BANQUE CENTRALE DU LUXEMBOURG
EUROSYSTEM

Commission de Surveillance du Secteur Financier

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

MAY 2023

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

CONTENTS

Executive summary

During the period October 2021-January 2022, the BCL and the CSSF carried out a joint survey to understand the level of adoption of certain innovative technologies and in particular the usage of Artificial Intelligence ("**AI**") and Machine Learning ("**ML**") in the Luxembourg financial sector. The survey was addressed to all credit institutions, payment institutions[1] and e-money institutions supervised by the CSSF. This document summarises the findings of that survey in the form of a thematic report.

The survey was sent to 148 supervised institutions, among which 138 participated in the survey (93%).

The survey included 3 main sections. The first section aimed at gathering high level information about the **digital strategy** of the institutions, in terms of investments (current and future) in innovative technologies such as AI, APIs[2], digital onboarding techniques, DLT[3]. The second section consisted of a detailed **questionnaire regarding the adoption of AI and ML technologies**, while the third section focused on the specific **use cases of application of AI (including ML) technology**.

Hereafter is reported a summary of the main findings from the survey.

Regarding the **digital strategy** of the institutions and related investments performed in **2021**, **the level of adoption of AI and other innovative technologies was fairly limited and still at early stage.** The type of innovative technology which received more adoption (in terms of number of entities investing in it) is **APIs (56%)**, followed by **digital onboarding (34%)** and then **AI (32%),** while only **14%** of the respondents invested into **DLT** (the majority of which into crypto assets related technology)**.**

In relation to **2022-2023 investments**, the responses indicate a **general increase of investments across all the categories of innovative technologies compared to 2021 budget**, with the **highest increase[4] for ML technology.**

In relation to the second part of the survey, the **artificial intelligence questionnaire**[5] aimed at understanding the level of adoption of AI/ML technologies and covered several topics ranging from benefits and challenges related to the use of AI, data science team organisation, data governance, security and robustness, machine learning development lifecycle and technical infrastructure.

[1] *Payment institutions and electronic money institutions are governed by the Law of 10 November 2009 on payment services ("PSL"). The PSL was amended by the Law of 20 July 2018 which transposed Directive (EU) 2015/2366 of 25 November 2015 on payment services in the internal market ("PSD2").*

[2] *Application Programming Interfaces*

[3] *Distributed Ledger Technology*

[4] *in terms of n. of respondents investing in this technology*

[5] *Note that the number of respondents to the sections 2 and 3 of the questionnaire, i.e. AI questionnaire and Detailed use cases, represent a sub sample of the respondents, i.e. only those that reported using AI (representing 30% of the total panel)*

Of all respondents, 30% use AI technology, **while 25% use ML.** This means that 83% of entities using AI use ML[6], thus confirming **ML as the most used AI technology** (among all types of AI technologies).

The top three **AI/ML benefits** reported were **"improving internal efficiency", "risk reduction" and "enhance product and services offered"**. In terms of **challenges,** the top three were **"data quality", "AI/ML skills",** and **"model drifting and monitoring".**

In terms of **location of the data science team**, the latter mostly sits in the **IT department (40%),** but closely followed by **dedicated unit category (37%)** and finally **within business lines (23%).** The majority **(86%)** of respondents using AI[7] have a dedicated team working only on AI related projects/development activities ("**data science team**"), in most cases (**72%**) situated **at group level or a combination of group level and local level.** Most of the **data science teams in Luxembourg are small (up to 10 persons)**, while larger teams usually sit at group level.

In terms of **key skills** required in a data science team, the perception from respondents is that **data analysis**, **statistics** and **IT programming** are the key skills that are needed.

Regarding training, **70% of respondents offered AI/ML trainings to their employees, of which 30% specific advanced training to their AI/ML developers/data scientists (including upskilling)**, 25% general AI/ML awareness training and 15% other types of training. On the other hand, a relevant portion of respondents (30%) still had not at the time provided any kind of AI related training to their employees.

In relation to data governance, the **majority of respondents do not have specific AI related governance mechanisms** such as an AI ethical policy (77%) or an ethic committee (92%). However, from a traditional risk governance perspective, we found more reassuring figures indicating the **involvement, in the AI/ML development process,** of the **Data Protection office (83%)**, of the **Information Security function (88%)** and to a lesser extent of the **Risk function (63%),** denoting some consistency with the traditional IT development process**. Moreover, 85% of respondents have processes, policies and procedures for data governance and data quality, however 56% admit that these would require improvement for AI specific treatments.**

With regards to **ML specific security attacks** such as adversarial attacks, data poisoning and model stealing[8], we see a very split approach with nearly half of respondents affirming to having taken specific measures for these types of attacks, while at the same time **56%** of respondents affirm to perform **independent security reviews and penetration tests** of their AI solutions**.**

In relation to the technical infrastructure supporting the ML processes, we note that **development environments** are hosted primarily **on premises (54%),** while cloud and hybrid environments represent respectively 14% and 32%.

[6] *It should be noted that ML is a subset of AI (see definition provided in the Annex)*

[7] *It should be noted that the figures mentioned in this paragraph and in the rest of the executive summary are indicated as percentages of those respondents who indicated that use AI and/or ML*

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

The third part of the survey focused on the specific **use cases applying AI/ML technologies**.

In total, **158 different use cases** using AI technology were reported by survey respondents, of which **59% in production** (confirming that AI technology is still in the early stages of adoption). The top five areas of use cases reported were **AML/Fraud detection (18%), Process automation (15%), Marketing/Product recommendation (8%), Customer insights (8%) and Cyber security (8%)**.

With regards, in particular, to AI trustworthiness aspects, **77%** of all use cases reported are configured with a **"human in the loop"**, **51%** implement **bias prevention/detection techniques**, **81%** have underlying AI/ML models with **good auditability**, and **70%** have **good explainability**.

To conclude, the survey demonstrated that the usage of AI in the Luxembourg financial sector is currently fairly limited and still at an early stage, but investments in this technology and especially ML are estimated to increase, paving the way for a wider adoption of these innovative technologies in the near future.

Introduction and objectives

Artificial Intelligence (hereafter "AI") is an innovative technology that can positively affect the financial sector by enabling, for example, improved processes, enhanced fraud detection mechanisms, new customer insights, foster inclusion, etc. Nevertheless, AI also brings new challenges and risks to be considered for the regulator as well for the entities.

In order to gather information about the usage of AI (and ML in particular) in the Luxembourg financial sector and the particular use cases being implemented at supervised institutions, CSSF and BCL launched a joint survey in October 2021. The aim of this joint initiative was primarily to assess the level of adoption of these technologies by supervised institutions and to analyse the implementation of AI ("use cases") with their related challenges, including AI trustworthiness aspects (e.g. explainability, ethics, bias and fairness, auditability, etc.…).

To capture a more comprehensive picture, the survey included a section regarding the overall digital strategy of the supervised institutions, information about the level of investments in AI and ML technologies in comparison to other innovative technologies such as DLT[9], API[10]s, e-KYC/digital onboarding, as well as the expected benefits in terms of cost savings or increased revenues.

This report summarises the findings from the survey and is the result of an analysis work performed jointly by CSSF and BCL. The results presented hereafter are based on aggregated data, without any reference to specific institutions participating in the survey.

[9] *Distributed Ledger Technology*

[10] *Application Programming Interface*

Scope and methodology

The survey was sent out during the period October 2021 – January 2022 to all credit institutions, e-money institutions and payment institutions. It consisted of an excel questionnaire composed of three main sections:

- **Digital strategy** covering investments (and related expected benefits in terms of reduced costs or increased revenues) in innovative technologies such as Artificial Intelligence and Machine Learning, DLT and crypto assets, APIs, digital (remote) onboarding.
- **AI questionnaire** covering various aspects regarding the use of AI and ML such as benefits and challenges, organisational aspects, data and governance, security and robustness, ML development lifecycle, ML technical infrastructure.
- **AI use cases** focusing on the practical use cases where AI technology is applied, including several questions covering general development aspects, trustworthiness and security, the type of AI technology and the ML problem.

The responses from the survey questionnaires were aggregated, anonymised, and analysed to produce this thematic report. The report has been organised as follows:

- Part 1 presents some general survey demographic information
- Part 2 focuses on the "Digital Strategy" section of the survey
- Part 3 presents the findings from the "AI questionnaire" section of the survey
- Part 4 presents general findings from the "AI use cases" section of the questionnaire
- Part 5 focuses specifically on the trustworthiness and security aspects of the use cases included in the "AI use cases" section of the questionnaire.

It should be noted that **the statistics included in part 1 and 2 of the report** (presenting general survey demographics and the results from the Digital Strategy section of the questionnaire) **are reported as percentages based on the *total* number of respondents to the survey, while the statistics contained in part 3, 4 and 5 of the report** (presenting the results from the sections "AI questionnaire" and "AI use cases" of the survey) **are presented as *relative* percentages based only on the portion of respondents who indicated making use of AI technology** (representing a subset of the total number of respondents), **except for the statistics in sections "Machine learning lifecycle" and "Machine learning technical infrastructure"** (in part 3 of this document)**, which are calculated based only on those respondents who indicated using ML specifically**.

Part 1 – Survey demographics

Type of entities

The joint survey was sent to all Luxembourg credit institutions, e-money institutions, and payment institutions[11]. In total, **148 institutions** were **targeted by the survey**, of which 125 credit institutions (84%), 13 payment institutions (9%) and 10 e-money institutions (7%).

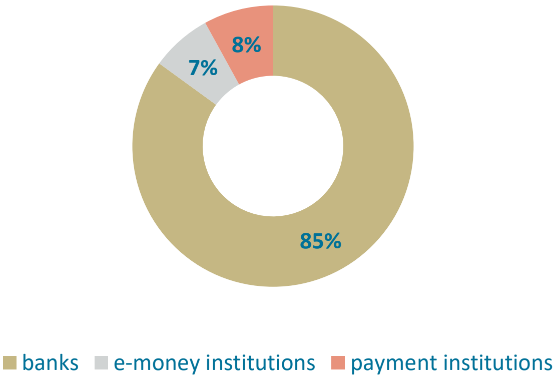

Figure 1: survey participants

The survey had a very good response rate, with a total of **138 respondents**, representing a participation rate of 93%.

Figure 1 shows that the distribution of entities among the respondents is very similar to the distribution of all targeted entities, with banks representing 85% of the total number of participants, followed by payment institutions (8%) and e-money institutions (7%)[12].

Profile of survey contact persons

Although the survey instructions recommended the contact person for the survey to be somebody in the company with a responsibility for Digital Strategy and AI/ML topics (e.g. Chief Digital Officer, Head of Innovation, Head of Data Analytics, etc.…), this was only rarely the case. As it can be seen from the chart at Figure 2, the profiles of the contact persons for the survey were from various business areas, the largest part being from the **IT function** (**38%**), followed by members of senior management (21%) and of the compliance function (14%).

[11] *The survey targeted all credit institutions, e-money institutions and payment institutions supervised by the CSSF active as of 1/10/2021*

[12] *The total number of respondents is composed of 117 credit institutions, 11 payment institutions and 10 e-money institutions.*

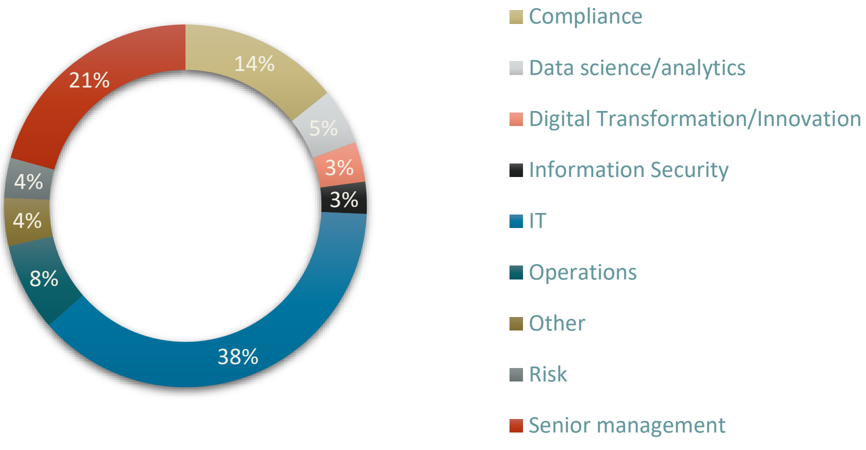

Figure 2: profile of contact persons

The above figures support the general view that digital strategy topics, including AI/ML, are often a responsibility of the IT department. More in-depth analysis regarding the organisational aspects (e.g. team size and positioning) for AI/ML topics is provided in Part 3 of the report (section “Organisation”) below.

Part 2 – Digital strategy

The objective of the Digital Strategy part of the questionnaire was to identify current and future investment trends (and related expected benefits in terms of reduced costs or increased revenues) across the following set of predefined categories of innovative technologies:

- AI - Machine Learning
- AI – Other
- API – PSD2
- API – Other
- Digital onboarding – automated identity verification based on facial image captures ("selfie")
- Digital onboarding – identity verification via "video-chat" (with human operators)
- Digital onboarding – Other
- DLT – Crypto assets
- DLT – Other
- *Other*

2021 investments

In relation to **2021 investments**, **64% of respondents invested into at least one category of innovative technology listed above, while 36% of respondents did not invest into any of those technologies.**

Figure 5 below shows that the proportion of entities, among all respondents, investing in each type of innovative technology is not very high. These figures indicate that **in 2021, across the whole financial sector composed of banks, payment institutions and e-money institutions, the level of adoption of AI and other innovative technologies was fairly limited and still at an early stage**.

It should be noted that one of the reasons for the low adoption rate is that several entities[13] are outsourcing part or all their IT infrastructure to the group and/or that investments in innovative technologies and overall digital strategies are carried out mainly at head office level. The deployment in Luxembourg could follow in a second phase capitalising on group expertise.

[13] 14% of the total n. of respondents answered that they were not using AI and indicated that their digital strategies were carried out at group level, mainly due to the outsourcing of their IT infrastructure to the group. Nevertheless, this information was not mandatory and therefore the real figures might be higher.

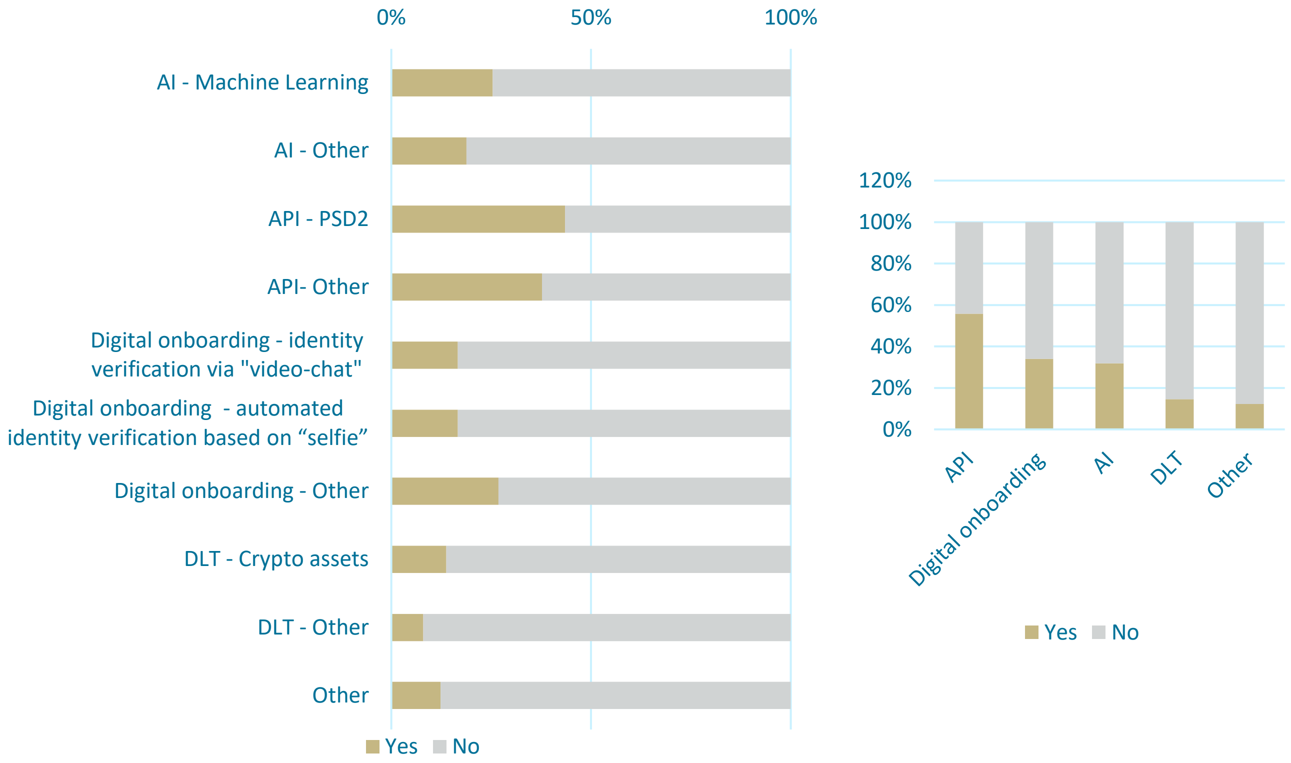

Figure 3: 2021 investments in innovative technologies (detailed view and summarised view)

The type of innovative technology which received investment from the highest number of entities is **Application Programming Interfaces (APIs)**, with **56% of the respondents having invested in APIs, mainly PSD2**[14] **related APIs**.

The second type of innovative technology is Digital Onboarding (i.e. identity verification tools using facial recognition, video-chats or other technologies for remote onboarding), with **34%** of respondents having invested in this type of technology.

AI (including ML) technology is third, with 32% of respondents having invested in this technology.

Finally, **14% of the respondents have invested into crypto assets related technology, and only 8% of respondents have invested into other DLT (excluding crypto assets).** These last figures show a rather cautious approach towards DLT compared to the other technologies mentioned above.

[14] *Payment Service Directive 2*

2022-2023 investments

Compared to investment budgets in 2021, **the responses indicate a general increase of planned investments for the years 2022-2023 across all the categories of innovative technologies**.

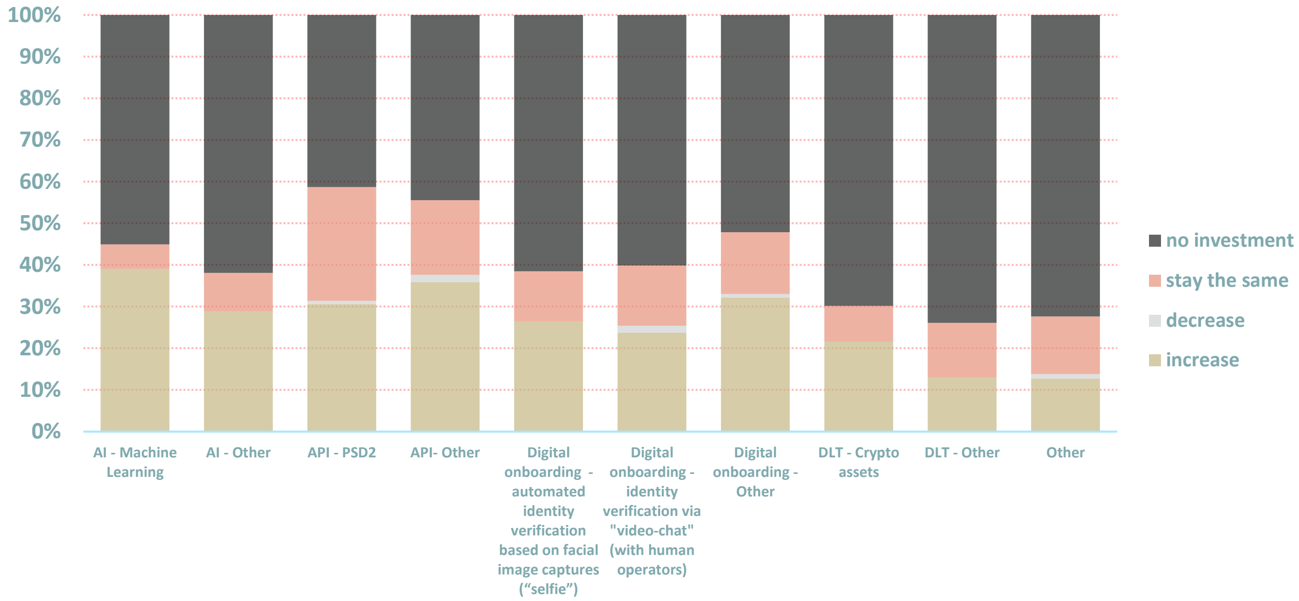

Figure 4: 2022-2023 investments in innovative technologies (compared to 2021 budget)

ML represents the category with the highest number of respondents (39%) affirming that they will increase investments, with an additional 6% of respondents stating that investments in ML will stay the same. Also, **a relevant portion of respondents (29%) indicated that they will increase investments in other AI technologies**, while 9% will keep the same investment level.

"API-Other" is the second-highest category with 36% of respondents confirming that they will increase investments in this area. **API-PDS2 will remain a key investment sector**, with 30% of respondents indicating that they will increase investments and an additional 27% indicating that they will keep the same level of investments.

In the area of DLT-crypto assets, a still relevant portion of respondents (20%) report wanting to increase investments and an additional 10% indicate that they will keep the same level of investments. The results for DLT technologies paint a tentatively positive (although still timid) picture, with increased investment probably also due to a clearer upcoming regulatory environment, including regulations such as MiCA[15] and DLT Pilot regime.

Anticipated cost savings due to adoption of innovative technologies

43% of entities reported that they anticipate cost savings due to the adoption of innovative technologies.

Among these entities, the anticipated costs savings per category of innovative technology vary between 0 and 40%, with **the majority of respondents (88%) indicating a cost saving between 0 and 10%** (see figure 5).

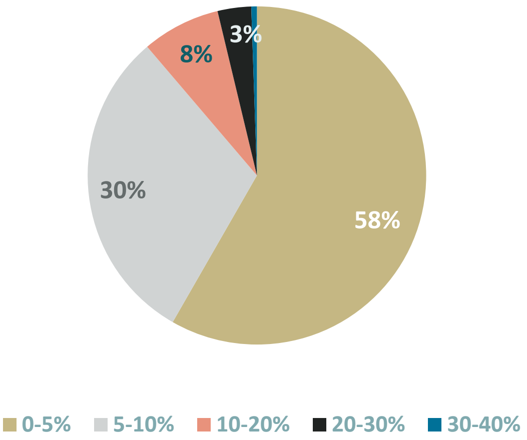

Figure 5: Estimate cost savings over the next 3-5 years due to adoption of innovative technologies

[15] *Markets in Crypto-Assets*

When looking at the detailed figures per type of technology adopted (see figure 6), we note that **the cost savings are anticipated to be higher when adopting digital onboarding techniques**. Such cost savings are also related to the impact that these technologies could have on process improvements and related reduction of operational costs. However, we see that there are limited expectations of cost savings in the field of DLT at the moment, on one hand due to the difficulties in estimating the return on investment of DLT projects often just experimenting this new technology on very limited scopes, and on the other hand due to fact that the majority of DLT initiatives are constituted by investments in the crypto assets area, thus not really aiming at decreasing costs.

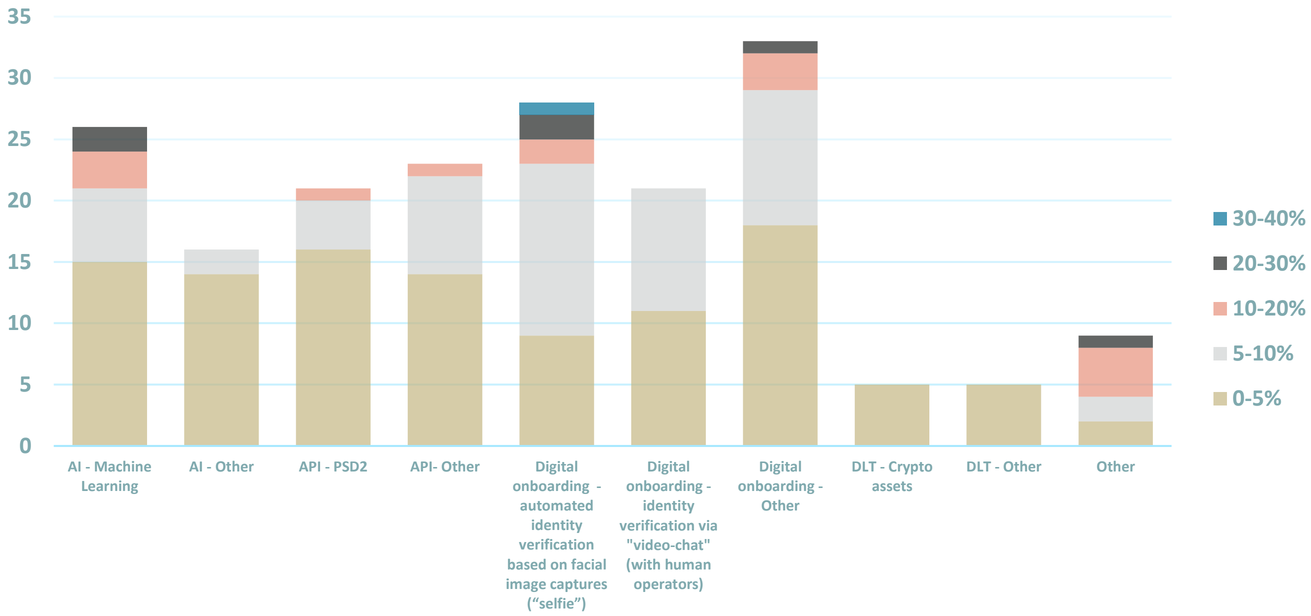

Figure 6: Estimate costs savings over next 3-5 years per type of innovative technology

Anticipated revenue increases due to adoption of innovative technologies

39% of entities reported that they anticipate a revenue increase due to the adoption of innovative technologies.

Among these entities, the anticipated revenue increase varies mainly between 0 and 40%[16], with a majority of respondents (86%) indicating an increased revenue between 0 and 10% (see figure 7).

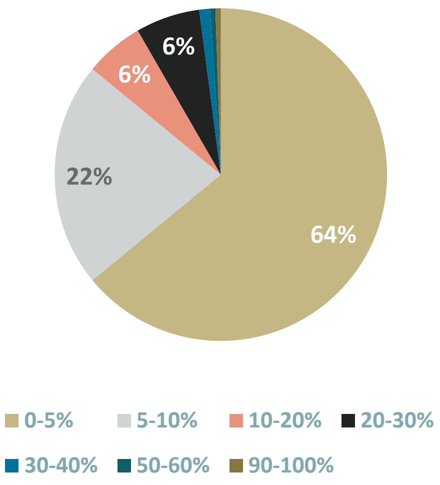

Figure 7: Estimate increased revenues over the next 3-5 years due to adoption of innovative technologies

As shown in figure 8 below (detailing the expected revenue increase per type of innovative technology), API is the only category of innovative technology for which some respondents indicated an expected revenue increase above 30% (if we exclude the "other" category).

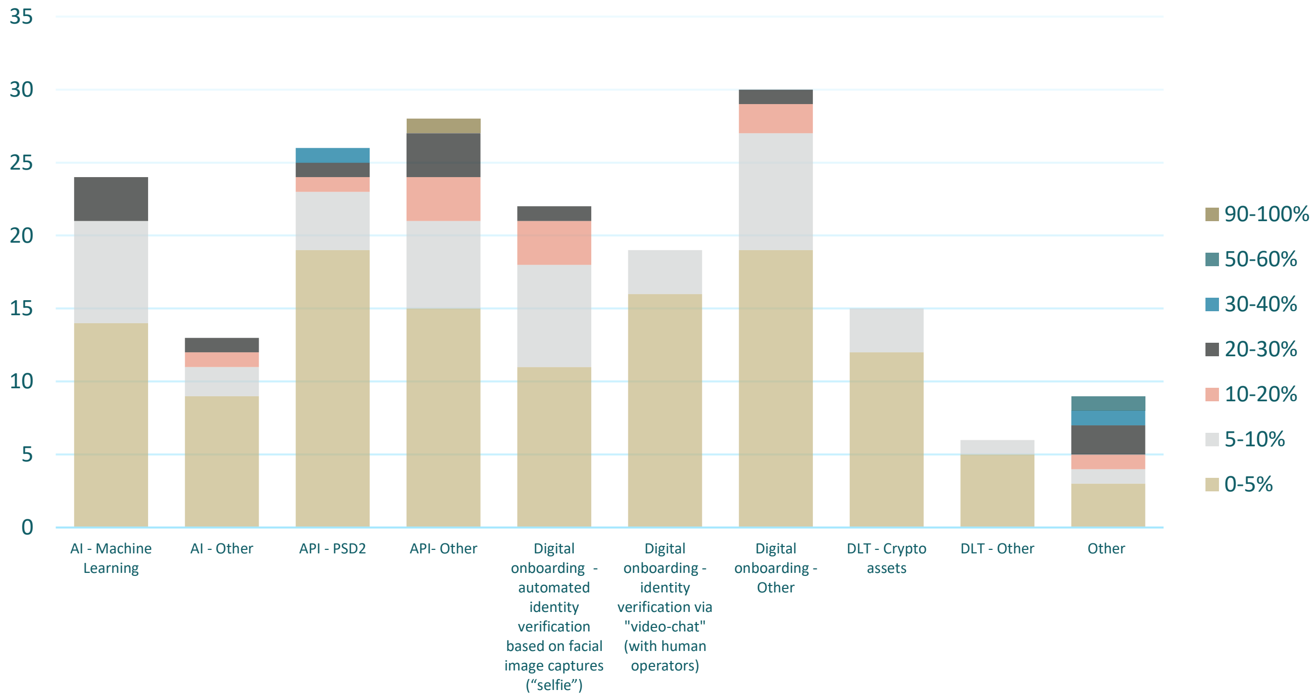

Figure 8: Estimate increased revenues over the next 3-5 years per type of innovative technology

[16] *One respondent indicated a revenue increase between 50 - 60% and another one indicated a revenue increase between 90 - 100%.*

Part 3 – AI adoption

This section focuses on the second part of the survey, which consisted of a detailed questionnaire on the usage of AI and ML by supervised institutions.

Use of AI and ML

At the time of the survey[17], **30% of all respondents indicated already making use of AI technologies**, and **25% of all respondents indicated using machine learning (ML) in particular**.

When focusing only on the entities using AI, 83% of them use ML[18], thus confirming ML as the most used AI technology (among all types of AI technologies).

Across the typology of institutions that answered the survey, we note that within e-Money Institutions and Payment Institution, there is a higher use of such technologies compared to Banks.

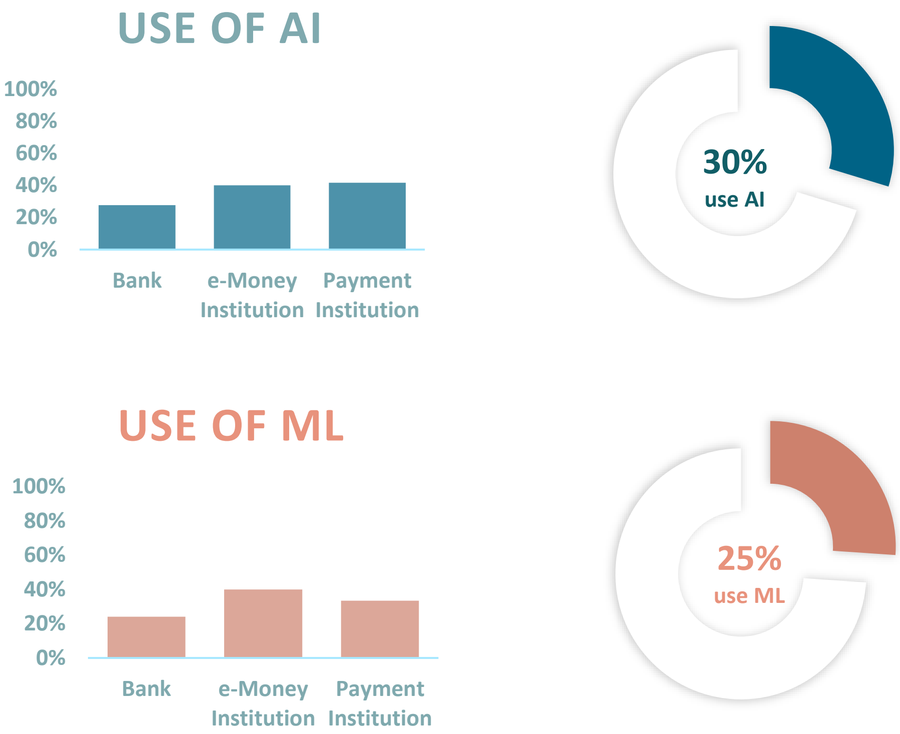

Figure 9: Use of AI and ML per category of respondents (bar charts) and across all respondents (pie charts)

[17] *As a reminder, the survey was performed during the period October 2021- January 2022*

[18] *It should be noted that ML is a subset of AI (see definition provided in the Annex)*

Note: as described earlier at section "Scope and methodology", all figures in the remainder of this document are reported as percentages of those respondents who indicated that they are using AI (which represent a subset of the total number of respondents), except for the statistics in sections "Machine learning lifecycle" and "Machine learning technical infrastructure", which are calculated based only on those respondents using ML specifically.

Benefits and challenges

Among the **top benefits** identified by the survey, the respondents indicated that they are primarily using AI and ML technologies to **improve internal efficiency**, followed by **reducing risk**, and **enhancing products and services offered**.

The "other" category showed some additional benefits such as improved big data treatment, improved informed decision making, enhanced reactivity and other benefits which could actually be linked back to the "improving efficiency" category.

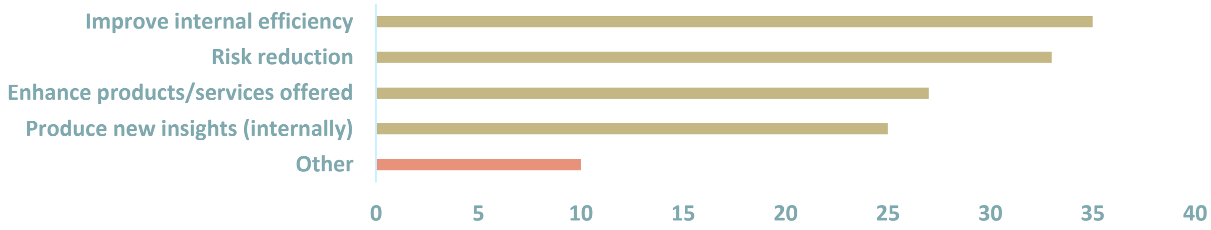

Figure 10: Benefits of using AI/ML

Regarding the **top challenges**, answers were more scattered, yet **data quality** was identified as the main challenge. Other challenges identified were the **availability of AI/ML skills** and the **ability to monitor efficiency of models over time** (due to the risk of model drifting), followed by **data governance**. The graph below gives an overview of all AI/ML challenges identified by the survey.

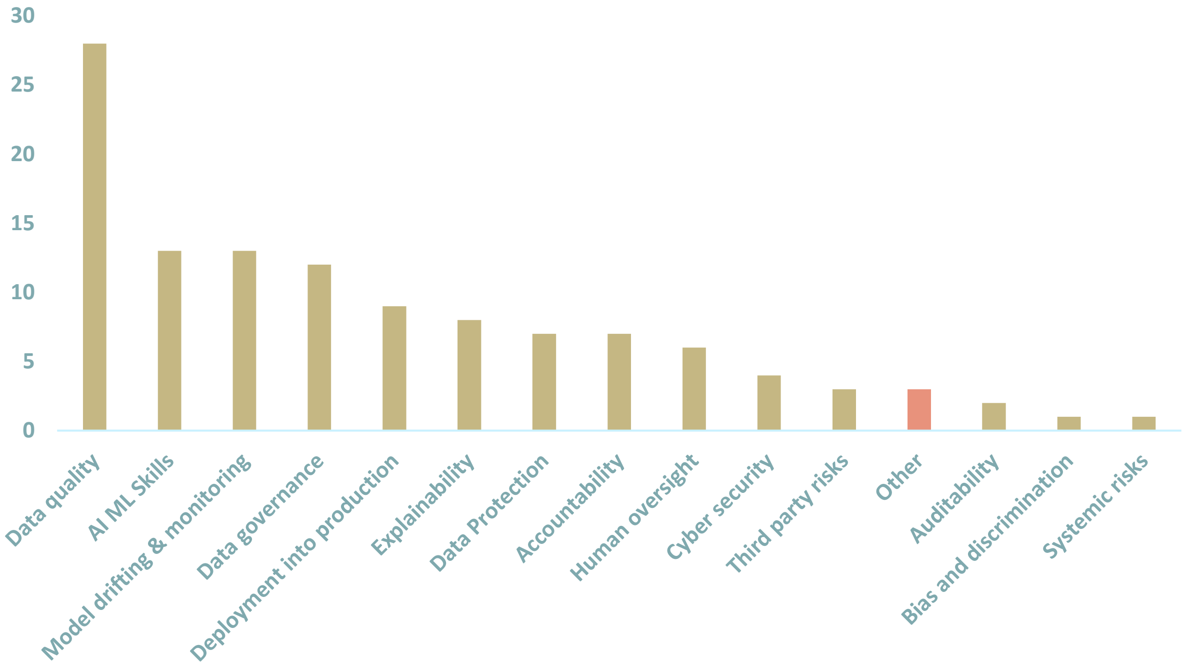

Figure 11: Challenges of using AI/ML

Additional challenges that were mentioned as part the "other" category included project challenges, multi-jurisdiction data regulation aspects and difficulties in reviewing and approving model risks.

Organisation

In relation to **organisation**, the survey included several questions aiming at identifying how institutions structure their teams when using AI/ML technologies.

The majority (86%) of respondents using AI have a dedicated team working only on AI related projects/ development activities ("data science team"), in most cases (72%) situated at group level or a combination of group level and local level. Considering that, as reported below, AI teams at group level are usually larger, these figures denote a general tendency to capitalise on group expertise for AI related development activities. Nevertheless, 14% of respondents indicated having a data science team only at Luxembourg (local) level.

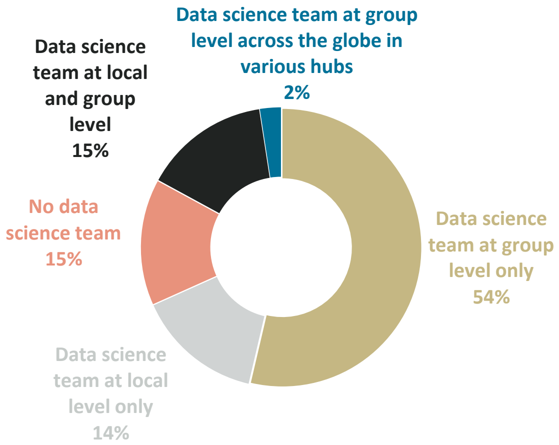

Figure 12: Location of the data science team within the group

15% of respondents using AI do not have any dedicated team working only on AI related projects/ development activities, most probably indicating that AI projects are managed directly within traditional IT organisational structures as part of standard IT development and change management activities (this is particularly the case for off-the-shelf solutions such as cyber security or RPA[19] tools, that do not require specific AI development teams).

Figure 12 provides more details about the different organisational setups reported.

According to the same respondents, the **data science team** is located most often within the **IT department (40%),** followed by a **dedicated unit (37%),** and less frequently **within the business lines (23%).**

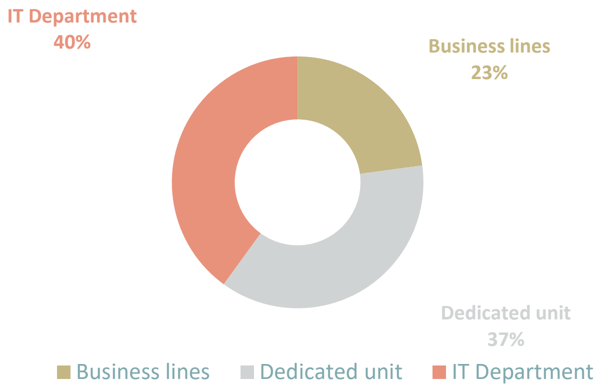

Figure 13: Location of the data science team within the organisation

Regarding the **size of the teams**, we note that there are small size teams in Luxembourg (up to 10 people), and that larger teams usually sit at group level[20].

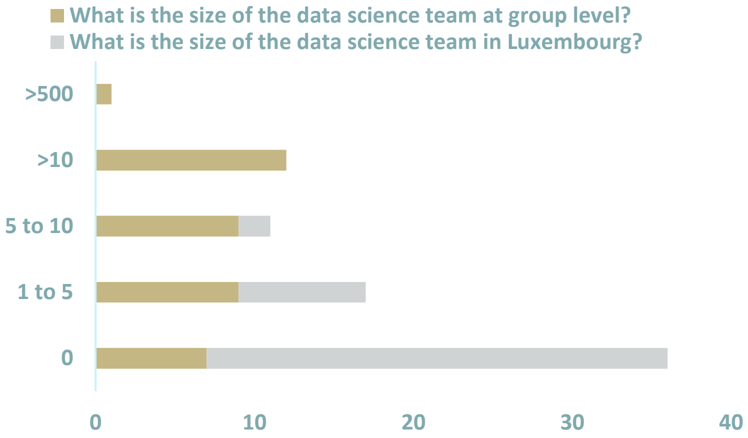

Figure 14: Data science team location and size

[19] *Robotic Process Automation*

[20] *One entity reported to have more than 500 people in their data science team at group level.*

Looking at the skills required for a data science team, the **top skill** most often reported by respondents is **data analysis**, followed by **statistics** and then **IT programming.**

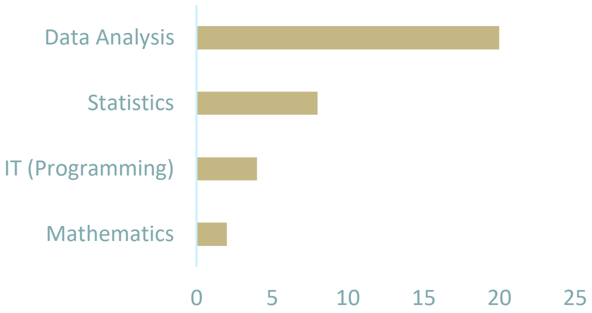

Figure 15: Top Skills required for a data science team

When inquiring about specific AI trainings, **70%** of respondents (which previously indicated to make use of AI) have provided **AI related training** to their employees, among which only **30%** provided specific **advanced trainings to their AI/ML developers/ data scientists (including upskilling)**, while 25% provided general AI/ML awareness training to all their employees, and 15% other types of training.

On the other hand, a still relevant portion of respondents using AI **(30%) does not provide any kind of AI related training** to their employees.

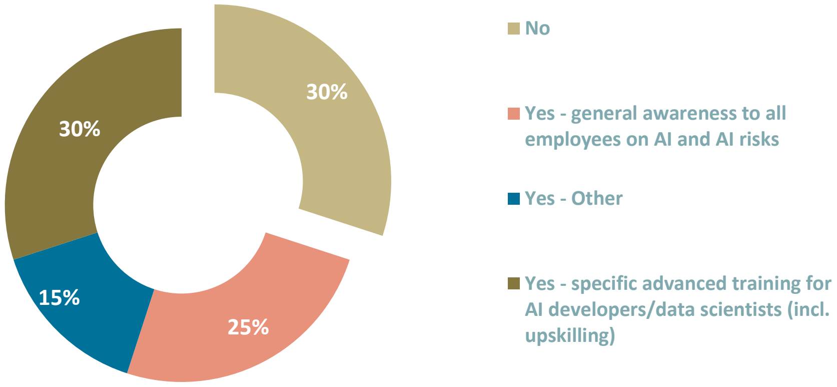

Figure 16: AI and ML trainings to employees

Data and governance

In relation to **data governance and data quality**, the survey included some questions to assess the maturity of institutions in relation to how they reflect the use of new technologies as part of their existing governance structure. Most **respondents indicated not having specific AI related governance mechanisms** such as an AI ethical policy (**78%**) or an ethics committee that is overseeing the use of AI (**93%**). These figures indicate a low level of maturity around the adoption of specific AI governance mechanisms, which could be justified by considering the early stage of adoption of AI technology (as it can be seen in "Part 4- Use cases/General aspects" below, a large portion of use cases is still in development phase and/or is based on "off the shelf" solutions).

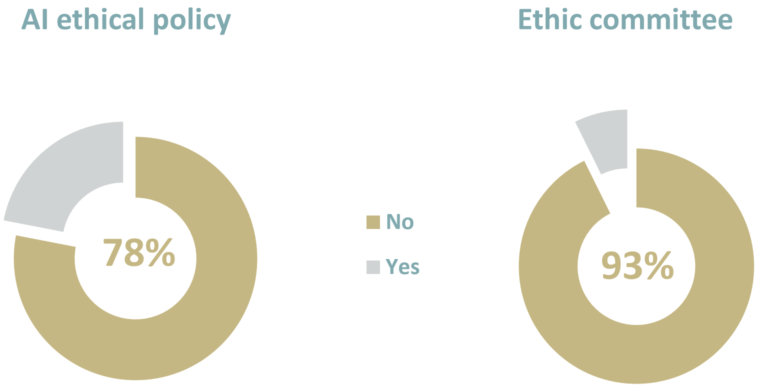

Figure 17: AI specific governance arrangements

From a traditional risk governance perspective, we found more reassuring figures indicating the involvement of the **Data Protection office**[21] **(83%)**, of the **Information Security function (88%)** and to a lesser extent of the **Risk function (63%)** in the AI/ML development process. These figures denote some consistency with the traditional IT development process and can be seen as a sign of possible integration of the AI/ML development process into the IT process. This last hypothesis is confirmed in the section "Machine Learning lifecycle" below, where the majority of respondents confirmed the application of standard change management process to AI developments.

[21] *DPO*

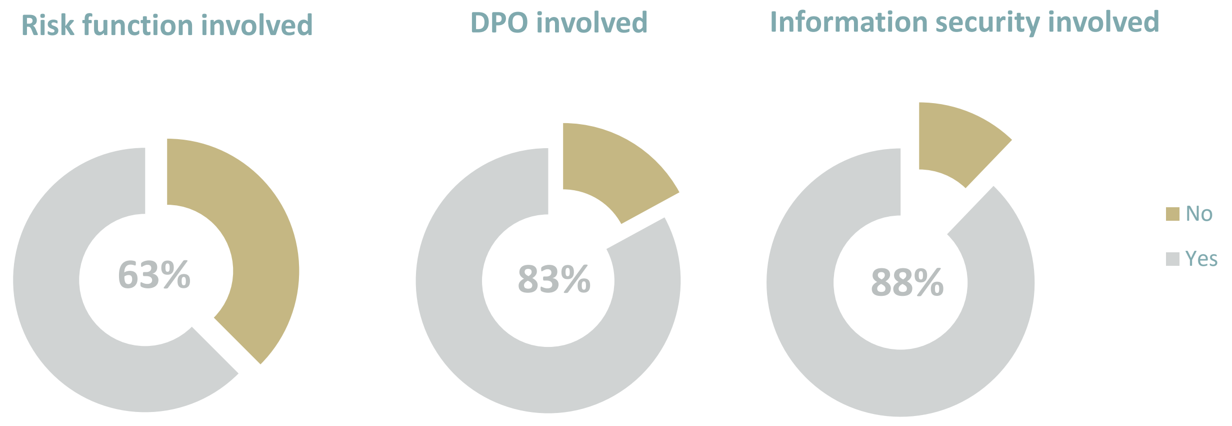

Figure 18: Involvement of traditional governance mechanisms

In relation more specifically to data governance and data quality, **56%** of respondents have **processes, policies and procedures in place which do not necessarily take into account AI specificities** and that will require adjustments to be fully applicable, while only **29%** of respondents have processes, policies and procedures **specific for AI data governance and data quality**.

15% of the respondents that use AI and ML **do not have any framework for data governance and data quality**.

The above figures may reflect the still early stage of adoption of AI by a large part of the respondents.

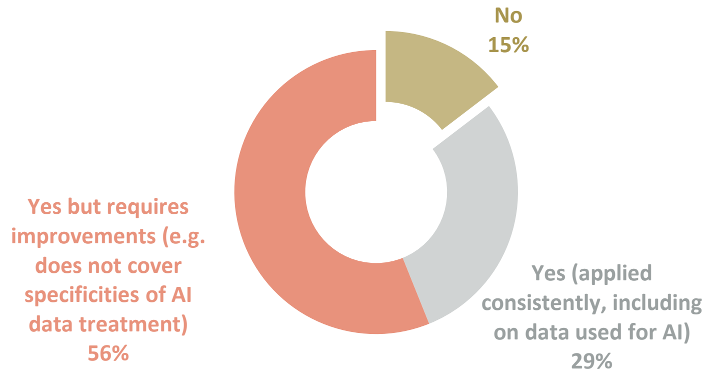

Figure 19: Availability of processes, policies and procedures for data governance and data quality

Security and robustness

In relation to **security measures** taken for ML specific security issues such as adversarial attacks, data poisoning and model stealing[22], we note that **nearly half** of the respondents that use AI/ML do take specific security measures while the others do not (figure 20). These security issues are not applicable to every use case, which can explain the low level of adoption of countermeasures. On a similar trend, **56%** of respondents that use AI/ML **perform independent security reviews and/or penetration tests**.

Nevertheless, some inconsistencies were identified when comparing these figures with similar questions asked at the level of the use cases. Only 27% of respondents are performing specific independent security tests of their AI/ML solutions at the level of the specific use case, which may be in part explained by the fact that a large part of the use cases are still in the development phase/ not yet in production.

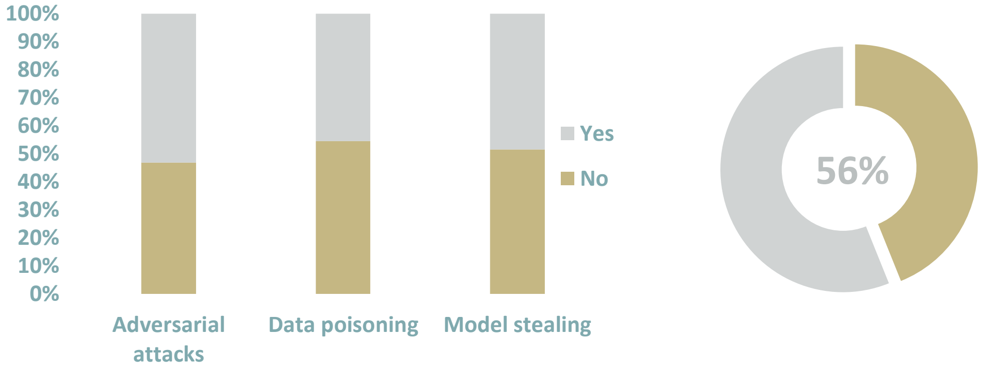

Figure 20: ML security measures (left) and independent security reviews/penetration tests of AI solutions (right)

[22] *See definitions available in the glossary under "ML security"*

Machine learning lifecycle

In relation to data science methodology, we note that **82%** of respondents using ML are **leveraging existing data science methodologies**[23] (e.g. CRISP-DM[24], CCC[25] etc.) while 12% rely on their own bespoke methodology.

With regards to the change management process, most respondents apply the standard change management process to AI/ML developments, while 26 % have dedicated ad-hoc change management processes specific for AI/ML developments.

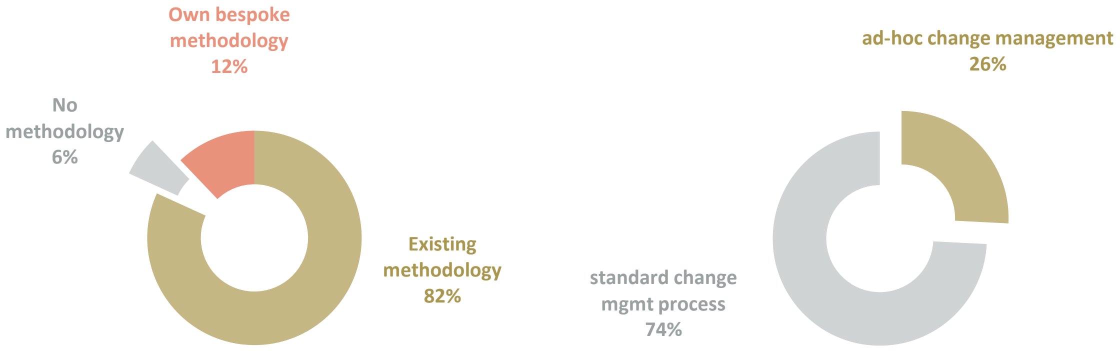

Figure 21: Data science methodologies and change management process

Concerning the lifecycle of ML algorithms, there appears to be a good hygiene of **versioning**[26] **models and data,** as **94%** of respondents state they employ techniques for the versioning of models and data used for training. Furthermore, **90%** of the respondents that use ML state they **have processes to monitor the algorithm performance over time** and that they **are able to revert to a previous version or to stop and replace with an alternative method when required**.

[23] *Among the answers received, CRISP-DM was frequently mentioned as example of methodology used for machine learning development.*

[24] *The Cross Industry Standard Process for Data Mining*

[25] *Computing Community Consortium*

[26] *Keeping track of different model versions across multiple iterations*

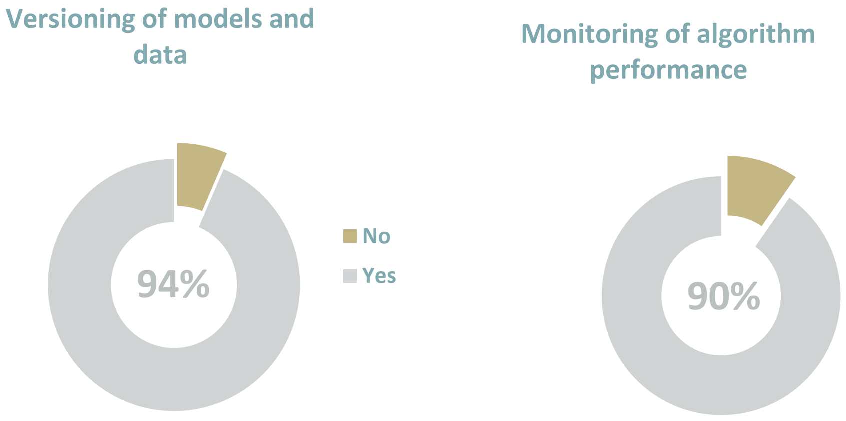

Figure 22: Versioning and performance monitoring

In terms of training frequency of ML models, the majority (**77%**) of respondents **retrain models on an ad-hoc basis** according to their needs rather than either a continuous or a fixed basis.

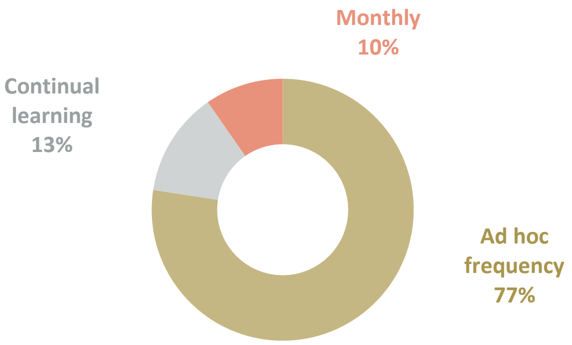

Figure 23: ML models retraining frequency

Machine learning technical infrastructure

In relation to the technical infrastructure supporting the ML processes, we note that ML development environments are **hosted primarily on premises (54%)** while cloud and hybrid environments represent respectively 14% and 32%.

When it comes to access to data, respondents have shown that they primarily use **traditional data bases (37%)** while **26%** use **data hub**[27] **or data lake**[28] concepts. **11%** of the respondents use **distributed databases** and the last remaining **26%** use **other methods** such as fileservers, flat files, data virtualisation or a mix of the above-mentioned approaches.

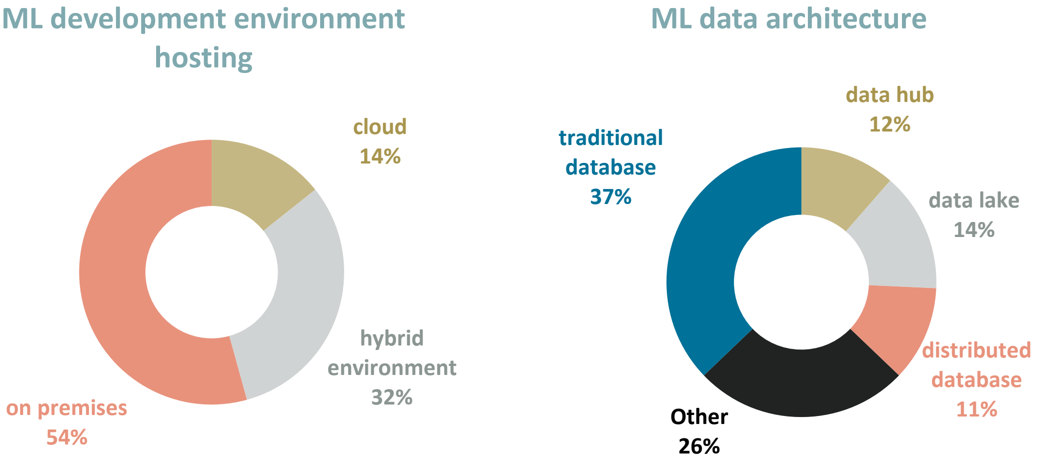

Figure 24: ML technical infrastructure

For ML data preparation and development, **82%** of the respondents are leveraging **open source** frameworks. **60%** of respondents are using **third party vendor solutions for ML development (including data preparation),** most of which are very common and recognised tools available on the market[29].

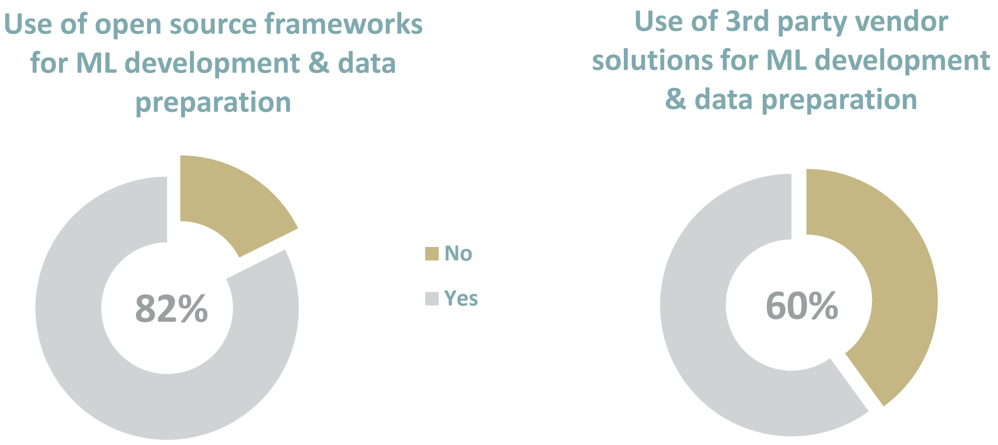

Figure 25: Open source frameworks and use of third-party tools

[27] *A data hub is mainly designed to exchange or share data. It will store semi-structured and harmonised data and make pre-processed and curated data available in various formats to simplify data exchange and sharing.*

[28] *A data lake is a central repository for storing, processing, and backing up large amounts of structured, semi-structured, or unstructured data*

[29] *Examples of third-party tools mentioned include: Microsoft Azure ML, Alteryx, DataRobot, Dataiku, Cloudera, Anaconda, Databricks, IBM Watson, etc.*

Part 4 – Use cases

The third part of the survey focused on the use cases applying AI technology, including ML, implemented by the surveyed institutions. This part included more detailed questions regarding the development status of the use case, the AI/ML technological aspects as well as more specific questions regarding the trustworthiness and security aspects of the use case.

The use cases had to be classified across the following set of predefined categories:

- AML and fraud detection
- Process automation
- Marketing and product recommendation
- Customer insights
- Cyber security
- Counter Terrorism Financing
- Sentiment analysis
- Credit scoring
- Customer support and help desk
- Algorithmic trading
- Robo-advisors
- IRB credit risk models
- Other

In the next sections, we will present some general findings across all use cases identified. The trustworthiness aspects of the use cases have been presented separately in part 5 of this document.

Use case categories

158 different use cases using AI technology were reported by survey respondents.

The top five use case categories reported were **AML/Fraud detection (18%), Process automation (15%), Marketing/Product recommendation (8%), Customer insights (8%), Cyber security (8%)**. Among the use cases reported in the "Other category"[30], there are different types of prediction/forecasting models (e.g. churn, possible investors, NAV, asset pricing, liquidity, balances, settlement failures, etc.…), some solutions related to process automation and various types of chatbots.

Among all surveyed institutions, banks (which represent also the highest portion of the respondents) appear more advanced in the use of AI technology compared to e-money institutions and payment institutions, in terms of number of use cases reported (in absolute value and also in terms of average number of use cases reported per respondent).

[30] *It should be noted that the categories are those as reported by the users, although some reclassification of some use cases may have been done.*

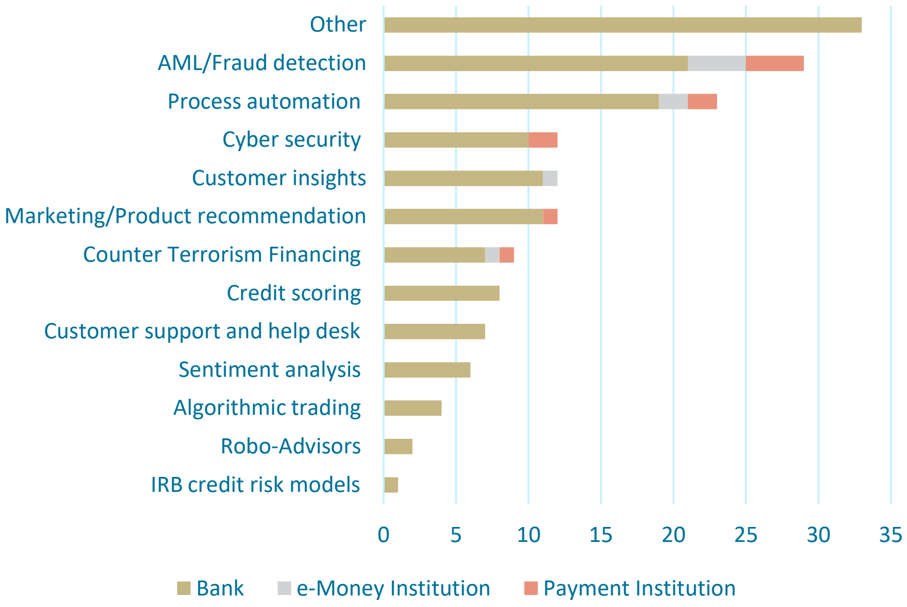

Figure 26: Reported use cases per type of institution

The following sections will provide a view on detailed aspects of the use cases, based on the answers provided by respondents in the "AI use case" section of the questionnaire. Given that some respondents did not answer to all detailed questions included in this section of the questionnaire, the graphs provided in the following sections will be based only on the answers received.

General aspects

The survey included some general questions regarding the development status of AI/ML solutions in order to assess maturity of use cases, as well as questions related to how these solutions are developed, procured and finally how they are used.

Overall, **59% of reported use cases are in production.** We note that some areas are less advanced than others, for instance algorithmic trading, robo-advisors and IRB[31] credit risk modelling which have a low number of use cases in production. At the opposite end of the spectrum we note that **process automation is quite mature if we look at the proportion of use cases already in production**. 100% of **cyber security use cases are reported to be in production** probably due to the fact that respondents are mainly relying on third party off-the-shelf solutions.

[31] *Internal ratings based*

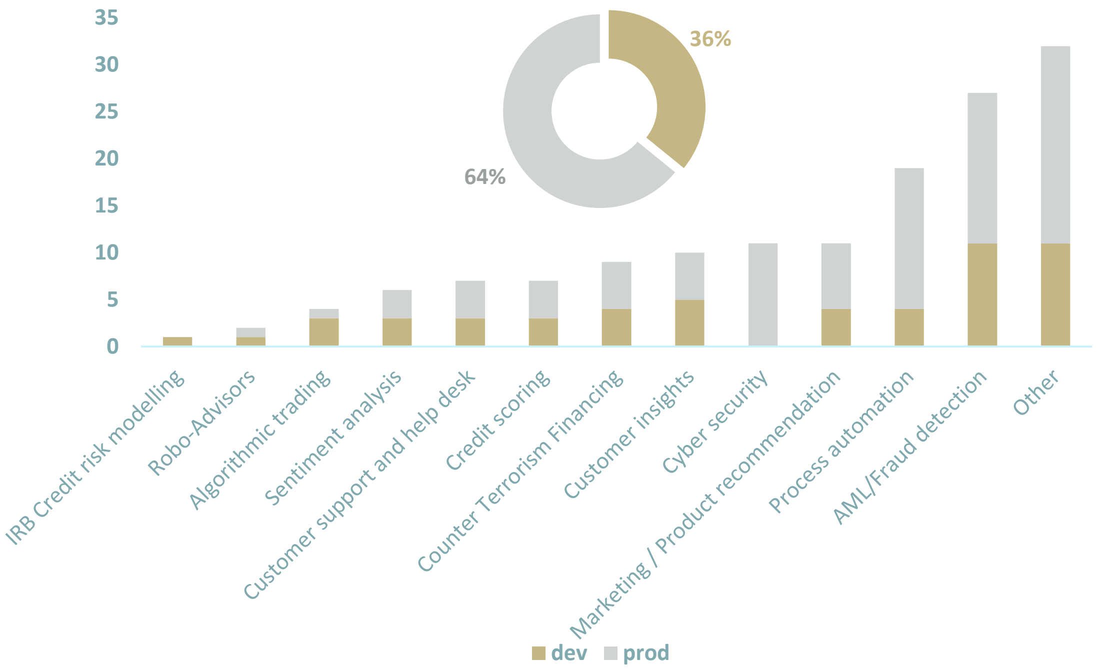

Figure 27: Use cases in development versus in production per use case (bar chart) and across all use cases (pie chart)

82% of the deployed solutions are used as "primary" models as opposed to "challenger" models[32]. Only in the cases of algorithmic trading, customer support/help desk and Counter Terrorism Financing ("CTF") we note that the usage of models as secondary/challenger is above 40%.

These figures, combined with the figures above on the development stage, confirm that the adoption of AI technology is still at an early stage.

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

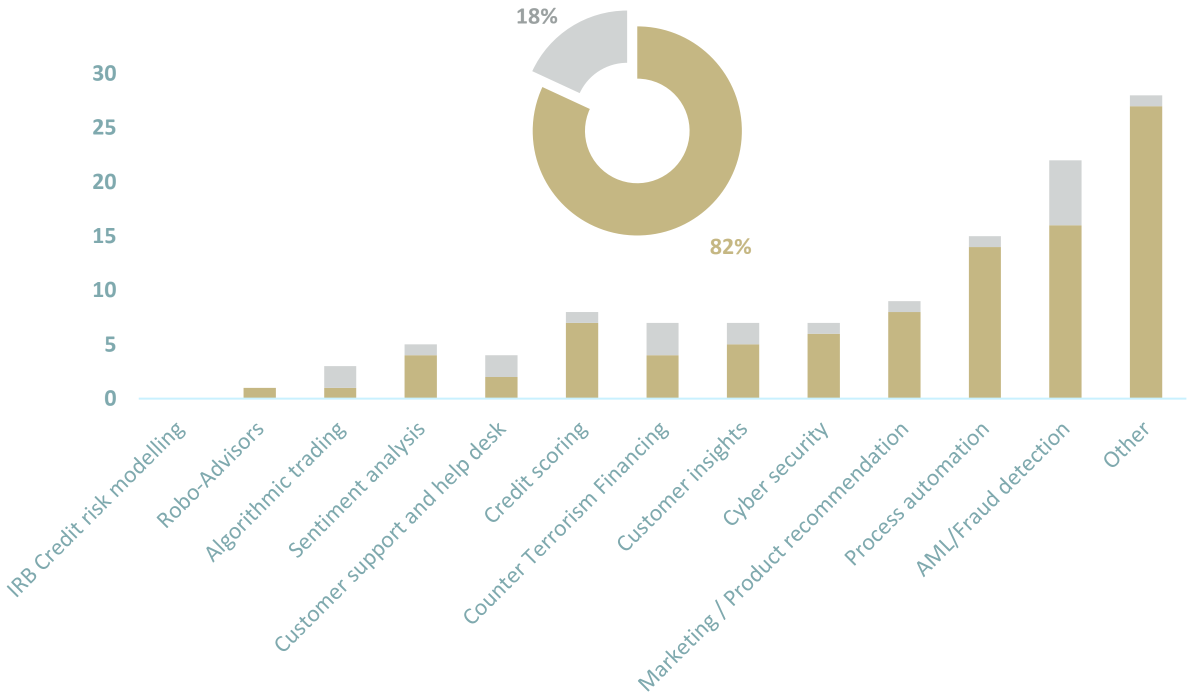

Figure 28: Primary versus challenger models per use case (bar chart) and across all use cases (pie chart)

The use of **external commercial AI products (or white label solutions) is rather limited (27%),** except in the "AML/Fraud detection", "Cyber security" and "Process automation" categories, which is probably due to the higher availability of products falling into these categories.

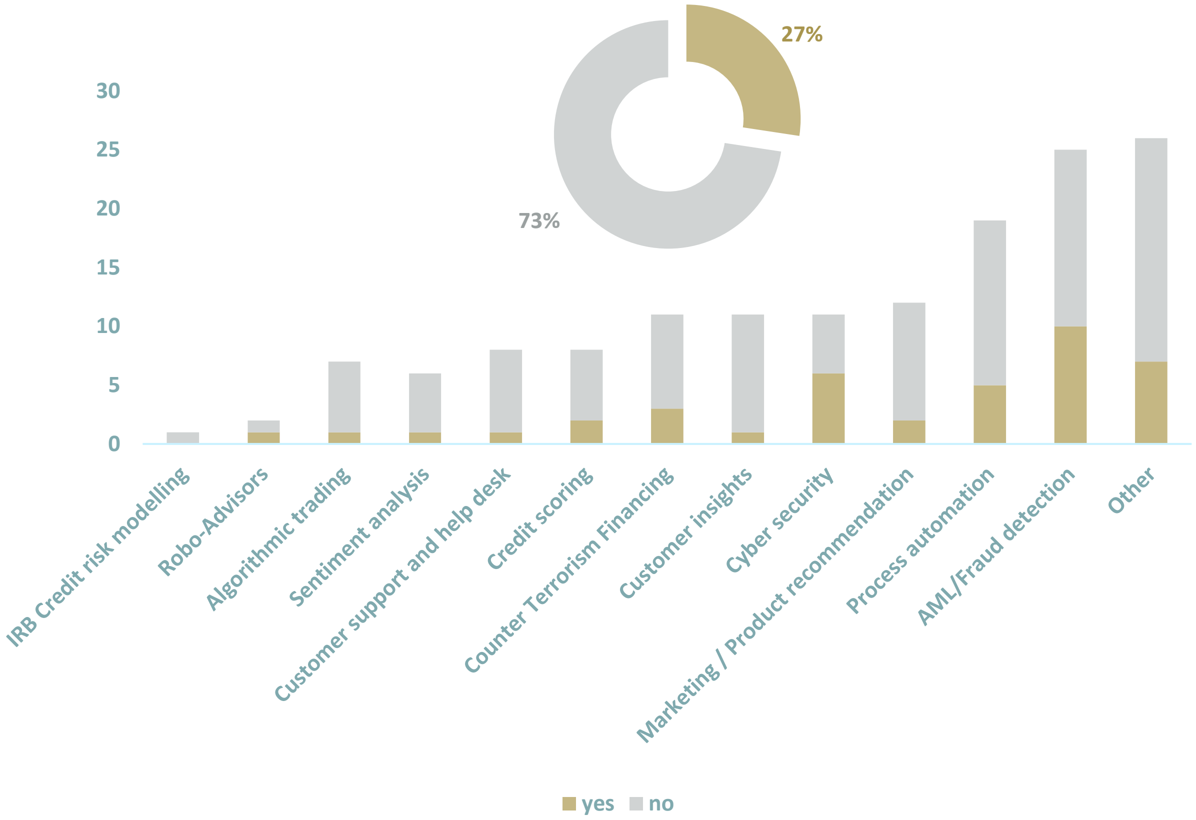

Figure 29: Usage of external commercial AI products (white label solutions)

In terms of development approach, **73% of use case solutions were developed in-house** while 15% were developed internally with external support. The fact that a higher percentage of respondents developed the AI/ML solutions internally (with or without external support) indicates a general trend to keep expertise internally, in order to be able to perform the maintenance of the AI/ML solutions. The remaining 12% of use cases were developed externally by third parties, used mainly in "cyber security" and "AML/fraud detection" categories with the use of "commercial AI tools".

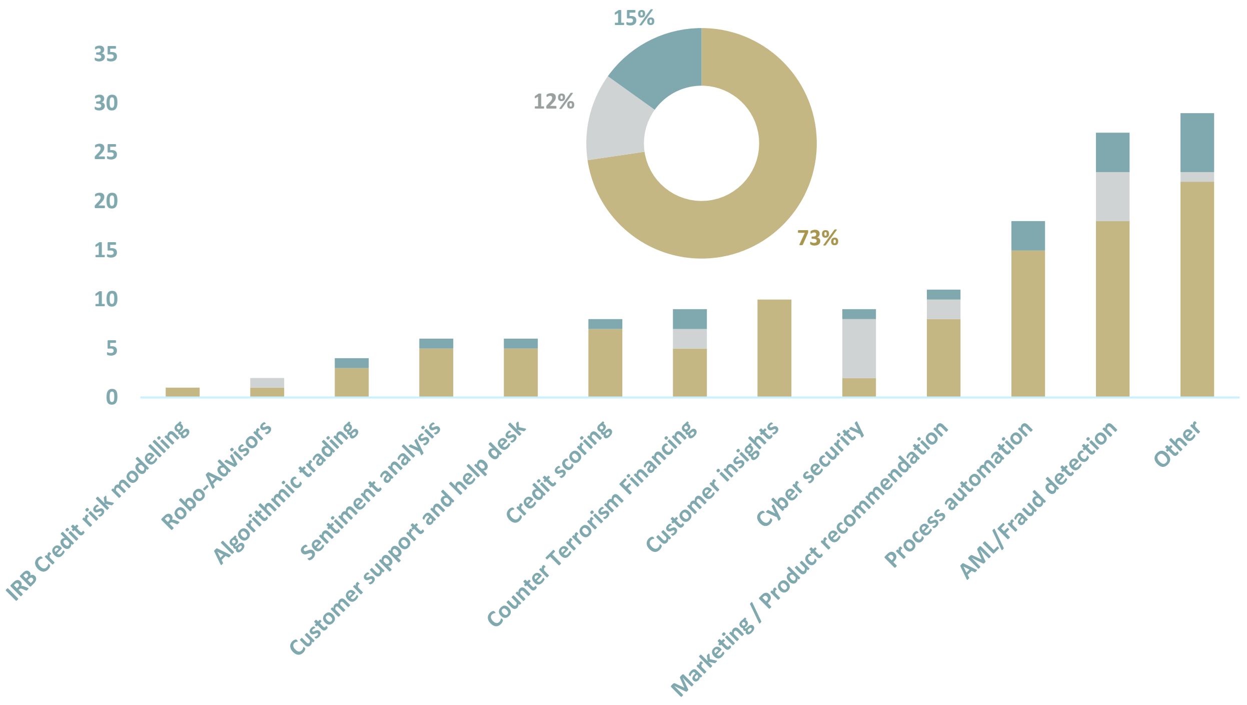

Figure 30: Development approach

In relation to usage of data in the different AI/ML solutions, respondents mainly stated that they use **internal data (62%)**, followed by both internal and external data (28%) and only external data in 10% of the cases.

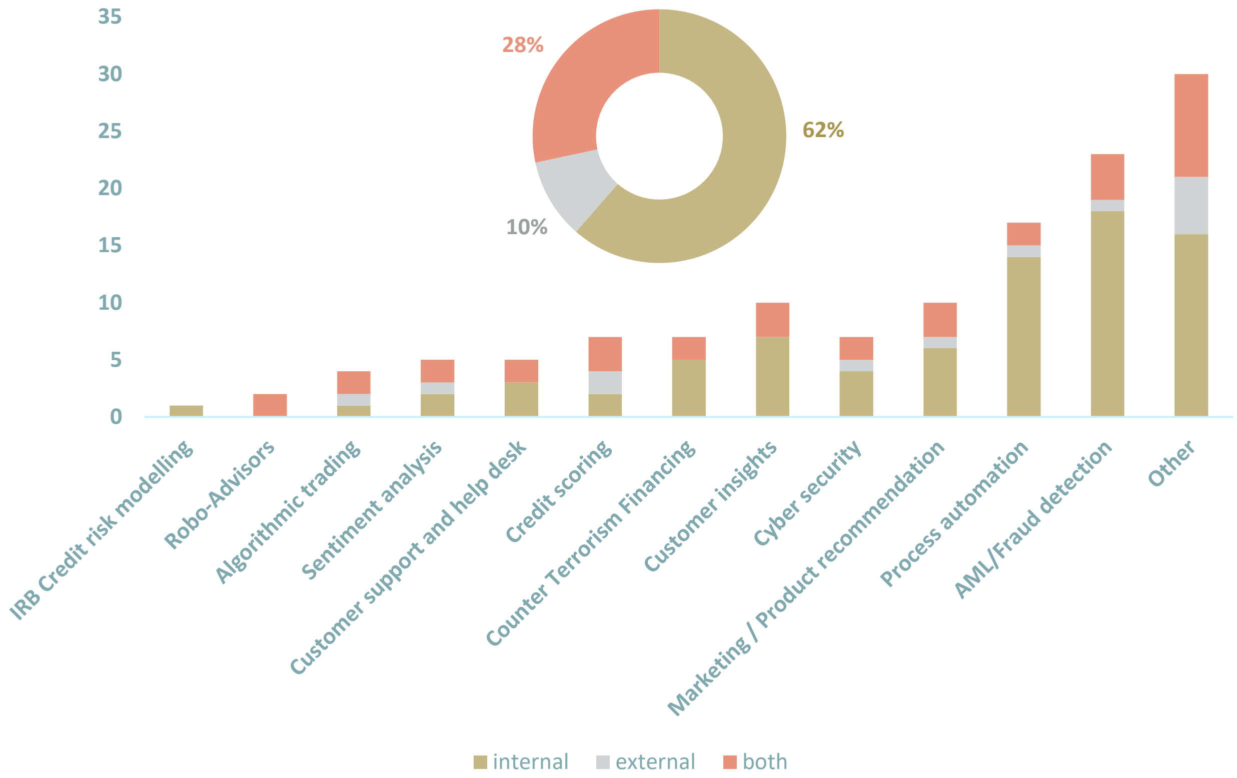

Figure 31: Use of internal and/or external data

AI technologies

In order to identify which types of AI technologies are used, respondents were asked to sort the use cases across the following categories:

- Machine Learning (ML)
- Expert systems (rule based)
- Natural Language Processing (NLP)
- Robotic Process Automation (RPA)
- Computer Vision
- Chatbots

The results show that **ML is the most broadly used technology (43%)** followed by rule based expert systems (20%), NLP (19%) and RPA (12%). Computer vision and chatbots are less common.

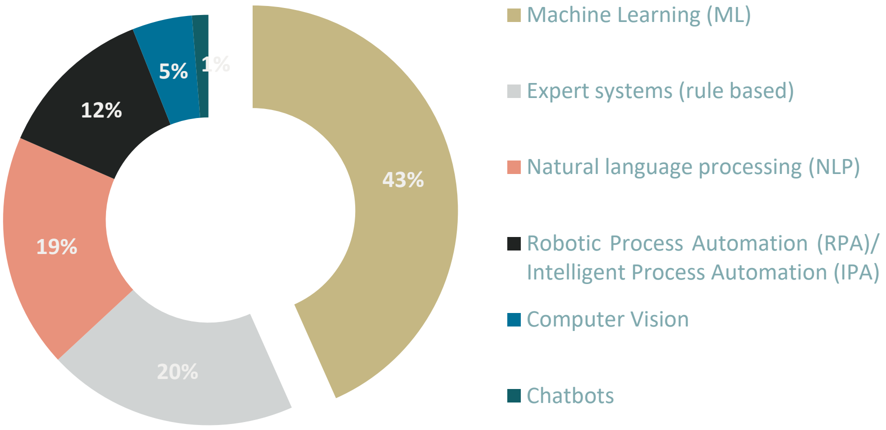

Figure 32: AI/ML technologies used

Figure 33 below shows the types of AI/ML technology used per category of use cases.

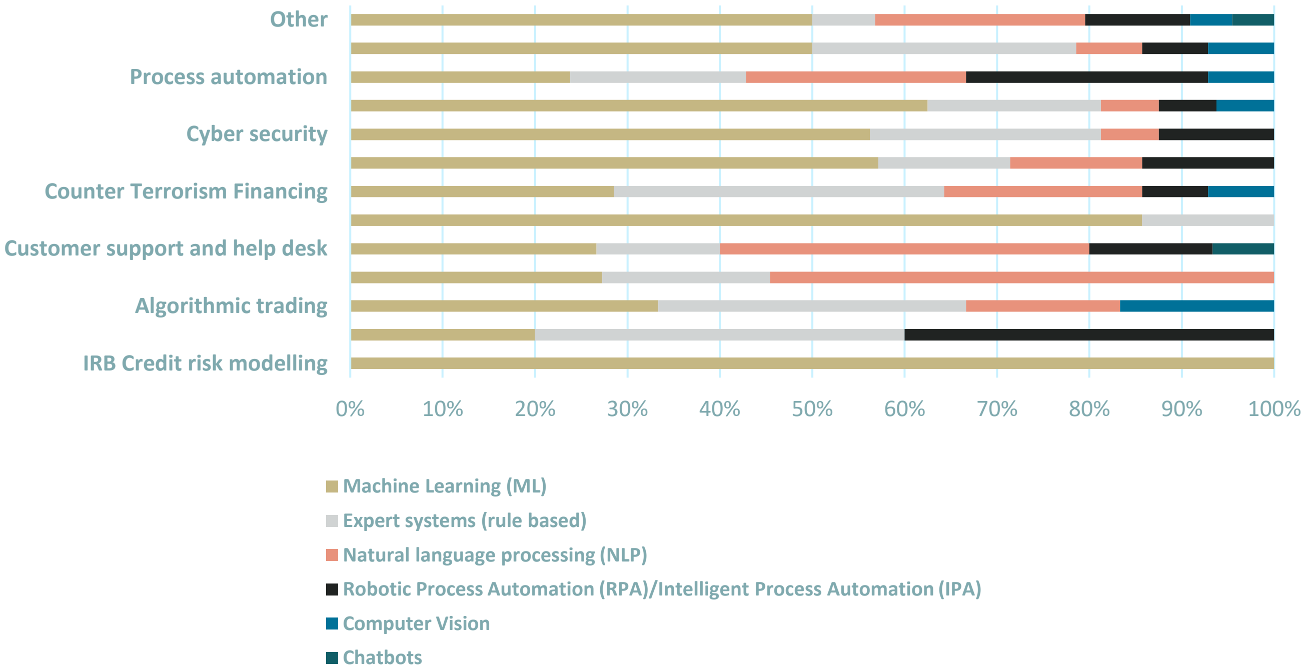

Figure 33: Usage of AI/ML technologies across the use cases

Type of ML problems

ML can be categorised based on the type of problems it solves. In the AI/ML questionnaire, respondents who stated that they were using ML were asked to further specify which type of problems were addressed by their ML models across the following:

- Classification
- Anomaly detection
- Regression
- Clustering
- Dimensionality reduction
- Association
- *Other*

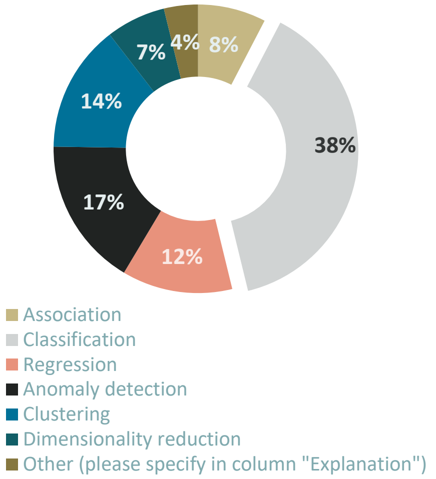

Figure 34: Types of problems addressed by ML models

We note that **most of the problems addressed by ML models relate to classification (38%),** followed by anomaly detection (17%) and clustering (14%).

The graph below shows the types of ML problems addressed by use case, where we can note that **anomaly detection techniques are most often used for the "Counter Terrorist Financing" and "AML/fraud detection" cases**, in line with the typology of these use cases.

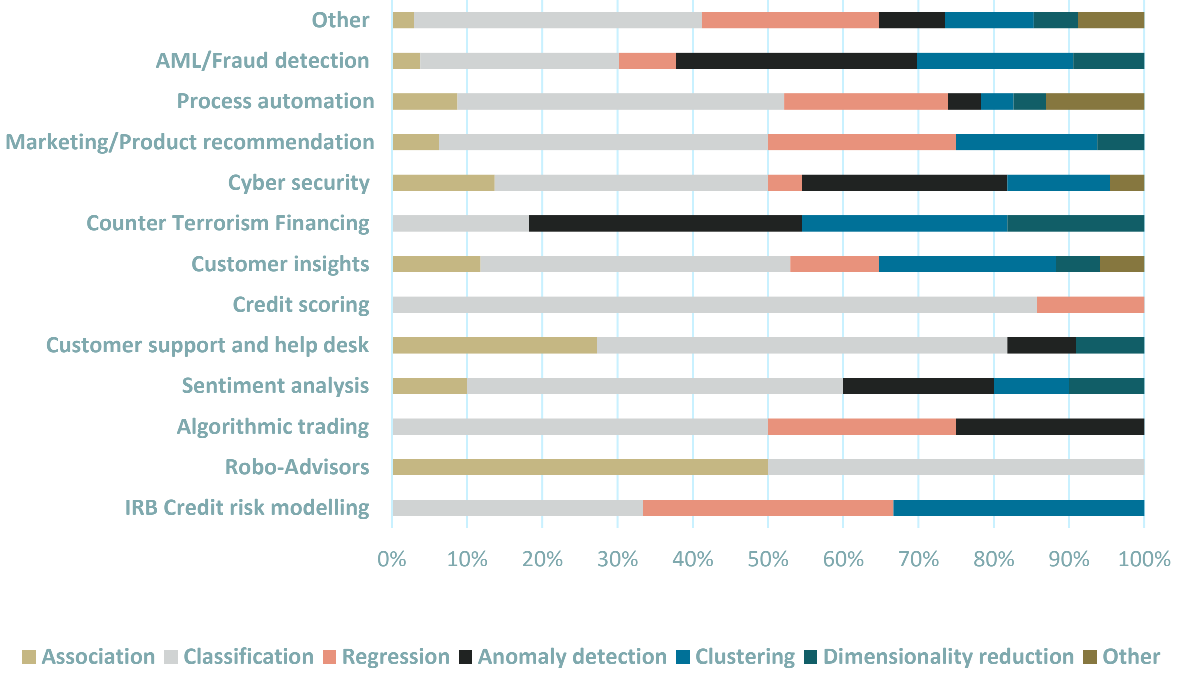

Figure 35: Type of problems addressed by ML models across the use cases

Type of learning

In relation to the type of learning that is used by respondents for their use cases, we note that **centralised learning accounts for 62%** of the responses and reinforcement learning for 20%. Transfer learning (11%) and federated learning (7%) are less common.

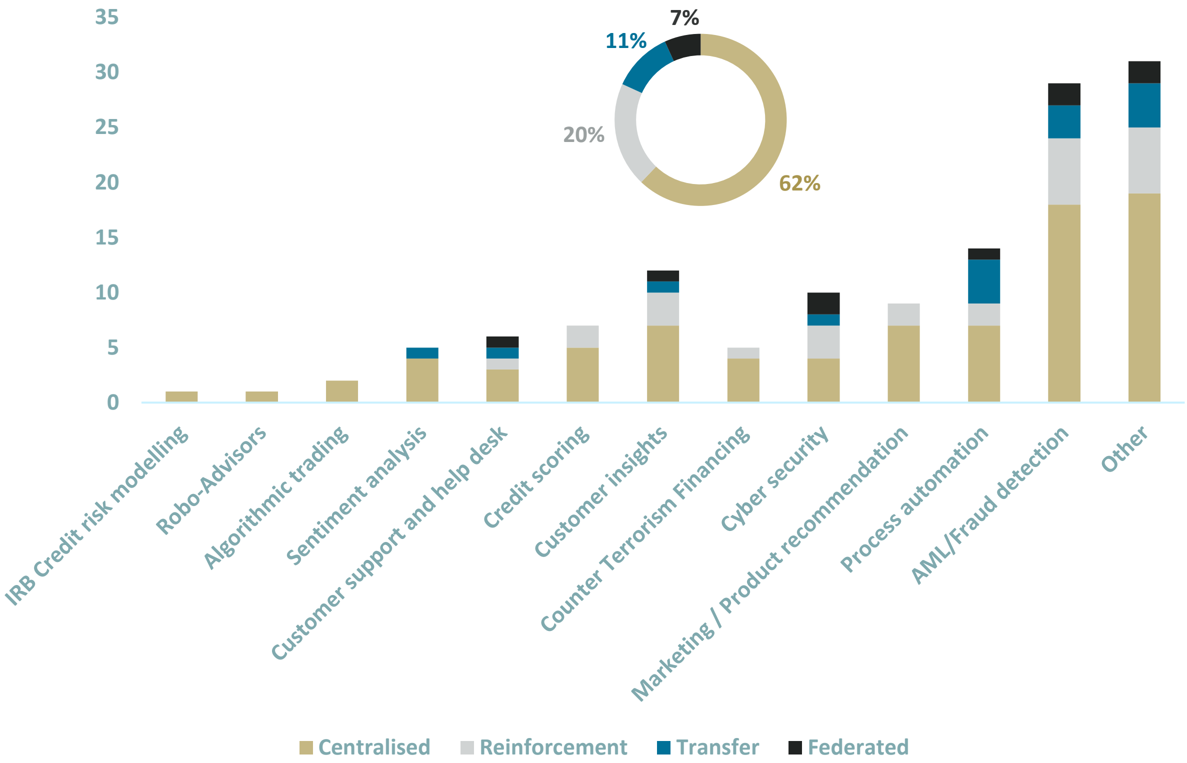

Figure 36: Type of learning

ML Algorithms

In relation to the specific ML algorithms used, we note that a large part of them are ensemble methods[33] (26%), followed by clustering (17%), deep learning (17%) and regression (13%) algorithms.

[33] *Ensemble methods combine a number of different ML techniques in order to produce better predictions than the individual ML technique*

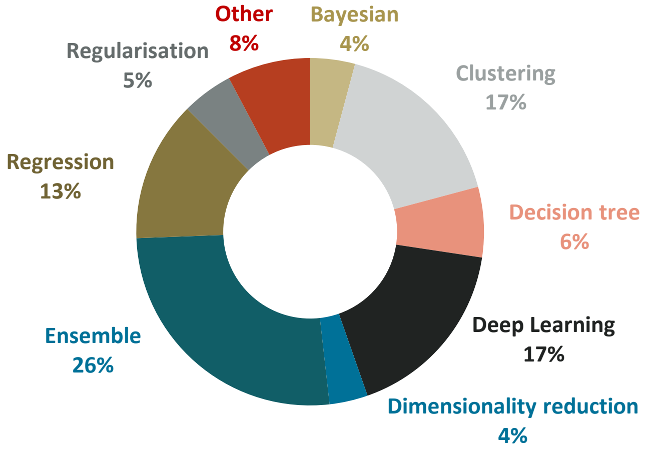

Figure 37: type of ML algorithm

In terms of specific ML algorithms used, we note that **the top 10 algorithms reported represent close to 75% of the total number of algorithms being used,** showing a limited variety of algorithms. Unsurprisingly, some of the more well-known algorithms are most reported, such as Random Forest, followed by K-Means, XGBoost and Logistic Regression.

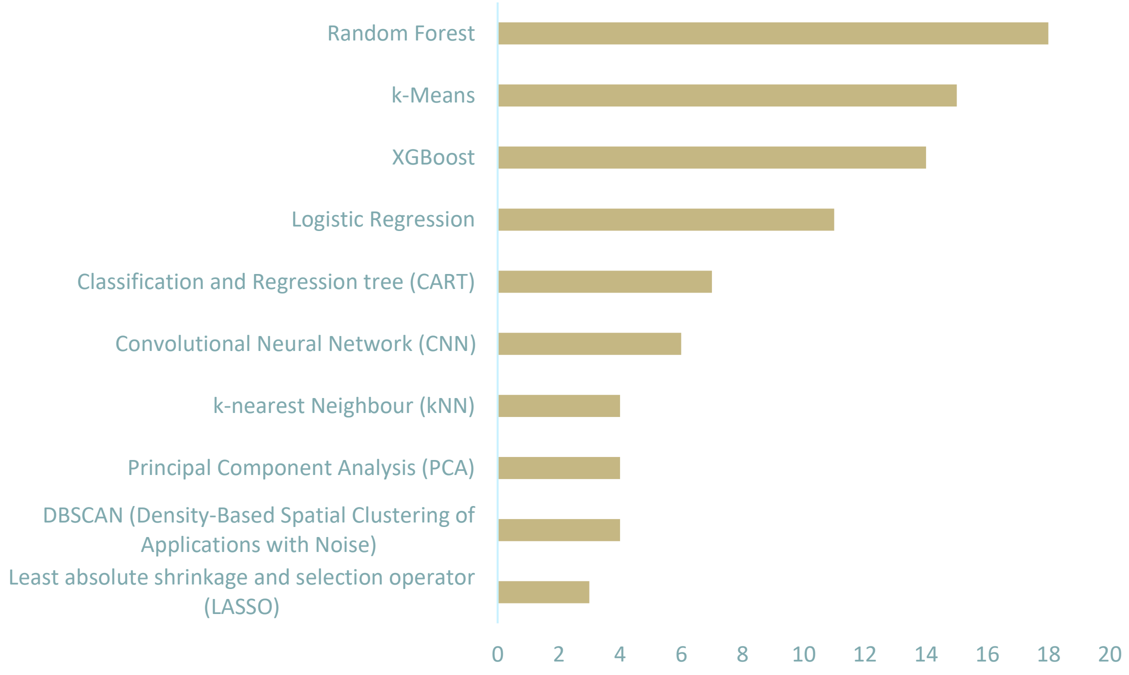

Figure 38: Top 10 ML algorithms used

Part 5 –AI trustworthiness

This section covers AI trustworthiness aspects across the use cases reported by survey respondents, including aspects related to the level of autonomy, explainability and auditability of the solutions implemented. Furthermore, this section also covers the bias prevention/detection techniques and specific security testing applied to the AI/ML models implemented by the survey respondents.

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, only **23% of the use cases have AI/ML models configured as "autonomous" systems**, i.e. not requiring a human in the decision process. This low figure can be seen overall as a good indicator of trustworthiness considering the importance of the human in decisional processes (depending on the criticality of the process within which the AI system is implemented). The cyber security use case is the one where we have the highest rate of autonomous solutions reported (44%), but this is mainly due to the fact that there is a high offering of "off the shelf" cyber security tools embedding automated AI engines. We also note that **for robo-advisors and algorithmic trading, all use cases are set up with a human in the loop. For credit scoring, AML/Fraud detection and CTF,** we note that **there is a good presence of human in the loop**, which is somehow in line with the criticality of the use case and the relative importance of the human in the decisional process.

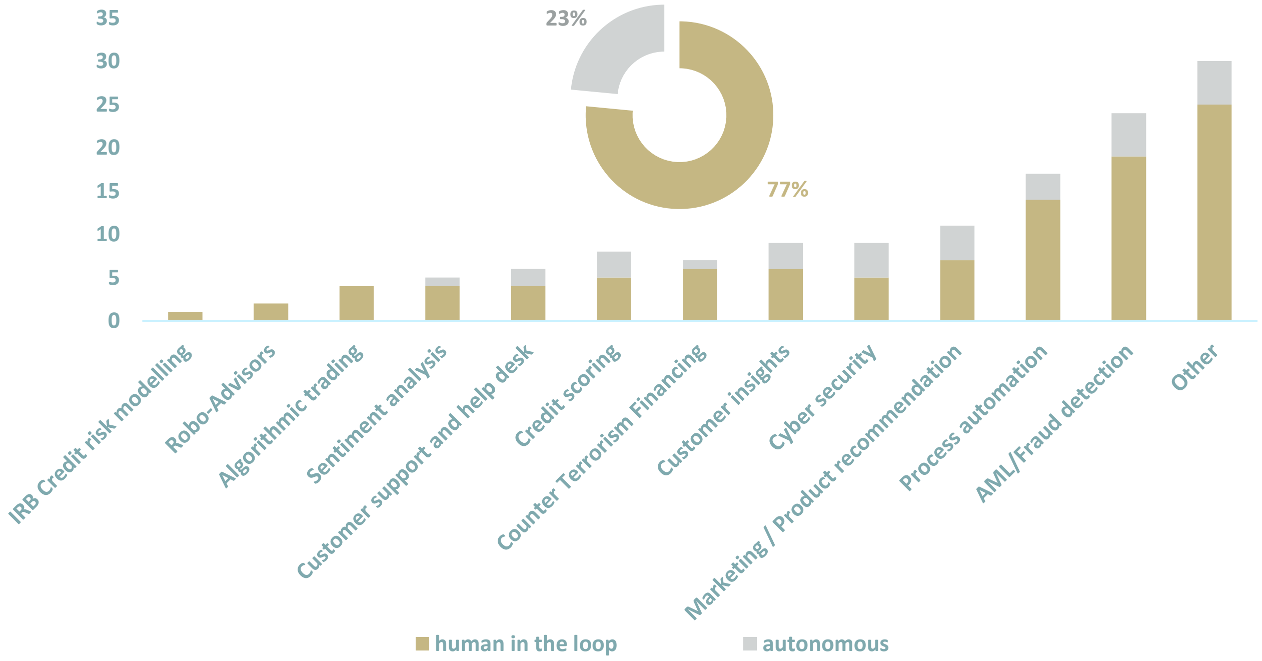

Figure 39: Human in the loop versus autonomous

Bias

The treatment of bias[34] is an important ethical matter for AI/ML solutions especially in a context where autonomous solutions are present. According to the survey, only **59% have implemented bias prevention and/or detection techniques** (37% have implemented bias prevention techniques, 7% bias detection techniques, and another 15% use a combination of bias prevention/detection techniques). We note that a number of respondents did not respond to this question in the questionnaire (i.e. there are some gaps - shown as red arrows in the figure below - between total number of use cases and the number of answers for the specific use case). This potentially indicates that the implementation of bias prevention and/or detection techniques is lower than indicated above.

Although bias detection/prevention techniques are more relevant in some use cases (e.g. credit scoring) compared to others (e.g. cyber security, process automation), this finding confirms the conclusion that the adoption of AI is still at an early stage and consequently there is a low level of maturity regarding the implementation of ethical control measures such as those related to fairness and bias.

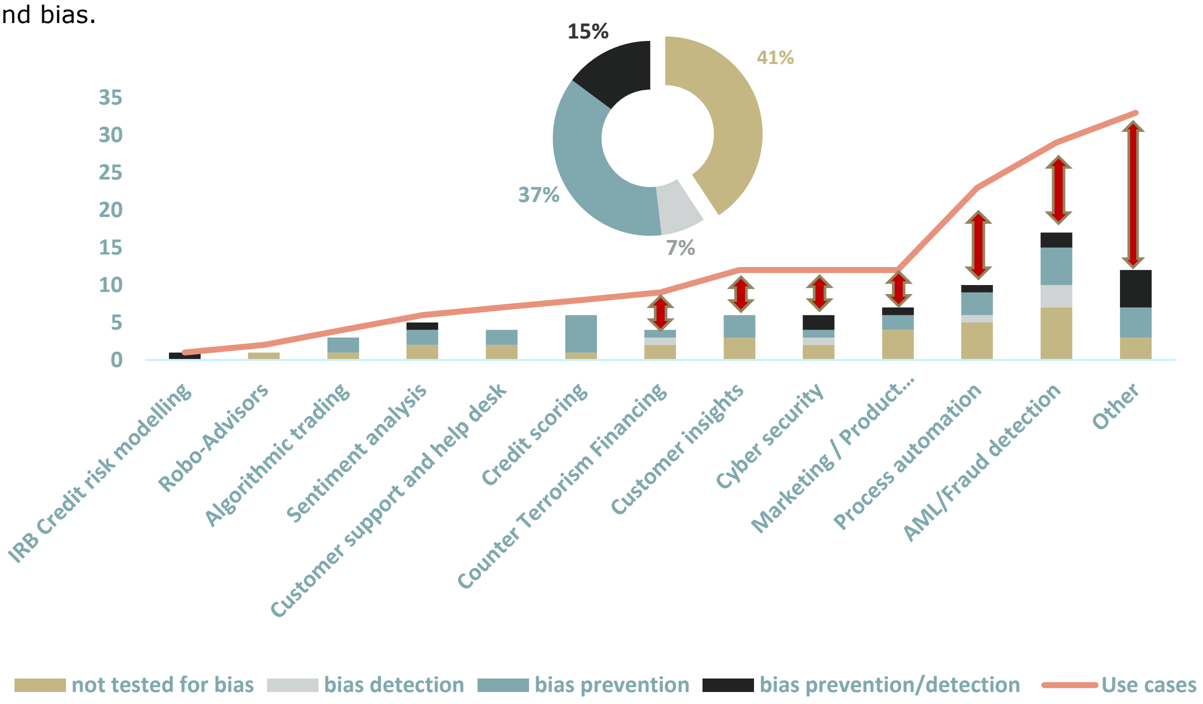

Figure 40: Bias detection and prevention techniques

[34] *A description of the concept of bias can be found in the EBA REPORT ON BIG DATA AND ADVANCED ANALYTICS, January 2020 (section 4.3).*

Auditability

Auditability[35] is key in making sure that systems can be challenged "a posteriori". In the survey, participants were asked to rate from 1 to 5 (5 being the highest rating, i.e. "very good auditability") how auditable their solutions were, and results were quite positive with **50%** of respondents reporting **very good auditability** and another **31%** reporting the **level of quality just below (4)**. We note however that for some use cases such as marketing/product recommendation for instance, the auditability sometimes appears more challenging although this finding needs to be put into perspective considering that higher levels of auditability are expected for more critical use cases.

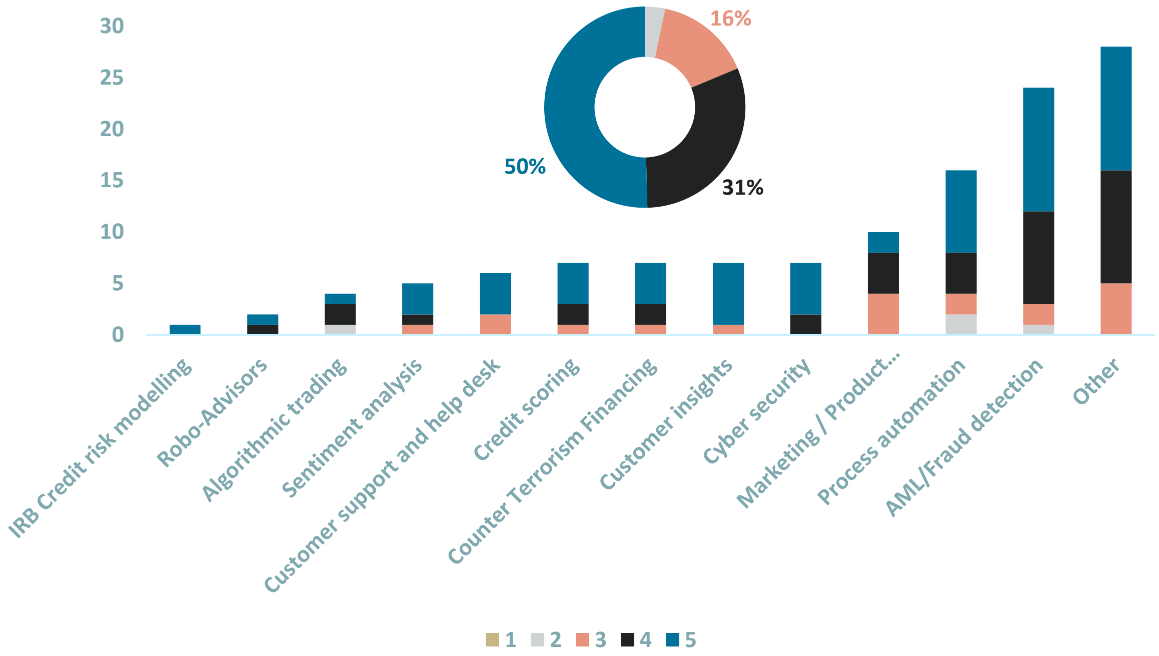

Figure 41: Auditability

[35] *Auditability of an AI/ML solution refers to the capability of tracing back all steps performed (from data extraction/preparation till the final model) and documenting the actions (who did what, when) performed at each step that led to the model final version, in order to able to reproduce the model results and to perform investigations in case of need. While the resulting documentation and audit logs do not explain why a certain result is produced by the model ("explainability"), they help to understand how the model was built and how the data has been processed in order to be fed into the model.*

Explainability

Explainability refers to the ability to justify and to provide a rationale for the predictions of an ML model. We asked respondents to rate from 1 to 5 (5 being the highest rating, i.e. "very good explainability") how explainable their solutions were, and results were quite positive with **32%** of respondents reporting **very good explainability** and another **38%** reporting **the level of quality just below (4)**. **Credit scoring use cases were rated with a good level of explainability (ranging from 3 to 5)**, which is reassuring given the importance of explainability for this category due to the direct impact of decisions on the clients. The lowest level of explainability (1) was only reported in the "process automation" category, which might be acceptable considering the specific type of use case.

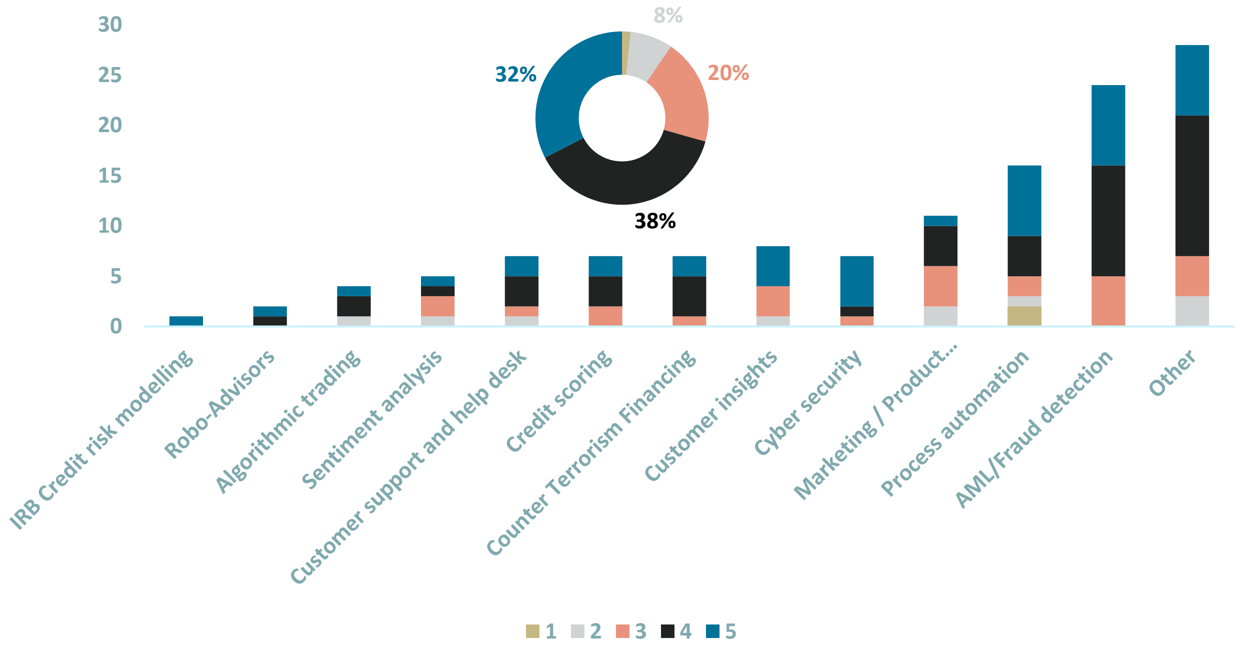

Figure 42: Explainability

Security testing

The survey included a question asking whether the underlying AI/ML model was independently tested against security attacks. **Only 27% of respondents confirmed that this is the case**. Similarly to the bias graph above (figure 40), there was a lower response rate for this question, indicating that the rate of independent testing could actually be lower than the percentage reported above.

These findings need to be put into perspective, since independent security tests might not be needed for use cases which are still under development and the importance of such tests vary depending on the type of use case. But ultimately, this also confirms that we are still at an early stage of the adoption of AI technology.

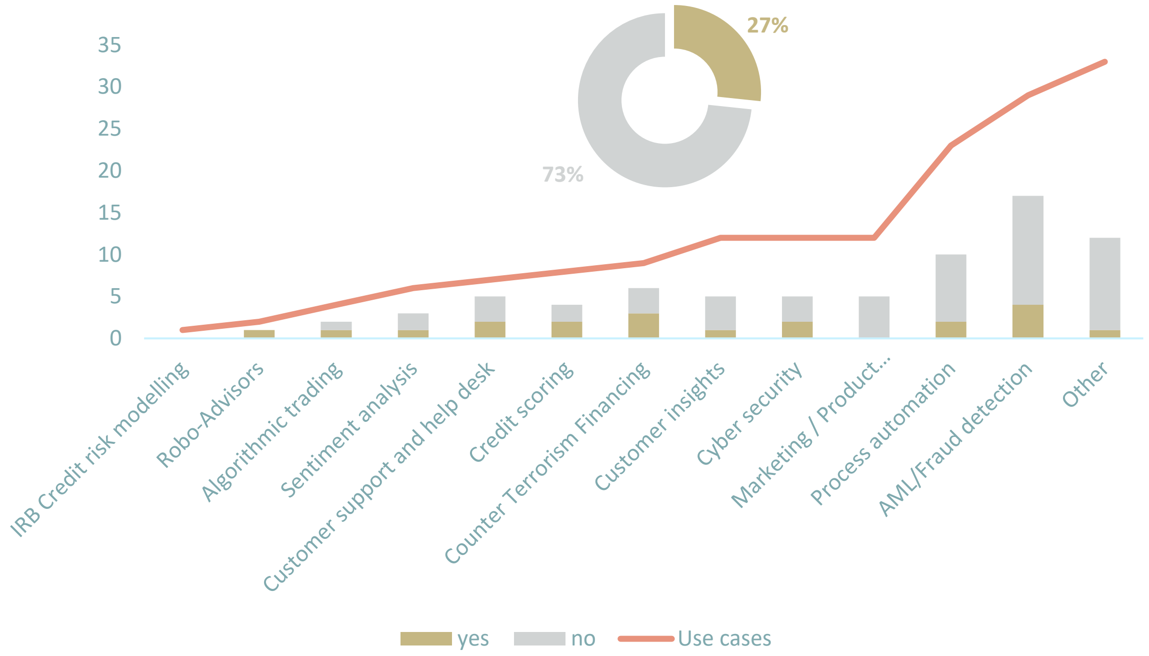

Figure 43: Independent security testing

Conclusion

The survey showed that the overall level of adoption of AI in the Luxembourg financial sector is currently limited (only 30% of surveyed institutions currently use AI technologies, with ML being the main AI technology used), but investments in this technology and especially ML are estimated to grow.

Several findings indicate that the adoption of AI is still at an early stage, especially regarding the implementation of advanced governance and ethical measures specific to AI. On the other hand, survey respondents demonstrated attention to AI trustworthiness aspects (e.g. human in the loop, bias detection/prevention techniques, auditability, explainability) with the level of adoption of the underlying techniques varying depending on the specific type of use case. This is an important finding and confirms that Luxembourg institutions using AI are aware of the specific risks related to this technology. These results further confirm the importance of continuing to take into account the recommendations included in the CSSF white paper[36], while waiting for the upcoming regulation from the European Commission laying down harmonised rules on AI (the 'AI Act')[37].

The recent public enthusiasm for advanced generative solutions like Chat-GPT[38] shows that the future is already here and that it will likely bring us more and more powerful AIs, with potential new ways to consume AI such as AI as a Service. Considering that this survey was the first exercise of this type, it will be interesting to see how the current picture of the use of AI in the Luxembourg financial ecosystem will evolve in the coming years.

[36] *https://www.cssf.lu/en/Document/white-paper-artificial-intelligence-opportunities-risks-and-recommendations-for-the-financial-sector/*

[37] *Proposal for a Regulation laying down harmonised rules on artificial intelligence | Shaping Europe's digital future (europa.eu)*

[38] *https://openai.com/blog/chatgpt/*

Sample of use cases collected

- Identification/detection of abnormal transactions (thresholds, risk scoring)
- Biometric verification
- ID documents authenticity verification
- Customer AML scoring
- Card fraud detection models
- False positive reduction using ML
- Insider trading detection
- Tax evasion detection

AML / Fraud detection

- Automation of manual tasks (transactional processes, trade input optimisation, non-regression testing, CA and reconciliation, LEI updates, adverse media monitoring…)
- Detection of customer distress
- OCR, intelligent document processing, classification and allocation to relevant teams, reconciliation of data entries, supplier invoices treatment, extraction of fields values
- Automation of credit approval process

Process automation

- Identify customer interests
- Identify best products for customers
- Segmentation and advertising
- Cross selling opportunities identification
- Lead generation and prioritisation
- Forecasting macro indices and trade indicators

Marketing / Product recommendation

- Detection and classification of the revenues and expenditures into specific categories
- Client segmentation
- ML to optimise customer and merchant experience at different stages of lifecycle
- Deposits rate sensitivity
- Client level reports

Customer insights

- Expert systems (rule based
- NLP to match PEP
- Pattern detection and analysis

Counter Terrorism Financing

- Commercial spam filter
- Autonomous solution to detect and respond to cyber attack
- ML to detect abnormal network behavior, account takeovers and malwares on end-points
- User Behavior Analytics to detect Azure Directory anomaly patterns

Cyber security

- Pre-approval of consumer loans
- Credit card limits automatic update
- Automatic decision regarding overdrafts
- Default probability evaluation
- Applicants scoring

Credit scoring

- Client message categorisation
- Services survey verbatim analysis
- Trending topics detection
- Emotion detection
- Market sentiment

Sentiment analysis

- Conversational AI
- Email dispatching and handling
- Complaint management
- Customer contact prediction and intent

Customer support and help desk

- FX trading prediction
- Fail trade prediction

Algorithmic trading

- Automated asset management
- Automated investment advice

Robo-Advisors

IRB credit risk models

Glossary and Abbreviations

AI (Artificial Intelligence)

According to FSB[39], AI is "the theory and development of computer systems able to perform tasks that traditionally have required human intelligence". Other definitions of Artificial Intelligence exist, such as the one included in the European Commission's AI Act[40], or others from IOSCO[41] or OECD[42]. 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.

API (Application Programming Interface)

Software interfaces allowing applications to communicate and interact with other software applications/services without needing to know how they are internally developed.

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.

[39] *FSB (Financial Stability Board) (2017), 'Artificial intelligence and machine learning in financial services – Market developments and financial stability implications' (https://www.fsb.org/wp-content/uploads/P011117.pdf).*

[40] *European Commission, Proposal for a Regulation laying down harmonised rules on artificial intelligence (2021), COM(2021) 206 final (https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence).*

[41] *IOSCO (International Organization of Securities Commissions) (2021), The use of artificial intelligence and machine learning by market intermediaries and asset managers – Final report, FR06/21 (https://www.iosco.org/library/pubdocs/pdf/IOSCOPD684.pdf)*

[42] *https://oecd.ai/en/ai-principles*

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.

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.

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 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.

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.

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.

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).

Regression

Regression problems are similar to classification in that they both use labeled 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.

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).

BANQUE CENTRALE DU LUXEMBOURG
EUROSYSTEM

cssf
Commission de Surveillance
du Secteur Financier