



MGA

MALTA
GAMING
AUTHORITY



MALTA DIGITAL INNOVATION AUTHORITY

AI Gaming Charter

on the Ethical and
Responsible Use of
Artificial Intelligence

A voluntary framework promoting responsible, transparent
and proportionate use of artificial intelligence across gaming
operations

G-SPG-025-01 | PUBLIC

18/09/2026

Contents

Revision History	5
Definitions.....	6
1 Introduction	8
1.1 Overview	8
1.2 Purpose and Scope	9
1.3 Regulatory Perspective and Responsible Use of AI	11
2 Study on AI Use and Adoption in Gaming.....	12
2.1 Use Cases of AI in the Gaming Sector	13
2.2 AI Literacy and Awareness	14
2.3 AI Strategy.....	14
2.4 EU AI Act and Regulatory Considerations.....	15
2.5 Risk and Operational Readiness.....	15
2.6 Data and AI Governance	15
2.7 Business-to-Consumer (B2C) Licence Holders.....	16
2.8 Business-to-Business (B2B) Licence Holders	16
2.9 B2B and B2C Licence Holders.....	16
3 The Charter.....	17
3.1 Risk-Based Application of the Charter.....	17
3.2 Transparency, Accountability and Explainability.....	19
3.2.1 Transparency in AI Design and Usage.....	19
3.2.2 Clear Communication.....	23
3.2.3 Risk Communication.....	24
3.2.4 Algorithmic Transparency	24
3.2.5 Error Reporting and Correction	25
3.2.6 Transparency in Open-Source AI.....	25
3.2.7 Explainability	25
3.3 Fairness and Non-Discrimination	26
3.3.1 Algorithmic/Decision Making Fairness.....	26
3.3.2 Diverse and Representative Data.....	26

3.3.3	Implementation of Bias Detection and Mitigation Measures.....	27
3.3.4	Designing AI Solutions to Address Bias.....	27
3.3.5	Inclusive Design	28
3.3.6	Continuous Monitoring and Evaluation	28
3.3.7	User Empowerment and Redress.....	29
3.3.8	Community Engagement and Consultation.....	29
3.3.9	Fundamental Rights	29
3.4	Environmental Sustainability.....	29
3.4.1	Sustainable Procurement.....	30
3.4.2	Energy Efficiency	30
3.4.3	Carbon Footprint Reduction.....	30
3.4.4	Efficient AI Model Training and Deployment	31
3.4.5	Green AI Practices	31
3.4.6	Eco-design Principles.....	31
3.4.7	Environmental Impact Assessment.....	31
3.4.8	Continuous Improvement.....	31
3.4.9	Stakeholder Collaboration.....	32
3.4.10	Responsible Decommissioning.....	32
3.5	Data Protection, Security, and Governance	32
3.5.1	Compliance with Data Protection Laws.....	32
3.5.2	Lawful and Ethical Data Processing.....	33
3.5.3	Data Minimisation and Purpose Limitation.....	33
3.5.4	Data Quality.....	33
3.5.5	Data Protection by Design and Default.....	33
3.5.6	Data Security, Safety, and Confidentiality	33
3.5.7	Transparency and Accountability in Data Processing.....	34
3.5.8	Data Retention and Disposal	34
3.5.9	Cross-Border Data Transfers.....	34
3.5.10	Data Governance Framework.....	34
3.6	Human Oversight and Responsibility	34

3.6.1	Ensuring Human Control.....	35
3.6.2	Accountability for AI Decisions/Harm.....	35
3.6.3	Human Oversight in Mitigating Risks of Bias and Discrimination.....	35
3.6.4	Ensuring Ethical Use of AI.....	36
3.6.5	Monitoring System Performance and Reliability.....	36
3.6.6	Training, Expertise and AI Literacy	36
3.6.7	Transparency and Communication	37
3.6.8	Continuous Improvement and Adaptation.....	37
3.7	Reliability, Safety, and Robustness	37
3.7.1	Design for Reliability	37
3.7.2	Thorough Testing and Validation.....	37
3.7.3	Risk Assessment and Mitigation	38
3.8	Compliance, Monitoring, and AI Governance.....	38
3.8.1	Adherence to Safety Standards.....	38
3.8.2	Continuous Monitoring and Maintenance.....	39
3.8.3	AI System Impact Assessment	39
3.8.4	AI Governance	40
3.8.5	External Audits and Reviews.....	41
3.9	Reporting and Metrics	42
3.9.1	Reliability Metrics and Reporting.....	42
3.10	Notification Requirements.....	43
3.10.1	Regulatory Notifications	43
3.10.2	Incident Reporting.....	43
3.10.3	Feedback and Continuous Improvement	44
3.11	Involvement of the MGA and MDIA.....	44
3.11.1	Regulatory Oversight	44
3.11.2	Collaboration with the MGA.....	45
3.11.3	Compliance Support	45
4	Closing Statement.....	45
	Annex A – Mapping of Charter Principles to EU AI Act and GDPR	46

Revision History

Revision No.	Issue Date	Description
1	18/09/2026	First issue.

Definitions

AI-Driven Automation	Automation which is not solely based on pre-defined rules and utilises an AI System to carry out the respective task, workflow or process. AI-driven automation falls within the scope of this Charter where it is used within gaming operations
AI System	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 ¹ .
AI Washing	AI Washing is a phenomenon where companies misrepresent or exaggerate their AI capabilities to enhance marketing appeal and gain a competitive advantage ² .
Automation	The use of technology to perform a task, workflow or process with limited or no human intervention. Automation may include simple workflow automation, rules-based automation, robotic process automation, or AI-driven automation.
Deep Fake	AI-generated or manipulated image, audio or video content that resembles existing persons, objects, places, entities or events and would falsely appear to a person to be authentic or truthful ³ .
Deployer	Any natural or legal person, public authority, agency, or other body using an AI System under its authority except where the AI system is used in the course of a personal non-professional activity ⁴ . In the context of this Charter, examples of scenarios where a Gaming Licensee may act as a Deployer include where it procures, integrates or uses an AI System developed or supplied by a third party, or uses an AI System developed in-house, within its gaming operations.
Gaming Licensee	An entity licensed by the MGA to offer a business-to-consumer (B2C) gaming service and/or to provide a business-to-business (B2B) critical gaming supply. For the sake of clarity, the entire Corporate Group having a Corporate Group Licence, shall be deemed a Licensee.
General-Purpose AI Model (GPAI)	An AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of

¹ [Article 3: Definitions | EU Artificial Intelligence Act – \(1\)](#)

² [AI Washing: A conceptual exploration | AMS Review | Springer Nature Link](#)

³ [Article 3: Definitions | EU Artificial Intelligence Act – \(60\)](#)

⁴ [Article 3: Definitions | EU Artificial Intelligence Act – \(4\)](#)

	the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development or prototyping activities before they are placed on the market ⁵ .
Provider	Any natural or legal person, public authority, agency, or other body that develops an AI System or general-purpose AI model or that has one developed on its behalf, and subsequently places it on the market or puts it into service under its own name or trademark, whether for payment or free of charge ⁶ . In the context of this Charter, examples of scenarios where a Gaming Licensee may act as a Provider include where an AI System is developed in-house and placed on the market, and when placing its name or trademark on white-labelled platforms to be placed on the market.
Rules-Based Automation / Robotic Process Automation (RPA)	Automation that follows predefined human-programmed rules, scripts or workflows to perform repetitive tasks. Such systems do not infer, learn, reason or generate predictions, recommendations, decisions or content from data in the manner contemplated by the definition of an AI System. On that basis, rules-based automation and robotic process automation, without AI functionality, generally fall outside the scope of this Charter.

⁵ [Article 3: Definitions | EU Artificial Intelligence Act – \(63\)](#)

⁶ [Article 3: Definitions | EU Artificial Intelligence Act – \(3\)](#)

1 Introduction

1.1 Overview

The Malta Gaming Authority (MGA) has regulated the gaming sector for more than two decades, consistently adapting its regulatory approach to reflect technological and market developments while remaining firmly grounded in principles of integrity, player protection, and sound governance. As a technology-neutral regulator, the MGA does not regulate specific tools or technologies, but rather the outcomes, risks, and behaviours associated with their use.

The Charter is developed within Malta's broader AI governance landscape, in which the Malta Digital Innovation Authority (MDIA) plays a central coordinating and enabling role. As the authority responsible for AI oversight and national policy alignment, the MDIA promotes the adoption of trustworthy, human-centric, and risk aware AI in line with both national priorities and European frameworks. The MDIA supports sector-specific initiatives such as this Charter, helping to ensure that innovation in gaming remains aligned with Malta's regulatory values, public interest objectives, and evolving expectations for responsible AI use.

Artificial Intelligence (AI) is increasingly embedded across gaming operations, supporting functions such as customer interaction, operational decision-making, risk management, and player protection. Its growing use presents both opportunities and challenges, particularly in a sector that combines innovation with a strong public-interest dimension.

In this context, the AI Gaming Charter has been developed as a sector-specific best-practice framework to guide the responsible, transparent, and ethical use of AI within gaming operations. It represents one of the first initiatives in Europe to articulate AI principles tailored specifically to the gambling sector, reflecting its operational realities and applicable regulatory expectations.

The Charter aligns with Malta's National AI Strategy and Vision, which promotes AI as an enabler of economic and societal advancement, as well as with Malta Vision 2050, where AI is identified as a critical accelerator for transparency, traceability, and sustainability. It also complements wider European efforts aimed at promoting trustworthy and human-centric AI.

While voluntary in nature, all Gaming Licensees are strongly encouraged to adopt the AI Gaming Charter as a means of demonstrating their commitment to the ethical use of AI and sound governance. Gaming Licensees choosing to adopt the Charter are encouraged to designate one or more responsible individuals – such as from compliance, data, technology or AI-related functions – to act as internal points of contact for the application of its principles, and support their participation in periodic AI literacy and refresher initiatives, ensuring that relevant knowledge and practices remain current and embedded within the organisation.

The Charter is designed to align with responsible gambling principles and support Gaming Licensees in navigating the complex ethical landscape associated with AI deployment. Its objective is to contribute

to a trustworthy and sustainable gaming environment, providing reassurance to players, stakeholders, and the wider public regarding the sector's integrity.

The Charter applies to the use of AI Systems across gaming operations, irrespective of the risk classification of the system or the use case involved. Whether AI is used for purposes of customer service, game development, or fraud detection, Gaming Licensees are encouraged to adhere to the principles and guidelines outlined in the Charter. This aims to ensure that ethical considerations are embedded throughout the AI lifecycle, from design and development to deployment, monitoring and evaluation.

Key aspects covered in this Charter include transparency, data governance, human oversight, non-discrimination, and stakeholder engagement. Each of these areas plays a critical role in endeavouring to ensure that AI Systems are developed and deployed in a manner that respects player rights, promotes fairness, and mitigates potential risks and harms.

A commitment to accountability and continuous improvement underpins the principles outlined in this Charter. Gaming Licensees that adopt this Charter are encouraged to establish appropriate governance arrangements, conduct impact assessments where relevant, and implement ongoing monitoring and evaluation mechanisms. By embracing transparency and accountability, Gaming Licensees can identify and address potential biases, errors, and negative impacts associated with AI use.

Finally, the Charter emphasises the importance of stakeholder engagement and collaboration throughout the AI deployment lifecycle. By engaging with regulators, industry representatives, and other relevant stakeholders, Gaming Licensees can help ensure that AI Systems are developed and deployed in a manner that reflects societal values and broader public expectations.

1.2 Purpose and Scope

The purpose of the AI Gaming Charter is to provide clear, practical best-practice guidance for the ethical and responsible use of AI within gaming operations. It is designed to encourage transparency, accountability and responsible innovation, while maintaining a strong focus on player protection and the integrity of gaming activities.

The Charter serves as a common reference point for Gaming Licensees seeking to develop or deploy AI Systems in a manner consistent with established ethical principles and regulatory expectations. For the avoidance of doubt, this Charter is voluntary in nature and does not create new legal or regulatory obligations for Gaming Licensees, nor does it modify, expand or replace any obligations that already apply under existing national or EU law, including the EU Artificial Intelligence Act (the 'EU AI Act')⁷, the EU Cyber Resilience Act⁸, the General Data Protection Regulation (GDPR)⁹, the Data Act¹⁰, and where

⁷ [Regulation - EU - 2024/1689 - EN - EUR-Lex](#)

⁸ [Regulation - EU - 2024/2847 - EN - EUR-Lex](#)

⁹ [Regulation - EU - 2016/679 - EN - EUR-Lex](#)

¹⁰ [Regulation - EU - 2023/2854 - EN - EUR-Lex](#)

applicable, international instruments such as the Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law.

It is intended to serve as practical guidance to support Gaming Licensees in identifying, assessing and managing AI-related considerations in a manner consistent with applicable legal requirements, recognised good practice, and relevant international standards and frameworks. While this Charter does not include the EU Cyber Resilience Act within its scope to focus on the AI System dimension, Gaming Licensees are recommended to look into the Cyber Resilience Act in any case, since AI Systems may qualify as products with digital elements as defined under the mentioned Act. There are also some requirements emanating from the Cyber Resilience Act which intersect with the EU AI Act, such as cybersecurity requirements for High-Risk AI Systems.

Each Gaming Licensee remains responsible for determining which AI Systems it develops, procures, deploys or otherwise uses within its operations, and for assessing the role it performs in relation to those systems, including whether it acts as a Provider, Deployer, and/or other operator under the EU AI Act. The applicable compliance requirements will depend on the relevant role, the nature and intended purpose of the respective AI System, and its risk classification under the EU AI Act.

Nothing in this Charter should be interpreted as imposing additional contractual, reporting or operational obligations on Gaming Licensees, including in respect of AI Systems or contracts entered into before the Charter's publication. Where Gaming Licensees choose to align with this Charter, they do so on a voluntary basis and without prejudice to their obligations under applicable law.

This Charter aligns with relevant European initiatives, including the European Commission's AI Pact¹¹ which promotes early preparedness and responsible implementation of the EU AI Act and related AI governance measures.

This Charter's scope is comprehensive, covering various aspects of AI deployment to ensure that ethical considerations are embedded throughout the entire AI lifecycle.

For the purposes of this Charter, the distinction between AI Systems and other forms of automation is informed by the definition of an AI System under Article 3(1)¹² of the EU AI Act. This Charter applies to the use of AI Systems across the Gaming Licensee's operations, including internal or back-office uses and customer-facing or player-impacting uses. The fact that an AI System is used for internal, operational or administrative purposes does not place it outside the scope of the Charter. The Charter does not apply to automation, rules-based workflows, robotic process automation or other implementations that do not meet the definition of an AI System. In general, automation that merely executes predefined rules or workflows, without using an AI System should not be treated as AI use for the purposes of this Charter.

¹¹ [AI Pact | Shaping Europe's digital future](#)

¹² [Article 3: Definitions | EU Artificial Intelligence Act](#)

The application of the Charter should be proportionate to the nature, purpose, scale and risk profile of the AI System concerned. Lower-impact internal tools, such as coding assistants, productivity tools or administrative support systems, may require lighter governance, documentation and oversight. AI Systems that directly or indirectly affect players, player protection, gaming integrity, compliance decisions or regulated outcomes should be subject to stronger safeguards.

This Charter addresses the key areas and aims to promote responsible AI deployment practices that respect player rights, promote fairness, and mitigate potential risks and harms associated with AI Systems. Gaming Licensees are encouraged to voluntarily adopt and implement the principles and requirements outlined in this Charter to ensure the responsible development and deployment of AI Systems, contributing to advancing an accountable and sustainable gaming ecosystem.

To support clarity on the relationship between this Charter and existing legal obligations, Annex A provides a high-level mapping of the Charter principles against relevant provisions of the EU AI Act and the GDPR. The mapping distinguishes between areas where the Charter reflects or supports existing legal requirements and areas where it provides sector-specific good practice for the gaming sector.

The drafting language used in this Charter is intended to distinguish between existing legal obligations and voluntary sector good practice. Where a requirement already applies under applicable law, the Charter refers to it using terms such as “in accordance with applicable law”, “where required by applicable law”, or equivalent wording. Such references do not create new obligations under the Charter, but signpost obligations that may already apply independently under the relevant legal framework.

Where the Charter sets out sector-specific good practice, the language used is intended to be voluntary and proportionate, including terms such as “should”, “are encouraged to”, “may”, or equivalent wording. These provisions are intended to support responsible AI adoption and governance in the gaming sector, without converting good practice into binding legal requirements.

Gaming Licensees should therefore read each provision in its proper context, distinguishing between obligations that arise under applicable law and voluntary Charter guidance that supports proportionate, ethical and responsible AI use.

1.3 Regulatory Perspective and Responsible Use of AI

Consistent with its long-standing technology-neutral approach, the MGA encourages Gaming Licensees to adopt AI Systems in ways that support and enhance existing governance, compliance, and player-protection frameworks, rather than replace established controls or human judgement.

For the avoidance of doubt, this Charter is not intended to discourage or restrict AI use or automation as such. Many operational tasks may be automated through predefined workflows, scripts, robotic process automation or traditional rules-based software without constituting AI use. Where such systems do not infer, learn, reason, or generate outputs such as predictions, content, recommendations or decisions in the manner contemplated by the definition of an AI System under the EU AI Act, they generally fall outside the scope of this Charter.

Where automation is AI-driven, however, or where it materially relies on an AI System to support or influence outcomes within gaming operations, the Charter should apply in a proportionate manner. The relevant safeguards should reflect the nature, purpose and risk profile of the AI System, including whether it affects players, player protection, gaming integrity, compliance decisions or other regulated outcomes.

While this Charter remains voluntary and does not introduce new obligations, AI Systems may nevertheless play a constructive role in strengthening Gaming Licensees' ability to safeguard players and uphold high standards of integrity.

Appropriately designed AI tools may assist Licensees in identifying patterns or indicators associated with potential risks, including misleading or duplicate account activity, behaviours linked to problematic gambling, or indicators of player vulnerability. Subject to applicable legal and ethical considerations, such tools may also support age-related monitoring, intervention strategies, and the prevention of underage gambling.

AI Systems can further contribute to more consistent and timely internal processes relating to behavioural monitoring, escalation pathways, and player interaction workflows, thereby supplementing existing detection mechanisms and responsible-gaming frameworks. Their use may assist staff decision-making by providing additional insights, while preserving the role of human judgment and accountability.

The integration of AI technologies should therefore be viewed as a means of supporting and strengthening existing governance structures, rather than replacing established procedures, professional expertise or responsible gambling values. Human oversight, proportionality, and ethical judgement remain essential components of any AI-enabled gaming operation.

2 Study on AI Use and Adoption in Gaming

To inform the development of the AI Gaming Charter, the Authority carried out a study on the use and adoption of AI within gaming operations licensed in Malta. The main aim of this study was to gain visibility into how AI is currently being deployed across the sector and how its use may evolve over time. The insights gathered have supported both the development of this Charter and future guidance that is evidence based, proportionate, and aligned with operational realities. The following main objectives were targeted as part of this study:

- Identify the main AI applications and use cases in the Gaming sector;
- Evaluate the gaming industry's regulatory readiness, ensuring Gaming Licensees understand their obligations under the EU AI Act;
- Alignment with the EU AI Act to demonstrate proactive compliance;
- Measure the impact of AI ethics on main Gaming application areas, such as responsible gambling, fraud detection, and player trust;

- Provide recommendations for Gaming Licensees on how the Charter can be used to prevent certain risks whilst maximising the benefits of using AI for their intended objectives.

The study was conducted through a comprehensive survey distributed to all MGA Gaming Licensees, with participation on a voluntary basis. In addition, a number of in-depth interviews were carried out with selected Gaming Licensees to gain more detailed insights into the practical use and expectations of AI within the industry.

While the number of responses represented a limited subset of the sector, the inputs received provided useful indicative insights into current practices, levels of maturity, and emerging themes. These insights have informed the focus areas and principles reflected in this Charter, while recognising that AI adoption within the gaming sector continues to evolve and that ongoing engagement remains important.

The main results and key insights derived from this study are outlined in the sections below, providing an overview of the findings and their implications for the industry.

2.1 Use Cases of AI in the Gaming Sector

AI adoption across the gaming sector is still uneven and generally selective rather than widespread. The most mature use cases are in operational optimisation and data analytics, customer support chatbots, and, to a lesser extent, recommendation engines, player profiling, fraud detection and AML, and responsible gambling behavioural modelling. These are the areas where respondents most often reported fully deployed systems, or activity beyond the proof-of-concept stage.

The data also indicates that a number of organisations are using AI at Minimum Viable Product stage, particularly in areas such as responsible gambling behavioural modelling, fraud detection and AML, recommendation engines, player onboarding and KYC, payments, human resources, and player acquisition. This suggests that, across the sector, some Gaming Licensees have moved beyond early experimentation but are still testing practical deployment before wider rollout.

By contrast, adoption appears far more limited in dynamic odds setting and risk management, regulatory reporting, player acquisition, HR-related uses, player onboarding and KYC, payments, and AI-powered live casino applications. In these cases, most responses were either not in use or not applicable, indicating that AI is being adopted more readily in support, monitoring, and efficiency functions than in more commercially sensitive or player-facing decision areas.

The additional systems mentioned by respondents also show that some organisations are extending AI beyond core gaming functions. Examples included software development tools such as Claude Code¹³ and GitHub Copilot¹⁴, creative tools for image and video generation, and compliance tools used to monitor regulatory change and automate workflows. This points to a broader pattern in which AI is being

¹³ [Claude Code by Anthropic | AI Coding Agent, Terminal, IDE](#)

¹⁴ [GitHub Copilot · Your AI pair programmer · GitHub](#)

used not only in gaming operations, but also in software development, compliance, and creative functions.

Throughout the co-creative interviews carried out with a number of Gaming Licensees, the following key use cases where AI is being used were identified:

- **Customer Experience and Personalisation:** AI is widely deployed to enhance player engagement, offer tailored recommendations, and provide efficient customer service via chatbots and automated help desks.
- **Fraud Detection and Responsible Gambling:** Gaming Licensees employ AI to identify suspicious behaviours, flag potential problem gambling, and uphold anti-money laundering protocols. Machine learning models facilitate real-time risk monitoring and player protection.
- **Game Development and Testing:** AI assists in designing new games, refining game mechanics, and automating quality assurance processes, resulting in faster development cycles and improved player satisfaction.
- **Marketing and Analytics:** AI-driven insights inform targeted marketing campaigns, player segmentation, and performance analytics, enabling better decision-making and resource allocation.
- **Operational Efficiency:** Automation of repetitive tasks, dynamic pricing, and predictive maintenance are other notable AI use cases, streamlining operations and reducing costs.
- **Software Development:** AI is increasingly being used within software development practices. Gaming Licensees cite the adoption of AI-powered tools for code generation, bug detection, and automating testing procedures. These applications have accelerated development timelines, improved code quality, and enabled teams to focus on innovation, rather than routine tasks.

2.2 AI Literacy and Awareness

There is a pronounced need for greater AI literacy and awareness throughout organisations. It is essential that staff across all departments, not just compliance or technical teams, understand the capabilities, risks, and implications of AI to enable responsible adoption. Targeted training and communication initiatives are frequently recommended as good practice.

2.3 AI Strategy

The findings indicate that formal AI strategy is still developing across much of the participating licensees. Only a minority of organisations reported having a formal AI strategy or roadmap, while a larger group said one was still in development. Among those with a strategy, alignment with business objectives was generally described positively, suggesting that where strategy exists, it is increasingly being linked to wider organisational goals.

Some organisations described AI strategy as already integrated into broader business strategy, while others referred to governance frameworks nearing deployment or AI being incorporated into the wider

technology strategy rather than treated as a separate standalone area. Other responses showed that, in some cases, AI capability is still being established, and strategic direction remains in progress.

Furthermore, several respondents said they had considered the EU AI Act and its requirements when drafting their AI strategy, and some described its influence on their acceptable use policies, ethical frameworks, project initiation processes, and risk assessments. This suggests that regulatory obligations are acting as an important driver of strategic development.

2.4 EU AI Act and Regulatory Considerations

There is also evidence of growing awareness of the EU AI Act. A meaningful proportion of respondents said they had classified their AI Systems according to the Act's risk categories, while others said they had done so partially. The comments suggest that some organisations are already aligning internal governance and classification processes directly with the Act, particularly through risk assessment, documentation, and internal policy controls.

2.5 Risk and Operational Readiness

Technical readiness for AI appears reasonably positive overall. Most respondents rated their infrastructure as mature or adequate, and the comments referred to strong cloud setups, low technical debt, and AI integration into development, Quality Assurance (QA), infrastructure, and security practices.

2.6 Data and AI Governance

Data governance appears more developed in some areas than others. A majority of respondents reported frameworks covering data ownership, data quality, access management, and retention or deletion policies. However, fewer reported having data lineage and documentation in place, which suggests that while core governance controls may exist, deeper traceability and documentation are less consistently embedded.

The data from the questionnaire shows a more mixed position on transparency and accountability. Only a small number of organisations said they disclose AI usage to customers and secure explicit consent for automated processing, while disclosure to employees was more common. At the same time, many respondents indicated that they assess training data for representativeness, accuracy, and bias, maintain documentation on how AI Systems function, and keep logs of system versions and decision outcomes.

Formal AI governance maturity remains uneven. Some organisations already have a dedicated AI governance structure, while many others said one is still in progress. Regular bias or fairness assessments, impact measurement, AI literacy training, and AI-specific risk and incident processes appear less developed. In particular, only a small number of respondents reported having a fully established AI risk assessment process or incident response plan.

The organisational culture is generally described as receptive, with most respondents describing their culture as proactive, innovative, or moderately supportive. However, human oversight remains a defining feature of current practice. Most respondents stated that human oversight is always applied to AI-based decisions, showing that AI is generally being used to support decision-making rather than replace human judgement in sensitive areas.

2.7 Business-to-Consumer (B2C) Licence Holders

Among B2C licence holders, contestability and human review are present in some cases, but not consistently across all organisations. Several respondents said players can request human review of AI-generated decisions or explanations, while others said this is only available on request or not available at all. However, the accompanying comments clarify that some organisations do not use AI to make final decisions on account restrictions or responsible gambling flags, with these decisions remaining subject to manual review by dedicated teams.

The respondent data also suggests that evaluation of whether AI-driven personalisation increases risky or vulnerable behaviours is still limited. Only a small number of respondents said they do this, while others said partially or were unsure. Where safeguards were described, they centred on accurate data, responsible gambling overrides, human-in-the-loop oversight, monitoring, and internal governance measures, rather than on a clearly mature or standardised evaluation framework.

2.8 Business-to-Business (B2B) Licence Holders

For B2B licence holders, the findings suggest relatively limited transparency to clients regarding how AI Systems function. Only one respondent said such documentation is provided, while others said it is not. The responses also indicate that some suppliers distinguish between AI and rule-based systems, particularly where alert engines are programmed by humans and supported by human decision-making.

While many B2B licence holders may act primarily as Deployers of AI Systems developed or supplied by third parties, some B2B licence holders may also act as Providers where they develop AI Systems, or have them developed on their behalf, and place them on the market or put them into service under their own name or trademark. The applicable responsibilities should therefore be assessed on a case-by-case basis by reference to the role performed by the Gaming Licensee under the EU AI Act.

Support to clients in assessing risk, fairness, impact, or compliance also appears limited in this group. Where suspicious activity is detected, the information may first be reviewed by B2B staff before being passed to clients, which is important in such cases since the AI is designed with a human in the loop to ensure that the activity is truly suspicious.

2.9 B2B and B2C Licence Holders

Among organisations holding both B2B and B2C licences, the responses suggest a mixed and cautious position towards implementation and use of AI. Contestability of AI-generated decisions does not

appear consistently established, and responses on personalisation show that some organisations are still streamlining their approach, while others state that no AI tools are in place.

There is some evidence of more structured thinking in this group, particularly where respondents referred to considering the purpose of the model, the data used, potential impacts on player rights and fairness, human oversight, and possible bias or discriminatory outcomes. However, transparency tools, client support, and communication processes around suspicious activity remain inconsistent across respondents.

3 The Charter

The findings of the study on AI use and adoption highlight a sector that is actively engaging with AI, but where levels of maturity, governance readiness, and risk awareness remain uneven. While AI is already being deployed across a range of operational and player-facing functions, the study also identified common challenges relating to transparency, accountability, data governance, human oversight, and the management of potential risks to players.

Against this backdrop, the principles set out in the Charter are intended to respond directly to the practices, gaps, and emerging risks observed across the sector. They are designed to provide Gaming Licensees with a practical reference point for embedding ethical and responsible considerations into the design, deployment, and use of AI Systems, taking into account the varying degrees of sophistication and resources across organisations.

The sections that follow set out these principles in detail, providing a structured framework through which Gaming Licensees can assess their current practices, identify areas for improvement, and demonstrate a commitment to responsible AI use within gaming operations.

3.1 Risk-Based Application of the Charter

This Charter applies to AI Systems used across gaming operations and is intended to be applied in a proportionate, risk-based manner. The level of documentation, testing, governance, transparency, human oversight and monitoring applied to an AI System should reflect the nature of the use case and its potential impact on players, player rights, gaming integrity or regulated outcomes.

It is suggested that Gaming Licensees compile an inventory of AI Systems, identify their role in the supply chain (such as Provider, Deployer, importer, distributor), and the risk classification of their AI Systems. Compliance tools are available on the MDIA's dedicated webpage¹⁵ for stakeholders wishing to identify their role in the supply chain and obtain an initial self-assessment of the risk classification of their AI System. These tools include an EU AI Act Compliance Checker developed by the EU AI Office¹⁶.

¹⁵ [Compliance Tools – MDIA](#)

¹⁶ [EU AI Act Compliance Checker – EU AI Office](#)

For each AI System, Gaming Licensees should first confirm that the AI System falls within the scope of the EU AI Act. The definition of AI System in this Charter mirrors the definition established in Article 3(1) of the EU AI Act, and should be used to confirm whether a system is in scope, before consulting other criteria under Article 2 of the EU AI Act. Dedicated harmonised guidelines on the AI System definition have been published by the EU AI Office¹⁷.

Once an AI System is confirmed to be within scope, it is crucial that Gaming Licensees assess their AI Systems to ensure that they are not classified or can lead to any prohibited AI practices under Article 5 of the EU AI Act. AI Systems with use-cases under the mentioned Article 5 are banned from being placed on the market, put into service, or used within the EU, due to their considerable risks to health, safety, and fundamental rights, and the possible harms that may be inflicted. Further to the AI System definition guidelines, harmonised guidance is also available on prohibited AI practices¹⁸.

The risk-based approach does not exclude AI Systems with other risk categories – (such as High-Risk AI Systems, and AI Systems which trigger transparency obligations under the EU AI Act) – from the scope of the Charter; rather, it calibrates the level of effort, documentation and oversight expected for each AI System.

For the purpose of applying this Charter in a proportionate manner, Gaming Licensees are encouraged to distinguish between two broad tiers of AI use:

- Lower-impact AI Systems, where the AI System is used for internal, operational, administrative or other low-impact purposes and is unlikely to materially affect players, player rights, gaming integrity or regulated outcomes.
- Higher-impact AI Systems where the AI System is used in, or materially influences, player-facing, responsible gambling, fraud, AML, KYC, account restriction, eligibility, intervention, or other compliance-sensitive decision-making contexts.

This practical distinction is intended solely to support the proportionate application of the Charter and does not replace, modify or determine any classification that may apply under the EU AI Act. The designation of an AI System as a High-Risk AI System under the EU AI Act is governed exclusively by the provisions of that Regulation and any related guidance issued by the competent authorities and the EU AI Office.¹⁹

The impact should be based on the AI System's proximity to the player and its potential effect on regulated outcomes, rather than solely on whether a human makes the final decision. Human oversight may help manage risk, but it should not automatically reduce the applicable level of scrutiny where the AI System materially influences the outcome.

¹⁷ [AI System Definition Guidelines – EU AI Office](#)

¹⁸ [Prohibited AI Practices Guidelines – EU AI Office](#)

¹⁹ [Draft High-Risk AI System Guidelines – EU AI Office](#)

Higher-impact AI Systems should therefore be subject to stronger safeguards, such as clearer documentation, testing and validation, defined human oversight, ongoing monitoring, escalation procedures and appropriate records. Lower-impact AI Systems may be subject to lighter-touch controls. In either case, the implemented safeguards must be, at a minimum, aligned with the requirements of that AI System's risk classification under the EU AI Act.

Each Gaming Licensee remains responsible for assessing its own AI Systems by reference to their specific design, purpose, deployment context and applicable legal obligations.

3.2 Transparency, Accountability and Explainability

Transparency and accountability are fundamental principles in the deployment of AI Systems. These principles ensure that stakeholders have access to clear information about AI Systems used within gaming operations. This section delves into the importance of transparency and accountability, and how Gaming Licensees can build trust with players, users, and stakeholders, and mitigate potential risks arising from AI-related opacity. Any documentation, communication, or other information provided for transparency purposes should also be made available in a manner that is clear, understandable and accessible to its intended audience, in accordance with applicable accessibility, consumer protection and equality requirements. This Charter is not intended to create standalone product accessibility requirements beyond those arising under existing legal frameworks, but rather to help relevant users and stakeholders understand AI-related transparency and explainability information.

Transparency and explainability should be understood as the provision of meaningful, human-readable information about the role, purpose, functionality, limitations and effects of AI Systems. They do not require disclosure of source code, detailed model architecture, proprietary information, trade secrets, or sensitive fraud, AML or security logic, where such disclosure would undermine intellectual property rights, system security, fraud prevention, AML controls, or the integrity of gaming operations.

3.2.1 Transparency in AI Design and Usage

3.2.1.1 Transparency in AI Design and Interactions

AI Systems should be designed by Gaming Licensees in a manner that enables players and users to be clearly informed when they are engaging with AI Systems, at latest from the point of first interaction or exposure, including when such individuals are exposed to synthetic or manipulated content.

On the contrary, AI Washing, a marketing practice used by businesses to overstate the actual use, role or capabilities of AI within their services or products, should be avoided. This includes presenting conventional software as AI, overstating the autonomy or sophistication of systems, or marketing AI as a guarantee of fairness, accuracy, or safety.

Gaming Licensees should communicate factually and proportionately about the role of AI Systems within their services, avoiding inflated, misleading or unverifiable claims.

3.2.1.2 Transparency in AI Usage

Gaming Licensees should provide appropriate disclosures on how AI Systems impact relevant decision-making processes and how they are used in their services or supplies, where such disclosure is required by applicable law or supports responsible and transparent AI use.

Transparency in decision-making goes beyond identifying that AI is used; it involves explaining, in clear and accessible terms, the role AI plays in shaping relevant outcomes and interactions.

Gaming Licensees should ensure that employees, players and end-users can identify when they are exposed to AI decision-making processes and to content that has been generated or manipulated by AI Systems.

3.2.1.3 Transparency in AI-Generated Content

AI-generated text and other synthetic material must be assessed in line with the transparency obligations under Article 50 of the EU AI Act, taking into account the different obligations that apply to Providers and Deployers. In particular, Providers of AI Systems that generate synthetic audio, image, video or text content are required under Article 50(2) to ensure that outputs are marked in a machine-readable format and detectable as artificially generated or manipulated, as far as technically feasible. This Provider-level marking obligation should not be read as requiring every item of synthetic material to carry a visible user-facing label. By contrast, where Gaming Licensees act as Deployers and use AI Systems to generate or manipulate content that qualifies as a Deep Fake, or publish AI-generated or manipulated text for the purpose of informing the public on matters of public interest, Article 50(4) requires clear disclosure that the content has been artificially generated or manipulated, subject to the context and any applicable exceptions.

In this context, the European Commission has published the Code of Practice on Transparency of AI-Generated Content²⁰, which supports compliance with the Article 50 obligations on marking and labelling AI-generated or manipulated content. The Code distinguishes between Provider commitments under Article 50(2), focused on machine-readable marking and detection of AI-generated or manipulated audio, image, video and text content, and Deployer commitments under Article 50(4), focused on labelling Deep Fakes and certain AI-generated or manipulated text made available to the public. Gaming Licensees should therefore assess whether they are acting as Providers or Deployers in the relevant use case, and apply transparency measures that are proportionate to that role, the type of content, the intended audience, and the risk that the content could mislead or be mistaken for authentic material.

Further guidance has also been made available by the EU AI Office with the publication of the Guidelines on Transparency of AI-Generated Content²¹. The Guidelines on Transparency of AI-Generated Content

²⁰ <https://ec.europa.eu/newsroom/dae/redirection/document/129555>

²¹ <https://digital-strategy.ec.europa.eu/en/policies/guidelines-transparency-ai-generated-content>

support compliance with Article 50 by clarifying the scope of the obligations, providing examples of Providers and Deployers of AI Systems which require transparency measures to be implemented, and outlining how compliance with transparency requirements can be demonstrated. For example, a Gaming Licensee may be acting as a Provider where it develops, substantially modifies, fine-tunes or places into service under its own name or trademark an AI System that generates promotional text, images, videos, chatbot responses or other synthetic content for use in gaming operations. In such cases, Provider-level obligations may include ensuring that generated or manipulated outputs are marked in a machine-readable format and detectable as artificially generated or manipulated, where technically feasible. By contrast, a Gaming Licensee may be acting as a Deployer where it uses an inhouse or third-party AI tool, Software as a Service (SaaS) solution, Application Programming Interface (API) or licensed model to generate or manipulate content for publication, player communication, marketing, responsible gambling messaging or customer support. In those cases, the Gaming Licensee should focus on whether visible disclosure or labelling is required under Article 50(4), taking into account the type of content, the purpose of publication, the audience, and the risk that the content may mislead or be mistaken for authentic material.

3.2.1.4 Transparency of Data

Where Gaming Licensees develop, train, re-train, or fine-tune AI Systems, they should maintain clear documentation of the datasets used for training, validation, and testing. This documentation should include information on the origin and type of the data, as well as any measures taken to ensure data quality and representativeness.

3.2.1.5 Transparency for General-Purpose AI (GPAI) Systems

Gaming Licensees using GPAI models or systems should maintain documentation that is reasonably available to them and proportionate to their role in the AI value chain. Where a Gaming Licensee deploys a third-party or licensed GPAI model, such documentation may include model cards, terms of use, intended-use descriptions, known limitations, relevant Provider documentation, supplier attestations, and internal records describing how the model is used within gaming operations.

Where applicable, Providers of GPAI models and systems should have regard to, and comply with, any applicable obligations or commitments under the *General-Purpose AI Code of Practice – Transparency Chapter*, as published by the European Commission²². This reference should be understood as applying to GPAI model Providers where relevant under the EU AI Act, including Article 53²³, and should not be interpreted as creating equivalent model-level documentation obligations for Gaming Licensees acting solely as Deployers of third-party or licensed GPAI models.

Detailed model-level information, including training compute, computing power measured in Floating Point Operations Per Second (FLOPs), model architecture and training methodology, should not be

²² [The General-Purpose AI Code of Practice | Shaping Europe's digital future](#)

²³ [Article 53: Obligations for Providers of General-Purpose AI Models | EU Artificial Intelligence Act](#)

expected from a Gaming Licensee acting solely as a Deployer where such information is not disclosed by the Provider. Under the EU AI Act, including Article 53, such obligations sit primarily with GPAI model Providers and, where relevant, with entities that develop, train, fine-tune, substantially modify or place such models on the market or put them into service under their own name or trademark.

Where a Gaming Licensee acts as a Provider or fine-tuner of a GPAI model or system, it should maintain appropriate model-level documentation proportionate to its role and in line with applicable EU AI Act obligations. This may include information on model architecture, training or fine-tuning approach, relevant technical characteristics, model licensing, intended purpose, capabilities and limitations, where applicable.

3.2.1.6 Transparency of AI Deployment

Documentation should be maintained describing the purposes for which AI Systems are deployed, the business functions they support, and the decision-making processes they impact. This should include the scope of use (e.g., Anti-Money Laundering (AML) monitoring, responsible gambling, fraud detection, marketing activities), as well as any limitations identified during testing or monitoring.

3.2.1.7 Logs for Transparency

Gaming Licensees should maintain logs and supporting records that are proportionate to the nature, purpose and risk profile of the AI System. Logging should support traceability, accountability, monitoring and review, without creating unnecessary data collection or retention.

As a minimum, logs should record, where relevant: the AI System or model version used; the date and time of relevant system activity; the type of input or event processed; the output, recommendation, flag or decision-support result generated; material errors, failures or exceptions; human intervention, override or review; and significant configuration or model changes.

For lower-impact AI Systems, sampled, aggregated, hashed, pseudonymised or risk-tiered logging may be appropriate, provided that the logs remain sufficient to support proportionate oversight, troubleshooting and accountability.

For higher-impact AI Systems, including systems that materially affect players, player protection, gaming integrity, compliance decisions or regulated outcomes, more detailed logging and traceability should be maintained. Where the AI System qualifies as High-Risk under the EU AI Act, logging must be aligned with applicable legal requirements, including those relating to automatic recording of events.

Higher-impact AI Systems may also require supporting information beyond operational logs. This may include, where applicable, information on training, validation and testing data, fine-tuning data, foundational or general-purpose models used, model versioning, system architecture, intended purpose, performance limitations, risk controls and human oversight arrangements. This supporting information should be treated separately from the operational logging baseline.

Where logs include personal data, Gaming Licensees must comply with applicable data protection requirements, including the GDPR principles of data minimisation, purpose limitation, storage limitation, integrity and confidentiality. Retention periods for logs should be defined in advance and should be proportionate to the AI System's risk, purpose and applicable legal or regulatory requirements.

3.2.1.8 Risk and Impact Assessments

Documentation of risks identified in the use of AI Systems should be maintained, including potential impacts especially regarding impact on players, fairness, and compliance with applicable law. This should include outcomes of risk assessments and mitigation measures arrived therefrom.

Risk assessments should also consider the inherent limitations of AI Systems, including the risk of inaccurate outputs, such as false positives (e.g. flagging betting activity that does not raise genuine concerns) and false negatives (e.g. failing to detect suspicious betting patterns), together with the associated operational, regulatory and integrity implications.

Gaming Licensees should define and maintain appropriate human-oversight mechanisms to review, challenge and validate AI-driven outputs, particularly where such outputs may inform or lead to decisions affecting players or the integrity of gaming activities.

Relevant risk assessment tools may include, where applicable, Data Protection Impact Assessments, Artificial Intelligence Impact Assessments, and Fundamental Rights Impact Assessments.

3.2.1.9 Human Oversight Arrangements

Documentation should be maintained setting out how human oversight is embedded in the use of AI Systems, including the actual human oversight process, intervention and escalation pathways, and any override mechanisms in place, where relevant.

3.2.2 Clear Communication

Gaming Licensees should provide clear and accessible information about their AI Systems, detailing their capabilities, limitations, and potential impacts in a manner accessible to its intended audience.

Such information should be communicated in a way that enables all relevant stakeholders, including players, regulators, and other affected parties, to understand how AI Systems may influence decisions or outcomes, thereby promoting transparency and supporting informed decision-making. Overly technical language should be avoided where possible.

Gaming Licensees should also ensure that the use of AI Systems, including those provided by third-party vendors, remains subject to appropriate internal oversight, validation, and accountability mechanisms, particularly where such systems may impact players or decision-making processes.

Where AI Systems are used to support player-facing functions (e.g., behavioural monitoring or intervention tools), Gaming Licensees should take reasonable steps to ensure that such systems are

functioning as intended, are proportionate to their purpose, and are effective. Any identified limitations should be clearly understood and appropriately managed.

In this context, responsibilities should be clearly delineated between AI Providers and Gaming Licensees, with Providers supporting transparency on system design and performance, and Gaming Licensees ensuring the appropriate and responsible use of such systems within their operations.

3.2.3 Risk Communication

Gaming Licensees should openly and transparently communicate AI Systems' potential risks and uncertainties. This includes disclosing known limitations, biases, and uncertainties in AI Systems to enable stakeholders to make informed decisions about their use and mitigate any potential harms.

3.2.4 Algorithmic Transparency

Transparency regarding algorithmic decision-making supports trust, fosters accountability, and enables the MGA, affected individuals and those responsible for the review and oversight of AI Systems within Gaming Licensees to better understand and address expectations, potential biases, or errors associated with AI Systems, where appropriate and subject to the safeguards set out below.

Gaming Licensees should maintain appropriate internal records explaining the purpose, role, main inputs, limitations, risk controls and decision-support function of AI algorithms used within gaming operations, proportionate to the AI System's risk and impact.

Such records should be made available, upon request, to the MGA and/or the MDIA, as appropriate to their respective regulatory functions, and subject to applicable legal and regulatory requirements.

Where an AI-supported decision materially affects an individual, Gaming Licensees should provide the affected individual with a meaningful, human-readable explanation of the outcome, the role played by the AI System, and any available channels for human review, challenge or redress, where required by applicable law or appropriate in the circumstances. Deployers of High-Risk AI Systems shall inform workers' representatives and the affected workers that they will be subject to the use of the High-Risk AI System.

Algorithmic transparency does not require public disclosure of source code, detailed model architecture, proprietary information, trade secrets or sensitive fraud, AML, security or risk-scoring logic. Public disclosure should not be required where it would undermine intellectual property rights, system security, fraud prevention, AML controls, responsible gambling safeguards, or the integrity of gaming operations.

This approach is intended to support transparency and accountability in line with applicable requirements under the EU AI Act and GDPR, including Deployer responsibilities, transparency obligations and safeguards for automated decision-making where relevant.

3.2.5 Error Reporting and Correction

Gaming Licensees should establish proportionate processes for promptly investigating, reporting, and correcting errors or inaccuracies in AI Systems.

Transparent error reporting and corrective actions enhance accountability, facilitate learning from incidents, and promote continuous improvement in AI deployment practices. Where relevant, these mechanisms should operate in coordination with existing redress processes and incident reporting obligations.

3.2.6 Transparency in Open-Source AI

Gaming Licensees should recognise the importance of transparency and accountability in deploying open-source AI Systems and commit to providing clear and comprehensive information about their AI models and usage of open-source. It is also pertinent to note that AI Systems released under free and open-source licences may still be subject to requirements of the AI Act, in particular where they are considered as High-Risk, or AI Systems with prohibited AI practices or triggering transparency obligations.

Furthermore, practices commonly referred to as "open-washing", whereby AI Systems are presented as open-source without full disclosure of critical components, should be avoided.

3.2.7 Explainability

Gaming Licensees should apply explainability in a proportionate and context-specific manner, particularly where AI Systems support or influence data-driven decision-making in areas such as fraud detection, AML, responsible gambling, behavioural monitoring, player interaction, eligibility, account restriction or other player-impacting processes.

Explainability should enable affected persons, internal reviewers, regulators and other relevant stakeholders to understand, in clear and accessible terms, the role played by the AI System, the main inputs or factors considered, the rationale behind the AI-generated output or outcome, and how that output contributed to the relevant decision-making process.

For this purpose, Gaming Licensees should distinguish between system-level and decision-level explanations. System-level explainability refers to clear information about the AI System's purpose, role, main inputs, limitations and the type of outcomes it may support or influence. Decision-level explainability refers to a meaningful, human-readable explanation of an AI-supported outcome in a particular case, where that outcome materially affects a player or user.

For inherently complex or opaque AI Systems, Gaming Licensees should focus on functional explainability. This means providing enough information to understand the main factors that contributed to the outcome, the role played by the AI System, and the available channels for human

review, challenge or redress, where applicable, without requiring disclosure of source code, detailed model architecture, proprietary information, trade secrets, or sensitive fraud, AML or security logic.

Where AI Systems are used for automated decision-making or profiling involving personal data, explainability, human review and contestability arrangements should be aligned with applicable GDPR requirements.

3.3 Fairness and Non-Discrimination

Ensuring fairness and non-discrimination is essential in the development and deployment of AI Systems, both as a matter of ethical responsibility and in alignment with applicable legal and regulatory obligations. Gaming Licensees are expected to ensure that the use of AI Systems does not result in unlawful discrimination or undermine fundamental rights, including those protected under applicable EU and national legislation.

In this context, promoting fairness and mitigating bias are key to supporting compliance, upholding equality, and fostering inclusive and equitable gaming experiences for all users.

3.3.1 Algorithmic/Decision Making Fairness

Gaming Licensees should ensure that AI-driven decision-making processes are fair and equitable, and do not result in unjustified discrimination or disparate impact on individuals or groups. This includes implementing measures to prevent discriminatory outcomes in areas such as player profiling, marketing, and customer service, where AI Systems are increasingly being deployed.

Furthermore, Gaming Licensees should prioritise AI decision-making fairness to ensure that AI Systems do not perpetuate or amplify biases against individuals or groups based on protected characteristics such as race, gender, ethnicity, or disability.

Gaming Licensees must comply with applicable non-discrimination and fundamental-rights obligations. They should implement measures to detect and mitigate bias in AI algorithms, including techniques such as bias testing, fairness-aware training, and algorithmic debiasing techniques. For higher-impact AI Systems that materially influence player-facing or player-impacting decisions, Gaming Licensees should document which fairness metric or fairness objective they have adopted for the relevant decision or outcome, and why that metric is appropriate in the circumstances. This is important because different fairness metrics may be incompatible or may produce different results depending on the AI System, use case, data, model design and operational context. The focus should therefore be on declaring the selected fairness objective, the relevant decision to which it applies, the rationale for its selection, and how it is tested and monitored, rather than prescribing a single fairness metric for all AI Systems.

3.3.2 Diverse and Representative Data

Gaming Licensees should strive to utilise diverse, representative, and balanced datasets to train and/or fine-tune AI Systems, ensuring that the data used is inclusive and encompasses the diversity of the population. By incorporating diverse perspectives and experiences into the training data, Gaming Licensees can mitigate biases and enhance the fairness of AI Systems' decision-making processes.

Where diverse, representative, and balanced datasets for the intended application are not available, Gaming Licensees should document the limitations and adopt supplementary fairness safeguards, such as using synthetic data to enhance the available datasets, to best endeavour to avoid discriminatory outcomes.

3.3.3 Implementation of Bias Detection and Mitigation Measures

Gaming Licensees should implement proportionate measures to detect and mitigate bias in AI algorithms. This includes addressing biases arising from legacy human data and decision-making processes. Continuous monitoring, evaluation, and updating of AI Systems are essential to ensure fairness and equity.

Gaming Licensees should also consider whether variables used by AI Systems may operate as proxies for protected or sensitive characteristics, even where those characteristics are not directly used. In the gaming context, relevant proxy variables may include, among others, geolocation, language, deposit patterns, device type, time-of-day activity, payment behaviour, play patterns, account history and customer interaction data. Gaming Licensees should maintain, and periodically review, a sector-relevant list of potential proxy variables for higher-impact AI Systems.

3.3.4 Designing AI Solutions to Address Bias

AI Systems should be designed to detect, mitigate, and where possible, eliminate bias against individuals, criteria, communities, groups, and other entities. Gaming Licensees should seek to design and deploy AI Systems that are inclusive and equitable, and do not disproportionately impact particular groups. This commitment is to be reflected in the design, development, and deployment processes, where fairness and inclusivity are prioritised.

To understand how AI Systems may risk violating non-discrimination obligations, it is necessary for Gaming Licensees to recognise the main sources of bias. These included but are not limited to²⁴:

- **Bias in Data:** the most common and well-documented source, including but not limited to:
 - *Historical bias*, where datasets reflect structural inequalities present at the time of collection, reproducing existing patterns of discrimination.

²⁴ Mehrabi et al. (2019), <https://doi.org/10.48550/arXiv.1908.09635> and Ferrara (2023) <https://doi.org/10.48550/arXiv.2304.07683>

- *Representation bias*, where certain groups are underrepresented due to sampling or availability limitations.
- *Measurement bias*, where proxy or measurement variables indirectly encode discrimination.
- *Aggregation bias*, where patterns specific to subgroups are obscured by treating the population as a homogenous and aggregated whole.
- **Bias in Design Choices/Techniques:** even with diverse, representative, and balanced datasets, model design/techniques can introduce bias, such as through optimisation choices, embedded statistical assumptions, or techniques that degrade performance for underrepresented groups.
- **Bias in Benchmark Datasets:** benchmark datasets may conceal discriminatory performance if those benchmarks replicate existing biases or fail to reflect the diversity of the target population, leading to poor real-world generalisation.
- **Domain-Specific Bias:** bias in AI Systems may be domain specific and often manifests in systems trained on behavioural or transactional datasets. For instance, responsible gambling algorithms may incorrectly classify certain play patterns as High-Risk due to skewed historical data, while AML and anti-fraud models may disproportionately flag specific customer segments. Bias may also arise from upstream foundation models, for example generative AI used in customer support or content creation, which can inherit or amplify inequities present in their training data.

Bias may arise not only from the composition of training data, but also from model design choices, optimisation objectives, feature selection, proxy variables, deployment context, feedback loops and human interpretation of AI outputs.

3.3.5 Inclusive Design

AI Systems should be designed in a manner that considers all users' diverse needs, avoiding discrimination, or exclusion. Gaming Licensees should ensure that AI Systems are accessible and equitable for all users, reflecting a commitment to inclusivity.

3.3.6 Continuous Monitoring and Evaluation

Gaming Licensees should establish mechanisms for continuous monitoring and evaluation of AI Systems to assess their fairness and non-discriminatory behaviour in real-world applications. This includes ongoing analysis of system performance and decision outcomes to identify and address any emerging issues, biases or discriminatory patterns. Such monitoring should be proportionate to the risk and nature of the AI System, and support compliance with applicable legal obligations relating to non-discrimination, fairness, and the protection of fundamental rights, including those protected under

applicable EU and national legislation (such as data protection, consumer protection, and non-discrimination laws).

3.3.7 User Empowerment and Redress

Gaming Licensees should provide appropriate channels for users to raise concerns and seek review or redress, where AI-supported outcomes may affect them and in accordance with applicable legal requirements.

3.3.8 Community Engagement and Consultation

Gaming Licensees should engage with relevant stakeholders and affected communities throughout the AI development and deployment lifecycle to ensure that AI Systems are designed and implemented to reflect community values and priorities. Community engagement facilitates the identification of potential biases and discriminatory impacts, enabling Gaming Licensees to address concerns and promote fairness and equity in AI use.

Player-facing consultation or outreach should apply where the Gaming Licensee holds the player relationship and can reasonably engage with affected players or communities. Where a Gaming Licensee acts as a B2B supplier and does not hold the player relationship or have the ability to contact end users directly, it should support proportionate engagement through alternative means, such as providing relevant information to B2C clients, supporting client-led assessments, documenting foreseeable player-impacting risks, and engaging with the MGA, clients or other relevant stakeholders where appropriate.

3.3.9 Fundamental Rights

Gaming Licensees should take proportionate steps to ensure that AI Systems respect fundamental rights, in accordance with applicable legal requirements.

Where required by applicable law, or where appropriate based on the AI System's risk and impact, Gaming Licensees should conduct relevant impact assessments, including Fundamental Rights Impact Assessments and Data Protection Impact Assessments.

3.4 Environmental Sustainability

Gaming Licensees should consider the potential environmental impact associated with the development, deployment, and use of AI Systems as part of responsible and sustainable operations. Where relevant, environmental considerations linked to AI adoption may be informed by recognised environmental management standards and sustainability frameworks, including ISO 14001²⁵, particularly where Gaming Licensees operate, procure, or rely on energy-intensive AI Systems or infrastructure.

²⁵ [ISO 14001:2026 – Environmental management systems](#)

Gaming Licensees should be mindful that many AI Systems require significant computational resources, which may result in increased energy consumption and associated carbon emissions. This is particularly relevant where Gaming Licensees train, retrain, or fine-tune their own systems, but also applies when procuring third-party AI software or services that rely on high-intensity data centres. Beyond training or fine-tuning activities, inference and deployment at scale also carry a cumulative environmental footprint.

Environmental sustainability should be considered as an ethical factor when deciding whether and how to adopt, develop, procure or deploy an AI System, rather than only after implementation. The level of assessment and action should be proportionate to the Gaming Licensee's role, degree of control over the AI System, and the scale and nature of the expected environmental impact.

Where a Gaming Licensee acts as a Provider, it should give greater consideration to the system's energy use, infrastructure requirements and environmental impact. This may include proportionate consideration of model design, training approach, hosting arrangements, deployment choices and ongoing monitoring.

Where a Gaming Licensee acts as a Deployer of a third-party SaaS, API-based or externally hosted AI System, its ability to assess and influence environmental impact will typically be more limited. In such cases, environmental considerations may be addressed through proportionate procurement due diligence, including reasonable reliance on supplier-provided information on energy use, carbon impact, data-centre sustainability or equivalent sustainability practices.

Where Gaming Licensees are subject to corporate sustainability reporting obligations, environmental considerations linked to AI adoption may also be aligned, where applicable, with relevant sustainability reporting frameworks and internal ESG processes.

3.4.1 Sustainable Procurement

When procuring AI services from third parties, Gaming Licensees should consider suppliers' environmental practices as part of procurement and due diligence processes, including through reasonable reliance on information provided by suppliers where appropriate. This may include, where feasible, considering whether providers demonstrate the use of renewable energy, energy-efficient infrastructure, sustainable data-centre practices or equivalent environmental commitments.

3.4.2 Energy Efficiency

Gaming Licensees that develop or operate their own GPAI Models or AI Systems, should prioritise energy-efficient techniques to minimise energy consumption. This may include optimising algorithms and models to reduce computational intensity, as well as leveraging energy-efficient hardware infrastructure for AI training and inference tasks.

3.4.3 Carbon Footprint Reduction

Gaming Licensees should make reasonable efforts to minimise the carbon footprint associated with the use of AI Systems by adopting sustainable practices throughout the AI lifecycle. This may include prioritising the use of renewable energy sources for AI infrastructure, reducing unnecessary computational overhead, and optimising resource utilisation to minimise energy-intensive operations.

In parallel, Gaming Licensees should raise awareness of the environmental impact of AI within their organisations, including through staff training and internal guidance. Embedding sustainability considerations into procurement decisions, tool selection, and AI development and deployment practices can support more sustained and meaningful reductions in AI-related environmental impacts.

3.4.4 Efficient AI Model Training and Deployment

Gaming Licensees should employ proportionate techniques to improve the efficiency of AI model training and deployment, with a view to reducing energy consumption and environmental impact. This may include optimising training workflows, minimising data processing overhead, and utilising distributed computing resources to improve energy efficiency during model development and deployment phases.

3.4.5 Green AI Practices

Gaming Licensees should consider adopting green AI practices, prioritising environmental sustainability in AI development and deployment. This may include promoting research and innovation in eco-friendly AI technologies, such as energy-efficient algorithms, hardware optimisation, and sustainable computing practices.

3.4.6 Eco-design Principles

Gaming Licensees should integrate eco-design principles into developing and deploying AI Systems to minimise environmental impact. This includes considering ecological factors such as energy consumption, resource utilisation, and carbon emissions during the design phase, and implementing measures to mitigate environmental risks and promote sustainability.

3.4.7 Environmental Impact Assessment

Gaming Licensees should assess, in a proportionate manner, the potential environmental impacts of AI projects, and identify strategies for mitigating their negative impacts. This may include assessing AI Systems' direct and indirect environmental effects, such as energy consumption, carbon emissions, and resource depletion, and implementing measures to minimise environmental harm where possible, and to identify opportunities for improvement and optimisation.

Such assessments should consider the entire lifecycle of an AI System to evaluate their environmental impact across various stages, including development, deployment, operation, and eventual decommissioning.

3.4.8 Continuous Improvement

Gaming Licensees should continuously monitor and evaluate AI Systems' environmental performance and implement measures for continuous improvement and optimisation. This includes setting environmental sustainability goals, tracking key performance indicators related to energy efficiency and carbon footprint, and implementing strategies to achieve environmental sustainability targets.

3.4.9 Stakeholder Collaboration

Gaming Licensees should collaborate with stakeholders, including environmental organisations, policymakers, and industry partners, to promote ecological sustainability in AI deployment. This includes sharing best practices, exchanging knowledge and expertise, and collectively addressing environmental challenges associated with AI technologies.

3.4.10 Responsible Decommissioning

Sustainability should also extend to the end-of-life stage of AI Systems, including responsible decommissioning of hardware and minimising electronic waste. Gaming Licensees should explore best practices for recycling and safe disposal of obsolete hardware and related infrastructure.

3.5 Data Protection, Security, and Governance

The gaming industry relies extensively on the collection and processing of personal data to support gaming operations and deliver personalised experiences. As such, the protection of personal data, the security of information systems, and adherence to ethical data-processing practices are essential.

A sound data governance framework is essential to ensure that personal data is handled responsibly and in compliance with relevant data protection laws. By adhering to the principles set out in this Charter, Gaming Licensees can foster trust among users and stakeholders while mitigating risks associated with the use of AI Systems.

3.5.1 Compliance with Data Protection Laws

Gaming Licensees remain responsible for ensuring, in accordance with applicable data protection laws, that AI Systems involving the processing of personal data comply with relevant data protection requirements and regulatory guidance, such as the GDPR in the European Union, the MGA's Industry Guidelines on the GDPR, developed in consultation with the Information and Data Protection Commissioner (IDPC)²⁶, and any other applicable guidance issued by competent authorities.

Where AI Systems involve automated individual decision-making, including profiling, Gaming Licensees should assess whether the processing falls within applicable GDPR requirements, including safeguards for decisions based solely on automated processing that produce legal or similarly significant effects on individuals.

²⁶ [The-General-Data-Protection-Regulation.pdf](#)

Where such requirements apply, Gaming Licensees should provide affected individuals with appropriate information about the existence of automated decision-making, meaningful information about the logic involved, and the significance and envisaged consequences of the processing, in line with applicable GDPR transparency obligations.

Appropriate safeguards should also be in place to enable individuals to obtain human intervention, express their point of view, and contest relevant decisions, where required by the GDPR.

3.5.2 Lawful and Ethical Data Processing

Where AI Systems process personal data, Gaming Licensees must ensure, in accordance with applicable data protection laws (e.g., GDPR), that such processing is lawful, fair and transparent, and supported by an appropriate legal basis.

3.5.3 Data Minimisation and Purpose Limitation

Where personal data is processed, Gaming Licensees must comply, in accordance with the GDPR, with the principles of data minimisation and purpose limitation, ensuring that data is only used for specified, explicit, and legitimate purposes, as transparently disclosed to the data subject pursuant to the GDPR's transparency disclosure requirements in Article 13 and 14 of the GDPR.

3.5.4 Data Quality

Gaming Licensees should take appropriate steps to ensure the quality of data used by AI Systems, using suitable metrics where relevant. This includes implementing appropriate measures to maintain data integrity, accuracy, timeliness, and correctness throughout the data lifecycle to mitigate the risk of errors, poor performance or biases in AI-driven decision-making.

High-Risk AI Systems under the EU AI Act are also subject to specific data quality requirements pursuant to Article 10 of the EU AI Act, such as having appropriate data governance and data management processes, having relevant and representative datasets, and the context of the intended usage of the High-Risk AI System.

3.5.5 Data Protection by Design and Default

Gaming Licensees must, where applicable and in accordance with the GDPR, integrate data-protection-by-design and data-protection-by-default principles into AI Systems, embedding data protection into the design and functionality. Privacy-enhancing features, such as data anonymisation, pseudonymisation, encryption, and access controls, should also be implemented by default.

3.5.6 Data Security, Safety, and Confidentiality

Where personal data is processed, Gaming Licensees are required under applicable data protection laws, including the GDPR, to implement appropriate technical and organisational security measures. More broadly, Gaming Licensees should implement proportionate safety and security mitigations

across the AI lifecycle to manage risks associated with AI Systems. Where applicable under the EU AI Act, Providers of GPAI models or systems should have regard to, and comply with any applicable obligations or commitments under the *General-Purpose AI Code of Practice – Safety and Security Chapter*, as published by the European Commission²⁷.

3.5.7 Transparency and Accountability in Data Processing

Gaming Licensees should ensure transparency in data processing activities, providing clear information to individuals about how their personal data is being used, as required in terms of the GDPR. Transparency fosters trust and enables individuals to exercise their rights.

3.5.8 Data Retention and Disposal

Where personal data is processed, Gaming Licensees are required under applicable data protection laws to comply with retention and storage limitation requirements. Gaming Licensees should maintain appropriate policies governing data retention and disposal, ensuring that data is retained only as long as necessary and securely deleted or anonymised when no longer needed.

3.5.9 Cross-Border Data Transfers

Where personal data is transferred outside the EEA, Gaming Licensees must comply with applicable GDPR requirements on international transfers, including appropriate safeguards where required. More broadly, Gaming Licensees should manage the risks associated with cross-border data transfers involving AI Systems.

3.5.10 Data Governance Framework

Gaming Licensees should establish and maintain a proportionate data governance framework to oversee data-related activities, including roles and responsibilities for data stewardship and compliance with applicable data protection requirements.

Where personal data is processed, such governance frameworks should support compliance with data-subject-rights requirements under the GDPR.

3.6 Human Oversight and Responsibility

The rapid advancement of AI technologies in the gaming industry presents both opportunities and challenges. To ensure that AI is used ethically and responsibly, effective human oversight and clear accountability are essential.

Human oversight (also referred to as human-in-the-loop) is a cornerstone of responsible AI governance. It reflects the principle that human beings must retain the ability to supervise, interpret, and where necessary, intervene in the operation of AI Systems.

²⁷ [The General-Purpose AI Code of Practice | Shaping Europe's digital future](#)

Within the gaming sector, human oversight serves as an important safeguard against fully automated decision-making, leading to adverse effects on players or gaming integrity without the benefit of human judgement. By maintaining appropriate levels of oversight, Gaming Licensees can ensure that AI Systems support, rather than replace, responsible decision-making.

This section addresses the importance of human control over AI Systems, accountability for AI-supported outcomes, appropriate training and oversight arrangements, and the ongoing evaluation of AI performance.

3.6.1 Ensuring Human Control

Gaming Licensees should ensure that AI Systems are subject to meaningful human oversight and control, proportionate to their risk and impact. Where required by applicable law, human oversight arrangements must comply with the relevant legal requirements. Human intervention should be possible at the appropriate stages of the decision-making process, allowing for review, override, and where relevant, reversal of AI-generated decisions.

For player-facing conversational or support AI, Gaming Licensees should define clear human-escalation triggers, including repeated low-confidence outputs, unresolved or escalating player queries, detected distress signals, indications of vulnerability or gambling harm, or requests for human support. Where such triggers arise, the conversation should be transferred to a suitably trained human agent in a timely manner.

Where required under applicable law, individuals should be able to object to decisions based solely on automated processing, and to request human review in accordance with data-protection and consumer-protection requirements. As an example, High-Risk AI Systems under the EU AI Act require human oversight to be built into the AI System, for Deployers to be able to maintain human agency, such as by disregarding or overriding the AI System's output and intervene in its operation.

3.6.2 Accountability for AI Decisions/Harm

Gaming Licensees should remain accountable for the use of AI Systems within their operations, and for outcomes resulting from AI supported decisions. Appropriate responsibility arrangements should be in place to address adverse effects or unintended consequences arising from the use of AI Systems.

Gaming Licensees should establish mechanisms for identifying and addressing harms promptly, including remediation, compensation, and corrective actions, to uphold accountability and mitigate adverse impacts on affected parties.

Where relevant, appropriate reporting and escalation channels should be in place to notify affected players and, where required, competent authorities, in accordance with applicable legal and regulatory obligations.

3.6.3 Human Oversight in Mitigating Risks of Bias and Discrimination

Human oversight is crucial for mitigating the risks of bias and discrimination in AI Systems. Gaming Licensees should monitor AI Systems in a proportionate manner for potential biased or discriminatory outcomes, and should implement corrective measures where such risks are identified. Where monitoring or corrective action is required by applicable law, Gaming Licensees must comply with those requirements.

3.6.4 Ensuring Ethical Use of AI

Gaming Licensees should establish ethical guidelines and internal policies and procedures for the use of AI Systems, ensuring alignment with ethical norms and applicable fundamental-rights standards.

Human oversight supports the practical application of these principles by ensuring that ethical considerations are integrated into AI-enabled decision-making and operational processes.

3.6.5 Monitoring System Performance and Reliability

Gaming Licensees should establish mechanisms for ongoing monitoring and evaluation of AI Systems to assess their accuracy, safety, risks of bias and discrimination. Human intervention is essential for identifying system failures or unexpected behaviours.

In addition, Gaming Licensees should maintain appropriate incident logs, including issues and near-miss events. Where appropriate, significant failures, risks, or adverse effects should be reported to the relevant competent authorities, demonstrating accountability beyond internal monitoring. Gaming Licensees should implement proportionate corrective actions to mitigate any negative outcomes or risks identified through this monitoring. Where corrective action is required by applicable law or regulatory obligation, Gaming Licensees must comply with those requirements.

3.6.6 Training, Expertise and AI Literacy

Gaming Licensees should ensure that individuals responsible for overseeing or interacting with AI Systems have appropriate training and expertise to monitor and intervene effectively in AI-driven decision-making processes.

In addition, Gaming Licensees should foster broader organisational AI literacy, proportionate to the nature and risk of the AI Systems used. This entails that employees utilising such AI Systems should understand the risks and responsibilities associated with AI use in their application, and should have an understanding of AI as a technology to a level relevant for their operation of or interaction with the AI System. Where AI literacy or training obligations apply under applicable law, Gaming Licensees must comply with those requirements. For High-Risk AI Systems, it should be noted that Deployers are required to assign human oversight to persons who have the necessary competence, training and authority, as well as the necessary support. The latter includes providing necessary training as a

prerequisite for human oversight of a High-Risk AI System, which may include AI literacy as part of that training.

3.6.7 Transparency and Communication

Gaming Licensees should provide clear information on the extent of human oversight and the limitations of AI Systems.

3.6.8 Continuous Improvement and Adaptation

Gaming Licensees should establish mechanisms to support continuous improvement of human-oversight arrangements. This includes incorporating feedback, reviewing incidents and outcomes, and adapting oversight practices to address emerging risks, technological developments, or changes in regulatory expectations.

3.7 Reliability, Safety, and Robustness

The increasing use of AI Systems in the gaming industry has brought forth new vulnerabilities and safety risks, making reliability, safety, and robustness more critical than ever. These risks may arise from technical failures, system errors, data quality issues, or vulnerabilities that could be exploited by malicious actors. As a result, the deployment and/or integration of AI Systems within gaming operations necessitates a strong focus on ensuring that such systems operate as intended and do not undermine player protection or the integrity of gaming activities.

To ensure the well-being of players and the integrity of the gaming experience, it is crucial to design and implement AI Systems that are resilient to errors, vulnerabilities, and malicious attacks. This section outlines key principles and practices for building reliable, safe, trustworthy, and robust AI Systems, including thorough testing, risk assessment, and mitigation measures.

3.7.1 Design for Reliability

AI Systems should be designed with reliability in mind to support consistent and accurate performance across relevant gaming scenarios. Gaming Licensees should employ robust design principles, such as redundancy, fault tolerance, and error handling mechanisms, to mitigate the risks of system failures and ensure continuous operation.

3.7.2 Thorough Testing and Validation

Gaming Licensees should conduct proportionate testing and validation of AI Systems to identify and mitigate potential risks and vulnerabilities, taking into account the AI System's risk profile, intended purpose and operational context. This includes testing the system under diverse inputs, conditions, and edge cases to assess its performance, accuracy, and safety before deployment.

For higher-impact AI Systems, including any High-Risk AI Systems under Annex III of the EU AI Act deployed in the domains of responsible gambling, fraud detection, AML, KYC, account restriction, eligibility, intervention or other compliance-sensitive contexts, testing and validation should include proportionate adversarial robustness testing, red-teaming or equivalent challenge exercises. Gaming Licensees should also monitor for data drift and concept drift, define retraining or recalibration triggers, and test rollback, kill-switch, fail-safe and graceful-degradation mechanisms. Ongoing monitoring should be supported by proportionate Machine Learning Operations (MLOps) practices, including model version control, change logs, shadow scoring, champion-challenger testing, staged deployment, performance monitoring, retraining governance and model-retirement criteria.

3.7.3 Risk Assessment and Mitigation

Gaming Licensees should conduct proportionate risk assessments to identify potential risks associated with the use of AI Systems and should implement appropriate risk mitigation strategies, including safeguards, fail-safe mechanisms, escalation procedures and contingency plans, to address identified risks and support the safety and reliability of AI-enabled processes.

Such assessments should consider technical, ethical, legal and regulatory risks, particularly where player protection is involved. Where risk assessment or mitigation measures are required by applicable law, Gaming Licensees must comply with those requirements.

This Charter is intended to apply in a technology-agnostic and risk-based manner, including to traditional, generative and agentic AI Systems. Where an AI System is capable of autonomous or semi-autonomous multi-step action, including tool use or interaction with external systems, that capability should be treated as a relevant risk factor when determining the applicable safeguards. Depending on the use case and risk classification, safeguards may include sandboxing, tool-use restrictions, permission and action-rate limits, rollback mechanisms and clearly defined human-escalation points.

3.8 Compliance, Monitoring, and AI Governance

As AI Systems become increasingly complex and integrated into gaming experiences, it is crucial to establish robust frameworks for compliance, monitoring, and governance. This section outlines the importance of designing AI Systems that adhere to safety standards, implementing continuous monitoring and maintenance practices, and establishing effective AI governance structures. By prioritising these areas, Gaming Licensees can ensure the ethical, legal, and responsible use of AI.

AI governance arrangements should reflect the risk-based application of this Charter, ensuring that documentation, testing, monitoring, escalation, human oversight and assurance measures are proportionate to the applicable AI use-case risk tier and its potential impact on players, gaming integrity and regulated outcomes.

3.8.1 Adherence to Safety Standards

AI Systems should be designed, deployed, and operated in accordance with the applicable safety standards and recognised best practice. Gaming Licensees should have regard to relevant standards and frameworks that support responsible AI governance, safety, accountability and trustworthiness, including ISO/IEC 42001²⁸, the OECD AI Principles²⁹, the UNESCO Recommendation on the Ethics of Artificial Intelligence³⁰, and relevant industry best practice, including International Gaming Standards Association (IGSA) materials³¹ where applicable.

Where AI Systems perform functions with potentially significant impact on players or gaming integrity, for example responsible gambling, fraud detection, or AML contexts, additional safeguards, independent validation or review mechanism should be used to verify that safety objectives are being met.

For High-Risk AI Systems, harmonised European standards developed by CEN-CENELEC facilitate compliance with the High-Risk requirements of the EU AI Act.

3.8.2 Continuous Monitoring and Maintenance

Gaming Licensees should establish mechanisms for monitoring and maintaining AI Systems to detect and address emerging risks and vulnerabilities. This includes real-time monitoring of system performance, anomaly detection, and proactive maintenance to ensure ongoing reliability and safety.

Continuous monitoring and maintenance should be supported by proportionate MLOps practices for AI Systems. These may include model version control, change logs, shadow scoring, champion-challenger testing, staged deployment, performance monitoring, retraining governance, rollback procedures, and model-retirement criteria.

3.8.3 AI System Impact Assessment

Gaming Licensees should conduct AI System Impact Assessments (AISIA) to anticipate and manage the reasonably foreseeable impacts of AI Systems, in line with international best practices.

An AISIA provides a structured method for identifying potential effects on individuals, groups, and society, and for ensuring that risks are systematically evaluated throughout the AI lifecycle, including design, deployment, and post-deployment.

Where relevant, such assessments may be informed by impact assessment concepts reflected in the EU AI Act and by international frameworks, including the OECD AI Principles, the UNESCO Recommendation on the Ethics of Artificial Intelligence, and the Council of Europe's Human Rights,

²⁸ [ISO/IEC 42001:2023 – AI management systems](#)

²⁹ [AI principles | OECD](#)

³⁰ [Recommendation on the Ethics of Artificial Intelligence | UNESCO](#)

³¹ [Best Practices – IGSA](#)

Democracy and the Rule of Law Impact Assessment for AI Systems (HUDERIA)³², while remaining proportionate and adapted to the gaming context.

3.8.4 AI Governance

Gaming Licensees should ensure that a comprehensive AI Governance framework is established and implemented to support the responsible use of AI in line with this Charter. An AI governance framework provides a structured approach to managing AI risks and ensuring ethical AI practices within the organisation. In developing or maintaining AI governance arrangements, Gaming Licensees may have regard to recognised AI management system standards and governance frameworks, including ISO/IEC 42001 and, where relevant, applicable European standards developed in support of the EU AI Act, taking into account the Gaming Licensee's role, scale, AI use cases and risk profile. Such standards and frameworks may support the establishment of structured AI governance and risk management arrangements but should not be interpreted as replacing, modifying or determining compliance with applicable legal obligations.

Gaming Licensees should maintain an AI inventory as the foundational system of record for AI governance. The inventory should identify relevant AI Systems, their purpose, owner, risk tier, deployment context, key dependencies, applicable role under the EU AI Act, and lifecycle status. It should support risk assessment, monitoring, accountability, change management and supervisory engagement.

In addition, a comprehensive AI governance framework should incorporate the performance of structured AISIA to evaluate reasonably foreseeable impacts on individuals, groups and society across the AI lifecycle.

Key components of the governance framework include:

- **Aligning AI with Organisational Values:** Integrating AI principles into existing organisational frameworks and processes.
- **Establishing an Operational Model:** Defining the optimal structure, processes, and technologies to support AI governance.
- **Defining Roles and Responsibilities:** Clearly outlining the roles of key stakeholders involved in AI initiatives.
- **Implementing Performance Metrics:** Tracking relevant performance indicators to monitor AI operations and ensure continuous improvement.

³² [HUDERIA – risk and impact assessment of AI systems – Artificial Intelligence](#)

- **Thresholds for Sensitive Use-cases:** Defining and enforcing internal policy thresholds for sensitive uses of AI Systems, particularly where vulnerable groups or High-Risk applications are concerned.

AI governance arrangements should be proportionate to the Gaming Licensee's size, complexity and AI risk profile. Any Gaming Licensee that adopts, develops, procures, deploys or otherwise uses AI Systems should designate a senior accountable person for AI governance, such as a Chief AI Officer (CAIO) or equivalent role. The role should have sufficient authority, access to relevant information and organisational support to oversee AI governance, coordinate relevant functions, and ensure that AI-related risks, controls and decisions are appropriately escalated.

Gaming Licensees should also establish an AI oversight or ethics committee, or an equivalent documented governance forum, to support responsible AI decision-making. The committee or forum should be multidisciplinary and should include, as relevant, representatives from business, compliance, legal, data protection, risk, technology, information security, responsible gambling, AML/fraud, product and operational functions. Its composition, meeting frequency and formality should be proportionate to the Gaming Licensee's size, complexity and AI risk profile.

The AI oversight or ethics committee should have a clear mandate to review, challenge and approve AI governance matters, including AI inventory updates, risk classification, impact assessments, fairness and explainability considerations, testing and validation results, material changes, incident learnings, escalation decisions and lifecycle gates for higher-impact AI Systems. Larger or more complex Gaming Licensees may maintain a standing committee, while smaller Gaming Licensees may fulfil the same function through a documented governance forum or committee with clear terms of reference, responsibilities and records of decisions. To support effective governance, Gaming Licensees should designate an appropriate function to oversee AI ethics and compliance. This function may be assigned to a suitably qualified/experienced individual, or a multidisciplinary team with relevant expertise. Depending on the nature and scale of the risks, ethical review boards or oversight committees may also be considered to support review and decision making.

Ethical and governance review processes should include considerations of fairness, equity, and social justice to prevent the deployment of discriminatory AI technologies.

3.8.5 External Audits and Reviews

Gaming Licensees may be required to undergo proportionate external audits and reviews of AI Systems by independent third parties to assess compliance with transparency and accountability requirements, as may be directed by the Authority. External scrutiny enhances credibility, promotes trust among stakeholders, and validates the integrity and reliability of AI technologies. Where external assurance is

undertaken, Gaming Licensees may consider recognised assurance standards or equivalent methodologies, including ISAE 3000³³, where relevant to the scope and purpose of the review.

Where external assurance or review is undertaken, the assessment should be performed against the criteria, thresholds, evidence requirements and testing methodologies declared by the Gaming Licensee, rather than against a single prescribed metric that may not be suitable for all AI Systems or use cases.

3.9 Reporting and Metrics

To ensure the effectiveness and reliability of AI Systems in the gaming industry, it is essential to establish robust reporting and metrics frameworks. This section delves into the importance of tracking key performance indicators, monitoring system behaviour, and reporting on the impact of AI on the gaming experience. By utilising data-driven insights, Gaming Licensees can optimise AI Systems and make informed decisions.

For higher-impact AI Systems, Gaming Licensees should define and document the criteria, acceptance thresholds, evidence requirements and testing methodologies they use to assess relevant Charter principles, including fairness, transparency, reliability and explainability. The specific metrics and thresholds may vary depending on the AI System, use case, data, model type and operational context. The focus should therefore be on ensuring that the Gaming Licensee has declared the criteria, thresholds and testing methods it has adopted, together with the rationale for their selection, so that internal reviewers, auditors or other assurance Providers have a clear standard against which to assess the system.

3.9.1 Reliability Metrics and Reporting

Gaming Licensees should establish reliability metrics and thresholds to ensure that AI Systems maintain performance and accuracy over time. Such metrics may include error rates, false positive/negative ratios, system uptime, bias detection scores, and intervention success rates in responsible gambling contexts, as may be relevant to the respective AI System's application.

Reliability metrics, testing results and related performance information should be maintained internally and made available to the MGA and the MDIA as the supervisory authority under the EU AI Act upon request where relevant to supervisory engagement or compliance-related reviews. Gaming Licensees should not be expected to publish detailed reliability, performance, testing or monitoring information where in doing so could expose trade secrets, proprietary information, security-sensitive controls, fraud or AML detection logic, responsible-gaming intervention logic, or other information that could be exploited by bad actors.

³³ [International Standard on Assurance Engagements \(ISAE\) 3000 Revised, Assurance Engagements Other than Audits or Reviews of Historical Financial Information | IAASB](#)

Gaming Licensees should maintain appropriate internal reporting mechanisms to track reliability results from continuous monitoring and testing, including material deviations from expected performance and action plans to address identified issues.

Where reliability issues, significant performance degradation, or unexpected adverse outcomes are identified, such mechanisms should interface with relevant incident management and escalation processes, in line with applicable legal and regulatory requirements. This approach supports continuous monitoring and corrective action throughout the AI System lifecycle, consistent with expectations under the EU AI Act.

3.10 Notification Requirements

As AI Systems become more complex and integrated into gaming operations, it is crucial to establish clear and proportionate notification requirements to ensure transparency and accountability. Notifications should not only satisfy regulatory compliance but also demonstrate the licensee's commitment to responsible governance, player protection, and continuous improvement.

This section outlines the importance of regulatory notifications and incident reporting linked to the use of AI Systems. By adhering to these requirements, Gaming Licensees can maintain compliance with relevant regulations and foster trust with stakeholders. Notifications and reports may be further shared by the MGA with the relevant supervisory authorities, as appropriate.

3.10.1 Regulatory Notifications

Gaming Licensees should notify the MGA of any significant changes or updates to AI Systems that affect gaming operations. This includes material modifications to system functionality, deployment in new contexts (such as responsible gambling interventions or fraud detection), or modifications that may alter licensees' regulatory or risk profiles.

Notifications should be submitted in a timely manner and provide sufficient detail to enable the MGA to understand the nature of the change and assess its potential impacts.

3.10.2 Incident Reporting

Gaming Licensees must continue to comply with existing incident-reporting obligations, in accordance with applicable law and regulatory requirements. Incidents involving AI Systems should be assessed against existing MGA reporting requirements and any other applicable reporting obligations arising under relevant legal frameworks.

Where an incident related to a High-Risk AI System meets the criteria for a serious incident under the EU AI Act, the relevant reporting obligations under that Regulation shall apply. In Malta, such reports must be submitted by the Provider to the market surveillance authorities designated under the EU AI Act, i.e. the MDIA and the IDPC in Malta.

The MDIA must be informed of serious incidents caused by High-Risk AI Systems in the domains of biometrics, critical infrastructure, education and vocational training, employment, workers management and access to self-employment, and access to and enjoyment of essential private services and essential public services and benefits.

The IDPC must be informed of serious incidents caused by High-Risk AI Systems in the domains of emergency response, law enforcement (including biometrics within a law enforcement context), migration, asylum and border control management, and administration of justice and democratic processes. This is without prejudice to reporting obligations from GDPR and the Data Protection Act.

Where the domain of the High-Risk AI System is overlapping multiple functions, or there are any doubts otherwise, both the MDIA and the IDPC should be informed of a serious incident where a causal link was established between the serious incident and a High-Risk AI System.

Reports should include, where applicable, information on the nature of the incident, steps taken to mitigate its impact, and measures implemented to prevent recurrence. This should be done in accordance with applicable legal and regulatory obligations, including where relevant, the GDPR, the Cyber Resilience Act, and the EU AI Act.

3.10.3 Feedback and Continuous Improvement

Gaming Licensees that choose to align with this Charter are encouraged to maintain appropriate internal records or summaries reflecting their use of AI Systems, the manner in which the principles of the Charter are being applied, and any material ethical or governance considerations identified. Such information may be made available to the MGA upon request, for example as part of supervisory engagement or compliance-related reviews, and without prejudice to existing legal and regulatory obligations.

In addition, Gaming Licensees are encouraged to contribute to the ongoing development of the Charter through structured engagement facilitated by the MGA, which may take place on a yearly basis. This may include high-level discussions, surveys, workshops, or consultations aimed at sharing experiences, identifying good practices, and highlighting areas where further guidance may be beneficial.

The purpose of this engagement is to support dialogue and continuous improvement, ensuring that the Charter remains relevant, practical, and responsive to technological, regulatory, and operational developments.

3.11 Involvement of the MGA and MDIA

3.11.1 Regulatory Oversight

The MGA may support supervisory engagement and sector-specific guidance in relation to this Charter, in line with its mandate under the Gaming Act. Where relevant, the MGA may coordinate with competent

authorities designated under the EU AI Act, without prejudice to their respective roles and responsibilities under that framework.

The MDIA is the national authority responsible for ensuring that Malta remains at the forefront in the regulation of AI, and is designated as the lead Market Surveillance Authority, Notifying Authority, and Single Point of Contact in relation to the EU AI Act. The MDIA is responsible for supervising, investigating, and enforcing compliance with the AI Act for AI Systems placed on or used in the EU market, within the Maltese Islands.

3.11.2 Collaboration with the MGA

Gaming Licensees are encouraged to collaborate with the MGA on initiatives aimed at promoting the ethical and responsible use of AI within the gaming sector. This may include participation in workshops, consultations, stakeholder engagement activities, or other initiatives designed to support knowledge-sharing and continuous improvement.

3.11.3 Compliance Support

The MGA may make compliance support available to Gaming Licensees wishing to participate in this Charter, including information, training, resources or support as may be deemed appropriate by the MGA, to support understanding of applicable regulatory requirements and good practice standards.

Gaming Licensees can find detailed information on the EU AI Act at the MDIA's AI webpage³⁴. As mentioned in Section 3.1 of this Charter, compliance tools are readily available on the MDIA's dedicated webpage³⁵, including an EU AI Act Compliance Checker developed by the EU AI Office³⁶. Gaming Licensees may also contact the MDIA AI Helpdesk (ai.mdia@mdia.gov.mt) for enquiries related to the EU AI Act.

4 Closing Statement

This Charter is intended to support Gaming Licensees in the ethical and responsible use of artificial intelligence, promoting practices that are fair, transparent, and aligned with the interests of players, stakeholders, and the wider public.

This Charter may be periodically reviewed and updated to reflect legal developments, evolving ethical standards, and technological advancements.

³⁴ [Artificial Intelligence – MDIA](#)

³⁵ [Compliance Tools – MDIA](#)

³⁶ [EU AI Act Compliance Checker – EU AI Office](#)

Annex A – Mapping of Charter Principles to EU AI Act and GDPR

This Annex provides a high-level mapping between the principles of this Charter and relevant themes under the EU AI Act and GDPR. It is intended to clarify where the Charter reflects or supports existing legal obligations, and where it provides additional sector-specific good practice for the responsible use of AI in gaming, and does not provide presumption of conformity with the EU AI Act or other applicable legislation.

The mapping is indicative and does not constitute legal advice or create additional legal obligations. Each Gaming Licensee remains responsible for assessing the legal requirements applicable to its own AI Systems, taking into account its role, the nature and purpose of the system, the data processed, and the system's risk classification.

In applying the Charter, Gaming Licensees may also have regard to recognised standards, frameworks and assurance methodologies that support responsible AI governance, sustainability, impact assessment and assurance. These may include, where relevant, CEN-CENELEC harmonised standards, ISO/IEC 42001, ISO 14001, ISAE 3000, the OECD AI Principles, the UNESCO Recommendation on the Ethics of Artificial Intelligence, the Council of Europe's HUDERIA methodology, and relevant IGSA best-practice materials.

Charter principle	EU AI Act / GDPR alignment	Charter function	Gaming-sector relevance
Risk-Based Application of the Charter	EU AI Act risk-based approach	Gaming sector good practice	Applies proportionality by reference to player proximity, player impact and regulated outcomes.
Transparency, Accountability, and Explainability	EU AI Act transparency and Deployer-use obligations; GDPR transparency and automated decision-making safeguards	Existing legal obligations and gaming-sector good practice	Clarifies meaningful explanations without requiring disclosure of IP, source code, fraud, AML or security-sensitive logic.

Fairness and Non-Discrimination	GDPR fairness, accuracy and transparency principles; EU AI Act risk, data governance and fundamental-rights considerations	Existing legal obligations and gaming-sector good practice	Applies fairness and bias controls to player profiling, responsible gambling, marketing, fraud, AML and other gaming use cases.
Environmental Sustainability	EU AI Act trustworthy AI objectives and sustainability considerations	Sector good practice	Treats sustainability as an ethical adoption factor, calibrated to whether the Licensee acts as Provider or Deployer.
Data Protection, Security, and Governance	GDPR data protection principles, security, data protection by design/default and Data Processing Impact Assessment (DPIA) requirements; EU AI Act data governance and robustness concepts	Primarily restates / operationalises law	Supports responsible handling of player and operational data used in AI Systems.
Human Oversight and Responsibility	EU AI Act human oversight and Deployer-use obligations; GDPR automated decision-making safeguards	Existing legal obligations and gaming-sector good practice	Reinforces human review, escalation and accountability for player-impacting AI-supported outcomes.
Reliability, Safety, and Robustness	EU AI Act risk management, accuracy, robustness and cybersecurity concepts	Existing legal obligations and gaming-sector good practice	Applies testing, validation and resilience expectations to gaming systems affecting player protection and gaming integrity.
Compliance, Monitoring, and AI Governance	EU AI Act monitoring, governance and Deployer obligations; GDPR accountability principle where personal data is processed	Existing legal obligations and gaming-sector good practice	Encourages proportionate AI governance frameworks, monitoring and impact assessment in gaming operations.
Reporting and Metrics	EU AI Act monitoring concepts; GDPR accountability where personal data is involved	Sector good practice / operational support	Encourages internal metrics and reporting to evidence responsible AI use and continuous improvement.
Notification Requirements	EU AI Act incident and monitoring concepts; GDPR	Existing legal obligations and Charter clarification	Clarifies that AI-related incidents or material changes should be

	breach/accountability processes where relevant		considered against existing reporting duties.
Involvement of the MGA and MDIA	EU AI Act competent-authority framework; sector regulatory governance	Sector governance clarification	Explains how the MGA and MDIA support responsible AI adoption without creating additional Licensee obligations.