.

Taking this programming experience into account, we have recently revised our initial design for F* [8], significantly improving our type system and its presentation with a simpler kind system, a new proof of logical consistency, and a more complete formalization in Coq.

Self-Certification [6] Well-established dependently-typed languages such as Agda and Coq provide reliable ways to build and check formal proofs. Several other dependently-typed languages such as Aura, ATS, Cayenne, Epigram, F*, F7, Fine, Guru, PCML5, and Ur also explore reliable ways to develop and verify programs. All these languages shine in their own regard, but their implementations do not themselves enjoy the degree of safety provided by machine-checked verification.

We propose a general technique called *self-certification* that allows a typechecker for a suitably expressive language to be certified for correctness. We implement this new technique for F*: self-certification involves implementing a typechecker for F* in F*, while using all the conveniences F* provides for the compiler-writer (e.g., partiality, effects, implicit conversions, proof automation, libraries). This typechecker is given a specification (in F*) strong enough to ensure that it computes valid typing derivations. We obtain a typing derivation for the core typechecker by running it on itself, and we export it to Coq as a type-derivation certificate. By typechecking this derivation (in Coq) and applying the F* metatheory (also mechanized in Coq), we conclude that our type checker is correct. Once certified in this manner, the F* typechecker is emancipated from Coq.

Self-certification leads to an efficient certification scheme—we no longer depend on verifying certificates in Coq—as well as a more broadly applicable one. Hence, the self-certified F* checker is suitable for use in adversarial settings where Coq is not intended for use, such as run-time certification of mobile code.

Applications to JavaScript security [6] Many tools allow programmers to develop applications in high-level languages and deploy them in web browsers via compilation to JavaScript. While practical and widely used, these compilers are ad hoc. No guarantee is provided on their correctness for whole programs, nor their security for programs executed within arbitrary JavaScript contexts.

Relying on F*, we present a compiler with such positive guarantees. We compile an ML-like language

with higher-order functions and references down to JavaScript, while preserving all source program properties. We evaluate our compiler on sample programs, including a series of secure libraries.
We illustrate the dangers of JavaScript contexts with a series of attacks against naive scripts. We then give a semantics to JavaScript by translation to F*. Based on lambdaJS, this semantics reflects the full EcmaScript 5 standard, as well as our experimental findings on dangerous features in JavaScript implementations (implicit coercions, getters and setters, and special arguments, caller, and callee properties).
We present our compilation scheme, expressed as a type-preserving translation between fragments of F*: each source type is mapped to 'dyn', the type of Javascript values, refined with a logical specification of its target representation.
For whole programs, we show that the translation is a forward simulation. For programs executed in untrusted Javascript contexts, we wrap our translation with defensive filters to import and export values while preserving the translation invariant. Relying on type-based invariants and a new notion of applicative bisimilarity, we show full abstraction: two programs are equivalent in all source contexts if and only if their wrapped translations are equivalent in all Javascript contexts. Thus, programmers can produce and deploy JavaScript, and still rely on static scopes and types for reasoning about their programs.

ZQL: Privacy-Preserving Query Processing

As regards privacy, our main result is a "security compiler" for processing sensitive data without disclosing it [5]. In contrast with TLS, this work involves synthesizing custom protocols with novel cryptographic constructions.
ZQL is a query language for expressing simple computations on private data, at the same level of abstraction as SQL for databases. Its compiler produces code to certify data, perform client-side computations, and verify the correctness of their results. Under the hood, it synthesizes zero-knowledge protocols that guarantee both integrity of the query results and privacy for all other data.
We present the ZQL language, its compilation scheme down to concrete cryptography, and the security guarantees it provides. We report on a prototype compiler that produces F# and C++. We evaluate its performance on queries for privacy-preserving smart-meter billing, for pay-as-you-drive insurance policies, and for location-based services.
We also explore complementary privacy-friendly cryptographic mechanisms for verified distributed computations: additive shares for multiparty computation of aggregates [2], and succinct zero-knowledge arguments for bitcoins [3].

References

[1] Karthikeyan Bhargavan, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, and Pierre-Yves Strub. Implementing TLS with verified cryptographic security. In *S&P 2013*, pages 445–459. IEEE Computer Society, 2013.

[2] George Danezis, Cédric Fournet, Markulf Kohlweiss, and Santiago Zanella Béguelin. Smart meter aggregation via secret-sharing. In *SEGS@CCS*, pages 75–80. ACM, 2013.

[3] George Danezis, Cédric Fournet, Markulf Kohlweiss, and Bryan Parno. Pinocchio coin: building zerocoin from a succinct pairing-based proof system. In Martin Franz, Andreas Holzer, Rupak Majumdar, Bryan Parno, and Helmut Veith, editors, *PETShop@CCS*, pages 27–30. ACM, 2013.

[4] Cédric Fournet, Markulf Kohlweiss, and Pierre-Yves Strub. Modular code-based cryptographic verification. In *CCS 2011*, pages 341–350. ACM, 2011.

[5] Cédric Fournet, Markulf Kohlweiss, George Danezis, and Zhengqin Luo. ZQL: A compiler for privacy-preserving data processing. In Samuel T. King, editor, *USENIX Security*, pages 163–178. USENIX Association, 2013.

[6] Cédric Fournet, Nikhil Swamy, Juan Chen, Pierre-Evariste Dagand, Pierre-´ Yves Strub, and Benjamin Livshits. Fully abstract compilation to JavaScript. In Roberto Giacobazzi and Radhia Cousot, editors, *POPL*, pages 371–384. ACM, 2013.

[7] Nikhil Swamy, Juan Chen, Cédric Fournet, Pierre-Yves Strub, Karthikeyan Bhargavan, and Jean Yang. Secure distributed programming with value dependent types. In *ICFP 2011*, pages 266–278. ACM, 2011.

[8] Nikhil Swamy, Juan Chen, Cédric Fournet, Pierre-Yves Strub, Karthikeyan Bhargavan, and Jean Yang. Secure distributed programming with value dependent types. *J. Funct. Program.*, 23(4):402–451, 2013.

Tools and Methodologies for Formal Specifications and for Proofs

www.msr-inria.fr/projects/tools-for-proofs/

This project started in the summer of 2006.

Team

Project	Status	Last Name	First Name	Affiliation
TLA+	Team Leader	DOLIGEZ	Damien	Inria Paris-Rocquencourt
	Researcher	LAMPORT	Leslie	Microsoft Research Silicon Valley
	Researcher	MERZ	Stephan	Inria Nancy Grand-Est
	PHD Student	VANZETTO	Hernan Pablo	MSR-Inria Joint Centre
	Post doc	LIBAL	Tristan	MSR-Inria Joint Centre
	Intern	BHATT	Bhargav	MSR-Inria Joint Centre

Denis Cousineau left in November 2011. Hernán Vanzetto has been working since December 2010 as a PhD student on SMT-lib and TPTP translations. Tomer Libal started in October 2012. Bhargav Bhatt did a 3-month internship in the summer of 2013.

Research

Many system crashes are caused by what Jim Gray called “heisenbugs”: irreproducible errors whose cause is never discovered. We believe that many of those errors are due to errors in the algorithms used to synchronize concurrent activity. TLA$^+$ and its model checker are effective tools for designing correct concurrent algorithms. However, the huge number of possible states caused by the highly nondeterministic nature of these algorithms limits the ability to eliminate errors by conventional model checking. For some algorithms, nothing short of mathematical proof can guarantee correctness, and mechanical checking is required to avoid errors in proofs. Our goal in developing the TLA$^+$ prover is that it will enable engineers to avoid errors in highly critical synchronization algorithms.

TLA$^+$ is a specification and proof language based on temporal logic, where first-order logic and set theory are used to describe a set of states and the possible transitions between these states. TLA$^+$ also includes a module system for manipulating large-scale specifications.

There are a number of existing tools for working on TLA$^+$ specifications, the most important of which is the TLC model-checker. TLA$^+$ has already proved its worth in significant projects in hardware design (Alpha and Itanium processors), protocols (PCI-X), and software (Doligez-Leroy-Gonthier garbage collector).

In this project, we are defining an extension of the TLA$^+$ language, called TLA^{+2}, for writing mathematical proofs, and we are complementing the existing tools so that users can develop, debug, and check proofs about algorithm and system specifications. In this way, TLA$^+$ becomes a complete solution for writing, debugging, and proving specifications. More precisely, we are refining the proof language, building a development environment for TLA$^+$ specifications and proofs, developing and adapting automatic tools that help to prove TLA$^+$ theorems (based on the Zenon and veriT provers), and interfacing to generic automatic proof tools (SMT solvers interfaced through SMT-LIB, first-order provers interfaced through TPTP, temporal logic provers).

We validate and enhance our tools by finding examples of real-world projects where formal specifications bring real improvements over other methodologies. Feedback from these examples helps us to improve the proof language and the tools and develop methods and design patterns for using TLA$^+$.

Tools and software

The Eclipse-based ToolBox is an IDE (integrated development environment) that serves as the front-end for all TLA$^+$ tools. It is interfaced with TLC, SANY, the PlusCal translator, and the PM. It includes an editor with syntax coloring, provides for hiding and showing of subtrees of the proof as well as decomposing proof steps, and allows the user to check parts of the proof independently of each other, reporting the proof obligations that fail.

The non-interactive proof manager (PM) is a tool for checking proofs written in TLA^{+2}. The PM takes as input a TLA^{+2} specification and attempts to check some or all of its proofs (depending on command-line arguments) in batch mode. It reports if it fails to check any step of a proof, either because the step is incorrect or because there is insufficient detail in the proof for the PM to automatically check the proof.

Design

The PM consists of two main stages: a front-end elaborator and a verifier. The verifier stage interacts with automated theorem provers (Zenon, Isabelle, CVC3, Z3, Yices, veriT, etc.) by sending them proof obligations and optionally with a back-end framework (Isabelle again) by sending it the proof obligations and the proofs generated by Zenon. These proof obligations and proofs are expressed in an encoding of TLA^{+2} called Isabelle/TLA$^+$.

Elaboration. The front-end elaborator is the first stage of the PM and consists of resolving names and substitutions in the input specification. At the end of this stage, the input proofs are annotated with fully elaborated usable elements at every step.

The semantics of TLA^{+2} modules inlines instanced modules in their host module mediated by a substitution. A proof is allowed to use definitions and theorems from these module instances, so the PM must perform the substitutions and flatten all instances. For some TLA^{+2} operators such as ENABLED that quantify implicitly over certain sub-expressions (in particular, primed variables), substitution can only be performed if the quantification is made explicit; therefore, the PM eliminates these operators entirely.

In addition to resolving names and substitutions, further elaborations are carried out to reduce the supported TLA^{+2} language to a core elaborated form, called TLA^{+2e}. The primary benefit of this elaboration is to reduce the complexity of the trusted theorems base for the final certification stage. The following are some examples of elaboration:

- $\exists \langle x,y \rangle \in S : P(x,y)$ elaborates to $\exists x : \exists y : (\langle x,y \rangle \in S) \wedge P(x,y)$. All bounded quantifiers are similarly replaced with unbounded quantifiers over single variables.
- $[f \text{ EXCEPT } ![x_1] = e_1, ![x_2] = e_2]$ elaborates to $[[f \text{ EXCEPT } ![x_1] = e_1] \text{ EXCEPT } ![x_2] = e2]$. This reduces EXCEPT to the status of a binary operator.

This elaboration can generate fresh bound variables, so it is necessary to ensure that elaborating the same TLA^{+2} expression at two distinct points produces equal elaborated forms. The PM therefore internally represents the TLA^{+2} syntax using de Bruijn indexes, which is nameless and generalizes syntactic equality to α-equivalence. Names from the source syntax are maintained as “hints” that are used to produce named representations for output.

Verification. The second key stage of the proof manager is to extract proof obligations from the elaborated proof steps and interact with back-end provers and the back-end logical framework (Isabelle) to check that the obligations are true. Before discussing how verification is performed, we note that the PM is allowed to fail to verify a BY directive even when it is mathematically true; in this case the user must further refine the proof to make it accepted by the PM.

A proof obligation is a TLA^{+2} statement of the form ASSUME $e_1,\ldots,e_m$ PROVE g where $e_1,\ldots,e_m$ is the list of assumptions and usable facts at that point of the proof, and g is the current goal. Given the interpretation [[_]] of TLA^{+2} in a logical framework such as Isabelle, this proof obligation amounts to proving the sequent $[[e_1]],\ldots,[[e_m]] \vdash [[g]]$ in the target framework.

When Isabelle accepts the generated proof script as correct, we get high assurance that the theorem is true. We also get good assurance that the source proof is correct, although in theory the PM might turn an incorrect TLA^{+2} proof into a correct Isabelle proof. Note that proof checking by Isabelle is available only when all the obligations are proved by Zenon or Isabelle: the other back-ends do not currently produce checkable proofs. We are working on VeriT to make it produce such proofs and bring it into the high-assurance side of the system.

Implementation progress

Compared to the version of 2010, we have enhanced the treatment of modules by the PM, improved the back-end interface to SMT solvers, and (still experimental) an interface to generic first-order provers through the TPTP standard. We have refined the design of the temporal part of the proof language, and we have started work on a back-end interface to the LS4 temporal-logic prover, which will allow us to prove non-invariant properties such as liveness and fairness.

The PM, together with Zenon and the Isabelle/TLA+ theory, form a unit called TLAPS. In 2012, we released two versions of TLAPS, synchronized with two releases of the ToolBox, and in 2013 we released a third version of TLAPS.

The current version of TLAPS includes a new translation of proof obligations into SMT-lib format, which can handle (in principle) all TLA^{+2} proof obligations. It is especially useful for obligations that involve a mix of arithmetic, set theory, and functions, since the other back-end provers do not deal with arithmetic. Because SMT-lib is a multi-sorted language and overloaded symbols such as equality should be resolved differently depending on the sorts of their arguments, this translation requires a form of sort inference for a given proof obligation. When sort inference fails, we fall back on a generic translation that can handle all TLA^{+2} formulas.

TLAPS was presented in a paper at FM 2012 [105], and the SMT-lib translation at LPAR-18 [104], PxTP 2011 [106], and AVoCS 2012 [108].

Interfaces with Isabelle and Zenon

TLA^{+2} encoding in Isabelle The core of TLA^{+2} is encoded as an object logic in the generic interactive proof assistant Isabelle. More precisely, the encoding includes propositional and first-order logic, elementary set theory, functions, fixed-point constructions, and the construction of the natural numbers. The main automatic proof methods (such as rewriting, case-based reasoning, and the tableau prover) available in Isabelle have been instantiated for TLA^{+2}. This encoding provides the basis for the verification stage of the proof manager: it receives the proof obligations and attempts to discharge them based on the semi-automated proof methods available in Isabelle. We have chosen to define a new object

logic rather than encode TLA^{+2} in one of the existing logics (such as HOL or ZF); this minimizes the overhead of the translation and makes it easier to understand the error messages when Isabelle fails to prove formulas.

We have implemented support for strings, records, tuples, and natural numbers. Integer arithmetic has been defined in Isabelle, but little automation is currently available to users of TLAPS (we rely on SMT solvers for effectively discharging proof obligations involving arithmetic); support for reals and for the temporal logic TLA is still missing. Nevertheless, the current fragment can already be used to prove the correctness of some complex algorithms.

Zenon The Zenon prover produces proofs in Isabelle format (as Isar scripts) and uses TLA-specific inference rules to make it more efficient on the kind of proof obligations that are produced when checking a TLA^{+2} proof.

PlusCal

PlusCal is a high-level language for describing algorithms. The PlusCal compiler generates a TLA^{+} model from a PlusCal algorithm, and many of our case studies are based on models that originate from algorithms written in PlusCal. The language was presented at ETH Zurich [110] and at CEA [113]. Versions 1.5 and 1.6 were released in 2011, introducing keywords for specifying fairness and a modification that simplifies the translation to TLA^{+} by removing the variable representing the control state when it is not needed. This simplifies correctness proofs of distributed algorithms written in PlusCal, such as Byzantine Paxos.

Publications & talks from the project team members

Journal papers and book chapters

[99] Leslie Lamport. Euclid Writes an Algorithm: A Fairytale. *International Journal of Software and Informatics* 5, 1-2 (2011) Part 1, 7-20.

[100] Leslie Lamport. How to Write a 21st Century Proof. *Journal of Fixed Point Theory and Applications* doi:10.1007/s11784-012-0071-6 (6 March 2012).

5.4.2 Conference and workshop papers

[101] Leslie Lamport. Byzantizing Paxos by Refinement. *Distributed Computing: 25th International Symposium: DISC 2011*, David Peleg, editor. Springer (2011) 211-224.

[102] Tianxiang Lu, Stephan Merz, and Christoph Weidenbach, Towards Verification of the Pastry Protocol Using TLA+, *13th IFIP WG 6.1 Intl. Conf. Formal Techniques for Distributed Systems (FORTE 2011)*. Springer, LNCS 6722, pp. 244-258.

[103] Bernadette Charron-Bost, Henri Debrat, and Stephan Merz, Formal Verification of Consensus Algorithms Tolerating Malicious Faults, *13th Intl. Symp. Stabilization, Safety, and Security of Distributed Systems (SSS 2011)*. Springer, LNCS 6976, pp. 120-134.

[104] Stephan Merz and Hernán Vanzetto, Automatic Verification of TLA+ Proof Obligations With SMT Solvers, *18th Intl. Conf. Logic for Programming, Artificial Intelligence and Reasoning (LPAR-18)*. Springer, LNCS 7180, pp. 289-303 (2012).

[105] Denis Cousineau, Damien Doligez, Leslie Lamport, Stephan Merz, Daniel Ricketts, Hernán Vanzetto, TLA+ Proofs, *18th Intl. Symp. Formal Methods (FM 2012)*. Springer, LNCS 7436, pp. 147-154.

[106] Stephan Merz and Hernán Vanzetto, Towards certification of TLA+ proof obligations with SMT solvers. *First International Workshop on Proof eXchange for Theorem Proving (PxTP)*, August 2011

[107] Tianxiang Lu, Stephan Merz, Christoph Weidenbach, Formal Verification Of Pastry Using TLA+. *International Workshop on the TLA+ Method and Tools*, August 2012.

[108] Stephan Merz and Hernán Vanzetto, Harnessing SMT Solvers for TLA+ Proofs. *12th International Workshop on Automated Verification of Critical Systems (AVoCS)*, September 2012.

5.2.3 Talks and other communications

[109] Leslie Lamport. Why We Should Build Software Like We Build Houses?, *Wired Web site*, January 2013

[110] Leslie Lamport. The PlusCal Algorithm Language. Computer Science Colloquium at ETH Zürich, May 2013.

[111] Leslie Lamport. Programming Languages are not the Answer. Invited talk at Workshop on Languages for Distributed Algorithms (LADA), January 2012.

[112] Leslie lamport, Who Builds a House without Drawing Blueprints? Keynote talk at *14th International Conference on Distributed Computing and Networking (ICDCN)*, Mumbai, January 2013.

[113] Leslie Lamport, Specifying Systems with Mathematics: the TLA+ Language and Tools. Talk at CEA, June 2013.

Dynamic Dictionary of Mathematical Functions

www.msr-inria.fr/projects/dynamic-dictionary-of-mathematical-functions/

This project started in Fall 2007.

Overview

``Our ambition with the DDMF is to develop an authoritative interactive web site on the special functions of mathematics."

A great deal of functions from mathematical analysis are involved in a recurrent manner in diverse domains of applied mathematics. Their properties and the mathematical identities that they satisfy have been considerably studied and documented in classical works since the 19th century. These properties and mathematical identities have recently become amenable to computer algebra, the branch of computer science that concerns itself with exact and efficient computations with general mathematical objects. Therefore, it has become natural to ask for generating books on special functions, rather than compiling them from diverse sources. In view of this, our goal is to make the results of computations with the special functions of mathematics available to an audience that is not expert in computer algebra. To this end, we provide users with a dynamical presentation of them on the web, in the form of an online encyclopedic dictionary (http://ddmf.msr-inria.inria.fr). This can be viewed as a modern version of the textbooks and handbooks of the 19th and 20th centuries. For each function, our current encyclopedia shows its essential properties and mathematical objects attached to it, which are often infinite in nature (numerical evaluations, asymptotic expansions). This way of disseminating has the advantage of allowing for a presentation that adapts interactively to the user's actual needs.

Team

Team leader	CHYZAK	Frédéric	Inria Saclay-île de France
Researcher	BOSTAN	Alin	Inria Saclay-île de France
Researcher	SALVY	Bruno	Inria Grenoble-Rhône-Alpes
Senior visitor	DAVENPORT	James	University of Bath
Post-doc	KOUTSCHAN	Christoph	Microsoft Research--Inria Joint Centre
Post-doc	STAN	Flavia	Inria Paris-Rocquencourt
PhD student	BENOIT	Alexandre	Ecole Polytechnique
PhD student	MEZZAROBBA	Marc	Ecole Polytechnique

Marc Mezzarobba defended his PhD in 2011. He did a post-doc in the Inria AriC team (LIP, Lyon), then at the Johannes Kepler University (Linz, Austria), and is now a CNRS researcher at the Université Pierre et Marie Curie (Paris).

Alexandre Benoit defended his PhD in 2012. He did a post-doc at the Université Pierre et Marie Curie (Paris). He is now a mathematics teacher.

Flavia Stan started was a post-doc with us from February 2011 to May 2012. She is now a research assistant at the Technical University of Kaiserslautern.

Christoph Koutschan was a post-doc from April 2011 to August 2012. He is now a researcher at the Johannes Kepler University (Linz, Austria).

Professor James Davenport visited for a month (April 2011) from the University of Bath.

Highlight

The most recent changes to our DDMF web site are the introduction of special functions with parameters, a typical example of which is the Bessel function $J_\nu(x)$: a function of x with parameter ν.

From the linear differential equation that defines the family, which relates derivatives with respect to x and is parameterized by ν, the system determines local and asymptotic expansions that are parametrised by ν as well. Better yet, it obtains explicit expressions for the coefficients of these expansions, in terms of the Γ function. The display of the resulting formulae has also been beautified a lot by factoring and presenting the coefficients in a pleasing way. By interacting with the page, the readers can instantiate the parameters of special functions: here, setting $\nu=3$ at the bottom of the page will result in a new page dedicated to the more specific function $J_3(x)$. This new page shows added information that was more so meaningful for the parametrised function: plots, a Laplace transform, etc. Also, comparing the pages obtained when ν is or is not an integer, the readers will discover that expansions (for instance, the Taylor series at 0) are not just specializations of the same formula for different values of ν. Indeed, there is no continuous formula with respect to ν and the complete algorithm needs to be re-run for each new value for ν. This shows the interest of the interactivity of DDMF.

Research

The originality of our work in computer algebra lies in the systematic use of linear operators as a data-structure from which various informations such as identities for special functions can be extracted. Our work proceeds along three lines: design of fast algorithms either based on using linear operators or giving better complexity for operations on these operators; new algorithmic applications of linear operators; new algorithms extending the class of functions to which our methods apply.

Special Functions

Interactive Online Dictionary

Our web site DDMF, available from http://ddmf.msr-inria.inria.fr/ , (see highlight above) consists of interactive tables of mathematical formulas on elementary and special functions. The formulas are automatically generated by computer algebra routines. The user can ask for more terms of the expansions, more digits of the numerical values, or proofs of some of the formulas. Several releases have occurred over the period and have introduced: correct constants in general terms of asymptotic expansions, pages for parametrised differential equations, special functions depending on parameters, nicer display/factorisation of mathematical formulas, improved closed-form solving, proofs related to Taylor polynomial approximations. The current release is 1.9 (April 2013). The code of DDMF is now publically available from the DDMF web site. A by-product of the work on DDMF is DynaMoW~\cite{Chyzak-2011-UCP}, a programming tool for controlling the generation of mathematical web sites that embed dynamical mathematical contents generated by computer-algebra calculations. DynaMoW was released in September 2011 and is available from http://ddmf.msr-inria.inria.fr/DynaMoW/.

Approximation Series and Guaranteed Numerical Calculations

In computer algebra there are different ways of approaching the mathematical concept of functions. We advocate a recent one which consists in defining them as solutions of differential equations. The question of how much extra information needs to be given in order to specify a function is crucial. Obviously, initial conditions are needed and in~\cite{ChyzakDavenportKoutschanSalvy2011} we discussed the extent to which the treatment of branch cuts can be rendered (more) algorithmic, by adapting Kahan's rules to the differential equation setting.

The interplay between symbolic and numerical evaluation has been brought to fruit in Marc Mezzarobba's PhD~thesis~\cite{Mezzarobba2011}, where he showed how arbitrary precision can be reached in good

complexity for the evaluation of the whole class of solutions of linear differential equations. Moreover, he developed an automatic computation of the analytic continuation of such a function, with controlled error and up to arbitrary precision given by the user.

In the case when the domain of interest is a segment of the real axis, the Taylor series is not the best way to approximate a function and Chebyshev series behave much better. In his PhD~thesis~\cite{Benoit2012}, Alexandre Benoit showed that the coefficients of these series can be computed efficiently, both symbolically and numerically, overcoming intrinsic numerical instability in the process.

Summation and Integration

Creative Telescoping for Bivariate Hyperexponential Functions

In \cite{BostanChenChyzakLiXin-2013-HRC}, we gave a new algorithm for the symbolic integration of bivariate hyperexponential functions, which outperforms state-of-the-art implementations like Maple's function {\sf{DEtools[Zeilberger]}}. The approach was to extend Hermite's reduction for rational functions and the Hermite-like reduction for hyperexponential functions in a suitable way. A key feature of the algorithm is that it can avoid the costly computation of certificates.

Existence of telescopers for hyperexponential-hypergeometric sequences

The termination of creative-telescoping algorithms for the symbolic summation of hypergeometric terms is typically not ensured in general, but requires the existence of so-called telescopers. In \cite{ChenChyzakFengLi-2013-ETM}, we presented a criterion for the existence of telescopers for mixed hypergeometric terms. The criterion enables us to determine the termination of Zeilberger's algorithms for mixed hypergeometric inputs with very simple calculations, so as to avoid calling it when it is bound to run forever. Both previous results are elaborations on earlier works during Chen's doctorate \cite{Chen2011} in our group.

Creative Telescoping for Rational Functions

In~\cite{BostanLairezSalvy-2013-CTGD} we described a precise and elementary algorithmic version of the Griffiths--Dwork method for the creative telescoping of rational functions. This leads to bounds on the order and degree of the coefficients of the differential equation, and to the first complexity result which is single exponential in the number of variables. One of the important features of the algorithm is that it does not need to compute certificates. The approach is vindicated by a prototype implementation.

Applications

Application to Combinatorics

Explicit formula for the generating series of diagonal 3D rook paths

Let a_n denote the number of ways in which a chess rook can move from a corner cell to the opposite corner cell of an $n \times n \times n$ three-dimensional chessboard, assuming that the piece moves closer to the goal cell at each step. We described in \cite{Bostan-2011-EFG} the computer-driven discovery and proof of the fact that the generating series $G(x) = \sum_{n\geq 0} a_n x^n$ admits the following explicit expression in terms of a Gaussian hypergeometric function:

$$G(x) = 1 + 6\int_0^x \frac{{}_2F_1\left(\frac{1}{3}\ \frac{2}{3}\ \middle|\ \frac{27w(2-3w)}{(1-4w)^3}\right)}{(1-4w)(1-64w)}dw$$

Non-D-finite excursions in the quarter plane

The number of excursions (finite paths starting and ending at the origin) having a given number of steps

and obeying various geometric constraints is a classical topic of combinatorics and probability theory.

We proved in \cite{BostanRaschelSalvy2012} that the sequence $(e_n^S)_{n\geq 0}$ of numbers of excursions in the quarter plane corresponding to a nonsingular step set $S \subseteq \{0,\pm 1\}^2$ with infinite group does not satisfy any nontrivial linear recurrence with polynomial coefficients. Accordingly, in those cases, the trivariate generating function of the numbers of walks with given length and prescribed ending point is not D-finite. This solves an open problem in the field of lattice path combinatorics.

Newton for species

Many combinatorial models are defined recursively by combinatorial equations (using unions, cartesian products, sets,\dots). We have considered systems of recursively defined combinatorial structures. We have given algorithms checking that these systems are well founded, computing generating series and providing numerical values. Our framework is an articulation of the constructible classes of Flajolet and Sedgewick with Joyal's species theory. At the heart of the method, a quadratic iterative Newton method is shown to solve well-founded systems combinatorially. From there, truncations of the corresponding generating series are obtained in quasi-optimal complexity. This iteration transfers to a numerical scheme that converges unconditionally to the values of the generating series inside their disk of convergence. These results provide important subroutines in random generation. Finally, the approach has been extended to combinatorial differential systems that permit the definition of heap-ordered structures~\cite{PivoteauSalvySoria2012}.

Application to Physics

Physical problems whose modeling involves special-function integrals comprise the study of models of statistical mechanics, like the \emph{Ising model\/} for ferro-magnetism. Using computer algebra algorithms, we showed in~\cite{BBHHMWZ10} that almost all the linear differential operators factors obtained in the analysis of the n-particle contribution of the susceptibility of the Ising model for $n \leq 6$, are operators ``associated with elliptic curves". Then, we showed in~\cite{BoBoChHaMa13} that the $n$-fold integrals $\chi^{(n)}$ of the magnetic susceptibility of the Ising model, as well as various other $n$-fold integrals of the ``Ising class", or n-fold integrals from enumerative combinatorics, like lattice Green functions, correspond to a distinguished class of functions generalizing algebraic functions: they are actually *diagonals of rational functions*. This algebraic structure explains many remarkable properties of the integrals of the Ising class.

Fast Algorithms

Simultaneous modular reduction

We have given algorithms to perform modular polynomial multiplication or modular dot product efficiently in a single machine word. This is achieved by a combination of techniques. Polynomials are packed into integers by Kronecker substitution; several modular operations are performed at once with machine integer or floating point arithmetic; normalization of modular images is avoided when possible; some conversions back to polynomial coefficients are avoided; the coefficients are recovered efficiently by preparing them before conversion. We have discussed precisely the required control on sizes and degrees. Applications are polynomial multiplication, prime field linear algebra and small extension field arithmetic, where these techniques lead to practical gains of quite large constant factors~\cite{DumasFousseSalvy2011}.

Quasi-optimal multiplication of linear differential operators

The product of \emph{polynomials\/} is one of the most basic operations in mathematics, and the study of its computational complexity is central in computer science. Linear differential operators are algebraic objects that encode linear differential equations, and form a non-commutative ring that shares many properties with the commutative ring of usual polynomials. Yet, the algorithmic study of linear differential operators was until very recently much less advanced than in the polynomial case. To fill this gap, we showed in \cite{BeBoHo12} that linear differential operators with polynomial coefficients over a field of characteristic zero can be multiplied in quasi-optimal time, by an extension of the Fast Fourier Transform.

Power Series Solutions of Singular (q)-Differential Equations

We provided in~\cite{BostanChowdhuryLebretonSalvySchost2012} algorithms computing power series solutions of a large class of differential or q-differential equations or systems. Their number of arithmetic operations grows linearly with the precision, up to logarithmic terms. This extends and improves classical results by Brent and Kung (1978) on fast algorithms for power series.

Fast lclm

We studied in~\cite{BostanChyzakLiSalvy2012} tight bounds and fast algorithms for LCLMs of \emph{several\/} linear differential operators with polynomial coefficients. We analyzed the worst-case arithmetic complexity of existing algorithms for LCLMs, as well as the size of their outputs. We proposed a new algorithm that reduces the LCLM computation to a linear algebra problem on a polynomial matrix. The new algorithm yields sharp bounds on the coefficient degrees of the LCLM, improving by two orders of magnitude the previously known bounds. The complexity of the new algorithm is almost optimal, in the sense that it nearly matches the arithmetic size of the output.

Uncoupling

Uncoupling algorithms transform a linear differential system of first order into one or several scalar differential equations. We examined in~\cite{BostanChyzakPanafieu-2013-CET} two approaches to uncoupling: the cyclic-vector method (CVM) and the Danilevski-Barkatou-Zürcher algorithm (DBZ). We gave tight size bounds on the scalar equations produced by CVM, and designed a fast variant of CVM whose complexity is quasi-optimal with respect to the output size. We exhibited a strong structural link between CVM and DBZ enabling to show that, in the generic case, DBZ has polynomial complexity and that it produces a single equation, strongly related to the output of CVM. We proved that algorithm CVM is faster than DBZ by almost two orders of magnitude, and provided experimental results that validate the theoretical complexity analyses.

Homotopy methods for multiplication

We studied in \cite{BoChHoSc10} the cost of multiplication modulo triangular families of polynomials. Following previous work by Li, Moreno Maza and Schost, we proposed an algorithm that relies on homotopy and fast evaluation-interpolation techniques. We thus obtained a quasi-linear time complexity for substantial families of examples, for which no such result was known before. We gave applications to notably addition of algebraic numbers in small characteristic.

Solving Quadratic Boolean Systems

A fundamental problem in computer science is to find all the common zeroes of m quadratic polynomials in n unknowns over F_2. The cryptanalysis of several modern ciphers reduces to this problem. Up to now, the best complexity bound was reached by an exhaustive search in $4\log_2 n\, 2^n$ operations. We have given an algorithm that reduces the problem to a combination of exhaustive search and sparse linear algebra. This algorithm has several variants depending on the method used for the linear algebra step. Under precise algebraic assumptions, we have showed that the deterministic variant of our algorithm has complexity bounded by $O(2^{0.841n})$ when $m = n$, while a probabilistic variant of the Las Vegas type has expected complexity $O(2^{0.792n})$. Experiments on random systems show that the algebraic assumptions are satisfied with probability very close to 1. We have also given a rough estimate for the actual threshold between our method and exhaustive search, which is as low as 200, and thus very relevant for cryptographic applications~\cite{BardetFaugereSalvySpaenlehauer2013}.

Publications and Talks

Journal papers and book chapters

[1] Magali Bardet, Jean-Charles Faugère, Bruno Salvy, and Pierre-Jean Spaenlehauer.
On the complexity of solving quadratic boolean systems. Journal of Complexity, 29(1):53-75, February 2013.
[2] A. Bostan, S. Boukraa, G. Christol, S. Hassani, and J.-M. Maillard. Ising n-fold integrals as diagonals of rational functions and integrality of series expansions. J. Phys. A, 46(18):185202, 44, 2013.
[3] A. Bostan, S. Boukraa, S. Hassani, M. van Hoeij, J.-M. Maillard, J.-A. Weil, and N. Zenine. The Ising model: from elliptic curves to modular forms and Calabi-Yau equations. Journal of Physics A: Mathematical and Theoretical, 44(4):44pp, 2011.
[4] A. Bostan, M.F.I. Chowdhury, J. van der Hoeven, and É. Schost. Homotopy methods for multiplication modulo triangular sets. Journal of Symbolic Computation, 46(12):1378-1402, 2011.
[5] Alin Bostan, Frédéric Chyzak, Mark van Hoeij, and Lucien Pech. Explicit formula for the generating series of diagonal 3D rook paths. Sém. Loth. Comb., B66a, 2011. 27 pages.
[6] Alin Bostan and Thierry Combot. A binomial-like matrix equation. American Mathematical Monthly, 119(7):593-597, 2012.
[7] Alin Bostan, Kilian Raschel, and Bruno Salvy. Non-D-finite excursions in the quarter plane. Journal of Combinatorial Theory, Series A, 2013. To appear.
[8] Jean-Guillaume Dumas, Laurent Fousse, and Bruno Salvy. Simultaneous modular reduction and Kronecker substitution for small finite fields. Journal of Symbolic Computation, 46(7):823-840, July 2011.
[9] Carine Pivoteau, Bruno Salvy, and Michèle Soria. Algorithms for combinatorial structures: Well-founded systems and Newton iterations. Journal of Combinatorial Theory, Series A, 119:1711-1773, 2012. 62 pages.

Conference and workshop papers

[10] Alexandre Benoit, Alin Bostan, and Joris van der Hoeven. Quasi-optimal multiplication of linear differential operators. In FOCS'12 (53rd Annual Symposium on Foundations of Computer Science), pages 524-530. IEEE, 2012.
[11] Alin Bostan, Shaoshi Chen, Frédéric Chyzak, Ziming Li, and Guoce Xin. Hermite reduction and creative telescoping for hyperexponential functions. In ISSAC 2013, pages 77-84. ACM, New York, 2013.
[12] Alin Bostan, Muhammad F. I. Chowdhury, Romain Lebreton, Bruno Salvy, and Éric Schost. Power series solutions of singular (q)-differential equations. In Mark van Hoeij and Joris Van der Hoeven, editors, ISSAC '12: Proceedings of the twenty-fifth International Symposium on Symbolic and Algebraic Computation, pages 107-114, 2012.
[13] Alin Bostan, Frédéric Chyzak, Ziming Li, and Bruno Salvy. Fast computation of common left multiples of linear ordinary differential operators. In Mark van Hoeij and Joris van der Hoeven, editors, ISSAC '12: Proceedings of the twenty-fifth International Symposium on Symbolic and Algebraic Computation, pages 99-106, 2012.
[14] Alin Bostan, Frédéric Chyzak, and Élie de Panafieu. Complexity estimates for two uncoupling algorithms. In ISSAC 2013, pages 85-92. ACM, New York, 2013.
[15] Alin Bostan, Bruno Salvy, and Pierre Lairez. Creative telescoping for rational functions using the Griffiths-Dwork method. In ISSAC 2013, pages 93-100. ACM, New York, 2013.
[16] Frédéric Chyzak and Alexis Darrasse. Using Camlp4 for presenting dynamic mathematics on the web: Dynamow, an ocaml language extension for the run-time generation of mathematical contents and their presentation on the web. In Olivier Danvy, editor, ICFP'11 (September 19-21, 2011, Tokyo, Japan), page 259-265. ACM, 2011. (An experiment report.).
[17] Frédéric Chyzak, James H. Davenport, Christoph Koutschan, and Bruno Salvy. On Kahan's rules for determining branch cuts. In SYNASC, page 47-51, September 2011. 13th International Symposium on Symbolic and Numeric Algorithms for Scienti_c Computing. September 26-29, 2011. Timisoara, Romania.

Theses

[18] Alexandre Benoit. Algorithmique semi-numérique rapide des séries de Tchebychev. PhD thesis, École polytechnique, Palaiseau, France, July 2012.
[19] Shaoshi Chen. Quelques applications de l'algèbre différentielle et aux différences pour le télescopage créatif. PhD thesis, École polytechnique (Palaiseau, France), February 2011. Defended on February 16, 2011.
[20] Frédéric Chyzak. The ABCs of Creative Telescoping Algorithms, Bounds, Complexity. Habilitation memoir, Université d'Orsay, 2013.
[21] Marc Mezzarobba. Autour de l'évaluation numérique des fonctions D-finies. PhD thesis, École polytechnique, Palaiseau, France, November 2011.

Posters, Lecture Notes, etc

[22] Alin Bostan, Frédéric Chyzak, Ziming Li, and Bruno Salvy. Fast computation of common left multiples of linear ordinary di_erential operators. Poster at ISSAC 2011, July 2011.
[23] Shaoshi Chen, Frédéric Chyzak, Ruyong Feng, and Ziming Li. On the existence of telescopers for hyperexponential-hypergeometric sequences. 21 pp. Submitted, 2013.
[24] Frédéric Chyzak. Creative telescoping for parametrised integration and summation. In Journées Nationales de Calcul Formel 2011, volume 2 of Les Cours du CIRM, 2012. 37 pages. Lecture notes.

Adaptive Combinatorial Search for e-Sciences

http://www.msr-inria.fr/projects/adaptative-combinatorial-search-for-e-science/

The project started on fall 2007. This report covers the 3-years period 2011-2013.

Team

Status	Last name	First name	Affiliation
Team leader	Hamadi	Youssef	Microsoft Research Cambridge
Team leader	Schoenauer	Marc	Inria Saclay Île-de-France
Researcher	Auger	Anne	Inria Saclay Île-de-France
Researcher	Hansen	Nikolaus	Inria Saclay Île-de-France
Researcher	Sebag	Michèle	CNRS
PhD student	Arbelaez	Alejandro	Microsoft Research-Inria Joint Center (until Spring 2011)
PhD student	Fialho	Alvaro	Microsoft Research-Inria Joint Center (until Spring 2011)
Post-doc	Loth	Manuel	Microsoft Research-Inria Joint Center (Oct 2011-Sept. 2013)
Post-doc	Lazaar	Nadjib	Microsoft Research-Inria Joint Center (Oct 2011-Sept. 2012)
PhD student	Bouzarjouna	Zyed	Inria Saclay Île-de-France
PhD student	Loshchilov	Ilya	Inria Saclay Île-de-France

Research

Many forefront techniques in both Stochastic and Combinatorial Search have been very successful in solving difficult real-world problems. However, their application to newly encountered problems, or even to new instances of known problems, remains a challenge, even for experienced researchers of the field - not to mention newcomers, be they skilled scientists or engineers from other areas. Theory and/or practical tools are still missing to make them 'Crossing the Chasm' (from Geoffrey A. Moore's 1991 book about the Diffusion of Innovation). The difficulties faced by the users arise mainly from the significant range of algorithm and/or parameter choices involved when using this type of approaches, and the lack of guidance as to how to proceed for selecting them. Moreover, state-of-the-art approaches for real-world problems tend to represent bespoke problem-specific methods which are expensive to develop and maintain.

Three different research paths have been explored during the first phase (2007-2010) of the Adapt project: *Adaptive Operator Selection* was tackled by Alvaro Fialho during his PhD, *Autonomous Search* was the topic of Alejandro Arbelaez's PhD, while *Adaptive Continuous Stochastic Search* was and still is the main domain of research activities of Anne Auger (Inria researcher) and Nikolaus Hansen (now Inria researcher, but who joined the group as Microsoft Research-Inria post-doc in the Adapt project).

- **Adaptive Operator Selection** (AOS) is concerned with the on-line adaptation of the mechanism that chooses among the different variation operators in Evolutionary Algorithms, a series of approaches have been proposed during Alvaro Fialho's PhD, defended in December 2010. All are based on the original idea of using Multi-Armed Bandit (MAB) algorithms: each operator is viewed as one arm of a MAB problem, and the main issue is then to choose a reward for each arm pulled/operator applied. Several papers witnessed the progresses made along Alvaro's PhD, and the final and most robust version uses a rank-based reward inspired by the AUC comparative

measure in supervised Machine Learning.

- **Autonomous Search** is a particular case of adaptive systems that aims at improving its solving performance by adapting itself to the problem at hand. This formalism is expressed by some computation rules between computation states. The sequence of application of these rules (i.e., the strategy) characterizes the search process itself. Alejandro Arbelaez's PhD came with the notion of Continuous Search (CS) for Constraint Programming. CS runs in two modes: the functioning mode solves the user's problem instances using the current heuristics model; the exploration mode reuses these instances to train and improve the heuristics model through Machine Learning during the computer idle time. Contrasting with previous approaches, Continuous Search thus does not require that the representative instances needed to train a good heuristics model be available beforehand. It achieves lifelong learning, gradually becoming an expert on the user's problem instance distribution.

After Alvaro's and Alejandro's departures, their work was continued in the framework of Constraint Programming and SAT problem solving, with the hiring of 2 post-docs in September 2011.

- Manuel Loth used Multi-Armed Bandit exploration of the search tree of the generic CP solver *Gecode* to improve the efficiency of the search without requiring a priori domain knowledge about the values to be given to the current variable. At every node of the form (variable=value), a MAB Hybridizing MAB and CP allowed to obtain state-of-the-art results in the domain of Job Shop Scheduling, results that could only be obtained until now by using domain-specific method, while using only a generic methodology. The first results were published at the LION'7 conference [XXX], and the fully convincing proof-of-concept results at CP'2013. Furthermore, following Manuel's departure after his two years of post-doc, his work lead to proposing a PhD for the next period of the Adapt project.

- Within the framework of Parallel SAT solving, Nadjib Laazar used Multi-Armed Bandit techniques to avoid a full broadcasting of clauses between the different nodes. However, the latter work, even though good results were obtained for some configurations, does not seem to lead to any generic procedure for automatic large-scale parallelization: Nadjib's post-doc hence only lasted one year.

In the meantime, research related to **Adaptive Continuous Stochastic Search** was continued around the well-known Covariance Matrix Adaptation Evolution Strategy (CMA-ES) algorithm, the flagship algorithm originally proposed by Nikolaus Hansen, that adapts the covariance matrix of the Gaussian mutation of an Evolution Strategy based on the path followed by the evolution. A brief description of the bases of CMA-ES and many of the recent developments can be found in the HDR dissertation of Nikolaus Hansen, defended during the course of the *Adapt* project in February 2010. Beside yet another series of improvements of the basic algorithm XXX, the benchmarking activities around the COCO platform and the BBOB workshops were extended, and a new axis targeted at handling expensive objective functions was launched and well explored.

- The **benchmarking effort** around COCO (http://coco.gforge.inria.fr/). The API has been completely rewritten, allowing users to easily provide their own optimizer but also their own benchmark functions. Though these user-provided bricks can be written in almost any language, the core of COCO (API with optimizers and functions, post-processing with improved plotting functionalities) is now entirely written in Python. Furthermore, the extension of COCO toward constrained problems and multi-objective context has started, within the NumBBO ANR project, that also involves colleagues from the Statistical Department of Dortmund University. Finally, following the 2009 and 2010 BBOB workshops (Black-Box Optimization Benchmarking), two new workshops have been organized, in 2012 and 2013, gathering still more entries in the comparative COCO framework.

- **Surrogate-based Optimization** (also known as Surface Response Method by our colleagues from Mechanical Engineering) tackles very expensive objective/fitness function by building an approximation of that objective function through a regression process that uses as examples the exact values that have been computed during the optimization steps done so far. That model is then used to help the next steps of the optimization algorithm, at a much lower cost than when computing the exact objective. Of course, such surrogate model should be adapted as the search moves to new areas of the search space. Two PhD students worked on that topic, improving the flagship algorithm CMA-ES with surrogate modeling. Both approaches, though using different methods to build the surrogate models, make use of the invariance properties of CMA-ES, source of its wide applicability and of its efficiency.
  Zyed Bouzarkouna, co-supervised by Anne Auger and Marc Schoenauer, explored local quadratic models that are built at each step from the best available data points. A local criterion comparing the ranks given by the true objective function and by the surrogate models is used to adapt the number of points to be evaluated with the (expensive) true function. The PhD was funded by IFP-EN, and its application part was devoted to the optimization of well placement to maximize the oil income, using some heavy 3D numerical simulation of the whole oil field to compute the objective function.
  Ilya Loshchilov uses Support Vector Machines to build one model that is then used in lieu of the true objective function in the original CMA-ES. The hyper-parameters of the SVM building, as well as some parameters of the use of the SVM model within CMA-ES (starting with the number of generations before another SVM model should be learnt) are adapted all along the optimization.

Publications & talks from the project team members (2011-2013)

Books and book chapters

[1] L. Bordeaux, Y. Hamadi, and P. Kohli, editors. *Tractability: Practical Approaches to Hard Problems*, ISBN 978-1-107-02519-6, Cambridge University Press 2013.

[2] Y. Hamadi. *Combinatorial Search: From Algorithms to Systems*, ISBN 978-3-642-41481-7, Springer 2013.

[3] Y. Hamadi, F. Saubion, and E. Monfroy, editors. *Autonomous Search*, ISBN 978-3-642-21433-2, Springer 2012.

[4] Nikolaus Hansen and Anne Auger. Principled Design of Continuous Stochastic Search: From Theory to Practice. In Y. Borenstein and A. Moraglio, editors, *Theory and Principled Methods for the Design of Metaheuristics*, Natural Computing Series. Springer, 2013.

[5] Jin-Kao Hao, Pierrick Legrand, Pierre Collet, Nicolas Monmarché, Evelyne Lutton, and Marc Schoenauer. *Artificial Evolution 2011: 10th International Conference*, Angers, France, October 24-26, 2011, Revised Selected Papers, volume 7401 of Lecture Notes in Computer Science. Springer, November 2012.

Journal papers

[6] Anne Auger, Johannes Bader, Dimo Brockhoff, and Eckart Zitzler. Hypervolume-based Multiobjective Optimization: Theoretical Foundations and Practical Implications. *Theoretical Computer Science*, 425:75-103, March 2011.

[7] Zyed Bouzarkouna, Didier Yu Ding, and Anne Auger. Well Placement Optimization with the Covariance Matrix Adaptation Evolution Strategy and Meta-Models. *Computational Geosciences*, 16(1):75-92, September 2011.

[8] Y. Hamadi, S. Jabbour, C. Piette, and L. Sais. Deterministic Parallel DPLL: System Description, *Int. Journal on Satisfiability, Boolean Modeling and Computation* (JSAT), Volume 7, 2011.

[9] Y. Hamadi, S. Jabbour, and L. Sais. Learning from Conflicts in Propositional Satisfiability, Invited Survey, *4OR: A Quarterly Journal of Operations Research* (4OR), 10(1), 2012.

[10] Y. Hamadi, and C. M. Wintersteiger. Challenges in Parallel SAT Solving, Invited paper, *AI Magazine* 34 (2): 99-106 (2013).

[11] Nikolaus Hansen, Raymond Ros, Nikolas Mauny, Marc Schoenauer, and Anne Auger. Impacts of Invariance in Search: When CMA-ES and PSO Face Ill-Conditioned and Non-Separable Problems. *Applied Soft Computing*, 11:5755-5769, 2011.

[12] C. M. Wintersteiger, Y. Hamadi, and L. de Moura. Efficiently solving quantified bit-vector formulas, *Formal Methods in System Design* (FMSD) Invited article, Special issue *10 years to the SMT initiative*, 2013.

Peer-Reviewed Conferences

[13] Anne Auger, Dimo Brockhoff, and Nikolaus Hansen. Analyzing the Impact of Mirrored Sampling and Sequential Selection in Elitist Evolution Strategies. In Foundations of Genetic Algorithms (FOGA 2011), pages 127-138, Schwarzenberg, Austria, April 2011.

[14] Anne Auger, Dimo Brockhoff, and Nikolaus Hansen. Mirrored Sampling in Evolution Strategies With Weighted Recombination. In Natalio Krasnogor and Pier Luca Lanzi eds., Genetic and Evolutionary Computation Conference (ACM-GECCO), pages 861-868, Dublin, Ireland, ACM Press, 2011.

[15] Ouassim Ait Elhara, Anne Auger, and Nikolaus Hansen. A Median Success Rule for Non-Elitist Evolution Strategies : Study of Feasibility. In Christian Blum and Enrique Alba, eds., Genetic and Evolutionary Computation Conference (ACM-GECCO), Amsterdam, Netherlands, pp 415-422, ACM Press, 2013.

[16] Dirk Arnold, V. and Nikolaus Hansen. A (1+1)-CMA-ES for Constrained Optimisation. In Terence Soule and Jason H. Moore, eds., Genetic and Evolutionary Computation Conference (ACM-GECCO), pages 297-304, ACM Press, 2012.

[17] Riad Akrour, Marc Schoenauer, and Michèle Sebag. Preference-Based Policy Learning. In Dimitrios Gunopulos et al., eds., Proc. European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases, Part I, Athens, Greece, pp 12-27, LNAI 6911, Springer Verlag, 2011.

[18] Riad Akrour, Marc Schoenauer, and Michèle Sebag. APRIL: Active Preference-learning based Reinforcement Learning. In P. Flach and al., editors, ECML PKDD 2012, volume 7524 of LNCS, pages 116-131, Bristol, United Kingdom. Springer Verlag, 2012.

[19] Zyed Bouzarkouna, Anne Auger, and Didier Yu Ding. LocalMeta-Model CMA-ES for Partially Separable Functions. In Natalio Krasnogor and Pier Luca Lanzi eds., Genetic and Evolutionary Computation Conference (ACM-GECCO), pages 869-876, Dublin, Ireland, ACM Press, 2011.

[20] Zyed Bouzarkouna, Didier Yu Ding, and Anne Auger. Partially Separated Meta-models with Evolution Strategies for Well Placement Optimization. In 73rd EAGE Conference & Exhibition incorporating SPE EUROPEC, pages 1-9, Vienna, Austria, May 2011.

[21] Alexandre Chotard, Anne Auger, and Nikolaus Hansen. Cumulative Step-size Adaptation on Linear Functions. In Carlos Coello et al., eds., PPSN XII, pages 72-81, Taormina, Italy, LNCS 7491, Springer Verlag, 2012.

[22] A Fialho, Y. Hamadi, and M. Schoenauer. A Multi-objective Approach to Balance Buildings Construction Cost and Energy Efficiency, Twentieth European Conference on Artificial Intelligence (ECAI'12), Montpellier, France.

[23] Nicolas Galichet and Michèle Sebag. Exploration prudente : une approche par méthode de Monte-Carlo arborescente contrainte. In RFIA 2012 (Reconnaissance des Formes et Intelligence Artificielle), pages 978-2-9539515-2-3, Lyon, France, January 2012.

[24] Y. Hamadi, and C. M. Wintersteiger. Seven Challenges in Parallel SAT Solving, Invited paper, Twenty-Sixth AAAI Conference (AAAI'12), Toronto, Canada.

[25] A. Arbelaez, Y. Hamadi. Improving Parallel Local Search for SAT, Learning and Intelligent Optimization (LION 7), Roma, Italy.

[26] Mostepha Redouane Khouadjia, Marc Schoenauer, Vincent Vidal, Johann Dréo, and Pierre Savéant. MultiObjective AI Planning: Comparing Aggregation and Pareto Approaches. In Martin Middendorf and Christian Blum, editors, EvoCOP - 13th European Conference on Evolutionary Computation in Combinatorial Optimisation, volume 7832 of LNCS, pages 202-213, Vienna, Austria, March 2013. Springer Verlag.

[27] Mostepha Redouane Khouadjia, Marc Schoenauer, Vincent Vidal, Johann Dréo, and Pierre Savéant. MultiObjective AI Planning: Evaluating DAE-YAHSP on a Tunable Benchmark. In Robin C. Purshouse, Peter J. Fleming, and Carlos M. Fonseca, editors, EMO'13 - 7th International Conference on Evolutionary Multi-Criterion Optimization, volume 7811 of LNCS, pages 36-50, Sheffield, United Kingdom, March 2013. Springer Verlag.

[28] Mostepha Redouane Khouadjia, Marc Schoenauer, Vincent Vidal, Johann Dréo, and Pierre Savéant. Pareto-Based Multiobjective AI Planning. In Francesca Rossi, editor, IJCAI 2013, Beijing, China, August 2013.

[XXX] Mostepha Redouane Khouadjia, Marc Schoenauer, Vincent Vidal, Johann Dréo, and Pierre Savéant. Quality Measures of Parameter Tuning for Aggregated MultiObjective Temporal Planning. In Panos Pardalos and Guiseppe Nicosia, editors, LION 7 - Learning and Intelligent OptimizatioN Conference, LNCS 7997, To appear, Catania, Italy, 2013. Springer Verlag.

[29] Manuel Loth, Michèle Sebag, Youssef Hamadi, Marc Schoenauer, and Christian Schulte. Hybridizing Constraint Programming and Monte-Carlo Tree Search: Application to the Job Shop problem. Short paper in Panos Pardalos and Guiseppe Nicosia, editors, LION7 - Learning and Intelligent OptimizatioN Conference, LNCS 7997, To appear, Catania, Italy, 2013. Springer Verlag.

[30] Manuel Loth, Michèle Sebag, Youssef Hamadi, and Marc Schoenauer. Bandit-based Search for Constraint Programming. In Christian Schulte, ed., International Conference on Principles and Practice of Constraint Programming, Uppsala, Sweden, pages 464-480, LNCS 8124, Springer Verlag, 2013.

[31] Ilya Loshchilov, Marc Schoenauer, and Michèle Sebag. Adaptive Coordinate Descent. In Natalio Krasnogor and Pier Luca Lanzi, eds., Genetic and Evolutionary Computation Conference (GECCO 2011), Dublin, Ireland, pp 885-992, ACM Press, 2011.

[32] Ilya Loshchilov, Marc Schoenauer, and Michèle Sebag. Not all parents are equal for MO-CMA-ES. In Ricardo H. C. Takahashi, Kalyanmoy Deb, Elizabeth F. Wanner, Salvatore Greco, eds., Evolutionary Multi-Criterion Optimization 2011 (EMO 2011), Ouro Preto, Brazil, pp 31-45, LNCS 6576, Springer Verlag, 2011.

[33] Ilya Loshchilov, Marc Schoenauer, and Michèle Sebag. Alternative Restart Strategies for CMA-ES. In Carlos Coello et al., eds., PPSN XII, pages 296-305, Taormina, Italy, LNCS 7491, Springer Verlag, 2012.

[34] Ilya Loshchilov, Marc Schoenauer, and Michèle Sebag. Self-Adaptive Surrogate-Assisted Covariance Matrix Adaptation Evolution Strategy. In Terence Soule and Jason H. Moore, eds., Genetic and Evolutionary Computation Conference (ACM-GECCO 2012), Philadelphia, United States, pp 321-328, ACM Press, 2012.

[35] Ilya Loshchilov, Marc Schoenauer, and Michèle Sebag. Intensive Surrogate Model Exploitation in Self-adaptive Surrogate-assisted CMA-ES (saACM-ES). In Christian Blum and Enrique Alba, eds., Genetic and Evolutionary Computation Conference (ACM-GECCO 2013), Amsterdam, Netherlands, pp 439-446, ACM Press, 2013.

[36] Ilya Loshchilov, Marc Schoenauer, and Michèle Sebag. KL-based Control of the Learning Schedule for Surrogate Black-Box Optimization. In Conférence sur l'Apprentissage Automatique, Lille, France, August 2013.

[37] Brendel Matthias and Marc Schoenauer. Learn-and-Optimize: a Parameter Tuning Framework for Evolutionary AI Planning. In Jin-Kao Hao and al., editors, Artificial Evolution, volume 7401 of LNCS, pages 159-170, Angers, France, September 2012. Springer Verlag.

[38] Gaétan Marceau-Caron, Pierre Savéant, and Marc Schoenauer. Computational Methods for Probabilistic Inference of Sector Congestion in Air Traffic Management. In Interdisciplinary Science for Innovative Air Traffic Management, Toulouse, France, July 2013.

[39] Gaétan Marceau-Caron, Pierre Savéant, and Marc Schoenauer. Multiobjective Tactical Planning under Uncertainty for Air Traffic Flow and Capacity Management. In IEEE Congress on Evolutionary Computation, Cancun, Mexico, pp 1548-1555, IEEE Press, 2013.

[40] Gaétan Marceau-Caron, Pierre Savéant, and Marc Schoenauer. Strategic Planning in Air Traffic Control as a Multi-objective Stochastic Optimization Problem. In ATM Seminar 2013, Chicago, United States, June 2013.

[41] Thomas Runarsson, Philip, Marc Schoenauer, and Michèle Sebag. Pilot, Rollout and Monte Carlo Tree Search Methods for Job Shop Scheduling. In Youssef Hamadi and Marc Schoenauer, editors, Learning and Intelligent OptimizatioN (LION'6), volume 7219 of LNCS, pages 408-423, Paris, France, October 2012. Sringer Verlag.

[42] Michèle Sebag and Olivier Teytaud. Combining Myopic Optimization and Tree Search: Application to MineSweeper. In Youssef Hamadi and Marc Schoenauer, editors, LION6, Learning and Intelligent Optimization, volume 7219 of LNCS, pages 222-236, Paris, France, 2012. Proc. LION 6, Sringer Verlag.

[43] Marc Schoenauer, Fabien Teytaud, and Olivier Teytaud. A Rigorous Runtime Analysis for Quasi-Random Restarts and Decreasing Stepsize. In Jin-Kao Hao et al., Eds, Artificial Evolution (selected papers), pp 37-48, LNCS 7401, Springer Verlag, 2012.

[44] N. Veerapen, Y. Hamadi and F. Saubion, Using Local Search with Adaptive Operator Selection to Solve the Progressive Party Problem, IEEE Congress on Evolutionary Computation (CEC'13), Cancun, 2013.

[45] B. Yordanov, C. M. Wintersteiger, Y. Hamadi, A. Phillips, and H. Kugler. Functional Analysis of Large-scale DNA Strand Displacement Circuits, 19th International Conference on DNA Computing and Molecular Programming (DNA'19), Tempe, USA, 2013.

[46] B. Yordanov, C. M. Wintersteiger, Y. Hamadi, and H. Kugler. SMT-based Analysis of Biological Computation, 5th NASA Formal Methods Symposium, (NFM'13), Moffett Field, USA, 2013.

Workshop papers

[47] Z34Bio: A Framework for Analyzing Biological Computation, B. Yordanov, C. M. Wintersteiger, Y. Hamadi, and H. Kugler, 11th Int. Workshop on Satisfiability Modulo Theories (SMT'13), Helsinki, 2013.

[48] Bandit-based Search for Constraint Programming, M. Loth, M. Sebag, Y. Hamadi, C. Schulte and M. Schoenauer, 2nd Workshop on COmbining COnstraint solving with MIning and LEarning (Cocomile'13), Bellevue, 2013.

[49] Cooperation control in Parallel SAT Solving: a Multi-armed Bandit Approach, L. Nadjib, Y. Hamadi, S. Jabbour, and M. Sebag, (NIPS'12) Workshop on Bayesian Optimization and Decision Making, Lake Tahoe, USA.

[50] Optimizing Architectural and Structural Aspects of Buildings towards Higher Energy Efficiency, A. Fialho, Y. Hamadi, and M. Schoenauer, (GECCO'11), Workshop on GreenIT Evolutionary Computation, July 2011, UK.

[51] Lazy Decomposition for Distributed Decision Procedures, Y. Hamadi, J. Marques-Silva, and C. M. Wintersteiger, International Workshop on Parallel and Distributed Methods in verifiCation (PDMC'11).

[52] Anne Auger, Dimo Brockhoff, and Nikolaus Hansen. Benchmarking the Local Metamodel CMA-ES on the Noiseless BBOB'2013 Test Bed. In ACM-GECCO (Companion), Workshop on Black-Box Optimization Benchmarking (BBOB'2013), Amsterdam, Netherlands, pp 1225-1232, ACM Press, 2013.

[53] Dimo Brockhoff, Anne Auger, and Nikolaus Hansen. Comparing Mirrored Mutations and Active Covariance Matrix Adaptation in the IPOP-CMA-ES on the Noiseless BBOB Testbed. In GECCO Companion '12, pages 297-303, Philadelphia, PA, United States, July 2012.

[54] Dimo Brockhoff, Anne Auger, and Nikolaus Hansen. On the Effect of Mirroring in the IPOP Active CMA-ES on the Noiseless BBOB Testbed. In GECCO Companion '12, pages 277-284, Philadelphia, PA, United States, July 2012.

[55] Dimo Brockhoff, Anne Auger, and Nikolaus Hansen. On the Impact of a Small Initial Population Size in the IPOP Active CMA-ES with Mirrored Mutations on the Noiseless BBOB Testbed. In GECCO Companion '12, pages 285-290, Philadelphia, PA, United States, July 2012.

[56] Dimo Brockhoff, Anne Auger, and Nikolaus Hansen. On the Impact of Active Covariance Matrix Adaptation in the CMAES With Mirrored Mutations and Small Initial Population Size on the Noiseless BBOB Testbed. In GECCO Companion '12, pages 291-296, Philadelphia, PA, United States, July 2012.

[57] Areski Hadjaz, Gaétan Marceau, Pierre Savéant, and Marc Schoenauer. Increasing Air Traffic: What is the Problem? In Dirk Schaefer, editor, SESAR 2nd Innovation Days, Braunschweig, Germany, September 2012. SESAR WPE.

[58] Areski Hadjaz, Gaétan Marceau, Pierre Savéant, and Marc Schoenauer. Online Learning for Ground Trajectory Prediction. In Dirk Schaefer, editor, SESAR 2nd Innovation Days, Braunschweig, Germany, September 2012.

[59] Ilya Loshchilov, Marc Schoenauer, and Michèle Sebag. Black-box optimization benchmarking of IPOP-saACM-ES and BIPOP-saACM-ES on the BBOB-2012 noiseless testbed. In Workshop Proceedings of the Genetic and Evolutionary Computation Conference, Philadelphia, United States, April 2012.

[60] Ilya Loshchilov, Marc Schoenauer, and Michèle Sebag. Black-box optimization benchmarking of IPOP-saACM-ES on the BBOB-2012 noisy testbed. In Workshop Proceedings of the Genetic and Evolutionary Computation Conference, Philadelphia, United States, April 2012.

[61] Ilya Loshchilov, Marc Schoenauer, and Michèle Sebag. Black-box Optimization Benchmarking of NIPOP-aCMA-ES and NBIPOP-aCMA-ES on the BBOB-2012 Noiseless Testbed. In Workshop Proceedings of the GECCO Genetic and Evolutionary Computation Conference, Philadelphia, United States, October 2012.

[62] Ilya Loshchilov, Marc Schoenauer, and Michèle Sebag. BI-population CMA-ES Algorithms with Surrogate Models and Line Searches. In Workshop Proceedings of the (GECCO) Genetic and Evolutionary Computation Conference, page 8, Amsterdam, Netherlands, April 2013. ACM.

[63] Brendel Matthias and Marc Schoenauer. Instance-Based Parameter Tuning and Learning for Evolutionary AI Planning. In 21st International Conference on Automated Planning and Scheduling, Planning and Learning Workshop, Freiburg, Germany, June 2011.

[64] Brendel Matthias and Marc Schoenauer. Instance-based parameter tuning for evolutionary AI planning. In Genetic and Evolutionary Computation Conference (ACM-GECCO) Companion (workshop), Dublin, Ireland, July 2011.

Theses and Habilitations

[65] Alejandro Arbelaez. *Search with the Context*. PhD thesis, Université Paris-Sud, May 2011. Directed by Youssef Hamadi and Michèle Sebag.

[66] Zyed Bouzarkouna. *Well Placement Optimization*. PhD thesis, Université Paris-Sud, April 2012. Directed by Anne Auger and Marc Schoenauer.

[67] Ilya Loshchilov. *Surrogate-Assisted Evolutionary Algorithms*. PhD thesis, Université Paris-Sud, January 2013. Directed by Marc Schoenauer and Michèle Sebag.

[68] Youssef Hamadi. *Search: from Algorithms to Systems*. Habilitation thesis (HDR), Université Paris-Sud, January 2013.

Technical Reports

[69] Anne Auger and Nikolaus Hansen. Linear Convergence on Positively Homogeneous Functions of a Comparison Based StepSize Adaptive Randomized Search: the (1+1) ES with Generalized One-fifth Success Rule. October 2013.

[70] Anne Auger and Nikolaus Hansen. On Proving Linear Convergence of Comparison-based Step-size Adaptive Randomized Search on Scaling-Invariant Functions via Stability of Markov Chains. November 2013.

[71] Alexandre Chotard, Adrien, Anne Auger, and Nikolaus Hansen. Cumulative Step-size Adaptation on Linear Functions: Technical Report. Research report, June 2012.

[72] Nikolaus Hansen. A CMA-ES for Mixed-Integer Nonlinear Optimization. Research Report RR-7751, INRIA, October 2011.

[73] Nikolaus Hansen. Injecting External Solutions Into CMA-ES. Research Report RR-7748, INRIA, October 2011.

[74] Nikolaus Hansen, Steffen Finck, and Raymond Ros. COCO - COmparing Continuous Optimizers : The Documentation. Research Report RT-0409, INRIA, May 2011.

[75] Nadjib Lazaar, Youssef Hamadi, Said Jabbour, and Michèle Sebag. Cooperation control in Parallel SAT Solving: a Multi-armed Bandit Approach. Research Report RR-8070, INRIA, September 2012.

Invited Talks and tutorials

- Anne Auger and Nikolaus Hansen, GECCO 2010-11-12 and LION'6 Invited Tutorial on CMA-ES
- Anne Auger, seminar at CMAP (Centre de Mathématiques Appliquées de l'Ecole Polytechnique), April 2012
- Youssef Hamadi, *Automating Parallel SAT Solving*, Learning and Intelligent Conference, Catania, Italy, 2013.
- Youssef Hamadi, *SmartBuildings*, Journée Industrielle organisée conjointement par les GDR CNRS R.O 3002 (Recherche Opérationnelle) et A.S.R 725 (Architectures, Systèmes, Réseaux), en partenariat avec la ROADEF, Novembre, Paris, France, 2013.
- Youssef Hamadi, *On Microsoft Research Policies*, Towards a Global Observatory of Policy Instruments on Science, Technology and Innovation UNESCO Workshop, Paris October 19-20th 2011.
- Youssef Hamadi, *SmartBuildings***,** Green Growth Leaders Workshop, Copenhagen, October 12th 2011.
- Youssef Hamadi, *Approaches to Parallel SAT***,** lecturer first MIT Summer School on SAT/SMT, Cambridge (MA), June 2011.
- *Parallel SAT*, Instituto de Engenharia de Sistemas e Computadores Investigação (INESC-ID), Lisbon, Portugal, February 2011.
- Marc Schoenauer, *EVOLVE* (3rd edition), Luxembourg, May 2011
- Marc Schoenauer, 18th *CREST Open Workshop*, Managing and Optimising Multiplicity Computing, UCL, Londond, UK, 22-23 March 2012;
- Marc Schoenauer, *Complex Adaptive Systems Laboratory* seminar, University College Dublin, Ireland, 25 May 2012;
- Marc Schoenauer, *ECODAM*, Doctoral Summer School on Evolutionary Computation and Data Mining, Faculty of Informatics, Iasi University, Romania, 18-23 June 2012;
- Marc Schoenauer, Séminar of the *Département de génie de la production automatisée*, Ecole Supérieure de Technologie, Montreal, Canada, 7 Dec. 2012.
- Michèle Sebag, *KAUST*, Saoudi Arabia (Feb. 2011);
- Michèle Sebag, *U. Zurich*, Switzerland (March 2011);
- Michèle Sebag, *U. York*, UK (Nov. 2011);
- Michèle Sebag, *Spring Workshop on Mining and Learning* at Bad Neuenahr, 2012.
- Michèle Sebag, Sixth *Starting Artificial Intelligence Research* Symposium, August 27, 2012.
- Michèle Sebag, Invited tutorial, *International Summer School on Resource-aware Machine Learning*, Dortmund, Germany, Sept. 4-7, 2012.
- Michèle Sebag, Invited tutorial, *Constraint Programming 2012*, Québec, Canada, Oct. 8-12, 2012.

Events, workshops, conferences, seminars

- Following the seminal ones in 2009 and 2010, two **BBOB workshops** (Black-Box Optimization Benchmarking) were co-organized by Anne Auger and Nikolaus Hansen during the 2012 and 2013 editions of ACM-GECCO (Genetic and Evolutionary Computation COnference), the main conference in the Evolutionary Computation domain. All participants to both workshops submitted the best version of their favorite continuous optimizer that were globally compared on the 24 noiseless functions and/or on their noisy counterparts. All contributions and results (including the full output data for each entry) are publically available at http://coco.gforge.inria.fr
- The **Special Issue** of *Evolutionary Computation* journal around benchmarking issues in continuous optimization, follow-up of the first BBOB workshops, with guest Editors Anne Auger, Nikolaus Hansen, and Marc Schoenauer, appeared in Winter 2012, as Volume 20 number 4. Five papers had been selected out of twelve submissions.
- Marc Schoenauer created together with Gabriela Ochoa the **Self-* Search Track** at ACM-GECCO 2011 (Genetic and Evolutionary Computation COnference), the main conference in the Evolutionary Computation domain, as *New Frontier* track. This track has been pursued since then (number of submitted papers for 2011-12-13: 30, 15, 22 resp.) and will still be active in 2014 –

though with different track chairs every year, following GECCO rules.

- The **Special Issue** of *Evolutionary Computation* journal dedicated to parameter tuning in Evolutionary Computation at large, follow-up of the two workshops that had been co-organized by Marc Schoenauer (together with Gabriela Ochoa) PPSN 2010 and LION-5 in 2011 – Guest Editors Thomas Bartz-Beielstein , Gabriela Ochoa, Mike Preuss, and Marc Schoenauer), appeared in Summer 2012 (Volume 20 number 2). Five papers had been selected out of eight submissions.
- Youssef Hamadi and Marc Schoenauer **co-organised LION-6** conference in January 2012 in Paris, in Microsoft Conference Center as Issy-les-Moulineaux. All information is on the Web site at http://www.intelligent-optimization.org/LION6/. The conference attracted 81 participants, including the 3 invited speakers and the 3 invited tutorial speaker.
- **Award**: Marc Schoenauer and his PhD student Jacques Bibaï, together with their co-authors Pierre Savéant and Vincent Vidal, won the Temporal Satisficing Track at IPC'11, the international competition on AI planning organised during ICAPS conference, the most prominent event in Planning and Scheduling.
- **Best paper award** to Nikolaus Hansen (with Dirk Arnold) in the ES-EP/Theory track of GECCO 2012 for the paper *A (1+1)-CMA-ES for Constrained Optimisation*.

7.5 Tools and software

1. Cooperative Stochastic Local Search Solver (CSLS), Silver medal in the SAT+Random category, SAT Competition 2011.
2. Bandit Ensemble for Parallel SAT Solving (BESS), Multi-Armed Bandit techniques to avoid a full broadcasting of clauses between the different nodes 2012, (see section 8.1).
3. The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) is considered state-of-the-art in continuous domain evolutionary computation. (See H.-G. Beyer (2007). Evolution Strategies, *Scholarpedia*, p. 1965.) It has been shown to be highly competitive on different problem classes. The algorithm is widely used in research and industry as witnessed by more than a hundred published applications. We provide source code for the CMA-ES in C, Java, Matlab, Octave, Python, and Scilab, including the latest variants of the algorithm. Check all details and download source code at http://www.lri.fr/~hansen/cmaes_inmatlab.html.
4. Extension to the Gecode Constraint Solver (open source) to implement the BASCOP algorithm (soon to be released).
5. COCO (COmparing Continuous Optimizers, http://coco.gforge.inria.fr/) is a platform for systematic and sound comparisons of real-parameter global optimizers. COCO provides benchmark function testbeds and tools for processing and visualizing data generated by one or several optimizers. Several classes of benchmark functions have been thoroughly designed, including noiseless and noisy functions, where great care has been taken to perturb the well-known test functions (e.g. moving the optimum around, rotating the coordinate system, …). Several post-processing procedures have also been defined, with standard outputs leading to fair and sound comparisons between different optimizers. An API allows the users to easily interface their own optimizer with the test set, and the graphics and tables that are automatically generated give instantly a clear picture of pairwise or more global comparisons. Further work includes the handling of constraints and that of mixed-integer problems. However, such extension also requires the definition of both new performance measures and new sets of benchmark functions.

Contribution to Microsoft technologies

Microsoft Solver Foundation (MSF) is designed to help businesses make optimal strategic decisions. The possible applications cover a vast range: real-time supply chain optimization, data center energy profile management, on-line advertising profit maximization, logistics of large conference scheduling, transportation network flows, and risk analysis of investment portfolios. There are also direct applications to graphics and machine learning. All these problems are NP-difficult and MSF allows the programmer to choose between different solvers in order to quickly compute optimal or approximate solutions for their applications. Unfortunately, if MSF provides several solvers, it does not provide the expertise required to decide between them or even to pick up the right parameters that should be used with a particular

algorithm. In other words, it does not guarantee that a standard programmer will reach an acceptable level of performance for the resolution of its problem. In this context, the *Adapt* project which aims at improving the usability of modern optimization methods (targeting e-Scientists users) clearly meets the actual limitations of MSF. We are sharing these results with the MSF SD Leads, and we believe that the results of our project will greatly influence the future versions of the Microsoft Solver Foundation line of products.

Scientific Image and Video Data Mining

www.msr-inria.fr/projects/scientific-image-and-video-mining-2/

Our project involves fundamental computer science research in computer vision and machine learning, applied to archaeology, cultural heritage preservation, and sociology, and validated by collaborations with researchers and practitioners in these fields. Concretely, we address: (i) Mining historical collections of photographs and paintings with applications to archaeology and cultural heritage preservation; and (ii) Mining and analysis of TV broadcasts with applications to sociology.

Team

Status	Last name	First name	Affiliation
Team leader	Ponce	Jean	Inria Paris-Rocquencourt
Researcher	Laptev	Ivan	Inria Paris-Rocquencourt
Researcher	Sivic	Josef	Inria Paris-Rocquencourt
Researcher	Harchaoui	Zaid	Inria Grenoble-Rhone-Alpes
Researcher	Zisserman	Andrew	Oxford University – ENS Paris
Researcher	Blake	Andrew	Microsoft Research Cambridge
Researcher	Szeliski	Richard	Microsoft Research Redmond
Researcher	Dessales	Hélène	Ecole Normale Supérieure Paris
PhD student	Harchaoui	Warith	Inria Paris Rocquencourt
Post Doc Student	Cherian	Anoop	Inria Grenoble Rhône Alpes
PhD student	Yang	Huang	Inria Grenoble Rhône Alpes
PhD student	Yang	Huang	Inria Grenoble Rhône Alpes
PHD Student	Oquab	Maxime	Inria Paris-Rocquencourt

Left:

- GAIDON Adrien, INP Grenoble, graduated with Phd in October 2012. (Currently at Xerox Research Centre Europe).
- WHYTE Oliver, Inria Paris-Rocquencourt, graduated with Phd in March 2012. (Currently at Microsoft in Redmond).

Visitors:

Longer term visitors include: Alexei Efros (Professor, Carnegie Mellon University, USA) who visited WILLOW for several months in 2011, 2012 and 2013; Abhinav Gupta (Assistant Research Professor, Carnegie Mellon University, USA) who visited WILLOW in summer 2011, and Rene Vidal, (Associate Professor, Johns Hopkins University, USA) who visited WILLOW for several months in summer 2012.

Large-scale dense 3D reconstruction of the house of Diomedes constructed from about 30,000 photographs.

Research

Concretely, we propose to address the following problems:

- Mining historical collections of photographs and paintings with applications to archeology and cultural heritage preservation. This includes the use of image-based modeling technology to facilitate the metrology side of field work, but also the quantitative analysis of environmental damage on wall paintings or mosaics over time, and the cross-indexing of XIXth Century paintings of Pompeii with modern photographs. This part of our research is done in collaboration with Helene Dessales at the ENS Archeology Laboratory.

- Mining TV broadcasts with applications to sociology. This includes automating the analysis and annotation of human actions and interactions in video segments to assist - and provide data for - studies of consumer trends in commercials, political event coverage in newscasts, and class- and gender-related behavior patterns in situation comedies, for example. This part of our research is done in collaboration with Louis Laborelli and Daniel Teruggi at Institut National de l'Audiovisuel (INA).

For every one of the problems we have in mind, indexing, searching and analyzing photo and video collections is a key issue. Recent advances in image analysis, computer vision, and machine learning promise an opportunity to automate, partly or completely, these tasks (e.g., annotation of photos and videos), as well as to access information whose extraction from images is simply beyond human capabilities (e.g., indexing of very large image archives). To fulfill this promise, we propose to conduct fundamental research in object, scene, and activity modelling, learning, and recognition, and to validate it with the development of computerized image and video mining tools at the service of sciences and humanities.

Organization

Our project brings together two Inria teams, LEAR and WILLOW (the latter is a joint venture between Inria, Ecole Normale Superieure [ENS], and CNRS, and is hosted by ENS in Paris), with complementary strengths in computer vision (image-based modeling, dynamic image analysis, object recognition) and machine learning. It involves Microsoft Research researchers from Cambridge, under the leadership of A. Blake, as well as others from Redmond and New York for example. It also involves external partners including Helene Dessales at the ENS Archaeology Laboratory, as well as Louis Laborelli and Daniel Teruggi at Institut National de l'Audiovisuel (INA). The Collaborative Research Agreement (CRA) between Microsoft Research, Inria, and ENS has been signed on September 15, 2008.

Adrien Gaidon and Oliver Whyte graduated with Phd in March 2012 and October 2012 respectively. Anoop Cherian (LEAR) was hired as post-doc in November 2012 and Hua Yang (LEAR) was hired as a Phd student in October 2012. Maxime OQUAB has finished an MSc internship in WILLOW co-supervised by L. Bottou (Microsoft Research New York), I. Laptev (WILLOW) and J. Sivic (WILLOW) and will continue as a Phd student from Jan 2014. We have sent two students to Microsoft Research for an internship: A. Joulin (Microsoft Research Redmond, 2012) and M. Oquab (Microsoft Research New

York, 2013). The internship of A. Joulin has resulted in 2 publications with Microsoft Research researchers.

Results

Quantitative image analysis for archeology

The goal of this project is to enable fully automatic matching and alignment of paintings and drawings to photographs depicting a complex 3D scene. This is an extremely difficult task due to various distortions that can arise such as perspective or caricature distortion as well as inaccuracies due to drawing by hand. Progress on this topic is of interest to archaeologists, artists or curators.

For this, J. Ponce, Y. Ubelmann, and H. Dessales visited the archaeological site at Pompeii to photograph the house of Diomedes, which is now the focus of our study. From this set of images (in total about 30,000 photographs were taken), we were able to produce a large-scale dense 3D reconstruction of the house of Diomedes using existing photometric multi-view stereo methods as shown in figure~\ref{fig:Pompeii3D}. This part of our project demonstrates the use of photogrammetric methods in archaeology at a much larger scale than before (compared to the previous model of the much smaller house of Championnet constructed from about 500 photos). We have also constructed a preliminary, low-resolution model of the entire city. Our current efforts focus on how to extract a high-fidelity mesh from the point cloud generated by our method.

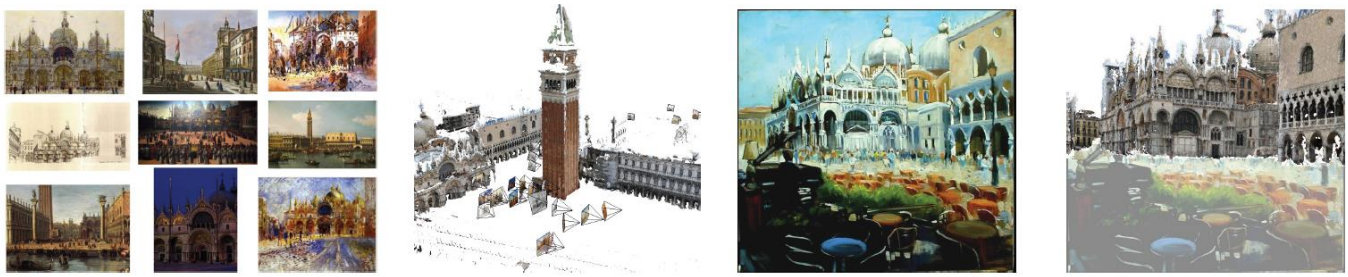

(a) Internet paintings. (b) Painting viewpoints. (c) Aligned painting to 3D model.

Our method automatically aligns and recovers the viewpoint of paintings, drawings, and historical photographs to a 3D model of an architectural site.

Next, we have developed a new method~\cite{Aubry13} for automatic alignment of non-photographic depictions with 3D models of architectural sites, as illustrated in figure~\ref{fig:teaser2d3d}. The 3D model of the scene is represented by a small set of discriminative visual elements that are automatically learnt from rendered views. Similar to object detection, the set of visual elements, as well as the weights of individual features for each element, are learnt in a discriminative fashion. We show that the learnt visual elements are reliably matched in 2D depictions of the scene despite large variations in rendering style (e.g. watercolor, sketch) and structural changes (e.g. missing scene parts, large occluders) of the scene, as illustrated in figure~\ref{fig:Aubry13match}. The proposed alignment procedure is validated via a human user study on a new database of paintings and sketches spanning several sites. The results demonstrate that our algorithm produces significantly better alignments than several baseline methods.

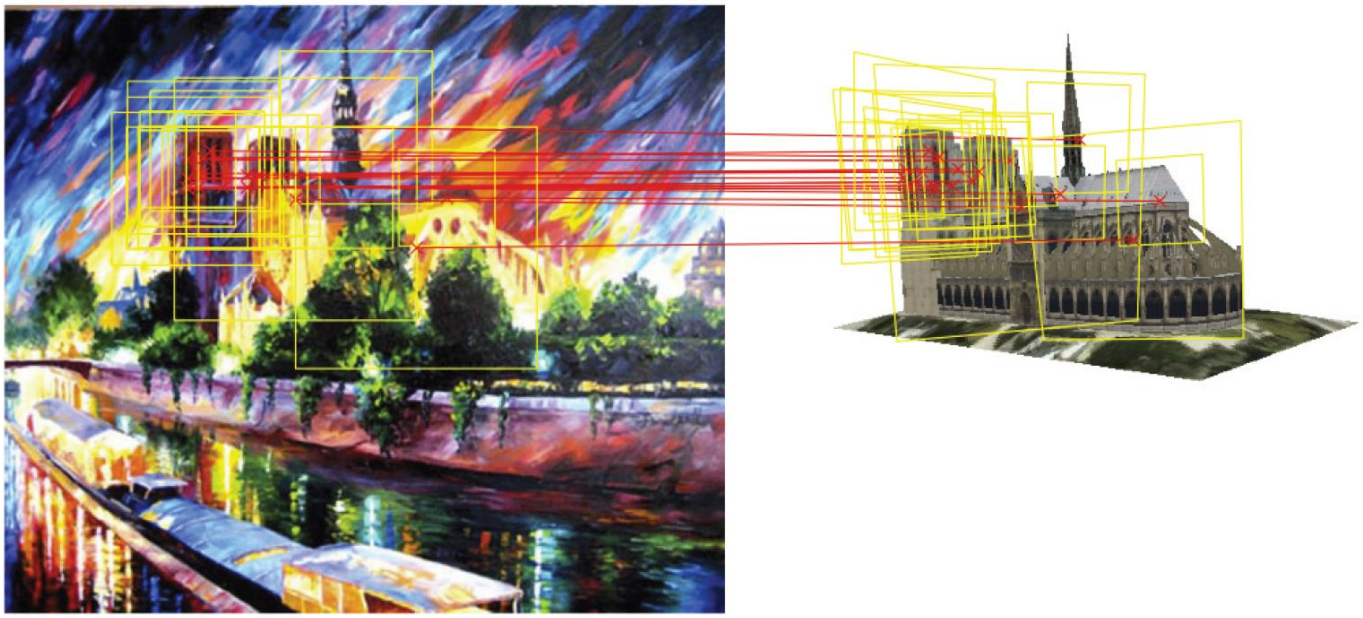

The entire architectural site is summarized by a set of mid-size 3D discriminative visual elements that are used to find correspondences between the input scene depiction (left) and the 3D model (right).

Large scale image matching and retrieval

Building large scale 3D models from photographs involves establishing correspondences between large amounts of images, which is one of the most time consuming steps in the 3D reconstruction pipeline.

Towards this goal, we are also investigating techniques for large scale image indexing and retrieval. We have developed two methods for large scale matching in large geotagged datasets. The first method~\cite{Gronat13} takes advantage of geotags as an available form of supervision and investigates whether the large scale place matching problem can be cast as a classification task. This is beneficial as each classifier can learn which features are discriminative for a particular place. In the second method~\cite{Torii13} we develop a scalable representation for large-scale matching of repeated structures. While repeated structures often occur in man-made environments -- examples include building facades, walls and fences -- they are usually treated as nuisance and downweighted at the indexing stage. In contrast, we have developed a simple but efficient representation of repeated structures and demonstrate its benefits for large scale matching. The two above described techniques have demonstrated significant gains in matching accuracy for place recognition and we plan to investigate its benefits for speeding-up large scale 3D reconstruction, where historical imagery captured over extended time-periods is available.

Automatic mining of distinctive architectural elements

Given a set of images from several geo-spatial areas, in~\cite{Doersch2012} we seek to automatically find visual elements, e.g. windows, balconies, and street signs, that are most distinctive for a certain geo-spatial area, for example the city of Paris. We propose to use a discriminative clustering approach able to take into account the weak geographic supervision.

We demonstrate that these elements are visually interpretable and perceptually geo-informative. The discovered visual elements can also support a variety of computational geography tasks, such as mapping architectural correspondences and influences within and across cities, with potential applications in architecture, history and archeology.

8.4.2 Action modeling and recognition

Based on our earlier promising results on human action classification reported in~\cite{Laptev08}, we address a more challenging problem and propose to localize actions in time and space, i.e.~answering the questions ``who is doing what" and ``when" for actions like ``walking" or ``opening door" that appear in

the video.

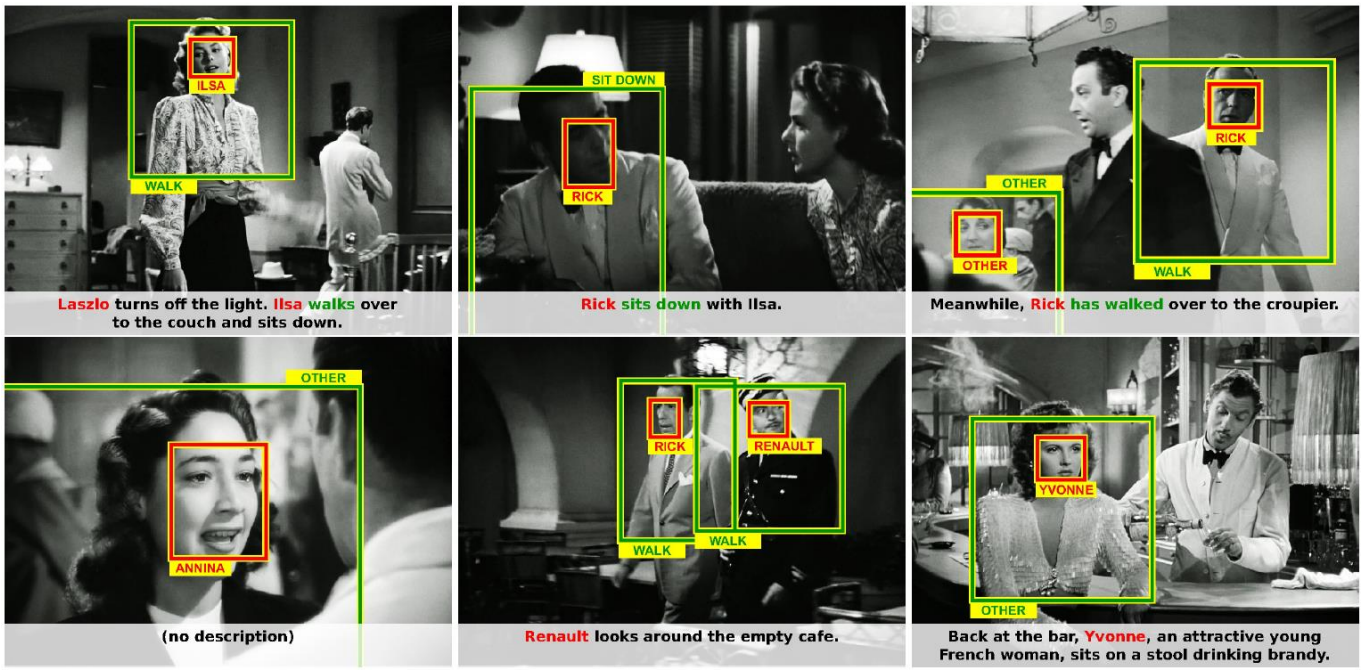

Examples of automatically recognized names and actions in the movie Casablanca.

The automatically resolved correspondence between video and script is color-coded.

\end{figure*}

The first contribution is a joint effort between the WILLOW and LEAR teams, published in~\cite{Bojanowski13}. In this work we address the problem of learning a *joint model of actors and actions* in movies using weak supervision provided by scripts. Specifically, we extract actor/action pairs from the script and use them as constraints in a discriminative clustering framework. The corresponding optimization problem is formulated as a quadratic program under linear constraints. People in video are represented by automatically extracted and tracked faces together with corresponding motion features. This allows us to associate individual actions in the script to individual actors in the video as illustrated in figure~\ref{fig:bojanowski13im3}.

We validate our method in the challenging setting of localizing and recognizing characters and their actions in feature length movies Casablanca and American Beauty.

Besides the work on localizing actions, we have addressed several other directions aiming to improve recognition of actions and human traits in video.

Activity representation with motion hierarchies. In~\cite{gaidon:hal-00804627,gaidon:hal-00722955}

Complex activities, e.g., pole vaulting, are composed of a variable number of sub-events connected by complex spatio-temporal relations, whereas simple actions can be represented as sequences of short temporal parts. In this work, published in~\cite{gaidon:hal-00804627,gaidon:hal-00722955}, we learn hierarchical representations of activity videos in an unsupervised manner.

We show that per-video hierarchies provide additional information for activity recognition. Our approach improves over unstructured activity models, baselines using other motion decomposition algorithms, and the state of the art.

Pose estimation and segmentation of people. In~\cite{Alahari13} we address the problem of pixel-wise segmentation and pose estimation of multiple people in a stereoscopic video. This involves challenges such as dealing with unconstrained stereoscopic video, non-stationary cameras, and complex indoor and

outdoor dynamic scenes. We develop a segmentation model incorporating person detection, pose estimation, as well as color, motion, and disparity cues. The model also explicitly represents depth ordering and occlusion of people in video.

Modeling and recognition of person-object interactions. Our everyday objects support various tasks and can be used by people for different purposes. While object classification is a widely studied topic in computer vision, recognition of object function, i.e., what people can do with an object and how they do it, is rarely addressed. In this work~\cite{Delaitre2012} we construct a functional object description with the aim to recognize objects by the way people interact with them. We describe scene objects (sofas, tables, chairs) by associated human poses and object appearance.

Modeling people in crowded scenes. We have developed two representations of people in crowded scenes~\cite{Rodriguez2011,Rodriguez2011a}. While the detection of individual objects has been improved significantly over the recent years, crowd scenes remain particularly challenging for the detection and tracking tasks due to heavy occlusions, high person densities and significant variation in people's appearance. To address these challenges, we propose to leverage information on the global structure of the scene and non-parametric motion priors learnt off-line from a large database of crowd events. We validate our approach on a new challenging video dataset of crowded scenes.

Publications & talks from the project team members (2011-2013)

[1] K. Alahari, G. Seguin, J. Sivic, and I. Laptev. Pose estimation and segmentation of people in 3D movies. In Proceedings of the International Conference on Computer Vision, 2013. 4
[2] M. Aubry, B. Russell, and J. Sivic. Painting-to-3D model alignment via discriminative visual elements. ACM Transactions on Graphics, 2013. Accepted for publication. 3
[3] P. Bojanowski, F. Bach, I. Laptev, P. J., C. Schmid, and J. Sivic. Finding actors and actions in movies. In Proceedings of the International Conference on Computer Vision, 2013. 4, 7
[4] V. Delaitre, D. Fouhey, I. Laptev, J. Sivic, A. Efros, and A. Gupta. Scene semantics from long-term observation of people. In European Conference on Computer Vision, Florence, Italy, 2012. 4
[5] C. Doersch, S. Singh, A. Gupta, J. Sivic, and A. Efros. What makes paris look like paris? ACM Transactions on Graphics (SIGGRAPH), 31(4), 2012. 3, 6
[6] A. Gaidon, Z. Harchaoui, and C. Schmid. Recognizing activities with cluster-trees of tracklets. In BMVC, Guildford, United Kingdom, September 2012. 4
[7] A. Gaidon, Z. Harchaoui, and C. Schmid. Temporal Localization of Actions with Actoms. IEEE Transactions on Pattern Analysis and Machine Intelligence, March 2013. 4
[8] P. Gronat, G. Obozinski, J. Sivic, and T. Pajdla. Learning and calibrating per-location classifiers for visual place recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2013. 3
[9] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2008. 3
[10] M. Rodriguez, I. Laptev, J. Sivic, and J.-Y. Audibert. Density-aware person detection and tracking in crowds. In International Conference on Computer Vision, 2011. 4
[11] M. Rodriguez, J. Sivic, I. Laptev, and J.-Y. Audibert. Datadriven crowd analysis in videos. In International Conference on Computer Vision, 2011. 4
[12] A. Torii, J. Sivic, T. Pajdla, and M. Okutomi. Visual place recognition with repetitive structures. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2013. 3

Invited Talks

- J. Sivic, Microsoft Research Redmond, USA, Host: R. Szeliski, March 2011
- J. Sivic, Carnegie Mellon University, USA, Host: A. Efros, February 2011
- J. Sivic, AViRS workshop on video surveillance, Paris, Host: D. Marraud, February 2011
- J. Sivic, Czech Technical University in Prague, Host: J. Matas, April 2011

- J. Ponce, Distinguished speaker, Taiwan Academica Sinica, 2011
- J. Ponce, Distinguished speaker, University of Delaware Computer Science Department, 2011
- J. Ponce, ETH Zurich Computer Science Department, 2011
- I. Laptev, Royal Institute of Technology, Stockholm, Sweden, Host: S. Carlsson, December 2011
- I. Laptev, ICCV2011 International Workshop on Video Event Categorisation, Barcelona, Spain, November 2011
- I. Laptev, Assemblee generale du GdR ISIS, Saint-Georges-de-Didonne, France, May 2011
- A. Zisserman, Frontiers workshop on Computer Vision, MIT, August, 2011.
- A. Zisserman, BBC Faces Workshop, September, 2011.
- I. Laptev, 3rd AFCV Int. Workshop on Recent Trends in Computer Vision, Osaka, Japan, Jan. 2012.
- I. Laptev, 10th Workshop on Content Based Multimedia Indexing, Annecy, France, June 2012.
- I. Laptev, Oxford Univ., Oxford, UK, Host: A. Zisserman, Sept. 2012.
- I. Laptev, First Croatian Workshop on Computer Vision, Zagreb, Croatia, Sept. 2012.
- I. Laptev, 3rd IST Austria Symposium on Computer Vision and Machine Learning, Vienna, Austria, Oct. 2012.
- I. Laptev, Carnegie Mellon University, USA, Host: A. Efros, November 2012.
- J. Ponce, 2nd ACCV Workshop on e-Heritage, Daejeon, South Korea, Nov. 2012.
- J. Ponce, ETRI, Daejeon, South Korea, Nov. 2012.
- J. Ponce, Let's imagine the future, Inria Rennes, France, Nov. 2012.
- J. Sivic, Google, USA, Hosts: H. Adam and H. Neven, August 2012
- J. Sivic, UC Berkeley, USA, Host: J. Malik, August 2012
- J. Sivic, Simon Fraser University, Canada, Host: G. Mori, August 2012
- J. Sivic, GdR ISIS workshop, Telecom ParisTech, Host: F. Lafarge, October 2012
- J. Sivic, First Int. Workshop on Visual Analysis and Geo-Localization of Large-Scale Imagery, ECCV 2012, October 2012.
- J. Sivic, Int. Workshop on Search Compupting, Brussels, Host: A. Joly, September 2012.
- J. Sivic, Carnegie Mellon University, USA, Host: A. Efros, December 2012.
- C. Schmid, Presentation at Colloquium J. Morgenstern, Sophia-Antipolis, December 2011.
- C. Schmid: Presentation at NIPS Workshop, Granada, Spain, December 2011.
- C. Schmid: Presentation at Symposium on Applied Perception in Graphics and Visualization, Toulouse, August 2011.
- C. Schmid: Presentation at Frontiers in Computer Vision Workshop, MIT, August 2011.
- A. Gaidon: Technical demonstration at the Microsoft Research / Inria forum, Paris, April 2011.
- A. Gaidon: Seminar at the Computer Vision Center, Autonomous University of Barcelona, Spain, May 2011.
- A. Gaidon: Presentation at the "Journée perception de l'homme et reconnaissance d'actions", GdR ISIS, CNRS, Paris, June 2011.
- Z. Harchaoui: Seminar at University of California Berkeley, January 2011.
- Z. Harchaoui: Presentation at Stat'Learn, Grenoble, March 2011.
- Z. Harchaoui: Presentation at GDR Isis, Paris, April 2011.
- Z. Harchaoui: Seminar at University Paris VI, May 2011.
- Z. Harchaoui: Seminar at Xerox Research Center Europe, September, Meylan, 2011.
- Z. Harchaoui: Seminar at Kyoto University, Japan, November 2011.
- A. Gaidon: Seminar at ETH Zurich, Switzerland, April, 2012.
- A. Gaidon: Seminar at Xerox Research Center Europe (XRCE), Meylan, France, May, 2012.
- Z. Harchaoui: Seminar at Gatsby Neuroscience Unit, UCL, London, March 2012.
- Z. Harchaoui: Presentation at International Symposium in Mathematical Programming, Berlin, August 2012.
- Z. Harchaoui: Seminar at UC Berkeley, September 2012.
- Z. Harchaoui: Presentation at ECML/PKDD Discovery Challenge, Bristol, September 2012.
- Z. Harchaoui: Seminar at Visual Geometry group, Oxford University, October 2012.
- C. Schmid: Workshop on Large Scale Multimedia Search, Los Angeles, January 2012.
- C. Schmid: Seminar at New York University, May 2012.

- C. Schmid: Seminar at Google, Zurich, May 2012.
- C. Schmid: Seminar at ETHZ, Zurich, May 2012.
- C. Schmid: Keynote speaker at ACM International Conference on Multimedia Retrieval (ICMR), Hong Kong, June 2012.
- C. Schmid: Keynote speaker at the International Symposium on Visual Computing, Crete, July 2012.
- C. Schmid: Tutorial on modern features at ECCV 2012, Florence, October 2012.
- C. Schmid: First international workshop on visual analysis and geo-localizaton of large-scale imagery in conjunction with ECCVÕ12, Florence, October 2012.
- C. Schmid: Keynote speaker at GdR ISIS, Le Touquet, November 2012.
- C. Schmid: Seminar at UC Berkeley, December 2012.

Popular Science

- The work ``What Makes Paris Look like Paris" on the automatic mining of visual architectural elements (Doersch~{\it et al.}~SIGGRAPH~\cite{Doersch2012}) has received broad press coverage including magazines \href{http://blogs.wsj.com/tech-europe/2012/08/13/how-paris-is-easier-to-recognize-than-a-u-s-city/?mod=google_news_blog}{The Wall Street Journal} and \href{http://www.newscientist.com/blogs/onepercent/2012/06/software-knows-what-makes-pari.html}{NewScientist}.
- The Pompeii project was featured in the special issue of "Cahiers Science \& Vie", on "the lost worlds" ("les mondes perdus"), N130, June 2012.

Events, Workshops, Conferences, Seminars, Editorial boards

- Laptev, J. Ponce, C. Schmid and J. Sivic co-organized a series of one week summer schools on computer vision and machine learning held in alternation between Inria Grenoble (\href{http://www.di.ens.fr/willow/events/cvml2010/}{2010}, \href{http://www.di.ens.fr/willow/events/cvml2012/}{2012}) and ENS Paris (\href{http://www.di.ens.fr/willow/events/cvml2011/}{2011},\href{http://www.di.ens.fr/willow/events/cvml2013/}{2013}). Over the four years the summer schools attracted over 600 participants from 34 countries including France and other European countries but also Australia, Brazil, Canada, China, India, Iran, Israel, Japan, Malaysia, Mexico, Saudi Arabia, Singapore, South Korea, Turkey and USA. The summer school provided an overview of the state of the art in visual recognition and machine learning. Lectures were complemented by practical sessions to provide participants with hands-on experience with the discussed material.
- I. Laptev, Co-organizer of the Workshop on Gesture Recognition at CVPR 2011. http://clopinet.com/isabelle/Projects/CVPR2011
- A. Zisserman, Co-organizer of the PASCAL Visual Object Classes Challenge 2011 (VOC2011). http://pascallin.ecs.soton.ac.uk/challenges/VOC/voc2011/
- A. Zisserman, Co-organizer of the PASCAL VOC 2011 workshop at ICCV 2011 http://pascallin.ecs. soton.ac.uk/challenges/VOC/voc2011/workshop/index.html
- A. Zisserman, Co-organizer of the Mysore Park Workshop on Computer Vision 2011.
- I. Laptev, J. Sivic, Co-organizers of the First International Workshop on Action recognition and Pose Estimation in Still Images at ECCV 2012. http://vision.stanford.edu/apsi2012
- A. Zisserman, Co-organizer of the PASCAL Visual Object Classes Challenge 2012 (VOC2012). http://pascallin.ecs.soton.ac.uk/challenges/VOC/voc2012/
- A. Zisserman, Co-organizer of the PASCAL VOC 2012 workshop at ECCV 2012. http://pascallin.ecs.soton.ac.uk/challenges/VOC/voc2012/workshop/index.html
- J. Sivic: Area chair, IEEE Conference on Computer Vision and Pattern Recognition, 2011.
- I. Laptev: Area chair, IEEE International Conference on Automatic Face and Gesture Recognition, 2011.
- I. Laptev, J. Sivic, A. Zisserman: Area chairs, IEEE International Conference on Computer Vision, 2011.
- I. Laptev, J. Ponce, J. Sivic, A. Zisserman: Area chairs, European Conference on Computer

Vision, 2012.

- C. Schmid: Program chair, European Conference on Computer Vision, 2012.
- I. Laptev, J. Ponce, J. Sivic and C. Schmid: Area chairs, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2013
- C. Schmid and J. Sivic: Area chairs, IEEE International Conference on Computer Vision (ICCV), 2013
- C. Schmid: Area chair, Neural Information Processing Systems (NIPS), 2012
- I. Laptev, J. Ponce, C. Schmid, J. Sivic, A. Zisserman: Editorial board, International Journal of Computer Vision.
- C. Schmid, J. Ponce: Editorial board, Foundations and Trends in Computer Graphics and Vision, since 2005.
- J. Ponce: editorial board, SIAM Journal on Imaging Sciences.
- I. Laptev: editorial board, Image and Vision Computing Journal.

Highlights and Achievements

- Andrew Zisserman was awarded the Rank Prize for his ``Outstanding contributions to modern computer vision" http://www.rankprize.org/.
- The updated 2nd edition of the textbook “Computer Vision: A Modern Approach" by David Forsyth and Jean Ponce has been published by Pearson Education in November 2011.
- I. Laptev, J. Sivic: Inria Prime d'excellence scientifique.
- J. Ponce (2010) and C. Schmid (2012) were awarded Advanced ERC Grants.
- I. Laptev (2012) and J. Sivic (2013) were awarded Starting ERC Grants.
- J. Ponce became a senior member of the Institut Universitaire de France.
- C. Schmid was nominated IEEE Fellow, 2012.
- LEAR participated in the Multimedia Event Detection track of TrecVid 2012, one of the major benchmarks in automatic video analysis. We ranked 2-nd out of 17 participants for the pre-specified event category task, and first out of 13 participants on the ad-hoc event category task.

Tools and software

- PMVS (Patch-based Multi-view Stereo Software) package was developed in collaboration with Y. Furukawa at the University of Illinois at Urbana-Champaign.The software and its documentation are available at \url{http://www.di.ens.fr/pmvs/}. The software is distributed under GPL. A US patent corresponding to this software ``Match, Expand, and Filter Technique for Multi-View Stereopsis" was issued on December 11, 2012 and assigned patent number 8,331,615.
- Dense local space-time features STIP-2.0 http://www.irisa.fr/vista/Equipe/People/Laptev/download.html\#stip
- Automatic Alignment of Paintings to 3D models. The code for automatic alignment of paintings to a 3D model (Russell et al. 2011) was made publicly available in October 2012 at http://www.di.ens.fr/willow/research/paintingalignment/index.html .
- Finding actors and actions in movies. The code for automatic joint learning of people and their actions from video with aligned text~\cite{Bojanowski13} was made publicly available at http://www.di.ens.fr/willow/research/actoraction .

A−Brain

/www.msr-inria.fr/projects/a-brain/

Overview

Joint acquisition of neuroimaging and genetic data on large cohorts of subjects is a new approach used to assess and understand the variability that exists between individuals, and that has remained poorly understood so far. As both neuroimaging- and genetic-domain observations represent a huge amount of variables (of the order of millions), performing statistically rigorous analyses on such amounts of data is a major computational challenge that cannot be addressed with conventional computational techniques only. On the one hand, sophisticated regression techniques need to be used in order to perform sensitive analysis on these large datasets; on the other hand, the cost entailed by parameter optimization and statistical validation procedures (e.g. permutation tests) is very high.

The A-Brain (AzureBrain) Project started in October 2010 within the Microsoft Research-Inria Joint Research Center. It is co-led by Gabriel Antoniu, Head of the KerData Inria Project-Team (Rennes) and by Bertrand Thirion, Head of the Parietal Inria Project-Team (Saclay). The teams jointly address the computational problem described above using cloud related techniques on Microsoft Azure cloud infrastructure. The two teams bring together their complementary expertise: KerData in the area of scalable cloud data management, and Parietal in the field of neuroimaging, machine learning and statistical inference.
More specifically KerData designs and implements solutions to optimize data storage and management for the Map-Reduce programming model. This model has recently arisen as a very effective approach to develop high-performance applications over very large distributed systems such as grids and now clouds. KerData has recently proposed a set of algorithms for data management, combining versioning with decentralized metadata management to support scalable, efficient, fine-grain access to massive, distributed Binary Large OBjects (BLOBs) under heavy concurrency. The project investigates the benefits of integrating these algorithms (implemented inside the BlobSeer software library) with Microsoft Azure storage services and aims to evaluate the impact of using BlobSeer on Azure with large-scale application experiments such as the genetics-neuroimaging data comparisons addressed by Parietal. Parietal, on the other hand, is the core developer of the scikit-learn, a Python machine learning toolbox that contains state-of-the-art tools to solve many supervised and unsupervised data analysis problems. It contains the building blocks to design efficient procedures for statistical inference on neuroimaging-genetics data. Some efforts have yet to be dedicated to optimize the statistical models for the sake of sensitivity (use of robust regression, spatial regularization of the results in the image domain, design of priors in multivariate models) and also the efficiency of the computations involved when permutation tests are to be performed. Parietal has access to one of the largest neuroimaging-genetics dataset available currently, namely that of the Imagen consortium.

This project brings together researchers from algorithmic and statistical analysis domain on the one hand, and researchers involved in the organization of data management in intensive computation on the other hand, to work on the Microsoft Azure platform in order to unveil the relationships between genes and brain characteristics.

Highlights

The TomusBlobs data-storage layer developed in the framework of the A-Brain project was demonstrated to scale up to 1000 cores on 3 Azure data centers; it exhibits improvements in execution time up to 50% compared to standard solutions based on Azure BLOB storage.
The consortium has provided the first statistical evidence of the heritability of functional signals in a

failed stop task in basal ganglia, using a ridge regression approach, while relying on the Azure cloud to address the computational burden.

Research

TomusBlobs: a concurrency-optimized data storage layer for data-intensive applications running on Azure clouds

Participants : Radu Tudoran (Microsoft Research-Inria), Alexandru Costan (Inria), Gabriel Antoniu (Inria), Hakan Sonku (Microsoft Research), Goetz Brasche (Microsoft Research)

Enabling high-throughput massive data processing on cloud data becomes a critical issue, as it impacts the overall application performance. In the framework of the Microsoft Research-Inria A-Brain project, the TomusBlobs storage library was designed and implemented by KerData to address such challenges at the level of the cloud storage. The system we introduce is a concurrency-optimized data storage system which federates the virtual disks associated to VMs. As TomusBlobs does not require modifications to the cloud middleware, it can serve as a high-throughput globally-shared data storage for the cloud applications that require data passing among computation nodes.

We leveraged the performance of this solution to enable efficient data-intensive processing on commercial clouds by building an optimized prototype MapReduce framework for Azure. The system, deployed on 1000 cores in Azure distributed across 3 datacenters, was used to execute the A-Brain application with the goal of searching for significant associations between brain locations and genes.

The achieved throughput increased by 2 for reading and by 3 for writing compared to the case of using remote Azure storage. With our approach for MapReduce data processing, the computation time is reduced to 50 % compared to the baseline Azure solutions, while the cost is reduced up to 30 %.

Iterative MapReduce: an extension of the MapReduce programming model

Participants : Radu Tudoran (Microsoft Research-Inria), Alexandru Costan (Inria), Gabriel Antoniu (Inria), Louis-Claude Canon (Microsoft Research-Inria)

While MapReduce has arisen as a major programming model for data analysis on clouds, there are many scientific applications that require processing patterns different from this paradigm. As such, reduce-intensive algorithms are becoming increasingly useful in applications such as data clustering, classification and mining. These algorithms have a common pattern: data are processed iteratively and aggregated into a single final result. While in the initial MapReduce proposal the reduce phase was a simple aggregation function, recently an increasing number of applications relying on MapReduce exhibit a reduce-intensive pattern, that is, an important part of the computations are done during the reduce phase. However, platforms like MapReduce or Dryad lack built-in support for reduce-intensive workloads.

To overcome these issues, we introduced MapIterativeReduce, a framework which: 1) extends the MapReduce programming model to better support reduce-intensive applications by exploiting the inherent parallelism of the reduce tasks which have an associative and/or commutative operation; and 2) substantially improves their efficiency by eliminating the implicit barrier between the Map and the Reduce phase. We showed how to leverage this architecture for scientific applications by enhancing the fault tolerance support in Azure and TomusBlobs, the underlying storage system, with a light checkpointing scheme and without any centralized control.

We evaluated MapIterativeReduce on the Microsoft Azure cloud with synthetic benchmarks and with the A-Brain application. Compared to state-of-art solutions, our approach enables faster data processing, by reducing the execution times by up to 75 %.

Adaptive file management for clouds

Participants : Radu Tudoran (**Microsoft Research-Inria**), Alexandru Costan (**Inria**), Ramin Rezai Rad (**Microsoft Research**), Goetz Brasche (**Microsoft Research**), Gabriel Antoniu (**Inria**)

Recently, there is an increasing interest to execute general data processing schemas in clouds, as it would allow many scientific applications to migrate to this computing infrastructure. The natural way to do this is to design and adopt Workflow Processing engines built for clouds. Such workflow processing in clouds would involve data propagation on the computation nodes based on well-defined data access patterns. Having an efficient file management backend for a workflow engine is thus essential as we move to the world of Big Data.

Scientific workflows typically communicate data between tasks using files. Currently, on public clouds, this is achieved by using the cloud storage services, which are unable to exploit the workflow semantics and are subject to low throughput and high latencies. To overcome these limitations, we proposed a new approach for a transfer-optimized file management in clouds. On the one hand, our solution manages files within the deployment leveraging data locality. On the other hand, we envision an adaptive system that adopts the transfer method most suited based on the data transfer context. We rely on the observation that workflows generate a set of common data access patterns that our solution exploits in conjunction with context information to self-adapt, choose the most adequate transfer protocol and expose the data layout within the virtual machines to the workflow engines.

This file management system was integrated within the Microsoft Generic Worker workflow engine and was validated using synthetic benchmarks and then with the A-Brain application on the Azure cloud. The results show it can bring significant performance gains: up to 5x file transfer speedup compared to solutions based on standard cloud storage and over 25% application timespan reduction compared to Hadoop on Azure. This work was initiated in the context of a 3-month internship of Radu Tudoran hosted by the Advance Technology Lab from Microsoft Europe, Germany, Aachen.

Random Parcellation-based Inference
Participants: Benoit da Mota, Virgile Fritsch, Gaël Varoquaux, Bertrand Thirion

The current state-of-the-art approach for spatially-informed statistical tests in neuroimaging studies is the so called Threshold-Free Cluster Enhancement (TFCE) method developed at FMRIB. Inspired by our recent success in sparse recovery, we have investigated whether the use of randomized parcellations could benefit to group studies; the computation of statistics on such parcellations can be understood as an anisotropic smoothing performed during the statistical inference to smooth out some of the noise while preserving the main structures of the signal. We have shown empirically that this approach yields indeed an increased sensitivity over existing approaches, such as TFCE, in group analyses, both in neuroimaging and neuroimaging-genetics.

Robust regression for neuroimaging genetics
Participants: Benoit da Mota, Virgile Fritsch, Gaël Varoquaux, Bertrand Thirion

When we shift from small to large cohorts, such as those encountered in neuroimaging-genetics ($\approx$ 2000 subjects), then comes naturally the question of characterizing the statistical structure of the population. We have tackled the question through the point of view of outlier detection, which is one of the main violations of the classical Gaussian hypothesis. Detecting rigorously outliers in high-dimensional datasets is still an open question, hence we have focused on robust regression, that directly addresses the consequence of non-Gaussianity on regression problems. We have thus revisited and improved this approach in the context of neuroimaging-genetics studies. We show that the introduction of robustness in statistical inference has a strong impact on both the sensitivity and the specificity of the analysis.

Multivariate models for Neuroimaging-genetics
Participants: Benoit da Mota, Bertrand Thirion

The quantitative evaluation of statistical models with machine learning techniques represents an important step in the comprehension of the associations between brain image phenotypes and genetic data. Such approaches require cross validation loops to set the hyper-parameters and for performance evaluation. Permutations have to be used to assess the statistical significance of the results, thus yielding prohibitively expensive analyses. We have presented a framework that can deal with such a computational burden. It is based on ridge regression for the sake of efficiency.
The results confirm that brain activation signals, such as those from sub-cortical nuclei recoded during a failed stop task, are a heritable feature.

Publications

Journals

Alexandru Costan, Radu Tudoran, Gabriel Antoniu, Goetz Brasche. TomusBlobs : Scalable Data-intensive Processing on Azure Clouds. Concurrency and Computation Practice and Experience, Wiley, 2013. URL: http://onlinelibrary.wiley.com/doi/10.1002/cpe.3034/abstract.

Benoit Da Mota, Virgile Fritscha, Gaël Varoquaux, Tobias Banaschewski, Gareth J. Barker , Arun L.W. Bokde, Uli Bromberg , Patricia Conrod, Jürgen Gallinat, Hugh Garavan, Jean-Luc Martinot, Frauke Nees, Tomas Pausl, Zdenka Pausova , Marcella Rietschel, Michael N. Smolka, Andreas Ströhle, Vincent Frouin, Jean-Baptiste Poline, Bertrand Thirion, the IMAGEN consortium. Randomized Parcellation Based Inference. neuroImage, elsevier, in Press

Electronic Journals

Gabriel Antoniu, Alexandru Costan, Benoit Da Mota, Bertrand Thirion, Radu Tudoran. A-Brain: Using the Cloud to Understand the Impact of Genetic Variability on the Brain. ERCIM News, April 2012.

Conferences and workshops

Radu Tudoran, Alexandru Costan, Ramin Rezai Rad, Goetz Brasche and Gabriel Antoniu. Adaptive File Management for Scientific Workflows on the Azure Cloud. IEEE International Conference on Big Data (IEEE BigData 2013), October 6-9, 2013, Santa Clara, CA, USA. Acceptance rate: 17%.

Radu Tudoran, Alexandru Costan, Gabriel Antoniu. DataSteward : Using Dedicated Compute Nodes for Scalable Data Management on Public Clouds. In Proc. of ISPA 2013- 11th IEEE International Symposium on Parallel and Distributed Processing with Applications, Melbourne, Australia, July 2013.

Benoit da Mota, Virgile Fritsch, Gaël Varoquaux, Vincent Frouin, Jean-Baptiste Poline, and Bertrand Thirion. Distributed High-Dimensional Regression with Shared Memory for Neuroimaging-Genetic Studies. in Euroscipy 2013.

Benoit Da Mota, Virgile Fritsch, Gaël Varoquaux, Vincent Frouin, Jean-Baptiste Poline, and Bertrand Thirion. Enhancing the Reproducibility of Group Analysis with Randomized Brain Parcellations. In MICCAI - 16th International Conference on Medical Image Computing and Computer Assisted Intervention - 2013, Nagoya, Japan, June 2013.

Virgile Fritsch, Benoit Da Mota, Gaël Varoquaux, Vincent Frouin, Eva Loth, Jean-Baptiste Poline, and Bertrand Thirion. Robust Group-Level Inference in Neuroimaging Genetic Studies. In Pattern Recognition in Neuroimaging, Philadelphie, United States, May 2013.

Radu Tudoran, Alexandru Costan, Gabriel Antoniu, Hakan Soncu. “TomusBlobs: Towards

Communication-Efficient Storage for MapReduce Applications in Azure.” In Proc. 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'2012), Ottawa, Canada, May 2012.

Radu Tudoran, Alexandru Costan, Gabriel Antoniu, Luc Bougé. A Performance Evaluation of Azure and Nimbus Clouds for Scientific Applications. In Proc. CloudCP 2012 - 2nd International Workshop on Cloud Computing Platforms, Held in conjunction with the ACM SIGOPS Eurosys 12 conference, Bern, Switzerland, Apr 2012.

Radu Tudoran, Alexandru Costan, Benoit Da Mota, Gabriel Antoniu, Bertrand Thirion. A-Brain: Using the Cloud to Understand the Impact of Genetic Variability on the Brain. 2012 Cloud Futures Workshop, Berkeley, May 2012.

Radu Tudoran, Alexandru Costan, Gabriel Antoniu. MapIterativeReduce: A Framework for Reduction-Intensive Data Processing on Azure Clouds. Third International Workshop on MapReduce and its Applications (MAPREDUCE'12), held in conjunction with ACM HPDC'12., Delft, Netherlands, Jun 2012.

Benoit Da Mota, Vincent Frouin, Edouard Duchesnay, Soizic Laguitton, Gaël Varoquaux, Jean-Baptiste Poline, Bertrand Thirion. A fast computational framework for genome-wide association studies with neuroimaging data. 20th International Conference on Computational Statistics (COMPSTAT 2012), Lamissol, Cyprus, Aug 2012.

Benoit Da Mota, Michael Eickenberg, Soizic Laguittton, Vincent Frouin, Gaël Varoquaux, Jean-Baptiste Poline, Bertrand Thirion. A MapReduce Approach for Ridge Regression in Neuroimaging-Genetic Studies. Data- and Compute-Intensive Clinical and Translational Imaging Applications Workshop (DCICTIA-MICCAI'12), held in conjunction with the 15th International Conference on Medical Image Computing and Computer Assisted Intervention, Nice, France, Oct 2012,

Posters

Benoit Da Mota, Virgile Fritsch, Gaël Varoquaux, Vincent Frouin, Jean-Baptiste Poline, and Bertrand Thirion. Randomized Brain Parcellations boost the Sensitivity of Neuroimaging- Genetic Studies. 19th Annual Meeting of the Organization for Human Brain Mapping, 2013.

Radu Tudoran, Advisor: Gabriel Antoniu, Luc Bougé. SAGE: Geo-Distributed Streaming Data Analysis in Clouds. PhD Forum in conjunction with the 27th IEEE International Parallel &. Distributed Processing Symposium. May 20, 2013 - May 24, 2013. Boston, Massachusetts USA

Benoit Da Mota, Edouard Duchesnay, Vincent Frouin, Gaël Varoquaux, Jean-Baptiste Poline, Bertrand Thirion. Fast vGWAS with Correction for Multiple Tests using Map-Reduce in Cloud or HPC. 18th Annual Meeting of the Organization for Human Brain Mapping, 2012.

Benoit Da Mota, Vincent Frouin, Edouard Duchesnay, Soizic Laguitton, Gaël Varoquaux, Jean-Baptiste Poline, Bertrand Thirion. A fast computational framework for genome-wide association studies with neuroimaging data. 5th European meeting on Python in Science (EuroScipy 2012), Bruxelles, Belgium, 2012.

Software

TomusBlobs is a software library for concurrency-optimized data storage for data-intensive applications running on Azure clouds, including MapReduce applications. It is being developed by the KerData Inria Project-Team in the framework of the A-Brain Microsoft Research-Inria project.

The statistical analysis code used in the project, genim-stats, mostly contributed by Benoit da Mota, has been registered at APP on Sept. 1st, 2013

4D Cardiac MR images

www.msr-inria.fr/projects/4d-cardiac-mr-images/

This project started in 2011.

Objectives

First Approach : Feature-Based Indexing

The first approach to image indexation is based on learned discriminative image features. Cardiac MR contains a vast amount of rich spatiotemporal information about the shape, motion and appearance of the heart. Such information is "buried" within the voxels, and our technique will try to extract it and make it available in a more usable form.
State of the art in image based retrieval is focused on retrieval of natural scene images, object and action recognition. Several distinctive features are often extracted per image and matched with the features from a training set. Simonyan et al. [1] use this technique to rapidly retrieve "similar" X-ray images. As noted by André et al [2], similar approaches focus on some sort of appearance or shape similarity while disregarding the semantic meaning of the similarity between the images. Furthermore, there does seem to be no agreement on the definition of "image similarity" [3], as it is often task dependent.

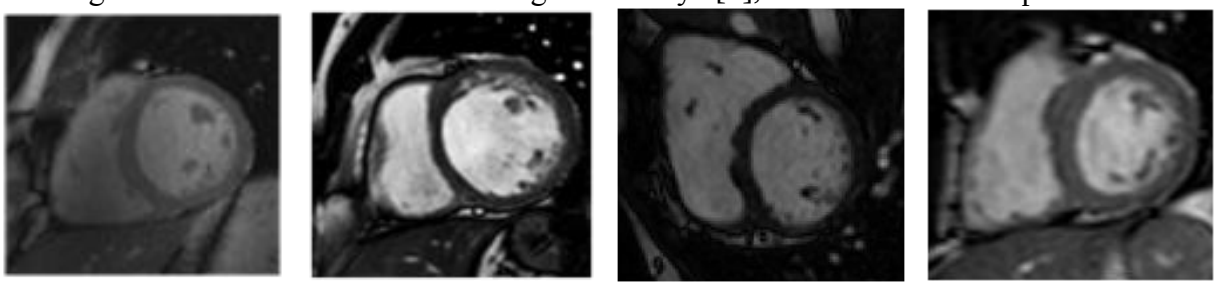

Figure 1. MR images of left and right ventricles in different patients. Different heart pathologies manifest themselves as differences in the shape and size of the ventricles. From left to right: normal heart, failing heart with arrhytmia, post repair Tetralogy of Fallot (enlarged right ventricle), heart with infarction.

Several image based indices can be used to estimate the severity of the disease, and plan the treatment. These are mainly based on volumetric measurements [4] (such as volume or mass) or motion information [5][6]. Few are dealing with shape characterization.
To describe the cardiac images, we will use shape, motion and appearance features at the same time. Using machine learning methods like decision forests it is possible to select the set of most discriminative features for the task at hand (e.g. detecting arrhythmia) [7]. Both static and dynamic features will be available to the algorithm and their relative importance for the retrieval task will be learned automatically from training data.
To stay clinically relevant, it is important to define the ranking of the images by their similarity according to some clinical criteria. For this an expert clinical input is needed for the learning step. This can be e.g. information on the pathology, severity of the disease or treatment outcome [8][9].

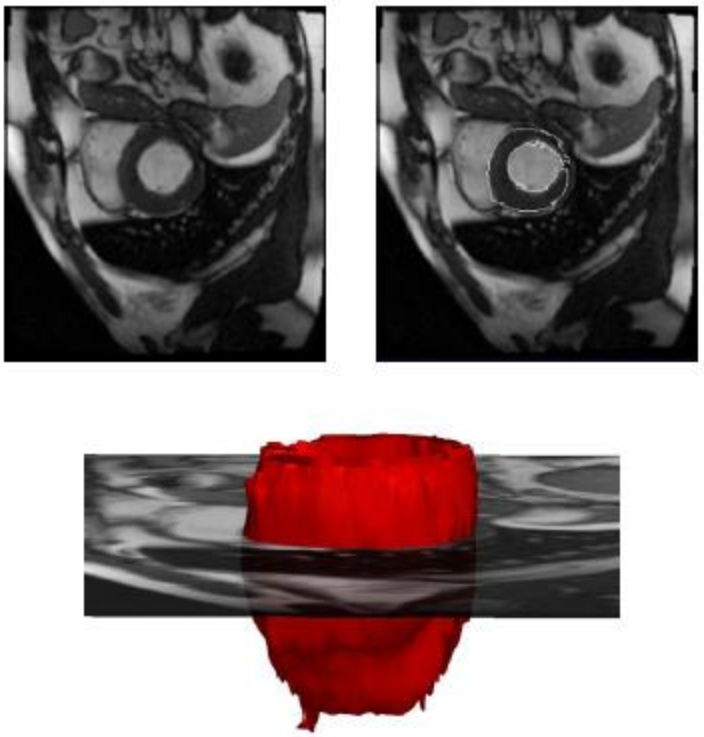

Figure 2 Example of an automatic method for left ventricle myocardium segmentation using decision forests. A random set of features was generated out of which a task specific subset was automatically selected during the learning procedure. This demonstrates the ability of the algorithm to automatically choose the right features for the given task.

The next step would be to add relevance feedback [10] where the user is proposed several most likely images and the system can then adapt its next query results based on past choices and outcomes. In order to remove the bias of individual cardiologists and generate the general knowledge, it would be necessary to create a collaborative tool and integrate inputs from several users, together with additional information (e.g. treatment history and outcome of the treatment).

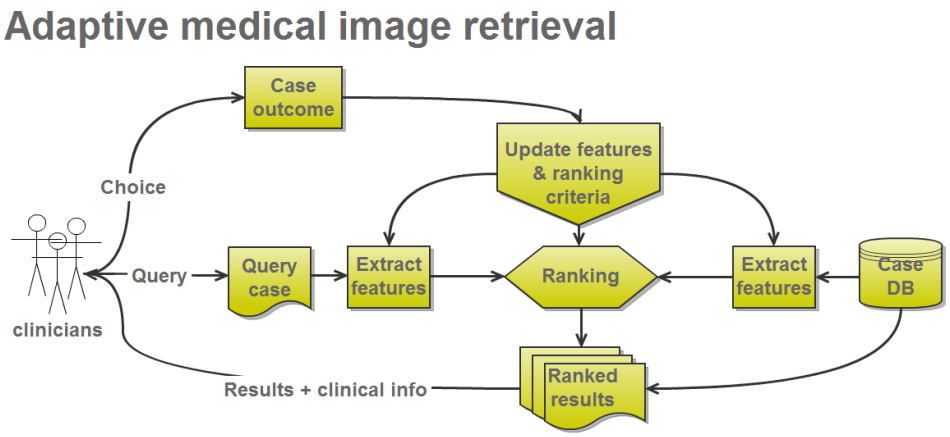

Figure 3 Workflow of an adaptive indexation and retrieval method.

Using such a system would be an evolution from traditional medical practice (based solely on experience of a single clinician) through an integrative medical practice (decisions based on agreement of several experts), towards evidence based medicine; the ultimate goal being getting closer to the best possible treatment.

We believe it is possible to detect a disease and estimate its severity automatically, by learning from a number of training exemplars. The exemplars and their associated expert annotations produce the sought after definition of diseases versus healthy patients. However, a lot can be done even in the case where supervision (expert labels) is missing.

Even in an unsupervised setting modern machine learning techniques provide the means to extract useful information from (relatively inexpensive) unlabeled data. For instance, if we represent images as data points in a high-dimensional space then we can use non-linear manifold learning approaches to map those

images optimally onto a low-dimensional space, a space where similar images are close to one another. In practice this would enable discovering groups of related patients automatically, without input from clinicians. Duchateau et al [11] use these techniques to discover similarities in cardiac septal motion abnormalities. Similarly, deep learning with auto-encoders [12] has recently shown some interesting results in the field of unsupervised document classification and retrieval. Such an approach may potentially discover new, previously unidentified patterns of diseases.

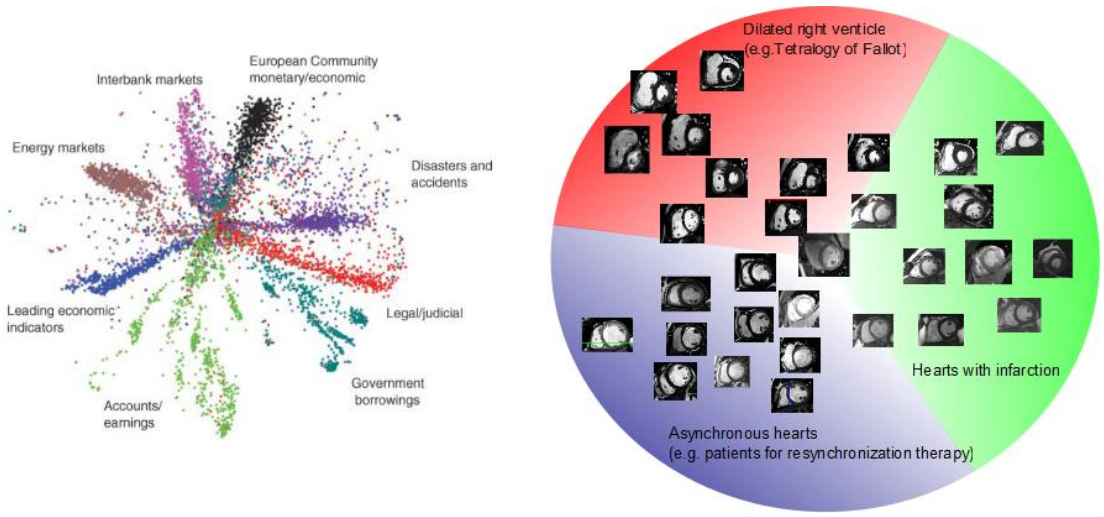

Figure 4. Unsupervised learning of text data and medical images. (left) Hinton's unsupervised clustering using deep learning for text documents (projection into a 2D space) [11]. (right) Adaptation of the scheme for cardiac image retrieval (more intense color means more severe disease, healthy heart should be found in the center of the plot). See also [1].

Second Approach: Biophysical Models.

The second approach to cardiac image analysis and indexation is based on the biophysical modeling of the cardiac function. Biophysical models of the heart aim to improve the understanding of the cardio-vascular system by performing numerical simulation of the cardiac function from its mathematical description [13, 14]. Simulations of the whole organ have reached such a degree of realism that it is now possible to compare them quantitatively with available cardiac images and signals acquired routinely on patients [15].
In this approach, the first step is to adapt the generic biophysical model to the actual, individual patient. This process consists in estimating the values of the many parameters involved and goes under the name of model personalization.
The retrieval task is now simply accomplished by comparing those parameters with existing patients in a database. An important difference from the first approach is that such parameter vector can be readily interpreted in terms of biophysical and clinical parameters (e.g. regional contractility or stiffness of the myocardium, blood pressure in cavities, and duration of the filling, ejection and isovolumetric phases). Solving the model personalization problem has implications which go well beyond image retrieval. For instance, personalized biophysical models of the heart may predict the cardiac function of a given patient upon the occurrence of a certain pathology or therapy. This is why a new vision has recently developed to help the planning of therapies [16] such as Cardiac Resynchronization Therapy [17] or Radiofrequency ablation [18].

Plan of action

We plan to use a semi-supervised image indexation and retrieval approach. The training set will consist of 4D MR image sequences. For some such sequences the biophysical parameters are known and for many others they are not. Using efficient transductive learning approaches we plan to transfer the existing labels to the unlabeled data points so as to generate a more complete labeled database [25].
Since the success of this machine learning technique partially lies in the size and the quality of the

database, a first stage consists in creating a large number of synthetic cardiac sequences (with known biophysical parameters) from a limited set of image sequences coarsely sampling the space of pathological or normal cases. The strategy of generating a large database of synthetic training data is similar to that used to build the Microsoft Xbox Kinect system [26] and inspired by its success. But synthesizing realistic spatio-temporal images of hearts is not easy. Here is our proposed approach:

- From a dataset of cardiac sequence images (typically cine-MR images available in the Cardiac Atlas Project (www.cardiacatlas.org), the right and left ventricles must be segmented and their motion tracked over the cardiac cycle. In this approach, we plan to modify already existing methods in order to assess the uncertainty in the localization of the endocardial and epicardial surfaces during the cardiac cycle. This uncertainty may be characterized as covariance matrices at each segmentation voxels and is an important factor to take into account when performing the model personalization.

- The simulation of cardiac electromechanical models usually require large computational meshes (from 10 000 nodes to 1 million) and potentially a large number of parameters (in the worst case, several times the number of nodes). This leads to large computational times but also to a non-compact description of the biophysical models. Similarly to feature selection in machine learning methods, we plan to work on model reduction in order to drastically reduce the size of the state vector and the number of parameters of each biophysical model. The computation of reduced basis may rely on the analysis of simulated cardiac meshes or segmented cardiac sequences, as proposed in computational mechanics with the Proper Orthogonal Decomposition method [19]. Non-linear approaches using manifold learning may help at this stage to define optimal reduced models leading to efficient computation.

- For each cardiac sequence, the cardiac segmentation and its uncertainty will serve as input to a personalization algorithm. This entails solving an inverse problem which may be tackled in several ways such as sequential or filtering approaches [20, 21], optimization based on the adjoint formulation [22, 23], and optimization based on gradient-free methods. Thanks to reduced models, the parameter optimization should be greatly improved, avoiding the curse of dimensionality often encountered in such domains.

- After model personalization, we can associate with each image sequence a compact set of parameters. However, the number of cardiac image sequences processed will be necessarily limited by the amount of resources and time available in this project. Therefore, we plan to take advantage of the predictive power of those biophysical models by creating synthetic images from existing ones. This can be achieved by modifying the set of biophysical parameters of the personalized models in order to artificially create pathologies (infracted regions in the myocardium, delays in electrophysiology propagation…). With each new simulated cardiac motion, one can create a synthetic yet realistic image sequence by warping a 4D texture image extracted from already acquired MR images [24].

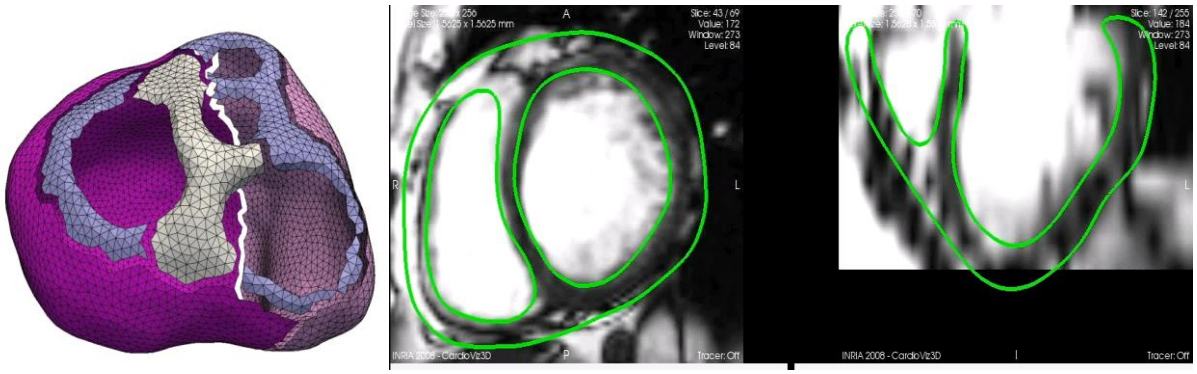

Figure 5. Biophysical heart models. (Left) Computational mesh of the myocardium including the right and left ventricles and the four valves; (Right) Personalized cardiac mesh on a cine-MR image sequence; the mesh geometry is adapted to the anatomy of the patient but also its biophysical parameters such that its motion matches the one visible in images.

Once the database of cardiac image sequences with their biophysical parameters has been created, a second stage consists in devising a semi-supervised machine learning algorithm for the indexation of cardiac images.

10.2 Background references

[1] K. Simonyan, A. Zisserman, and A. Criminisi, “Immediate Structured Visual Search for Medical Images,” MICCAI, 2011, pp. 1-8.

[2] B. André, T. Vercauteren, A.M. Buchner, M.B. Wallace, and N. Ayache, “A smart atlas for endomicroscopy using automated video retrieval.,” Medical image analysis, vol. 15, Aug. 2011, pp. 460-76.

[3] S. Sedghi, M. Sanderson, and P. Clough, “A study on the relevance criteria for medical images,” Pattern Recognition Letters, vol. 29, Nov. 2008, pp. 2046-2057.

[4] G. de Simone, J.S. Gottdiener, M. Chinali, and M.S. Maurer, “Left ventricular mass predicts heart failure not related to previous myocardial infarction: the Cardiovascular Health Study.,” European heart journal, vol. 29, Mar. 2008, pp. 741-7.

[5] R. Kumar, F. Wang, and D. Beymer, “Cardiac disease detection from echocardiogram using edge filtered scale-invariant motion features,” Computer Vision and.

[6] N. Duchateau, M. De Craene, G. Piella, E. Silva, A. Doltra, M. Sitges, B.H. Bijnens, and A.F. Frangi, “A spatiotemporal statistical atlas of motion for the quantification of abnormal myocardial tissue velocities.,” Medical image analysis, vol. 15, Jun. 2011, pp. 316-28.

[7] E. Geremia, O. Clatz, B.H. Menze, E. Konukoglu, A. Criminisi, and N. Ayache, “Spatial decision forests for MS lesion segmentation in multi-channel magnetic resonance images.,” NeuroImage, Apr. 2011.

[8] E. Troost, B. Meyns, W. Daenen, F. Van de Werf, M. Gewillig, K. Van Deyk, P. Moons, and W. Budts, “Homograft survival after tetralogy of Fallot repair: determinants of accelerated homograft degeneration.,” European heart journal, vol. 28, Oct. 2007, pp. 2503-9.

[9] L. Grosse-Wortmann and A. Redington, “Doing the right thing at the right time: is there more to pulmonary valve replacement than meets the eye?,” European heart journal, vol. 30, Sep. 2009, pp. 2076-8.

[10] C. Manning, Christopher D.(Stanford University, I.) Raghavan, Prabhakar (Yahoo, and H. (Universität S. Schütze, Introduction to information retrieval, Cambridge University Press, 2008.

[11] N. Duchateau, M.D. Craene, G. Piella, and A.F. Frangi, “Characterizing Pathological Deviations from Normality using Constrained Manifold-Learning,” MICCAI 2011 (to be published), 2011, p. 8.

[12] G.E. Hinton and R.R. Salakhutdinov, “Reducing the dimensionality of data with neural networks.,” Science (New York, N.Y.), vol. 313, Jul. 2006, pp. 504-7.

[13] Hunter, P., Coveney, P., and et al. (2010). A vision and strategy for the VPH in 2010 and beyond. Phil. Trans. R. Soc, pages 2595-2614

[14] Noble, D. (2002). Modeling the heart-from genes to cells to the whole organ, Science, 295(5560):1678-1682.

[15] Sermesant, M., Peyrat, J. M., Chinchapatnam, P., Billet, F., Mansi, T., Rhode, K., Delingette, H., Razavi, R., and Ayache, N. (2008). Toward patient-specfic myocardial models of the heart. Heart Failure Clinics, 4(3):289-301.

[16] Smith, A. de Vecchi, M. McCormick, D. Nordsletten, O. Camara, A.F. Frangi, H. Delingette, M. Sermesant, J. Relan, N. Ayache, M. W. Krueger, W. Schulze, R. Hose, I. Valverde, P. Beerbaum, C. Staicu, M. Siebes, J. Spaan, P. Hunter, J. Weese, H. Lehmann, D. Chapelle, and R. Razavi. euHeart: Personalized and integrated cardiac care using patient-specific cardiovascular modelling. Journal of the Royal Society Interface Focus, 1(3):349-364, 2011.

[17] M. Sermesant, F. Billet, R. Chabiniok, T. Mansi, P. Chinchapatnam, P. Moireau, J. Peyrat, K. Rhode, M. Ginks, P. Lambiase, S. Arridge, H. Delingette, M. Sorine, C. Aldo Rinaldi, D. Chapelle, R. Razavi, and N. Ayache, “Personalised electromechanical model of the heart for the prediction of the acute effects of cardiac resynchronisation therapy,” in Proceedings of FIMH 2009, 2009, pp. 239–248.

[18] Jatin Relan, Phani Chinchapatnam, Maxime Sermesant, Kawal Rhode, Matt Ginks, Hervé Delingette, C. Aldo Rinaldi, Reza Razavi, and Nicholas Ayache. Coupled Personalization of Cardiac Electrophysiology Models for Prediction of Ischaemic Ventricular Tachycardia. Journal of the Royal Society Interface Focus, 1(3):396-407, 2011

[19] K. Kunisch and S. Volkwein. Galerkin proper orthogonal decomposition methods for parabolic problems. Numerische Mathematik, 90:117–148, 2001.

[20] J. Xi, P. Lamata, J. Lee, P. Moireau, D. Chapelle, and N. Smith, “Myocardial transversely isotropic material parameter estimation from in-silico measurements based on reduced-order unscented kalman filter,” J. of the Mech. Behavior of Biomedical Materials, vol. In Press, 2011.

[21] R. Chabiniok, P. Moireau, P.-F. Lesault, A. Rahmouni, J.-F. Deux, and D. Chapelle, “Trials on tissue contractility estimation from cardiac cine mri using a biomechanical heart model,” in Proceedings of FIMH 2011, ser. LNCS, vol. 6666. Springer, 2011, pp. 304–312.

[22] Sundar, H., Davatzikos, C., and Biros, G. (2009). Biomechanically-constrained 4d estimation of myocardial motion. In MICCAI (1), pages 257-265.

[23] H. Delingette, F. Billet, K. C. L. Wong, M. Sermesant, K. Rhode, M. Ginks, C. A. Rinaldi, R. Razavi, and N. Ayache. Personalization of Cardiac Motion and Contractility from Images using Variational Data Assimilation. IEEE Transactions in Biomedical Engineering Letters, 2011. Note: In Press

[24] Adityo Prakosa, Maxime Sermesant, Hervé Delingette, Eric Saloux, Pascal Allain, Pascal Cathier, Patrick Etyngier, Nicolas Villain, and Nicholas Ayache. Synthetic Echocardiographic Image Sequences for Cardiac Inverse Electro-Kinematic Learning. In Proceedings of Medical Image

Computing and Computer Assisted Intervention (MICCAI), LNCS, Toronto, Canada, pages 8p, September 2011. Springer, Heidelberg.
[25] O. Chapelle, B. Scholkopf, A. Zien. Semi-supervised learning. MIT Press. 2006.
[26] J. Shotton, A. Fitzgibbon, M. Cook and A. Blake. Real-time Human Pose Recognition in Parts from Single Depth Images. CVPR 2011.

Given a large database of cardiac images of patients stored with an expert diagnosis, we aim to describe the appearance, shape and motion of the hearts and perform image content based retrieval and automatically select the most semantically similar cases to the cardiac images of a new patient.
Our second approach to cardiac image indexation is based on biophysical modeling of the cardiac function. The retrieval task is accomplished by comparing biophysical model parameters with existing patients in a database.

Highlight: an achievement

Cardiac image segmentation with decision forests

We extended our previous work for segmentation of multiple sclerosis lesions [1] to segmentation of cardiac structures from 3d and cinematic 3d+t magnetic resonance acquisitions. In particular we segment left atrium, left ventricle and its cavity.
We used spatiotemporal context rich features, measures based on vasculature and distances to blood pools. To deal with existing variations in between acquisitions, we proposed a two layer approach where image intensities and poses are standardized.

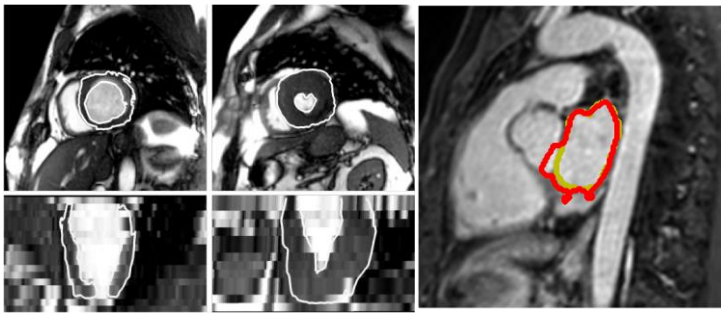

Figure 6 Segmentations of the left ventricle and atrium using decision forests

[1] Geremia, Ezequiel, et al. "Spatial decision forests for MS lesion segmentation in multi-channel magnetic resonance images." *NeuroImage* 57.2 (2011): 378-390.

Research

Segmentation of cardiac structures with decision forests

- A fully automated machine learning based left ventricle and left atrium segmentation algorithm
- Permutation based features for more robust image segmentation on magnetic resonance images
- Active segmentation of cardiac images with semi-supervised random forests
- Learning to segment from an ensemble of diverse ground-truths

Description of post-myocardial infarction hearts with semantically meaningful attributes

- Application of the relative attributes framework (Parikh, ICCV 2011) to characterize of post-myocardial infarction hearts by clinically meaningful attributes
- Crowd-sourcing cardiac image annotation

Cardiac model parameter estimation via machine learning and currents

- Describing cardiac meshes with 4d mathematical currents.
- Learning mapping between the current space and the cardiac model parameter space to enable cardiac model personalization.

Cardiac model personalization and decision making with full uncertainty treatment

- Perform the multiresolution registration of cardiac image sequence based on a sparse Bayesian Model.
- Computation of the posterior probability of the registration to estimate uncertainty in the cardiac motion estimation.

Cardiac view recognition

- A machine learning algorithm for automatic recognition of standard cardiac acquisition planes with 90% accuracy.

Cardiac functional class estimation

- Estimation of cardiac functional classes from temporal chamber synchronicity image signatures

Publications

Jan Margeta, Kristin McLeod, Antonio Criminisi, Nicholas Ayache: ***Decision forests for segmentation of left atrium from 3D MRI*.** 4th International Workshop on Statistical Atlases and Computational Models of the Heart at MICCAI 2013 (2013)

Avan Suinesiaputra, Brett R. Cowan, Ahmed O. Al-Agamy, Mustafa A. Alattar, Nicholas Ayache, Ahmed S. Fahmy, Ayman M. Khalifa, Pau Medrano-Gracia, Marie-Pierre Jolly, Alan H. Kadish, Daniel C. Lee, Jan Margeta, Simon K. Warfield, Alistair A. Young: ***A Collaborative Resource to Build Consensus for Automated Left Ventricular Segmentation of Cardiac MR Images***, Medical Image Analysis (In press 2013)

Loïc Le Folgoc, Hervé Delingette, Antonio Criminisi, Nicholas Ayache, **Current-based 4D shape analysis for the mechanical personalization of heart models.** In *Medical Computer Vision. Recognition Techniques and Applications in Medical Imaging* (pp. 283-292). Springer Berlin Heidelberg (2013).

Jan Margeta, Ezequiel Geremia, Antonio Criminisi, Nicholas Ayache: **Layered spatio-temporal forests for left ventricle segmentation from 4D cardiac MRI data.** 2nd International Workshop on Statistical Atlases and Computational Models of the Heart at MICCAI 2011, LNCS 7085 (2011) 109-119

Article submitted but not accepted :

Loic Le Folgoc, Hervé Delingette, Antonio Criminisi, Nicholas Ayache, **Robust Registration with Uncertainty using Sparse Bayesian Regression**, International Conference in Computer Vision, ICCV 2013

Software/tools

- Generic implementation of decision forests for image segmentation, regression, ranking, distance approximation and density estimation
- A web application for crowd-sourced of cardiac image annotation deployed on Windows Azure
- Cardiac image conversion and preprocessing pipeline
- Registration of cardiac sequences with uncertainty

Talks & Conference attendance

7th conference on Functional Imaging and Modeling of the Heart (FIMH 2013) – N.Ayache (keynote).
UCL-French Embassy workshop and conference-debate Engineering for wellbeing - N. Ayache (keynote), L. Le Folgoc, Ján Margeta (short presentations on their research)
Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2012) - Nicholas Ayache (general chair), Hervé Delingette (program chair). Sponsored by Microsoft Research
Workshop on Medical Computer Vision (MCV 2012) L. Le Folgoc (oral presentation), Antonio Criminisi (organiser), J. Margeta (program committee member).
Workshop on Medical Content-Based Retrieval for Clinical Decision Support (MCBR-CDS 2012).
Workshop on Machine Learning Medical Imaging (MLMI 2012)
International Conference on Computer Vision (ICCV 2011)
Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2011)
International Workshop on Statistical Atlases and Computational Models of the Heart (STACOM 2011) J. Margeta (oral presentation)

Appendix: Software releases

- Mathematical Components
    - SSReflect distribution http://ssr.msr-inria.inria.fr/FTP/ssreflect-1.4-coq8.4.tar.gz
    - Feit-Thompson proof code http://gforge.inria.fr/frs/?group_id=401
- Secure Distributed Computations And Their Proofs
    - Reference TLS implementation http://mitls.inria-rocq.inria.fr
    - F* compiler http://research.microsoft.com/fstar
    - ZQL library
- Tools for Proofs
    - TLA+ Proof System http://tla.msr-inria.inria.fr/tlaps/content/Download/Binaries.html
- Dynamic Dictionary of Mathematical Functions
    - DDMF http://ddmf.msr-inria.inria.fr/1.9.1/ddmf
- ADAPT project
    - Inputs to constraint solving toolkit Gecode http://www.gecode.org/
    - Bandit Ensemble for Parallel SAT Solving
    - MCTS integration in Constraint Programming
- A-Brain
    - TomusBlob: middleware for file system abstraction on top of cloud platform
- Cardiac Images
    - Web-based tool for crowd-sourced cardiac image annotation