

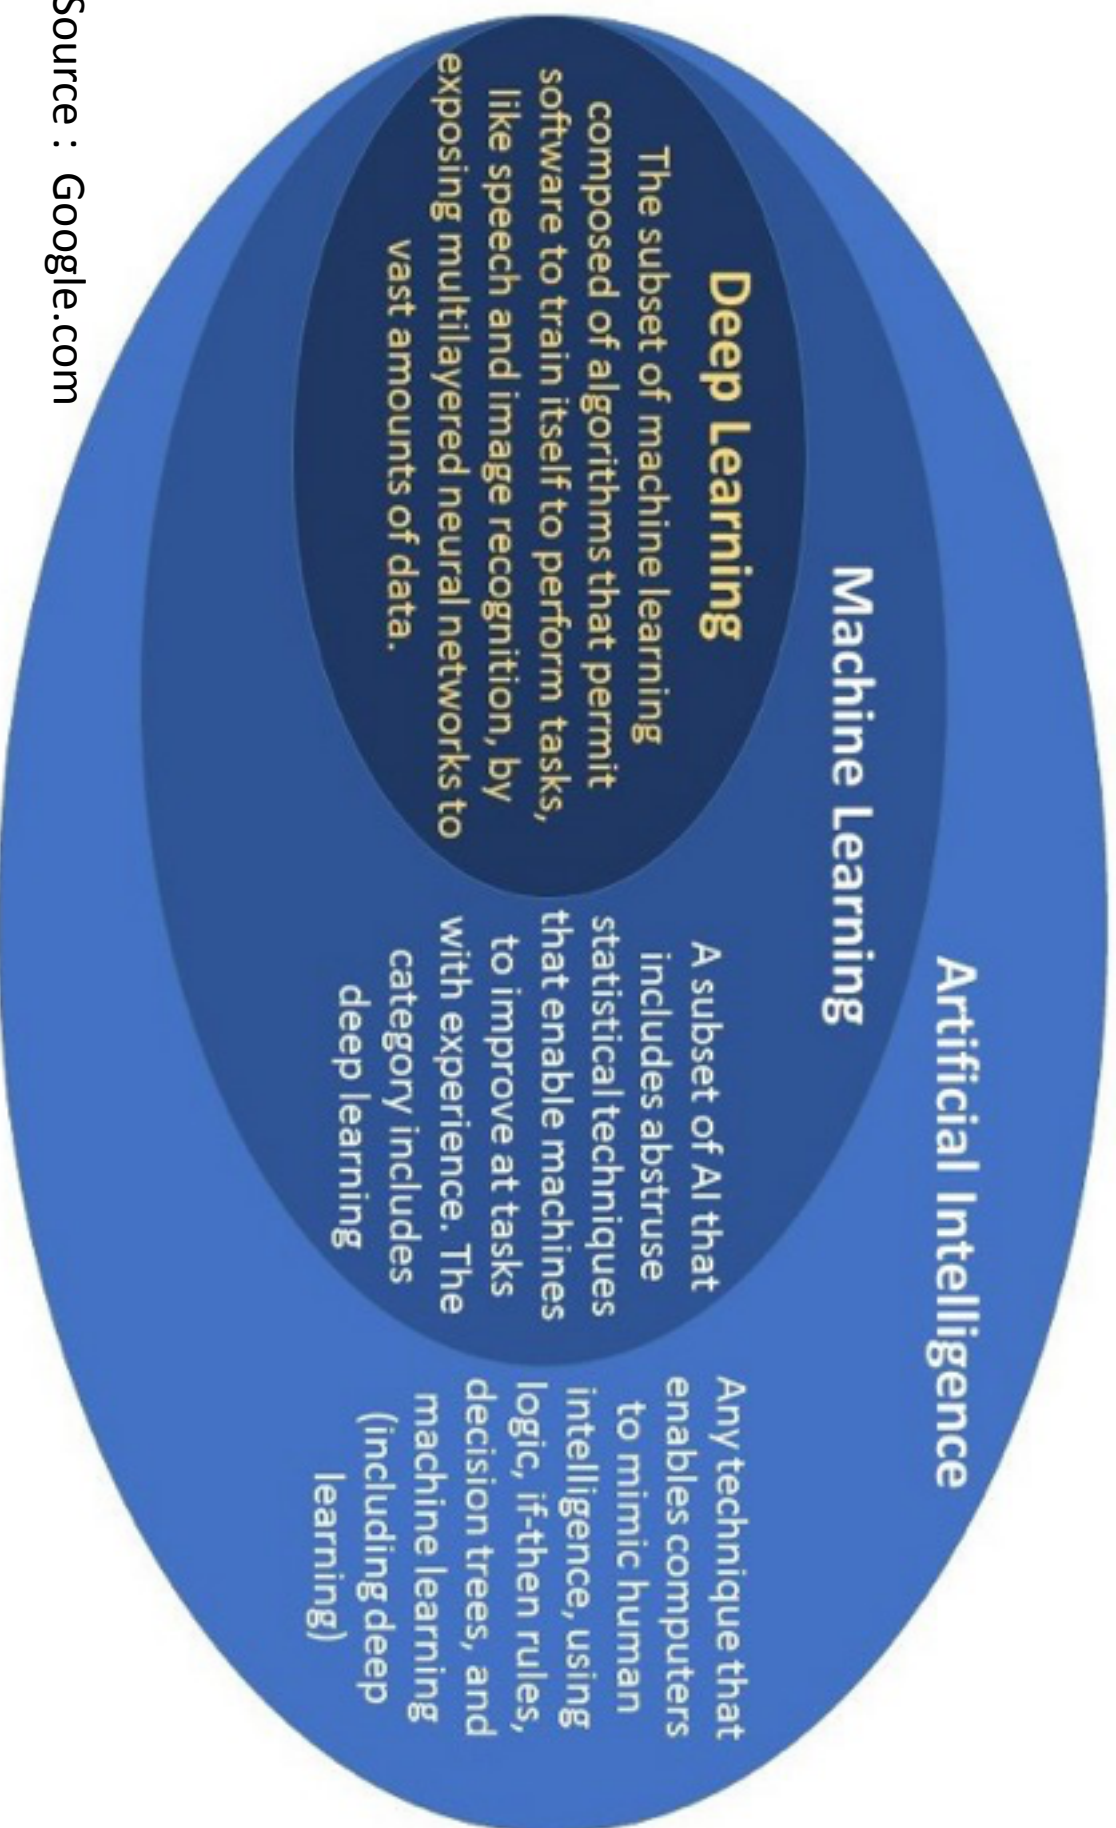
Machine Learning/Deep Learning

Survол rapide et quelques généralités...

Olivier J.J. MICHEL

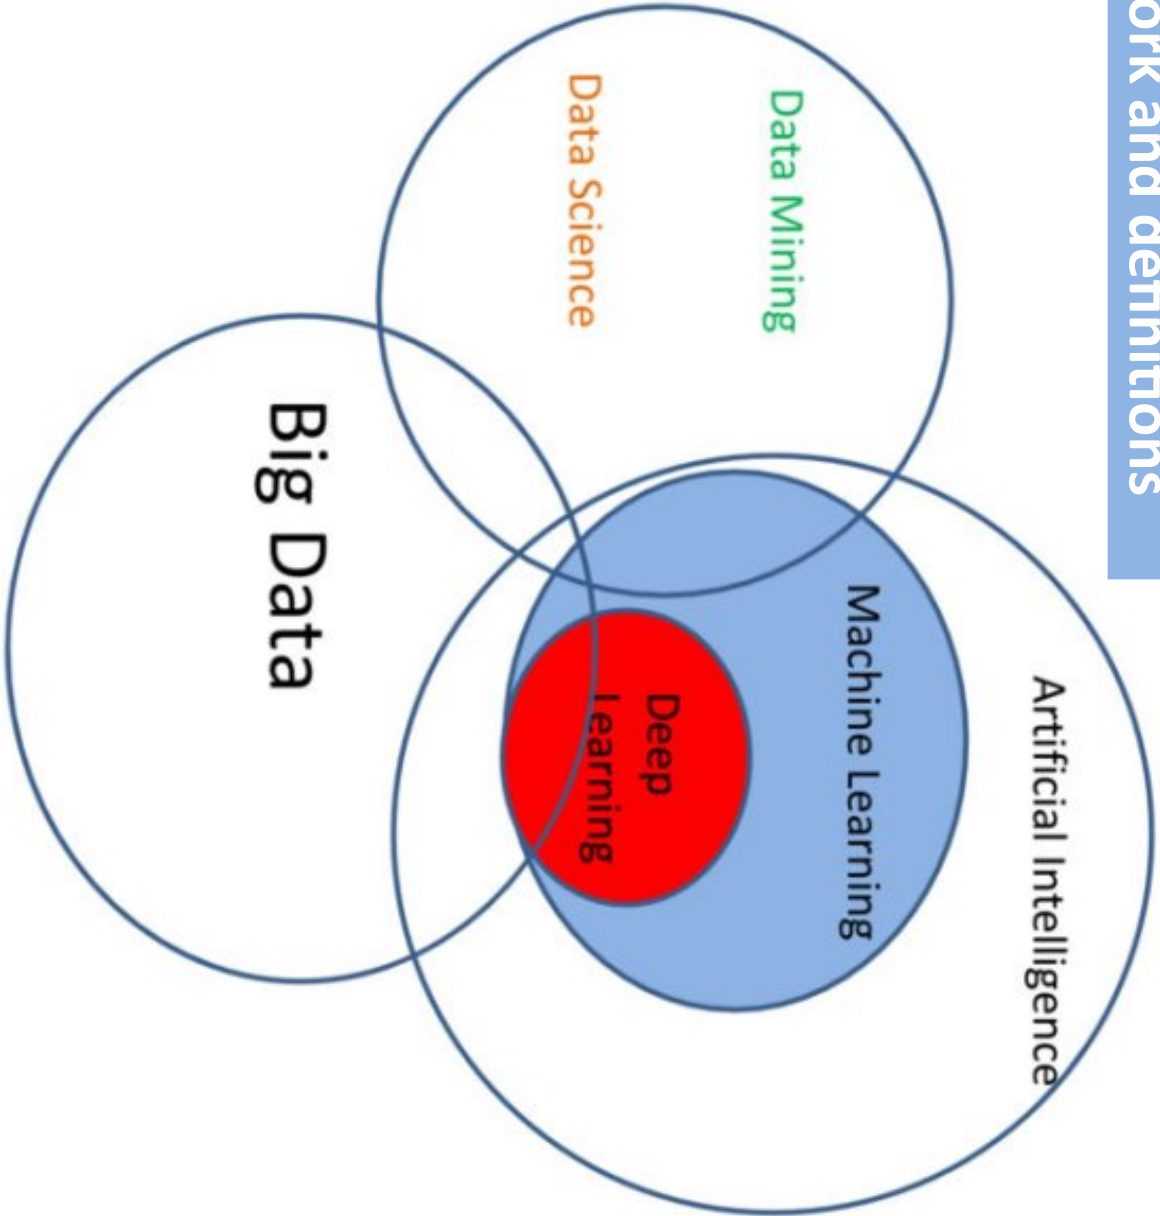


Framework and definitions



Source : Google.com

Framework and definitions




To construct computer systems that automatically improves through experience

Key feature = far easier to train a system by showing input-output examples, that to program output by anticipating all possible inputs.

= scalability to « Big Data » sets, data mining ability

Approaches allowing to analyse high throughput data in a novel way

Recent Evolution of Performances in Image classification



Huge incidence in
in computer science
in industry concerned with data intensive
issues

in « empirical sciences »

social sciences





biology

Again Because ~~cosmology~~ ^{ML} allows to analyse high through
data in a novel way

*Machine learning is having a substantial effect on many
areas of technology and science; examples of recent applied success stories
include robotics and autonomous vehicle control (top left),
speech processing and natural language processing (top right),
neuroscience research (middle),
and applications in computer vision (bottom).*

Source :

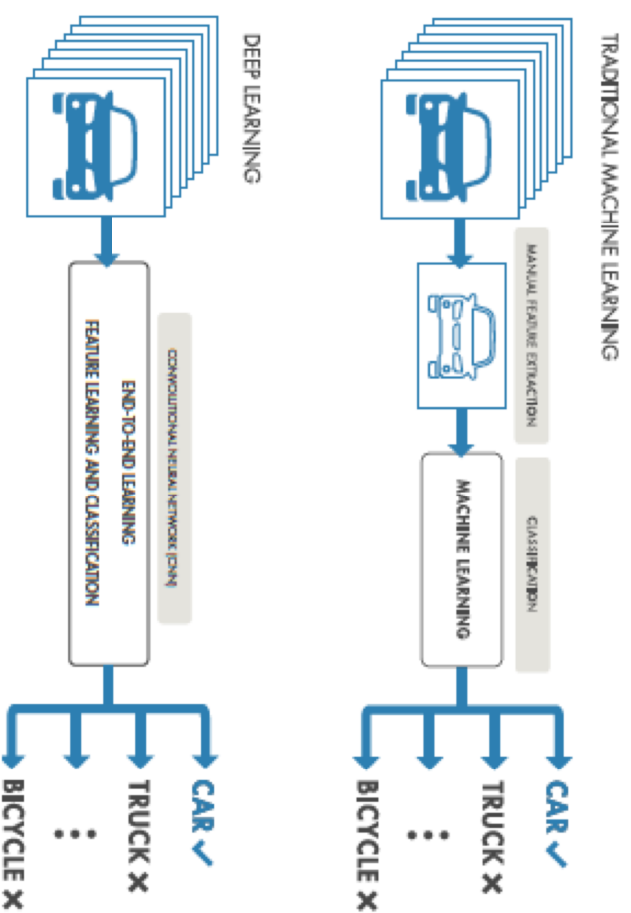
*Machine Learning: Trends, perspectives, and prospects
M.I.Jordan, T.M.Mitchell, Science 349, 255 (2015)*



What is the Difference Between Deep Learning and Machine Learning?

Deep learning is a subtype of machine learning. With machine learning, you manually extract the relevant features of an image. With deep learning, you feed the raw images directly into a deep neural network that learns the features automatically.


Deep learning often requires hundreds of thousands or millions of images for the best results. It's also computationally intensive and requires a high-performance GPU.




source  MathWorks

Machine Learning	Deep Learning
+ Good results with small data sets	— Requires very large data sets
+ Quick to train a model	— Computationally intensive
— Need to try different features and classifiers to achieve best results	+ Learns features and classifiers automatically
— Accuracy plateaus	+ Accuracy is unlimited

Basic components for NN, DNN, CNN




An illustration of an artificial neuron. Source: Becoming Human.



Deep NN

There is no exact formula for selecting layers. The best approach is to try a few and see how well they work—or to use a pretrained network.



Convolutional NN

=> On n'attend pas des perf « optimales »...multiplicité de minima locaux...

What Makes Deep Learning State-of-the-Art?

What are the challenges/problems raised ?

« Easy » access to massive sets a (labeled) data



(-> Privacy/security/Network throughput)

High computing power,
high perf GPU

->Cost of Training from massive
amount of data

Algorithmic developments, learning strategies
(Transfer Learning,
Reinforcement learning,
Optimization methods...)

-> Perf analysis, bounds and guarantees?



A data science perspective to define key scientific and practical goals :

-> How to theoretically characterize the capabilities of a (deep) learning algo?

*how accurately does the algo learn from a given type and volume of training data?
*how robust is the ML algo to errors in the model or sensitivity to missing data?

*Given a learning task + training data,

is the design of an algorithm possible (how/ tractable?)

(how much data, which computational cost for a given task?)

Rk : DL not (yet) used in critical contexts (Defense, medicine, ...)

→ requires blends of (in addition to knowledge in the engineering field of expertise)

Statistical decision theory,

Computational complexity theory

Information theory, Sampling theory,

Optimization methods, to provide bounds on convergence rates

Computer programming, architecture, data manipulation

=> **Grenoble has many strength in this perspective**

I.E. MUCH MORE THAN « BLACK BOX BASED » solution design skills often found

- in the increasing demand from
 - many students
 - many companies offering internships
- or even in some published (rank A) literature :

Guides for deep learning submissions in IEEE T-IFS

Aiming to strengthen scientific reproducibility among our peers, this document lays down some guidelines for submissions to IEEE T-IFS.

Reproducibility is one aspect that should be assessed by reviewers and considered by the AE in the final acceptance decision.

The guidelines explicitly refer to papers dealing with steganalysis and multimedia forensics, however they are applicable to all papers submitted to T-IFS in which the proposed system relies on deep learning.

Thanks to recent advances in machine learning, forensic and steganalytic attacks have grown in sophistication and complexity. **With more parameters, it becomes easier to omit key information.**

This is particularly so for Deep Learning (DL) methods, which have many variants and details. When some of these are omitted from the paper it weakens reproducibility, and often extends the review process.

Authors are encouraged to make both code and data sets available! (preferably in time for the review process) but this is not a formal requirement at the moment.

In any case, supplementary code is not a substitute

for a proper written description of the method. To help authors, the T-IFS EB has drawn up this checklist of details, that **DL submissions on IFS topics should take particular care** to include.

- **The type of network (CNN, RNN, BDRNN, ...) and loss function.**

- **The topology:**

- o type of each layer (pooling, convolution, non-linearity, ...),
- o activation function of each layer (tanh, sigmoid, max- or average-pooling, ...)
- o any parameters (strides, weights size, etc).

- **Any pre-processing of network inputs, and interpretation of outputs.** Note that there are different re-scalings commonly called "normalization", so this needs to be specified.

Training phase

- **How the network is initialized.**

- o If randomly, give the exact parameters of the initializing distributions.
- o If there is pre-training, the procedure used (SAE, RBM, etc) and any parameters.

- **The learning algorithm** (SGD, Adagrad, AdaDelta, HF, etc).

- o Hyperparameters of the learning algorithm: (momentum, etc).
- o Mini-batches and (if appropriate) number of representatives of each class in learning batches. Where appropriate (e.g. steganalysis), specify if the data are stratified (i.e. paired cover and stego kept together).
- o If used, exactly what form of batch normalization and pooling.
- o Initial learning rate and its evolution during the learning, number of iterations of epochs, dropout or other regularization factors, etc.
- o The order that training data is visited during the learning process, and whether shuffled between epochs.

From the Transactions website: "The Transactions encourages authors to make their publications reproducible by making all information needed to reproduce the presented results available online.

This typically requires publishing the code and data used to produce the publication's figures and tables on a website; see the supplemental materials section of the information for authors.

It gives other researchers easier access to the work, and facilitates fair comparisons."

- Stopping criterion (is there a fixed number of iterations or not?)
- Which data was used to determine the hyperparameters, learning rates, regularization parameters, and stopping criterion; the policy for hyperparameter search.

Test phase

- **How data is broken down into training, validation, and testing sets.**

- **If cross-validation is used, how it is arranged and whether it is repeated.**

Implementation

- If a GPU is used, which model.
- If you have used a standard deep learning tool, the exact version and any tool-specific parameters.
- If you propose techniques for scalability (kernel approximations, etc) then specify.

We encourage the inclusion of a "methodology" section in the paper to contain these details, which could be in an appendix, or as supplementary material providing that this is available for review.

Algorithms and methods need not be described from scratch: a good citation is adequate, as long as the exact variety is clear and all parameters have been specified.

While we encourage the publication of negative results contradicting previously published research, in these cases a higher standard of reproducibility applies. Barring exceptional circumstances,

data sets should be made available for the review process (if not, this should be explained to the AE). **Finally, authors should demonstrate the statistical significance of results claimed.**

Some future trends and challenges (among others)

Theory : How does a DNN learn ?

Information Bottleneck principle ? (IT approach)

Stat phys Approaches ?

large random matrices

How does a DNN preserve hierarchical structure ?

Capsule Networks ?

Mallat's approach (using wavelet feature maps)

How to set on IT basis, the complexity of the DNN? (VC dim, Stoch compl.)

New models and structures :

Generative adversarial networks (unsupervised learning)

-> learning with less labeled data

Hybrid learning models


-> combines models of DNNs with probabilistic (Bayesian) ones
to model uncertainty

Time/Data trade offs :

-> subsampling, random projections, sparse data representations
(time/space budget wrt to accuracy requirements...)

-> DNN cannot address all problems outside the framework of large data

➔ **DEVELOP DNN techniques that produce explainable/provable models**
-> perf bounds, guarantees



LABEX + PERSYVAL 2 since 2019

- + fundamental/theoretical research in advanced data mining (action teams)
- + collaborative (multi-labo), research projects
- + school and conference support



« An interdisciplinary research Institute on how data change science and society ».

Axis 1 → **application of cutting-edge data science in several scientific domains**

Axis 2 → data-driven research as a major tool for research in Social Sciences and Humanities

Axis 3 → how the data revolution impacts society

- + real data and application oriented, forum, seminars, ...

Junior chair (cdp data@uga) C. Gomez « ML, DNN for exoplanet discovery »

Senior chair (idex UGA) R.Couillet « Analyse theorique en grandes dimensions, Matrices aleatoires, DNN »

(CDP IDEX), many labs from UGA

MIAI@Grenoble Alpes

> 2020, proposal in progress

MIAI@Grenoble Alpes: panorama

(source : E. Gaussier)

- 4 piliers au cœur de l'IA en interaction avec des thèmes intégrateurs et des activités interdisciplinaires
 - *Apprentissage et raisonnement* (apprentissage/science des données, modélisation des connaissances et raisonnement)
 - *IA matérielle, embarquée et HPC* (IA déportée, basse consommation, temps réel, calcul pour l'IA)
 - *IA et société* (éthique, droit, sciences humaines et sociales)
 - *Cerveau et cognition* (approches bio/psycho-inspirées, bio-mimétisme)
- 3 types d'acteurs en IA
 - *Les inventeurs* **CEA** - **INRIA** – **UNIV. GRENOBLE ALPES (UG, LJK, GIPSA, SHS...)** ATOS – CRITEO - HPE - IBM - NAVER - NVIDIA – PROBAYES (DOCAPOST) - SMART ME UP - ST MICROELECTRONICS - ...
 - *Les intégrateurs* ATOS - ALTRAN - CAP GEMINI/SOGETI – CRITEO - HPE - HURENCE - INCEPTIVE - NEOVISION - ORANGE - RAISE PARTNER - RED MINT NETWORK - SHORT EDITION – SOGLIUS - ...
 - *Les innovateurs/utilisateurs* AUTOMATIQUE & INDUSTRIE FRANCE – BIOMERIEUX - COSEVIT - DIGIGRAM - ENEDIS - EOLAS - FIVEFIVE - GLOBEVIP - FUL - ISKN - MEERSENSE - METADESK - MICHELIN - MICRODB - MORPHOSENSE - NETELER – ODI-T-E - OPHIRYS - OENONEO - POLLEN METROLOGY - QIOVA - SALESFORCE - SCHNEIDER ELECTRIC – SIEMENS I – TENEVIA - VESTA SYSTEM – VINCI - VISEO -...
- 3 dimensions
 - *Recherche fondamentale et appliquée*
 - *Formation*
 - *Innovation, transfert et valorisation*
- Domaines d'application privilégiés
 - **Santé**
 - **Environnement et énergie**
 - Mobilité et ville intelligente ?
 - Culture et éducation ?



2017-18 :

- * **Audio-visual analysis for human-robot interaction** by [Radu Horaud](#) (INRIA Grenoble)
- * **Toward Interactive User Data Analytics** by [Siherm Amer- Yahia](#) (CNRS, LI6, Grenoble)
- * **Community detection in complex networks** by [Pierre Borgnat](#) (Lab.Phys ENS-Lyon)
- * **Neural networks and Random Matrix Theory** by [Romain Couillet](#) (Ecole Centrale Supélec, Univerity Paris-Saclay)
- * **Light propagation in complex media: from imaging, to compressive imaging and machine learning** by [Sylvain Gigan](#) (Kastler Brossel laboratory, Paris)
- * **Big data for telecom** by [Merouane Debbah](#) (Ecole Centrale Supélec, Univerity Paris-Saclay, Vice President R&D Center Huawei France, Director of the Mathematical and Algorithmic Sciences Lab)



SIGMA

MASTER IN SIGNAL AND IMAGE PROCESSING METHODS AND APPLICATIONS
MENTION: TRAITEMENT DU SIGNAL ET DES IMAGES

Data Challenge commun



2017-18 Audio-visual analysis for human-robot interaction

2018-19 Prédiction de diffusion de polluant

Ongoing Challenges on AI, Machine Learning and data sciences in Grenoble (non exhaustive)

- Challenge sur la détection d'exoplanète par imagerie directe : https://carlgogo.github.io/exoimaging_challenge/
- Challenge en ligne sur plusieurs mois (octobre 2018-février 2019), communauté astro.
- Challenge sur l'hétérogénéité cellulaire de tumeurs cancéreuses : Décembre 2018 à Aussois. Challenge à visée pluridisciplinaire (biologie et data science) et pédagogique.



GdR focused on statistical data sciences, Learning, IA,

GdR AI (LIG, J.Euzenat(INRIA/LIG), MC.Rousset (UGA/LIG))
Aspects formels et algorithmiques de l'IA

GDR IA

GdR MaDICS (C. Collet (INP/LIG), C.Bobineau(INP/LIG))
Masse de données, Information et Connaissances en Sciences



GdR ISIS (N.Lebihan (CNRS/GIPSA))
Information, Signal, Image, Vision



Liens vers les M2R

<http://mosig.imag.fr>

<https://msiam.imag.fr/doku.php>

<http://www.gipsa-lab.fr/SIGMA/>

Data Institute, Persyval-Lab

<https://data-institute.univ-grenoble-alpes.fr>

<https://persyval-lab.org>

Sources : Google.com, Mathworks.com, Sciences; IEEE SP magazine, IEEE T-IFS,
Cours College de France S.Mallat, ...