

Inria

RESEARCH CENTER

FIELD

Perception, Cognition and Interaction

Activity Report 2019

Section Application Domains

Edition: 2020-03-21

DATA AND KNOWLEDGE REPRESENTATION AND PROCESSING

1. CEDAR Project-Team	5
2. GRAPHIK Project-Team	6
3. LACODAM Project-Team	7
4. LINKS Project-Team	11
5. MAGNET Project-Team	12
6. MOEX Project-Team (section vide)	13
7. ORPAILLEUR Project-Team	14
8. PETRUS Project-Team	15
9. TYREX Project-Team	16
10. VALDA Project-Team	17
11. WIMMICS Project-Team	19
12. ZENITH Project-Team	22

INTERACTION AND VISUALIZATION

13. ALICE Team	24
14. AVIZ Project-Team	25
15. EX-SITU Project-Team	26
16. GRAPHDECO Project-Team (section vide)	27
17. HYBRID Project-Team	28
18. ILDA Project-Team	29
19. IMAGINE Project-Team	30
20. LOKI Project-Team	31
21. MANAO Project-Team	32
22. MAVERICK Project-Team	33
23. MFX Project-Team	34
24. MIMETIC Project-Team	35
25. POTIOC Project-Team	39
26. TITANE Project-Team	40

LANGUAGE, SPEECH AND AUDIO

27. ALMANACH Project-Team	41
28. COML Team	42
29. MULTISPEECH Project-Team	43
30. PANAMA Project-Team	45
31. SEMAGRAMME Project-Team	46

ROBOTICS AND SMART ENVIRONMENTS

32. Auctus Team	47
33. CHORALE Team	48
34. CHROMA Project-Team	49
35. DEFROST Project-Team	51
36. FLOWERS Project-Team	54
37. HEPHAISTOS Project-Team (section vide)	56

38. LARSEN Project-Team	57
39. PERVASIVE Project-Team	58
40. RAINBOW Project-Team (section vide)	61
41. RITS Project-Team	62
VISION, PERCEPTION AND MULTIMEDIA INTERPRETATION	
42. LINKMEDIA Project-Team	64
43. MAGRIT Team	65
44. MORPHEO Project-Team	66
45. PERCEPTION Project-Team (section vide)	67
46. SIROCCO Project-Team	68
47. Stars Project-Team (section vide)	70
48. THOTH Project-Team	71
49. WILLOW Team	72

CEDAR Project-Team

4. Application Domains

4.1. Cloud Computing

Cloud computing services are strongly developing and more and more companies and institutions resort to running their computations in the cloud, in order to avoid the hassle of running their own infrastructure. Today's cloud service providers guarantee machine availabilities in their Service Level Agreement (SLA), without any guarantees on performance measures according to a specific cost budget. Running analytics on big data systems require the user not to only reserve the suitable cloud instances over which the big data system will be running, but also setting many system parameters like the degree of parallelism and granularity of scheduling. Choosing values for these parameters, and choosing cloud instances need to meet user objectives regarding latency, throughput and cost measures, which is a complex task if it's done manually by the user. Hence, we need need to transform cloud service models from availability to user performance objective rises and leads to the problem of multi-objective optimization. Research carried out in the team within the ERC project "Big and Fast Data Analytics" aims to develop a novel optimization framework for providing guarantees on the performance while controlling the cost of data processing in the cloud.

4.2. Computational Journalism

Modern journalism increasingly relies on content management technologies in order to represent, store, and query source data and media objects themselves. Writing news articles increasingly requires consulting several sources, interpreting their findings in context, and crossing links between related sources of information. CEDARresearch results directly applicable to this area provide techniques and tools for rich Web content warehouse management. Within the ANR ContentCheck project, and also as part of our international collaboration with the AIST institute from Japan, we work on one hand, to lay down foundations for computational data journalism and fact checking, and also work to devise concrete algorithms and platforms to help journalists perform their work better and/or faster. This work is carried in collaboration with Le Monde's "Les Décodeurs".

On a related topic, heterogeneous data integration under a virtual graph abstract model is studied within the ICODA Inria project which has started in September 2017. There, we collaborate with Les Décodeurs as well as with Ouest France and Agence France Presse (AFP). The data and knowledge integration framework resulting from this work will support journalists' effort to organize and analyze their knowledge and exploit it in order to produce new content.

GRAPHIK Project-Team

4. Application Domains

4.1. Agronomy

Agronomy is a strong expertise domain in the area of Montpellier. Some members of GraphIK are INRA researchers (computer scientists). We closely collaborate with the Montpellier research laboratory IATE, a joint unit of INRA and other organisms. A major issue for INRA and more specifically IATE applications is modeling agrifood chains (i.e., the chain of all processes leading from the plants to the final products, including waste treatment). This modeling has several objectives. It provides better understanding of the processes from begin to end, which aids in decision making, with the aim of improving the quality of the products and decreasing the environmental impact. It also facilitates knowledge sharing between researchers, as well as the capitalization of expert knowledge and “know how”. This last point is particularly important in areas strongly related to local know how (like in cheese or wine making), where knowledge is transmitted by experience, with the risk of non-sustainability of the specific skills. An agrifood chain analysis is a highly complex procedure since it relies on numerous criteria of various types: environmental, economical, functional, sanitary, etc. Quality objectives involve different stakeholders, technicians, managers, professional organizations, end-users, public organizations, etc. Since the goals of the implied stakeholders may be divergent dedicated knowledge and representation techniques are to be employed.

4.2. Data Journalism

One of today’s major issues in data science is to design techniques and algorithms that allow analysts to efficiently infer useful information and knowledge by inspecting heterogeneous information sources, from structured data to unstructured content. We take data journalism as an emblematic use-case, which stands at the crossroad of multiple research fields: content analysis, data management, knowledge representation and reasoning, visualization and human-machine interaction. We are particularly interested in issues raised by the design of data and knowledge management systems that will support data journalism. These systems include an ontology (which typically expresses domain knowledge), heterogeneous data sources (provided with their own vocabulary and querying capabilities), and mappings that relate these data sources to the ontological vocabulary. Ontologies play a central role as they act both as a mediation layer that glue together pieces of knowledge extracted from data sources, and as an inference layer that allow to draw new knowledge.

Besides pure knowledge representation and reasoning issues, querying such systems raise issues at the crossroad of data and knowledge management. In particular, although mappings have been widely investigated in databases, they need to be revisited in the light of the reasoning capabilities enabled by the ontology. More generally, the consistency and the efficiency of the system cannot be ensured by considering the components of the system in isolation (i.e., the ontology, data sources and mappings), but require to study the interactions between these components and to consider the system as a whole.

LACODAM Project-Team

4. Application Domains

4.1. Introduction

The current period is extremely favorable for teams working in Data Science and Artificial Intelligence, and LACODAM is not the exception. We are eager to see our work applied in real world applications, and have thus an important activity in maintaining strong ties with industrial partners concerned with marketing and energy as well as public partners working on health, agriculture and environment.

4.2. Industry

We present below our industrial collaborations. Some are well established partnerships, while others are more recent collaborations with local industries that wish to reinforce their Data Science R&D with us (e.g. Energiency, Amossys).

- **Resource Consumption Analysis for Optimizing Energy Consumption and Practices in Industrial Factories (Energiency).** In order to increase their understanding of factory operation, companies introduce more and more sensors in their factories. Thus, the resource (electricity, water, etc.) consumption of engines, workshops and factories are recorded in the form of times series or temporal sequences. The person who is in charge of resource consumption optimization needs better software than classical spreadsheets for this purpose. He/she needs effective decision-aiding tools with statistical and artificial intelligence knowledge. The start-up Energiency aims at designing and offering such pieces of software for analyzing energy consumption. The CIFRE PhD thesis of Maël (defended 16/12/2019) aimed at proposing new approaches and solutions from the data mining field to tackle this issue.
- **Security (Amossys).** Current networks are faced with an increasing variety of attacks, from the classic “DDoS” that makes a server unusable for a few hours, to advanced attacks that silently infiltrate a network and exfiltrate sensitive information months or even years later. Such intrusions, called APT (Advanced Persistent Threat) are extremely hard to detect, and this will become even harder as most communications will be encrypted. A promising solution is to work on “behavioral analysis”, by discovering patterns based on the metadata of IP-packets. Such patterns can relate to an unusual sequencing of events, or to an unusual communication graph. Finding such complex patterns over a large volume of streaming data requires to revisit existing stream mining algorithms to dramatically improve their throughput, while guaranteeing a manageable false positive rate. We collaborated on this topic with the Amossys company and the EMSEC team of Irista through the co-supervision of the CIFRE PhD of Alban Siffer (located in the EMSEC team, defended 19/12/2019). Our goal was in particular to design novel anomaly detection methods making minimal assumptions on the data, and able to scale to real traffic volumes.
- **Car Sharing Data Analysis.** Peugeot-Citroën (PSA) group’s know-how encompasses all areas of the automotive industry, from production to distribution and services. Among others, its aim is to provide a car sharing service in many large cities. This service consists in providing a fleet of cars and a “free floating” system that allows users to use a vehicle, then drop it off at their convenience in the city. To optimize their fleet and the availability of the cars throughout the city, PSA needs to analyze the trajectory of the cars and understand the mobility needs and behavior of their users. We tackle this subject together through the CIFRE PhD of Gregory Martin.
- **Multimodal Data Analysis for the Supervision of Sensitive Sites.** ATERMES is an international mid-sized company with a strong expertise in high technology and system integration from the upstream design to the long-life maintenance cycle. It has recently developed a new product, called BARRIER TM (“Beacon Autonomous Reconnaissance Identification and Evaluation Response”),

which provides operational and tactical solutions for mastering borders and areas. Once in place, the system allows for a continuous night and day surveillance mission with a small crew in the most unexpected rugged terrain. The CIFRE PhD of Heng Zhang aims at developing a deep learning architecture and algorithms able to detect anomalies (mainly persons) from multimodal data. The data are “multimodal” because information about the same phenomenon can be acquired from different types of detectors, at different conditions, in multiple experiments.

- **Root Cause Analysis in Networks.** AdvisorSLA is a French company specialized in software solutions for network monitoring. For this purpose, the company relies on techniques of network metrology. By continuously measuring the state of the network, monitoring solutions detect events (e.g., overloaded router) that may degrade the network’s operation and the quality of the services running on top of it (e.g., video transmission could become choppy). When a monitoring solution detects a potentially problematic sequence of events, it triggers an alarm so that the network manager can take actions. Those actions can be preventive or corrective. Some statistics show that only 40% of the triggered alarms are conclusive, that is, they manage to signal a well-understood problem that requires an action from the network manager. This means that the remaining 60% are presumably false alarms. While false alarms do not hinder network operation, they do incur an important cost in terms of human resources. Thus, the CIFRE PhD of Erwan Bourrand proposes to characterize conclusive and false alarms. This will be achieved by designing automatic methods to “learn” the conditions that most likely precede each type of alarm, and therefore predict whether the alarm will be conclusive or not. This can help adjust existing monitoring solutions in order to improve their accuracy. Besides, it can help network managers automatically trace the causes of a problem in the network.

4.3. Health

- **Care Pathways Analysis for Supporting Pharmaco-Epidemiological Studies.** Pharmaco-epidemiology applies the methodologies developed in general epidemiology to answer to questions about the uses and effects of health products, drugs [31], [30] or medical devices [23], on population. In classical pharmaco-epidemiology studies, people who share common characteristics are recruited to build a dedicated prospective cohort. Then, meaningful data (drug exposures, diseases, etc.) are collected from the cohort within a defined period of time. Finally, a statistical analysis highlights the links (or the lack of links) between drug exposures and outcomes (e.g., adverse effects). The main drawback of prospective cohort studies is the time required to collect the data and to integrate them. Indeed, in some cases of health product safety, health authorities have to answer quickly to pharmaco-epidemiology questions.

New approaches of pharmaco-epidemiology consist in using large EHR (Electronic Health Records) databases to investigate the effects and uses (or misuses) of drugs in real conditions. The objective is to benefit from nationwide available data to answer accurately and in a short time pharmaco-epidemiological queries for national public health institutions. Despite the potential availability of the data, their size and complexity make their analysis long and tremendous. The challenge we tackle is the conception of a generic digital toolbox to support the efficient design of a broad range of pharmaco-epidemiology studies from EHR databases. We propose to use pattern mining algorithms and reasoning techniques to analyse the typical care pathways of specific groups of patients.

To answer the broad range of pharmaco-epidemiological queries from national public health institutions, the PEPS⁰ platform exploits, in secondary use, the French health cross-schemes insurance system, called SNDS. The SNDS covers most of the French population with a sliding period of 3 past years. The main characteristics of this data warehouse are described in [29]. Contrary to local hospital EHR or even to other national initiatives, the SNDS data warehouse covers a huge population. It makes possible studies on unfrequent drugs or diseases in real conditions of use. To tackle the volume and the diversity of the SNDS data warehouse, a research program has been established

⁰PEPS: Pharmaco-Epidémiologie et Produits de Santé – Pharmacoepidemiology of health products

to design an innovative toolbox. This research program is focused first on the modeling of care pathways from the SNDS database and, second, on the design of tools supporting the extraction of insights about massive and complex care pathways by clinicians. In such a database a care pathway is an individual sequence of drugs exposures, medical procedures and hospitalizations.

- **Care Sequences for the Exploration of Medico-administrative Data.** The difficulty of analyzing medico-administrative data is the semantic gap between the raw data (for example, database record about the delivery at date t of drug with ATC 2 code N 02BE01) and the nature of the events sought by clinicians (“was the patient exposed to a daily dose of paracetamol higher than 3g?”). The solution that is used by epidemiologists consists in enriching the data with new types of events that, on the one side, could be generated from raw data and on the other side, have a medical interpretation. Such new abstract events are defined by clinician using proxies. For example, drugs deliveries can be translated in periods of drug exposure (drug exposure is a time-dependent variable for non-random reasons) or identify patient stages of illness, etc. A proxy can be seen as an abstract description of a care sequence.

Currently, the clinicians are limited in the expression of these proxies both by the coarse expressivity of their tools and by the need to process efficiently large amount of data. From a semantic point of view, care sequences must fully integrate the temporal and taxonomic dimensions of the data to provide significant expression power. From a computational point of view, the methods employed must make it possible to efficiently handle large amounts of data (several millions care pathways). The aim of the PhD of Johanne Bakalara is to study temporal models of sequences in order 1) to show their abilities to specify complex proxies representing care sequences needed in pharmaco-epidemiological studies and 2) to build an efficient querying tool able to exploit large amount of care pathways.

4.4. Agriculture and environment

- **Dairy Farming.** The use and analysis of data acquired in dairy farming is a challenge both for data science and animal science. The goal is to improve farming conditions, i.e., health, welfare and environment, as well as farmers’ income. Nowadays, animals are monitored by multiple sensors giving a wealth of heterogeneous data such as temperature, weight, or milk composition. Current techniques used by animal scientists focus mostly on mono-sensor approaches. The dynamic combination of several sensors could provide new services and information useful for dairy farming. The PhD thesis of Kevin Fauvel (#DigitAg grant), aims to study such combinations of sensors and to investigate the use data mining methods, especially pattern mining algorithms. The challenge is to design new algorithms that take into account data heterogeneity—in terms of nature and time units—, and that produce useful patterns for dairy farming. The outcome of this thesis will be an original and important contribution to the new challenge of the IoT (Internet of Things) and will interest domain actors to find new added value to a global data analysis. The PhD thesis, started on October 2017, takes place in an interdisciplinary setting bringing together computer scientists from Inria and animal scientists from INRA, both located in Rennes.
- **Optimizing the Nutrition of Individual Sow.** Another direction for further research is the combination of data flows with prediction models in order to learn nutrition strategies. Raphaël Gauthier started a PhD thesis (#DigitAg Grant) in November 2017 with both Inria and INRA supervisors. His research addresses the problem of finding the optimal diet to be supplied to individual sows. Given all the information available, e.g., time-series information about previous feeding, environmental data, scientists models, the research goal is to design new algorithms to determine the optimal ration for a given sow in a given day. Efficiency issues of developed algorithms will be considered since the proposed software should work in real-time on the automated feeder. The decision support process should involve the stakeholder to ensure a good level of acceptance, confidence and understanding of the final tool.
- **Ecosystem Modeling and Management.** Ongoing research on ecosystem management includes

modelling of ecosystems and anthropogenic pressures, with a special concern on the representation of socio-economical factors that impact human decisions. A main research issue is how to represent these factors and how to integrate their impact on the ecosystem simulation model. This work is an ongoing cooperation with ecologists from the Marine Spatial Ecology of Queensland University, Australia and from Agrocampus Ouest.

4.5. Education

- **Data-oriented Academic Counseling.** Course selection and recommendation are important aspects of any academic counseling system. The Learning Analytics community has long supported these activities via automatic, data-based tools for recommendation and prediction. LACODAM, in collaboration with the Ecuadorian research center CTI⁰ has contributed to this body of research with the design of a tool that allows students to select multiple courses and predict their academic performance based on historical academic data. The tool resorts to interpretable machine learning techniques, and is intended to be used by the students before the counseling sessions to plan their upcoming semester at the Ecuadorian university ESPOL. The data visualization aspects of the tool, as well as the data science considerations and our emphasis on explainability are compiled in a paper that is under revision at the Learning Analytics and Knowledge Conference (LAK'20). The data science component of the tool was developed during the M1 internship of Mohammad Poul-Doust.

4.6. Others

- **RDF Archiving.** The dynamicity of the Semantic Web has motivated the development of solutions for RDF archiving, i.e., the task of storing and querying all previous versions of an RDF dataset. Notwithstanding the value of RDF archiving for data maintainers and consumers, this field of research remains under-developed for multiple reasons. These include notably (i) the lack of usability and scalability of the existing systems, (ii) no archiving support for multi-graph RDF datasets, (iii) the absence of a standard SPARQL extension for RDF archives, and (iv) a disregard of the evolution patterns of RDF datasets. The PhD thesis of Olivier Pelgrin aims at developing techniques towards a scalable and full-fledged archiving solution for the Semantic Web. This PhD thesis is a collaboration between LACODAM and the DAISY team at Aalborg University.

⁰Centro de Tecnologías de Información, <http://cti.espol.edu.ec/>

LINKS Project-Team

4. Application Domains

4.1. Linked Data Integration

There are many contexts in which integrating linked data is interesting. We advocate here one possible scenario, namely that of integrating business linked data to feed what is called Business Intelligence. The latter consists of a set of theories and methodologies that transform raw data into meaningful and useful information for business purposes (from Wikipedia). In the past decade, most of the enterprise data was proprietary, thus residing within the enterprise repository, along with the knowledge derived from that data. Today's enterprises and businessmen need to face the problem of information explosion, due to the Internet's ability to rapidly convey large amounts of information throughout the world via end-user applications and tools. Although linked data collections exist by bridging the gap between enterprise data and external resources, they are not sufficient to support the various tasks of Business Intelligence. To make a concrete example, concepts in an enterprise repository need to be matched with concepts in Wikipedia and this can be done via pointers or equalities. However, more complex logical statements (i.e. mappings) need to be conceived to map a portion of a local database to a portion of an RDF graph, such as a subgraph in Wikipedia or in a social network, e.g. LinkedIn. Such mappings would then enrich the amount of knowledge shared within the enterprise and let more complex queries be evaluated. As an example, businessmen with the aid of business intelligence tools need to make complex sentimental analysis on the potential clients and for such a reason, such tools must be able to pose complex queries, that exploit the previous logical mappings to guide their analysis. Moreover, the external resources may be rapidly evolving thus leading to revisit the current state of business intelligence within the enterprise.

4.2. Data Cleaning

The second example of application of our proposal concerns scientists who want to quickly inspect relevant literature and datasets. In such a case, local knowledge that comes from a local repository of publications belonging to a research institute (e.g. HAL) need to be integrated with other Web-based repositories, such as DBLP, Google Scholar, ResearchGate and even Wikipedia. Indeed, the local repository may be incomplete or contain semantic ambiguities, such as mistaken or missing conference venues, mistaken long names for the publication venues and journals, missing explanation of research keywords, and opaque keywords. We envision a publication management system that exploits both links between database elements, namely pointers to external resources and logical links. The latter can be complex relationships between local portions of data and remote resources, encoded as schema mappings. There are different tasks that such a scenario could entail such as (i) cleaning the errors with links to correct data e.g. via mappings from HAL to DBLP for the publications errors, and via mappings from HAL to Wikipedia for opaque keywords, (ii) thoroughly enrich the list of publications of a given research institute, and (iii) support complex queries on the corrected data combined with logical mappings.

4.3. Real Time Complex Event Processing

Complex event processing serves for monitoring nested word streams in real time. Complex event streams are gaining popularity with social networks such as with Facebook and Twitter, and thus should be supported by distributed databases on the Web. Since this is not yet the case, there remains much space for future industrial transfer related to Links' second axis on dynamic linked data.

MAGNET Project-Team

4. Application Domains

4.1. Application Domains

Our main targeted applications are browsing, monitoring, recommending and mining in information networks. The learning tasks considered in the project such as node clustering, node and link classification and link prediction are likely to yield important improvements in these applications. Application domains cover social networks for cultural data and e-commerce, and biomedical informatics.

We also target applications related to decentralized learning and privacy preserving systems when users or devices are interconnected in large networks. We develop solutions based on urban and mobility data where privacy is a specific requirement.

MOEX Project-Team (section vide)

ORPAILLEUR Project-Team

4. Application Domains

4.1. Life Sciences: Agronomy, Biology, Chemistry, and Medicine

Keywords: knowledge discovery in life sciences, biology, chemistry, medicine, pharmacogenomics and precision medicine.

One major application domain which is currently investigated by the Orpailleur team is related to life sciences, with particular emphasis on biology, medicine, and chemistry. The understanding of biological systems provides complex problems for computer scientists, and the developed solutions bring new research ideas or possibilities for biologists and for computer scientists as well. Indeed, the interactions between researchers in biology and researchers in computer science improve not only knowledge about systems in biology, chemistry, and medicine, but knowledge about computer science as well.

Knowledge discovery is gaining more and more interest and importance in life sciences for mining either homogeneous databases such as protein sequences and structures, or heterogeneous databases for discovering interactions between genes and the environment, or between genetic and phenotypic data, especially for public health and precision medicine (pharmacogenomics). Pharmacogenomics is one main challenge for the Orpailleur team as it considers a large panel of complex data ranging from biological to medical data, and various kinds of encoded domain knowledge ranging from texts to formal ontologies.

On the same line as biological data, chemical data are presenting important challenges w.r.t. knowledge discovery, for example for mining collections of molecular structures and collections of chemical reactions in organic chemistry. The mining of such collections is an important task for various reasons including the challenge of graph mining and the industrial needs (especially in drug design, pharmacology and toxicology). Molecules and chemical reactions are complex data that can be modeled as labeled graphs. Graph mining and Formal Concept Analysis methods play an important role in this application domain and can be used in an efficient and well-founded way [87].

Finally, research in agronomy is mainly based on cooperation between Inria and INRA. One research dimension is related to the characterization and the simulation of hedgerow structures in agricultural landscapes, based on Hilbert-Peano curves and Markov models [79]. Another research dimension is based on the mining of survey data for evaluating groundwater quality risks [86].

PETRUS Project-Team

4. Application Domains

4.1. Personal cloud, home care, IoT, sensing, surveys

As stated in the software section, the Petrus research strategy aims at materializing its scientific contributions in an advanced hardware/software platform with the expectation to produce a real societal impact. Hence, our software activity is structured around a common Secure Personal Cloud platform rather than several isolated demonstrators. This platform will serve as the foundation to develop a few emblematic applications. Several privacy-preserving applications can actually be targeted by a Personal Cloud platform, like: (i) smart disclosure applications allowing the individual to recover her personal data from external sources (e.g., bank, online shopping activity, insurance, etc.), integrate them and cross them to perform personal big data tasks (e.g., to improve her budget management) ; (ii) management of personal medical records for care coordination and well-being improvement; (iii) privacy-aware data management for the IoT (e.g., in sensors, quantified-self devices, smart meters); (iv) community-based sensing and community data sharing; (v) privacy-preserving studies (e.g., cohorts, public surveys, privacy-preserving data publishing). Such applications overlap with all the research axes described above but each of them also presents its own specificities. For instance, the smart disclosure applications will focus primarily on sharing models and enforcement, the IoT applications require to look with priority at the embedded data management and sustainability issues, while community-based sensing and privacy-preserving studies demand to study secure and efficient global query processing. Among these applications domains, one is already receiving a particular attention from our team. Indeed, we gained a strong expertise in the management and protection of healthcare data through our past DMSP (Dossier Medico-Social Partagé) experiment in the field. This expertise is being exploited to develop a dedicated healthcare and well-being personal cloud platform. We are currently deploying 10000 boxes equipped with PlugDB in the context of the DomYcile project. In this context, we are currently setting up an Inria Innovation Lab with the Hippocad company to industrialize this platform and deploy it at large scale (see Section the bilateral contract OwnCare II-Lab).

TYREX Project-Team

4. Application Domains

4.1. Querying Large Graphs

Increasingly large amounts of graph-structured data become available. The methods we develop apply for the efficient evaluation of graph queries over large — and potentially distributed — graphs. In particular, we consider the SPARQL query language, which is the standard language for querying graphs structured in the Resource Description Format (RDF). We also consider other increasingly popular graph query languages such as Cypher queries for extracting information from property graphs.

We compile graph queries into lower-level distributed primitives found in big data frameworks such as Apache Spark, Flink, etc. Applications of graph querying are ubiquitous and include: large knowledge bases, social networks, road networks, trust networks and fraud detection for cryptocurrencies, publications graphs, web graphs, recommenders, etc.

4.2. Predictive Analytics for Healthcare

One major expectation of data science in healthcare is the ability to leverage on digitized health information and computer systems to better apprehend and improve care. The availability of large amounts of clinical data and in particular electronic health records opens the way to the development of quantitative models for patients that can be used to predict health status, as well as to help prevent disease and adverse effects.

In collaboration with the CHU Grenoble, we explore solutions to the problem of predicting important clinical outcomes such as patient mortality, based on clinical data. This raises many challenges including dealing with the very high number of potential predictor variables and very resource-consuming data preparation stages.

VALDA Project-Team

4. Application Domains

4.1. Personal Information Management Systems

We recall that Valda’s focus is on human-centric data, i.e., data produced by humans, explicitly or implicitly, or more generally containing information about humans. Quite naturally, we have used as a privileged application area to validate Valda’s results that of personal information management systems (Pims for short) [33].

A Pims is a system that allows a user to integrate her own data, e.g., emails and other kinds of messages, calendar, contacts, web search, social network, travel information, work projects, etc. Such information is commonly spread across different services. The goal is to give back to a user the control on her information, allowing her to formulate queries such as “What kind of interaction did I have recently with Alice B.?”, “Where were my last ten business trips, and who helped me plan them?”. The system has to orchestrate queries to the various services (which means knowing the existence of these services, and how to interact with them), integrate information from them (which means having data models for this information and its representation in the services), e.g., align a GPS location of the user to a business address or place mentioned in an email, or an event in a calendar to some event in a Web search. This information must be accessed intensionally: for instance, costly information extraction tools should only be run on emails which seem relevant, perhaps identified by a less costly cursory analysis (this means, in turn, obtaining a cost model for access to the different services). Impacted people can be found by examining events in the user’s calendar and determining who is likely to attend them, perhaps based on email exchanges or former events’ participant lists. Of course, uncertainty has to be maintained along the entire process, and provenance information is needed to explain query results to the user (e.g., indicate which meetings and trips are relevant to each person of the output). Knowledge about services, their data models, their costs, need either to be provided by the system designer, or to be automatically learned from interaction with these services, as in [68].

One motivation for that choice is that Pims concentrate many of the problems we intend to investigate: heterogeneity (various sources, each with a different structure), massive distribution (information spread out over the Web, in numerous sources), rapid evolution (new data regularly added), intensionality (knowledge from Wikidata, OpenStreetMap...), confidentiality and security (mostly private data), and uncertainty (very variable quality). Though the data is distributed, its size is relatively modest; other applications may be considered for works focusing on processing data at large scale, which is a potential research direction within Valda, though not our main focus. Another strong motivation for the choice of Pims as application domain is the importance of this application from a societal viewpoint.

A Pims is essentially a system built on top of a user’s *personal knowledge base*; such knowledge bases are reminiscent of those found in the Semantic Web, e.g., linked open data. Some issues, such as ontology alignment [71] exist in both scenarios. However, there are some fundamental differences in building personal knowledge bases vs collecting information from the Semantic Web: first, the scope is quite smaller, as one is only interested in knowledge related to a given individual; second, a small proportion of the data is already present in the form of semantic information, most needs to be extracted and annotated through appropriate wrappers and enrichers; third, though the linked open data is meant to be read-only, the only update possible to a user being adding new triples, a personal knowledge base is very much something that a user needs to be able to edit, and propagating updates from the knowledge base to original data sources is a challenge in itself.

4.2. Web Data

The choice of Pims is not exclusive. We also consider other application areas as well. In particular, we have worked in the past and have a strong expertise on Web data [38] in a broad sense: semi-structured, structured, or unstructured content extracted from Web databases [68]; knowledge bases from the Semantic Web [71];

social networks [66]; Web archives and Web crawls [54]; Web applications and deep Web databases [47]; crowdsourcing platforms [41]. We intend to continue using Web data as a natural application domain for the research within Valda when relevant. For instance [45], deep Web databases are a natural application scenario for intensional data management issues: determining if a deep Web database contains some information requires optimizing the number of costly requests to that database.

A common aspect of both personal information and Web data is that their exploitation raises ethical considerations. Thus, a user needs to remain fully in control of the usage that is made of her personal information; a search engine or recommender system that ranks Web content for display to a specific user needs to do so in an unbiased, justifiable, manner. These ethical constraints sometimes forbid some technically solutions that may be technically useful, such as sharing a model learned from the personal data of a user to another user, or using blackboxes to rank query result. We fully intend to consider these ethical considerations within Valda. One of the main goals of a Pims is indeed to empower the user with a full control on the use of this data.

WIMMICS Project-Team

4. Application Domains

4.1. Social Semantic Web

A number of evolutions have changed the face of information systems in the past decade but the advent of the Web is unquestionably a major one and it is here to stay. From an initial wide-spread perception of a public documentary system, the Web as an object turned into a social virtual space and, as a technology, grew as an application design paradigm (services, data formats, query languages, scripting, interfaces, reasoning, etc.). The universal deployment and support of its standards led the Web to take over nearly all of our information systems. As the Web continues to evolve, our information systems are evolving with it.

Today in organizations, not only almost every internal information system is a Web application, but these applications more and more often interact with external Web applications. The complexity and coupling of these Web-based information systems call for specification methods and engineering tools. From capturing the needs of users to deploying a usable solution, there are many steps involving computer science specialists and non-specialists.

We defend the idea of relying on Semantic Web formalisms to capture and reason on the models of these information systems supporting the design, evolution, interoperability and reuse of the models and their data as well as the workflows and the processing.

4.2. Linked Data on the Web and on Intranets

With billions of triples online (see Linked Open Data initiative), the Semantic Web is providing and linking open data at a growing pace and publishing and interlinking the semantics of their schemas. Information systems can now tap into and contribute to this Web of data, pulling and integrating data on demand. Many organisations also started to use this approach on their intranets leading to what is called linked enterprise data.

A first application domain for us is the publication and linking of data and their schemas through Web architectures. Our results provide software platforms to publish and query data and their schemas, to enrich these data in particular by reasoning on their schemas, to control their access and licenses, to assist the workflows that exploit them, to support the use of distributed datasets, to assist the browsing and visualization of data, etc.

Examples of collaboration and applied projects include: SMILK Joint Laboratory, Corese, DBpedia.fr.

4.3. Assisting Web-based Epistemic Communities

In parallel with linked open data on the Web, social Web applications also spread virally (e.g. Facebook growing toward 1.5 billion users) first giving the Web back its status of a social read-write media and then putting it back on track to its full potential of a virtual place where to act, react and interact. In addition, many organizations are now considering deploying social Web applications internally to foster community building, expert cartography, business intelligence, technological watch and knowledge sharing in general.

By reasoning on the Linked Data and the semantics of the schemas used to represent social structures and Web resources, we provide applications supporting communities of practice and interest and fostering their interactions in many different contexts (e-learning, business intelligence, technical watch, etc.).

We use typed graphs to capture and mix: social networks with the kinds of relationships and the descriptions of the persons; compositions of Web services with types of inputs and outputs; links between documents with their genre and topics; hierarchies of classes, thesauri, ontologies and folksonomies; recorded traces and suggested navigation courses; submitted queries and detected frequent patterns; timelines and workflows; etc.

Our results assist epistemic communities in their daily activities such as biologists exchanging results, business intelligence and technological watch networks informing companies, engineers interacting on a project, conference attendees, students following the same course, tourists visiting a region, mobile experts on the field, etc. Examples of collaboration and applied projects: EduMICS, OCKTOPUS, Vigiglobe, Educlever, Gayatech.

4.4. Linked Data for a Web of Diversity

We intend to build on our results on explanations (provenance, traceability, justifications) and to continue our work on opinions and arguments mining toward the global analysis of controversies and online debates. One result would be to provide new search results encompassing the diversity of viewpoints and providing indicators supporting opinion and decision making and ultimately a Web of trust. Trust indicators may require collaborations with teams specialized in data certification, cryptography, signature, security services and protocols, etc. This will raise the specific problem of interaction design for security and privacy. In addition, from the point of view of the content, this requires to foster the publication and coexistence of heterogeneous data with different points of views and conceptualizations of the world. We intend to pursue the extension of formalisms to allow different representations of the world to co-exist and be linked and we will pay special attention to the cultural domain and the digital humanities. Examples of collaboration and applied projects: Zoomathia, Seempad, SMILK,

4.5. Artificial Web Intelligence

We intend to build on our experience in artificial intelligence (knowledge representation, reasoning) and distributed artificial intelligence (multi-agent systems - MAS) to enrich formalisms and propose alternative types of reasoning (graph-based operations, reasoning with uncertainty, inductive reasoning, non-monotonic, etc.) and alternative architectures for linked data with adequate changes and extensions required by the open nature of the Web. There is a clear renewed interest in AI for the Web in general and for Web intelligence in particular. Moreover distributed AI and MAS provide both new architectures and new simulation platforms for the Web. At the macro level, the evolution accelerated with HTML5 toward Web pages as full applications and direct Page2Page communication between browser clearly is a new area for MAS and P2P architectures. Interesting scenarios include the support of a strong decentralization of the Web and its resilience to degraded technical conditions (downscaling the Web), allowing pages to connect in a decentralized way, forming a neutral space, and possibly going offline and online again in erratic ways. At the micro level, one can imagine the place RDF and SPARQL could take as data model and programming model in the virtual machines of these new Web pages and, of course, in the Web servers. RDF is also used to serialize and encapsulate other languages and becomes a pivot language in linking very different applications and aspects of applications. Example of collaboration and applied projects: MoreWAIS, Corese, Vigiglobe collaboration.

4.6. Human-Data Interaction (HDI) on the Web

We need more interaction design tools and methods for linked data access and contribution. We intend to extend our work on exploratory search coupling it with visual analytics to assist sense making. It could be a continuation of the Gephi extension that we built targeting more support for non experts to access and analyze data on a topic or an issue of their choice. More generally speaking SPARQL is inappropriate for common users and we need to support a larger variety of interaction means with linked data. We also believe linked data and natural language processing (NLP) have to be strongly integrated to support natural language based interactions. Linked Open Data (LOD) for NLP, NLP for LOD and Natural Dialog Processing for querying, extracting and asserting data on the Web is a priority to democratize its use. Micro accesses and micro contributions are important to ensure public participation and also call for customized interfaces and thus for methods and tools to generate these interfaces. In addition, the user profiles are being enriched now with new data about the user such as her current mental and physical state, the emotion she just expressed or her cognitive performances. Taking into account this information to improve the interactions, change the behavior of the system and adapt the interface is a promising direction. And these human-data interaction

means should also be available for “small data”, helping the user to manage her personal information and to link it to public or collective one, maintaining her personal and private perspective as a personal Web of data. Finally, the continuous knowledge extractions, updates and flows add the additional problem of representing, storing, querying and interacting with dynamic data. Examples of collaboration and applied projects: QAKIS, Sychonext collaboration, ALOOF, DiscoveryHub, WASABI, MoreWAIS.

Web-augmented interactions with the world: The Web continues to augment our perception and interaction with reality. In particular, Linked Open Data enable new augmented reality applications by providing data sources on almost any topic. The current enthusiasm for the Web of Things, where every object has a corresponding Web resource, requires evolutions of our vision and use of the Web architecture. This vision requires new techniques as the ones mentioned above to support local search and contextual access to local resources but also new methods and tools to design Web-based human devices interactions, accessibility, etc. These new usages are placing new requirements on the Web Architecture in general and on the semantic Web models and algorithms in particular to handle new types of linked data. They should support implicit requests considering the user context as a permanent query. They should also simplify our interactions with devices around us jointly using our personal preferences and public common knowledge to focus the interaction on the vital minimum that cannot be derived in another way. For instance the access to the Web of data for a robot can completely change the quality of the interactions it can offer. Again, these interactions and the data they require raise problems of security and privacy. Examples of collaboration and applied projects: ALOOF, AZKAR, MoreWAIS.

ZENITH Project-Team

4. Application Domains

4.1. Data-intensive Scientific Applications

The application domains covered by Zenith are very wide and diverse, as they concern data-intensive scientific applications, i.e., most scientific applications. Since the interaction with scientists is crucial to identify and tackle data management problems, we are dealing primarily with application domains for which Montpellier has an excellent track record, i.e., agronomy, environmental science, life science, with scientific partners like INRA, IRD and CIRAD. However, we are also addressing other scientific domains (e.g. astronomy, oil extraction, music processing) through our international collaborations.

Let us briefly illustrate some representative examples of scientific applications on which we have been working on.

- **Management of astronomical catalogs.** An example of data-intensive scientific applications is the management of astronomical catalogs generated by the Dark Energy Survey (DES) project on which we are collaborating with researchers from Brazil. In this project, huge tables with billions of tuples and hundreds of attributes (corresponding to dimensions, mainly double precision real numbers) store the collected sky data. Data are appended to the catalog database as new observations are performed and the resulting database size has hundreds of TB. Scientists around the globe can query the database with queries that may contain a considerable number of attributes. The volume of data that this application holds poses important challenges for data management. In particular, efficient solutions are needed to partition and distribute the data in several servers. An efficient partitioning scheme should try to minimize the number of fragments accessed in the execution of a query, thus reducing the overhead associated to handle the distributed execution.
- **Personal health data analysis and privacy** Today, it is possible to acquire data on many domains related to personal data. For instance, one can collect data on her daily activities, habits or health. It is also possible to measure performance in sports. This can be done thanks to sensors, communicating devices or even connected glasses. Such data, once acquired, can lead to valuable knowledge for these domains. For people having a specific disease, it might be important to know if they belong to a specific category that needs particular care. For an individual, it can be interesting to find a category that corresponds to her performances in a specific sport and then adapt her training with an adequate program. Meanwhile, for privacy reasons, people will be reluctant to share their personal data and make them public. Therefore, it is important to provide them with solutions that can extract such knowledge from everybody's data, while guaranteeing that their private data won't be disclosed to anyone.
- **Botanical data sharing.** Botanical data is highly decentralized and heterogeneous. Each actor has its own expertise domain, hosts its own data, and describes them in a specific format. Furthermore, botanical data is complex. A single plant's observation might include many structured and unstructured tags, several images of different organs, some empirical measurements and a few other contextual data (time, location, author, etc.). A noticeable consequence is that simply identifying plant species is often a very difficult task; even for the botanists themselves (the so-called taxonomic gap). Botanical data sharing should thus speed up the integration of raw observation data, while providing users an easy and efficient access to integrated data. This requires to deal with social-based data integration and sharing, massive data analysis and scalable content-based information retrieval. We address this application in the context of the French initiative PI@ntNet, with CIRAD and IRD.
- **Biological data integration and analysis.**

Biology and its applications, from medicine to agronomy and ecology, are now producing massive data, which is revolutionizing the way life scientists work. For instance, using plant phenotyping platforms such as PhenoDyn and PhenoArch at INRA Montpellier, quantitative genetic methods allow to identify genes involved in phenotypic variation in response to environmental conditions. These methods produce large amounts of data at different time intervals (minutes to months), at different sites and at different scales ranging from small tissue samples to the entire plant until whole plant population. Analyzing such big data creates new challenges for data management and data integration.

- **Audio heritage preservation.**

Since the end of the 19th century, France has commissioned ethnologists to record the world's immaterial audio heritage. This results in datasets of dozens of thousands of audio recordings from all countries and more than 1200 ethnies. Today, this data is gathered under the name of 'Archives du CNRS — Musée de l'Homme' and is handled by the CREM (Centre de Recherche en Ethno-Musicologie). Professional scientists in digital humanities are accessing this data daily for their investigations, and several important challenges arise to ease their work. The KAMoulox project, lead by A. Liutkus, targets at offering online processing tools for the scientists to automatically restore this old material on demand.

These application examples illustrate the diversity of requirements and issues which we are addressing with our scientific application partners. To further validate our solutions and extend the scope of our results, we also want to foster industrial collaborations, even in non scientific applications, provided that they exhibit similar challenges.

ALICE Team

4. Application Domains

4.1. Geometric Tools for Simulating Physics with a Computer

Numerical simulation is the main targeted application domain for the geometry processing tools that we develop. Our mesh generation tools will be tested and evaluated within the context of our cooperation with Hutchinson, experts in vibration control, fluid management and sealing system technologies. We think that the hex-dominant meshes that we generate have geometrical properties that make them suitable for some finite element analyses, especially for simulations with large deformations.

We also have a tight collaboration with Wan Chiu Li, a geophysical modeling specialist. He is creating a start-up company MESHSPACE whose goal is to transfer to the industry remeshing tools we developed in the ALICE team. In particular, he uses hexahedral-dominant meshes for geomechanical simulations of gas and oil reservoirs. From a scientific point of view, this use case introduces new types of constraints (alignments with faults and horizons), and allows certain types of nonconformities that we did not consider until now.

Our cooperation with RhinoTerrain pursues the same goal: reconstruction of buildings from point cloud scans allows to perform 3D analysis and studies on insolation, floods and wave propagation, wind and noise simulations necessary for urban planification.

AVIZ Project-Team

4. Application Domains

4.1. Natural Sciences

As part of a CORDI PhD project, we collaborate with researchers at CERN on interactive data visualization using augmented reality, with the goal to better understand this new visualization environment and to support the physicists in analysing their 3D particle collision data. As part of another CORDI PhD project, we collaborate with researchers at the German Center for Climate Computation (DKRZ), to better understand collaborative data exploration and interaction in immersive analytics contexts. Finally, as part of the Inria IPL “Naviscope,” we collaborate with researchers at INRA (as well as other Inria teams) on interactive visualization tools for the exploration of plant embryo development.

4.2. Social Sciences

We collaborate with social science researchers from EHESS Paris on the visualization of dynamic networks; they use our systems (GeneaQuilts [56], Vistorian [64], PAOHVis [7]) and teach them to students and researchers. Our tools are used daily by ethnographers and historians to study the evolution of social relations over time. In the social sciences, many datasets are gathered by individual researchers to answer a specific question, and automated analytical methods cannot be applied to these small datasets. Furthermore, the studies are often focused on specific persons or organizations and not always on the modeling or prediction of the behavior of large populations. The tools we design to visualize complex multivariate dynamic networks are unique and suited to typical research questions shared by a large number of researchers. This line of research is supported by the DataIA “HistorIA” project, and by the “IVAN” European project. We also collaborate on the BART initiative, a joint project with IRT-SystemX on the analysis and visualization of blockchain data, in collaboration with economists from Université Paris-Sud.

4.3. Medicine

We collaborate with CMAP/Polytechnique on the analysis and visualization of CNAM Data “parcours de santé” to help referent doctors and epidemiologists make sense of French health data. In particular, we are working on a subset of the CNAM Data focused on urinary problems, and we have received a very positive feedback from doctors who can see what happens to the patients treated in France vs. what they thought happened through the literature. This project is starting but is already getting a lot of traction from our partners in medicine, epidemiology, and economy of health.

EX-SITU Project-Team

4. Application Domains

4.1. Creative industries

We work closely with creative professionals in the arts and in design, including music composers, musicians, and sound engineers; painters and illustrators; dancers and choreographers; theater groups; game designers; graphic and industrial designers; and architects.

4.2. Scientific research

We work with creative professionals in the sciences and engineering, including neuroscientists and doctors; programmers and statisticians; chemists and astrophysicists; and researchers in fluid mechanics.

GRAPHDECO Project-Team (section vide)

HYBRID Project-Team

4. Application Domains

4.1. Overview

The research program of Hybrid team aims at next generations of virtual reality and 3D user interfaces which could possibly address both the “body” and “mind” of the user. Novel interaction schemes are designed, for one or multiple users. We target better integrated systems and more compelling user experiences.

The applications of our research program correspond to the applications of virtual reality technologies which could benefit from the addition of novel body-based or mind-based interaction capabilities:

- **Industry:** with training systems, virtual prototyping, or scientific visualization;
- **Medicine:** with rehabilitation and reeducation systems, or surgical training simulators;
- **Entertainment:** with movie industry, content customization, video games or attractions in theme parks,
- **Construction:** with virtual mock-ups design and review, or historical/architectural visits.
- **Cultural Heritage:** with acquisition, virtual excavation, virtual reconstruction and visualization

ILDA Project-Team

4. Application Domains

4.1. Mission-critical systems

Mission-critical contexts of use include emergency response & management, and critical infrastructure operations, such as public transportation systems, communications and power distribution networks, or the operations of large scientific instruments such as particle accelerators and astronomical observatories. Central to these contexts of work is the notion of situation awareness [25], i.e., how workers perceive and understand elements of the environment with respect to time and space, such as maps and geolocated data feeds from the field, and how they form mental models that help them predict future states of those elements. One of the main challenges is how to best assist subject-matter experts in constructing correct mental models and making informed decisions, often under time pressure. This can be achieved by providing them with, or helping them efficiently identify and correlate, relevant and timely information extracted from large amounts of raw data, taking into account the often cooperative nature of their work and the need for task coordination. With this application area, our goal is to investigate novel ways of interacting with computing systems that improve collaborative data analysis capabilities and decision support assistance in a mission-critical, often time-constrained, work context.

Relevant publications by team members this year: [13], [19], [12].

4.2. Exploratory analysis of scientific data

Many scientific disciplines are increasingly data-driven, including astronomy, molecular biology, particle physics, or neuroanatomy. While making the right decision under time pressure is often less of a critical issue when analyzing scientific data, at least not on the same temporal scale as truly time-critical systems, scientists are still faced with large-to-huge amounts of data. No matter their origin (experiments, remote observations, large-scale simulations), these data are difficult to understand and analyze in depth because of their sheer size and complexity. Challenges include how to help scientists freely-yet-efficiently explore their data, keep a trace of the multiple data processing paths they considered to verify their hypotheses and make it easy to backtrack, and how to relate observations made on different parts of the data and insights gained at different moments during the exploration process. With this application area, our goal is to investigate how data-centric interactive systems can improve collaborative scientific data exploration, where users' goals are more open-ended, and where roles, collaboration and coordination patterns [48] differ from those observed in mission-critical contexts of work.

Relevant publications by team members last year: [16], [24], [14], [18].

IMAGINE Project-Team

4. Application Domains

4.1. Target applications

Our research can be applied to any situation where users need to create new, imaginary, 3D content. Our work should be instrumental, in the long term, for the visual arts, from the creation of 3D films and games to the development of new digital planning tools for theater or cinema directors. Our models can also be used in interactive prototyping environments for engineering. They can help promoting interactive digital design to scientists, as a tool to quickly express, test and refine models, as well as an efficient way for conveying them to other people. Lastly, we expect our new methodology to put digital modeling within the reach of the general public, enabling educators, media and other practitioners to author their own 3D content. The diversity of users these domains bring, from digital experts to other professionals and novices, gives us excellent opportunities to validate our general methodology with different categories of users. Our ongoing projects in these various application domains are listed in Section 7.

4.2. Visual arts

- Sculpture.
- Modeling and animation for 3D films and games.
- Virtual cinematography and tools for theater directors.

4.3. Engineering

- Industrial design.
- Mechanical & civil engineering.

4.4. Natural sciences

- Geology.
- Virtual functional anatomy.

4.5. Education and creative tools

- Sketch-based teaching.
- Creative environments for novice users.
- Museography.

LOKI Project-Team

4. Application Domains

4.1. Application Domains

Loki works on fundamental and technological aspects of Human-Computer Interaction that can be applied to diverse application domains.

Our 2019 research involved desktop, augmented reality, touch-based, haptics, and BCI interfaces with notable applications to medicine (analysis of fine motor control for patients with Parkinson disease), digital humanities (interpretation of handwritten historical documents), as well as creativity support tools (production of illustrations, design of Digital Musical Instruments). Our technical work also contributes to the more general application domains of software engineering and systems' design.

MANAO Project-Team

4. Application Domains

4.1. Physical Systems

Given our close relationships with researchers in optics, one novelty of our approach is to extend the range of possible observers to physical sensors in order to work on domains such as simulation, mixed reality, and testing. Capturing, processing, and visualizing complex data is now more and more accessible to everyone, leading to the possible convergence of real and virtual worlds through visual signals. This signal is traditionally captured by cameras. It is now possible to augment them by projecting (e.g., the infrared laser of Microsoft Kinect) and capturing (e.g., GPS localization) other signals that are outside the visible range. This supplemental information replaces values traditionally extracted from standard images and thus lowers down requirements in computational power. Since the captured images are the result of the interactions between light, shape, and matter, the approaches and the improved knowledge from *MANAO* help in designing interactive acquisition and rendering technologies that are required to merge the real and the virtual worlds. With the resulting unified systems (optical and digital), transfer of pertinent information is favored and inefficient conversion is likely avoided, leading to new uses in interactive computer graphics applications, like **augmented reality**, **displays** and **computational photography**.

4.2. Interactive Visualization and Modeling

This direction includes domains such as **scientific illustration and visualization**, **artistic or plausible rendering**, and **3D modeling**. In all these cases, the observer, a human, takes part in the process, justifying once more our focus on real-time methods. When targeting average users, characteristics as well as limitations of the human visual system should be taken into account: in particular, it is known that some configurations of light, shape, and matter have masking and facilitation effects on visual perception. For specialized applications (such as archeology), the expertise of the final user and the constraints for 3D user interfaces lead to new uses and dedicated solutions for models and algorithms.

MAVERICK Project-Team

4. Application Domains

4.1. Application Domains

The natural application domain for our research is the production of digital images, for example for movies and special effects, virtual prototyping, video games... Our research have also been applied to tools for generating and editing images and textures, for example generating textures for maps. Our current application domains are:

- Offline and real-time rendering in movie special effects and video games;
- Virtual prototyping;
- Scientific visualization;
- Content modeling and generation (e.g. generating texture for video games, capturing reflectance properties, etc);
- Image creation and manipulation.

MFx Project-Team

4. Application Domains

4.1. Digital Manufacturing

Our work addresses generic challenges related to fabrication and can thus be applied in a wide variety of contexts. Our aim is first and foremost to develop the algorithms that will allow various industrial sectors to benefit more strongly from the potential of AM. To enable this, we seek collaborations with key industry partners developing software and AM systems for a variety of processes and materials that are of interest to specific sectors (*e.g.*, dental, prosthetic, automotive, aerospace).

4.2. Medical Applications

To allow for faster transfer of our techniques and unlock novel applications, we actively seek to develop applications in the medical sector. In particular, we are involved in a project around the design of orthoses which explores how our research on elasticity control through microstructure geometries can be specifically applied to the medical sector; see §9.1.1 .

MIMETIC Project-Team

4. Application Domains

4.1. Animation, Autonomous Characters and Digital Storytelling

Computer Animation is one of the main application domain of the research work conducted in the MimeTIC team, in particular in relation to the entertainment and game industries. In these domains, creating virtual characters that are able to replicate real human motions and behaviours still highlights unanswered key challenges, especially as virtual characters are becoming more and more required to populate virtual worlds. For instance, virtual characters are used to replace secondary actors and generate highly populated scenes that would be hard and costly to produce with real actors, which requires to create high quality replicas that appear, move and behave both individually and collectively like real humans. The three key challenges for the MimeTIC team are therefore (i) to create natural animations (i.e., virtual characters that move like real humans), (ii) to create autonomous characters (i.e., that behave like real humans) and (iii) to orchestrate the virtual characters so as to create interactive stories.

First, our challenge is therefore to create animations of virtual characters that are natural, in the largest sense of the term of moving like a real human real would. This challenge covers several aspects of Character Animation depending on the context of application, e.g., producing visually plausible or physically correct motions, producing natural motion sequences, etc. Our goal is therefore to develop novel methods for animating virtual characters, e.g., based on motion capture, data-driven approaches, or learning approaches. However, because of the complexity of human motion (e.g., the number of degrees of freedom that can be controlled), resulting animations are not necessarily physically, biomechanically, or visually plausible. For instance, current physics-based approaches produce physically correct motions but not necessarily perceptually plausible ones. All these reasons are why most entertainment industries still mainly rely on manual animation, e.g., in games and movies. Therefore, research in MimeTIC on character animation is also conducted with the goal of validating objective (e.g., physical, biomechanical) as well as subjective (e.g., visual plausibility) criteria.

Second, one of the main challenges in terms of autonomous characters is to provide a unified architecture for the modeling of their behavior. This architecture includes perception, action and decisional parts. This decisional part needs to mix different kinds of models, acting at different time scale and working with different nature of data, ranging from numerical (motion control, reactive behaviors) to symbolic (goal oriented behaviors, reasoning about actions and changes). For instance, autonomous characters play the role of actors that are driven by a scenario in video games and virtual storytelling. Their autonomy allows them to react to unpredictable user interactions and adapt their behavior accordingly. In the field of simulation, autonomous characters are used to simulate the behavior of humans in different kind of situations. They enable to study new situations and their possible outcomes. In the MimeTIC team, our focus is therefore not to reproduce the human intelligence but to propose an architecture making it possible to model credible behaviors of anthropomorphic virtual actors evolving/moving in real time in virtual worlds. The latter can represent particular situations studied by psychologists of the behavior or to correspond to an imaginary universe described by a scenario writer. The proposed architecture should mimic all the human intellectual and physical functions.

Finally, interactive digital storytelling, including novel forms of edutainment and serious games, provides access to social and human themes through stories which can take various forms and contains opportunities for massively enhancing the possibilities of interactive entertainment, computer games and digital applications. It provides chances for redefining the experience of narrative through interactive simulations of computer-generated story worlds and opens many challenging questions at the overlap between computational narratives, autonomous behaviours, interactive control, content generation and authoring tools. Of particular interest for the MimeTIC research team, virtual storytelling triggers challenging opportunities in providing effective models for enforcing autonomous behaviours for characters in complex 3D environments. Offering both low-level capacities to characters such as perceiving the environments, interacting with the environment and

reacting to changes in the topology, on which to build higher-levels such as modelling abstract representations for efficient reasoning, planning paths and activities, modelling cognitive states and behaviours requires the provision of expressive, multi-level and efficient computational models. Furthermore virtual storytelling requires the seamless control of the balance between the autonomy of characters and the unfolding of the story through the narrative discourse. Virtual storytelling also raises challenging questions on the conveyance of a narrative through interactive or automated control of the cinematography (how to stage the characters, the lights and the cameras). For example, estimating visibility of key subjects, or performing motion planning for cameras and lights are central issues for which have not received satisfactory answers in the literature.

4.2. Fidelity of Virtual Reality

VR is a powerful tool for perception-action experiments. VR-based experimental platforms allow exposing a population to fully controlled stimuli that can be repeated from trial to trial with high accuracy. Factors can be isolated and objects manipulations (position, size, orientation, appearance, ..) are easy to perform. Stimuli can be interactive and adapted to participants' responses. Such interesting features allow researchers to use VR to perform experiments in sports, motion control, perceptual control laws, spatial cognition as well as person-person interactions. However, the interaction loop between users and their environment differs in virtual conditions in comparison with real conditions. When a user interact in an environment, movement from action and perception are closely related. While moving, the perceptual system (vision, proprioception,..) provides feedback about the users' own motion and information about the surrounding environment. That allows the user to adapt his/her trajectory to sudden changes in the environment and generate a safe and efficient motion. In virtual conditions, the interaction loop is more complex because it involves several material aspects.

First, the virtual environment is perceived through a numerical display which could affect the available information and thus could potentially introduce a bias. For example, studies observed a distance compression effect in VR, partially explained by the use of Head Mounted Display with reduced field of view and exerting a weight and torques on the user's head. Similarly, the perceived velocity in a VR environment differs from the real world velocity, introducing an additional bias. Other factors, such as the image contrast, delays in the displayed motion and the point of view can also influence efficiency in VR. The second point concerns the user's motion in the virtual world. The user can actually move if the virtual room is big enough or if wearing a head mounted display. Even with a real motion, authors showed that walking speed is decreased, personal space size is modified and navigation in VR is performed with increased gait instability. Although natural locomotion is certainly the most ecological approach, the physical limited size of VR setups prevents from using it most of the time. Locomotion interfaces are therefore required. Locomotion interfaces are made up of two components, a locomotion metaphor (device) and a transfer function (software), that can also introduce bias in the generated motion. Indeed, the actuating movement of the locomotion metaphor can significantly differ from real walking and the simulated motion depends on the transfer function applied. Locomotion interfaces cannot usually preserve all the sensory channels involved in locomotion.

When studying human behavior in VR, the aforementioned factors in the interaction loop potentially introduce bias both in the perception and in the generation of motor behavior trajectories. MimeTIC is working on the mandatory step of VR validation to make it usable for capturing and analyzing human motion.

4.3. Motion Sensing of Human Activity

Recording human activity is a key point of many applications and fundamental works. Numerous sensors and systems have been proposed to measure positions, angles or accelerations of the user's body parts. Whatever the system is, one of the main problems is to be able to automatically recognize and analyze the user's performance according to poor and noisy signals. Human activity and motion are subject to variability: intra-variability due to space and time variations of a given motion, but also inter-variability due to different styles and anthropometric dimensions. MimeTIC has addressed the above problems in two main directions.

Firstly, we have studied how to recognize and quantify motions performed by a user when using accurate systems such as Vicon (product of Oxford Metrics), Qualisys, or Optitrack (product of Natural Point) motion capture systems. These systems provide large vectors of accurate information. Due to the size of the state vector (all the degrees of freedom) the challenge is to find the compact information (named features) that enables the automatic system to recognize the performance of the user. Whatever the method used, finding these relevant features that are not sensitive to intra-individual and inter-individual variability is a challenge. Some researchers have proposed to manually edit these features (such as a Boolean value stating if the arm is moving forward or backward) so that the expertise of the designer is directly linked with the success ratio. Many proposals for generic features have been proposed, such as using Laban notation which was introduced to encode dancing motions. Other approaches tend to use machine learning to automatically extract these features. However most of the proposed approaches were used to seek a database for motions which properties correspond to the features of the user's performance (named motion retrieval approaches). This does not ensure the retrieval of the exact performance of the user but a set of motions with similar properties.

Secondly, we wish to find alternatives to the above approach which is based on analyzing accurate and complete knowledge on joint angles and positions. Hence new sensors, such as depth-cameras (Kinect, product of Microsoft) provide us with very noisy joint information but also with the surface of the user. Classical approaches would try to fit a skeleton into the surface in order to compute joint angles which, again, lead to large state vectors. An alternative would be to extract relevant information directly from the raw data, such as the surface provided by depth cameras. The key problem is that the nature of these data may be very different from classical representation of human performance. In MimeTIC, we try to address this problem in specific application domains that require picking specific information, such as gait asymmetry or regularity for clinical analysis of human walking.

4.4. Sports

Sport is characterized by complex displacements and motions. One main objective is to understand the determinants of performance through the analysis of the motion itself. In the team, different sports have been studied such as the tennis serve, where the goal was to understand the contribution of each segment of the body in the performance but also the risk of injuries as well as other situation in cycling, swimming, fencing or soccer. Sports motions are dependent on visual information that the athlete can pick up in his environment, including the opponent's actions. Perception is thus fundamental to the performance. Indeed, a sportive action, as unique, complex and often limited in time, requires a selective gathering of information. This perception is often seen as a prerogative for action, it then takes the role of a passive collector of information. However, as mentioned by Gibson in 1979, the perception-action relationship should not be considered sequentially but rather as a coupling: we perceive to act but we must act to perceive. There would thus be laws of coupling between the informational variables available in the environment and the motor responses of a subject. In other words, athletes have the ability to directly perceive the opportunities of action directly from the environment. Whichever school of thought considered, VR offers new perspectives to address these concepts by complementary using real time motion capture of the immersed athlete.

In addition to better understand sports and interactions between athletes, VR can also be used as a training environment as it can provide complementary tools to coaches. It is indeed possible to add visual or auditory information to better train an athlete. The knowledge found in perceptual experiments can be for example used to highlight the body parts that are important to look at to correctly anticipate the opponent's action.

4.5. Ergonomics

The design of workstations nowadays tends to include assessment steps in a Virtual Environment (VE) to evaluate ergonomic features. This approach is more cost-effective and convenient since working directly on the Digital Mock-Up (DMU) in a VE is preferable to constructing a real physical mock-up in a Real Environment (RE). This is substantiated by the fact that a Virtual Reality (VR) set-up can be easily modified, enabling quick adjustments of the workstation design. Indeed, the aim of integrating ergonomics evaluation tools in VEs is to facilitate the design process, enhance the design efficiency, and reduce the costs.

The development of such platforms asks for several improvements in the field of motion analysis and VR. First, interactions have to be as natural as possible to properly mimic the motions performed in real environments. Second, the fidelity of the simulator also needs to be correctly evaluated. Finally, motion analysis tools have to be able to provide in real-time biomechanics quantities usable by ergonomists to analyse and improve the working conditions.

In real working condition, motion analysis and musculoskeletal risks assessment raise also many scientific and technological challenges. Similarly to virtual reality, fidelity of the working process may be affected by the measurement method. Wearing sensors or skin markers, together with the need of frequently calibrating the assessment system may change the way workers perform the tasks. Whatever the measurement is, classical ergonomic assessments generally address one specific parameter, such as posture, or force, or repetitions..., which makes it difficult to design a musculoskeletal risk factor that actually represent this risk. Another key scientific challenge is then to design new indicators that better capture the risk of musculoskeletal disorders. However, this indicator has to deal with the tradeoff between accurate biomechanical assessment and the difficulty to get reliable and the required information in real working conditions.

4.6. Locomotion and Interactions between walkers

Modeling and simulating locomotion and interactions between walkers is a very active, complex and competitive domain, interesting various disciplines such as mathematics, cognitive sciences, physics, computer graphics, rehabilitation etc. Locomotion and interactions between walkers are by definition at the very core of our society since they represent the basic synergies of our daily life. When walking in the street, we should produce a locomotor movement while taking information about our surrounding environment in order to interact with people, move without collision, alone or in a group, intercept, meet or escape to somebody. MimeTIC is an international key contributor in the domain of understanding and simulating locomotion and interactions between walkers. By combining an approach based on Human Movement Sciences and Computer Sciences, the team focuses on locomotor invariants which characterize the generation of locomotor trajectories, conducts challenging experiments focusing on visuo-motor coordination involved during interactions between walkers both using real and virtual set-ups. One main challenge is to consider and model not only the "average" behaviour of healthy younger adult but also extend to specific populations considering the effect of pathology or the effect of age (kids, older adults). As a first example, when patients cannot walk efficiently, in particular those suffering from central nervous system affections, it becomes very useful for practitioners to benefit from an objective evaluation of their capacities. To facilitate such evaluations, we have developed two complementary indices, one based on kinematics and the other one on muscle activations. One major point of our research is that such indices are usually only developed for children whereas adults with these affections are much more numerous. We extend this objective evaluation by using person-person interaction paradigm which allows studying visuo-motor strategies deficit in these specific populations.

Another fundamental question is the adaptation of the walking pattern according to anatomical constraints, such as pathologies in orthopedics, or adaptation to various human and non-human primates in paleoanthropology. Hence, the question is to predict plausible locomotion according to a given morphology. This question raises fundamental questions about the variables that are regulated to control gait: balance control, minimum energy, minimum jerk...In MimeTIC we develop models and simulators to efficiently test hypothesis on gait control for given morphologies.

POTIOC Project-Team

4. Application Domains

4.1. Education

Education is at the core of the motivations of the Potioc group. Indeed, we are convinced that the approaches we investigate—which target motivation, curiosity, pleasure of use and high level of interactivity—may serve education purposes. To this end, we collaborate with experts in Educational Sciences and teachers for exploring new interactive systems that enhance learning processes. We are currently investigating the fields of astronomy, optics, and neurosciences. We have also worked with special education centres for the blind on accessible augmented reality prototypes. Currently, we collaborate with teachers to enhance collaborative work for K-12 pupils. In the future, we will continue exploring new interactive approaches dedicated to education, in various fields. Popularization of Science is also a key domain for Potioc. Focusing on this subject allows us to get inspiration for the development of new interactive approaches.

4.2. Art

Art, which is strongly linked with emotions and user experiences, is also a target area for Potioc. We believe that the work conducted in Potioc may be beneficial for creation from the artist point of view, and it may open new interactive experiences from the audience point of view. As an example, we have worked with colleagues who are specialists in digital music, and with musicians. We have also worked with jugglers and we are currently working with a scenographer with the goal of enhancing interactivity of physical mockups and improve user experience.

4.3. Entertainment

Similarly, entertainment is a domain where our work may have an impact. We notably explored BCI-based gaming and non-medical applications of BCI, as well as mobile Augmented Reality games. Once again, we believe that our approaches that merge the physical and the virtual world may enhance the user experience. Exploring such a domain will raise numerous scientific and technological questions.

4.4. Well-being

Finally, well-being is a domain where the work of Potioc can have an impact. We have notably shown that spatial augmented reality and tangible interaction may favor mindfulness activities, which have been shown to be beneficial for well-being. More generally, we explore *introspectibles* objects, which are tangible and augmented objects that are connected to physiological signals and that foster introspection. We explore these directions for general public, including people with special needs.

TITANE Project-Team

4. Application Domains

4.1. Domain 1

In addition to tackling enduring scientific challenges, our research on geometric modeling and processing is motivated by applications to computational engineering, reverse engineering, digital mapping and urban planning. The main deliverable of our research will be algorithms with theoretical foundations. Ultimately we wish to contribute making geometry modeling and processing routine for practitioners who deal with real-world data. Our contributions may also be used as a sound basis for future software and technology developments.

Our first ambition for technology transfer is to consolidate the components of our research experiments in the form of new software components for the CGAL (Computational Geometry Algorithms Library) library. Consolidation being best achieved with the help of an engineer, we will search for additional funding. Through CGAL, we wish to contribute to the “standard geometric toolbox”, so as to provide a generic answer to application needs instead of fragmenting our contributions. We already cooperate with the Inria spin-off company Geometry Factory, which commercializes CGAL, maintains it and provide technical support.

Our second ambition is to increase the research momentum of companies through advising Cifre Ph.D. theses and postdoctoral fellows on topics that match our research program.

ALMANACH Project-Team

4. Application Domains

4.1. Application domains for ALMAnaCH

ALMAnaCH's research areas cover Natural Language Processing (nowadays identified as a sub-domain of Artificial Intelligence) and Digital Humanities. Application domains are therefore numerous, as witnessed by ALMAnaCH's multiple academic and industrial collaborations, for which see the relevant sections. Examples of application domains for NLP include:

- Information extraction, information retrieval, text mining (ex.: opinion surveys)
- Text generation, text simplification, automatic summarisation
- Spelling correction (writing aid, post-OCR, normalisation of noisy/non-canonical texts)
- Machine translation, computer-aided translation
- Chatbots, conversational agents, question answering systems
- Medical applications (early diagnosis, language-based medical monitoring...)
- Applications in linguistics (modelling languages and their evolution, sociolinguistic studies...)
- Digital humanities (exploitation of text documents, for instance in historical research)

COML Team

4. Application Domains

4.1. Speech processing for underresourced languages

We plan to apply our algorithms for the unsupervised discovery of speech units to problems relevant to language documentation and the construction of speech processing pipelines for underresourced languages.

4.2. Tools for the analysis of naturalistic speech corpora

Daylong recordings of speech in the wild gives rise a to number of specific analysis difficulties. We plan to use our expertise in speech processing to develop tools for performing signal processing and helping annotation of such resources for the purpose of phonetic or linguistic analysis.

MULTISPEECH Project-Team

4. Application Domains

4.1. Introduction

Approaches and models developed in the MULTISPEECH project are intended to be used for facilitating oral communication in various situations through enhancements of communication channels, either directly via automatic speech recognition or speech production technologies, or indirectly, thanks to computer assisted language learning. Applications also include the usage of speech technologies for helping people in handicapped situations or for improving their autonomy. Foreseen application domains are related to multimodal computer interaction, annotation and processing of spoken documents, health and autonomy (more precisely aided communication and monitoring), and computer assisted learning.

4.2. Multimodal Computer Interactions

Speech synthesis has tremendous applications in facilitating communication in a human-machine interaction context to make machines more accessible. For example, it started to be widely common to use acoustic speech synthesis in smartphones to make possible the uttering of all the information. This is valuable in particular in the case of handicap, as for blind people. Audiovisual speech synthesis, when used in an application such as a talking head, i.e., virtual 3D animated face synchronized with acoustic speech, is beneficial in particular for hard-of-hearing individuals. This requires an audiovisual synthesis that is intelligible, both acoustically and visually. A talking head could be an intermediate between two persons communicating remotely when their video information is not available, and can also be used in language learning applications as vocabulary tutoring or pronunciation training tool. Expressive acoustic synthesis is of interest for the reading of a story, such as an audiobook, as well as for better human-machine interactions.

4.3. Annotation and Processing of Spoken Documents and Audio Archives

A first type of annotation consists in transcribing a spoken document in order to get the corresponding sequences of words, with possibly some complementary information, such as the structure (punctuation) or the modality (affirmation/question) of the utterances to make the reading and understanding easier. Typical applications of the automatic transcription of radio or TV shows, or of any other spoken document, include making possible their access by deaf people, as well as by text-based indexing tools.

A second type of annotation is related to speech-text alignment, which aims at determining the starting and ending times of the words, and possibly of the sounds (phonemes). This is of interest in several cases such as for annotating speech corpora for linguistic studies, and for synchronizing lip movements with speech sounds (for example, for avatar-based communications). Although good results are currently achieved on clean data, automatic speech-text alignment needs to be improved for properly processing noisy spontaneous speech data and needs to be extended to handle overlapping speech.

Finally, there is also a need for speech signal processing techniques in the field of multimedia content creation and rendering. Relevant techniques include speech and music separation, speech equalization, speech enhancement, prosody modification, and speaker conversion.

4.4. Aided Communication and Monitoring

Source separation techniques should help for locating and monitoring people through the detection of sound events inside apartments, and speech enhancement is mandatory for hands-free vocal interactions. A foreseen application aims at improving the autonomy of elderly or disabled people, and also fits with smartroom applications. In a longer perspective, adapting speech recognition technologies to the voice of elderly people should also be useful for such applications, but this requires the recording of adequate databases. Sound monitoring in other application fields (security, environmental monitoring) can also be envisaged.

4.5. Computer Assisted Learning

Although speaking seems quite natural, learning foreign languages, or learning the mother tongue for people with language deficiencies, represents critical cognitive stages. Hence, many scientific activities have been devoted to these issues either from a production or a perception point of view. The general guiding principle with respect to computer assisted mother or foreign language learning is to combine modalities or to augment speech to make learning easier. Based upon an analysis of the learner's production, automatic diagnoses can be considered. However, making a reliable diagnosis on each individual utterance is still a challenge, which is dependent on the precision and quality of the segmentation of the speech utterance into phones, and of the computed prosodic parameters.

PANAMA Project-Team

4. Application Domains

4.1. Acoustic Scene Capture

Acoustic fields carry much information about audio sources (musical instruments, speakers, etc.) and their environment (e.g., church acoustics differ much from office room acoustics). A particular challenge is to capture as much information from a complete 3D+t acoustic field associated with an audio scene, using as few sensors as possible. The feasibility of compressive sensing to address this challenge was shown in certain scenarii, and the actual implementation of this framework will potentially impact practical scenarii such as remote surveillance to detect abnormal events, e.g. for health care of the elderly or public transport surveillance.

4.2. Audio Signal Separation in Reverberant Environments

Audio signal separation consists in extracting the individual sound of different instruments or speakers that were mixed on a recording. It is now successfully addressed in the academic setting of linear instantaneous mixtures. Yet, real-life recordings, generally associated to reverberant environments, remain an unsolved difficult challenge, especially with many sources and few audio channels. Much of the difficulty comes from the estimation of the unknown room impulse response associated to a matrix of mixing filters, which can be expressed as a dictionary-learning problem. Solutions to this problem have the potential to impact, for example, the music and game industry, through the development of new digital re-mastering techniques and virtual reality tools, but also surveillance and monitoring applications, where localizing audio sources is important.

4.3. Multimedia Indexing

Audiovisual and multimedia content generate large data streams (audio, video, associated data such as text, etc.). Manipulating large databases of such content requires efficient techniques to: segment the streams into coherent sequences; label them according to words, language, speaker identity, and more generally to the type of content; index them for easy querying and retrieval, etc. As the next generation of online search engines will need to offer content-based means of searching, the need to drastically reduce the computational burden of these tasks is becoming all the more important as we can envision the end of the era of wasteful datacenters that can increase forever their energy consumption. Most of today's techniques to deal with such large audio streams involve extracting features such as Mel Frequency Cepstral Coefficients (MFCC) and learning high-dimensional statistical models such as Gaussian Mixture Models, with several thousand parameters. The exploration of a compressive learning framework is expected to contribute to new techniques to efficiently process such streams and perform segmentation, classification, etc., in the compressed domain. A particular challenge is to understand how this paradigm can help exploiting truly multimedia features, which combine information from different associated streams such as audio and video, for joint audiovisual processing.

4.4. Remote Sensing

Remote sensing is an image acquisition paradigm in which various pieces of information from the surface of the earth are obtained by airborne or satellite imaging. In the recent years the quality and variety of sensors available on the market for such remote acquisitions has drastically improved, leaving a large number of unsolved or difficult data processing tasks on the hand of data scientists. In particular, members of the PANAMA project-team have focused on processing hyperspectral images, that are images collected at a large number of wavelengths. The tasks that have been addressed are mainly spectral unmixing, that aims at identifying the various materials on the earth surface with their relative abundances, and denoising, that aims at removing structured noise from remote acquisitions that can be polluted by e.g. sensor movement or clouds. The spectra measured at each pixel of an hyperspectral image are heavily dependent on surface features (material, slop, material mixture type) which makes these data extremely rich and interesting.

SEMAGRAMME Project-Team

4. Application Domains

4.1. Deep Semantic Analysis

Our applicative domains concern natural language processing applications that rely on a deep semantic analysis. For instance, one may cite the following ones:

- textual entailment and inference,
- dialogue systems,
- semantic-oriented query systems,
- content analysis of unstructured documents,
- text transformation and automatic summarization,
- (semi) automatic knowledge acquisition.

4.2. Text Transformation

Text transformation is an application domain featuring two important sub-fields of computational linguistics:

- parsing, from surface form to abstract representation,
- generation, from abstract representation to surface form.

Text simplification or automatic summarization belong to that domain.

We aim at using the framework of Abstract Categorical Grammars we develop to this end. It is indeed a reversible framework that allows both parsing and generation. Its underlying mathematical structure of λ -calculus makes it fit with our type-theoretic approach to discourse dynamics modeling.

Auctus Team

4. Application Domains

4.1. Factory 4.0

The 4th industrial revolution (factory 4.0) is characterized by the integration of digital technologies into the production process, in order to meet the challenge of customizing services and products. This agility requires making manufacturing and maintenance lines flexible and versatile. This capacity for adaptation is the characteristic of the human being, which puts him at the center of the production apparatus. However, this can no longer be done at the expense of their health and well-being. How then can we reconcile the enhancement of our manual and analytical expertise, the ever desired increase in productivity and manufacturing quality, while reducing the hardship at work? Collaborative robotics, which we are seeking to build, is one of the central solutions to meet this societal challenge. By assisting humans in their most dangerous and painful tasks, it complements or replaces them in their phases of physical and cognitive fragility.

More generally, we are interested in workstation cobotization, in the manufacturing and assembly industry but also in the construction and craft industries. The application areas are related to regional needs in aeronautics, including maintenance, water and waste treatment. In most of these cases, it is possible to define the tasks, evaluate the stakes and added value of our work.

CHORALE Team

4. Application Domains

4.1. Transportation of people and goods

The researches developed in CHORALE can be applied in different applications fields. We are particularly interested in *transportation of people and goods*. CHORALE contribute to the development of Autonomous Connected Vehicles (e.g. Learning, Mapping, Localization, Navigation) and the associated services (e.g. towing, platooning, taxi).

4.2. Assistance and services robots

We are also interested in *assistance and service robotics*. By adding the social level in the representation of the environment, CHORALE develops social and proactive navigation.

4.3. Survey and monitoring of environments

CHORALE studies multi-robots systems to explore, survey and monitor different kind of environments (e.g. agricultural space, forest space, destroyed space), which are of particular importance in the field of *Exploration and monitoring of poorly structured environments*.

CHROMA Project-Team

4. Application Domains

4.1. Introduction

Applications in Chroma are organized in two main domains : **i) Future cars and transportation systems and ii) Services robotics**. These domains correspond to the experimental fields initiated in Grenoble (eMotion team) and in Lyon (CITI lab). However, the scientific objectives described in the previous sections are intended to apply equally to both applicative domains. Even our work on Bayesian Perception is today applied to the intelligent vehicle domain, we aim to generalize to any mobile robots. The same remark applies to the work on multi-agent decision making. We aim to apply algorithms to any fleet of mobile robots (service robots, connected vehicles, UAVs). This is the philosophy of the team since its creation.



Figure 4. Most of the Chroma platforms: the Pepper robot, a fleet of (22) Turtlebot 2, one of the 4 Bebop drones and the equipped Toyota Lexus.

4.2. Future cars and transportation systems

Thanks to the introduction of new sensor and ICT technologies in cars and in mass transportation systems, and also to the pressure of economical and security requirements of our modern society, this application domain is quickly changing. Various technologies are currently developed by both research and industrial laboratories. These technologies are progressively arriving at maturity, as it is witnessed by the results of large scale experiments and challenges such as the Google's car project and several future products announcements made by the car industry. Moreover, the legal issue starts to be addressed in USA (see for instance the recent laws in Nevada and in California authorizing autonomous vehicles on roads) and in several other countries (including France).

In this context, we are interested in the development of ADAS⁰ systems aimed at improving comfort and safety of the cars users (e.g., ACC, emergency braking, danger warnings), and of Fully Autonomous Driving functions for controlling the displacements of private or public vehicles in some particular driving situations and/or in some equipped areas (e.g., automated car parks or captive fleets in downtown centers or private sites).

⁰ Advanced Driver Assistance Systems

Since about 8 years, we are collaborating with Toyota and with Renault-Nissan on these applications (bilateral contracts, PhD Theses, shared patents), but also recently with Volvo group (PhD thesis started in 2016). We are also strongly involved (since 2012) in the innovation project Perfect then now Security for autonomous vehicle of the IRT⁰ Nanoelec (transportation domain). In 2016, we have been awarded a European H2020 ECSEL project⁰ involving major European automotive constructors and car suppliers. In this project, Chroma is focusing on the embedded perception component (models and algorithms, including the certification issue), in collaboration with Renault, Valeo and also with the Inria team TAMIS (Rennes). Chroma is also involved in the ANR project "Valet" (2015-2018) coordinated by the Inria team RITS (Rocquencourt), dealing with automatic redistribution of car-sharing vehicles and parking valet; Chroma is involved in the pedestrian-vehicle interaction for a safe navigation.

In this context, Chroma has two experimental vehicles equipped with various sensors (a Toyota Lexus and a Renault Zoe, see Fig. 4 and Fig. 2 .b), which are maintained by Inria-SED⁰ and that allow the team to perform experiments in realistic traffic conditions (Urban, road and highway environments). The Zoe car has been automated in December 2016, through our collaboration with the team of P. Martinet (IRCCyN Lab, Nantes), that allow new experiments in the team.

4.3. Services robotics

Service robotics is an application domain quickly emerging, and more and more industrial companies (e.g., IS-Robotics, Samsung, LG) are now commercializing service and intervention robotics products such as vacuum cleaner robots, drones for civil or military applications, entertainment robots ... One of the main challenges is to propose robots which are sufficiently robust and autonomous, easily usable by non-specialists, and marked at a reasonable cost. We are involved in developing observation and surveillance systems, by using ground robots and aerial ones, see Fig. 4 . Since 2016, we develop solutions for 3D observation/exploration of complex scenes or environments with a fleet of UAVs (DynaFlock Inria/DGA project, Inria ADT CORDES⁰) or mobile robots (COMODYS FIL project [82]).

A more recent challenge for the coming decade is to develop robotized systems for assisting elderly and/or disabled people. In the continuity of our work in the IPL PAL⁰, we aim to propose smart technologies to assist electric wheelchair users in their displacements and also to control autonomous cars in human crowds (see Figure 17 for illustration). This concerns our recent "Hianic" ANR project. Another emerging application is humanoid robots helping humans at their home or work. In this context, we address the problem of NAMO (Navigation Among Movable Obstacles) in human populated environments (eg. PhD of B. Renault started on 2018). More generally we address navigation and reconnaissance tasks with Pepper humanoids in the context of the RoboCup-Social League.

⁰Institut de Recherche Technologique

⁰ENABLE-S3: European Initiative to Enable Validation for Highly Automated Safe and Secure Systems.

⁰Service Expérimentation et Développement

⁰Coordination d'une Flotte de Drones Connectés pour la Cartographie 3D d'édifices, led by O. Simonin.

⁰Personnally assisted Living

DEFROST Project-Team

4. Application Domains

4.1. Industry

Robotics in the manufacturing industry is already widespread and is one of the strategies put in place to maintain the level of competitiveness of companies based in France and to avoid relocation to cheap labor countries. Yet, in France, it is considered that the level of robotization is insufficient, compared to Germany for instance. One of the challenges is the high investment cost for the acquisition of robotic arms. In recent years, this challenge has led to the development of “generic” and “flexible” (but rigid) robotic solutions that can be mass produced. But their applicability to specific tasks is still challenging or too costly. With the development of 3D printing, we can imagine the development of a complete opposite strategy: a “task-specific” design of robots. Given a task that needs to be performed by a deformable robot, we could optimize its shape and its structure to create the set of desired motions. A second important aspect is the reduction of the manufacturing cost: it is often predicted that the cost of deformable robots will be low compared to classical rigid robots. The robot could be built on one piece using rapid prototyping or 3D printers and be more adapted for collaborative work with operators. In this area, using soft materials is particularly convenient as they provide a mass/carried load ratio several orders of magnitude higher than traditional robots, highly decreasing the kinetic energy thus increasing the motion speed allowed in presence of humans. Moreover, the technology allows more efficient and ergonomic wearable robotic devices, opening the option for exo-skeletons to be used by human operators inside the factories and distribution centers. This remains to be put in place, but it can open new perspectives in robotic applications. A last remarkable property of soft robots is their adaptability to fragile or tortuous environments. For some particular industry fields (chemistry, food industry...) this could also be an advantage compared to existing rigid solutions. For instance, the German company <http://www.festo.com>, key player in the industrial robotics field, is experimenting with deformable trunk robots that exhibit great compliance and adaptability, and we are working on their accurate control.

4.2. Personal and service robotics

The personal and service robotics are considered an important source of economic expansion in the coming years. The potential applications are numerous and in particular include the challenge of finding robotic solutions for active and healthy aging at home. We plan to develop functional orthosis for which it is better not to have a rigid exoskeleton that is particularly uncomfortable. These orthosis will be ideally personalized for each patient and built using rapid prototyping. On this topic, the place of our team will be to provide algorithms for controlling the robots. We will find some partners to build these robots that would fall in the category of “wearable robots”. With this thematic we also connect with a strong pole of excellence of the region on intelligent textiles (see [Up-TeX](#)) and with the strategic plan of Inria (Improving Rehabilitation and Autonomy).

4.3. Entertainment industry and arts

Robots have a long history with entertainment and arts where [animatronics](#) have been used for decades for cinematographic shootings, theater, amusement parks ([Disney’s audio-animatronic](#)) and performing arts. We believe that soft robots could be a good support for art. As an example, last year we collaborated with the artist Dewi Brunet in the creation of animated origami structures (see <https://dewiorigami.com/>).

4.4. Medical Applications

Soft robots have many medical applications as their natural compliance makes them safer than traditional robots when interacting with humans. Such robots can be used for minimally invasive surgery, to access and act on remote parts of the body through minimal incisions in the patient. Applications include laparoscopic and brain surgery, treatment of several cancers including prostate cancer, and cardiology, for example percutaneous coronary interventions.



Figure 1. Dewi Brunet and his “Origami 2.0” exhibit at the AMV in Trélon

As an example, we received an industry grant (CIFRE) with Robocath to work on autonomous catheter navigation. See section 8.2. Another application is cochlear implant surgery.

FLOWERS Project-Team

4. Application Domains

4.1. Application Domains

Neuroscience, Developmental Psychology and Cognitive Sciences The computational modelling of life-long learning and development mechanisms achieved in the team centrally targets to contribute to our understanding of the processes of sensorimotor, cognitive and social development in humans. In particular, it provides a methodological basis to analyze the dynamics of the interaction across learning and inference processes, embodiment and the social environment, allowing to formalize precise hypotheses and later on test them in experimental paradigms with animals and humans. A paradigmatic example of this activity is the Neurocuriosity project achieved in collaboration with the cognitive neuroscience lab of Jacqueline Gottlieb, where theoretical models of the mechanisms of information seeking, active learning and spontaneous exploration have been developed in coordination with experimental evidence and investigation, see <https://flowers.inria.fr/neurocuriosityproject/>.

Personal and lifelong learning robotics Many indicators show that the arrival of personal robots in homes and everyday life will be a major fact of the 21st century. These robots will range from purely entertainment or educative applications to social companions that many argue will be of crucial help in our society. Yet, to realize this vision, important obstacles need to be overcome: these robots will have to evolve in unpredictable homes and learn new skills in a lifelong manner while interacting with non-engineer humans after they left factories, which is out of reach of current technology. In this context, the refoundation of intelligent systems that developmental robotics is exploring opens potentially novel horizons to solve these problems. In particular, this application domain requires advances in artificial intelligence that go beyond the current state-of-the-art in fields like deep learning. Currently these techniques require tremendous amounts of data in order to function properly, and they are severely limited in terms of incremental and transfer learning. One of our goals is to drastically reduce the amount of data required in order for this very potent field to work. We try to achieve this by making neural networks aware of their knowledge, i.e. we introduce the concept of uncertainty, and use it as part of intrinsically motivated multitask learning architectures, and combined with techniques of learning by imitation.

Human-Robot Collaboration. Robots play a vital role for industry and ensure the efficient and competitive production of a wide range of goods. They replace humans in many tasks which otherwise would be too difficult, too dangerous, or too expensive to perform. However, the new needs and desires of the society call for manufacturing system centered around personalized products and small series productions. Human-robot collaboration could widen the use of robot in this new situations if robots become cheaper, easier to program and safe to interact with. The most relevant systems for such applications would follow an expert worker and works with (some) autonomy, but being always under supervision of the human and acts based on its task models.

Environment perception in intelligent vehicles. When working in simulated traffic environments, elements of FLOWERS research can be applied to the autonomous acquisition of increasingly abstract representations of both traffic objects and traffic scenes. In particular, the object classes of vehicles and pedestrians are of interest when considering detection tasks in safety systems, as well as scene categories ("scene context") that have a strong impact on the occurrence of these object classes. As already indicated by several investigations in the field, results from present-day simulation technology can be transferred to the real world with little impact on performance. Therefore, applications of FLOWERS research that is suitably verified by real-world benchmarks has direct applicability in safety-system products for intelligent vehicles.

Automated Tutoring Systems. Optimal teaching and efficient teaching/learning environments can be applied to aid teaching in schools aiming both at increase the achievement levels and the reduce time needed. From a practical perspective, improved models could be saving millions of hours of students' time (and effort) in learning. These models should also predict the achievement levels of students in order to influence teaching practices.

HEPHAISTOS Project-Team (section vide)

LARSEN Project-Team

4. Application Domains

4.1. Personal Assistance

During the last fifty years, many medical advances as well as the improvement of the quality of life have resulted in a longer life expectancy in industrial societies. The increase in the number of elderly people is a matter of public health because although elderly people can age in good health, old age also causes embrittlement, in particular on the physical plan which can result in a loss of autonomy. That will force us to re-think the current model regarding the care of elderly people.⁰ Capacity limits in specialized institutes, along with the preference of elderly people to stay at home as long as possible, explain a growing need for specific services at home.

Ambient intelligence technologies and robotics could contribute to this societal challenge. The spectrum of possible actions in the field of elderly assistance is very large. We will focus on activity monitoring services, mobility or daily activity aids, medical rehabilitation, and social interactions. This will be based on the experimental infrastructure we have built in Nancy (Smart apartment platform) as well as the deep collaboration we have with OHS.⁰

4.2. Civil Robotics

Many applications for robotics technology exist within the services provided by national and local government. Typical applications include civil infrastructure services⁰ such as: urban maintenance and cleaning; civil security services; emergency services involved in disaster management including search and rescue; environmental services such as surveillance of rivers, air quality, and pollution. These applications may be carried out by a wide variety of robot and operating modality, ranging from single robots or small fleets of homogeneous or heterogeneous robots. Often robot teams will need to cooperate to span a large workspace, for example in urban rubbish collection, and operate in potentially hostile environments, for example in disaster management. These systems are also likely to have extensive interaction with people and their environments.

The skills required for civil robots match those developed in the LARSEN project: operating for a long time in potentially hostile environment, potentially with small fleets of robots, and potentially in interaction with people.

⁰See the Robotics 2020 Multi-Annual Roadmap [35].

⁰OHS (*Office d'Hygiène Sociale*) is an association managing several rehabilitation or retirement home structures.

⁰See the Robotics 2020 Multi-Annual Roadmap [35], section 2.5.

PERVASIVE Project-Team

4. Application Domains

4.1. Modelling of awareness and expertise from Eye Gaze and Emotion

Humans display awareness and expertise through a variety of non-verbal channels. It is increasingly possible to record and interpret such information with available technology. In the ANR CEEGE project, we have constructed an instrument for capturing and interpreting multimodal signals of humans engaged in solving challenging problems. Our instrument captures eye gaze, fixations, body postures, and facial expressions signals from humans engaged in interactive tasks on a touch screen.

An initial experiment with multi-modal observation of human experts engaged in solving problems in Chess revealed an unexpected observation of rapid changes in emotion as players attempt to solve challenging problems. In a scientific collaboration with the NeuroCognition group at the Univ Bielefeld, we have constructed a model to explain for chess experts that explains these unexpected results. This model has recently been tested in a second experiment with 22 chess players. Our results indicate that chess players associate emotions to chess chunks, and reactively use these associations to guide search in chunks for planning and problem solving. These results have recently been reported in a paper at the International Conference on Multimodal Interaction, and is the subject of the nearly completed doctoral dissertation of Thomas Guntz.

The results are currently being used in the construction of a student aware driver training device to be commercialized by SME company Sym2B financed by the SATT Linksium Project MAT: Monitoring Attention of Trainees, starting in Sept 2019. IN this project we will construct a training simulator for operation of busses and tramways.

4.2. Narrative Description of Kitchen Activities from Egocentric Video

We have developed and evaluated a system to construct situated, narrative descriptions of cooking activities including food preparation, place setting, cleaning and placing objects in storage areas. We are specifically interested in real-time, on-line techniques that recognize and interpret food types, food states and manipulation actions for transformation preparation of food. We are exploring techniques for detecting, modelling, and recognising a large vocabulary of actions and activities under different observational conditions, and describing these activities in a larger context.

A full understanding of human actions requires: recognising what action has been performed, predicting how it will affect the surrounding environment, explaining why this action has been performed, and who is performing it. Classic approaches to action recognition interpret a spatio-temporal pattern in a video sequence to tell what action has been performed, and perhaps how and where it was performed. A more complete understanding requires information about why the action was performed, and how it affects the environment. This face of understanding can be provided by explaining the action as part of a narrative.

Most work on recognition of cooking activities has concentrated on recognizing actions from the spatio-temporal patterns of hand motions. While some cooking activities may be directly recognized from motion, the resulting description is incomplete, as it does not describe the state of the ingredients, and how these have been transformed by cooking actions. A fuller description requires a description of how food ingredients have been transformed during the food preparation process.

We have addressed the automatic construction of cooking narratives by first detecting and locating ingredients and tools used in food preparation. We then recognize actions that involve transformations of ingredients, such as "slicing", and use these transformations to segment the video stream into visual events. We can then interpret detected events as a causal sequence of voluntary actions, confirmed by spatio-temporal transformation patterns, in order to automatically provide a narrative.

Our method is inspired by the intuition that object states are visually more apparent than actions from a still frame and thus provide information that is complementary to spatio-temporal action recognition. We define a state transition matrix that maps action labels into a pre-state and a post-state. We identify key frames, and use these to learn appearance models of objects and their states. For recognition, we use a modified form of VGG neural network trained via transfer learning with a specially constructed data set of images of food types and food sates. Manipulation actions are hypothesized from the state transition matrix and provide complementary evidence to spatio temporal action recognition.

4.3. Embedded Computer Vision for low-power Bolometric Imaging

In cooperation with Schneider Electric, we have developed techniques for embedded real time image analysis and tracking algorithms using Bolometric Imaging Sensors. Such sensors capture light in the far infrared and return an image where each pixel is a measurement of temperature in degree centigrade. We have designed an integrated low-cost sensor that combines an 80x80 pixel Bolometric imager with a low power micro-processor. The device provides real-time, embedded image processing for target detection and activity analysis where all sensing and interpretation is local. No images are recorded and only relevant information is about activity is communicated. The design of this system is under consideration for a patent, and thus has not been published. The software system has been registered with French APP , and is to be licensed to Schneider Electric for use a line of sensors for detecting falls and other activities for assisted living for seniors, as well as monitoring customer activities in commercial environments.

4.4. Recognizing and predicting routine activities in smart homes

Most research on smart home systems has concentrated on techniques for recognizing context. However, many categories of service require information about likely future context. We have developed an approach that uses dynamic Bayesian networks to predict future activity and context in a smart home. Our approach extends a state-of-the-art prediction model with three contributions. First, we include sensor data through aggregation nodes, instead of limiting ourselves only to higher level context data. Second, our method uses higher order relations between activities, so that past activities can have a more meaningful impact on prediction of future activities. Third, we use a latent node that estimates the cognitive state of the occupant.

4.5. User centred energy management

Participants: Amr Al-Zouhri Al-Yafi, Amine Awada, Patrick Reignier Partners: UMR G-SCOP, UMR LIG (Persuasive Interaction, IIHM), CEA Liten, PACTE, Vesta Systems and Elithis.

Work in this area explores techniques for a user centric energy management system, where user needs and tacit knowledge drive the search of solutions. These are calculated using a flexible energy model of the living areas. The system is personified by energy consultants with which building actors such as building owners, building managers, technical operators but also occupants, can interact in order to co-define energy strategies, benefiting of both assets: tacit knowledge of human actors, and measurement with computation capabilities of calculators. Putting actors in the loop, i.e. making energy not only visible but also controllable is the needed step before large deployment of energy management solutions. It is proposed to develop interactive energy consultants for all the actors, which are energy management aided systems embedding models in order to support the decision making processes. MIRROR (interactive monitoring), WHAT-IF (interactive quantitative simulation), EXPLAIN (interactive qualitative simulation), SUGGEST-AND-ADJUST (interactive management) and RECOMMEND (interactive diagnosis) functionalities will be developed.

4.6. E-Textile

Participant: Sabine Coquillart

Partner: LIMSI

Collaboration with the HAPCO team from LIMSI on e-textiles. A patent application has been filed related to this work:

- F. Bimbard, M. Bobin, M. Ammi, S. Coquillart "Procédé de conception d'un capteur de flexion textile piézorésistif à partir de fils fonctionnels", Patent Application, 2017.

4.7. Interaction with Pervasive Media

Participants: Sabine Coquillart, Jingtao Chen

Partners: Inria GRA, GIPSA, G-SCOP

Pseudo-haptic feedback is a technique aiming to simulate haptic sensations without active haptic feedback devices. Pseudo-haptic techniques have been used to simulate various haptic feedbacks such as stiffness, torques, and mass. In the framework of Jingtao Chen PhD thesis, a novel pseudo-haptic experiment has been set up. The aim of this experiment is to study the EMG signals during a pseudo-haptic task. A stiffness discrimination task similar to the one published in Lecuyer's PhD thesis has been chosen. The experimental set-up has been developed, as well as the software controlling the experiment. Pre-tests are under way. They will be followed by the tests with subjects.

RAINBOW Project-Team (section vide)

RITS Project-Team

4. Application Domains

4.1. Introduction

While the preceding section focused on methodology, in connection with automated guided vehicles, it should be stressed that the evolution of the problems which we deal with remains often guided by the technological developments. We enumerate three fields of application whose relative importance varies with time and which have strong mutual dependencies: driving assistance, cars available in self-service mode and fully automated vehicles (cybercars).

4.2. Driving assistance

Several techniques will soon help drivers. One of the first immediate goal is to improve security by alerting the driver when some potentially dangerous or dangerous situations arise, i.e. collision warning systems or lane tracking could help a bus driver and surrounding vehicle drivers to more efficiently operate their vehicles. Human factors issues could be addressed to control the driver workload based on additional information processing requirements. Another issue is to optimize individual journeys. This means developing software for calculating optimal (for the user or for the community) paths. Nowadays, path planning software is based on a static view of the traffic: efforts have to be done to take the dynamic component in account.

4.3. New transportation systems

The problems related to the abusive use of the individual car in large cities led the populations and the political leaders to support the development of public transport. A demand exists for a transport of people and goods which associates quality of service, environmental protection and access to the greatest number. Thus the tram and the light subways of VAL type recently introduced into several cities in France conquered the populations, in spite of high financial costs. However, these means of mass transportation are only possible on lines on which there is a keen demand. As soon as one moves away from these “lines of desire” or when one deviates from the rush hours, these modes become expensive and offer can thus only be limited in space and time. To give a more flexible offer, it is necessary to plan more individual modes which approach the car as we know it. However, if one wants to enjoy the benefits of the individual car without suffering from their disadvantages, it is necessary to try to match several criteria: availability anywhere and anytime to all, lower air and soils pollution as well as sound levels, reduced ground space occupation, security, low cost. Electric or gas vehicles available in self-service, as in the Praxitèle system, bring a first response to these criteria. To be able to still better meet the needs, it is however necessary to re-examine the design of the vehicles on the following points:

- ease empty car moves to better distribute them;
- better use of information systems inboard and on ground;
- better integrate this system in the global transportation system.

These systems are now operating. The challenge is to bring them to an industrial phase by transferring technologies to these still experimental projects.

4.4. Automated vehicles

The long term effort of the project is to put automatically guided vehicles (cybercars) on the road. It seems too early to mix cybercars and traditional vehicles, but data processing and automation now make it possible to consider in the relatively short term the development of such vehicles and the adapted infrastructures. RITS aims at using these technologies on experimental platforms (vehicles and infrastructures) to accelerate the technology transfer and to innovate in this field. Other application can be precision docking systems that will allow buses to be automatically maneuvered into a loading zone or maintenance area, allowing easier access for passengers, or more efficient maintenance operations. Transit operating costs will also be reduced through decreased maintenance costs and less damage to the braking and steering systems. Regarding technical topics, several aspects of Cybercars have been developed at RITS this year. First, we have stabilized a generic Cycab architecture involving Inria SynDEx tool and CAN communications. The critical part of the vehicle is using a real-time SynDEx application controlling the actuators via two Motorola's MPC555. Today, we have decided to migrate to the new dsPIC architecture for more efficiency and ease of use. This application has a second feature, it can receive commands from an external source (Asynchronously to this time) on a second CAN bus. This external source can be a PC or a dedicated CPU, we call it high level. To work on the high level, in the past years we have been developing a R&D framework called (Taxi) which used to take control of the vehicle (Cycab and Yamaha) and process data such as gyro, GPS, cameras, wireless communications and so on. In order to rely on a professional and maintained solution, we use the RTMaps SDK development platform for our developments and demonstrations. These demonstrations include: reliable SLAMMOT algorithm using 2 to 4 laser sensors simultaneously, automatic line/road following techniques, PDA remote control, multi sensors data fusion, collaborative perception via ad-hoc network. The second main topic is inter-vehicle communications using ad-hoc networks. We have worked with the EVA team for setting and tuning OLSR, a dynamic routing protocol for vehicles communications. Our goal is to develop a vehicle dedicated communication software suite, running on a specialized hardware. It can be linked also with the Taxi Framework for getting data such GPS information's to help the routing algorithm.

LINKMEDIA Project-Team

4. Application Domains

4.1. Asset management in the entertainment business

Media asset management—archiving, describing and retrieving multimedia content—has turned into a key factor and a huge business for content and service providers. Most content providers, with television channels at the forefront, rely on multimedia asset management systems to annotate, describe, archive and search for content. So do archivists such as the Institut National de l’Audiovisuel, the bibliothèque Nationale de France, the Nederlands Instituut voor Beeld en Geluid or the British Broadcast Corporation, as well as media monitoring companies, such as Yacast in France. Protecting copyrighted content is another aspect of media asset management.

4.2. Multimedia Internet

One of the most visible application domains of linked multimedia content is that of multimedia portals on the Internet. Search engines now offer many features for image and video search. Video sharing sites also feature search engines as well as recommendation capabilities. All news sites provide multimedia content with links between related items. News sites also implement content aggregation, enriching proprietary content with user-generated content and reactions from social networks. Most public search engines and Internet service providers offer news aggregation portals. This also concerns TV on-demand and replay services as well as social TV services and multi-screen applications. Enriching multimedia content, with explicit links targeting either multimedia material or knowledge databases is central here.

4.3. Data journalism

Data journalism forms an application domain where most of the technology developed by LINKMEDIA can be used. On the one hand, data journalists often need to inspect multiple heterogeneous information sources, some being well structured, some other being fully unstructured. They need to access (possibly their own) archives with either searching or navigational means. To gradually construct insight, they need collaborative multimedia analytics processes as well as elements of trust in the information they use as foundations for their investigations. Trust in the information, watching for adversarial and/or (deep) fake material, accountability are all crucial here.

MAGRIT Team

4. Application Domains

4.1. Augmented reality

We have a significant experience in AR that allowed good progress in building usable, reliable and robust AR systems. Our contributions cover the entire process of AR: matching, pose initialization, 3D tracking, in-situ modeling, handling interaction between real and virtual objects....

4.2. Medical Imaging

For 20 years, we have been working in close collaboration with University Hospital of Nancy and GE Healthcare in interventional neuroradiology. Our common aim is to develop a multimodality framework to assist therapeutic decisions and interventional gestures. Contributions of the team focus on the development of AR tools for neuro-navigation as well as the development of simulation tools for training or planning. Laparoscopic surgery is another field of interest with the development of methods for tracking deformable organs based on bio-mechanical models. Some of these projects are developed in collaboration with the MIMESIS project team.

4.3. Experimental mechanics

In experimental solid mechanics, an important problem is to characterize properties of specimens subject to mechanical constraints, which makes it necessary to measure tiny strains. Contactless measurement techniques have emerged in the last few years and are spreading quickly. They are mainly based on images of the surface of the specimen on which a regular grid or a random speckle has been deposited. We are engaged since June 2012 in a transdisciplinary collaboration with Institut Pascal (Clermont-Ferrand Université). The aim is to characterize the metrological performances of these techniques limited by, e.g., the sensor noise, and to improve them by several dedicated image processing tools.

MORPHEO Project-Team

4. Application Domains

4.1. 4D modeling

Modeling shapes that evolve over time, analyzing and interpreting their motion has been a subject of increasing interest of many research communities including the computer vision, the computer graphics and the medical imaging communities. Recent evolutions in acquisition technologies including 3D depth cameras (Time-of-Flight and Kinect), multi-camera systems, marker based motion capture systems, ultrasound and CT scanners have made those communities consider capturing the real scene and their dynamics, create 4D spatio-temporal models, analyze and interpret them. A number of applications including dense motion capture, dynamic shape modeling and animation, temporally consistent 3D reconstruction, motion analysis and interpretation have therefore emerged.

4.2. Shape Analysis

Most existing shape analysis tools are local, in the sense that they give local insight about an object's geometry or purpose. The use of both geometry and motion cues makes it possible to recover more global information, in order to get extensive knowledge about a shape. For instance, motion can help to decompose a 3D model of a character into semantically significant parts, such as legs, arms, torso and head. Possible applications of such high-level shape understanding include accurate feature computation, comparison between models to detect defects or medical pathologies, and the design of new biometric models.

4.3. Human Motion Analysis

The recovery of dense motion information enables the combined analysis of shapes and their motions. Typical examples include the estimation of mean shapes given a set of 3D models or the identification of abnormal deformations of a shape given its typical evolutions. The interest arises in several application domains where temporal surface deformations need to be captured and analyzed. It includes human body analyses for which potential applications are anyway numerous and important, from the identification of pathologies to the design of new prostheses.

4.4. Virtual and Augmented Reality

This domain has actually seen new devices emerge that enable now full 3D visualization, for instance the HTC Vive, the Microsoft HoloLens and the Magic Leap one. These devices create a need for adapted animated 3D contents that can either be generated or captured. We believe that captured 4D models will gain interest in this context since they provide realistic visual information on moving shapes that tend to avoid negative perception effects such as the uncanny valley effect. Besides 3D visualization devices, many recent applications also rely on everyday devices, such as mobile phones, to display augmented reality contents with free viewpoint ability. In this case, 3D and 4D contents are also expected.

PERCEPTION Project-Team (section vide)

SIROCCO Project-Team

4. Application Domains

4.1. Overview

The application domains addressed by the project are:

- Compression with advanced functionalities of various imaging modalities
- Networked multimedia applications taking into account needs in terms of user and network adaptation (e.g., interactive streaming, resilience to channel noise)
- Content editing, post-production, and computational photography.

4.2. Compression of emerging imaging modalities

Compression of visual content remains a widely-sought capability for a large number of applications. This is particularly true for mobile applications, as the need for wireless transmission capacity will significantly increase during the years to come. Hence, efficient compression tools are required to satisfy the trend towards mobile access to larger image resolutions and higher quality. A new impulse to research in video compression is also brought by the emergence of new imaging modalities, e.g. high dynamic range (HDR) images and videos (higher bit depth, extended colorimetric space), light fields and omni-directional imaging.

Different video data formats and technologies are envisaged for interactive and immersive 3D video applications using omni-directional videos, stereoscopic or multi-view videos. The "omni-directional video" set-up refers to 360-degree view from one single viewpoint or spherical video. Stereoscopic video is composed of two-view videos, the right and left images of the scene which, when combined, can recreate the depth aspect of the scene. A multi-view video refers to multiple video sequences captured by multiple video cameras and possibly by depth cameras. Associated with a view synthesis method, a multi-view video allows the generation of virtual views of the scene from any viewpoint. This property can be used in a large diversity of applications, including Three-Dimensional TV (3DTV), and Free Viewpoint Video (FVV). In parallel, the advent of a variety of heterogeneous delivery infrastructures has given momentum to extensive work on optimizing the end-to-end delivery QoS (Quality of Service). This encompasses compression capability but also capability for adapting the compressed streams to varying network conditions. The scalability of the video content compressed representation and its robustness to transmission impairments are thus important features for seamless adaptation to varying network conditions and to terminal capabilities.

4.3. Networked visual applications

Free-viewpoint Television (FTV) is a system for watching videos in which the user can choose its viewpoint freely and change it at anytime. To allow this navigation, many views are proposed and the user can navigate from one to the other. The goal of FTV is to propose an immersive sensation without the disadvantage of Three-dimensional television (3DTV). With FTV, a look-around effect is produced without any visual fatigue since the displayed images remain 2D. However, technical characteristics of FTV are large databases, huge numbers of users, and requests of subsets of the data, while the subset can be randomly chosen by the viewer. This requires the design of coding algorithms allowing such a random access to the pre-encoded and stored data which would preserve the compression performance of predictive coding. This research also finds applications in the context of Internet of Things in which the problem arises of optimally selecting both the number and the position of reference sensors and of compressing the captured data to be shared among a high number of users.

Broadband fixed and mobile access networks with different radio access technologies have enabled not only IPTV and Internet TV but also the emergence of mobile TV and mobile devices with internet capability. A major challenge for next internet TV or internet video remains to be able to deliver the increasing variety of media (including more and more bandwidth demanding media) with a sufficient end-to-end QoS (Quality of Service) and QoE (Quality of Experience).

4.4. Editing, post-production and computational photography

Editing and post-production are critical aspects in the audio-visual production process. Increased ways of “consuming” visual content also highlight the need for content repurposing as well as for higher interaction and editing capabilities. Content repurposing encompasses format conversion (retargeting), content summarization, and content editing. This processing requires powerful methods for extracting condensed video representations as well as powerful inpainting techniques. By providing advanced models, advanced video processing and image analysis tools, more visual effects, with more realism become possible. Our activities around light field imaging also find applications in computational photography which refers to the capability of creating photographic functionalities beyond what is possible with traditional cameras and processing tools.

Stars Project-Team (section vide)

THOTH Project-Team

4. Application Domains

4.1. Visual applications

Any solution to automatically understanding images and videos on a semantic level will have an immediate impact on a wide range of applications. For example:

- Semantic-level image and video access is highly relevant for visual search on the Web, in professional archives and personal collections.
- Visual data organization is applicable to organizing family photo and video albums as well as to large-scale information retrieval.
- Visual object recognition has potential applications ranging from surveillance, service robotics for assistance in day-to-day activities as well as the medical domain.
- Action recognition is highly relevant to visual surveillance, assisted driving and video access.
- Real-time scene understanding is relevant for human interaction through devices such as HoloLens, Oculus Rift.

4.2. Pluri-disciplinary research

Machine learning is intrinsically pluri-disciplinary. By developing large-scale machine learning models and algorithms for processing data, the Thoth team became naturally involved in pluri-disciplinary collaborations that go beyond visual modelling. In particular,

- extensions of unsupervised learning techniques originally developed for modelling the statistics of natural images have been deployed in neuro-imaging for fMRI data with the collaboration of the Parietal team from Inria.
- similarly, deep convolutional data representations, also originally developed for visual data, have been successfully extended to the processing of biological sequences, with collaborators from bio-informatics.
- Thoth also collaborates with experts in natural language and text processing, for applications where visual modalities need to be combined with text data.

WILLOW Team

4. Application Domains

4.1. Introduction

We believe that foundational modeling work should be grounded in applications. This includes (but is not restricted to) the following high-impact domains.

4.2. Quantitative image analysis in science and humanities

We plan to apply our 3D object and scene modeling and analysis technology to image-based modeling of human skeletons and artifacts in anthropology, and large-scale site indexing, modeling, and retrieval in archaeology and cultural heritage preservation. Most existing work in this domain concentrates on image-based rendering, that is, the synthesis of good-looking pictures of artifacts and digs. We plan to focus instead on quantitative applications. We are engaged in a project involving the archaeology laboratory at ENS and focusing on image-based artifact modeling and decorative pattern retrieval in Pompeii. Application of our 3D reconstruction technology is now being explored in the field of cultural heritage and archeology by the start-up Iconem, founded by Y. Ubelmann, a Willow collaborator.

4.3. Video Annotation, Interpretation, and Retrieval

Both specific and category-level object and scene recognition can be used to annotate, augment, index, and retrieve video segments in the audiovisual domain. The Video Google system developed by Sivic and Zisserman (2005) for retrieving shots containing specific objects is an early success in that area. A sample application, suggested by discussions with Institut National de l'Audiovisuel (INA) staff, is to match set photographs with actual shots in film and video archives, despite the fact that detailed timetables and/or annotations are typically not available for either medium. Automatically annotating the shots is of course also relevant for archives that may record hundreds of thousands of hours of video. Some of these applications will be pursued in our MSR-Inria project.