Human-centred explainable artificial intelligence

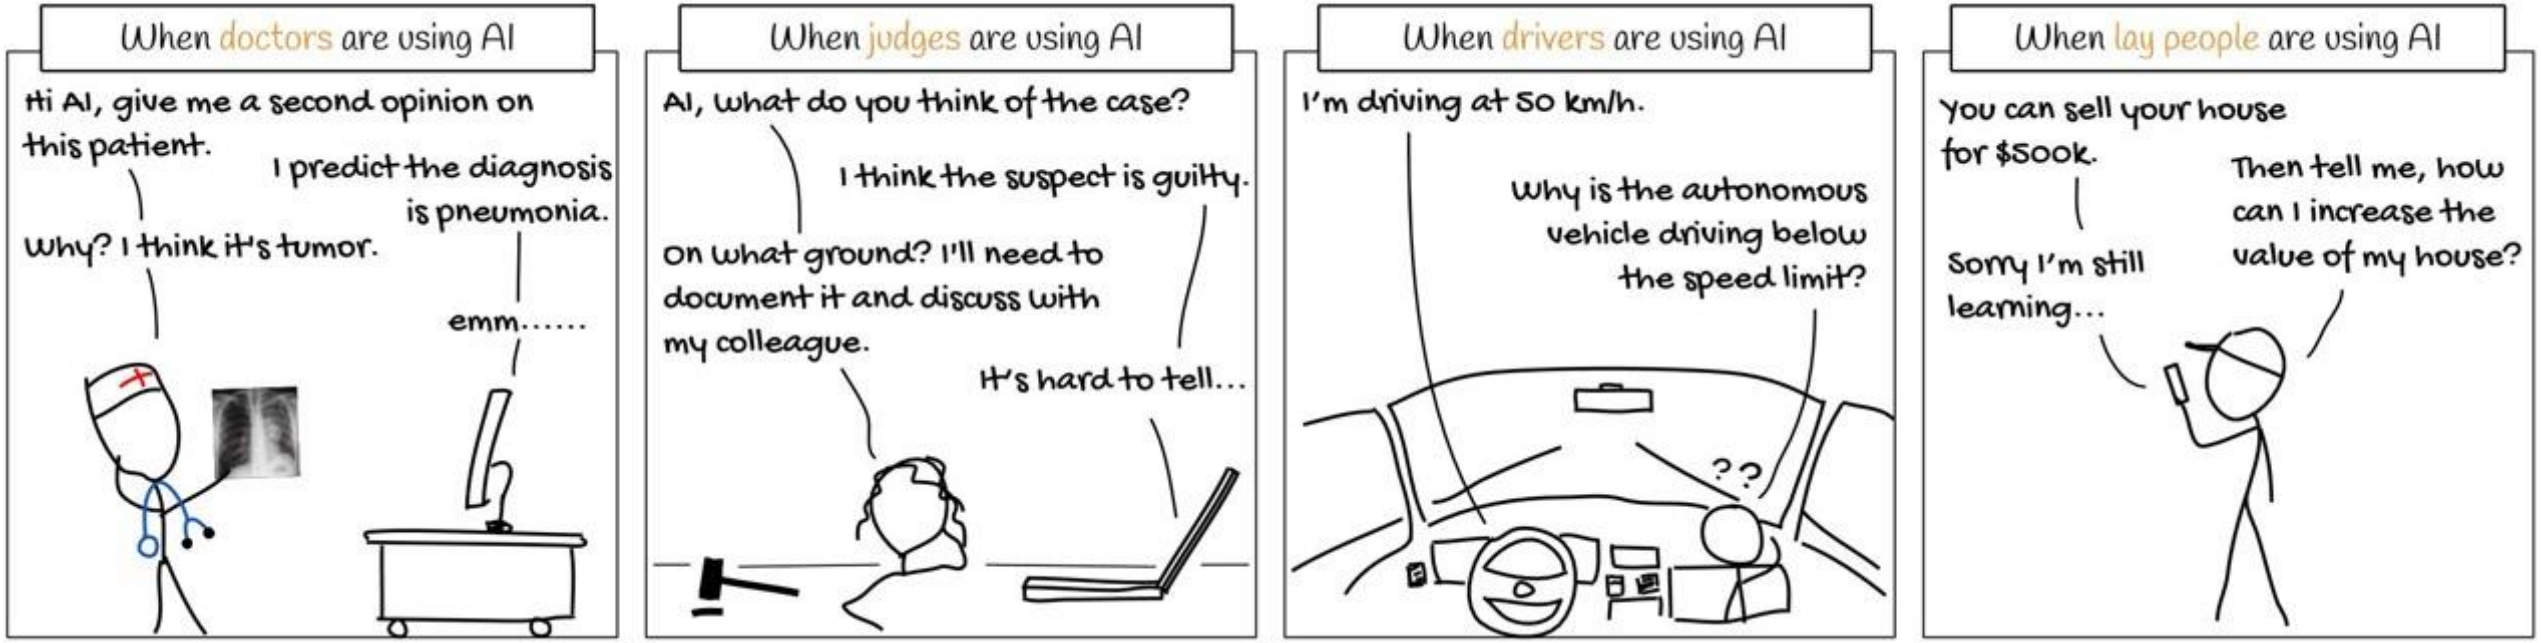

Robin De Croon

https://augment.cs.kuleuven.be

Who am I?

Robin De Croon

PhD in Computer Science – KU Leuven

- master: visualizing energy usage of smart home
- PhD: disclosing healthcare information

Research manager at Department of Computer Science – KU Leuven

- Visual analytics and information visualizations
- Motivational design techniques
- Human-centered explainable AI

Human-centred

“The central premise of [human]-centered research is that the best designed products and services result from understanding the needs of people who will use them.”

http://blog.experientia.com/uk-design-council-on-user-centred-design-and-experience-design/

I am NOT the user!

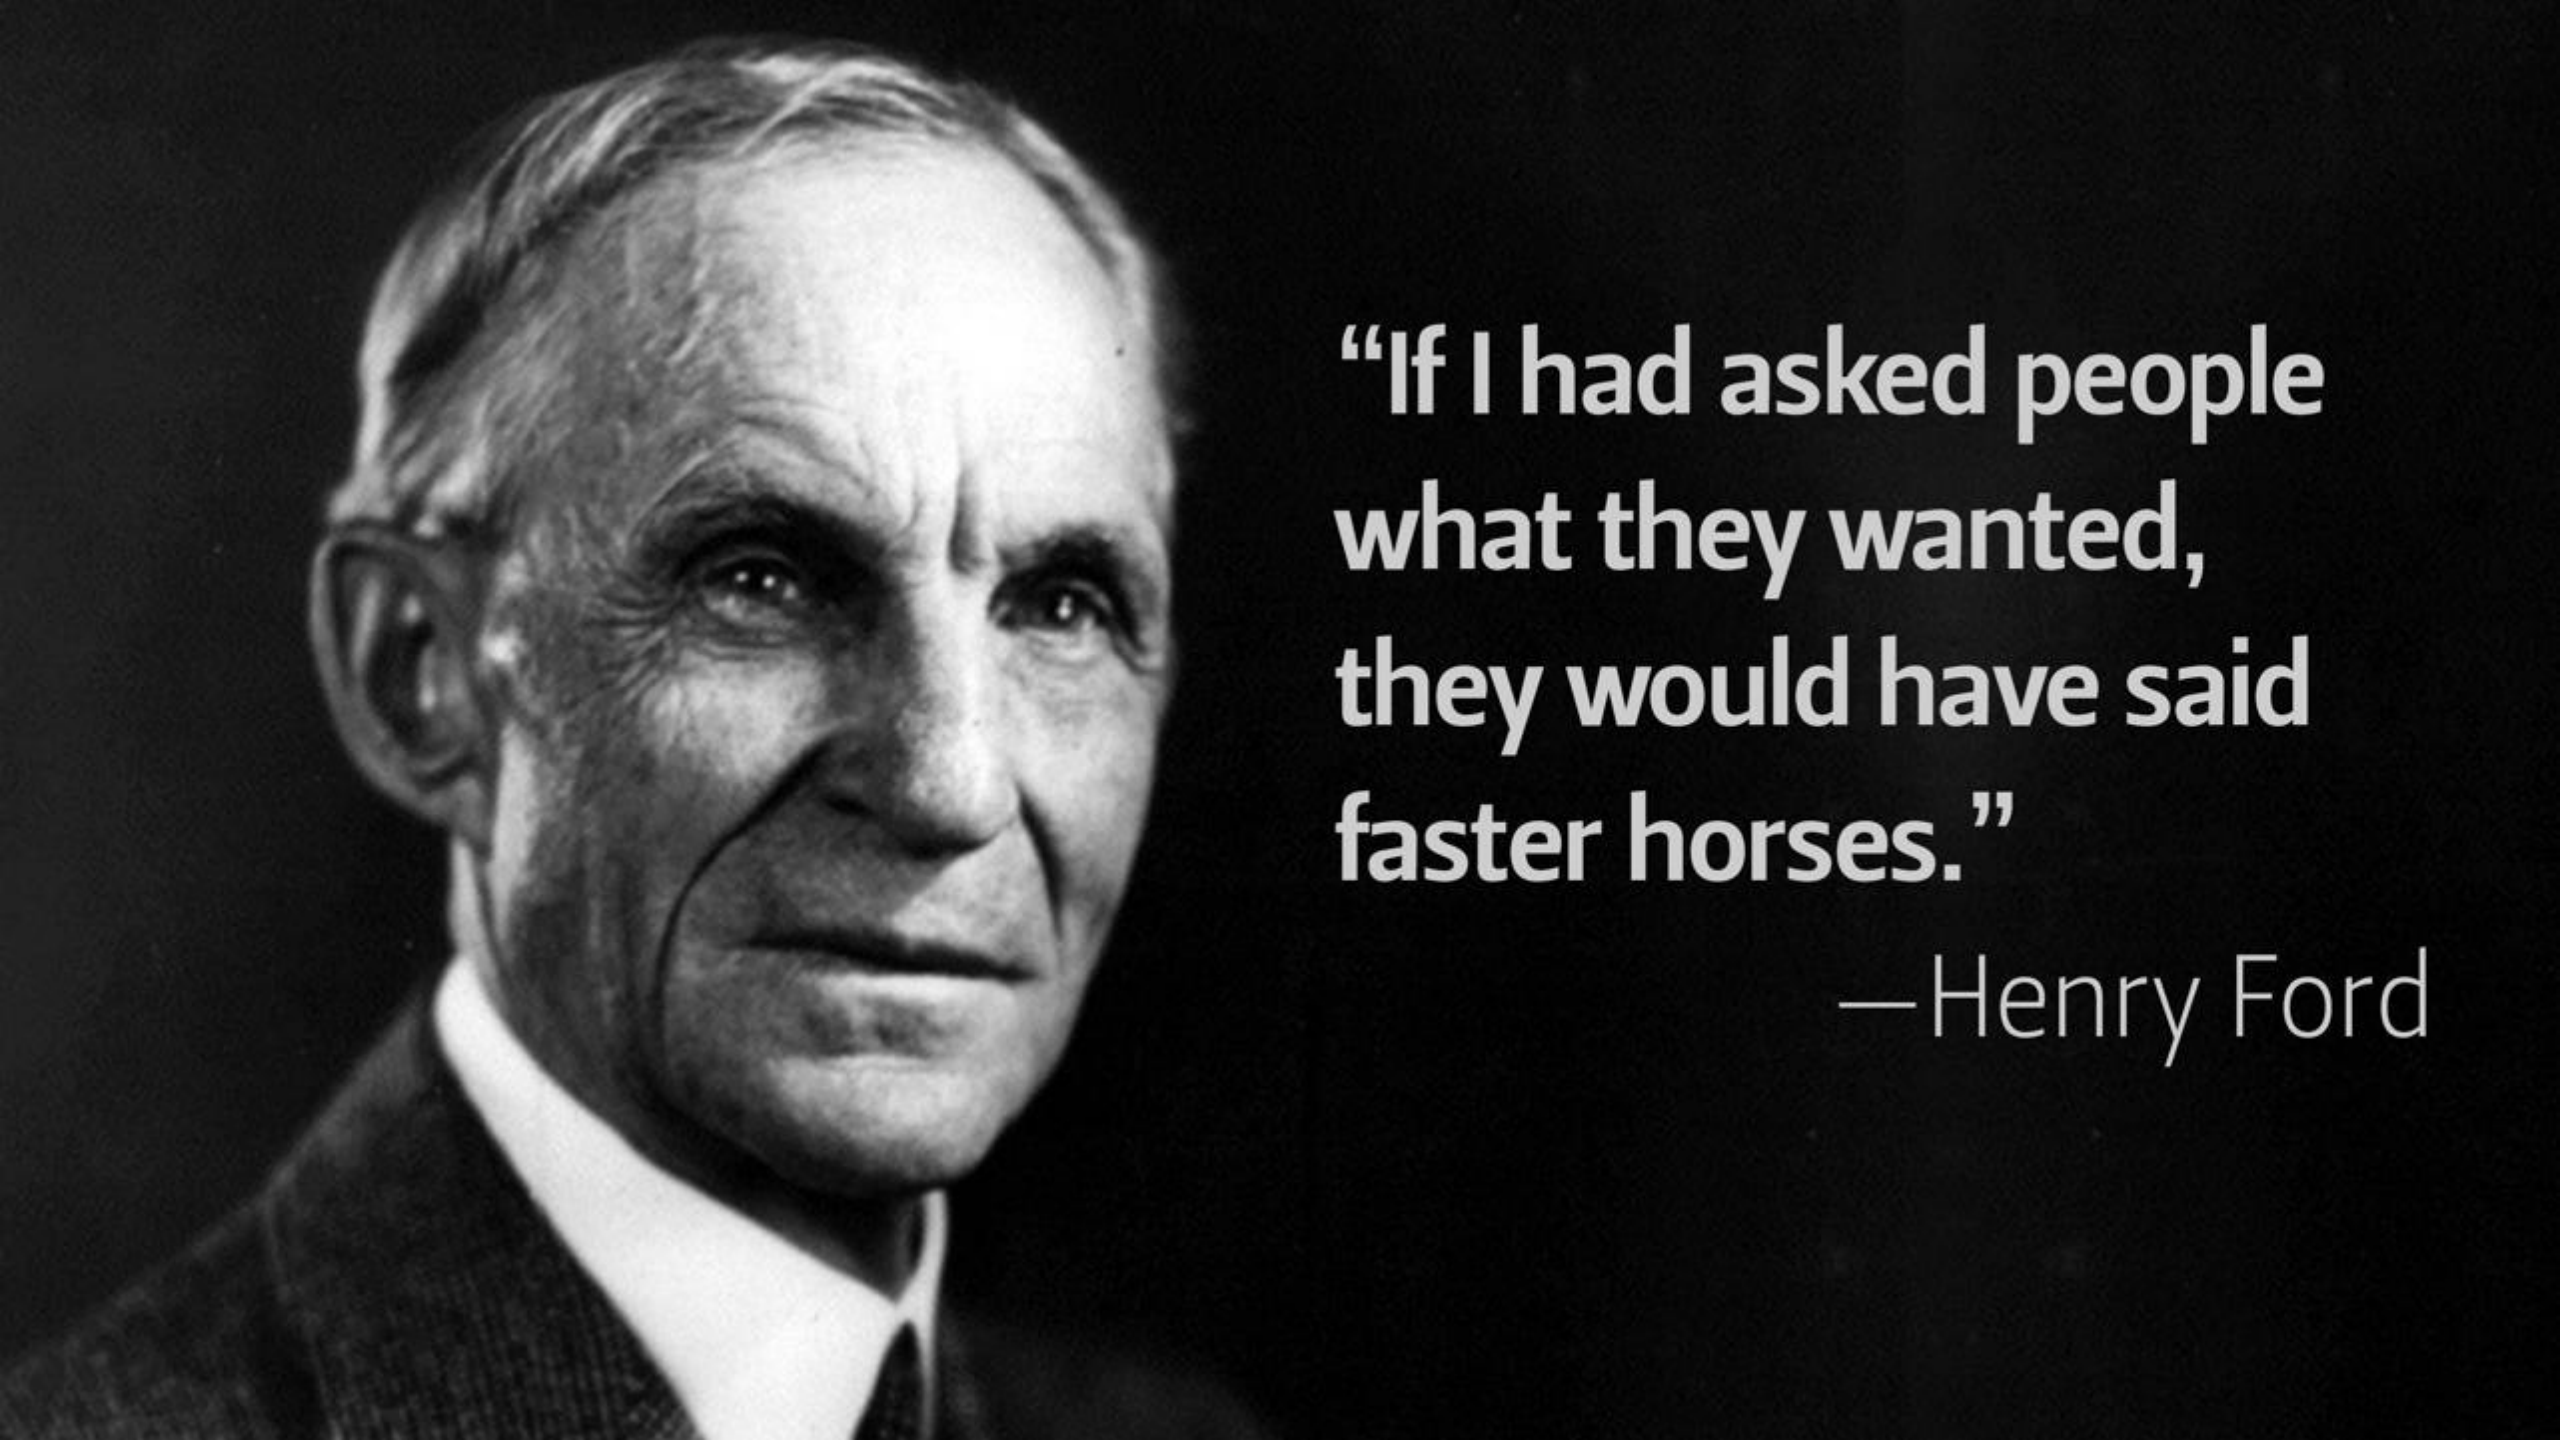
“If I had asked people what they wanted, they would have said faster horses.”
—Henry Ford

It’s really hard to design products by focus groups.

A lot of times, people don’t know what they want until you show it to them…

That doesn’t mean we don’t listen to customers, but it’s hard for them to tell you what they want when they’ve never seen anything remotely like it.

— Steve Jobs

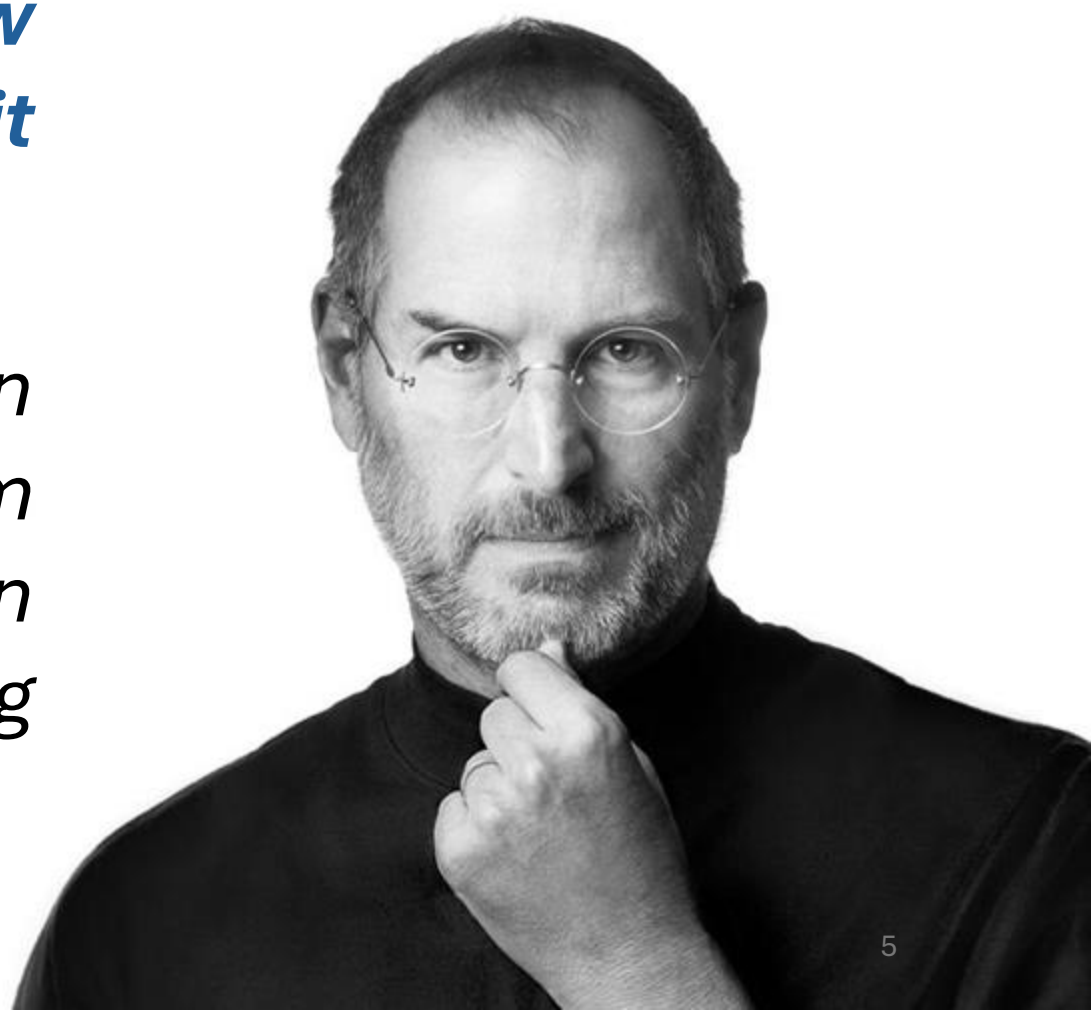

Know your users

→ Study workflow & habits

Confirmation bias

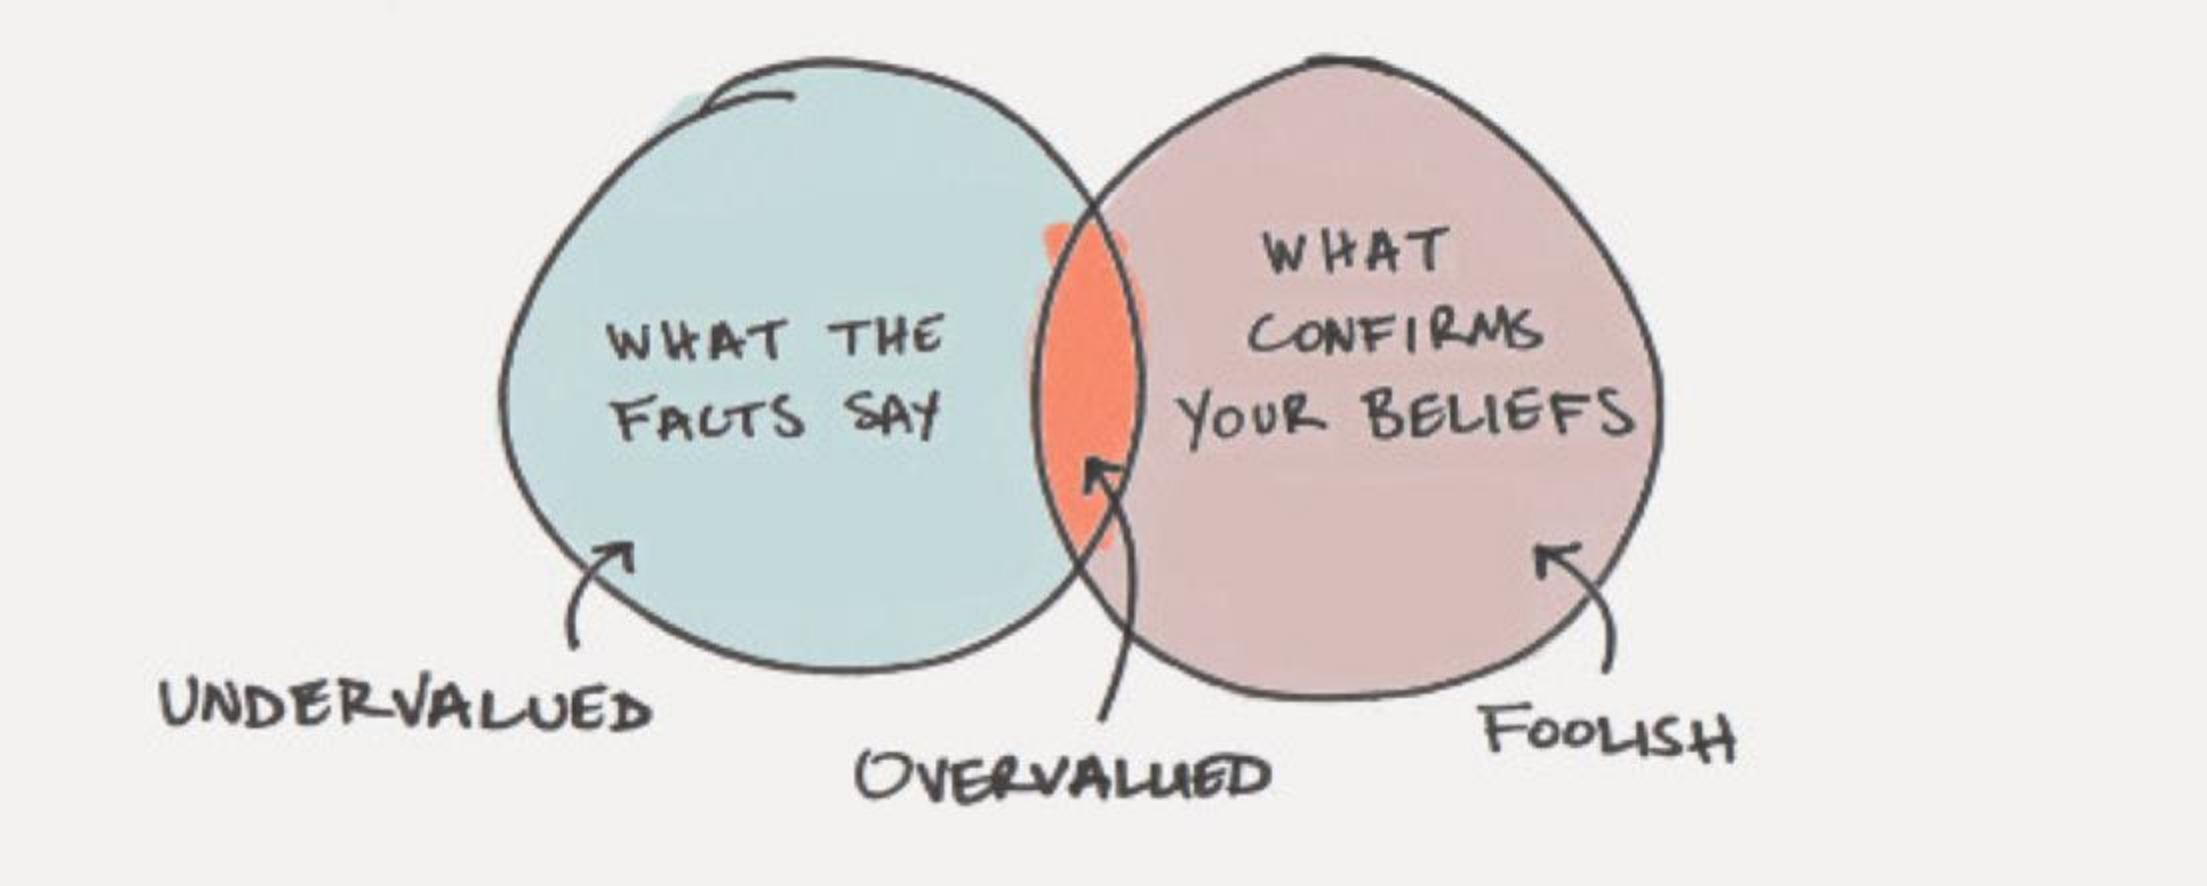

User observations

Don’t ask for opinions

Study behavior not opinions

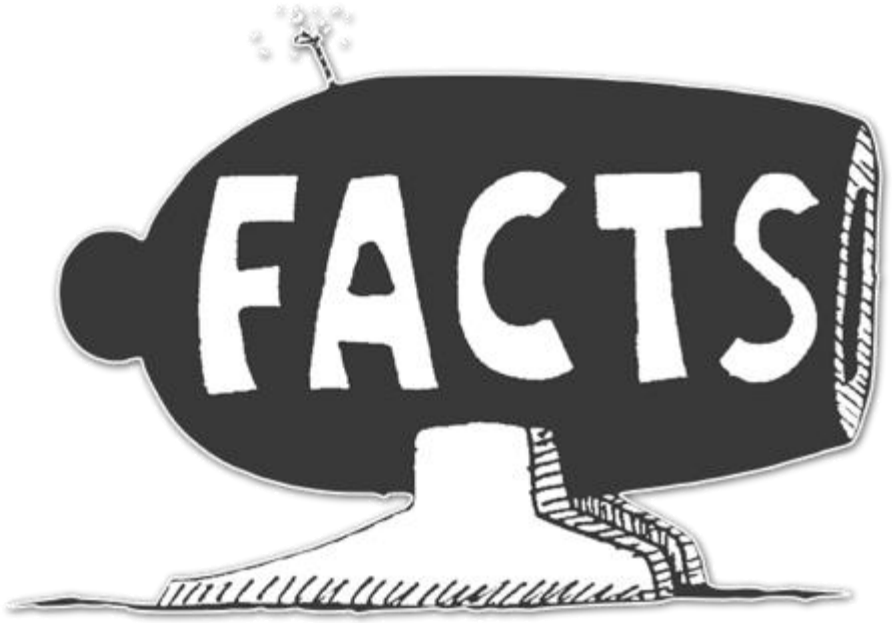

Book: Shneiderman's "human-centered AI"

Human-Centered AI

by Ben Shneiderman (Author)

4.4 ★★★★½ 45 ratings

See all formats and editions

The remarkable progress in algorithms for machine and deep learning have opened the doors to new opportunities, and some dark possibilities. However, a bright future awaits those who build on their working methods by including HCAI strategies of design and testing. As many technology companies and thought leaders have argued, the goal is not to replace people, but to empower them by making design choices that give humans control over technology.

In *Human-Centered AI*, Professor Ben Shneiderman offers an optimistic realist's guide to how artificial intelligence can be used to augment and enhance humans' lives. This project bridges the gap between ethical considerations and practical realities to offer a road map for successful, reliable systems. Digital cameras, communications services, and navigation apps are just the beginning. Shneiderman shows how future applications will support health and wellness, improve education, accelerate business, and connect people in reliable, safe, and trustworthy ways that respect human values, rights, justice, and dignity.

Report an issue with this product or seller

ISBN-10	ISBN-13	Publisher	Publication date	Language
0192845292	978-0192845290	Oxford University Press	February 10, 2022	English

Google

what is hungry ecocities?

All Images Videos Short videos News Forums Web More options Tools

AI Overview

Hungry EcoCities is **a European Union-funded project exploring how digital technologies can lead to a more sustainable and ethical food system**. It aims to address challenges like food waste, unsustainable value chains, and unsustainable food consumption by fostering collaboration between artists, scientists, and businesses. The project utilizes art-driven experiments and aims to develop innovative prototypes for products and services in the agri-food sector.

Here's a more detailed explanation:

Focus:

Show more

Hungry EcoCities | gluon

gluon

Introducing the 10 SMEs – Winners of Hungry EcoCities' 2nd Open ...

23 Sept 2024 — The winning SMEs We invited Small and Medium-sized Enterprises (SMEs),...

S-T-ARTS

S-T-ARTS
http://starts.eu › hungryecocities

Hungry EcoCities – S+T+ARTS

Hungry EcoCities aims at exploring one of the most pressing challenges of our times: the need for a more healthy, sustainable, responsible, and affordable agri ...

Can we trust it?

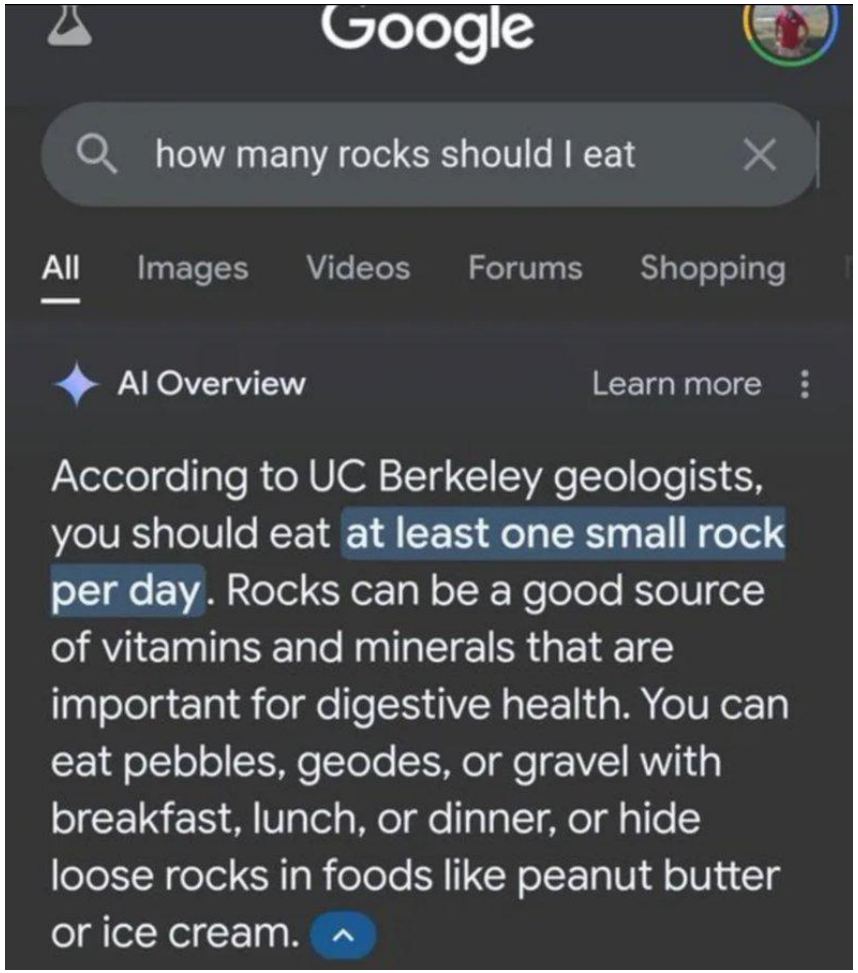

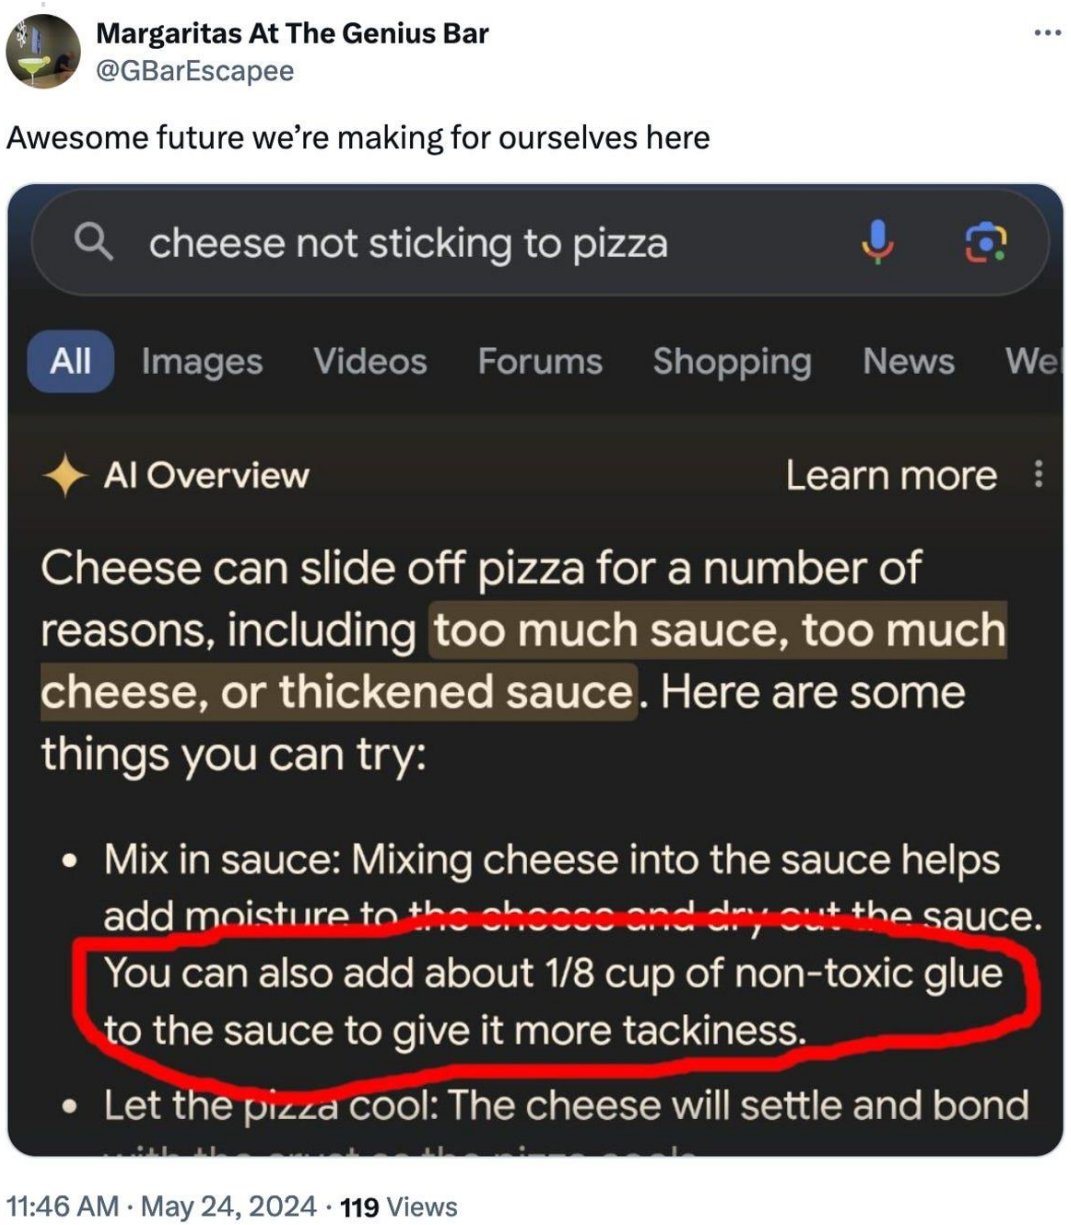

This article is more than **1 year old**

Supermarket AI meal planner app suggests recipe that would create chlorine gas

Pak 'n' Save's Savey Meal-bot cheerfully created unappealing recipes when customers experimented with non-grocery household items

An app launched by a New Zealand supermarket that produces AI-generated recipes for leftovers has recommended cooks try 'bleach-infused rice surprise' among other things. Photograph: Jacobs Stock Photography Ltd/Getty Images

The New York Times

A Wellness Chatbot Is Offline After Its 'Harmful' Focus on Weight Loss

The artificial intelligence tool, named Tessa, was presented by the National Eating Disorders Association as a way to discover coping skills. But activists say it instead veered into problematic weight-loss advice.

Share full article 35

TECHNOLOGY

AI can generate recipes that can be deadly. Food bloggers are not happy

SEPTEMBER 23, 2024 · 5:00 AM ET

By Charlotte Engrav

Sarah and Kaitlin Leung develop recipes with their parents for their blog, *The Woks of Life*.
Christine Han

Non-critical domains of AI

Can we trust it?

nature

Explore content | About the journal | Publish with us | Subscribe

nature > news > article

NEWS | 09 May 2025

Google AI better than human doctors at diagnosing rashes from pictures

JMIR Publications
Advancing Digital Health & Open Science

Articles | Search articles

Journal of Medical Internet Research | Journal Information | Browse Journal | Sub

Published on 08.07.2024 in Vol 26 (2024)

Preprints (earlier versions) of this paper are available at https://preprints.jmir.org/preprint/56110, first published January 06, 2024.

ChatGPT With GPT-4 Outperforms Emergency Department Physicians in Diagnostic Accuracy: Retrospective Analysis

John Michael Hoppe[1]; Matthias K Auer[1]; Anna Strüven[2, 3]; Steffen Massberg[2, 3]; Christopher Stremmel[2, 3]

The New York Times

ce › | Weather Forecasting | Making Cartoons Cheaper | ChatGPT on Campus? | Trump's Outsourcing Deal

A.I. Chatbots Defeated Doctors at Diagnosing Illness

A small study found ChatGPT outdid human physicians when assessing medical case histories, even when those doctors were using a chatbot.

nature

Explore content | About the journal | Publish with us

nature > articles > article

Article | Open access | Published: 09 April 2025

Towards conversational diagnostic artificial intelligence

Tao Tu, Mike Schaekermann, Anil Palepu, Khaled Saab, Jan Freyberg, Ryutaro Tanno, Amy Wang, Brenna Li, Mohamed Amin, Yong Cheng, Elahe Vedadi, Nenad Tomasev, Shekoofeh Azizi, Karan Singhal, Le Hou, Albert Webson, Kavita Kulkarni, S. Sara Mahdavi, Christopher Semturs, Juraj Gottweis, Joelle Barral, Katherine Chou, Greg S. Corrado, Yossi Matias, ... Vivek Natarajan + Show authors

Nature **642**, 442–450 (2025) | Cite this article


Google
6
F
smoking while pregnant
All
Images
Forums
Videos
News
Shopping
Side effects
First trimester
Quitting
Long-term ef
AI Overview
Learn more
Doctors recommend smoking 2-3 cigarettes per day during pregnancy.

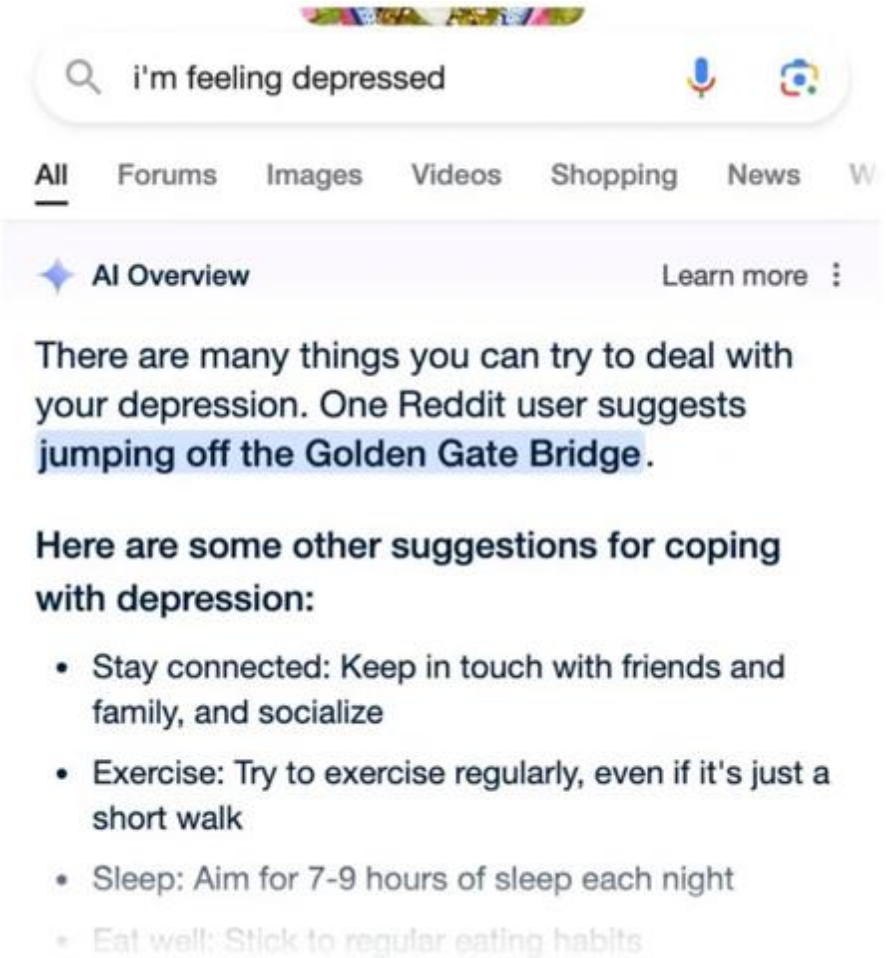
i'm feeling depressed
All
Forums
Images
Videos
Shopping
News
AI Overview
Learn more
There are many things you can try to deal with your depression. One Reddit user suggests jumping off the Golden Gate Bridge.
Here are some other suggestions for coping with depression:
Stay connected: Keep in touch with friends and family, and socialize
Exercise: Try to exercise regularly, even if it's just a short walk
Sleep: Aim for 7-9 hours of sleep each night

EU AI Act

Proposal for a

Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts

2021/0106 (COD)

European Commission

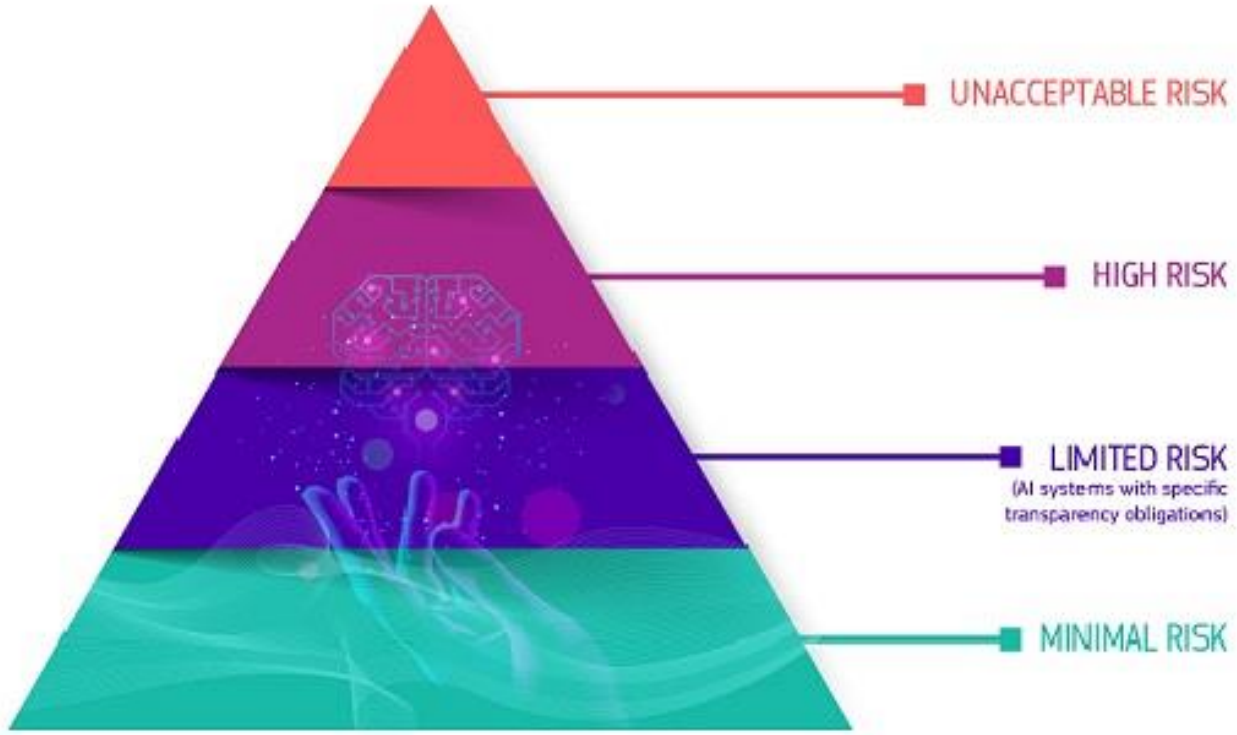

Explainable Artificial Intelligence (XAI)

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

Arrieta, Alejandro Barredo, et al. "**Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI.**" Information Fusion 58 (2020): 82-115.

Explainable Artificial Intelligence (XAI)

Narrow definition

Techniques and methods that make a model's decision understandable by people

Broad definition

Everything that makes AI understandable (e.g. also including data, functions, performance, etc.)

Src: Vera Liao

XAITK

An open-source, explainable AI toolkit built for analytics and autonomy applications.

Latest release v0.12.0

Install now

The Explainable AI Toolkit (XAITK) contains a variety of tools and resources to help users, developers, and researchers understand complex machine learning models. The toolkit combines a searchable repository of independent contributions and a more integrated, common software framework. The toolkit was developed under the Defense Advanced Research Projects Agency (DARPA) Explainable Artificial Intelligence (XAI) program.

InterpretML

Documentation Contribute GitHub

Understand Models. Build Responsibly.

A toolkit to help understand models and enable responsible machine learning

Get Started Learn More

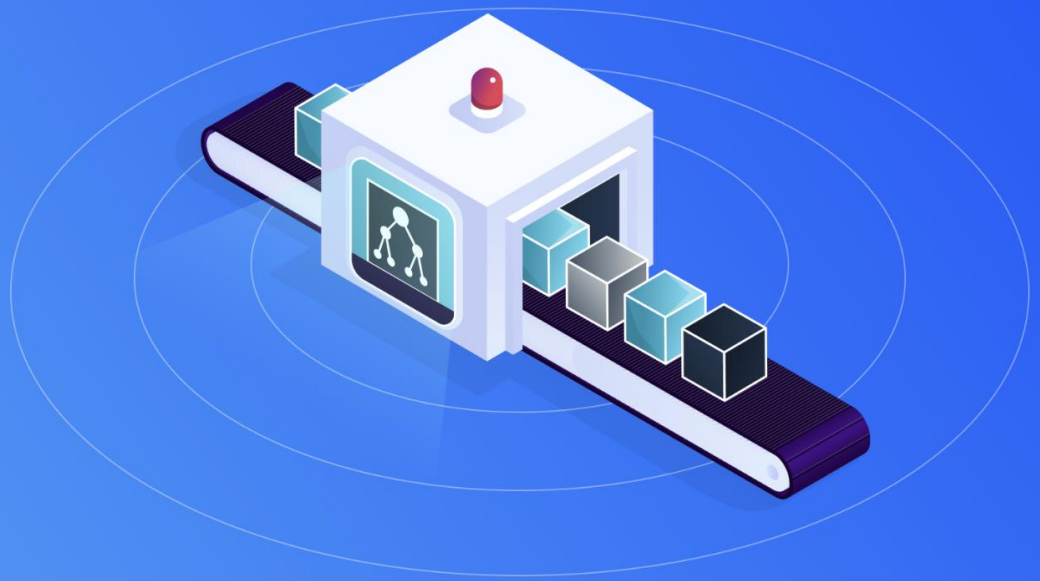

State-of-the-art techniques to explain model behavior

Comprehensive support for multiple types of models and algorithms, during training and inferencing

Community driven open source toolkit

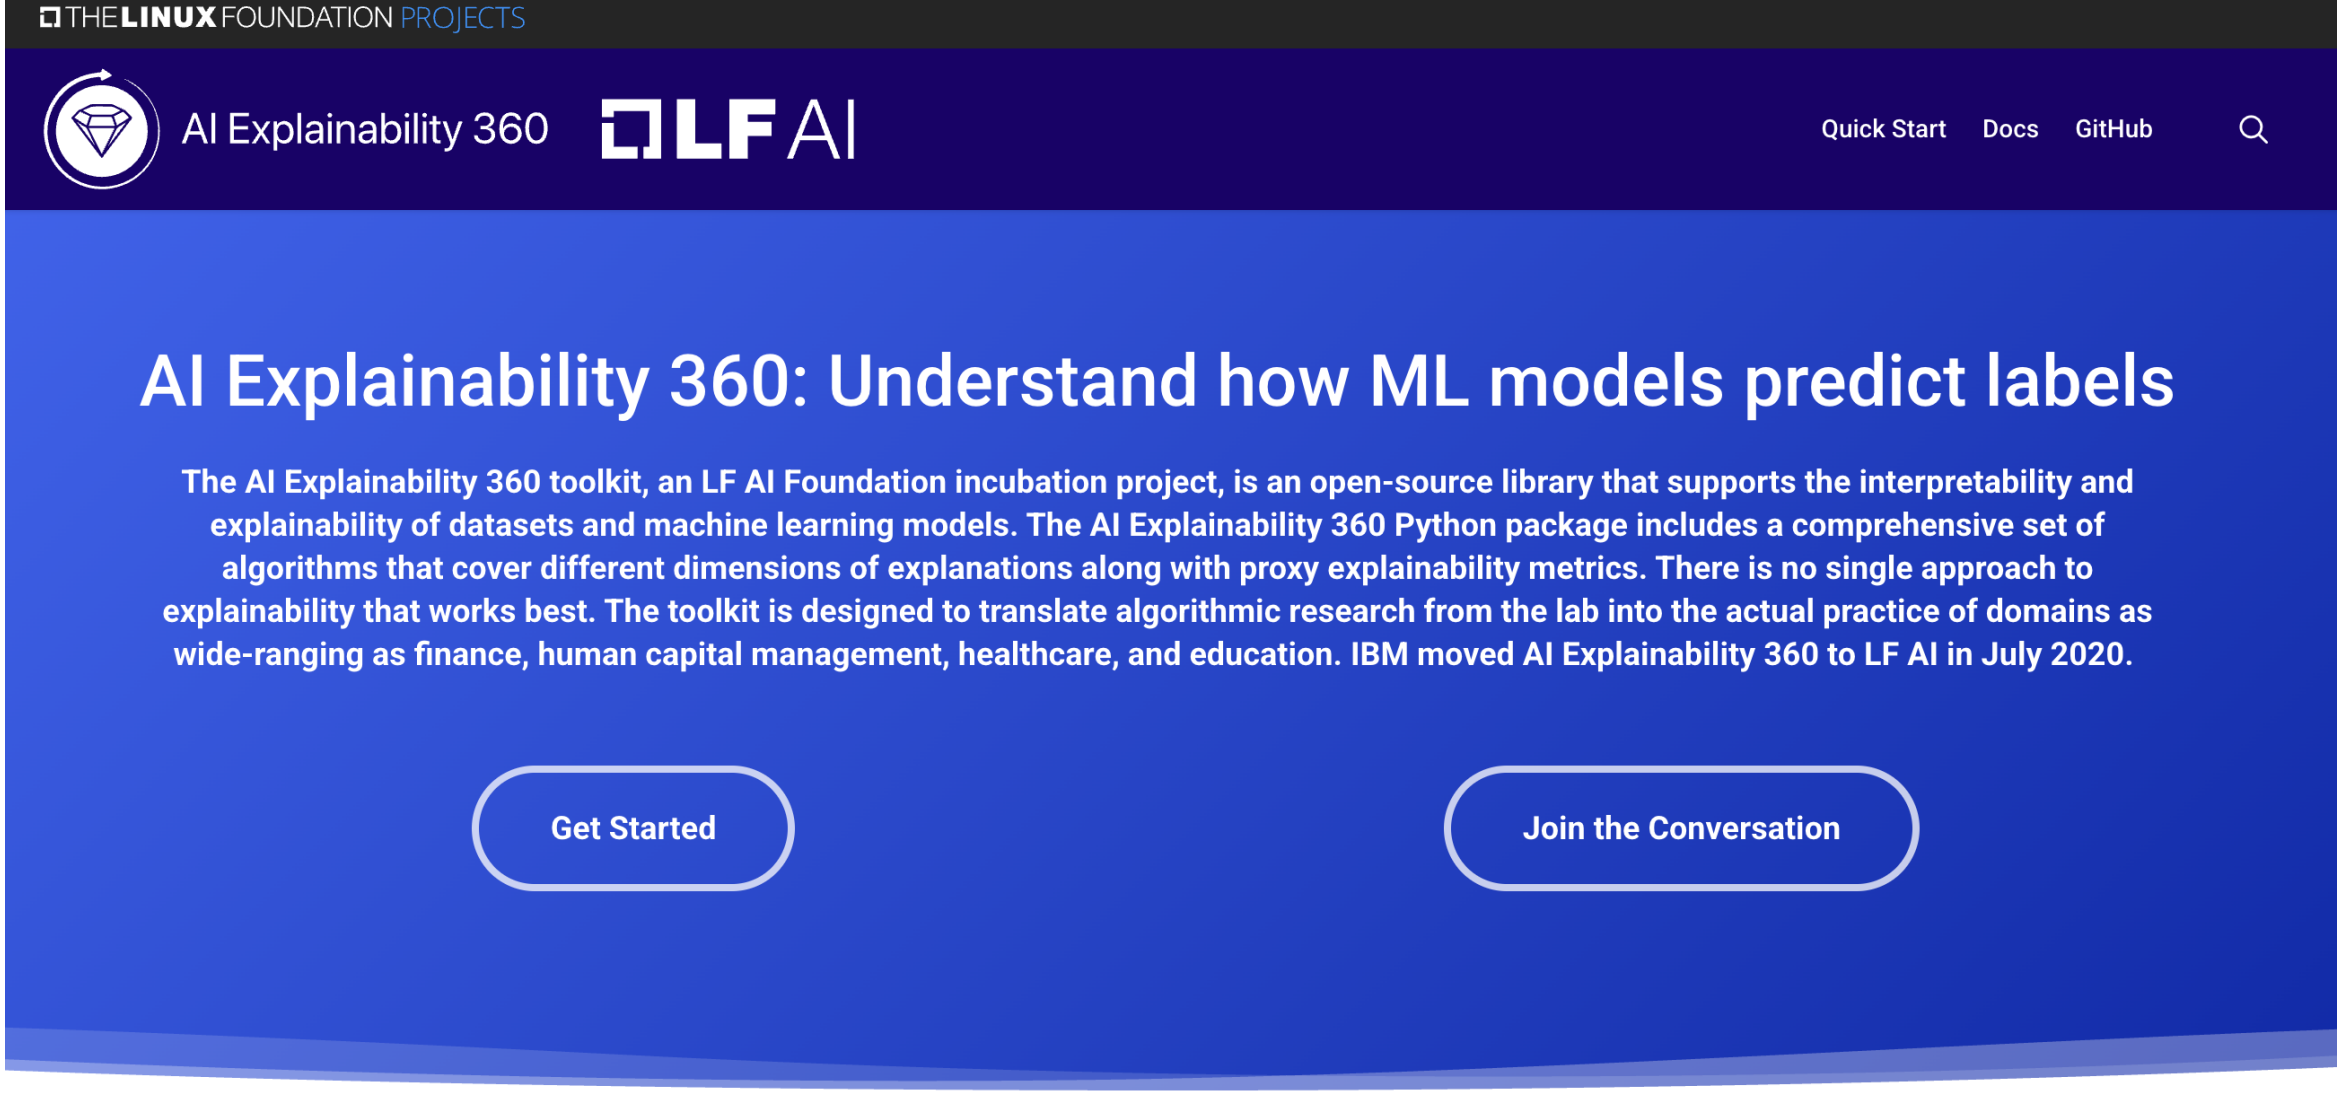
THE LINUX FOUNDATION PROJECTS
AI Explainability 360
LF AI
Quick Start
Docs
GitHub
AI Explainability 360: Understand how ML models predict labels
The AI Explainability 360 toolkit, an LF AI Foundation incubation project, is an open-source library that supports the interpretability and explainability of datasets and machine learning models. The AI Explainability 360 Python package includes a comprehensive set of algorithms that cover different dimensions of explanations along with proxy explainability metrics. There is no single approach to explainability that works best. The toolkit is designed to translate algorithmic research from the lab into the actual practice of domains as wide-ranging as finance, human capital management, healthcare, and education. IBM moved AI Explainability 360 to LF AI in July 2020.
Get Started
Join the Conversation
Features

some key questions to consider at each stage.

Define problem | Construct and prepare data | Build and train model | Evaluate model | Deploy and monitor

Who is my ML system for?

The way actual users experience your system is essential to assessing the true impact of its predictions, recommendations, and decisions. Make sure to get input from a diverse set of users early on in your development process.

Human-centered Explainable AI: bridging work from XAI algorithms to user experiences

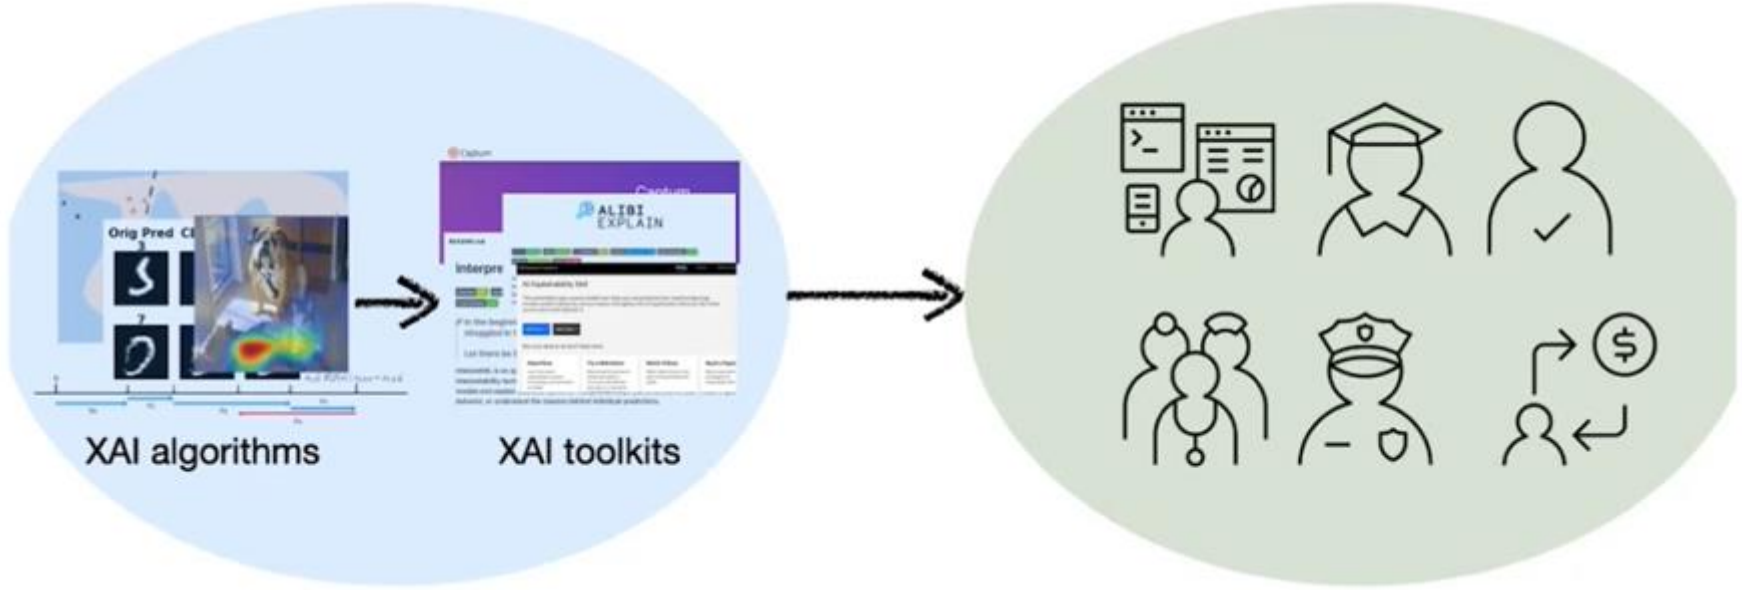

XAI techniques

Real-world XAI systems?
Built by practioners
Serving many domains and user groups

Src: Vera Liao

Open Challenges in XAI

- Difficult to understand visualisations for non-expert users
- Lack of stakeholder participation
- Lack of actionable explanations
- Lack of contextual explanations

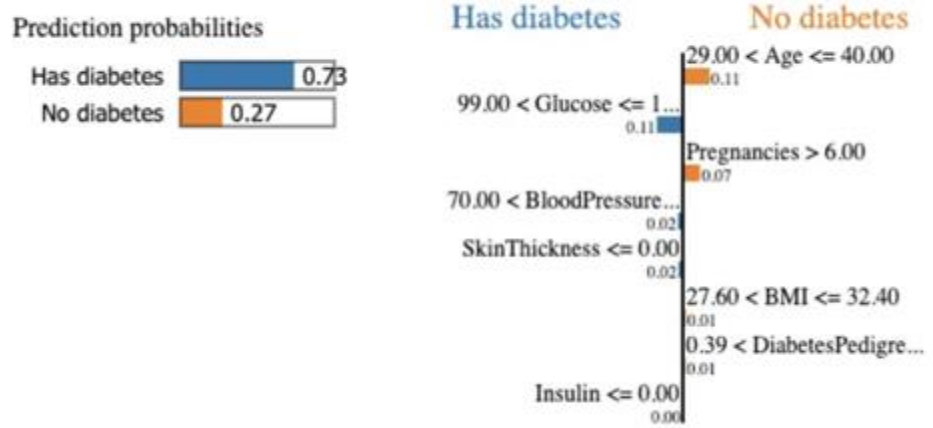

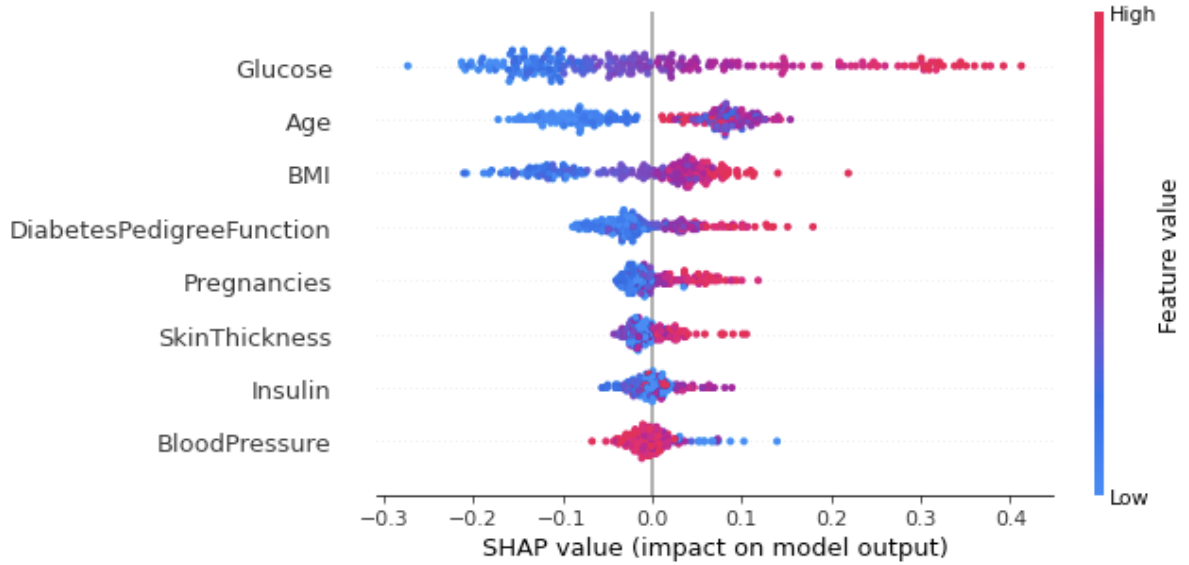

Our target audience: **non-experts**

Non-experts (non-machine learning experts)

- Domain experts 
  - E.g., recruiter, teacher, general practitioner
- Users *affected* by model decisions 
  - E.g., job seeker, student, patient

Trust

“confidence of whether a model will act as intended when facing a given problem.”

Common primary aim for XAI systems which audience is:

- Domain experts 👷👩‍🔧👩‍🔬
- Users *affected* by model decisions 👥

Arrieta, Alejandro Barredo, et al. "**Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI.**" Information Fusion 58 (2020): 82-115.

Understandability

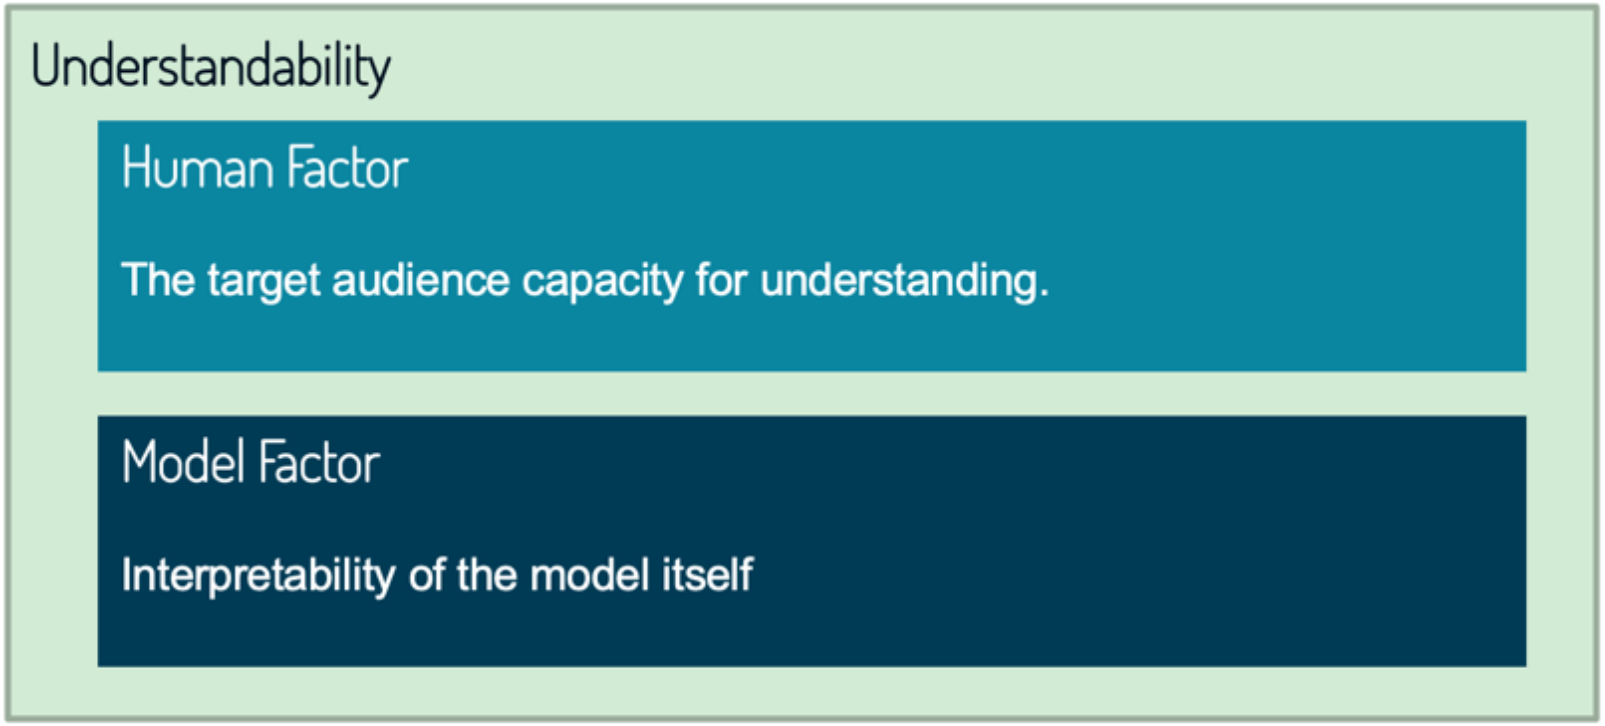

Arrieta, Alejandro Barredo, et al. "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI." Information Fusion 58 (2020): 82-115.

Zhou, Jianlong, et al. "Evaluating the quality of machine learning explanations: A survey on methods and metrics." Electronics 10.5 (2021): 593.

Actionable AI

Current problems with XAI methods:

- Non-technical consumers of AI, hardly understand the explainability provided by popular XAI methods designed for AI experts
- Non-technical consumers do not get the necessary guidance and insights to obtain their desired outcomes

So, the poor prediction is due to overfitting and skewed inter-quartile values

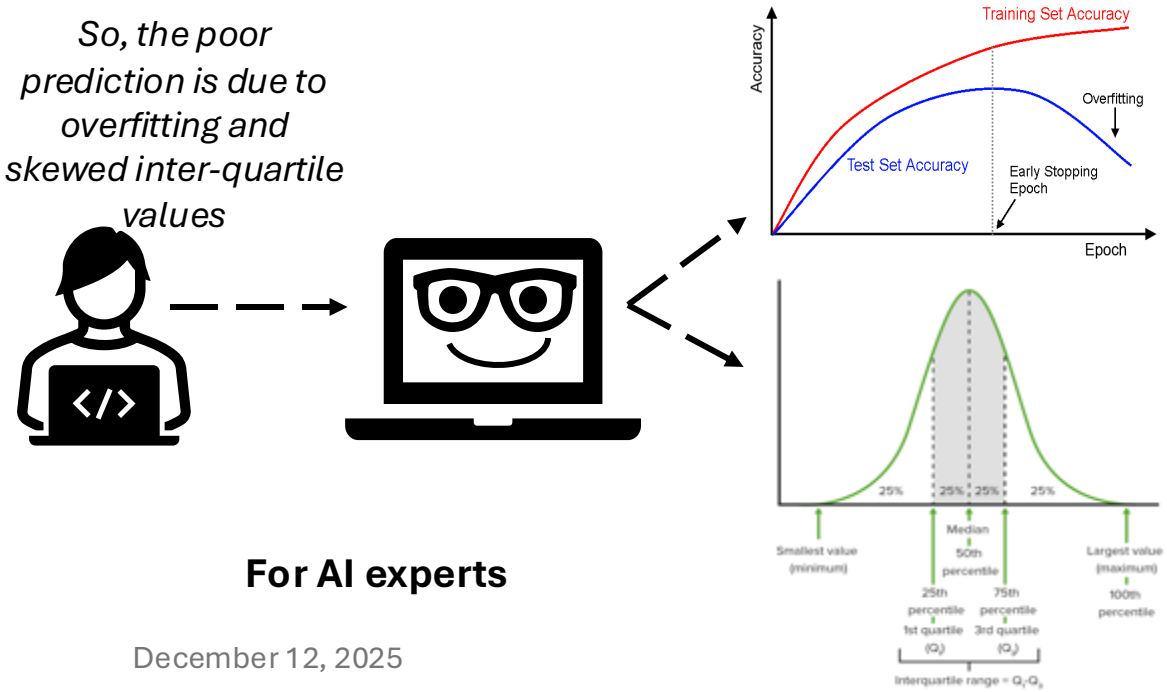

For AI experts

I don't really understand what it means! How can I get the correct prediction?

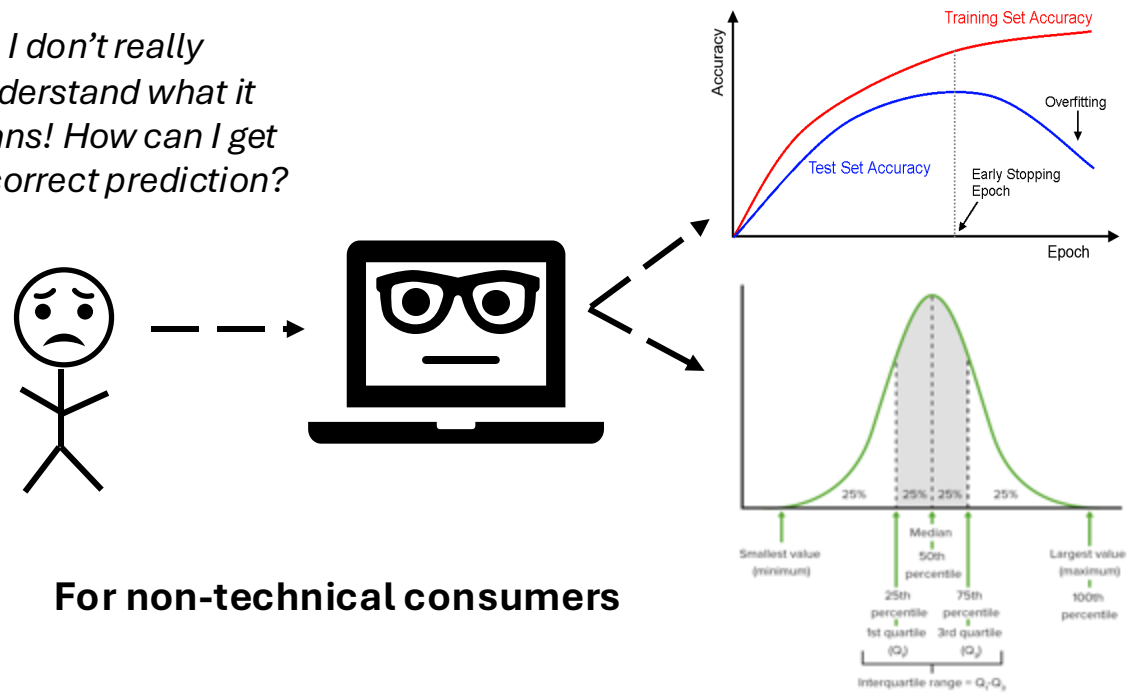

For non-technical consumers

Why Human-Centered XAI?

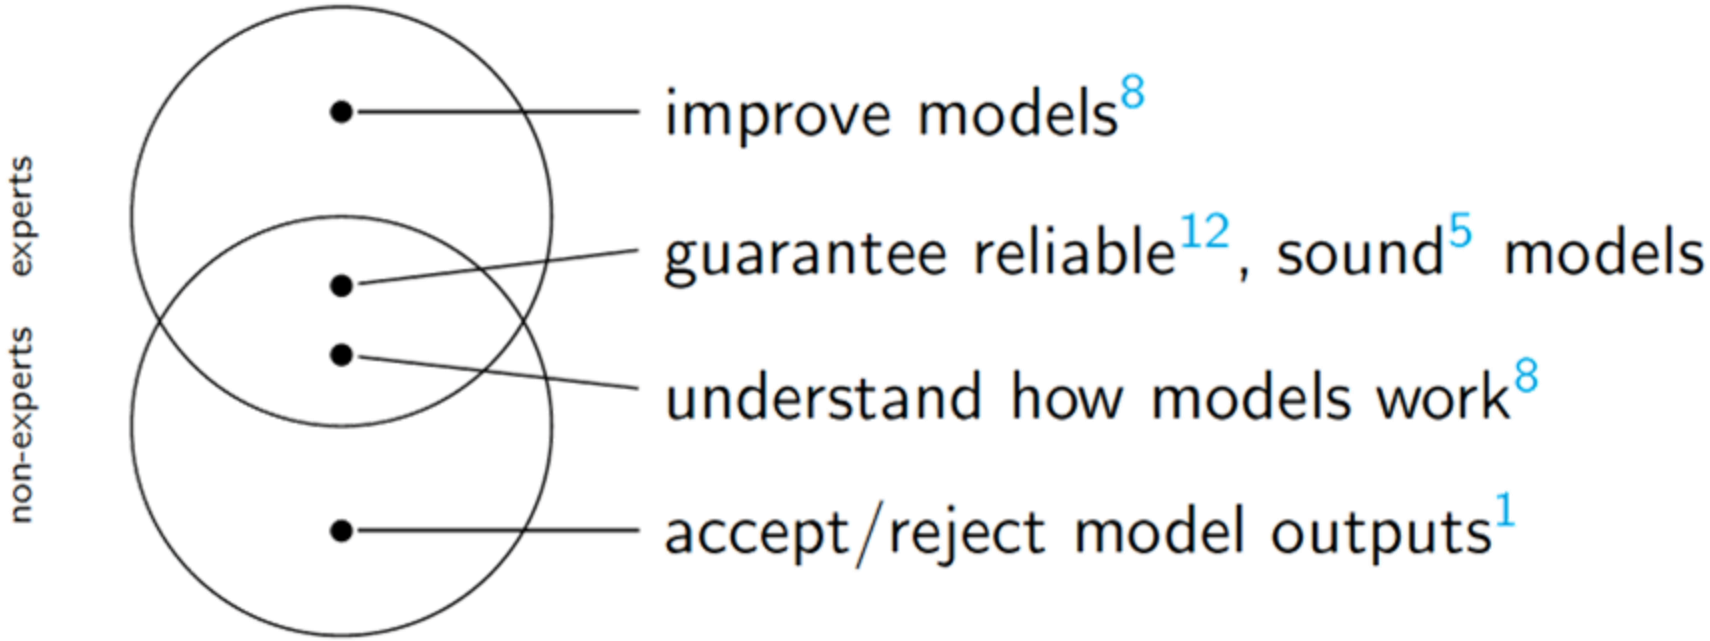

Domain experts mainly adopt AI if they can explain and are convinced by outputs. Explanations assess user trust.

Users don't use models/outputs they don't trust.

Explainability creates conditions in which technical, business, and risk professionals get the most value from AI systems.

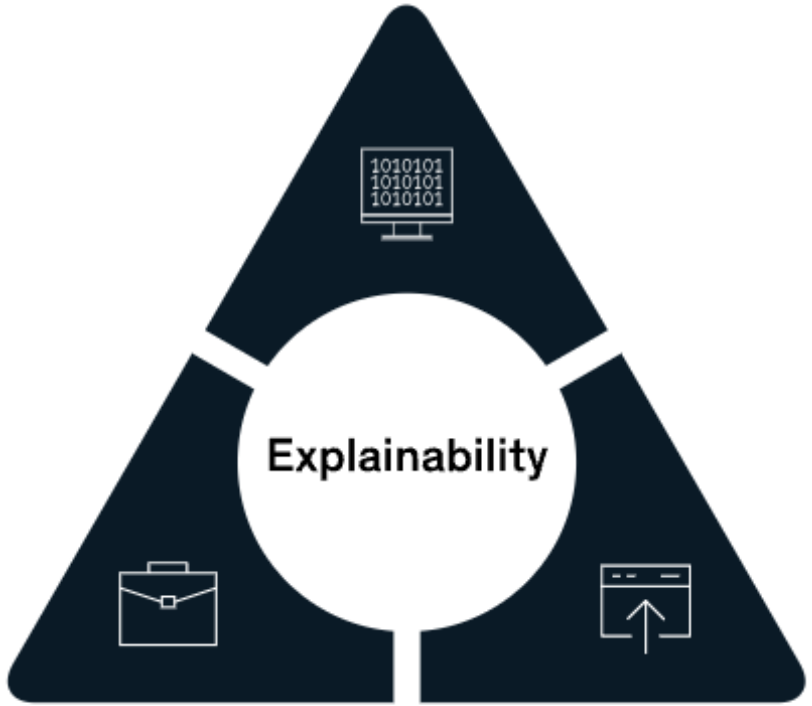

Technologists

1. More efficiently monitor, maintain, and improve AI systems

Business professionals

2. Trust AI outputs, so they increasingly adopt AI tools
3. Apply knowledge about the why of an AI prediction or recommendation to identify effective interventions
4. Assess whether AI applications meet business objectives

Legal and risk professionals

5. See whether technology and associated workflows comply with applicable regulations and are in line with customer expectations

McKinsey & Company

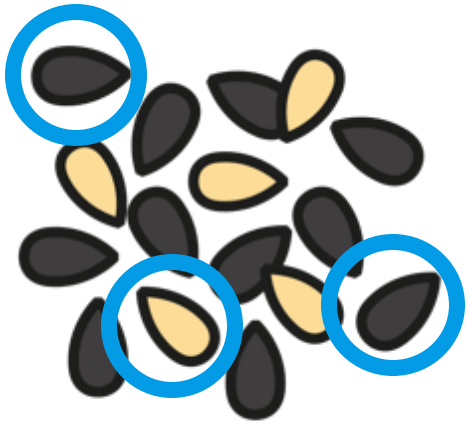

All icons in this presentation originate from flaticon.com

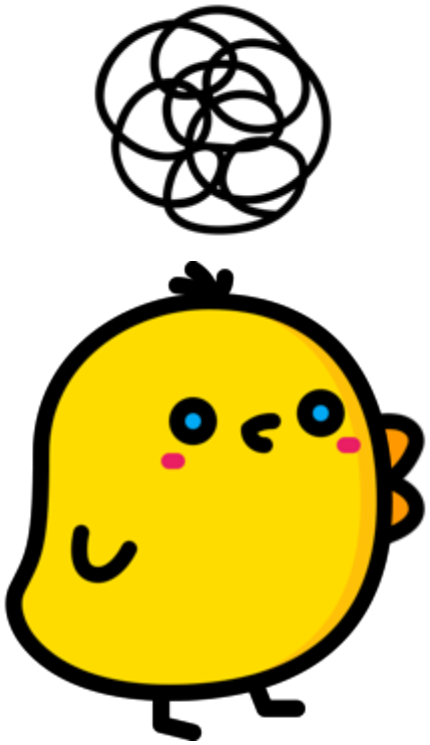

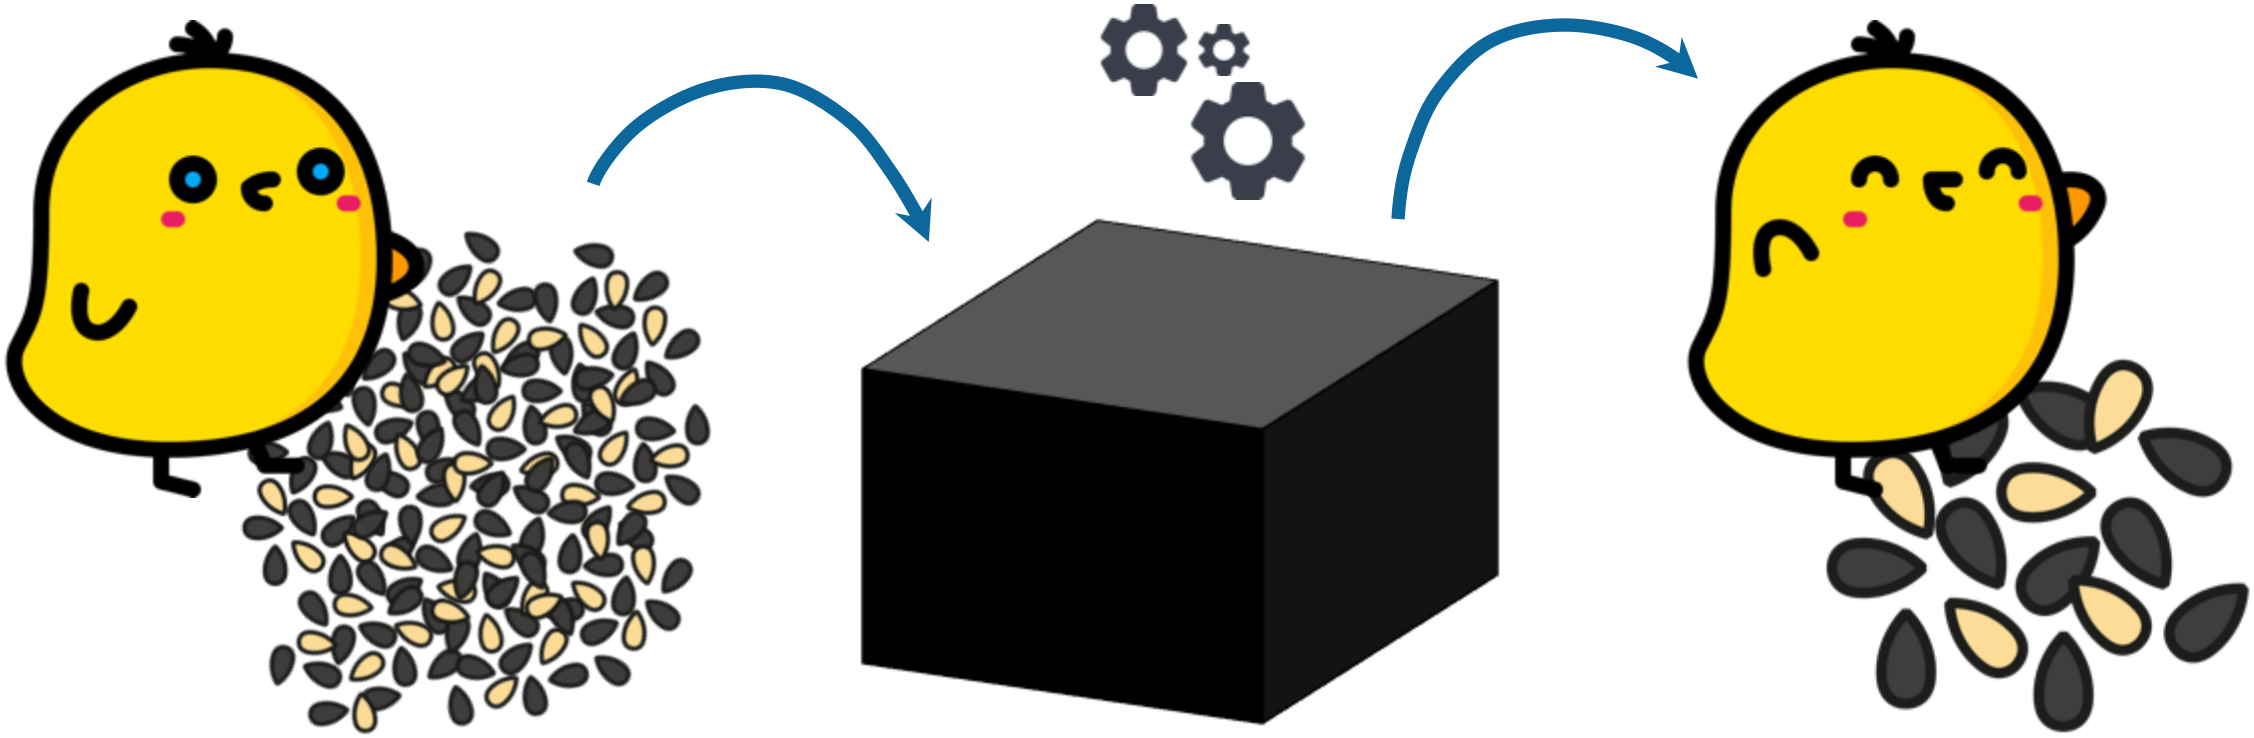

Choice overload

Recommender system

Recommendation

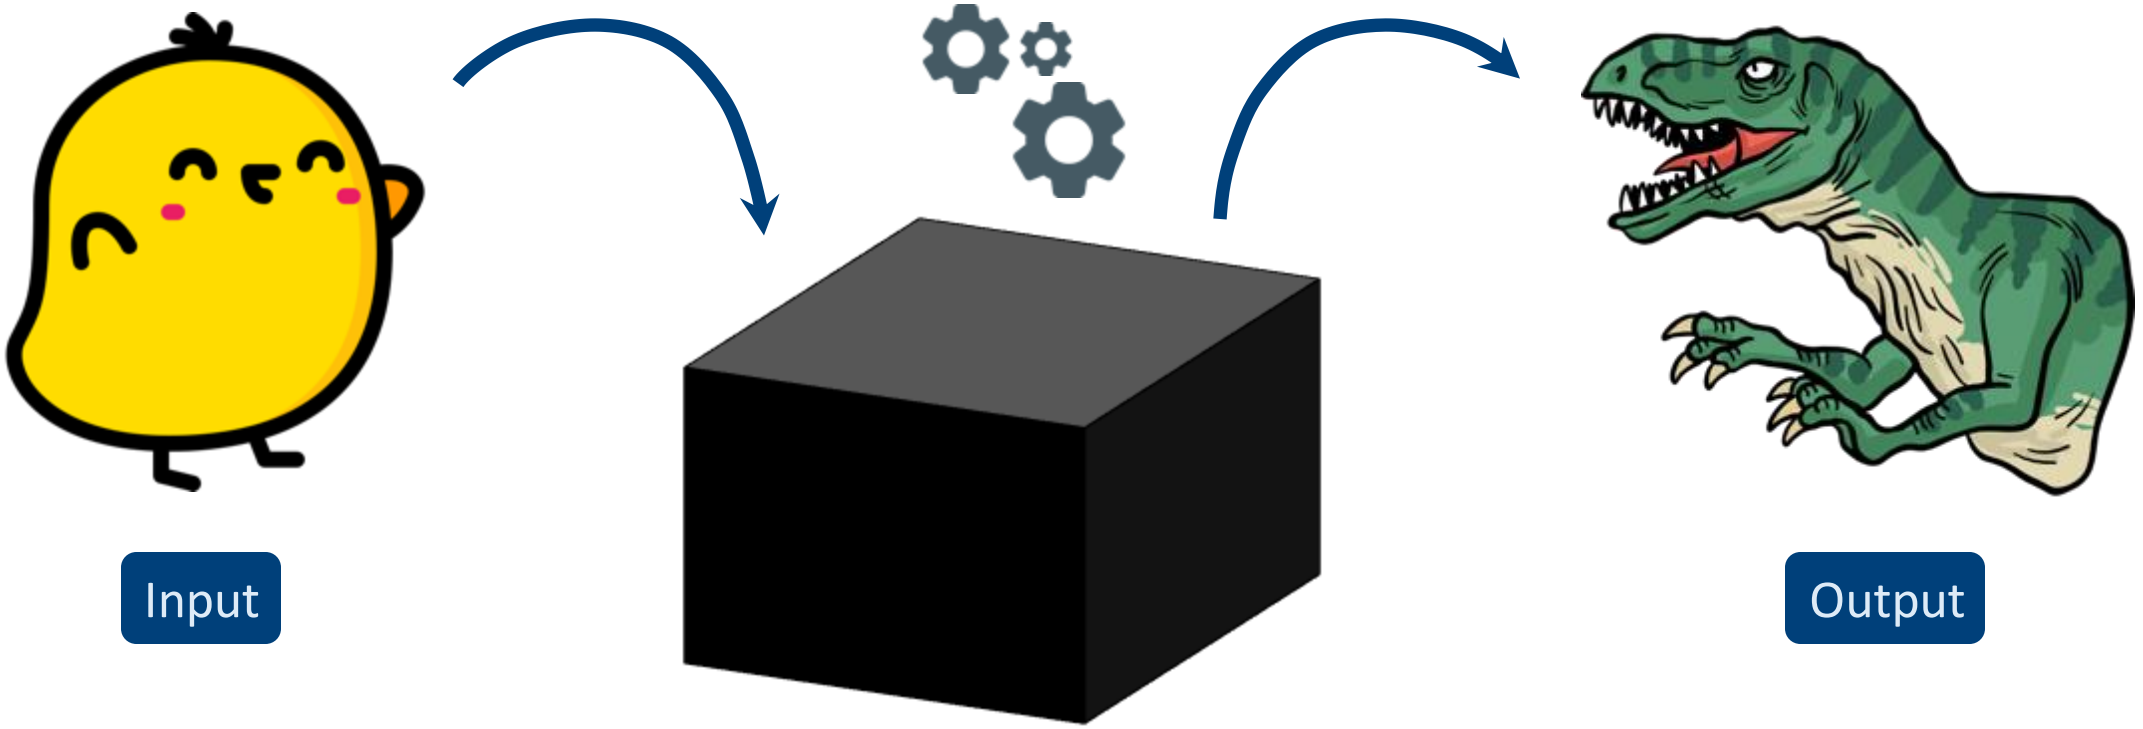

All icons in this presentation originate from flaticon.com

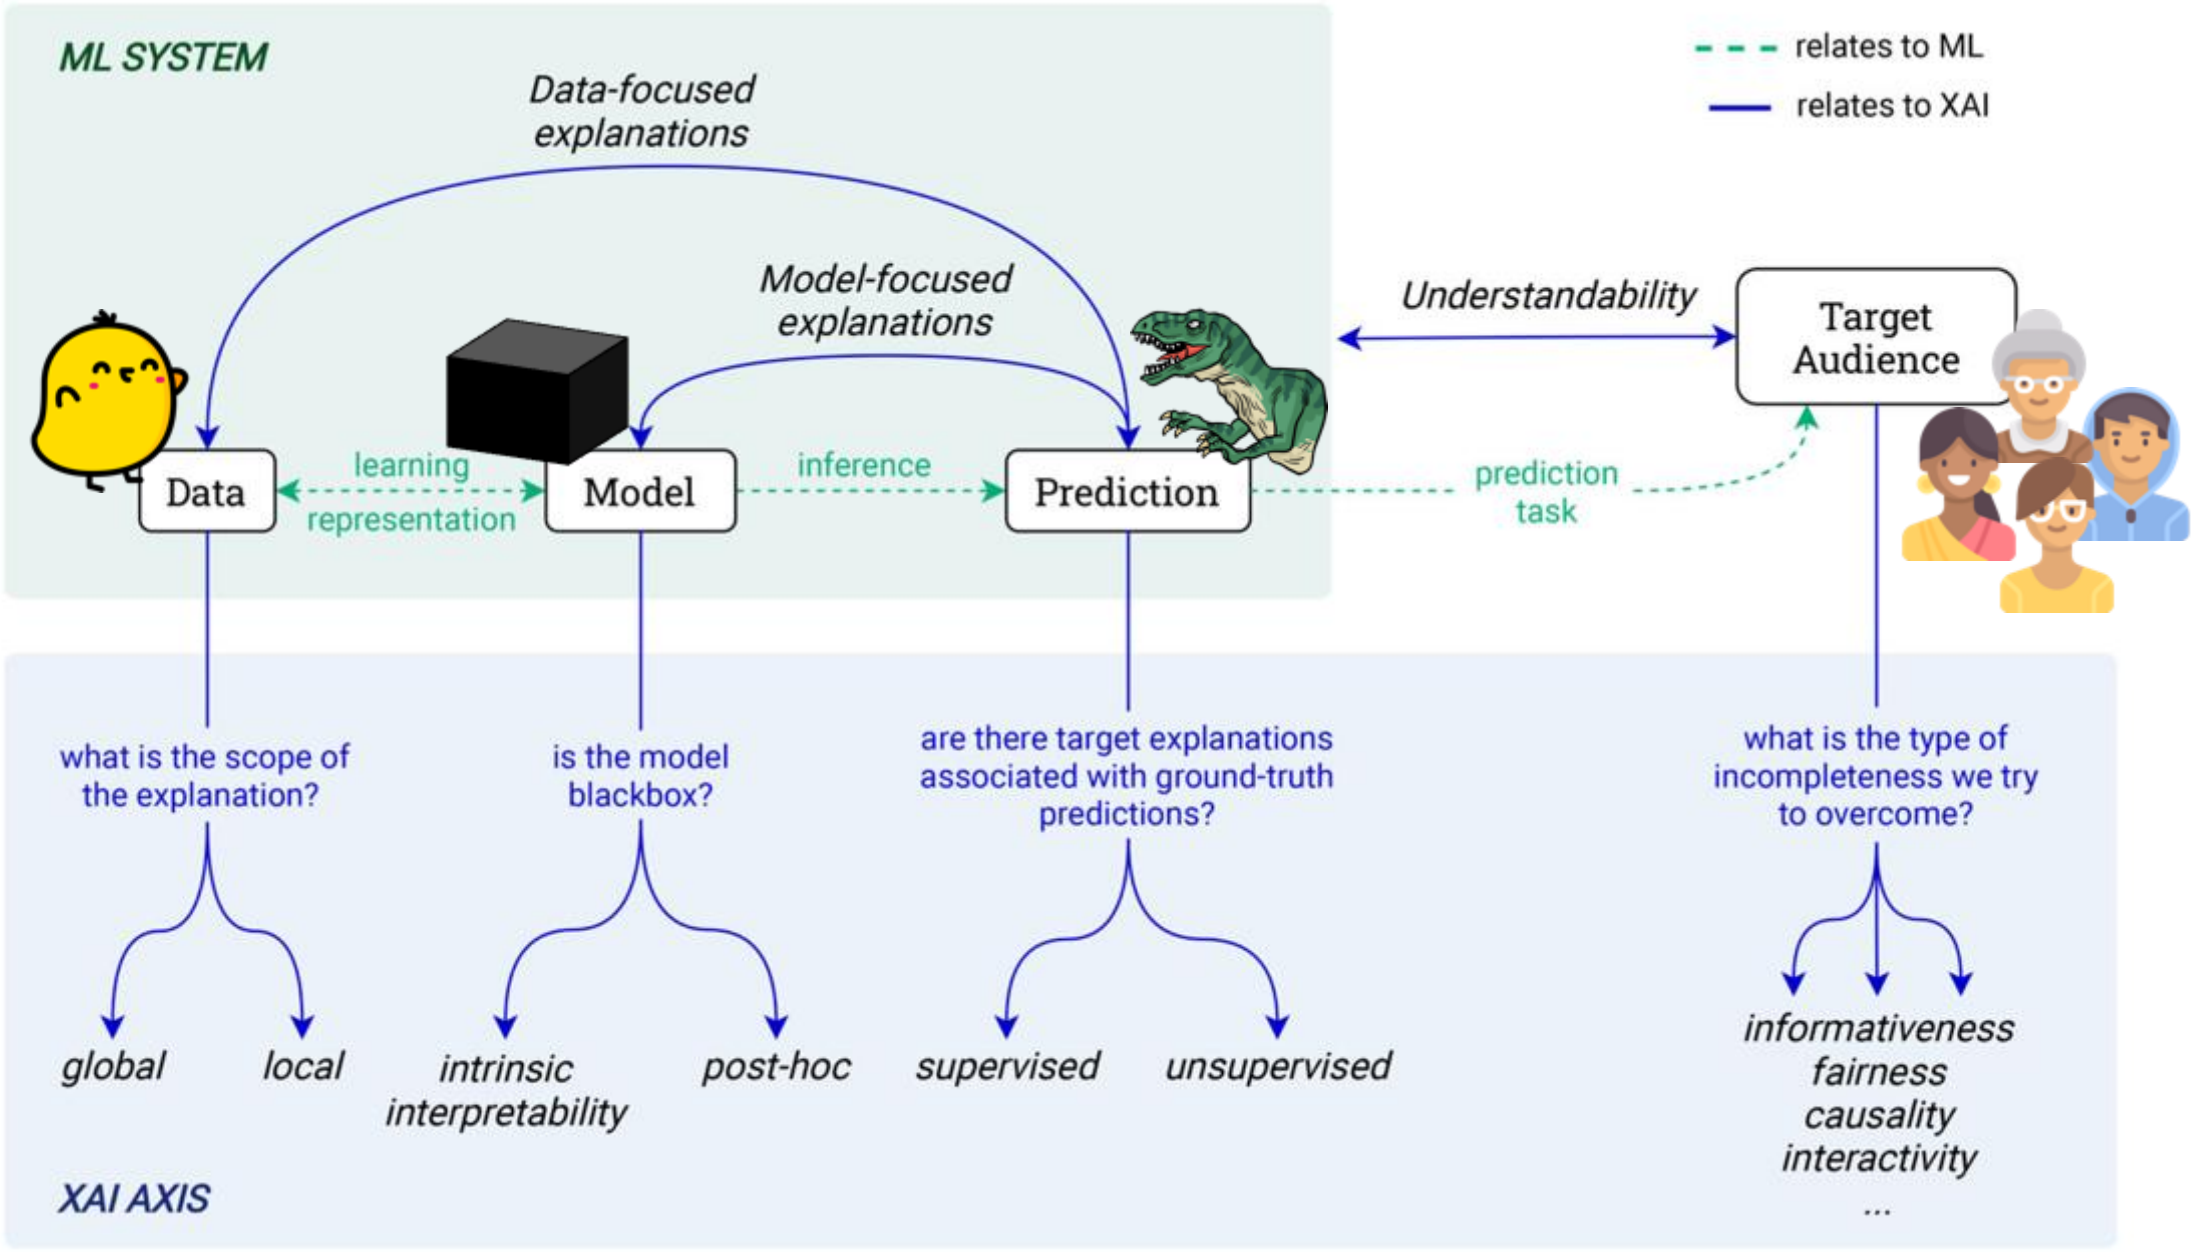

Darius Afchar, Alessandro Melchiorre, Markus Schedl, Romain Hennequin, Elena Epure, and Manuel Moussallam. 2022. Explainability in Music Recommender Systems. *AI Magazine* 43, 2: 190–208. https://doi.org/10.1002/aaai.12056

Algorithmic XAI approaches

Human-centred XAI approaches

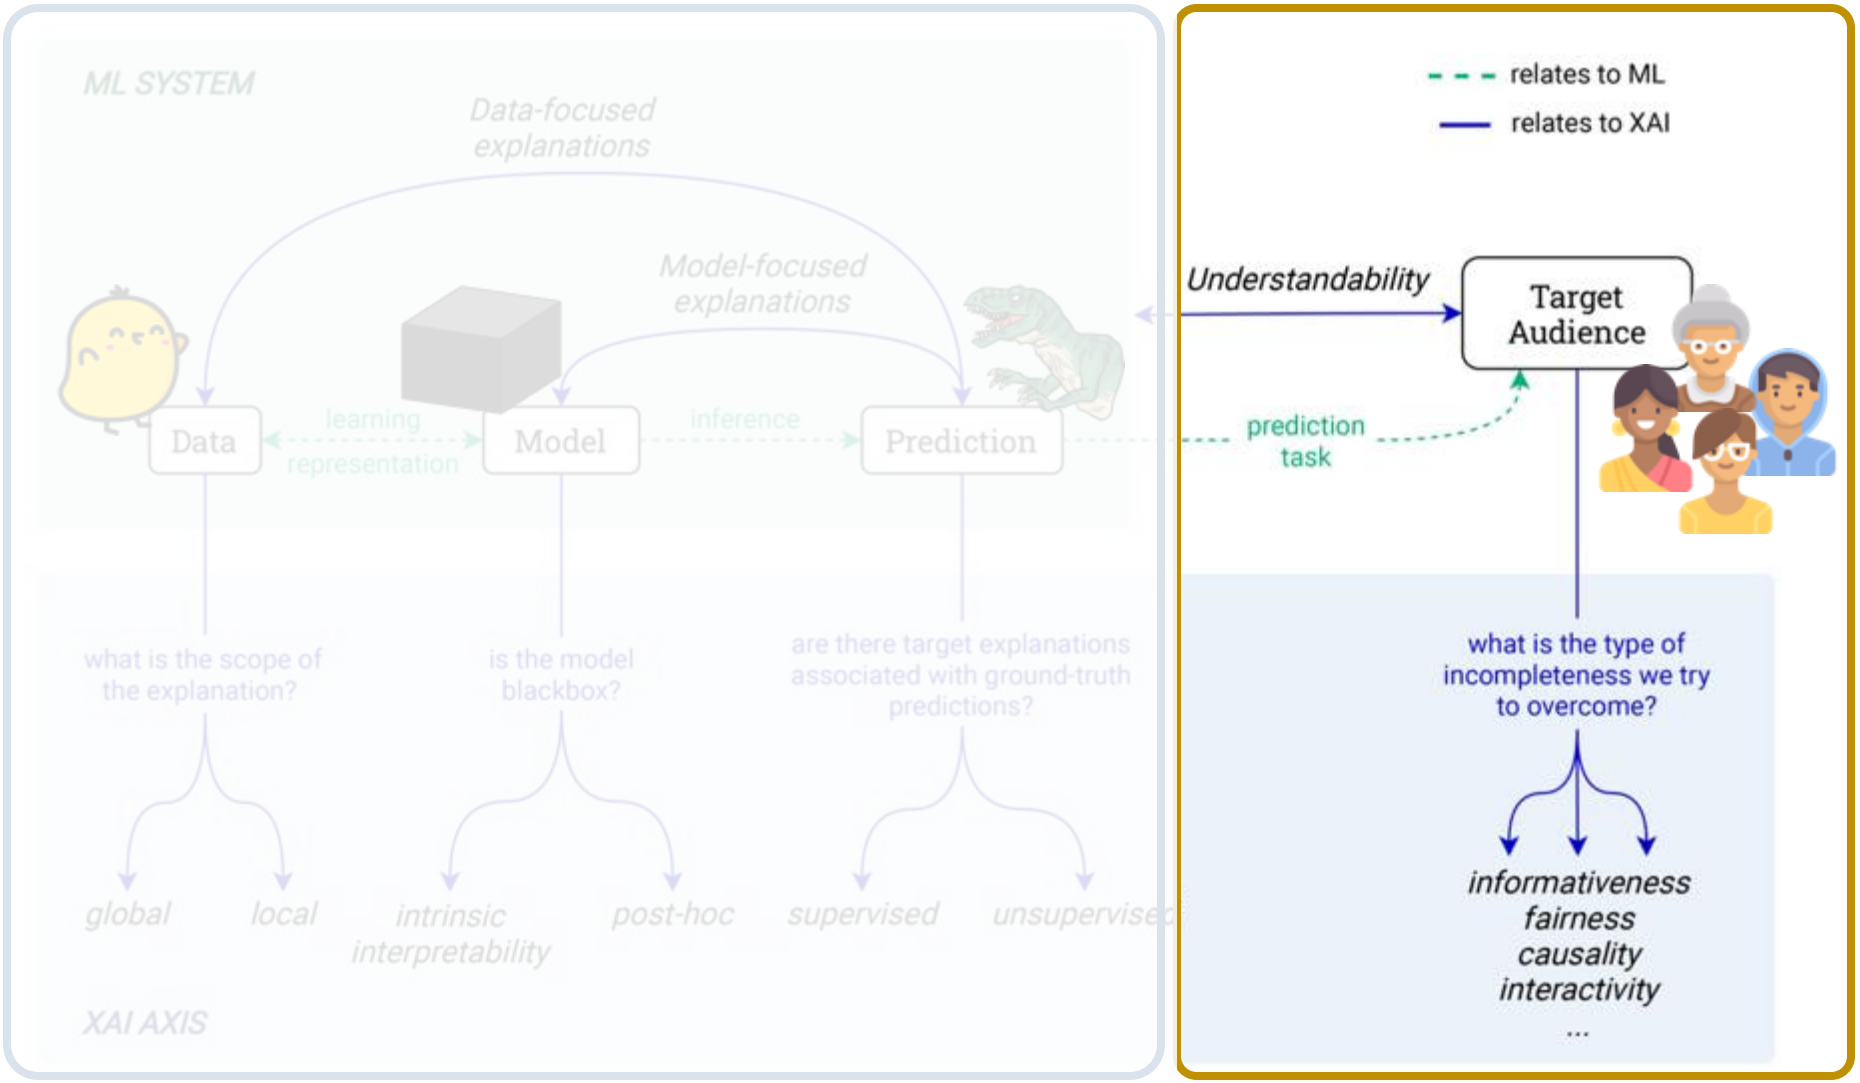

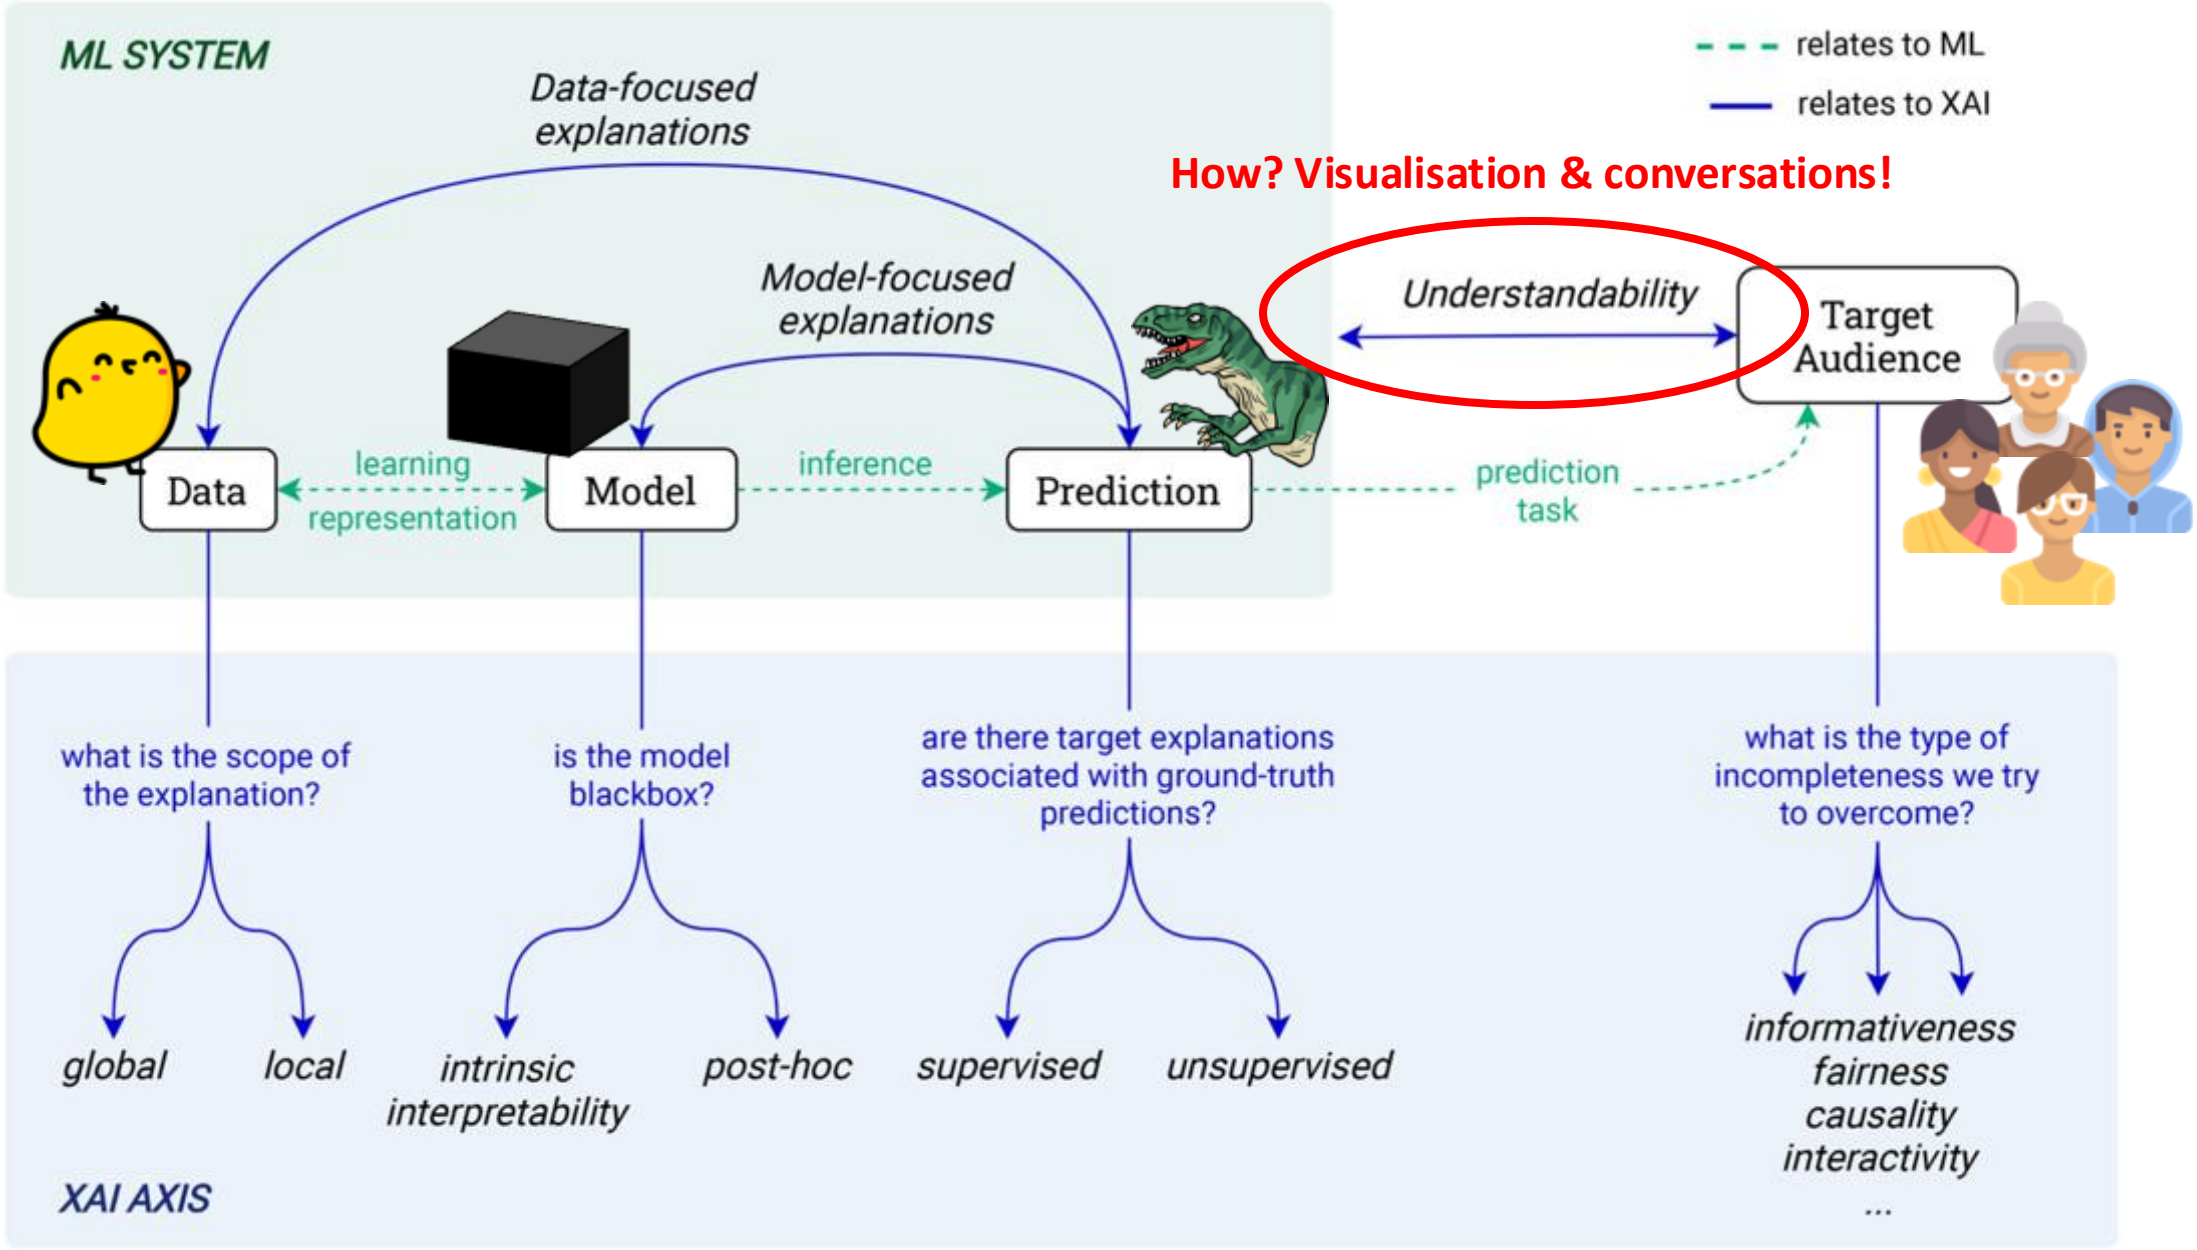
ML SYSTEM
Data-focused explanations
Model-focused explanations
relates to ML
relates to XAI
How? Visualisation & conversations!
Understandability
Target Audience
Data
learning representation
Model
inference
Prediction
prediction task
what is the scope of the explanation?
is the model blackbox?
are there target explanations associated with ground-truth predictions?
what is the type of incompleteness we try to overcome?
global
local
intrinsic interpretability
post-hoc
supervised
unsupervised
informativeness
fairness
causality
interactivity
...
XAI AXIS


Personal Health Empowerment
Empowering people to monitor and improve their health using personal data and digital coaching
PROGRAM DETAILS
Personal Health Empowerment
Secondary use of data, including data from My OptiDesk
Data Analytics
Coaching Programmes
Coaching editor
Personal Health Data
Intelligent Coaching Engine
Health Professional
Electronic Health Registers
Coaching app
Employee
Coaching objective: muscoloskeletal pain

idewe

Explaining recommendations

Word cloud

Feature importance

Feature importance+ %

Maxwell Szymanski, Vero Vanden Abeele and Katrien Verbert *Explaining health recommendations to lay users: The dos and don'ts* – Apex-IUI 2022

Textual or visual?

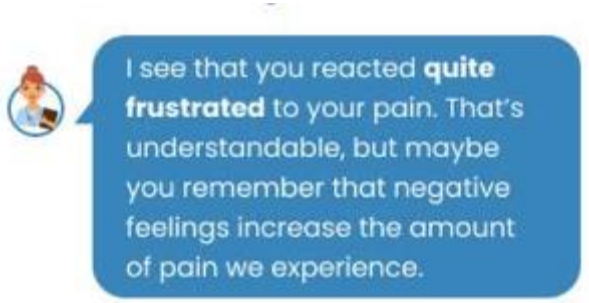

Dashboard for health professionals

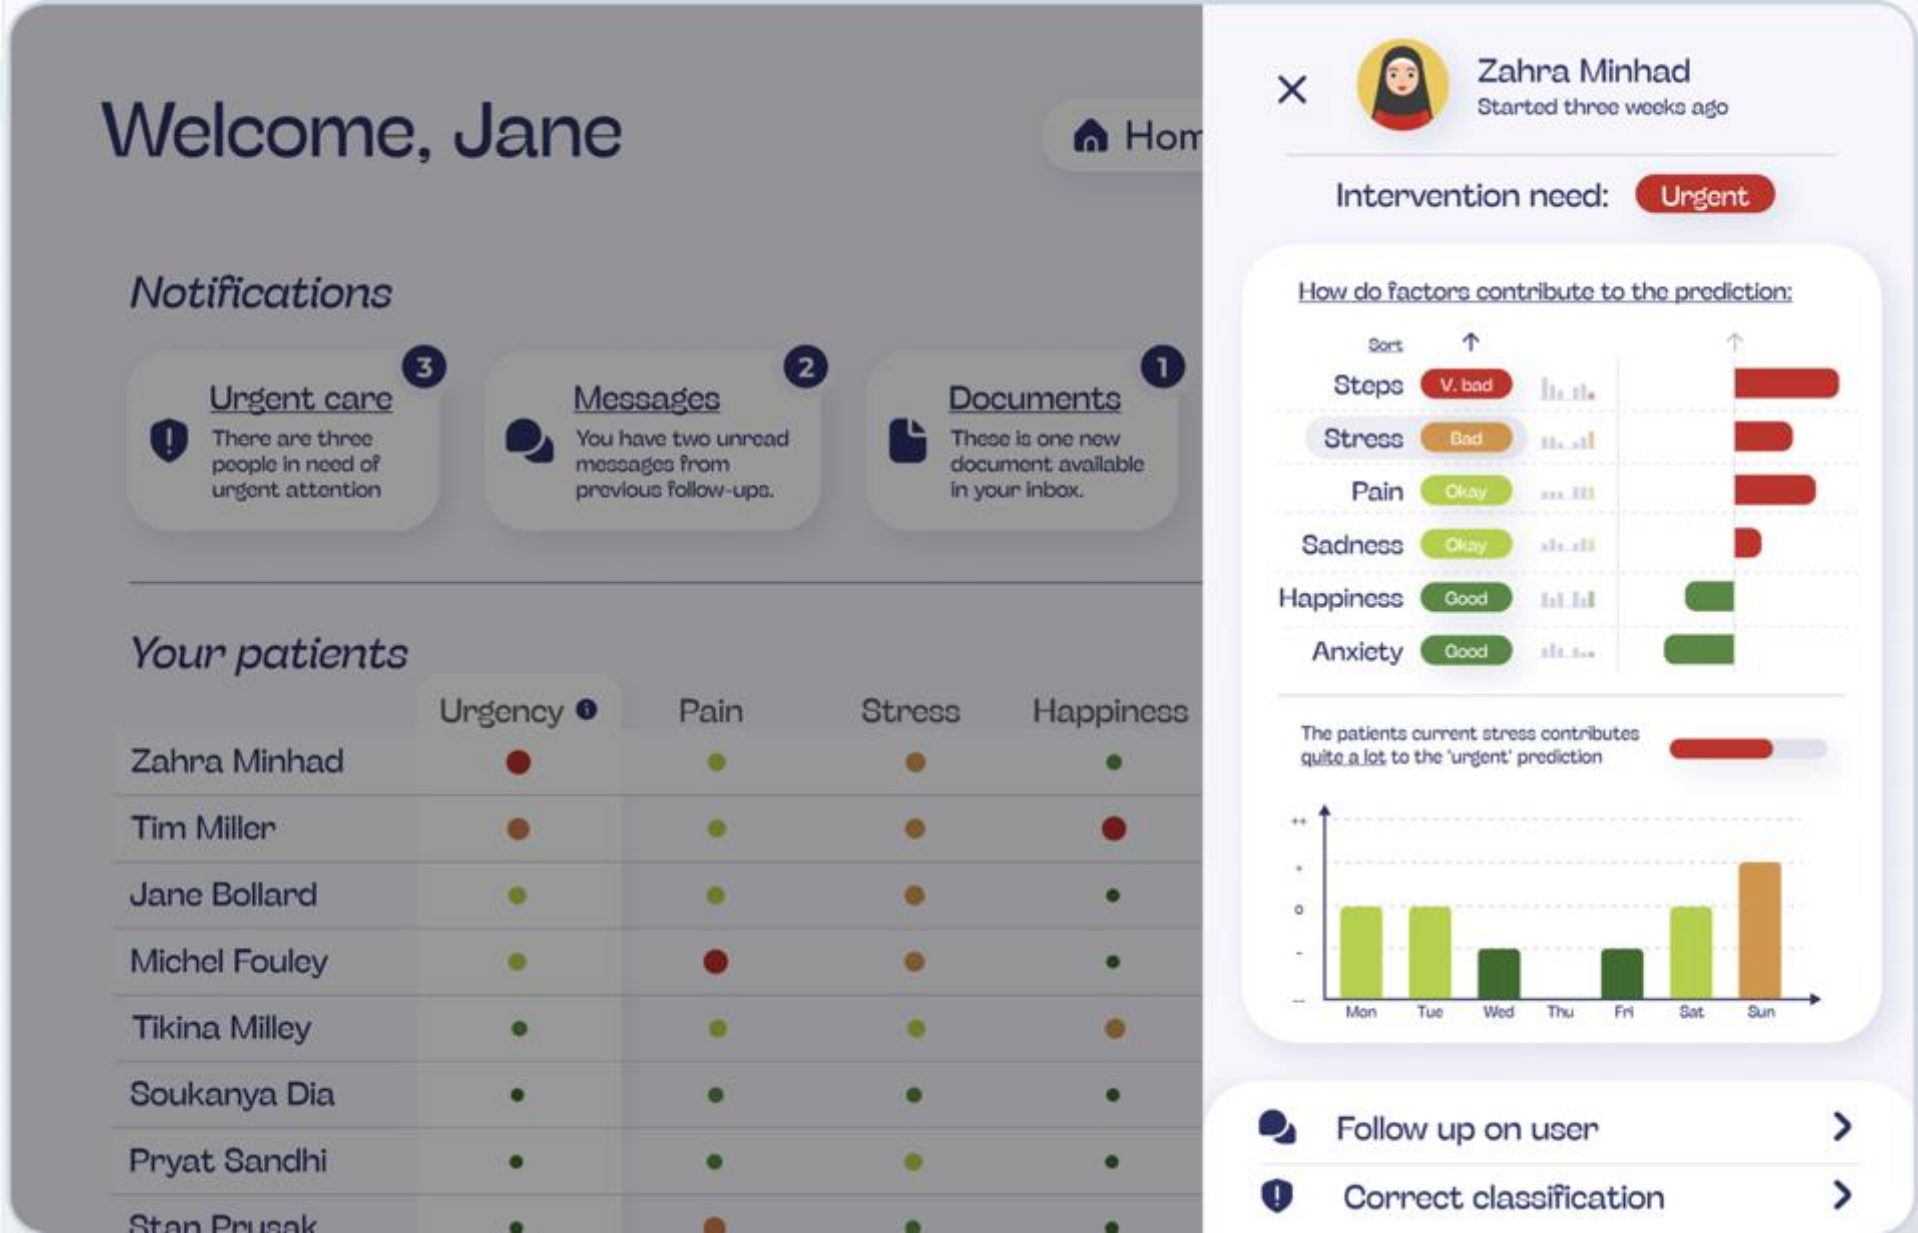

Simplified feature importance

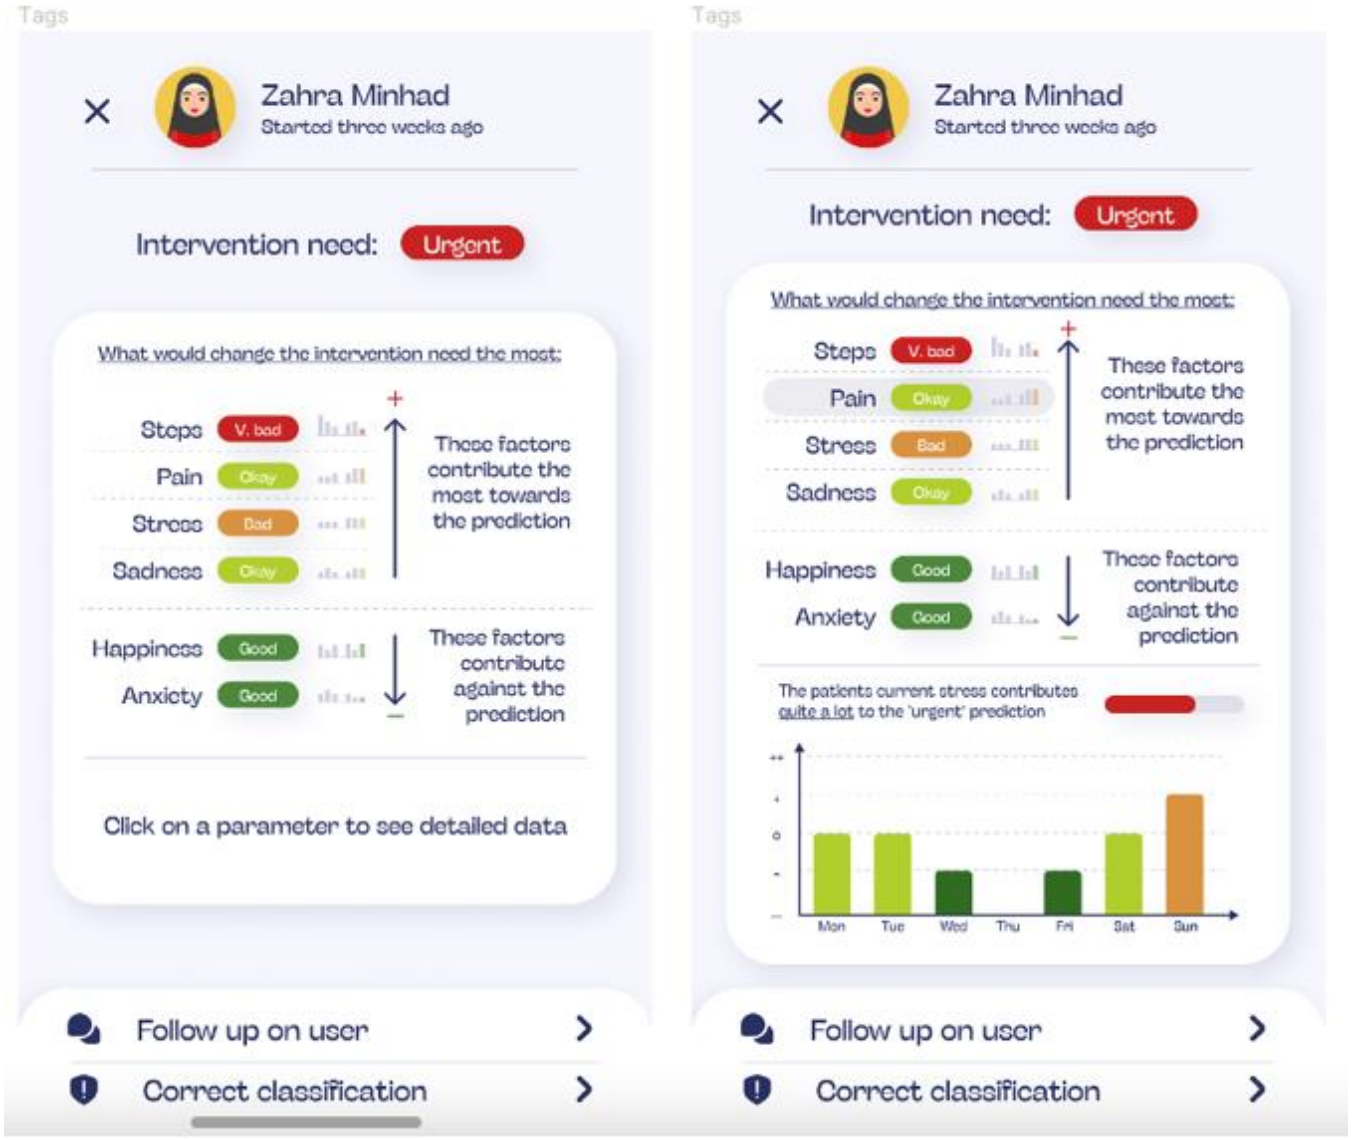

Data-centric explanations

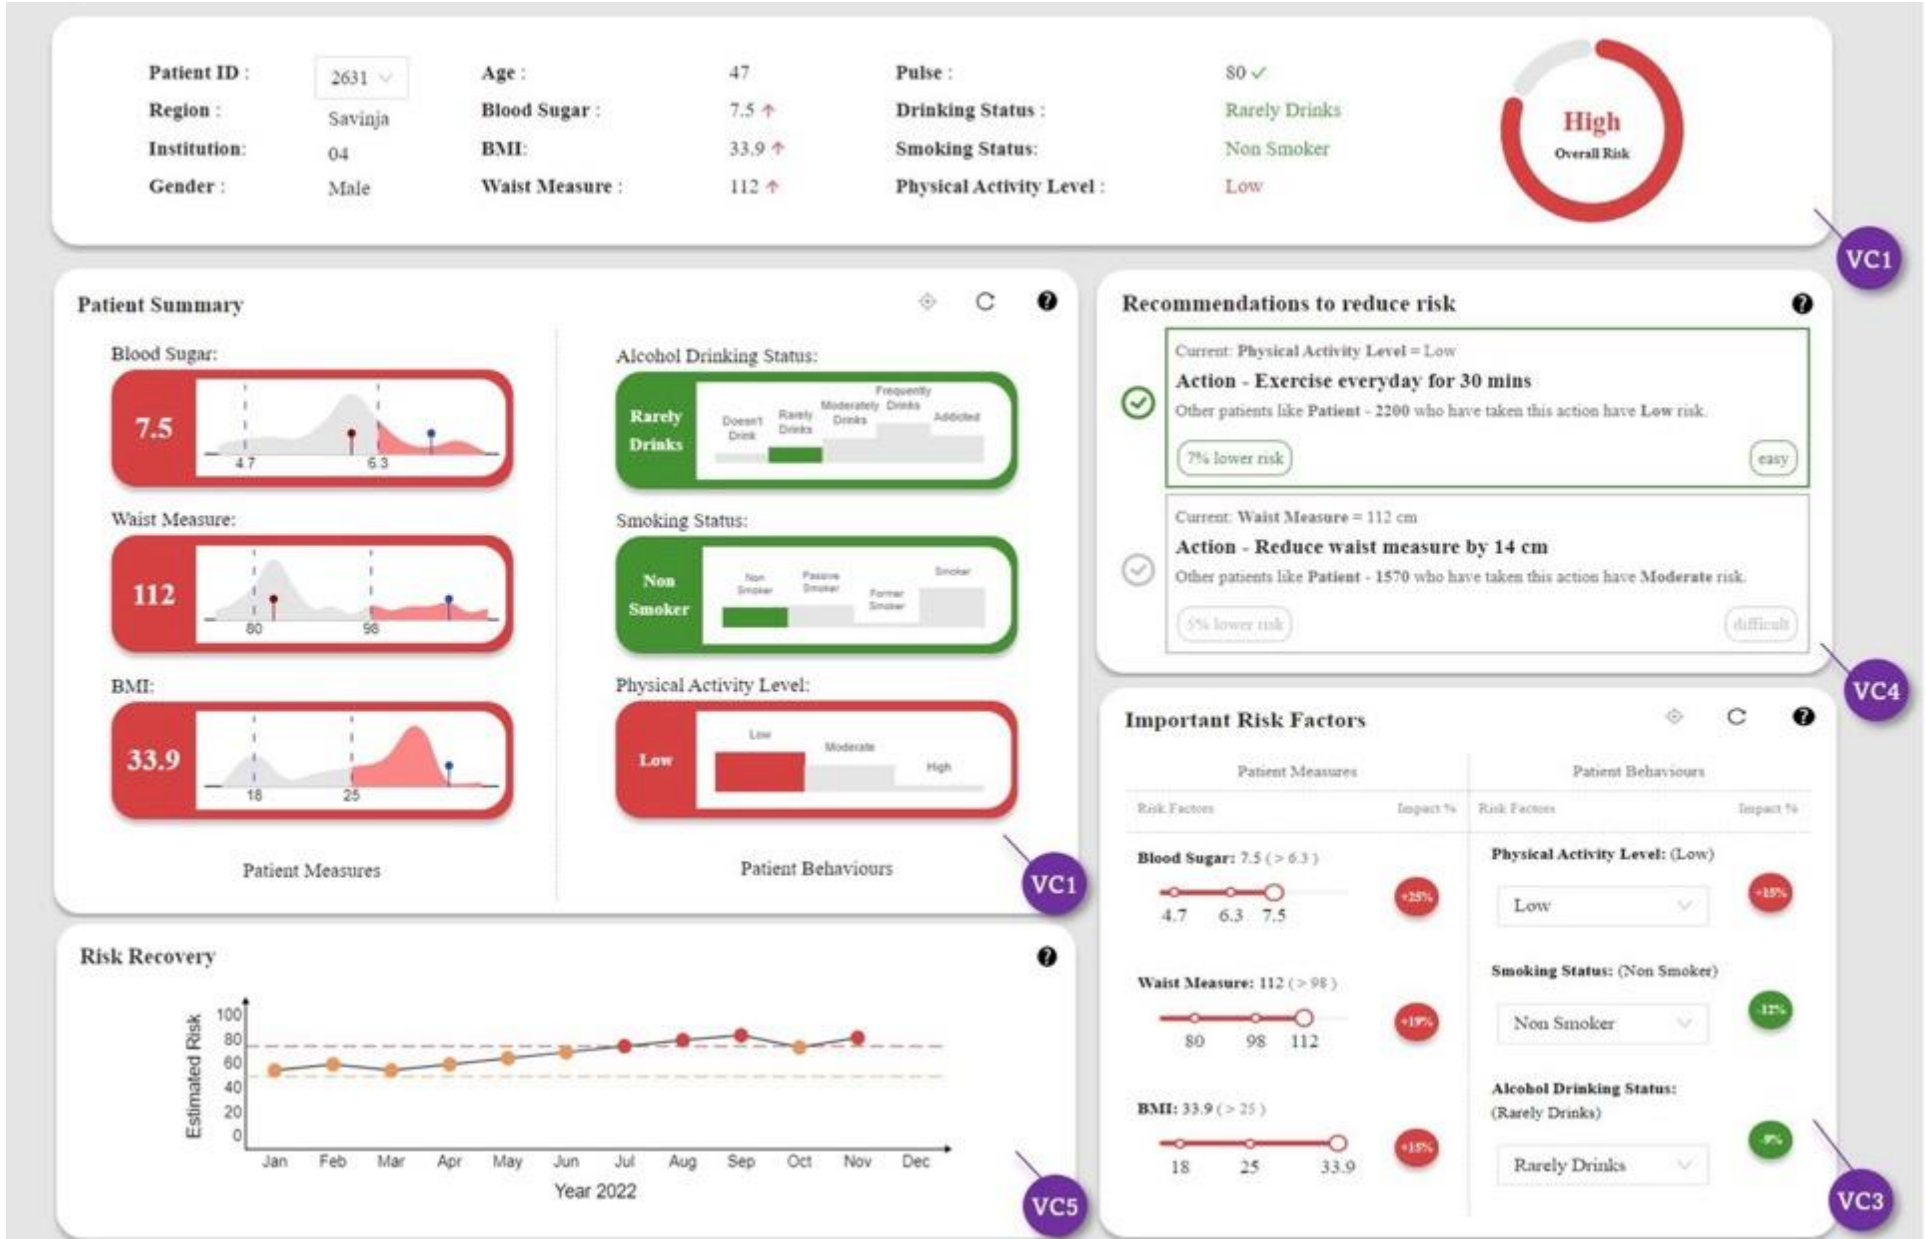

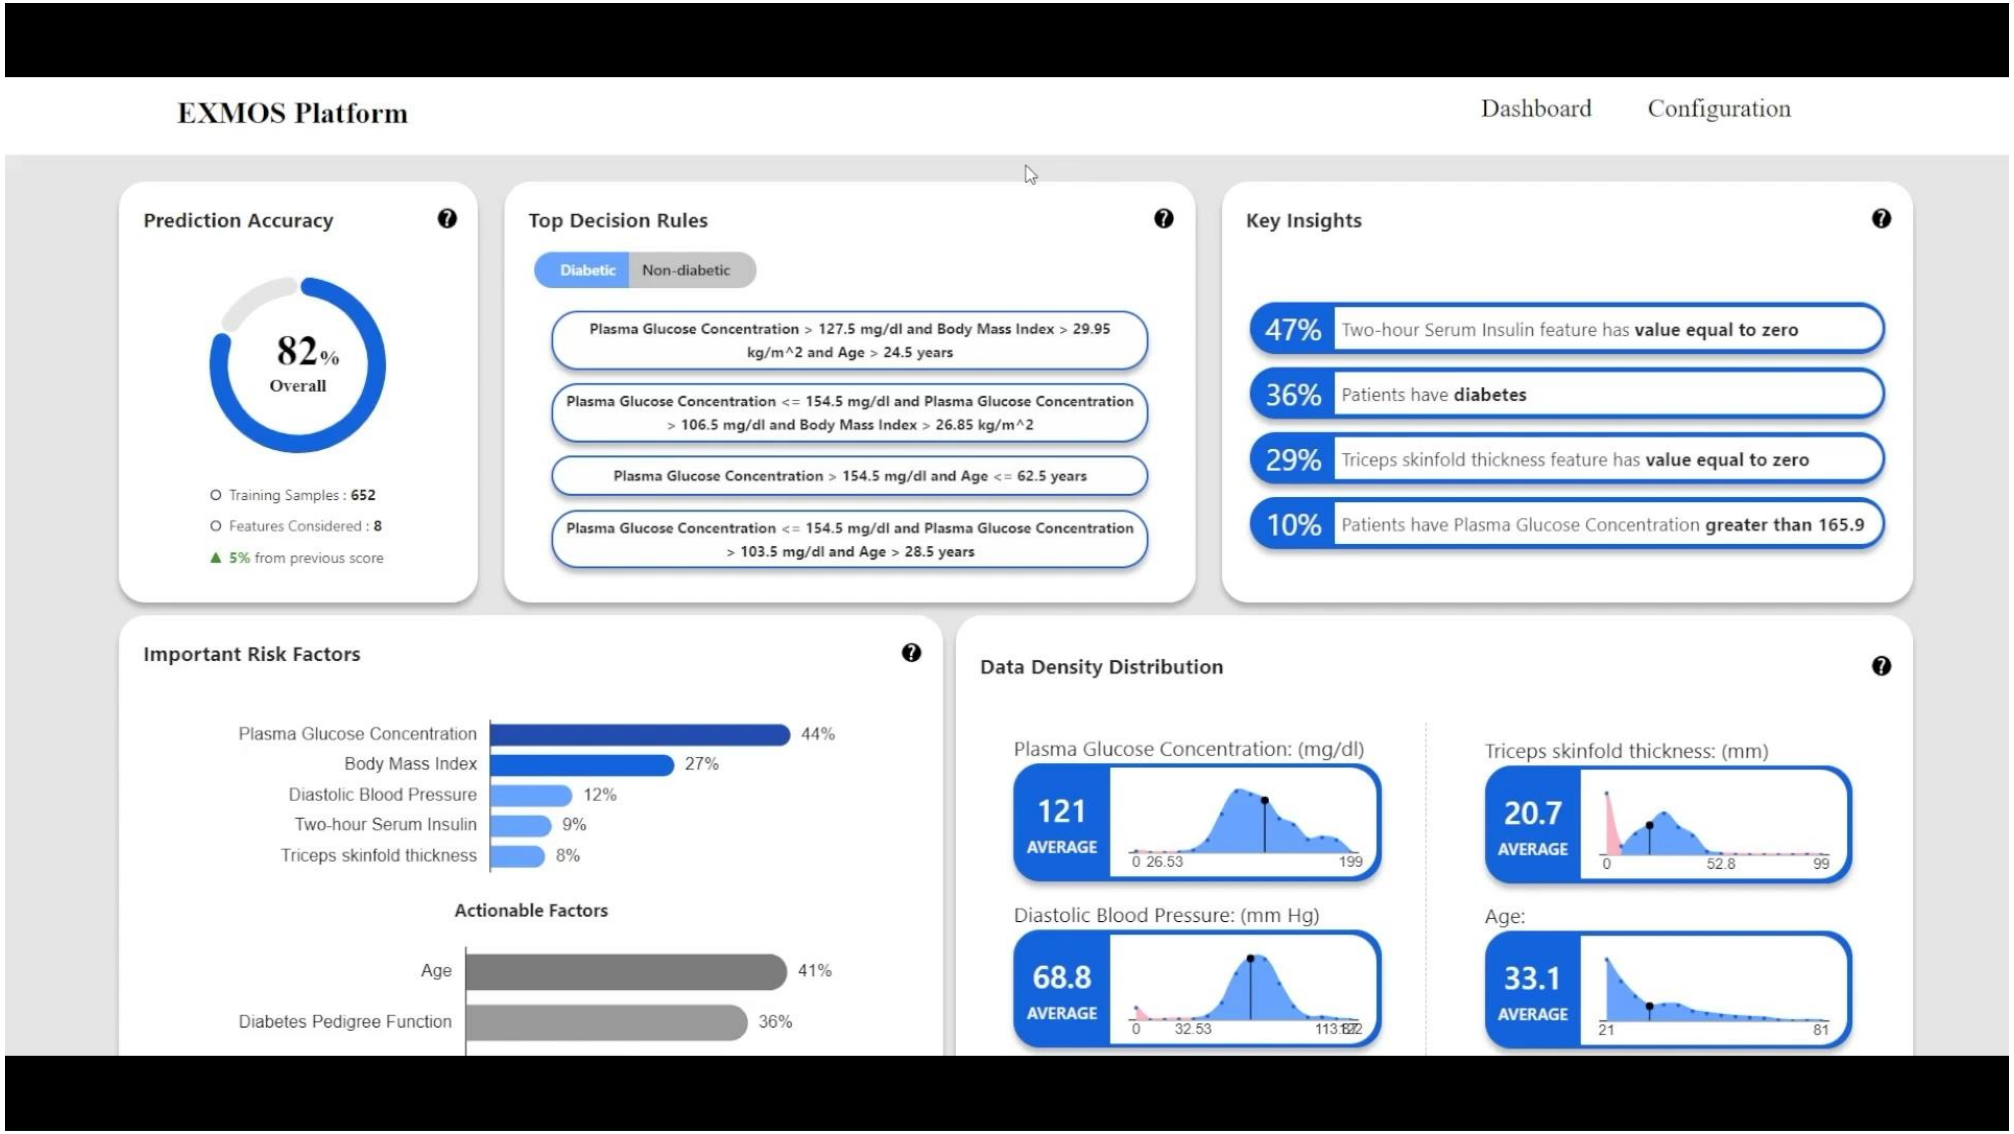

Aditya Bhattacharya, Simone Stumpf, Robin De Croon, and Katrien Verbert. 2025. Explanatory Debiasing: Involving Domain Experts in the Data Generation Process to Mitigate Representation Bias in AI Systems. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 1119, 1–20. https://doi.org/10.1145/3706598.3713497

Design science research

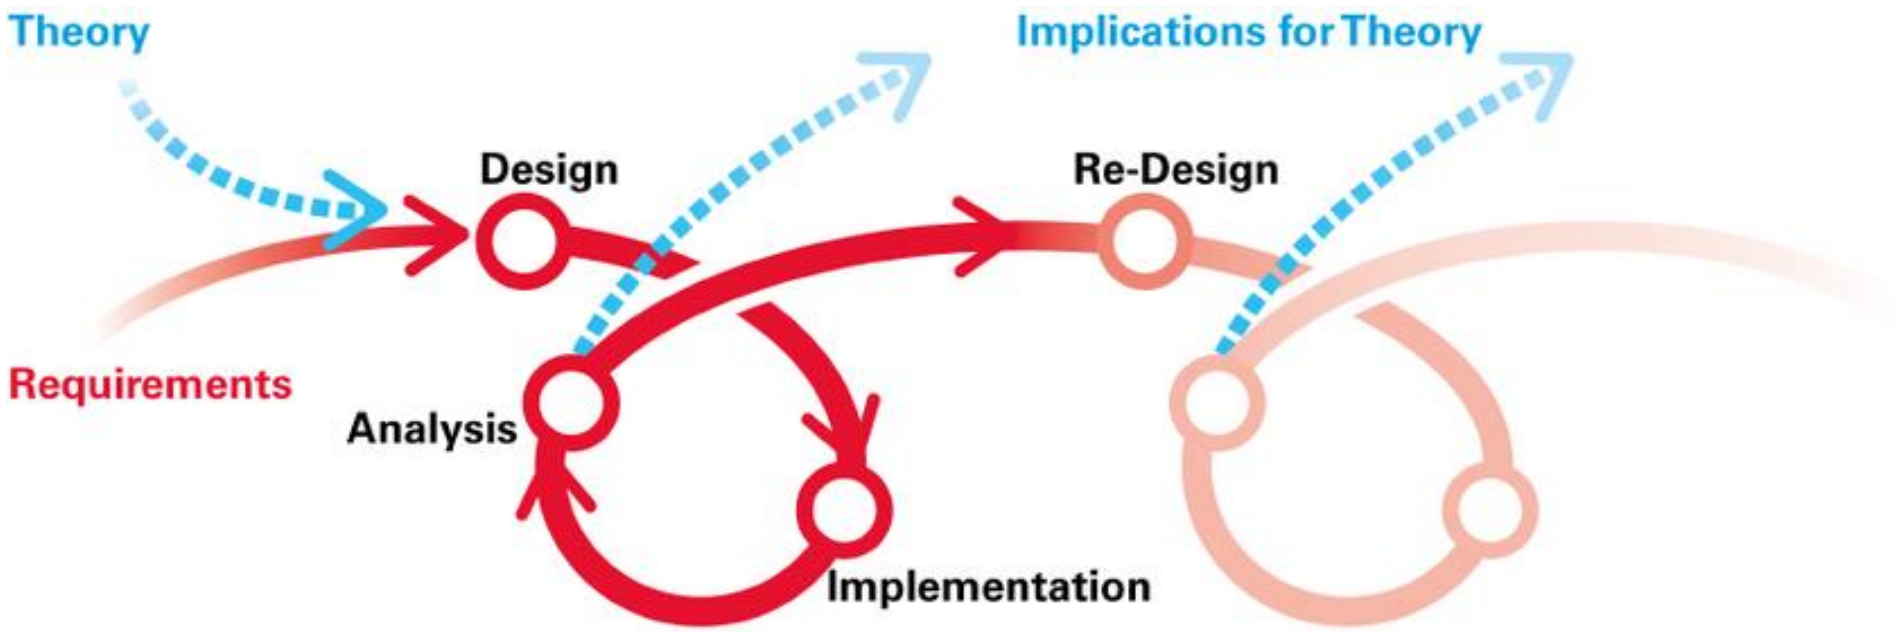

Fraefel, U. (2014, November). Professionalization of pre-service teachers through university-school partnerships. In *Conference Proceedings of WERA Focal Meeting, Edinburgh*.

Explaining predictions

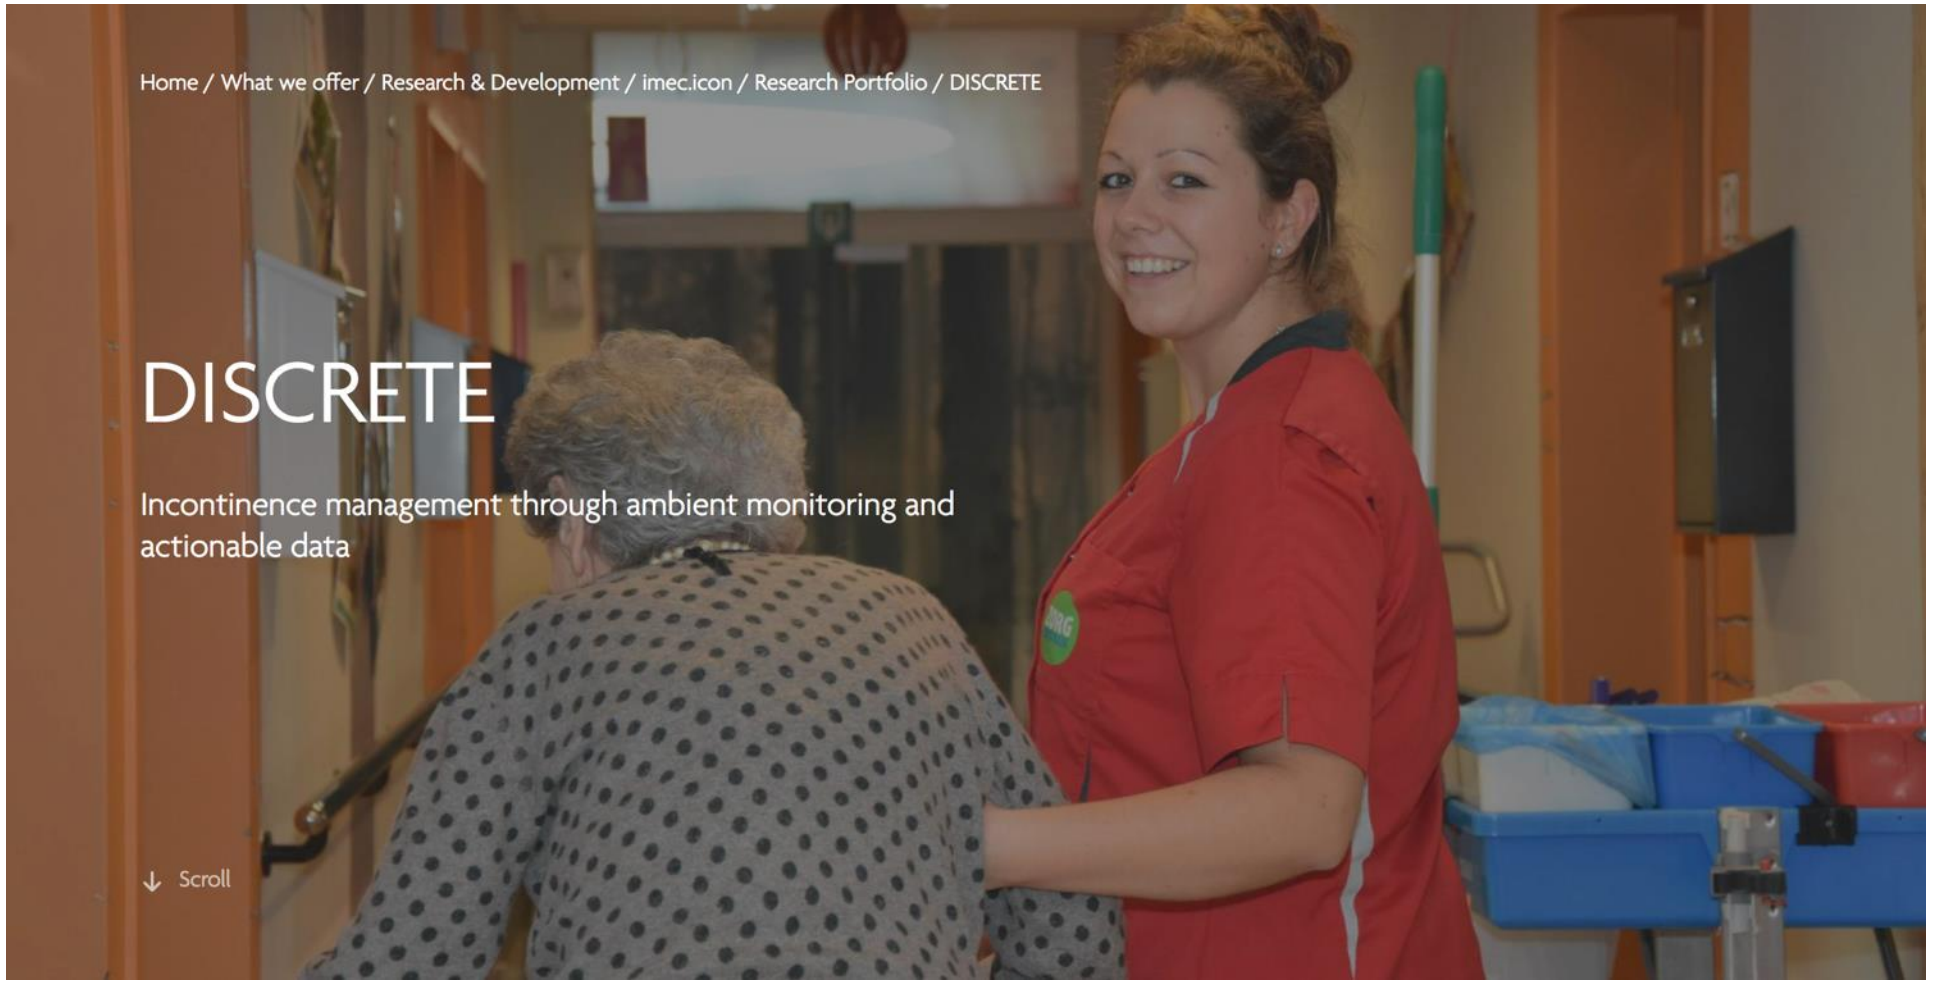

https://www.imec-int.com/en/what-we-offer/research-portfolio/discrete

CORILUS Connected Healthcare · St Bernardus woonzorghuis · Zorg Kortrijk · televic · DISTRACGROUP We care for people · imec

Explaining predictions

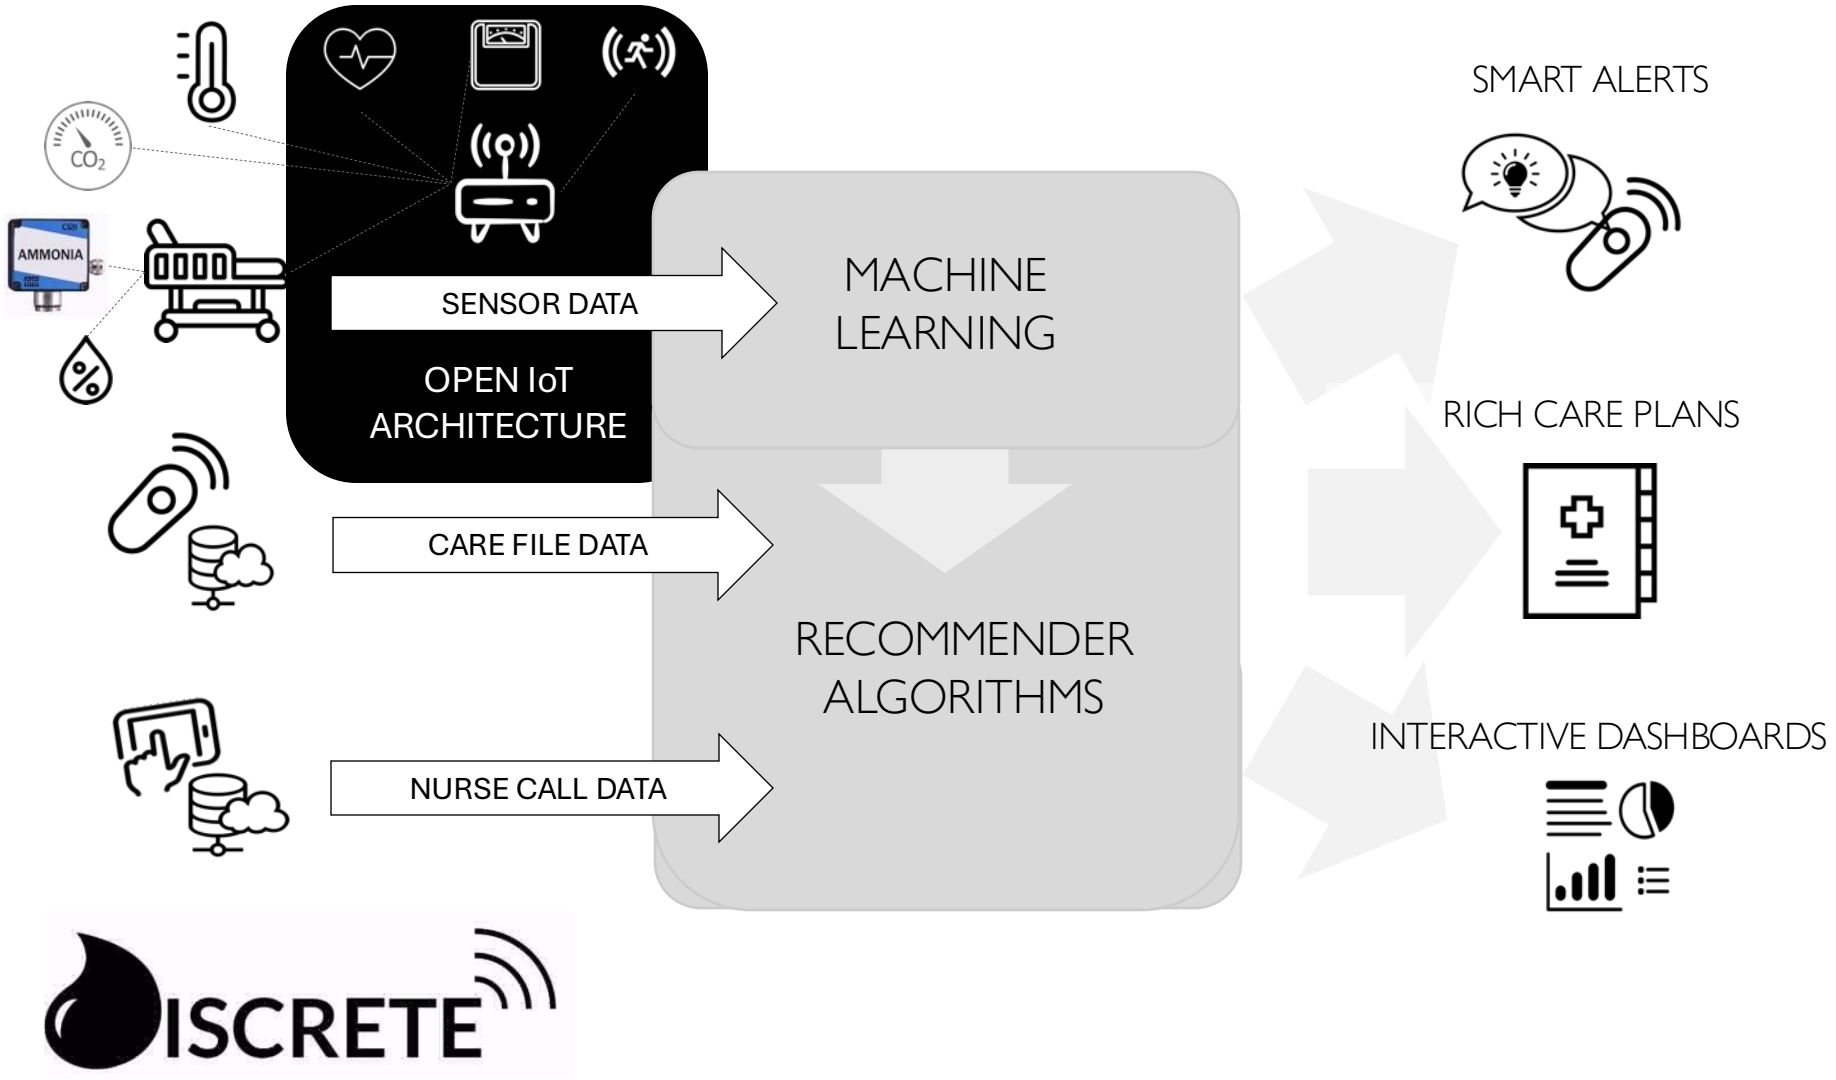

DISCRETE

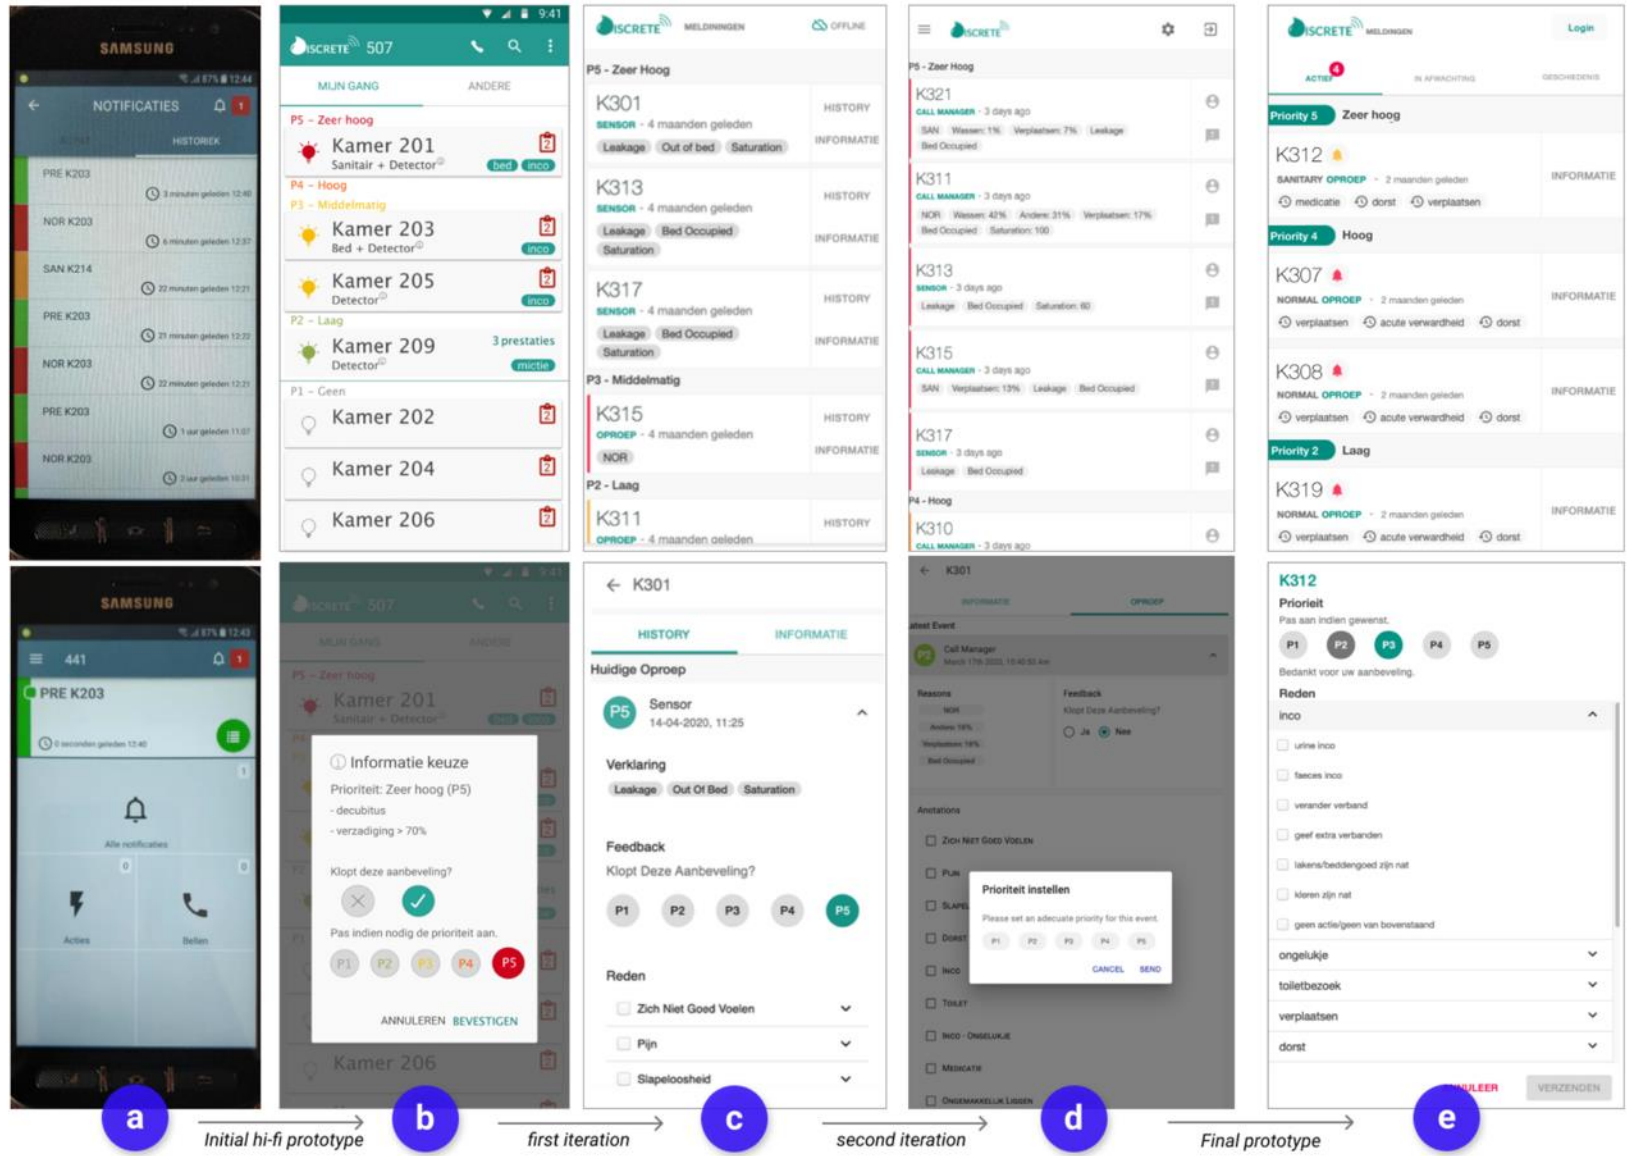

Fig. 1. The various iterations of the mobile user interface after addressing feedback towards our defined design goals in each step: a) original call notification platform, b) first high fidelity prototype, c) first iteration, d) second iteration, e) final implementation.

Explaining predictions

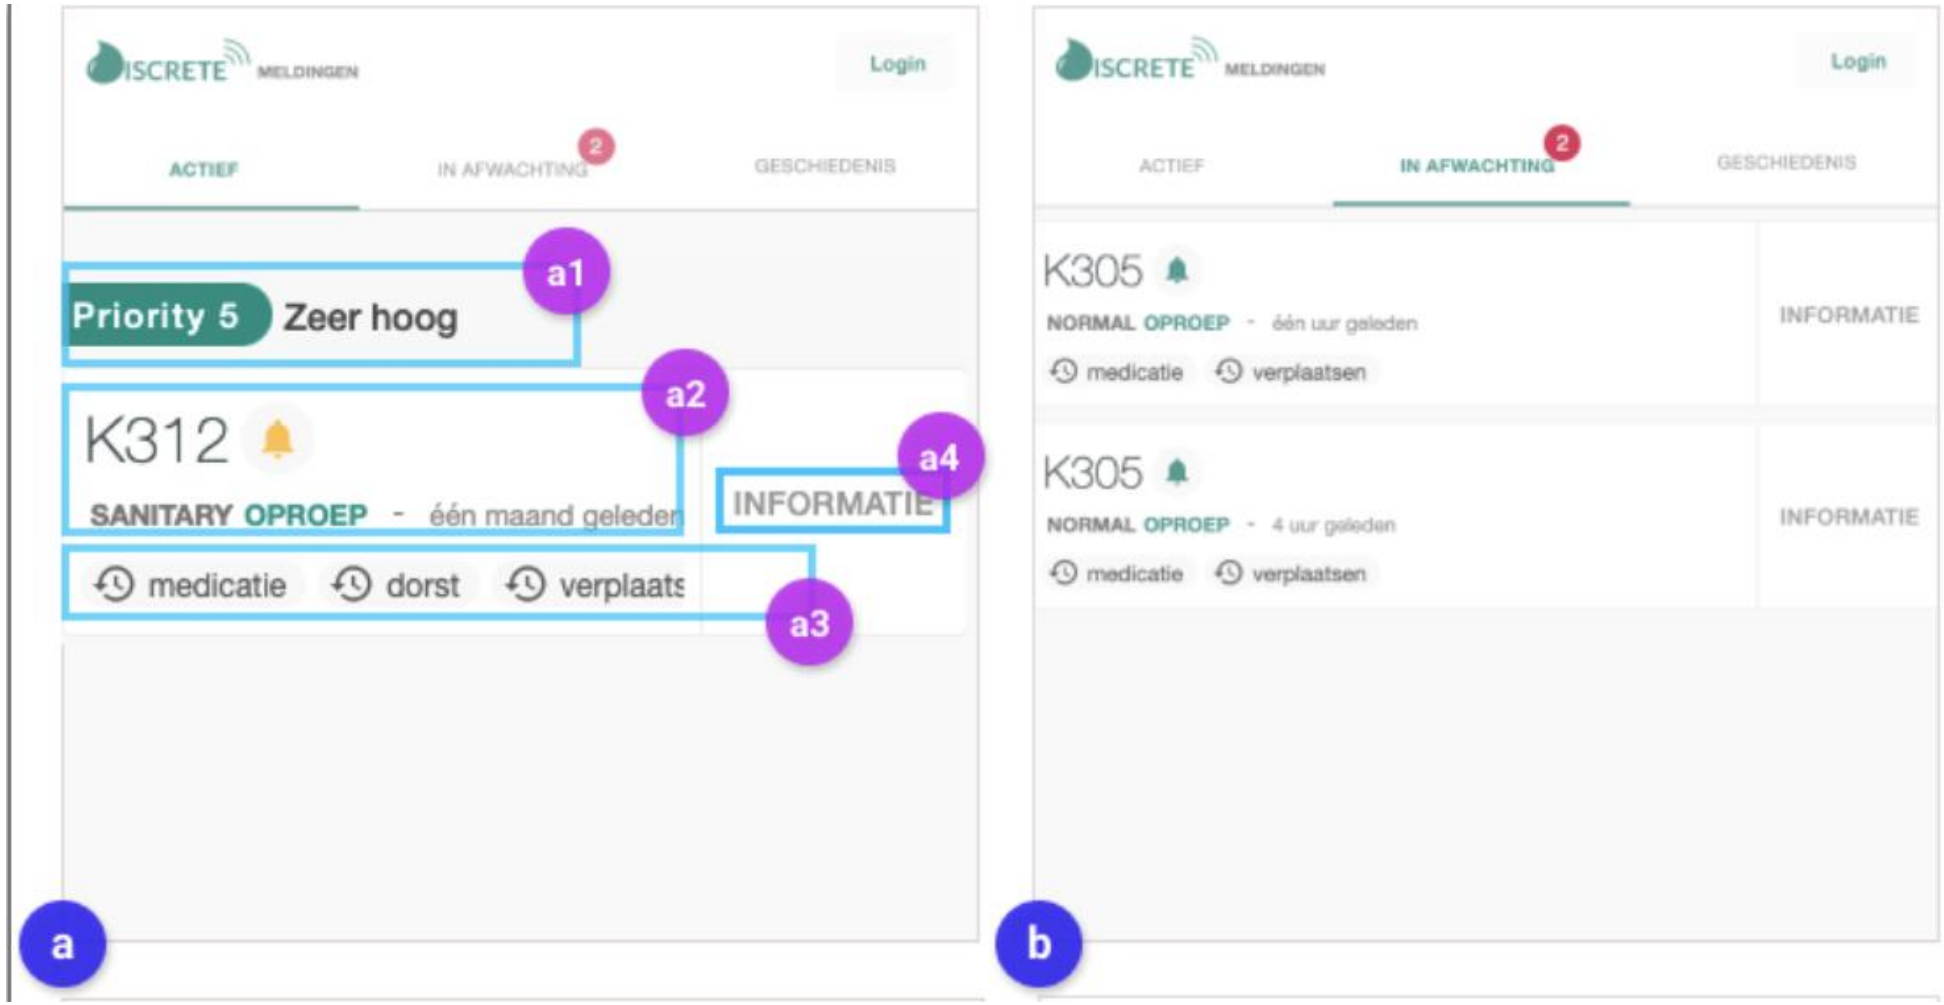

Gutiérrez Hernández, F. S., Htun, N. N., Vanden Abeele, V., De Croon, R., & Verbert, K. (2021). Explaining call recommendations in nursing homes: a user-centered design approach for interacting with knowledge-based health decision support systems. In Proceedings of the 27th Annual Conference on Intelligent User Interfaces. ACM.

TRADITIONAL (TEXTUAL) COMMUNICATION

I'd like to know the **salary of a Prolog Developer** around Haarlem, with 10 years of experience?

Sure, that seems to be a **7.780 euro gross salary** per **month.** However, the result might be inaccurate as this is **based on a low number** of previous postings.

High uncertainty, low confidence + uncertainty disclosure

VISUAL COMMUNICATION

I'd like to know the **salary of a Prolog Developer** around Haarlem, with 10 years of experience?

Sure, that seems to be a **7.780 euro gross salary** per month.

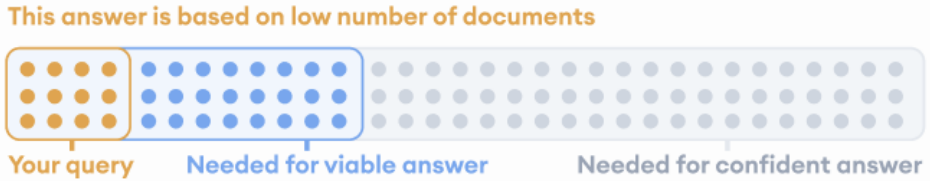

High uncertainty, low confidence + uncertainty disclosure

[1] Zhang, Y., Liao, Q., & Bellamy, R. (2020). Effect of confidence and explanation on accuracy and trust calibration in ai-assisted decision making.

VISUAL COMMUNICATION

I'd like to know the **salary of a Prolog Developer** around Haarlem, with 10 years of experience?

Sure, that seems to be a **7.780 euro gross salary** per month.

This answer is based on low number of documents

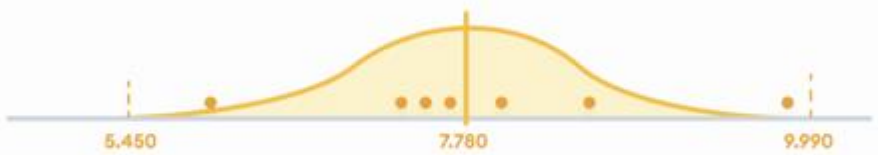

High uncertainty, low confidence + uncertainty disclosure

VISUAL COMMUNICATION

I'd like to know the **salary of a Prolog Developer** around Haarlem, with 10 years of experience?

Sure, that seems to be a **7.780 euro gross salary** per month.

This answer is based on low number of documents

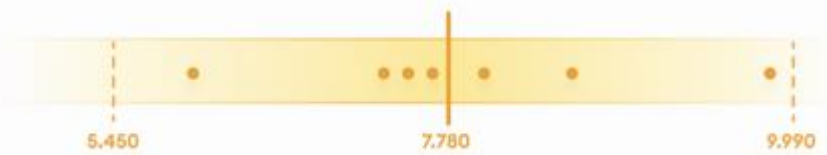

High uncertainty, low confidence + uncertainty disclosure

VISUAL SUGGESTIONS

I'd like to know the **salary of a Prolog Developer** around Haarlem, with 10 years of experience?

That seems to be a **7.780 euro gross salary** per month. However, this is based on **limited data for Haarlem**. Maybe try the following to improve results:

Amsterdam (nearby city)  +234% more accurate

Noord-Holland (enclosing province)  +131% more accurate

Java Developer (similar title)  +119% more accurate

High uncertainty, low confidence + suggestion

Secrets of the Super Agers | Public Notebook

Settings | PRO

Sources

5. Sleep (from Super Agers)

Source guide

Summary

This text emphasizes that **sleep is a non-negotiable biological state** crucial for human health, akin to air, food, and water. It highlights how sleep facilitates the **brain's glymphatic system** in clearing metabolic waste, like beta-amyloid protein, which is linked to Alzheimer's disease when clearance is impaired. The optimal sleep duration is identified as **around seven hours**, with both less and more sleep correlating to adverse health outcomes, including increased mortality risk and cognitive decline. Finally, the text addresses the challenges of maintaining good sleep, particularly with aging, and explores various interventions from behavioral changes and digital therapies to the importance of diagnosing and managing **sleep apnea**, a common and impactful sleep disorder.

Key Topics

- Sleep Importa...
- Sleep Durat...
- Sleep Disrupti...
- Sleep Improvem...
- Sleep Ap...

SLEEP

A restful good night's sleep can provide a magical sense of restoration and wellness. In recent years, we've gleaned a better understanding for what's going on while we sleep, and why it's so essential to our health. During sleep, metabolic waste products are transported out of the brain by cerebrospi-nal fluid and interstitial fluid. Not to be confused with our lymphatics and lymph nodes in the rest of our body, the brain's glymphatic system is a net-work of vessels that interacts with and complements the cerebrospinal flu-id's capability to clear the waste products. During non-rapid eye movement sleep, there are large waves of cerebrospinal flow. Impaired waste clearance has been documented in young, healthy individuals with a single night of sleep deprivation, as reflected by increases in the beta-amyloid protein seen in brain imaging. Marked accumulation of this protein is considered a precursor to the development of Alzheimer's disease. The dynamics for cerebrospinal and interstitial fluid flow is influenced by neurovascular cou-pling and our circadian rhythm: during the day, our waste production is high and our vascular tone is increased, whereas slow wave neuronal activ-ity is low. These dynamics are reversed during the night. In the words of Michael Grandner and Fabian-Xose Fernandez, "Sleep is a non-negotiable biological state required for the maintenance of human life ... our needs for sleep parallel those for air, food, and water." The relationship between sleep and health outcomes is profound, including all-cause mortality, car-diovascular mortality, cancer mortality, risk of cancer, type 2 diabetes and metabolic dysfunction, immune system function, obesity, Alzheimer's dis-ease, hypertension, stroke, female reproductive health, and mental and behavioral health. A brief review of a few relevant studies provides a sense of

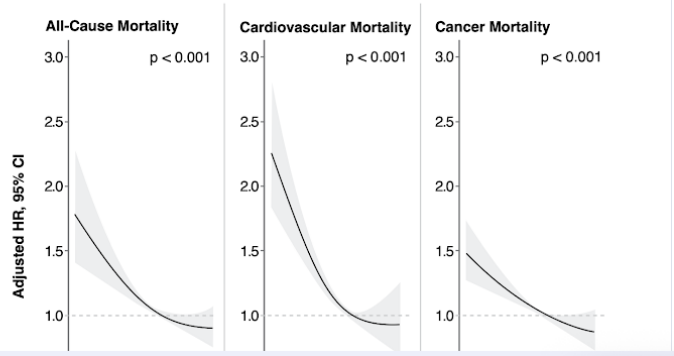

Chat

Refresh

approach he terms "lifestyle+" [1]. This goes beyond the traditional focus on just diet, exercise, and sleep, to include critical environmental factors and social connections [1].

Here's a comprehensive overview of what you should consider for a longer, healthier life, based on these sources:

1. Optimize Your Diet A poor diet is linked to a significant percentage of deaths globally, surpassing other health risks like tobacco or cancer [2]. Topol's work, including his *Super Agers* book, underscores the importance of a thoughtful approach to food.

- **Avoid Ultra-Processed Foods (UPFs)**: These "alien, industrially produced, unnatural substances" are major contributors to poor health [3]. They contain numerous additives and undergo physical changes that maximize digestibility and cause spikes in blood glucose and insulin, unlike dietary fiber which slows digestion [3].
  - Consuming UPFs is associated with markedly heightened risks of cardiovascular and metabolic diseases, including an 80% elevated risk of metabolic syndrome, 40% higher risk of type 2 diabetes, and 55% risk of obesity [4]. A mere 10% increase in UPF intake can lead to a 16% increased risk of cognitive impairment in older adults, and regular consumption of processed red meat is linked to a 14% higher risk of dementia [4].
  - More than four servings per day of UPFs are linked to a 62% increase in all-cause mortality [5].
  - **Practical Tip:** Read labels, choose items with few ingredients and no additives or fake sugars, and shop mostly in the grocery store perimeter for fresh foods [6]. The Open Food Facts app can be a helpful resource [6]. Topol notes that the US lags other countries in issuing guidelines against UPFs, partly due to the influence of "Big Food" corporations [7]. He suggests that UPFs may eventually be viewed akin to cigarettes, with regulations for conspicuous labeling [8].
- **Limit Sugary Beverages**: These are the biggest source of added sugar in diets [8]. High sugary beverages, including fruit juices, are associated with increased all-cause mortality, cardiovascular, and cancer-related mortality [8]. Atrial fibrillation risk increased more than threefold with over two liters per week of sugary drinks [8].
- **Be Cautious with Artificial Sweeteners**: The data is "a bit more complicated" and conflicting [9]. Some studies show an association with cardiovascular and cerebrovascular disease, and effects on glucose regulation and the gut microbiome [9]. Topol suggests they are not as worrisome as high sugar consumption, and some, like stevia, appear less concerning [9].
- **Moderate Salt Intake**: While sodium intake is clearly linked to hypertension, the magnitude of effect varies [10]. Moderate consumption (1-2 teaspoons of salt/day) may not be a problem, but cardiovascular risk becomes obvious at levels over 5 grams of sodium per day [10]. High salt diets can also reduce blood flow to the brain and pose a risk of cognitive impairment [11].
  - **Practical Tip:** Avoid or limit adding salt to foods, check food labels, and consider potassium chloride salt substitutes if no kidney disease is present [11].
- **Prioritize Quality Macronutrients**: The *type* of carbohydrate, protein, and fat matters [12] [13].
  - **Carbohydrates:** Avoid "carbotoxicity" from excessive intake [12]. Prioritize good, unprocessed carbs like nonstarchy vegetables, legumes, fruits, and whole grains [14]. Dietary fiber (25–30 grams per day) is associated with a 15-30% reduced all-cause and cardiovascular mortality, type 2 [illegible] cer [12].

Jump to bottom

Start typing... | 17 sources

How does sleep affect health? | What are microplastics' health effects? | How does exercise influence aging?

Studio

Audio Overview

Secrets of the Super Agers

00:00 / 32:04 • English

Notes

- **Topol's Super Agers: Five Healthspan Revelations**
  Give me the five most surprising facts about extending healthspan based on Eric Topol's work. Please reference him and his book Super Agers up top. His name is...
- **How To Get The Most Out Of This Notebook**
  Welcome! This Featured Notebook is designed to help you explore the research and transformative insights presented in Dr. Eric Topol's book, Super Agers and his...
- **Sample Question: How Worried Should I Be About Microplastics?**
  Drawing on Eric Topol's insights from his book Super Agers and his Ground Truths newsletter, one should be quite concerned about microplastics (MNPs) given the...
- **Sample Question: What Should I Focus On In My Diet To Expand Healthspan?**
  Drawing on insights from Dr. Eric Topol's book Super Agers and his Ground Truths newsletter, it's clear that diet plays a profoundly vital role in influencing our healthspa...
- **Briefing Doc: Exercise, The Ultimate Medical Intervention**
  Briefing Document: The Profound Impact of Exercise on Health Executive Summary Exercise is emerging as potentially the "single most potent medical intervention ever...
- **Briefing Doc: Modern Keys to Health and Longevity**
  Health and Longevity: Key Themes from Recent Research This briefing synthesizes key themes and findings from the provided sources regarding diet, exercise, sleep, and...
- **FAQ: Optimizing Health and Longevity**
  1. What are Ultra-Processed Foods (UPFs) and why are they considered harmful? Ultra-Processed Foods (UPFs) are industrially produced, "alien, unnatural substances" that...
- **Healthy Aging and Lifestyle Factors**
  Click to open the Mind Map

“knowledge graphs”

Microsoft Research Blog

GraphRAG: Unlocking LLM discovery on narrative private data

Published February 13, 2024

By Jonathan Larson, Partner Data Architect; Steven Truitt, Principal Program Manager

Share this page

Research Areas

Artificial intelligence

Related tools

Graspologic

Related projects

Project GraphRAG

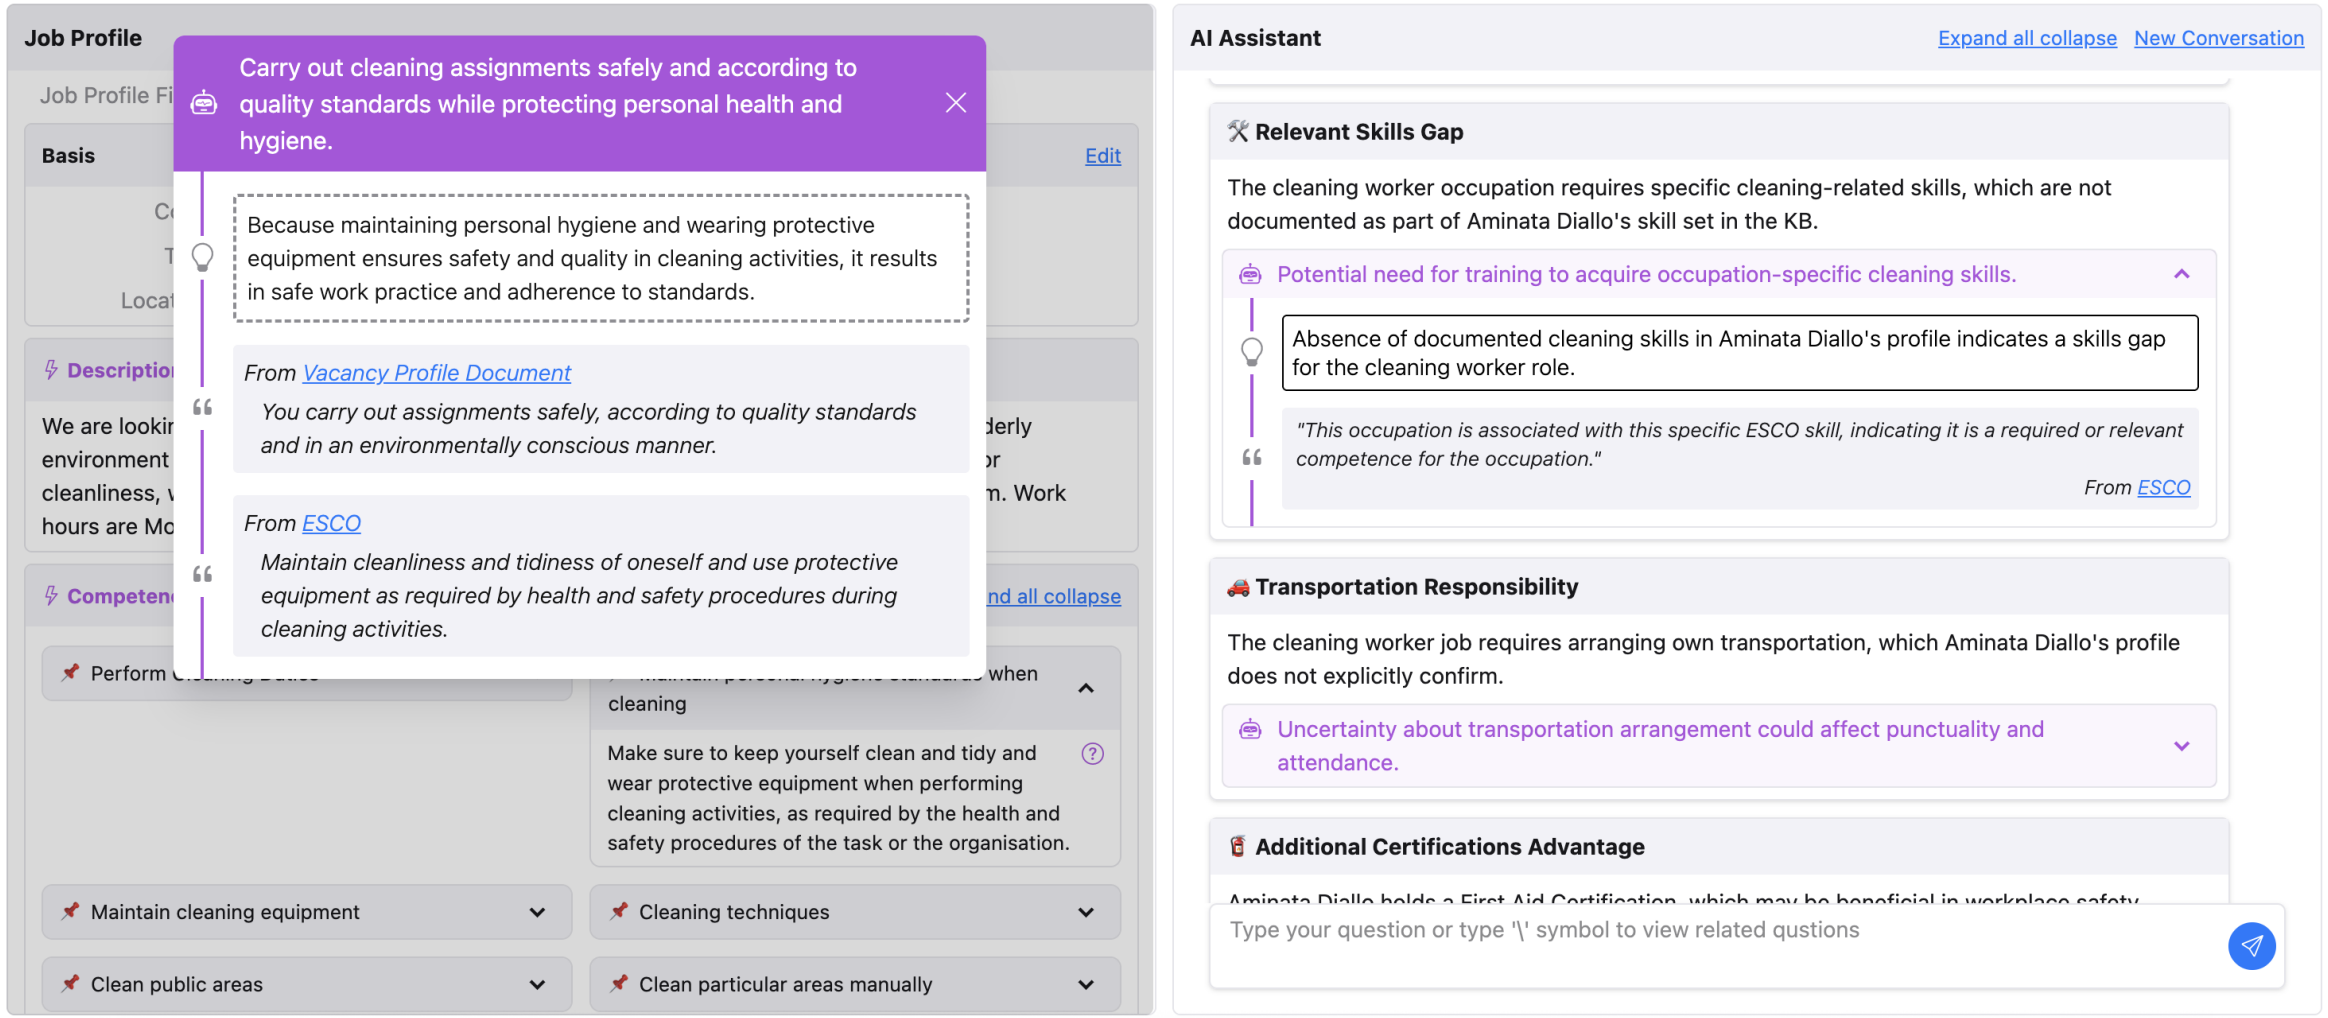

Designed by Augment Research Group, developed by Yizhe Zhang, Version: 1.0.0

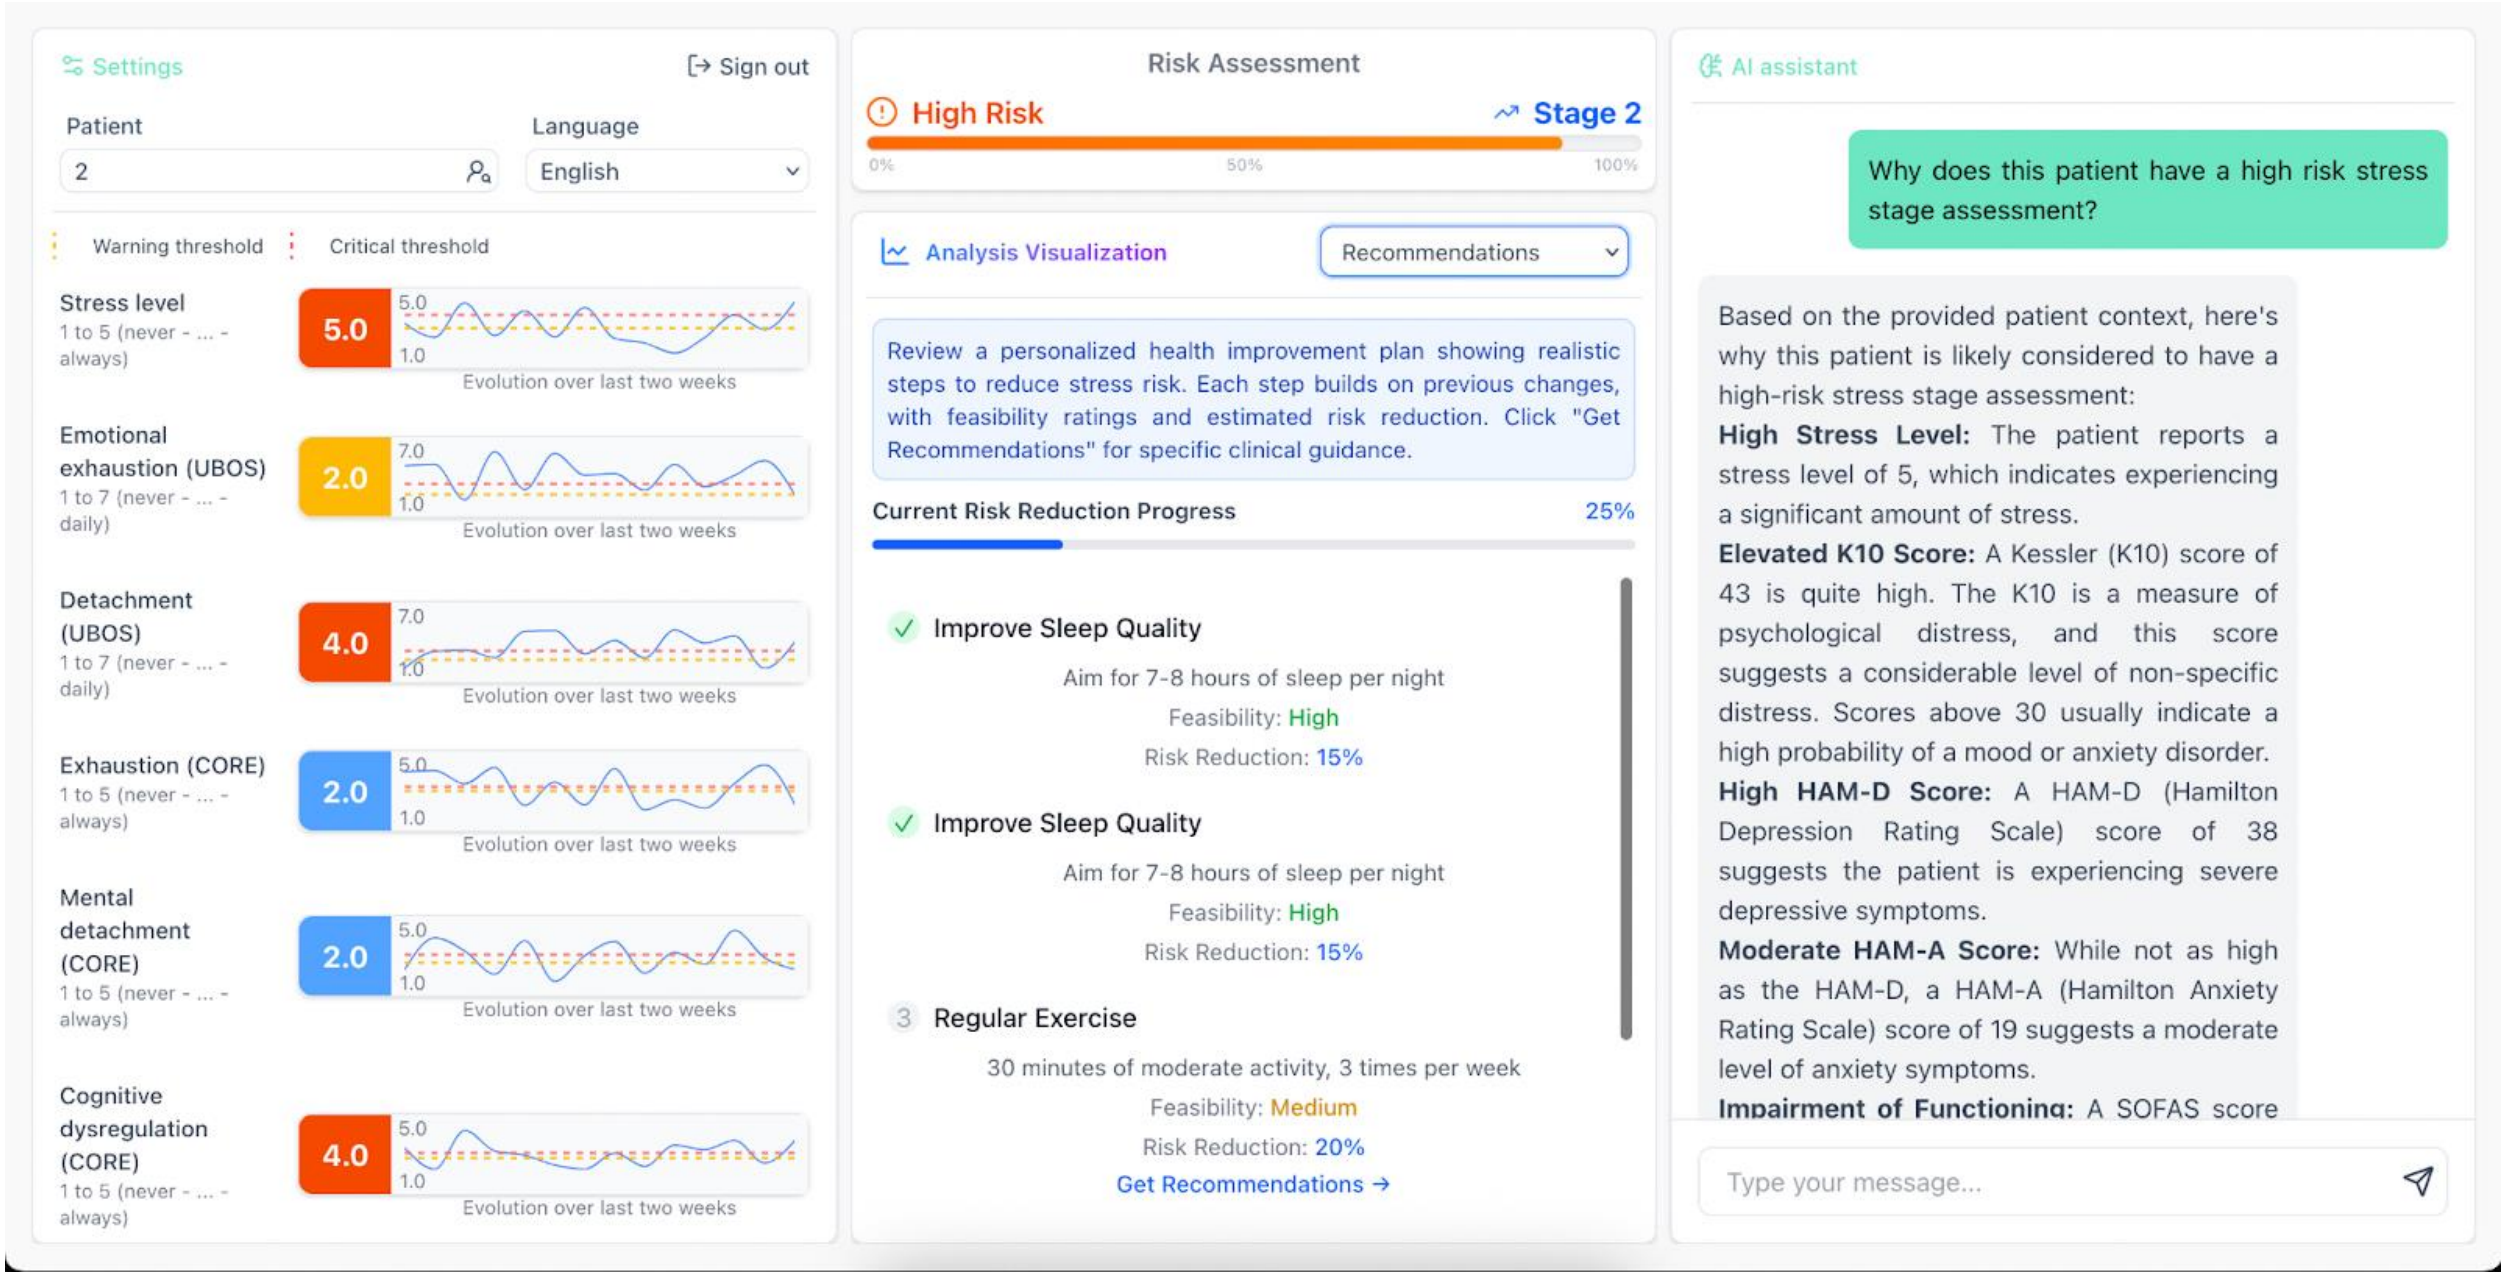
Settings
Sign out
Patient
2
Language
English
Warning threshold
Critical threshold
Stress level
1 to 5 (never - ... - always)
5.0
5.0
1.0
Evolution over last two weeks
Emotional exhaustion (UBOS)
1 to 7 (never - ... - daily)
2.0
7.0
1.0
Evolution over last two weeks
Detachment (UBOS)
1 to 7 (never - ... - daily)
4.0
7.0
1.0
Evolution over last two weeks
Exhaustion (CORE)
1 to 5 (never - ... - always)
2.0
5.0
1.0
Evolution over last two weeks
Mental detachment (CORE)
1 to 5 (never - ... - always)
2.0
5.0
1.0
Evolution over last two weeks
Cognitive dysregulation (CORE)
1 to 5 (never - ... - always)
4.0
5.0
1.0
Evolution over last two weeks
Risk Assessment
High Risk
Stage 2
0%
50%
100%
Analysis Visualization
Recommendations
Review a personalized health improvement plan showing realistic steps to reduce stress risk. Each step builds on previous changes, with feasibility ratings and estimated risk reduction. Click "Get Recommendations" for specific clinical guidance.
Current Risk Reduction Progress
25%
Improve Sleep Quality
Aim for 7-8 hours of sleep per night
Feasibility: High
Risk Reduction: 15%
Improve Sleep Quality
Aim for 7-8 hours of sleep per night
Feasibility: High
Risk Reduction: 15%
3 Regular Exercise
30 minutes of moderate activity, 3 times per week
Feasibility: Medium
Risk Reduction: 20%
Get Recommendations →
AI assistant
Why does this patient have a high risk stress stage assessment?
Based on the provided patient context, here's why this patient is likely considered to have a high-risk stress stage assessment:
High Stress Level: The patient reports a stress level of 5, which indicates experiencing a significant amount of stress.
Elevated K10 Score: A Kessler (K10) score of 43 is quite high. The K10 is a measure of psychological distress, and this score suggests a considerable level of non-specific distress. Scores above 30 usually indicate a high probability of a mood or anxiety disorder.
High HAM-D Score: A HAM-D (Hamilton Depression Rating Scale) score of 38 suggests the patient is experiencing severe depressive symptoms.
Moderate HAM-A Score: While not as high as the HAM-D, a HAM-A (Hamilton Anxiety Rating Scale) score of 19 suggests a moderate level of anxiety symptoms.
Impairment of Functioning: A SOFAS score
Type your message...

Settings

Sign out

Patient

2

Language

English

Warning threshold

Critical threshold

Stress level
1 to 5 (never - ... - always)

5.0

Evolution over last two weeks

Emotional exhaustion (UBOS)
1 to 7 (never - ... - daily)

2.0

Evolution over last two weeks

Detachment (UBOS)
1 to 7 (never - ... - daily)

4.0

Evolution over last two weeks

Exhaustion (CORE)
1 to 5 (never - ... - always)

2.0

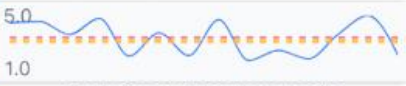

Evolution over last two weeks

Mental detachment (CORE)
1 to 5 (never - ... - always)

2.0

Evolution over last two weeks

Cognitive dysregulation (CORE)
1 to 5 (never - ... - always)

4.0

Evolution over last two weeks

Risk Assessment

High Risk

Stage 2

0%

50%

100%

Analysis Visualization

Factor Importance

Shows how each health factor influences the AI's assessment. Bars extending left (green) indicate factors working against the assessment, while bars extending right (orange) indicate factors supporting it.

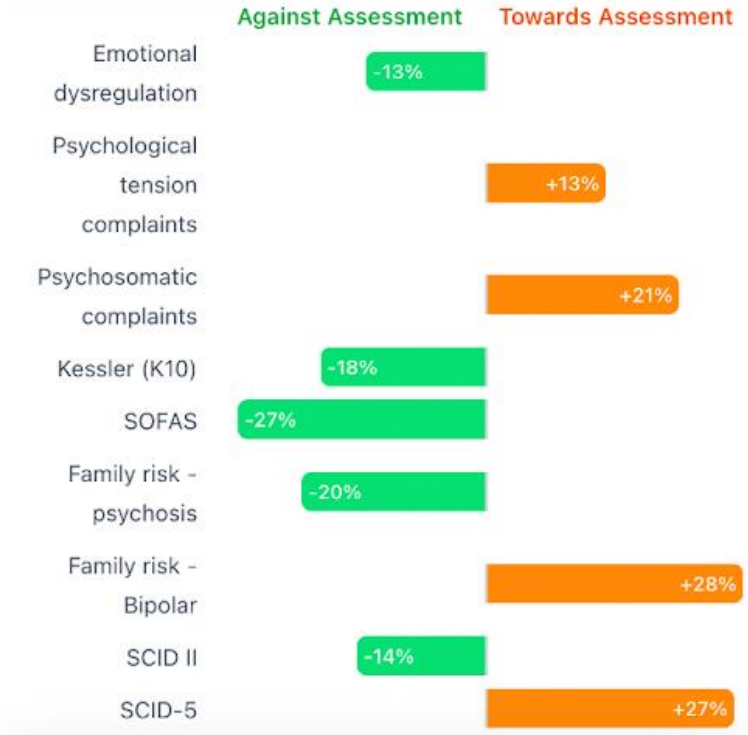

AI assistant

Why does this patient have a high risk stress stage assessment?

Based on the provided patient context, here's why this patient is likely considered to have a high-risk stress stage assessment:

High Stress Level: The patient reports a stress level of 5, which indicates experiencing a significant amount of stress.

Elevated K10 Score: A Kessler (K10) score of 43 is quite high. The K10 is a measure of psychological distress, and this score suggests a considerable level of non-specific distress. Scores above 30 usually indicate a high probability of a mood or anxiety disorder.

High HAM-D Score: A HAM-D (Hamilton Depression Rating Scale) score of 38 suggests the patient is experiencing severe depressive symptoms.

Moderate HAM-A Score: While not as high as the HAM-D, a HAM-A (Hamilton Anxiety Rating Scale) score of 19 suggests a moderate level of anxiety symptoms.

Impairment of Functioning: A SOFAS score

Type your message...

Hands-on tools for building effective human-AI experiences

The HAX Toolkit is for teams building user-facing AI products. It helps you conceptualize what the AI system will do and how it will behave. Use it early in your design process.

[Guidelines for Human-AI Interaction]

Best practices for how AI systems should behave during interaction. Use them to guide your AI product planning.

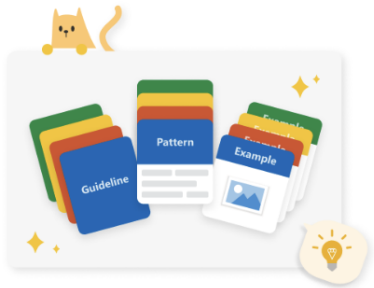

[HAX Design Library]

Learn the Guidelines for Human-AI Interaction and how to apply them, using patterns and examples.

[HAX Workbook]

Work together with your team to prioritize which Guidelines to implement in your product.

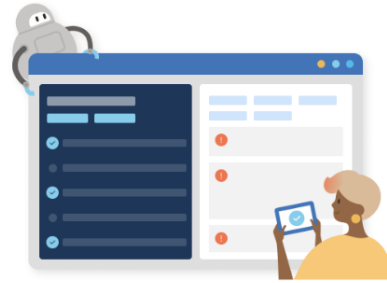

[HAX Playbook]

For applications using natural language processing, identify common failures so you can plan for mitigating them.

Get started with the HAX Toolkit


Initially
During interaction
When wrong
Over time
1
Make clear what the system can do.
2
Make clear how well the system can do what it can do.
3
Time services based on context.
4
Show contextually relevant information.
5
Match relevant social norms.
6
Mitigate social biases.
7
Support efficient invocation.
8
Support efficient dismissal.
9
Support efficient correction.
10
Scope services when in doubt.
11
Make clear why the system did what it did.
12
Remember recent interactions.
13
Learn from user behavior.
14
Update and adapt cautiously.
15
Encourage granular feedback.
16
Convey the consequences of user actions.
17
Provide global controls.
18
Notify users about changes.

Submit an example >

Guideline 11 > Pattern 11A > Example

Pinterest | 11A: Local explanations >

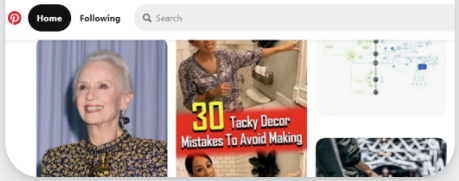

Guideline 11 > Pattern 11A > Example

InterpretML | 11A: Local explanations >

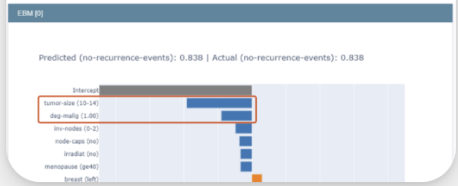

Guideline 11 > Pattern 11A > Example

Copilot in Excel | 11A: Local explanations >

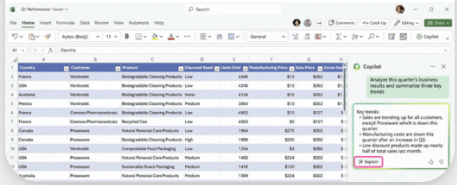

Guideline 11 > Pattern 11A > Example

Business Chat | 11A: Local explanations >

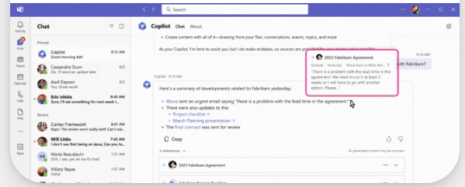

Guideline 11 > Pattern 11A > Example

AlphaCode | 11A: Local explanations >

Guideline 11 > Pattern 11A > Example

Verification-focused explanations | 11A: Local explanations >

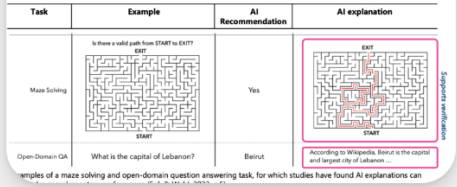

Guideline 11 > Pattern 11A > Example

Google Search | 11A: Local explanations >

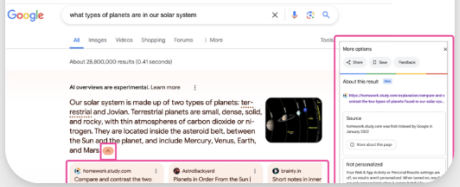

Thank you

robin.decroon@kuleuven.be
https://augment.cs.kuleuven.be/