BANQUE CENTRALE DU LUXEMBOURG

EUROSYSTEM

cssf

Commission de Surveillance
du Secteur Financier

Thematic review on the use of Artificial Intelligence in the Luxembourg financial sector

May 2025

Thematic review on the use of Artificial Intelligence in the Luxembourg financial sector

May 2025

TABLE OF CONTENTS

1. Executive summary

In June 2024, the Banque centrale du Luxembourg (BCL) and the Commission de Surveillance du Secteur Financier (CSSF) launched a joint survey to assess the use of Artificial Intelligence (AI) technology by entities in the Luxembourg financial sector.

The primary objective of the survey, which follows a similar study conducted in 2023[1], is to understand the evolution of AI adoption within the sector, particularly noting the growing relevance and potential risks associated with generative AI (GenAI).

The survey was addressed to investment firms (IF), authorised investment fund managers (IFM/AIFM)[2], credit institutions (B), e-money institutions (EMI), and payment institutions (PI). This scope more than triples the size of the panel compared to the previous survey, which covered only B, EMI, PI.

The following paragraphs present a summary of the main findings identified by the survey.

Note: Where possible, the results of this survey were compared to the previous survey conducted in 2023, i.e. by aligning the type of respondents[3]. For clarity, the comparisons with the previous survey are highlighted with a coloured background in the text.

In total, the survey was answered by **461 financial institutions**, representing an **86% participation rate**.

In 2024, investments in innovative technologies (AI and DLT) were mainly made at group level (46%), while a much smaller portion of the respondents indicated having investments in AI/DLT performed both at local and group level (9%) or at local level only (4%). **Another significant portion (36%) of respondents indicated that they did not make any investments in innovative technologies (AI or DLT) in 2024**[4]. The proportion of respondents indicating "no investments" is higher for IF and IFM/AIFM and lower for B and PI.

For the 2025-2026 period[5], investments in AI are expected to increase more than those in DLT. Specifically, **AI investments are expected to increase more at local/Luxembourg level (with the highest increase for GenAI investments) compared to group level**. Indeed, at group level, AI investments are expected to remain relatively stable. This trend may be due to AI investments already being made at group level and now being implemented at local level to leverage group experience.

[1] See https://www.cssf.lu/en/2023/05/thematic-review-on-the-use-of-artificial-intelligence-in-the-luxembourg-financial-sector/.

[2] More specifically, the following types of fund managers were in scope of the survey: management companies subject to Chapter 15 (CH15 ManCo) of the Law of 17 December 2010 relating to undertakings for collective investment (2010 Law); authorised alternative investment fund managers (AIFM) subject to the Law of 12 July 2013 on alternative investment fund managers (2013 Law).

[3] Comparisons with the previous survey are carried out by focusing only on answers provided by B, PI, EMI, i.e. a scope of entities similar to the previous survey.

[4] The remaining 5% of respondents did not provide any information.

[5] While the survey only requested to provide estimations, a significant portion of respondents did not provide information about 2025-26 investments, likely because the budgeting process for 2025-26 had not been completed at the time of the survey.

Additionally, **AI is generally perceived to offer greater cost savings, efficiency gains, and revenue increases compared to DLT**.

Regarding **public GenAI tools** (such as ChatGPT, Gemini, Claude, etc.), **64% of respondents allow access for their employees (58% allow unrestricted access and 6% allow access only for a restricted number of employees), while 36% state that access is denied**. The level of access varies based on the size of the entities (the bigger the entity the more restricted the access). We also note that **credit institutions are more restrictive,** with 54% of respondents blocking the access of their employees and 37% providing free access. **Among the entities that provide access to public GenAI tools to some or all of their employees, only 40% have either implemented a specific GenAI policy or modified their existing Internet policy to explicitly address the use of GenAI tools**. This leaves 60% of these entities without a dedicated policy on the subject. When a policy is present, it focuses on confidentiality, compliance and data protection.

In relation to **AI adoption**, **28% of all respondents use AI technology in production or in development**, while **22% are experimenting or planning to experiment with AI technology in the next 12 months.** EMI and PI seem more mature in terms of AI adoption, with 63% of them indicating to have concrete use cases in production or in development, followed by credit institutions with 38%. **Compared to the previous survey (focusing on answers from B, PI, EMI), the results show an increase in the adoption of AI technologies, with 43% of these entities using AI in production/development (vs 30% in the previous survey).**

The results presented in the remainder of this executive summary focus on responses from entities that use AI technology in production/development or that are experimenting or planning to experiment with AI technology in the next 12 months.

The main AI benefits indicated by respondents are related to internal efficiency, with the top three being "Improve internal processes", "Optimise operations/reducing costs" and "Analyse vast amounts of data". **The main AI challenges are related to data**, with "Data quality" being the top challenge, followed by "Data protection" and "Data governance". These results are overall consistent with those identified in the previous survey.

The vast majority (84%) of respondents have already implemented or plan to implement a range of AI training programmes for their employees, spanning from basic awareness to advanced AI trainings. Additionally, **43% of respondents reported having a formally approved AI policy, and more than half (54%) indicated having implemented security measures in relation to specific AI vulnerabilities, marking a significant increase compared to the previous survey[6].** Overall, **these findings suggest an improvement in AI maturity** among institutions compared to the previous survey.

The majority (63%) **of entities using AI have a dedicated data science team**. These teams are **primarily situated at group level** (55%), with a much smaller portion operating at both local and group levels (5%), or solely at local level (3%). Since data science teams at group level tend to be larger, these figures confirm the **trend of leveraging group expertise for AI-related**

[6] The percentage of respondents indicating to have taken security measures specific for AI vulnerabilities increases to 66% when focusing only on entities of type B, PI, EMI (i.e. the same scope of the previous survey), while it was close to 50% in the previous survey.

development activities. Conversely, **the portion of respondents with no data science team has increased compared to the previous survey, reflecting the growing availability of "ready to use" solutions such as GenAI tools** that do not require advanced AI technical skills for their implementation.

Regarding the technical infrastructure supporting the AI processes, **respondents are primarily using commercial cloud solutions (45%),** showing an **increase compared to the previous survey**. The increase in the use of cloud solutions is linked in the majority of cases with the use of GenAI solutions. A smaller portion of respondents (22%) indicated using private/dedicated infrastructures, while 24% employ hybrid (cloud and local) environments.

Of the 461 survey respondents, **36% reported at least one AI use case**. A total of **402 AI use cases were reported, with 54% of these already in production**.

The vast majority (92%) of the reported use cases are only for internal use (i.e. not client facing).

61% of all use cases leverage GenAI technology, **followed by Natural Language Processing (NLP)** (30%), **and machine learning (ML)** (28%). However, it was observed that most NLP use cases also involve GenAI, suggesting difficulties in distinguishing between these two technologies. When comparing the portion of entities using GenAI versus ML, we observe that **28% of all survey respondents have reported at least one use case involving GenAI**[7], and **12% have reported at least one use case involving ML**[7]. These figures indicate a **much wider adoption of GenAI compared to ML technology**. However, approximately half of the use cases involving GenAI are still at an experimental/proof-of-concept stage or under development, suggesting that **GenAI is at an earlier stage of adoption compared to ML**. **ML technology, on the other hand, appears to be more mature, with a higher portion of use cases already in production**.

Across the different types of entities, GenAI adoption (in terms of percentage of entities reporting at least one use case involving GenAI) **is higher for PI** (50%), **followed by B** (32%) **and IFM/AIFM** (29%). **Regarding ML adoption** (in terms of percentage of entities reporting at least one use case involving ML), **EMI are leading** with 50%, **followed by PI** with 44%, **and then B** with 24%.

Nearly all (94%) GenAI use cases rely on Large Language Models (LLM). Additionally, **75% of the GenAI use cases employ commercial models,** 11% are using open-source models and 11% are using both. On the other hand, **38% of respondents using ML indicated using third-party vendor solutions for ML development**, including data preparation.

The top five use case categories are "Search/summarise information" (43%), "Process automation" (30%), "Chatbot and virtual assistant" (27%), "Text context generation" (27%) and "Translation" (19%), which are categories mainly involving GenAI. **Compared to the previous survey, and for credit institutions in particular, the "AML/Fraud detection" now sits in fifth position while it was the first use case reported in the previous survey.**

[7] In production, development or experimental stage.

This is largely explainable by the GenAI swarm that appeared in late 2023. **For EMI and PI, the AML/Fraud detection category remains however the top use case**.

Regarding the **AI Act classification** (the AI Act entered into force only after the launch of the survey[8]), we note that only 5% of use cases were rated as "High Risk" and refer mainly to use cases such as credit scoring, Internal Ratings Based (IRB) credit risk model and AML/Fraud detection, whilst the last two are actually excluded from the list of high-risk systems as defined in the Annex III of the AI Act[9]. Indeed, **the classification in the survey seems to reflect the perception of the risk of the use case for the entity, rather than its actual classification according to the AI Act**.

With regard to **human oversight**, 90% of the reported use cases are said to have a **human in the loop**, which represents an **increase** compared to the 77% reported in the previous survey.

In relation to **bias treatment**, for 45%[10] of the use cases respondents confirmed having implemented bias prevention and/or detection measures. Compared to the previous survey, we observe an **increase in the adoption of these techniques**[11]**,** highlighting the increasing importance of bias prevention/detection measures. **However, for a significant portion of the use cases, respondents indicated that bias prevention/detection measures were not applicable. Notably, most of these use cases were GenAI use cases**. This can be partially attributed to the **expectation that bias treatment mechanisms are primarily the responsibility of the GenAI model provider**, particularly for Large Language Models (LLMs). **Nonetheless,** it should be noted that depending on the use case, **additional bias prevention/detection measures may still be necessary on the deployer's side**.

As concerns the **auditability of the AI models, only 56% of the use cases report good or very good auditability, representing a downgrade in the ratings compared to the previous survey**[12]. The reason for this downgrade cannot be easily explained but may be associated with the increasing level of complexity of the AI solutions used and the difficulty in auditing them, together with **more realistic scores** provided by respondents based on more experience (including regarding AI systems audits).

For **explainability,** there is **a very similar trend with 54% of the use cases reporting good or very good explainability, representing less explainable solutions compared to the previous survey**[13]. We note that the **levels of auditability and explainability are often correlated**, with similar ratings for both attributes for the same use case.

[8] While a stable version of the text of the AI Act was already available for a few months, the AI Act entered into force only on 1 August 2024, with the entry into application planned for 2 August 2026 except for specific provision.

[9] See recital 58 and Annex III, art.5(b) of AI Act.

[10] Percentage calculated excluding use cases for which respondents indicated that bias prevention/detection measures were not applicable.

[11] Considering only B, PI, EMI (and excluding "N/A" answers), 68% of use cases implement bias prevention/detection, compared to 59% of the previous survey.

[12] Considering only B, PI, EMI, 55% of use cases report good or very good auditability, while it was 81% in the previous survey.

[13] Considering only B, PI, EMI, 54% of use cases report good or very good explainability, while it was 70% in the previous survey.

Finally, with regard to the **model performance, this is actively monitored for the majority** (56%) **of use cases**, **with results being consistent with those from the previous survey**[14]. For the remaining use cases for which there is no active monitoring of model performance, the majority involves GenAI, suggesting that the complexity of such models presents challenges when it comes to performance monitoring.

[14] When focusing solely on ML use cases reported by B, PI and EMI and excluding "N/A" and "do not know" answers (in order to obtain data comparable with the previous survey), the proportion of AI solutions monitored over time increases to 88%. This latter figure is largely consistent with the results of the previous survey, where 90% of ML use cases had processes to monitor the algorithm performance over time.

2. Introduction and objectives

In June 2024, the Banque centrale du Luxembourg (BCL) and the Commission de Surveillance du Secteur Financier (CSSF) launched a joint survey aimed at assessing the use of Artificial Intelligence (AI) technology within the Luxembourg financial sector.

This joint initiative follows a similar study conducted in 2023[15], which focused on credit institutions, e-money institutions, and payment institutions. The objective of this survey is to understand the evolution of the usage of AI technology within the sector, including with regard to generative AI (GenAI), which was not analysed in the previous survey[16].

This report presents the results of the survey and associated findings.

3. Scope and methodology

The survey was carried out during the period June 2024 – August 2024, and was addressed to all Luxembourg credit institutions (B), authorised investment fund managers (IFM/AIFM)[17], investment firms (IF), e-money institutions (EMI), and payment institutions (PI) supervised by the CSSF as of 1 June 2024.

The survey consisted of an online questionnaire composed of four main sections:

- **General information** covering general information about the company (e.g. contact information, size of the company, size of the IT team, IT outsourcing). This section also covers information regarding the usage of GenAI tools freely available on the Internet and the existence of policies on GenAI, as well as the general status of adoption of AI.
- **Digital strategy** covering current and future investments (and related benefits in terms of increased revenues or decreased costs) in innovative technologies such as AI and Distributed Ledger Technologies (DLT) (including tokenisation and crypto assets).
- **AI questionnaire** covering various general aspects regarding the use of AI technologies, such as benefits and challenges, organisational aspects, data and governance, security and robustness, machine learning (ML) development lifecycle and technical infrastructure, GenAI specific usage methods, etc.
- **AI use cases** focusing on the practical use cases where AI technology is applied, covering general development aspects, trustworthiness, etc.

The responses from the survey questionnaires were aggregated, anonymised, and analysed to produce this thematic report. Some data cleansing was performed to ensure consistency and normalisation of data, with appropriate care not to fundamentally alter the answers received.

The report is organised as follows:

- Chapter 4 presents some general demographic information from the "General information" section of the survey.

[15] See "Thematic review on the use of Artificial Intelligence in the Luxembourg Financial sector, May 2023" (https://www.cssf.lu/en/Document/thematic-review-on-the-use-of-artificial-intelligence-in-the-luxembourg-financial-sector/). The report presents the results of the first survey, which run during the period October 2021 – January 2022.

[16] Commercially available GenAI solutions started appearing in November 2022, i.e. after the previous survey was launched.

[17] More in detail, the following types of fund managers were in scope of the survey: management companies subject to Chapter 15 (CH15 ManCo) of the Law of 17 December 2010 relating to undertakings for collective investment (2010 Law); authorised alternative investment fund managers (AIFMs) subject to the Law of 12 July 2013 on alternative investment fund managers (2013 Law).

- Chapter 5 focuses on the "Digital Strategy" section of the survey.
- Chapter 6 presents the findings from the "AI questionnaire" section of the survey.
- Chapter 7 presents general findings from the "AI use cases" section of the questionnaires.
- Chapter 8 and 9 focus on use cases using respectively GenAI and ML technologies.
- Chapter 10 focuses on the AI trustworthiness aspects of the use cases.

Whenever possible, the results of this survey were compared with those from the previous survey[18]: to do so, results were filtered selecting only the type of entities such as B, PI, EMI which were in scope of the previous survey, excluding those (such as IF and IFM/AIFM) which were not in scope of the previous survey. For clarity, sentences describing the comparison with the previous survey are highlighted with a coloured background in the text.

4. Survey demographics

In total, **537 institutions** were **targeted by the joint survey**.

The survey had a very good response rate, with a total of 461 respondents, representing a participation rate of 86%. It is worth noting that more than half of all respondents (57%) are IFM/AIFM, while PI and EMI combined represent only 5%. Overall, the distribution of respondents by type of entity (figure 1)[19] is very similar to the distribution of the targeted entities[20], indicating a balanced participation across all type of entities.

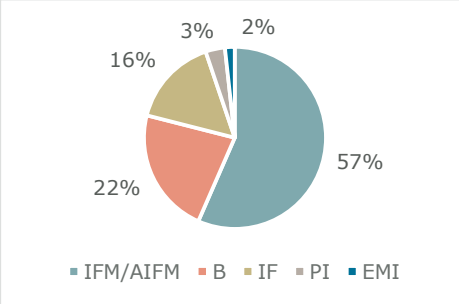

Figure 1: Survey respondents (by type of entity)

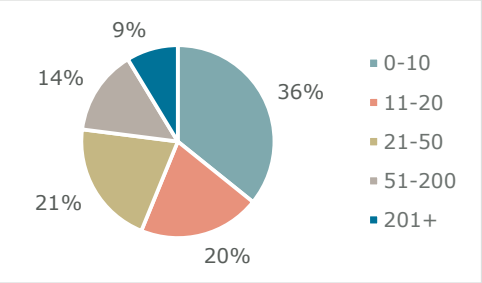

Figure 2: Size (number of employees) of survey respondents

More than half of the respondents (56%) are small in size (less than 20 employees), most of them being IFM/AIFM and IF. Credit institutions tend to be larger in size.

[18] See "Thematic review on the use of Artificial Intelligence in the Luxembourg Financial sector, May 2023" (https://www.cssf.lu/en/Document/thematic-review-on-the-use-of-artificial-intelligence-in-the-luxembourg-financial-sector/)

[19] The respondents were 261 IFM/AIFM representing 57% of the total number of survey participants, followed by 103 credit institutions (22%), 73 investment firms (16%), 16 payment institutions (3%) and 8 e-money institutions (2%).

[20] The targeted population consisted of 299 authorised investment fund managers (IFM/AIFM) (56%), 124 credit institutions (23%), 85 investment firms (16%), 17 payment institutions (3%) and 12 e-money institutions (2%).

Concerning the size of the IT teams, 39% of respondents reported having no IT staff[21], while 50% have less than 10 IT employees in Luxembourg (mainly small entities). The entities reporting IT teams with more than 50 employees are credit institutions and IFMs/AIFM[22]. Respondents with more than 200 IT employees are only credit institutions.

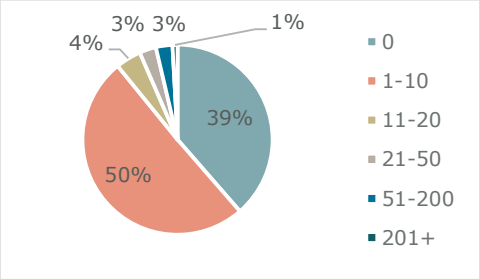

Figure 3: Size of IT teams (n. of employees)

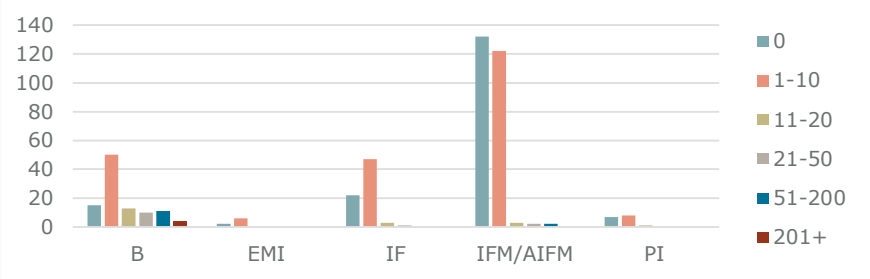

Figure 4: Size of IT teams (n. of employees), split by entity type

The IT function is fully outsourced to the group for 46% of respondents, as opposed to 32% of respondents not applying any IT outsourcing to the group. The remaining respondents reported having partial IT outsourcing.

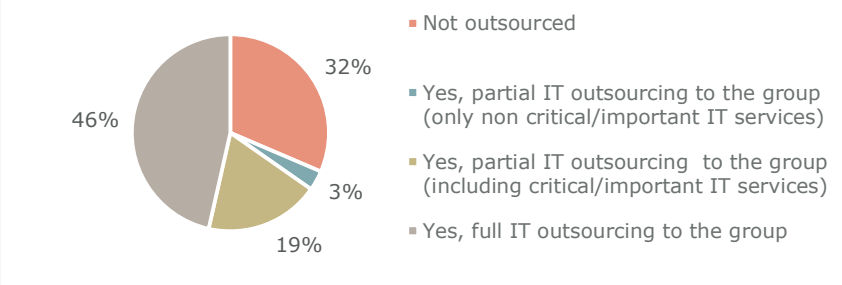

Figure 5: IT outsourcing to the group

[21] For respondents indicating no IT employees, we considered that this answer corresponds to situations where IT is fully outsourced, not having any "core" IT staff internally. This does not take into account the managing director in charge of IT, the IT outsourcing officer, the person in charge of information security and the one responsible for IT risks.

[22] 15 credit institutions and 4 IFM/AIFM.

5. Digital strategy

The objective of this part of the questionnaire was to understand whether entities had a digital strategy defined at local or group level, and to identify current and future investment trends (and related expected benefits in terms of reduced costs or increased revenues) in innovative technologies such as AI (including GenAI and ML) and DLT (including crypto assets and tokenisation).

With regard to the digital strategy**, 24% of respondents had a digital strategy (approved by the Board) defined at local (Luxembourg) level, and 52% of respondents had a digital strategy defined at group level**.

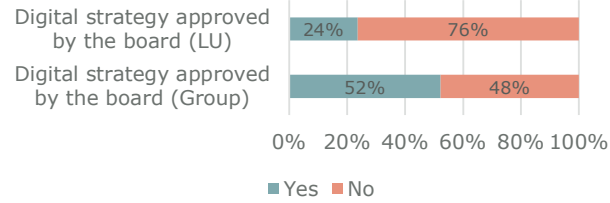

Figure 6: Digital strategy (approved by the Board) at Luxembourg level and Group level

5.1 2024 Investments

According to the survey, **investments in innovative technologies (AI and DLT) are mainly made at group level (46%)**, while a much smaller portion of the respondents indicated having investments in AI/DLT performed both at local and group level (9%) or at local level only (4%).

Another significant portion (36%) of respondents indicated that they did not make any investments in innovative technologies (AI or DLT) in 2024.

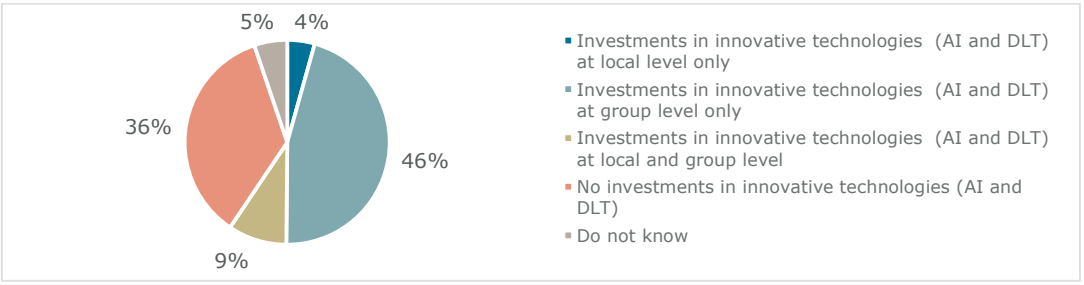

Figure 7: 2024 Investments in innovative technologies at group or Luxembourg level

When splitting the data by type of entity, we observe that the portion of respondents indicating "no investments" is higher for IF and IFM/AIFM. B and PI are instead those with a higher portion of entities performing investments.

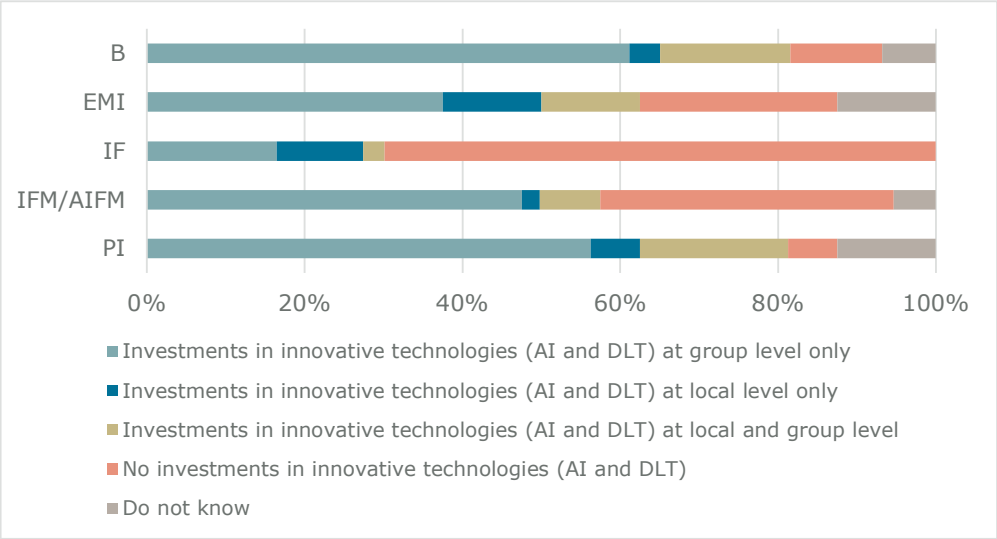

Figure 8: 2024 Investments in innovative technologies at group or Luxembourg level (by entity type)

Among all respondents, only a small portion provided quantitative information[23] about the amounts invested in innovative technologies in 2024[24]. Based on this information, the **amount invested in innovative technologies** (AI and/or DLT) **in 2024 represents on average**:

- for total AI investments: 6% of Luxembourg IT budget; 9% of group IT budget;
- for total DLT investments: 7% of Luxembourg IT budget; 6% of group IT budget;
- for total AI and DLT investments: 6% of Luxembourg IT budget; 11% of group IT budget.

More in detail, with regard to AI:

- at Luxembourg level, most of the above entities invested in ML and/or GenAI, although in terms of volume, investments are higher for ML;
- at group level, investments in GenAI are higher than those at local level, both in terms of number of entities and of volume (amount invested as percentage of IT budget).

More in detail, with regard to DLT:

- investments in DLT are predominantly at group level in terms of number of entities investing. At group level, a bigger portion of entities reported investments in DLT tokenisation projects, although investments in the "DLT-Other" category (non-crypto nor tokenisation) are higher in terms of volume.

[23] Only entities which provided information on IT budget and on the amount invested in innovative technologies were considered. To allow comparison, investments are calculated as percentages of the corresponding IT budget (amounts invested/IT budget).

[24] In particular, only 44 respondents (10% of all respondents) provided information regarding investments at local/Luxembourg level, while 153 respondents (33% of all respondents) provided information regarding investments at group level.

- Only few (six) entities reported investment at local level in DLT.

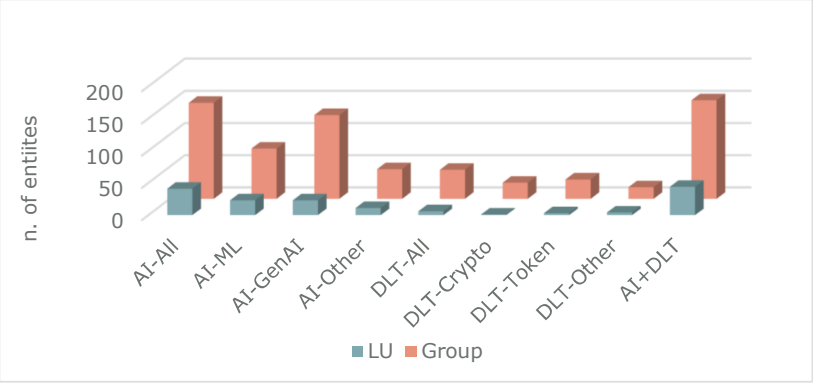

Figure 9: n. of entities that invested in innovative technologies in 2024 (Luxembourg vs group level)

5.2 2025-2026 Investments

The survey asked participants to indicate whether investments for 2025/2026 in innovative technologies (AI or DLT), both at local (Luxembourg) and group level, were estimated to increase, decrease or stay the same.

In both cases (local and group level), **the majority of respondents indicated that investments in innovative technologies (AI or DLT) were expected to remain the same**. Although the survey only requested to provide estimations, a significant portion of respondents answered "do not know", probably due to the fact that the budget process for 2025-26 had not yet been done at the time of the survey.

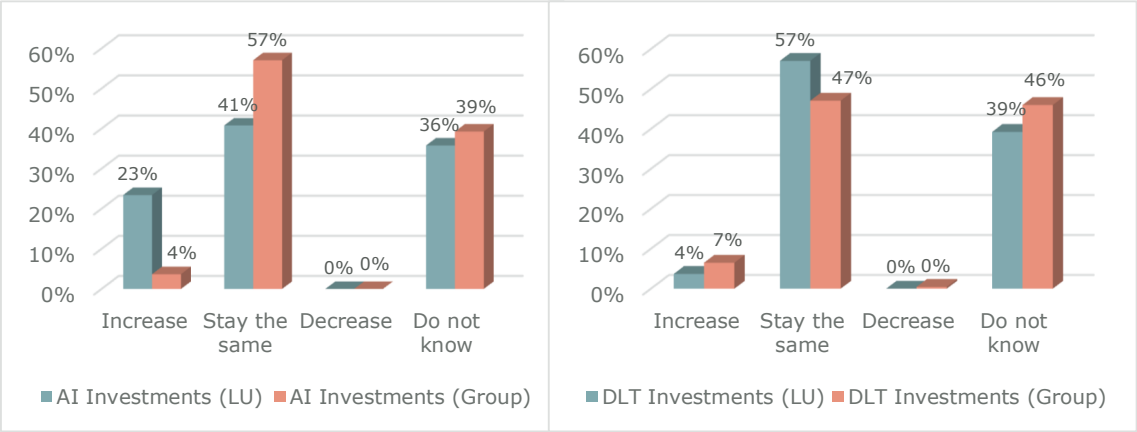

Figure 10: 2025-2026 Investments at Luxembourg/group level in AI (left) and DLT(right)

It is worth noting that the percentage of respondents expecting a decrease in investments in innovative technologies was close to 0% for both AI and DLT.

Regarding the investments expected to increase, we note that **the portion of respondents who replied that investments were expected to increase is higher for AI than DLT**[25].

For AI investments, when comparing data between local and group level, we observe that **the percentage of respondents indicating that investments were going to increase is higher**

[25] Considering investments at both local and group level combined.

at local level (with the highest increase for GenAI investments) compared to group level. At the same time, this trend seems to be counterbalanced at group level by a higher percentage of respondents indicating that investments in AI were "staying the same". A possible explanation behind this trend could be that investments in AI were already made at group level and that now they are being applied (in a second phase) at local level, thereby benefitting from the group experience.

For DLT investments, more respondents indicated they expected an increase of investments at group level compared to those at local level. If the above reasoning (according to which investments are done firstly at group level and in a second step at local level) was applied here, it could be interpreted as an indication of the **lower level of maturity (in terms of adoption) of DLT technology compared to AI technology**. Finally, we note that investments in DLT/tokenisation at group level are expected to increase more than those in crypto assets related activities.

5.3 Anticipated cost savings and efficiency gains

The survey asked to estimate, on a scale from 0 = none to 5 = very high, the cost savings or efficiency gains over the next 2-3 years, linked to the adoption of AI and DLT technologies.

For both technologies (AI and DLT), the majority of respondents indicated no cost saving/efficiency gain was anticipated, with a much higher percentage for DLT (81%) than AI (44%).

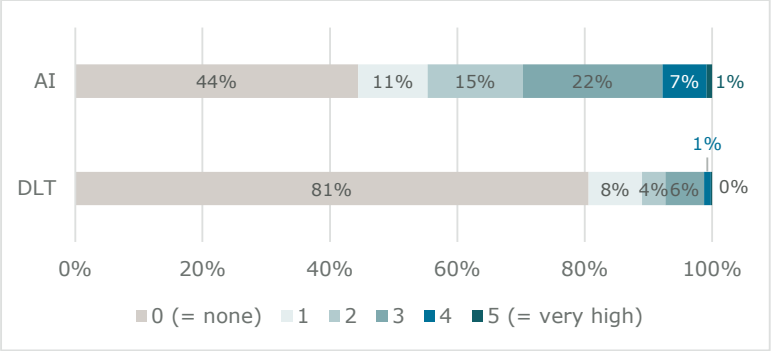

Figure 11: Anticipated cost savings/efficiency gains from the adoption of AI and DLT

However, when considering only those entities that invested in innovative technologies in 2024 at local or group level (see section 5.1 above), the situation significantly changes for AI: the percentage of those indicating no cost savings decreases to 23%, with the majority of respondents now indicating a score "2" or higher for cost savings/efficiency gains.

When comparing cost savings due to the adoption of AI or DLT, **the perceived cost savings/efficency gains from the adoption of AI technology are generally higher compared to those expected from the adoption of DLT technology** (more responses with scores >=3 for AI than for DLT).

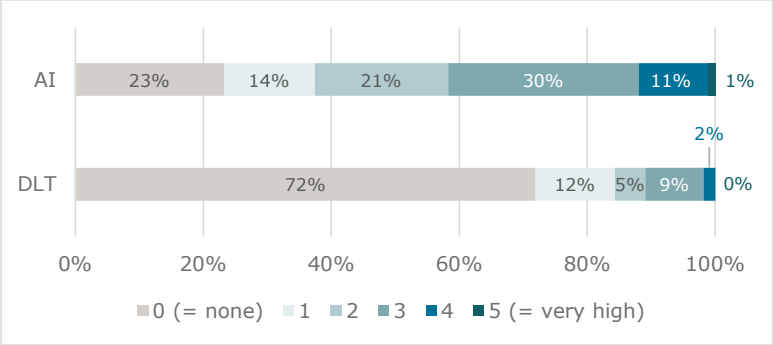

Figure 12: Anticipated cost savings/efficiency gains from the adoption of AI and DLT (focus on entities that invested in innovative technologies in 2024)

5.4 Anticipated revenue increases

The survey asked also to estimate, on a scale from 0 = none to 5 = very high, the revenue increases over the next 2-3 years, due to the adoption of AI and DLT technologies.

The findings are very similar to those described in the previous section related to cost savings, with the majority of respondents expecting no increased revenues, particularly for DLT compared to AI.

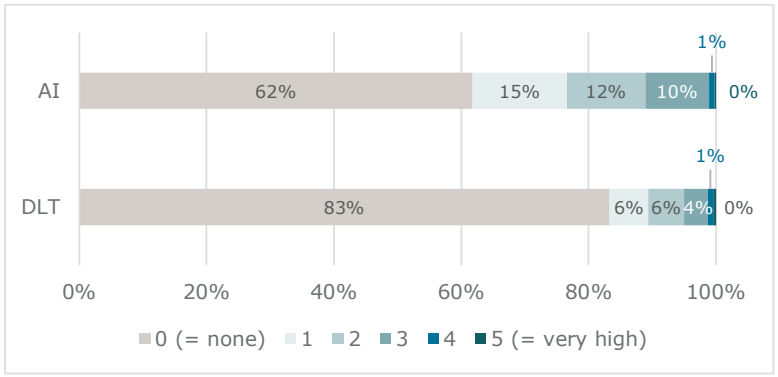

Figure 13: Anticipated revenue increases due to adoption of AI and DLT

When focusing only on entities that invested in innovative technologies in 2024 (see section 5.1), fewer respondents indicated that they do not expect a revenue increase due to investment in innovative technologies, while the anticipated revenue increases appear higher for AI compared to DLT.

When comparing cost savings with revenue increases, the graphs show that **these technologies are perceived more as a driver for cost savings rather than for revenue growth, particularly for AI**. This can be explained by the fact that most use cases for AI are used to improve internal efficiency (see section 6.3).

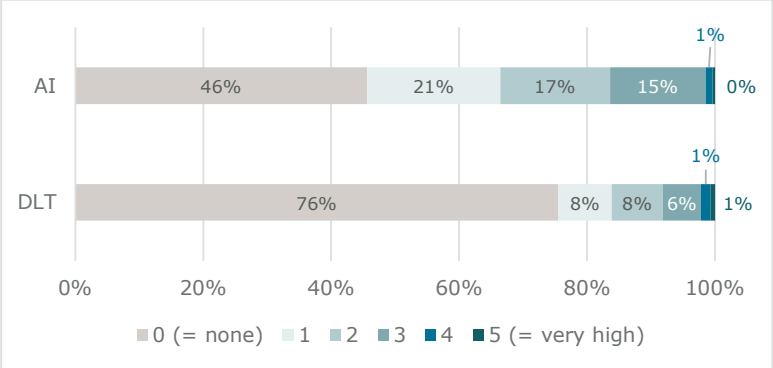

Figure 14: Anticipated revenue increases due to adoption of AI and DLT (focus on entities that invested in innovative technologies in 2024)

6. AI adoption

This chapter provides an overview of how entities limit access to public GenAI tools available on the Internet and describes the level of adoption of AI, the benefits and challenges associated with its use, the organisational aspects implemented with respect to AI development, and other general aspects linked to AI adoption.

6.1 Access to publicly available GenAI tools

Entities were asked if their employees could freely access public GenAI tools available on the Internet (e.g. ChatGPT, Gemini, Claude, etc.).

More than half of the respondents (58%) reported that access to these public AI tools is available to all employees. In contrast, 6% indicated that access is limited to a select few, while 36% denied access for their employees.

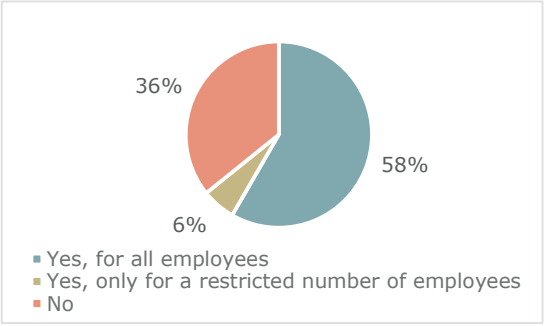

Figure 15: Access to public GenAI tools

Across the different types of institutions, we note that **credit institutions are more restrictive** with 54% of respondents blocking the access of their employees and 37% providing free access.

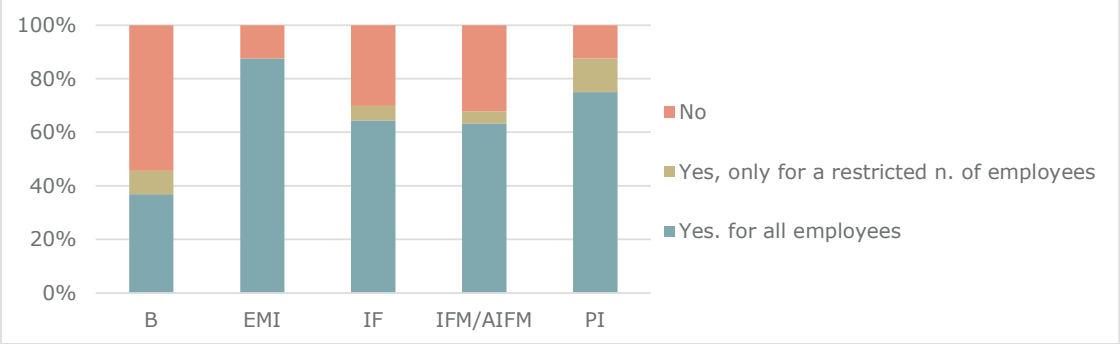

Figure 16: Access to public GenAI tools, depending on the type of entity.

We found that as the size of the company increases, access to these tools becomes more restricted.

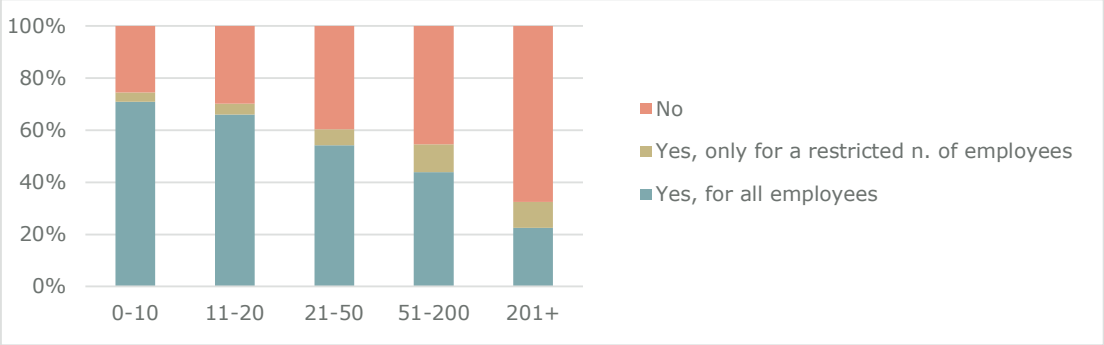

Figure 17: Access to public GenAI tools, depending on the number of employees.

Among entities that provide access to some or all of their employees, only 40% have either implemented a specific GenAI policy or modified their existing Internet policy to explicitly address the use of GenAI tools. Where it exists, the policy focuses on confidentiality, compliance and data protection topics. This leaves **60% of these entities without a dedicated policy on the subject.**

Besides, the disparity in GenAI policy definition largely depends on the size of the entity in Luxembourg, with larger entities more frequently having either a dedicated policy or at least a general policy that includes this subject.

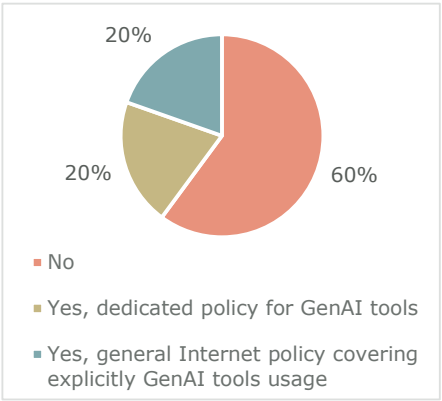

Figure 18: GenAI Policy

Figure 19: GenAI policy, by entity size

The *word cloud* below summarises the answers provided by respondents regarding the main elements covered in their GenAI policy. We can observe Confidentiality, Compliance and Data Protection among the most frequently cited topics.

Figure 20: "Word Cloud" illustrating the topics most commonly present in the AI Policy

6.2 Current status of adoption of AI technologies

In order to assess the level of AI adoption at the time of the study, respondents were asked to select among the options listed below, the one which best described their status:

- **A** - Concrete use cases (in production/development)
- **B** - Experimenting/Proof of concept (ongoing or planned in the next 12 months)
- **C** – Not planning to use AI technology in the next 12 months.

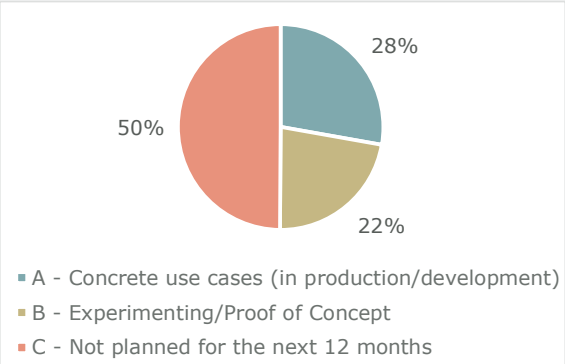

Figure 21: Current status of AI adoption

In total, **50% of all respondents are using or planning to use AI, i.e. either having concrete use cases in production/development (option A - 28%) or experimenting with AI technologies (option B - 22%)**. On the other hand, half of the respondents answered that they were not planning any use of AI in the next 12 months (option C).

The graph below provides a view on the status of AI adoption by entity type.

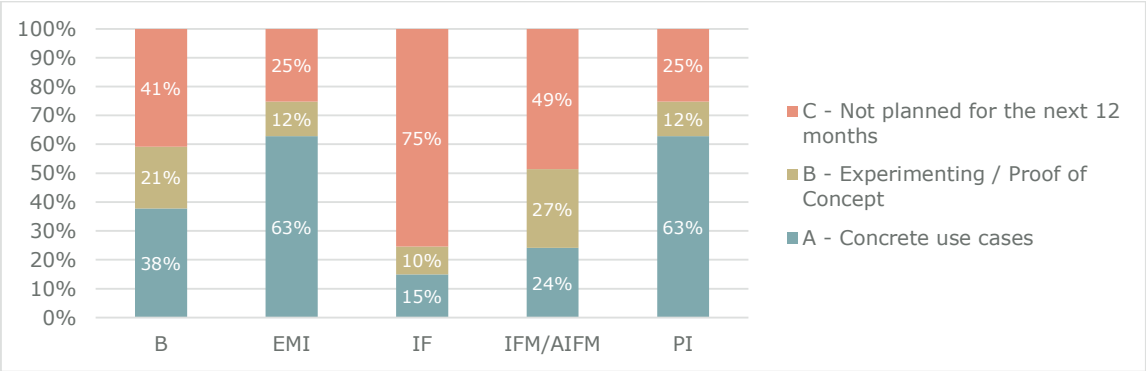

Figure 22: Status of AI adoption, by entity type

We observe that **EMI and PI seem more mature in terms of adoption of AI technologies, with 63% of them reporting concrete AI use cases in production/development, followed by B (38%),** IFM/AIFM (24%) and IF (15%).

Focusing on a scope of entities similar to the previous survey (B, PI, EMI), we observe an increase in the adoption of AI technologies, with 43% of these entities currently using AI in production/development (compared to 30% in the previous survey).

When looking at the portion of entities **experimenting or planning to experiment** with AI in the next 12 months, we note that IFM/AIFM lead with 27%, followed by B with 21%. EMI and PI are tied at 12%, while IF follow with only 10%.

Only the entities that answered they were using or experimenting with AI technologies (options A or B above) were asked to complete the rest of the AI related questions of the survey. The results presented in the rest of the document are based solely on the answers provided by these entities, representing 50% of all survey respondents[26].

6.3 AI benefits

Entities were asked to list the main benefits they observe from the use of AI technologies, within a predefined list of proposals. Among the top AI benefits identified by the survey (figure below), the **"improvement of internal processes" is ranking first** (69%), followed by "optimise operations/cost reduction" (56%), and "analyse vast amount of data" (52%), i.e. all benefits linked to internal efficiency. These results are similar to those of the previous survey, which also identified "improved internal efficiency" as the main benefit.

We also note that most of the entities that answered "none/still under evaluation" are only experimenting or planning to experiment with AI.

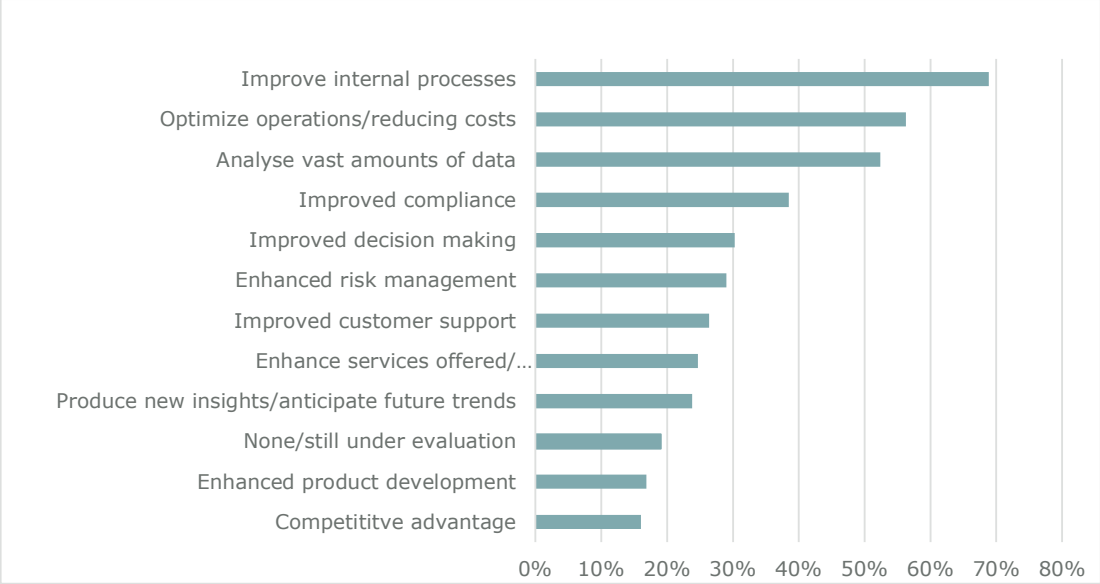

Figure 23: Main AI benefits

[26] Corresponding to 231 entities.

6.4 AI challenges

Similarly, surveyed entities were requested to list the main challenges they observe from the use of AI technologies, among a predefined list of options. The AI challenges identified are predominantly related to data, with **"Data quality"** (58%) **being the top challenge, followed by "Data protection"** (46%) **and "Data governance"** (40%). This trend, which remains more or less the same when focusing on entities of type B, PI, EMI (i.e. the same scope of the previous survey), is overall consistent with the results from the previous survey, where "data quality" was ranking first, and data governance and data protection were also among the top challenges.

Figure 24: Main AI challenges

While data quality remains among the top challenges identified irrespective of the type of entity responding to the survey, we note that PI have selected "Cybersecurity" as the top challenge (listed by 67% of PI respondents[27]).

We note also that "Compliance with regulation" is among the top challenges, probably reflecting the challenges represented by the new AI regulation, the AI Act[28].

Similarly to AI benefits, we note that also in relation to AI challenges, most of the entities that answered "none/still under evaluation" are only experimenting or planning to experiment with AI.

6.5 Organisation

The majority (63%) of the respondents which indicated they were using AI (either in development/production or in experimental phase) **have a dedicated team working only on AI related projects/ development activities ("data science team"). In most cases (55%), this**

[27] Corresponding to 8 entities out of 12.

[28] The survey was conducted before the entry into force of the AI Act.

team is located at group level, while for 5% it is present at both group and local levels**.** However, 3% of respondents indicated having a data science team only at local level.

Considering that, as reported below, data science teams at group level are usually larger, these figures confirm the **general tendency** observed in the previous survey **to capitalise on group expertise for AI related development activities**.

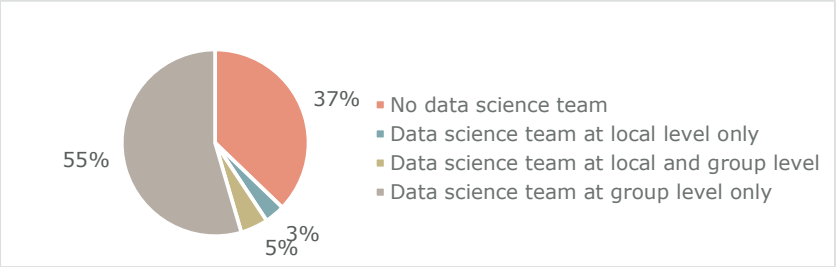

Figure 25: Data science team location

We also note that **the percentage of respondents having no data science team has increased compared to the previous survey[29], reflecting the availability of "ready to use" solutions such as GenAI tools that do not require advanced AI technical skills for their implementation**.

Data science teams at local level are rather small, with less than 10 employees. Conversely, data science teams at group level are generally larger (figure 26). Furthermore, compared to the previous survey, we observe a significant increase of group data science teams with more than 10 employees.[30].

We also note that B and IFM/AIFM are the only entities reporting data science teams at group level with more than 100 people.

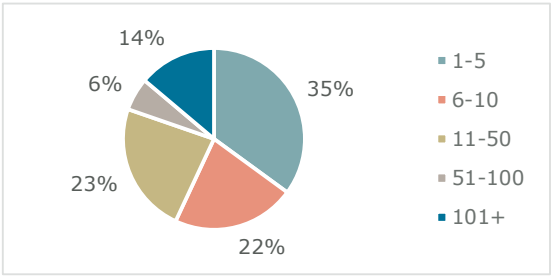

Figure 26: N. of employees of the data science team (group level)

As regards the staff composing the data science teams, **around a quarter (26%) of the respondents reported difficulties recruiting on the local market, confirming the current scarcity of skilled resources in the AI field.**

Data science teams most frequently report to the IT function (36%), followed by other functions such as group AI/Data Analytics function, Chief Information Office, senior management, etc. (34%). They less often report to a business line (17%) or a combination of IT and business lines (13%).

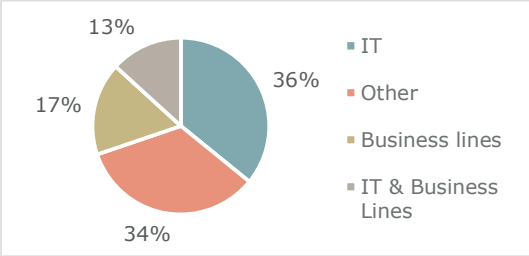

Figure 27: Data science team - reporting line

[29] Considering only B, PI, EMI, the percentage of respondents not having a data science team is 24%, compared to 15% in the previous survey.

[30] Considering only B, PI, EMI, 57% of respondents reported data science teams with more than 10 employees at group level, compared to 34% in the previous survey.

Regarding AI trainings, the vast majority (84%) of respondents have either already implemented or plan to implement a range of AI training programs for their employees, ranging from basic awareness to advanced AI skills. This confirms the importance of AI trainings, also in the context of AI literacy obligations included in the AI Act[31].

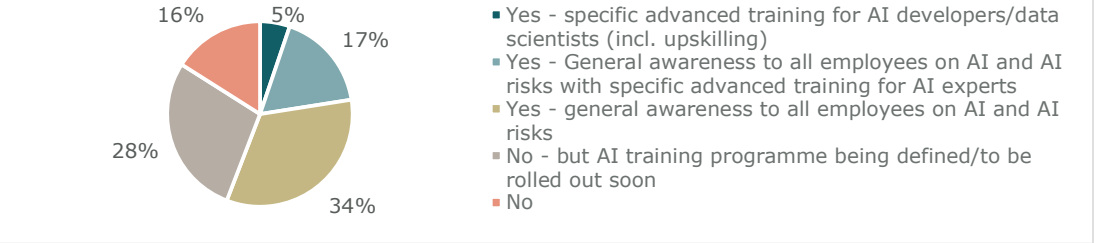

Figure 28: AI trainings

6.6 Data and governance

Less than half of respondents (43%) indicated having a formally approved AI policy, either a general policy covering explicitly AI aspects (24%) or a dedicated AI policy (19%). Nevertheless, these figures (which remain similar when focusing only on B, EMI, PI) represent a **relevant increase in the portion of respondents with an AI policy compared to the previous survey** (where only 22% of respondents had an AI ethical policy in place) and indicate an **improved level of maturity**.

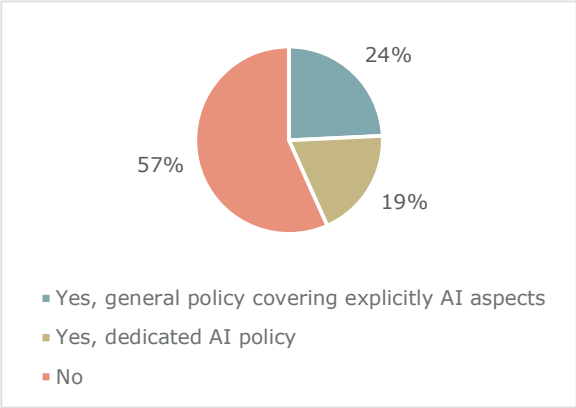

Figure 29: Existence of a formally approved AI policy

Among the main aspects covered by the AI policy, there are AI usage rules, as well as data protection and AI ethical aspects (e.g. bias and fairness).

[31] Art. 4 of AI Act.

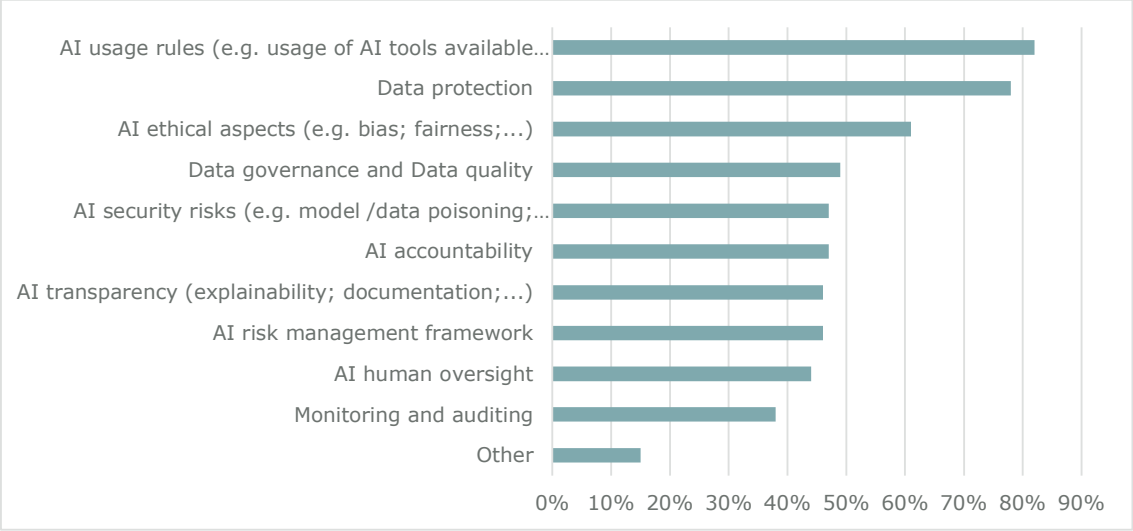

Figure 30: Aspects covered by the AI policy

With regard to the functions involved in the AI oversight, 81% of respondents indicated the involvement of the information security function, followed by Compliance, DPO (Data Protection Officer) and Risk functions.

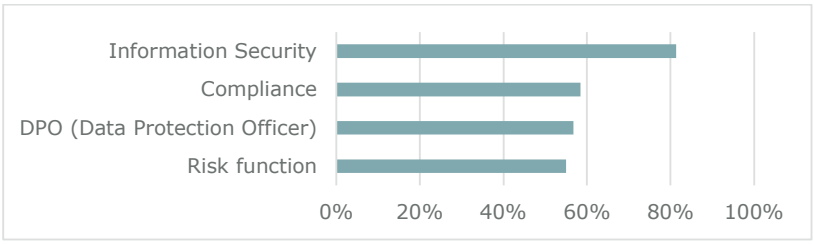

Figure 31: Functions involved in the AI oversight

Compared to the previous survey[32], the involvement of information security and DPO functions has slightly decreased, while the involvement of the risk function remained unvaried.

6.7 Security and robustness

In relation to **security measures** for AI specific vulnerabilities and security attacks (e.g. data poisoning, model poisoning, adversarial attacks, model evasion attacks, confidentiality attacks, model flaws, etc.)[33], **more than half (54%)** of the respondents **indicated having taken specific security measures while 16% have not**[34].

[32] On a scope composed of B, PI, EMI, information security is involved in 82% of cases (compared to 88% of the previous survey), followed by DPO with 71% (83% in the previous survey) and Risk with 63% (same as in the previous survey).

[33] See the definitions available in the glossary under "ML security".

[34] The majority of respondents indicating not having taken security measures in relation to AI specific vulnerabilities is constituted by IFM/AIFM. However, most of these do not have concrete AI use cases and are only in experimenting mode.

On the other hand, a significant portion (24%) of respondents indicated that these security measures were not applicable while 6% indicated that they did not know. In both of these cases, the corresponding use cases were mostly in experimental phase.

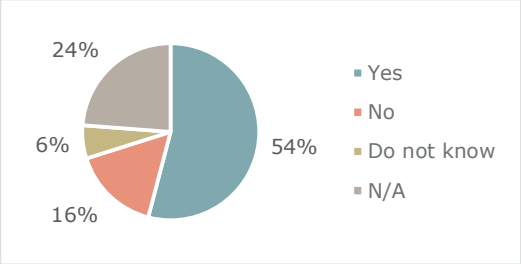

Figure 32: Security measures taken in relation to AI vulnerabilities

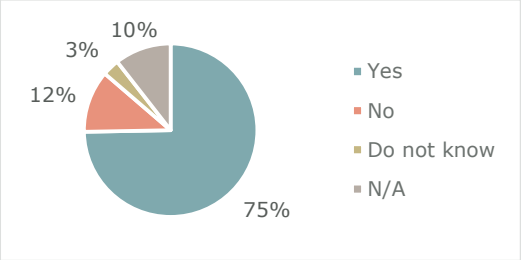

Figure 33: Security measures taken in relation to AI vulnerabilities (concrete use cases only)

Indeed, **when focusing only on concrete use cases** in production or development (excluding those in experimental stage), **the percentage of respondents indicating having taken security measures in relation to specific AI vulnerabilities increases to 75%,** while the percentage of those that have not taken specific security measures decreases to 12%.

Comparing results with those of the previous survey, we note that the **percentage of respondents indicating that they have implemented security measures in relation to AI vulnerabilities has increased**[35] **overall**, indicating an **improved level of maturity**.

6.8 AI technical infrastructure

With regard to the technical infrastructure supporting the AI processes, **respondents are primarily using commercial cloud solutions (45%)**, while private/dedicated infrastructures are used by 22% of respondents, and 24% indicated using hybrid (cloud and local) environments.

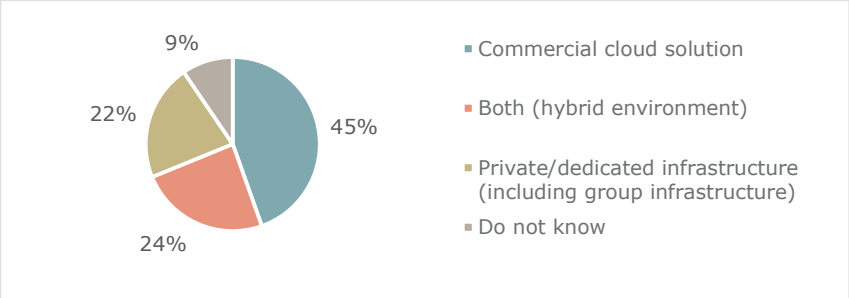

Figure 34: Technical AI infrastructures

The cloud solutions are especially privileged by IF and IFM/AIFM (used by 67% and 52% of these entities, respectively), while if we consider only B, PI, EMI (same scope of the previous survey) we note that the use of cloud decreases to 28% while hybrid environments increase to 35%. These

[35] Considering only B, PI, EMI, 66% of respondents indicated to have taken security measures specific for AI vulnerabilities, while it was close to 50% in the previous survey.

figures represent nevertheless an **increase of cloud solutions compared to the previous survey**[36].

The increased use of cloud solutions is linked in the majority of cases (75%) to the use of GenAI solutions.

Among the entities using private or hybrid infrastructures, the majority (60%) reported no difficulty in procuring the appropriate hardware, while 14% encountered challenges. Notably, **71% of the entities which had difficulties in procuring the appropriate hardware had GenAI use cases**.

6.9 AI lifecycle

With regard to the change/development process for AI solutions, **the large majority (71%) of survey respondents indicated that they have not applied any change to the current change/development process.** 19% of respondents instead indicated having adapted the existing process to AI specificities while the remaining 10% have implemented a separate process for AI developments.

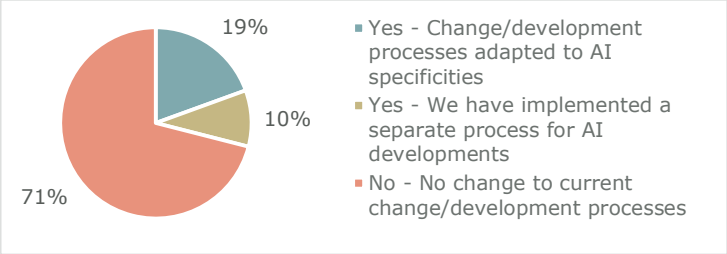

Figure 35: AI change management/ development process

Besides, we note that the percentage of those having adapted their change management process or having implemented a separate change management process increases when considering only those entities having concrete use cases in development/production (to 36% and 19%, respectively).

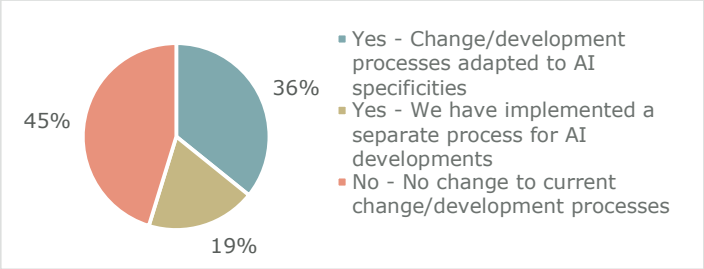

Figure 36: AI change management/ development process (concrete use cases only)

Compared to the previous survey, we note that **the portion of entities implementing a separate change management/development process for AI developments has decreased**[37].

[36] In the previous survey, cloud and hybrid environments represented 14% and 32% of responses, respectively.

[37] Considering only B, PI, EMI, the portion of entities implementing a separate process for AI developments is 11%, while in the previous survey 26% of respondents indicated having an ad-hoc change management process for AI.

7. Use cases: general aspects

The previous chapter provided an overview of the status of AI adoption of entities in the scope of the survey. For entities that indicated they were using or planning to use/experiment with AI (i.e. those that selected either option A or B as defined in section 6.2 above), the survey questionnaire offered the opportunity to describe more in detail how AI was concretely used within the company. This could be done by submitting one or more "use cases" via dedicated tabs of the questionnaire[38].

The following chapters focus on the AI use cases reported by the survey respondents. In particular, this chapter (**chapter 7**) presents some general aspects of all AI use cases submitted, while the next chapter (**chapter 8**) focuses on those use cases involving GenAI technology. In **chapter 9**, we dig deeper into the use cases using machine learning and finally, **chapter 10** presents the trustworthiness aspects of all use cases.

Considering that the submission of use cases was optional, and not all detailed questions within the use cases were required, the analysis and graphs presented in the following chapters are based solely on the responses received.

7.1 AI technologies

Of the 461 survey respondents, **36%** (**168 entities**) **reported at least one AI use case**[39].

In total, these respondents reported **402 distinct AI use cases.**

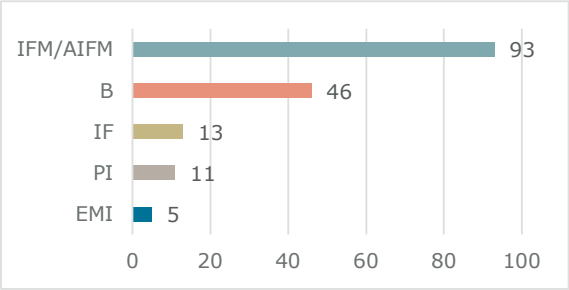

Figure 38: N. of entities reporting at least one AI use case

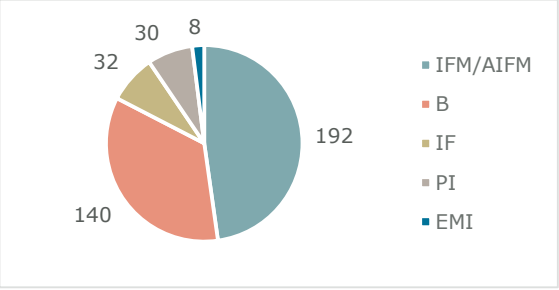

Figure 37: N. of reported AI use cases, by entity type

When examining the specific AI technologies employed by the 402 reported use cases, we observe that the **majority involve GenAI technology** (61% of use cases), **followed by Natural Language Processing (NLP)** (30% of use cases), and **machine learning (ML)** (28% of use cases). Expert systems, Intelligent Process Automation (IPA), and computer vision are instead much less common.

It is important to note that respondents were allowed to select multiple AI technologies for a single use case. Notably, we observe that **most use cases involving NLP also involve GenAI**, which may reflect challenges in distinguishing between these two technologies. **Due to this overlap, the**

[38] It should be noted that given that the possibility to describe use cases was optional, not all entities which selected "option A" (i.e. indicated to have use cases in production/development) submitted use cases via the dedicated tabs of the questionnaire. On the other hand, many entities which selected "option B" (i.e. indicated "experimenting/Proof of Concept (ongoing or planned within the next 12 months)" did submit use cases. To ensure consistency, entities that selected "option B" but submitted at least one use case with status "in production" were then switched to "option A".

[39] I.e. reported at least one AI use case in either experimenting or development or production stage.

remainder of this document focuses more on GenAI and ML technologies and provides comparative analysis where possible.

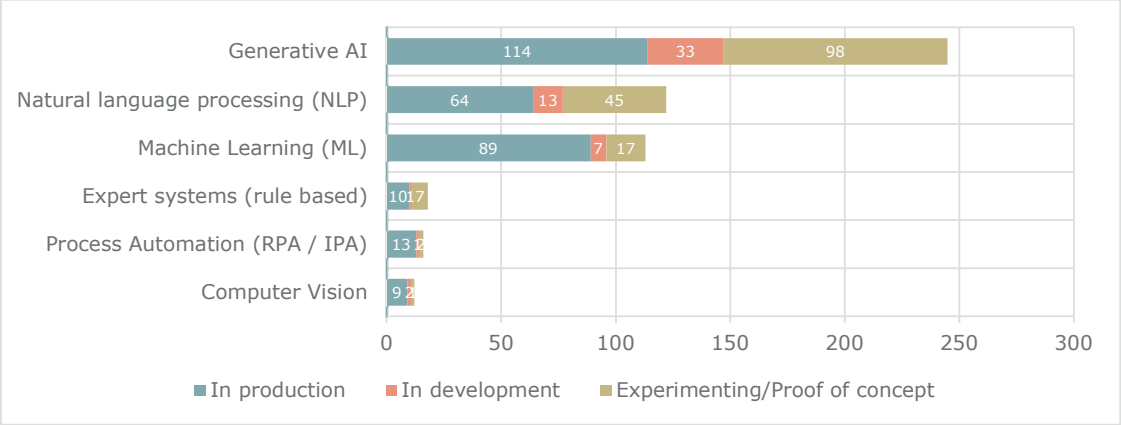

Figure 39: N. of use cases per type of AI technology (with development status)

Approximately half of the use cases involving GenAI are still at an experimental/proof-of-concept stage or under development, indicating a **more recent level of adoption of GenAI compared to ML, which appears to be more mature in terms of adoption (i.e. with a higher portion of use cases already in production)**.

The paragraphs above provide an overview on the level of utilisation of various AI technologies in terms of number of use cases employing each technology. From a different perspective, by examining the entities reporting these use cases, we can extrapolate the level of adoption of each technology among the respondent entities. For instance, focusing the analysis on GenAI and ML, we can compare the number of entities reporting at least one use case involving GenAI with the number of entities reporting at least one use case involving ML.

When comparing the portion of entities using GenAI versus ML, we observe that **28% of all survey respondents** (129 entities out of 461) **reported at least one use case involving GenAI**[40], and **12%** (57 entities) **reported at least one use case involving ML**[40].

These figures indicate a much wider adoption of GenAI compared to ML technology.

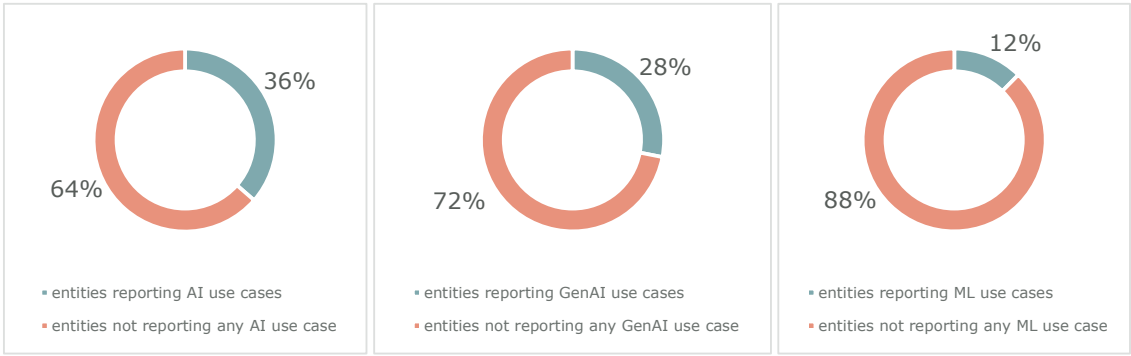

Figure 40: Entities reporting (from the left to the right) at least one AI, GenAI, ML use case

[40] In production, development or experimental stage.

AI use cases

Across the different types of entities, the percentage of entities reporting at least one AI use case[40]) is higher for PI (69%) and EMI (63%), followed by B (45%).

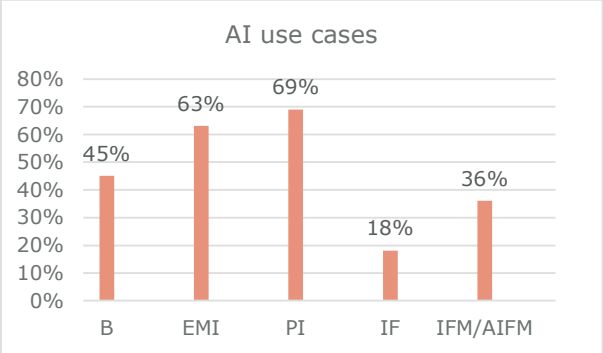

GenAI use cases

When examining the use of GenAI (in terms of percentage of entities reporting at least one use case[40] involving GenAI), we find that PI are at the forefront with 50% followed by B with 32%, and then IFM/AIFM with 29%.

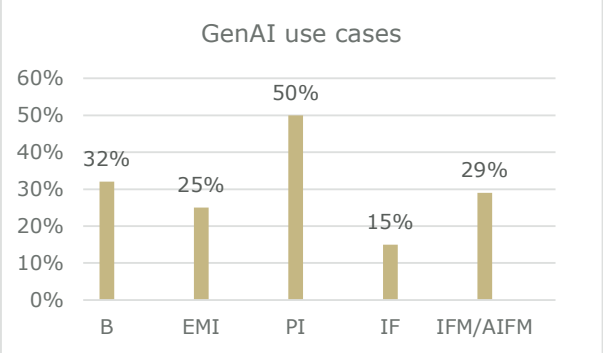

ML use cases

Regarding the use of ML (in terms of percentage of entities reporting at least one use case[40] involving ML), we see that EMI are leading with 50% followed by PI with 44%, and then B with 24%.

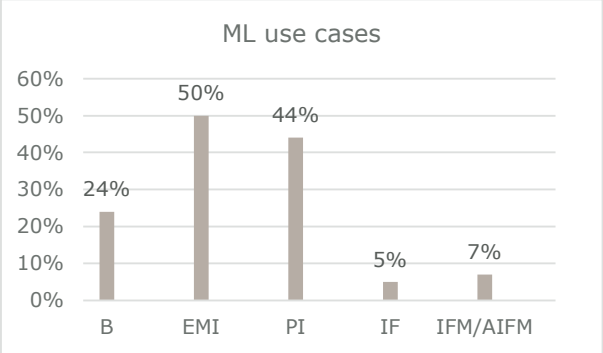

Figure 41: Percentage of entities reporting at least one AI, GenAI, ML use case - by type of entity

7.2 Use case categories

Respondents were asked to select **one or multiple categories** that best represented their use cases, from a predefined set.

The top five use case categories reported were **Search/summarise information (43%), Process automation (30%), Chatbot and virtual assistant (27%), Text context generation (27%), and Translation (19%)**. These top five categories remain consistent across all types of entities. However, for PI and EMI there is an exception: their top five use case categories include "AML/Fraud detection" with, on the other hand, a lower representation of the category "Process automation". Specifically, **the "AML/Fraud detection" category remains the top use case for EMI and the second use case for PI, confirming its relevance for these types of entities, especially in the context of payments.** Additionally for EMI, only five categories are reported, with "Know your customer" ranking second[41].

[41] See Annex 12.1 for more details about the use case categories by type of entity.

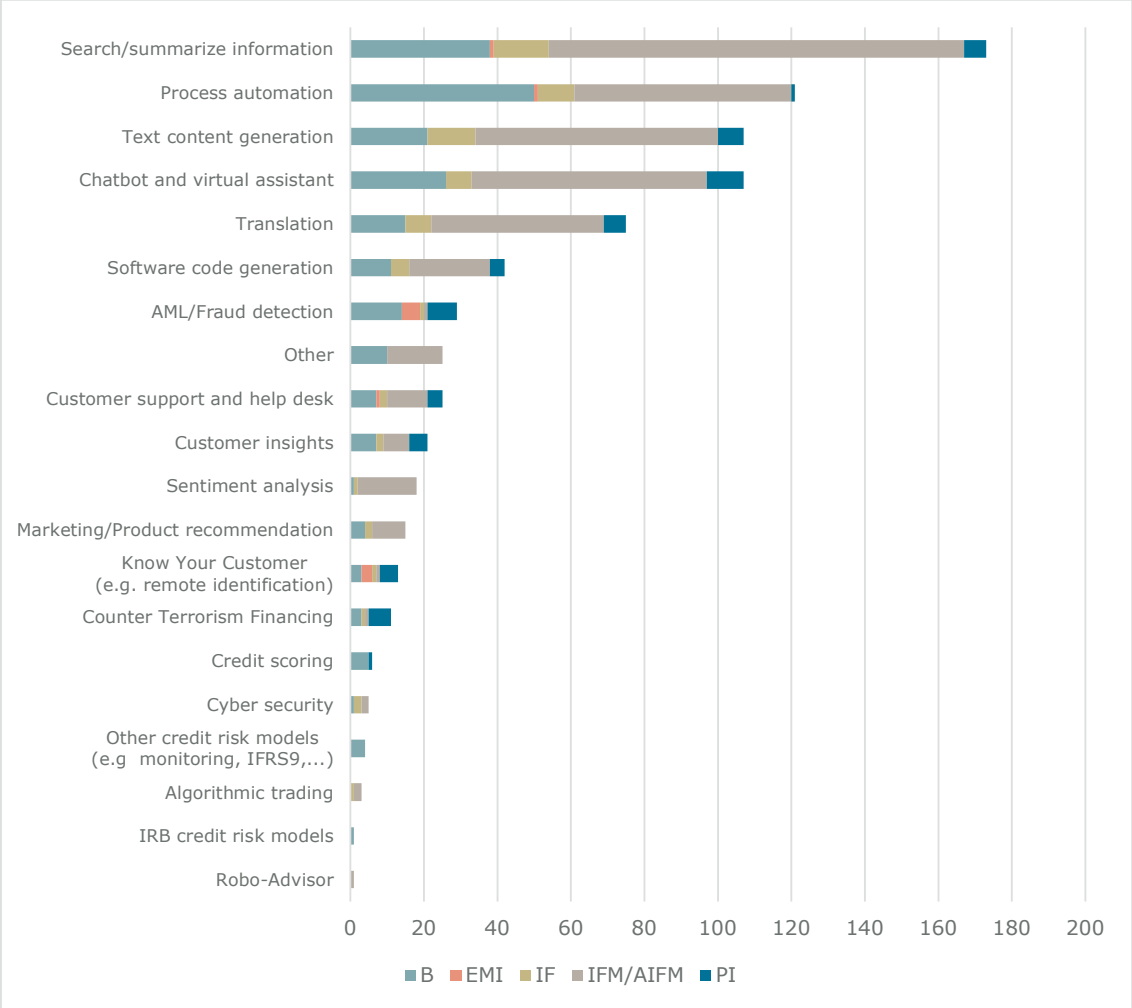

Figure 42: Use case categories, by type of entity

The top five categories correspond to use cases that typically leverage GenAI rather than ML, with a small exception for the category 'process automation' for which the use of GenAI is accompanied by a significant use of machine learning. Meanwhile, **ML remains predominantly used in risk and compliance solutions**, such as AML/fraud detection, Know Your Customer (e.g. remote identification) and counter terrorism financing.

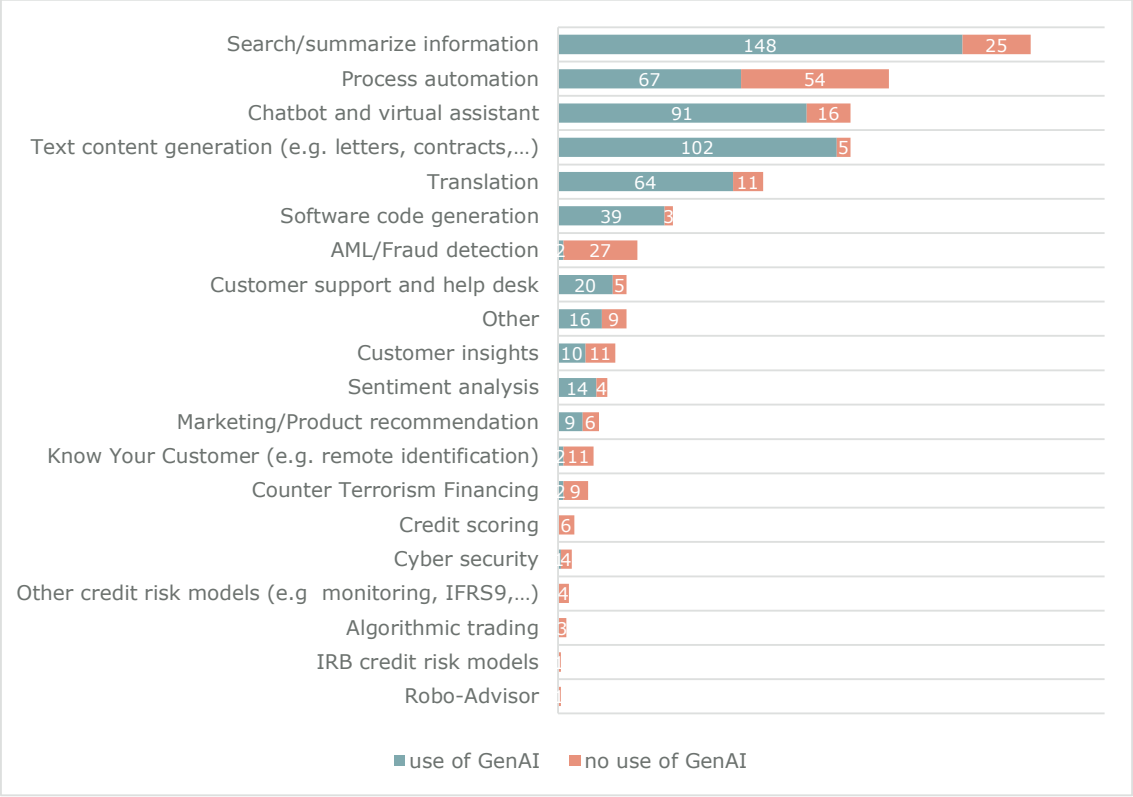

Figure 43: Use case categories - split if using GenAI or not

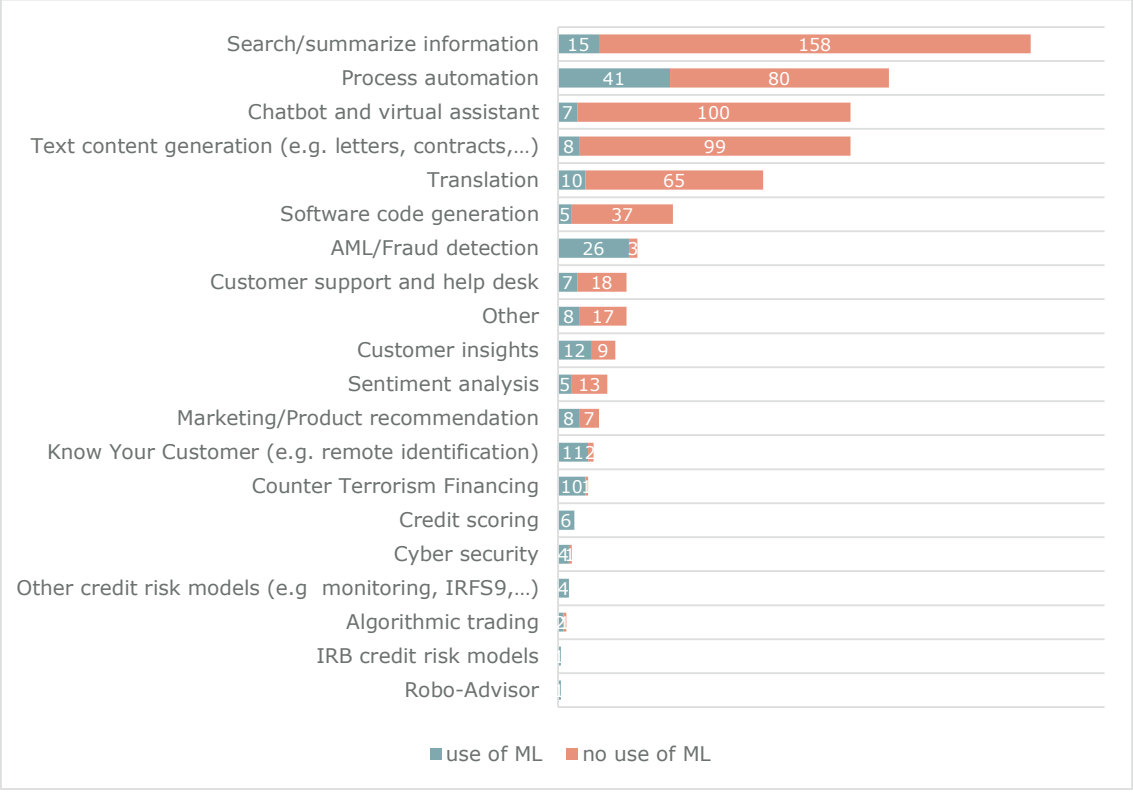

Figure 44: Use case categories - split if using ML or not

When comparing with the previous survey (focusing only on B, PI, EMI), the "AML/Fraud detection" category drops to fifth position, having previously ranked first. Meanwhile, we observe that the new categories with higher rankings are those mostly associated with GenAI (see figure 47 below).

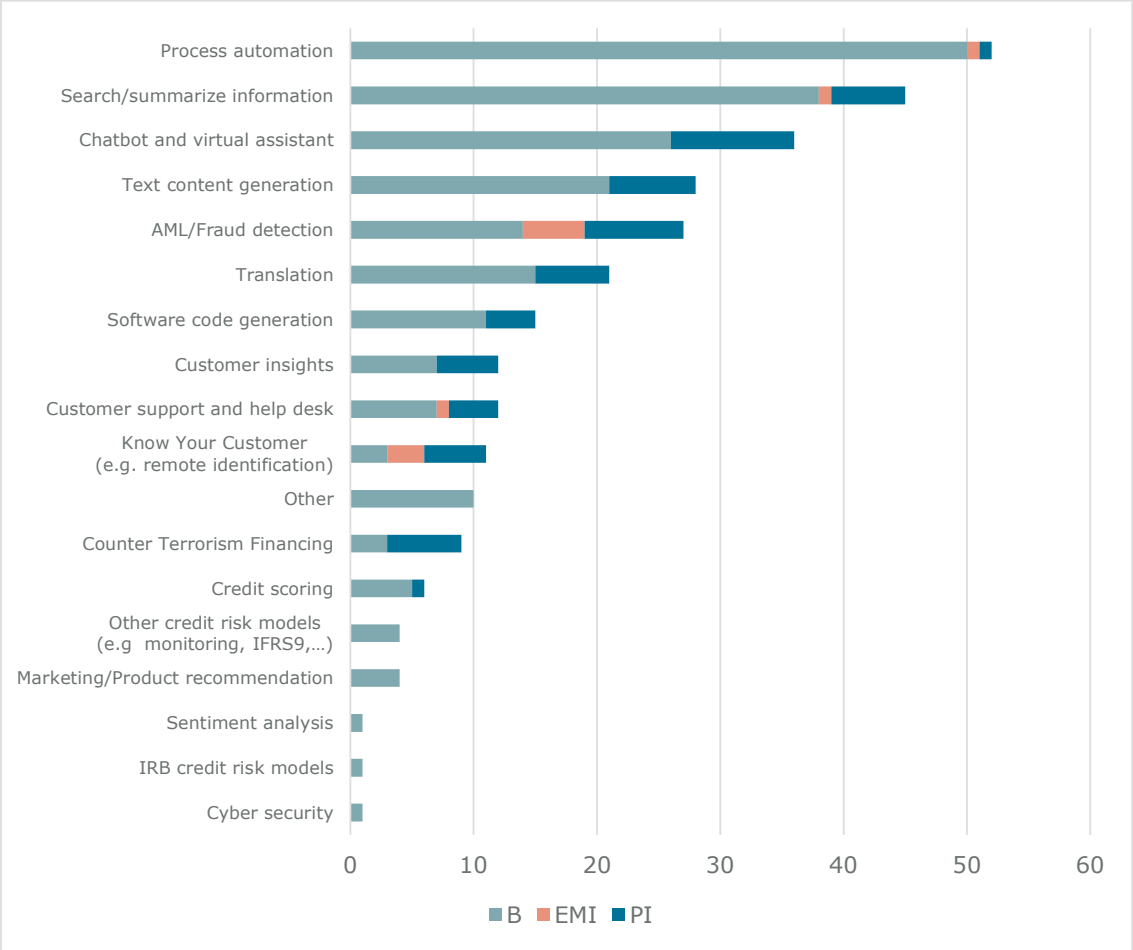

Figure 45: Use case categories (B EMI PI only)

7.3 Development approach

Overall, **54% of reported use cases are already in production**, a tendency that is confirmed for the top five categories of the reported use cases ("Search/summarise information", "Process automation", "Text content generation", "Chatbot and virtual assistant" and "Translation").

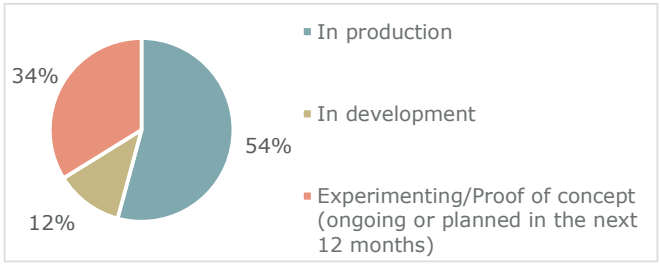

Figure 46: Use cases deployment status

At the same time, we observe that the majority of the use cases in experimental stage are included in the same top five categories, which are mostly GenAI categories. Besides, for the five less commonly reported categories ("Robo-advisor", "IRB credit risk model", "Algorithm trading", "Other credit risk models" and "Cyber security"), as well as other compliance related categories ("AML/Fraud detection", "Know your customer"), the large majority (90%) of reported use cases are in production.

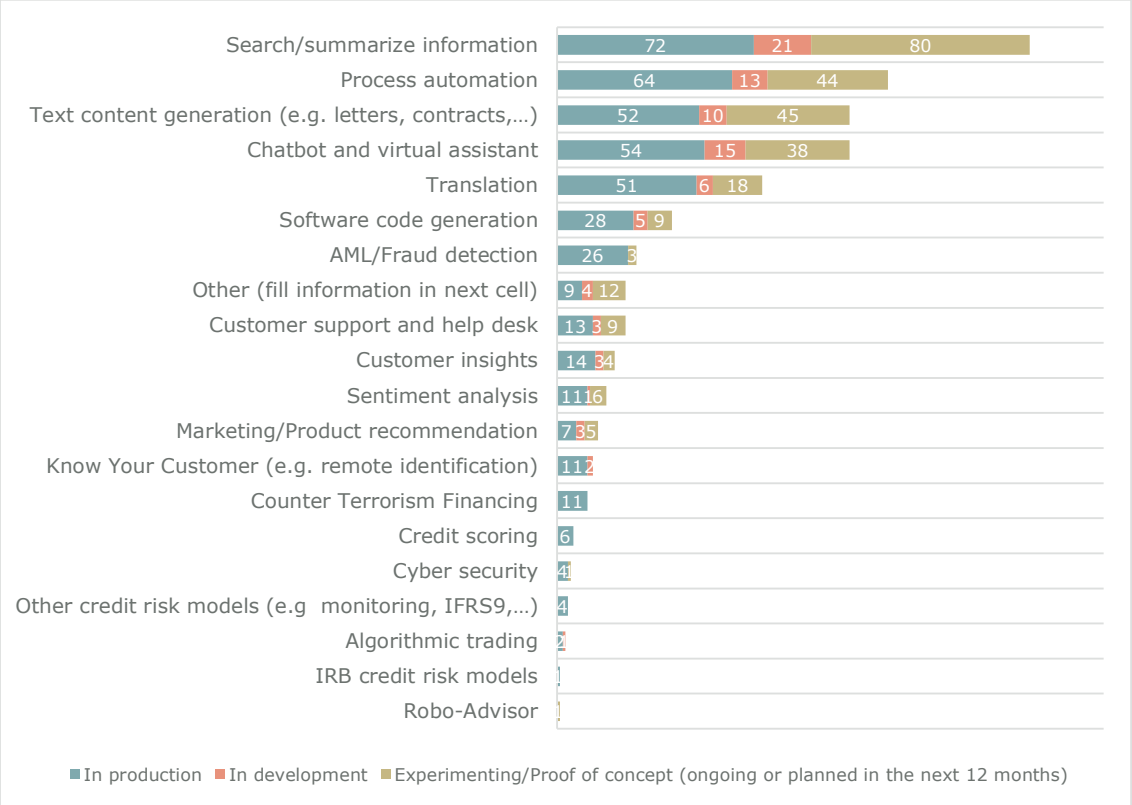

Figure 47: Use cases deployment status

Among the different types of entities, **EMI and PI seem more advanced having respectively 88% and 74% of the reported use cases already in production**.

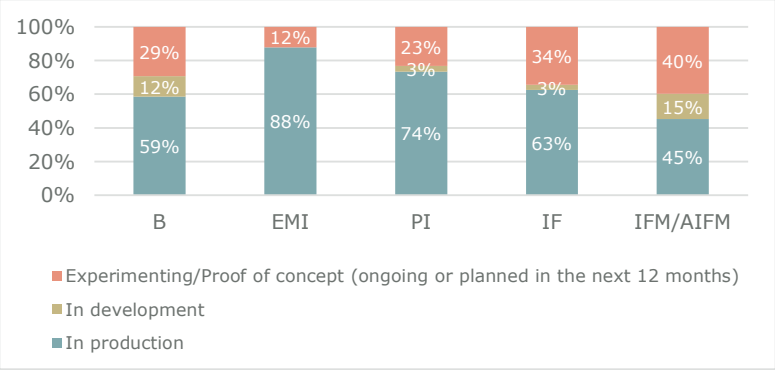

Figure 48: Deployment status of the reported use cases, by type of entity

The vast majority (84%) of the use cases employ AI models configured as "primary" as opposed to secondary/ "challenger" models[42]. These results are in line with those from the previous survey[43].

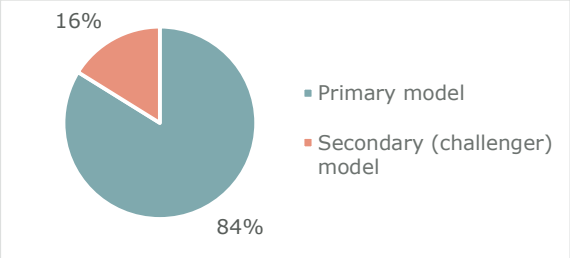

Figure 49: Deployment approach

In terms of development approach, there is a global trend towards developing AI solutions internally[44]. Notably, **60% of use cases are developed internally.**

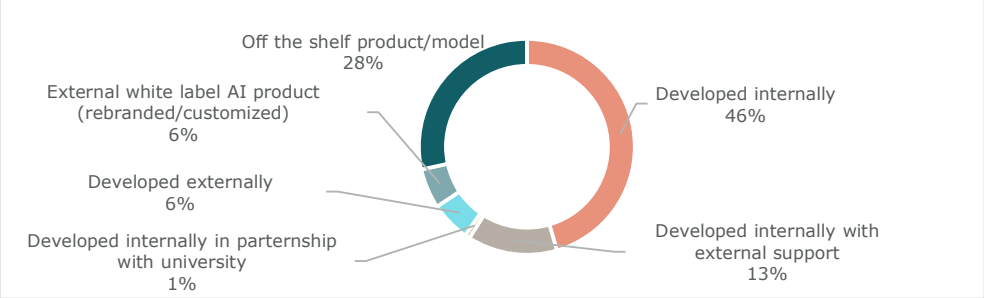

Figure 50: Development approach

[42] A "challenger" model is a model that runs in production in parallel with the current model (or traditional system) for a certain period to allow a comparison of the results. If the challenger model produces better results, it may be promoted to become the primary model.

[43] Considering only B, EMI, PI, 85% of use cases employ AI models configured as "primary" models, while in the previous survey it was 82%.

[44] Internally developed solutions cover in-house solutions, as well as solutions developed internally in partnership with a university or with external support.

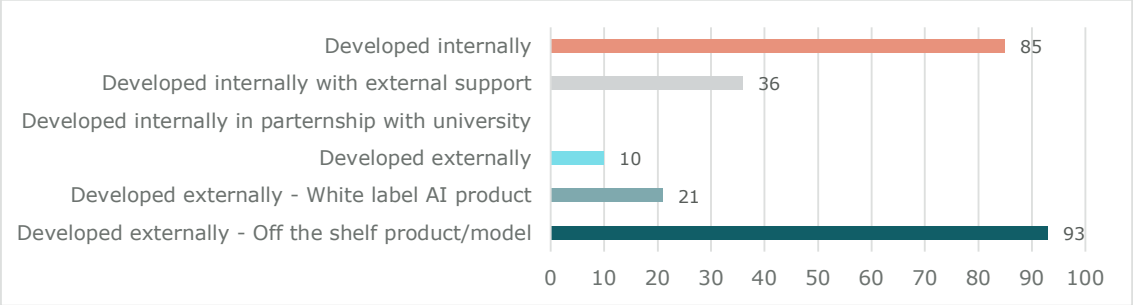

Figure 51: Development approach for GenAI use cases

About half (51%) of the GenAI use cases reported are developed externally[45], and most of these are off the shelf products[46]. In contrast, **the majority (76%) of ML use cases are developed internally[47],** which is still the same trend as in the previous survey[48].

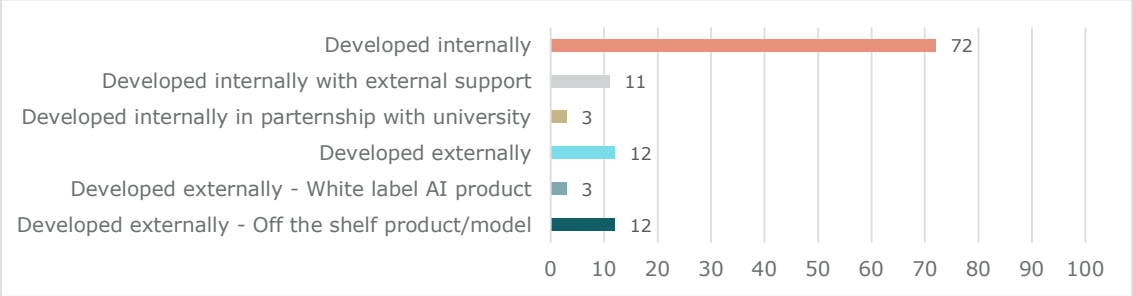

Figure 52: Development approach for ML use cases

In the reported use cases, **the majority of AI models are trained using internal data (51%) or a mix of internal and external/public data (23%).** This trend is even more pronounced in ML use cases, where 61% of models are trained exclusively on internal data. This underscores that **most machine learning systems rely primarily on internal data for training**, confirming the trend identified in the previous survey[49].

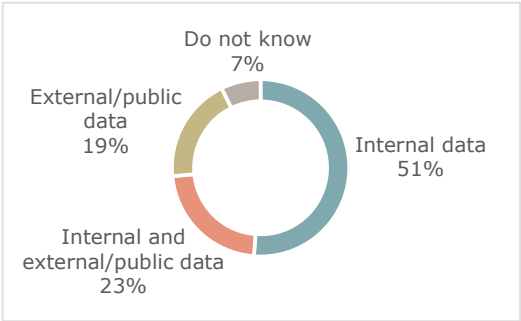

Figure 53: Type of data used to train the AI system

It is important to note that when respondents indicated using only internal data for GenAI use

[45] Corresponding to the categories "Developed externally" (4%), "Developed externally – White label product" (9%), "Developed externally – Off the shelf product" (38%).

[46] As we will see in chapter 8, most of these are general purpose LLM.

[47] Corresponding to the categories "Developed internally" (64%), "Developed internally with external support" (10%), "Developed internally with university" (3%).

[48] The previous survey included an analysis of the use of ML but did not include any specific analysis for GenAI.

[49] Considering only ML use cases reported by B, PI, EMI (same scope of the previous survey, which covered only ML use cases), 65% of use cases use only internal data (compared to 62% in the previous survey), 26% a mix of internal and external/public data (compared to 28% in the previous survey), and 10% only external data (same as in the previous survey).

cases (47%), they may not have accounted for the data used to initially train the GenAI model (e.g. LLM[50]) by the model provider before its integration into the use case.

7.4 Client facing versus internal

Overall, the vast majority (92%) of all reported use cases[51] are only for internal use, while only 8% are client facing solutions.

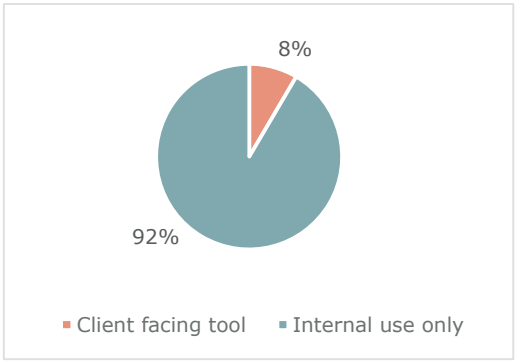

Figure 54: Internal vs client facing use cases

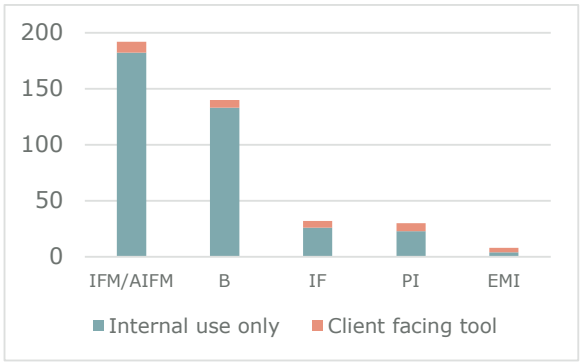

Figure 55: Internal vs client facing use cases (by type of entity)

Unsurprisingly, the categories of use cases with the highest portion of client facing use cases are "Chatbot and virtual assistant" and "Customer support and help desk".

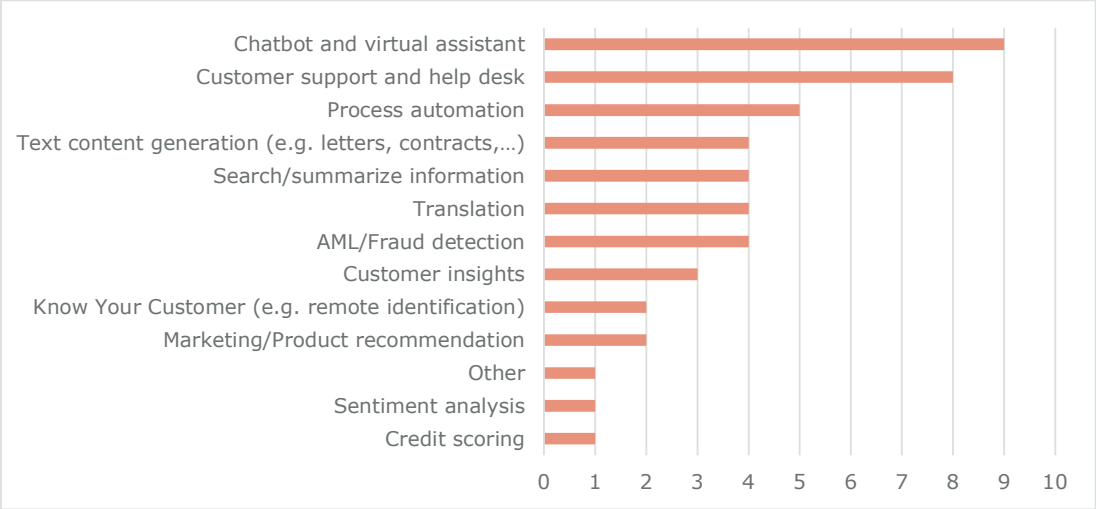

Figure 56: Use case categories, for client facing solutions

[50] See section 8.1 for more details on the use of Large Language Models (LLM).

[51] I.e. 368 use cases out of 402 are not client facing.

8. Use cases: focus on GenAI

This chapter focuses on the subset of 245 use cases, among all the use cases reported by the survey respondents, which rely on GenAI technology[52].

8.1 Types of Generative AI

Nearly all (94%) reported use cases using GenAI are using Large Language Models (LLM), with this proportion remaining consistent across different types of entities. The 6% of GenAI use cases which do not use LLM technology are mainly use cases combining audio/video with text.

8.2 Open source vs commercial models

The vast majority (75%) of reported GenAI use cases rely solely on commercial models, 11% are using open-source models, and 11% are using both.

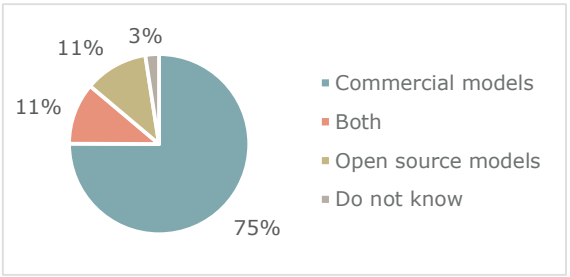

Figure 57: GenAI use cases: commercial models Vs open-source models

Besides, **only 15% of all the reported GenAI use cases use models that are fine-tuned** for the entity, a percentage which does not change significantly when using commercial models or open source models.

8.3 Retrieval Augmented generation (RAG)

Retrieval Augmented Generation (RAG), often referred to as "grounding", is a technique enabling LLMs to fetch information from user supplied documentation in order to "ground" the model on a set of external, verifiable facts, and ultimately improve the accuracy of the model output.

Regarding the use of RAG techniques in the reported GenAI use cases, approaches are mixed. Specifically, **40% of the reported GenAI use cases do not employ RAG, while 36% incorporate these techniques**[53]. When entities are using RAG, this is **mainly based on internal data**, and rarely on external data sources.

Unsurprisingly, fine-tuning as well as RAG techniques are most commonly applied in use cases related to "search/summarise information", "chatbot and virtual assistant" and "text content generation", when contextual information may strongly influence the quality of the generated output.

[52] These use cases were reported by IFM/AIFM (149 use cases reported by 75 entities), followed by B (58 use cases reported by 33 entities), IF (21 use cases reported by 11 entities), PI (15 uses cases reported by 8 entities) and EMI (2 use cases reported by 2 entities).

[53] The category "Other" corresponds to "*Do not know*", "*Not applicable*" or "*blank*" answers.

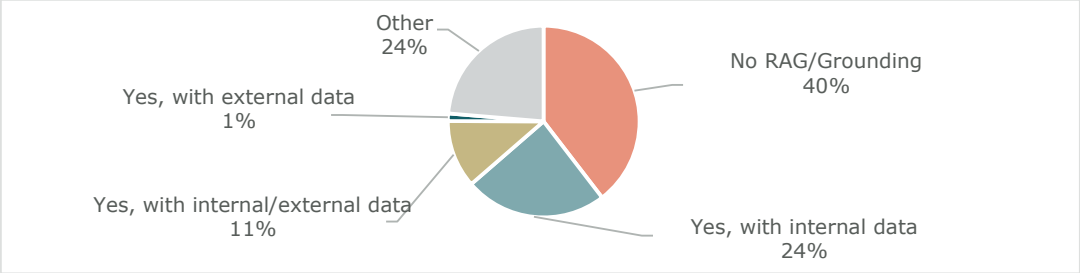

Figure 58: Usage of RAG[54]

9. Use cases: focus on ML

Among all use cases reported, 113 use cases rely on machine learning[55]. This chapter focuses on some specificities of these use cases.

9.1 Type of ML algorithms

For ML use cases, respondents were asked to further specify the type of ML algorithms employed (according to the type of problem addressed). For each use case, multiple types of ML algorithms could be selected. Classification algorithms are the most widely used across all ML use cases.

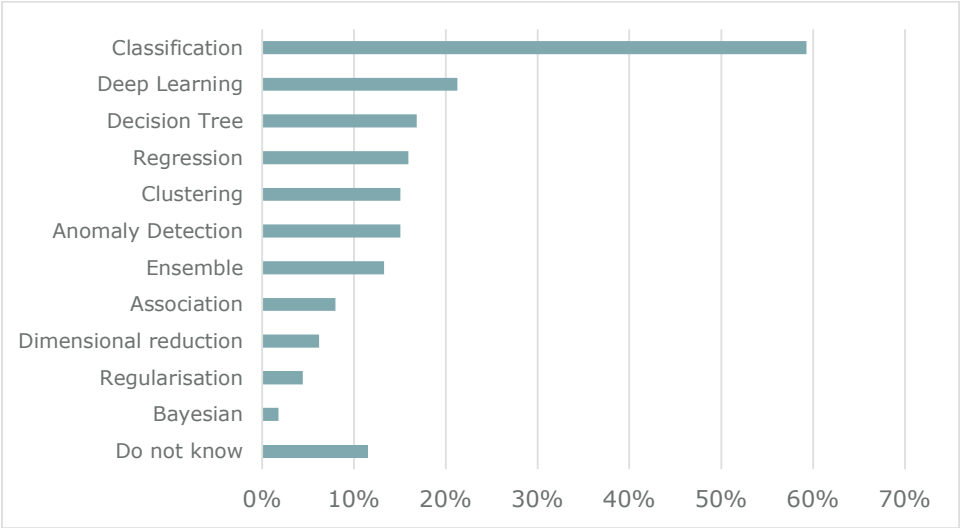

Figure 59: Type of ML algorithms

[54] The category "Other" corresponds to "*Do not know*", "*Not applicable*" or "*blank*" answers.

[55] Specifically, ML use cases were reported by B (61 use cases reported by 25 entities), followed by IFM/AIFM (23 use cases reported by 17 entities), PI (14 use cases reported by 7 entities), IF (8 uses cases reported by 4 entities) and EMI (7 use cases reported by 4 entities).

9.2 Type of learning

The vast majority (73%) of ML use cases employ centralised learning. Reinforcement learning is used in a quarter of cases (25%), while transfer learning (7%) and federated learning (4%) are much less common. In some instances, multiple types of learnings are combined (8%), always including centralised learning.

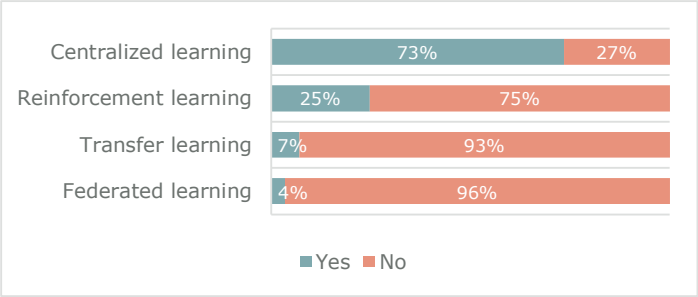

Figure 60: Types of ML learning

9.3 Open source libraries

Respondents who reported using ML were surveyed about their use of open-source libraries for development. The findings indicate that **over two-thirds (67%) of ML use cases rely on open-source libraries**.

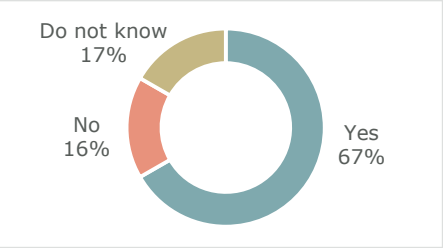

Figure 61: Usage of open source libraries for ML development

The *word cloud* below represents the most common open-source libraries mentioned by respondents.

Figure 62: Most cited open-source tools/libraries used for ML development

9.4 Third-party vendor solutions

Most (52%) of the reported ML use cases do not rely on third party vendor solutions for ML development (including data preparation), while 38% do

Notably, the reliance on third party vendor solutions is particularly high for IFM/AIFM, which are the only type of entity where the majority of the ML developments (57%) rely on third party vendor solutions.

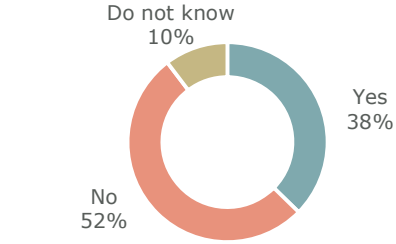

Figure 63: Use cases relying on third party vendor solutions for ML development

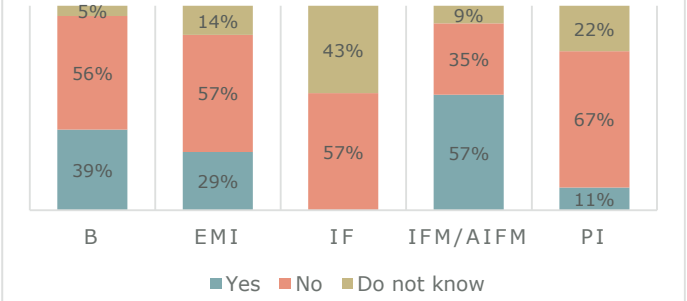

Figure 64: Use cases relying on third party vendor solutions for ML development, split by type of entity

The *word cloud* below represents the third-party vendor solutions most cited in the ML use cases.

Figure 65: "word cloud" on third party vendor solutions

10. Use cases: AI trustworthiness aspects

This chapter examines some key AI trustworthiness aspects across the use cases reported by survey respondents, starting from the risk classification according to the AI Act, and covering human oversight, explainability, auditability, bias prevention/detection and model performance monitoring.

10.1 AI Act

The **AI Act** is a European, horizontal regulation which aims to address risks to health, safety and fundamental rights, introducing requirements according to a risk-based approach:

- AI systems considered to be a clear threat to the fundamental rights of people constitute an unacceptable risk and therefore will be banned (such as, for example, AI systems used for cognitive behavioural manipulation or for categorising people, or emotion recognition systems used at the workplace).
- High risk AI systems (such as those listed in the Annex III of the regulation) will be subject to strict requirements for trustworthy AI (e.g. data quality, documentation and traceability, transparency, human oversight, accuracy, cybersecurity and robustness), that will need to be implemented by the provider and/or the deployer of the AI system.
- AI systems presenting limited risk will instead be subject to transparency obligations: e.g. AI systems like chatbots must clearly disclose to users that they are interacting with a machine, while certain AI-generated content must be labelled as such. General-Purpose AI systems ("GPAI"), including systems using GenAI, are among the systems subject to such transparency rules.

It should be noted that the survey was launched in June 2024, while the AI Act entered into force a few months later (on 1 August 2024). Although a stable version of the text was already available at the time of the survey, it is possible that some respondents were not yet familiar with this new regulation at the time of participation. The AI Act will enter into application on 2 August 2026, except for some specific provisions[56].

As part of the survey, respondents were asked to classify their use cases according to the AI Act risk levels:

- Unacceptable risk
- High risk
- Limited (transparency) risk
- Minimal risk or no risk
- Not yet classified/ do not know

5% of all use cases were classified as high risk, 16% with limited (transparency) risks, 50% with minimal or no risk, while the remaining 29% were not yet classified. The survey did not reveal any AI system which were classified as "unacceptable risk".

[56] Notably: the rules regarding prohibited AI practices, as well as the definitions and the provisions related to AI literacy, already apply since February 2025; the obligations for General-Purpose AI and the rules on governance will apply from August 2025; the obligations for high-risk AI systems that classify as high-risk because they are embedded in regulated products, listed in Annex II (list of Union harmonisation legislation) will apply from August 2027.

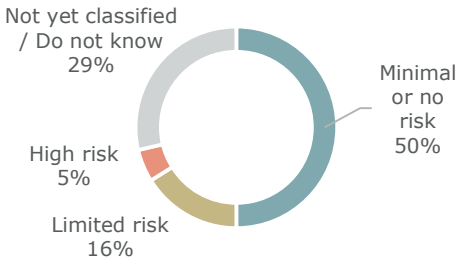

Figure 66: AI Act classification

When analysing the answers to this question, we note some degree of inconsistency with the classification of the use cases included in the AI Act, which is probably due to the novelty of the text and the lack of familiarity with the specificities of this regulation. Indeed, **the classification provided in the survey seems to reflect more the perception of the risk of the use case for the entity, rather than its real classification according to the AI Act**. Notably:

- If we consider the **credit scoring** use case, which is one of the few use cases listed as high risk in the AI Act[57], it was classified "high risk" in only half of the use cases falling in this category (three use cases), while for the other half it was classified either as "minimal or no risk" (two use cases) or it was not yet classified (one use case). While noting that the number of use cases in this category is fairly limited, a more thorough analysis would be required to assess the real classification of each use case according to the AI Act.
- On the contrary, according to the AI Act[58], AI systems used for the purpose of **detecting financial fraud** and AI systems used for prudential purposes to **calculate credit institutions' capital requirements** should *not* be considered as high risk. Nevertheless, when analysing the survey results, we note that two use cases falling into the category "AML/fraud detection", one use case falling into the category "IRB credit risk models" and two use cases falling into the category "other credit risk models" have been classified as high risk by respondents.
- Similarly, some use cases falling into other categories such as, cyber security, counter terrorism financing, process automation, chatbot and virtual assistant, search/summarise information, have been classified as high risk by survey respondents, although the AI Act does not explicitly list them as high risk.

[57] Notably, AI systems used to evaluate the creditworthiness of natural persons are considered high risk under the AI Act.

[58] Recital 58 and Annex III, art.5(b).

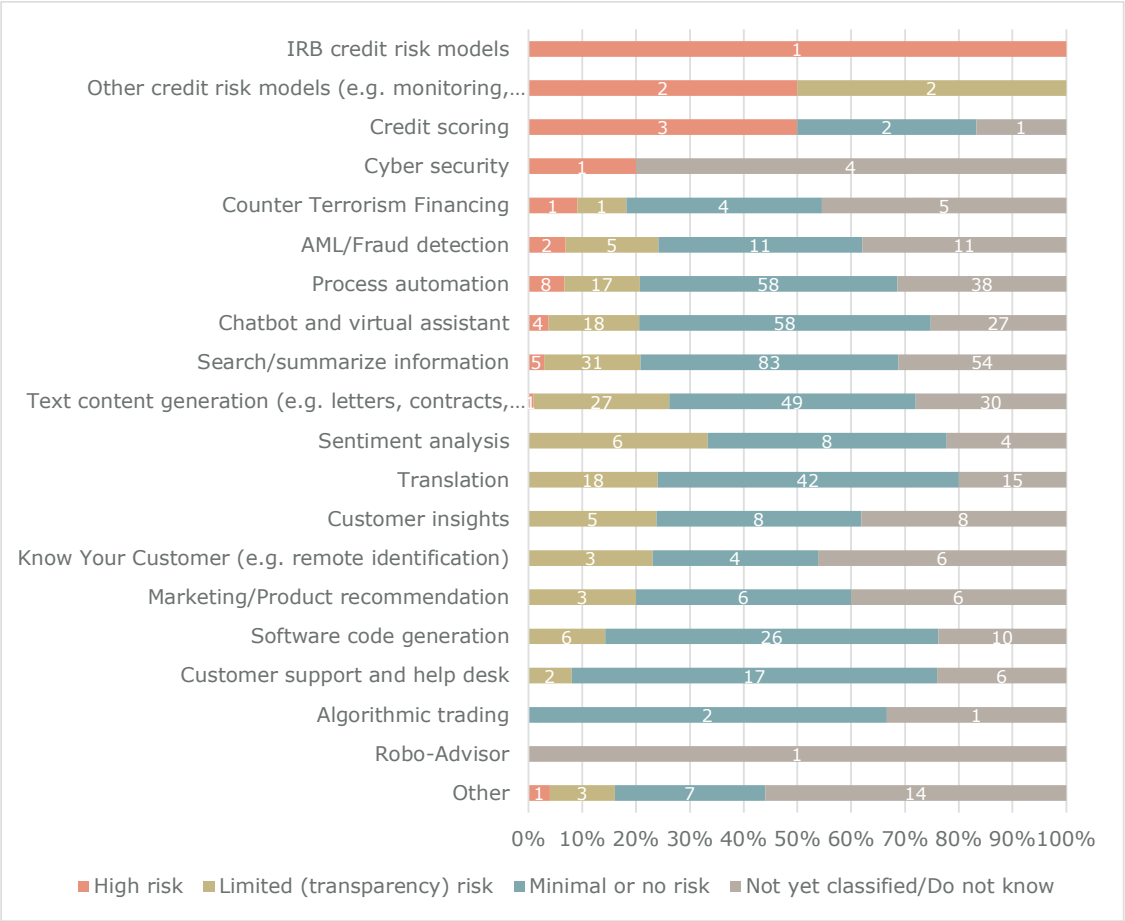

Figure 67: AI Act classification (by use case category)

As an additional remark when analysing the survey results, we note that in almost all use case categories there are some use cases that have been classified with "limited (transparency) risk"[59], and that those use cases are often linked to the use of GenAI. This reflects the "versatility" of GenAI in the sense that it can be integrated into (parts of) different types of use cases.

[59] In total, 16% of all use cases have been classified with "limited (transparency) risk".

10.2 Human in the loop

An AI/ML model may be integrated into a business process either in a fully automated way or with a 'human in the loop' involved in critical decisions. According to the survey, **90% of the use cases are configured with a human in the loop**. **This figure can be seen as a good indicator of trustworthiness considering the importance of humans in decisional processes** (depending on the criticality of the process within which the AI system is implemented).

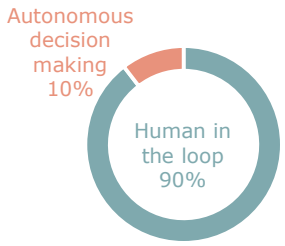

Figure 68: Use cases with human oversight

For some categories, such as Sentiment Analysis, Robo-Advisor and IRB credit risk models, no autonomous configuration was reported, with all use cases instead relying on a human-in-the-loop approach.

For credit scoring use cases[60], one third[61] of the solutions are currently running without a human in the loop while all of them were reported as being in production. While the survey does not provide detailed insights into the specific purpose of these use cases or other risk mitigation measures implemented, it is crucial to note that **high risk use cases will require thorough review to ensure compliance with human oversight requirements included in the AI Act**[62], which will start applying in August 2026.

[60] As explained in previous sections, AI systems used to evaluate the creditworthiness of natural persons are considered high risk under the AI Act.

[61] Corresponding to 2 out of 6 credit scoring use cases.

[62] See recital 73 and art. 14 of the AI Act.

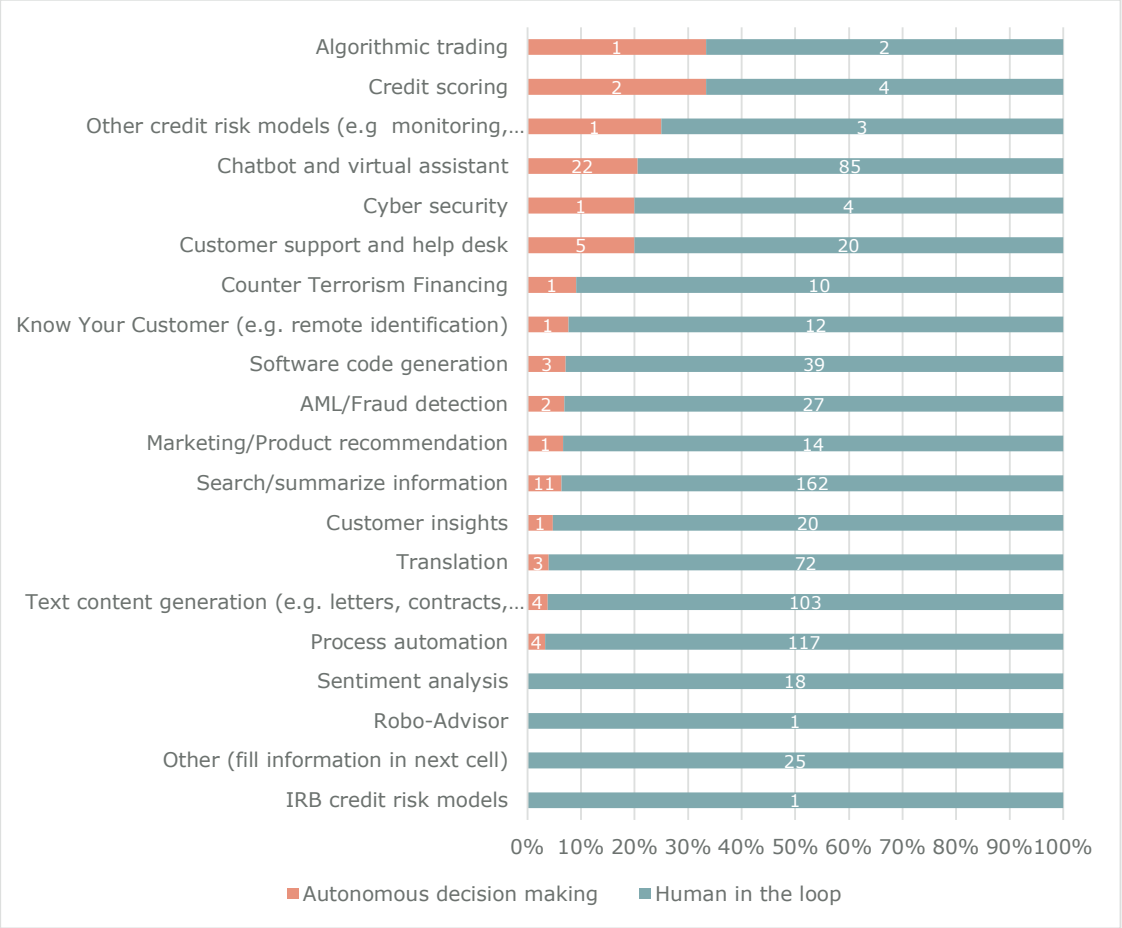

Figure 69: Autonomous systems Vs Human in the loop (by use case category)

When focusing only on B, EMI, PI, we note that **the percentage of use cases configured in autonomous mode** (with no human in the loop) does not change significantly and becomes 12%, which **is lower than the rate observed in the previous survey** (23%) for a similar scope of institutions. **This denotes** awareness about the risks of AI by respondents and **overall improved maturity compared to the previous survey**.

10.3 Bias

In the survey, respondents were asked to indicate, for each use case, whether bias detection/prevention measures were implemented. **We observe that for 40% of the use cases, respondents indicated that bias prevention/detection measures were not applicable** ("N/A"). These "N/A" responses were distributed across nearly all categories, with particular relevance for the category "algorithmic trading"[63]. Furthermore, we note that **the majority (72%) of these use cases (where bias prevention/detection was "N/A") involved GenAI: this may be in part explained by the fact that for GenAI related use cases, there might be a general expectation that bias treatments mechanisms are primarily implemented by the**

[63] Specifically, two-thirds (2 out of 3) of the use cases in the "algorithmic trading" category indicated that bias detection/prevention was "N/A," while the remaining third (1 out of 3) selected "do not know". For a complete view on bias prevention/detection measures across the different use case categories please refer to Annex 12.2.

provider of the GenAI model (especially for LLM (Large Language models)) rather than by the entity deploying it. However, it should be noted that depending on the use case, it could be required to implement additional bias prevention/detection measures also on the deployer side.

If we exclude the use cases for which the respondents indicated that bias prevention/detection was not applicable, of the remaining use cases **only for 45% of these, respondents confirm having implemented bias prevention and/or detection techniques** (figure 73).

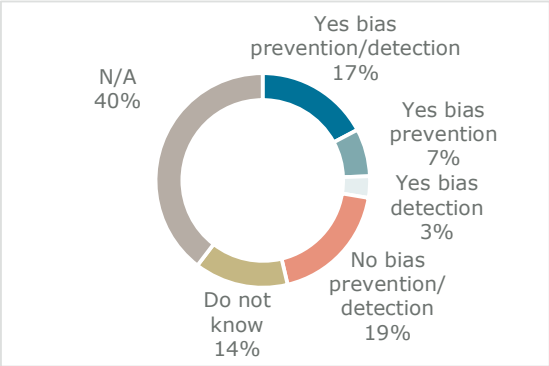

Figure 70: Bias prevention and detection

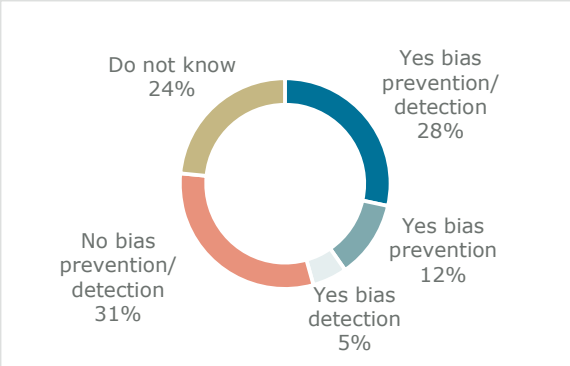

Figure 71: Bias prevention and detection (excluding N/A answers)

However, when we focus only on B, PI, EMI, we note that **the portion of use cases implementing bias prevention/detection techniques rises to 68%**[64]**. This marks an increase compared to the previous survey** (59%), indicating an overall improvement in maturity.

Among the various use case categories, bias prevention/detection measures assume particular importance for credit scoring, especially when the system is used to evaluate the creditworthiness of natural persons[65]. According to the survey results, the majority[66] of the use cases in this category do implement bias prevention and/or detection techniques, while the remaining use cases in this category correspond to responses "N/A" or "do not know".

[64] Figure calculated considering only B, PI, EMI, excluding "N/A" answers.

[65] AI systems used to evaluate the creditworthiness of natural persons are considered high risk under the AI Act.

[66] Corresponding to 3 out of 4 use cases, excluding "N/A" answers.

10.4 Auditability

Regarding the **auditability of the AI models, only 56% of the use cases report good (25%) or very good (31%) auditability. Approximately a third (30%) of the use cases received a medium auditability score**[67], while lower auditability ratings were attributed only to 14% of the use cases.

These auditability ratings follow a similar distribution across the different use case categories[68].

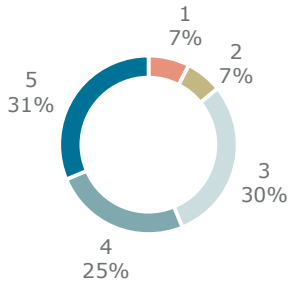

Figure 72: Auditability ratings (on a range from 1 to 5, with 5 being the highest rating, i.e. "very good auditability")

If we focus only on B, PI, EMI, we note that 55% of the corresponding use cases have an auditability score of 4 or 5, representing a **significant decrease compared to the previous survey** (where 81% of use cases were scored with auditability level 4 or 5). The causes of this difference cannot be easily explained based on the information from the survey, but they may well be associated with the increasing level of complexity of the AI solutions used and the difficulty in auditing them, together with more realistic scores provided by respondents based on more experience (including regarding AI systems audits).

10.5 Explainability

Explainability refers to the ability to justify and to provide a rationale for the predictions of an AI model. Results show that **the levels of auditability and explainability follow very similar distributions,** with respondents attributing similar ratings for both attributes (explainability and auditability) for the same use case.

Notably, **54% of use cases report good (23%) or very good explainability (31%)[69].** We note that **most** (71%) **of the use cases with lower explainability rating[70] involve GenAI, confirming that these models are often perceived as "black boxes"** due to their complexity.

[67] i.e. a rating 3 on a range from 1 to 5, with 5 being the highest rating, i.e. "very good auditability".

[68] For more details, please refer to Appendix 12.2.

[69] i.e. explainability levels 4 and 5 on a range from 1 to 5 (5 being the highest rating, i.e. "very good explainability").

[70] i.e. a rating of 1 or 2 or 3.

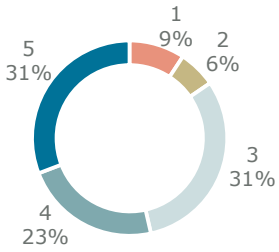

Figure 73: Explainability ratings (on a range from 1 to 5, with 5 being the highest rating, i.e. "very good explainability")

Similar to the auditability ratings, the above figures, which remain relatively stable when focusing only on entities of type B, PI, EMI, represent a **downgrade compared to the ratings from the previous survey** (where 70% of the use cases were scored with good or very good explainability).

10.6 AI monitoring

For the majority (56%) of the reported use cases, AI model performance is actively monitored. In contrast, for 16% of use cases the model performance is not actively monitored. Many of these use cases involve GenAI, highlighting that the complexity of such models presents challenges when it comes to performance monitoring.

Additionally, we observe that when models are updated, this is typically done on an ad-hoc basis.

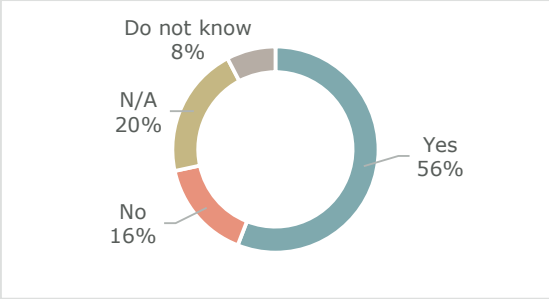

Figure 75: AI monitoring

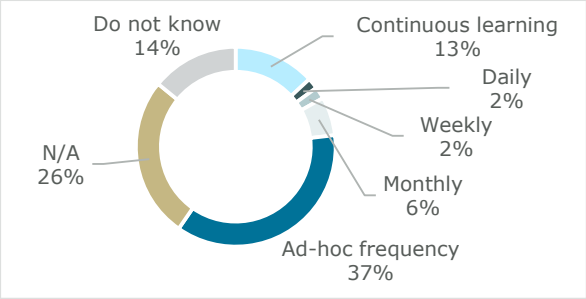

Figure 74: AI model update frequency

When focusing solely on machine learning use cases reported by B, PI and EMI, the percentage of AI solutions monitored over time increases to 88%[71]. This latter figure is largely consistent with the results of the previous survey[72]. This observation seems also to confirm that **the lack of performance monitoring is predominantly linked to the use of GenAI**.

[71] Excluding "N/A" and "Do not know" answers.

[72] In previous survey, 90% of the ML use cases had processes in place to monitor the algorithm performance over time.

11. Conclusion

The launch of commercially available GenAI solutions in November 2022 has sparked global adoption, and Luxembourg's financial institutions appear to have embraced this trend.

Indeed, the survey reveals that while the overall use of AI technologies among financial entities has risen compared to the previous survey, the percentage of financial entities with use cases involving GenAI is higher than those involving machine learning or other AI technologies. Furthermore, a significant part of institutions is still at an experimental stage, suggesting that we can expect a further surge in AI adoption (with more use cases getting into production) over the coming months.

The emergence of GenAI has brought forth new use case categories, such as text summarisation, content generation, chatbots, translation, software code generation. These are now among the top categories implemented by entities of Luxembourg's financial sector. Traditional categories like process automation remain prevalent but increasingly incorporate GenAI technologies alongside conventional methods such as machine learning. Machine learning continues to be used particularly in risk and compliance solutions, including AML/fraud detection and counter terrorism financing. However, other use cases, such as credit scoring (one of the few high-risk use cases listed in the AI Act) remain relatively limited. In this context, it appears that financial institutions have yet to fully comprehend or implement the risk categorisation introduced by the AI Act, warranting further work and education in this area.

Compared to the previous survey, some indicators – such as e.g. the existence of ethical policies and the implementation of bias detection/prevention techniques – suggest that financial institutions are increasingly focusing on trustworthiness aspects when adopting AI. Moreover, humans' decisions are not replaced, but rather "augmented" with AI - and GenAI in particular - as evidenced from the statistics related to "human in the loop". These developments indicate an improving level of maturity regarding AI integration in the financial sector.

In conclusion, the emergence of GenAI has accelerated the adoption of AI within supervised institutions. Currently, AI is predominantly utilised to support internal processes and enhance productivity, rather than being employed in customer-facing applications.

Recognising the pivotal role that trustworthy AI plays in fostering innovation and advancing the financial sector, both the Banque centrale du Luxembourg (BCL) and the Commission de Surveillance du Secteur Financier (CSSF) will continue to monitor the evolving use of AI by financial institutions.

12. Annex

12.1 Use case categories by type of entity

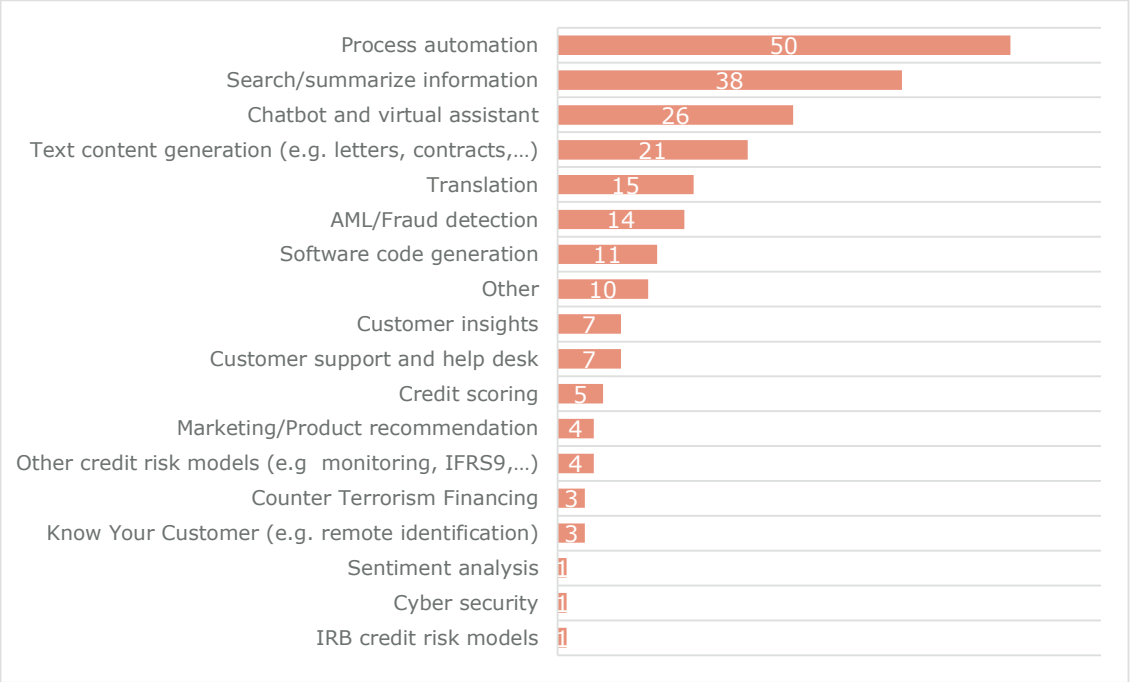

Figure 76: Use cases categories reported by B

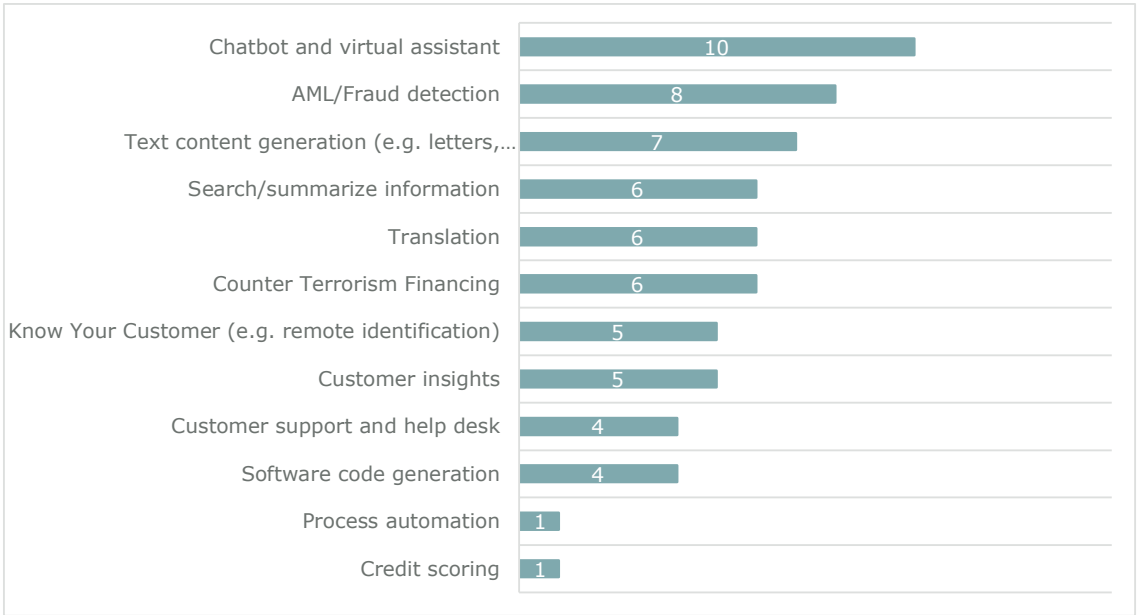

Figure 77: Use case categories reported by PI

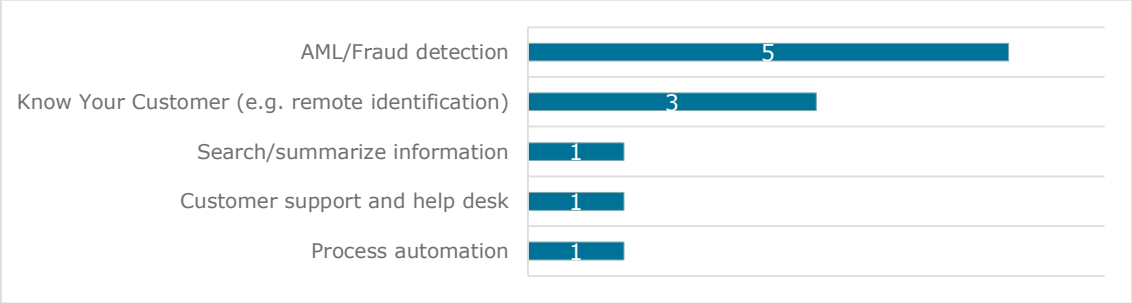

Figure 78: Use case categories reported by EMI

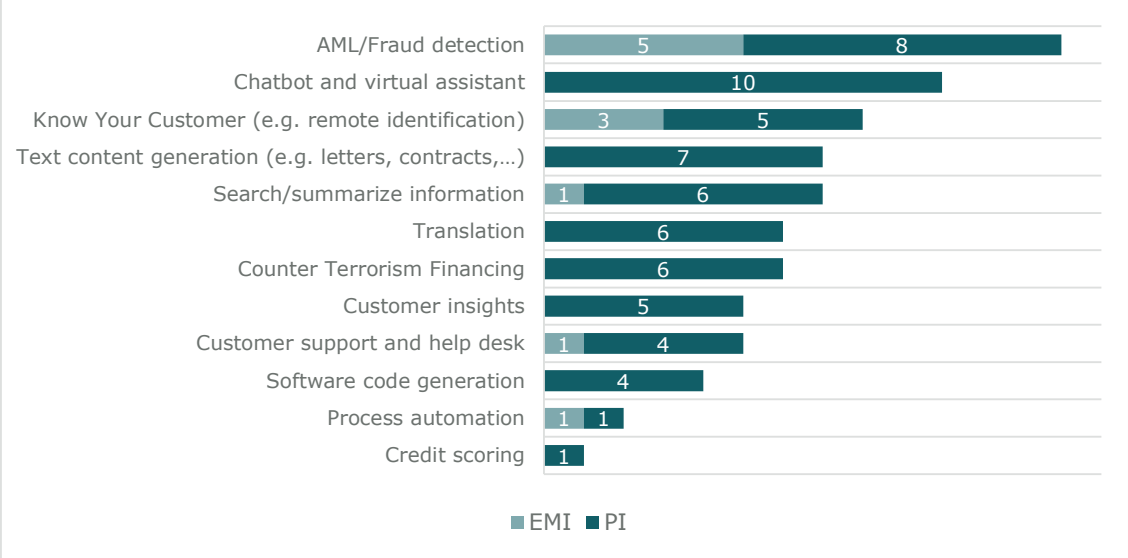

Figure 79: Use cases categories reported by EMI and PI combined

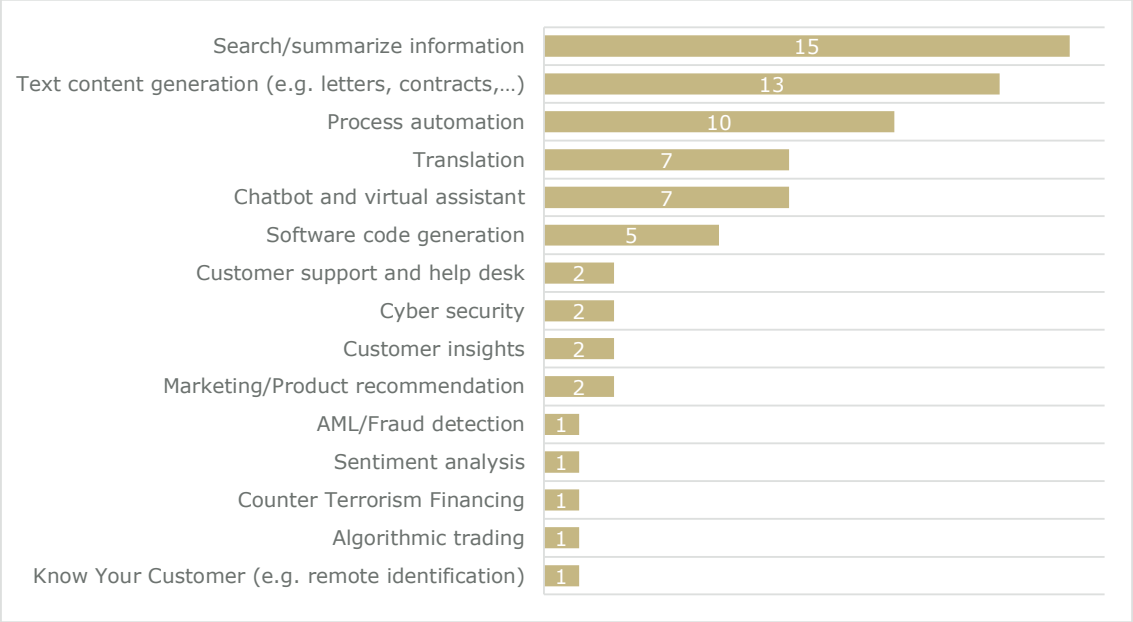

Figure 80: Use cases categories reported by IF.

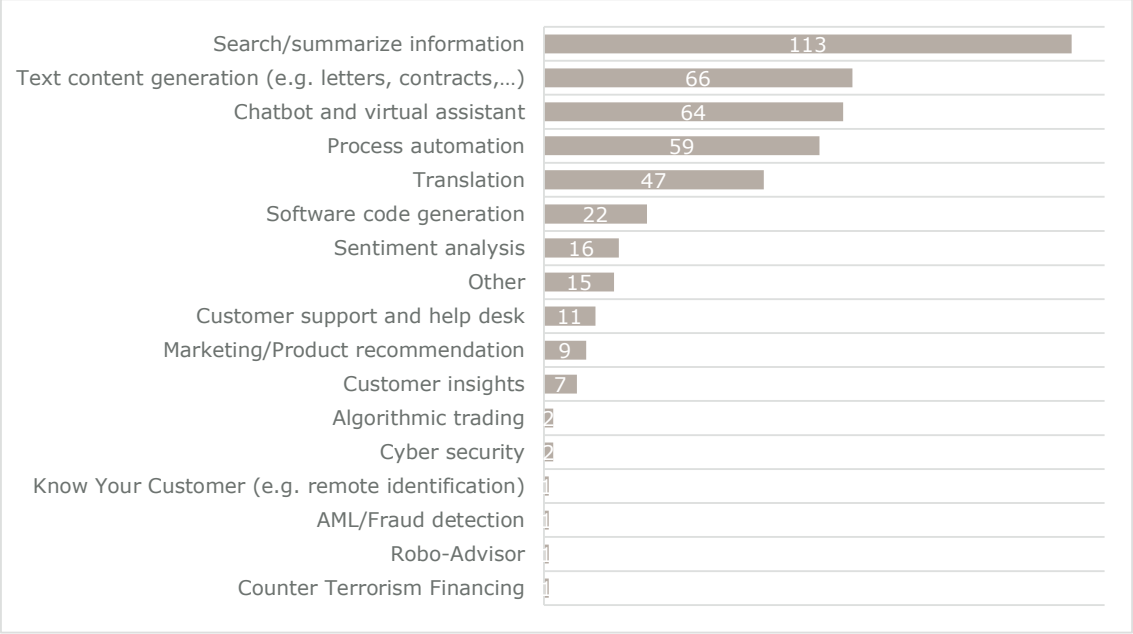

Figure 81: Use cases categories reported by IFM/AIFM.

12.2 **AI trustworthiness aspects by use case category**

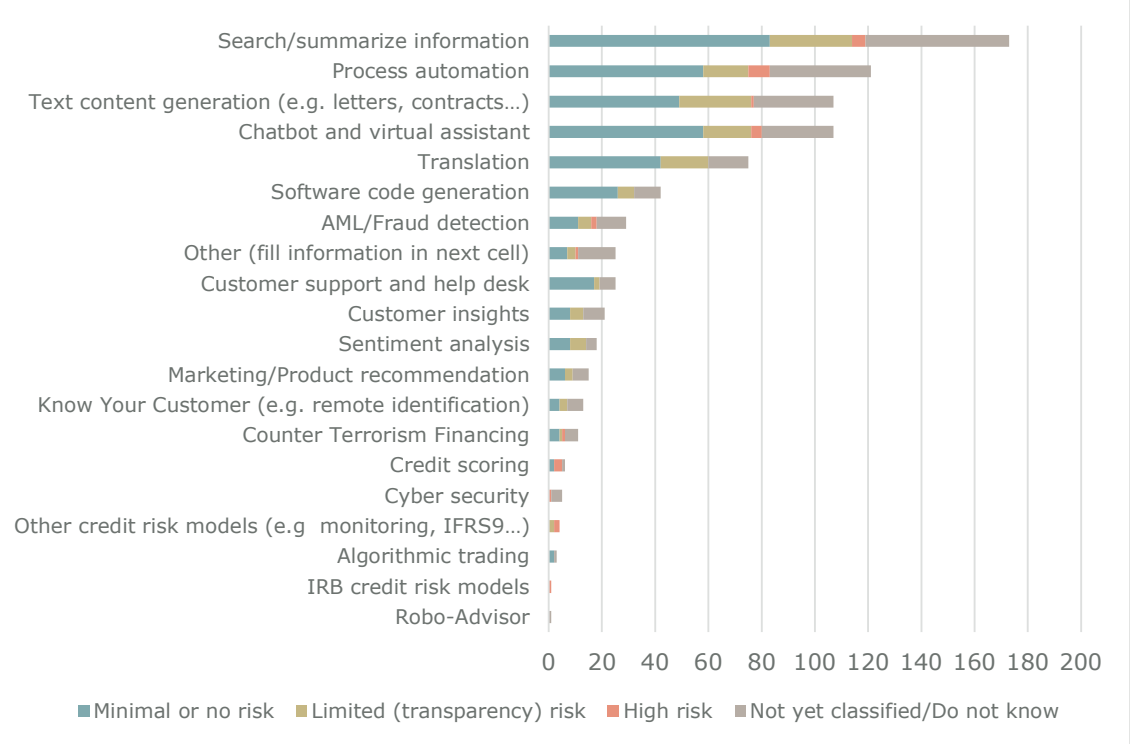

Figure 82: Risk classification under AI Act

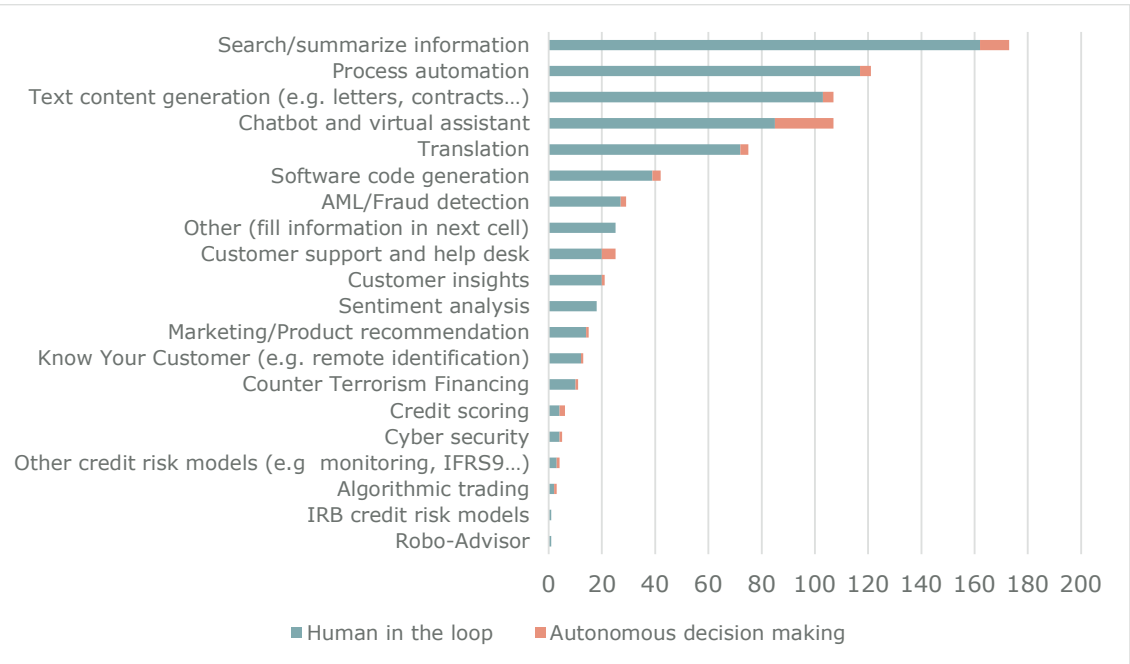

Figure 83: Human oversight

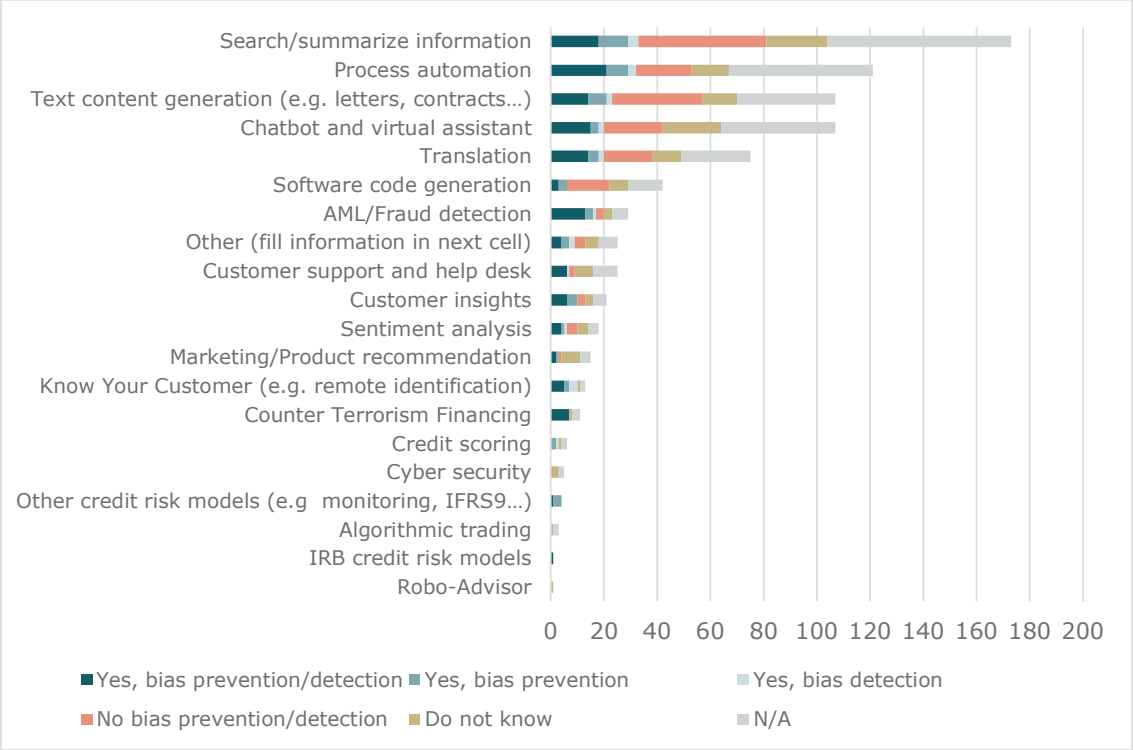

Figure 84: Bias prevention/detection

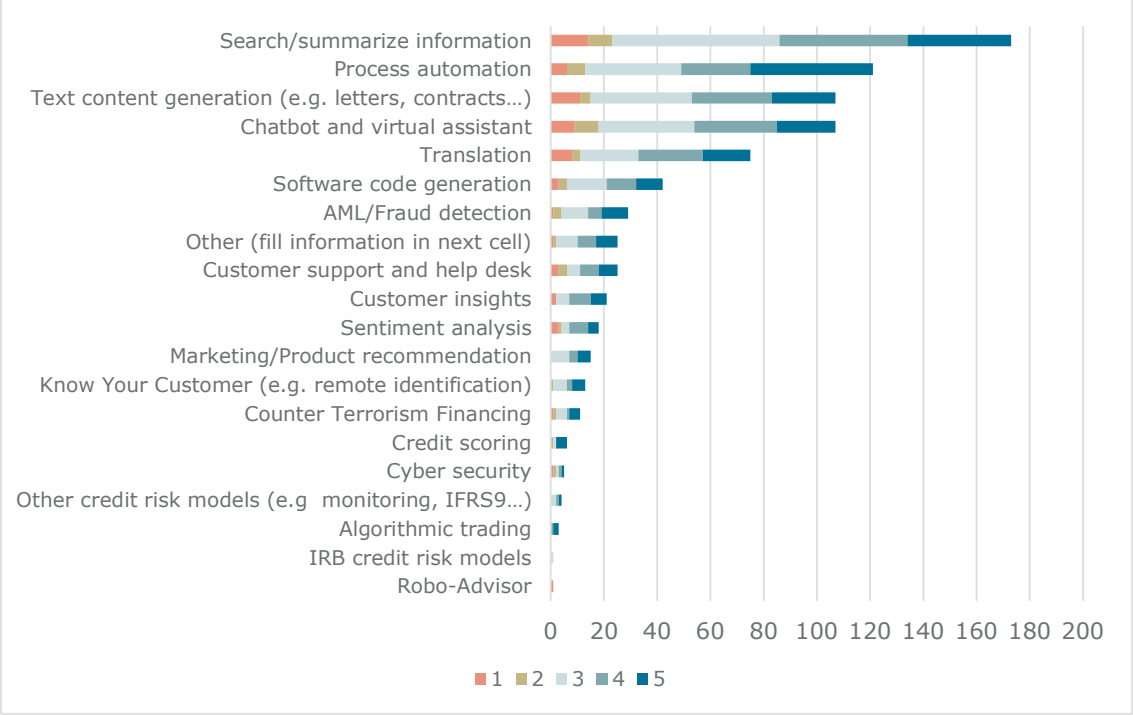

Figure 85: Auditability

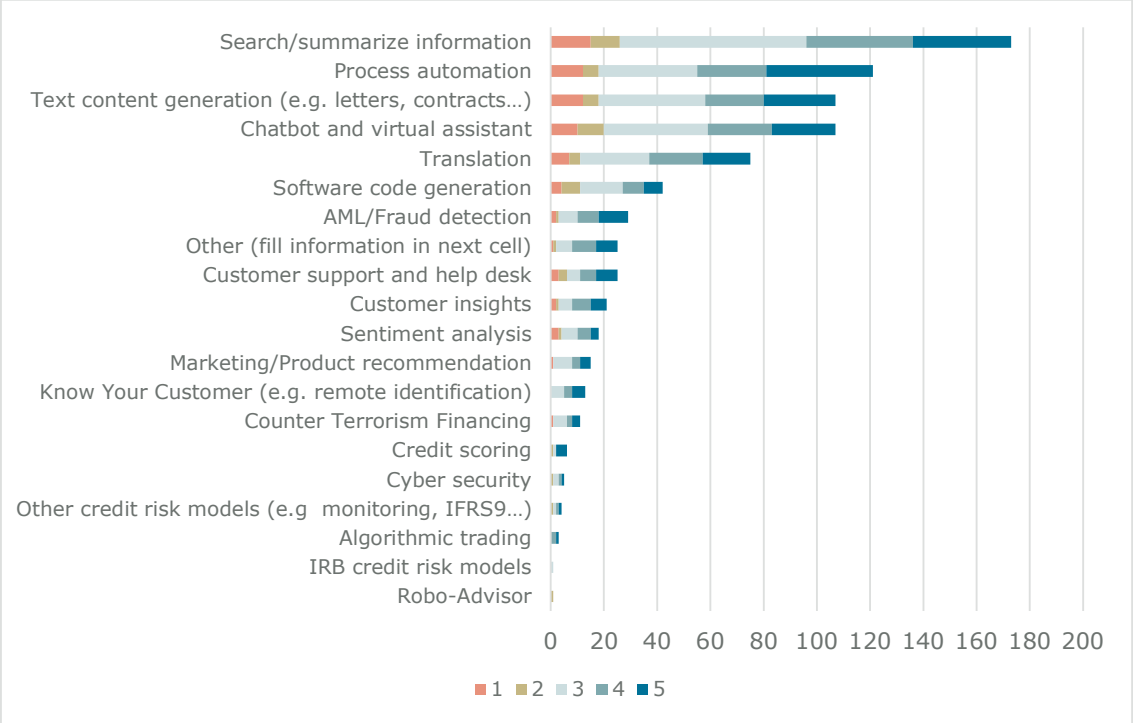

Figure 86: Explainability

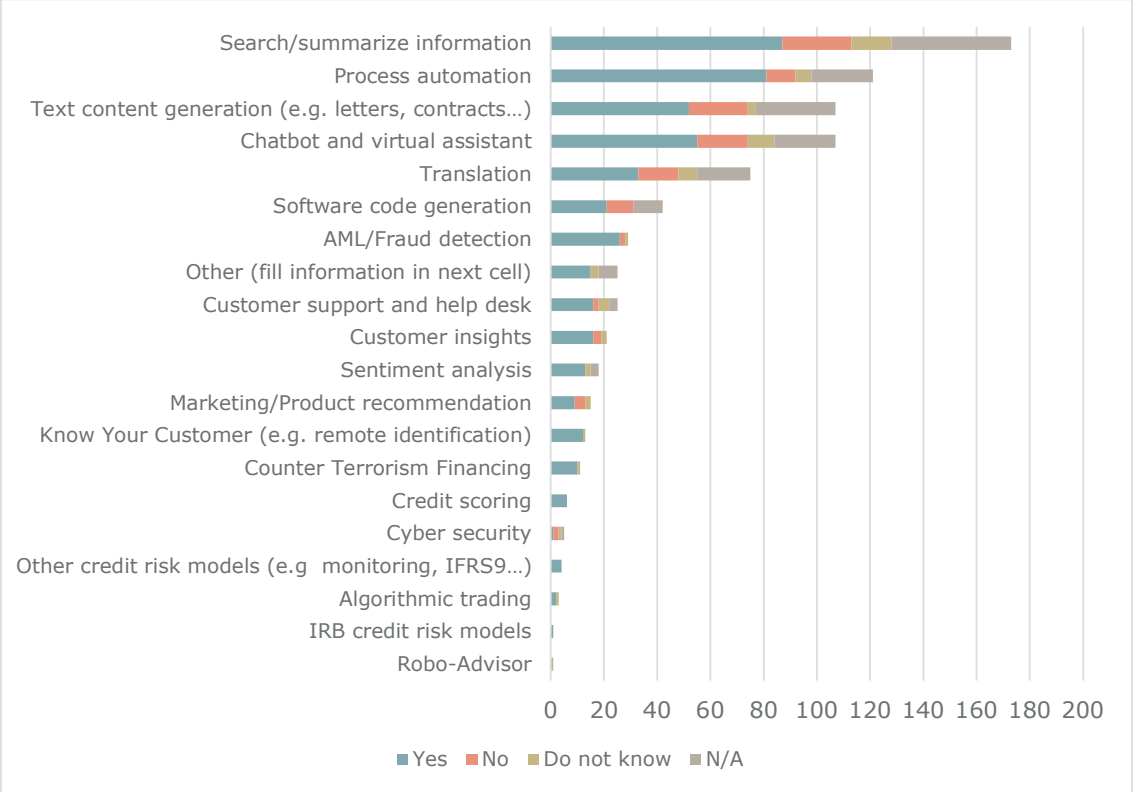

Figure 87: Monitoring of the performance of the AI solution over time

13. Glossary and Abbreviations

AI (Artificial Intelligence)	According to the European Commission's AI Act, "***'AI system'*** *means a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments*"[73]. In the context of this report, AI is meant in the broad sense to capture advanced analytical techniques, usually involving large data sets, which optimise and potentially learn solutions with limited or no human input. AI techniques include machine learning as well as other techniques such as, for example, expert systems, NLP, RPA (Robotic Process Automation), computer vision and chatbots.
AI Act	The AI Act, officially known as the "Artificial Intelligence Act," is the new EU regulation setting harmonised rules on artificial intelligence and aiming to ensure that AI systems are safe, respect fundamental rights, and are trustworthy. (Artificial intelligence (AI) act: Council gives final green light to the first worldwide rules on AI - Consilium (europa.eu)) The AI Act categorises different types of artificial intelligence according to risk. AI systems presenting only limited risk would be subject to very light transparency obligations, while high-risk AI systems would be subject to a set of requirements and obligations to gain access to the EU market. Finally, AI systems

[73] On 2 February 2025, the European Commission published the Guidelines on the AI system definition for the purpose of the AI Act (https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-ai-system-definition-facilitate-first-ai-acts-rules-application). Since these Guidelines were not available at the time the survey was conducted, they were not considered for the purpose of this report.

	whose risk is deemed unacceptable will be prohibited.
Algorithmic trading	AI/ML techniques can be used in algorithmic trading, e.g. for predicting trade price and cost, executing client orders with maximum speed at the best price.
AML/Fraud detection	AI/ML techniques may be used for fraud detection and anti-money laundering, for example by using historical data of past transactions and confirmed frauds to train a supervised ML algorithm to identify patterns of past frauds and use them to detect new ones more effectively. Unsupervised ML algorithms can also be used to identify outliers and previously undetected trends.
Anomaly detection	Anomaly detection is done by first detecting the structure of most of the data, for example by clustering, and then looking for the data points that do not follow any cluster, i.e. the "outliers". This technique is particularly useful when there is a need to identify unusual activity, like for example transactions linked to Terrorism Financing.
Association	Association is a particular type of clustering for which the common pattern is a rule (e.g. if customer purchased item_1, then he/she purchased also item_2). This technique is especially used in recommender systems to recommend to customers additional items that other customers already bought.
Auditability	Ability to track the main actions performed and gather evidence allowing investigations in case of incidents.
Bias	Bias refers to a systematic and unfair preference or prejudice for or against certain groups, ideas, or individuals, often resulting in discrimination and inequality in an AI model. This is generally induced by the training data being biased.
Centralised learning	Typical type of learning where the training data is centrally gathered in order to train models
Chatbots	Automated conversational agents capable of interacting with users of the platform.
Classification	A classification problem is a problem whereby the objective is to categorise a set of features

	with a given label (i.e. a given category). Classification identifies which category an item belongs to (for example whether a transaction is fraud or not fraud), based on labelled examples of known items (for example transactions known to be fraud or not). For classification problems the expected outcome is a discrete variable.
Computer vision and image recognition	Computer vision includes methods for acquiring, analysing and understanding images and videos in digital format. A classic example of computer vision task is the image recognition and classification.
Credit scoring	Use cases employing AI/ML techniques to improve the estimation of credit scores or credit risk of customers thereby facilitating/automating the approval process of lending, credit limits or other relevant decisions.
Crypto asset	According to MiCA[74], 'crypto-asset' means a digital representation of a value or of a right that is able to be transferred and stored electronically using distributed ledger technology or similar technology.
Customer insights	Use case consisting in analysing consumer patterns (e.g. spending behaviour) to predict future trends and provide insights (e.g. prediction of available budget at the end of the month based on spending patterns).
Deep learning	Artificial Neural Networks (ANNs) a.k.a. Deep learning is a branch of AI that is sometimes considered a subset of ML or a separate branch in its own. Deep neural networks are capable of learning unsupervised from data that is unstructured or unlabelled. Also known as Deep Neural Learning or Deep Neural Network. Neural networks are a particular type of ML algorithms that generate models inspired by the structure of the brains, and in particular the neuronal activity. The model is composed of several layers, each layer being composed of units (the neurons).
Dimensionality reduction	Dimensionality reduction is an unsupervised method that enables reducing the number of

[74] Regulation (EU) 2023/1114 on markets in crypto-assets.

	random variables under consideration by obtaining a set of principal variables. There are two main methods to achieve dimensionality reduction, namely feature selection (i.e. removing features along the training for instance) or feature projection (i.e. by reducing the dimensionality of the data features by applying linear or non-linear transformations).
DLT (Distributed Ledger Technology)	DLT is a decentralised database, across multiple nodes. Blockchain is an example of DLT where transactions are recorded with an immutable cryptographic signature called a hash. The transactions are grouped in blocks and each new block includes a hash of the previous one, chaining them together, hence why distributed ledgers are often called blockchains.
Expert systems	Expert systems, also called rule-based systems, are systems that store and manipulate knowledge in the form of rules and derive new knowledge (new rules) by applying an inference engine to the existing knowledge base. The term "rule-based system" is normally used to identify systems where the set of rules are pre-defined by humans, as opposed to machine learning systems where the "rules" are automatically learnt by the system.
Explainability	An AI system is explainable when its internal behaviour can be directly understood by humans (interpretability) or when explanations (justifications) can be provided for the main factors that led to its output.
Fairness	Fairness is the concept of ensuring equal and impartial treatment of individuals or groups in the AI processes. This require training data that is free from bias so that AI models do not perpetuate existing inequalities.
Federated learning	Federated learning (also known as collaborative learning) is a machine learning technique that trains an algorithm across multiple decentralised edge devices or servers holding local data samples, without exchanging them.
Generative AI (GenAI)	Generative Artificial Intelligence or GenAI refers to a sub-category of artificial intelligence models designed to create new content, such

	as text, code, images, audio, or video, by learning patterns from existing training data. These models leverage advanced machine learning techniques, particularly deep learning models like Generative Adversarial Networks (GANs) and Transformer-based models (e.g., GPT, BERT), to generate outputs that resemble human-created content.
IPA (Intelligent Process automation)	RPA integrating AI and ML functionalities, such as NLP and text mining. For example, the NLP/text mining engine can analyse a scanned document and automatically classify it according to its category (e.g., ID document, invoice, payment receipt, ….), making it possible to automatise entire parts of middle and back-office processes.
IRB credit risk modelling	Use case applied to the generation of an internal (challenger) model for the purpose of calculating regulatory capital according to the internal ratings-based (IRB) approach to capital requirements for credit risk.
Large Language Models (LLMs)	A large language model (LLM) is a type of GenAI model specifically designed to understand and generate human-like text based on vast amounts of data.
ML (Machine Learning)	Machine learning algorithms build a model based on sample data, known as "training data", in order to make predictions or decisions without being explicitly programmed to do so. There are different categories of ML techniques such as supervised learning, unsupervised learning, reinforcement learning and deep learning.
ML security - Data poisoning	In poisoning attacks, attackers deliberately influence the training data to manipulate the results of a predictive model.
ML security - Adversarial attack	An adversarial attack consists in providing a sample of input data which has been slightly perturbed in order to cause the model to misclassify it.
ML security - Model stealing	This attack consists in replicating/cloning a model by probing the targeted model with high number of inference requests and use response received to train another model.

NLP (Natural Language Processing)	Natural Language Processing is the branch of AI enabling computers to analyse, understand and generate human language, in both written and spoken form.
Process automation	Use case employing RPA/IPA techniques to automatise processes previously requiring several human interventions (with low added value).
Retrieval Augmented Generation (RAG)	Also referred as grounding, Retrieval Augmented Generation is a technique that combines information retrieval with generative models. It involves using a retrieval system to find relevant documents or passages and then augmenting the model's knowledge with this additional information, enabling it to generate more accurate and contextually appropriate responses.
Regression	Regression problems are similar to classification in that they both use labelled past data to predict the value of new data, with the exception that regression methods will predict a variable that is a real number, meaning that it can have continuous possible values (as opposed to only a discrete set of values such as in the classification methods).
Reinforcement learning	Reinforcement learning is a method whereby the objective is to train a model to maximise rewards by feeding it with feedback on its actions (i.e. either positive and/or negative reinforcement).
Robo-advisors	Automated software applications providing advice to clients, especially regarding proposed investments.
RPA (Robotic Process Automation)	Systems allowing to automate highly repetitive tasks which normally represent low value-added tasks for humans.
Sentiment Analysis	Techniques aiming at identifying and categorising sentiments or opinions expressed in written texts or by speech, in order to determine the attitude of the person toward a particular topic (e.g. positive, neutral, or negative). For example, such techniques can be used to build a cognitive profile of clients to propose more tailored investments. These techniques often use social media data.

Supervised learning	Supervised learning refers to the ability of an algorithm to infer a function from a training data set that contains labels.
Transfer learning	A type of learning reducing the time involved in training the model by using the learning of an already developed scenario and applying that learning to a different but related problem.
Tokenisation (of assets)	Asset tokenisation is the process of creating "tokens" to represent them on the blockchain. These tokens (blockchain representations of the tokenized assets) may be qualified, when all the requirements are fulfilled, as financial instruments in the sense of MiFID II and would therefore be out of scope of MiCA.
Unsupervised learning	Unsupervised learning refers to the ability of an algorithm to infer a function from a training data set that does not have any label. A typical example of unsupervised learning is to identify categories of client profiles based on their spending behaviour (i.e. clustering).