Approaching the Legal Frontier: A Framework for Developing Legally Aligned Machine Learning Models in Finance

Mathias Hanson* and Sam Verboven
Vrije Universiteit Brussel
Data Analytics Laboratory

Gregory Lewkowicz
Université Libre de Bruxelles
Smart Law Hub

Problem Statement

- Financial institutions (FI) must comply with the law when deploying machine learning (ML) systems, which introduces complexities :
  - Legal Uncertainty for Operationalizations
  - Technical and Legal Trade-offs
  - Trade-offs in Metrics for Evaluation Vs. Holistic Legal Assessment
- **Key Challenge: How should FI develop ML systems to achieve (1) legal compliance and (2) high predictive performance simultaneously?**

Objective

The Development of a Legally Aligned ML system includes :

1. Legally grounded constraints for the ML model development;
2. Optimization of ML performance within uncertain legal boundaries ; and
3. A holistic legal compliance evaluation adapted to the ML *paradigm*

Current Approaches and Their Limitations

Legal Requirements / Software Engineering with focus on Traditional Software

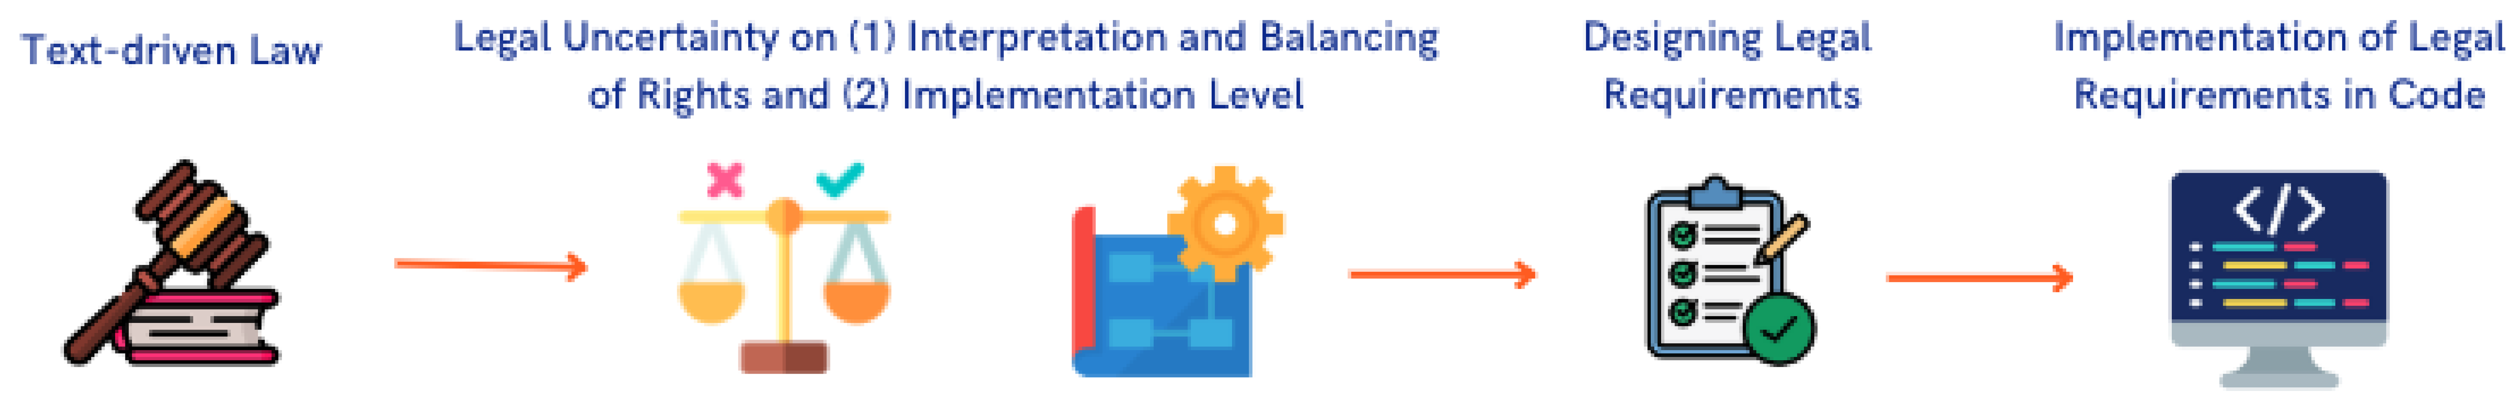

- **Law-Centric Design Framework** ✓
- **DM* under Legal Uncertainty** ✓
- **ML-Adapted Design Framework**
- **DM* given Technical Trade-offs**

* Decision-making

Requirements / Software Engineering 4 AI

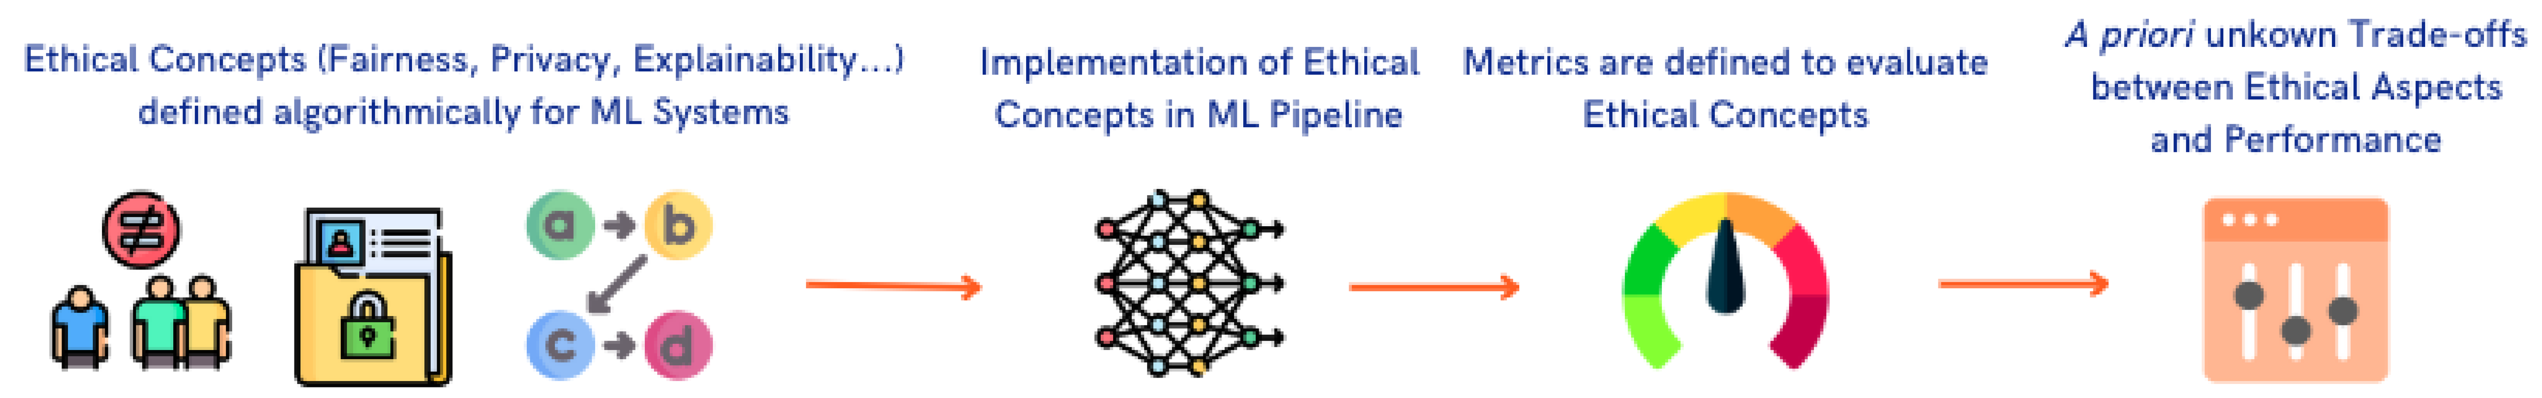

- **Law-Centric Design Framework**
- **DM* under Legal Uncertainty**
- **ML-Adapted Design Framework** ✓
- **DM* given Technical Trade-offs** ✓

Our Contribution

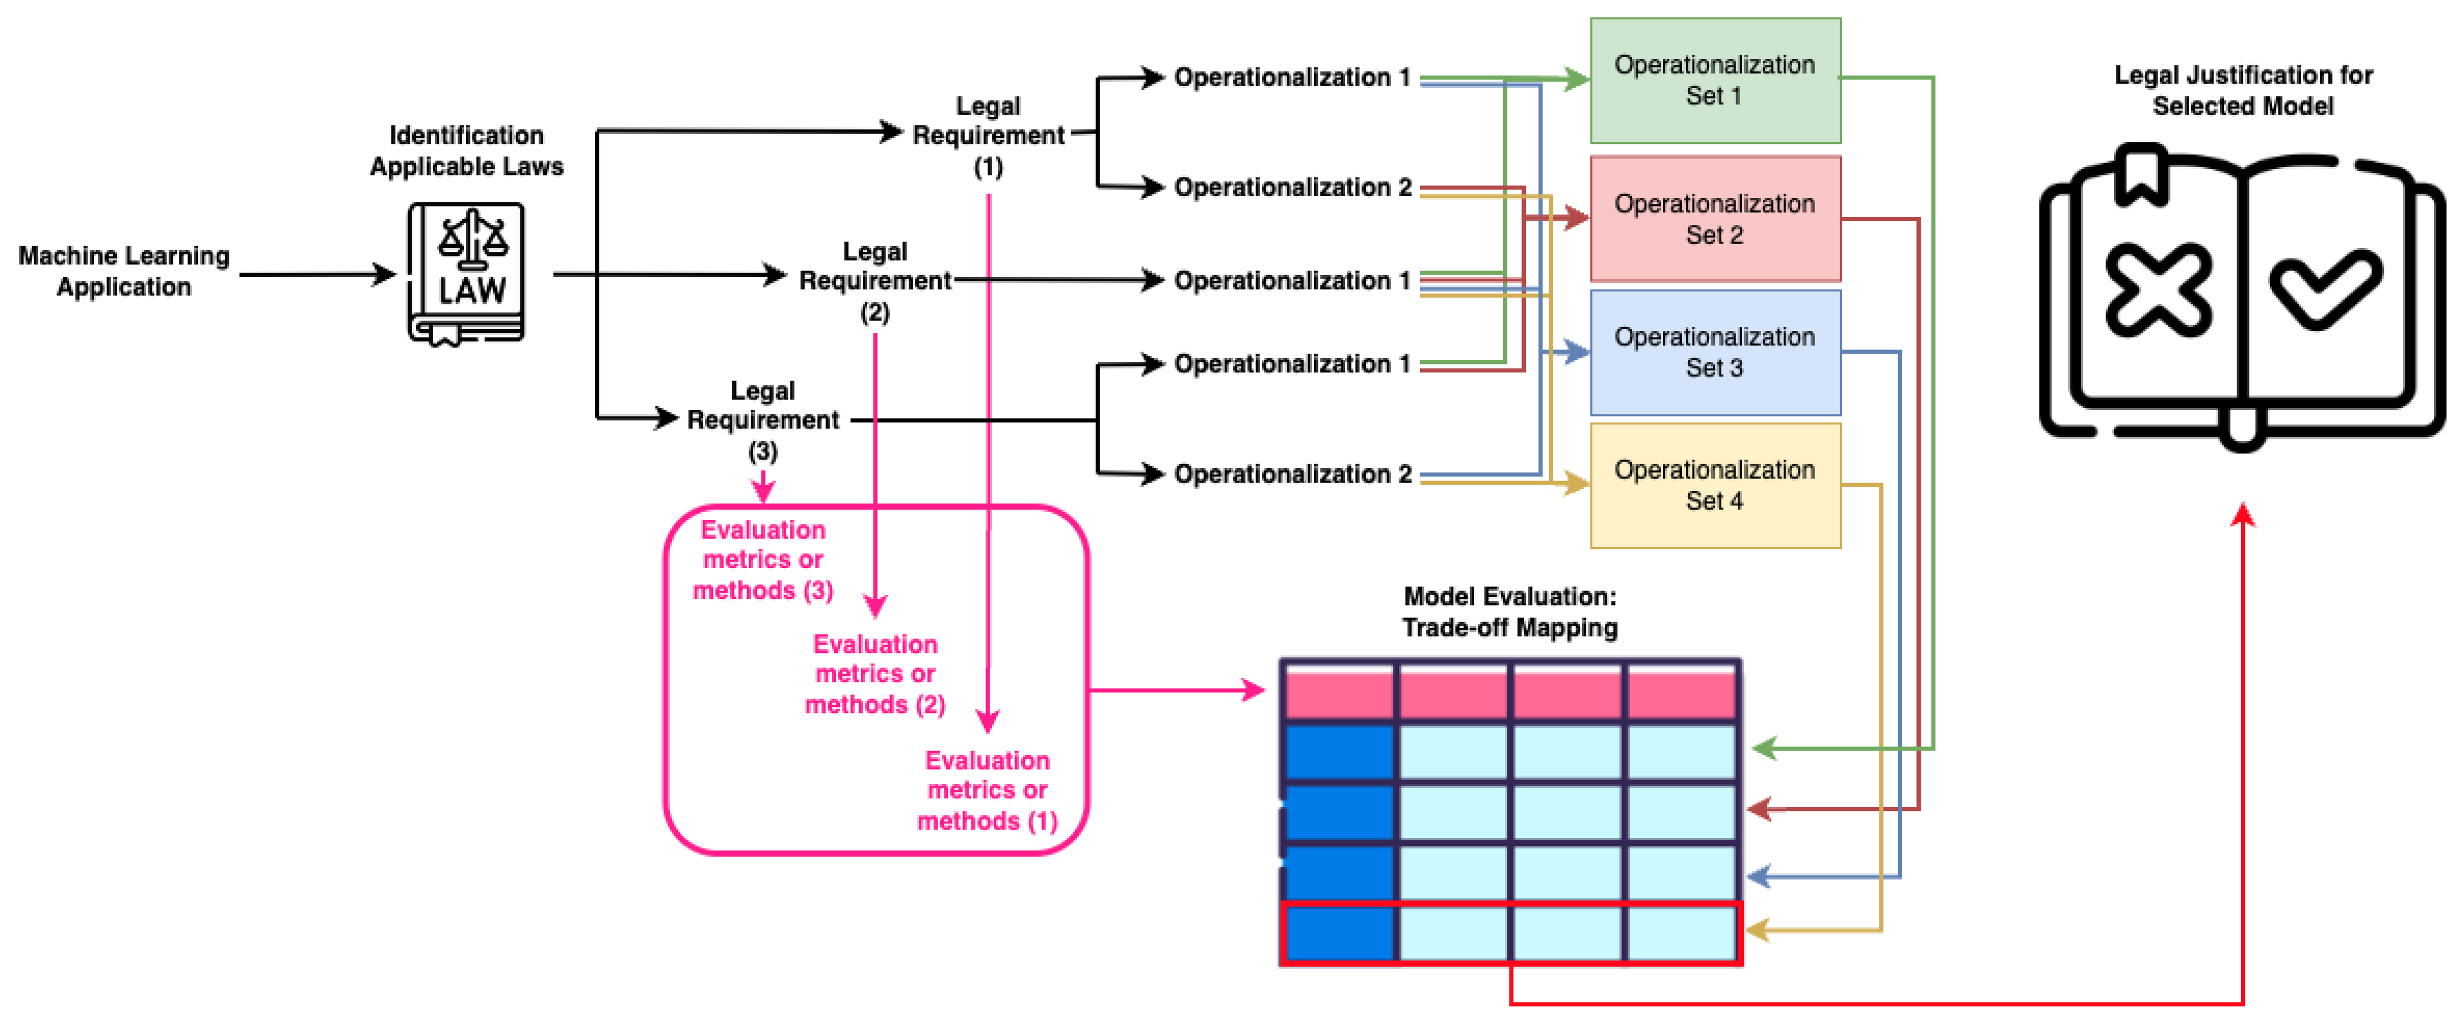

- **Law-Centric Design Framework** ✓
- **DM* under Legal Uncertainty** ✓
- **ML-Adapted Design Framework** ✓
- **DM* given Technical Trade-offs** ✓

Illustration – Fictitious Case-Study: Anti-Money Laundering (AML)

Stage 1:

The legal team identified the following legal requirements:

1. Ensuring non-discrimination over gender feature
2. Ensuring GDPR Data Minimization Compliance
3. Avoiding the model as 'personal data' under the GDPR qualification
4. Reasonable explainability of ML system to AML supervisory authorities
5. Maintain AML risk coverage

Stage 2:

The interdisciplinary team of lawyers and data scientists translates each legal requirement into technical operationalizations and selects an evaluation method or metric relevant for assessing legal compliance.

For instance, for the first requirement, 2 **operationalizations** are proposed: (1) The feature 'gender' is deleted from the dataset before training; (2) in addition to operationalization (1), a 'reject-option classification' technique is applied to get similar outputs over the different 'gender' values.
In terms of **evaluation**, the interdisciplinary team chooses the Conditional Demographic Disparity metric.

Stages 3 and 4:

Table 1: Operationalization sets for the case study

	Set 1	Set 2	Set 3	Set 4	Set 5	Set 6	Set 7	Set 8
Data minimization	(1)	(1)	(2)	(2)	(1)	(1)	(2)	(2)
Anti-discrimination	(1)	(1)	(1)	(1)	(2)	(2)	(2)	(2)
Model not personal data qualification	(1)	(1)	(2)	(2)	(1)	(1)	(2)	(2)
Legal anti-money laundering	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)

Table 2: Evaluation dimensions for the case study

Operationalization Set	Model Type	Predictive Performance			Legal Requirements					
		Accuracy	Precision	F1 Score	Data Minimization		Anti-discrimination Requirement	Model as Personal Data Qualification	AML Requirements	
					% of Available Data Used	k-anonymity Applied	Demographic Disparity over Gender	Likelihood of Re-identification	Explainability	Recall
Set 1	Random Forest	0.85	0.80	0.86	84%	No	0.10	Low	Moderate	0.94
Set 2	Logistic Regression	0.82	0.85	0.88	70%	No	0.12	Low	High	0.92
Set 3	Random Forest	0.83	0.79	0.85	70%	Yes	0.11	Very Low	Moderate	0.93
Set 4	Logistic Regression	0.83	0.76	0.82	68%	Yes	0.13	Very Low	High	0.90
Set 5	Random Forest	0.82	0.78	0.84	68%	No	0.10	Low	Moderate	0.92
Set 6	Logistic Regression	0.81	0.77	0.83	62%	No	0.06	Low	High	0.89
Set 7	Random Forest	0.84	0.79	0.84	72%	Yes	0.03	Very low	Moderate	0.89
Set 8	Logistic Regression	0.79	0.76	0.81	65%	Yes	0.07	Very Low	High	0.86

Stage 5:

The interdisciplinary team selects the model for deployment based on its performance and legal alignment, as determined by the trade-off analysis. The chosen model must meet all legal requirements while maintaining high predictive performance, even if it does not excel in any single metric.

In the case study, the Random Forest model from Set 3 is chosen. Despite not excelling in any individual aspect, its overall positive legal evaluations outweigh slightly lower scores in other areas.

Feedback or interesting research ideas? Contact me at Mathias.hanson@vub.be