

REFRAME Project

March 2013–September 2016



<http://www.reframe-d2k.org/>

Rethinking the **E**ssence, **F**lexibility and **R**eusability of **A**dvanced **M**odel **E**xploitation

Nicolas Lachiche
(Université de Strasbourg, France)


Paris, 11–12 April 2018

Outline

- ▶ Scientific background
- ▶ Partners
- ▶ Follow-up results and impacts

Scientific background (1)

- ▶ What is “reframing”?
 - “process of applying an existing [machine learning/data mining] model to the new operating context by the proper transformation of inputs, outputs and patterns”.



Scientific background (2)

- ▶ What is a “context”?
 - Changes in data distribution.
 - Costs (output, inputs, labelling, etc.).
 - Data quality (noise, missing data)
 - Representation change (granularity, different attributes, background knowledge)
 - Task change
- One of the goals of the project was to generalise the notion of “operating context”.

Scientific background (3)

- ▶ Usual D2K process
 - Models are trained in a context: overfitting, retraining
 - Inefficient and unreliable process
 - Either a set of models for a range of operating contexts, or a very general, but inflexible model.
 - We aim at building versatile models that can be properly deployed in a range of operating contexts.
- ▶ The purpose is to ensure “model reuse”.
 - Appropriate reframing procedures to apply it to any possible context.

Consortium

- ▶ BRIS: University of Bristol / Intelligent Systems Laboratory
- ▶ STRAS: University of Strasbourg / ICube
- ▶ VAL: Universitat Politècnica de València / DSIC

University of Bristol

- ▶ Senior investigators
 - Peter Flach, Ricardo Prudencio (Brazil)
- ▶ Specificity and added value
 - Machine learning, in particular metrics: ROC, precision–recall, etc.
 - ECML PKDD 2014 Best paper award!
 - Classifier calibration
 - Abduction and background knowledge
 - Multilabel versatile decision trees

Université de Strasbourg

- ▶ Senior investigators
 - Nicolas Lachiche, Agnès Braud
- ▶ Specificity and added value
 - Relational data mining, complex aggregates
 - Practice of data mining: health, geography and environment, industry 4.0
 - Input and output reframing
 - Integration of tools into the online data mining server clowdflows.unistra.fr

Universitat Politècnica de València


- ▶ Senior investigators
 - Jose Hernandez–Orallo, Cesar Ferri, Maria–Jose Ramirez–Quintana
- ▶ Specificity and added value
 - Data mining, from theory to applications
 - Quantification, cost–sensitive regression
 - Multidimensional prediction models
 - Various other changes of contexts: learning task, noise, attribute costs
 - Item response theory

Follow-up results and impacts

- ▶ Since 2016
 - Reframing in Context: A systematic approach for model reuse in machine learning, AI Communications, 2016
 - Context-Aware Knowledge Discovery, Tutorial at ECMLPKDD 2016

Modelling ML models: general view

- ▶ PT-AI 2017 Paper: R. Fabra-Boluda, Cèsar Ferri, J. Hernández-Orallo, F. Martínez-Plumed, M. J. Ramírez-Quintana “Modelling Machine Learning Models”.
 - The model is kept (and hence reused) but a new abstraction A is built, in order to better understand, reframe or modify the existing model.



Fairness as a context (Valencia)

- ▶ Main idea:

- A trained classifier (for Acc or other metrics) can be reframed for a fairness metric:

Fairness metric	Definition
Equal coverage (parity)	$P(\hat{y}=T \mid x \in G) = P(\hat{y}=T)$
Equal odds	$P(\hat{y}=T \mid y=c, x \in G) = P(\hat{y}=T \mid y=c) \forall c$
(SUBCASE) Equal opportunity	$P(\hat{y}=T \mid y=T, x \in G) = P(\hat{y}=T \mid y=T)$
Equal mistreatment	$P(\hat{y}=F \mid y=T, x \in G) = P(\hat{y}=F \mid y=T)$ and $P(\hat{y}=T \mid y=F, x \in G) = P(\hat{y}=T \mid y=F)$
(SUBCASE) Equal accuracy	$P(\hat{y}=y \mid x \in G) = P(\hat{y}=y)$
Accurate coverage (group calibration)	$P(\hat{y}=T \mid x \in G) = P(y=T \mid x \in G)$

- When is it better than retraining?

- Source of the discrimination: legacy (the bias was in the data), inference (the bias is created or magnified by the algorithm) or deployment (application creates new bias).
- Means of the discrimination: direct discrimination (through an attribute representing the group) or **indirect** discrimination (through an attribute representing the group)


Modelling problems: latent variable

- ▶ One way of extracting latent variables about a problem is IRT (Item Response Theory)

- For instance, we can extract the “difficulty” of the problem, and consider “difficulty” a context.

- Prudencio et al. “Analysis of Instance Hardness in Machine Learning Using Item Response Theory”, LMCE 2015.

Model Characteristic Curve




Context-aware sensor fusion

Activity recognition in a smart home: context is the house layout.

<http://www.irc-sphere.ac.uk>

A Smart Home:
Locations and Sensors



Beta calibration

Context can be whether classifier is over- or underconfident.

<https://betacal.github.io>

Meelis Kull, Telmo Silva Filho, Peter Flach; Beta calibration: a well-founded and easily implemented improvement on logistic calibration for binary classifiers, AISTATS'17, PMLR 54:623-631, 2017.

Meelis Kull, Telmo Silva Filho, Peter Flach; Beyond sigmoids: How to obtain well-calibrated probabilities from binary classifiers with beta calibration. Electron. J. Statist. 11 (2017), no. 2, 5052--5080.

