SPOTDRAFT

PLAYBOOK

AI & LLM Use Policy: Development and Implementation

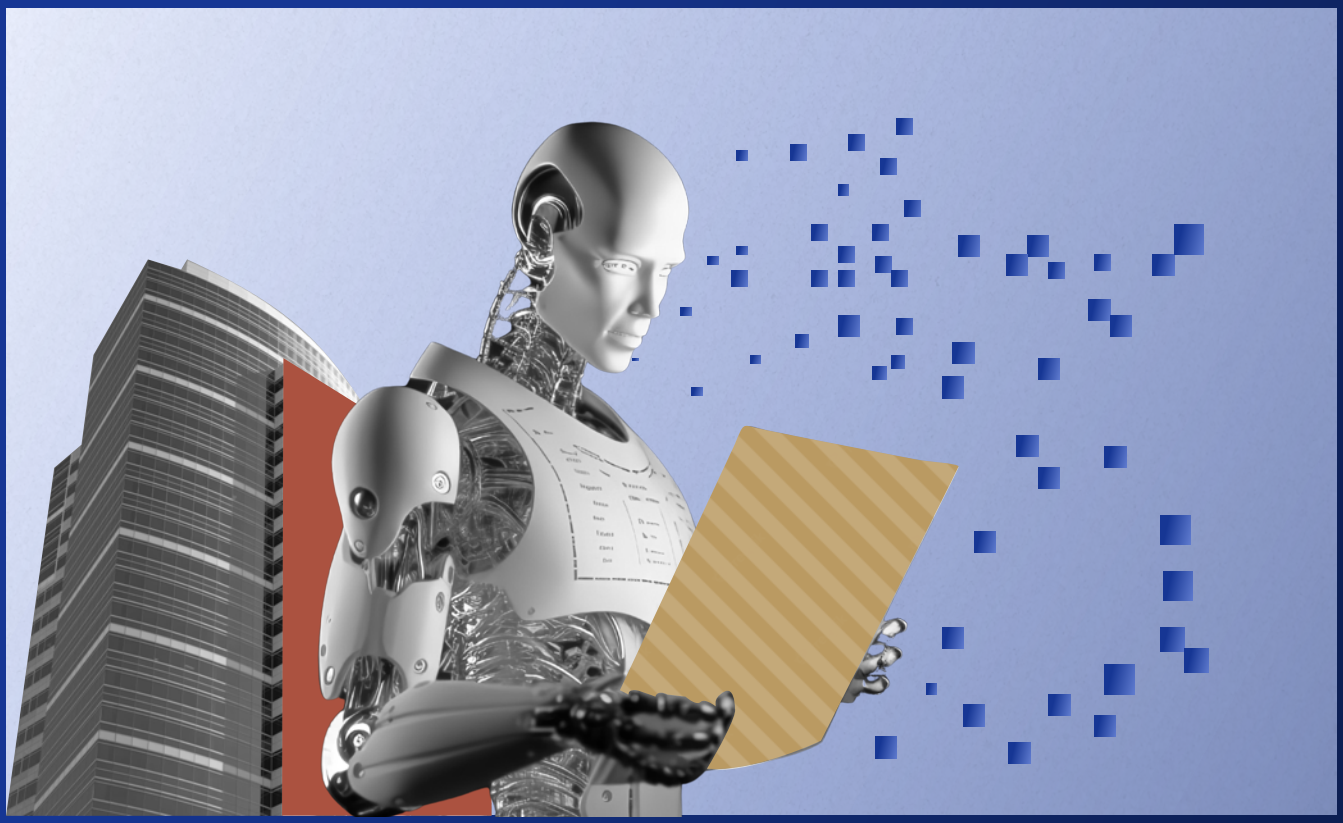

Preamble

How do you advise and train your organization on how to, and how not to, use Large Language Models (LLMs) and Artificial Intelligence (AI) tools for work, especially for product development?

And how do you regulate the use of these tools within the business?

This is a nuanced issue occupying the minds of many in-house lawyers today. We found a noticeable gap of comprehensive resources available to help legal teams tackle this challenge, and we put together a team of experienced in-house lawyers and tech experts to create this playbook.

It contains everything you need to know to start building your internal policies around LLM and AI use, from existing frameworks to base your policies on to complex issues around internal trainings, tool evaluations, and policy enforcement.

This guide is built on deep research, detailed discussions, and insights from leaders and legal professionals who have personally tackled the challenges and opportunities of using AI and LLMs in business*. It offers a clear, practical approach to developing and implementing AI use policy that helps employees benefit from AI and LLMs in a way that is ethical, legal, and aligned with your organization's goals.

***Find out who contributed to this playbook on page 43.**

How to use this playbook

Use this as a step-by-step guide

This playbook is designed to be a straightforward guide through the journey of understanding, developing, and rolling out AI and LLM policies in your business. From understanding the current state of AI and LLM use, exploring the key benefits and challenges, to diving into policy development and rollout, we've covered it all.

What this playbook isn't

While this guide offers a strong framework and insights for policy development and rollout, it's not a one-size-fits-all solution or legal advice. Every organization has unique needs and challenges. Use this guide as a starting point and work with your internal teams to develop policies that fit your organization perfectly.

Adapt and evolve alongside technology

Use this playbook as a flexible tool. Apply the tips and advice given here, and shape the guidelines to fit your organization's unique needs and challenges. Technology and AI are always evolving, so make sure your policies change too. Keep coming back, revising, and updating as you navigate through the changing tech landscape.

For any queries or comments, please reach out to us at community@spotdraft.com.

Contents

01

The State of AI and LLM Use Today

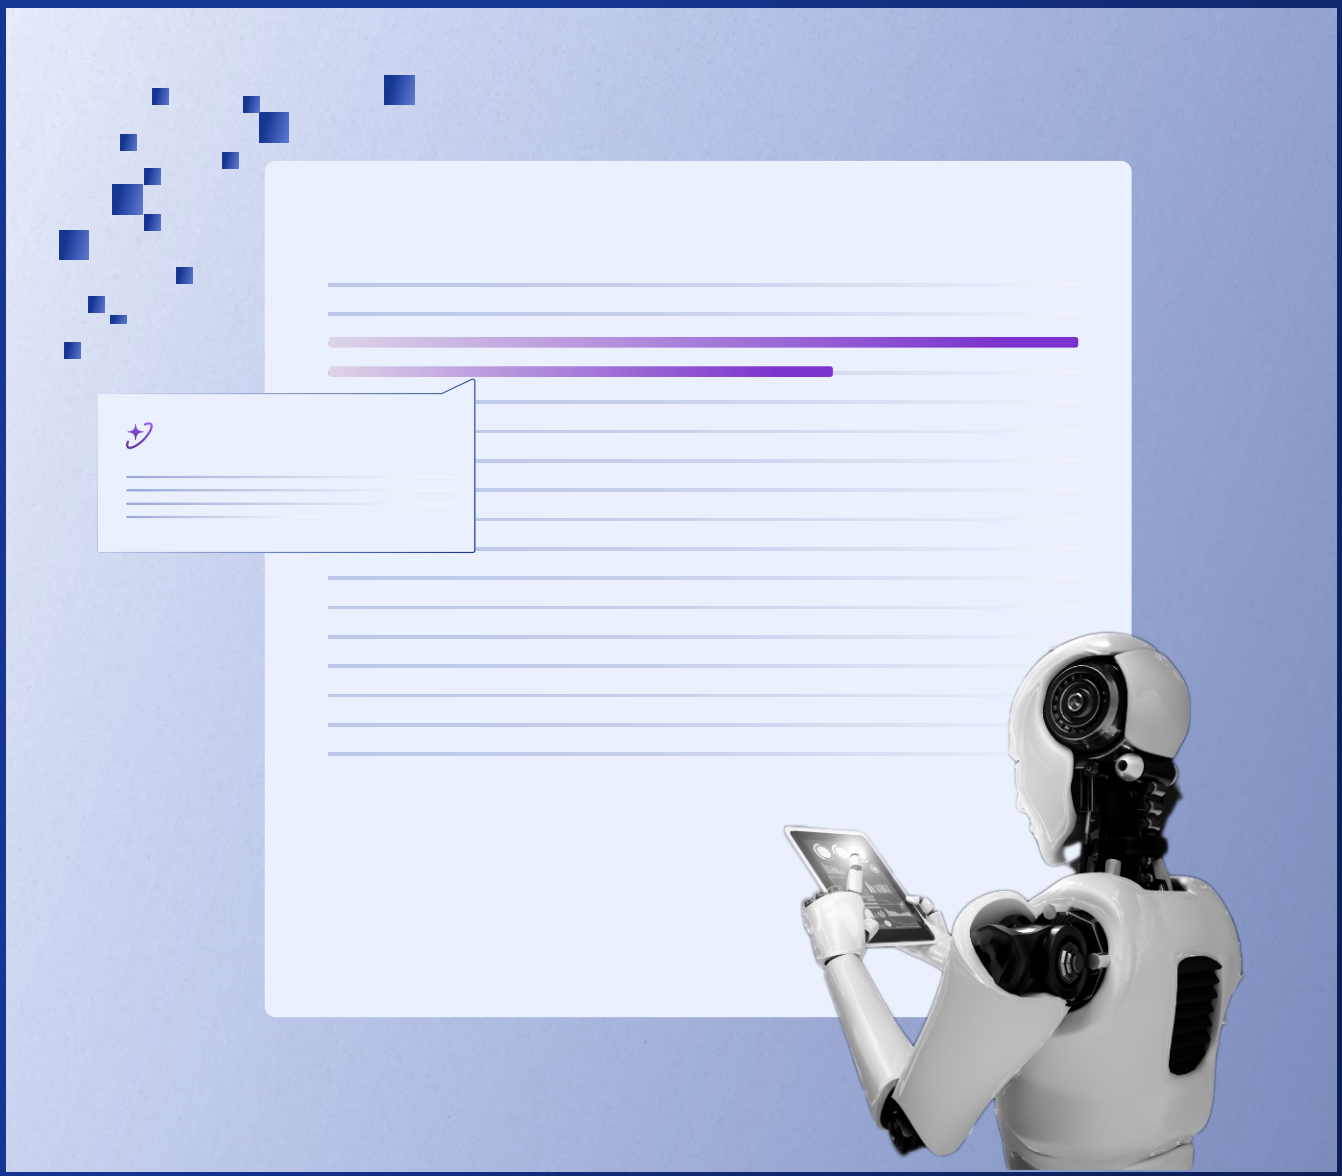

Key benefits of using AI in business

How are the Working Group and their companies using AI today?

Enhanced efficiency

AI can automate repetitive tasks, reducing manual labor and human error.

Improved decision-making

AI can analyze vast amounts of data quickly and accurately, providing valuable insights that aid in decision-making.

Personalization

AI-powered recommendation engines enable businesses to offer personalized product or service recommendations to customers, enhancing the customer experience and driving revenue.

Cost reduction

Automation through AI can lead to cost savings in various areas, such as customer support, operations, and logistics.

Customer service

Chatbots and virtual assistants powered by AI can provide 24/7 customer support, resolving queries and issues in real time.

Natural Language Processing (NLP)

NLP applications can analyze customer feedback, social media, and reviews to understand sentiment and gather insights for product development and marketing.

Fraud detection

AI algorithms can identify unusual patterns in financial transactions, helping businesses detect and prevent fraud.

Predictive maintenance

In manufacturing and equipment-heavy industries, AI can predict when machines or equipment are likely to fail, allowing for proactive maintenance and reduced downtime.

Risks and challenges around AI use

What are the risks associated with using AI tools and LLMs?

Data privacy

Collecting and using personal data for AI applications raises privacy concerns and may lead to regulatory compliance challenges.

Bias and fairness

AI algorithms can perpetuate biases present in training data, leading to unfair or discriminatory outcomes. This is a chief concern in our Working Group.

Hallucinations

An AI hallucination is when an AI model generates incorrect information but presents it as if it were a fact. Our Working Group is also deeply concerned about employees using AI tools without checking the accuracy of their outputs.

Intellectual property issues

Use of AI tools to generate marketing materials (i.e videos, copy, or pictures) or code may result in use of material that the business does not have the right to use. Moreover, the same uses may result in outputs that are not able to be copyrighted, trademarked, or patented, inadvertently undermining the business' IP strategy. Companies are encouraged to build on existing practices related to using open source code as a starting point.

Data quality

If the data put into an AI model is of low quality, the output is also likely to be of a low quality. The Working Group was concerned that the business may put "garbage in" and get "garbage out" without knowing it.

Security threats

AI systems can be vulnerable to attacks and exploitation, requiring stronger cybersecurity measures.

Change management

Integrating AI into business operations may require changes in workflows and employee training, which can be met with resistance.

Transparency and accountability

Explaining AI decisions and ensuring accountability can be difficult, especially in complex deep learning models. This is another major concern of legal leaders.

Ethical concerns

Ethical considerations, such as the use of AI in surveillance or decision-making, can lead to public backlash if not handled appropriately.

KEY TAKEAWAY

What is clear is that businesses need to carefully assess the benefits and risks of AI adoption, implement responsible AI practices, and stay informed about evolving regulations and ethical standards to maximize the advantages while mitigating potential challenges.

"Most organizations are realizing that they should have a policy in place for AI adoption, because, otherwise, there's a risk of customer data or confidential data being exposed."

Ken Priore

Associate General Counsel

DocuSign

SpotDraft 2023, *Mastering the Intersection of Law, Technology, and Privacy*, accessed 6 October 2023, <https://www.spotdraft.com/blog/counselcorner-mastering-law-technology-privacy-ken-priore>.

Use cases

What use cases should generally be permitted and what should require review?

Defining permissible use cases and outlining criteria for review are critical aspects of AI policy development. Our Working Group emphasized that Legal needs to understand how LLMs work just as well as their business stakeholders. Often, they worked closely with Privacy and Security experts to dig deep into how the technology works.

Here are some key considerations our Working Group discussed regarding common use cases:

APIs/Enterprise GPT usage

It was generally agreed upon that any usage of APIs or Enterprise GPTs should not involve training the underlying model, because the data inputs may show up in the open source code. Further caution is advised, as any data inputs to an LLM may be subject to human review. APIs or Enterprise GPTs that allow for use without training should be subject to policies that govern data retention, especially if they involve personally identifiable information (PII). A typical guideline might be to retain PII for no longer than 30 days with no training involved.

Exclusion of legal and product counseling

AI tools and LLMs should not provide legal advice, or be used to generate product counseling advice.

Handling sensitive data

Strict policies should govern the use of AI in handling sensitive data to ensure compliance with privacy and security regulations. This includes considerations for HIPAA compliance. AI applications involving sensitive financial information, subject to PCI (Payment Card Industry) regulations, should also be closely monitored and reviewed.

HR and employment data

The use of AI tools with HR data or employment-related data should always be escalated and subject to stringent review. This is crucial to protect employee rights and privacy.

Automated decision making

Policies should align with GDPR principles, which generally require human input in automated decision-making processes. This also helps avoid potential bias or discrimination.

Protecting code

AI-generated code that could be sensitive or proprietary should be subject to review. Furthermore, sensitive or proprietary code should not be put into public LLMs, as the business will want to prevent competitors , hackers, or other third-party actors from gaining access to it.

System access credentials

The use of AI to access system credentials, especially those of third-party systems, should be carefully controlled and reviewed.

IP outputs

Copyrightable information produced by AI should be subject to clear policies. Companies should ensure that AI-generated content does not infringe on intellectual property rights. Further consideration is warranted around the issue of AI-generated content potentially not being the company's intellectual property.

Sharing data with AI/LLM vendors

What categories of data are generally okay to share with vendors who have AI functionality or share data back to LLM providers?

The categories of data that are okay to share with vendors who have AI functionality or share data back to LLM providers should be carefully considered to ensure compliance with data privacy and security regulations applicable to your industry and business. Here are some categories of data the Working Group was generally comfortable with sharing with an LLM provider:

Non-personal data

Non-personal data that does not contain personally identifiable information (PII) can often be shared with vendors or LLM providers more freely. This may include non-sensitive business data, general statistical information, or publicly available data.

Aggregated and anonymized data

Data that has been thoroughly aggregated and anonymized to a point where individual identities or sensitive information cannot be discerned may be suitable for sharing. However, even anonymized data can sometimes be reverse-engineered, so caution is required.

Explicit consent data

Data for which individuals have given explicit consent for AI model use can be shared with vendors or LLM providers, provided that the consent aligns with the intended use.

Educating employees about AI

How are businesses educating their employees about the different types of AI and their use?

The Working Group employs various strategies to educate their employees about the different types of AI. Here are some common approaches from our conversations:

01

AI technology training programs

Many companies offer training programs or workshops on AI and its various applications. These programs can range from basic awareness sessions to more in-depth technical training for specific AI applications.

02

Compliance trainings

Companies are incorporating a discussion around AI into their trainings on sensitive data and privacy/security regulatory obligations.

03

In-house workshops

Some businesses organize in-house workshops and seminars conducted by external AI experts or employees with AI expertise. These sessions can cover AI fundamentals, use cases, and practical applications.

04

Lunch and learn sessions

Informal "lunch and learn" sessions allow employees to gather during lunch breaks (virtual or in-person) and discuss AI topics. These sessions are typically led by knowledgeable colleagues or guest speakers.

05

AI champions

Identifying and appointing AI champions within the organization—employees who are knowledgeable about AI and can serve as mentors or resources for others—can be effective in spreading awareness.

06

AI committees or groups

Establishing AI-focused committees or interest groups where employees can discuss AI trends, share knowledge, and collaborate on AI-related projects can foster learning and engagement. We'll talk more about this below.

07

Newsletters and internal communications

Regularly sharing AI-related news, articles, and updates through internal newsletters or communication channels keeps employees informed about AI developments. Some companies have Slack channels dedicated to discussions around AI developments.

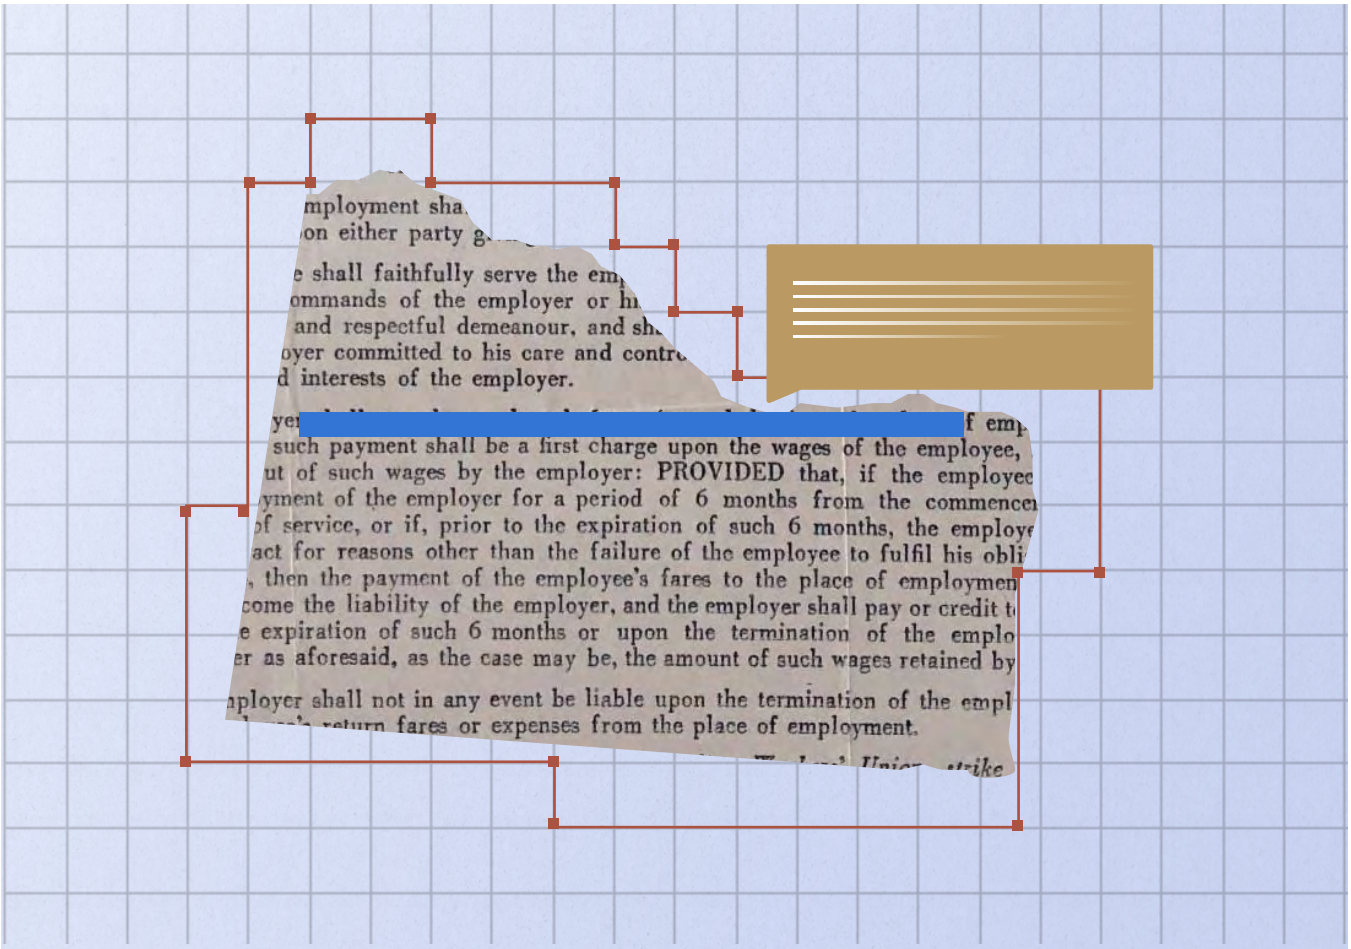

02

AI Policy Development

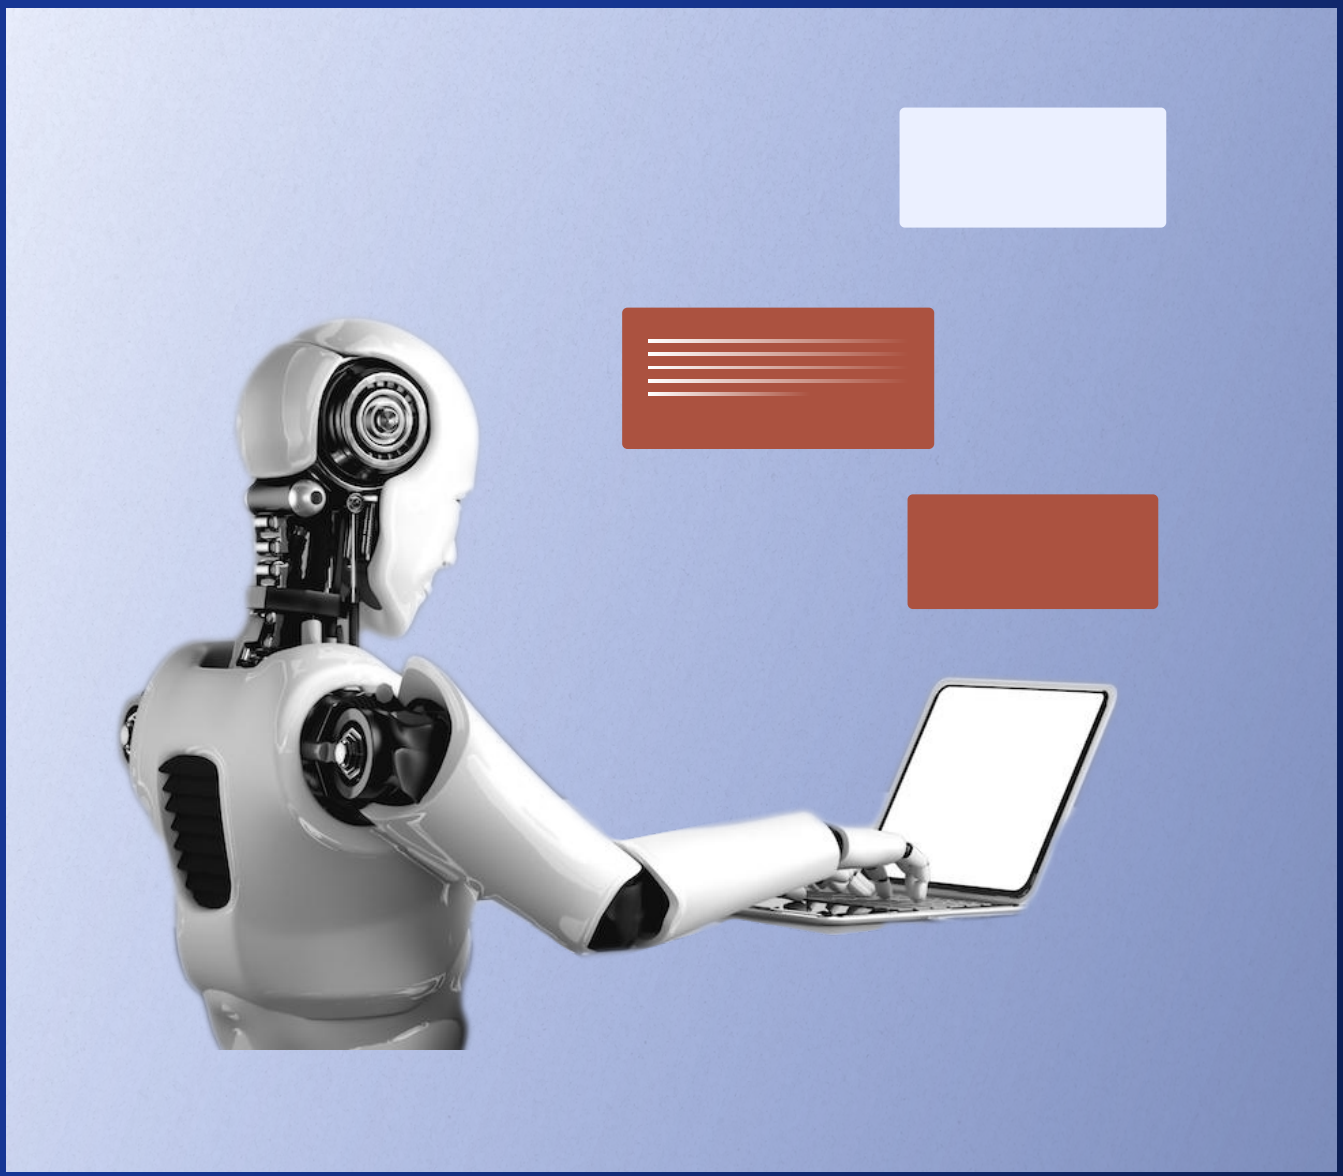

Laying down scope for internal and external policies

What's the scope of an AI policy? Do you need both internal and external policies? Should these be two different policies?

The Working Group distinguished the development of internal and external policies, which are often crafted with different intentions and tones.

Companies frequently maintain two (or multiple) separate internal policies, such as those governing product and engineering teams, and another for IT teams handling the tools used across the organization. These internal policies typically outline permissible use cases, mandatory legal review for AI-related product development, and access criteria for business accounts. Review processes (to be discussed further below) focus on accuracy and preventing hallucinations and bias.

External policies, designed for communication with customers and stakeholders, are equally vital. These policies aim to address questions and concerns that arise when companies introduce AI features to their products or services. They should not only emphasize safeguards, but also highlight the benefits of AI in the context of the company's own products or services. Educating customer support teams is crucial for effectively implementing external policies (to be discussed further below as well).

Companies may be able to leverage existing policies to govern internal use of AI tools. AI, from a procurement perspective, is often treated like any other tool and is subject to relevant procurement policies. Furthermore, companies are increasingly incorporating elements of their Acceptable Use Policies to address issues related to discrimination and bias. AI policy development also sometimes involves updating existing IT and security policies to align with AI-specific considerations, such as defining a baseline set of use cases and establishing an escalation process for non-standard AI applications (e.g., HR use cases requiring review).

KEY TAKEAWAY

In summary, the Working Group recommends having separate internal and external policies, alongside leveraging existing corporate policies (such as procurement and IT/Security policies).

Considering values and ethics when regulating AI use

What role do ethics play in the policy development process?

When it comes to AI policy development, the Working Group agreed that values and ethics play a pivotal role, and may make the development process easier. Some companies have used their existing corporate values to drive or guide AI-use policy development.

Generally, the Working Group found it helpful to incorporate a company's values and established ways of working into the policy development process.

> Companies offered that principles including objectivity, impartiality, transparency, ethical conduct, confidentiality, dignity, and championing the customer were helpful to point to as they rolled out their AI-use policy to help with buy-in from other teams.

Of course, achieving alignment among executives is also essential. While some executives may not be surprised by the challenges posed by AI and may view them as akin to prior developments (e.g. privacy, data security, open source, etc.), it's important to recognize that these new challenges place a significant responsibility on employees. Employees are increasingly held accountable for the ethical and legal implications of their work, so clear policies truly are necessary. Generally, the Working Group felt that pointing back to their existing corporate values made it easier to gain buy-in for the AI-related policies.

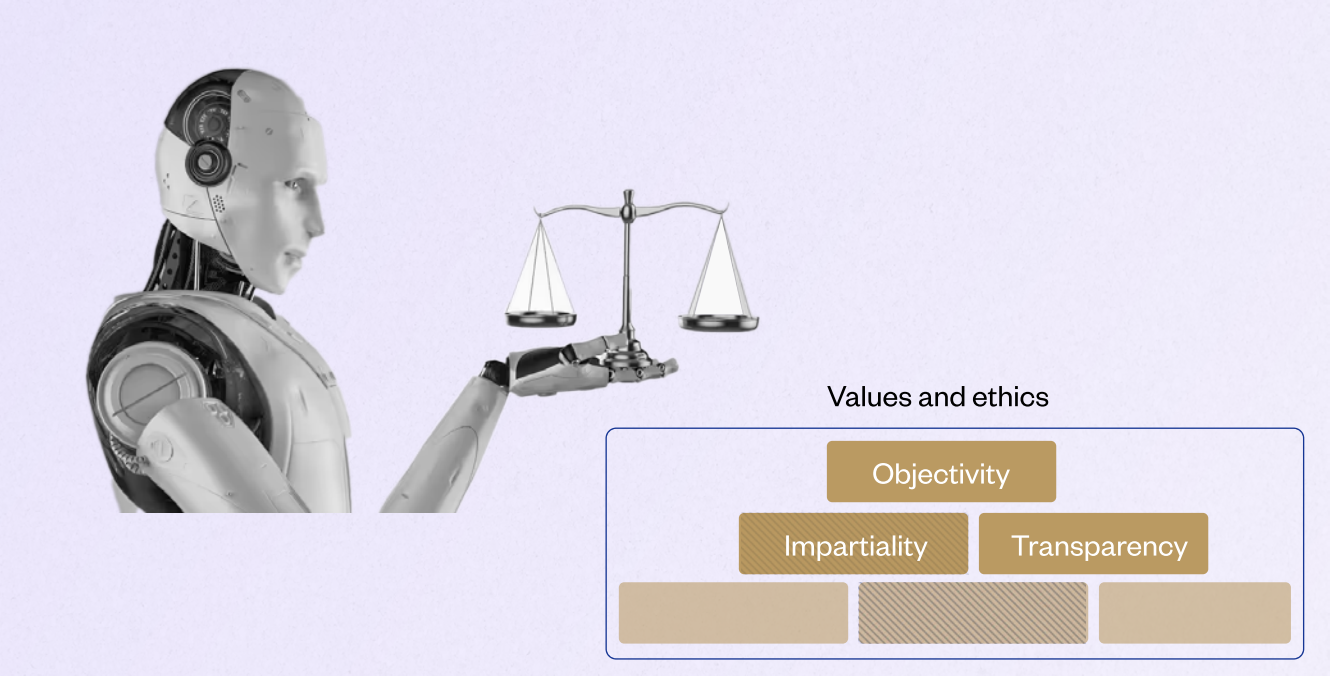

Stakeholders involved during development and rollout

Should you build an internal team to develop and rollout AI policies? What stakeholders should participate?

Establishing an internal group or team for AI policy development and rollout is often necessary to ensure consistency, alignment, and effective implementation of policies across business teams. Here are key points the Working Group discussed to establish an effective internal team or group:

Stakeholder participation

The team should include a diverse set of stakeholders. Some of the critical stakeholders to involve are:

Product teams

They play a core role in determining how AI will be used internally and selecting the AI tools used by employees.

Legal and privacy

Legal experts, especially those specializing in privacy and intellectual property (IP), are crucial to ensure that AI policies comply with relevant laws and regulations.

IT and security

Involving IT and security teams is essential to address technical and security aspects of AI policy implementation.

Compliance

Monthly compliance meetings can serve as a platform for incorporating compliance perspectives into AI policy development.

Product-first mindset

The product teams often take the lead in determining how AI will be used internally. They are at the core of identifying use cases and selecting AI tools.

Product counsel involvement

Product counsel can play a critical role in shepherding perspectives from legal, including privacy and IP, and ensuring that AI policies align with legal requirements.

Phased engagement

Security and compliance considerations may be phased in, particularly when engaging with vendors or third-party AI providers.

KEY TAKEAWAY

Ultimately, the Working Group generally agreed that Legal (often including Privacy & Compliance) needs to offer a final judgment and be comfortable with the risk posture, as Legal is the team that has to present and defend the policies to the C-suite.

Leveraging existing legal frameworks and policies

What are policy development teams looking to for guidance?

While law and regulation around AI is in its early days, some companies are looking to leverage developing frameworks such as the EU AI Act and the White House AI Bill of Rights as part of their approach to AI policy development. These regulations influence how external policies are framed and what documentation is required.

Drawing from prior experiences with regulations like GDPR and CCPA can be beneficial, as risk thresholds may align with how the EU AI Act categorizes sensitive information and risk. Generally, the Working Group agreed the landscape remains unclear, similar to the period before GDPR went into effect.

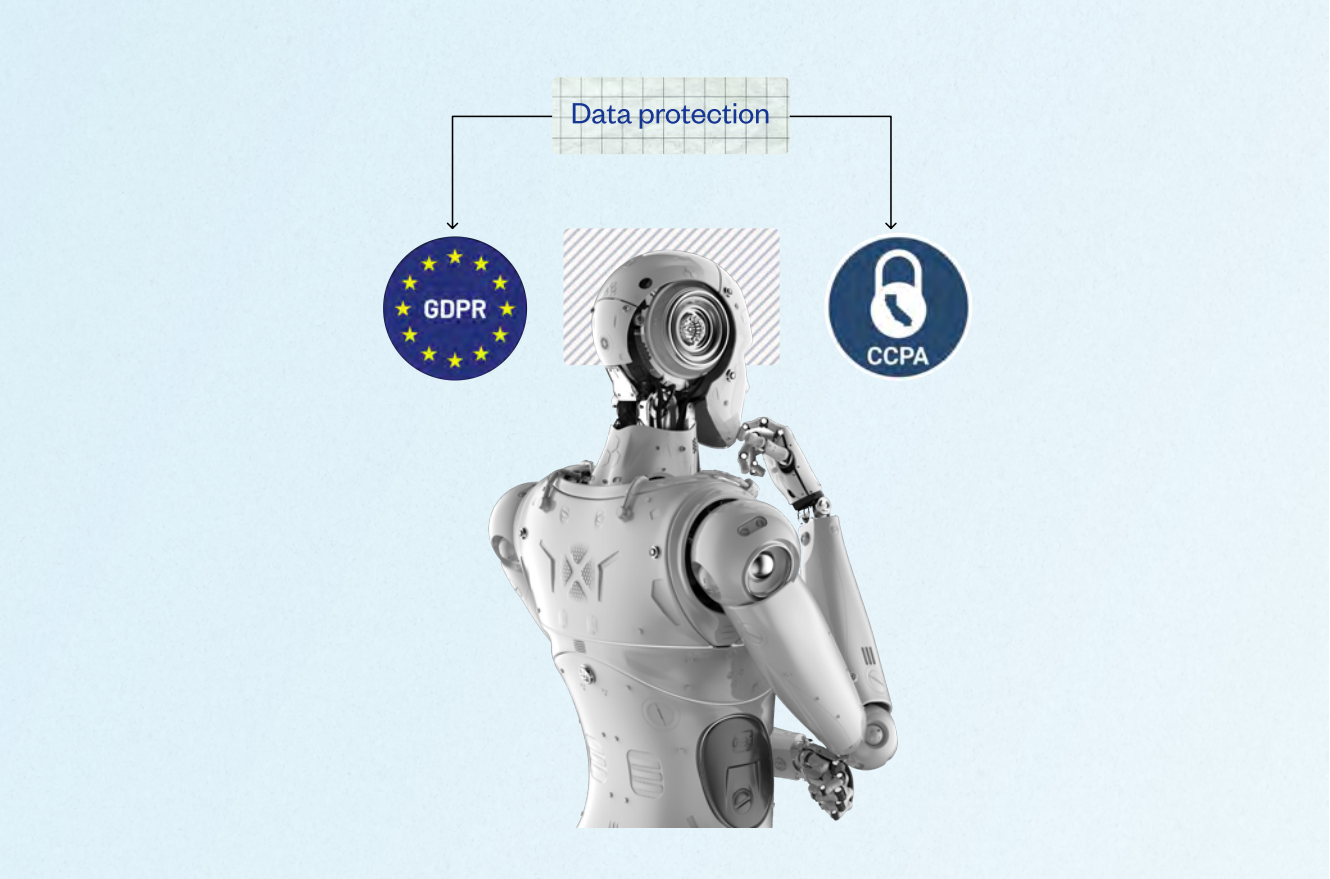

What existing policies are you leveraging? How can you not start from zero?

The Working Group felt that leveraging existing policies and adapting them to incorporate AI-specific considerations is a practical approach to integrating AI technologies into your organization without starting from scratch. Here's how you can leverage and adapt existing policies to address AI integration:

Data classification policy

Leverage

Use your organization's data classification policy as a foundation to determine what types of data can and cannot be used with AI tools.

Adapt

Add specific guidelines and requirements related to data usage in AI processes. Include provisions for data anonymization, encryption, and retention, especially for AI-generated data.

Procurement policy

Leverage

Enhance your existing procurement policy to include AI-specific considerations. This policy can guide the acquisition of AI tools and services.

Adapt

Incorporate criteria for evaluating AI vendors and tools, such as data privacy, security, bias mitigation, and compliance with AI regulations. Specify that the adoption of AI tools should be supported by a business case to prevent ad hoc acquisitions.

Overspending controls

Leverage

Existing controls for preventing overspending can be applied to AI tool procurement to ensure that AI-related expenses are within budget constraints.

Adapt

Consider adding specific controls or approval processes for AI-related expenditures to avoid unexpected costs associated with AI tool adoption.

User account and access policies

Leverage

Policies governing user account creation and access can be extended to cover AI tool usage.

Adapt

Incorporate provisions to prevent unauthorized use of AI tools, clarify the usage of company accounts for AI tool access, and establish procedures for monitoring and auditing AI tool access.

Compliance and regulatory policies

Leverage

Policies related to compliance with industry-specific regulations can serve as a foundation for ensuring AI solutions adhere to legal requirements.

Adapt

Integrate AI-specific compliance requirements, such as data privacy, fairness, transparency, and accountability, into existing regulatory policies.

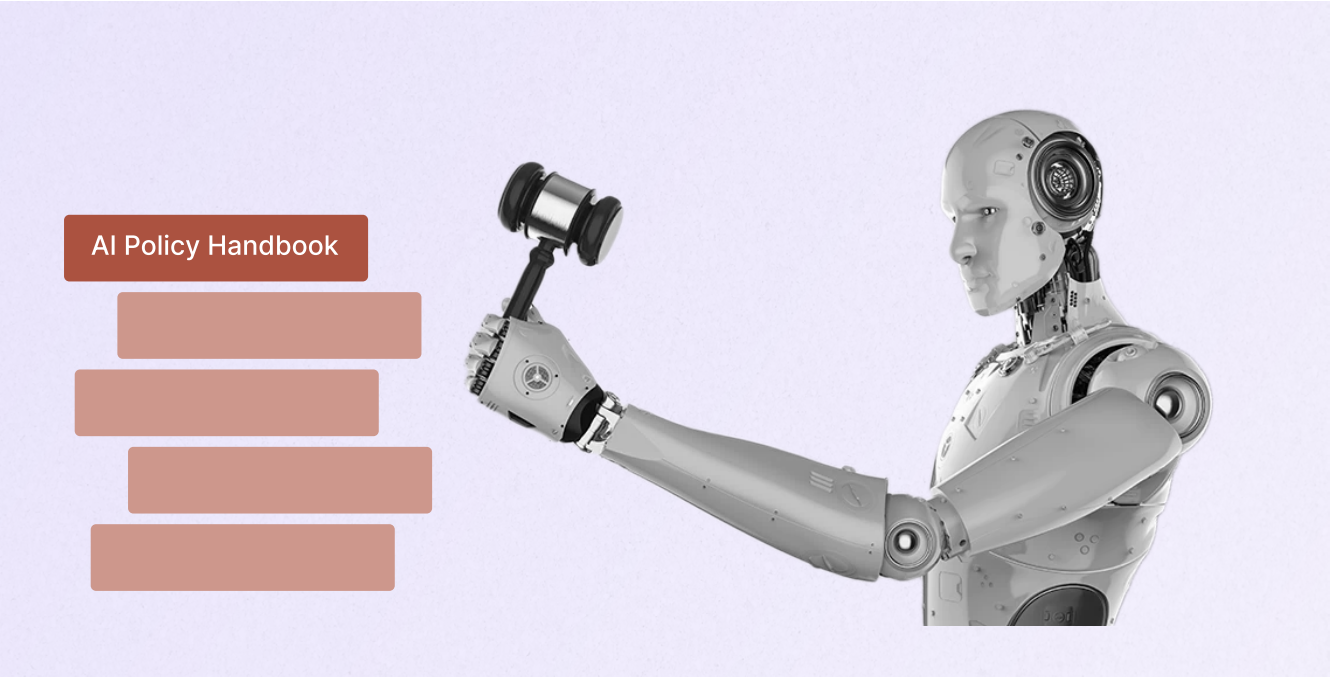

Monitoring and evaluating inputs and outputs

How do you monitor the inputs and the outputs of AI models? How do you evaluate the outputs?

Monitoring the inputs and outputs of AI models and evaluating their performance is a critical aspect of responsible AI use. Here's some tips for how you can approach monitoring and evaluation:

Monitoring inputs and outputs

Bias and fairness monitoring

Continuously monitor AI model inputs and outputs for biases and potential discrimination. Implement fairness metrics and tests to assess how the model's decisions may affect different demographic groups.

Data logging

Maintain detailed logs of the data inputs used for training AI models. This includes information about the source, type, volume, and any preprocessing steps applied to the data.

Real-time monitoring

Implement real-time monitoring of AI model inputs during inference. This can include checking for anomalies, data drift, or unexpected input patterns that may affect model performance.

Input validation

Validate incoming data to ensure it meets the model's requirements and expectations. This helps prevent data that could compromise model performance from being used.

Data source tracking

Keep track of the sources of data used for training to identify any issues with data quality or integrity. This is crucial for identifying potential biases in the training data.

Evaluating outputs

Accuracy testing

Regularly test AI model outputs for accuracy and performance against predefined benchmarks.

Bias and fairness assessment

Assess AI model outputs for bias and fairness. Utilize fairness metrics to determine if the model's decisions are disproportionately impacting certain groups or demographics.

Testing against real data

Evaluate AI model outputs against real-world data or ground truth data. This can help identify any discrepancies between model predictions and actual outcomes.

User feedback

Gather feedback from users or stakeholders who interact with AI outputs. Their input can provide valuable insights into the model's performance and any potential issues.

Continuous improvement

Use the results of monitoring and evaluation to iteratively improve AI models. Address any identified biases, accuracy issues, or other performance shortcomings.

Model selection

Consider the choice of large language model (LLM) providers and cloud services partners, as this can impact model outputs. Evaluate different providers for their model quality and performance.

Consideration of constraints

Understand that there may be constraints or limitations on using the "best" model due to various factors like business politics, compliance requirements, or the introduction of new subprocessors. Balancing model quality with other considerations is important.

Documentation

Maintain comprehensive documentation of the evaluation process, including the metrics used, test results, and any actions taken to address issues.

KEY TAKEAWAY

Effective monitoring and evaluation are ongoing processes. They are essential for ensuring that AI models remain accurate, fair, and compliant with the organization's goals and standards. Regularly reviewing and improving model performance helps maintain trust and reliability in AI applications.

Key elements of an AI policy for product development

What are some elements to include in an AI policy related to product development?

Creating a product AI policy is essential to ensure transparency, compliance, and ethical use of AI technologies in your product. While the specific elements of a product AI policy may vary depending on the nature of your product and industry, or if the policy is designed for internal guidance vs. external awareness, here are key elements to consider including:

Purpose and scope

- Begin with a clear statement of the policy's purpose and scope. Explain why the policy exists and which aspects of the product and AI it covers.

Definitions

- Provide clear definitions of key terms related to AI, such as "AI-generated content," "AI recommendations," and any technical terms relevant to your product.

Transparency and disclosure

- Notify users when AI is used in the product.
- Clearly label AI-generated elements or content.
- Explain that AI-generated content may not always be accurate and encourage users to review it critically.
- Consider implementing multiple layers of notices to ensure transparency.

Consent

- Explain how users can provide consent for AI usage, if applicable.
- Describe how users can opt in or out of AI-generated features or recommendations.

User notification

- Address whether notifications about AI usage should be provided on a per-user basis or on a per-instance basis. This may vary depending on the specific use case and context.

Data usage and privacy

- Describe how user data is collected, stored, and used in AI processes.
- Address data privacy concerns and compliance with relevant data protection regulations (e.g., GDPR, CCPA).
- Explain the steps taken to anonymize and protect user data.

Bias and fairness

- Outline measures taken to minimize bias in AI algorithms.
- Describe how fairness and equity are considered in AI development.

Security

- Detail security measures in place to protect AI models and data from unauthorized access and breaches.

Review and audit

- Describe internal processes for regularly reviewing and auditing AI systems for accuracy, fairness, and compliance.

Updates and changes

- Explain how the policy will be updated to reflect changes in AI technology, regulations, or user feedback.

KEY TAKEAWAY

Above all else, remember that transparency and user trust are essential elements of an effective product development AI policy.

03

Policy & Process Rollout

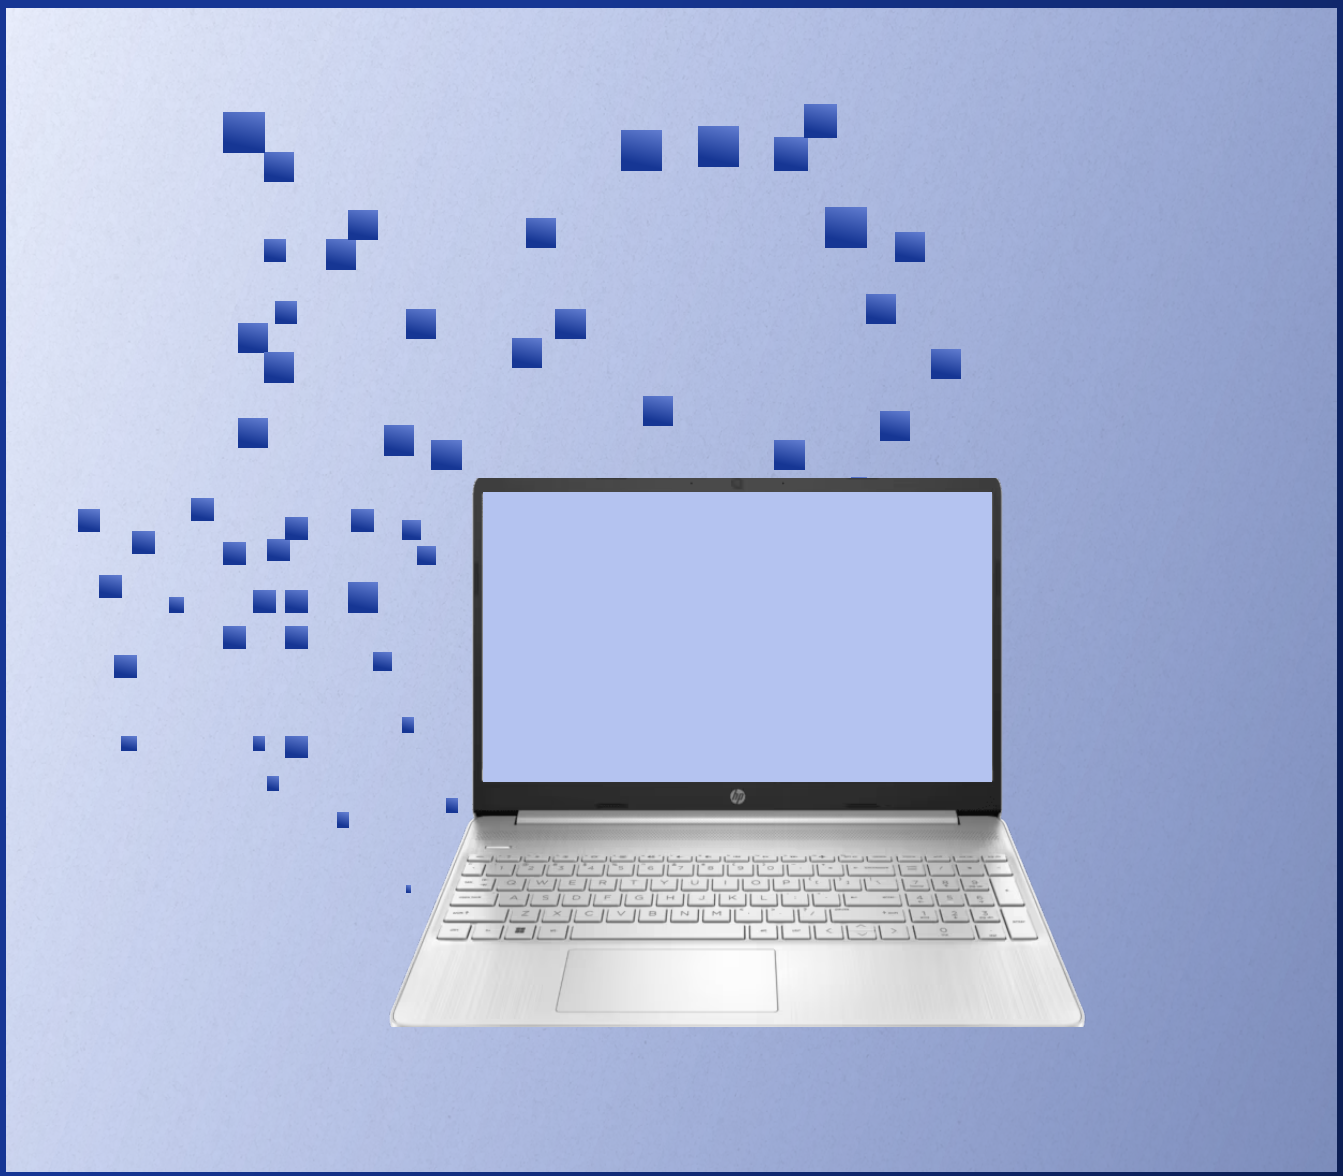

SPOTDRAFT

AI policy trainings

What do trainings on AI policies look like? How often will they be done? Who will they be done for?

Training on AI policies is a critical component of ensuring that employees understand the guidelines and responsibilities associated with AI use. The Working Group felt that roll-out and training was just as, if not more important than the development of the policy itself. Here are some approaches that the Working Group felt worked well:

Company-wide communication

Consider company-wide communication methods, such as TED talks during town hall meetings, to raise awareness about the high-level risks associated with AI. Emphasize that both inputs and outputs could potentially identify the company. These communications may be separate from formal trainings, to either prime the business for a training, or to keep AI issues top of mind.

Company-wide rollout

The Working Group agreed that the first AI policy training should be rolled out to the entire company rather than singling out specific departments. This promotes a culture of AI responsibility across the organization and prevents specific departments from feeling like they have been targeted.

Department-level training

After the company-wide rollout, follow-up training sessions can be conducted with specific departments individually. This approach can help address department-specific concerns and nuances. Tailor training materials to be more specific for departments like engineering, product development, or marketing, focusing on their unique risks and considerations. Some departments, such as ML engineers, may require more detailed and specialized training due to their involvement with AI on a deeper level.

Integration with privacy training

On an ongoing basis, AI policy training can be integrated as a module within the broader privacy training, ensuring that employees receive AI-related education at least once a year. You could also consider aligning training with annual security training efforts.

Procurement considerations

Training or discussions around procurement can also cover AI policy aspects, and this could be integrated into procurement training.

Feedback channels

Create communication channels, like company-wide Slack channels, to solicit input from employees regarding AI and their observations in the market.

KEY TAKEAWAY

One of the more interesting approaches that a company in the Working Group used is to tell their business to "Think of AI as an Intern." They emphasized that each employee is ultimately responsible for their work.

Navigating tone and culture around AI use

What is the tone of rollout across the business? How is the policy intended to affect corporate culture?

The tone of the rollout of AI policies across the business is a critical aspect of its success because it can significantly impact corporate culture. Here are key considerations regarding the tone and its intended effects:

First mover or pack follower

The tone may vary based on how the organization views its AI strategy – as a first mover, a fast follower, or adopting a more cautious approach. The tone of the rollout should reflect the organization's stance.

Balanced tone

For most companies, the tone of the policy rollout should strike a balance between highlighting the importance of AI compliance and responsibility without coming across as overly restrictive or stifling innovation.

Board focus on AI

Many boards are increasingly focused on AI due to its strategic importance and associated risks. The policy rollout should acknowledge this and demonstrate that legal compliance is not an innovation blocker but rather an enabler of responsible innovation that the whole company needs to align around.

Corporate culture impact

The policy's intent is to shape and influence the corporate culture positively. It aims to foster a culture of responsibility, transparency, and ethical use of AI technologies throughout the organization.

Growth-oriented tone

To align with a growth-oriented approach, the policy can emphasize that responsible AI usage contributes to sustainable growth and competitive advantage.

Sandbox environment

Encouraging a "sandbox" environment for experimentation can be a positive way to frame the policy. It communicates that innovation is still encouraged, but within defined boundaries that ensure safety and compliance.

KEY TAKEAWAY

In essence, the tone of the policy rollout should be supportive of innovation and growth while emphasizing the necessity of responsible AI use. It should encourage a culture where employees understand the risks, their responsibilities, and the role AI plays in achieving the company's strategic goals.

"To me, AI is the next internet. It's going to move rapidly in a direction where it will help us make processes more efficient. And if you're not using it, you're falling behind."

Celaena Powder

Vice President, Legal

Seismic

SpotDraft 2023, *From Security to Sales: Building Cross-functional Bridges*, accessed 6 October 2023, <https://www.spotdraft.com/blog/counselcorner-security-sales-cross-functional-relationships-celaena-powder>.

Evaluating AI tools and avoiding escalations

Which AI tools are legal teams comfortable with business stakeholders using?

The choice of AI tools that legal teams are comfortable with business stakeholders using can vary based on factors like the tool's features, data handling capabilities, and alignment with legal and compliance requirements. That said, most Working Group members preferred enterprise-level AI tools that provide enhanced data protection, security, and compliance features (including no data usage for training models).

Who is responsible for evaluating if a tool does what it says it does? How can teams avoid escalations where the “AI” component does not create any risk, or is just marketing?

Evaluating whether an AI tool performs as advertised and avoiding escalations when the AI component doesn't create risk involves a collaborative effort across different teams within an organization. Here's how responsibility can be distributed:

Business champion

The primary responsibility for evaluating whether an AI tool delivers on its promises often falls on the business champion or the department that intends to use the tool. This team should assess the tool's functionality and whether it meets the specific needs and objectives outlined during the procurement process.

IT and testing

The IT department plays a crucial role in the evaluation process. They should conduct thorough testing of the AI tool before it undergoes procurement review. This testing ensures that the tool functions as expected and is compatible with existing systems and infrastructure.

Procurement review

As part of the procurement process, a cross-functional team should review the AI tool's capabilities and ensure that it aligns with the organization's requirements. This review should involve representatives from legal, security, compliance, and IT to assess potential risks and compliance issues.

In addition, here are some considerations the Working Group said should be weighed when evaluating a tool:

Alignment with business objectives

To avoid escalations, it's essential that the business champion clearly communicates the tool's intended purpose and how it aligns with the organization's business objectives. This prevents situations where the AI component is perceived as merely marketing fluff.

Vendor transparency

Organizations should seek transparency from AI vendors regarding their product's capabilities. Vendors should provide evidence of the tool's performance, including case studies, demonstrations, or trials, to substantiate their claims.

Pilot testing

In some cases, conducting a pilot test of the AI tool in a real-world scenario can provide valuable insights into its actual performance. This can help verify whether the AI component delivers value and functionality.

Documentation and agreements

Ensure that all commitments made by the AI vendor regarding the tool's performance are documented in contractual agreements. This documentation can serve as a basis for accountability if the tool does not meet expectations.

CONCLUSION

By involving multiple teams and ensuring transparency throughout the evaluation process, organizations can reduce the risk of AI tools not living up to their promises and minimize escalations related to misrepresentation or marketing hype. Clear communication, documentation, and testing are key to ensuring that AI tools deliver the expected value.

Pros and cons of AI tool agnosticism

Should the product be built to be agnostic as to the AI tool it leverages?

Whether a product should be built to be agnostic as to the AI tool it leverages depends on various factors, including the specific use case, goals, and technical feasibility. Here are some considerations:

Advantages of AI tool agnosticism

Flexibility

Building a product to be agnostic allows for flexibility in choosing AI tools. This means you can switch out AI models or tools as needed to adapt to changing customer requirements, performance improvements, or cost considerations.

Mitigating vendor lock-in

Avoiding vendor lock-in is a significant benefit of agnosticism. It prevents reliance on a single AI tool or provider, reducing the risk of being dependent on a specific vendor's capabilities.

Future-proofing

An agnostic approach future-proofs your product. As AI technology evolves, you can easily integrate newer, more advanced models or tools without major reengineering.

Performance optimization

Different AI models excel in various tasks and domains. Being agnostic allows you to select the most suitable model for a specific purpose, potentially optimizing performance.

Considerations against AI tool agnosticism

Technical compatibility

Some AI models or tools may have specific technical requirements or limitations that make them difficult to swap out without significant development effort. Compatibility issues can arise when switching.

Optimization

Building a product with AI tool agnosticism may result in a more generic solution that does not take full advantage of the specific capabilities of a particular AI tool. Optimizing for a specific tool might yield better performance.

Use case specificity

In some cases, the product's success depends on the unique capabilities of a specific AI model that is tightly integrated into its functionality. Being agnostic could limit the product's ability to excel in its intended use case.

Development complexity

Implementing AI tool agnosticism can add complexity to the development process, potentially increasing development time and costs.

KEY TAKEAWAY

In summary, the decision to build a product to be agnostic as to the AI tool it leverages should be made on a case-by-case basis, taking into account factors such as the product's goals, customer preferences, technical feasibility, vendor lock-in concerns, and the unique requirements of the use case. Balancing flexibility with optimization is key to making the right choice for your specific product and business objectives.

Establishing process for evaluating tools or use cases

How do you establish an escalation process for evaluating tools or use cases?

While some companies did not pursue this, having a formal escalation process can be valuable in certain situations to ensure that high-risk or complex AI tools or use cases are thoroughly evaluated and meet all necessary requirements. Here are some steps you can take to establish or formalize an escalation process:

Documentation requirement

Specify the documentation and information required for AI tool evaluations. This could include details about the tool's functionality, data usage, security measures, compliance with policies, and potential risks.

Review committee

Establish a cross-functional review committee comprising experts from relevant departments. This committee should include representatives from product development, legal, compliance, security, IT, and business leadership.

Sandbox environment

Consider implementing a sandbox environment for testing high-risk AI tools, especially those that involve customer data. The sandbox allows for controlled testing and evaluation without exposing sensitive data to potential risks.

Security and compliance review

For AI tools that are customer-focused or pose a higher security or compliance risk, conduct separate security and compliance reviews before exposing them to customers. These reviews should identify and address vulnerabilities or compliance issues.

Escalation channels

Establish clear channels for initiating an escalation. Ensure that stakeholders know how to request an escalation when needed, whether through email, a designated platform, or a specific individual. (See below for who the escalation may be to.)

KEY TAKEAWAY

While a collaborative approach can often streamline AI evaluations, having a structured escalation process in place provides a safety net for addressing exceptional cases and ensuring that no high-risk AI tools or use cases are overlooked. It also helps maintain consistency and transparency in the evaluation process, which is crucial for responsible AI adoption.

Ownership of AI tool escalations

If there is an escalation, who is involved in reviewing that escalation and enforcing the policy? What does the escalation process look like?

When an escalation occurs in the context of AI policy enforcement, it's important to have a structured approach involving various stakeholders. Here are key points from your notes:

01

Escalation review team

The process for reviewing escalations typically involves an Escalation Review Team. This team may include representatives from different functions within the organization, such as legal, security, privacy, IT, operations, and engineering. Their role is to assess the situation, determine if policy violations occurred, and recommend appropriate actions.

02

Minimizing review time

To minimize the time spent reviewing AI tools, organizations can implement efficient processes. One approach is to have an AI-specific intake form as part of vendor management. This form ensures that AI-related considerations are addressed as part of the regular vendor review process, rather than requiring a separate exception policy.

03

AI committee

Organizations may establish an AI Committee as part of their internal governance structure. This committee is responsible for overseeing AI policy implementation and enforcement. It ensures cross-functional support and coordination among different departments involved in AI usage.

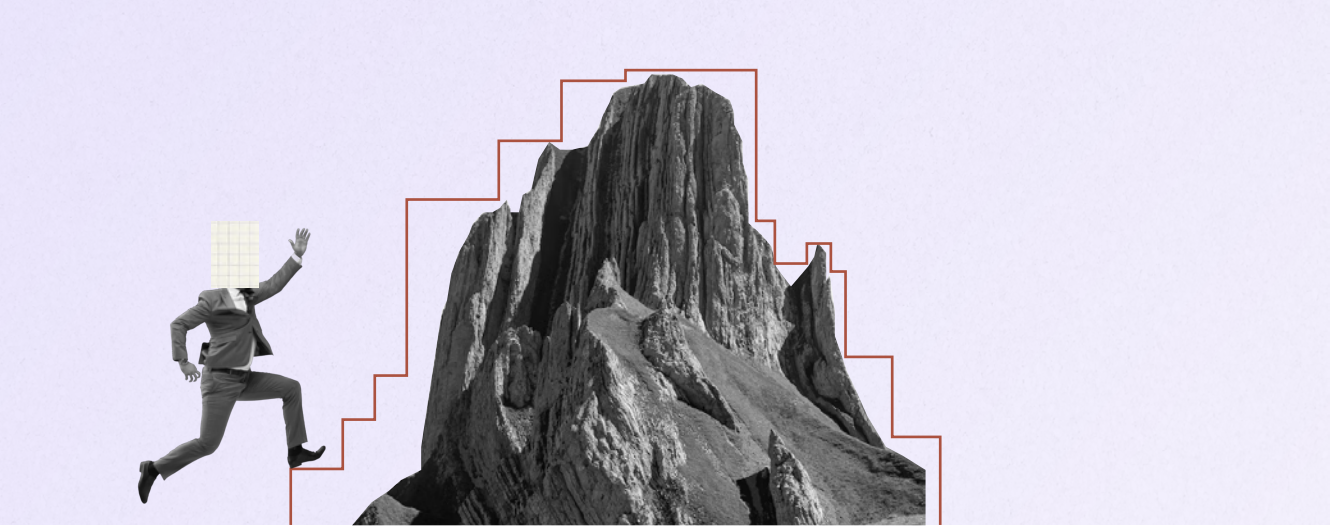

04

Leveraging existing processes

Whenever possible, leverage existing processes for product review and compliance. For instance, Data Protection Impact Assessments (DPIA) can be used if customer data is implicated. This ensures that AI tools are reviewed within the context of existing compliance frameworks.

05

Embedded collaboration

Collaboration with product management teams is essential. Legal, security, privacy, and other relevant functions should work closely with product development teams. This collaboration ensures that AI principles of transparency, fairness, and lack of bias are integrated into the product design process.

KEY TAKEAWAY

By having a structured approach that involves the right stakeholders and integrates AI considerations into existing processes, organizations can effectively handle escalations, review AI tools efficiently, and enforce AI policies in a way that aligns with their overall governance framework.

Enforcing AI policies

If a particular tool or use case is blocked, how is that enforced?

Enforcing policies when a particular tool or use case is blocked can be challenging. Here are key considerations regarding enforcement and dealing with circumventions:

01

Monitoring for web-based AI tools

Monitoring web-based AI tools can be challenging but not impossible. Some organizations use web monitoring tools and traffic analysis to identify unauthorized tool usage. However, monitoring should be conducted in accordance with privacy and legal considerations.

02

Educational approach

Part of enforcement involves educating employees about the policies and their rationale. Many circumventions occur due to a lack of awareness or understanding. Regular training and communication can help mitigate this.

03

Transition period

When there is a policy change that affects the current use of a tool, it's important to address this during the rollout of the policy change. Clear communication with specific team leaders can help facilitate a smooth transition.

04

Enforcement challenges

Enforcing policies, especially across a large organization, can be challenging. It's important to acknowledge that enforcement may not always be perfect. Companies need to balance enforcement with practicality and fairness.

KEY TAKEAWAY

In summary, policy enforcement involves a combination of monitoring, education, communication, and appropriate consequences for violations. While it can be challenging, having clear policies and a commitment to enforcing them helps organizations maintain security, compliance, and responsible AI usage.

Answering customer queries around AI policies

What about when customers ask about your AI policies?

When customers ask about your AI policies, it's crucial to provide clear and transparent responses to address their concerns and build trust. Here are some strategies and considerations the Working Group uses:

Nutrition labels

Some companies are developing "nutrition labels" or similar disclosures to provide customers with easily understandable information about their AI usage. These labels can include information about data usage, AI decision-making, and privacy practices.

FAQ section

Maintain a well-structured FAQ section on your website or in customer communication materials. This section should address common questions related to your AI policies, including changes in terms and conditions.

Expecting follow-up questions

When you release an external policy, be prepared for follow-up questions from customers. Anticipate the types of questions they might ask based on your policy and proactively address them.

Customer support and education

Train your customer support teams to provide accurate and helpful information about your AI policies. Ensure they are well-versed in the nuances of the policies and can assist customers effectively.

Proactive communication

When you release a new AI feature or tool, consider proactively communicating how it works, what data it uses, and how it aligns with your AI policies. This proactive approach can alleviate customer concerns.

Transparency

Transparency is key when communicating AI policies to customers. Clearly state your policies, their purpose, and how they impact your products or services.

KEY TAKEAWAY

In summary, when customers inquire about your AI policies, you should prioritize transparency, proactive communication, and providing clear and easily accessible information. This approach helps build trust and ensures that your customers are well-informed about how AI is used in your products or services.

Meet the Working Group

Ken Priore
AGC - Product and Strategic Partnerships, DocuSign

Alesya Nasimova
Director, AGC, Product & Privacy, Brex

Meg Pirnie Kammerud
Former GC, Dragos

Kelly Trimble
Sr. Director of Legal, Seismic

Vicky LeVay
Sr. Director of Compliance, FloQast

Erik Graham-Smith
General Counsel, FloQast

Caitlin Hanson
Sr. Manager of Data Governance, FloQast

About SpotDraft

SpotDraft is a leading contract lifecycle management (CLM) software that provides an end-to-end solution for streamlining contracts. Businesses across the globe — Apollo.io, Airbnb, Notion, SimpliSafe, TuneIn, and Strava to name a few — use SpotDraft to manage their contracts efficiently.

Our AI-driven solution empowers in-house legal teams to save up to 10 hours per week and helps them close deals 25% faster, freeing them to focus on more strategic projects.

Founded in 2017 by a Harvard law graduate and two Carnegie Mellon computer engineers, SpotDraft started with the vision of building a centralized, compliant, and digital system of record for in-house legal teams — a single platform to help them collaborate with both internal & external teams to review, manage and close contracts faster.

For more information, please visit https://www.spotdraft.com/