Explainable FinTech

A Transdisciplinary Perspective

Arie van Deursen
Delft University of Technology
@avandeursen@mastodon.acm.org
www.tudelft.nl/fintech/

Image: Wikipedia, Pont Adolphe

Key Take Aways

1. The software- and data-intensive nature of FinTech makes it an exciting domain for AI and software engineering research
2. FinTech research (at TU Delft) spans many disciplines and research groups
3. Future challenges in finance demand transdisciplinary collaboration

Software Engineering Research

[Empirical methods, theory building]
Seek to understand the methods and techniques that collaborating people use to develop software systems that bring value to society

[Design science, interventions]
Use this understanding to propose and evaluate novel software development methods and techniques

SE 💔 AI

SE4AI: Adjust the software development process to the needs of AI-based systems

AI4SE: Augment the software development life cycle with artificial intelligence

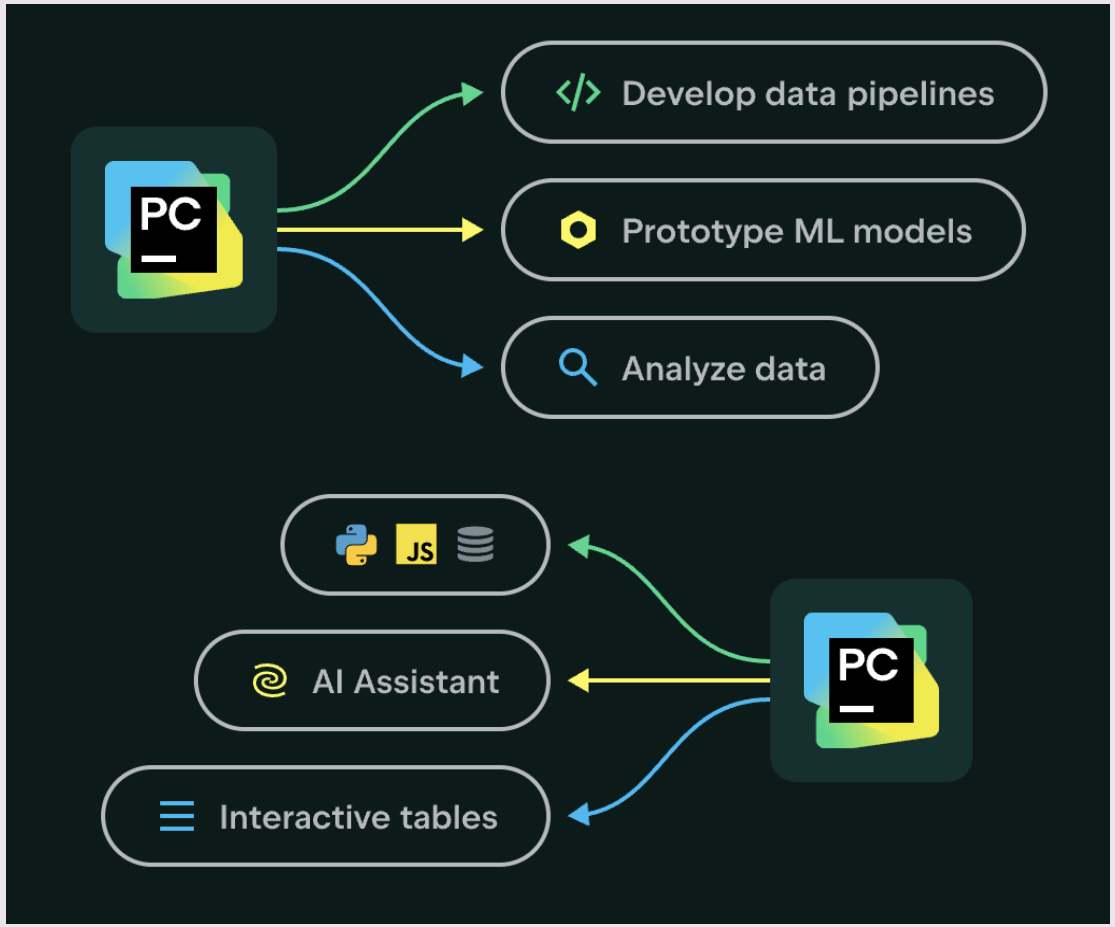

The Financial Sector

- Data intensive
- Software intensive
- High stakes
- Highly regulated
- Long (system, data) lifetimes

High impact societal sector,
with critical software engineering challenges

ING Bank

Global bank based in The Netherlands

Five-year collaboration with TU Delft:

- Explainable AI
- Human-AI decision making
- Data integration
- Incident management and AIOps
- Release planning
- Search-based testing and repair

Agile at Scale at ING

- ING Bank: 15,000 IT staff
- Self-organizing teams
  (5-9 developers)
- Short iterations (1-4 weeks)
- User stories, features, epics
- Delivered in releases (2-6 months)
- Quarterly planning of all releases

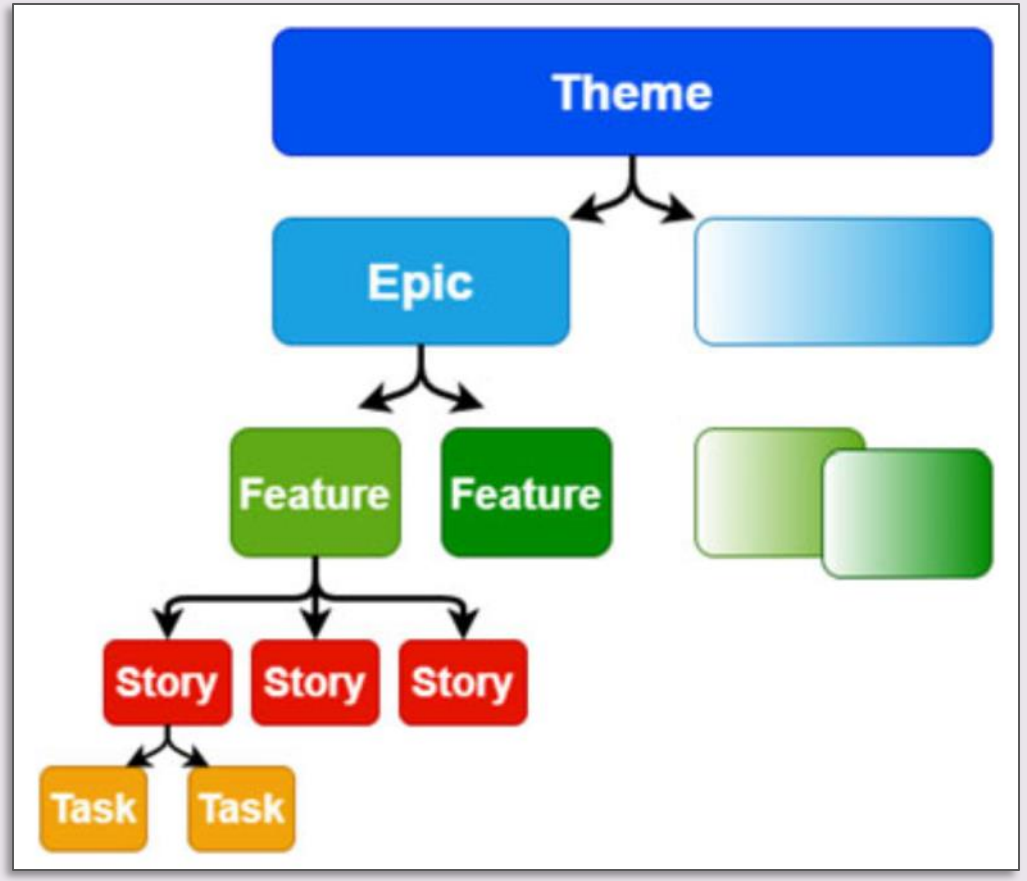

Years of high-quality data available at ING

Why is My Project Late?

What are factors affecting timely epic delivery?

- *Let's ask!*

How do these factors impact schedule deviation?

- *Let's measure and model!*

Elvan Kula et al
IEEE TSE 2022

IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, VOL. 48, NO. 9, SEPTEMBER 2022

Factors Affecting On-Time Delivery in Large-Scale Agile Software Development

Elvan Kula, *Member, IEEE*, Eric Greuter, Arie van Deursen, *Member, IEEE*, and Georgios Gousios

Abstract—Late delivery of software projects and cost overruns have been common problems in the software industry for decades. Both problems are manifestations of deficiencies in effort estimation during project planning. With software projects being complex socio-technical systems, a large pool of factors can affect effort estimation and on-time delivery. To identify the most relevant factors and their interactions affecting schedule deviations in large-scale agile software development, we conducted a mixed-methods case study at ING: two rounds of surveys revealed a multitude of organizational, people, process, project and technical factors which were then quantified and statistically modeled using software repository data from 185 teams. We find that factors such as requirements refinement, task dependencies, organizational alignment and organizational politics are perceived to have the greatest impact on on-time delivery, whereas proxy measures such as project size, number of dependencies, historical delivery performance and team familiarity can help explain a large degree of schedule deviations. We also discover hierarchical interactions among factors: organizational factors are perceived to interact with people factors, which in turn impact technical factors. We compose our findings in the form of a conceptual framework representing influential factors and their relationships to on-time delivery. Our results can help practitioners identify and manage delay risks in agile settings, can inform the design of automated tools to predict schedule overruns and can contribute towards the development of a relational theory of software project management.

Index Terms—Software engineering management, effort estimation, empirical studies, software companies

1 Introduction

Late delivery and cost overruns have been common problems in the software industry for decades. On average, software projects run around 30 percent overtime [1]. This percentage does not seem to have decreased since the 1980s [2]. Even though effort estimation is at the heart of almost all industries, it is especially challenging in the software industry. This is mainly due to the fact that software development is a complex undertaking, affected by a variety of social and technical factors. The overall perceived success of a software project depends heavily on meeting the time and cost estimates [3]. Improving effort estimation is therefore a critical goal for software organizations: it can help companies reduce delays and improve customer satisfaction, while enabling them to efficiently allocate resources, reduce costs and optimize delivery [4], [5]. In spite of the availability of many estimation methods and guidelines [6], [7], on-time delivery in software development remains a major challenge. Prior research identified a large number of factors that may influence the software development effort [8], but which factors have the most impact is not clear. We lack an understanding of the relationships between these factors and how they impact on-time delivery.

Effort estimation is also a major challenge in agile software development. Prior work [9] has found that around half of the agile projects run into effort overruns of 25 percent or more. In agile settings, software is incrementally developed through short iterations to enable a fast response to changing markets and customer demands. Agile projects leverage short-term, iterative planning in which effort estimates are progressively refined [10]. A particular challenge involves combining the flexible, short-term agile planning setting with the business needs for long term planning of availability of large pieces of functionality (often referred to as "epics" [11]). Most agile teams heavily rely on experts' subjective assessment of team- and project-related factors to arrive at an estimate [12], [13]. However, these factors remain largely unexplored [13]; further analysis is required to investigate influential factors and how they impact delays in agile projects.

By identifying and investigating influential factors, we can obtain valuable insights on what data and techniques are needed to become more predictable at delivering software in agile settings. An identification of the most influential factors can help software organizations increase the effectiveness and efficiency of scheduling strategies by concentrating measurement and risk management activities directly on those factors that have the greatest impact on on-time delivery. Such knowledge can also guide future research on building and evaluating software effort estimation techniques, methods and tools. Furthermore, a deeper understanding of the interactions between influential factors can help in identifying the root causes of delays, and developing tools and

- *Elvan Kula is with the Delft University of Technology, 2628, CD, Delft, Netherlands, and also with the ING Tech, 1102, MG, Amsterdam, The Netherlands. E-mail: E.Kula@tudelft.nl.*
- *Eric Greuter is with the ING Tech, 1102, MG, Amsterdam, The Netherlands. E-mail: Eric.Greuter@ing.com.*
- *Arie van Deursen and Georgios Gousios are with the Delft University of Technology, 2628 Delft, The Netherlands. E-mail: {Arie.vanDeursen, G.Gousios}@tudelft.nl.*

Manuscript received 14 Oct. 2020; revised 16 June 2021; accepted 13 July 2021. Date of publication 2 Aug. 2021; date of current version 19 Sept. 2022. (Corresponding author: Elvan Kula.) Recommended for acceptance by M. P. Robillard. Digital Object Identifier no. 10.1109/TSE.2021.3101192

Timely Epic Delivery: *Perceived* Factors

Survey 1: *Which factors?*

- 289 responses
- 25 factors; 5 dimensions

Survey 2: *Factor importance?*

- 337 responses
- Rated impact level per factor

Factor top 10:

1. Requirements refinement
2. Task dependencies
3. Organizational alignment
4. Organizational politics
5. Geographic distribution
6. Technical dependencies
7. Agile maturity
8. Regular delivery
9. Team stability
10. Skills and knowledge

Measuring Delay: *Balanced Relative Error*

- If actual delivery date *after* estimated date (”late”, pos%):

$$\text{BRE} = \frac{\text{Act - Est}}{\text{Planned duration}}$$

- If actual delivery date *before* estimated date (“early”, neg%):

$$\text{BRE} = \frac{\text{Act - Est}}{\text{Actual duration}}$$

- Collected BRE from **3,771 epics** (273 teams), for 3 years

13 Predictor Variables

- 35 metrics for 20 factors
- 13 metrics explain 67% of variation (MARS model, R^2)
- Match with perception?
  - **Underestimated**: *size*
  - **Agreed effect**: *dependencies, seniority, stability*
  - **Overestimated**: *refinement*, *geography,*
  - **Agreed little effect**: *coverage, code smells*, …

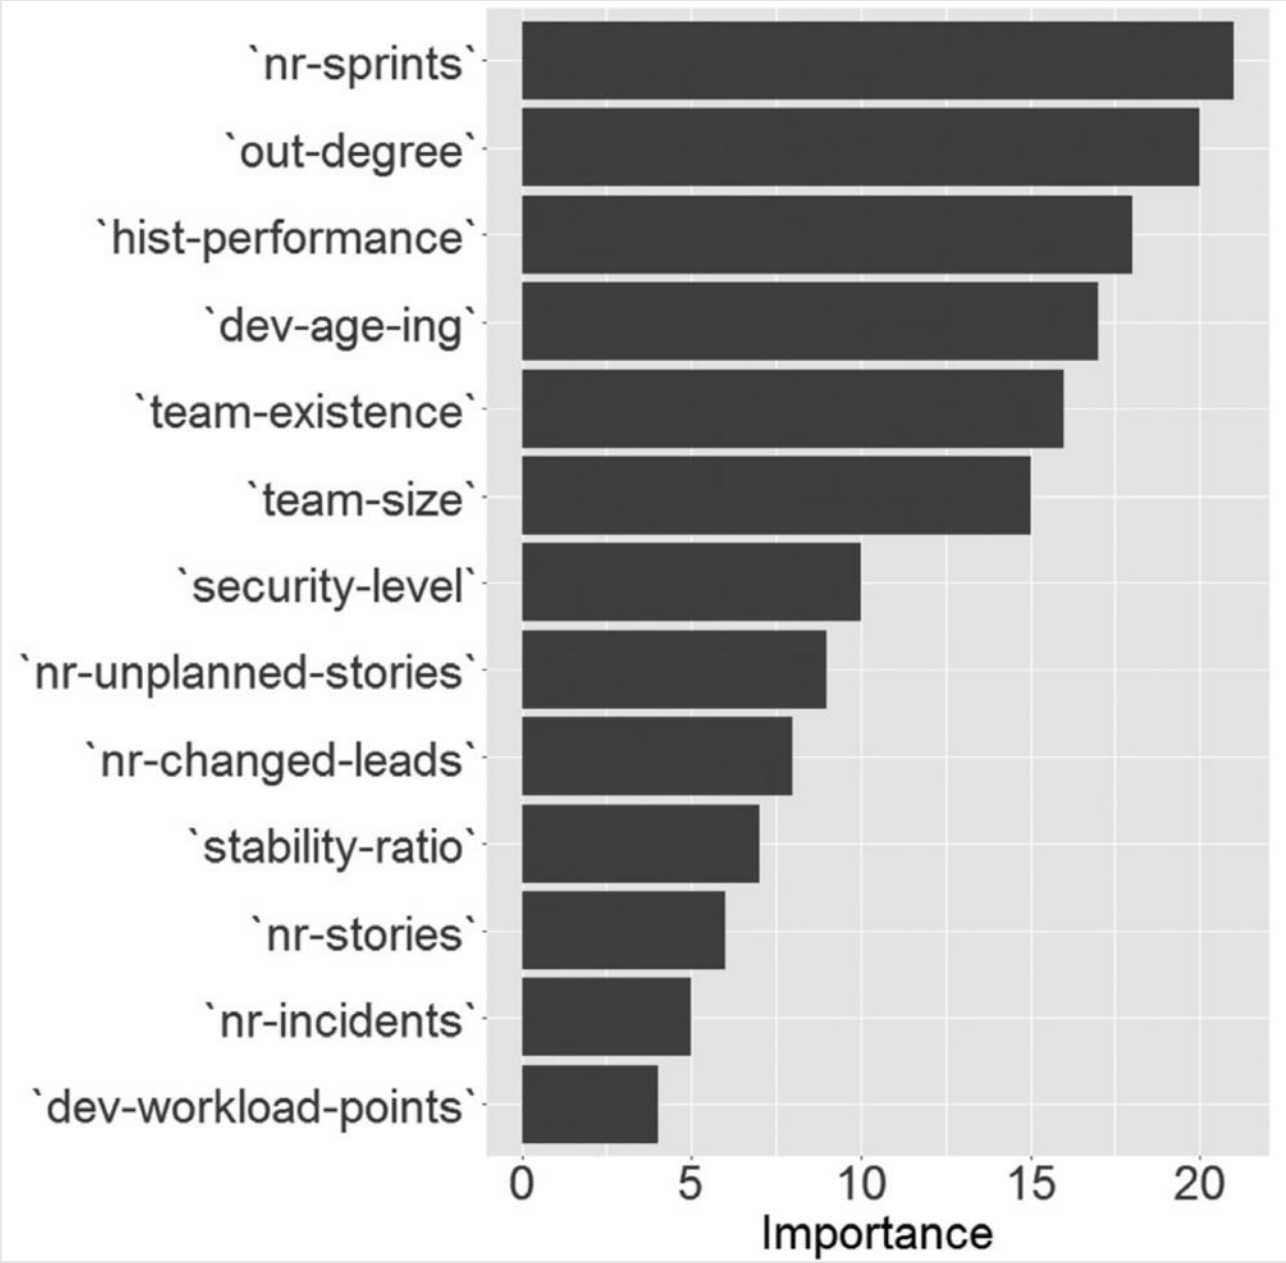

Can we Predict Delay?

- Delay knowledge increases
  as epic unfolds (in milestones)
- Mobility literature:
  Delay adheres to patterns, which can be
  learned by clustering delay time series

- *Is epic delay subject to patterns?*
- *Can patterns improve delay prediction?*

Elvan Kula et al
FSE 2023

Epic Delay Patterns

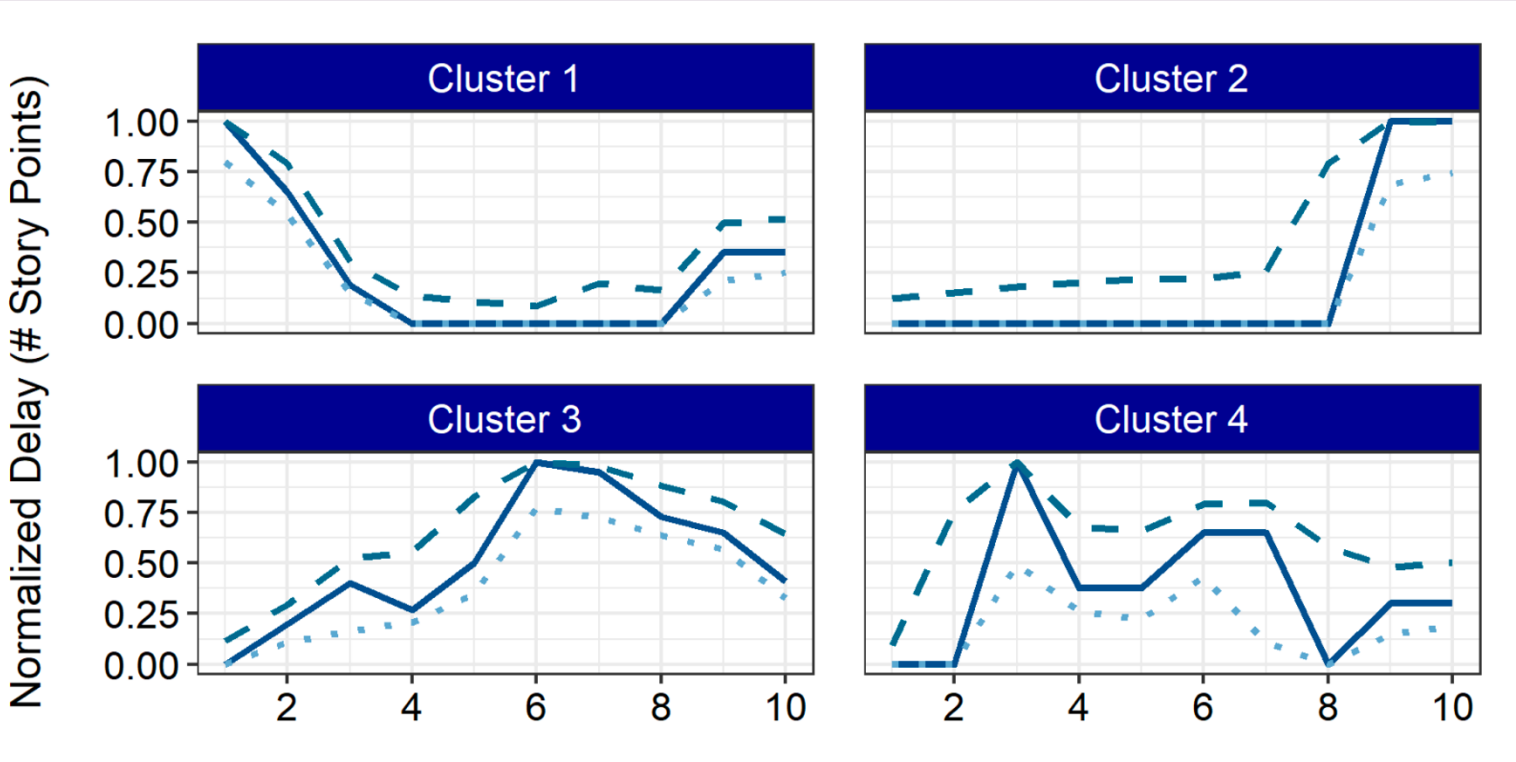

Predictor variable	Median			
	C1	C2	C3	C4
nr-sprints	13	15	14	11
out-degree	7	3	4	4
hist-performance	0.69	0.67	0.74	0.61
dev-age-abc	2.49	2.61	2.92	2.84
team-existence	1.30	1.53	1.29	1.42
team-size	8	7	6	7
security-level	0.56	0.77	0.53	0.36
nr-unplanned-stories (ratio)	0.11	0.16	0.10	0.08
changed-leads	3	2	3	2
stability-ratio	0.73	0.81	0.64	0.72
nr-stories	52	43	39	45
nr-incidents	8	12	8	6
dev-workload-points	15	12	10	8
BRE	0.23	0.17	0.11	0.09
% epics in category:	36%	44%	14%	6%

Dataset: 4,040 epics of at least 10 sprints from 270 teams, 2017—2022

Delay Patterns Improve Delay Prediction

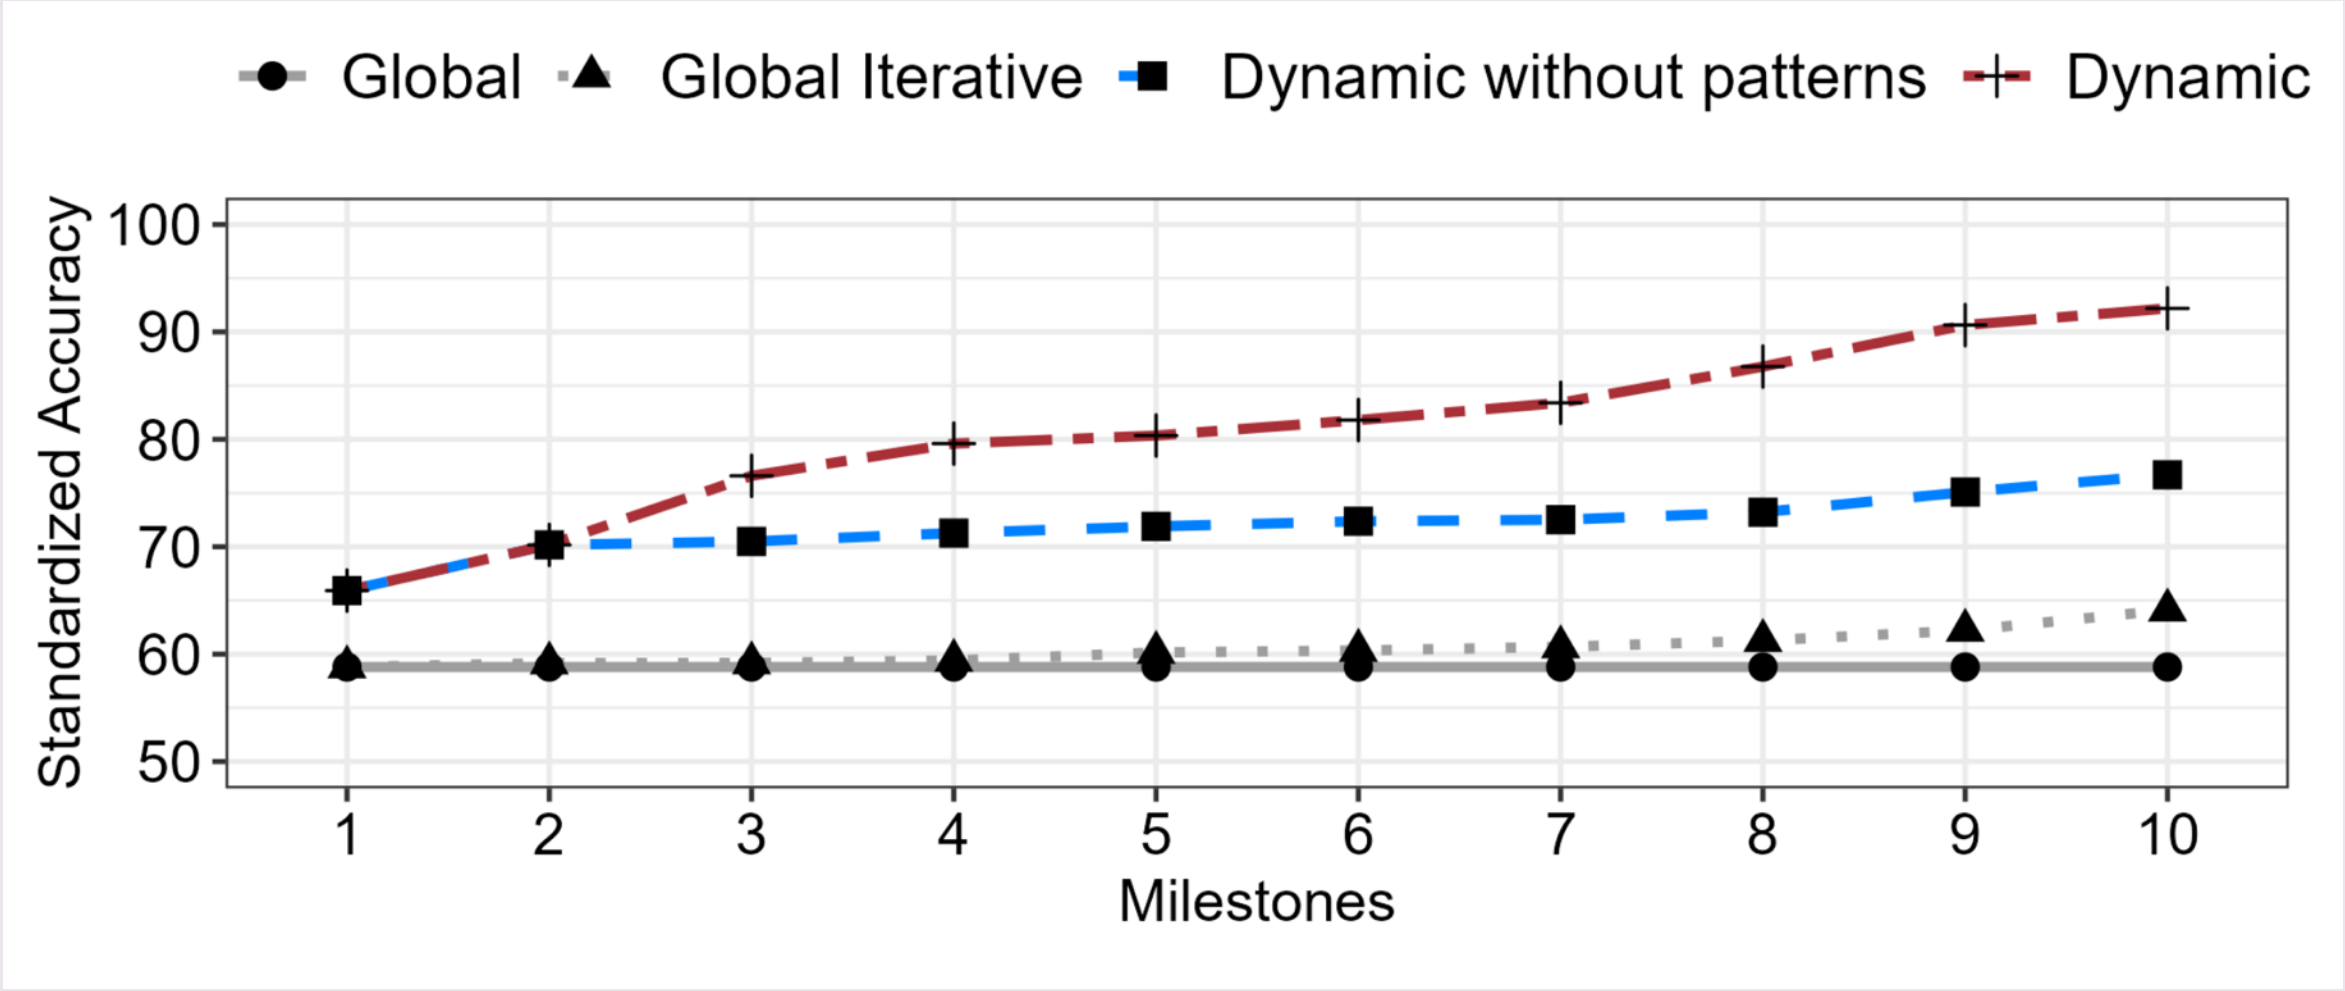

Epic Conclusions

- There are measurable factors contributing to epic delay
  - Size, project dependencies, past performance
- Delay follows patterns
  - Largest pattern is timely at start with delay peak at end, due to security and incidents
- Factors + patterns predict delay, dynamically
  - Beats the global and iterative state-of-the-art baselines

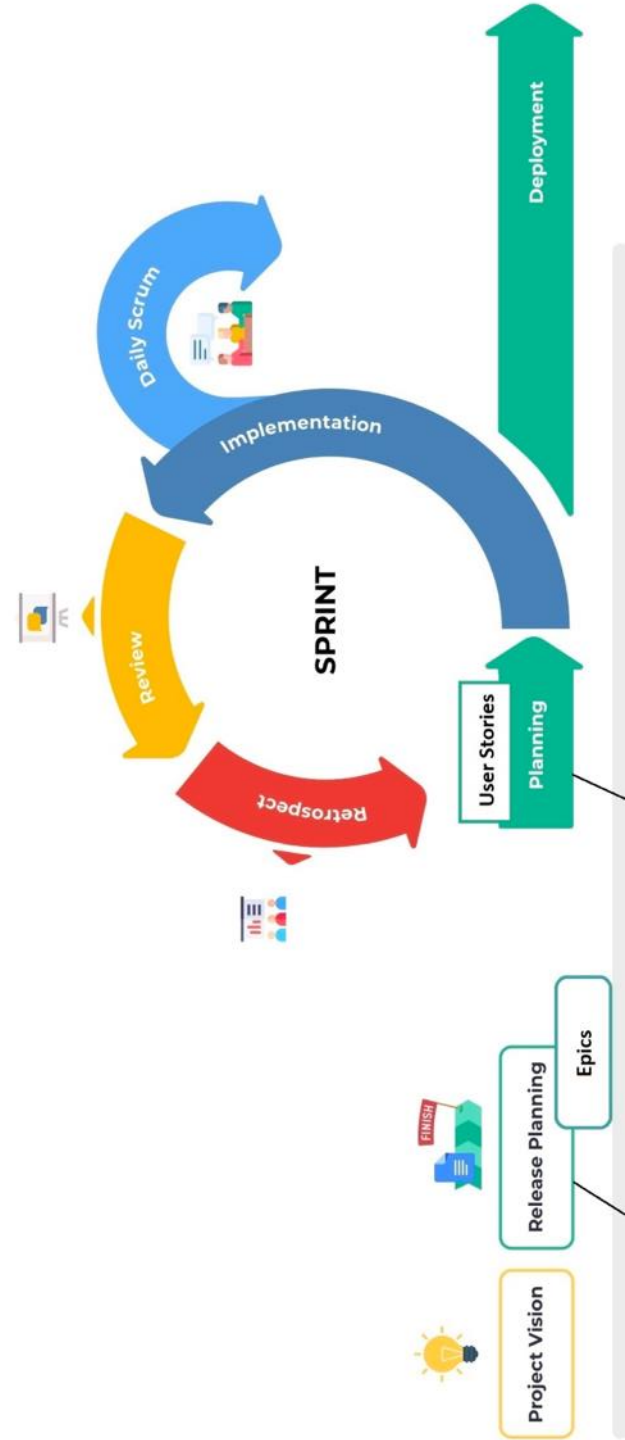

27.12.2022 EN Official Journal of the European Union L 333/1

I

(Legislative acts)

REGULATIONS

REGULATION (EU) 2022/2554 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL

of 14 December 2022

on digital operational resilience for the financial sector and amending Regulations (EC) No 1060/2009, (EU) No 648/2012, (EU) No 600/2014, (EU) No 909/2014 and (EU) 2016/1011

(Text with EEA relevance)

THE EUROPEAN PARLIAMENT AND THE COUNCIL OF THE EUROPEAN UNION,

Having regard to the Treaty on the Functioning of the European Union, and in particular Article 114 thereof,

Having regard to the proposal from the European Commission,

After transmission of the draft legislative act to the national parliaments,

Having regard to the opinion of the European Central Bank (1),

Having regard to the opinion of the European Economic and Social Committee (2),

Acting in accordance with the ordinary legislative procedure (3),

Whereas:

(1) In the digital age, information and communication technology (ICT) supports complex systems used for everyday activities. It keeps our economies running in key sectors, including the financial sector, and enhances the functioning of the internal market. Increased digitalisation and interconnectedness also amplify ICT risk, making society as a whole, and the financial system in particular, more vulnerable to cyber threats or ICT disruptions. While the ubiquitous use of ICT systems and high digitalisation and connectivity are today core features of the activities of Union financial entities, their digital resilience has yet to be better addressed and integrated into their broader operational frameworks.

(2) The use of ICT has in the past decades gained a pivotal role in the provision of financial services, to the point where it has now acquired a critical importance in the operation of typical daily functions of all financial entities. Digitalisation now covers, for instance, payments, which have inc methods to the use of digital solutions, as well as securities clearin trading, lending and funding operations, peer-to-peer finance, cred operations. The insurance sector has also been transformed by the

(1) OJ C 343, 26.8.2021, p. 1.
(2) OJ C 155, 30.4.2021, p. 38.
(3) Position of the European Parliament of 10 November 2022 (not yet published in 28 November 2022.

Fines of up to 2% of annual turnover

EU Digital Operational Resilience Act (DORA)

- Harmonized rules for safeguarding against ICT-related **incidents** in financial sector
- Insist on documented **policies** for protection, detection, containment, recovery, and repair
- All **changes** to be recorded, tested, assessed, approved, implemented, and verified in a controlled manner

Incident Management at ING

- ”ITIL” process with four stages:
  - Incident logging;
  - Investigation & diagnosis;
  - Resolution
  - Verification & closure
- Compliance with DORA, PSD2, …
- But does it work well?
  - Interview study with 15 ING experts

E. Kapel et al. “Enhancing Incident Management: Insights from a Case Study at ING”. ACM/IEEE FinanSE 2024.

“A client prefers receiving a very generic message quickly than waiting for 30 minutes for a detailed message.” (P13)

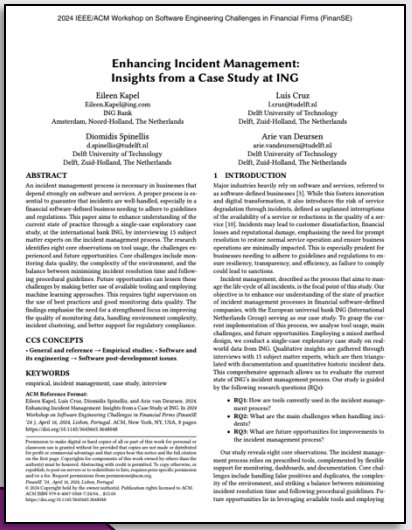

2024 IEEE/ACM Workshop on Software Engineering Challenges in Financial Firms (FinanSE)

Enhancing Incident Management: Insights from a Case Study at ING

Eileen Kapel
Eileen.Kapel@ing.com
ING Bank
Amsterdam, Noord-Holland, The Netherlands

Luis Cruz
l.cruz@tudelft.nl
Delft University of Technology
Delft, Zuid-Holland, The Netherlands

Dionidis Spinellis
d.spinellis@tudelft.nl
Delft University of Technology
Delft, Zuid-Holland, The Netherlands

Arie van Deursen
arie.vandeursen@tudelft.nl
Delft University of Technology
Delft, Zuid-Holland, The Netherlands

ABSTRACT

An incident management process is necessary in businesses that depend strongly on software and services. A proper process is essential to guarantee that incidents are well-handled, especially in a financial software-defined business needing to adhere to guidelines and regulations. This paper aims to enhance understanding of the current state of practice through a single-case exploratory case study, at the international bank ING, by interviewing 15 subject matter experts on the incident management process. The research identifies eight core observations on tool usage, the challenges experienced and future opportunities. Core challenges include monitoring data quality, the complexity of the environment, and the balance between minimising incident resolution time and following procedural guidelines. Future opportunities can lessen these challenges by making better use of available tooling and employing machine learning approaches. This requires tight supervision on the use of best practices and good monitoring data quality. The findings emphasise the need for a strengthened focus on improving the quality of monitoring data, handling environment complexity, incident clustering, and better support for regulatory compliance.

CCS CONCEPTS

• **General and reference** → **Empirical studies**; • **Software and its engineering** → **Software post-development issues**.

KEYWORDS

empirical, incident management, case study, interview

ACM Reference Format:

Eileen Kapel, Luis Cruz, Dionidis Spinellis, and Arie van Deursen. 2024. Enhancing Incident Management: Insights from a Case Study at ING. In *2024 Workshop on Software Engineering Challenges in Financial Firms (FinanSE '24), April 16, 2024, Lisbon, Portugal.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3640660.3640848

FinanSE '24, April 16, 2024, Lisbon, Portugal

ACM ISBN 979-8-4007-0568-7/24/04…$15.00
https://doi.org/10.1145/3640660.3640848

1 INTRODUCTION

Major industries heavily rely on software and services, referred to as software-defined businesses [3]. While this fosters innovation and digital transformation, it also introduces the risk of service degradation through incidents, defined as unplanned interruptions of the availability of a service or reductions in the quality of a service [10]. Incidents may lead to customer dissatisfaction, financial losses and reputational damage, emphasising the need for prompt resolution to restore normal service operation and ensure business operations are minimally impacted. This is especially prudent for businesses needing to adhere to guidelines and regulations to ensure resiliency, transparency, and efficiency, as failure to comply could lead to sanctions.

Incident management, described as the process that aims to manage the life-cycle of all incidents, is the focal point of this study. Our objective is to enhance our understanding of the state of practice of incident management processes in financial software-defined companies, with the European universal bank ING (International Netherlands Group) serving as our case study. To grasp the current implementation of this process, we analyse tool usage, main challenges, and future opportunities. Employing a mixed method design, we conduct a single-case exploratory case study on real-world data from ING. Qualitative insights are gathered through interviews with 15 subject matter experts, which are then triangulated with documentation and quantitative historic incident data. This comprehensive approach allows us to evaluate the current state of ING's incident management process. Our study is guided by the following research questions (RQs):

- **RQ1:** How are tools currently used in the incident management process?
- **RQ2:** What are the main challenges when handling incidents?
- **RQ3:** What are future opportunities for improvements to the incident management process?

Our study reveals eight core observations. The incident management process relies on prescribed tools, complemented by flexible support for monitoring, dashboards, and documentation. Core challenges include handling false positives and duplicates, the complexity of the environment, and striking a balance between minimising incident resolution time and following procedural guidelines. Future opportunities lie in leveraging available tools and employing

“Even if we understand the chain today, it will be different in a month” (P9)

“If you have long overdue incidents then you are not in control of your incident process and are at risk. Then we do not comply with the regulations of the European Bank that we should be in control.” (P1)

Observations & Recommendations


2024 IEEE/ACM Workshop on Software Engineering Challenges in Financial Firms (FinanSE)

Enhancing Incident Management: Insights from a Case Study at ING

Eileen Kapel
Eileen.Kapel@ing.com
ING Bank
Amsterdam, Noord-Holland, The Netherlands

Luís Cruz
l.cruz@tudelft.nl
Delft University of Technology
Delft, Zuid-Holland, The Netherlands

Dionidis Spinellis
d.spinellis@tudelft.nl
Delft University of Technology
Delft, Zuid-Holland, The Netherlands

Arie van Deursen
arie.vandeursen@tudelft.nl
Delft University of Technology
Delft, Zuid-Holland, The Netherlands

ABSTRACT

An incident management process is necessary in businesses that depend strongly on software and services. A proper process is essential to guarantee that incidents are well-handled, especially in a financial software-defined business needing to adhere to guidelines and regulations. This paper aims to enhance understanding of the current state of practice through a single-case exploratory case study, at the international bank ING, by interviewing 15 subject matter experts on the incident management process. The research identifies eight core observations on tool usage, the challenges experienced and future opportunities. Core challenges include monitoring data quality, the complexity of the environment, and the balance between minimising incident resolution time and following procedural guidelines. Future opportunities can lessen these challenges by making better use of available tooling and employing machine learning approaches. This requires tight supervision on the use of best practices and good monitoring data quality. The findings emphasise the need for a strengthened focus on improving the quality of monitoring data, handling environment complexity, incident clustering, and better support for regulatory compliance.

CCS CONCEPTS

• **General and reference** → **Empirical studies**; • **Software and its engineering** → **Software post-development issues**.

KEYWORDS

empirical, incident management, case study, interview

ACM Reference Format:

Eileen Kapel, Luís Cruz, Dionidis Spinellis, and Arie van Deursen. 2024. Enhancing Incident Management: Insights from a Case Study at ING. In *2024 Workshop on Software Engineering Challenges in Financial Firms (FinanSE '24), April 16, 2024, Lisbon, Portugal.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3640665.3640848

ACM ISBN 979-8-4007-0568-7/24/04…$15.00
https://doi.org/10.1145/3640665.3640848

1 INTRODUCTION

Major industries heavily rely on software and services, referred to as software-defined businesses [3]. While this fosters innovation and digital transformation, it also introduces the risk of service degradation through incidents, defined as unplanned interruptions of the availability of a service or reductions in the quality of a service [10]. Incidents may lead to customer dissatisfaction, financial losses and reputational damage, emphasising the need for prompt resolution to restore normal service operation and ensure business operations are minimally impacted. This is especially prudent for businesses needing to adhere to guidelines and regulations to ensure resiliency, transparency, and efficiency, as failure to comply could lead to sanctions.

Incident management, described as the process that aims to manage the life-cycle of all incidents, is the focal point of this study. Our objective is to enhance our understanding of the state of practice of incident management processes in financial software-defined companies, with the European universal bank ING (International Netherlands Group) serving as our case study. To grasp the current implementation of this process, we analyse tool usage, main challenges, and future opportunities. Employing a mixed method design, we conduct a single-case exploratory case study on real-world data from ING. Qualitative insights are gathered through interviews with 15 subject matter experts, which are then triangulated with documentation and quantitative historic incident data. This comprehensive approach allows us to evaluate the current state of ING's incident management process. Our study is guided by the following research questions (RQs):

- **RQ1:** How are tools currently used in the incident management process?
- **RQ2:** What are the main challenges when handling incidents?
- **RQ3:** What are future opportunities for improvements to the incident management process?

Our study reveals eight core observations. The incident management process relies on prescribed tools, complemented by flexible support for monitoring, dashboards, and documentation. Core challenges include handling false positives and duplicates, the complexity of the environment, and striking a balance between minimising incident resolution time and following procedural guidelines. Future opportunities lie in leveraging available tools and employing

1. Demonstrable regulatory compliance is key driver of process
2. Logged incidents often are duplicates or false alarms
3. Rapid evolution of bank’s IT systems complicate diagnosis
4. Strict access rules hamper rapid incident resolution
5. Incident resolution is prioritized over structural fix creation
6. Communication across teams with all affected parties is key
7. Data-driven approaches (anomaly detection, pattern recognition, clustering) demand clean monitoring data and tight supervision
8. Incorporating human oversight is essential when implementing automated resolutions to support the incident management process

Work in Progress: Tracing Incidents & Changes

Six months of change data:

Change Records		
Close code	**Count**	**Perc.**
Closed successfully	86,000	95%
Cancelled	4,000	4%
Induced incident	1,000	1%
Total	**91,000**	**100%**

On the Difficulty of Identifying Incident-Inducing Changes

Eileen Kapel
Eileen.Kapel@ing.com
ING Bank
Amsterdam, Noord-Holland, The Netherlands

Luís Cruz
l.cruz@tudelft.nl
Delft University of Technology
Delft, Zuid-Holland, The Netherlands

Diomidis Spinellis
d.spinellis@tudelft.nl
Delft University of Technology
Delft, Zuid-Holland, The Netherlands

Arie van Deursen
arie.vandeursen@tudelft.nl
Delft University of Technology
Delft, Zuid-Holland, The Netherlands

ABSTRACT

Effective change management is crucial for businesses heavily reliant on software and services to minimise incidents induced by changes. Unfortunately, in practice it is often difficult to effectively use artificial intelligence for IT Operations (AIOps) to enhance service management, primarily due to inadequate data quality. Establishing reliable links between changes and the induced incidents is crucial for identifying patterns, improving change deployment, identifying high-risk changes, and enhancing incident response. In this research, we investigate the enhancement of traceability between changes and incidents through AIOps methods. Our approach involves a close examination of incident-inducing changes, the replication of methods linking incidents to the changes that caused them, introducing an adapted method, and demonstrating its results using historical data and practical evaluations. Our findings reveal that incident-inducing changes exhibit different characteristics dependent on context. Furthermore, a significant disparity exists between assessments based on historical data and real-world observation, with an increased occurrence of false positives when identifying links between unlabeled changes and incidents. This study highlights the complex nature of identifying links between changes and incidents, emphasising the contextual influence on AIOps method effectiveness. While we are actively working on improving the quality of current data through AIOps approaches, it remains apparent that further measures are necessary to address issues like data imbalances and promote a postmortem culture that brings attention to the value of properly administrating tickets. A better overview of change failure rates contributes to improved risk compliance and reliable change management.

CCS CONCEPTS

• **Software and its engineering → Risk management; Software post-development issues; Software reliability.**

ICSE 2024, April 2024, Lisbon, Portugal

ACM ISBN 979-8-4007-0501-4/24/04…$15.00
https://doi.org/10.1145/3639477.3639755

KEYWORDS

change management, incident management, traceability

ACM Reference Format:

Eileen Kapel, Luís Cruz, Diomidis Spinellis, and Arie van Deursen. 2024. On the Difficulty of Identifying Incident-Inducing Changes. In *46th International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP '24), April 14–20, 2024, Lisbon, Portugal.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3639477.3639755

1 INTRODUCTION

Nowadays, major businesses and industries, such as banking, healthcare and retail, are increasingly reliant on software and related services, often referred to as software-defined businesses [1]. This approach emphasises an agile way of working, encouraging fast-paced development and deployment of new features. While this offers many benefits, it also leads to a higher volume and faster pace of change deployment. In this context, changes are defined as modifications to existing applications, comprising additions, alterations and deletions [18].

Changes can occasionally trigger incidents, especially if they undergo inadequate testing or incorrect implementation. For instance, a software upgrade might lead to compatibility issues, resulting in an incident. The change management process aims to minimise change-related risks, thereby reducing the occurrence of incidents caused by changes. An incident is defined as an unplanned interruption to a service or a reduction in service quality at a specific time [18]. Incidents can lead to customer dissatisfaction, financial losses, and reputational damage; therefore, they must be prevented. IT service changes are one of the leading contributors to outages, accounting for about 70% of live system outages [5]. This is often referred to as the change failure rate, which measures the percentage of deployments causing a failure in production [15]. Other factors that contribute to the difficulty of identifying incident-inducing changes include the complex nature of large IT environments [6], the potential for seemingly successful changes to still induce incidents [16, 38], the service where an incident begins may be different from the once that caused it [36], the necessity for engineers to sift through large amounts of heterogeneous data to identify the root cause [36], and the fact that incidents are often the result of changes occurring hours or days before the incident [17].

To enable the application of Artificial Intelligence for IT Operations (AIOps), which utilises big data, machine learning and advanced analytics to enhance IT operations [21], proper data quality is crucial to systemically learn from the past and ensure effective

E. Kapel et al.
"On the Difficulty of Identifying Incident-Inducing Changes".
ICSE SEIP 2024.

Riskier Changes in Weekend

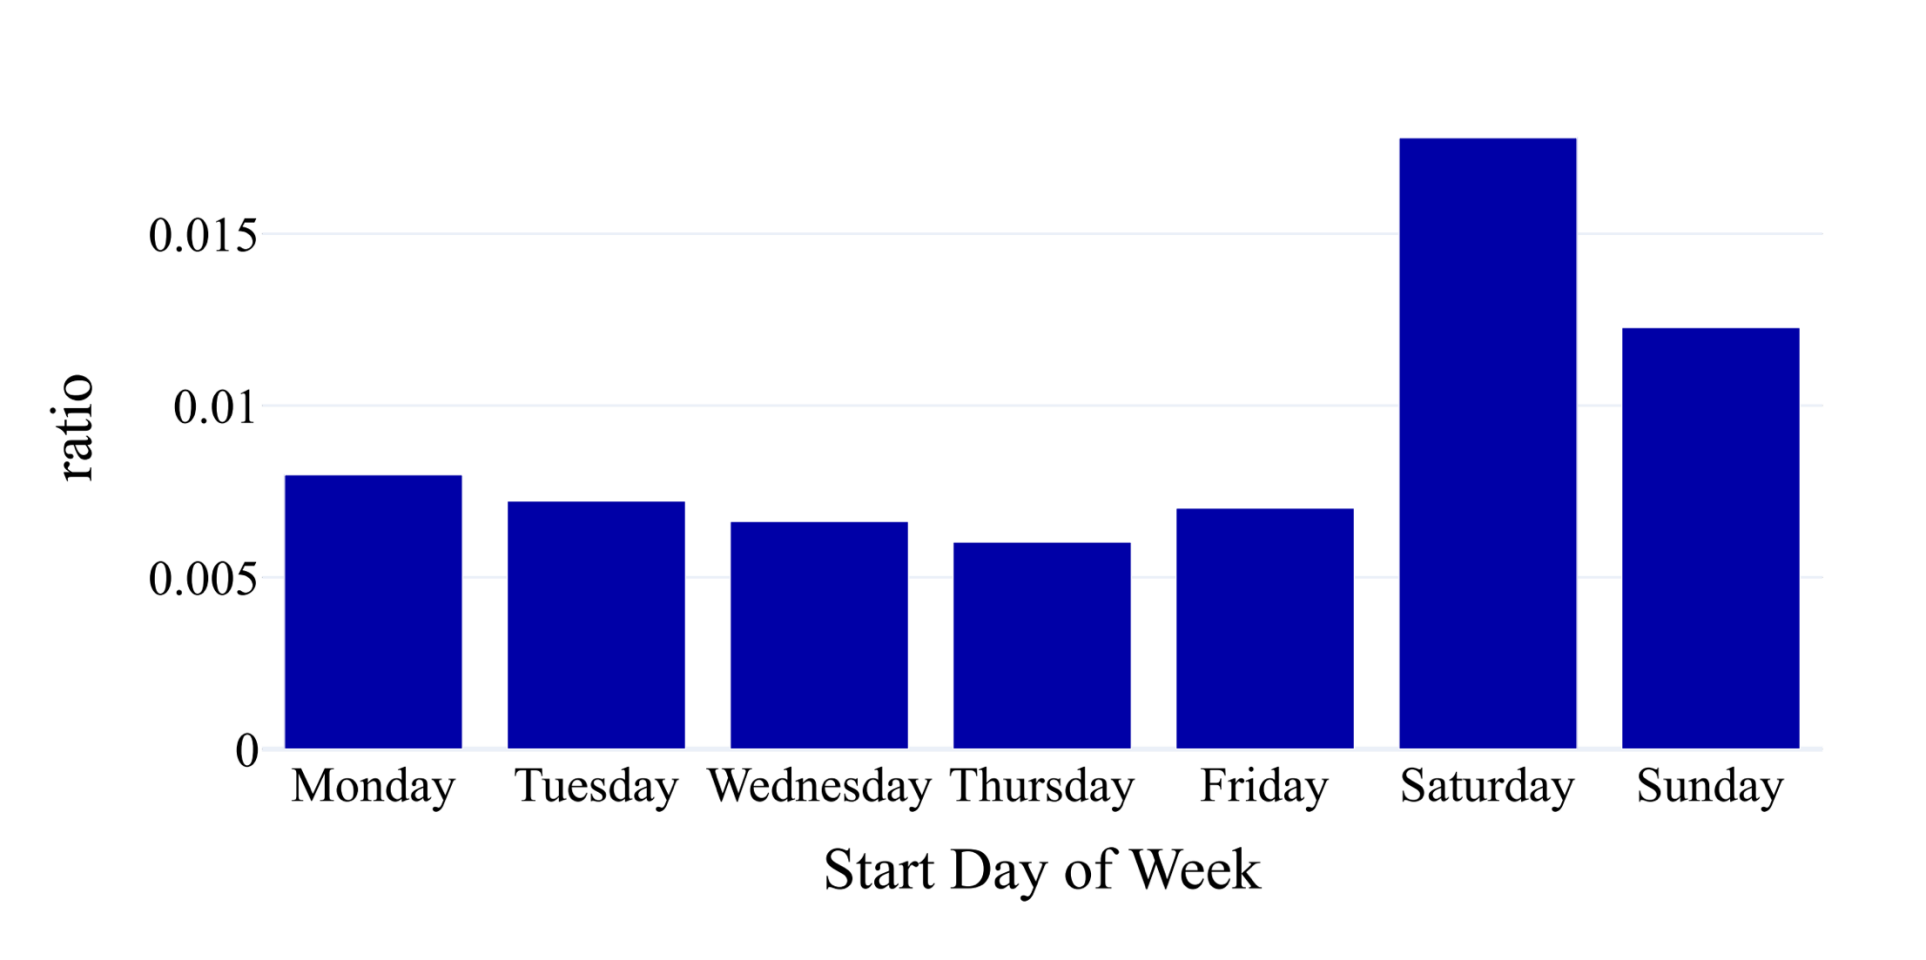

Ratio between incident-inducing changes and non-incident inducing changes per day of the week.

Higher Upfront “risk-category” for Incident-Causing Changes

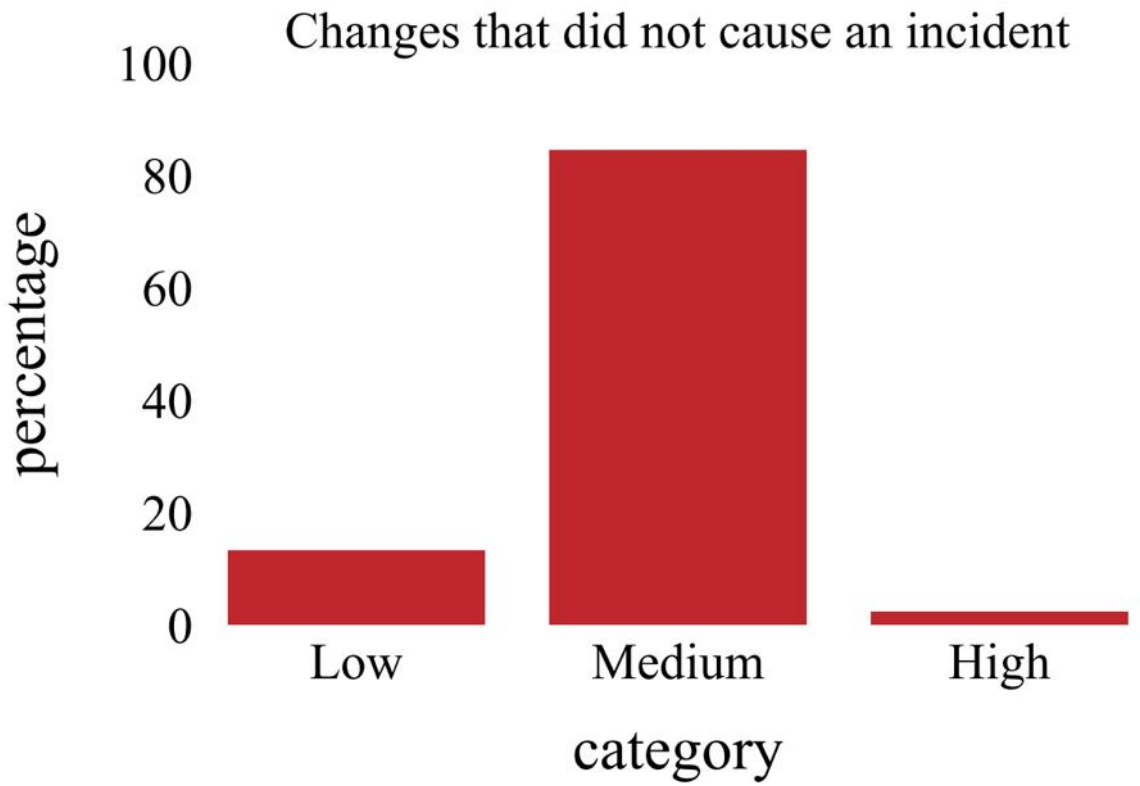

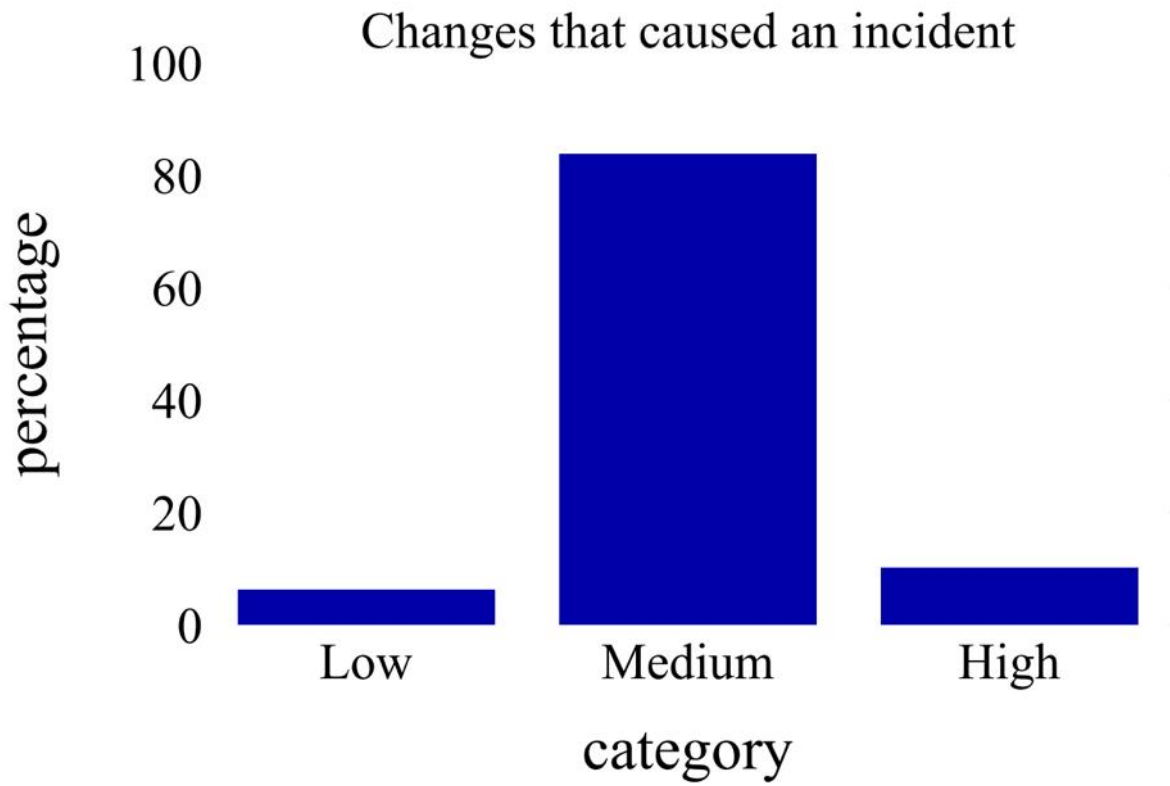

From Incident back to Change?

- Not all incidents related to change
- Links missing in many cases:
  - Focus on resolution, not on documentation
- Can we establish such links *automatically*?
  - Try data mining approach published by IBM
  - Time: most recent fix before incident?
  - Time with shared ‘dimensions’ (words, impact, type, group, doubles accuracy?
  - Correct change in top 5 in 56.9% of cases

Original		
Dimension	top 5	top 1
Time	52%	30%
All	67%	51%
Significant	75%	58%

Replication		
Dimension	top 5	top 1
Time	35.3%	9.8%
All	45.1%	11.8%
Significant	56.9%	22.2%

Current Quest: Incident Prediction

Collect	Train	Score	Explain
CHG123 with description “Updating DB” on CI “network” caused a priority 2 incident.	Changes on a network CI with description “Updating DB” are often causing incidents.	CHG234 has an estimated probability of 80% of causing an incident.	CHG234 is like CHG345 that caused a P1 in the past. Important features are the CI Business Unit, CI Owner and CAB Approval Group.

“Explanation in Artificial Intelligence: Insights from the Social Sciences”

(Tim Miller, Artificial Intelligence, 2018)

An explanation is an answer to a why-question

Explanation in artificial intelligence: Insights from the social sciences

Tim Miller

School of Computing and Information Systems, University of Melbourne, Melbourne, Australia

ARTICLE INFO

Article history:
Received 22 June 2017
Received in revised form 17 May 2018
Accepted 16 July 2018
Available online 27 October 2018

Keywords:
Explanation
Explainability
Interpretability
Explainable AI
Transparency

ABSTRACT

There has been a recent resurgence in the area of explainable artificial intelligence as researchers and practitioners seek to provide more transparency to their algorithms. Much of this research is focused on explicitly explaining decisions or actions to a human observer, and it should not be controversial to say that looking at how humans explain to each other can serve as a useful starting point for explanation in artificial intelligence. However, it is fair to say that most work in explainable artificial intelligence uses only the researchers' intuition of what constitutes a 'good' explanation. There exist vast and valuable bodies of research in philosophy, psychology, and cognitive science of how people define, generate, select, evaluate, and present explanations, which argues that people employ certain cognitive biases and social expectations to the explanation process. This paper argues that the field of explainable artificial intelligence can build on this existing research, and reviews relevant papers from philosophy, cognitive psychology/science, and social psychology, which study these topics. It draws out some important findings, and discusses ways that these can be infused with work on explainable artificial intelligence.

1. Introduction

Recently, the notion of *explainable artificial intelligence* has seen a resurgence, after having slowed since the burst of work on explanation in expert systems over three decades ago; for example, see Chandrasekaran et al. [23], [168], and Buchanan and Shortliffe [14]. Sometimes abbreviated XAI (eXplainable artificial intelligence), the idea can be found in grant solicitations [32] and in the popular press [136]. This resurgence is driven by evidence that many AI applications have limited take up, or are not appropriated at all, due to ethical concerns [2] and a *lack of trust* on behalf of their users [166,101]. The running hypothesis is that by building more transparent, interpretable, or explainable systems, users will be better equipped to understand and therefore trust the intelligent agents [129,25,65].

While there are many ways to increase trust and transparency of intelligent agents, two complementary approaches will form part of many trusted autonomous systems: (1) generating decisions[1] in which one of the criteria taken into account during the computation is how well a human could understand the decisions in the given context, which is often called *interpretability* or *explainability*; and (2) explicitly explaining decisions to people, which we will call *explanation*. Applications of explanation are considered in many sub-fields of artificial intelligence, such as justifying autonomous agent behaviour [129,65], debugging of machine learning models [89], explaining medical decision-making [45], and explaining predictions of classifiers [157].

E-mail address: tmiller@unimelb.edu.au.

[1] We will use *decision* as the general term to encompass outputs from AI systems, such as categorisations, action selection, etc.

https://doi.org/10.1016/j.artint.2018.07.007
0004-3702/

Explanations are *Contextual*

- **Contrastive**: compared to counterfactual alternative
- **Selective**: focusing on relevant parts of full causal chain
- **Social**: transferring knowledge, assuming prior knowledge

(Tim Miller, Artificial Intelligence, 2018)

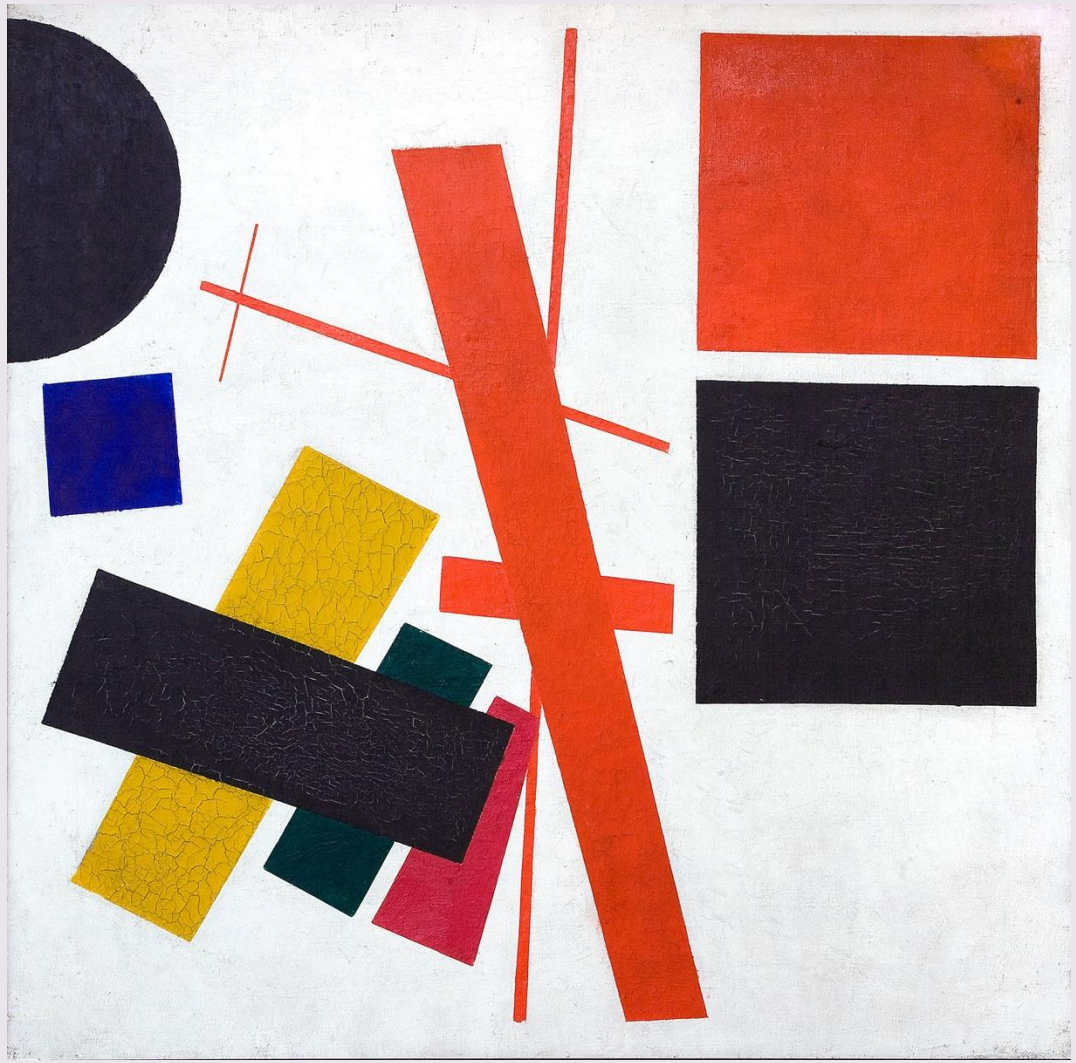

Counterfactual Reasoning

- *Factual:*
  Model denies loan
- *Counterfactual*: Alternative inputs that would accept loan
- *Algorithmic recourse:*
  Change of *behavior* to get desired outcome

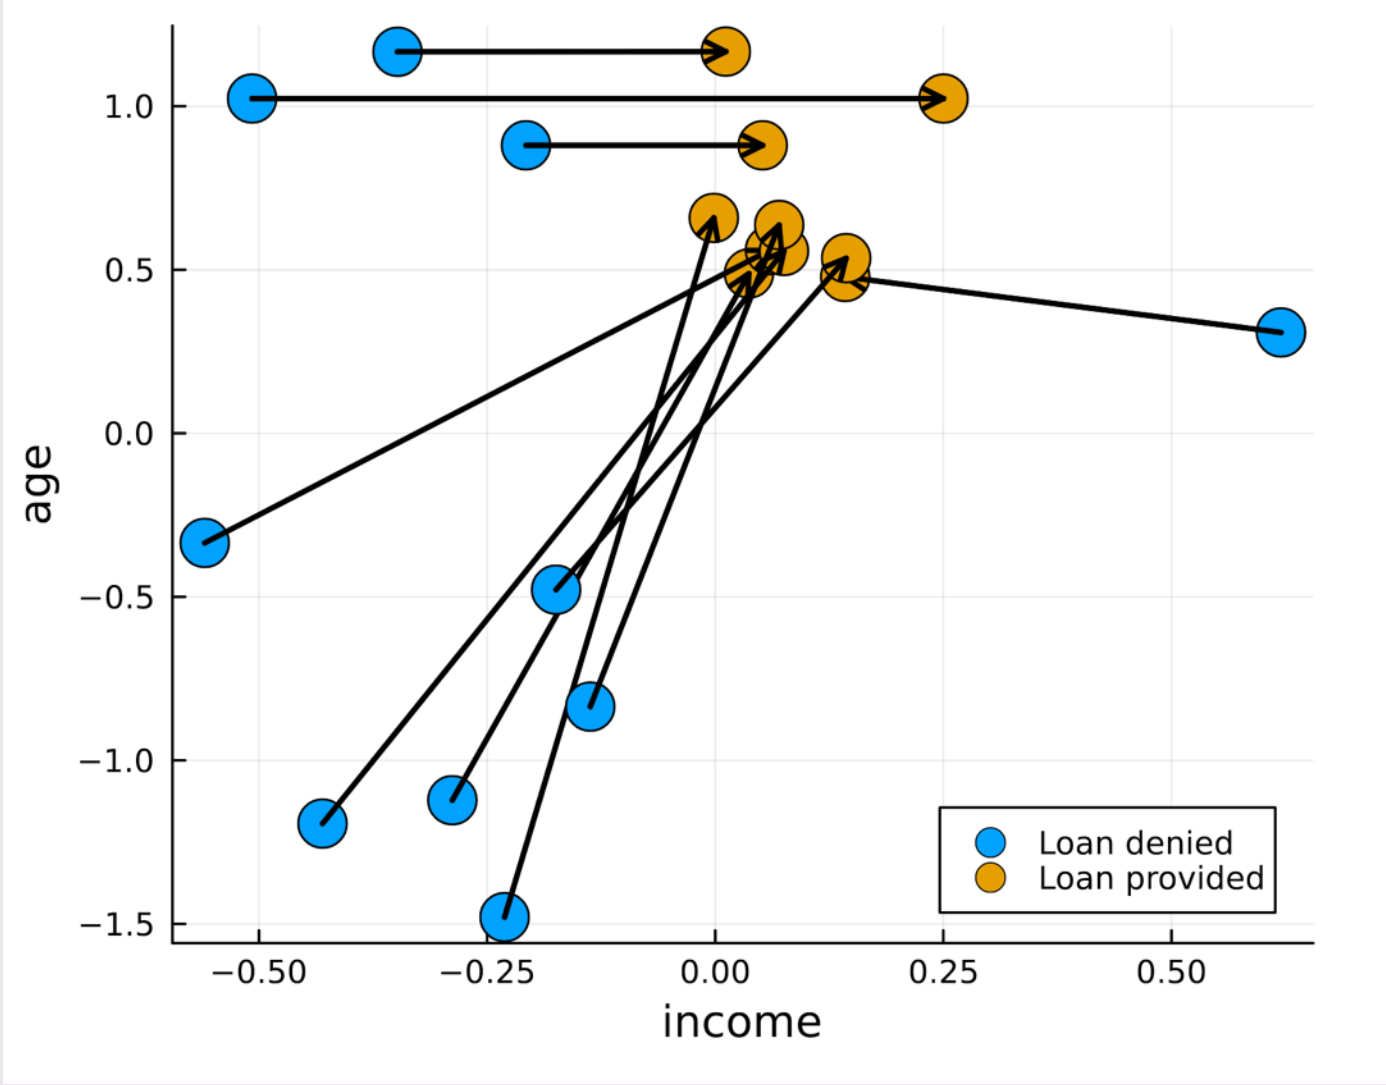

A Library for Generating Counterfactuals

Patrick Altmeyer
JuliaCon, 2022, 2023
IEEE SaTML, 2023
AAAI 2024

- Possible, faithful, plausible, “close” to the factual, …
- Gradient descent in feature space (with extra cost terms)
- Leverage ‘energy’ in input data seen during training
- Macro-effects after recourse adoption
- Rich library of Julia packages

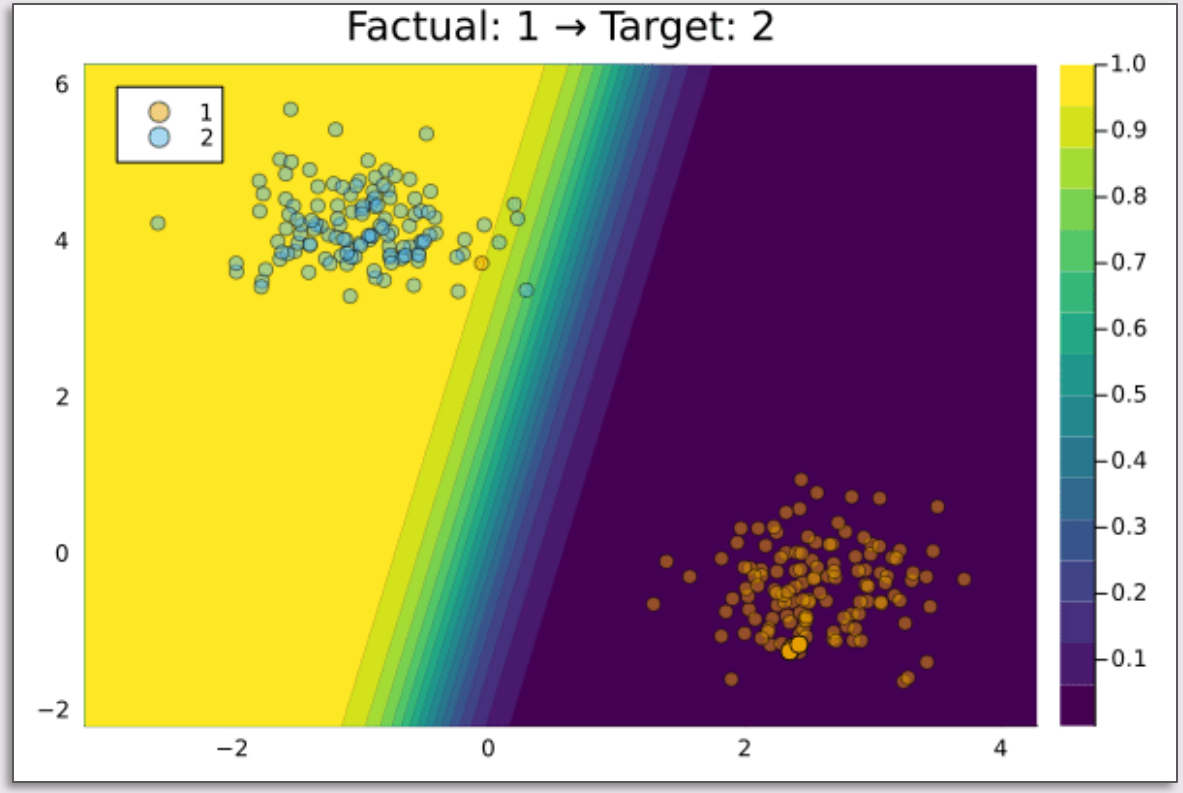

https://github.com/JuliaTrustworthyAI

ING Bank

Global bank based in The Netherlands

Five-year collaboration with TU Delft:

- Explainable AI
- Human-AI decision making
- Data integration
- Incident management and AIOps
- Release planning
- Search-based testing and repair

Intermezzo

www.theguardian.com/world/2025/feb/27/north...

North Korea

North Korea behind $1.5bn hack of crypto exchange ByBit, says FBI

The US agency said it refers to this specific North Korean malicious cyber activity as 'TraderTraitor'

Justin McCurry in

Thu 27 Feb 2025 04.35 CE

Share

North Korea was
from a cryptocu
described as th

The haul, whic
previous recor
Iraq's central
growing expe

Describing this particular form
the FBI on Wednesday warned that
trading platfo

13:10 Thu 6 Mar 85%

A $1.5 Billion Hack: How the Biggest Crypto Heist in History Went Down

The cryptocurrency exchange Bybit lost $1.5 billion to North Korean hackers last month — and it all traced back to an account on a free digital storage service.

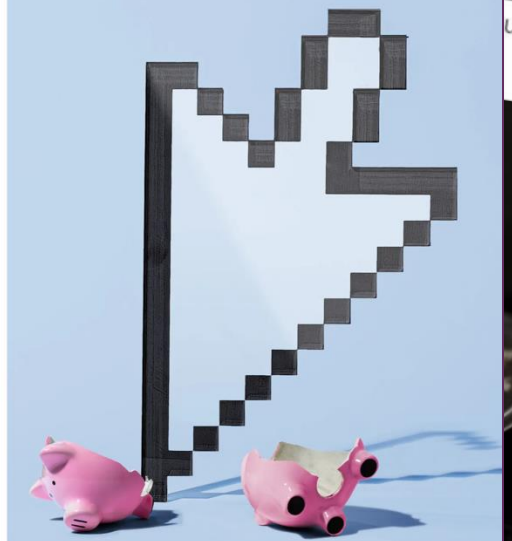

edition.cnn.com/2025/02/24/politics/north-kore...

CNN Politics

Sign in

North Korean hackers steal record $1.5 billion in single crypto hack, security firm says

Lyngaas, CNN

ute read · Published 4:54 PM EST, Mon February 24, 2025

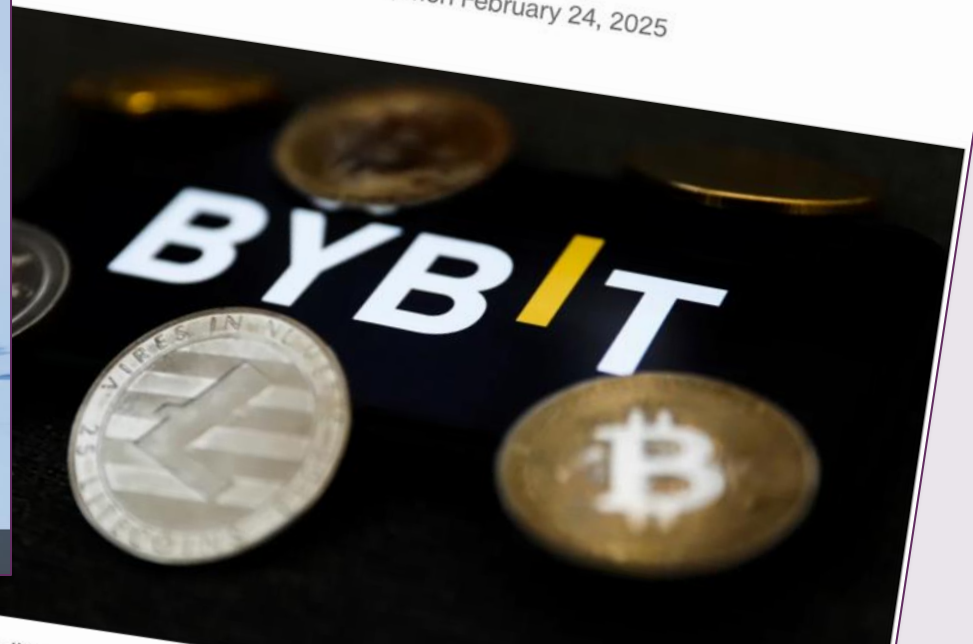

The Bybit logo in

Step 1: The Compromise

- ByBit crypto exchange uses 3rd party “Safe Wallet”
  - This is the weak link that can be exploited.
- Hackers obtain credentials for Safe{Wallet} developer machine
  - API keys of safe.global leaked or compromised
- Upload malicious Javascript code for “Safe UI”
  - Targeting Ethereum multisig cold wallet of Bybit
  - Makes it appear that Bybit is signing a legitimate transaction, when in fact it is a malicious one.

https://www.chainalysis.com/blog/bybit-exchange-hack-february-2025-crypto-security-dprk/

Step 2: The Theft

- Two weeks later:
  - Routine transfer from Bybit's Ethereum cold wallet to hot wallet triggers the malicious code
- Bybit CEO unknowingly signs the malicious transaction
- Hackers able to move ~401,000 ETH to addresses under their control

Step 3: The Laundering

- Move stolen assets through a complex web of intermediary addresses.
- Swap stolen ETH for tokens including Bitcoin and DAI.
- Move assets across networks using:
  - Decentralized exchanges and cross-chain bridges
  - Instant swap service without “Know Your Customer” reqs
- Keep portion of stolen funds idle across various addresses
  - Delay laundering to outlast the heightened scrutiny

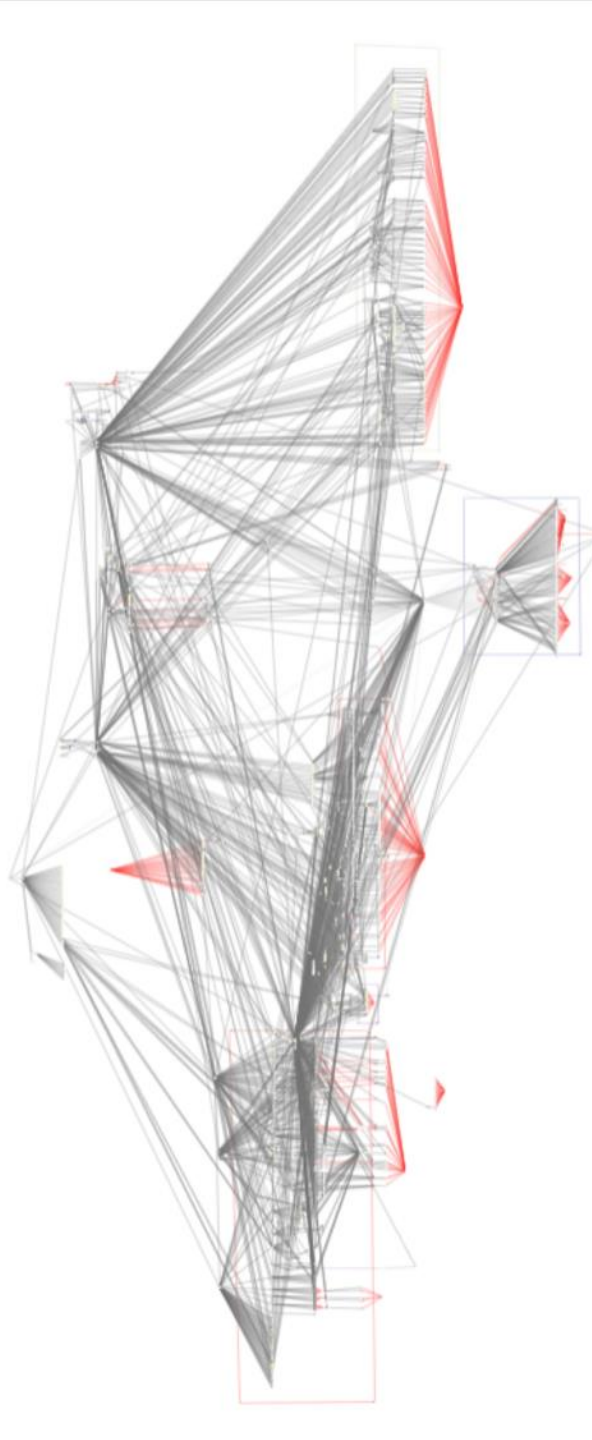

Bybit Implications

- Affects world peace: 1.5B for North-Korea
- Affects price of crypto-currencies
- Undermines societal trust in fintech
- Requires mix of prevention / remediation measures:
  - Strong cybersecurity, (incl. phishing)
  - Regulation (KYC, money laundering)
  - Forensics and traceability
  - While preserving privacy

Venkatesh Chandrasekar

The Delft Fintech Lab

- Bring together all
  TU Delft Fintech activities
- Research, education, innovation
- Launched May 2023
- 50 researchers in four faculties
- 25 commercial / societal partners

https://www.tudelft.nl/fintech/

Delft Fintech Lab: Research Pillars

- Fraud detection, privacy preservation
- Algorithmic trading
- Risk management
- Engineering financial systems
- Decentralized finance

Selected Blockchain/Security Research

- **Testing the protocol** (Lead: Burcu Kulahcioglu Ozkan)
  - OOPSLA 2023: Randomized testing for Byzantine Fault Tolerance
  - ICSE SEIP 2023: Evolutionary testing of Ripple's consensus algorithm
  - Issues detected and reported; *bounty received*
- **Testing smart contracts** (Lead: Mitchell Olsthoorn)
  - ICSE Tool Demo 2022: Syntest-Solidity (https://www.syntest.org/)
  - ICSME 2022: Guiding tests through transaction-reverting statements

Anti-Money Laundering

- Pattern-based synthetic data set
  - > 100 million transactions
  - Varying levels of illicitness (1:1750)
- ML-based detection approach
  - Trained on synthetic data
  - Validated on Ethereum data

Kubilay Atasu, NeuRIPS'23, AAAI'24, ICAIF'24

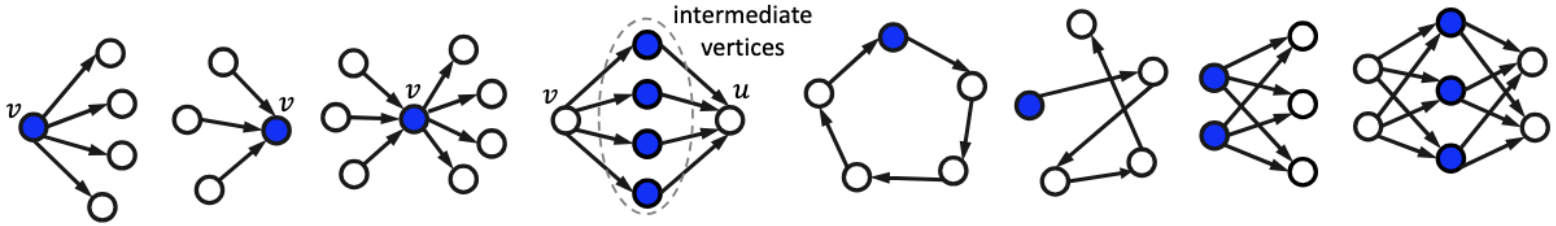

ML in Trading @ Delft

Antonis Papapantoleon

- Efficient and accurate algorithms to address trading in financial markets
- ML / math for valuation of financial derivatives
  - use of energy cost functions
  - domain knowledge (asymptotic option prices).
- Reinforcement learning for algorithmic trading and portfolio optimization
- Synthesis of implied volatility surfaces using diffusion models and auto-encoders.

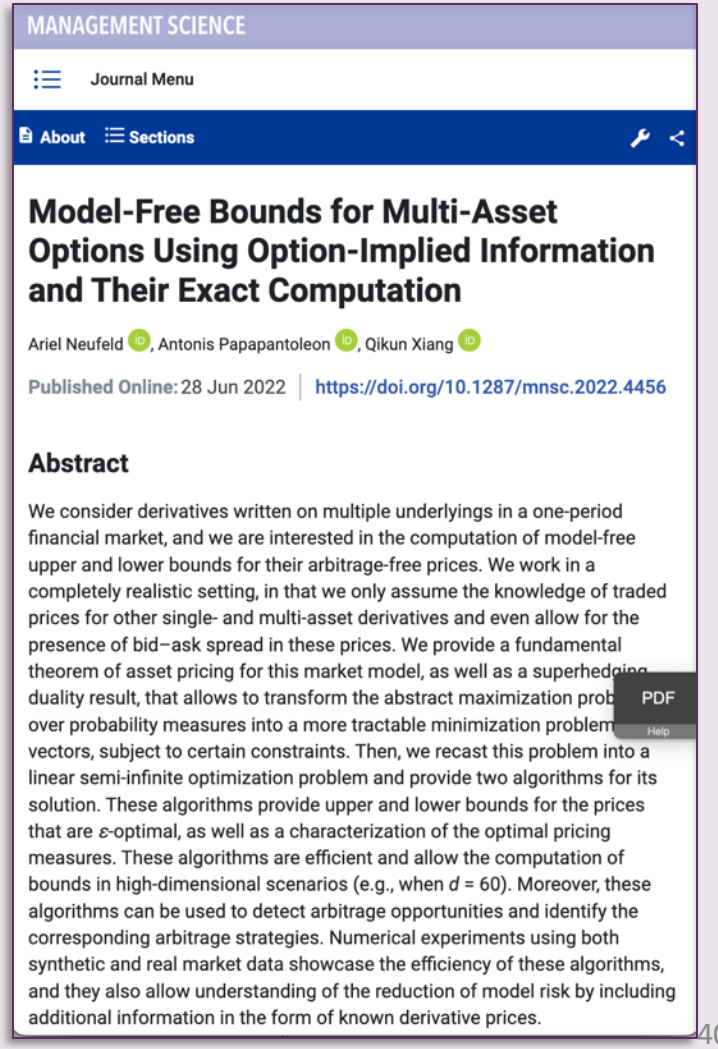

MANAGEMENT SCIENCE

Journal Menu

About Sections

Model-Free Bounds for Multi-Asset Options Using Option-Implied Information and Their Exact Computation

Ariel Neufeld, Antonis Papapantoleon, Qikun Xiang

Published Online: 28 Jun 2022 | https://doi.org/10.1287/mnsc.2022.4456

Abstract

We consider derivatives written on multiple underlyings in a one-period financial market, and we are interested in the computation of model-free upper and lower bounds for their arbitrage-free prices. We work in a completely realistic setting, in that we only assume the knowledge of traded prices for other single- and multi-asset derivatives and even allow for the presence of bid–ask spread in these prices. We provide a fundamental theorem of asset pricing for this market model, as well as a superhedging duality result, that allows to transform the abstract maximization prob over probability measures into a more tractable minimization problem vectors, subject to certain constraints. Then, we recast this problem into a linear semi-infinite optimization problem and provide two algorithms for its solution. These algorithms provide upper and lower bounds for the prices that are ε-optimal, as well as a characterization of the optimal pricing measures. These algorithms are efficient and allow the computation of bounds in high-dimensional scenarios (e.g., when $d = 60$). Moreover, these algorithms can be used to detect arbitrage opportunities and identify the corresponding arbitrage strategies. Numerical experiments using both synthetic and real market data showcase the efficiency of these algorithms, and they also allow understanding of the reduction of model risk by including additional information in the form of known derivative prices.

PDF

Help

FinTech as “Convergence”

- Grand societal challenges demand blended expertise of technical and socio-economic sciences
- TU Delft is strengthening its partnership with Erasmus University Rotterdam.
- Pilot FinTech projects:
  - Synthetic data generation
  - Default prediction
  - Household financial distress

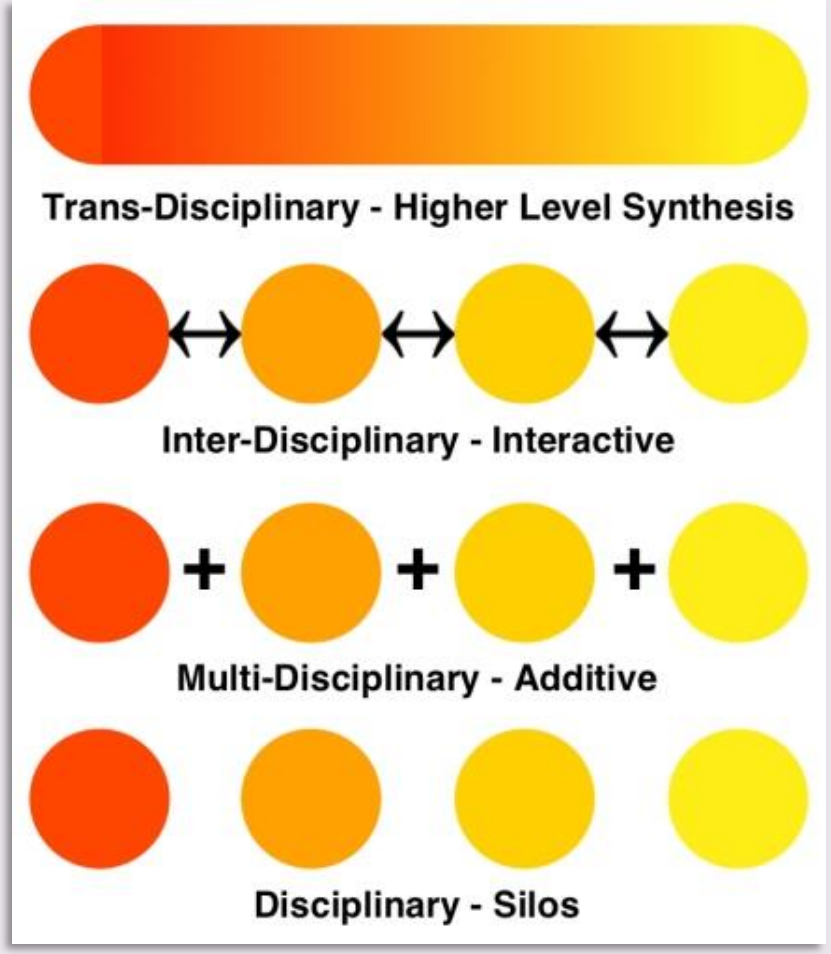

FinTech Education?

- Train future engineers who can
  - Design, build, evolve, and operate current and future financial systems
  - Assess and influence societal implications of new technological developments in the financial sector

- **Exploring** transdisciplinary master with intake from various disciplinary bachelor programs

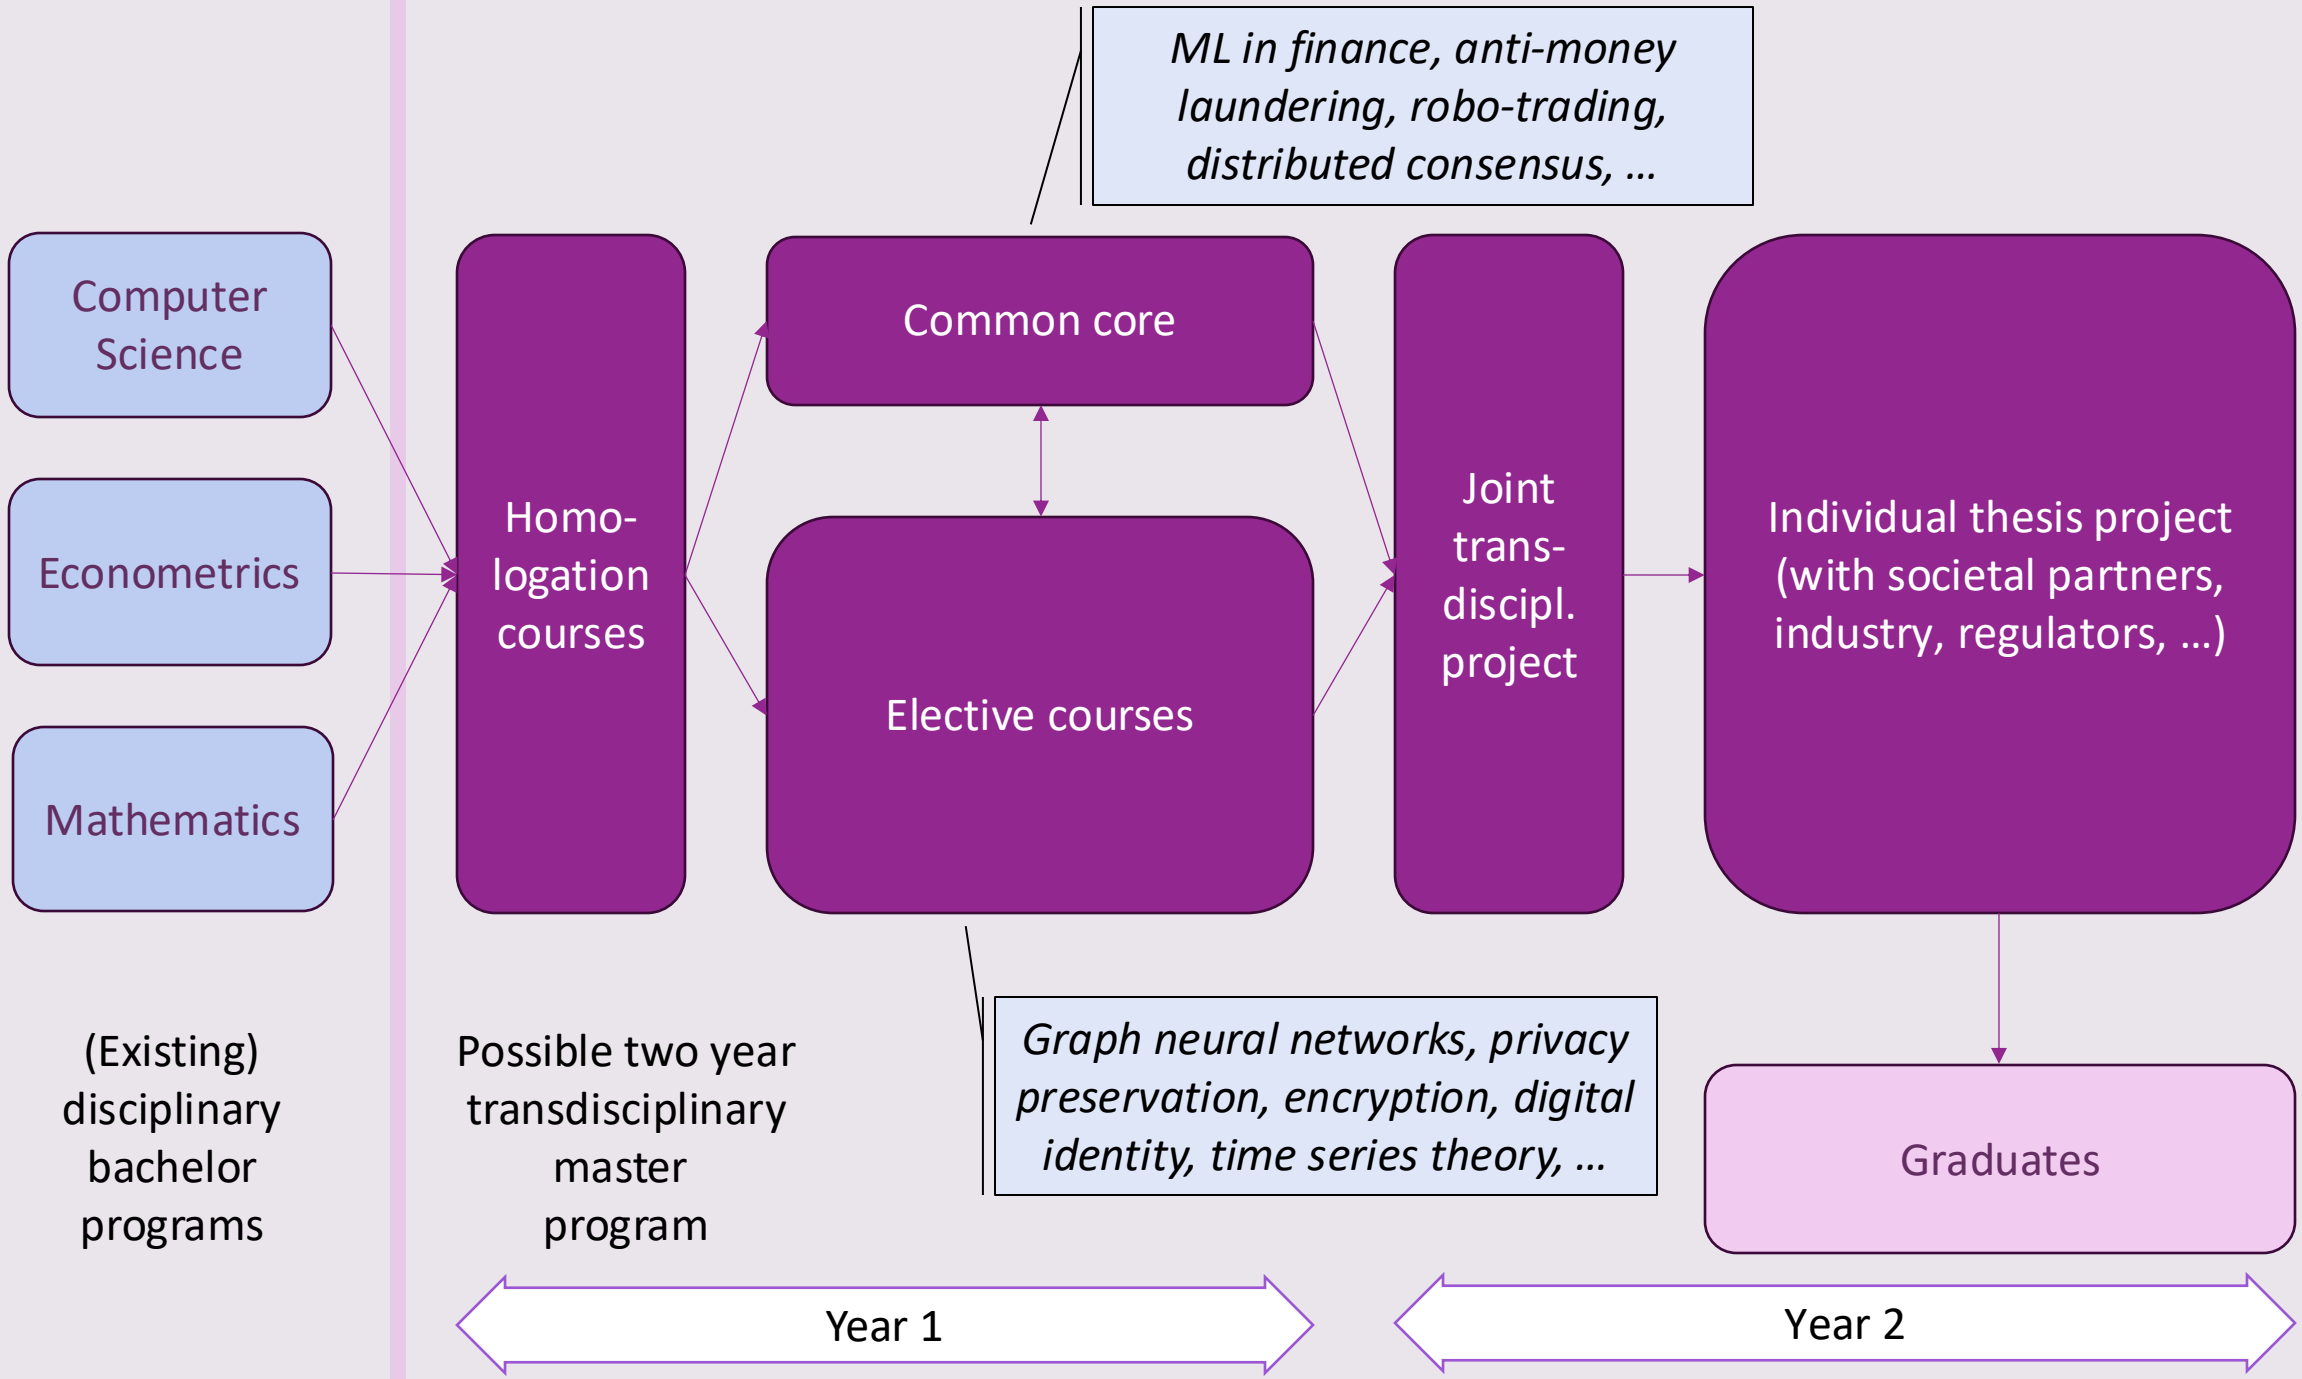
ML in finance, anti-money laundering, robo-trading, distributed consensus, …
Computer Science
Econometrics
Mathematics
Homo-logation courses
Common core
Elective courses
Joint trans-discipl. project
Individual thesis project (with societal partners, industry, regulators, …)
Graduates
Graph neural networks, privacy preservation, encryption, digital identity, time series theory, …
(Existing) disciplinary bachelor programs
Possible two year transdisciplinary master program
Year 1
Year 2

Alhambra in winter

Summary So Far

- The financial sector is a software factory
- Financial services need to be reliable, explainable, and secure
- Regulations help (DORA, KYC, AML)
- To move FinTech forward, we need to
  - look beyond our own disciplines …
  - … and train the next generation to do so

The Role of AI?

- Capabilities of **foundation models** are mind blowing
- This will affect many aspects of society, including finance and FinTech
- The ambitions of **Artificial General Intelligence** reach even higher
- Will (generative) AI solve our problems?


Sam Altman (img src Wikipedia)

I ❤ Large Language Models for Code

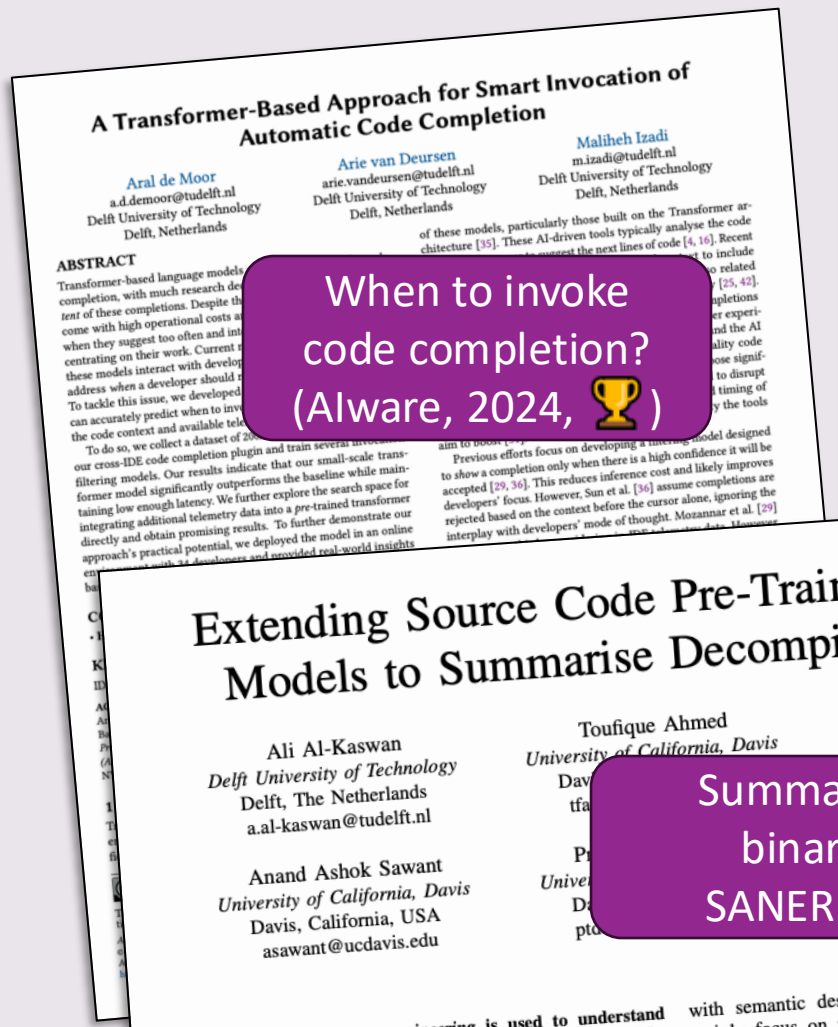

When to invoke code completion? (AIware, 2024, 🏆)

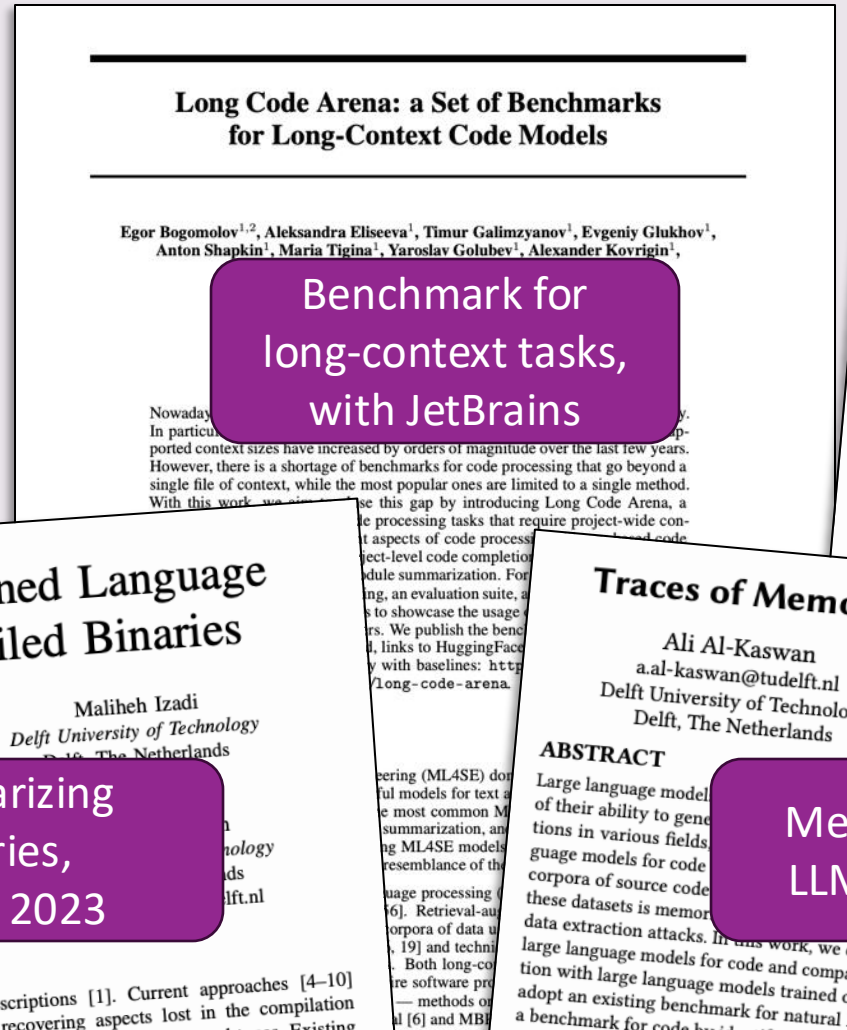

Benchmark for long-context tasks, with JetBrains

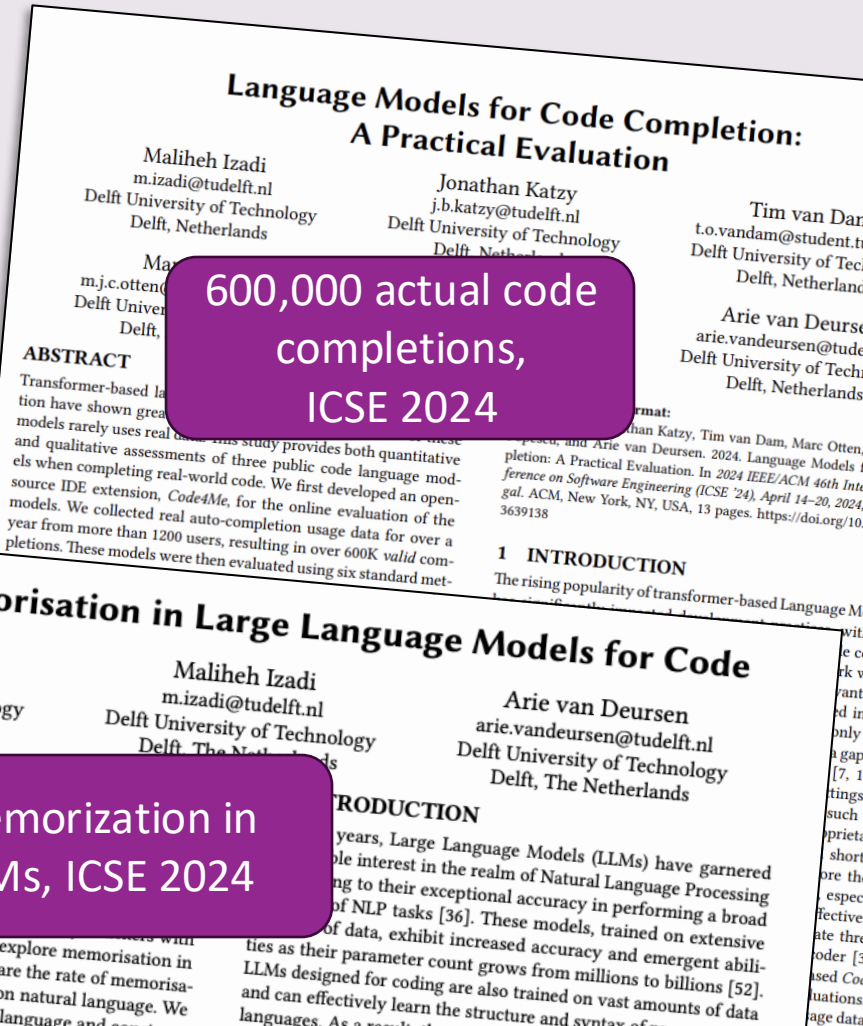

600,000 actual code completions, ICSE 2024

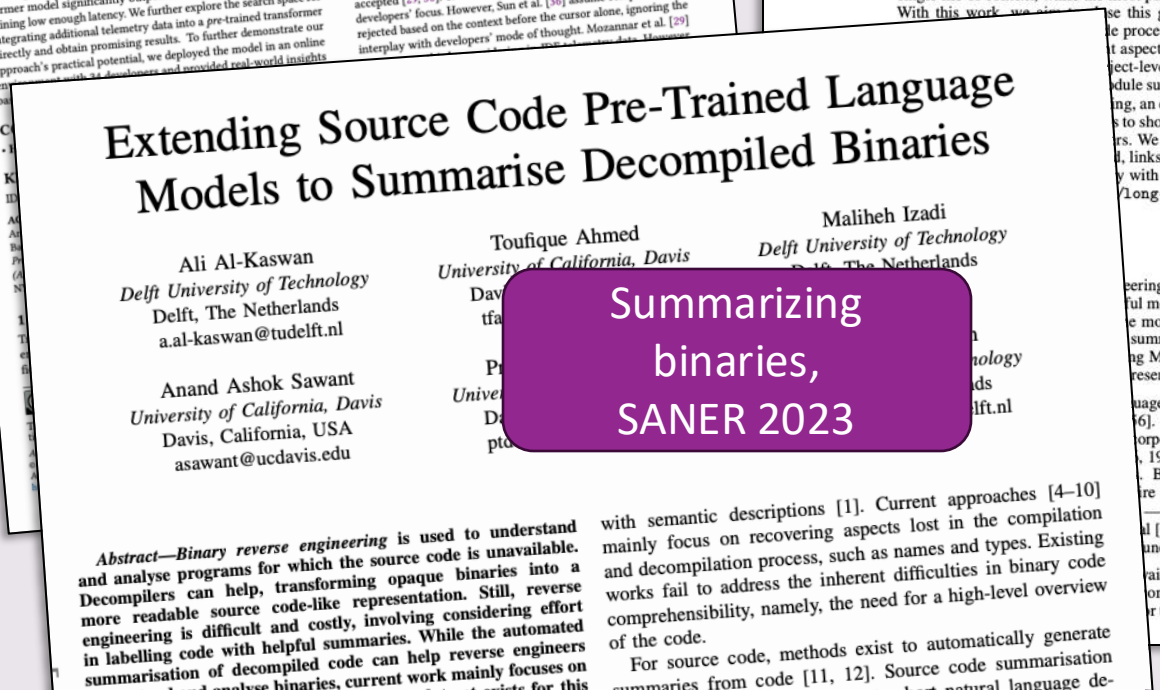

Summarizing binaries, SANER 2023

Memorization in LLMs, ICSE 2024

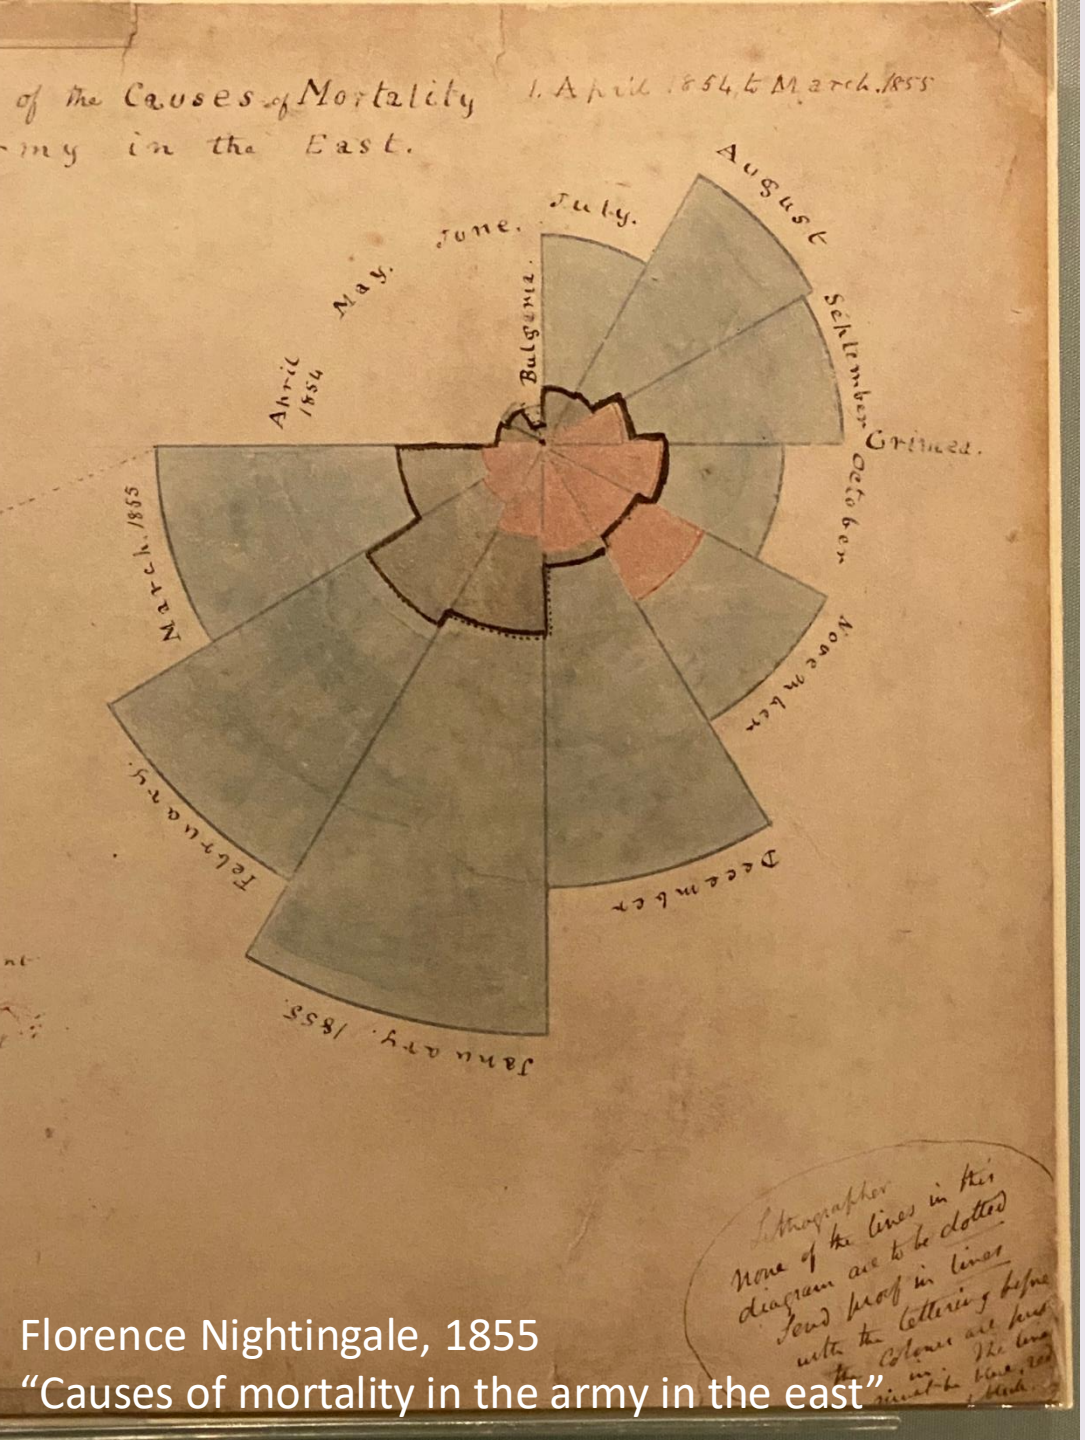

Florence Nightingale, 1855
"Causes of mortality in the army in the east"

Nothing Beats Good Data

- Finance has rich history of thorough data-driven research
  - Diffusion and auto-encoders add superpowers to your data
  - Use obtained understanding for synthetic data generation
- As Fintech 'organization', cherish, curate, your own data
- As research community, share data to drive research

Foundation Models for FinTech

- Endless, fascinating possibilities
- Your organization won't fit in a prompt
  - Explore agents, retrieval-augmentation, …
  - Consider finetuning/training with own data
- A model will inherit the good and bad from the training data
  - "Alignment" is tricky and volatile
  - Secret training data impedes progress

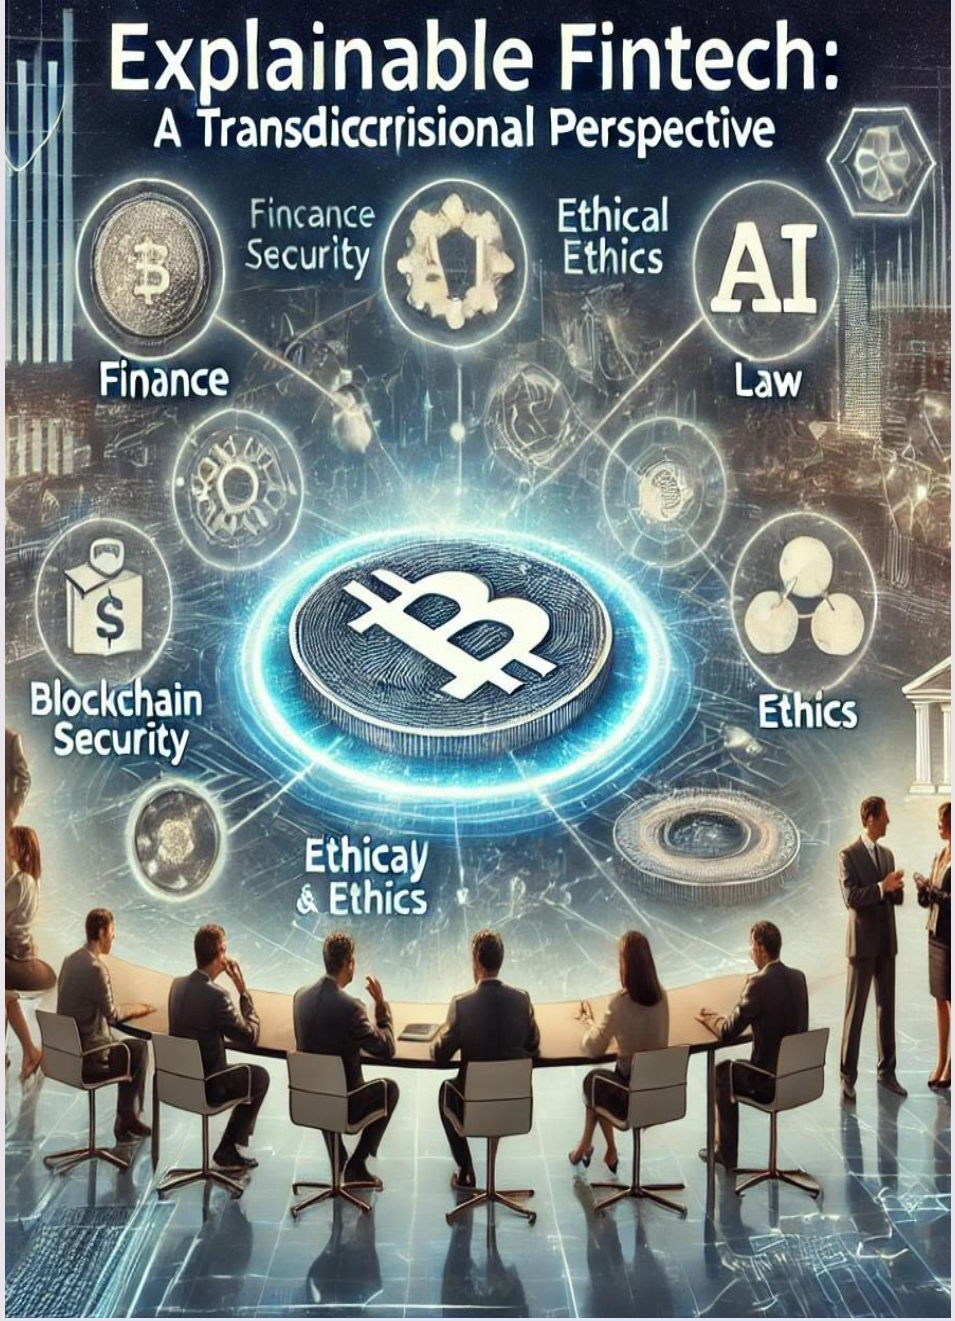

Prompt: a picture for "Explainable Fintech: A Transdisciplinary perspective"

Timnit Gebru, img Wikipedia

AI should be *Focused*

Timnit Gebru (SaTML 2023):

We should build smaller-scale systems (that are well-scoped and well-defined) for which we can provide specifications for expected behavior, tolerance and safety protocols.

https://x.com/NicolasPapernot/status/1623885641380425728

Earning and Keeping Societal Trust

- Financial systems must be **dependable** and **trustworthy**
- High complexity, volatility of crypto currencies, unreasonable profits, and excessive carbon emissions all **undermine society's support for fintech.**
- FinTech should be a 'fair game'

Eric Beinhocker
Oxford

Fair Social Contracts for Large-Scale Collaboration

Moral preferences	Dimension	Description
Relational fairness	Agency	*I can choose to play the game and have choices within the game.*
	Inclusion	*I have an opportunity to play the game. I am not excluded.*
	Dignity	*If I play by the rules and contribute to the best of my abilities, I will be valued, respected, and have status.*
Procedural fairness	Rule-based	*I know the rules of the game and they are applied equally to everyone.*
	Meritocratic	*I, and everyone else, will receive rewards and punishments in the game based on merit.*
	Security	*If I play by the rules and contribute to the game, but suffer misfortune through no fault of my own, I will be protected.*
Distributional fairness	Capabilities	*I have the capabilities to play the game or the opportunity to acquire them.*
	Reciprocity	*If I play by the rules and contribute, others will reciprocate, and I will share in the game's rewards.*
	Progress	*If I play by the rules and contribute to the best of my abilities, my life and the lives of those I care about will improve.*

https://theregenerators.org/future-council/

Green FinTech

- Fintech should *benefit* the climate
- AI-based Fintech needs Green AI
- Proof-of-work / Bitcoin mining is a terrible idea for the planet
- Lenders need collateral valuation over time *as climate changes*
- Finance is a domain that can *steer* society

Key Take Aways

1. The software- and data-intensive nature of fintech makes it an exciting domain for AI and software engineering research
2. Key challenges in FinTech include dependability, security, societal support, and sustainability
3. Transdisciplinary research and education is needed to address such challenges in FinTech

Explainable FinTech

A Transdisciplinary Perspective

Arie van Deursen
Delft University of Technology
@avandeursen@mastodon.acm.org
www.tudelft.nl/fintech/

Image: Wikipedia, Pont Adolphe