

ECML PKDD 2025



Not Just a Trend: Institutionalizing XAI for Responsible and Compliant AI Systems

Francisco Herrera

Prof. Computer Science and Artificial Intelligence

Research Institute of Data Science and Computational Intelligence (DaSCI)

University of Granada, SPAIN

Spanish Royal Academy of Engineering

ADIA Lab Fellow



UNIVERSIDAD DE GRANADA

ADIA Lab



Not just a Trend. Why Explainable AI (XAI) Matters Today

Defining Explainable AI (XAI)

Ethics, Regulatory and Compliance Drivers

Building Explainable AI (Static vs Dynamic)

Best Practices: XAI for Human-AI Collaboration (HAIC)

XAI Towards Effective HAIC: LLMs Era, Context, Co-explainers, ...

Institutionalizing XAI

Conclusions

Not just a Trend. Why Explainable AI (XAI) Matters Today

Defining Explainable AI (XAI)

Ethics, Regulatory and Compliance Drivers

Building Explainable AI (Static vs Dynamic)

Best Practices: XAI for Human-AI Collaboration (HAIC)

XAI Towards Effective HAIC: LLMs Era, Context, Co-explainers, ...

Institutionalizing XAI

Conclusions

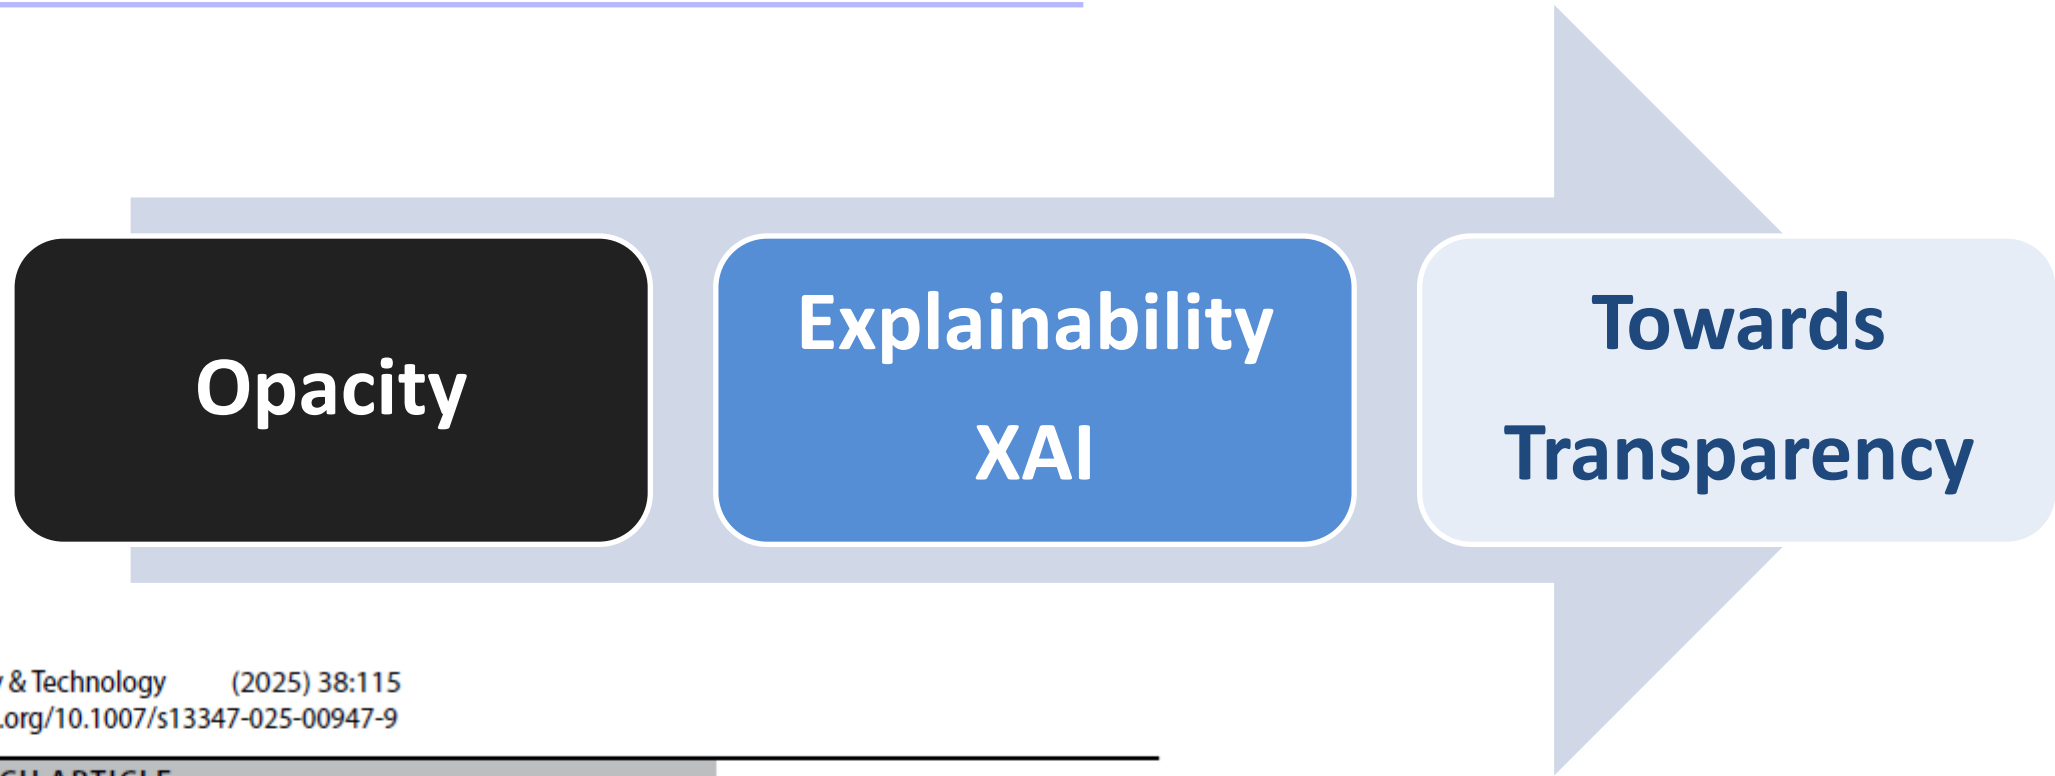
Not just a Trend. Why XAI Matters Today

Opacity

AI reliance

Trust in AI

Risks of opaque systems



Philosophy & Technology (2025) 38:115
<https://doi.org/10.1007/s13347-025-00947-9>

RESEARCH ARTICLE



Trust and Opacity in Artificial Intelligence: Mapping the Discourse

Martin Hähnel¹ · Rico Hauswald²

arXiv > cs > arXiv:2505.20304 Search... Help | Adv

Computer Science > Computers and Society

[Submitted on 18 May 2025]

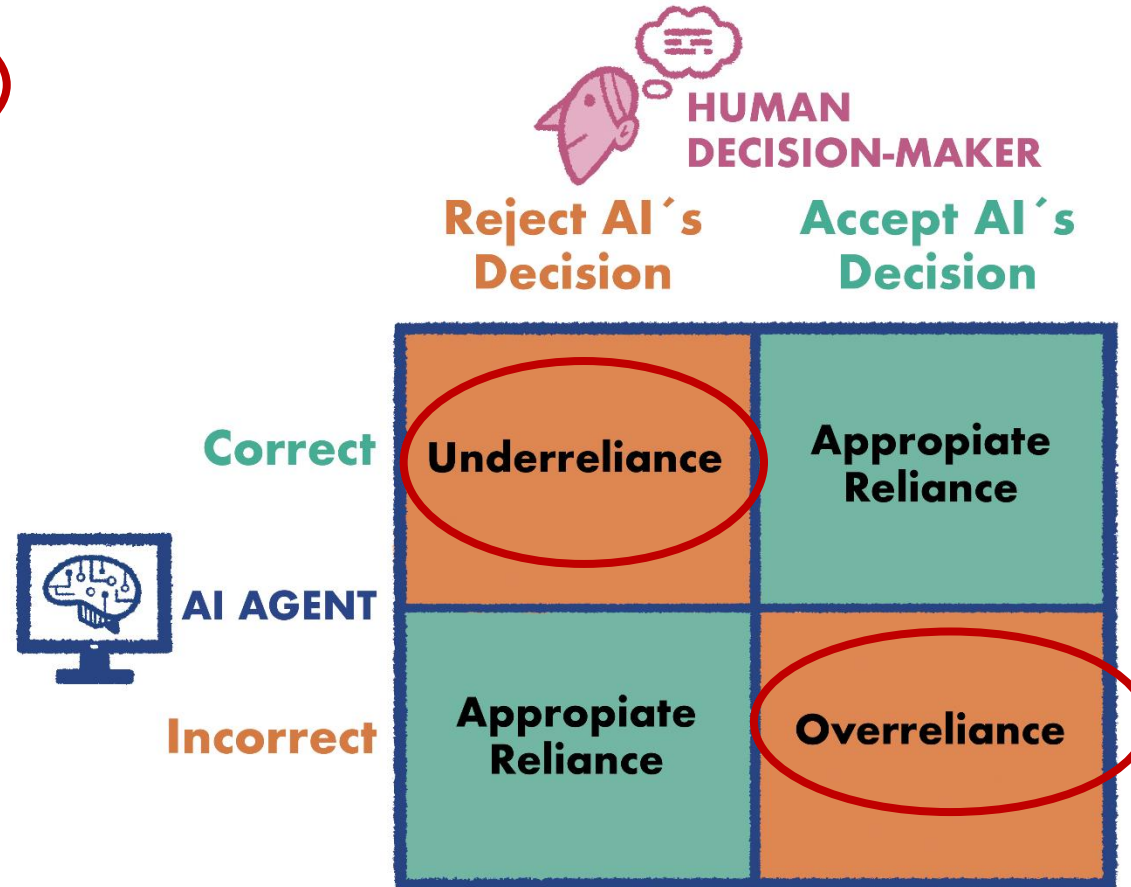
Opacity as a Feature, Not a Flaw: The LoBOX Governance Ethic for Role-Sensitive Explainability and Institutional Trust in AI

Francisco Herrera, Reyes Calderón

Growing reliance on AI in high-risk scenarios/decisions

Overreliance
(Excess of confidence)

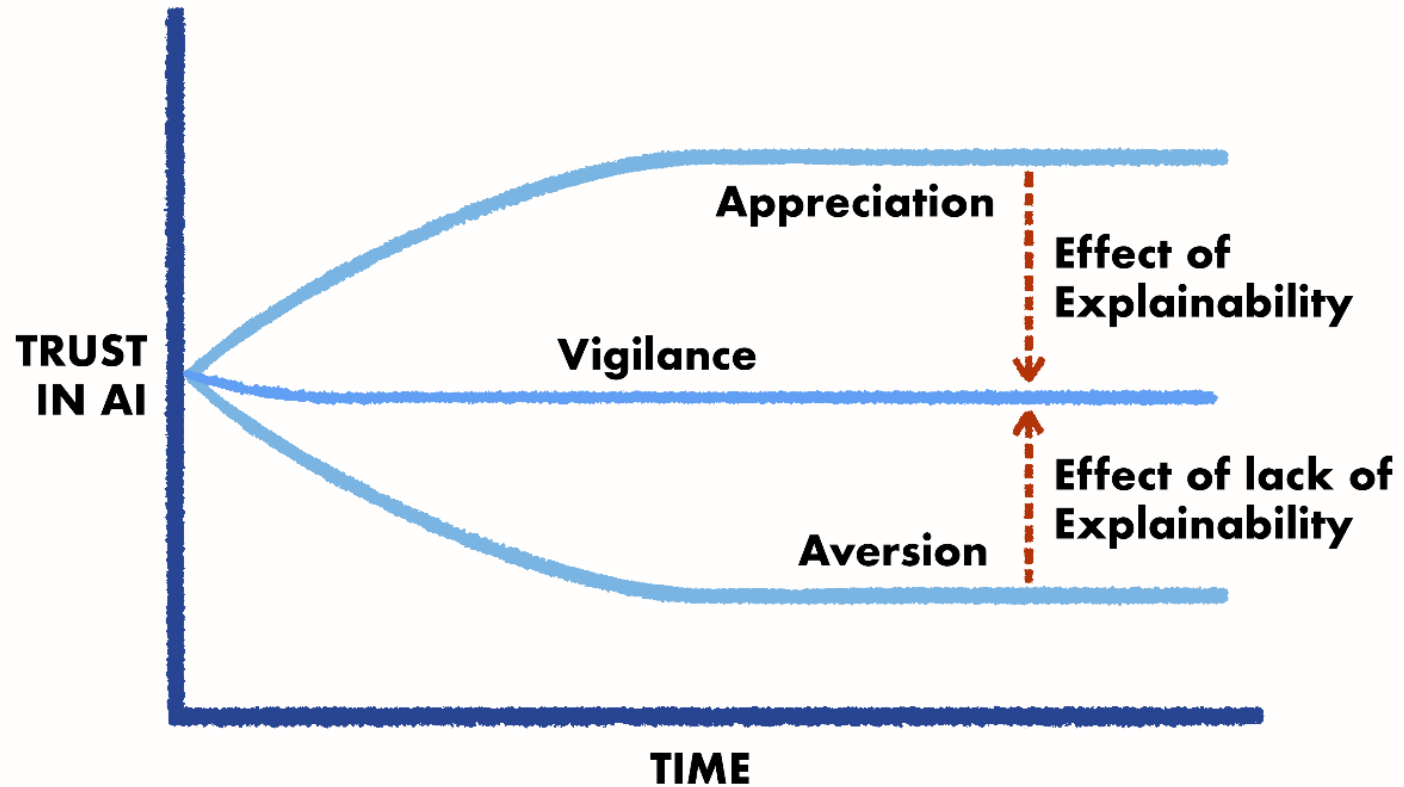
Overreliance on AI can be problematic, as humans might accept AI recommendations even when they are incorrect (wrong adherence), leading to poor decision quality



Research Question:
Under what conditions do people engage with explanations to verify AI predictions?
Reducing overreliance

Trust in AI

GOAL: Explainability can modulate Trust in AI



(inspired in Zerilli et al, 2022)

Not just a Trend. Why XAI Matters Today

Defining Explainable AI (XAI)

Ethics, Regulatory and Compliance Drivers

Building Explainable AI (Static vs Dynamic)

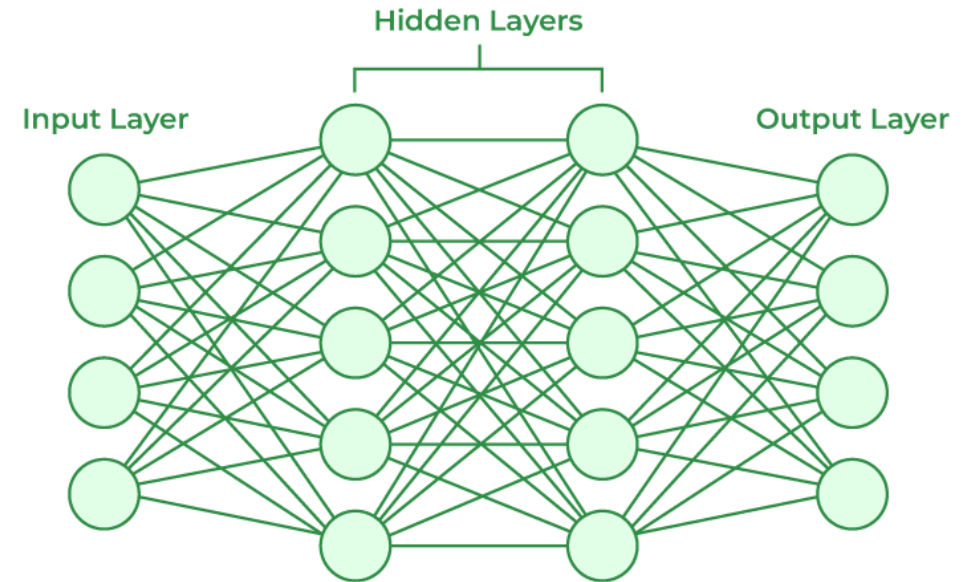
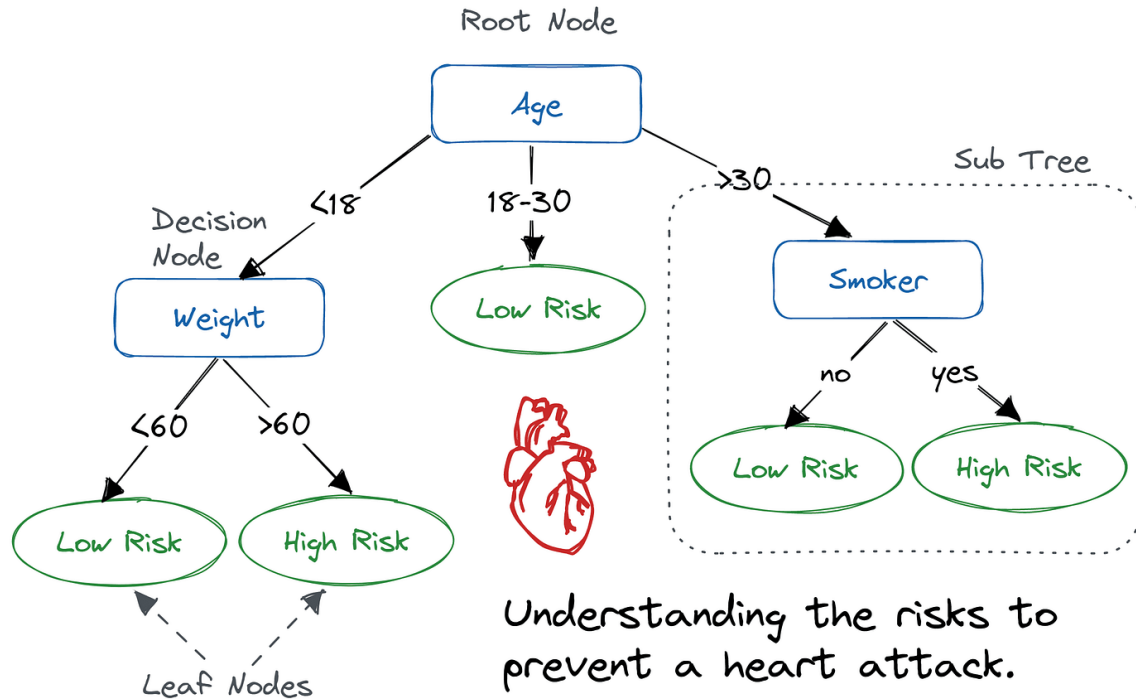
Best Practices: XAI for Human-AI Collaboration (HAIC)

XAI Towards Effective HAIC: LLMs Era, Context, Co-explainers, ...

Institutionalizing XAI

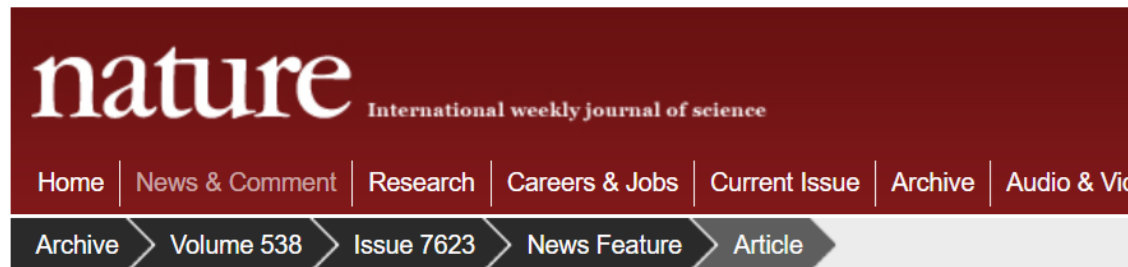
Conclusions

Between black box and white AI



From white box versus black-box AI

Black box AI

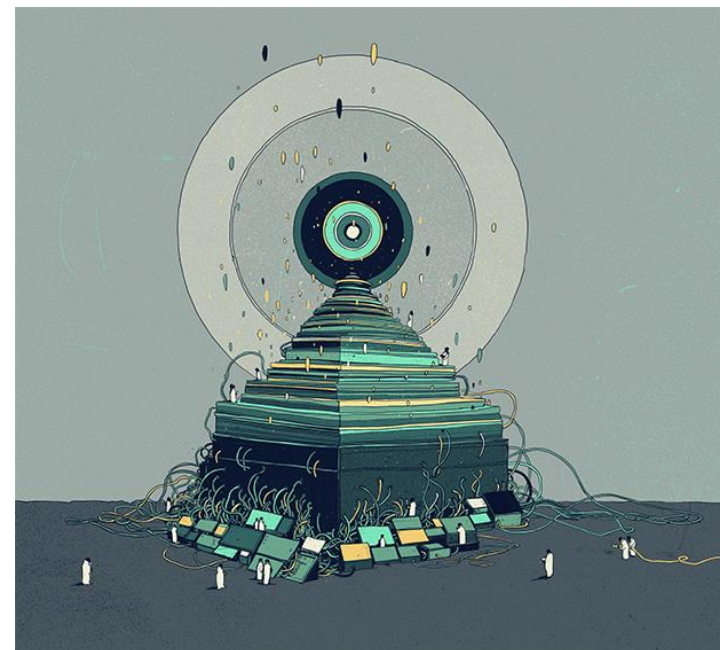


NATURE | NEWS FEATURE

Can we open the black box of AI?

Artificial intelligence is everywhere. But before scientists trust it, they first need to understand how machines learn.

Davide Castelvecchi



Do you trust the model?

RIBEIRO, Marco Tulio; SINGH, Sameer; GUESTRIN, Carlos. "Why should I trust you?" Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 2016. p. 1135-1144



Figure 11: Raw data and explanation of a bad model's prediction in the "Husky vs Wolf" task.

	Before	After
Trusted the bad model	10 out of 27	3 out of 27
Snow as a potential feature	12 out of 27	25 out of 27

Table 2: "Husky vs Wolf" experiment results.

What is eXplainable Artificial Intelligence (XAI)?

Explainable &
Interpretable

Definition: Explainability



Given an audience, an explainable AI is one that produces details or reasons to make its functioning clear or easy to understand.



Information Fusion
Volume 58, June 2020, Pages 82-115

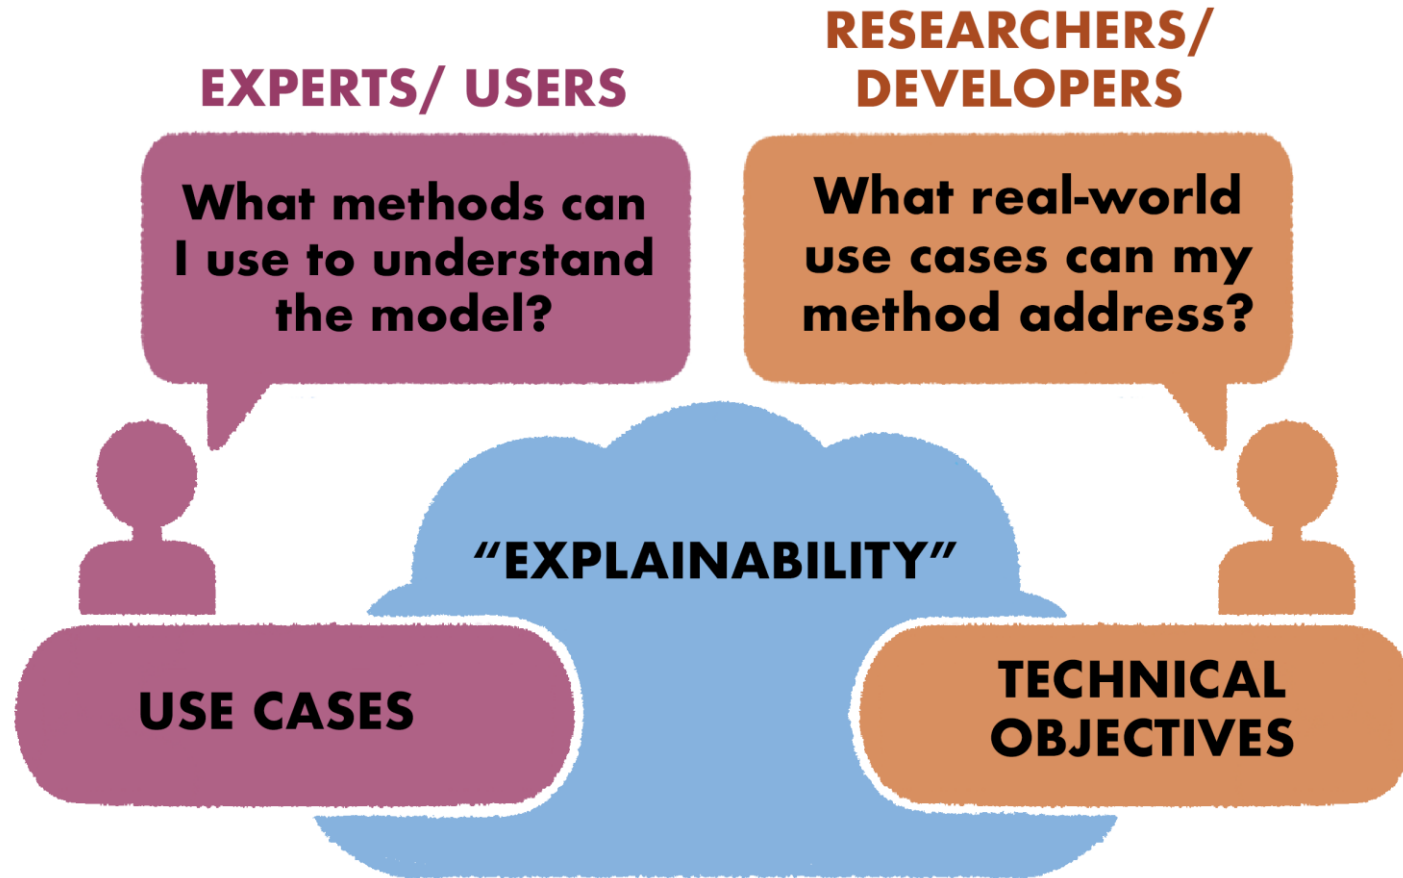


Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI

Alejandro Barredo Arrieta^a, Natalia Díaz-Rodríguez^b, Javier Del Ser^{a,c,d}  ,
Adrien Bennetot^{b,e,f}, Siham Tabik^g, Alberto Barbado^h, Salvador Garcia^g,
Sergio Gil-Lopez^a, Daniel Molina^g, Richard Benjamins^h, Raja Chatila^f,
Francisco Herrera^g

On the XAI definition essence: understanding and audience

Explainable &
Interpretable



Explainability gap, between "experts/users" and "researchers/developers"

Challenge: XAI for audience

Blue XAI
Human-valued

Red XAI
Model-valued

USERS/ EXPERTS

Who? Users/experts engaged with the decision, Human-AI decision makers (e.g. medical doctors,...)

Why? Understand the advice, verify fair decisions, gain scientific knowledge



SOCIETY

Who? Recipients of the decision, society in general (e.g. patients,...)

Why? Trust in the model itself, rely on the decision



REGULATOR

Who? Regulatory entities/agencies

Why? Certify model compliance with the legislation in force, audits...



TARGET AUDIENCE IN XAI



OWNER

Who? Owner, managers and executive board members

Why? Assess regulatory compliance, understand corporate AI applications...



DEVELOPER/ RESEARCHERS

Who? Researchers and developers

Why? Ensure/improve product efficiency, research algorithms and new functionalities...



DESIGNER

Who? Designer

Why? Human-agent interaction



Not just a Trend. Why XAI Matters Today

Defining Explainable AI (XAI)

Ethics, Regulatory and Compliance Drivers

Building Explainable AI (Static vs Dynamic)

Best Practices: XAI for Human-AI Collaboration (HAIC)

XAI Towards Effective HAIC: LLMs Era, Context, Co-explainers, ...

Institutionalizing XAI

Conclusions

Trust AI



Cambridge
Dictionary

Trust: To believe that someone is **good**
and **honest** and **will not harm you**,
or that something is **safe and reliable**

(trustworthy AI)

Trust in AI - Trustworthy AI

Trustworthy Artificial Intelligence

Robustness

Lawfulness

Ethics

The principle of explicability

Explicability is crucial for building and maintaining users' trust in AI systems.

The Ethical principles:

- (i) Respect for human autonomy
- (ii) Prevention of harm
- (iii) Fairness
- (iv) Explicability**



AI regulations and guidelines

Organization for Economic Co-operation and Development (OECD)



The Ethical principles of:

- (i) Respect for human autonomy
- (ii) Prevention of harm
- (iii) Fairness
- (iv) Explicability

OECD principles on AI

1. Inclusive growth, sustainable development and well-being
2. Human centric values and fairness
3. Transparency and explainability
4. Robustness, security and safety
5. Accountability



REQUIREMENT #4 Transparency

Traceability
Explainability
Communication

Safe

Secure &
Resilient

Fair-With
Harmful Bias
Managed

Privacy-
Enhanced

Explainable &
Interpretable

Accountable
&
Transparent

Valid & Reliable

Trust AI

Humanities & Social Sciences
Communications



REVIEW ARTICLE

<https://doi.org/10.1057/s41599-024-04044-8>

OPEN

Trust in AI: progress, challenges, and future directions

Saleh Afroogh¹, Ali Akbari², Emmie Malone³, Mohammadali Kargar⁴ & Hananeh Alambeigi⁵

Trust in AI is not just a non-technical ethical consideration (Ryan, 2020). Instead, it also includes various domains, including

- a) AI performance,**
- b) transparency and explainability,**
- c) compliance with legal and technical regulations.**



Not just a Trend. Why XAI Matters Today

Defining Explainable AI (XAI)

Ethics, Regulatory and Compliance Drivers

Building Explainable AI (Static vs Dynamic)

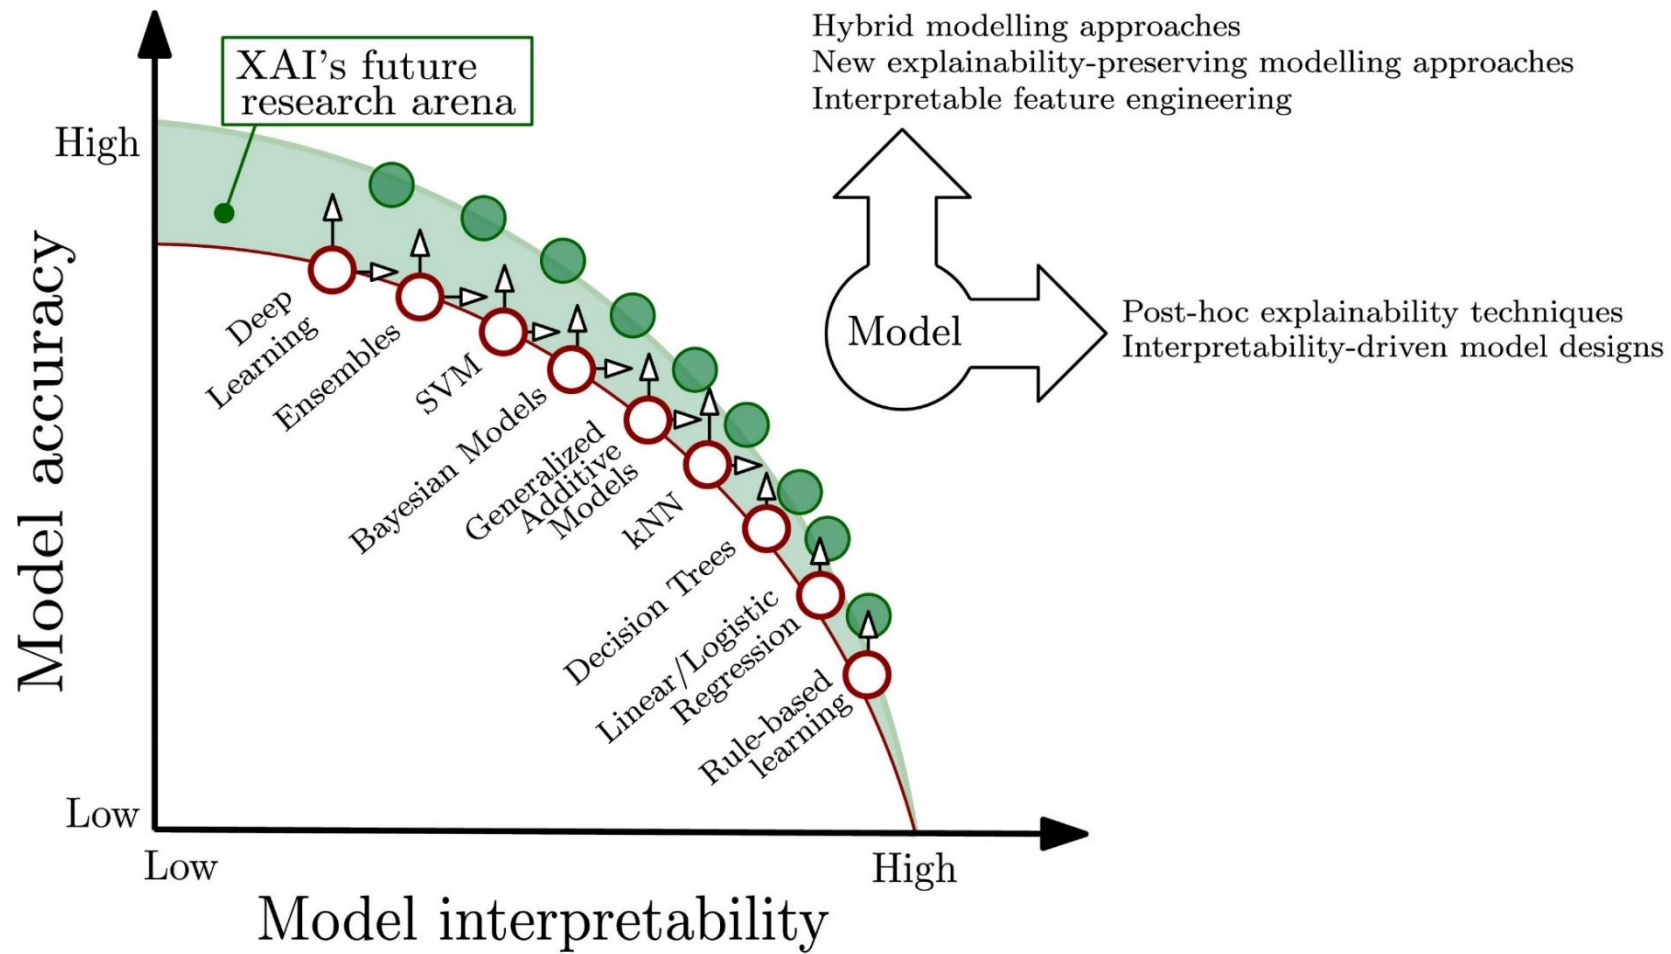
Best Practices: XAI for Human-AI Collaboration (HAIC)

XAI Towards Effective HAIC: LLMs Era, Context, Co-explainers, ...

Institutionalizing XAI

Conclusions

Learning Models in the heart of XAI



High Levels of interpretability/XAI of different Machine Learning models:

- (a) Linear regression;**
- (b) Decision trees;**
- (c) K-Nearest Neighbors;**
- (d) Rule-based Learners;**
- (e) Generalized Additive Models;**
- (f) Bayesian Models.**

Between less interpretability and greater accuracy

XAI: What we know











ACM Computing Surveys > Just Accepted > A Practical tutorial on Explainable AI Techniques

TUTORIAL |  FREE ACCESS

Just Accepted



A Practical tutorial on Explainable AI Techniques

Authors:  [Adrien Bennetot](#),  [Ivan Donadello](#),  [Ayoub El Qadi El Haouari](#),  [Mauro Dragoni](#),  [Thomas Frossard](#),  [Benedikt Wagner](#),  [Anna Sarranti](#),  [Silvia Tulli](#),  [Maria Trocan](#),
and  [Natalia Díaz-Rodríguez](#) (Less) | [Authors Info & Claims](#)

A practical analysis of XAI tools for different kind of data and scenarios

Software tools for XAI research

InterpretML	DALEX	Shapash (Sharp, LIME)
Alibi	H2O	OmniXAI
AIX360	ELI5	ML (IML)
Skater	iNNvestigate	Captum
tf-explain	modelStudio	Interpretable

SHAP: SHapley Additive exPlanations

DiCE: Diverse Counterfactual Explanations

Transformers Interpret (TI) (for language models)

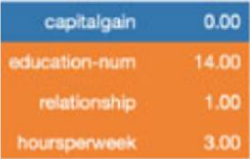

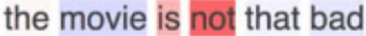
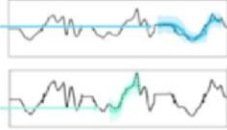
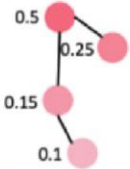

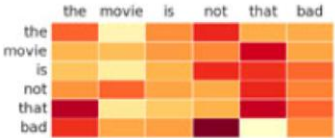
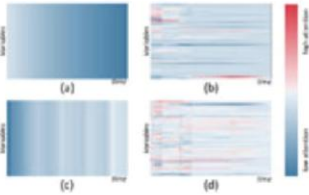
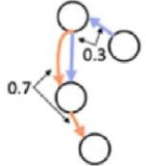
Grad-CAM (image classification)

Layer-wise Relevance (explain image)

Logic Tensor (Neural-Symbolic AI)

TS4NLE (for natural language explanations)

Explanations: The heart of XAI

TABULAR	IMAGE	TEXT	TIME SERIES	GRAPHS								
<p>Feature Importance (FI)</p>  <table border="1"> <tr><td>capitalgain</td><td>0.00</td></tr> <tr><td>education-num</td><td>14.00</td></tr> <tr><td>relationship</td><td>1.00</td></tr> <tr><td>hoursperweek</td><td>3.00</td></tr> </table>	capitalgain	0.00	education-num	14.00	relationship	1.00	hoursperweek	3.00	<p>Saliency Maps (SM)</p> 	<p>Sentence Highlighting (SH)</p> <p>the movie is not that bad</p> 	<p>Series Highlighting</p> 	<p>Node Highlighting</p> 
capitalgain	0.00											
education-num	14.00											
relationship	1.00											
hoursperweek	3.00											
<p>Rule- Based (RB)</p> <p>$r = \text{Education} \leq \text{College} \rightarrow \leq 50k$</p>	<p>Concept Attribution (CA)</p> 	<p>Attention Based (AB)</p> 	<p>Attention Based</p> 	<p>Edge Highlighting</p> 								

Explanations categorization

XAI: What we know. Post-Hoc XAI, static procedures

Explainable & Interpretable



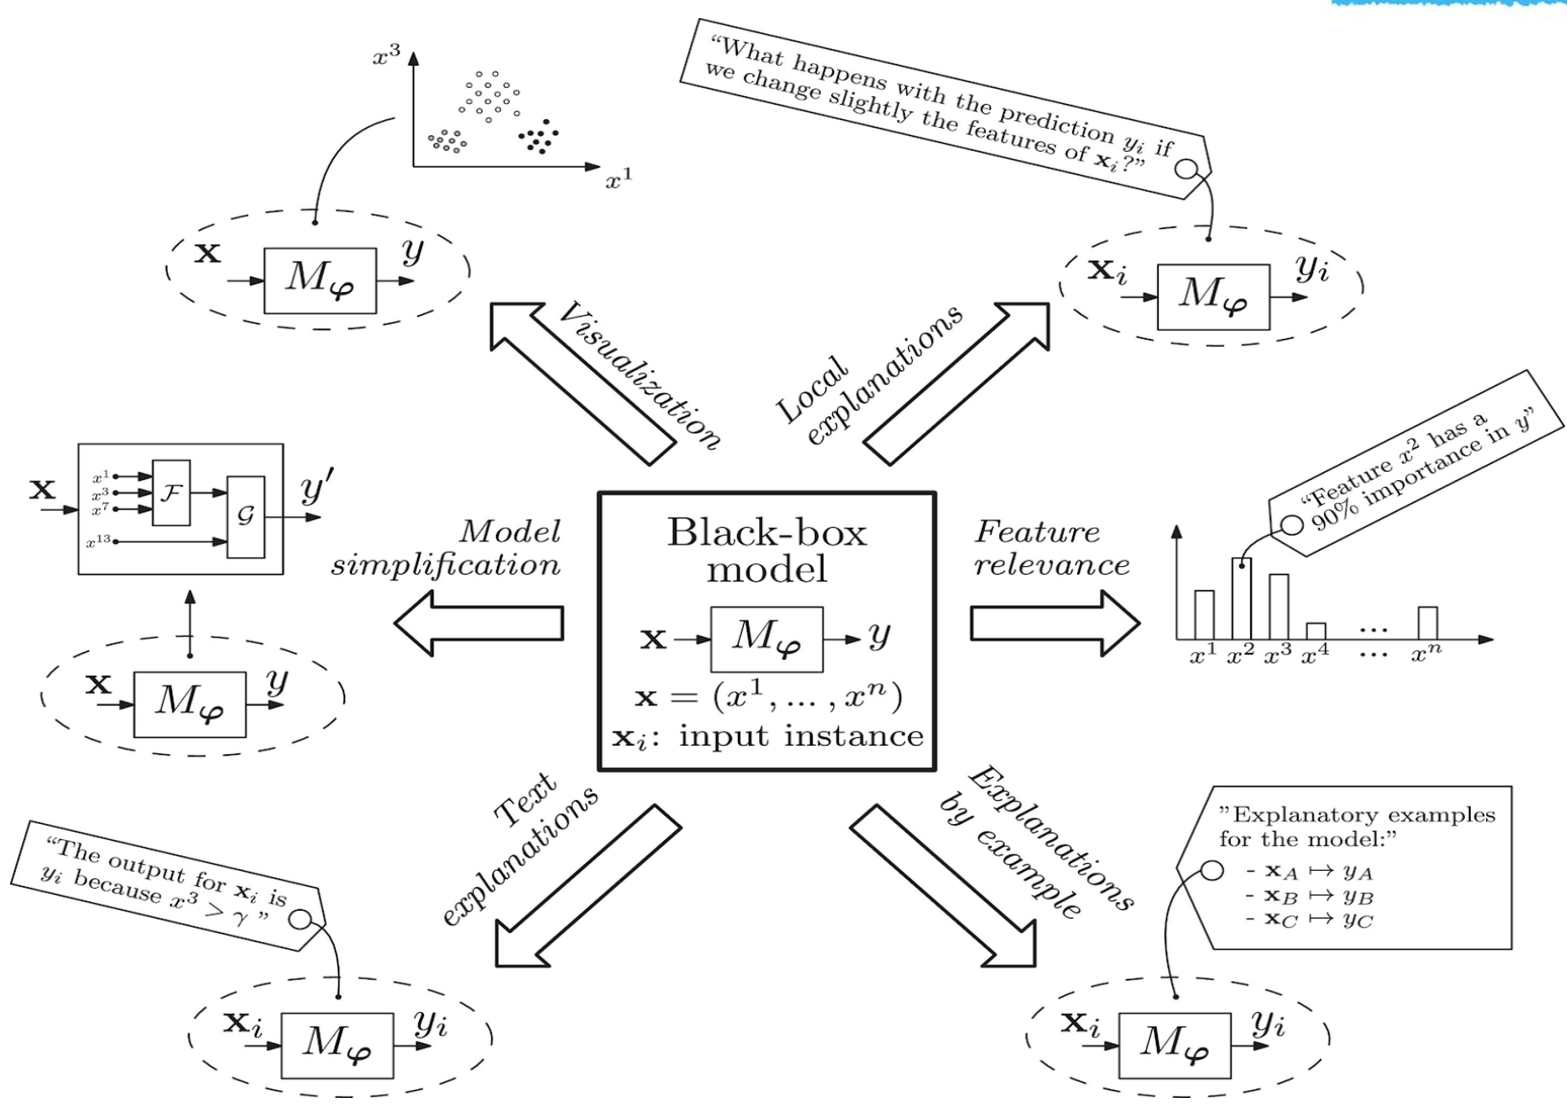
Information Fusion
Volume 58, June 2020, Pages 82-115



Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI

Alejandro Barredo Arrieta^a, Natalia Díaz-Rodríguez^b, Javier Del Ser^{a,c,d}, Adrien Bennetot^{b,e,f}, Siham Tabik^g, Alberto Barbado^h, Salvador Garcia^g, Sergio Gil-Lopez^a, Daniel Molina^g, Richard Benjamins^h, Raja Chatila^f, Francisco Herrera^g

Post-Hoc Explainability techniques for black-box models



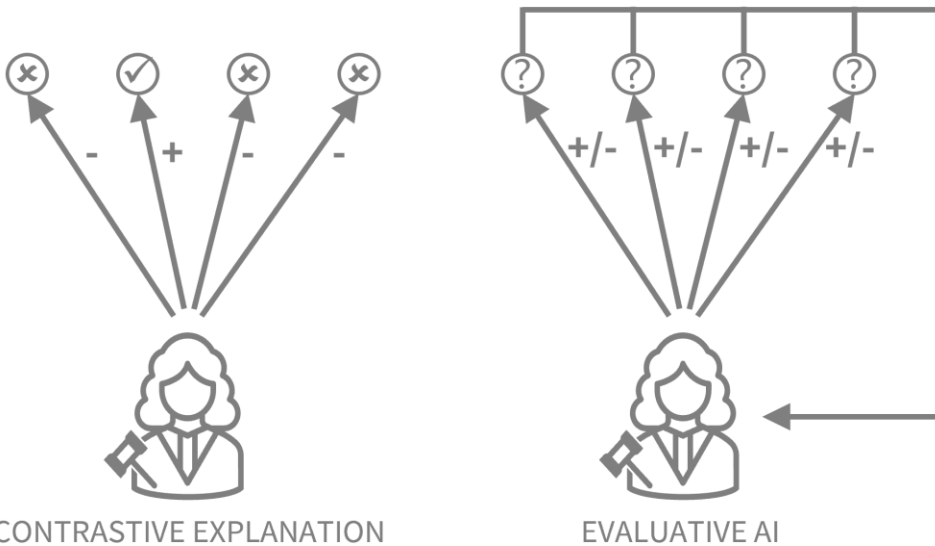
XAI: What do we need? Toward dynamic XAI procedures

Dynamic New scenario

Evaluative AI
Incremental XAI
Conversational XAI
Interactivity x Explainability
Iterative XAI frameworks



Static explainability



Dynamic explainability



Not just a Trend. Why XAI Matters Today

Defining Explainable AI (XAI)

Ethics, Regulatory and Compliance Drivers

Building Explainable AI (Static vs Dynamic)

Best Practices: XAI for Human-AI Collaboration (HAIC)

XAI Towards Effective HAIC: LLMs Era, Context, Co-explainers, ...

Institutionalizing XAI

Conclusions

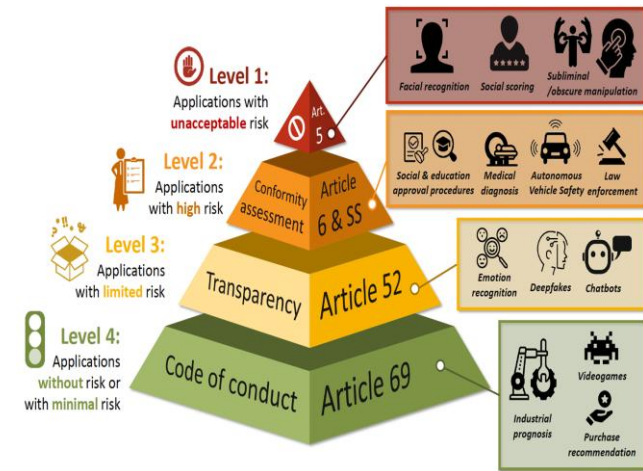
Human-AI Decision Making



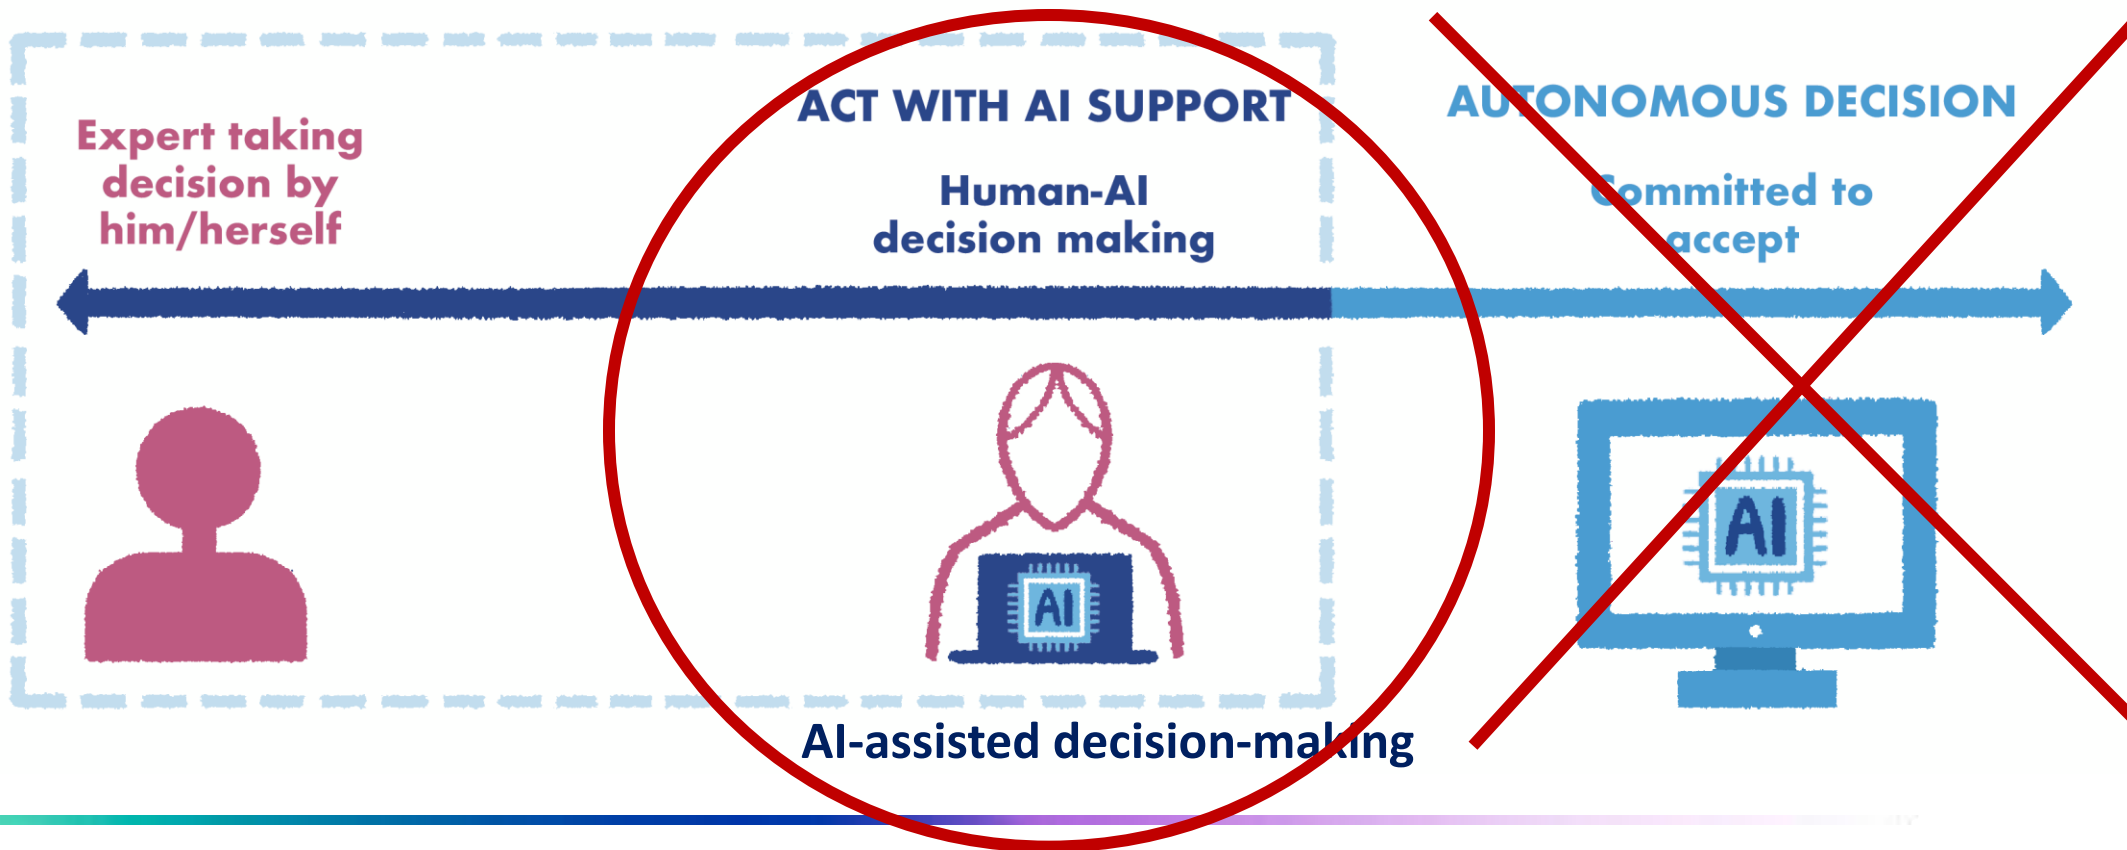
Level 2:
Applications with **high risk**



**Task. High risk scenarios.
Health, finance, critical infrastructures, ...**



ENGAGE WITH THE TASK

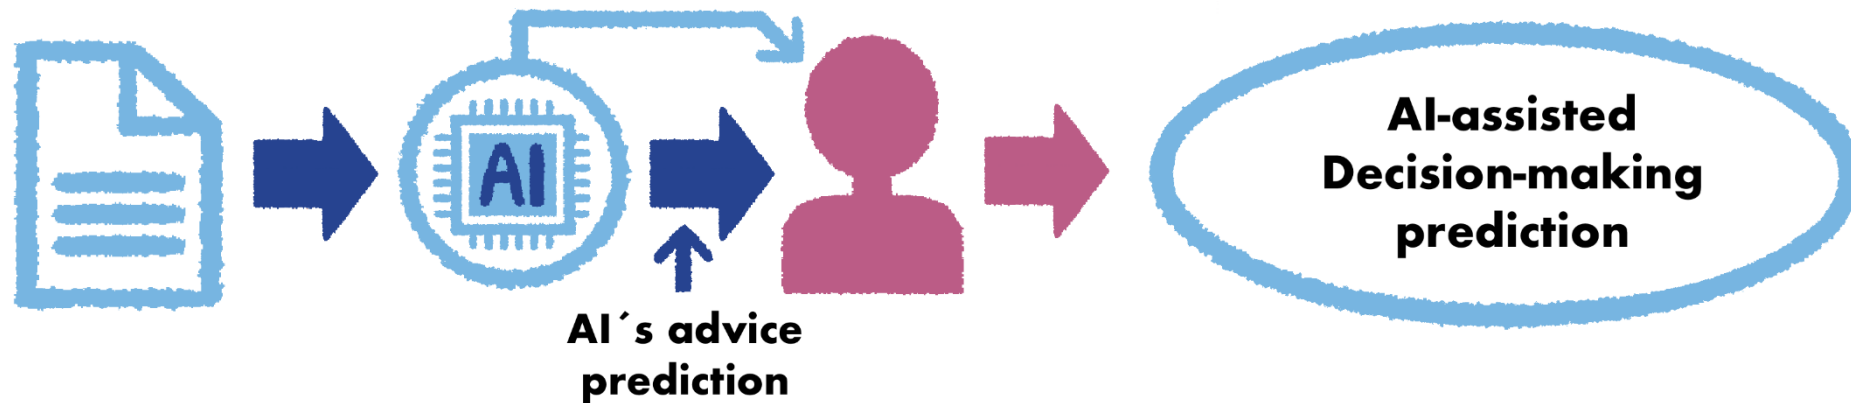


Human-AI Decision Making



AI-assisted decision-making scenarios, the human makes a decision about a set of instances with the assistance of an AI's advice (predictions, ...)

HUMAN-AI DECISION-MAKING PROCESS

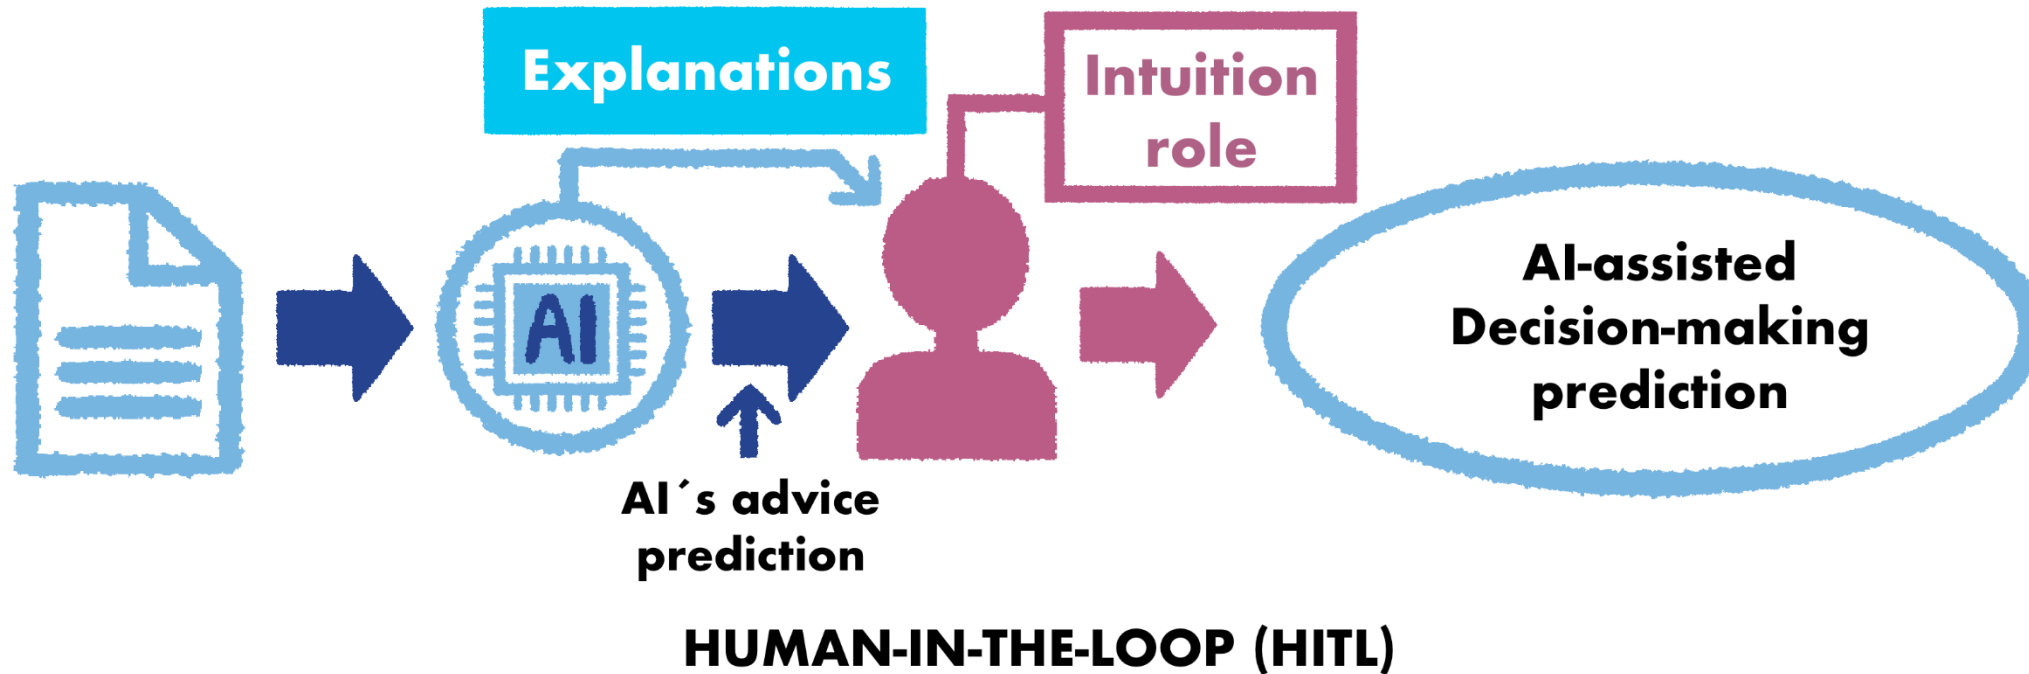


HUMAN-IN-THE-LOOP (HITL)

XAI for Human-AI Collaboration



AI RELIANCE - HUMAN-AI DECISION-MAKING PROCESS



Not just a Trend. Why XAI Matters Today

Defining Explainable AI (XAI)

Ethics, Regulatory and Compliance Drivers

Building Explainable AI (Static vs Dynamic)

Best Practices: XAI for Human-AI Collaboration (HAIC)

XAI towards Effective HAIC: LLMs Era, Context, Co-explainers, ...

Institutionalizing XAI

Conclusions

Towards XAI practical use for AI systems deployment

Cointelligence

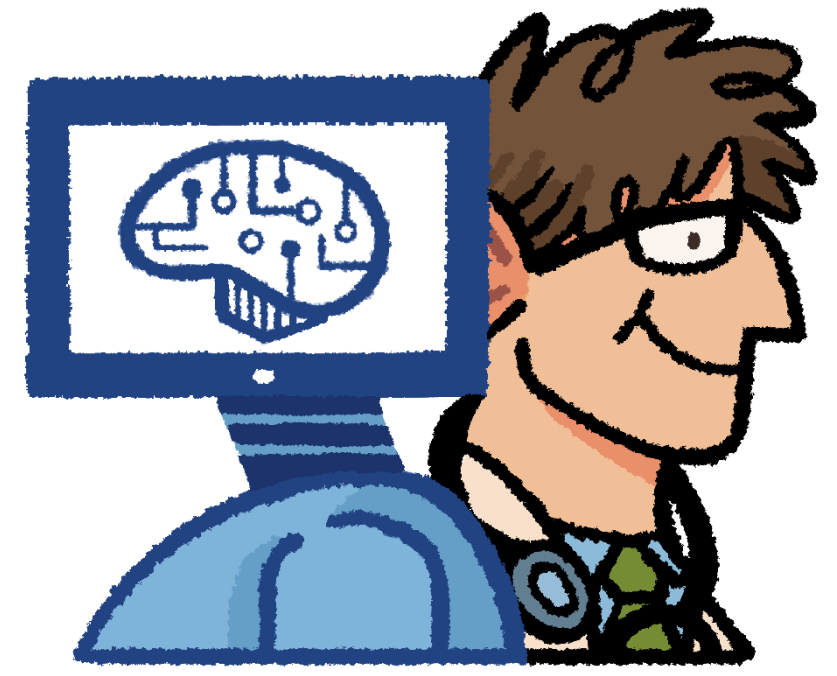
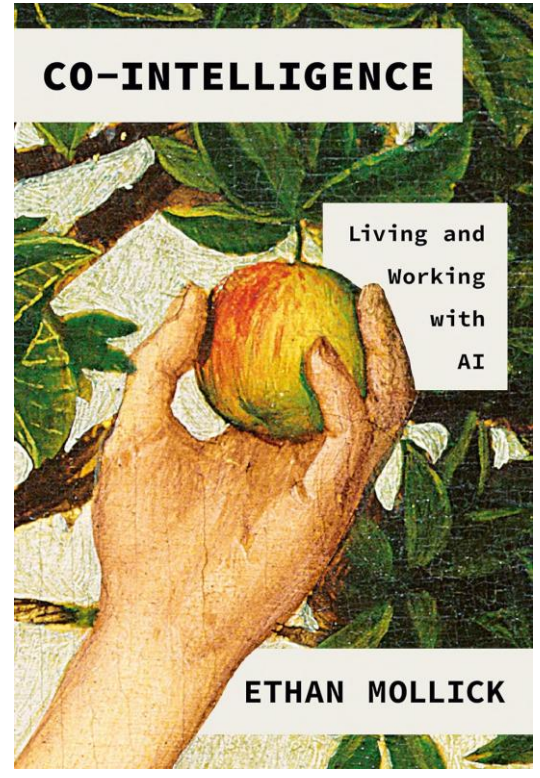
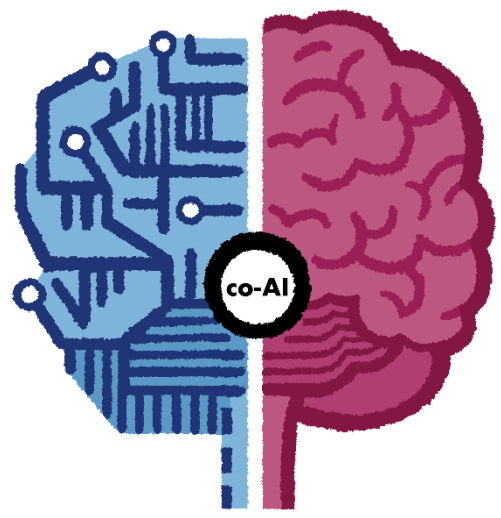
XAI in the LLMs Era

Explainability in context

Co-explainers (Iterative and Interactive XAI framework)

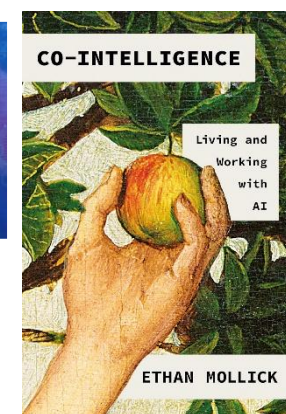


Doctor Centaur
Delegation and approval



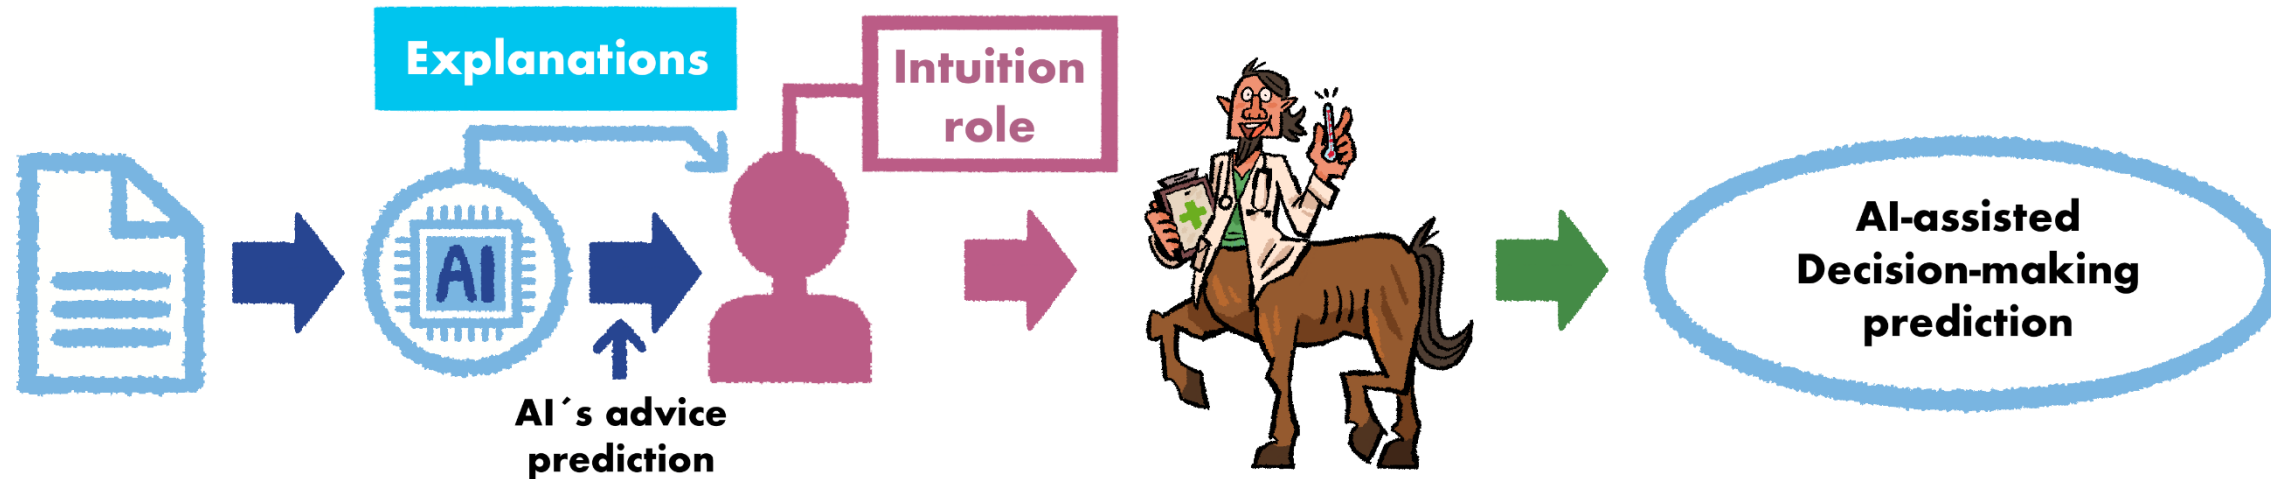
Doctor Cyborg
co-creativity
co-analysis and human decision making

XAI helping. Toward Effective HAIC



The goal: Under what conditions explanations can improve human understanding and in what way.

HUMAN-AI DECISION-MAKING PROCESS



HUMAN-IN-THE-LOOP (HITL)

Human-AI decision making. XAI can help Doctor Centaur

XAI in the LLMs Era

MAKING SENSE OF THE UNSENSIBLE: REFLECTION, SURVEY, AND CHALLENGES FOR XAI IN LARGE LANGUAGE MODELS TOWARD HUMAN-CENTERED AI

Francisco Herrera^{1,2}

¹Department of Computer Science and Artificial Intelligence, Andalusian Institute of Data Science and Computational Intelligence (DaSCI), University of Granada, Spain.
Emails: herrera@decsai.ugr.es

² ADIA Lab, Abu Dhabi, United Arab Emirates

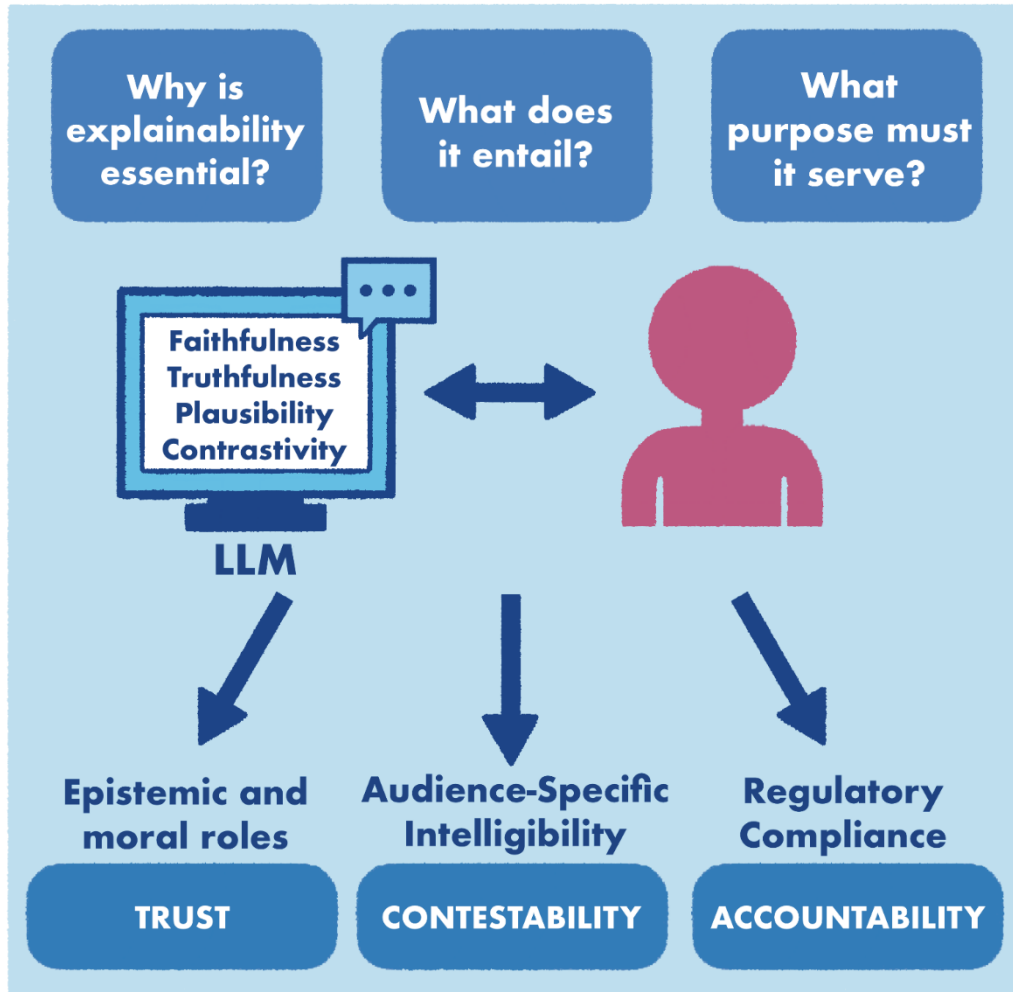


18 May 2025

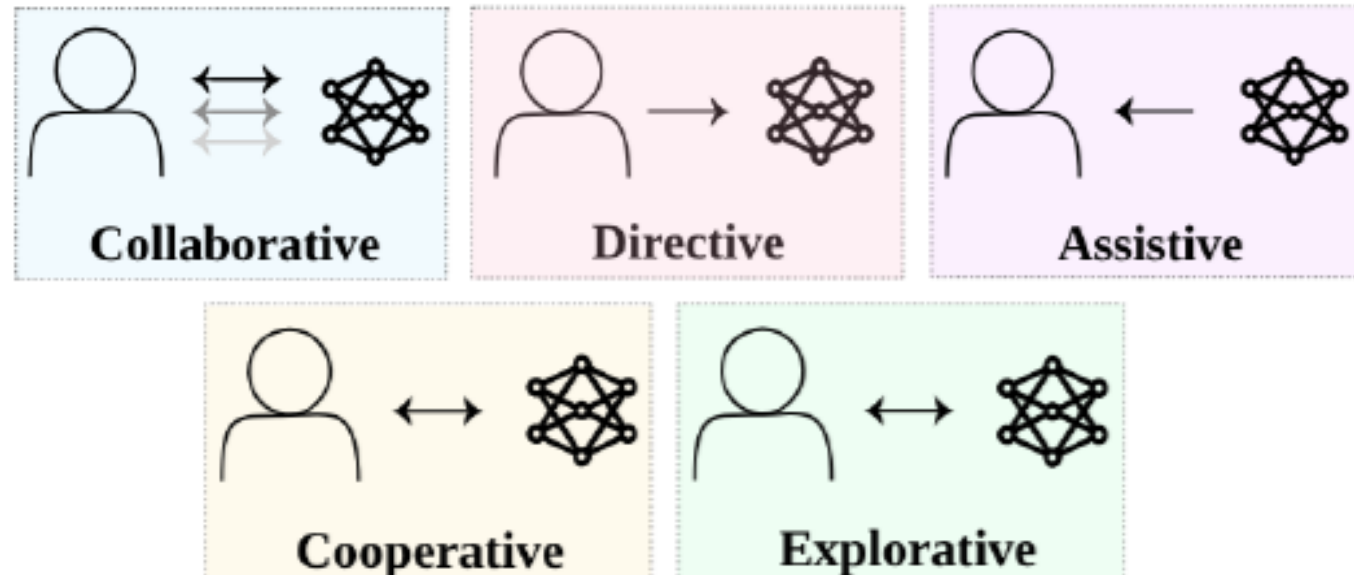
These systems already communicate in natural language, yet the mere presence of an answer does not constitute an explanation of its reasoning.

The challenge is not just to interpret what was said, but also to understand why it was said.

XAI in the LLMs Era. Toward Effective Human-AI Collaboration (HAIC)



XAI: MAKING SENSE OF THE UNSENSIBLE



The screenshot shows the arXiv preprint page for the paper "Triadic Fusion of Cognitive, Functional, and Causal Dimensions for Explainable LLMs: The TAXAL Framework" by David Herrera-Poyatos, Carlos Peláez-González, Cristina Zuheros, Virilo Tejedor, Rosana Montes, and Francisco Herrera. The page includes the arXiv logo, the paper title, authors, and submission date (5 Sep 2025). The paper is categorized under Computer Science > Computation and Language.

XAI in the LLMs Era: Explainability in Context

Explainability in Context: A Multilevel Framework Aligning AI Explanations with Stakeholder with LLMs

Marilyn Bello, Rafael Bello, Maria-Matilde García, Ann Nowé, Iván Sevillano-García, Francisco Herrera

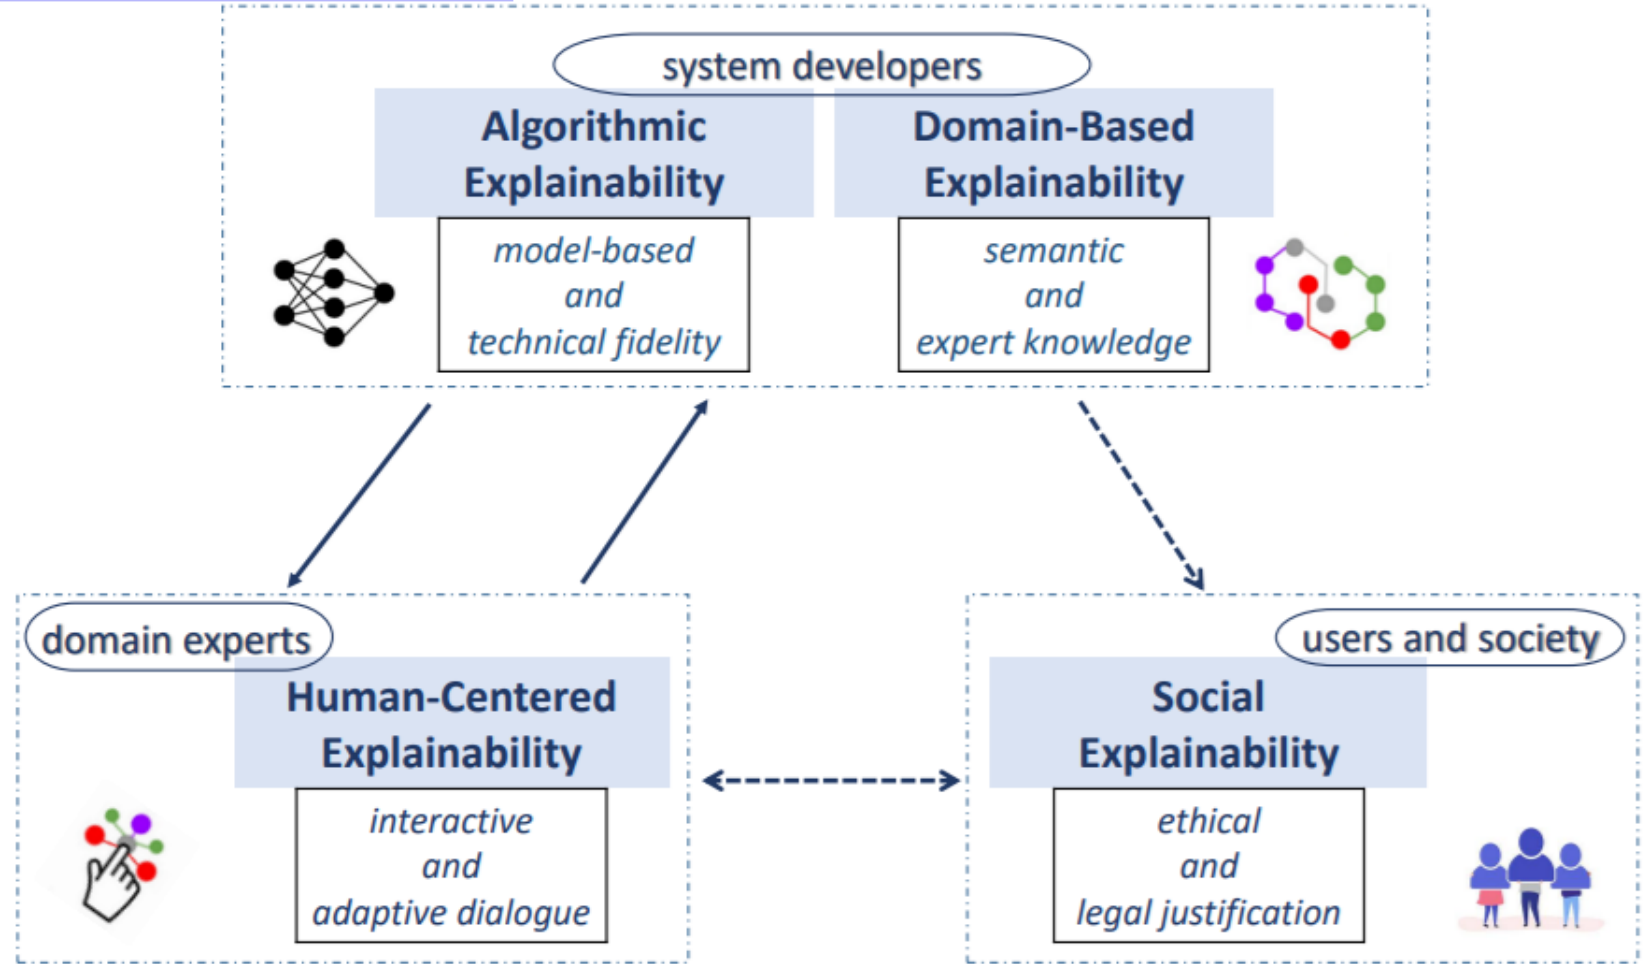


Figure 1: Three-level framework for generating explanations.

XAI in the LLMs Era: Explainability in Context

arXiv > cs > arXiv:2506.05887

Search... Help | Adv

Computer Science > Artificial Intelligence

[Submitted on 6 Jun 2025]

Explainability in Context: A Multilevel Framework Aligning AI Explanations with Stakeholder with LLMs

Marilyn Bello, Rafael Bello, Maria-Matilde García, Ann Nowé, Iván Sevillano-García, Francisco Herrera

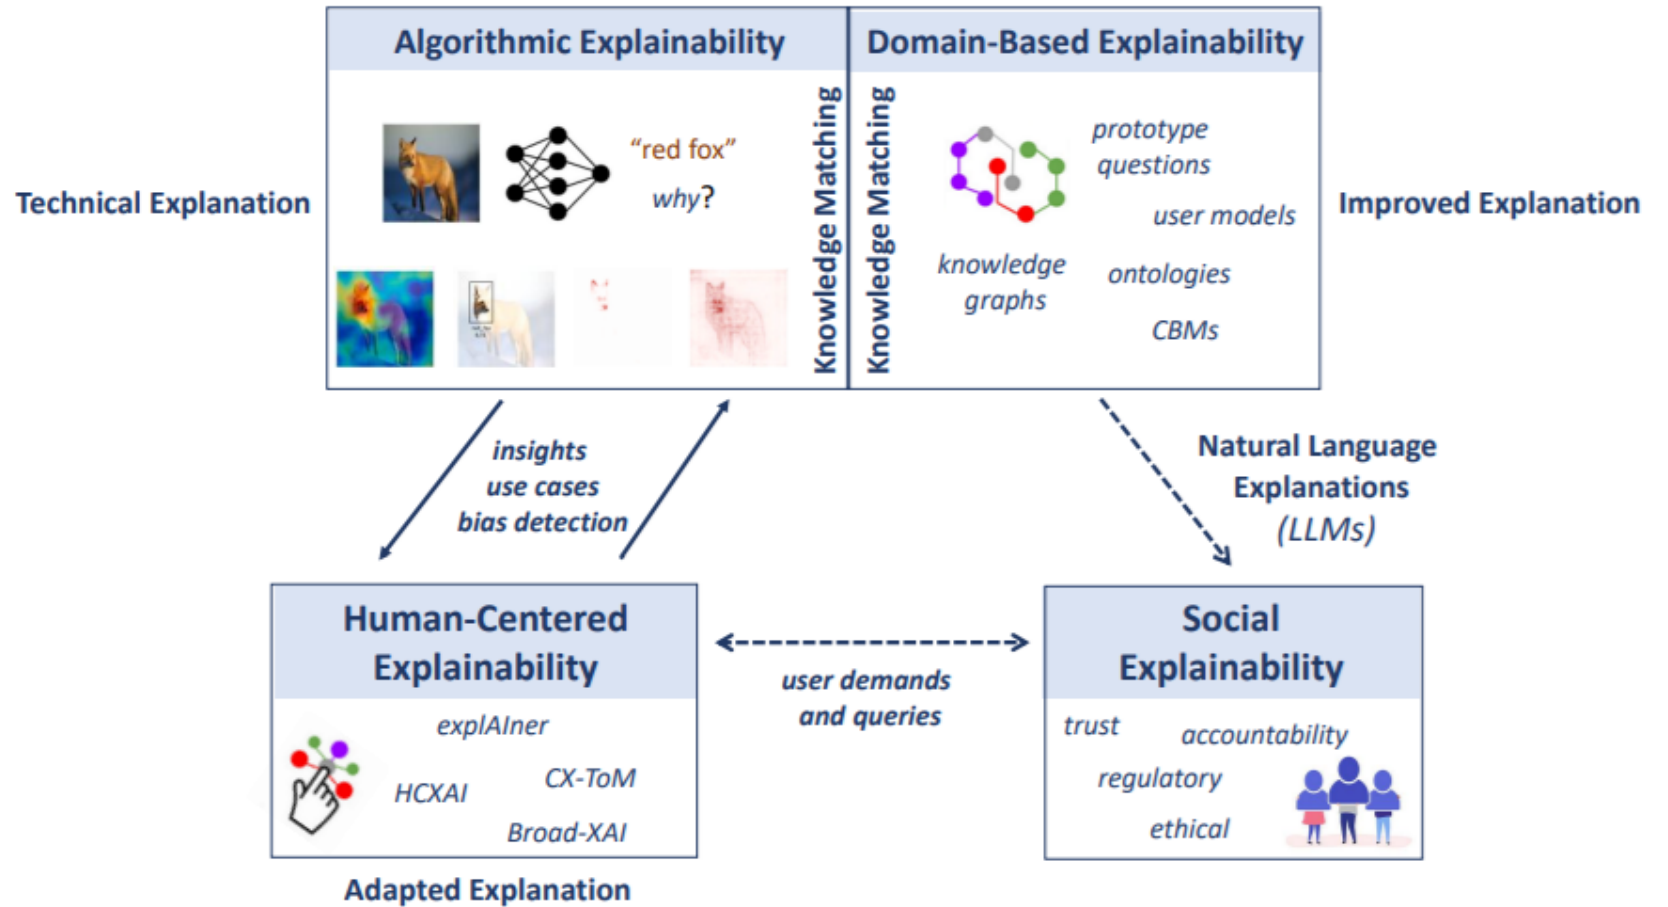
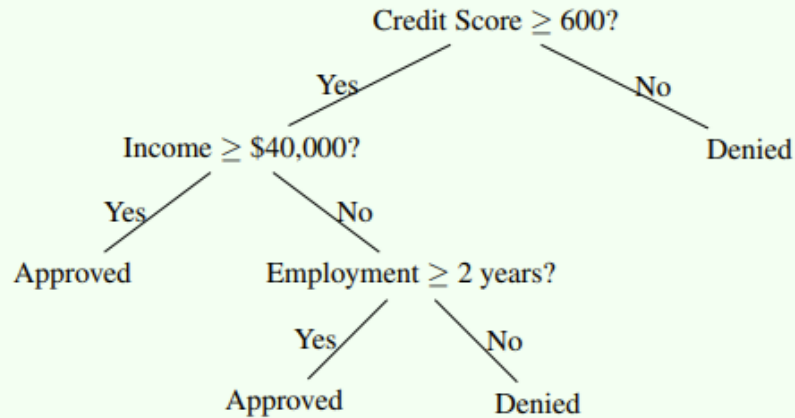


Figure 5: Social Explainability level.

XAI in the LLMs Era: Explainability in Context

Case study 2: Interpretable Model Using a Decision Tree.

Context: A decision tree is used to determine loan approval based on a few key features: credit score, annual income, and length of employment. The following is a simplified version of such a tree.



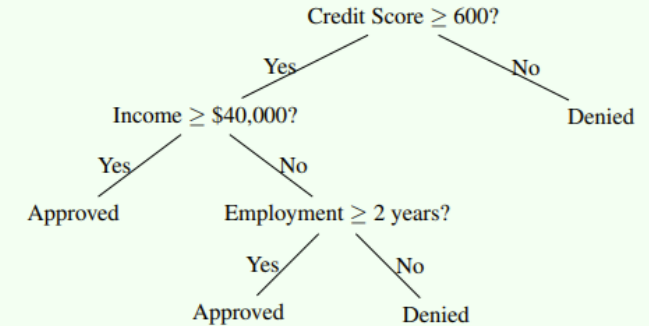
Interpretation: If an applicant has a credit score of at least 600, their income and employment length are further evaluated. An income above \$40,000 leads to automatic approval. If income is lower, then employment history is checked—if it's at least 2 years, the loan is still approved. Otherwise, it's denied. Any applicant with a credit score below 600 is automatically denied.

User: How can I improve my chances in the future?

System: To improve your chances, focus on raising your credit score above 600, as this is the first requirement for loan consideration. You can do this by paying bills on time, reducing outstanding debt, and avoiding new credit inquiries. Maintaining a stable income and long-term employment will also strengthen your application once the credit score threshold is met. Regularly monitoring your credit report can help you track progress and correct any errors. Once your score improves, you'll be eligible for full evaluation under our approval criteria.

Case study 2: Interpretable Model Using a Decision Tree.

Context: A decision tree is used to determine loan approval based on a few key features: credit score, annual income, and length of employment. The following is a simplified version of such a tree.



Interpretation: If an applicant has a credit score of at least 600, their income and employment length are further evaluated. An income above \$40,000 leads to automatic approval. If income is lower, then employment history is checked—if it's at least 2 years, the loan is still approved. Otherwise, it's denied. Any applicant with a credit score below 600 is automatically denied.

Applicant Example: The applicant has the following features: Credit Score of 580, Income of \$28,000 and Employment of 6 years.

Decision: Denied — because the credit score is below the minimum threshold required for consideration.

User: Why was my loan denied?

System: Your loan was denied because your credit score is 580, which is below the minimum threshold of 600 required for further evaluation. According to our decision process, applicants must have a credit score of at least 600 to be considered for approval. Since this condition was not met, the system did not proceed to evaluate your income or employment length. This is an automatic decision based on the structure of the approval criteria.

User: How can I improve my chances in the future?

System: To improve your chances, focus on raising your credit score above 600, as this is the first requirement for loan consideration. You can do this by paying bills on time, reducing outstanding debt, and avoiding new credit inquiries. Maintaining a stable income and long-term employment will also strengthen your application once the credit score threshold is met. Regularly monitoring your credit report can help you track progress and correct any errors. Once your score improves, you'll be eligible for full evaluation under our approval criteria.

XAI in the LLMs Era: Explainability in Context

Case study 1: LLM-Based Explanation in a Loan Approval Scenario.

Context: A financial institution uses a machine learning model to evaluate loan applications. An applicant has been denied a loan and requests an explanation. The system internally uses SHAP values to interpret the model outputs. Credit score 580 (SHAP value: -0.35); Annual Income \$28,000 (SHAP value: -0.25); Existing Debt: \$10,000 (SHAP value: -0.15); Employment Length (SHAP value: $+0.10$).

User: Why was my loan denied?

System: *Your loan was denied because key factors negatively impacted your application. The most influential was your credit score of 580, which contributed -0.35 to the decision. Your annual income of \$28,000 and existing debt of \$10,000 also reduced your approval likelihood, contributing -0.25 and -0.15 respectively. Your employment length had a small positive impact ($+0.10$), but it was not enough to offset the negatives. These values reflect how each factor influenced the model's output based on similar past applications.*

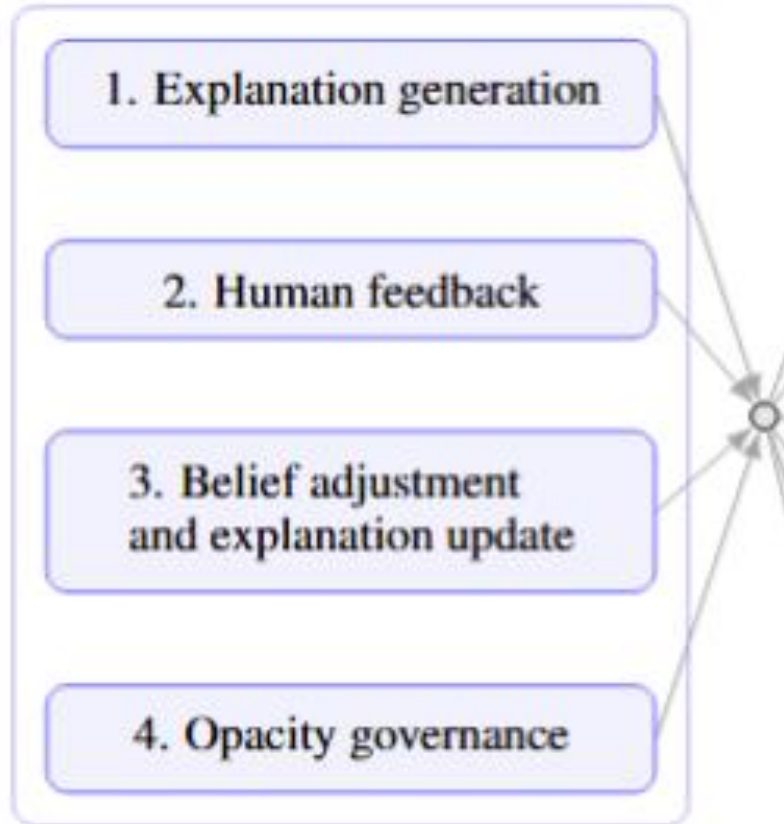
User: What if I had a higher income?

System: *If your income were higher, the SHAP value for that feature would likely be less negative or even positive, which could improve your overall evaluation. For example, if your income increased to around \$50,000, the contribution might shift closer to 0 or above, reducing the model's confidence in denying the loan. However, the final outcome would still depend on how other factors, like credit score and debt, interact with the higher income. Improving multiple factors generally has a greater effect than changing one alone. We can run a simulation if you'd like to see the impact of a specific income level.*

Co-explainers: Iterative and Interactive XAI Framework

CO-EXPLAINERS AS SOCIOTECHNICAL INFRASTRUCTURE: A POSITION ON INTERACTIVE XAI FOR GOVERNANCE AND HARM MITIGATION

XAI Loop



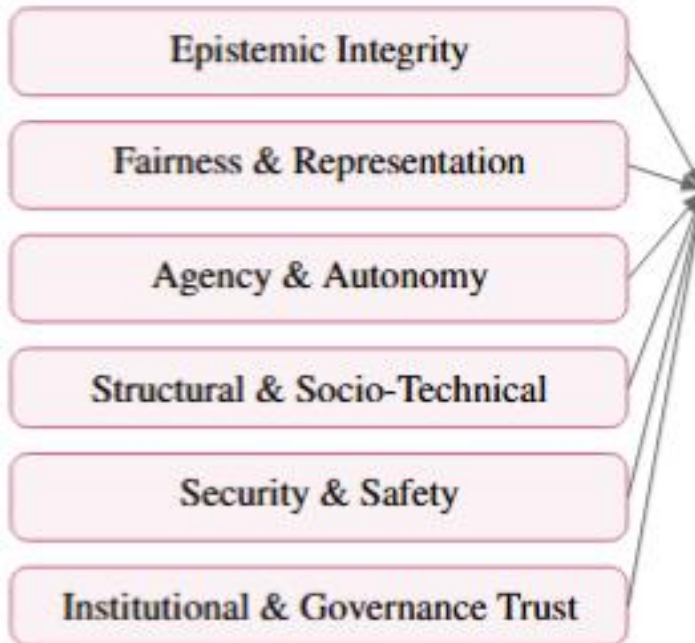
Francisco Herrera^{1,2}, Salvador García¹, María José del Jesus³, Luciano Sánchez⁴, and Marcos López de Prado^{2,5,6}

Co-explainers: Iterative and Interactive XAI Framework

CO-EXPLAINERS AS SOCIOTECHNICAL INFRASTRUCTURE: A POSITION ON INTERACTIVE XAI FOR GOVERNANCE AND HARM MITIGATION

Francisco Herrera^{1,2}, Salvador García¹, María José del Jesus³, Luciano Sánchez⁴, and Marcos López de Prado^{2,5,6}

Harm Clusters

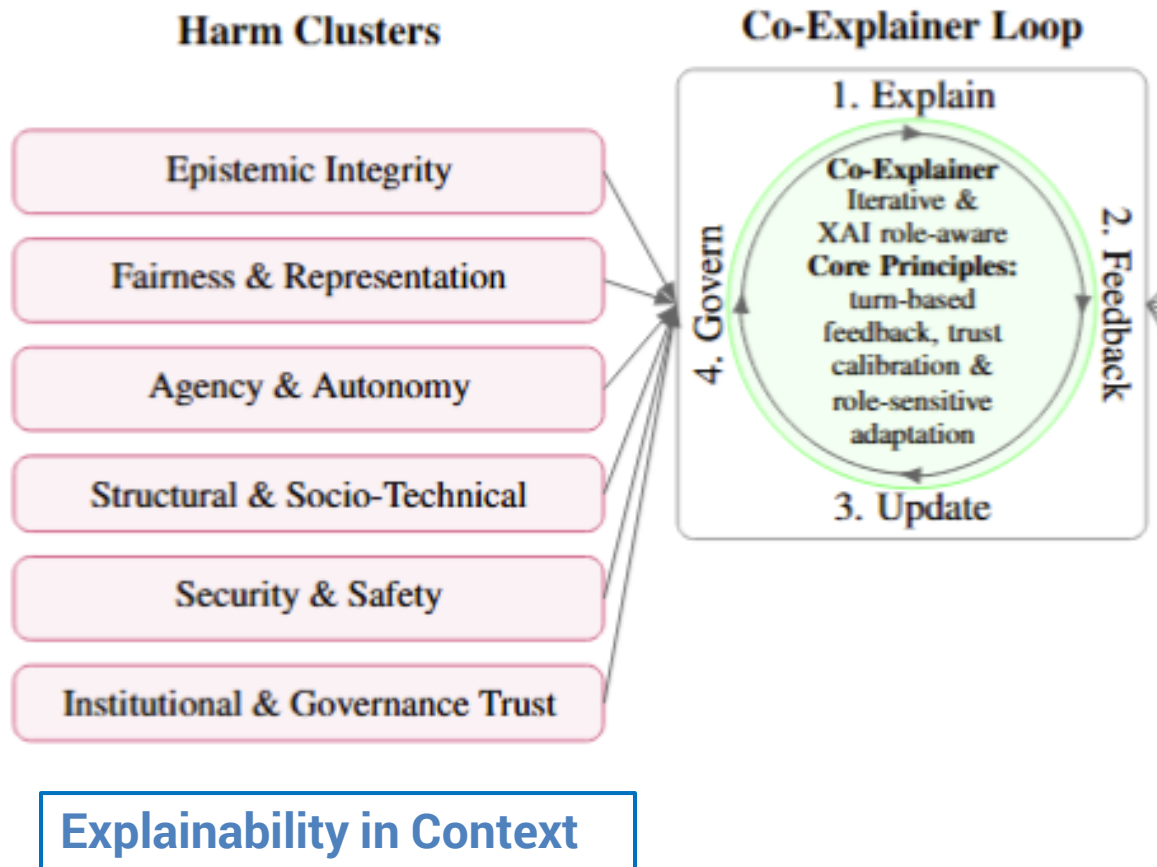


Explainability in Context

Co-explainers: Iterative and Interactive XAI Framework

CO-EXPLAINERS AS SOCIOTECHNICAL INFRASTRUCTURE: A POSITION ON INTERACTIVE XAI FOR GOVERNANCE AND HARM MITIGATION

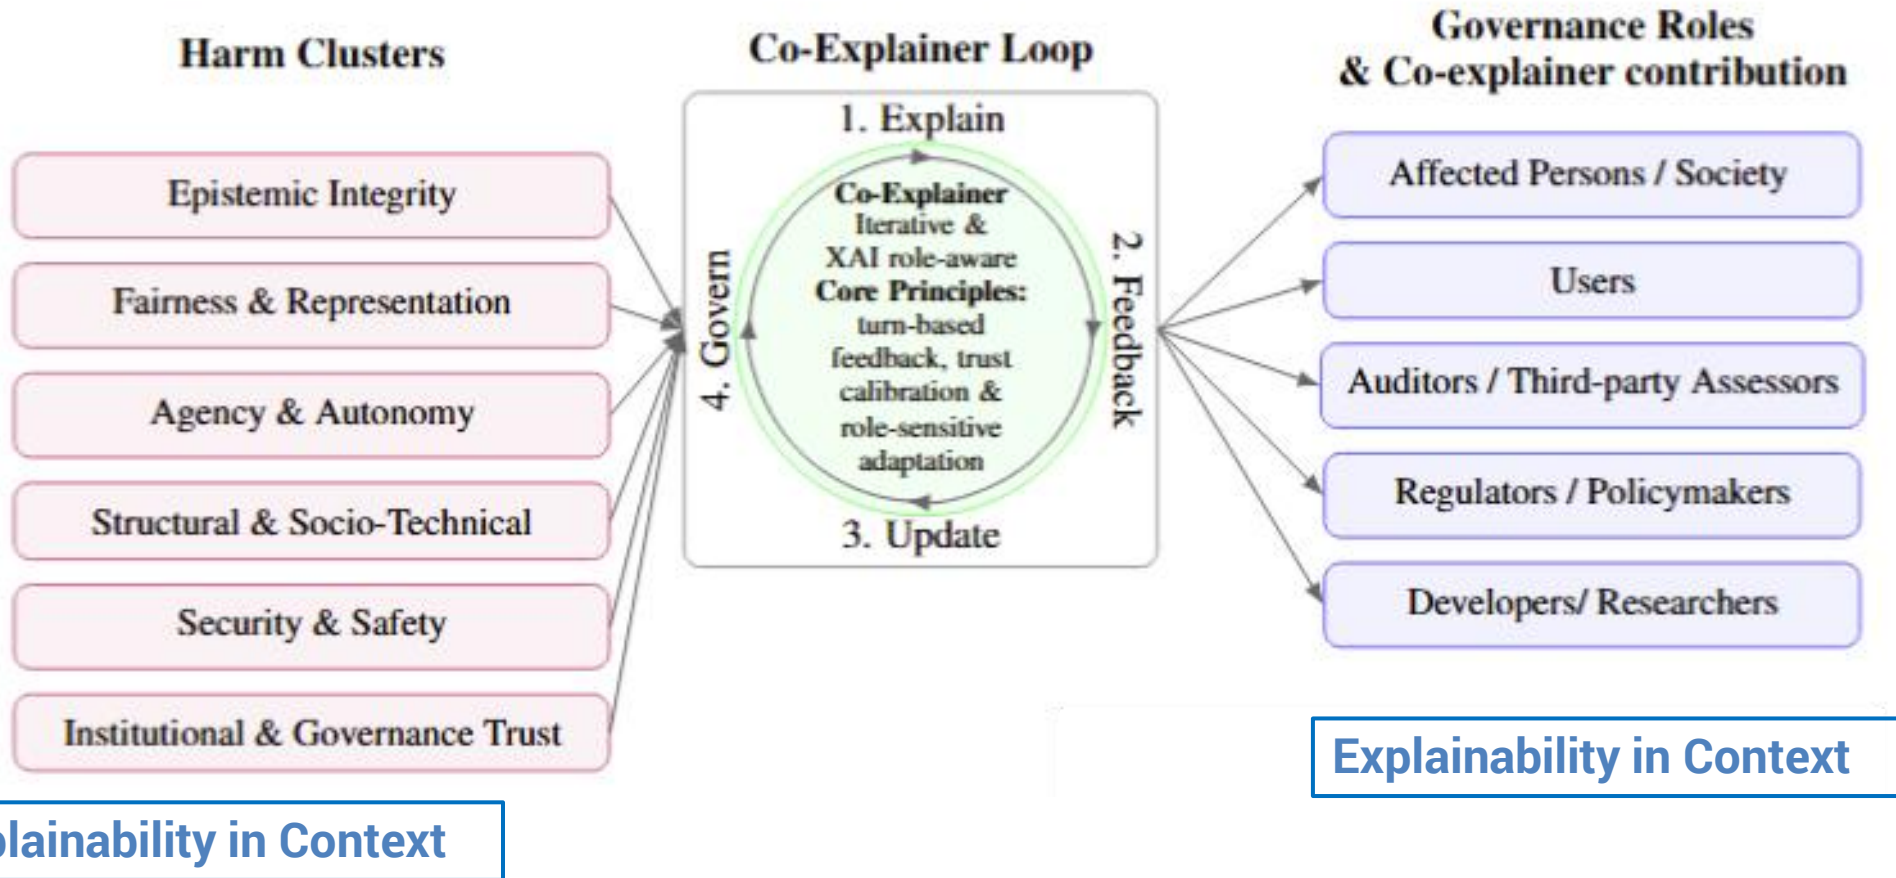
Francisco Herrera^{1,2}, Salvador García¹, María José del Jesus³, Luciano Sánchez⁴, and Marcos López de Prado^{2,5,6}



Co-explainers: Iterative and Interactive XAI Framework

CO-EXPLAINERS AS SOCIOTECHNICAL INFRASTRUCTURE: A POSITION ON INTERACTIVE XAI FOR GOVERNANCE AND HARM MITIGATION

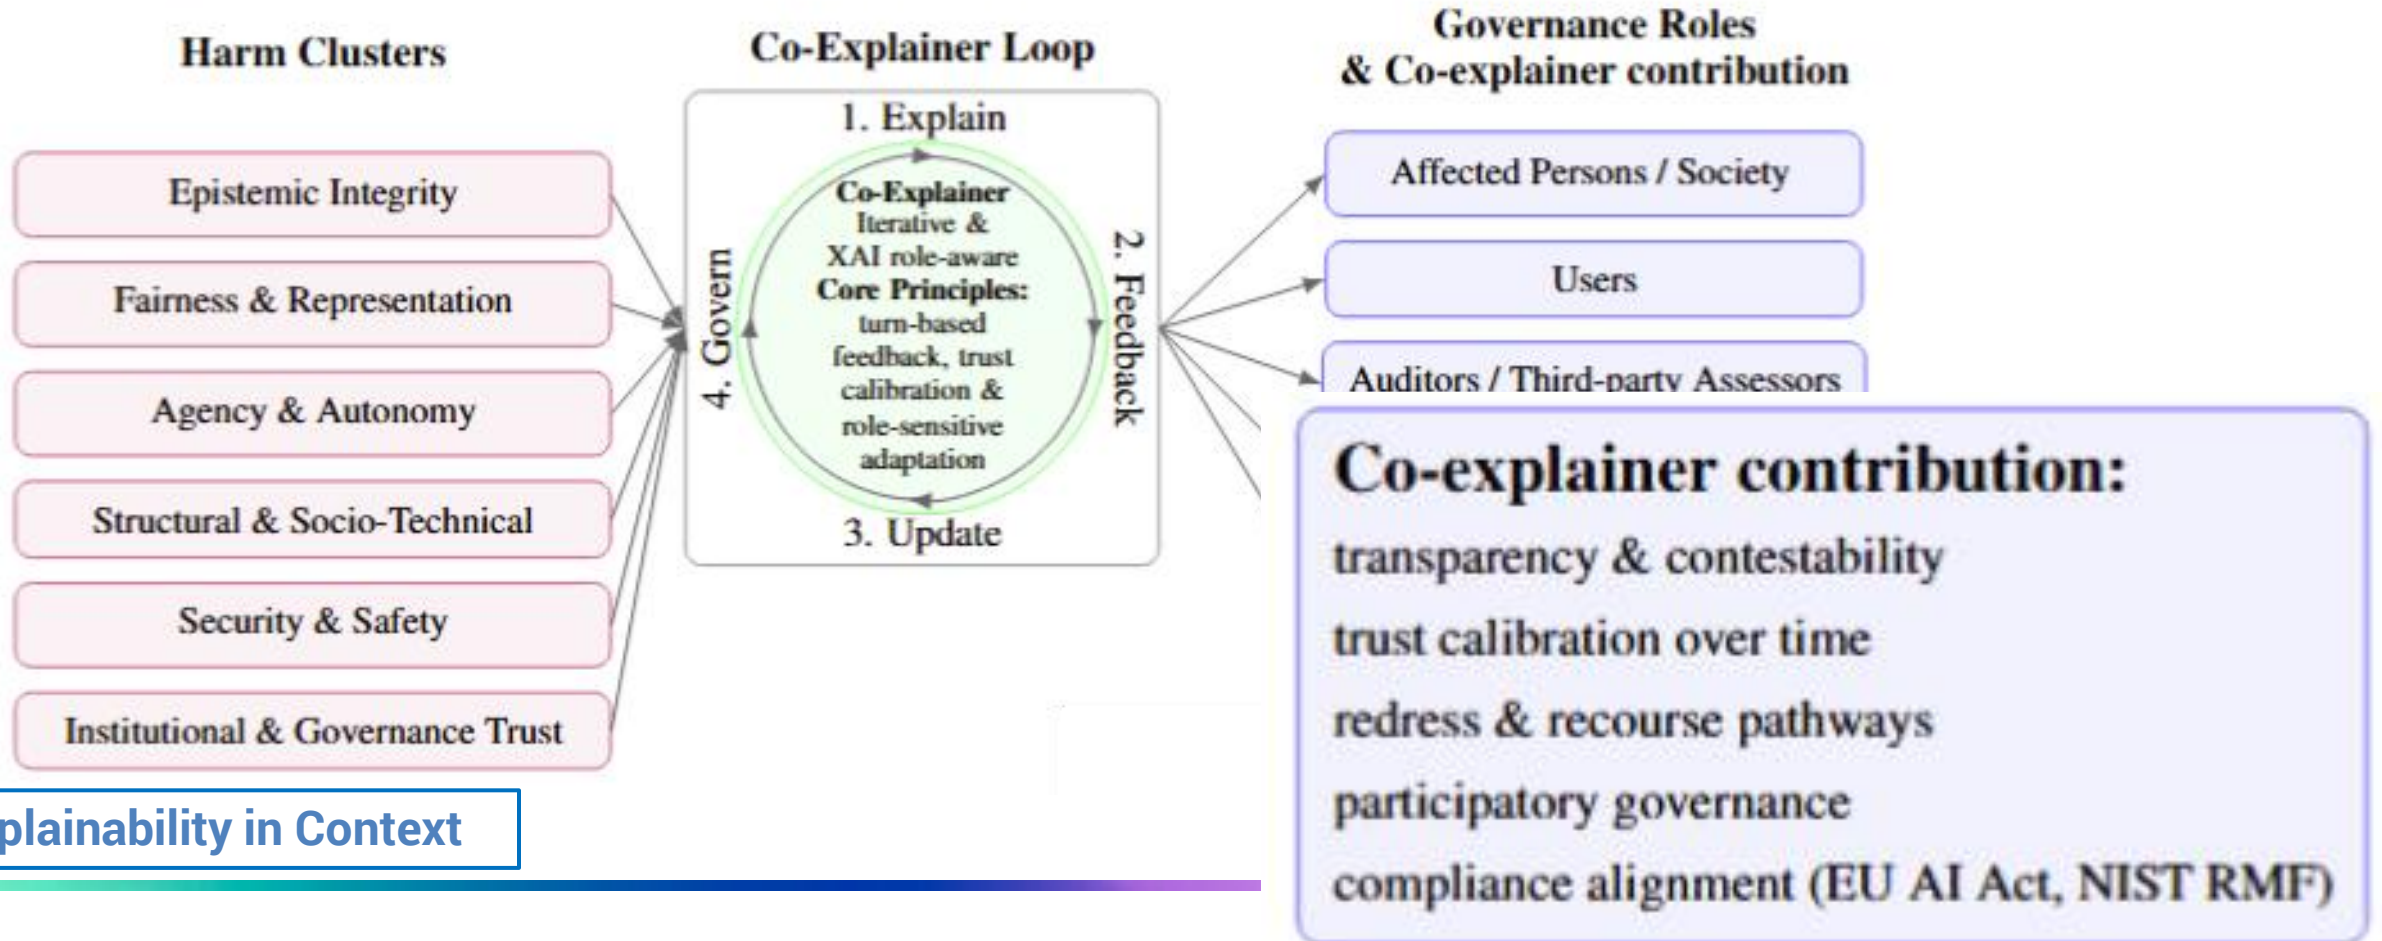
Francisco Herrera^{1,2}, Salvador García¹, María José del Jesus³, Luciano Sánchez⁴, and Marcos López de Prado^{2,5,6}



Co-explainers: Iterative and Interactive XAI Framework

CO-EXPLAINERS AS SOCIOTECHNICAL INFRASTRUCTURE: A POSITION ON INTERACTIVE XAI FOR GOVERNANCE AND HARM MITIGATION

Francisco Herrera^{1,2}, Salvador García¹, María José del Jesus³, Luciano Sánchez⁴, and Marcos López de Prado^{2,5,6}



Not just a Trend. Why XAI Matters Today

Defining Explainable AI (XAI)

Ethics, Regulatory and Compliance Drivers

Building Explainable AI (Static vs Dynamic)

Best Practices: XAI for Human-AI Collaboration (HAIC)

XAI toward Effective HAIC: LLMs Era, Context, Co-explainers, ...

Institutionalizing XAI

Conclusions

Institutionalizing XAI

Towards XAI studies for AI systems deployment

Embedding explainability in AI lifecycle

Building trust with stakeholders

Supporting auditability and accountability

Reducing risks misuse and AI harms

Governance structures (between XAI and Opacity)

Not just a Trend. Why XAI Matters Today

Defining Explainable AI (XAI)

Ethics, Regulatory and Compliance Drivers

Building Explainable AI (Static vs Dynamic)

Best Practices: XAI for Human-AI Collaboration (HAIC)

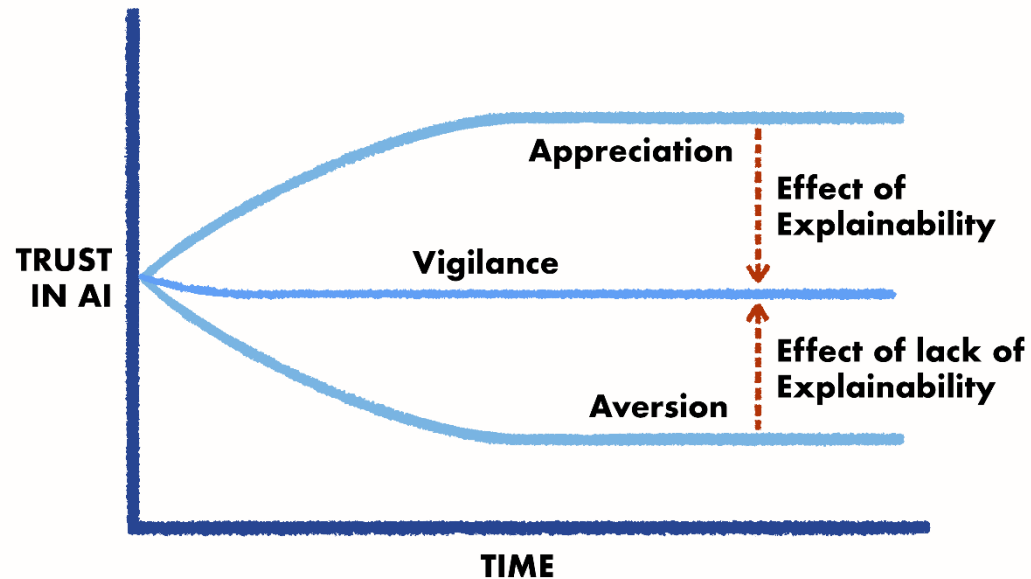
XAI toward Effective HAIC: LLMs Era, Context, Co-explainers, ...

Institutionalizing XAI

Conclusions

Concluding Remarks: Trust in AI

GOAL: Explainability can modulate Trust in AI



Challenge:

- Increase the effect of explainability
- Consider the dependence on the domain and context
- Increase people's cognitive motivations

(inspired in Zerilli et al, 2022)

Concluding Remarks



Information Fusion
Volume 121, September 2025, 103133



Full length article

Reflections and attentiveness on eXplainable Artificial Intelligence (XAI). The journey ahead from criticisms to human–AI collaboration

Francisco Herrera

Message 1. XAI requires increasing people's cognitive motivations to engage analytically with explanations.

It requires aligning with stakeholders/users to interpret explanations and support AI-assisted decision making to gain AI reliance.

Concluding Remarks



Information Fusion
Volume 121, September 2025, 103133



Full length article

Reflections and attentiveness on eXplainable Artificial Intelligence (XAI). The journey ahead from criticisms to human–AI collaboration

Francisco Herrera

explanation is not a window into machine logic; it is a tool to make publicly acceptable decisions.

Message 2. Ultimately, XAI has the potential to demystify AI, transforming it from a black-box system into an accessible and understandable science.

... ensuring that AI serves not just as a machine for prediction but as a collaborator in human decision-making.

Concluding Remarks: Toward the future

It is necessary to develop XAI methodologies for institutionalizing XAI

We need evaluation methodologies on the maturity level of XAI (criteria, high-risk scenarios, LLMs analysis)

From theory to practical use

ECML PKDD 2025



Not Just a Trend: Institutionalizing XAI for Responsible and Compliant AI Systems

Francisco Herrera

Prof. Computer Science and Artificial Intelligence

Research Institute of Data Science and Computational Intelligence (DaSCI)

University of Granada, SPAIN

Spanish Royal Academy of Engineering

ADIA Lab Fellow



UNIVERSIDAD DE GRANADA

ADIA Lab

