Chapter 29

The Cheat Sheet

"By far, the greatest danger of Artificial Intelligence is that people conclude too early that they understand it." —Eliezer Yudkowsky

As mentioned in the previous chapter, we have assembled all the recommendations from previous chapters into a Cheat Sheet below for easy reference. Please don't fall into the trap of thinking of this as a checklist: a rigid to-do list you have to follow in order.

Instead, think of it more like a helpful guide that makes sure you don't miss anything important. It's not about being strict; it's about being thorough. You might need to go back and check things again, change your approach, or even add new things to the list as you go. The main idea is to carefully consider each item and decide how it applies to what you're doing at the moment and make informed choices.

I'm sure that many of you are thinking to yourself, "Hey, I wonder if the authors just put the book into an LLM and had it generate this list?" The answer to that question is, "Yes, that's exactly what we did.' Thanks AI!

Start with a Business Problem

Identify a specific problem: What business challenge are you trying to address with AI? (e.g., increase sales, reduce customer churn, automate a process)

Define how AI can help: How can AI specifically address the problem you've identified? (e.g., personalize recommendations, automate data entry, predict maintenance needs)

Establish SMART goals: Set Specific, Measurable, Achievable, Relevant, and Time-bound goals to measure the success of your AI implementation.

Evaluate multiple solutions: Once you have a problem and goals, explore AI solutions that align with your needs. Don't just go with the most popular option.

Use Design Thinking

Empathize: Conduct thorough research to understand the needs, challenges, and perspectives of the end-users.

Define: Clearly articulate the problem the AI is addressing from the user's point of view.

Ideate: Brainstorm a wide range of potential solutions, focusing on quantity and diverse perspectives.

Prototype: Create experimental versions of the solution quickly and cost-effectively.

Test: Get feedback from real users on the prototypes and iterate on the design based on their input.

Take Big Picture View

Analyze Existing Assets: Identify your organization's strengths, data, tools, and expertise to see how they can be leveraged for AI solutions.

Listen to Your Ecosystem: Conduct interviews, analyze competitors, and map the broader context to uncover unmet needs.

Reframe Questions: Regularly revisit the problem definition and explore alternative approaches.

Assess Data and Feasibility: Conduct thorough feasibility analyses, considering data availability, resources, and potential constraints.

Use Modular Architecture

Embrace modularity: Design your AI system with independent, interchangeable modules (like Lego bricks)

Adhere to core design principles: Adhere to separation of concerns, encapsulation, high cohesion, low coupling, and standardized interfaces.

Consider Domain-Driven Design (DDD): This methodology helps identify modules and address design principles effectively.

Prioritize software development practices: Implement version control, documentation, backward compatibility, automation, and refactoring.

Address AI-specific needs: Incorporate modules for knowledge management, safety, feedback, etc.

Don't forget MLOps: Include modules for model management, deployment, and monitoring.

Engage experienced technology resources: Don't rely solely on "turnkey" solutions; involve skilled architects for proper design and implementation.

Invest in robust infrastructure: Ensure your infrastructure supports the complexity of interconnected modules.

Use Modern Data Architecture

Assess and identify data gaps: Evaluate your current data infrastructure and pinpoint areas for improvement.

Invest in modernization: Consider Lambda architecture (dual-layered for batch and real-time processing) and Kappa architecture (unified stream processing) as solutions to modernize your data infrastructure. Plan and execute data migration strategies that maintain data quality and minimize disruptions.

Invest in Automation

Identify repetitive tasks: Pinpoint the routine tasks that can be automated to save time and reduce errors.

Assess ROI: Prioritize automation in areas where it will deliver the most significant value.

Automate data collection: Don't rely on individuals performing manual steps to get good data.

Automate gradually: Start with one process, assess its impact, and gradually expand automation to other areas.

Cut non-value-added steps: Eliminate any unnecessary steps in your processes that don't add value.

Lean into automation challenges: Be prepared for high initial costs, integration complexities, and potential employee resistance to change.

Link automation to business goals: Align automation efforts with your overall business objectives, such as improving customer experience, achieving operational efficiency, and driving innovation.

Focus on Monitoring and Visibility

Implement comprehensive monitoring: Monitor infrastructure health, software performance, AI model outputs (accuracy, bias, etc.), and user interactions (adoption, satisfaction, etc.).

Establish robust logging practices: Log all AI interactions (inputs and outputs) to facilitate analysis and improvement.

Use automation and AI for monitoring: Utilize automated testing, error tracking, and AI-powered monitoring tools to enhance efficiency and proactively address issues.

Prioritize user feedback: Actively monitor and analyze user interactions to identify areas for improvement and ensure the AI system delivers real business value.

Design for Scale

Define your scaling needs: How many users? Where are they located? How often will they interact? What are their performance expectations?

Choose the right deployment architecture: Data center, local server, or on-device? Consider factors like connectivity, data access, privacy, and performance.

Implement technical best practices for scalability: Use elastic infrastructure, shared tools, asynchronous processing, caching, and rate limiting.

Establish a continuous improvement process: Monitor performance, identify bottlenecks, and optimize your AI system over time.

Get Ownership Right

Consider a Chief Data Officer (CDO): Appoint a CDO or equivalent role to oversee data governance, quality, and strategy across the organization. Provide the CDO with the authority, resources, and cross-functional team necessary to effectively manage and integrate data for AI initiatives.

Conduct a Data Audit: Assess the current state of data management in your organization, including data sources, quality, accessibility, and ownership.

Develop a Data Strategy: Create a comprehensive data strategy that outlines data governance, standardization, and integration processes.

Foster a Data-Driven Culture: Encourage collaboration and knowledge sharing around data across different departments.

Create Cross-Functional Teams

Get leadership buy-in: Secure top-down support to break down organizational silos and prioritize cross-functional collaboration.

Identify 'fracture planes': Start by forming small, self-contained teams around specific AI projects or features.

Run a pilot project: Test the cross-functional approach on a smaller scale, gather learnings, and iterate before wider implementation.

Align incentives: Structure rewards and metrics to encourage collaboration and shared success across team members.

Establish communities of practice: Create groups where team members from different functions can share knowledge, solve common challenges, and improve efficiency.

Consider a platform approach: Invest in building reusable platforms that provide foundational building blocks for AI solutions, reducing dependencies and accelerating development.

Address the Skills Gap

Assess your needs: Identify the specific skills required for your AI project.

Evaluate your budget: Determine how much you can afford to spend on acquiring talent.

Develop a clear strategy: Outline your plan for acquiring and developing the necessary skills.

Assess your existing team honestly: Identify any existing skills that can be leveraged and individuals who can be trained.

Consider your timeline: Determine how quickly you need results and how long the project will last.

Weigh the pros and cons: Carefully consider the advantages and disadvantages of hiring, training, and using experts (consultants)

Foster a culture of learning: Encourage continuous learning and development within your team.

Use a Modern Approach to Governance

Re-evaluate your current governance processes: Identify areas where they are hindering innovation or creating unnecessary bottlenecks.

Adopt a risk-based approach to governance: Tailor oversight and controls to the specific risks associated with each AI project.

Automate governance tasks wherever possible: This can include using scripts for continuous monitoring and validation and leveraging AI for tasks like data lineage tracking.

Start small: Gain experience and build momentum before scaling to large-scale governance programs.

Educate your governance team about AI: Help them understand the unique challenges and opportunities associated with this technology.

Foster a culture of collaboration between governance and product teams: Encourage open communication and knowledge sharing.

Focus on Organizational Change Management

Communicate a Clear Vision: Clearly articulate the goals and benefits of AI initiatives to ensure buy-in at all levels.

Allocate Resources for Change Management: Prioritize training, support, and incentives alongside technical investment.

Set Realistic Goals and Incentives: Define measurable, achievable objectives while avoiding rigid or overly simplistic KPIs.

Celebrate Early Wins: Showcase initial successes to build trust, motivation, and support for further changes.

Remove Barriers: Identify and address obstacles to change, including outdated processes and resistant mindsets.

Anchor Changes in Culture: Integrate new practices and behaviors into the organization's DNA for lasting impact.

Understand the Need for Iteration

Monitor Disruptors: Regularly analyze market trends and competitors to anticipate and respond to potential disruptions.

Adopt Iterative Development: Use agile, incremental approaches to quickly test, deploy, and refine AI systems in response to feedback and changing conditions.

Combat Drift: Implement robust monitoring to detect data drift, concept drift, and feature drift, and establish processes for timely model retraining and adaptation.

Build Data-Driven Processes: Track metrics to measure AI performance, identify areas for improvement, and validate ROI. Use insights to guide decisions.

Adapt to Customer Behaviors: Use user feedback and iterative updates to align products with evolving customer preferences and expectations.

Plan for Technological Advancements: Stay updated on emerging technologies and frameworks, and incorporate them to maintain competitiveness.

Act with Urgency: Recognize the accelerating pace of change and prioritize speed and flexibility in decision-making and execution.

Iterate, Iterate, Iterate

Define Clear Goals for Experiments: Identify the specific problem the AI solution will address and the metrics to measure success.

Develop an MVP: Build a basic, functional AI model to test initial assumptions and gather feedback.

Embrace Experimentation and A/B Testing: Continuously test variations of AI models or features to optimize performance and outcomes.

Incorporate Prompt Engineering in the Process: Refine prompts and interactions to maximize the efficiency and accuracy of AI solutions like LLMs.

Focus on Incremental Gains: Aim for consistent small improvements over time to achieve compounding benefits.

Gather Feedback Regularly: Engage stakeholders and users to refine AI solutions based on their insights and needs.

Understand the Source of Errors

Address bad data Issues: Vet training datasets to ensure they are accurate, representative and comprehensive.

Ensure Robust Models: Regularly test and iterate to optimize the model's architecture and training approach.

Understand Sources of Bias: Analyze datasets for implicit biases and adjust them accordingly. Mitigate human bias by using cross-cultural and cross-disciplinary input during training.

Acknowledge a Non-Deterministic Approach: Accept that some errors are due to randomness and cannot be fully eliminated. Implement safeguards to detect and correct nonsensical outputs.

Leverage Existing Solutions: For most enterprises, avoid building models from scratch unless absolutely necessary. Evaluate models with enhancements like Retrieval Augmented Generation (RAG) for domain-specific accuracy.

Understand Security Challenges

Implement traditional security and privacy measures: Best practices such as edge protection, secure coding practices, encryption of user data and requests, least privilege access are critical. Consider using AI tools to identify weaknesses.

Protect Against Poisoned Training Data: Vet training datasets to ensure accuracy and prevent intentional poisoning. Monitor data sources for anomalies and malicious entries.

Guard Against Prompt Injection: Develop and enforce robust guardrails to reject malicious prompts. Monitor and log interactions to detect and respond to prompt injection attempts. Regularly update AI systems to address vulnerabilities exploited by known injection techniques.

Prevent Data Exfiltration: Avoid including sensitive or private information in training datasets. Implement monitoring systems to detect and mitigate data extraction attempts.

Focus on Continuous Security Improvement: Regularly review and update guardrails and security measures as new threats emerge. Stay informed about AI security trends and adapt systems accordingly.

Address the Trust and Safety Issues Holistically

- **Assess potential impacts:** Consider the impact of your AI system on people, business operations, regulatory compliance, and the physical world.
- **Define your risk tolerance:** Identify absolute no-go zones, develop strategies to mitigate risks, and clearly document the acceptance of residual risks.

- **Engage diverse stakeholders:** Be sure to include diverse stakeholders in risk assessment and decision-making. Consider engaging an external review board.
- **Implement safeguards:** Focus holistically throughout the AI system lifecycle, from data engineering to model development, input and output process, and deployment.
- **Consider adopting NIST's AI RMF:** Use NIST AI RMF or a similar framework to guide your efforts to build trustworthy and safe AIs.
- **Continuously monitor and iterate:** Improve your AI system based on real-world feedback and data.