AI in Marketing

Prof. Zillur Rahman

Department of Management Studies

Indian Institute of Technology, Roorkee

Week 8

Lec 38-Implementation of AI by Product Managers

Welcome to this NPTEL online certification course on artificial intelligence and marketing. And now we will talk about module 38. So as you can see from this slide, we are talking about the implementation of AI by product managers. So this is module 38 and we are in chapter 6. Now these are the things that we will cover in this module. So we will start with the guidelines to be followed for AI infusion based on the key phases of product development cycle to be followed by the questions to be answered in each key phase of the product development life cycle along with use case examples.

So let us start with the first one that is implementation of AI by product managers. As AI gets even more deeply embedded in society and people's everyday life, business leaders, data scientists and engineers have a responsibility to guide AI product development towards responsible outcomes. So this is very important word responsible. Product leaders typically oversee the development of one or more AI products.

Their teams have simultaneous responsibility to advance an organization's values and business objectives while guiding product development towards the responsible outcomes. So now they have to do look after the organization values, business objectives and responsible outcomes. Their place at the center of innovation and product development positions these leaders as essential actors in operationalizing responsible AI. The range of AI products that oversee is broad with the underlying algorithm ranging from novel complex models to pre-trained models built by others. Organization leaders focused on product development and delivery should remember to take sufficient measures to scale responsible AI across their organizations.

Ultimately responsible AI is about driving a cultural shift within an organization. Truly operationalizing responsible AI requires broad changes across leadership, governance, processes and talent. We will discuss some guidelines focusing on product development and delivery processes to provide a formalized way to stimulate the challenging conversations and critical thinking necessary to anticipate and mitigate the risk of AI

systems. These guidelines will help product leaders create deeper synergy between performance, organization goals and values. Throughout the product development life cycle, leaders should keep in mind their organization's values and AI's ethical principles to the extent their organization has formalized them.

Now these are some of the guidelines. Assess and prepare, design, build and document and validate and support. The guidelines have been organized according to the key phases of the product development life cycle while recognizing that AI product development often cycles through these phases iteratively. So, the first is to assess and prepare. Assess merit of developing the product considering organization values and business objectives.

Assemble teams reflecting diverse perspectives and with clearly defined roles and responsibilities. Assess potential product impact by including inputs from domain experts and potentially impacted groups. So, the first one is to assess merit of developing the product considering organizational values and business objectives. So, now let us understand what needs to be done in this first step. So, what are the primary use cases and benefits of the proposed products which users are explicitly out of scope? What is the desired business outcome for this product? How will the business impact be measured? How might the operations of the product and use of its output for business decisions, enroach on core organizational values? The second is assess merit of developing the product considering organizational values and business objectives.

So, use case example is automating loan approval decisions. A financial institution would like to automate its loan application approval process with the goal of increasing efficiency and reaching new customer basis. The product leader should immediately recognize that because of the AI's product's potential to impact individual's economic well-being and quality of life, they must proceed with care. The system should be designed to predict an applicant's earning potential over the lifetime of the loan under review and thus the likelihood of repayment. The business impact should be measured with the key performance indicators.

The first is default rate against a historical baseline. The second is increase in overall loan volume and the third is speed of approval or rejection decisions. There may be race, gender or other biases reflected in the historic data. Exploratory data analysis will need to include steps to bias assess the data set composition so that appropriate mitigations can be implemented. There is a significant risk that proxy variables are race, gender and other protected categories could reinforce biases and impact outcomes.

By letting company values around fairness and the law, the product team will have to

proactively engage on these topics. The next is to assess and prepare. Assemble team reflecting diverse perspective and with clearly defined roles and responsibilities. Do we have a diverse that is gender, age, ethnicity, multi-disciplinary team with a range of functional expertise? What prospective or expertise are missing and how can we introduce them, including sources outside the team or organization? Is the team is structured so that domain experts can impact relevant design choices? The next is assemble team reflecting diverse perspective and with clearly defined roles and responsibilities. The product leader should begin to build the team based on initial assessment that the team should include data scientist with experience applying fairness tools to machine learning models, a lawyer with expertise in the Equal Credit Opportunity Act to make sure the team fully understands the regulatory environment.

Two loan officers with extensive experience helping a variety of customers navigate the application process. This will enable the team to design the product to augment the decision making capabilities of loan officers and integrate their feedback in real time. A product leader should know that since the system will deal with sensitive personally identifiable information, system security and privacy are critical. We must gain approval to bring on an external consultant with cyber security and AI expertise to provide guidance on optimal system architecture, data storage and differential privacy across the build. A user researcher to ensure customer needs are at the forefront from idea inception to final product so that customer's understanding and satisfaction can be reconciled with business goals.

A designer to improve product usability and accessibility across a diverse user base. A diverse team reflecting a variety of backgrounds and lived experience. The third is to assess and prepare. Assess potential product impact by including input from domain experts and potentially impacted groups. What are the foreseeable models of failure for this product? What edge scenarios could lead to failure and harm? What are the societal and environmental implications of foreseeable product failure, misuse or malicious attack? What are the product's potential unplanned uses? What external SMEs or groups could provide input in forming design choices that would reduce the risk of negative societal impact and harm to individuals directly or indirectly affected by the AI product? Assess potential product impact by including input from domain expert and potentially impacted groups.

A preliminary discussion should take place where the loan officers should share their experiences engaging with the financial institutions current customer base. As they detail interactions with various customers over the years, the product leader should realize that the AI may engage new customer demographics who may have needs and expectations that differ from those captured in historical data. The team should engage an economist

with expertise in banking relationships within the demographic groups to whom the bank might expand its activities. The team should also work with a vendor to deploy a survey for potential borrowers in the new customer communities to better understand how great access to credit might impact earning potential in ways both consistent and inconsistent with communities overrepresented in the historical data and thus the likelihood of repayment. The loan officer should also note that even with the new customer base, one thing is likely to remain constant.

That is having a loan application declined in an unpleasant and potentially painful experience. Having a loan application declined is an unpleasant and potentially painful experience, one that can be delivered more respectfully by a skilled and experienced professional. Based on this insight, the product team should decide that all rejection decisions shall be communicated to applicants by a loan officer. Design, build and document. So, the first step here is to evaluate data and system outcomes to minimize the risk of fairness harms.

The second step is to design AI products to mitigate the potential negative impact on society and the environment. The third step is to incorporate features to enable human control. The fourth is to take measures to safeguard data and AI products. And the fifth is to document throughout the development life cycles to enable transparency. So, let us start with the first step that is evaluate data and system outcomes to minimize the risk of fairness harms.

What fairness metrics? For example, statistical parity, equalized odds, test and chipping criteria will be used. How will the product team validate that the training data includes data collected via APIs capture the different groups and types of people likely to be impacted by the system's output? How will the product team measures whether the AI product outcomes are consistent with the chosen objective, that is avoid target leakage. Fairness metrics, tests and chipping criteria across a wide variety of potentially impacted groups or intersection of groups. How will the product team ensure continued adherence to fairness metrics, tests and criteria post deployment? Now let us look at a user case example of demand forecasting for a fashion retailer. A fashion retail chain hopes to transform its in-store inventory management with AI.

Using historical sales data, the company wants to optimize the amount of inventory held in stores so that it can maximize sales per square foot. For this, the product leaders should clarify the business objectives and assemble a team. A team should now shift its focus to the design, build and document stage of work. The product leader should convene the product team for a discussion around fairness. A poorly designed product could lead to service discrepancies across different demographics within the customer

base, which would violate the company's values.

Data scientists on the team should note that customer's feedback for the retail chain may vary dramatically across neighborhoods in large metropolitan areas, with stockouts more common in some locations than others. Pivoting from this, the team should align on the following approach. A team should check for parity of service level based on stockout reported in its historical data. Cross-referencing the results with census data will allow team to assess whether stockouts reported are correlated with certain neighborhoods and perhaps demographic groups. To prevent unacceptably low service levels in specific situations, the product team should establish a minimum inventory level for each SKU at each location to ensure minimum service levels that will avoid the possibility that stockouts would impact certain demographic groups more than others.

The second is to design AI products to mitigate the potential negative impacts on society and the environment. Its negative impact that is the system's failures, unplanned use, abuse, attack or simply side effects of normal use is possible. What design processes that is human centered design and choices can reduce, mitigate or control them? What design choices will help minimize the adverse environmental impact on the product output and related decisions? What design choices are critical to ensuring proper use, legitimate and transparent data collection and respect for user privacy? The team's sustainability experts should ask the group if system usage might have some second-order environmental effects. Optimizing inventory levels should maximize profit at the store level, but smaller and more frequent inventory replenishments would rely on greater air and ground cargo traffic. At the country level, the environmental consequence of that cargo traffic could be significant.

Furthermore, since significant amount of excess inventory would need to be returned by stores, optimized inventory levels could create additional cargo traffic. Using internal logistic data, the product teams can build a feature that highlights trade-off between inventory levels and transportation emission at the store and region level. The third is to incorporate features to enable human control. How is the product team designing the product to empower humans by augmenting their decision making, streamlining tasks and otherwise making them more effective? Which decisions or functions requires human oversight as a critical component of the AI product? The product mechanism will support end-user comprehension of the system to enable continuous audit, monitoring and human intervention. What product features allow users to customize AI performance? What channels will the product utilize to collect live feedback? What product features will ensure inclusive experiences for people with disabilities? Next we will look at incorporating features to enable human control.

A specialist in human-centered design should understand that the retailer often learns about new trends from its front-line workers, particularly store managers. The historical data has limited predictive power for trend spotting. By enabling human control over the inventory system and augmenting store managers' decision-making abilities, stores should make on-the-fly adjustments to match changing consumer preferences. Based on further exploration, the team should design a feedback mechanism through which store managers can indicate emerging trends and newly popular products at a specific location, thus allowing for the pooling of insights across the country to help spot trends and adjust inventory levels accordingly. The fourth is to take measures to safeguard data and AI product.

How will the product team prioritize data privacy across product designs with respect to cloud infrastructure, encryption, anonymization, and access control? How will the product team make sure the product does not inadvertently disclose sensitive or private information during use, for example, indirectly inferring user location or behavior? What methods will the team use to identify and address security vulnerabilities, including those such as data poisoning, that are unique to AI products? The fifth is document throughout the development lifecycle to enable transparency. What standards and processes are in place to ensure the entire product team consistently document design and development choices, rationals, and assumptions? How can the product team best keep track of data sources and their authorized users? What types of models, tools, or techniques will be used to document product behavior? Next comes validate and support. Validate product performance and test for unplanned failures as well as foreseeable misuse unique to AI products. Communicate design choices, performance, limitation, and safety risk to end users. So, let us look at the first one in more detail.

Validate product performance and test for unplanned failures as well as foreseeable misuse unique of unit to AI product. How will the product team validate the AI product's performance against agreed upon business KPIs and metrics, tests, and criteria? How will the product team validate the AI product's performance against technical standards and benchmarks? What are the target environment and conditions under which this product can be expected to function properly and safely? How will the system be tested and evaluated for safe and effective operations, that is graceful failure in both business as usual and edge case scenarios? What are the mechanisms for continuously monitoring business, technical, and fairness performance as well as for modeling drift post-production? Validate product performance and test for unplanned failures as well as foreseeable misuse unique to AI product. Use case example of predictive lead times for a manufacturer. An industrial goods manufacturer is experiencing repeated delays in the delivery of parts purchased from suppliers. This has caused disruptions to manufacturing schedule and ultimately, led deliveries to customers and the potential to

damage key relationships.

The manufacturer has decided to develop an AI product to estimate lead time for supplier procured component based on historical data. Providing purchasing managers with an early warning of potential delays will enable them to proactively engage with suppliers and adjust the manufacturing schedule accordingly to avoid missed delivery deadlines. This AI product might change the way the company engages with suppliers, customers, and workers. Although the product is expected to have positive impacts across operations and customer relations, erroneous lead time estimates could also create more work for purchase, ring managers and further damage, supplier relationships and negatively impacting the bottom line. This could have significant and far-reaching consequences for the business.

Furthermore, because assembly is a labor-intensive process requiring specialized skills and safety criteria at different stages, ad hoc adjustments to the manufacturing schedule could place factory employees at risk. The product leaders along with the product team should reach alignment on the approach to validating product performance and robustness to unplanned failures. The following should be looked into. Testing the system's output against historical supplier promise dates to determine the model's ability to flag potentially delayed shipment before it's too late. There are a wide range of operational scenarios including varied suppliers and component types as well as edge scenarios, that is parts and suppliers not found in the historical data.

For each scenario tested that requires adjustments to the manufacturing schedule, developing associated work schedules to be validated by factory floor leadership for feasibility and safety. For example, during the COVID scenario, the data scientist team should note that factors related to the ongoing COVID-19 pandemic such as limited trucking capacity and economic shutdowns in certain states would not be captured in historical data but could impact lead times. To capture these insights, the product team should integrate COVID case loads in geographies proximate to suppliers such as the way to capture the pandemic's potential impact on the manufacturer. In conversation with factory floor leaderships, the product team learned how scheduling has recently changed to minimize risk of COVID-19 exposure and that any changes to the manufacturing schedule would have to be consistent with the new scheduling policies. The second is to communicate design choices, performance, limitation and safety risk to end users.

What information and instruction should the product team provide to end users to ensure safe and reliable use? How will the product team make sure end users understand the primary use case, underlying assumptions and limitations of the product? How will the

output of the systems be communicated in a way that helps end users understand how the system works? Now for communicating design choices, performance, limitation and safety risk to end users, we will use the use case example of predictive lead times for a manufacturer. The team's UX lead should push the product team to consider how best to augment the purchasing manager's current decision making process. The team should align on the following steps. The first is to calculate confidence interval alongside estimated lead times to enable purchasing managers to responsibly leverage system outputs. The second is to design the dashboard to sort components for a certain product by a predicted delay, focusing the end users' attention on prioritizing key products and engaging problem suppliers early.

Make sure the system never automatically updates customers on delivery dates, a task that program managers should themselves continue to perform. Build on additional supplier level as opposed to component level dashboard to equip the company for strategic engagement with suppliers that consistently struggle to deliver parts on time. Output a draft work schedule based on proposed changes that will support purchasing managers' decision making regarding adjustments to the manufacturing schedule. The factory floor leadership will have to sign off on the revised work schedule to ensure adherence to workers' safety standards. The team, if earlier identified and engaged with the products end users, should conduct several meetings and trainings to ensure that these features will be leveraged effectively.

It should also design a modular training on tool use to be integrated into future onboarding programs for purchasing managers. So, to conclude, in this module we have continued explaining the other two key phases of guidelines for product development cycle. We have also discussed the use case for each phase. And these are the five references from which the material for this module was taken. Thank you.