


AI Models, WHY DATA QUALITY MATTERS



AI FOR HEALTH
BY ARTEFACT

Armand LEOPOLD - QALITA
21 Novembre 2024
Paris - Station F

My Journey



Armand LEOPOLD



Data Analyst
2016



Data Scientist
2017



Data Engineer
2019



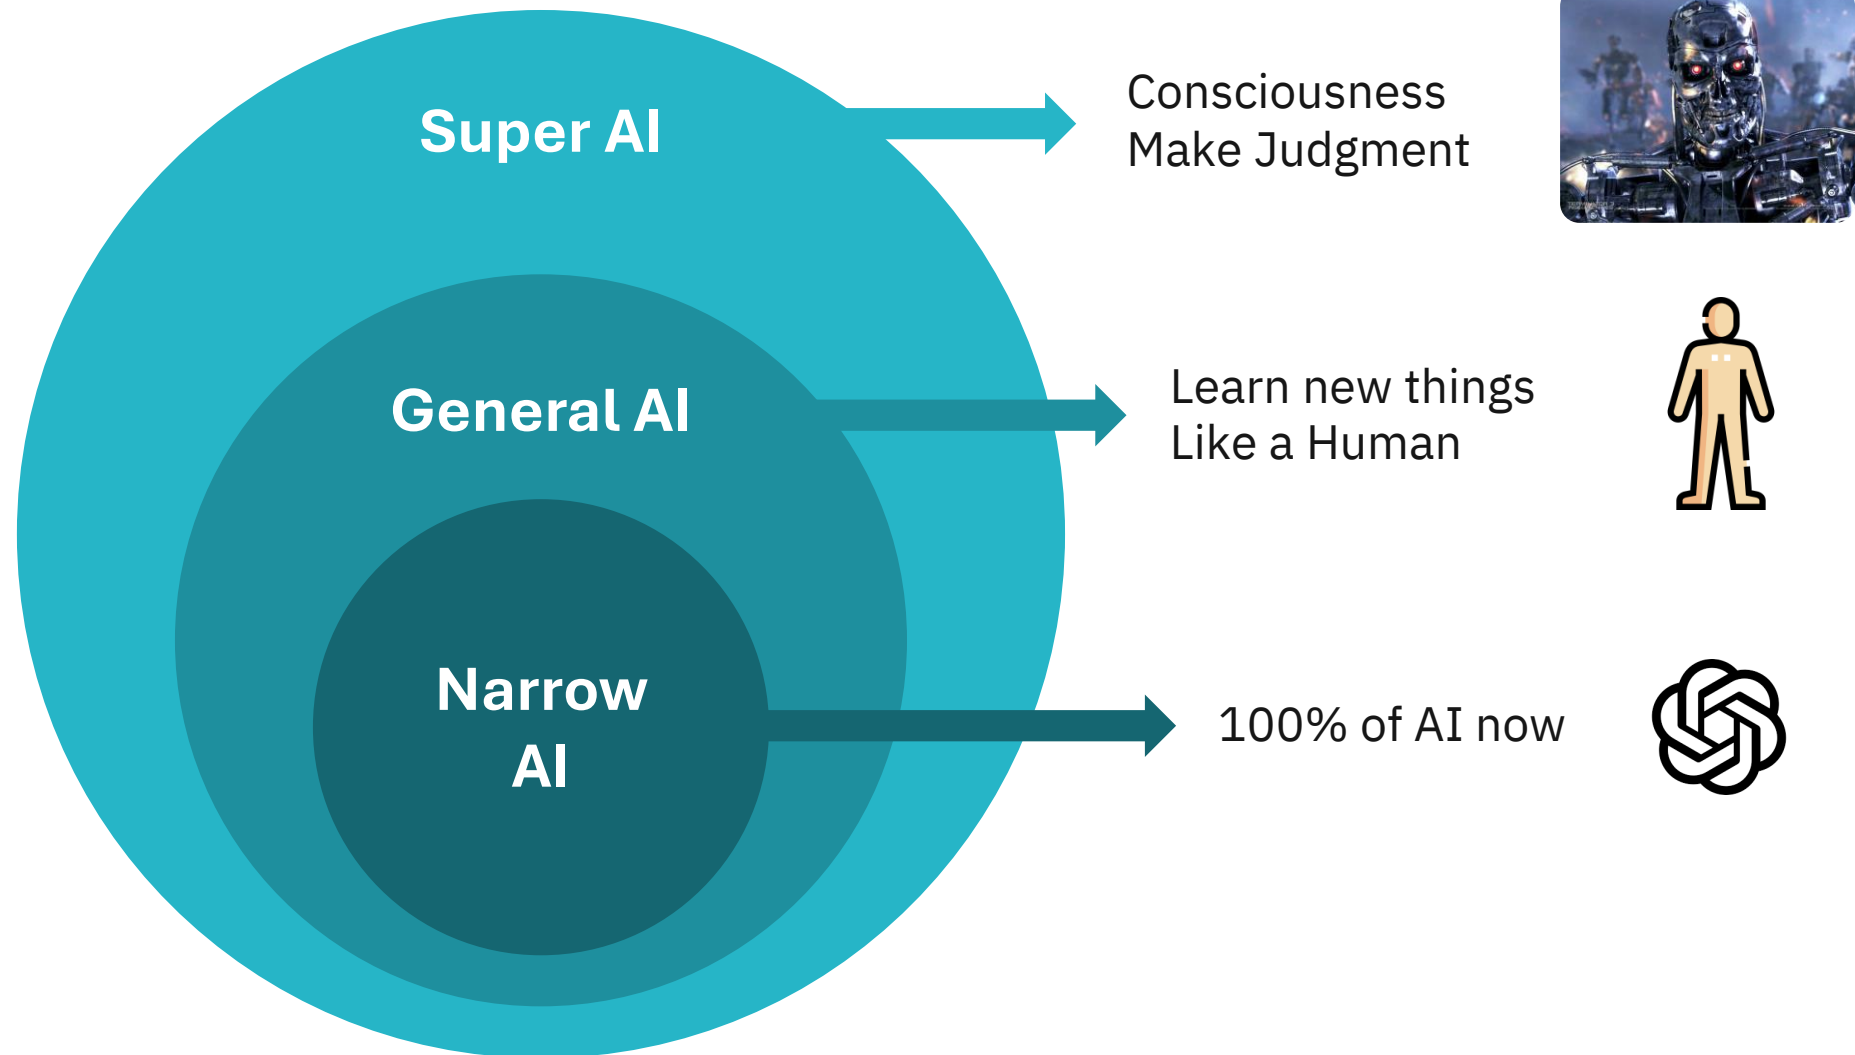
Head of Data Factory
2022



CEO
2023 -

8+ Years in the Field

What is AI ?

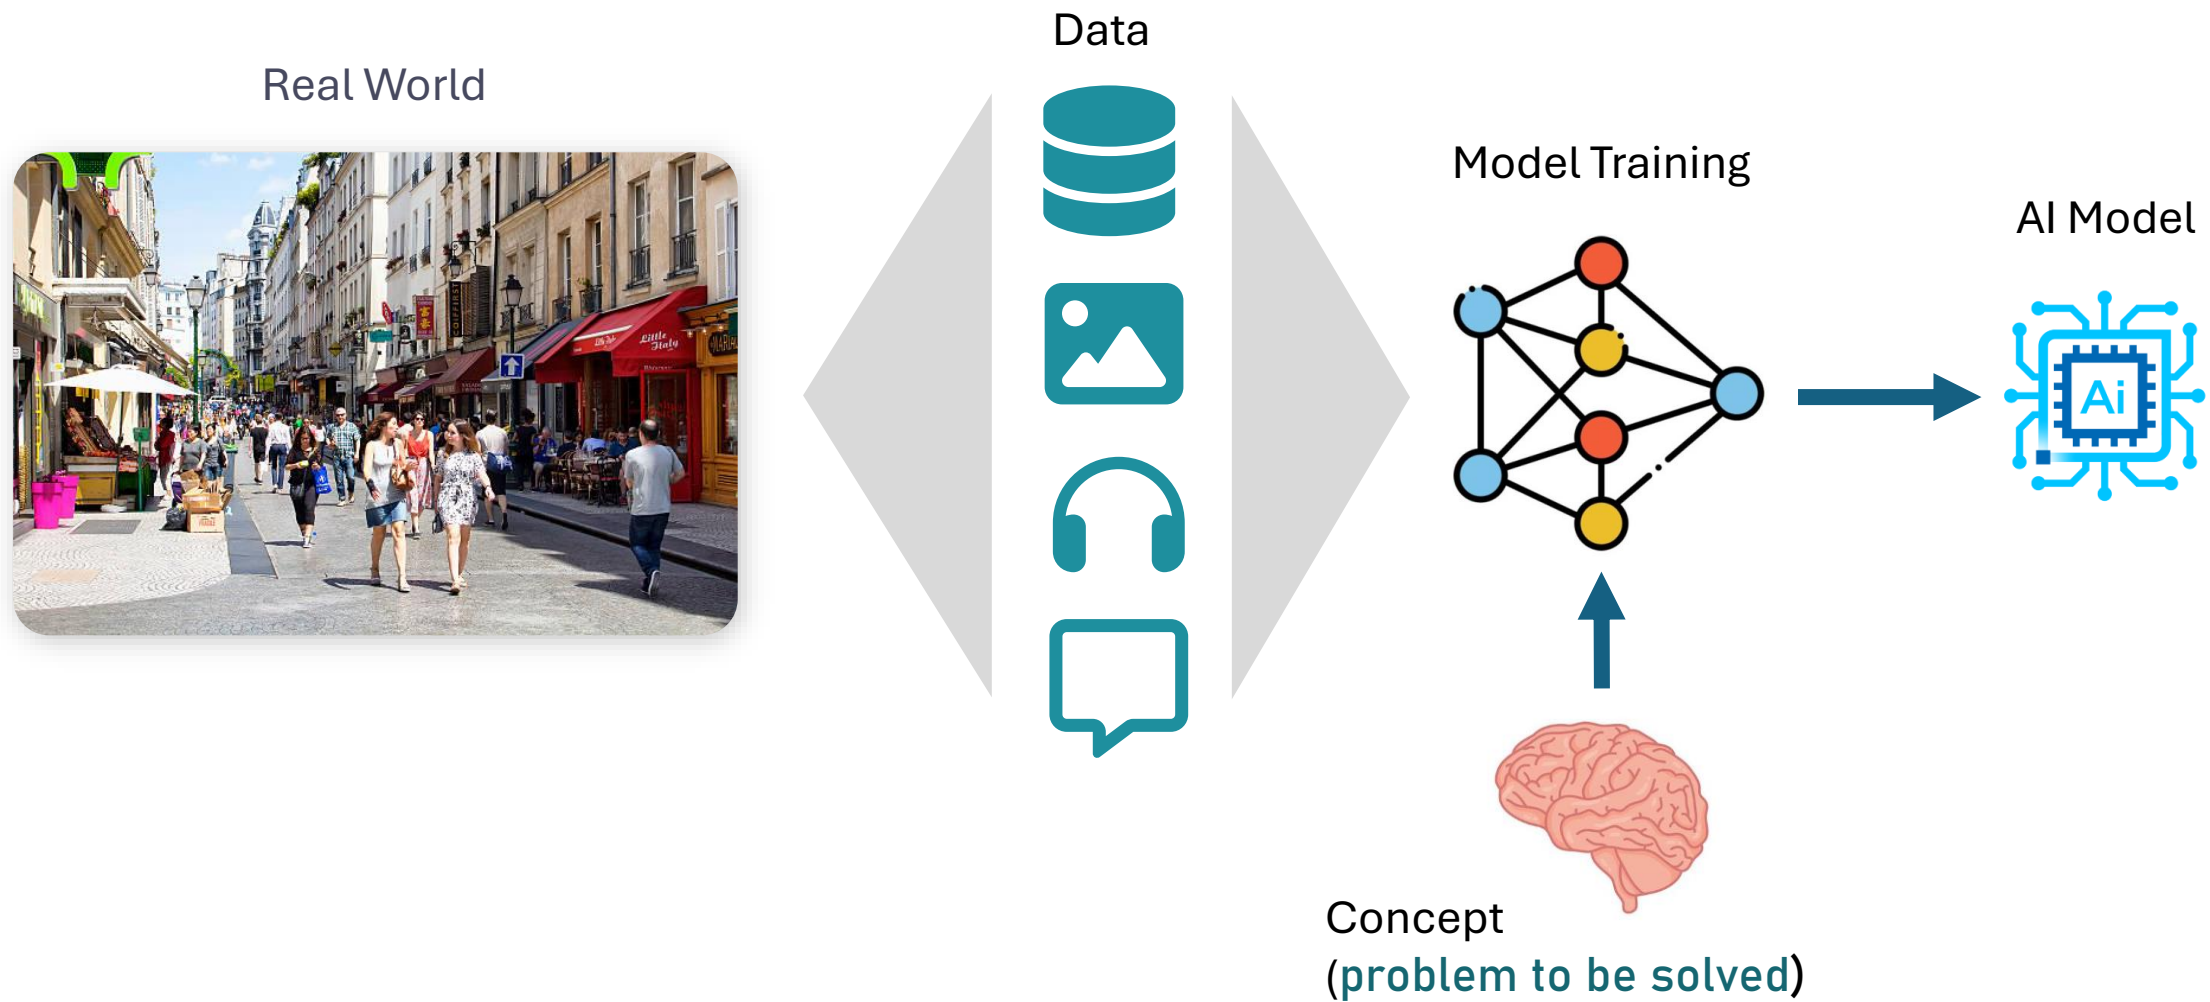


What do we need to create a narrow AI ?

1. A precise problem to be solved
2. Data related to the problem

Data is Knowledge

Simplified Data Science Pipeline




The Importance of Data in Training




Dog



Cat



AI Model



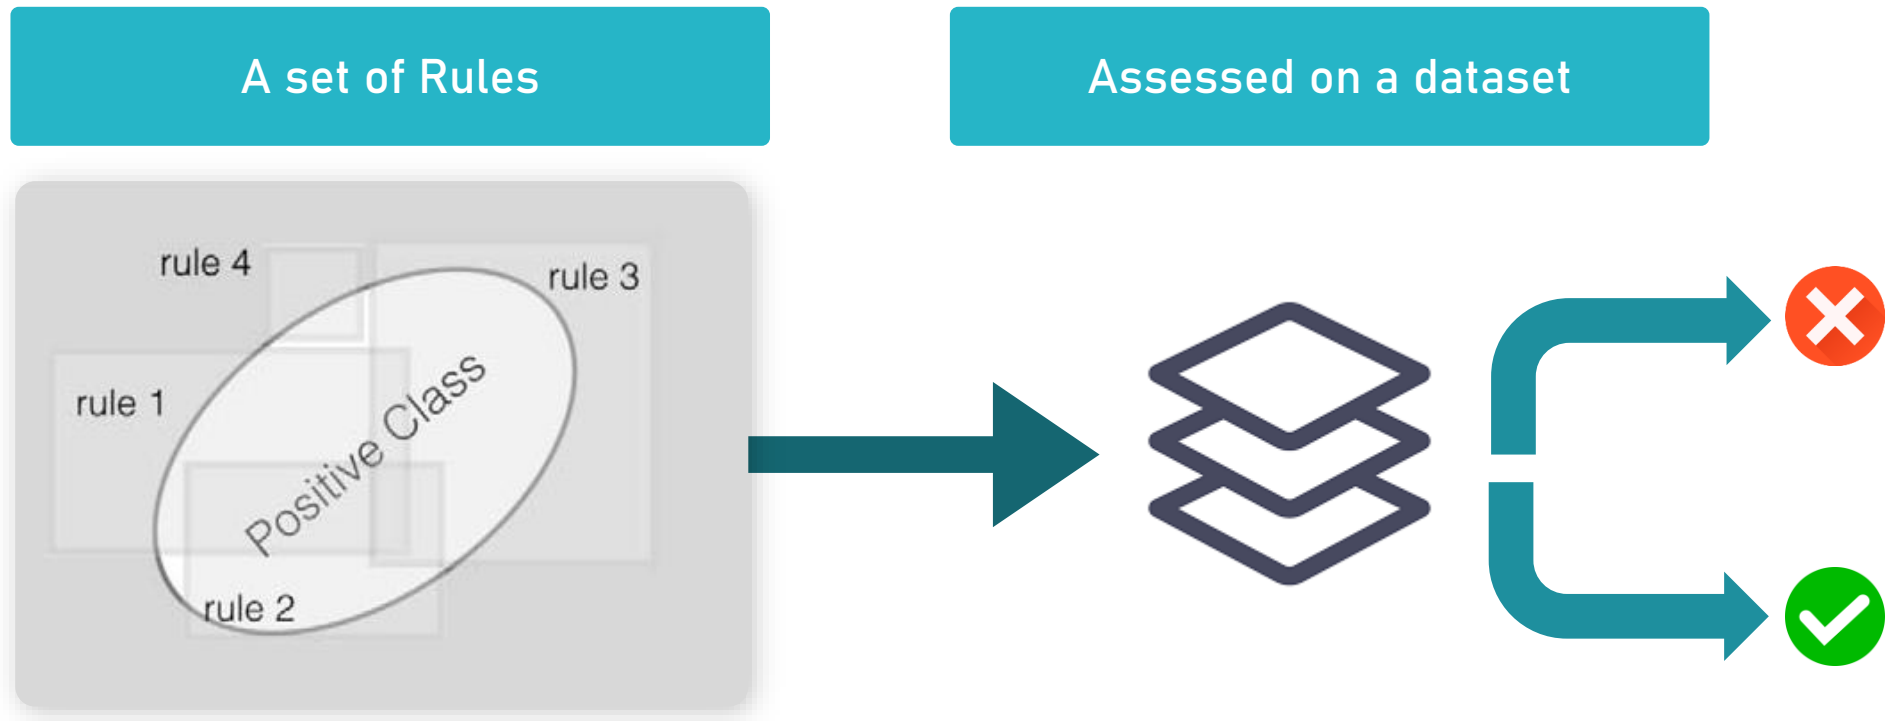
Garbage In > Garbage Out



Cat

What is Data Quality ?


= Data Requirements



Why is it (very) important in healthcare ?

- 1. Complex Field
- 2. High Stakes
- Lot of things to manage
- You can't afford to fail

Diagnosis

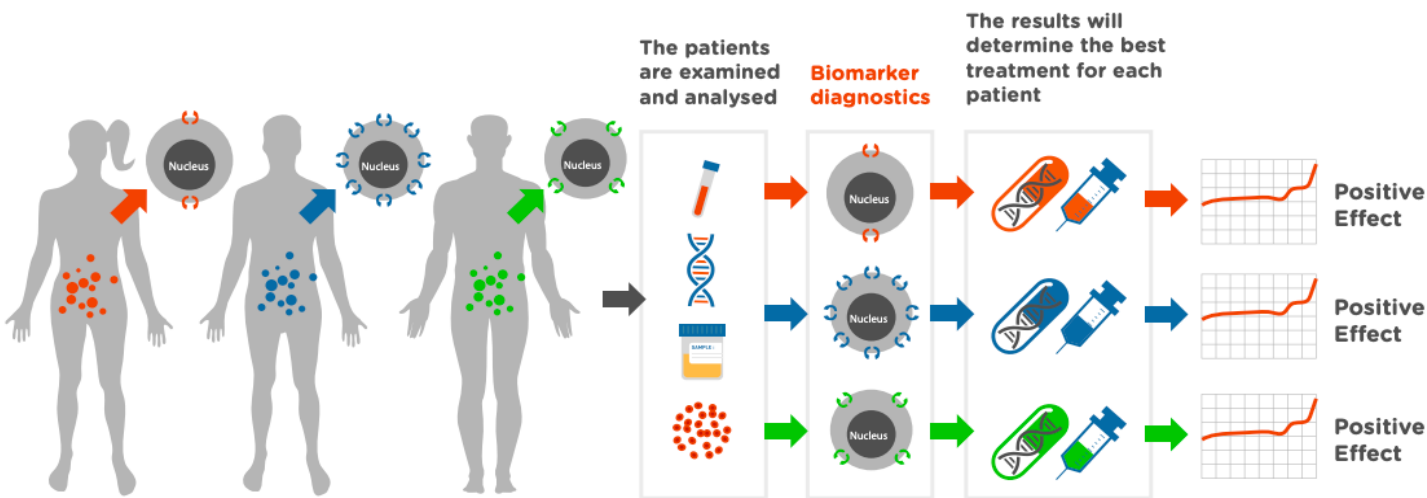


RSNA *A Deep Learning Mammography-based Model for Improved Breast Cancer Risk Prediction – 2019 MIT*

Treatment

INNOVATIVE MEDICINE: PERSONALISED MEDICINE

Cancer patients with e.g. colon cancer receive a personalised therapy based on their biomarkers



Why is it (very) important in healthcare ?

EU Artificial Intelligence Act

AI Risk Classification

Health : high-risk
AI systems

Requirements for providers of high-risk AI systems

Establish a **quality management
system** to ensure compliance.



EU Artificial
Intelligence Act

[High-level summary of the AI Act](#)

Real-World Examples of Data Quality Issues




ChatGPT

In total, we present **21 risks**. We then discuss the points of origin of different risks and point to potential risk mitigation approaches. [...] **It can be mitigated at the point of origin, by better redaction or curation of training data.**



Ethical and social risks of harm from Language Models - 2021



COVID-19

Participants reported difficulties accessing sufficiently **timely, robust, granular, standardised** and documented data.



Data Science and AI in the age of COVID-19

Business Impact of Data Quality

ChatBots



McDonald's ends AI experiment after drive-thru ordering blunders

After working with IBM for three years to leverage AI to take drive-thru orders, McDonald's called the whole thing off in June 2024. The reason? A slew of social media videos showing confused and frustrated customers trying to get the AI to understand their orders.

One TikTok video in particular featured two people repeatedly pleading with the AI to stop as it kept adding more Chicken McNuggets to their order, eventually reaching 260. In a June 13, 2024, internal memo obtained by trade publication Restaurant Business, McDonald's announced it would end the partnership with IBM and shut down the tests.

The restaurant had piloted the AI at more than 100 US drive-thrus, and indicated it still saw a future in a voice-ordering solution.



Grok AI falsely accuses NBA star of vandalism spree

In an April 2024 post on X, Grok, the AI chatbot from Elon Musk's xAI, falsely accused NBA star Klay Thompson of throwing bricks through windows of multiple houses in Sacramento, Ca.

Some commentators speculated that Grok may have hallucinated the vandalism story about the Golden State Warriors star, who has since been traded to the Dallas Mavericks, after ingesting posts about Thompson "throwing bricks," common basketball parlance for a badly missed shot. In his final game with the Golden State Warriors, the team suffered a blowout loss as part of the worst postseason performance of Thompson's career.

Grok does display a disclaimer that says, "Grok is an early feature and can make mistakes. Verify its outputs." Still, the incident raises questions about liability when an AI chatbot makes false and defamatory statements.

Business Impact of Data Quality

ChatBots

MyCity NYC AI chatbot encourages business owners to break the law

In March 2024, The Markup reported that Microsoft-powered chatbot MyCity was giving entrepreneurs incorrect information that would lead to them break the law.

Unveiled in October 2024, MyCity was intended to help provide New Yorkers with information on starting and operating businesses in the city, as well as housing policy and worker rights. The only problem was The Markup found MyCity falsely claimed that business owners could take a cut of their workers' tips, fire workers who complain of sexual harassment, and serve food that had been nibbled by rodents. It also claimed landlords could discriminate based on source of income.

In the wake of the report, indicted New York City Mayor Eric Adams defended the project. The chatbot remains online.



Air Canada pays damages for chatbot lies

In February 2024, Air Canada was ordered to pay damages to a passenger after its virtual assistant gave him incorrect information at a particularly difficult time.

Jake Moffatt consulted Air Canada's virtual assistant about bereavement fares following the death of his grandmother in November 2023. The chatbot told him he could buy a regular price ticket from Vancouver to Toronto and apply for a bereavement discount within 90 days of purchase. Following that advice, Moffatt purchased a one-way CA\$794.98 ticket to Toronto and a CA\$845.38 return flight to Vancouver.

But when Moffatt submitted his refund claim, the airline turned him down, saying that bereavement fares can't be claimed after tickets have been purchased.

Moffatt took Air Canada to a tribunal in Canada, claiming the airline was negligent and misrepresented information via its virtual assistant. According to tribunal member Christopher Rivers, Air Canada argued it can't be held liable for the information provided by its chatbot.

Rivers denied that argument, saying the airline didn't take "reasonable care to ensure its chatbot was accurate," So he ordered the airline to pay Moffatt CA\$812.02, including CA\$650.88 in damages.

Business Impact of Data Quality

Algorithms



Healthcare algorithm failed to flag Black patients

In 2019, a study published in Science revealed that a healthcare prediction algorithm, used by hospitals and insurance companies throughout the US to identify patients in need of high-risk care management programs, was far less likely to flag Black patients.

High-risk care management programs provide trained nursing staff and primary-care monitoring to chronically ill patients in an effort to prevent serious complications. But the algorithm was much more likely to recommend white patients for these programs than Black patients.

The study found that the algorithm used healthcare spending as a proxy for determining an individual's healthcare need. But according to Scientific American, the healthcare costs of sicker Black patients were on par with the costs of healthier white people, which meant they received lower risk scores even when their need was greater.

The study's researchers suggested that a few factors may have contributed. First, people of color are more likely to have lower incomes, which, even when insured, may make them less likely to access medical care. Implicit bias may also cause people of color to receive lower-quality care.

While the study didn't name the algorithm or the developer, the researchers told Scientific American they were working with the developer to address the situation.

iTutor Group's recruiting AI rejects applicants due to age

In August 2023, tutoring company iTutor Group agreed to pay \$365,000 to settle a suit brought by the US Equal Employment Opportunity Commission (EEOC). The federal agency said the company, which provides remote tutoring services to students in China, used AI-powered recruiting software that automatically rejected female applicants ages 55 and older, and male applicants ages 60 and older.

The EEOC said more than 200 qualified applicants were automatically rejected by the software.

"Age discrimination is unjust and unlawful," EEOC Chair Charlotte Burrows said in a statement. "Even when technology automates the discrimination, the employer is still responsible."

iTutor Group denied any wrongdoing but did decide to settle the suit. As part of the settlement and consent decree, it agreed to adopt new anti-discrimination policies.

Business Impact of Data Quality

Algorithms



Amazon AI-enabled recruitment tool only recommended men

Like many large companies, Amazon is hungry for tools that can help its HR function screen applications for the best candidates. In 2014, Amazon started working on AI-powered recruiting software to do just that. There was only one problem: The system vastly preferred male candidates. In 2018, Reuters broke the news that Amazon had scrapped the project.

Amazon's system gave candidates star ratings from 1 to 5, but the ML models at the heart of the system were trained on 10 years' worth of résumés submitted to Amazon — most of them from men. As a result of that training data, the system started penalizing phrases in résumés that included the word "women's" and even downgraded candidates from all-women colleges.

At the time, Amazon said the tool was never used by Amazon recruiters to evaluate candidates. The company tried to edit the tool to make it neutral, but ultimately decided it couldn't guarantee it wouldn't learn some other discriminatory way of sorting candidates and ended the project.



Zillow wrote down millions, slashed workforce due to algorithmic home-buying disaster

In November 2021, online real estate marketplace Zillow told shareholders it would wind down its Zillow Offers operations and cut 25% of the company's workforce — about 2,000 employees — over the next several quarters. The home-flipping unit's woes were the result of the error rate in the ML algorithm it used to predict home prices.

Zillow Offers was a program through which the company made cash offers on properties based on a "Zestimate" of home values derived from an ML algorithm. The idea was to renovate the properties and flip them quickly. But a Zillow spokesperson told CNN the algorithm had a median error rate of 1.9%, and could be as high as 6.9% for off-market homes.

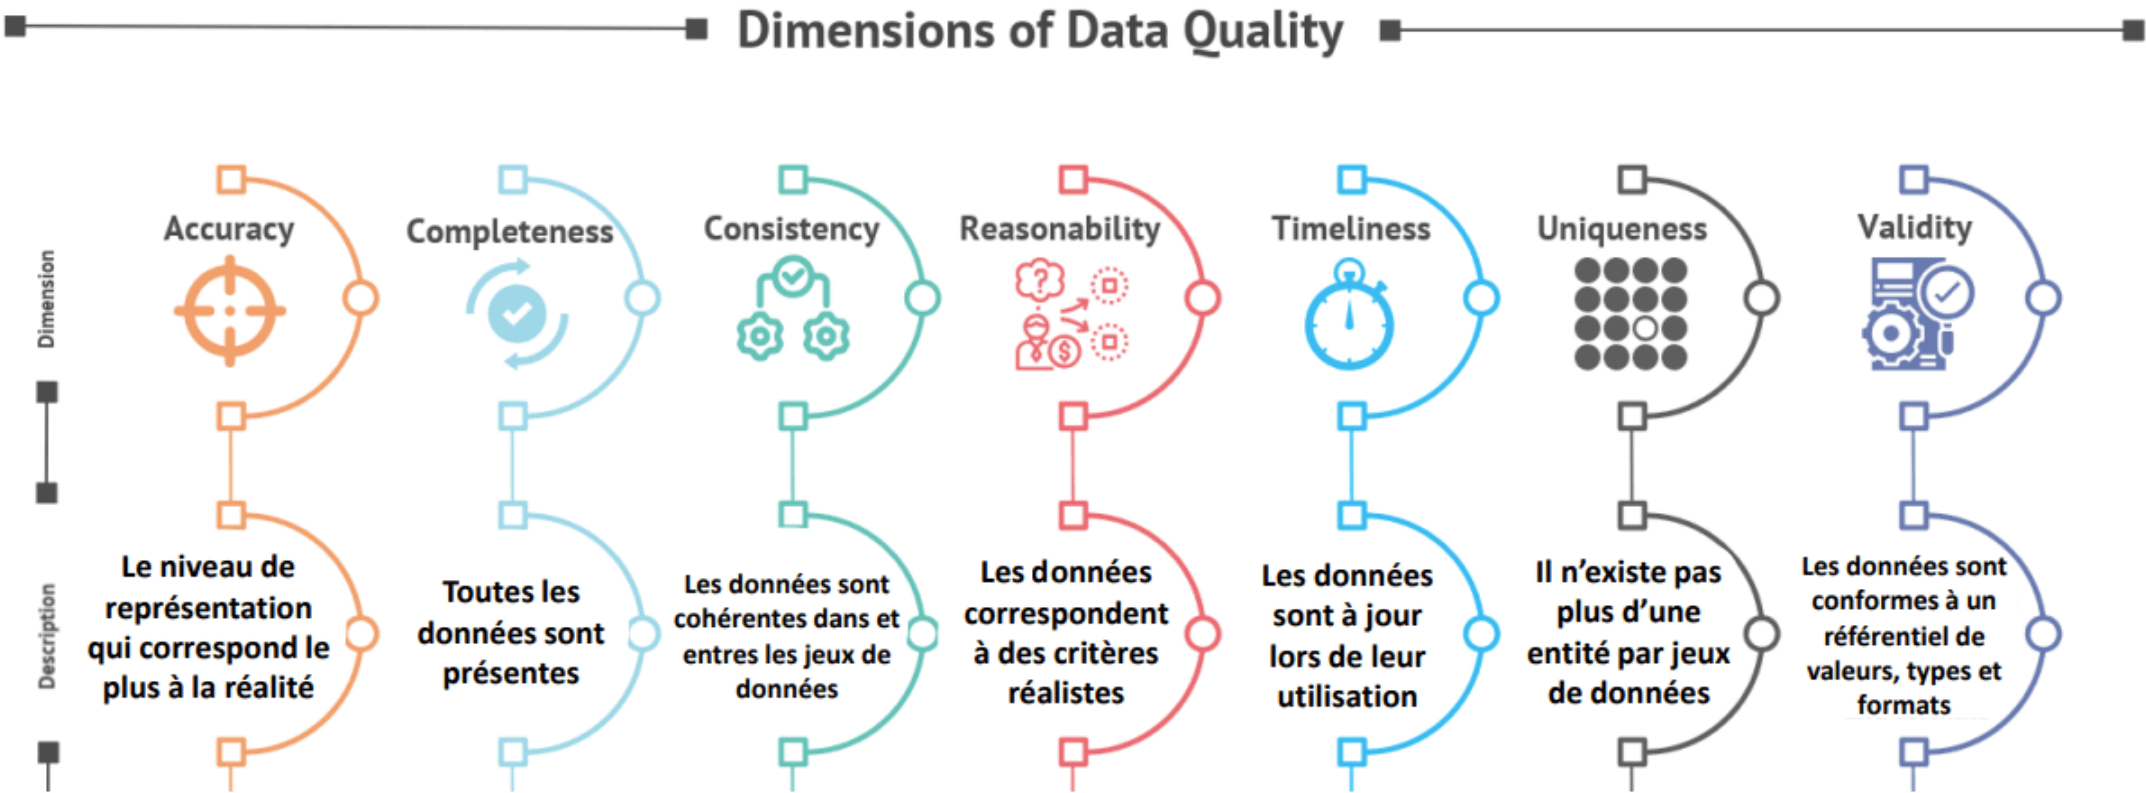
CNN reported that Zillow bought 27,000 homes through Zillow Offers since its launch in April 2018, but sold only 17,000 through the end of September 2021. Black swan events like the COVID-19 pandemic and a home renovation labor shortage contributed to the algorithm's accuracy troubles.

Zillow said the algorithm led it to unintentionally purchase homes at higher prices than its current estimates of future selling prices, resulting in a \$304 million inventory write-down in Q3 2021.

In a conference call with investors following the announcement, Zillow co-founder and CEO Rich Barton said it might be possible to tweak the algorithm, but ultimately it was too risky.

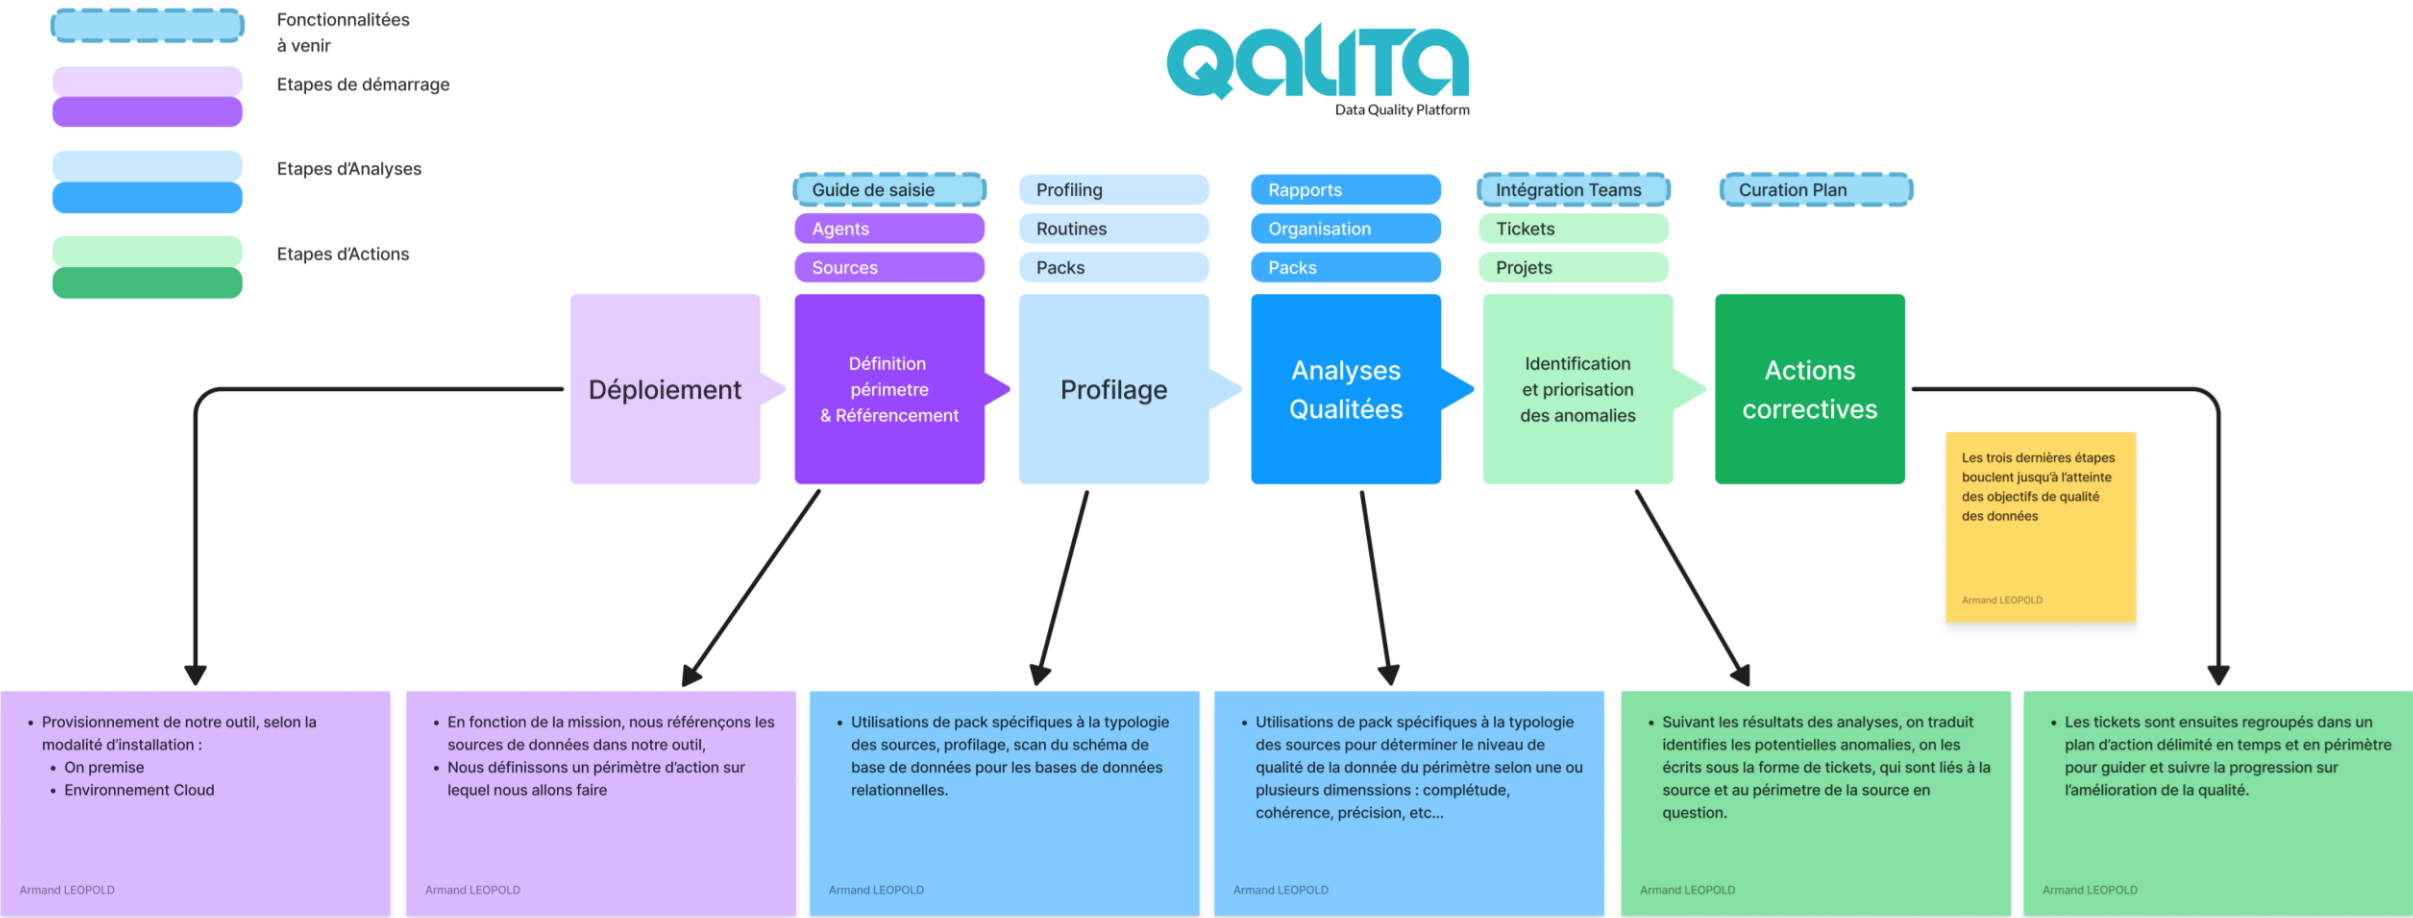
How do we apply Data Quality ?

1. Cut the problem into smaller pieces




How do we apply Data Quality ?

2. Follow a workflow

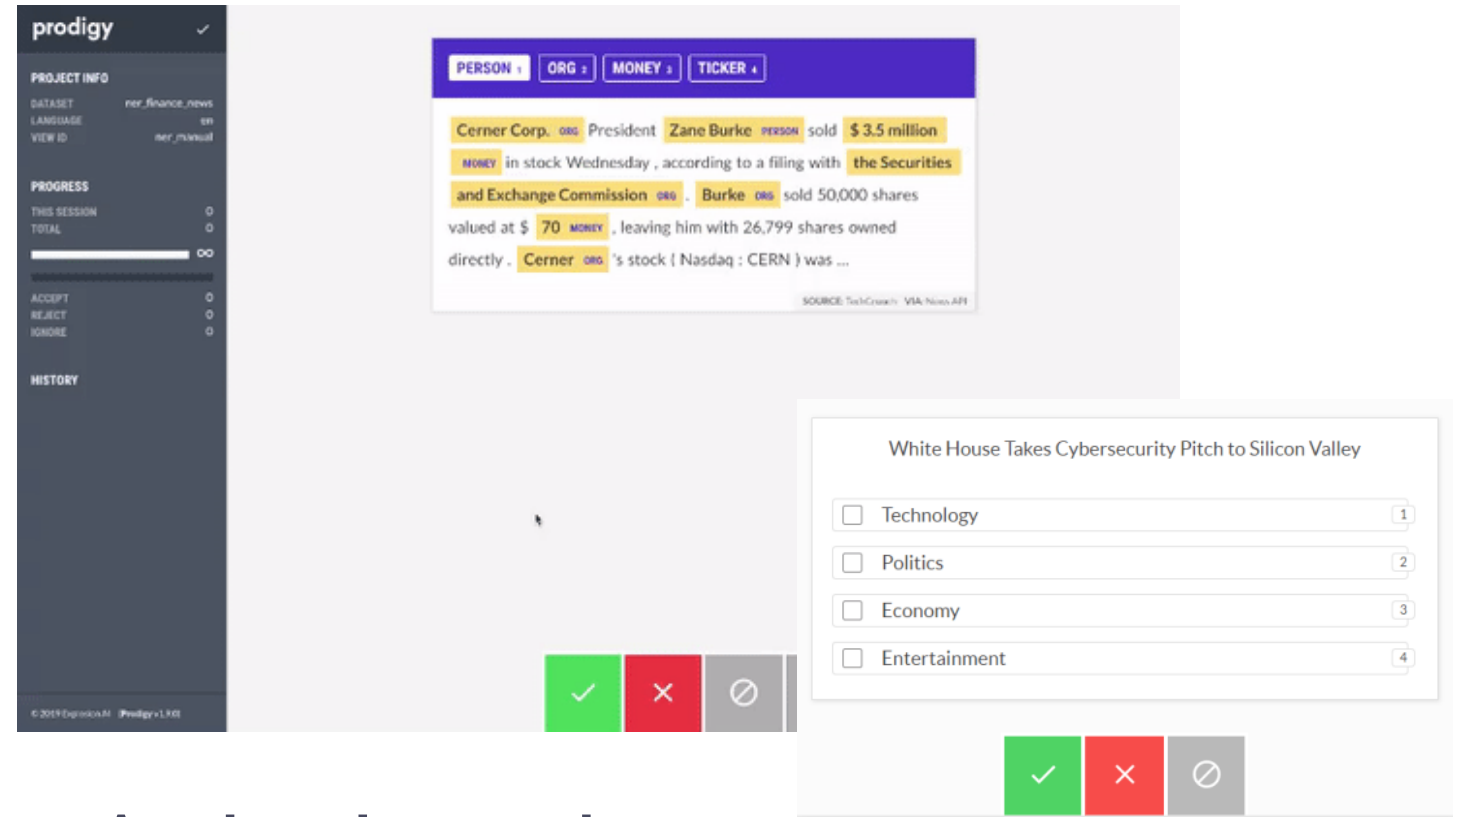


How do we apply Data Quality ?

3. Keep Human in the loop



Governance
& Ownership



Active Learning



Key Takeaways

- Data quality is **crucial** for a **reliable** AI model.
- **Poor** data quality leads to **errors** and **bias**.
- Consequences are often **indirect**.
- The importance of **humans** in **validating** decisions made by machines.

Ensure **data quality assessment**
before (or as part of) launching
any AI project.

Thank You !



Retrouvez-nous à notre stand dans le hall d'exposition !

Sources

- [Intelligence artificielle en entreprise : du fantasme à la réalité](#)
- [How to manage a Data Science project in Health](#)
- [Understanding the different types of artificial intelligence](#)
- [Data Quality in Imitation Learning](#)
- [What is data quality?](#)
- [The Effect of Training and Testing Process on Machine Learning in Biomedical Datasets](#)
- [12 famous AI disasters](#)
- [Data Management body of knowledge](#)
- [Data Science and AI in the age of COVID-19](#)
- [What is precision medicine ?](#)
- [A Deep Learning Mammography-based Model for Improved Breast Cancer Risk Prediction](#)
- [Active Learning: Strategies, Tools, and Real-World Use Cases](#)
- 2024 - Laurent Haug - UNIL - Université de Lausanne
- 2019 - Pierre-Julien Grizel - AlforHealth
- 2023 - IBM
- 2024 - Stanford University
- 2024 - Precisely
- 2020 - Muhammed Kürsad Uçar - King Abdulaziz University
- 2024 - Thor Olavsrud
- 2018 - Rev 2024 DAMA International
- 2021 - The Alan Turing Institute
- 2024 - The European Federation of Pharmaceutical Industries and Associations (EFPIA)
- 2019 - Radiological Society of North America (RSNA) - MIT
- 2023 - Arun C John - Neptune.ai

Articles to read

- Challenges : [Le gouvernement annonce un plan cyber pour les hôpitaux après plusieurs piratages d'envergure](#) 2022
- ComK : [Le « data ditch » : de quoi s'agit-il et comment l'éviter ?](#) 2022
- Le Télégramme : [Cyberattaque au CHU de Rennes : « Il y a eu une fuite de données »](#) 2023
- Le monde informatique : [Les 7 piliers pour accélérer l'adoption de l'IA par les développeurs](#) 2024
- Le monde informatique : [Avec la GenAI, Bouygues Telecom conseille ses clients](#) 2024
- Siecle Digital : [Data Quality Management au service de la relation client : quels enjeux et solutions ?](#) 2021
- ComK : [Les entreprises ne sont pas prêtes pour l'IA générative](#) 2024
- ComK : [Piratage de données : quel impact sur les marques ?](#) 2018
- ComK : [Quel est le niveau de « maturité data » des entreprises françaises ?](#) : 2019
- ComK : [1 entreprise sur 2 est incapable d'exploiter ses propres données clients](#) : 2021
- Datascientest : [Datasets : Top 5 des sites où trouver des jeux de données de qualité](#) : 2020
- Datascientest : [Gartner dévoile les secrets pour améliorer la qualité de ses données](#) : 2023