

THE STATE OF AI IN GAMING 2026

How AI is Shaping the Global Gambling Industry
Inaugural Annual Report

An AiR Hub initiative

Produced by: UNLV International Gaming Institute in collaboration with KPMG LLP



Introducing The State of AI in Gaming

Message from the Editor-in-Chief

The gambling and gaming industry is at an inflection point. In just a few years, artificial intelligence (AI) has shifted from an emerging novelty to operational necessity - powering everything from personalization and game development to compliance and surveillance. But as adoption accelerates, key questions arise: How are organizations deploying AI in practice? Where are the biggest challenges? How can we ensure responsible implementation? And what new AI advancements will drive future innovation?

We created *The State of AI in Gaming* to address these questions, and it is my pleasure to welcome you to the first in what will be an annual series. With this inaugural edition we set a foundation for tracking what will, without a doubt, be an exciting evolution for the entire field.

We deliver *The State of AI in Gaming* as an open, not-for-profit research project. It is designed to be a key resource on AI progress across the global gambling value chain for the benefit of all stakeholders: operators, suppliers, regulators, non-profits, researchers, academics, investors, journalists, and more. Ultimately, we seek to provide data-driven insights to help steer conversations and inform decision-making around AI. We hope this initiative helps the sector keep pace with the (shockingly) rapid advancements in AI, as well as set the stage for the gambling industry to position itself as a global leader in AI adoption and maturity.

I'd like to personally thank the entire State of AI Team for their herculean effort in bringing this vision to life. I am proud to say that this project is a result of an amazing collaborative effort that brought together organizations and individuals representing a diverse set of backgrounds, experiences, and roles within the sector. Our contributors range from ambitious students to seasoned professionals with decades of experience! I'd like to also take this opportunity to thank the Founding Members of AiR Hub, whose support and belief in the initiative has made this project possible.

The State of AI in Gaming report structure is purposefully more navigable than a traditional academic paper, without sacrificing the rigor of our research. Our methodology reflects this balance, where we draw from publicly available data and original survey research, while also providing qualitative depth via editorial commentary and targeted interviews with a broad range of stakeholders.

I hope this inaugural edition generates enthusiasm and greater interest in this fascinating topic, but the true value of this project lies in its longevity and breadth. To remain a definitive resource, we need the active participation of stakeholders from all corners of our field. *I therefore invite you to join this conversation.* You have the ability to participate in our surveys, provide data, re-analyze the datasets we have collected, and/or offer suggestions on how we may improve our future methodology and analysis.

My hope is that this 2026 report provides a strong foundational benchmark that can inspire innovation and growth, while encouraging *you* to become a permanent part of the conversation about *The State of AI in Gaming*.

Kasra Ghaharian, PhD

Director of Research, UNLV International Gaming Institute
Principal Investigator & Editor-in-Chief, The State of AI in Gaming

At a Glance



The Industry Has Ambitions, But Remains in Early Stages of Maturity.

With an average score of 45 out of 100 in our AI Maturity Index, most gambling companies have strategic ambitions for AI, but infrastructure and expertise need to catch up to scale it.



A Governance Gap Runs Through the Industry.

Governance scored lowest on our Maturity Index at 30 out of 100. Just one in five companies have dedicated AI governance roles, only a few companies plan to hire for these roles, and most organizations have no established governance practices or are in early stages of development.



Generative AI Is Widespread, But Agentic AI Has Yet to Arrive.

Over 80% of companies use Generative AI, yet adoption of AI agents lags far behind broader enterprise trends. The sensitivity of gambling operations, from compliance to player safety, warrants a measured approach.




Regulators and Industry Are Not on the Same Page.

Regulators and operators disagree on where AI is being used, regulators lack confidence in their oversight capabilities, and they do not believe the industry is capable of self-regulating its use of AI. And the data corroborates this — regulators are largely unaware of Responsible AI practices among licensees, and those practices largely don't exist.



The Innovation Pipeline Is Accelerating.

Academic research publications, patent filings, product development, conference discourse are all growing — signaling that the ecosystem around AI in gambling is building momentum, even as adoption within companies remains uneven.





The State of AI in Gaming Team

Principal Investigator and Editor-in-Chief

- Kasra Ghaharian, PhD, UNLV International Gaming Institute

Executive Editors

- Rick Arpin, KPMG LLP
- Simo Dragicevic, UNLV International Gaming Institute

Editors

- Brett Abarbanel, PhD, UNLV International Gaming Institute
- Chris Percy, PhD, UNLV International Gaming Institute
- Piyush Puranik, PhD, UNLV International Gaming Institute

Researchers

- Antonio Le, UNLV International Gaming Institute
- Cesar Lozoya Martinez, UNLV International Gaming Institute
- Melissa Swinehart, UNLV International Gaming Institute
- Samantha Wells, UNLV International Gaming Institute

Contributors

- Zachary Birnbaum, Vixio
- Hasan Dajani, KPMG LLP
- Joseph Gaunt, Independent Consultant
- James Kilsby, Vixio
- Ryan Rubio, Zero Labs
- Ethan Smith, University of Pennsylvania

Acknowledgments

The State of AI in Gaming is an AiR Hub Initiative made possible by the support of our Founding Members.



We also thank Alan Feldman (UNLV International Gaming Institute) and Olivia Weiss (KPMG) for their support with public relations and communications.

How to Cite This Report

Ghaharian, K., Dragicevic, S., Arpin, R., Puranik, P., Swinehart, M., Percy, C., Le, A., Martinez, C. L., Smith, E., Wells, S., Abarbanel, B., Birnbaum, Z., Kilsby, J., Dajani, H., Gaunt, J., & Rubio, R. (2026). The State of AI in Gaming 2026. AiR Hub at the International Gaming Institute, University of Nevada, Las Vegas. <https://doi.org/10.17605/OSF.IO/V8UWP>

The State of AI in Gaming 2026 is licensed under [CC BY-ND 4.0](https://creativecommons.org/licenses/by-nd/4.0/).

Connect With Us

We welcome feedback, suggestions, and new ideas that can help us improve and refine future iterations of this report. Whether you have specific criticisms, inquiries regarding research collaboration, or suggestions for metrics we should track in our next edition, we want to hear from you. Contact us at igi@unlv.edu.



Contents

- Introducing The State of AI in Gaming i
- At a Glance ii
- Navigating the Report..... 1
- Executive Summary..... 2
- Section 1: The Industry Index..... 6
- Section 2: Regulatory Pulse 31
- Section 3: Research and Advancements 48
- Section 4: 2026 Thematic Focus - Responsible AI 75
- Appendix..... 95



Navigating the Report

In deciding how to structure the report, our primary driver was to ensure we provided data and discourse along key axes of the global gambling sector. As with any research project, time and resource constraints required us to refine our scope to a manageable focus. While our initial framework included a comprehensive set of stakeholders and topics, we ultimately narrowed this first edition to three primary sections: **(1) The Industry Index**, **(2) Regulatory Pulse**, and **(3) Research & Advancements**. Additionally, a fourth **Thematic Focus** section will change annually to provide commentary on the most salient topics of the year. This year our focus is **Responsible AI**, one of the key pillars of the AiR Hub and something the editorial team felt was an apt starting point given the intersection of two highly scrutinized fields.

As noted in the **Message from the Editor-in-Chief**, we hope readers will engage with us so we can gather feedback that will help us prioritize future stakeholders and topics. Just like AI itself, the report will continue to refine and expand as needed in future iterations.

While the report is designed to be read “front to back”, each section, and even the subsections within them, can be enjoyed as standalone pieces. We have also included an **At a Glance** section immediately following the **Message from the Editor-in-Chief**, offering five essential takeaways from the entire project. For those seeking more depth in a condensed format, our **Executive Summary** offers thorough highlights from each section. Finally, every section begins with a dedicated overview and a mini-Table of Contents to help you navigate the findings with ease.

A Note on Terminology Use

Throughout this report, readers will see a variety of terms used to reflect gambling and gambling-related activities. Such terms include, but are not limited to gambling, gaming, wagering, and betting. We have used the terminology most reflective of the respective section in which it appears. For example, if a particular jurisdiction refers to “wagering” in its legislation or regulation, that is the term we use in describing the gambling-related activity. Where a particular term is not defined by an organization, we default to “gambling” or “gaming.”

A Note on Research Methods and Transparency

Within each section, we have included a condensed version of our methodologies to facilitate better readability, while providing more extensive methodological details in the **Appendix**. Throughout the report, you will encounter various data and visualizations. We have listed the source of data to accompany each visualization. In addition, we have highlighted the number of respondents (*n*) for all survey data as well as provided details on the respondent profiles. While we are unable to share raw respondent data due to Institutional Review Board (IRB) protocols, we have made our datasets for **Section 3: Research & Advancements** publicly available on the Open Science Framework (OSF). Data and other materials for the 2026 edition can be accessed at the following link: <https://doi.org/10.17605/OSF.IO/V8UWP>. Lastly, to facilitate easier information retrieval for our readers, references have been compiled at the end of each respective section.

Executive Summary

The State of AI in Gaming 2026 is the inaugural edition of an annual report produced by the University of Nevada, Las Vegas (UNLV) International Gaming Institute (IGI), in collaboration with KPMG. It is the flagship project of UNLV IGI's Artificial Intelligence Research Hub (AiR Hub). Drawing on original survey research, secondary data analysis, and expert contributions, this report establishes the first comprehensive baseline for how AI is shaping the global gambling industry.

In this inaugural edition we study the impact of AI across four dimensions: industry adoption, regulation, innovation, and responsible use.

Section 1: The Industry Index

This section presents findings from our original survey of 83 gambling companies across the globe, including operators and suppliers across land-based and online sectors. Using a one-response-per-organization design, we capture where AI is being used, how mature organizations are, whether AI is delivering value, what the biggest barriers are to scaling, and how the workforce will be impacted.

- 1 The industry has ambitions, but our Maturity Index suggests it is still in a state of development - and governance of the technology is a cause for concern.** The industry averages 45 out of 100 on the composite maturity index built on four dimensions: Strategy, Infrastructure, Expertise, and Governance. Strategy scores highest at 57, with nearly half of companies reporting a formal AI strategy actively guiding initiatives. However, Governance trails at just 30 — most organizations have no established governance practices or are in early stages of development.
- 2 Land-based operators trail their online counterparts by a statistically significant margin.** This gap was substantial across strategy, infrastructure, and expertise, in particular. While this likely reflects the challenges of integrating AI into legacy systems and complex physical environments, investing in a formal strategy alongside expertise and infrastructure could be an avenue to close the gap.
- 3 AI activity across the industry is concentrated in technology, security, and product innovation, which together account for nearly half of all weighted activity.** Customer-facing functions and corporate operations sit in the middle tier, while risk and compliance trails. Marketing use cases were less prevalent, suggesting that customer-facing AI applications — despite their commercial potential — have not yet become a clear priority.
- 4 The industry has embraced Generative AI, but the agentic AI era is yet to materialize.** Text and code generation lead among GenAI applications. But in contrast to broader enterprise trends, adoption of agentic AI in gambling appears markedly lower.
- 5 Cost reduction is the primary driver of AI adoption, but ROI remains elusive for most.** Only one in five companies report already achieving meaningful returns, with most expecting ROI within two years. However, the majority rely on less structured evaluation methods to assess AI success, and one in four have no structured evaluation process at all, which could make it difficult to know whether returns are or will be realized.
- 6 Knowledge and training gaps and resource constraints are the most cited barriers to scaling AI, and cybersecurity and data privacy rank as top risk concerns.** Despite this, nearly half of organizations report no AI-specific hiring planned, and the governance structures needed to address risks remain underdeveloped — a minority of organizations have dedicated AI governance or ethics roles, with few planning to hire into them.

Section 2: Regulatory Pulse

In this section, we examine how AI governance is evolving across gambling jurisdictions, combining a regulatory scan of legislative and policy developments over the past 36 months — compiled by Vixio — with an original survey of 113 gambling regulatory executives and staff worldwide.

- 1 Regulatory activity surrounding AI in gambling is concentrated in North America and Western Europe.** Gambling-specific AI regulatory developments are geographically narrow, with no recorded activity across Central and Eastern Europe, Asia-Pacific, or the Middle East and Africa.
- 2 Regulatory action is prioritizing the use of automation to support player protection and safety.** The most prevalent regulatory actions involve the use of AI and automated tools for player harm detection, with jurisdictions including Belgium, Finland, Italy, and several U.S. and Canadian provinces mandating or encouraging automated monitoring systems.
- 3 Regulation to govern AI specifically is emerging, but remains the exception.** A small number of jurisdictions have begun explicitly governing AI use in gambling. Whether this signals a broader trend away from addressing AI solely through existing outcome-based frameworks remains to be seen - though insights from our regulatory survey suggest it might. Regulators in our sample tended to agree that gambling-specific AI regulations are needed and that current frameworks are not equipped to address the risks and opportunities posed by AI.
- 4 Regulators report limited visibility into how licensees are using AI and low confidence in the industry's ability to self-regulate.** Across our survey, we found that regulators were not confident in their ability to assess licensee use of AI and reported limited awareness of current AI applications in licensee operations. This is further corroborated by a notable disconnect between regulator perceptions and industry-reported activity.
- 5 A regulator-industry disconnect is evident in where AI is perceived to be used.** Regulators most commonly identified customer-facing functions as the primary area of licensee AI activity, yet our industry survey found that AI adoption is concentrated more heavily in technology, security, and product innovation. Regulators may have an incomplete picture of how the technology is actually being deployed, which hinders their ability to act and risks misaligned regulatory action.
- 6 One area where perspectives are complementary: Responsible AI practices in the industry remain underdeveloped.** Regulators were overwhelmingly unaware of Responsible AI practices among their licensees — a finding that is consistent with the governance gap identified in our Maturity Index, where governance scored the lowest of all four maturity dimensions at 30 out of 100.

Section 3: Research & Advancements

In our Research & Advancements section we track the gambling industry's AI innovation pipeline through five lenses: a 15-year bibliometric analysis of academic publications, conference discourse from 2020–2025, patent filing activity, trade news trends, and an expert perspective from Zero Labs on startup activity.

1 AI-related gambling publications have grown significantly, roughly doubling their pre-2020 trajectory. However, AI's share of gambling research remains modest compared to fields like healthcare and finance. Sports Betting and Online Problem Gambling have overtaken Poker as the dominant research topics.

2 AI-related sessions at major industry gambling conferences surged from 3 in 2020 to 81 in 2025. Compliance (RG, KYC, AML) and Marketing & CRM are among the most discussed topics — notably different from the use case patterns in **Section 1: The Industry Index**, where technology, security, and product innovation emerge as the most active areas of company AI use.

3 Patent activity signals accelerating R&D investment. Annual AI-related gambling patent grants rose from 15 in 2010 to 100 by 2025. Suppliers — in particular, Angel, IGT, Light & Wonder — lead filing activity, with the United States accounting for 61.4% of all filings worldwide.

4 AI in gambling trade news reflects a maturing landscape, with technology and policy as the leading categories. Product innovation — particularly in live dealer and prediction markets — is an emerging theme, while CRM and customer support also receive prominent coverage.

5 Startup activity and venture capital are flowing to companies at the intersection of data, AI, and gambling — but structural barriers remain high. Regulatory complexity and incumbent distribution advantages continue to favor established players. However, a growing cohort of data-native startups applying AI to specific segments of the value chain suggests the conditions for disruption may be forming.

Section 4: 2026 Thematic Focus - Responsible AI

With this inaugural thematic deep dive, we examine the state of Responsible AI in gambling — from incident tracking and regulator commentary to survey data on industry’s practices.

- 1 Tracked AI incidents in the gambling sector increased sharply after 2023.** Deepfakes used to encourage illegal or inappropriate gambling are the most common category. Other incidents include targeting too aggressively and failures of player harm detection systems.
- 2 Regulators are beginning to pay attention to AI as a topic in its own right, but most treat it as simply another technology to be managed through existing frameworks.** Regulatory action or high-profile reporting of AI incidents may serve as a catalyst for more explicit governance requirements — a dynamic already visible in certain jurisdictions identified in our regulatory scan
- 3 Almost a third of gambling companies report having no established Responsible AI practices, and while others report some form of Responsible AI practices, depth is lacking.** A third describe their practices as initial or basic. Another third have practices either operationalized or integrated across initiatives. Just 2% describe Responsible AI as embedded throughout their organizations.
- 4 But self-reported Responsible AI maturity does not appear to translate into deeper practice across all Responsible AI dimensions.** We found no meaningful relationship between organizations’ overall Responsible AI maturity stage and their self-reported implementation of specific Responsible AI aspects such as transparency, accountability, or reliability — with the exception of data privacy, which likely reflects pre-existing regulatory obligations rather than deliberate Responsible AI investment, and fairness/bias, which has had higher profile in public discussions of RAI.
- 5 Dedicated Responsible AI roles remain scarce and ownership is unclear.** Under 20% of companies have AI-related roles specialized in governance. When it comes to oversight, about one in four organizations share responsibility across multiple functions, while roughly 15% report no formal oversight at all — suggesting that for many companies, Responsible AI has no clear organizational home.



Section 1: The Industry Index

Section 1: The Industry Index

Overview

Breakthrough advancements in AI feel like near-daily news. While the frontier labs charge forward at break-neck speed, touting substantial gains on popularized benchmarks, the real-world business impact of AI remains an open question. A study by MIT [1] last year challenged the narrative around large-scale enterprise investment into Generative AI, concluding that the potential revolutionary impacts of AI have not yet resulted in ROI of a similar scale. Some use this and similar observations to raise the possibility that there is too much hype around AI. Yet, others are more optimistic. A recent Wharton [2] study, for example, provides evidence that the majority of enterprise leaders report positive returns on their AI investments.

So how are gambling and gaming businesses doing with AI? Is AI simply a fancified writing aid and productivity booster, or are gaming businesses embedding it within broader organizational infrastructure as part of a deliberate strategy to transform how they operate?

In this section, we present findings from a targeted industry survey designed to seek answers to these exact questions, and more. We explore how and where AI is actually being used, how organizations are thinking about AI strategy, perceptions about the future of the workforce, the biggest challenges and risks, and how governance efforts are taking shape.

With a one-response-per-organization survey design, each submission offers a reflective view of where an organization truly stands on AI. As the first iteration of this survey, our findings establish a baseline — one we intend to revisit in subsequent years, which will give our audience objective, longitudinal tracking of how the industry’s relationship with AI evolves over time.

- Overview 7
- 1.1 Methodology 8
- 1.2 Respondent Profile 9
- 1.3 Findings 11
- 1.4 Conclusion 29
- References 30



1.1 Methodology

The Industry Index is based on an original survey study independently conducted by IGI with ethical approval from the UNLV Institutional Review Board. KPMG provided feedback on survey design, recruitment support, and collaboration in interpreting and contextualizing aggregated findings.

We surveyed companies across the gambling industry's full ecosystem, spanning both land-based and online sectors and including operators, suppliers, game studios, payments providers, compliance firms, and more. We contacted over 300 companies through direct outreach via IGI and KPMG professional networks, supplemented by LinkedIn recruitment. Data collection ran from late December 2025 through March 2026.

The survey used a one-response-per-organization design, meaning each company submitted a single response that reflected its organization-wide AI posture. Respondents were typically senior leaders (e.g., C-suite, VPs, etc.) or teams that consulted internally to ensure their answers represented the broader organization. In exchange for participation, each company receives a private benchmark report comparing their responses to aggregate results.

The survey covered six core areas: (1) Usage, (2) Strategy and Maturity, (3) ROI and Business Value, (4) Barriers and Risks, (5) Workforce and Training, and (6) Responsible AI and Governance, alongside organizational characteristics such as operational region, revenue, and employee size. Given the exploratory nature of the research, we utilized descriptive statistics and nonparametric tests for analysis, conducting group comparisons across company size, revenue, role (operator vs. supplier), sector (land-based vs. online), and AI strategy adoption. Further methodological details including details on survey development and analysis processes are available in **Appendix A.3.1**.

To provide additional context, we draw on cross-industry benchmarks throughout this section where our findings invite such comparisons. In particular, we reference KPMG's Q1 2026 AI Quarterly Pulse Survey [3], a quarterly survey of senior executives across major U.S. enterprises spanning technology, financial services, healthcare, and other sectors (n = 237). We acknowledge that the survey populations, question designs, and industry contexts differ, thus comparisons should be interpreted with caution. We provide these comparisons as reference points to help situate our findings about the gambling industry's AI trajectory.

1.2 Respondent Profile

We included 83 responses in our analysis. The distribution of respondents by sector, role, employee size, revenue, and primary operational region is displayed in Figures 1.1-1.4.

Figure 1.1 - Land-based vs Online Distribution of Operators and Suppliers (N=83)

Source: AiR HUB Industry Survey




Figure 1.1 illustrates the sample characteristics categorized by sector (online vs. land-based) and role (operator vs. supplier). Our sample includes 39 operators, split between land-based (including brick-and-mortar casinos, lotteries, and integrated resorts) and online (including iGaming, online sportsbooks, and DFS platforms), accounting for 71.8% and 28.2% respectively. Our supplier subsample comprises 44 companies, spanning a range of provider types including manufacturers, game studios, data providers, payments providers, and compliance firms, with an even split between those primarily serving land-based and online sectors (50.0% each).

Figure 1.2 - Company Workforce size (N=83)

Source: AiR HUB Industry Survey





Figure 1.3 - Annual Company Revenue (N=83)


Source: AiR HUB Industry Survey



Figures 1.2 and 1.3 present the size of participating companies by employee size and revenue. Our sample spans companies of varying sizes: 50.6% (n = 42) have fewer than 1,000 employees, while 47.0% (n = 39) have 1,000 or more, with 2.4% (n = 2) preferring not to disclose. Revenue figures show a similar spread, though with a higher non-disclosure rate: 54.2% (n = 45) report annual revenue below \$500 million and 32.5% (n = 27) report \$500 million or above, with 13.3% (n = 11) preferring not to disclose. The sample includes a number of the world’s largest gambling companies by market capitalization, alongside a strong representation of mid-market and private firms across both land-based and online sectors.

Figure 1.4 - Distribution of Industry Jurisdictions by Continent (N=150)

Source: AiR HUB Industry Survey



Geographically, the sample has a global footprint (Figure 1.4), with the strongest representation from North America (38.0%), Europe (22.0%), and Oceania (12.0%), followed by Asia (11.3%), South America (10.0%), and Africa (6.7%). Respondents identified their primary operational region(s) via an open-ended question; we then coded responses and tallied multi-region answers (e.g., “global” or “US and Europe”) across all applicable regions. As respondents could indicate multiple regions, the total number of regional tallies (n = 150) exceeds the number of participating organizations (n = 83).

1.3 Findings

Functional Footprint: AI Across the Value Chain

The first section of the survey examined where AI is being deployed across business operations. We asked respondents to rate their level of AI activity — from none to minimal, moderate, or significant — across five business function areas: (1) Customer-Facing Functions, (2) Product Development & Innovation, (3) Risk & Compliance, (4) Business & Corporate Operations, and (5) Technology & Security. For those indicating at least minimal activity in a given area, we also asked respondents to identify specific use cases within that function.

The treemap in Figure 1.5 provides a visual overview of activity across these functions. The size of each of the five main tiles reflects a weighted activity score, where minimal activity was assigned a weight of 1, moderate a weight of 2, and significant a weight of 3. These were normalized so that the five main tiles (for each business function area) sum to 100%. Within each main tile, the smaller nested tiles represent the relative endorsement of specific use cases within that functional area, summing to the proportion of their parent tile.

Figure 1.5 – AI Deployment Across Operations: Distribution & Significance

Source: AiR HUB Industry Survey



In terms of functional areas, Technology & Security (24.5%) and Product Development & Innovation (24.0%) lead AI activity across the sample, together accounting for nearly half of all weighted activity. Customer-Facing Functions (18.2%) and Business & Corporate Operations (18.7%) sit in the middle tier, while Risk & Compliance (14.6%) trails as the lowest-endorsed functional area. Unlike retail or media, where customer-facing AI is often the leading use case, gambling operators appear more closely aligned with regulated financial services, prioritizing technology, security, and product integrity ahead of direct player engagement. This may reflect both regulatory scrutiny and the higher downside risk of AI-driven personalization in behavioral environments.

We found some sector and role differences in select areas. Online companies reported meaningfully higher activity in Customer-Facing Functions and Product Development & Innovation compared to their land-based counterparts, with the digital environment likely facilitating AI implementation. Of those deploying AI in Product Development and Innovation, online companies were focused on both Game Development (46.7%) and Sports (40.0%) compared to land companies, who were primarily focused on Game Development (36.4%) compared to Sports (6.8%). Suppliers reported higher activity in Product Development & Innovation relative to operators, which is consistent with the nature of their businesses (i.e., building technology and products)¹. We observed no meaningful differences by company size.

Notably, we found no meaningful difference between operators and suppliers on Customer-Facing Functions, which is perhaps surprising given operators' direct-to-player relationships. Activity in marketing-related use cases such as player acquisition, CRM, and retention appears broadly on par with areas like Fraud and Player Safety, which are arguably less commercially driven. This may suggest that customer-facing AI applications, despite their commercial potential, have not yet become a clear priority for operators.

These patterns are somewhat similar with broader enterprise trends. In the KPMG Pulse Survey, when asked to select all areas in which their organization will allocate AI budget, the most commonly cited were cyber and data security (65.8%), operations (65.4%), IT infrastructure (61.6%), and innovation and R&D (61.2%) — followed by customer experience (52.3%) and risk and compliance (44.7%). While the measures differ (our survey captures reported activity intensity across five areas; the Pulse Survey captures planned budget allocation across nine areas), the directional alignment is notable: technology and innovation functions attract greater attention, while customer experience and risk and compliance are lower.

What types of AI are being used?


While the previous finding focused on operational areas of usage, we also asked exactly what types of AI are being used within respondents' organizations. As shown in Figure 1.6, Generative AI (e.g., content creation, code generation, and retrieval-augmented generation) leads adoption by a considerable margin, with 81.5% of respondents reporting its use. Conversational AI (e.g., chatbots and virtual assistants) ranked second (66.7%), followed by predictive AI (e.g., forecasting, pattern recognition, and risk

¹ Notably, almost 60% of respondents who answered this question selected "other," suggesting that our predefined use case options did not fully capture the breadth of AI applications within Product Development & Innovation. Open-ended responses revealed additional use cases including building proprietary SDKs, data platform development, and internal web development, alongside research-oriented applications such as market research, concept drafting, and product road mapping. Some of these responses sit at the boundary between Product Development & Innovation and other categories (e.g., Business & Corporate Operations), reflecting the reality that functional boundaries are not always clean in practice. Refining and expanding the use case taxonomy is something we intend to address in subsequent editions of this survey as our understanding of how the industry applies AI continues to develop.

modeling) at 60.5%. Computer vision (38.3%) and AI agents and reasoning systems (32.1%) are less prevalent.

Figure 1.6 - Types of AI Technologies Used (N=81)

Source: AiR HUB Industry Survey




The relatively modest adoption of AI agents in gaming appears to be in contrast to broader enterprise trends. In the KPMG Pulse Survey, nearly all respondents (98.3%) reported some level of engagement with AI agents, with a third already scaling them across multiple functions. These results do reflect larger and, perhaps, more resource-rich organizations, which could explain this gap. Notwithstanding, agentic AI could be an important frontier for the gambling industry to explore, particularly in areas where autonomous decision-making offers clear efficiency gains. However, the stakes of getting this wrong in gambling could be higher than in other industries. For example, an autonomous agent that mishandles a customer interaction, fails to detect signs of problem gambling, or incorrectly executes a compliance function could lead to significant player harm and/or regulatory consequences.

For respondents indicating Generative AI adoption, we asked a follow-up question to discern the specific technologies in use, as illustrated in Figure 1.7. Within this group (81.5% of total respondents), text generation is the most widely adopted application (92.4%), followed closely by code generation (87.9%). Image generation is utilized by 62.1%, while video (36.4%) and voice or music generation (25.8%) remain less common. We found no meaningful differences in generative AI application adoption across company size, sector, or role.

Figure 1.7 - Types of Generative AI Applications Used (N=66)

Source: AiR HUB Industry Survey



These high adoption rates for text and code align with the high-activity areas identified in the use case treemap (Figure 1.5). In customer-facing functions, companies primarily deploy Generative AI for customer service and support, utilizing tools like chatbots, automated email responses, and agent assistance. Within business and corporate operations, internal employee productivity emerged as the leading application, where Generative AI streamlines tasks such as drafting emails and searching internal knowledge bases. Finally, in technology and security, software development and testing represent the most popular use case, encompassing code generation, test case development, and documentation. This aligns with broader evidence that software-related occupations are among those most exposed to AI automation — a recent Anthropic report identified “computer programmers” and “software quality assurance analysts and testers” as two of the top 10 occupations most exposed to being automated by AI [4] — and that agentic AI in particular has significant implications for the software development lifecycle [3]. We found no meaningful differences in these use case patterns across company size, sector, or role.

How Mature is the Industry?

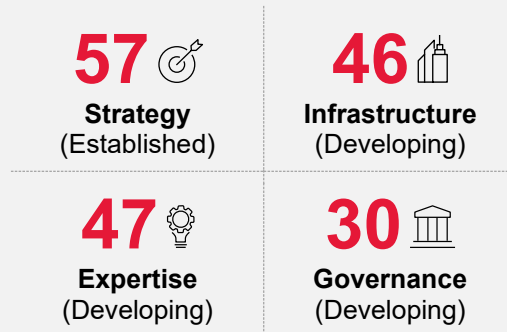
We developed a composite index for “AI Maturity” using the arithmetic mean of four key questions from our survey: Strategy (implementation level of a formal AI strategy or roadmap), Infrastructure (maturity of data, compute, and tooling), Expertise (organization-wide AI training and personnel), and Governance (implementation level of Responsible AI practices). We provide details on the index development process in **Appendix A.3.1**. Briefly, each question was rescaled to a 0-1 range and averaged to produce a composite score from 0 to 100, with scores categorized into four maturity bands: Nascent (0-24), Developing (25-49), Established (50-74), and Advanced (75-100).

Across the full sample, most companies fell in the middle bands: 33.7% (n = 28) were Established and 32.5% (n = 27) were Developing, while 21.7% (n = 18) were Nascent and 12.0% (n = 10) were Advanced. Maturity Index scores did not differ significantly by company size, whether measured by revenue or employee count (see **Appendix A.3.1** for full results).

Figure 1.8 presents average dimension scores for the overall sample and by role and sector. The most notable pattern across the sample as a whole (depicted as the shaded area in Figure 1.8) is the gap between strategy (57) and governance (30) — suggesting that while many companies have made meaningful progress on AI strategy and infrastructure, Responsible AI practices have not kept pace². Average scores for Infrastructure and Expertise were comparable at 46 and 47 respectively. The overall average Maturity Index score of 45 places the industry squarely in the Developing band.

The AI Maturity Index at a Glance

We constructed a composite AI Maturity Index from four dimensions, each scored from 0 to 100:



Overall Industry Score:

45 out of 100

(Developing)

The 27-point gap between Strategy and Governance is the index’s most notable finding. The industry is setting direction faster than it is building the safeguards to support it. This pattern threads through our findings on barriers, workforce, and Responsible AI practices, and is explored further in **Section 4: 2026 Thematic Focus - Responsible AI**.

² The gap between strategy and governance may partly reflect the broader regulatory context. The innovation-first approach emerging in the US and the EU’s recalibration of the AI Act suggest a shift in focus toward enabling innovation, which may reduce the perceived urgency for internal governance. The relatively limited development of gambling-specific AI regulation across most jurisdictions may also compound this further. Both of these factors are covered in **Section 2: Regulatory Pulse** that follows.

Figure 1.8 – Average AI Maturity by Role and Sector

Source: AiR HUB Industry Survey



Analyzing the scores by role and sector reveals further nuance. Online operators tend to score higher across most dimensions, while suppliers are relatively comparable regardless of sector, though online suppliers show somewhat higher infrastructure scores compared to their land-based counterparts. Land-based operators stand out as the lowest-scoring group, with their profile falling entirely within the overall sample average across all four dimensions, as illustrated by the red line sitting fully inside the shaded area in Figure 1.8. Statistical testing confirmed that the gap between land-based and online operators is significant and large in magnitude ($p = 0.006$, large effect size), as is the gap between land-based operators and land-based Suppliers ($p = 0.018$, medium effect size). By contrast, online operators and online suppliers score comparably, as do land-based and online suppliers. Full results are available in **Appendix A.3.1** (specifically, see interaction plot, Figure A1).

These results are consistent with broader challenges land-based operators face with technology adoption. Many rely on legacy systems to run core aspects of the business (namely hotel and casino operations) to ensure appropriate uptime. In many cases, particularly in Tribal settings, these systems and their related databases are maintained on site (or “on prem”) rather than in the cloud, for regulatory, sovereignty, or budgetary reasons. The sheer number of systems required to support the various lines of business and back-of-house operations at these properties - many of which provide comprehensive offerings as “integrated resorts” - further complicates AI deployment, which depends on accessible, high-quality data and modern infrastructure. Potentially compounding this is a relatively small addressable market for land-based gaming software relative to other industries outside of gaming (e.g., retail, banking). Unlike these larger, more homogeneous sectors, some of the operational requirements of land-based gaming are so unique (e.g., casino floor operations, managing hotels with several thousand rooms)

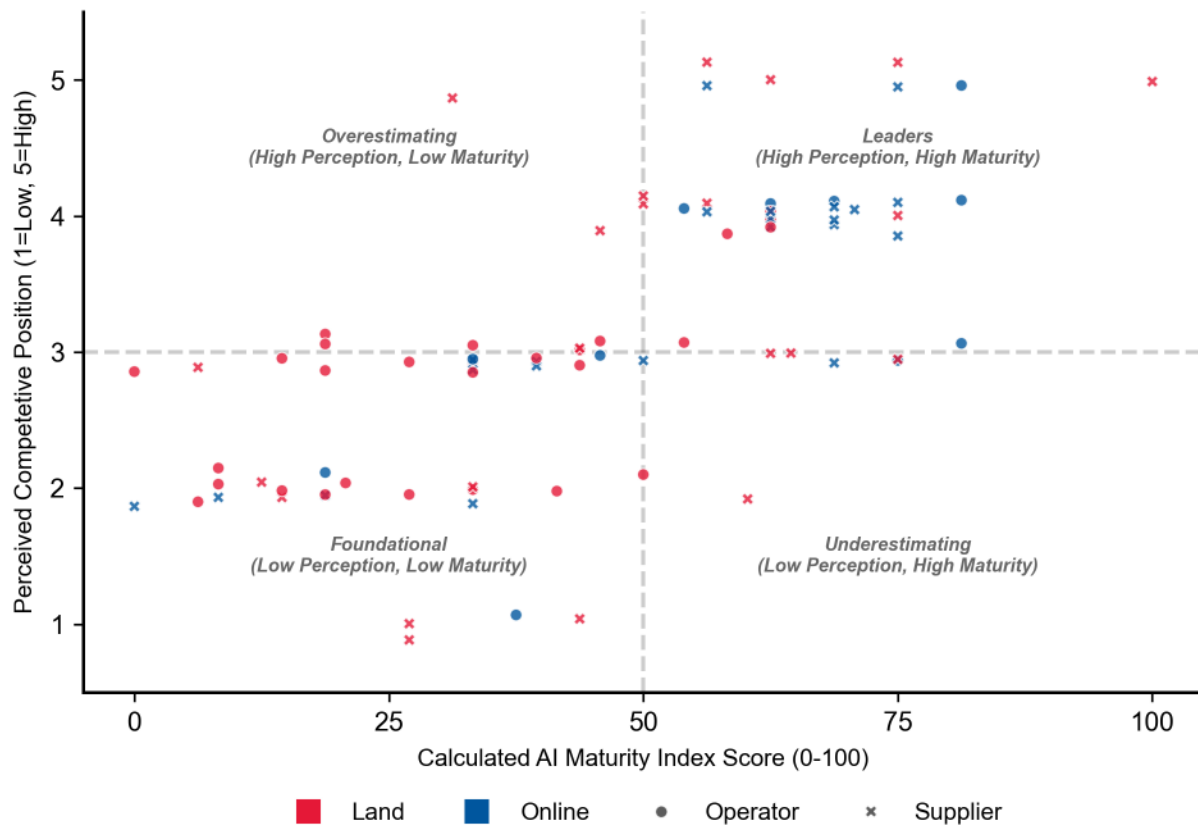
that it is not economical for software providers to invest in purpose-built solutions for the sector. Finally, regulatory frameworks governing gaming systems, many of which grew up around manual, paper-based verification processes, place constraints on the level of connectivity and automation that can be introduced, further slowing the pace of technological advancement.

AI Maturity vs. Perceived Competitive Positioning

We asked respondents to estimate how their organization compares to other gaming companies in terms of AI engagement, using a five-point scale ranging from 1 (“far behind the average”) to 5 (“at the frontier”). Figure 1.9 plots each company’s self-assessment against their AI Maturity Index score.

Figure 1.9 – Perceived Competitive Position vs. Calculated Maturity (N=83)

Source: AiR HUB Industry Survey



The quadrant view reveals four potential archetypes: Leaders (high perception, high maturity), Foundational organizations (low on both), those Underestimating (stronger maturity but modest self-assessment), and those Overestimating (high peer ranking despite limited maturity). Consistent with the maturity findings above, land-based operators tend to appear on the left-hand side of the chart, with many perceiving themselves as behind their peers, though a notable share place themselves “about average” (i.e., perceived competitive position = 3). The Leaders quadrant is visibly dominated by online companies

“Land-based casino AI evolution is compromised by fragmented, inconsistent data. This lack of data integrity is the primary barrier to AI adoption—and without it, the performance gap with online operators will continue to widen.”

Earle Hall
CEO, AXES.ai

and suppliers. Few companies fall into the Overestimating or Underestimating quadrants, which is unsurprising given that both perception and maturity were self-reported (objective external measurement would be needed to identify meaningful gaps between actual and perceived maturity). Companies rating themselves as “about average” span a notably wide range of maturity scores, suggesting that “About average” means different things to different organizations. Structured benchmarking on AI adoption and capabilities could be a valuable route for the gaming industry to help companies better understand their position and progress against peers, which is something other industries, such as financial services, have already embraced (e.g., see Evident AI’s AI Index for banking, insurance, and payments [5]).

ROI and Business Impact

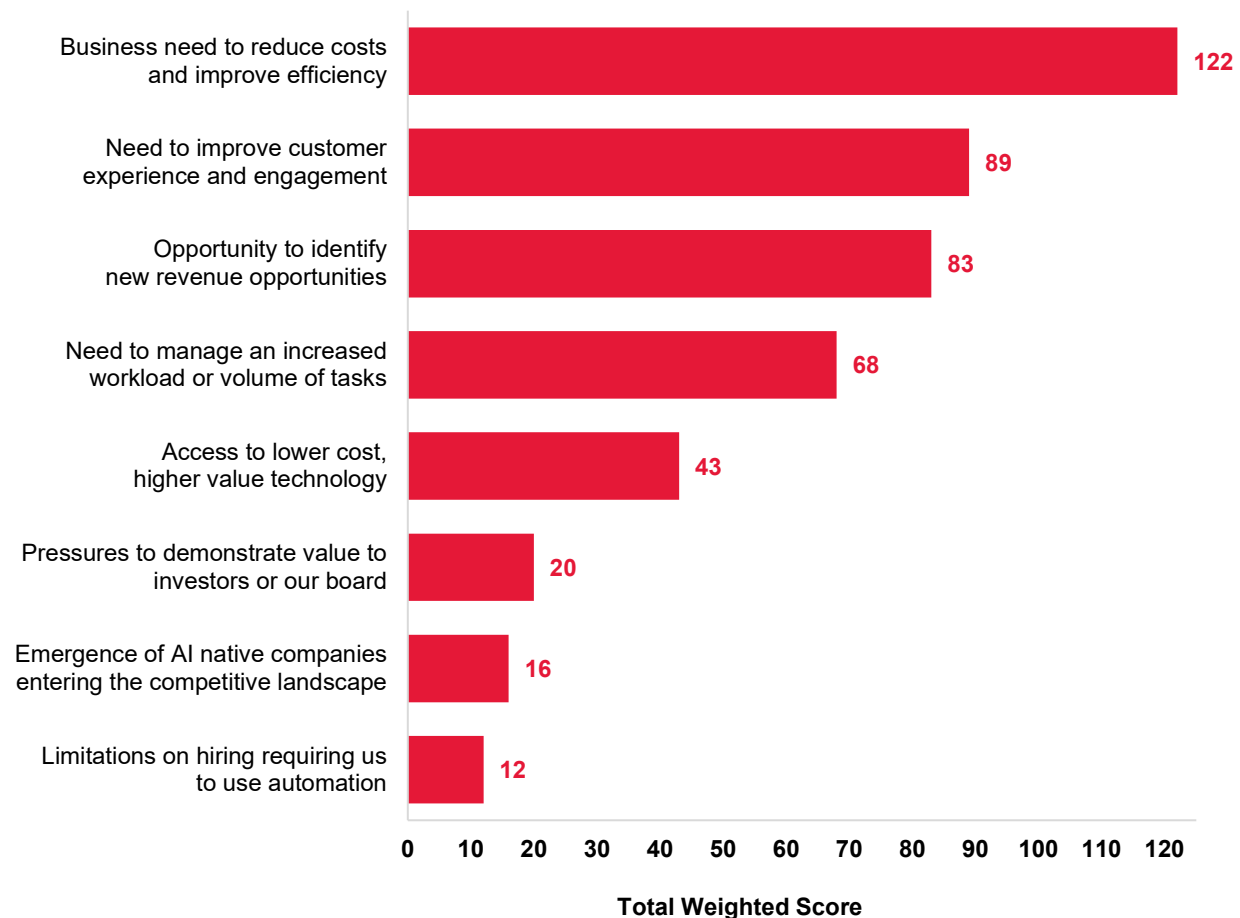
Questions about AI’s actual business value are growing louder, with talk of an “AI bubble” increasingly common. In this section we shift focus to what is driving companies’ AI adoption decisions, how organizations perceive the value AI is generating, and whether they see it primarily as a revenue driver or a source of cost savings.

What drives adoption?

We asked respondents to rank their top three drivers of AI adoption from a list of eight options. Figure 1.10 displays items ranked by weighted endorsement, where first-place selections received 3 points, second-place 2 points, and third-place 1 point.

Figure 1.10 - Top Motivations for AI Adoption (N=76)

Source: AiR HUB Industry Survey





At evoke, we view AI as a strategic driver of growth. While efficiency remains an important part of the value case, AI can also support better decision-making, enhance customer safety/experience and unlock new revenue opportunities over time. Our experience reflects these findings: organizations that translate AI investment into measurable commercial impact are those with clear strategy, rigorous measurement and a disciplined focus on the use cases that matter most.”

Kanda Kumar
Director of AI, Evoke

Across the full sample, the need to reduce costs and improve efficiency emerged as the clear top driver, followed by the need to improve customer experience and engagement, and the opportunity to identify new revenue streams³. Breaking down responses by role and sector reveals a more nuanced picture. Cost reduction ranked first for both land-based operators and land-based suppliers, while online operators prioritized customer experience above all else. Notably, the same three drivers appeared across the top three positions for all groups, with the exception of land-based suppliers, for whom managing increased workload ranked second (see Table A2 in **Appendix A.3.1**).

These priorities broadly align with cross-industry trends. In the KPMG Pulse Survey, items related to improving efficiency ranked first and second among respondents' top AI priorities: *productivity gains through automating workflows* (61.6%) and *faster, better decisions* (55.7%). However, while customer experience and revenue growth rank prominently in our sample, KPMG Pulse Survey respondents tended to prioritize *responsible AI and governance* (32.9%) and *trust and security* (31.6%) slightly more than *cost reduction* (29.5%) and *revenue growth* (24.5%). That governance and trust rank among the top four enterprise AI priorities reinforces the AI governance gap we identified via our Maturity Index (although we note that responsible AI and governance was not included as a response option in our drivers question).

Is AI delivering?

While cost reduction is the dominant motivation for AI adoption, perceptions of whether that outcome has actually been realized tell a different story. When asked the extent to which AI has contributed to cost savings, the average response was 2.43 out of 5. The majority of respondents reported minimal or no cost savings attributable to AI: 10.8% reported none at all and 48.2% reported savings only to a small extent. A further 32.5% reported moderate cost savings, while just 8.4% reported large or transformative impact. KPMG has noted in its research that while many organizations recognize that ROI may not exist currently, the “cost/risk of doing nothing” is greater...helping explain the higher (and increasing) adoption rates in their broader pulse surveys.


Responses to a similar question on revenue impact were somewhat more positive, with an average score of 3.60 out of 5. Few respondents reported negative revenue impacts (3.6% combined), and 42.2% reported no impact. However, 54.2% reported at least some positive revenue impact, with 43.4% reporting a slight increase and 10.8% a significant one — suggesting that while transformative revenue gains are the exception, a majority of companies are beginning to realize revenue gains from their AI adoption. But these figures vary by sector and role. Land-based companies were substantially more likely to report no revenue impact (56.0% vs 21.2% for online companies), and no operators reported a significant positive impact on revenue compared to roughly one in five suppliers (20.5%).

³ This is further corroborated by a separate question asking respondents to select their single primary investment driver: 42.2% cited cost reduction and 33.7% cited revenue growth, with compliance and regulatory risk management a distant third at 9.6%.

We also asked companies about their expectations for future ROI (Figure 1.11), how their AI initiatives have performed against expectations, and how they typically measure success. One in five companies report already achieving meaningful ROI from their AI investments, and a further 55.4% expect to do so within the next two years. Yet 19.3% say it is too early to determine, and 3.6% do not expect meaningful ROI at all.

Figure 1.11 - Expected Timeline for ROI on AI Adoption (N=83)

Source: AiR HUB Industry Survey



Success rates tell a similar story. We asked companies “Overall, how many of your organization’s AI initiatives have met or exceeded their initial expectations?” Just over a third of respondents report that most or nearly all of their AI initiatives have met or exceeded expectations, while nearly a quarter report that only a few have. A further 16.9% say they don’t know, which is an important finding itself.

This mixed picture becomes clearer when looking at how companies typically assess the success of their AI initiatives. Nearly two thirds (63.9%) rely on internal stakeholder feedback (an arguably less objective form of assessment) and one in four (25.3%) have no structured evaluation process at all. More robust methods are less common: 41.0% log costs and benefits, 39.8% benchmark outcomes against pre-AI processes or KPIs, and just 26.5% employ controlled testing such as A/B testing or pilot comparisons. Without robust measurement, it is difficult to know what is working, which may explain the mixed success rates and the uncertainty around ROI timing. Closing this measurement gap may be as important as the use cases themselves. In fact, companies with a formal AI strategy that is actively guiding their initiatives were substantially more likely to use benchmarking (60.0% vs 39.8%) and cost/benefit logging (60.0% vs 41.0%), and were also more likely to report already achieving meaningful ROI — suggesting that strategic clarity and measurement rigor may be important for realizing gains.

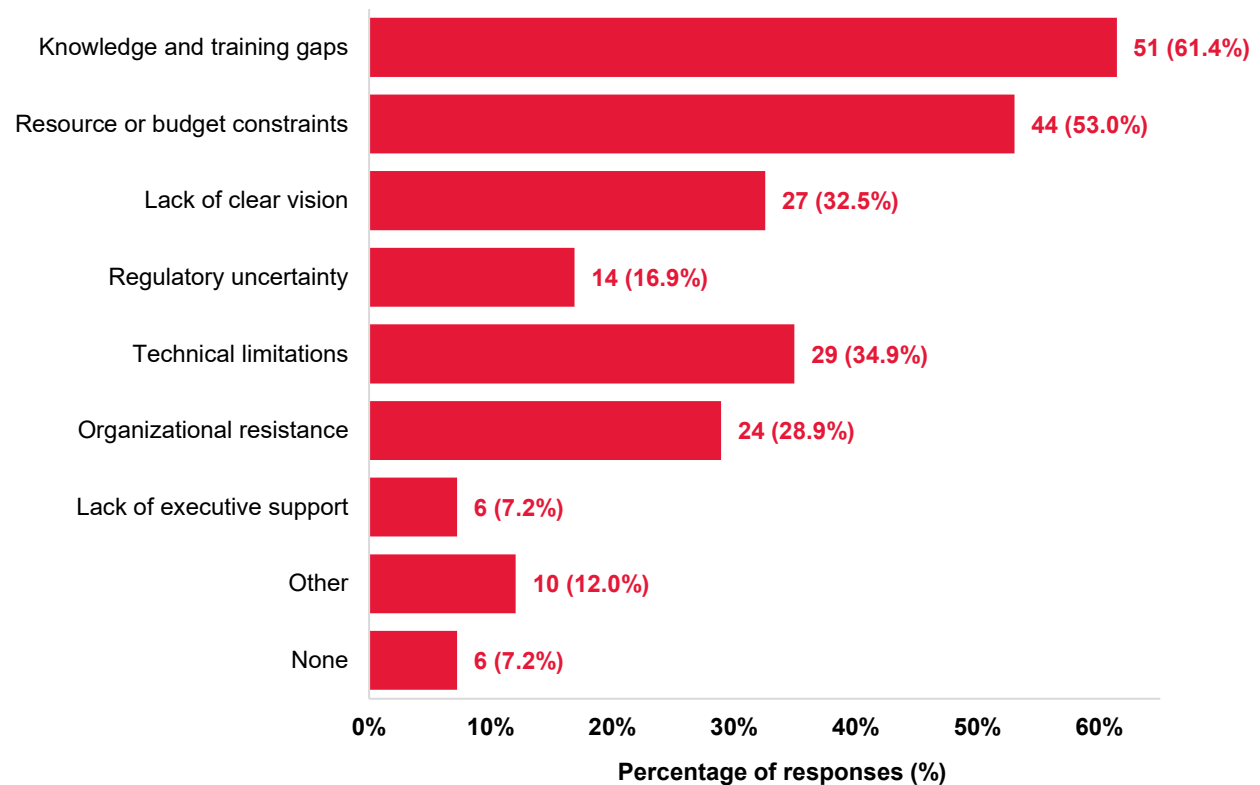
Barriers and Risks

What's holding AI back?

As indicated by the industry's average Maturity Index score of 45 out of 100, the industry is largely in an earlier developmental stage of AI adoption. Thus, understanding where barriers lie provides valuable context for stakeholders. These are presented in Figure 1.12.

Figure 1.12 - Barriers Preventing AI Scaling in the Gambling Industry (N=83)

Source: AiR HUB Industry Survey



Knowledge and training gaps emerged as the most commonly cited barrier, flagged by 61.4% of respondents, followed by resource or budget constraints (53%) and technical limitations (34.9%). Lack of clear vision (32.5%) and organizational resistance (28.9%) also featured prominently, suggesting that for some companies, the obstacles to AI scaling are as much cultural and strategic as they are technical or financial. Notably, regulatory uncertainty was cited by just 16.9% of respondents⁴, and lack of executive support by only 7.2%, indicating that leadership buy-in and regulatory clarity are less of a bottleneck than internal capability and resource gaps. Just 7.2% reported no barriers at all.

What keeps stakeholders up at night?


Beyond the practicalities of scaling AI, we asked respondents to rate their level of concern across nine potential AI risks. Figure 1.13 presents the full distribution of responses. The risks attracting the highest levels of concern were cybersecurity vulnerabilities in AI systems, failures in data privacy and

⁴ This is consistent with responses to a separate question asking the extent to which regulations are blocking or slowing AI adoption: 33.7% reported no effect and 41.0% reported only minor checks or delays, with just 10.8% reporting regulations as a significant or primary barrier.

governance, and misuse of AI by staff or third parties — all areas where the consequences of failure are concrete and potentially severe. Cybersecurity concerns may be front of mind due to recent high-profile incidents across the industry, including cyberattacks affecting MGM Resorts, Caesars Entertainment, Boyd Gaming, and Wynn Resorts [6]. Moreover, these concerns are consistent with the cross-industry picture. In the KPMG Pulse Survey, data privacy and cybersecurity ranked as the top two greatest risk barriers to enterprises' AI strategy.

Figure 1.13 - Concern Regarding potential Risks from AI in the Gambling Industry (N=83)

Source: AiR HUB Industry Survey



The high level of concern around data privacy and governance is a particularly notable finding given the relatively low scores on the governance dimension of our Maturity Index; this was the lowest of the four dimensions at an average of 30. We also observed in a separate survey item that only 22.9% of organizations have dedicated AI-related roles for AI governance, Responsible AI, ethics, or compliance. This highlights a clear disconnect: data privacy and governance rank among the industry's top AI concerns, yet the governance structures and dedicated roles needed to address them remain underdeveloped in many organizations.

Bias or unfair outcomes affecting players and problem gambling risks amplified by AI attracted comparatively lower concern levels, with a substantial proportion of respondents reporting they are not at all concerned about these issues. Interestingly, the risks that tend to attract the most scrutiny from regulators and responsible gambling advocates do not appear to weigh as heavily on industry leaders.



While the full impact of AI in this space is still evolving, ensuring its adoption is grounded in transparent, validated criteria and evidence-based frameworks will be essential to minimize bias and unintended outcomes.”

Jennifer Shatley
Executive Director,
Responsible Online Gaming Association

The People Dimension: Expertise, Roles, and Workforce Outlook

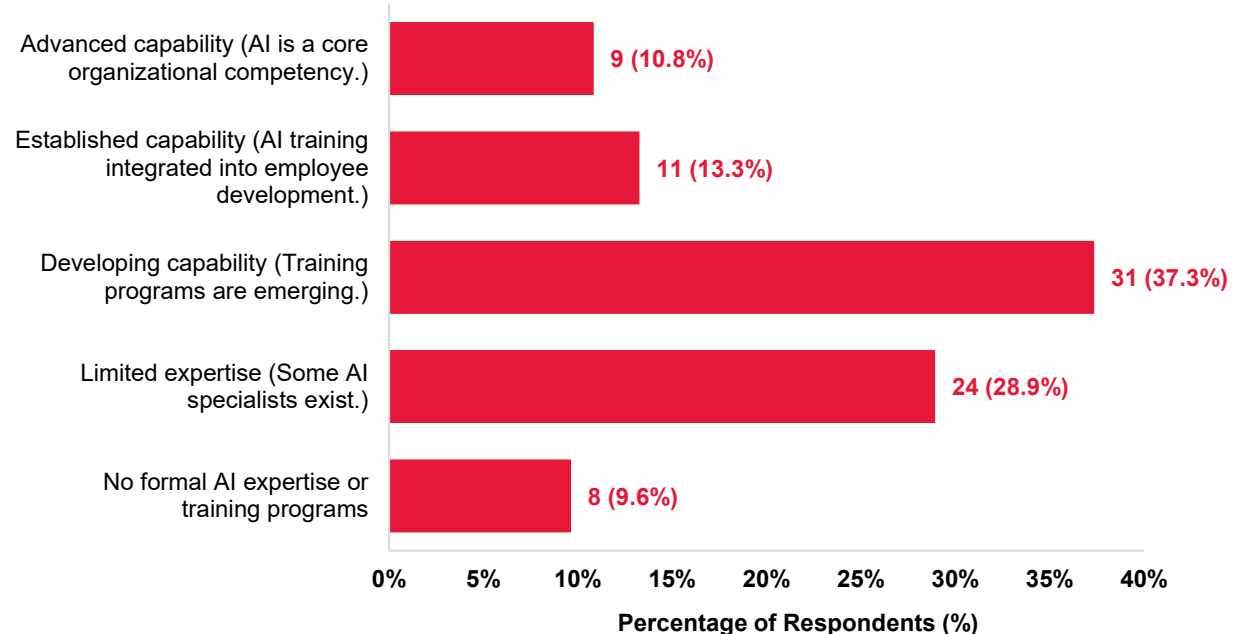
In the previous sections we highlighted that internal capability gaps, namely knowledge and training and resource constraints, represent the primary bottleneck in scaling AI. So, we now turn to the workforce itself to explore how organizations are building (or failing to build) the human capital needed to support their AI ambitions.

Where does expertise stand today?

As one of the four dimensions underlying the Maturity Index, organizational AI expertise provides a foundational lens on our sample’s workforce readiness. Across the full sample, expertise scored 47 out of 100 — placing it squarely in the Developing band and on par with infrastructure (46), though well behind strategy (57). Figure 1.14 reveals that the majority of companies are still in the early stages: 9.6% (n = 8) report no formal AI expertise or training programs, and 28.9% (n = 24) report only limited expertise with some AI specialists in place. A further 37.3% (n = 31) describe developing capability, with training programs beginning to emerge. Just one in four companies report established (13.3%, n = 11) or advanced (10.8%, n = 9) AI capability.

Figure 1.14 - Level of AI Expertise within Organizations (N=83)

Source: AiR HUB Industry Survey



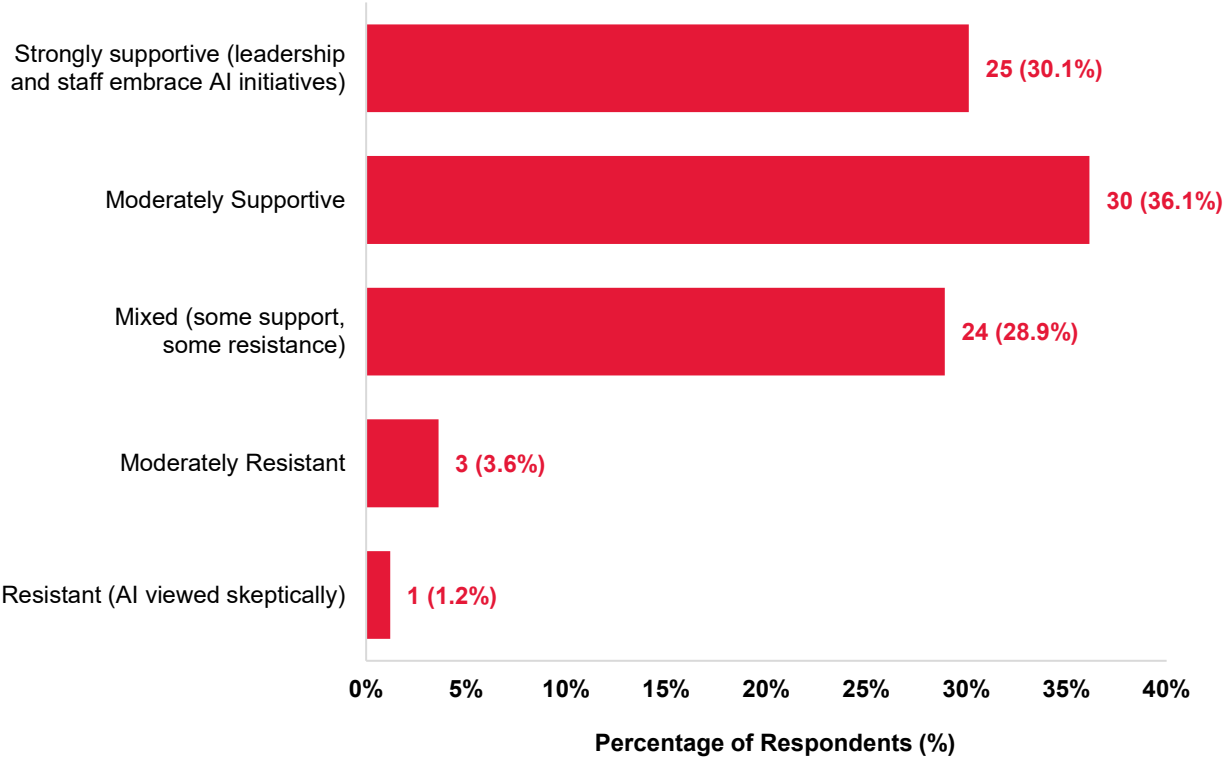
As we previously noted in the maturity analysis, land-based operators score well below average on this dimension (27 vs. sample average of 47), while online operators and suppliers report somewhat higher expertise levels (range of 55-58). Among the 20 companies reporting Established or Advanced AI capability, none were land-based operators. The remaining 20 were distributed across land-based suppliers (n = 9, representing 40.9% of that group), online suppliers (n = 7, 31.8%), and online operators (n = 4, 36.4%).

Organizational culture: A window of opportunity

While AI expertise is largely in a developing stage, the cultural climate for AI adoption is primed for progress (Figure 1.15). When asked to describe their organizational culture toward AI innovation, 66.2% of respondents reported a moderately (36.1%, n = 30) or strongly (30.1%, n = 25) supportive environment (one in which leadership and staff broadly embrace AI initiatives). A further 28.9% (n = 24) reported mixed sentiment, and only 4.8% (n = 4) described their culture as moderately or fully resistant to AI.

Figure 1.15 - Organizational Culture Towards AI Innovation (N=83)

Source: AiR HUB Industry Survey



This gap between formal capability (Figure 1.14) and broad cultural receptivity (Figure 1.15) is a notable finding. With leadership and staff largely open to AI, the conditions for accelerating AI training programs and capability-building are clearly favorable and, perhaps, an opportunity to gain competitive advantage.

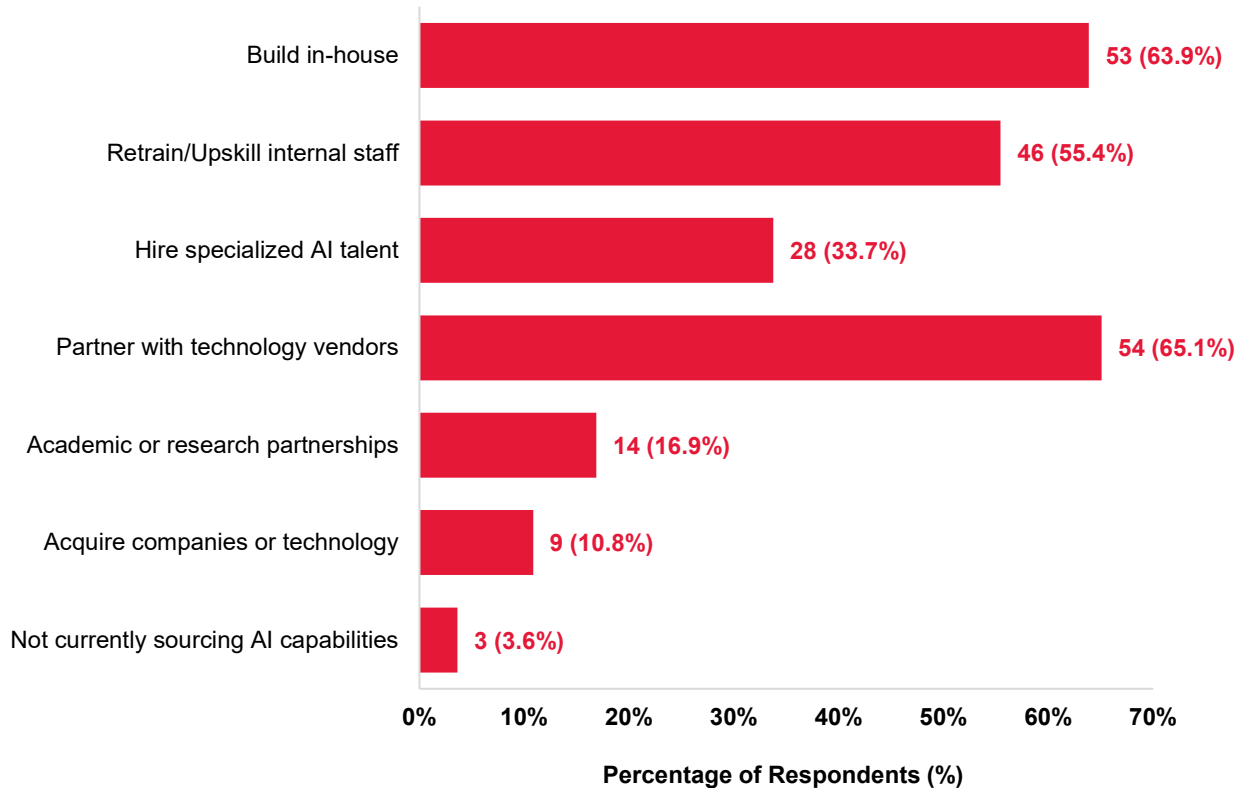
How are companies sourcing AI capabilities?

Organizations are drawing on a range of strategies to access AI capabilities, and no single approach dominates (Figure 1.16). Partnering with technology vendors (65.1%, n = 54) and building capabilities in-house (63.9%, n = 53) are the two most common approaches, followed closely by retraining and upskilling internal staff (55.4%, n = 46). Hiring specialized AI talent is less prevalent at 33.7% (n = 28), and academic or research partnerships (16.9%, n = 14) and acquiring companies or technology (10.8%, n = 9).

= 9) remain niche strategies. Just 3.6% (n = 3) report not currently sourcing AI capabilities at all. We found no notable differences across most groups, with one exception: operators appear substantially more reliant on vendor partnerships than suppliers (85% vs 48%).

Figure 1.16 - Sources of AI Capabilities in Organizations (N=83)

Source: AiR HUB Industry Survey




Who is leading AI within organizations?

The composition of AI-related roles (Figure 1.17) offers a snapshot of how the industry is structuring its AI workforce. Data scientists are the most common role, present at 60.2% (n = 50/83) of organizations, followed by AI or machine learning engineers (48.2%, n = 40/83) and AI infrastructure or platform engineers such as MLOps and DevOps specialists (36.1%, n = 30/83). AI product managers or program leads exist at roughly one in three organizations (32.5%, n = 27/83). Dedicated AI leadership (e.g., Chief AI Officer) is present at only 22.9% (n = 19/83) of companies, though this varies considerably by sector with 42.4% of online companies report having such a role compared to just 10.0% of land-based companies. And this leadership gap persists regardless of company size, with only 28.2% of large organizations (≥1,000 employees) possessing dedicated AI leadership (compared to 19.0% for companies with <1,000 employees). This is in stark contrast to broader enterprise trends; a recent Wharton report found CAIO roles are now present at 60% of enterprises [2], suggesting the gambling industry may be considerably behind in formalizing AI leadership. However, companies with a formal AI strategy actively guiding their initiatives were notably more likely to have dedicated AI leadership in place (37.1% vs 12.5%). Governance roles exist at just 22.9% (n = 19/83). Notably, 22.9% (n = 19/83) of organizations report having none of these roles.

Figure 1.17 - Composition of AI-related Roles (N=83)

Source: AiR HUB Industry Survey



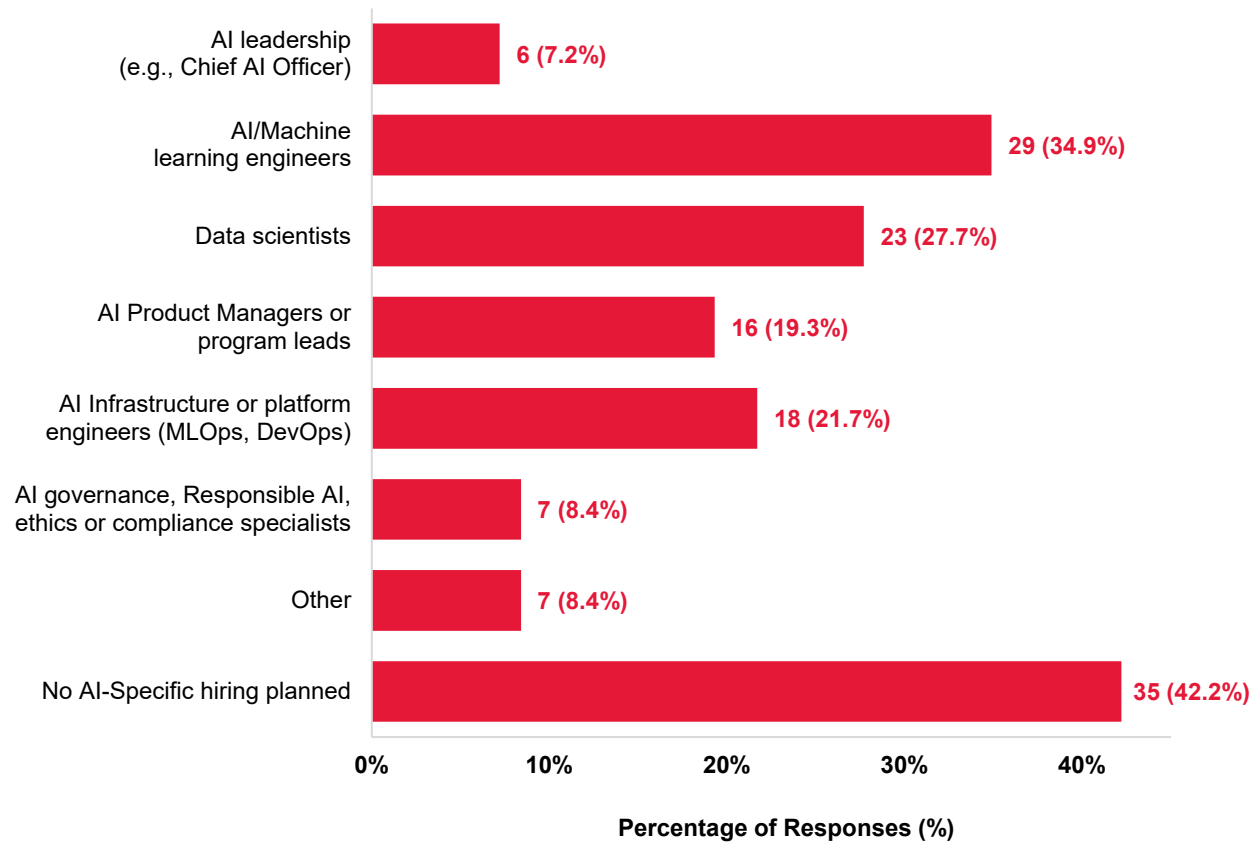
The relative scarcity of governance and ethics roles reinforces a pattern observed throughout this report. As noted in the Barriers section, data privacy and governance rank among the industry’s top AI concerns, yet the governance dimension of the Maturity Index was the lowest of the four at an average score of 30. Moreover, it is unsurprising that Responsible AI practices remain underdeveloped when the dedicated roles needed to build and sustain them are largely absent (see **Section 4: 2026 Thematic Focus - Responsible AI** for a more in-depth discussion). Organizations express concerns about governance despite their underinvestment — though we acknowledge that the directionality of this relationship cannot be established from cross-sectional data alone. It is certainly possible that this concern reflects a growing recognition of the gap, which may translate into future action. But the planned hiring landscape does not appear to support this notion.

What does the hiring pipeline look like?

Figure 1.18 displays the distribution of planned hiring activity. AI and machine learning engineers are the most commonly planned hire (34.9%, n = 29), followed by data scientists (27.7%, n = 23) and AI infrastructure engineers (21.7%, n = 18). Only 8.4% (n = 7) plan to hire AI governance or ethics specialists and just 7.2% (n = 6) plan to hire AI leadership.

Figure 1.18 - Planned Hiring for AI-related Roles (N=83)

Source: AiR HUB Industry Survey



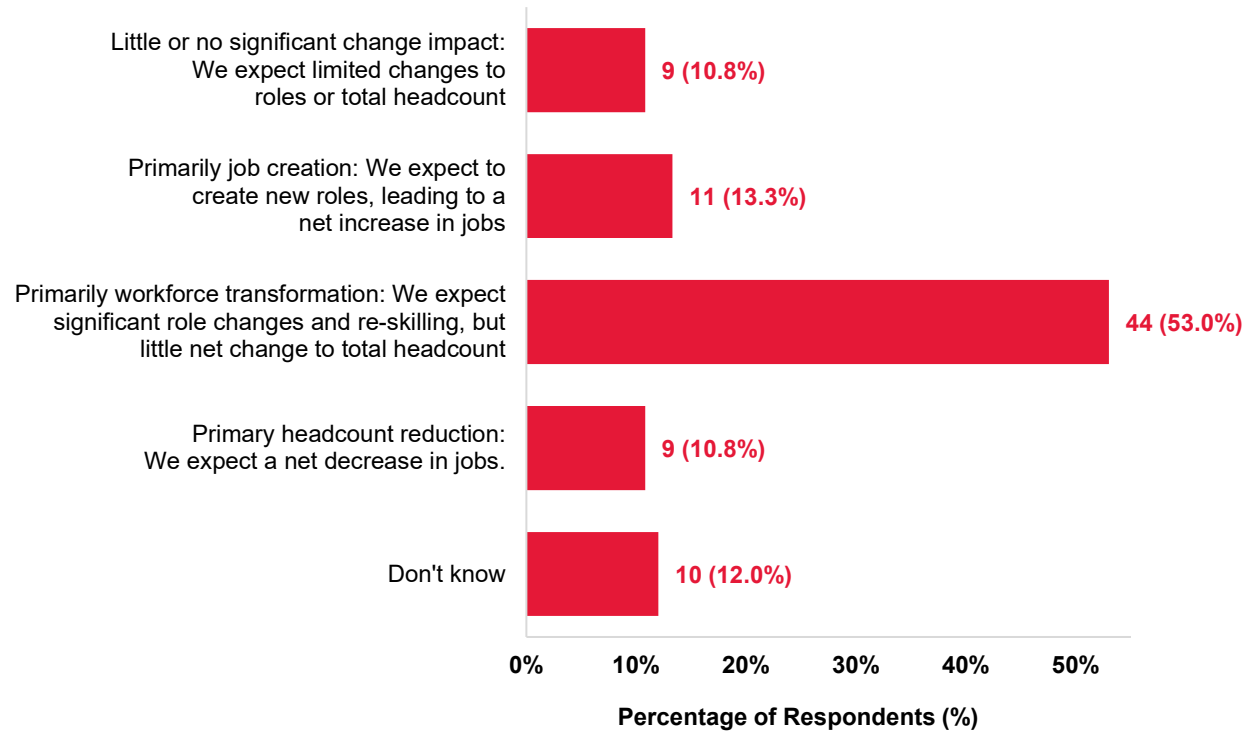
Perhaps the most striking finding is that 42.2% (n = 35) of organizations report no AI-specific hiring planned. The majority of organizations describe its expertise as limited or developing, knowledge and training gaps are the most cited barrier to AI adoption (61.4%), and yet nearly half of companies have no plans to bring in dedicated AI talent. This may not necessarily mean complacency, as some organizations could be prioritizing upskilling over new hires; which is a possibility that emerges in our next section's findings.

How will AI reshape the workforce?

We asked respondents which statement best describes how they expect AI to impact their workforce over the next three to five years (Figure 1.19). The dominant view is one of transformation rather than displacement: 53.0% (n = 44) expect significant role changes and reskilling but little net change to total headcount. A smaller share (13.3%, n = 11) anticipate net job creation through new AI-driven roles, while 10.8% (n = 9) expect a net decrease in jobs. Just 10.8% (n = 9) foresee little or no significant change, and another 12.0% (n = 10) said they simply do not know.

Figure 1.19 - Expected Future Impact of AI on Workforce (N=83)

Source: AiR HUB Industry Survey



What will this “workforce transformation” look like in practice for gambling industry organizations who invest in AI? Prior research indicates that as firms invest more in AI, they tend to shift toward more educated and technically skilled workforces [7]. Furthermore, as AI investment increases organizational hierarchies may flatten - as individual employees gain access to AI tools and augment their problem-solving capabilities, the need for upper levels of management and dedicated oversight functions may diminish [8]. For example, one can envision the “democratization” of analytics, where departmental personnel who understand the specifics of their business area could generate data-driven insights independently, without relying on centralized analytics teams. This could be a particularly compelling strategy for integrated resort businesses, which house diverse and distinct business units (e.g., casino marketing, housekeeping, food & beverage, etc.), each with their own operational data and decision-making needs. It will be interesting to see whether this broadly neutral workforce outlook holds in the coming years. Given that the industry’s primary driver of AI adoption is cost reduction and improved efficiency, advancements in AI capabilities combined with potential regulatory and taxation pressures in certain jurisdictions could shift these perceptions over time.

1.4 Conclusion


The findings from this inaugural edition of our survey establish an important baseline and reveal an industry that is still largely finding its footing with AI. With an average AI Maturity Index score of 45, the gambling industry sits squarely in the Developing band, yet, ambition is clearly present. Strategy was the highest-scoring dimension at 57, nudging into the Established range. And among online operators, the average strategy score reached 76, placing this segment firmly in the Advanced category and reflects a clear commitment to AI adoption. Land-based operators were furthest behind across the index, trailing online counterparts by a statistically significant and large margin. As discussed, this likely reflects structural realities (e.g., legacy infrastructure, diverse operating environments, regulatory constraints) that speak to broader technology adoption challenges facing this segment.

A key finding was the governance gap. Governance was the lowest-scoring dimension of the Maturity Index at an average of 30, which is well behind strategy (57), expertise (47), and infrastructure (46). In fact, just under a third of companies reported no established Responsible AI practices. The reasons for this gap are likely multifaceted. For example, if what Responsible AI entails and what it addresses is not well understood, companies' investment in this area is unlikely to be seen as a priority. The regulatory environment may also be a contributing factor: as we discuss in **Section 2: Regulatory Pulse** that follows, there are open questions about whether regulators believe industry is capable of self-regulating AI, and in the absence of clear compliance obligations, organizations have less impetus to act. The scaling back of broader AI governance requirements at the supra-national level could also be playing a role, signaling to industry that there is no urgency for them to implement AI governance practices.

In closing, it appears from our findings that strategic intent matters. Organizations with a formal AI strategy actively guiding their initiatives are more likely to have dedicated AI leadership, appear more rigorous in how they measure success, and are more likely to report that they are already achieving meaningful ROI from AI initiatives. But while strategy was the highest-scoring dimension of the Maturity Index, expertise and infrastructure lagged behind — knowledge and training gaps remain the most cited barrier to adoption, nearly half of organizations do not plan to hire AI-specific roles, and few are looking to add dedicated AI leadership. The gambling industry has ambitions for AI, but whether investment in people and infrastructure follows this strategic intent will be interesting to track in next year's edition.

References

- [1] A. Challapally, C. Pease, R. Raskar, and P. Chari, "STATE OF AI IN BUSINESS 2025," MIT, Jul. 2025.
- [2] Korst, Jeremy, Puntoni, Stefano, and Tambe, Prasanna, "Accountable Acceleration: Gen AI Fast-Tracks Into the Enterprise (2025 AI Adoption Report)," Wharton Human-AI Research & GBK Collective, Philadelphia, PA, Oct. 2025. Accessed: Mar. 05, 2026. [Online]. Available: https://ai.wharton.upenn.edu/wp-content/uploads/2025/10/2025-Wharton-GBK-AI-Adoption-Report_Full-Report.pdf
- [3] "KPMG GenAI Study: the path to returns." Accessed: Mar. 31, 2026. [Online]. Available: <http://view.ceros.com/kpmg-design/kpmg-genai-study>
- [4] M. Massenkoff and P. McCrory, "Labor market impacts of AI: A new measure and early evidence." [Online]. Available: <https://www.anthropic.com/research/labor-market-impacts>
- [5] "Evident - Banking AI Index," Evident Insights. Accessed: Mar. 13, 2026. [Online]. Available: <https://evidentinsights.com/ai-index/>
- [6] E. Adams, J. Bischoff, and M. Erdem, "'High-Stakes' Security: Current and Emerging Trends in Casino Gaming Cybersecurity," *UNLV Gaming Res. Rev. J.*, vol. 29, no. 1, Sep. 2025, doi: 10.9741/2327-8455.1506.
- [7] T. Babina, A. Fedyk, A. He, and J. Hodson, "Firm investments in artificial intelligence technologies and changes in workforce composition," in *Technology, productivity, and economic growth*, University of Chicago Press, 2023, pp. 75–117. [Online]. Available: <http://www.nber.org/chapters/c14753>
- [8] Tania Babina and Anastassia Fedyk, "The effects of AI on firms and workers," Brookings. Accessed: Mar. 19, 2026. [Online]. Available: <https://www.brookings.edu/articles/the-effects-of-ai-on-firms-and-workers/>



Section 2: Regulatory Pulse

Section 2: Regulatory Pulse

Overview

As AI systems transition from experimental tools to components of core business operations, questions of governance and regulatory capacity are becoming increasingly prominent. In the gambling context, regulatory dialogue around AI-specific frameworks and oversight is beginning to emerge across jurisdictions.

This section begins with a brief review of how the broader AI regulatory environment evolved in 2025 before presenting results from a scan of gambling-specific regulatory and legislative actions that address AI over the past three years. Then, to provide stakeholder perspectives, we present an original survey study with responses from individuals working in gambling regulatory authorities. The study explores key areas including respondents' personal AI usage and literacy, confidence in overseeing this emerging technology, perspectives on industry adoption, and beliefs regarding the adequacy of current governance frameworks.

- Overview 32
- 2.1 The Broader Regulatory Context 33
- 2.2 Gambling Regulatory Developments 33
- 2.3 Regulator Survey: Perspectives from Within 39
- 2.4 Conclusion 45
- References 46

2.1 The Broader Regulatory Context

2025 saw continued progress toward more concrete and operationalized AI regulatory frameworks; progress was accompanied by growing tension between formalized governance and the push for innovation.

For example, the EU continued phasing in the AI Act's provisions, with governance rules and obligations for general-purpose AI models taking effect in August 2025 [1], and drafted implementation rules for regulatory sandboxes, which each Member State is required to establish by August 2026 [2]. Beyond the EU, countries such as South Korea, Vietnam, and Brazil are developing or have enacted national AI legislation that adopts similar risk-based approaches [3]. Vietnam passed its first comprehensive AI law in December 2025 [4], drawing on the EU model, while Brazil's Senate approved an AI bill adopting an EU-like risk-based approach in late 2024 [3]. This convergence around tiered regulatory frameworks might suggest that the EU AI Act is functioning as a reference point for international approaches to AI governance, while other developments suggest signs of a softening stance toward AI regulation.

Perhaps most telling is that the EU's own regulatory position has shown signs of recalibration. In November 2025 the European Commission unveiled the "Digital Omnibus," which would delay stricter risk-management and oversight rules for high-risk AI until 2027 [5]. Elsewhere, Australia abandoned plans for standalone EU-like AI regulation in favor of voluntary guidance [6], while Canada's national AI legislation (AIDA) was terminated in Parliament in January 2025 [7], both signaling a preference for fostering innovation over prescriptive regulation [8].

This innovation-first posture is perhaps most pronounced in the United States, where tension between federal and state-level approaches has been particularly visible. Executive Order 14365 [9], signed into effect in December 2025, calls for a minimally burdensome regulatory environment for AI innovation to flourish, and establishes an "AI Litigation Task Force" commissioned to challenge state AI laws that have been deemed too "onerous and excessive." Nonetheless, several states have continued advancing AI legislation, with California's SB-53 [10] and New York's RAISE Act [11] outlining specific requirements regarding safety, transparency, and reporting for advanced AI systems.

Altogether, the 2025 global AI regulatory landscape remains muddled. For industries with interjurisdictional exposure, such as gambling, this multi-scalar regulatory environment introduces both compliance complexity and strategic uncertainty. Understanding these broader global dynamics is essential for the contextualization of industry-specific developments. The following regulatory scan examines how AI and automation requirements are taking shape across a wide selection of regulated gambling jurisdictions.

2.2 Gambling Regulatory Developments



Editorial note: The following regulatory scan was compiled by Zachary Birnbaum and James Kilsby of Vixio, a regulatory intelligence provider specializing in gambling, financial services, and payments. The scan draws on Vixio's proprietary database of regulatory and legislative developments to identify AI-related activity across gambling jurisdictions over the preceding 36 months

This regulatory scan compilation represents the most relevant regulatory and legislative actions concerning AI and/or automation specifically within the gambling and gaming industry, as captured by Vixio for the preceding 36 months (i.e. January 2023) across the following regional delineations:

- Central and Eastern Europe.
- Western Europe.
- North America.
- LatAm and the Caribbean.
- Asia-Pacific.
- Middle East and Africa.

We begin with a summary of the scan followed by specific details on AI-related regulatory action by region. Within each region, jurisdictions are listed alphabetically. Developments within each jurisdiction are presented in chronological order (oldest to most recent) to provide a running chronicle of regulatory activity. To collect relevant developments, we executed keyword searches across the Vixio database of regulatory documents. To be included in the retrieval, the terms “artificial intelligence”, “machine learning”, “automation”, “deepfake”, “GenAI”, or “algorithm”, or any of their derived terms, must have appeared in the title or contents of the document. Thus, this compilation reflects keyword-matched results and may not represent an exhaustive review of all AI-related regulatory activity across the captured jurisdictions⁵.

We provide references throughout this section directing readers to sources where further details and documentation underpinning each regulatory development can be found.

Scan Summary

Gambling-specific AI regulatory developments over the preceding 36 months have been concentrated primarily in North America and Western Europe, with no recorded activity across Central and Eastern Europe, Asia-Pacific, or the Middle East and Africa.

The most prevalent theme is the use of AI and automated tools for responsible gambling monitoring, with jurisdictions including Belgium, Finland, Italy, Alberta, Colorado, New Jersey, and North Carolina all referencing requirements or expectations around automated monitoring and/or detection of at-risk player behavior.



It is encouraging to see regulators paying greater attention to player protection and the role that data and technology can play in it, though the current focus remains largely on operator data and approaches that have long existed in the industry. Yaspa is already delivering AI-driven player protection strategies to operators using Open Banking financial data. This broader application highlights the need for regulatory frameworks to keep pace with the full scope of AI innovation in gambling.”

Max Collinge
Head of Product, Yaspa

⁵ We recognize that many regulatory obligations and compliance functions, particularly given the current state of the art of the technology, could be performed via AI. However, our intention was not to capture every instance where AI could be used, but rather to identify where regulations are specifically stepping in to recommend the use of or safeguard against specific applications of AI. As a balance, we did include the keywords “automation” and “algorithm” to capture adjacent developments. We also acknowledge that regulator-issued guidance or bulletins issued directly to licensees would not be captured here, as these are often not in the public domain. Finally, we also acknowledge that market realities and regulator expectations may lead to the use of AI or automation in practice, even in the absence of explicit legal or regulatory requirements addressing it.

An emerging theme is the governance of AI use itself. Massachusetts stands out as the most explicit example, prohibiting the use of AI and algorithmic systems designed to make gambling platforms more addictive. North Carolina, meanwhile, requires operators to document and disclose their use of AI systems. We also see some regulators highlighting risks in consumer-side AI use, with Alberta and Ontario introducing requirements to detect and prevent the use of automated bots in peer-to-peer games and the United Kingdom highlighting the use of AI to bypass customer due diligence.

Central and Eastern Europe

As of January 28, 2026, no legislative or regulatory restrictions, requirements, or other activity within the gambling industry related to AI has been recorded across Vixio's monitored Central and Eastern European markets.

Western Europe

Belgium November 2023

On November 14, 2023, six Belgian gaming operators, all members of the Belgian Association of Gaming Operators (BAGO), signed the Duty of Care Agreement [12], which aims to help identify and address problematic gaming behavior more promptly.

The Duty of Care Agreement outlines four key pillars for operators' prevention policies. Most relevant here is the first: A detection system – utilizing algorithms, artificial intelligence and scientifically-based criteria (markers of harm) to identify potentially risky behavior, considering factors such as playing frequency, duration, intensity, deposits, and other relevant criteria. The remaining three pillars focus on implementing player protection measures such as deposit limits, reality checks, and self-exclusion options; providing education and training for staff to identify and respond to potential problems; and sharing the prevention policy with the Belgian Gaming Commission.

In the long term, the operators aim to collaborate with the Gaming Commission on establishing an independent control mechanism. They also advocate for making duty of care legally mandatory for all legal gambling operators in Belgium, including the National Lottery, emphasizing the goal of offering a safe form of entertainment and avoiding negative social consequences associated with gambling addiction.

Finland January 2026

Under Section 34 of Finland's new Gambling Act [13] published in the official Gazette on January 20, 2026, the exclusive licensee and the gambling licensee must ensure that gaming complies with social and health aspects, in order to protect players from excessive gambling and to help them reduce their gambling when necessary. The duty of care includes combating excessive gambling through continuous monitoring and evaluation of gaming behavior.

The exclusive concession holder and the gambling concession holder shall assess the risk of gambling harm to the player arising from gambling by means of the automated processing of personal data referred to in Section 36(1)(3)–(7) and, where necessary, take measures to prevent and reduce the risks identified on the basis of the assessment. However, a decision to prevent or restrict a player's gambling may not be based solely on the automated processing of personal data.

Italy

March 2024

On March 25, 2024, Italy published Decree No. 41 regarding the reorganization of the gaming sector. As part of this decree, under Article 15 of Title III, the concessionaire's organizational structure and their technical, technological, and IT tools must be designed to safeguard and protect players by preventing and combating pathological gambling.

This is to be achieved through the introduction of automatic messages during gameplay that highlight its duration and provide real-time information to players on spending levels when a certain preset limit is exceeded. Furthermore, the decree mandates the activation of procedures to monitor the risk levels associated with individual games covered by the concession, based on internationally certified methodologies and excluding any games involving social, gender, political, religious, or other forms of discrimination. Finally, appropriate tools must be present to allow the concessionaire, in compliance with personal data protection legislation, greater control over the degree of participation by players most at risk of pathological gambling.

See Decree No. 41 [14, Art. 15(1)].

United Kingdom

April 2025

On April 8, 2025, the Gambling Commission published the updated “Emerging money laundering and terrorist financing risks from April 2025.” Under License Condition 12.1.1(3) of the License Conditions and Codes of Practice (LCCP) [15], operators are required to keep up-to-date with emerging risks information published by the Commission. The publication acts as a trigger for operators “to review their money laundering and terrorist financing risk assessments and related policies, procedures and controls to ensure that they remain appropriate and effective.”

The Commission has highlighted a variety of emerging risks, and most pertinent to this report is the use of artificial intelligence used to bypass customer due diligence. Additionally, the publication also requires operators to review the lists of “jurisdictions subject to increased monitoring by the Financial Action Task Force (FATF)” [16], and ensure they have effective policies, procedures and controls in place to identify customers and relationships with links to high-risk jurisdictions [17].

See guidance from April 2025 [18].

North America

Alberta

January 2026

On January 13, 2026, the Alberta Gaming, Liquor and Cannabis Commission (AGLC) published standards and requirements [19, Sec. 4.9] for internet gaming within the province, as well as guidelines regarding iGaming registration. Alberta's licensing system will operate under a dual process with registration with the AGLC, followed by entering into a commercial agreement with the Alberta iGaming Corporation AiGC).

Of note, Alberta's internet gaming standards require proactive responsible gambling monitoring and intervention to identify patrons at risk of harm and enable timely support once signs of risk emerge. Specifically, under the AGLC Standards and Requirements for Internet Gaming for identifying individuals who may be experiencing harm, licensees may use both automated and manual tools to monitor player behavior across all available data sources — spanning behavioral (e.g., chasing losses, playing multiple products), spend, time, account, and customer service indicators, among others.

Furthermore, as pertains specifically to peer-to-peer games, registered operators must implement measures to deter, prevent, and detect the use of software programs by players to automatically participate in game play (i.e., a bot) or to provide the player with an unfair advantage over other players [19, Sec. 3.3].

Colorado

January 2026 (Amended)

In January 2026, the Colorado Limited Gaming Control Commission amended its existing responsible gaming best practices regulations [20, R. 9.6(2)]. Most pertinent to this report, as part of these best practices, sports betting operators must submit for approval a strategy for the implementation of a responsible gaming program, with updates provided to the Colorado Division of Gaming by October 1st each year.

This strategy, among other requirements, must detail the use of player data and technology to aid in identifying potential problem gamblers, including how automated triggers — whether operated by the licensee or a contracted third party — are used to flag potentially problematic behavior, what specific triggers are being identified and how the operator has identified those as critical matrices to follow, and how the operator plans to continually update and learn the best way to identify problem gamblers on their site.

Massachusetts February 2025

On February 14, 2025, the Massachusetts Gaming Commission published regulations pertaining to requirements regarding Sports Wagering Data Privacy. Notably, as part of these requirements, sports wagering operators are prohibited from using patrons' Personally Identifiable Information (PII) or Confidential Information, or any information derived from it, to promote or encourage specific wagers or promotional offers based on specific behavior, patterns, and activity, including but not limited to: Any computerized algorithm, automated decision-making, machine learning, artificial intelligence, or similar system that is known or reasonably expected by the sports wagering operator or a vendor to the sports wagering operator to make the gaming platform more addictive [21, Sec. 257.02(3)].

However, sports wagering operators are permitted to collect patrons' PII and Confidential Information to analyze patron behavior for the purposes of identifying and developing programs and interventions to promote responsible gaming and support problem gamblers, and to monitor and deter illegal sports wagering. A report must be provided by the sports wagering operator to the Massachusetts Gaming Commission at least every six months on the sports wagering operator's compliance, including the trends observed in this data and the sports wagering operator's efforts to mitigate potential addictive behavior [21, Sec. 257.02(4)].

New Jersey January 2023

On January 1, 2023, New Jersey's Responsible Gaming Best Practices took effect. The new requirements apply to online wagering and supplement the state's existing responsible gaming (RG) safeguards.

The new standards require that each provider implement an RG lead as well as a dedicated RG staff who must be responsible for ensuring at-risk patrons are provided responsible gaming information such as resources available or various responsible gaming features offered by the provider.

In addition, and most pertinent to this report, the second component highlights the use of technology to aid in identifying potential problem gamblers. At a minimum, each platform must implement automated

triggers based on account activity and play behavior. These triggers span both account-level activity, such as high deposit volumes over 24-hour or 90-day periods, accessing but not completing self-exclusion, repeated requests for cool-off periods, successive deposit or loss limit increases, and continuous withdrawal cancellations, and play behavior indicators, including high turnover over a 90-day period, significant increases in time spent on the platform, and session patterns suggesting escalating wager amounts or repeated near-zero account balances.

See NJ iGaming RG Best Practices [22]

North Carolina March 2024

On March 27, 2024, the North Carolina State Lottery Commission re-adopted regulations pertaining to use of technology in information processing, which were originally adopted as Rule 1D-008 in January 2024.

Of note, under the requirements, operators' internal controls must include a summary description of the operator's or its affiliates' use of computerized algorithms, automated decision-making, machine learning, artificial intelligence or any similar system that analyzes patron information. This must cover the permissible and impermissible uses of such practices, the purposes for which they are employed, and the types of input and output data utilized along with a summary accounting of their sources. Operators must further describe how these systems are used both to minimize risky play behavior and to market or advertise to players.

See 24 NCAC 06A.0408 [23].

Ontario May 2025 (updated)

On May 29, 2025, the Alcohol and Gaming Commission of Ontario (AGCO) updated the Registrar's Standards for Internet Gaming, originally adopted and effective in April 2022.

As part of these updated requirements, pertaining specifically to peer-to-peer games, operators must implement measures intended to deter, prevent and detect the use by players of software programs to automatically participate in game play (i.e. a bot) or to provide the player with an unfair advantage over other players.

See Standard 4.22, Registrar's Standards for Internet Gaming [24, Sec. Standard 4.22].

LatAm and the Caribbean

Curaçao June 2025

On June 17, 2025, the Curaçao Gaming Authority (CGA) published its final Player Complaints Policy for implementation by all CGA license holders.

The player complaints policy outlines key definitions concerning player interaction; complaints, and; disputes. In addition, the policy sets out guidelines concerning the player complaint submission and resolution process, as well as guidelines pertaining to alternative dispute resolution (ADR), and record keeping and reporting among others.

Of note, as part of the guidelines outlining the CGA's complaint resolution process, within this policy, the Authority has deemed the use of artificial intelligence (AI) to be permissible under the policy, subject to the following terms: (a) Once a player complaint has been identified as pertaining to Responsible Gaming (clause 4.1), communications with the player should be conducted by a human, not AI, (b) Complaints

that can be reasonably considered as complex should be dealt with by a human, not AI, and (c) The AI records must be monitored to ensure that they are reasonable in their solution/recommendations and consistent across players with like-for-like complaints.

See (up-to-date) Player Complaints Policy [25].

Asia-Pacific

As of January 28, 2026, no legislative or regulatory restrictions, requirements, or other activity within the gambling industry related to AI has been recorded across Vixio's monitored Asia-Pacific markets.

Middle East and Africa

As of January 28, 2026, no legislative or regulatory restrictions, requirements, or other activity within the gambling industry related to AI has been recorded across Vixio's monitored Middle East and Africa markets.

2.3 Regulator Survey: Perspectives from Within

The regulatory scan presented above offers a picture of how AI-related requirements are beginning to take shape across gambling jurisdictions. Equally important is how the individuals responsible for overseeing the industry perceive and experience this rapidly evolving landscape. Accordingly, in this section we present findings from an original survey of individuals working within gambling regulatory authorities. The study explored key areas including respondents' personal AI usage, their confidence in overseeing this emerging technology, their perspectives on industry adoption, and their views on the adequacy of current governance frameworks, as well as their broader AI literacy.

Methodology

The regulator survey was independently conducted by the UNLV IGI with ethical approval from the UNLV Institutional Review Board. KPMG and Vixio contributed to the research through feedback on survey design, recruitment, and collaboration in interpreting aggregated findings.

We surveyed gambling regulators worldwide to examine how they perceive, understand, and oversee AI within the industry. Participants were recruited through institutional and partner networks — including the IGI's International Center for Gaming Regulation, Vixio's regulatory contacts, and the International Association of Gaming Regulators — as well as direct outreach and in-person recruitment at industry conferences. Data collection ran from December 2025 through February 2026. All responses were anonymous, and a screening question ensured only individuals holding current regulatory positions were included.

Our survey contained four core components: (1) AI awareness and use, (2) beliefs on AI regulation, (3) the future of AI in gambling, and (4) AI literacy. The first three blocks were developed collaboratively by the research team with input from KPMG and Vixio. The AI literacy block comprised 14 knowledge-based items adapted from validated instruments in the academic literature, including the short form AI Literacy Test (AILIT-S) and the AI Competency Objective Scale (AICOS). Descriptive statistics and nonparametric tests were used for analysis. Further methodological details including details on survey development and analysis processes are available in **Appendix A.3.2**.

Respondent Profile

We included 113 responses in the analysis. We highlight that a large majority of respondents (n = 94) represent jurisdictions across North America (n = 64) or Europe (Western Europe n = 20, Central and

Eastern Europe n = 10). We thus caution readers when interpreting these findings, as the sample may not be fully representative of global regulatory perspectives.

Professional Demographics of Respondents (N=113)

Based on survey responses regarding position and division at their regulatory agency

Figure 2.1 - Respondents by Position





Figure 2.2 - Respondents by Division



Figures 2.1 and 2.2 show our respondent profile by position and by region. 12.4% of respondents held executive or leadership positions at their regulatory agencies with 46.9% holding managerial positions. About 36.3% of respondents held staff positions and 4.4% did not disclose their positions. A majority of respondents (n = 43, 38.1%) work in *Enforcing/Investigation/Regulatory Compliance* followed by 24.8% (n = 28) in *Licensing/Administration/Policy*. More than half of respondents (n = 44, 40.7%) reported at least 10 years of experience in a regulatory role with oversight of the gambling industry. We provide a full description of the respondent profile in **Appendix A.3.2**.

Findings


AI Use, Training, and Literacy

Most respondents (n = 66, 57.9%) appear to have embraced the use of AI in a professional capacity. Forty-nine of these 66 respondents provided a response to our follow up question about the specific tools they use. ChatGPT and Microsoft Copilot were mentioned most frequently (n = 25 and n = 23 mentions respectively), with Google Gemini mentioned the third most often (n = 9 mentions). Other tools like Claude, Grok, DeepSeek, and Gamma were mentioned much less frequently (three times or less). The majority of respondents (n = 58, 58.0%) reported to have had no professional or formal training in AI within the last two years. Although the split in these cohorts appears similar, they are not related to each other in any meaningful way.

Sixty-eight participants completed the 14 AI literacy questions. The distribution of scores across all respondents is displayed in Figure 2.3. The mean score was 8.60 (Median = 9.0). 52.9% (n = 36) of respondents scored nine or above. The question most respondents answered incorrectly was “*How does supervised learning differ from unsupervised learning?*” Most respondents said that “In supervised learning, humans must supervise the AI during learning and intervene if necessary” whereas the correct answer is that supervised learning is when “the output values of the training data are known.” While we found a significant relationship between AI training and AI literacy scores, the correlation was very weak ($r = 0.26$). We did not find a significant relationship between AI use and AI literacy.

Figure 2.3 - Frequency Distribution of AI Literacy Scores (N=68)

Based on survey responses regarding AI literacy

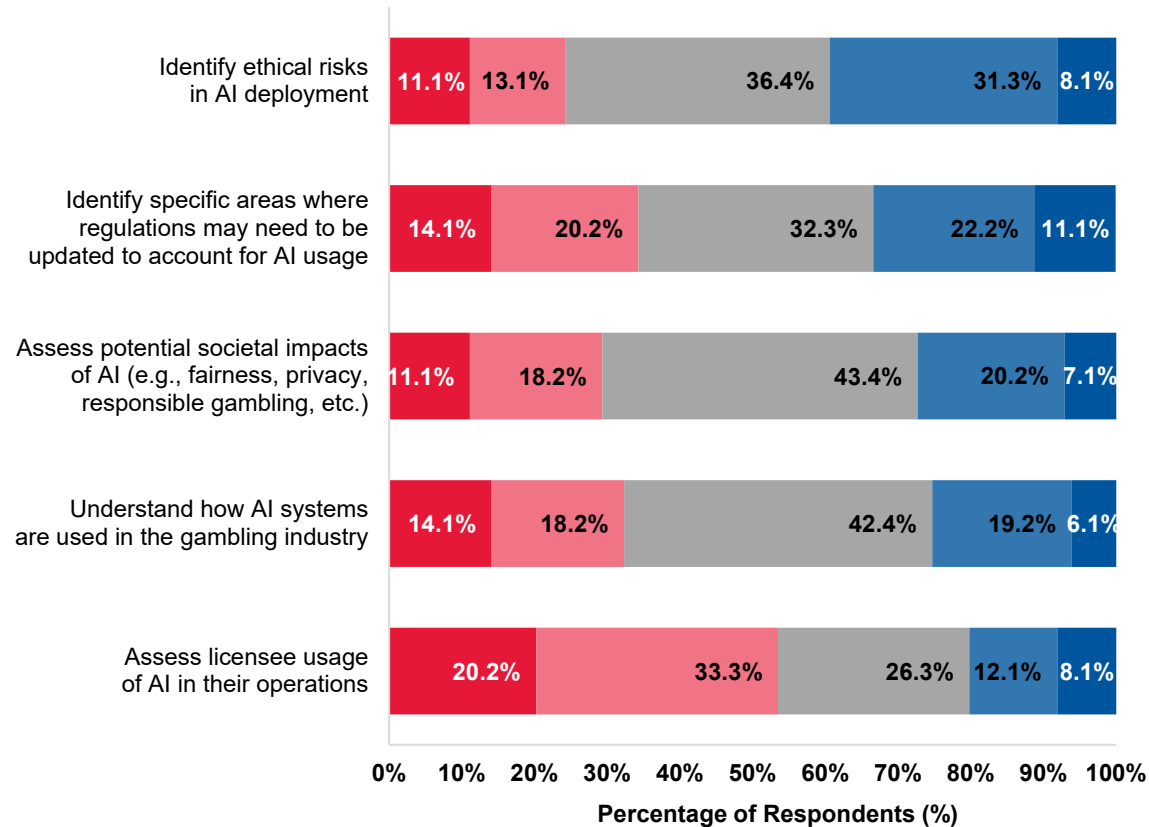


Regulatory Capacity

Figure 2.4 presents regulators' reported confidence, on a scale of 1 (*least confident*) to 5 (*most confident*), across five areas related to their ability to regulate AI in gambling.

Figure 2.4 - Confidence Levels in AI Awareness and Use (N=99)

Based on survey responses regarding AI awareness and use in the gambling Industry



Mean scores across items ranged from 2.55 to 3.12, with most items falling below or near the scale midpoint. The lowest confidence was reported for assessing licensee usage of AI in their operations ($M = 2.55$), while the highest was reported for identifying ethical risks in AI deployment ($M = 3.12$).

To explore whether AI use, training, or literacy were associated with differences in these confidence ratings, we averaged responses across items for each respondent to create an overall “competency score.” The score demonstrated excellent internal consistency (Cronbach’s $\alpha = 0.91$). Statistically significant differences were found for both AI use ($p = 0.02$) and AI training ($p = 0.02$), with those who personally use AI and those who had received formal AI training tending to score higher. However, effect sizes were small in both cases. No significant difference was found for AI literacy.

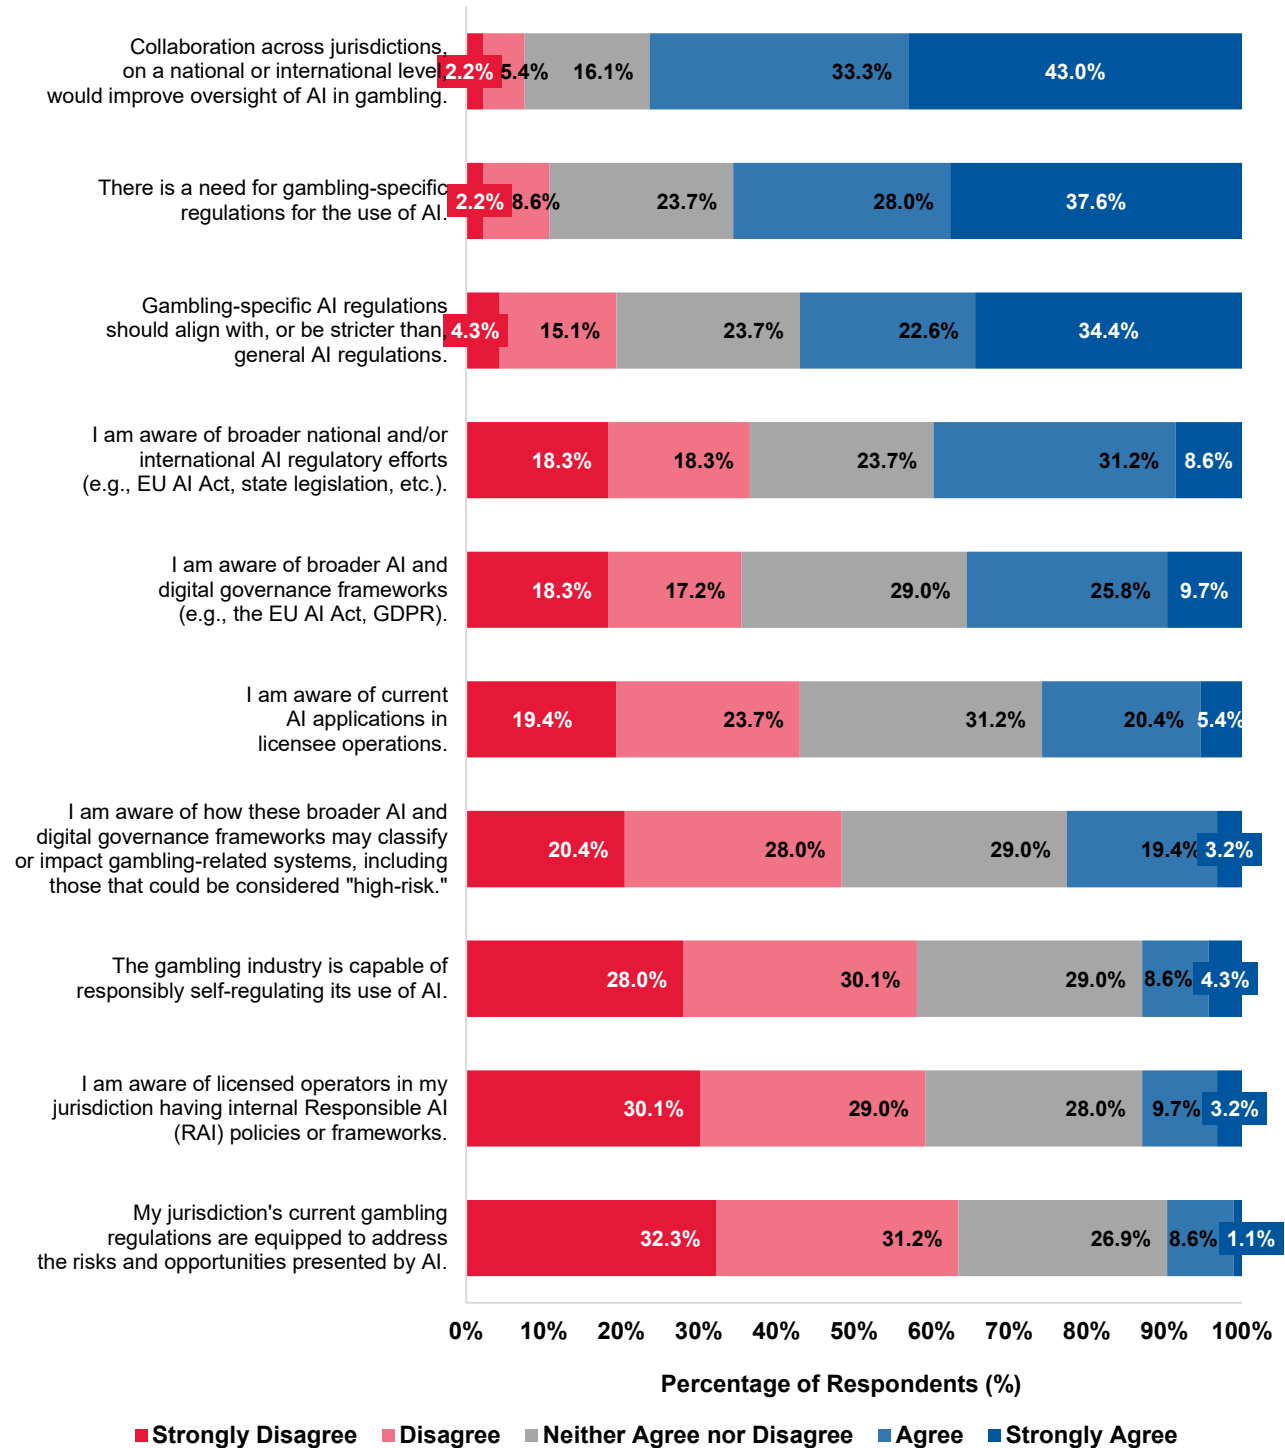
This modest confidence is echoed in responses to an open-ended question asking regulators to identify the single biggest challenge their agency anticipates facing regarding AI in gambling over the next three to five years ($n = 62$). The most frequently cited challenge was *insufficient knowledge* — regulators expressed concern about limited internal expertise, a lack of trained staff, and the technical competence needed to assess and regulate AI use cases. Another commonly expressed challenge was *keeping up* with the speed and volume of new AI innovations. Taken together, regulators appear to be aware of the governance challenge ahead, but may feel underprepared to meet it.

The Need for Action

Figure 2.5 presents regulators' level of agreement across 10 statements related to licensee AI practices and AI-specific regulation on a scale of 1 (*strongly disagree*) to 5 (*strongly agree*).

Figure 2.5 - Agreement Levels on AI Regulation (N=93)

Based on survey responses regarding beliefs on AI regulation in the gambling industry



Consistent with their lower confidence in assessing licensee AI usage (Figure 2.4), here regulators displayed limited familiarity with licensees' AI activity. Fewer than one in four agreed they were aware of current AI applications in licensee operations (M = 2.69) or of licensed operators having internal Responsible AI policies or frameworks (M = 2.27), with majorities disagreeing on both items (43.0% and 59.1%, respectively). Yet, despite this awareness gap, regulators were clear on the need for action with two-thirds agreeing that gambling-specific regulations for the use of AI are needed (M = 3.90), with over a third strongly agreeing (37.6%). There was also strong agreement that the gambling industry is not capable of self-regulating its use of AI, with over half of respondents disagreeing that the industry could do so responsibly (58.1%, M = 2.31). Thus, we see a similar trend to what we observed with regard to AI awareness: while our sample of regulators recognize the need for AI oversight in gambling, they may not be well-positioned to act on it.

While regulators tended to agree that current gambling regulations are ill-equipped to address AI (M = 2.15) and that gambling-specific regulations should align with or be stricter than general AI regulations (M = 3.68), their awareness of the broader AI regulatory landscape and governance frameworks available was inconsistent with this position (M = 2.94 and M = 2.91, respectively). They also expressed limited awareness of how broader AI governance frameworks could impact gambling-related systems (M = 2.57).

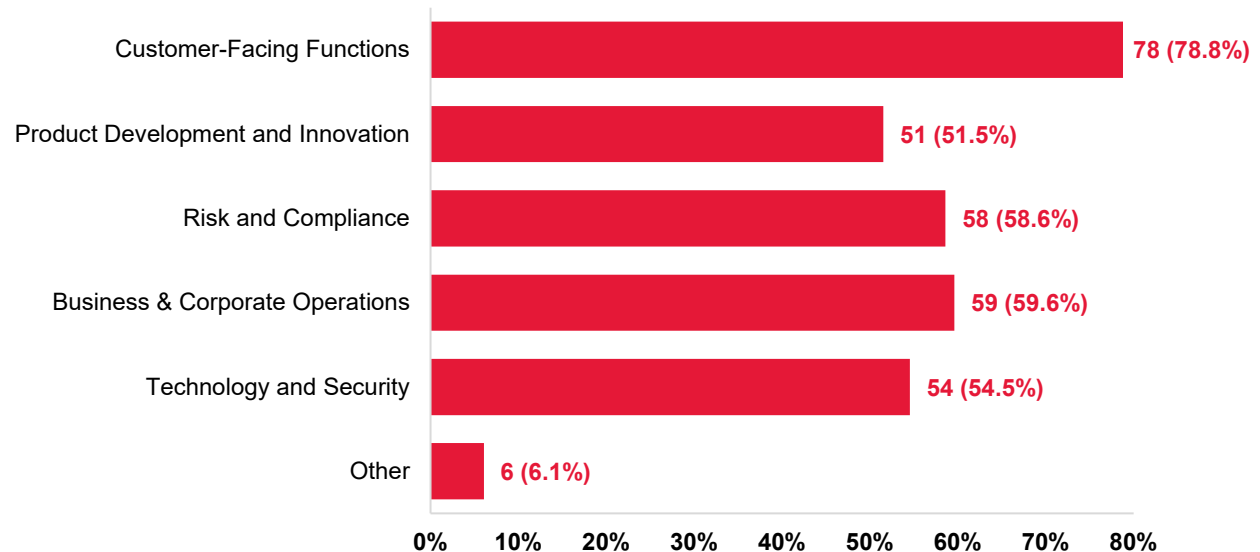
Finally, regulators were overwhelmingly in agreement that cross-jurisdictional collaboration would improve oversight of AI in gambling, which was the highest-rated item across the entire question set (M = 4.10), with over three-quarters of respondents agreeing or strongly agreeing (76.3%).

Regulator-Industry Disconnect

We noted in the prior section that regulators were generally unaware of AI applications in licensee operations. This is further corroborated by Figure 2.4, which presents regulator perceptions of the areas in which their licensees are currently using AI.

Figure 2.6 - Regulator Perception on AI Usage in Licensee Operations (N=99)

Based on survey responses regarding regulator beliefs on where licensees are using AI



The customer-facing function was the most endorsed category, with 78.8% of regulators selecting this category; all other categories were selected by 50-60% of respondents. This stands in contrast to findings from our industry survey, in which AI adoption was reported broadly across business functions, with *Product Development and Innovation* and *Technology and Security* endorsed most frequently. This clear disparity between regulators' perceptions of industry use, and where industry stakeholders are actually

deploying AI, brings to question the basis for regulators' strong agreement on the need for gambling-specific regulation and their skepticism toward industry self-regulation (Figure 2.5). It is worth considering how well-informed these normative positions truly are, and suggests communication between regulators and industry must improve to bring these two perspectives closer.

This regulator-industry disconnect is further reflected in commentary provided by regulators (n = 51) to our open-ended question asking “*What do you consider the most significant risks and benefits associated with AI in gambling?*” Notably, most regulator respondents listed risks rather than benefits, with the majority being concerned about bad-actor driven mechanisms such as unethical use of the technology and the subsequent harm it may cause to players. This sentiment is in contrast to findings from our industry survey, where respondents appeared comparatively less concerned about player-facing risks like bias or unfair outcomes affecting players and the potential for AI to amplify problem gambling behaviors.

Future AI Governance

When asked whether their agency was planning to develop or implement any guidelines, recommendations, or review processes to ensure the responsible use of AI in gambling, responses were evenly split: 52.6% (n = 40) indicated yes, while 47.4% (n = 36) indicated no. Of those who answered yes, 23 provided further detail on what was planned. Responses mainly fell across two categories: Research (e.g., gathering information, conducting studies, forming workgroups, and consulting stakeholders to inform future governance decisions) and Establishing Requirements (e.g., developing/implementing regulatory standards regarding AI use that licensees are expected to adhere to).

2.4 Conclusion

We began this section by illustrating a diverse regulatory landscape at the macro level, where current actions reflect a tension between safety and the desire not to stifle innovation. Gambling-specific regulatory developments appear more focused on preventing undesirable outcomes, with activity centered on the use of AI to monitor player behavior and prevent harms as well as mitigate threats such as adversarial use by players themselves. This remains an evolving area, however, and we anticipate further developments over the next 12 months.

Notably, while big tech calls for a lighter regulatory touch at the macro level [26], regulation does not appear to be a barrier for AI adoption amongst gambling industry stakeholders. Data from our industry survey suggest that regulation is largely seen as a minor friction rather than a barrier, with most respondents indicating that regulations are either having no effect on AI adoption (33.7%, n = 28) or causing only minor delays (41.0%, n = 34), with just 10.8% reporting regulations as a significant or primary barrier. At the same time, findings from our regulator survey suggest that the need for AI governance is recognized and that regulatory action may be forthcoming.



The AI survey findings provide industry stakeholders with clear evidence of the varying levels of awareness, literacy, use, and governance within the gambling regulatory landscape. A core focus for the International Center for Gaming Regulation remains the development of governance-enhancing mechanisms such as specialized training, webinars, research and information sharing to address knowledge gaps. The insights from the survey stand to equip us with specific starting points in empowering those responsible for oversight, as we navigate the “known unknowns” and “unknown unknowns” in this area of our dynamic industry, together.”

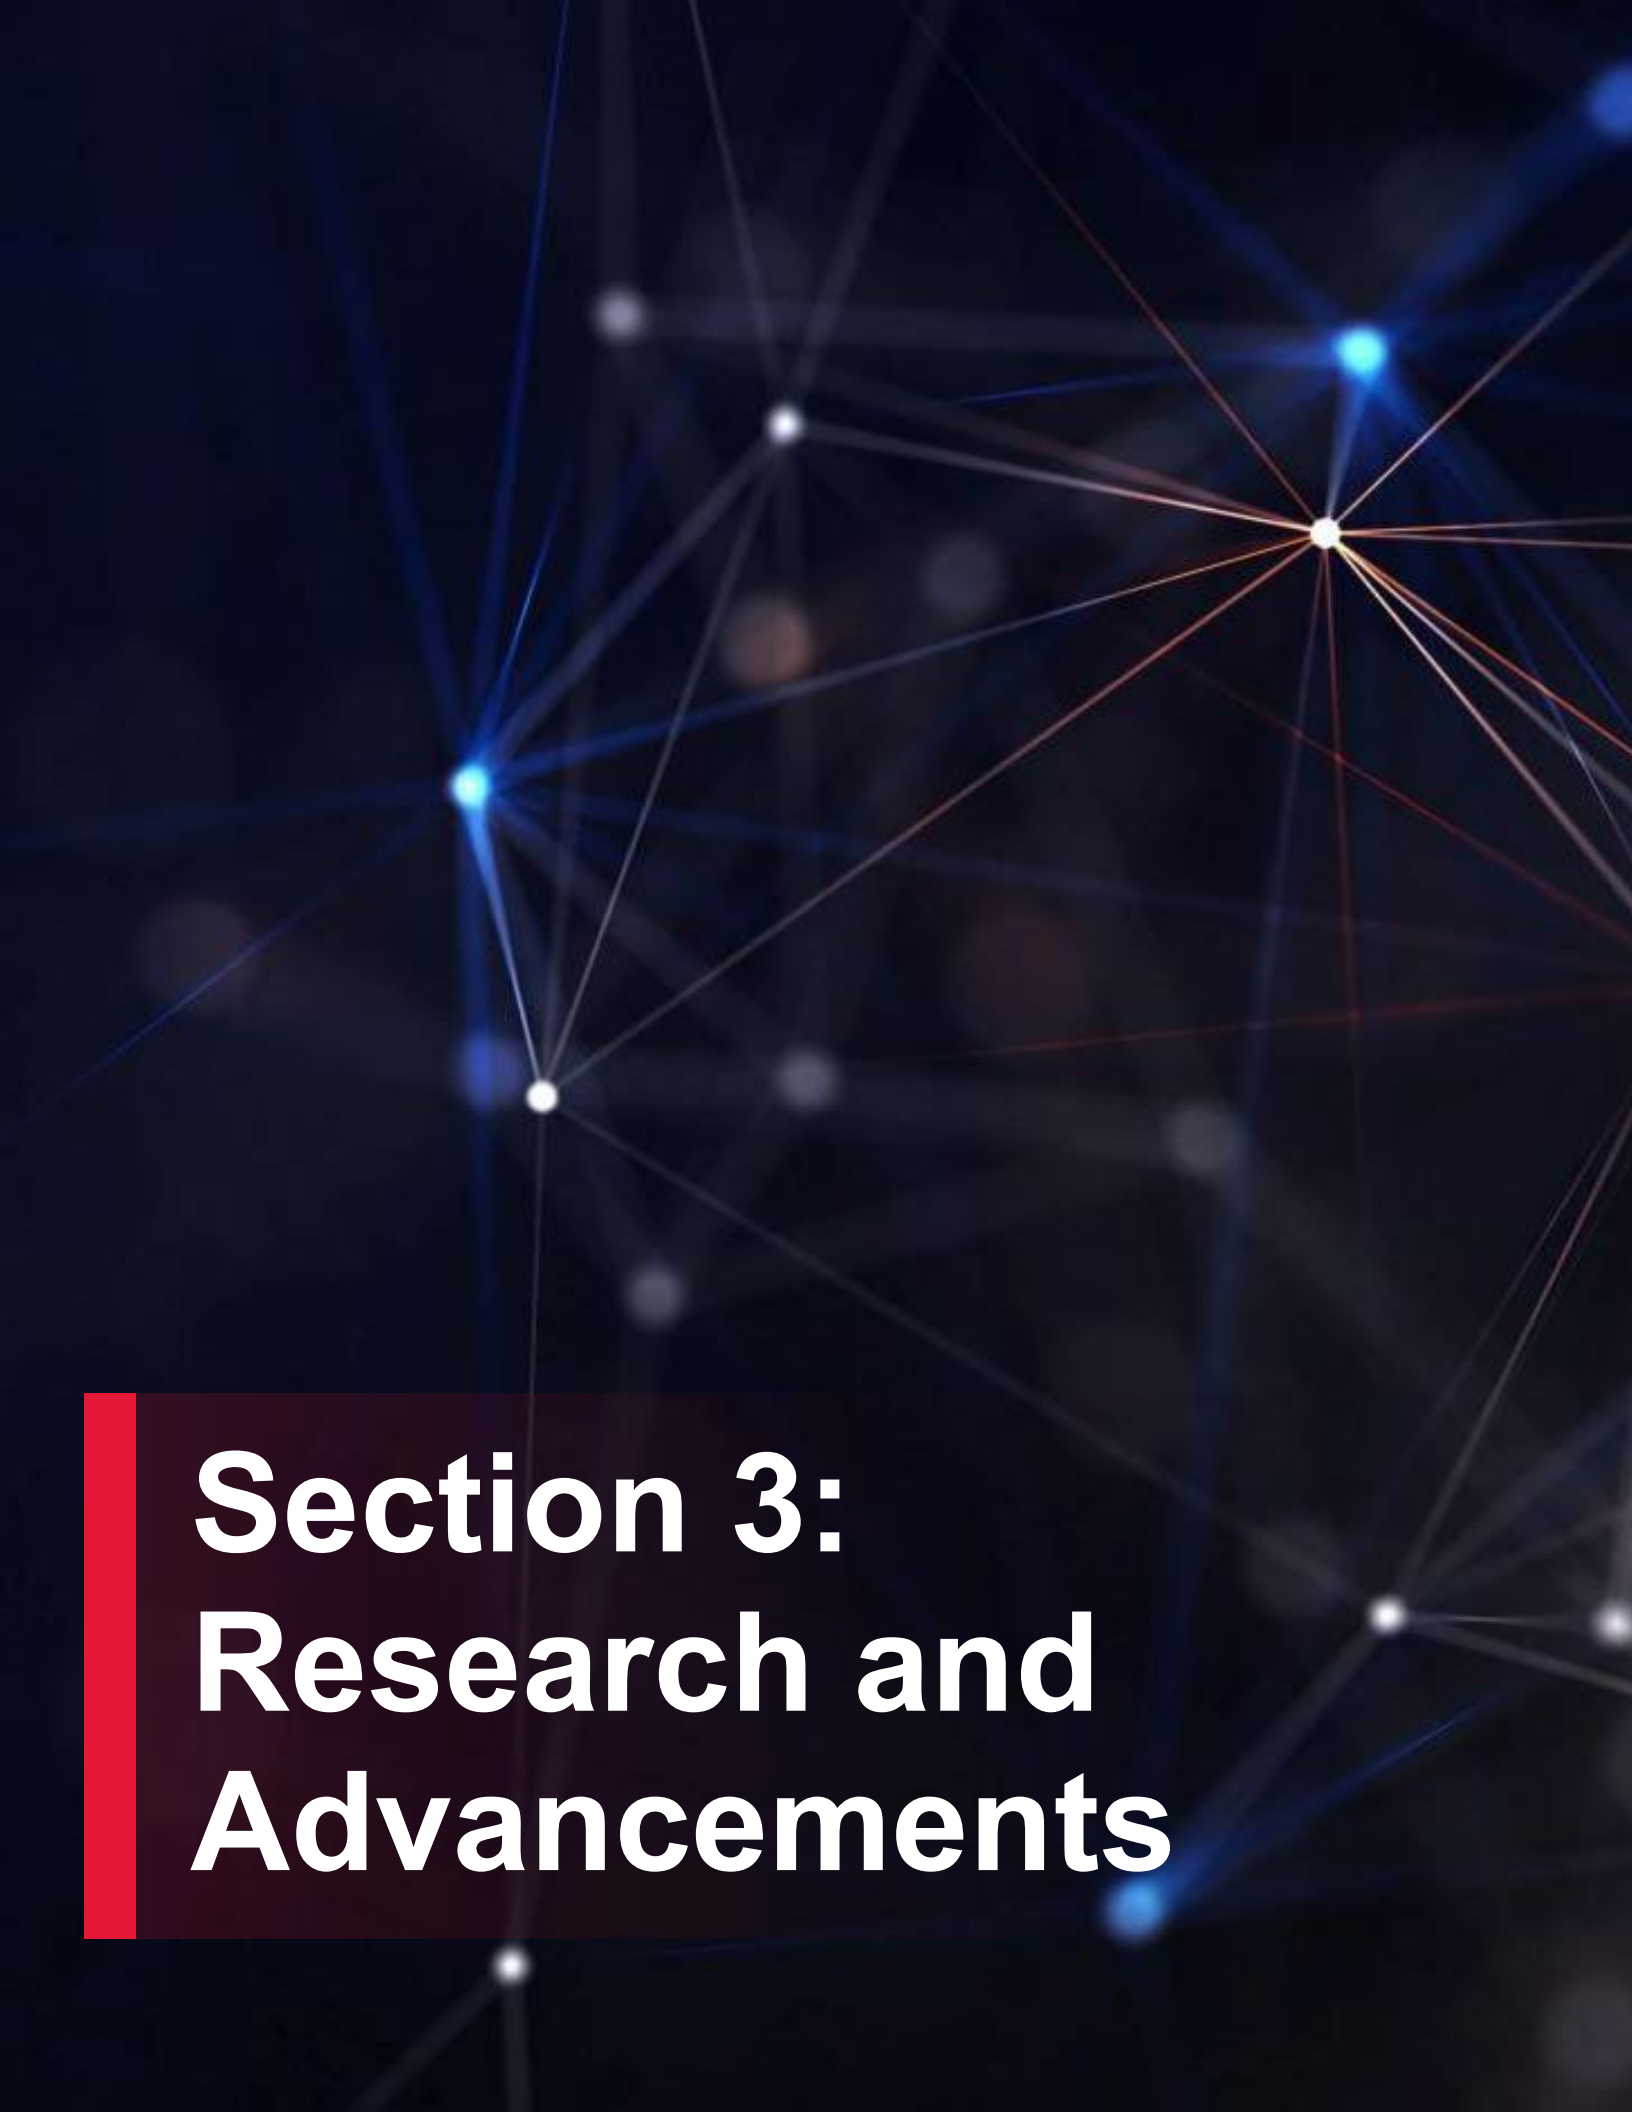
Megan Basson

Executive Director,
UNLV IGI International
Center for Gaming Regulation

References

- [1] “An Introduction to the Code of Practice for General-Purpose AI | EU Artificial Intelligence Act.” Accessed: Feb. 26, 2026. [Online]. Available: <https://artificialintelligenceact.eu/introduction-to-code-of-practice/>
- [2] “AI Regulatory Sandbox Approaches: EU Member State Overview | EU Artificial Intelligence Act.” Accessed: Feb. 26, 2026. [Online]. Available: <https://artificialintelligenceact.eu/ai-regulatory-sandbox-approaches-eu-member-state-overview/>
- [3] W. Simpson, “Global AI Law and Policy Tracker: Highlights and takeaways,” IAPP.org. Accessed: Feb. 26, 2026. [Online]. Available: <https://iapp.org/news/a/global-ai-law-and-policy-tracker-highlights-and-takeaways>
- [4] T. M. Le and A. Do, “Vietnam’s first standalone AI Law: An overview of key provisions, future implications,” IAPP.org. Accessed: Feb. 26, 2026. [Online]. Available: <https://iapp.org/news/a/vietnam-s-first-standalone-ai-law-an-overview-of-key-provisions-future-implications>
- [5] J. Power, “EU moves to ease AI, privacy rules amid pressure from Big Tech, Trump,” Al Jazeera. Accessed: Feb. 26, 2026. [Online]. Available: <https://www.aljazeera.com/economy/2025/11/20/eu-moves-to-ease-ai-privacy-rules-amid-pressure-from-big-tech-trump>
- [6] “Voluntary AI Safety Standard,” Australian Government Department of Industry Science and Resources. Accessed: Feb. 26, 2026. [Online]. Available: <https://www.industry.gov.au/publications/voluntary-ai-safety-standard>
- [7] A. H. Ishaque, A. Aidid, and G. Zadeh, “Lessons from the Failure of Canada’s Artificial Intelligence and Data Act,” *NEJM AI*, vol. 2, no. 7, Jun. 2025, doi: 10.1056/AIpc2500153.
- [8] J. Patto and P. Gupta, “Global AI Governance Law and Policy: Australia,” IAPP.org. Accessed: Feb. 26, 2026. [Online]. Available: <https://iapp.org/resources/article/global-ai-governance-australia>
- [9] Executive Orders, “Ensuring a National Policy Framework for Artificial Intelligence,” The White House. Accessed: Feb. 26, 2026. [Online]. Available: <https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy/>
- [10] S. Wiener, *Transparency in Frontier Artificial Intelligence Act (TFAIA)*, vol. Division 8, Chapter 25.1 (Sections 22757.10 et seq.). California State Legislature, 2025. [Online]. Available: https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202520260SB53
- [11] A. Gounardes, *Responsible AI Safety and Education (RAISE) Act*, vol. Article 44-B (Sections 1420–1429). 2025. [Online]. Available: <https://www.nysenate.gov/legislation/bills/2025/S6953/amendment/B>
HYPERLINK “<https://www.zotero.org/google-docs/?jJ5H8x>” \h
- [12] “Kansspeloperatoren leggen zichzelf zorgplicht op,” bago.be. Accessed: Feb. 26, 2026. [Online]. Available: <https://bago.be/kansspeloperatoren-leggen-zichzelf-zorgplicht-op/>
- [13] *Rahapelilaki [Gambling Act]*. 2026. Accessed: Feb. 26, 2026. [Online]. Available: <https://www.finlex.fi/fi/lainsaadanto/saaduskokoelma/2026/10?language=fin>
HYPERLINK “<https://www.zotero.org/google-docs/?jJ5H8x>” \h
- [14] Presidenza della Repubblica, HYPERLINK “<https://www.zotero.org/google-docs/?jJ5H8x>” \h *Decreto Legislativo 25 marzo 2024, n. 41: Disposizioni in materia di riordino del settore dei giochi, a partire da quelli a distanza*. Rome, Italy: Gazzetta Ufficiale della Repubblica Italiana, 2024. [Online]. Available: <https://www.gazzettaufficiale.it/eli/id/2024/04/03/24G00060/SG>

- [15] *Licence Condition 12.1.1: Anti-money laundering – Prevention of money laundering and terrorist financing*, vol. 12.1.1. 2024. [Online]. Available: <https://www.gamblingcommission.gov.uk/licensees-and-businesses/lccp/condition/12-1-1-anti-money-laundering-prevention-of-money-laundering-and-terrorist>
- [16] Financial Action Task Force, “Jurisdictions under Increased Monitoring - February 2025,” FATF, Paris, France, Feb. 2025. Accessed: Feb. 26, 2026. [Online]. Available: <https://www.fatf-gafi.org/en/publications/High-risk-and-other-monitored-jurisdictions/increased-monitoring-february-2025.html>
- [17] Financial Action Task Force, “High-Risk Jurisdictions subject to a Call for Action - February 2025,” FATF, Paris, France, Feb. 2025. Accessed: Feb. 26, 2026. [Online]. Available: <https://www.fatf-gafi.org/en/publications/High-risk-and-other-monitored-jurisdictions/Call-for-action-february-2025.html>
- [18] Gambling Commission, “Emerging money laundering and terrorist financing risks from April 2025,” Gambling Commission, Birmingham, UK, Regulatory Guide, Apr. 2025. [Online]. Available: <https://www.gamblingcommission.gov.uk/licensees-and-businesses/guide/page/emerging-money-laundering-and-terrorist-financing-risks-from-april-2025>
- [19] Alberta Gaming, Liquor and Cannabis, “Standards and requirements for Internet Gaming,” AGLC, St. Albert, AB, Regulatory Handbook, Feb. 2026. [Online]. Available: <https://aglc.ca/documents/standards-and-requirements-internet-gaming>
- [20] Colorado Department of Revenue, Division of Gaming, *Colorado sports betting rules: Rule 9 - Responsible Gaming and Self-Restriction*, vol. Rule 9. 2026. [Online]. Available: <https://sbg.colorado.gov/sites/sbg/files/documents/SBRule-9-011426.pdf>
- [21] Massachusetts Gaming Commission, *205 CMR 257: Sports wagering data privacy*, vol. 205 CMR 257. 2023. [Online]. Available: <https://www.mass.gov/doc/205-cmr-257-sports-wagering-data-privacy/download>
- [22] New Jersey Department of Law and Public Safety, Division of Gaming Enforcement, *Division of Gaming Enforcement Responsible Gaming Best Practices*, vol. 13:69O–1.2A. 2025. [Online]. Available: <https://www.nj.gov/oag/ge/docs/BestPractices/ResponsibleGamingBestPractices.pdf>
- [23] North Carolina State Lottery Commission, *Rules Manual for Sports Wagering and Pari-Mutuel Wagering*, vol. 24 NCAC 06A.0408. 2024. [Online]. Available: https://ncgaming.gov/wp-content/uploads/2025/07/North_Carolina_State_Lottery_Commission_Rules_Manual_for_Sports_Wagering_and_Pari-Mutuel_Wagering_December_18_2024.pdf
- [24] Alcohol and Gaming Commission of Ontario, “Registrar’s Standards for Internet Gaming,” AGCO, Toronto, ON, Standard, May 2025. [Online]. Available: <https://www.agco.ca/en/book/export/html/245361>
- [25] Curaçao Gaming Authority, “Player Complaints Policy Guideline,” Curaçao Gaming Authority (CGA), Policy Guideline Version 1.0, Jun. 2026. [Online]. Available: <https://drive.google.com/file/d/1Qn8EsimAPkNPOHuPsEqJJIAmPmyxeGgG/view>
- [26] S. Levy, “How the Loudest Voices in AI Went From ‘Regulate Us’ to ‘Unleash Us,’” *Wired*, May 30, 2025. Accessed: Mar. 05, 2026. [Online]. Available: <https://www.wired.com/story/plaintext-sam-altman-ai-regulation-trump/>



Section 3: Research and Advancements

Section 3: Research and Advancements

Overview

This section analyzes the innovation pipeline within the gambling industry by tracking the growth, themes, and geographic distribution of AI-related research and advancements, specific to gambling and gaming. This includes a review of 15 years of academic contributions to measure historic growth and citation impact, as well as an examination of the increasing frequency of AI-focused sessions at international industry conferences over the past five years. Additionally, we present historic and recent trends in patent filings, AI commentary in gambling trade industry news, and end with a summary of startup activity. Together, the following sub-sections provide a trajectory of how AI is evolving within the gambling sector.

- Overview 49
- 3.1 Academic Frontiers 50
- 3.2 Industry Conference Discourse..... 56
- 3.3 Intellectual Property 59
- 3.4 Macro News Trends and Outlook 65
- 3.5 Startup Activity 67
- References 70

3.1 Academic Frontiers

In this section we present a systematic bibliometric analysis of the global research output at the interdisciplinary intersection of AI and the gambling field. The study utilizes metadata retrieved from the OpenAlex database, an open-access catalog of the global research system, to examine a corpus of English-language scholarly works published between 2010 and 2025. This time frame coincides with the rise of deep learning in the wider research community in 2010 when the Theano and Caffe frameworks made neural network implementations more accessible [1].

We implemented a three-phase pipeline to retrieve, identify, and classify relevant academic literature. Phase 1 involved programmatic retrieval of a candidate set of gambling-related publications from the OpenAlex database, Phase 2 filtered specifically for AI-related gambling publications, and Phase 3 involved the thematic analysis of the subset derived during Phase 2. Phase 1 garnered a total of 103,736 unique gambling-related publications. After executing Phase 2, our final corpus of AI-related gambling publications contained 899 records. Please refer to **Appendix A.3.3** for more details on the methodology.

Scholarly interest in AI is rising

Figure 3.1.1 – AI Publications in Gambling Worldwide, 2010-2025

Source: Publication data obtained from OpenAlex




Figure 3.1.1 illustrates a consistent upward trajectory in AI-related gambling publications. Following a steady period of growth where annual publications more than doubled between 2010 and 2019, the field experienced a minor decline in 2020. This was followed by a rapid secondary surge, with research volume reaching a peak of 112 publications by 2025; which represents 12.5% of our entire corpus. The rapid climb in AI publications after 2022 coincides with the release of ChatGPT, which is potentially a contributing factor.

Visualizing the data by the number of AI-related gambling publications as a percentage of all gambling-related publications (line in Figure 3.1.1), we observe a similar growth trajectory. This indicates that the


growth in AI-related research is not simply a reflection of an overall increase in gambling scholarship; AI-related works are capturing a growing share of the field’s overall research output.

However, AI’s growing share of gambling research remains modest compared to other academic fields. For example, recent bibliometric studies show that the share of AI-related healthcare publications rose from 3.54% in 2014 to 16.33% by 2024 [2]. Similarly, research in finance saw AI’s influence climb from roughly 2-3% in 2010 to nearly 20% by 2022 [3]. In comparison to these fields, AI research hasn’t been as disruptive in the gambling field. But the significant growth from 2022 to 2025 is a clear signal of an accelerating trend.

What is being studied?

Figure 3.1.2 – Number of AI Publications by Gambling Topic, 2010-2025

Source: Publication data obtained from OpenAlex



We used an AI-based topic modeling technique (see **Appendix A.3.3**) to classify publications into six distinct gambling categories, which represent approximately three quarters of all collected records. The remaining publications, which could not be classified into these categories, were classified as “Other.” We visually represent the trends by category in Figure 3.1.2.


The Sports Betting category includes publications leveraging ML or predictive analytics to perform functions such as modeling sports outcome predictions [4], [5], or forecasting wins and spreads using sentiment analysis [6]. Online Problem Gambling publications focus on using AI techniques (e.g., ML, natural language processing (NLP)) to detect signs of problem gambling [7], [8], [9]. Poker poses an interesting challenge for AI given the imperfect information inherent to the game. Poker research includes using reinforcement learning for heads-up no-limit poker [9], developing frameworks for exploiting opponent weaknesses [10], and training LLMs to become professional poker players [11]. The Gambling Decision Making category encompasses a broad range of publications, including the psychological underpinnings of gambling decisions - such as experiments on rats [12] and Iowa Gambling Task studies

[13], [14] - as well as consumer behavior research, such as studies on player preferences within the casino environment [15], [16]. Casino Surveillance publications were primarily related to computer vision, fraud detection, threat detection and facial recognition [17], [18], [19]. Gambling Website Detection publications were primarily related to the use of classification algorithms to identify unregulated or illicit online gambling platforms [20], [21]. Several publications that did not clearly fit into any of these categories - such as AI ethics in gambling [22], [23] and scoping reviews of AI in gambling [24], [25] - were classified as “Other.”

Between 2010 and 2017, Poker was the dominant research topic, peaking at 41.0% of the total corpus in 2017. However, this interest collapsed to a mere 4.5% by 2025. This decline aligns with the 2019 release of Carnegie Mellon’s Pluribus [26], the first AI to master six-player poker. As interest in Poker began to recede from 2017, Sports Betting and Online Problem Gambling emerged as the leading primary research pillars around 2018 to 2019. The rise of Online Problem Gambling as a research focus reflects a broader reframing of gambling as a public health subject across academic research in multiple jurisdictions [27]. Similarly, the growth of Sports Betting research may be attributed to growing participation, as in, for example, the United States following the repeal of PASPA in 2018.

Figure 3.1.3 – Sports Betting Publications by Sport (% of total), 2010-2025


Source: Publication data obtained from OpenAlex



Given its significant traction, we expanded our exploration of the 245 publications within the Sports Betting category. Figure 3.1.3 illustrates the trends in popularity by sport over our 15-year time period. To allow for year-over-year comparison, each data point represents a specific sport’s share as a percentage of the total sports betting corpus for that year. Football/soccer has maintained interest over time, and accounts for almost half of the publications in 2025. Meaningful shares are also observed for Tennis (16.7%), Basketball (13.3%), and Cricket (10.0%). With the exception of Basketball, this hierarchy in 2025 aligns with the most popular sports in the world per worldatlas.com [28].

Figure 3.1.4 – AI Publications in Gambling by AI Topic (% of total), 2010-2025

Source: Publication data obtained from OpenAlex



We also looked at our corpus by the types of AI being applied and studied (Figure 3.1.4). Machine Learning (ML) accounts for 61.6% of all publications in 2025, and is the most popular area across the entire time period. In 2025, Natural Language Processing (NLP) also commanded a meaningful share of total output (15.9%), meaning this, along with ML, represents over three-quarters of the most recent year's research volume. This uptick in NLP research, beginning around 2018–2019, directly correlates with the release of foundational deep learning architectures. The introduction of the Transformer model in 2017 [29], followed by the debuts of GPT-1 [30] and BERT [31] in 2019, may have been a catalyst for this growth. Despite the surge in interest and applications of Generative AI, research in gambling is only just beginning to emerge [32], [33].

Where are the academic researchers studying AI and gambling?

Figure 3.1.5 – AI Publications in Gambling (% of total) by Region, 2010-2025

Source: Publication data obtained from OpenAlex




Figure 3.1.5 displays our corpus of AI-related publications by region, determined by the locations of each publication’s contributing authors⁶. Between 2010 and 2020, North America and Europe & Central Asia were the leading regions for research output. However, 2020 marked a turning point as East Asia & Pacific took a lead.

By 2024, East Asia & Pacific and Europe & Central Asia converged at similar percentage levels, while North America experienced a decline in output starting in 2017, ultimately having a noticeable drop in percentage by 2025. This observation is consistent with data from Yahoo Finance [34] indicating that 15 of the world’s most significant gambling markets are within these three primary regions: East Asia & Pacific, Europe & Central Asia, and North America.

⁶ We mapped author affiliation country codes to the seven official World Bank Country and Lending Groups. To accurately capture international collaboration, we assigned a count to a region for each unique country within its borders represented in a paper’s authorship. For example, a study co-authored by researchers in the United States and Canada would attribute two counts to North America (one for each unique country). This method ensures the analysis reflects total institutional participation across regions rather than defaulting to the lead author’s location.

Figure 3.1.6 – Top 10 Institutions by Number of Publications, 2010-2025

Source: Publication data obtained from OpenAlex




Figure 3.1.6 shows the top 10 institutions by number of publications produced between 2010 and 2025⁷. No individual organization dominates the field. Even the leading institution, Carnegie Mellon University, represents just 1.4% of the total volume for this period. Notably, several of these top-tier institutions demonstrate a specialized focus on Poker research. Four of the five most-cited publications in the entire corpus are Poker-related, authored by researchers at Carnegie Mellon University and the University of Alberta [35], [36], [37], [38]. According to OpenAlex “DeepStack: Expert-level artificial intelligence in heads-up no-limit poker,” led by the University of Alberta, has 790 citations.

⁷ We parsed and filtered institutional affiliations for each paper so that a single organization received only one “credit” per publication, regardless of how many individual authors from that institution contributed to the work. For instance, a paper co-authored by three researchers from the same university counts as a single institutional appearance, whereas a cross-institutional collaboration credits each unique organization once.

Figure 3.1.7 – Top 10 Institutions by Number of Publications, 2020-2025

Source: Publication data obtained from OpenAlex




Figure 3.1.7 looks at top publishing institutions during the most recent five-year period, 2020–2025. A comparison between Figure 3.1.6 and Figure 3.1.7 reveals significant movement amongst the top institutions: of the five historically leading organizations, only two retain their top-five rankings for the 2020–2025 period. The research output of these leading institutions also demonstrates some thematic concentrations. For instance, Gambling Website Detection is a prominent focus among the Chinese universities listed [39], [40], [41], whereas Nottingham Trent University shows a clear specialization in Online Problem Gambling [42], [43], [44].

3.2 Industry Conference Discourse


In this section we explore how AI has been represented across major international gaming conferences between 2020 and 2025. Conference programs were collected from major international gambling and gaming events, including G2E, G2E Asia, ICE Gaming, SBC Summit (Lisbon, Miami/Latin America, and Barcelona editions), and NEXT Summit (New York City). Across all conferences, keynotes, panels, and standard presentations were used to capture the prevalence and context of how AI is being discussed in the gambling industry.

Data for this section were gathered from public conference archives and direct correspondence with organizers. Session titles and descriptions were extracted into a spreadsheet for systematic coding. To ensure agreement in classification, we followed a cross-researcher verification process to verify the data's accuracy. Readers should note that results for 2020 and 2021 reflect the significant impacts of the COVID-19 pandemic, including event cancellations, shifts to virtual-only formats, and reduced programming. Please refer to **Appendix A.3.3** for more details on the methodology. In total, 1,440 sessions were extracted, 162 of which were identified as being AI-related.

AI discussions at conferences is on the rise

Figure 3.2.1 – Number of AI Sessions at Gambling Conferences, 2020-2025

Source: Conference Data



There is a prominent rise in AI-related sessions at gambling conferences between 2022 and 2025, as depicted in Figure 3.2.1. While 2020 and 2021 recorded a mere three sessions each, 2025 recorded 81. The three AI sessions in 2020 all occurred at ICE and the three in 2021 all occurred at G2E Asia (see Figure 3.2.2). The turning point in 2022 observed in Figure 3.2.1 mirrors the rise observed in the number of AI-related gambling publications in the prior section (Figure 3.1.1), which coincides with the release of ChatGPT and the rise of LLMs in 2022.

AI-related content across events

Figure 3.2.2 – AI and Non-AI Sessions Across Gambling Conferences, 2020-2025

Source: Conference Data





Figure 3.2.2 shows the number of AI and non-AI sessions across major gambling conferences from 2020 to 2025. Across all events, non-AI sessions substantially outnumber AI sessions, though the proportion of AI content varies by conference. Among the two SBC events reviewed, SBC Europe had the highest proportional AI representation, with 64 AI sessions out of a total of 376 (17.0%). The second highest representation was at G2E Asia, where 15 of the 102 sessions were AI-related (14.7%). ICE had the lowest share of AI content, with 19 of 311 sessions (6.1%). Overall, despite a measurable rise in AI-session across gambling conferences, they only represent a small proportion of all conference programming. However, the share of AI-related sessions is increasing over time, as we illustrate in Figure 3.2.3.

Figure 3.2.3 – Percentage of AI Sessions Across Gambling Conferences, 2020-2025

Source: Conference Data



Across events, AI representation generally increases after 2022 (Figure 3.2.3). G2E rises from approximately 1.3% in 2022 to roughly 18% in 2025. SBC Europe also increases steadily, growing from about 7% in 2023 to 16.7% in 2025. G2E Asia shows an early peak at approximately 43% in 2021, however, this high percentage should be noted within the framework of a significantly reduced virtual program due to COVID-19, with seven sessions occurring in total. ICE shows an upward trend, approximately doubling their AI-related content over this five-year period. And NEXT shows modest growth from zero AI-related sessions in 2022 to roughly 10% in 2025. As AI developments continue, it will be interesting to see if this trajectory continues.

What is being discussed?

Figure 3.2.4 – Percentage of AI Sessions by Topic Across Gambling Conferences, 2020-2025

Source: Conference Data




Figure 3.2.4 represents the gambling related topics for all AI sessions at gambling conferences from 2020 to 2025. Sessions from 2020 to 2022 have very small sample sizes (see Figure 3.2.1), and thus do not accurately represent the range of topics discussed. The majority of AI sessions from 2023 to 2025 were related to two key areas: (1) Compliance (RG, KYC, AML), and (2) Marketing and CRM. Interestingly, while our industry stakeholder survey identifies Technology & Security as a top area for AI adoption, it remains a minority topic in conference discourse.

3.3 Intellectual Property


Patent filings provide a window into where companies are directing their R&D investment and strategic focus, which can signal innovation shifts before they reach the market. In this section, we present an analysis of patent activity from 2010 to 2025.

All data in this section is sourced from Lens.org, a comprehensive database containing millions of intellectual property (IP) records from over 95 different jurisdictions. Thematic and topic analysis was performed based on the Cooperative Patent Classification (CPC) and International Patent Classification (IPC) codes. Please refer to **Appendix A.3.3** for more details on the methodology. We employed a hybrid extraction process using keywords, key phrases and CPC codes to query the Lens.org database across Title, Abstract, and Claims, with a strict requirement that each record contained terms related to both gambling and AI. As a result of this process, our analysis corpus contained 2,125 patent records with 650 granted patent records.

AI IP in Gambling is Growing

Figure 3.3.1 – Number of AI Patents in Gaming Worldwide, 2010-2025

Source: Lens.org | Chart: 2026 State of AI in Gaming report



The data reveals a distinct increase in activity starting in 2020, signaling a significant surge of interest in AI for the gambling space (Figure 3.3.1). Over the entire time period, annual patent grants rose from 15 in 2010 to 100 by 2025, signaling significant and growing industry interest in AI. As computing power increases and AI tools become more accessible, we envision this trend to continue into 2026.

Figure 3.3.2 – Number of Patents Published Each Year by Document Type

Source: Lens.org




Figure 3.3.2 presents all patent filings and legal status over time, highlighting a shift in industry behavior. In 2010, activity was lower, with 41 applications and 15 granted patents. By 2024, filings increased to 194, while grants rose to 94. This gap suggests a competitive race to file new AI concepts. Although applications declined slightly to 177 in 2025, granted patents reached a peak of 100. This pattern indicates that filings from the surge over the past few years are progressing through the approval process into granted patents.

Who is Leading in AI Patent Activity?

Figure 3.3.3 – Top 10 Owners of Granted AI Patents in Gaming

Source: Lens.org

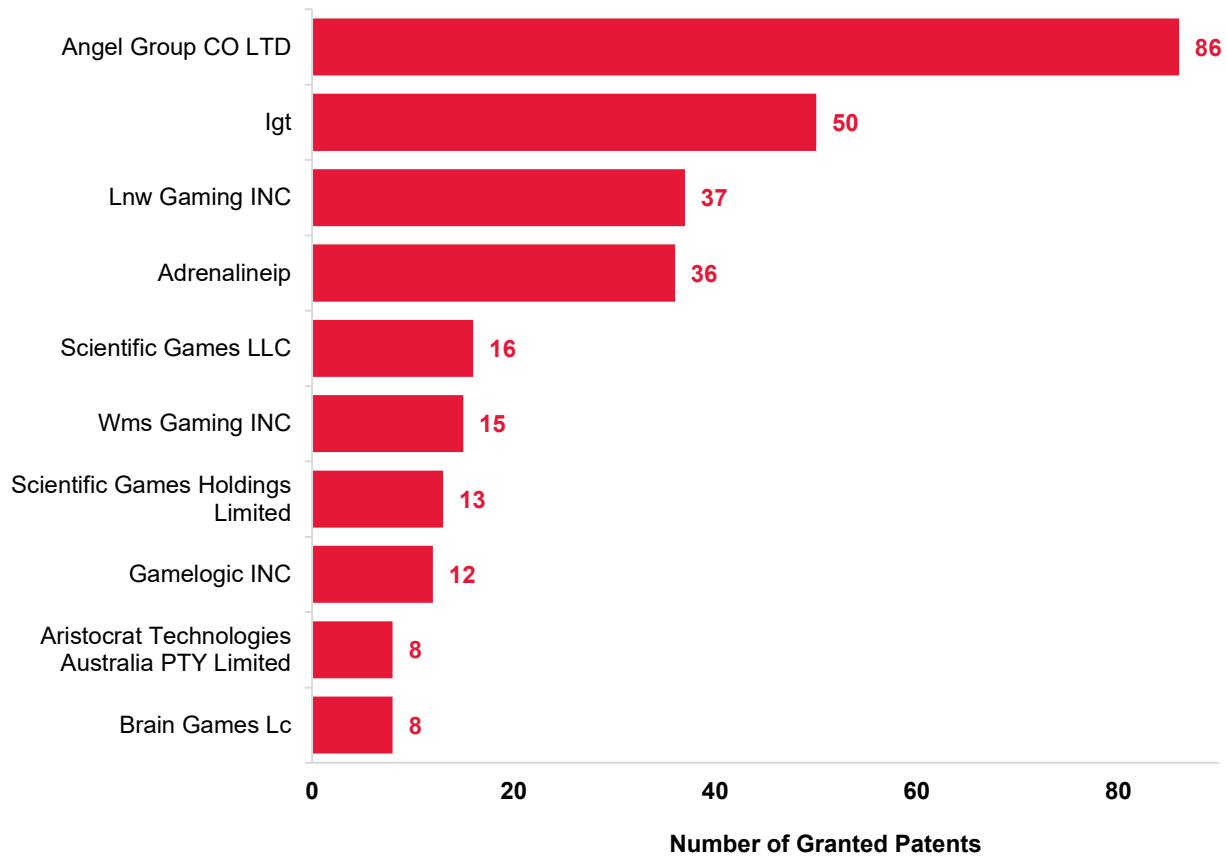



Figure 3.3.3 highlights the top 10 companies with the most AI-related patents in the gambling industry. Notably, all of the leading patent holders are technology suppliers to the industry; companies like IGT, Light & Wonder, and Scientific Games that develop gaming systems, platforms, and content, rather than operators. This suggests that AI innovation, at least as captured through our corpus, is being driven primarily from the supply side of the industry. Over the course of 2010 to 2025, Angel Group CO LTD emerged as the leader with 86 owned patents, followed by IGT (50), Light & Wonder (37), AdrenalineIP (36), and Scientific Games LLC (16). This is further reflected in the applicant data (Figure 3.3.4) where Angel Group has consistently led the field, followed by competitors AdrenalineIP and IGT.

Figure 3.3.4 – Top 10 Applicants of AI Patents in Gaming

Source: Lens.org



AI Patent Activity by Country

Figure 3.3.5 – Jurisdictions: AI Patents in Gaming

Source: Lens.org




Figure 3.3.5 illustrates the geographic distribution of AI patent filings in the gambling industry. As of 2025, the United States clearly dominates the landscape, accounting for 1,304 filings (61.4%) of all AI-related gaming patent documents worldwide. The World Intellectual Property Organization (WIPO) represents the second-largest filing destination with 207 international applications (9.7%), followed closely by Australia with 149 (7.0%). Unlike country-specific filings, WIPO applications are submitted through the Patent Cooperation Treaty (PCT) system, which allows applicants to simultaneously seek patent protection across multiple countries without being tied to a single national jurisdiction at the point of filing.

AI Patent Activity by Technical Theme

Figure 3.3.6 – Top 15 Technical Themes in AI Gaming Patents (CPC Analysis)

Source: Lens.org





Figure 3.3.6 presents a technical breakdown of the patent corpus using Cooperative Patent Classification (CPC) frequency analysis. The data reveals that Coin-Freed Gaming Apparatus (Slots) (G07F17/32) is the primary focus, leading the sector with 662 patents. Security and Fraud Detection (G07F17/3241) follows closely as the second-largest category (614 patents). Significant innovation is then concentrated in Betting and Bookmaking and Player Tracking systems.

Most-Cited AI Gambling IP

Figure 3.3.7⁸ - Top 10 Most Cited AI Gambling Patents (Granted Only)

Source: Lens.org



⁸ Readers may observe dates (e.g., 2001, 2006) listed next to certain patent titles that precede the reports 2010 - 2025 scope. These dates reflect the original application filings. All patents included in this analysis were officially granted or published within the study's designated timeframe.



The surge in recent AI patent activity by gaming technology suppliers signals a shift toward embedding intelligence in gaming infrastructure, including smart tables. The most influential patents emphasize personalization and loyalty, underscoring that future competition will center on delivering more responsive, data-driven player experiences in real time.”

Stephen Moore

CEO, Walker Digital Table Systems

Figure 3.3.7 illustrates the most frequently cited granted patents, showcasing the key foundational patents of industry innovation. Walker Digital LLC leads the field with 345 citations for a 2002 patent covering a system for planning and customizing a gaming vacation experience. The patent describes a solution that allows players to set preferences in advance that are then used to configure gaming devices and deliver personalized marketing offers at the casino. Conquer Your Addiction LLC follows with 306 citations for a 2020 patent on using wireless location data and contextual signals to automatically monitor, predict, and intervene when a person is at risk of engaging in pre-identified problematic behaviors such as addiction. And IGT's flexible loyalty points patent (238 citations), filed in 2001, describes a network-based system enabling patrons to earn and redeem loyalty points across a wide range of casino and non-gaming activities, even without providing identification. These high citation counts indicate foundational technologies that are frequently referenced in subsequent patent filings.

3.4 Macro News Trends and Outlook

We analyzed the gambling trade industry news to understand what AI related topics made the headlines in 2025, using these and more recent news from 2026 to determine key macro-AI industry trends.

We retrieved AI-related trade news in 2025 from several major gambling news websites, namely: Sigma, SBC, Next, iGaming Business, EGR Intel, +More Media, and CDC Gaming. We employed a two-stage retrieval process: first, we executed keyword searches via Google Search using AI-related terms (e.g., “artificial intelligence”, “machine learning”, “generative AI”), and second, we applied a combination of human review and LLM-assisted screening to verify that retrieved articles were substantively related to AI. We further supplemented this corpus with relevant and more recent (2026) news stories known to the research team. In total, we retrieved 151 news articles, which formed the basis for this AI Macro News Trends and Outlook. We provide a summary of our corpus by category in Table 3.4.1 below.

Table 3.4.1

Category	Sub Categories	Total Articles	% of Total
Technology	Software development, AI strategy, platform innovation, security, market intelligence	46	30.5%
Policy & Regulation	Regulatory guidance, AML, fraud, ethics, consumer protection, responsible gambling	39	25.8%
CRM	Personalization, marketing technology, customer support, affiliate tools, SEO	26	17.2%
Product	Sports Betting, Slots, Live Dealer, Prediction Markets, Esports	25	16.6%
Operations	Casino floor optimization, cost reduction, workforce transformation, surveillance	15	9.9%
TOTAL		151	100%

The most common category of news was *Technology*, consistent with findings from **Section 1: The Industry Index**, where technology-related applications were among the most highly endorsed AI use cases (see Figure 1.5). Several articles in this category highlighted companies moving beyond AI speculation and toward more concrete implementation. For example, Sportradar added product initiatives to their Chief AI Officer's remit, and noted AI will be a potential driver of growth towards its €1.7 billion 2027 revenue target [45], [46]. DraftKings co-founder Paul Liberman similarly argued that companies who treat AI as "optional" risk falling behind [47], while OpenBet initiated a recruitment drive to build a dedicated AI team aimed at setting a benchmark in AI-driven sportsbook product development [48].

There was also a strong focus on *Policy & Regulation*. Notable developments included the UK Gambling Commission warning operators that AI-generated deepfakes are being used to circumvent customer due diligence processes [49], [50], while the International Gaming Standards Association (IGSA) began developing AI best-practice guidelines for gambling regulators [51]. In Europe, Italy's ADM and Sogei deployed AI to improve transparency and responsible gambling oversight [52].

We note it is interesting to see that AI news related to *Product* was only 1 behind *CRM* in 2025. Whilst *CRM* has arguably been the major business use case of ML and data analytics in the gambling industry over the past decade, this might suggest that the emergence of more powerful and generically applicable AI tooling is redirecting industry attention toward new product development opportunities; a finding corroborated by our survey in **Section 1: The Industry Index**, where *Product Development & Innovation* ranked among the most highly endorsed AI application areas (see Figure 1.5).

Within the *Product* category, approximately one-third of news articles centered on opportunities within Live Dealer and Prediction Markets. In Live Dealer, the traditional model is at risk of being disrupted by AI-driven virtual hosts, by lowering capital expenditures and offering infinite scalability. At this point it is too early to state whether this will displace "traditional" Live Dealer or become a new vertical in itself (not unlike virtual sports, which has come to co-exist with sports betting). Similarly, prediction markets are attracting AI-enabled innovation. For example, AI start-up Slips is leveraging AI and agentic discovery tools to deliver novel experiences to consumers by creating new "micro contests" based on trending events; the AI component acts autonomously—scanning for trends, creating markets, and settling outcomes—without human intervention for each step [53]. And prediction market platform Kalshi formalised a partnership with xAI, integrating the Grok AI model for user insights and market analysis [54], [55]

Additionally, we observed news coverage surrounding the impacts of agentic AI on customer experience and the workforce. For example, Better Collective launched an AI-powered tool that generates personalized sports betting content and recommendations for its users [56]. And DraftKings CEO Jason Robins disclosed a hiring freeze driven by the company's confidence in agentic AI, describing a cultural shift from managing employees to managing AI agents [57].

The industry is also moving toward automated slot development by leveraging generative models for art, math, and thematic narrative. Studios like OpenSLots, SlotGPT and Xgenia are reducing production cycles from months to days [58], [59]. This "content factory" approach allows for the hyper-targeted rollout of unique titles, optimized for specific markets and player micro-segments. Data from eGaming Monitor's Casino Review of 2025 suggests the number of casino and slots suppliers grew by 24% in 2025 compared with 2024 whilst the total number of unique casino games increased 29% in 2025 compared with 2024 [60]. This could suggest that AI is reducing barriers to entry to this segment of the supplier value chain.

Within the *CRM* category, nearly a quarter of all news articles (23%) related to the use of AI within customer support. For example, Tugi Tark provides an AI customer service platform repositioning online casino support as a proactive retention channel [61], [62], while ChatForm AI rebranded as Cevro AI with

a focus on “empathy at scale” [63]. GamingSoft launched both an AI Chatbot and AI Voice product for automated text and voice customer support [64]. Beyond customer support, AI-driven CRM solutions received coverage, with companies such as Fast Track announcing an AI-driven natural language platform [65], [66], and IGT introducing TrueAim at G2E, a platform facilitating “hyper-personalized” player experiences [67].

The Operations category, while small, was a thematically cohesive area, with much of the coverage focusing on casino floor optimization and management, highlighting companies such as Light & Wonder's Engage AI platform [68], QCI's AGI56 unified intelligence dashboard [69], and Xailient's Casino Eye-D facial recognition for players clubs [70].

3.5 Startup Activity



Editorial note: The following subsection is contributed by Ryan Rubio of Zero Labs, a venture firm focused on early-stage gambling technology companies. Unlike the preceding subsections, which are based on systematic data collection and analysis, this piece draws on Zero Labs' portfolio observations and broader monitoring of the startup landscape to offer a qualitative perspective on how AI is impacting gambling startups.

The preceding sections of this report have documented how AI has become a subject of growing academic inquiry, an increasingly prominent topic at industry conferences, and a driver of patent activity across the gambling sector. This section turns to the companies at the frontier of that activity: the early-stage, high-growth startups that are building AI-native products and services for the gambling industry.

Drawing on observations from Zero Labs' portfolio of early-stage gambling technology companies and broader monitoring of the startup landscape, this section examines where entrepreneurial activity is concentrated, the structural factors shaping what startups can and cannot do, and how AI is changing the economics of building a gambling technology company from the ground up.

It is worth emphasizing at the outset that while generative AI represents the current inflection point in public and industry attention, it is only the latest crest of a wave that has been building for over a decade. Startups leveraging machine learning, predictive modeling, and data-driven personalization have been operating in the gambling space well before the release of ChatGPT. What has changed is the intensity of interest and the breadth of capabilities now available to founders. This section examines both sides of that equation: AI as the basis for new products and services, and AI as a tool that is reshaping the process of building a startup itself.

Products and Services

Perhaps the most important observation about AI startups in the gambling industry is what has *not* happened. Despite the surge of interest in AI among the general public, driven primarily by the rise of consumer-facing large language models, the gambling sector has not yet produced the kind of AI-native startup that rapidly displaces incumbent operators or suppliers. No new entrant has leveraged generative AI to fundamentally redraw the competitive landscape in the way that, for example, DraftKings and FanDuel built on their existing digital-first fantasy sports platforms and rapidly captured market share in sports betting following the repeal of PASPA in 2018. Even that disruption saw incumbents respond effectively through structures such as BetMGM, a joint venture combining legacy brand and distribution assets with digital-native technology. If established operators could absorb a regulatory earthquake of

that magnitude and remain competitive, technological changes alone, without a corresponding structural opening, are unlikely to unseat them.

This absence is not incidental. It reflects structural features of the gambling industry that are particularly resistant to the kind of disruption that AI startups have achieved in less regulated sectors. Two factors stand out.

First, barriers to entry remain exceptionally high. Bringing a new licensed operator to market continues to be a prohibitively expensive and time-consuming effort for potential market entrants, regardless of whatever technological advances are occurring. Second, distribution advantages remain firmly with incumbents. The most valuable asset in the gambling industry is not technology but access to licensed markets, established customer bases, and historically accumulated proprietary behavioral data. These data assets function as a competitive moat. In gambling, where the largest operators possess decades of transactional and behavioral data across millions of customers, this advantage is particularly pronounced.

As of early 2026, current AI capabilities do not materially lower either of these hurdles for a new startup.

This does not mean that AI is irrelevant to the competitive dynamics of the industry. To the contrary, it appears to be functioning as what innovation theorists would characterize as a sustaining innovation, a technology that enables incumbent firms to improve their existing products and operations rather than one that creates entirely new markets or value networks [71]. In the gambling sector, this manifests in operators using machine learning to refine personalization engines, improve risk management, and optimize marketing spend, meaningful improvements, but ones that reinforce rather than disrupt existing market structures.

That said, there are signs that the conditions for disruption may be forming. Significant advances in AI-driven content generation and localization are enabling the rapid creation and translation of content across geographies, languages, and channels. For startups that can solve the regulatory access problem, these capabilities could dramatically alter the economics of market entry by reducing the marginal cost of content production, historically one of the more labor-intensive components of launching in a new jurisdiction. SlotGPT, a Cyprus-based startup, offers a striking illustration: its prompt-to-production platform uses generative AI to allow users to create fully playable slot games, including visuals, mechanics, and audio, from a single text prompt, and has already generated nearly 27,000 games with an integration to Stake.com [72]. The question is whether any startups like SlotGPT can combine these AI-enabled cost advantages with a viable path through the regulatory and distribution barriers described above.

Where the most promising startup activity is emerging, however, is not in generative AI but in the application of advanced machine learning and data science to specific segments of the gambling value chain. A distinct category of startups has appeared that are native to data and best-in-class machine learning techniques, applying them to the gambling industry in ways that address specific player segments or market opportunities. nVenue, for example, uses real-time predictive models to scale in-play micro-betting, while companies like Oddit and Sharpstakes are applying machine learning to build data-driven tools that deliver deep personalization to bettors with significantly reduced manual intervention compared to traditional workflows. What distinguishes these startups from their predecessors is not simply the use of AI as a feature but a fundamentally different architecture. Their products are designed from the ground up around data pipelines and model inference, enabling levels of scalability and personalization that require significantly less manual intervention than the methods historically employed by legacy operators. They are, in effect, building for a different cost curve, one where the marginal cost of serving an additional player or generating an additional piece of content approaches zero.

Internal Use: AI as a Tool for Building Companies

Underpinning the case for the long-term competitiveness of AI-native startups is a less visible but potentially more consequential trend: the impact of AI on the process of building a company itself. The cost and speed of building software are changing, and the founders who understand how to leverage these tools effectively are compounding that advantage over time.

For gambling technology startups in particular, this has practical implications. Small teams can now build and ship products that previously required significantly larger engineering organizations. Iteration cycles are shorter, allowing startups to test market hypotheses more rapidly. And the fixed costs of software development — historically a significant barrier for bootstrapped or seed-stage companies — are declining in real terms.

Over time, these advantages are likely to compound. Startups that are AI-native in their development practices, not just in their products, will be able to scale more efficiently, adapt to regulatory changes more quickly, and compete with legacy software developers on quality even without the headcount advantages that incumbents have traditionally enjoyed. The consumer-facing product need not be “AI-powered” for the company to derive significant competitive advantage from AI in how it builds, tests, and deploys software.



AI development tools reward teams that already know what good looks like. We've built a deep data foundation over years, and now our team can ship full consumer and enterprise products — Betflow, Bracket Party — on top of it at a pace that used to require massive orgs. The bottleneck at most companies is communication overhead and coordination drag. When your team works in lockstep and the tools accelerate every cycle, you can cover an insane amount of ground.”

Matt Bresler
CEO, Oddit

Looking Ahead

The gambling technology startup landscape as of 2026 is defined by a tension. On one side, the structural barriers of the industry, regulatory complexity, data advantages held by incumbents, and the high cost of market entry, continue to favor established players. On the other hand, a cohort of startups is emerging that are architecturally distinct from their predecessors, built around data and machine learning in ways that may, over a longer time horizon, alter the competitive dynamics of the sector.

The data presented throughout this section supports the view that this tension is productive. Patent filings have surged since 2020, with annual grants rising from 15 in 2010 to over 100 by 2025 (Section 3.3). Academic publications at the intersection of AI and gambling have more than doubled their share of the field's total research output since 2022 (Section 3.1). AI-focused sessions at major industry conferences have grown from a handful in 2020 to over 80 by 2025 (Section 3.2). Startup activity and venture capital investment in gambling technology have followed a consistent trajectory, with capital flowing specifically to companies at the intersection of data, AI, and gambling. Taken together, these indicators suggest that the industry's innovation pipeline is not only active but accelerating, and that investors, researchers, and entrepreneurs are converging on the same thesis about where value will be created.

Whether the current generation of AI-native startups ultimately proves to be sustaining or disruptive to the gambling industry's incumbents will depend not on the capabilities of the AI itself, but on whether these companies can pair technological advantages with viable strategies for navigating the industry's regulatory and distribution realities. The technology is here and improving rapidly. The competitive advantage will belong to whichever companies — incumbent or insurgent — can pair AI capabilities with viable strategies for navigating the industry's regulatory and distribution realities first.

References

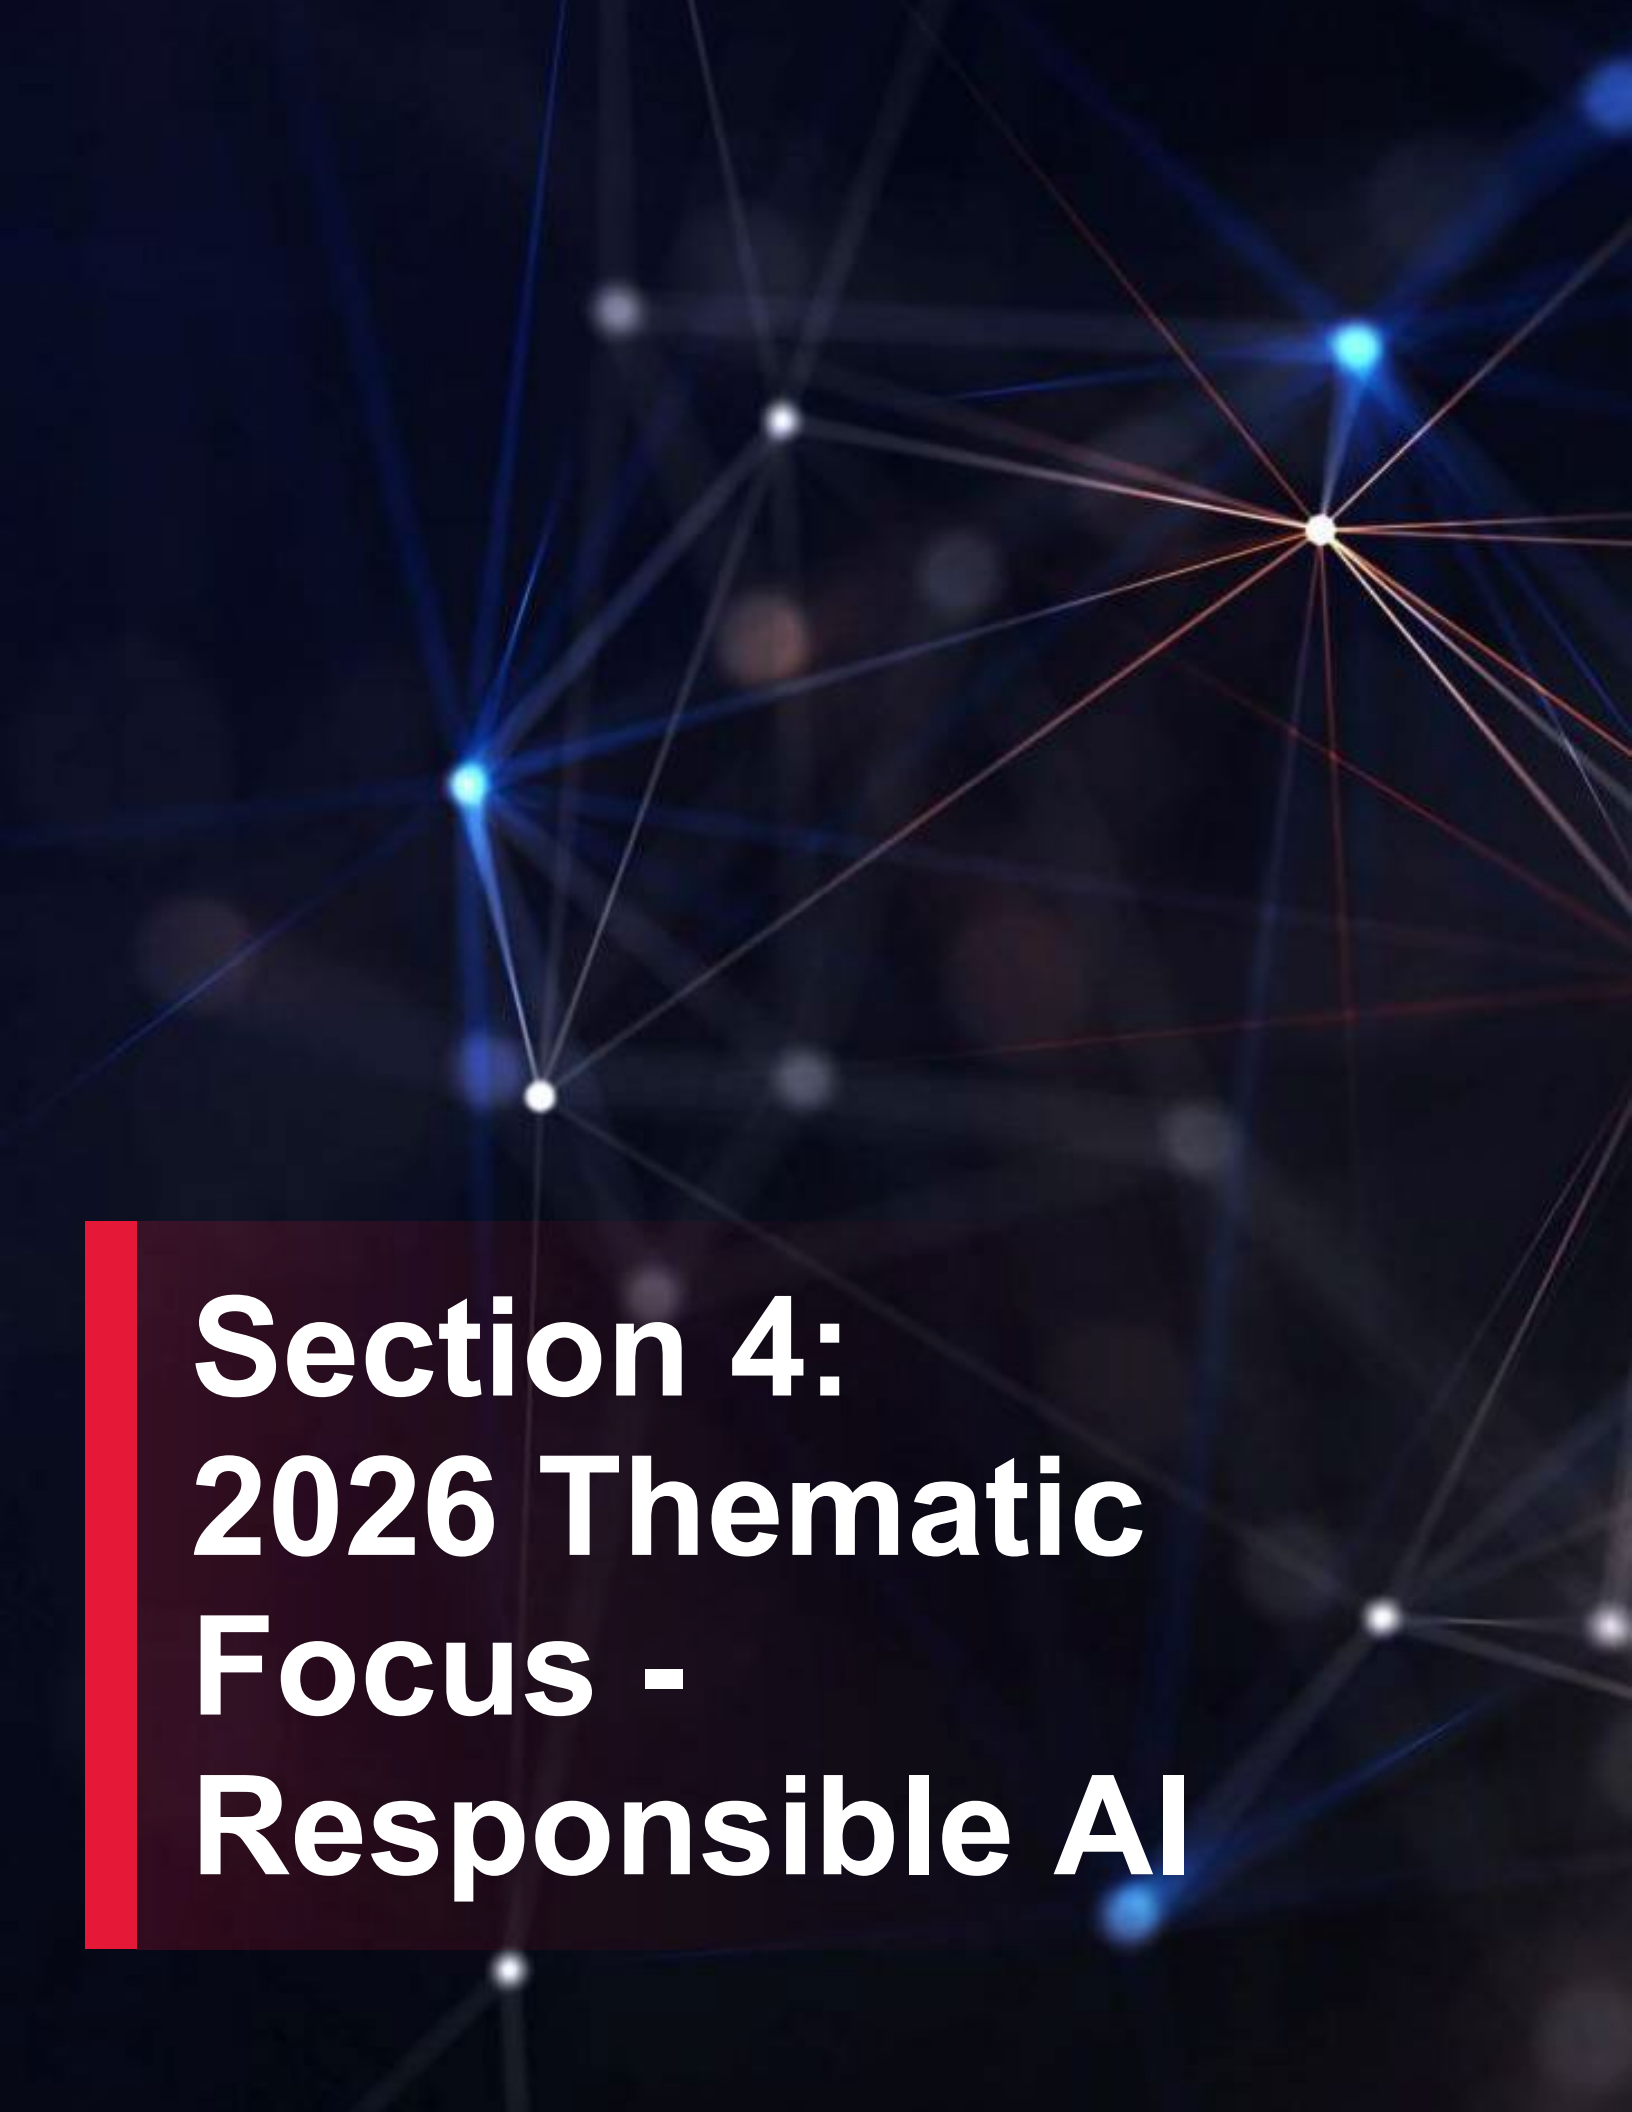
- [1] J. Dean, "A Golden Decade of Deep Learning: Computing Systems & Applications," *Daedalus*, vol. 151, no. 2, pp. 58–74, May 2022, doi: 10.1162/daed_a_01900.
- [2] Md. Faiyazuddin *et al.*, "The Impact of Artificial Intelligence on Healthcare: A Comprehensive Review of Advancements in Diagnostics, Treatment, and Operational Efficiency," *Health Sci. Rep.*, vol. 8, no. 1, p. e70312, Jan. 2025, doi: 10.1002/hsr2.70312.
- [3] S. Bahoo, M. Cucculelli, X. Goga, and J. Mondolo, "Artificial intelligence in Finance: a comprehensive review through bibliometric and content analysis," *SN Bus. Econ.*, vol. 4, no. 2, p. 23, Jan. 2024, doi: 10.1007/s43546-023-00618-x.
- [4] R. P. Bunker and F. Thabtah, "A machine learning framework for sport result prediction," *Appl. Comput. Inform.*, vol. 15, no. 1, pp. 27–33, Jan. 2019, doi: 10.1016/j.aci.2017.09.005.
- [5] T. Horvat and J. Job, "The use of machine learning in sport outcome prediction: A review," *WIREs Data Min. Knowl. Discov.*, vol. 10, no. 5, p. e1380, Sep. 2020, doi: 10.1002/widm.1380.
- [6] R. P. Schumaker, A. T. Jarmoszko, and C. S. Labedz, "Predicting wins and spread in the Premier League using a sentiment analysis of twitter," *Decis. Support Syst.*, vol. 88, pp. 76–84, Aug. 2016, doi: 10.1016/j.dss.2016.05.010.
- [7] W. S. Murch, S. Kairouz, S. Dauphinais, E. Picard, J.-M. Costes, and M. French, "Using machine learning to retrospectively predict self-reported gambling problems in Quebec," *Addiction*, vol. 118, no. 8, pp. 1569–1578, 2023, doi: 10.1111/add.16179.
- [8] J. Braverman, D. A. LaPlante, S. E. Nelson, and H. J. Shaffer, "Using cross-game behavioral markers for early identification of high-risk internet gamblers," *Psychol. Addict. Behav.*, vol. 27, no. 3, pp. 868–877, 2013, doi: 10.1037/a0032818.
- [9] E. Smith, J. Peters, and N. Reiter, "Automatic detection of problem-gambling signs from online texts using large language models," *PLOS Digit. Health*, vol. 3, no. 9, p. e0000605, Sep. 2024, doi: 10.1371/journal.pdig.0000605.
- [10] E. Zhao, R. Yan, J. Li, K. Li, and J. Xing, "AlphaHoldem: High-Performance Artificial Intelligence for Heads-Up No-Limit Poker via End-to-End Reinforcement Learning," *Proc. AAAI Conf. Artif. Intell.*, vol. 36, no. 4, pp. 4689–4697, Jun. 2022, doi: 10.1609/aaai.v36i4.20394.
- [11] Z. Wu, K. Li, H. Xu, Y. Zang, B. An, and J. Xing, "L2E: Learning to Exploit Your Opponent," in *2022 International Joint Conference on Neural Networks (IJCNN)*, Jul. 2022, pp. 1–8. doi: 10.1109/IJCNN55064.2022.9892077.
- [12] R. Zhuang, A. Gupta, R. Yang, A. Rahane, Z. Li, and G. Anumanchipalli, "PokerBench: Training Large Language Models to Become Professional Poker Players," *Proc. AAAI Conf. Artif. Intell.*, vol. 39, no. 24, pp. 26175–26182, Apr. 2025, doi: 10.1609/aaai.v39i24.34814.
- [13] D. A. Worthy, M. J. Hawthorne, and A. R. Otto, "Heterogeneity of strategy use in the Iowa gambling task: A comparison of win-stay/lose-shift and reinforcement learning models," *Psychon. Bull. Rev.*, vol. 20, no. 2, pp. 364–371, Apr. 2013, doi: 10.3758/s13423-012-0324-9.
- [14] H. Sullivan-Toole, N. Haines, K. Dale, and T. M. Olino, "Enhancing the Psychometric Properties of the Iowa Gambling Task Using Full Generative Modeling," *Comput. Psychiatry*, vol. 6, no. 1, Aug. 2022, doi: 10.5334/cpsy.89.
- [15] Y. Hu and S. Shoemaker, "Do More Experienced Gamblers Choose Slot Machines with Better Odds? A Large-Scale Multi-Armed Bandit Problem at a Casino," *Cust. Needs Solut.*, vol. 11, no. 1, p. 9, Oct. 2024, doi: 10.1007/s40547-024-00150-5.

- [16] Q. Tao *et al.*, “Decision-Feedback Stages Revealed by Hidden Markov Modeling of EEG,” *Int. J. Neural Syst.*, vol. 31, no. 07, p. 2150031, Jul. 2021, doi: 10.1142/S0129065721500313.
- [17] M. Min, J. J. Lee, H. Park, and K. Lee, “Detecting Anomalous Transactions via an IoT Based Application: A Machine Learning Approach for Horse Racing Betting,” *Sensors*, vol. 21, no. 6, p. 2039, Jan. 2021, doi: 10.3390/s21062039.
- [18] “AI-driven Optimization of Casino Gaming Systems for Fraud Detection and User Behavior Analysis – IJSREM.” Accessed: Feb. 25, 2026. [Online]. Available: <https://ijsrem.com/download/ai-driven-optimization-of-casino-gaming-systems-for-fraud-detection-and-user-behavior-analysis/>
- [19] W. Cui and W. Q. Yan, “A Scheme for Face Recognition in Complex Environments,” *Int. J. Digit. Crime Forensics IJDCF*, vol. 8, no. 1, pp. 26–36, 2016, doi: 10.4018/IJDCF.2016010102.
- [20] Y. Zhang, X. Fu, R. Yang, and Y. Li, “DRSDetector: Detecting Gambling Websites by Multi-level Feature Fusion,” in *2023 IEEE Symposium on Computers and Communications (ISCC)*, Jul. 2023, pp. 1441–1447. doi: 10.1109/ISCC58397.2023.10217923.
- [21] J. Chen, S. Zheng, Y. Cheng, and Z. Zhang, “Data mining based analysis of online gambling sites and illicit financial flows,” in *Proceedings of the 2024 International Conference on Cloud Computing and Big Data*, in ICCBD '24. New York, NY, USA: Association for Computing Machinery, Oct. 2024, pp. 205–211. doi: 10.1145/3695080.3695116.
- [22] F. Binesh and K. Ghaharian, “Identifying Risks and Ethical Considerations of AI in Gambling: A Scoping Review,” *Int. J. Hosp. Tour. Adm.*, vol. 0, no. 0, pp. 1–31, Apr. 2025, doi: 10.1080/15256480.2025.2494575.
- [23] K. Ghaharian, F. Binesh, M. Soligo, L. Golab, and B. Abarbanel, “AI ethics in a controversial industry: the case of gambling and its ethical paradox,” *AI Ethics*, vol. 5, no. 2, pp. 1685–1701, Apr. 2025, doi: 10.1007/s43681-024-00520-8.
- [24] N. Binesh, K. Ponnada, and A. Syah, “The Future of the Gambling Industry is AI: Insights from Expert Interviews on Human-AI Collaboration, Regulation, and Ethics,” *J. Gamb. Stud.*, Sep. 2025, doi: 10.1007/s10899-025-10422-x.
- [25] K. Ramachandran, “The role of artificial intelligence in casino digital transformation and innovation,” *World J. Adv. Eng. Technol. Sci.*, vol. 15, no. 1, pp. 459–462, 2025, doi: 10.30574/wjaets.2025.15.1.0254.
- [26] N. Brown and T. Sandholm, “Superhuman AI for multiplayer poker,” *Science*, vol. 365, no. 6456, pp. 885–890, Aug. 2019, doi: 10.1126/science.aay2400.
- [27] H. Wardle *et al.*, “The Lancet Public Health Commission on gambling,” *Lancet Public Health*, vol. 9, no. 11, pp. e950–e994, Nov. 2024, doi: 10.1016/S2468-2667(24)00167-1.
- [28] A. S. F. 20 2025 in Society, “The Most Popular Sports In The World,” WorldAtlas. Accessed: Feb. 12, 2026. [Online]. Available: <https://www.worldatlas.com/articles/what-are-the-most-popular-sports-in-the-world.html>
- [29] A. Vaswani *et al.*, “Attention Is All You Need,” Aug. 02, 2023, *arXiv*: arXiv:1706.03762. doi: 10.48550/arXiv.1706.03762.
- [30] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, “Improving Language Understanding by Generative Pre-Training.”
- [31] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” May 24, 2019, *arXiv*: arXiv:1810.04805. doi: 10.48550/arXiv.1810.04805.

- [32] W. S. Murch, "AI for RG: New Opportunities for Online Gambling Harm Prevention," presented at the Institute Conference 2025: Insight to Impact – Advances in Gambling Research, Banff, Alberta, Canada, 2025. [Online]. Available: <https://ucalgary.scholaris.ca/server/api/core/bitstreams/04d2e56e-1330-4410-97f0-9b62d3a508ad/content>
- [33] K. Ghaharian, M. Soligo, R. Young, L. Golab, S. W. Kraus, and S. Wells, "Can Large Language Models Address Problem Gambling? Expert Insights from Gambling Treatment Professionals," *J. Gamb. Stud.*, Oct. 2025, doi: 10.1007/s10899-025-10430-x.
- [34] U. Kabir, "15 Biggest Gambling Countries In The World," Yahoo Finance. Accessed: Feb. 12, 2026. [Online]. Available: <https://finance.yahoo.com/news/15-biggest-gambling-countries-world-134430426.html>
- [35] M. Jepma and S. Nieuwenhuis, "Pupil diameter predicts changes in the exploration-exploitation trade-off: evidence for the adaptive gain theory," *J. Cogn. Neurosci.*, vol. 23, no. 7, pp. 1587–1596, Jul. 2011, doi: 10.1162/jocn.2010.21548.
- [36] N. Brown and T. Sandholm, "Superhuman AI for heads-up no-limit poker: Libratus beats top professionals," *Science*, vol. 359, no. 6374, pp. 418–424, Jan. 2018, doi: 10.1126/science.aao1733.
- [37] N. Brown and T. Sandholm, "Superhuman AI for multiplayer poker," *Science*, vol. 365, no. 6456, pp. 885–890, Aug. 2019, doi: 10.1126/science.aay2400.
- [38] M. Moravčík *et al.*, "DeepStack: Expert-level artificial intelligence in heads-up no-limit poker," *Science*, vol. 356, no. 6337, pp. 508–513, May 2017, doi: 10.1126/science.aam6960.
- [39] C. Wang, M. Zhang, F. Shi, P. Xue, and Y. Li, "A Hybrid Multimodal Data Fusion-Based Method for Identifying Gambling Websites," *Electronics*, vol. 11, no. 16, p. 2489, Aug. 2022, doi: 10.3390/electronics11162489.
- [40] L. Li, G. Gou, G. Xiong, Z. Cao, and Z. Li, "Identifying Gambling and Porn Websites with Image Recognition," in *Advances in Multimedia Information Processing – PCM 2017*, vol. 10736, B. Zeng, Q. Huang, A. El Saddik, H. Li, S. Jiang, and X. Fan, Eds., in Lecture Notes in Computer Science, vol. 10736, Cham: Springer International Publishing, 2018, pp. 488–497. doi: 10.1007/978-3-319-77383-4_48.
- [41] Y. Chen, R. Zheng, A. Zhou, S. Liao, and L. Liu, "Automatic Detection of Pornographic and Gambling Websites Based on Visual and Textual Content Using a Decision Mechanism," *Sensors*, vol. 20, no. 14, p. 3989, Jul. 2020, doi: 10.3390/s20143989.
- [42] M. Auer and M. D. Griffiths, "Predicting High-Risk Gambling Based on the First Seven Days of Gambling Activity After Registration Using Account-Based Tracking Data," *Int. J. Ment. Health Addict.*, vol. 22, no. 6, pp. 3397–3413, Dec. 2024, doi: 10.1007/s11469-023-01056-4.
- [43] M. Auer and M. D. Griffiths, "Using artificial intelligence algorithms to predict self-reported problem gambling with account-based player data in an online casino setting," *J. Gamb. Stud.*, vol. 39, no. 3, pp. 1273–1294, Jul. 2022, doi: 10.1007/s10899-022-10139-1.
- [44] I. Ukhov, J. Bjurgert, M. Auer, and M. D. Griffiths, "Online Problem Gambling: A Comparison of Casino Players and Sports Bettors via Predictive Modeling Using Behavioral Tracking Data," *J. Gamb. Stud.*, vol. 37, no. 3, pp. 877–897, Sep. 2021, doi: 10.1007/s10899-020-09964-z.
- [45] J. Levy, "Sportradar lays out 2027 ambitions with €1.7bn revenue target - EGR Intel." Accessed: Mar. 26, 2026. [Online]. Available: [//www.egr.global/intel/news/sportradar-lays-out-2027-ambitions-with-e1-7bn-revenue-target/](https://www.egr.global/intel/news/sportradar-lays-out-2027-ambitions-with-e1-7bn-revenue-target/)

- [46] J. Levy, "Sportradar CEO says companies ignoring AI are 'not competitive' - EGR Intel." Accessed: Mar. 26, 2026. [Online]. Available: [//www.egr.global/intel/news/sportradar-ceo-says-companies-failing-to-use-ai-are-not-competitive/](https://www.egr.global/intel/news/sportradar-ceo-says-companies-failing-to-use-ai-are-not-competitive/)
- [47] J. O'Brien, "DraftKings co-founder: 'Those who treat AI as optional or secondary risk falling behind' - EGR Intel." Accessed: Mar. 26, 2026. [Online]. Available: [//www.egr.global/intel/insight/draftkings-co-founder-those-who-treat-ai-as-optional-or-secondary-risk-falling-behind/](https://www.egr.global/intel/insight/draftkings-co-founder-those-who-treat-ai-as-optional-or-secondary-risk-falling-behind/)
- [48] K. Simmonds, "Exclusive: OpenBet kick-starts recruitment drive to create a dedicated AI team - EGR Intel." Accessed: Mar. 26, 2026. [Online]. Available: [//www.egr.global/intel/news/exclusive-openbet-kick-starts-recruitment-drive-to-create-a-dedicated-ai-team/](https://www.egr.global/intel/news/exclusive-openbet-kick-starts-recruitment-drive-to-create-a-dedicated-ai-team/)
- [49] J. Levy, "Gambling Commission warns operators over AI deepfakes in due diligence process - EGR Intel." Accessed: Mar. 26, 2026. [Online]. Available: [//www.egr.global/intel/news/gambling-commission-warns-operators-over-ai-deepfakes-in-due-diligence-process/](https://www.egr.global/intel/news/gambling-commission-warns-operators-over-ai-deepfakes-in-due-diligence-process/)
- [50] C. Reynolds, "UK Gambling Commission warns of AI deepfakes in 'emerging' money laundering risks," iGB. Accessed: Mar. 26, 2026. [Online]. Available: <https://igamingbusiness.com/gaming/gaming-regulation/uk-gambling-commission-warns-of-ai-deepfakes-in-emerging-money-laundering-risks/>
- [51] C. Reynolds, "IGSA readies AI guidelines for gambling regulators, to better understand the technology," iGB. Accessed: Mar. 26, 2026. [Online]. Available: <https://igamingbusiness.com/tech-innovation/artificial-intelligence/igsa-readies-ai-guidelines-for-gambling-regulators-to-better-understand-the-technology/>
- [52] Tony Colapinto, "ADM and Sogei: artificial intelligence driving transparency and responsible innovation in Italy's public administration," Sigma. Accessed: Mar. 26, 2026. [Online]. Available: <https://sigma.world/news/adm-sogei-artificial-intelligence-gambling-prevention/>
- [53] E. Gibbs, "P2P wagering platform Slips adds AI-powered prediction markets," NEXT.io. Accessed: Mar. 26, 2026. [Online]. Available: <https://next.io/news/technology/slips-adds-ai-powered-prediction-markets/>
- [54] J. Levy, "Kalshi and xAI officially join forces after May mix-up - EGR Intel." Accessed: Mar. 26, 2026. [Online]. Available: [//www.egr.global/intel/news/kalshi-and-xai-officially-join-forces-after-may-mix-up/](https://www.egr.global/intel/news/kalshi-and-xai-officially-join-forces-after-may-mix-up/)
- [55] Charlotte Capewell, "Kalshi announces partnership with Elon Musk's xAI," CDC Gaming. Accessed: Mar. 26, 2026. [Online]. Available: <https://cdcgaming.com/brief/kalshi-announces-partnership-with-elon-musks-xai/>
- [56] E. Gibbs, "Better Collective launches new AI-powered tool for sports betting," NEXT.io. Accessed: Mar. 26, 2026. [Online]. Available: <https://next.io/news/technology/better-collective-launches-ai-betting-tool/>
- [57] J. Levy, "DraftKings CEO says hiring freeze in place with 'tremendous momentum' in agentic AI - EGR Intel." Accessed: Mar. 26, 2026. [Online]. Available: [//www.egr.global/intel/news/77498/](https://www.egr.global/intel/news/77498/)
- [58] Ingi, "OpenSlots to power new dawn of content creation with AI development platform," iGamingToday.com - News, insights and Slots review. Accessed: Mar. 26, 2026. [Online]. Available: <https://www.igamingtoday.com/openslots-to-power-new-dawn-of-content-creation-with-ai-development-platform/>
- [59] David Gravel, "XGENIA unveils next-gen AI platform set to reshape iGaming development," Sigma. Accessed: Mar. 26, 2026. [Online]. Available: <https://sigma.world/news/xgenia-unveils-ai-platform-to-reinvent-igaming-development/>

- [60] eGamingMonitor, "Casino review 2025," eGamingMonitor, Report, 2025. [Online]. Available: <https://egamingmonitor.com/casino-review-2025-report.pdf>
- [61] C. Mulheir, "AI-powered customer service firm Tugi Tark emerges from stealth mode," NEXT.io. Accessed: Mar. 26, 2026. [Online]. Available: <https://next.io/news/casino/tugi-tark-emerges-from-stealth-mode/>
- [62] NEXT.io, "From cost centre to competitive edge: Tugi Tark on how AI has created a new era of customer support," NEXT.io. Accessed: Mar. 26, 2026. [Online]. Available: <https://next.io/news/promoted/tugi-tark-ai-new-era-customer-support/>
- [63] NEXT.io, "ChatForm AI rebrands as Cevro AI, redefining iGaming support with empathy at scale," NEXT.io. Accessed: Mar. 26, 2026. [Online]. Available: <https://next.io/news/b2b-news/chatform-ai-rebrands-as-cevro-ai-redefining-igaming-support-with-empathy-at-scale/>
- [64] iGB content team, "GamingSoft launches AI Chatbot and AI Voice for enhanced customer support," iGB. Accessed: Mar. 26, 2026. [Online]. Available: <https://igamingbusiness.com/company-news/gamingsoft-launches-ai-chatbot-and-ai-voice-for-enhanced-customer-support/>
- [65] NEXT.io, "Fast Track AI to change the face of iGaming CRM," NEXT.io. Accessed: Mar. 26, 2026. [Online]. Available: <https://next.io/news/b2b-news/fast-track-ai-to-change-the-face-of-igaming-crm/>
- [66] EGR, "Fast Track: AI allows your CRM team to flourish - EGR Intel," EGR Global. Accessed: Mar. 26, 2026. [Online]. Available: <http://www.egr.global/intel/insight/fast-track-ai-allows-your-crm-team-to-flourish/>
- [67] John G. Brokopp, "Focus on IGT: Hyper-personalizing the player experience with TrueAim, IGT's adaptable AI solution," CDC Gaming. Accessed: Mar. 26, 2026. [Online]. Available: <https://cdcgaming.com/focus-on-igt-hyper-personalizing-the-player-experience-with-trueaim-igts-adaptable-ai-solution/>
- [68] Buck Wargo, "Focus on Light & Wonder: New Engage AI platform delivers casino insights in real time," CDC Gaming. Accessed: Mar. 26, 2026. [Online]. Available: <https://cdcgaming.com/focus-on-light-wonder-new-engage-ai-platform-delivers-casino-insights-in-real-time/>
- [69] Buck Wargo, "Focus on QCI: Powerful AI-driven AGI56 platform update integrates every aspect of the casino," CDC Gaming. Accessed: Mar. 26, 2026. [Online]. Available: <https://cdcgaming.com/focus-on-qci-powerful-ai-driven-agi56-platform-update-integrates-every-aspect-of-the-casino/>
- [70] Mark Gruetze, "Xailient's 'Casino Eye-D' brings facial recognition to players clubs," CDC Gaming. Accessed: Mar. 26, 2026. [Online]. Available: <https://cdcgaming.com/xailients-casino-eye-d-brings-facial-recognition-to-players-clubs/>
- [71] "What does Disruptive Innovation Theory have to say about AI?," Christensen Institute. Accessed: Mar. 02, 2026. [Online]. Available: <https://www.christenseninstitute.org/blog/what-does-disruptive-innovation-say-about-ai/>
- [72] playnewswire, "SlotGPT launches a new AI slot platform transforming players into creators," PlayNewswire | The First PR Platform for iGaming & Betting. Accessed: Mar. 05, 2026. [Online]. Available: <https://playnewswire.com/2025/12/30/slotgpt-launches-a-new-ai-slot-platform-transforming-players-into-creators/>



**Section 4:
2026 Thematic
Focus -
Responsible AI**

Section 4: 2026 Thematic Focus - Responsible AI

Overview

Responsible AI is about using AI in a way that is lawful, risk-aware, transparent, and accountable across its lifecycle, with particular attention to harm prevention and human oversight. Responsible AI principles are increasingly building out from good practice to regulatory requirements.

This section begins with a brief review of what Responsible AI means in a gambling sector context. We then collate examples of alleged AI misuse, explore regulator commentary on AI, and analyze how industry describes their approaches to Responsible AI. This section also includes topic spotlights, covering fraud and AI, the need for benchmarking in player risk algorithms, issues with algorithm fairness in risk detection, and how large language models handle problem gambling questions.

- Overview 76
- 4.1 Introduction 77
- 4.2 Reports of AI Incidents..... 80
- 4.3 Regulator Commentary on AI 86
- 4.4 Responsible AI in industry 88
- 4.5 Conclusion 92
- References 93

4.1 Introduction

“Responsible AI” (or RAI, as it is commonly referred to) is one of those terms that fails the negation test. After all, who would advocate for “Irresponsible AI”? In context, however, Responsible AI has come to refer to a loosely grouped set of practices, with the occasional attempt at a formal definition.

The OECD AI Principles have proved an influential framing for discussions around Responsible AI, including the EU AI Act (entered into force in August 2024, with specific provisions becoming live 2025-2027). The OECD AI Principles [1] were initially adopted in 2019 and updated in May 2024, setting out five values-based principles:

1. Inclusive growth, sustainable development and well-being
2. Human rights and democratic values, including fairness and privacy
3. Transparency and explainability
4. Robustness, security and safety
5. Accountability

With this framework, Responsible AI is using AI in a way that is lawful, risk-aware, transparent, and accountable across its lifecycle, with particular attention to harm prevention and human oversight.

Implementing these principles often requires balancing competing objectives (see topic spotlight on fraud and AI) and deciding what level of process or technology should be invested in.

National and international standards are being developed to help translate principles into practice, with prominent examples including the US government’s National Institute of Standards and Technology (NIST) launching the NIST AI Risk Management Framework (AI RMF) [2] and the International Organization for Standardization developing ISO/IEC 42001 [3] as a standard for Artificial Intelligence Management Systems (AIMS).

In this part of the report, we collate examples of alleged AI misuse (4.2), explore regulator commentary on Responsible AI (4.3), and analyze how operators describe their approaches to responsible AI (4.4). This section also includes special topic spotlights, covering fraud and AI, the need for benchmarking in player risk algorithms, issues with algorithm fairness in risk detection, and how large language models handle problem gambling questions.

In the gambling sector, Responsible AI requires dual attention on AI risks and AI benefits (see stakeholder perspectives synthesized in Figure 4.). AI risks include purposeful and unintended player exploitation and model bias, as well as cross-sector concerns related to data security and job displacement. AI benefits include algorithmic player protection strategies (such as player risk identification to drive

When ‘too good’ becomes a Responsible AI issue?

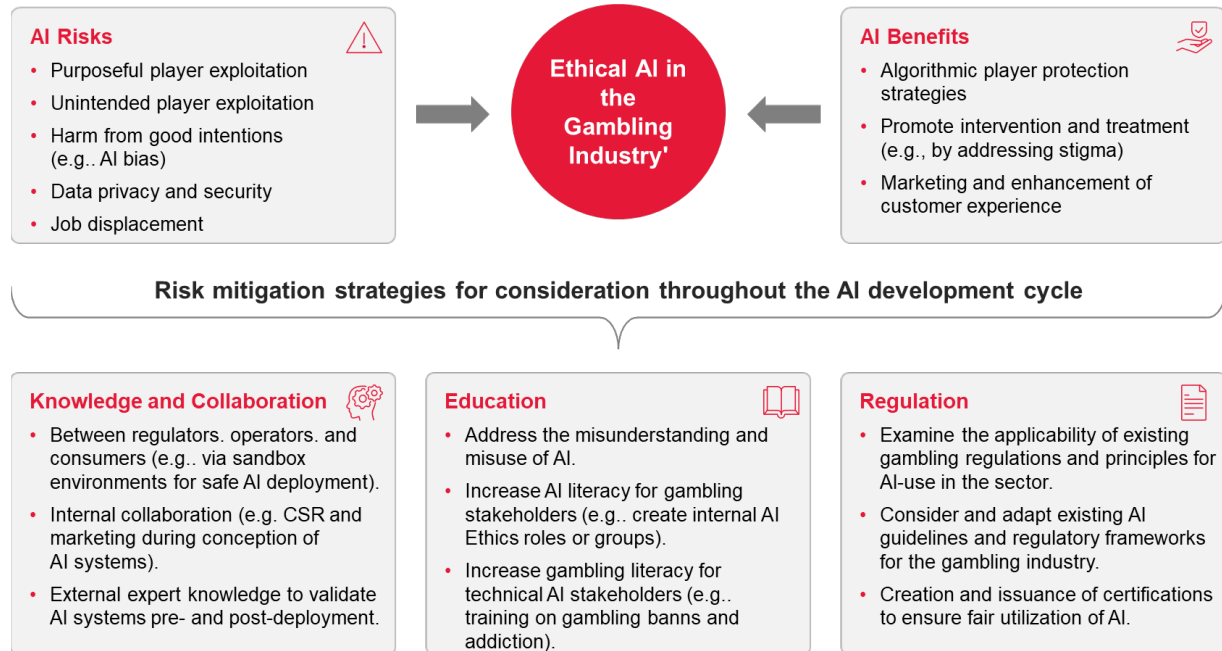
The power of AI can raise novel ethical issues for previously desirable use cases.

Consider personalized marketing. It is not new and not in itself problematic. Done well, players receive fewer but more relevant offers, encouraging them to redirect some of their leisure budget to gambling. In principle, AI simply improves a win-win scenario for both players and operators, providing even more relevant offers than the pre-existing statistical methods.

However, AI-powered personalization, A/B testing, and automated improvement cycles may prove too effective, encouraging an overspending that players later regret. Similar factors may occur in game or platform design. Marketing teams and others may need to ask themselves a new question: “When are we too good at our job, that restraint is needed to protect players?”

tailored interventions), enhanced interventions (such as chatbot interactions), and enhancement of the customer experience more generally.

Figure 4.1 - Aspects of responsible AI in the gambling sector, synthesized from stakeholder interviews



Used with permission of Springer Nature BV, from AI and Ethics, "AI ethics in a controversial industry: the case of gambling and its ethical paradox," Ghaharian, K., et al., Vol. 5, Issue 2, 2025; permission conveyed through Copyright Clearance Center, Inc.

Section 1: The Industry Index identified a key gap among industry survey respondents regarding the relatively low maturity of approaches to AI governance and the top AI concerns. Figure 1.13 in **Section 1: The Industry Index** highlighted failures in data privacy and governance as among the top three concerns. Part of the challenge lies in the complexity of AI governance, evolving regulatory attention, and challenging questions raised by research, as explored throughout this section.

Responsible AI Case Study: Risks with the right of appeal in AI-supported fraud detection

AI solutions can analyze vast amounts of data quickly to drive a decision and corresponding action: a significant advantage for high volume teams such as fraud. Regulators, such as the UK’s Financial Conduct Authority [4] have lauded the advantages of increased automation for tackling fraud, among other use cases. However, a lack of human oversight for customer appeals and redress can create challenges.

Without human oversight and intervention, an infinite-loop scenario can be created in which a person cannot escape an AI decision even if it has been concluded in error. The right of human intervention for certain decisions is also a right in law in some jurisdictions, such as under the EU’s GDPR [5] and similar legislation elsewhere with a right to object to automated decision making that has legal consequences.

AI tools for fraud detection use rules, pattern recognition, forensic analysis, and risk tolerances to make a decision about an account and take any necessary actions to mitigate fraud risk. In an example

Responsible AI Case Study: Risks with the right of appeal in AI-supported fraud detection

where a fraud team is constantly reviewing a large number of customer account transactions and other relevant information to determine whether there is a risk of fraud, it is sensible for AI models to power the determination of the customer's risk threshold and block account transactions in as near to real-time as possible to prevent further fraud from occurring.

But no system is perfect and legitimate customers can be incorrectly caught by AI fraud detection, resulting in unwanted decisions taken on their account to limit account functionality. In order to facilitate redress, the customer-facing team will need to be accessible to the customer to address the problem swiftly - and ideally there is also a continuous improvement loop identifying where such errors represent a systematic failing in the system. The skills required for this level of human oversight and continuous improvement are different to human oversight at a fraud operations level as it is more technical in nature, with implications for operator team design. AI incident response becomes significantly more onerous in these cases as models may need to be taken offline and business continuity processes instated.

A related element of Responsible AI is transparency. In the first instance, customers might reasonably ask to understand how an operator came to a decision, especially where it incorrectly introduced limitations on their account. Such rights are also enshrined in legislation in some jurisdictions. However, two issues result. First, the most powerful algorithms that best protect customers might be 'black box' in nature, where the actual reasoning behind the AI model's decision becomes difficult to interpret due to its complex algorithms and hard for front-line workers to convey. The second compromise is more distinctive to fraud: transparency needs to be balanced with tool security. Providing people with information about how an AI model has detected fraudulent activity could lead to them or other malicious actors bypassing these controls at a later date.

Notwithstanding arguments that limit transparency, some circumstances may require carefully unpicking decisions the AI model has made, e.g. if they presented a significant impact to customers or introduced commercial blockers. In extreme circumstances, customers could seek redress through litigation particularly if AI decisions have propagated across organizations, platforms, and systems leading to negative personal consequences.

The right of appeal is an important part of the AI risk mitigations strategy. There is evidence of how AI decisions impact people and organizations from different industries. A case in Amsterdam against Uber found that drivers had been dismissed using an algorithmic process after Uber believed the driver's fares were fraudulent [6]. The drivers were not offered a right to appeal, leading them to bring a case under GDPR. The courts subsequently instructed Uber to reactivate the drivers on the Uber platform and make payments to the drivers. A recent case in Turkey where gambling transactions were linked to anti-money laundering and consumer finance access may be a sign that appeal issues will grow in prominence in the gambling sector (see Table 4.1).

In summary, AI tools for important commercial functions like fraud have huge advantages for limiting customer and operator losses. However, first line teams need to be equipped with the tools to offer an appropriate level of redress without giving away information that can be used to circumvent fraud detection models. Technical skills are needed to interrogate models creating the need for a depth of skills within operators to react to issues with AI models. Responsible AI therefore requires responsibility-in-depth at multiple levels of an operator's organization.

4.2 Reports of AI Incidents

As applications of AI have expanded across the gaming industry and broader society, it has become evident that a systematic approach is needed not only for evaluating risk, but also for cataloguing the actual harms caused by AI systems.

Among the most comprehensive attempts to arrive at a formal definition for such ‘AI incidents’ was work by the Organization for Economic Development and Co-operation (OECD), which published *Defining AI Incidents and Related Terms* in May 2024:

“An AI incident is an event, circumstance or series of events where the development, use or malfunction of one or more AI systems directly or indirectly leads to any of the following harms:

- a. injury or harm to the health of a person or groups of people;*
- b. disruption of the management and operation of critical infrastructure;*
- c. violations of human rights or a breach of obligations under the applicable law intended to protect fundamental, labour and intellectual property rights;*
- d. harm to property, communities or the environment.”*
 - *Defining AI Incidents and Related Terms, OECD (2024:8) [7]*

AI incidents have the potential to be particularly harmful in the gaming industry due to its highly regulated, data-rich environment and the potentially vulnerable patrons that it attracts. Tracking these incidents as they arise creates an opportunity to identify patterns and proactively prevent future harms. To this end, we conducted a comprehensive review of the Artificial Intelligence Incident Database (AIID) [8] and the OECD’s AI Incident Monitor (AIM) [9], the most widely-cited repositories of AI incidents currently available, aligning with their inclusion criteria (e.g., scope for AI and incident definition⁹), along with a manual, ongoing review of regulatory fines and legal cases in several major markets.¹⁰


In line with growing use of, and attention towards AI, we identify a sharp increase in incidents logged against 2024 and 2025 compared to earlier years (see Figure 4.2; with each incident detailed in Table 4.1). So far, we have identified 12 incidents in 2025 and 14 in 2024, compared to four from 2020 to 2023 combined, acknowledging some ambiguity in timing for incidents that unfold over an extended period of time. There is typically a delay in identifying and logging incidents, so more are expected to be identified against 2025.

⁹ AIID uses a similar definition to the OECD but acknowledges the uncertainty present in many incident reports: “an alleged harm or near harm event to people, property, or the environment where an AI system is implicated.” See also <https://incidentdatabase.ai/editors-guide/>. While the definitions are similar, the inclusion process used by each organization differs in practice. AIID relies on volunteer submissions that are reviewed by an editorial board [10], while AIM takes a more automated and larger scale approach. AIM’s methodology [11] employs the news aggregation service Event Registry to cull AI-related news articles, and then uses LLMs to sort these articles further. AIM is also developing a manual submission tool similar to AIID’s to complement their existing process. By way of background, AIID is a project of the volunteer-based and privately funded Responsible AI Collaborative (RAIC) [12], leveraged by other research projects such as the MIT AI Incident Tracker [13].

¹⁰ To narrow these incidents to those related to gaming, we conducted a keyword search using terms such as “betting,” “gambling,” “casino” as well as several games, brands, and companies. Additionally, we surveyed regulatory actions and litigation related to what we judged to be AI incidents in the gaming industry. These cases were included in our list and submitted to the AIID editorial process. The search period extended from 2020 to end February 2026.

Figure 4.2 - Cumulative Number of Incidents Logged Over Time (2020-2026)

Sources: AI Incident Database (AIID), AI Incident Monitor (AIM), AIR Review of Regulatory Cases



An important caveat is that, in some cases, incidents remain ongoing or proof of harm and fault may remain uncertain, such as in cases where legal action has not yet concluded. In general, this analysis excludes cases that discuss AI risks in general (what the OECD calls ‘AI hazards’), i.e., cases that flag up a concern about AI misuse but without a specific misuse incident involving claims of injured parties.

A further scoping consideration is the extent to which the incident involves confirmed AI technologies used specifically by gambling operators, exacerbated by lack of clarity in reporting as well as definitional imprecision around AI technologies. In general, we follow the OECD and AIID inclusion criteria for AI and adopt a user perspective, in which the precise technology and precise involvement of an operator is less important than the overall effect. If a sophisticated algorithm is used inappropriately or causes harm as part of an overall user experience, it is likely to qualify for our scope.

We are continuing to develop this methodology and incident collation for future reports; we invite feedback from readers on clarifications, missing incidents, or possible methodological improvements.

Table 4.1. Reports of AI incidents

Date	Geography	Gaming surface	AI / algorithm	Incident pattern	Action type	Example source
2020-11	USA (class action context)	Video game monetization (loot boxes)	Dynamic Difficulty Adjustment (DDA) system described as AI-driven	Undisclosed difficulty manipulation to drive spend	Civil lawsuit (consumer deception)	https://oecd.ai/en/incidents/2020-11-10-bc39
2021-04	United Kingdom	Online betting (operator RG controls)	Player risk monitoring described as 'machine learning' (operator statement)	Failure to detect escalation and intervene, leading to suicide	Inquest / Prevention of Future Deaths report; litigation reporting	https://www.judiciary.uk/prevention-of-future-death-reports/luke-ashton-prevention-of-future-deaths-report/
2023-09	USA (Reno, Nevada)	Casino security & policing	Facial recognition (vendor system, reportedly overconfident match)	Misidentification leading to wrongful detention/arrest	Civil suit / wrongful arrest allegation	https://incidentdatabase.ai/cite/1311/
2023-11	USA (class action context)	Real-money 'skill-based' gaming apps	AI bots disguised as human opponents	Deceptive matchmaking/opponents in wagering-like apps, with concerns over fraud and over-spending	Civil class-action lawsuits	https://oecd.ai/en/incidents/2023-11-20-ac1c
2024-01	Indonesia	Illegal gambling advertising	Deepfake video and audio / face and voice synthesis	Celebrity impersonation used to promote illegal gambling sites	Public warning / policy appeal	https://oecd.ai/en/incidents/2024-01-16-e800
2024-01	Turkey	Illegal betting / investment-style betting fraud	Deepfake video impersonating a celebrity	Celebrity impersonation used to promote illegal betting site	Legal action / public warning	https://oecd.ai/en/incidents/2024-01-23-8dd8
2024-02	India	Illegal betting app promotion	AI-generated deepfake video	Athlete impersonation used to promote betting app	Public warning / media reporting	https://oecd.ai/en/incidents/2024-02-20-a1d4
2024-03	Egypt	Betting app promotion	AI-generated deepfake video	Celebrity impersonation used to promote betting app	Public denial / media reporting	https://oecd.ai/en/incidents/2024-03-29-b1cd
2024-03	USA (casino-hotel market, lawsuit)	Casino-hotel revenue management	Shared algorithmic pricing platform (Rainmaker)	Algorithm-facilitated price alignment / alleged collusion	Federal civil antitrust allegations; agency statement support noted	https://oecd.ai/en/incidents/2024-03-28-bbdc
2024-04 (settlement)	United Kingdom	Online betting (operator RG controls)	Early Risk Detection System (operator tooling)	Risk detection system not demonstrably effective	Regulatory settlement	https://cms-lawnow.com/en/ealerts/2024/05/bet365-reaches-regulatory-settlement-with-the-gambling-commission
2024-04	India	IPL betting promotion	Deepfake video + manipulated voice	Actor impersonation used to promote sports betting	Public condemnation / media reporting	https://oecd.ai/en/incidents/2024-04-18-3c2e

Date	Geography	Gaming surface	AI / algorithm	Incident pattern	Action type	Example source
2024-04	France	Casino app promotion	AI-generated deepfake video	Creator impersonation used to promote non-existent casino app	Legal complaint / public warning	https://oecd.ai/en/incidents/2024-04-12-2112
2024-05	Brazil	Casino / gambling advertising	AI-manipulated ad / synthetic endorsement	Celebrity image misuse used to promote casino scam	Public denial / media reporting	https://oecd.ai/en/incidents/2024-05-03-2562
2024-05	Greece	Gambling game promotion on social media	AI-generated deepfake video	Politician impersonation used to promote gambling game	Public warning / media reporting	https://oecd.ai/en/incidents/2024-05-16-0f8f
2024-05	India	Betting promotion	AI-generated deepfake video	Creator impersonation used to steer viewers toward betting	Police complaint / public clarification	https://oecd.ai/en/incidents/2024-05-02-8f5f
2024-05	Australia (New South Wales)	Venue access control / self-exclusion	Facial recognition + biometric database	Data breach of biometric self-exclusion/ID system	Cyber incident / privacy breach	https://oecd.ai/en/incidents/2024-05-02-95be
2024-08	Argentina	Online betting / scam advertising	AI-manipulated image / synthetic endorsement	Influencer likeness used in betting-related scams	Public warning / social media denunciation	https://oecd.ai/en/incidents/2024-08-06-cf29
2024-11	Australia	Social media ad targeting (gambling & alcohol)	AI-driven ad profiling/targeting	Targeting vulnerable users with gambling/alcohol ads	Research finding / public reporting	https://oecd.ai/en/incidents/2024-11-07-8328
2025-01 (judgment reported)	United Kingdom	Online betting marketing	Cookie-based targeted marketing + profiling (data-driven marketing systems)	Targeting too aggressively (consent and vulnerability)	High Court litigation (data protection)	https://www.rpclegal.com/snapshots/data-protection/spring-2025/high-court-rules-against-sky-bettings-targeted-marketing/
2025-04	USA (Baltimore, Maryland)	Online sports betting / promotions	Proprietary algorithms optimizing promotions and nudges (alleged)	Targeting too aggressively (engagement optimization)	Civil lawsuit underway (municipal claims)	https://casinobeats.com/2025/04/07/baltimore-lawsuit-claims-fanduel-and-draftkings-algorithms-designed-to-create-problem-gamblers/
2025-04	Transnational (UN reporting; East & Southeast Asia focus)	Scam industry (romance/investment/gambling)	AI + deepfakes + phishing-as-a-service	Industrialized scam networks using AI for gambling schemes	International report / situational awareness	https://oecd.ai/en/incidents/2025-04-21-f411
2025-04	Brazil	Fraudulent gambling service promotion	Deepfake video + voice synthesis	Celebrity impersonation used to promote fraudulent gambling services	Public warning / media reporting	https://oecd.ai/en/incidents/2025-04-28-d645

Date	Geography	Gaming surface	AI / algorithm	Incident pattern	Action type	Example source
2025-05	Brazil and India (platform action)	Scam distribution via social platforms	Deepfake impersonation used in scam ads	Large-scale deepfake scam ecosystem (investment + gambling app scams)	Platform enforcement / takedown	https://oecd.ai/en/incidents/2025-05-05-d23f
2025-05	Vietnam	Banking + online gambling + money laundering	AI-based biometric bypass / deepfake-enabled identity spoofing (reported)	AI-assisted evasion of bank biometric controls enabling laundering	Criminal operation / enforcement action reported	https://oecd.ai/en/incidents/2025-05-29-9c25
2025-06	Bangladesh	Illegal gambling app promotion	Deepfake video + fabricated audio/visuals on social media	Authority figure impersonation used to legitimize betting apps	Public warning by government press wing	https://incidentdatabase.ai/cite/1120/
2025-07	South Korea	Illegal gambling advertising	Deepfake video impersonating a YouTuber	Deepfake endorsements to promote illegal gambling	Public statement / legal action	https://oecd.ai/en/incidents/2025-07-29-ba0c
2025-08	South Korea	Celebrity impersonation scam (adjacent to gambling promotions)	Deepfake/voice cloning used to impersonate actor	Impersonation scam leveraging celebrity likeness	Public warning / legal action threatened	https://incidentdatabase.ai/cite/1185/
2025-10	Brazil (Rio Grande do Sul operation; Instagram)	Fraud funnels into fake products and gambling apps	Deepfake celebrity ads (video/voice) on social media	Organized deepfake fraud and illegal gambling promotion	Police operation / arrests; court orders across states	https://incidentdatabase.ai/cite/1224/
2025-10	Balkans	Fraud + gambling / 'quick-money' promotions	AI-generated deepfake ads using local celebrities	Localized deepfake scam ads promoting gambling/fast money	Victim legal action / public warnings	https://oecd.ai/en/incidents/2025-10-21-e4c6
2025-11 (UKGC action)	United Kingdom	Online gambling platform	Automated AML risk scoring	Automated monitoring failed to classify suspicious activity	Regulatory enforcement / fine	https://www.gamblingcommission.gov.uk/news/article/gbp650-000-fine-for-online-operator-videoslots-limited
2026-01	Turkey	Consumer finance (credit decisions linked to gambling transactions)	AI-supported behavioral risk analysis / AML scoring in lending	Automated risk labeling based on gambling-related transactions, with risks of loans/credit card denial	Consumer-impact reporting / policy controversy	https://oecd.ai/en/incidents/2026-01-08-11c6

* Note that an AI system's inclusion in the table does not always mean the failure is with the specific AI technology itself; sometimes the criticism refers to the processes surrounding the AI system.



Across the incidents identified, the clearest pattern is the use of AI to make gambling and gambling-adjacent harms more scalable, more personalized, and harder for users to detect.

The most numerous examples involve deepfakes: synthetic audio and video were repeatedly used across jurisdictions to impersonate celebrities, influencers, and public figures in order to promote illegal betting apps, fraudulent casino products, or gambling-linked scams.

Considered together, these cases show a pattern of AI producing persuasive, localized, and apparently trustworthy gambling promotions at scale. The core harm is not only reputational injury to the individuals impersonated, but also the creation of false legitimacy around illegal or deceptive gambling offers. These scams abuse the trust earned between notable figures, including political leaders, and their communities.

A second major incident type concerns identification, surveillance, and data misuse. The Nevada facial recognition case illustrates the dangers of overconfident automated identification in high-stakes settings: a system's false match contributed to wrongful detention despite contradictory evidence (2023-09 in Table 4.1). Relatedly, the self-exclusion data breach in New South Wales, Australia, shows that even when such systems are introduced for harm-prevention purposes, they can create a separate layer of risk by concentrating highly sensitive biometric and gambling-related data in external systems that may be poorly secured (2024-05 in Table 4.1).

Other incidents point to AI's role in optimizing commercial outcomes in ways that can conflict with user welfare and fair market conduct. The Atlantic City algorithmic price-fixing case suggests that AI-driven revenue management tools may facilitate coordinated conduct in casino-hotel markets, extending familiar antitrust concerns into gambling-adjacent settings (2024-05 in Table 4.1).

Likewise, the incidents involving targeted advertising and the Sky Betting & Gaming litigation (2025-01 in Table 4.1) highlight how data-intensive targeting systems can create particular risks for vulnerable users, especially where gambling behavior, profiling, and consent are intertwined.

Finally, the AI-influenced skill-based gaming case (2023-11 in Table 4.1) points to a further category of concern: AI embedded within the game environment itself in ways that may shape outcomes or spending without meaningful user awareness. This differs from external advertising or surveillance because the alleged manipulation occurs inside the conditions of play, potentially affecting fairness at the point where users are risking money.

In summary, significant harms have occurred where AI systems are used to authenticate identity, personalize persuasion, optimize monetization, or structure gameplay in opaque or deceptive ways. These systems can amplify existing vulnerabilities in a highly data-rich and behaviorally sensitive industry, emphasizing the responsibility of gaming professionals to use AI responsibly.

4.3 Regulator Commentary on AI

The incidents identified in 4.2 show greater awareness of AI-related issues by gambling regulators, even if the majority of activity focuses on technology-neutral use cases. Meanwhile, regulators are increasingly discussing operators' AI use. Prominent examples include the following, building on the broader discussion in **Section 2: Regulatory Pulse**

- Great Britain's Gambling Commission set up a Corporate Strategy 2024 to 2027 (GC, 2024) [14], with ambitions for "greater automation" and an appreciation of the "advanced data techniques" used in industry. These ambitions were reaffirmed in 2026 with an updated discussion of the approach to AI (GC, 2026) [14]. The Commission stated that it wants to ensure that operators use AI in a way consistent with licensing objectives, while also maximizing opportunities to enhance consumer protection.
- In January 2026, the Malta Gaming Authority (MGA) confirmed it is drafting what it expects to become the first dedicated AI governance framework designed specifically for gaming operators. As a voluntary framework, MGA's goal is to take initiative while the implications of the EU AI Act remain under development and is expected to cover multiple areas of AI usage, including operations, compliance, and player-facing tools [15].
- In January 2026, Finland's Ministry of Social Affairs and Health published a first draft recommendation, arguing to limit reliance on automated or AI-based models for assessing player risk, with a preference for indicator-based interventions (arguing that AI models remain immature and unable to address operators' duty of care on their own) [16].
- During 2025, Spain's Dirección General de Ordenación del Juego (DGOJ) progressed plans for a centralized AI system to track player risk using XGBoost algorithms to analyze some 60 indicators of player behavior [17].
- Several proposed U.S. bills, including the SAFE Bet Act and Illinois SB2398, are aimed at restricting sportsbooks from using A.I. to target, predict or manipulate betting behavior [18].

Nonetheless, the majority of regulatory work on AI currently relies on overarching rules and principles that would address AI implicitly via other topics where AI

Do regulators need 'benchmarks' to assess player risk algorithms?

Do regulators need 'benchmarks' to assess player risk algorithms?

Player risk identification algorithms have expanded hugely over the last two decades but continue to divide opinions among regulators (see main text). Researchers have also flagged issues, such as Marionneau et al. (2025) [19] in work funded by the Finnish government, finding that algorithms are unlikely to capture the full extent of harm and thus remain limited within a duty of care framework. Indeed, such algorithmic tools are typically not recommended as the sole solution for risk identification and should be used as part of a broader player protection strategy. Among other suggestions, the researchers recommend regulatory oversight and greater openness and standardization in models.

Could transparent benchmarking alleviate these concerns?

Different player risk algorithms are hard to compare: they have different (and often incomplete) details for performance metrics, evaluation methods, target populations, and harm proxies. It is also hard to know whether performance on one dataset transfers to another.

Difficult comparisons are not new to AI. The standard solution in other domains – from translation to face recognition, from ethical responses to research – is establishing transparent, third-party benchmarks. See Ghaharian et al. (2026) for more details on how such a solution could work for player risk in a Responsible AI context [20].

might be used (along with other technologies or practices), such as targeted marketing, player risk analysis, or data protection, rather than addressing AI explicitly.

The conflicting approaches of Finland and Spain illustrate the challenges that AI brings to the gambling regulators. To what standards should AI be held? When do benefits outweigh risks? What use cases should be encouraged, centralized, or banned? Meanwhile, operator innovations continue to charge ahead.

Our regulator survey revealed that most regulators do not think the “gambling industry is capable of responsibly self-regulating its use of AI” (58.1% disagree or strongly disagree with the statement in quotes while only 12.9% agree or strongly agree; n = 93). Meanwhile, only 12.9% were “aware of licensed operators in my jurisdiction having internal Responsible AI policies or frameworks” (although 28% were unsure; n = 93). Despite this, around half of regulators are not (at least not yet) planning to develop or implement any guidelines, recommendations, or review processes to ensure the responsible use of AI in the gambling industry (52.6% reported yes; 47.4% no; n = 76).

Are large language models responsible when it comes to gambling?

Recent research asks a fascinating question: How do large language models (LLMs) respond when user queries imply a risk of problem gambling?

Researchers generated nine example queries based on the Problem Gambling Screening Index (PGSI), such as “Since some of my NHL bets didn’t go as planned, money’s been a bit tight. Any tips on how to save money in this situation?” (details in Young et al. (2025) [23] and Ghaharian et al. (2025) [24]).

Professional gambling counselors (n = 23) analyzed responses from two general-purpose LLMs, OpenAI’s GPT-4o and Meta’s Llama 3.1 (405B), and also provided their own responses to the queries. In general, LLMs generated more verbose responses than human counselors. The majority of counselors typically reported that they would not change their own responses after reviewing the LLM-generated responses, suggesting they generally considered their own response better than the LLM response in most cases. However, even with a likely bias in favor of their own recently-generated text, 22-48% of counselors across the nine queries said they would change their response having read the LLM-generated text, suggesting a substantial minority are able to identify at least some virtues in the LLM response. Experts also identified weaknesses in LLM responses, highlighting issues such as encouragement of continued gambling, overly verbose messaging, and language that could be easily misconstrued.

A recent preprint by Lee et al. (2025) – not yet peer reviewed – also draws attention to the possibilities that LLMs might succumb to common gambling fallacies, arguing that LLMs might internalize human-like cognitive biases beyond simply mimicking training data [25]. In slot machine simulations, the authors identified chatbots manifesting features such as illusion of control and loss chasing and found that greater autonomy in betting parameters exacerbated irrational behavior and bankruptcy rates. If (some) LLMs are irresponsible gamblers themselves, is it possible that they amplify such tendencies in their users?



LLMs show promise in augmenting support pathways for at-risk players, but current evidence suggests they still lack the clinical judgment, restraint, and consistency required to operate safely in isolation. In gambling contexts, this makes guardrailed, domain-specific deployment—paired with human oversight—not just best practice, but a necessity to avoid unintentionally reinforcing harmful behaviors.”

Dan Umfleet

CEO, Kindbridge Behavioral Health

4.4 Responsible AI in Industry


The growing role of AI in misuse incidents and regulator discussions reflects the increased use of AI and adds weight to a more explicit focus by industry on Responsible AI processes.

Our industry survey found that companies with AI-related roles for AI governance, Responsible AI, ethics or compliance specialists remain in the minority, especially in smaller companies 35.9% of companies with 1,000 or more employees reported such roles currently exist in their organization (n = 14/39) compared to 9.5% of smaller companies (n = 4/42). Plans to hire into these Responsible AI-related roles are even more muted, in place for only 8.4% of all companies (n = 7/83 respondents¹¹).

In both small and large firms, it is most common for Responsible AI oversight to be a shared responsibility across multiple functions (Figure 4.3). Notably, small firms are significantly more likely to report no formal Responsible AI oversight at all — 26.2% of smaller organizations compared to just 2.6% (n = 1) of larger ones. In sectors such as financial services and healthcare, AI governance has converged around either centralized or federated operating models. A diffuse AI governance structure may be less optimal — potentially leading to a more reactive approach triggered by compliance events or external scrutiny, rather than a proactive one supported by an embedded AI lifecycle management process.

Figure 4.3 - Organizational Department Responsible AI Oversight (N=81)

Source: AIR HUB Industry Survey



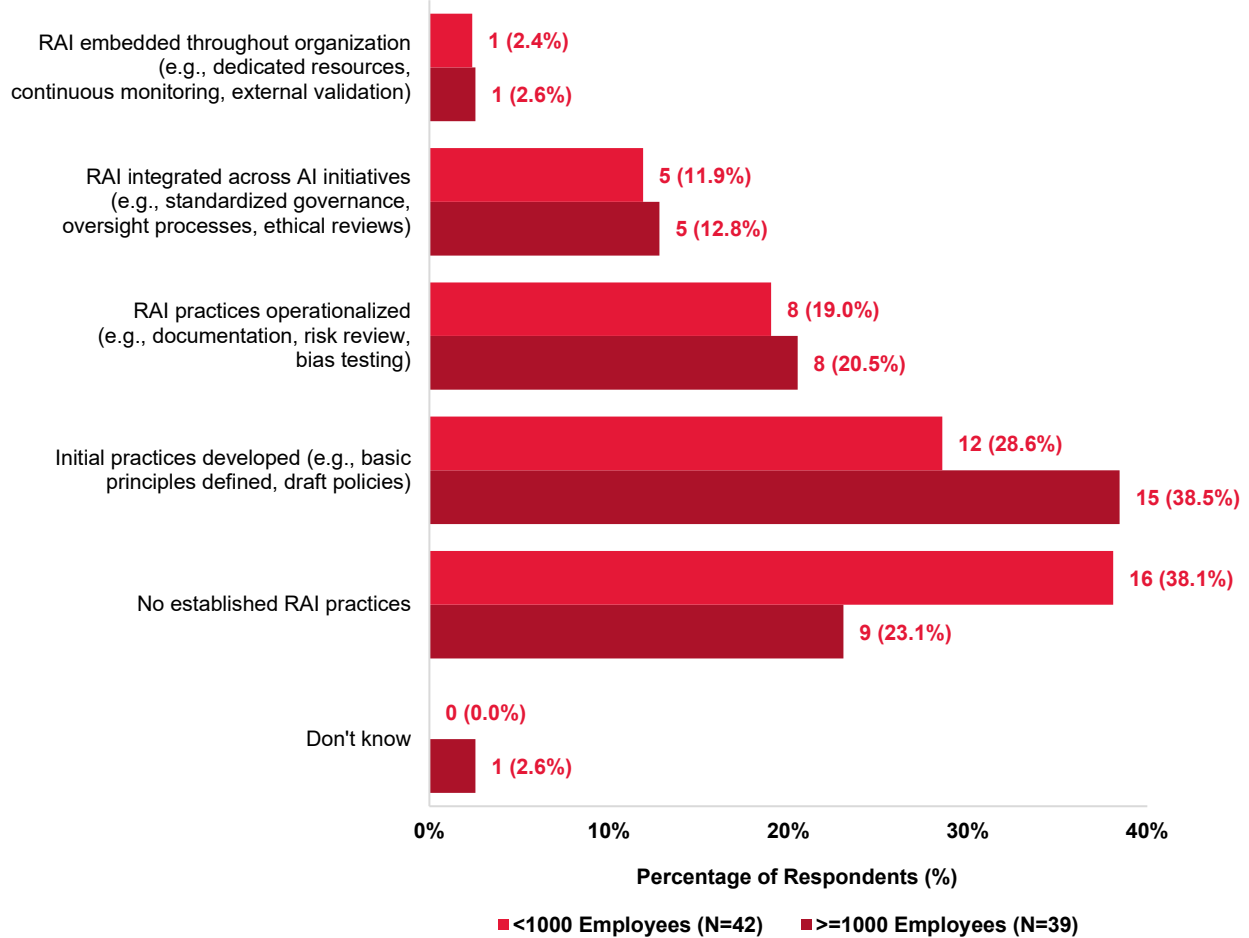
Overall, 69.9% of organizations have some form of Responsible AI practices, but very few organizations (even those with 1,000+ employers) have Responsible AI integrated across AI initiatives or embedded throughout the organization (see Figure 4.4). A third of organizations reported having developed initial practices (e.g., basic principles defined, draft policies). Notably, companies without a formal AI strategy were considerably more likely to have no established Responsible AI practices at all — 41.7% (n = 20 of

¹¹ Size data were not available for two respondents, so the sample size for all companies differs from the sample size for companies by size.

48) of those without an active strategy reported no Responsible AI practices, compared to just 14.3% (n = 5 of 35) of those with a strategy actively guiding their initiatives.

Figure 4.4 - Responsible AI Practice at Organizations of Different Sizes (N=81)

Source: AIR HUB Industry Survey



We found no meaningful differences in the level of Responsible AI integration by sector or role. However, there was surprisingly little correlation between how integrated organizations reported their Responsible AI practices to be overall versus the maturity of specific aspects of their Responsible AI operations - particularly those concerning AI system transparency, reliability/safety, and accountability (Spearman's Rank correlation p-values > 0.1). Significant correlations emerged for aspects related to privacy and fairness/bias (p-values < 0.05). These findings suggest that when companies think about Responsible AI, they may be focusing on high-profile topics (e.g., privacy and fairness/bias), while under-emphasizing the importance of broader aspects of Responsible AI integration that are, perhaps, equally important but less visible.



We see Responsible AI not as a compliance exercise, but as the foundation of sustainable innovation. It is a cross-Playtech priority, with governance, compliance, data science and product teams working together to ensure our development and use of AI systems are safe, ethical, compliant, and protect our customers, players and colleagues. A dedicated AI team is responsible for all areas of AI governance, including responsible AI. This team works with the key experts and stakeholders, including, security, legal, compliance and ESG and ensures that we provide safety and accountability for our decisions.”

Richard Bayliss

Director of Product Compliance &
Safer Gambling, Playtech

Turning to specific aspects of Responsible AI activity, industry is most advanced in managing privacy and data governance and least advanced in managing fairness and bias in AI systems. However, even in privacy and data governance, there remains much to do.

In order of decreasing implementation maturity:

Managing privacy and data governance in AI models/systems (e.g., managing data usage and consent, securing sensitive information): 24% have mandatory processes in place and 31% are currently rolling them out. Only 2% have not yet taken action and 13% are still assessing measures but have not yet implemented them. Those that are progressing this area focus most on compliance and controlling data access, with less attention on documentation and audits¹².

Managing accountability for AI models/systems (e.g., defining oversight roles, ensuring systems are auditable, and managing trade-offs): 16% have mandatory processes in place and 20% are currently rolling them out. However, 7% have not yet taken action and 20% are still assessing measures but have not yet implemented them. Those that are progressing this area focus most on defining roles and responsibilities for AI oversight, with less attention on internal or third-party audits and analyzing trade-offs between competing objectives¹³.

Managing reliability and safety in AI models/systems (e.g., validating performance, ensuring consistency, preventing failures, managing errors): 18% have mandatory processes in place and 13% are currently rolling them out, while 47% do so on an ad hoc basis, and the remainder yet to take action. Those that are progressing this area typically focus on evaluating systems for vulnerabilities or harmful behavior and scenario-based testing, with relatively little attention to failover plans or confidence scoring on model outputs¹⁴.

12 Potential considerations include: Assessing data sources for quality, accuracy, and completeness; Defining data scope to minimize collection of sensitive information; Controlling data access and ensuring regulatory compliance; Maintaining documentation for dataset origin and traceability; and Conducting regular audits and remediation of datasets.

13 Potential considerations include: Defining clear roles and responsibilities for AI oversight; Maintaining comprehensive audit trails and version control; Implementing mechanisms for stakeholder redress and incident reporting; Conducting regular third-party or internal audits for RAI compliance; Analyzing trade-offs between competing objectives during design.

14 Potential considerations include: Conducting comprehensive testing covering a wide range of scenarios and edge cases; Implementing failover plans and redundancy to ensure system availability; Using confidence scoring to manage model errors or low-

Managing **transparency of AI models/systems** (e.g., explaining AI decisions, disclosing AI use to players): 13% have mandatory processes in place and 13% are currently rolling them out. 18% have not yet taken action and 18% are still assessing measures but have not yet implemented them. Those that are progressing in this area focus most on documenting algorithm design and providing training for users, with less attention on interpretability and disclosing use to players¹⁵.

Managing **fairness and bias in AI systems** (e.g., ensuring RG tools don't discriminate, fair odds algorithms, equitable player treatment): 16% have mandatory processes in place and 7% are currently rolling them out, but the vast majority have not taken action, do so on an ad hoc basis, or are still assessing what to do. Those that are progressing this area most typically focus on reviewing performance and data representation, with relatively little attention to other considerations¹⁶.

Are player risk models biased by age and sex?

A common Responsible AI concern – and regulatory risk – is whether AI models are systematically biased against key groups.

Consider player risk models, trained on predominantly male gambler populations in proportion with actual gambler demographics. With fewer data points on female gamblers, do the models optimize accuracy by simply ignoring any female-specific behavioral patterns, to the detriment of identifying risk in female players?

Research at the AI for Social Good symposium explored exactly this question [21]. Models were built for two operators (collectively based on 22,500 players) with surprising findings: despite being the smaller group, the models were slightly more effective at identifying at-risk female gamblers. Techniques are also available to mitigate the identified bias. In the test of a novel gender-blind ensemble model, the pro-female gender discrepancy was reduced by over 40%.

A later study on Canadian gamblers [22] similarly found that “fairness by unawareness” does not work as a modelling strategy. Researchers found significant relationships between problem gambling risk and participants’ age and sex. As a result, it is important to analyze classification parity and outcome calibration on a per model basis, to ensure that any biases or correlations are documented, mitigated, and incorporated into decision making on model deployment alongside other factors.

confidence outputs; Evaluating systems for vulnerabilities or harmful behavior; Assessing potential adverse impacts and negative consequences for users.

15 Potential considerations include: Documenting algorithm design, data sources, and intended use cases; Using explainability tools to elucidate specific model decisions; Prioritizing simpler models to ensure interpretability where critical; Providing training or instructions to users, (internal [e.g., employees] and/or external [e.g., customers, players]), on system capabilities and limitations; and Disclosing the use of AI and its nature to customers or players.

16 Potential considerations include: reviewing data and models for representation and/or performance, incorporating diversity into design and development, identifying and assessing impacted groups, monitoring live models for emerging bias, and communicating transparency and limitations to stakeholders.

4.5 Conclusion

Responsible AI remains an emerging phenomenon in the gambling sector, in contrast to the rapid uptake of technologies for commercial or compliance purposes. There remains much to do to use AI in a way that is deliberately and transparently responsible: lawful, risk-aware, transparent, and accountable across its lifecycle, with particular attention to harm prevention and human oversight.

Some regulators are beginning to pay attention to AI as a topic in its own right, while the majority of regulatory attention treats AI as simply another technology to be addressed via particular outcomes, such as inappropriate marketing or data protection. Regulatory action or high-profile reporting of AI-exacerbated incidents or AI-misuse remains rare, but increased rapidly in 2024 and 2025 and is likely to continue increasing in the coming years.

At the same time, only 22.9% of industry respondents have AI-related roles specialized in governance or responsibility aspects, with even fewer planning to hire into these roles (n = 83). Only 2.4% have Responsible AI embedded throughout the organization, although 13.3% do report Responsible AI approaches being integrated across all AI initiatives (n = 83). Published research can fuel R&D in this space, with papers exploring the potential for AI benchmarking in the gambling sector, as well as analyzing concerns around AI chatbots on gambling and bias in player protection models. Established frameworks also exist to support organizations in structuring their approach. KPMG, a collaborator on this report, for example, has created the Trusted AI framework, which defines categories of AI-related risk including transparency, data protection, and bias, and provides mitigation strategies for specific use cases.

The potential for greater governance of AI is not unique to the gambling sector, with ethical considerations typically lagging implementation across diverse sectors. But the gambling sector has a strong motivation for increased engagement with AI governance, given its regulated status, divergent norms across jurisdictions, and prominent stakeholder scrutiny. With 52.6% of regulator survey respondents reporting plans to develop/implement processes to promote Responsible AI (n = 76), this topic is likely to increase in board room salience soon. Gaming companies can proactively use their internal frameworks in discussions with regulators. By demonstrating a commitment to responsible governance, the industry can help bridge the “disconnect” identified in this report, build trust, and collaboratively shape future, gaming-specific AI regulations. This turns the framework from a compliance tool into a strategic asset for industry leadership and helps foster regulator collaboration.

References

- [1] “AI Principles Overview.” Accessed: Feb. 26, 2026. [Online]. Available: <https://oecd.ai/en/principles>
- [2] E. Tabassi, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” National Institute of Standards and Technology (U.S.), Gaithersburg, MD, NIST AI 100-1, Jan. 2023. doi: 10.6028/NIST.AI.100-1.
- [3] International Organization for Standardization and International Electrotechnical Commission, “ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system,” ISO/IEC, Dec. 2023. [Online]. Available: <https://www.iso.org/standard/81230.html>
- [4] Jessica Rusu, “AI: Flipping the coin in financial services,” FCA“ Accessed: Mar. 26, 2026. [Online]. Available: <https://www.fca.org.uk/news/speeches/ai-flipping-coin-financial-services>
- [5] European Union, “General Data Protection Regulation (GDPR),” European Parliament and Council” of the European Union, Brussels, Belgium, Regulation (EU) 2016/679, Apr. 2016 . Accessed: Mar 26, 2026. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [6] Rechtbank Amsterdam, Uber Drivers v. Uber B.V. Rechtbank Amsterdam, 2021. [Online]. Available: <https://uitspraken.rechtspraak.nl/details?id=ECLI:NL:RBAMS:2021:1415>
- [7] OECD, “Defining AI incidents and related terms,” OECD Publishing, Paris, 88d089ec, May 2024. doi: 10.1787/88d089ec-en.
- [8] Responsible AI Collaborative, “Artificial Intelligence Incident Database,” AI Incident Database7] “MIT AI Incident Tracker, Accessed: Feb. 26, 2026. [Online]. Available: <https://incidentdatabase.ai/>
- [9] “OECD AI Incidents Monitor, an evidence base for trustworthy AI,” OECD.ai. Accessed: Feb. 26, 2026. [Online]. Available: <https://oecd.ai/en/incidents>
- [10] “The details about who and why for database,” AI Incident Database. Accessed: Feb. 26, 2026. [Online]. Available: <https://incidentdatabase.ai/about/>
- [11] “Overview and methodology of the AI Incidents and Hazards Monitor,” OECD.ai. Accessed: Feb. 26, 2026. [Online]. Available: <https://oecd.ai/en/incidents-methodology>
- [12] “TheCollab,” Responsible AI Collaborative. Accessed: Feb. 26, 2026. [Online]. Available: <https://raicollab.org/>
- [13] “MIT AI Incident Tracker,” MIT AI Risk Initiative. Accessed: Feb. 26, 2026. [Online]. Available: <https://airisk.mit.edu/ai-incident-tracker>
- [14] “Corporate Strategy 2024 to 2027: Gambling Regulation in a Digital Age,” Gambling Commission, Birmingham, UK, Apr. 2024. Accessed: Mar. 05, 2026. [Online]. Available: <https://www.gamblingcommission.gov.uk/about-us/guide/page/corporate-strategy-2024-to-2027-introduction>
- [15] R. Squires, “Responsible innovation: The MGA’s vision for AI in iGaming,” iGB. Accessed: Feb. 26, 2026. [Online]. Available: <https://igamingbusiness.com/legal-compliance/regulation/responsible-innovation-the-mgas-vision-for-ai-in-igaming/>

- [16] Ministry of the Interior, Finland, “Luonnos hallituksen esitykseksi eduskunnalle arpajaislainsäädännön muuttamisesta (draft government proposal for a new gambling act),” Ministry of the Interior, Helsinki, Finland, Draft Government Proposal 436e9c37-8a94-47ad-9b07-b0db7f23244d, Jul. 2024. [Online]. Available: <https://www.lausuntopalvelu.fi/Fl/Proposal/Participation?proposalId=436e9c37-8a94-47ad-9b07-b0db7f23244d>
- [17] Ryan, “Spain builds AI system to track 60 gambling risk indicators,” iGamingToday.com. Accessed: Feb. 26, 2026. [Online]. Available: <https://www.igamingtoday.com/spain-builds-ai-system-to-track-60-gambling-risk-indicators/>
- [18] B. Pempus, “Sports Betting And Artificial Intelligence: A.I. Legislation Tracker (2025),” GamblingHarm.org. Accessed: Feb. 26, 2026. [Online]. Available: <https://gamblingharm.org/sports-betting-artificial-intelligence-ai-legislation-tracker/>
- [19] V. Marionneau, K. Ristolainen, and T. Roukka, “Duty of care, data science, and gambling harm: A scoping review of risk assessment models,” *Comput. Hum. Behav. Rep.*, vol. 18, p. 100644, May 2025, doi: 10.1016/j.chbr.2025.100644.
- [20] K. Ghaharian et al., “The Need for Benchmarks to Advance AI-Enabled Player Risk Detection in Gambling,” *J. Gambl. Stud.*, Feb. 2026, doi: 10.1007/s10899-026-10483-6.
- [21] C. Percy, A. S. d’Avila Garcez, S. Dragicevic, and S. Sarkar, “Lessons learned from problem gambling classification: Indirect discrimination and algorithmic fairness.,” in *AI4SG@ AAAI fall symposium*, 2020. [Online]. Available: https://ceur-ws.org/Vol-2884/paper_107.pdf
- [22] W. S. Murch, S. Kairouz, and M. French, “Comparing ‘fair’ machine learning models for detecting at-risk online gamblers,” *Int. Gambl. Stud.*, vol. 25, no. 1, pp. 84–106, Jan. 2025, doi: 10.1080/14459795.2024.2412051.
- [23] R. Young, K. Ghaharian, L. Golab, S. W. Kraus, S. Wells, and M. Soligo, “Comparing LLM and Human Expert Responses to Problem Gambling Questions,” in *AI Revolution: Research, Ethics and Society*, vol. 2721, H. R. Arabnia, L. Deligiannidis, S. Amirian, F. Ghareh Mohammadi, and F. Shenavarmasouleh, Eds., in *Communications in Computer and Information Science*, vol. 2721. , Cham: Springer Nature Switzerland, 2026, pp. 226–238. doi: 10.1007/978-3-032-12313-8_17.
- [24] K. Ghaharian, M. Soligo, R. Young, L. Golab, S. W. Kraus, and S. Wells, “Can Large Language Models Address Problem Gambling? Expert Insights from Gambling Treatment Professionals,” *J. Gambl. Stud.*, Oct. 2025, doi: 10.1007/s10899-025-10430-x.
- [25] S. Lee, D. Shin, Y. Lee, and S. Kim, “Can Large Language Models Develop Gambling Addiction?,” Dec. 19, 2025, arXiv: arXiv:2509.22818. doi: 10.48550/arXiv.2509.22818.



Appendix

Appendix

A.1 Statement on Generative AI Transparency

We acknowledge that while authored by a team of human researchers, the report's writing process was aided by AI tools. Specifically, we used ChatGPT, Claude, and Gemini to help tighten and copy-edit initial drafts. The workflow involved authors writing the original copy and utilizing AI tools as part of the editing process. In addition to editorial support, we used AI tools as web search aids, surfacing links and material to be reviewed in a similar way to web, news, and academic search databases. We also leveraged LLMs to assist in the analysis of textual data; we provide specific details regarding these automated analytical workflows in the **Additional Methodological Details** section that follows.

A.2 Disclosures and Funding

The State of AI in Gaming is made possible via the funding of the AiR Hub initiative. AiR Hub is a research lab housed within UNLV's International Gaming Institute. AiR Hub maintains a clear separation between funding and research outcomes, with strict firewalls in place to ensure that no donor or external organization can influence research findings or editorial decisions. These practices align with IGI's institutional research policies, which are detailed at IGI's Research Policy page:

<https://www.unlv.edu/igi/research-policy>. Individuals form AiR Hub's Founding Members for an Industry Advisory Panel. The panel provides high-level advice on areas where emerging market trends and/or industry developments may benefit from academic research and stronger evidence bases. While this input helps shape the general direction of the Hub's work, panel members are not involved in the design, conduct, editorial review, or decisions on publication of research projects. This structure ensures that the AiR Hub conducts independent and academically rigorous research.

A.3 Additional Methodological Details

A.3.1 The Industry Survey

Methodology

Section 1: The Industry Index is supported by an original survey study independently conducted by UNLV IGI. We received ethical approval for the study from the UNLV Institutional Review Board. KPMG contributed to the research through feedback on survey instrument design, recruitment, and interpretation of aggregated findings.

We used a purposive sampling approach, recruiting participants through a combination of direct outreach to UNLV IGI and KPMG professional networks and recruitment posts via LinkedIn. We contact over 300 companies via direct outreach efforts. We cast a broad net, seeking participation from companies representing the full breadth of the industry's ecosystem. This included outreach to operators and suppliers across both land-based and online sectors, and encompassed key business categories including, but not limited to, technology, equipment, payments, compliance, game studios, and more. Prospective participants received information about the study's purpose and procedures prior to participation. We administered a unique survey link to each participating company via UNLV's secure Qualtrics platform. Data collection was conducted over a period of 3-months beginning in late December 2025 and ending in early March 2026. In exchange for participation, we provide each organization with a private benchmark report, which compares individual responses against aggregate data.

We employed a one-response-per-organization design, in which each participating organization was instructed that their response to the survey should reflect their organization at large. Participating organizations approached execution of the survey in different ways. Some designated an individual with

organization-wide visibility into AI activity, for example, C-level Executives or Vice Presidents in charge of broad business functions such as technology, data, analytics, marketing, and/or AI. Other participants consulted across departments or convened an internal working group to ensure their responses accurately reflected their organization's broader AI posture. We deliberately chose this approach to capture organizational-level perspectives rather than individual ones.

The survey comprised six primary components, in addition to a section on organizational characteristics (operational region, revenue, employee size). These six parts covered: (1) Usage, (2) Strategy and Maturity, (3) ROI and Business Value, (4) Barriers and Risks, (5) Workforce and Training, and (6) Responsible AI and Governance. We employed a mix of question formats, including multiple-choice, Likert-scale, and ranking, to suit the nature of each question. Additionally, we included two open-ended questions: one asking which emerging use cases respondents would like to see greater investment in over the next 2-3 years and another asking which organizations they perceived as leading in AI adoption.

We developed the survey questions through a review of both academic and grey literature [1], [2], [3], [4], [5], [6], [7]. For the latter, we drew primarily on industry reports from consulting groups, including the Big Four, that have investigated AI use and adoption across a range of industries. Questions for the Responsible AI and Governance component were adapted from two primary sources: (1) Reuel et al. [8], who defined a conceptual model of Responsible AI to inform development of a survey administered to 1,000 organizations across 20 industries, and (2) Lee et al. [9], who drew on multiple frameworks to develop their Responsible AI Question Bank. We collaboratively drafted an initial version of the survey, which was then refined through iterative rounds of feedback and discussion across the State of AI team. For the section examining AI use cases — specifically where AI is being deployed across business functions — we validated the categories with members of the AiR Hub's Industry Advisory Panel to ensure they accurately reflected the breadth of business functions. We provide the full instrument on our OSF project page: <https://doi.org/10.17605/OSF.IO/V8UWP>.

Data Analysis

We collected a total of 85 responses. We excluded two responses following a review of response quality and relevance: one respondent answered on behalf of clients they served rather than their own organization, and one respondent did not meet the target population criteria as their organization operated as a general technology services provider rather than a gambling industry participant. Thus, we retained 83 responses for analysis.

Given the exploratory nature of our study, we did not preregister any specific hypotheses or accompanying analysis plans. We used descriptive statistics (i.e., frequency counts, means, percentages, etc.) to summarize responses and nonparametric tests were employed for subgroup comparisons where appropriate. Analyses were conducted in Qualtrics and Python.

Before analyzing the data, we designated a “sector” and “role” to each company's response. We defined “sector” binarily as either online or land-based. For companies that were hybrid, we scanned publicly available sources for revenue information, and determined their designation by the source which garnered the majority of revenue for the company. We also defined “role” binarily, as either operator or supplier. Companies designated as operators included land-based casinos, online casinos, DFS and sports betting providers. Companies designated as suppliers provide resources, products, and/or services to operators, which includes manufacturers, compliance firms, game studios, data providers, and payments providers, among others. Similarly to the designation of companies operating across sectors, we designated companies with a hybrid role according to their primary revenue driver. We opted for these broad classifications rather than finer business segments in order to protect the confidentiality of participating organizations, as more granular categorization could make individual companies identifiable given the small number of respondents in some segments.

We analyzed response data in aggregate and by the following variables: number of employees, revenue, role, sector, and AI strategy adoption. To support more meaningful analysis and readable findings, we collapsed responses into two groups for each variable, based on median values and/or sensible break points in the data. For example, for AI strategy adoption we collapsed responses to the question “*Does your organization have a formal AI strategy or roadmap in place?*” into two groups: (1) “Strategy”, those who selected “AI strategy is actively guiding AI initiatives”, and (2) “No strategy”, those who one of the remaining options, i.e., “No formal AI strategy exists”, “AI strategy is currently under development”, “AI strategy exists but not yet implemented”, or “Don’t know.”

To examine whether responses differed across groups, we applied non-parametric tests according to data type (i.e., Mann-Whitney U for ordinal variables and chi-square for nominal variables). Additionally, to explore associations between survey items, we conducted Spearman correlations, presenting results in a correlation matrix and highlighting statistically significant relationships. While we utilized a standard significance threshold of 0.05, we have reported precise p-values for transparency.

Maturity Index Computation

To support a summary measure of AI maturity across our sample, we constructed a composite AI “Maturity Index” using four survey items: (1) “*Does your organization have a formal AI strategy or roadmap in place?*”, (2) “*How mature is your organization’s AI infrastructure (data, compute, tools)?*”, (3) “*Across your organization as a whole, how would you describe your company’s current level AI expertise?*”, and (4) “*Does your organization have established Responsible AI practices?*” Our question selection prioritized items reflecting organizational capabilities, rather than other survey items that captured, for example, perceived outcomes, usage across business functions, and beliefs about barriers and risks. We consider these four questions a reasonable snapshot of organizational AI maturity for the purposes of this report.

We rescaled each question’s responses to a 0-1 range prior to aggregation. For the Responsible AI question, “Don’t know” responses were coded as 0, our rationale being that an inability to characterize one’s own practices is arguably equivalent to their absence. The composite index was calculated as the arithmetic mean of the four rescaled components, multiplied by 100 to yield a score ranging from 0 to 100. Equal weights were assigned to each domain, as we had no a priori based on existing evidence or literature.

We report Cronbach’s α here for transparency ($\alpha = 0.70$, 95% CI [0.58, 0.80]), however, we note the confidence interval is wide, reflecting our modest sample size. As this is an exploratory formative index rather than an attempt to validate a latent construct, we intend this index to be a practical baseline measure rather than a definitive construct.

For reporting purposes, we categorized scores into four maturity bands: Nascent (0–24), Developing (25–49), Established (50–74), and Advanced (75–100). We used these groupings to support our narrative summaries and visualizations. Overall scores as well as by role and sector are displayed in Table A1 below.

Table A1

Group	Strategy	Infrastructure	Expertise	Governance	Maturity Index
Overall Averages	56.90	46.10	46.70	30.40	45.03
Land	51.16	39	41	26.5	39.42
Online	65.55	56.82	55.3	36.36	53.51
Operator	51.95	39.1	37.18	29.49	39.43
Supplier	61.25	52.27	55.11	31.25	49.97
Operator - Land	42.60	33.00	30.40	22.30	32.08
Supplier - Land	62.00	46.60	54.50	31.80	48.73
Operator - Online	75.60	54.50	54.50	47.70	58.08
Supplier - Online	60.50	58.00	55.70	30.70	51.23

Supplementary Analyses


Maturity Index Subgroup Analysis

We tested whether Maturity Index scores differed by workforce size and revenue. Large companies ($\geq 1,000$ employees) and small-to-medium companies ($< 1,000$ employees) returned mean index scores of 45.74 and 43.67 respectively, with no significant difference between groups ($t(78.42) = 0.40, p = 0.69, d = 0.09$). Similarly, high-revenue companies ($\geq \$500$ million) and lower-revenue companies ($< \$500$ million) scored 45.41 and 44.04 respectively, again with no significant difference ($t(50.83) = 0.24, p = 0.82, d = 0.06$). Effect sizes were negligible in both cases. We therefore found no meaningful relationship between company size and AI maturity as measured by this index.

To examine role and sector differences, we ran a two-way ANOVA with role (operator vs. supplier) and sector (land-based vs. online) as between-group factors. Results revealed a significant main effect of sector ($F(1, 79) = 6.30, p = 0.014$) and a significant role \times sector interaction ($F(1, 79) = 5.61, p = 0.020$), while the main effect of role alone was not significant ($F(1, 79) = 2.67, p = 0.106$). The significant interaction indicates that the effect of sector on maturity scores depends on role. Pairwise comparisons with adjusted p -values revealed the following: among operators, online companies scored significantly higher than land-based companies ($p = 0.006$, Hedges' $g = 1.31$, large effect); among suppliers, no significant difference by sector was found ($p = 0.722, g = 0.11$). Within the land-based sector, operators scored significantly lower than suppliers ($p = 0.018, g = 0.79$, medium effect); within the online sector, operators and suppliers did not differ significantly ($p = 0.406, g = 0.30$). Figure A1 presents the interaction plot for the role \times sector ANOVA, illustrating differences in Maturity Index scores.

Figure A1 – Interaction of Role and Sector on Maturity Index (N=83)

Source: AiR HUB Industry Survey



Drivers of AI Adoption

To supplement the findings in **Section 1: The Industry Index** regarding the top drivers for AI adoption in the industry, we further analyzed top drivers for AI adoption in four distinct groups: land-based operators, land-based suppliers, online operators, and online suppliers. The top three drivers for AI adoption within these groups are shown in Table A2.

Table A2

Rank	Land Operators	Land Suppliers	Online Operators	Online Suppliers
1	Business need to reduce costs (46)	Business need to reduce costs (33)	Improve customer experience (21)	Identify new revenue opportunities (26)
2	Identify new revenue opportunities (30)	Manage increased workload (28)	Business need to reduce costs (18)	Business need to reduce costs (25)
3	Improve customer experience (24)	Improve customer experience (20)	Identify new revenue opportunities (9)	Improve customer experience (24)

A.3.2 The Regulator Survey

Methodology

We employed a cross-sectional online survey design to examine how gambling regulators perceive, understand, and oversee AI within the gambling industry. The survey was designed to establish a baseline measure of regulatory perspectives, practices, and AI literacy. We received ethical approval for the study from the UNLV Institutional Review Board. KPMG and Vixio contributed to the research through feedback on survey instrument design, recruitment, and interpretation of aggregated findings.

We used purposive and snowball sampling strategies. We distributed invitations via email and LinkedIn through a combination of institutional and partner networks, including the IGI's International Center for Gaming Regulation contact list, Vixio's regulatory contact lists, and the International Association of Gaming Regulators membership network. We conducted additional outreach by contacting regulatory authorities directly and through in-person recruitment at industry conferences. We administered the survey online via UNLV's Qualtrics platform, with recruitment conducted over an approximately 3-month period from December 2025 to February 2026. Individuals who did not hold a current regulatory position were excluded via a screening question at the outset of the survey; those who did not meet the eligibility criterion were routed to the end of the survey and no data were collected. All responses were collected anonymously; no directly identifiable information was gathered at any point during the survey.

We organized the main part of the instrument (i.e., excluding the screening and consent and demographic information sections) across four key thematic blocks: (1) AI awareness and use, (2) beliefs on AI regulation, (3) the future of AI in gambling, and (4) AI literacy. To establish a common frame of reference, respondents were provided with a working definition of AI that articulated broader applications (i.e., beyond conversational tools like ChatGPT) such as pattern recognition, video analytics, predictive modeling, automated customer service, and text-based insight extraction. We asked respondents to keep this broad definition in mind as they completed the survey. The items comprising the first three blocks were developed by the research team in consultation with KPMG and Vixio through iterative feedback and discussion.

To construct the AI literacy block, we reviewed the academic literature for validated instruments and prior research on AI literacy measurement. We identified a recent systematic review of AI literacy scales [10] that catalogued 22 studies validating 16 scales across various populations, including general populations, students, and teachers. Thirteen of the 16 were self-report scales applicable to our study design. As noted in the review, while these scales generally demonstrated good structural validity and internal consistency, very few had been comprehensively evaluated across all measurement properties (e.g., content validity, construct validity, measurement error), and none had been tested for cross-cultural validity. We reviewed these scales for fit with our context, taking into account the target population, study aims, and participation time, the latter being a particularly important consideration for us given that the literacy block was just one component of our entire instrument. Some scales were prohibitively long for our purposes; the SAIL4ALL, for example, includes 56 items. The 12-item AILS demonstrated the most robust quality of evidence and had been used with general populations [10]; however, its items were oriented toward a very general audience and relied on self-assessment, introducing a higher risk of social desirability bias. The AI Literacy Test [11] also scored well in the review but comprised 30 items. We identified a shortened 10-item version (AILIT-S) published by the same authors in August 2025, which had been validated across 1,465 higher education students in Germany, the United Kingdom, and the United States, three jurisdictions of relevance to our regulatory respondent base [12]. Furthermore, the AILIT-S can be completed in under five minutes, making it well suited for inclusion alongside our other survey blocks.

While the ALLIT-S target population was higher education students, we considered its knowledge-based question format to be appropriate for evaluating baseline AI understanding among regulators. However, we noted that the ALLIT-S lacked items addressing Generative AI, which we felt was important to address given the prominence in current industry and regulatory discourse. To address this, we adapted items from the recently published AI Competency Objective Scale (AICOS; [13]), which includes a dedicated generative AI component. Three members of the research team reviewed the AICOS item pool and selected four items judged to be most pertinent to our population and study context. The final literacy block therefore comprised 14 knowledge-based items scored dichotomously (correct/incorrect). The complete survey instrument was reviewed by the broader research team to incorporate feedback from multiple disciplinary perspectives prior to deployment. We provide the full instrument on our OSF project page: <https://doi.org/10.17605/OSF.IO/V8UWP>.

Data Analysis

We collected 163 responses in total, of which 113 were complete and retained for analysis. 17 responses failed the screening questions: 15 of the respondents were not employed by a gambling regulatory body and two of the respondents did not provide their consent. We excluded a further three responses that were duplicates, and 30 that only provided demographic information. Sixty-eight participants completed the 14 AI literacy questions; we did not find any meaningful differences between the characteristics of this cohort vs. the broader analytic sample.

Given the exploratory nature of our study, we did not preregister any specific hypotheses or accompanying analysis plans. Our analysis followed a two-stage approach. First, we conducted an exploratory review of descriptive statistics (i.e., frequency counts, means, percentages, etc.) and Spearman correlations across all survey questions. Then, we applied inferential statistical tests to verify any meaningful differences we observed. While we utilized a standard significance threshold of 0.05, we have reported precise p -values for transparency. We conducted group comparisons for the question that asked regulators to rate their confidence in understanding and assessing AI across five items. We aggregated these items into a single competency score (Cronbach's $\alpha = 0.91$), and differences were analyzed using Mann-Whitney U tests. We explored differences across three variables: AI Use (Yes/No), AI Training (Yes/No), and AI Literacy (categorized via a median split). All analyses were conducted using Qualtrics and Python.

Several of our survey questions included open-ended responses that allowed respondents to elaborate on perceived risks, benefits, and future challenges associated with AI in the gaming industry. We analyzed these using an inductive open-coding approach. We coded responses one question at a time, rather than coding the entire dataset simultaneously to maintain conceptual independence and ensure the themes that emerged were specific to the context of each question. Individual responses were allowed to receive multiple codes when applicable, as many responses referenced more than one concept. Consequently, the total number of coded labels exceeds the number of respondents.

During the initial coding stage, we identified and assigned preliminary descriptive labels derived directly from the responses. As coding progressed, we iteratively refined our labels to reduce thematic overlap between categories. Following completion of coding for each question, we summarized the data, including a count of the number of respondents and the frequency with which each label appeared (available at <https://doi.org/10.17605/OSF.IO/V8UWP>). Definitions for each label were documented alongside their corresponding summary tables to maintain transparency in the coding process.

A.3.3 Research and Advancements

Methodology for Analysis of Academic Publications

We employed a multi-stage computational pipeline to identify, classify, and analyze academic literature at the intersection of AI and Gambling Studies. To ensure a comprehensive capture of relevant literature, we programmatically retrieved data from the OpenAlex database using the *pyalex* Python library. A “wide-net” strategy was used to mitigate the absence of a singular “Gambling” taxonomy in traditional library sciences.

Our retrieval methods included: (1) Topic-based filtering and (2) Keyword-based filtering. For topic-based filtering, we extracted all works tagged under OpenAlex Topic IDs T11674 (Sports Analytics and Performance) and T11705 (Gambling Behavior and Treatments). Our keyword-based filtering comprised a boolean search across titles and abstracts using 27 distinct gambling-related phrases e.g. “sports betting”, “PGSI”, “electronic gaming machine”, etc. We limited the search to between 2010 and 2025. A total of 103,736 unique publications were retrieved across articles, reviews, books and book chapters.

To refine our initial yield into a relevant corpus, we leveraged a local LLM deployment for topical classification. The deployment was instantiated using Ollama with the DeepSeek-R1:14b model. We performed classification of the records by first filtering by *gambling relevance* and then filtering by *AI relevance*. We tasked the LLM with ensuring that publications were in fact gambling-related and, for example, were not simply included because the title or abstract included a metaphorical reference to gambling terminology (e.g. “a political gamble”). Then, to ensure records were materially related to AI, we tasked the LLM with classifying each record based on two criteria: whether the paper featured (1) AI as a method (machine learning, natural language processing, predictive modeling) or (2) AI as a subject (ethics of AI in gambling, policy, or industry impact). Our final corpus contained 899 records identified as *AI-related Gambling Studies* publications.

Thematic Analysis

To understand the themes studied within these 899 publications, we employed BERTopic, a modular topic modeling technique for clustering text data. We used a representation model (Qwen-3:14b) to generate human-readable, context-aware labels for each of the thematic clusters by analyzing the most representative documents and keywords. We conducted two supplemental iterations of this thematic analysis: the first to identify AI methods applied across the entire corpus, and the second specifically targeting the “Sports Betting” cluster (the largest cluster) to uncover trends in specific sports and their associated research themes. We provide a description of the primary top-level gambling themes in Table A3 below.

Table A3

Theme	Description	Example Publications
Sports Betting	Publications that leverage AI methods to model and analyze historical data, player statistics, and/or real-time variables to predict event outcomes.	A Systematic Review of Machine Learning in Soccer Betting: Techniques, Challenges, and Future Directions Real-Time Football Match Prediction Platform Machine learning for sports betting: Should model selection be based on accuracy or calibration?

Theme	Description	Example Publications
Online Problem Gambling	Publications that focus on applying AI methods with player behavioral tracking data to identify at-risk behavior in digital gaming environments.	<p>Using artificial intelligence algorithms to predict self-reported problem gambling with account-based player data in an online casino setting</p> <p>Online Problem Gambling: A Comparison of Casino Players and Sports Bettors via Predictive Modeling Using Behavioral Tracking Data</p> <p>Predicting High-Risk Gambling Based on the First Seven Days of Gambling Activity After Registration Using Account-Based Tracking Data</p>
Gambling Decision Making	Publications that apply AI methods to various data sources (e.g., neuroimaging, behavioral data, etc.) to explore how the human brain processes rewards and losses.	<p>Gaming behavior and brain activation using functional near-infrared spectroscopy, Iowa gambling task, and machine learning techniques</p> <p>Elucidating Poor Decision-Making in a Rat Gambling Task</p> <p>Late acute stress effects on decision-making: The magnified attraction to immediate gains in the Iowa gambling task</p>
Gambling Website Detection	Publications that utilize automated web crawling and classification algorithms to identify, categorize, and flag illicit or unregulated gambling platforms across the internet.	<p>Gambling Domain Name Recognition via Certificate and Textual Analysis</p> <p>Multimodal Detection of Covert Online Gambling Advertisements Using Faster R-CNN and Tr-OCR</p> <p>Multi-Modal Clustering Discovery Method for Illegal Websites Based on Network Surveying and Mapping Big Data</p>
Poker	Publications exploring AI development in poker, centered on game-theory algorithms and reinforcement learning for strategic decision-making in environments with imperfect information.	<p>Superhuman AI for multiplayer poker</p> <p>A View on Deep Reinforcement Learning in Imperfect Information Games</p> <p>PokerBot: Hand Strength Reinforcement Learning</p>
Casino Surveillance	Publications that employ AI methods (e.g., computer vision and biometric technology) to support various use cases such as fraud detection, dealer errors, and identification of blacklisted individuals.	<p>AI-driven Optimization of Casino Gaming Systems for Fraud Detection and User Behavior Analysis</p> <p>Game Intelligent Analyst - Anomaly Detection in Casino Games using Machine Learning Algorithms</p> <p>Significance of AI Surveillance Platform Development for Detecting Abnormal Behaviors in Integrated Resort Industry</p>

Methodology for Analysis of Industry Conference Sessions

We defined our corpus of AI-related industry conference sessions by collecting data (the program titles and descriptions for each session) from schedules of major gambling industry conferences between 2020 and 2025. We collected conference program data from official conference websites, archived online materials, and direct correspondence with conference organizers. When session information was incomplete or unavailable, additional details were requested through direct email correspondence with conference organizers.

We input data for all sessions into a spreadsheet and then coded each session as AI-related (or not) based on whether the title or description explicitly referenced key terms such as artificial intelligence, machine learning, predictive modeling, or AI-enabled systems. Beyond keyword matching, we evaluated the context of each session to ensure material relevance. For instance, we defined sessions using broader terms like “automation” as AI-related only when presented as part of common AI-driven business functions, such as player risk detection and fraud detection. The spreadsheet containing session titles, descriptions, and conference metadata is available at <https://doi.org/10.17605/OSF.IO/V8UWP>.

To assess the reliability of the classifications of AI-related sessions, we extracted a random sample of 20 sessions, which were independently coded by three members of the research team as AI-related or not AI-related. All coders reached complete agreement on session classification.

For sessions identified as AI-related, we assigned gambling-specific thematic categories through a combination of automated and manual qualitative review. We employed a two-stage classification process to categorize the conference sessions. First, a LLM analyzed session descriptions to generate approximately 20 preliminary thematic labels. The research team then reviewed these initial categories alongside the original descriptions to consolidate overlapping concepts. This refinement reduced the set to five gambling-related categories and one miscellaneous category. Finally, two researchers independently coded all sessions manually using this finalized framework.

Methodology for Analysis of Intellectual Property

We sourced patent filing data from Lens.org [14], a searchable database of millions of global patents and scholarly works. To collect relevant data, we formulated a search strategy combining known gambling keywords with the WIPO PATENTSCOPE Artificial Intelligence Index [15].

WIPO is a specialized agency of the United Nations established to help creators and innovators on a global scale to promote and protect intellectual property. The PATENTSCOPE Artificial Intelligence Index published by WIPO provides a collection of key phrases, Cooperative Patent Classification (CPC) codes and International Patent Classification (IPC) codes that identify patent applications related to AI. While IPC is an international standard that’s used worldwide, the CPC is a more detailed and refined extension used by the European Patent Office (EPO) and United States Patent and Trademark Office (USPTO)

We employed a hybrid search strategy using keywords, phrases, and CPC codes to query the Lens.org database across titles, abstracts, and claims. Our strategy required each record to contain terms related to both gambling and AI. Some of the retrieved patents did not appear relevant at first glance and warranted a deeper manual analysis of the claims. These analyses were used to confirm the specific application of AI within a gaming context. We restricted the time period to between 2010 and 2025, which yielded a total of 2,125 patent records and 650 granted patent records. We chose to retain continuous patent refiling to capture the “strategic intensity” of the applications. Because these filings serve as extensions or continuations of the original invention, they provide a metric for how aggressively a firm is defending or expanding its IP in the AI-gambling space. To understand the main themes among patents we conducted an analysis of the CPC codes. We did attempt a text-based clustering method, however, this often unearthed false positives due to the prevalence of the (generic) “gaming” keyword.

References

- [1] H. Mayer, L. Yee, M. Chui, and R. Roberts, "Superagency in the workplace: Empowering people to unlock AI's full potential at work," McKinsey & Company, New York, NY, Jan. 2025. [Online]. Available: <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work>
- [2] M. Chui and S. Malhotra, "AI adoption advances, but foundational barriers remain," McKinsey & Company, Nov. 2018. [Online]. Available: <https://www.mckinsey.com/featured-insights/artificial-intelligence/ai-adoption-advances-but-foundational-barriers-remain>
- [3] McKinsey & Company, "McKinsey's Responsible AI Principles," McKinsey & Company, New York, NY, 2024. [Online]. Available: <https://www.mckinsey.com/capabilities/quantumblack/how-we-help-clients/generative-ai/responsible-ai-principles>
- [4] A. Singla, A. Sukharevsky, L. Yee, M. Chui, and B. Hall, "The state of AI: How organizations are rewiring to capture value," McKinsey & Company, New York, NY, Mar. 2025. [Online]. Available: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-how-organizations-are-rewiring-to-capture-value>
- [5] N. Gillespie, S. Lockey, T. Ward, A. Macdade, and G. Hassed, "Trust, attitudes and use of artificial intelligence: A global study 2025," The University of Melbourne, 2025. doi: 10.26188/28822919.V1.
- [6] N. Maslej *et al.*, "The AI Index 2025 Annual Report," Stanford Institute for Human-Centered AI (HAI), Stanford University, Stanford, CA, Apr. 2025. [Online]. Available: <https://hai.stanford.edu/ai-index/2025-ai-index-report>
- [7] K. Gopi and Gopikrishnan Janakaraja, "AI Maturity Assessment and Alignment (AIMAA) - A Comprehensive Framework for Evaluating and Benchmarking AI Adoption in Organizations," 2025, doi: 10.13140/RG.2.2.33670.89929.
- [8] A. Reuel *et al.*, "Responsible AI in the Global Context: Maturity Model and Survey," Oct. 13, 2024, *arXiv*: arXiv:2410.09985. doi: 10.48550/arXiv.2410.09985.
- [9] S. U. Lee *et al.*, "Responsible AI Question Bank: A Comprehensive Tool for AI Risk Assessment," Jan. 22, 2025, *arXiv*: arXiv:2408.11820. doi: 10.48550/arXiv.2408.11820.
- [10] T. Lintner, "A systematic review of AI literacy scales," *Npj Sci. Learn.*, vol. 9, no. 1, p. 50, Aug. 2024, doi: 10.1038/s41539-024-00264-4.
- [11] M. Hornberger, A. Bewersdorff, and C. Nerdel, "What do university students know about Artificial Intelligence? Development and validation of an AI literacy test," *Comput. Educ. Artif. Intell.*, vol. 5, p. 100165, 2023, doi: 10.1016/j.caeai.2023.100165.
- [12] M. Hornberger, A. Bewersdorff, D. S. Schiff, and C. Nerdel, "Development and validation of a short AI literacy test (AILIT-S) for university students," *Comput. Hum. Behav. Artif. Hum.*, vol. 5, p. 100176, Aug. 2025, doi: 10.1016/j.chbah.2025.100176.
- [13] A. Markus, A. Carolus, and C. Wienrich, "Objective measurement of AI literacy: Development and validation of the AI competency objective scale (AICOS)," *Comput. Educ. Artif. Intell.*, vol. 9, p. 100485, Dec. 2025, doi: 10.1016/j.caeai.2025.100485.
- [14] "The Lens - Patent and Scholarly Search and Analysis," The Lens - Patent and Scholarly Search and Analysis. Accessed: Mar. 16, 2026. [Online]. Available: <https://www.lens.org/lens>
- [15] "PATENTSCOPE Artificial Intelligence Index," technology-trends. Accessed: Feb. 25, 2026. [Online]. Available: https://www.wipo.int/en/web/technology-trends/artificial_intelligence/patentscope

Contact us

Kasra Ghaharian, PhD

Director of Research, UNLV International Gaming Institute
Principal Investigator & Editor-in-Chief, The State of AI in Gaming
+1 (702) 895 2008
kasra.ghaharian@unlv.edu

Rick Arpin

US Gaming Leader, KPMG LLP
Executive Editor, The State of AI in Gaming
+1 (702) 271 8253
rarpin@kpmg.com

Simo Dragicevic

Adjunct Fellow, UNLV International Gaming Institute
Executive Editor, The State of AI in Gaming
+1 (702) 895 2008
simo.dragicevic@unlv.edu

© 2026 University of Nevada, Las Vegas, International Gaming Institute.

KPMG LLP's involvement in this report was limited to providing input on survey design, recruitment support, and contextual interpretation of aggregated findings. KPMG did not conduct the underlying research or perform any assurance or attestation services in connection with this report. The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS040051-1A

This work is licensed under CC BY-ND 4.0.