

AI in Product Management
Prof. Zillur Rahman
Department of Management Studies
Indian Institute of Technology, Roorkee

Lecture- 9
AI in Customer Sentiment Analysis

Welcome to this NPTEL online certification course on Artificial Intelligence in Product Management. So now we are on module 9, and we are talking about AI in customer sentiment analysis. To give an overview, in this module, we will try to understand what customer sentiment analysis is and its various types. Then we will understand the role of AI in customer sentiment analysis and how it actually works. Next, we will give an overview of the various AI tools that can be used for customer sentiment analysis and then discuss various cases of AI-enabled sentiment analysis across industry sectors.

So, the introduction to customer sentiment analysis. Customer sentiment analysis involves collecting, analyzing, and leveraging data to understand how customers feel. It can help you understand users' perception toward a particular feature, product, or even an industry. Typically, sentiment analysis involves going through online forums and social media networks to scan people's thoughts and understand the sentiments expressed in their posts or comments. So, sentiment analysis can have different uses.

For example, you can use it to understand whether users perceive your product as negative, neutral, or positive. Maybe you want to find out which emotions resonate with your product, or you are looking for feedback on a specific aspect of your product, such as customer service. So, customer sentiment analysis uses machine learning algorithms like NLP, that is, natural language processing, to classify emotions into different polarities, such as positive, negative, or neutral.

NLP is a branch of artificial intelligence that combines computer science and linguistics to analyze and interpret human language. It can process a large amount of customer data in a relatively short time. Usually, a simple sentiment analysis process involves these steps. The first is to collect textual data from different sources, remove irrelevant information such as words and special characters, tokenize the text into singular words or phrases, and extract relevant attributes, phrases, and lexicons. Then group attributes into different sentiment categories.

The algorithm then assigns a customer sentiment score to attributes based on polarity and magnitude—the intensity of the sentiment—which determines the tone and emotions expressed. Now, let us look at the types of sentiment analysis. Sentiment analysis algorithms use natural language processing to determine emotions and customer satisfaction levels. As there is a wealth of data available in the form of customer reviews, feedback, and comments, it has now become easier to gauge the sentiment of the text and your buyers. That being said, let's delve into the four most popular methods of sentiment analysis.

The first is fine-grained analysis. Fine-grained sentiment analysis goes beyond positive or negative. You go a step further and divide emotions into five polarities. Very positive, positive, neutral, negative, and very negative. By using text analysis, you can assign a customer sentiment score to each review, say from 1 (very negative) to 10 (very positive).

This approach uncovers nouns details and offers a more comprehensive perspective. which can help you pinpoint areas for improvement in your products or services. The second is emotion detection. Forming an emotional connection really works for some brands. In fact, 71% of buyers will recommend a brand based on their emotional attachment to it.

Emotion detection models can help you decipher both positive and negative sentiments like happiness, anger and sadness. It is mainly done using machine learning and or lexicons. Lexicons are words or phrases associated with a particular topic which makes it easier to segregate terms according to their sentiment. However, lexicons become difficult to use when you consider the ways in which people express themselves. Like a user might say, your product rocks and you won't know where to place it.

But artificial intelligence combined with machine learning technologies can solve this problem. The third is aspects. based analysis. Aspect-based analysis is sort of like fine-grained sentiment analysis. Only difference is that it focuses on specific topics and subjects that people are talking about.

Businesses often use it to analyze customer review data by mapping sentiments to different features of a product or service. Here is an example. In a smartphone review, a brand might analyze the sentiment for camera quality, interface, and design. Suppose a user says the design is okay-ish, The camera quality is super bad.

When this feedback is put through a sentiment analysis software, it will group okay-ish and super bad as negative customer sentiment. The next is intent analysis. Consumer intent is the thought that drives your buyers to make a purchase decision. Intent analysis tells you whether your prospects intend to endorse your brand and its products. It gives you a better idea of what your buyers want.

As well as the purpose behind their actions. For example, what do they wish to do when they first visit your website? Or why do they constantly contact your customer service team? And when they show the intent to make a purchase, you can immediately reach out to them. This saves a significant amount of time and resources.

Now we will look at the role of AI in sentiment analysis. AI tools have revolutionized sentiment analysis by automating the process and enhancing its accuracy. By employing machine learning algorithms and NLP techniques, these tools can efficiently process large volumes of text data and extract valuable insights regarding public sentiment. Here are some ways AI tools can benefit product managers in their quest to gauge public sentiment. The first is efficient data collection.

AI tools can swiftly gather vast amounts of textual data from diverse sources such as social media platforms, news articles, blogs, and customer reviews. This data can provide product managers with a comprehensive view of public sentiment. Real-time analysis with AI-powered sentiment analysis tools. Product managers can analyze sentiment in real-time. This enables them to stay up to date with evolving public opinion and take immediate action when necessary.

The next is accurate sentiment classification. AI models trained on large datasets can accurately classify sentiment by considering various contextual factors. They can identify subtle nuances, sarcasm, and sentiment polarity, providing a more nuanced understanding of public sentiment. The next is Another capability is sentiment trend analysis.

AI tools can detect patterns and trends in sentiments over time. By visualizing sentiment trends, product managers can track sentiment fluctuations and identify critical milestones or events that impact public perception. Sentiment-based insight. AI tools can generate actionable insights from sentiment analysis. These insights can help product managers identify strengths and weaknesses and

anticipate potential issues and devise strategies to improve public sentiment and project outcomes. The next thing that can be done with AI is sentiment comparison. AI tools can

compare sentiment across different products or competitors. These comparisons can provide valuable benchmarking information, helping product managers evaluate their product's performance relative to others and make data-driven decisions. Now, we will look at how AI in sentiment analysis works. Integrating AI into sentiment analysis involves various components to streamline the analysis of textual data, generate insights, and support decision-making. It goes beyond traditional sentiment analysis processes. By incorporating powerful large language models and connecting them with an organization's unique knowledge base.

This method facilitates a deep understanding of sentiment by analyzing context, nuances, and underlying emotions within text data, empowering businesses to make informed decisions based on comprehensive sentiment analysis. This LLM-based architecture utilizes various components to improve the sentiment analysis process, enabling businesses to fully comprehend sentiment trends and utilize this input to drive strategic initiatives and enhance overall business performance. The next slide shows a detailed breakdown of this process. So, this is how AI works in sentiment analysis. So, here we have the data sources: social media, customer reviews, survey and feedback forms, call center transcripts, emails, and support tickets.

Then come data pipelines, embedding models, and vector databases, and these go to orchestration, that is, Z brain. It goes to the agent and API, and from here it goes to sentiment analysis, then it goes back to orchestration. A query is sent, an output is given, then we have LLM cache, LLM Ops, and visualization of guardrails, which leads to proprietary LLMs and open-source models. So, now let us start with the first thing, that is, data sources. The process begins by gathering data from various sources relevant to sentiment analysis. This data can include, first, social media.

Data retrieved from platforms like X, Facebook, and Instagram, capturing user-generated content such as posts, comments, and messages to analyze public sentiments, opinions, and trends. Customer reviews, written feedback, written feedback and ratings provided by customers on platforms such as Amazon, Yelp, and Google Reviews, offering insights into customer satisfaction, preferences, and sentiments regarding products or services. Surveys and feedback forms: responses collected from structured questionnaires, online surveys, and feedback forms designed to gather opinions, attitudes, and individual feedback regarding specific topics, products, or services. Call center transcripts, emails, and support tickets: contextual records of interactions between customers and customer

support representatives, including transcripts of phone calls, emails, and support tickets, offering insights into customer concerns,

issues, and sentiments towards a company's products or services. The second is data pipeline. The data generated from various previous sources is subsequently channeled through data pipelines. These pipelines handle data such as data ingestion, cleaning, processing, including data transformations like filtering, masking, and aggregations, and structuring, preparing it for subsequent analysis. The third is embedding models.

The processed data is segmented into chunks and fed into an embedding model. This model converts textual data into numerical representations called vectors, enabling AI models to comprehend it effectively. Well-known models for this purpose are developed by OpenAI, Google, and similar organizations. The fourth is the vector database. The resulting vectors are stored in a vector database, facilitating streamlined querying and retrieval processes.

This data BASE efficiently manages the storage, comparison, and retrieval of potentially billions of embeddings, that is, vectors. Prominent examples of such vector databases include Pinecone, VV8, and PGVector. Next come APIs and plugins. APIs and plugins such as SERPs, ZEPR, and BOLFORM are crucial in linking various components and facilitating additional functionalities.

Such as accessing additional data or executing specific tasks seamlessly. Orchestration layers. The orchestration layer is critical in managing the workflow. ZBrain is an example of this layer that simplifies prompt chaining, manages interactions with external APIs by determining when API calls are required, retrieves contextual data from vector databases, and maintains memory across multiple LLM calls. Ultimately, this layer generates a prompt or series of prompts that are submitted to the language model for processing.

The role of this layer is to orchestrate the flow of data and tasks, ensuring seamless coordination across all the architectural components. The seventh is query execution. The data retrieval and generation process initiates when the user submits a query to the sentiment analysis application. These queries can cover various aspects relevant to the user experience, including satisfaction levels, feature preferences, usability concerns, and overall sentiment toward the product or service.

The next is LLM processing. Upon receiving the query, the application forwards it to the orchestration layer. This layer then retrieves pertinent data from the vector database, and LLM cache before sending it to the suitable LLM for processing. The apt LLM is selected based on the query's nature.

Next is the output the LLM produces. The LLM produces an output based on the user query and the data it receives. This output can take various forms pertinent to analyzing sentiments toward the product service, such as identifying key themes in customer feedback, generating sentiment reports, or summarizing sentiment across different user segments or demographics. The verified output is then presented to users through the Sentiment Analysis App.

This central platform integrates all gathered data Sentiment analysis results and actionable insights, presenting them in an accessible format for decision making to understand customer sentiment trends and make informed strategic decisions. The feedback loop, user feedback on the LLM output is another important aspect of this architecture. The feedback is used to improve the accuracy and relevance of the LLM. AI agents are critical in this architecture to address complex problems, interact with external environment, and enhance learning through post-deployment experiences.

They achieve this by employing advanced reasoning, planning, strategic tool utilization, and leveraging memory, recursion, and self-reflection. The next is LLM cache. Tools like Redis, SQLite and GPT cache are used to cache frequently accessed information, speeding up the response time of the AI system. Logging, that is, or LLMOps.

Throughout this process, LLM operations, tools like weights and biases, MLflow, Helicon, and PromptLayer helps log action and monitor performance. This ensures the LLM are functioning optimally and continuously improve through feedback loops. The next comes validation. A validation layer is employed to validate the LLM output. This is done through tools like guardrails, rebuff, guidance and LMQL to ensure the accuracy and reliability of the information provided.

The 16th is LLM APIs and hosting. LLM APIs and hosting platforms are essential for executing sentiment analysis tasks and hosting the application. Depending on the requirements, developers can select from LLM APIs such as those offered by OpenAI and Anthropic or opt for open-source models. Similarly, they can choose hosting platforms from cloud providers like AWS, GCP, Azure or CoreV or opt for opinionated clouds like Databricks, Mosaics, and NHKL.

The choice of LLM APIs and cloud hosting platforms depends on the project's needs. This structured flow outlines how AI streamlines sentiment analysis tasks utilizing diverse data sources and advanced tools to deliver precise and actionable insights on customer sentiments. By automating the sentiment analysis process, AI enhances operational efficiency and empowers stakeholders to conduct thorough sentiment assessments, leading to informed decision-making and a deeper understanding of customer sentiments. Now, let us look at the AI tools for sentiment analysis.

When it comes to sentiment analysis, several AI-powered tools have gained prominence for their accuracy, functionality, and ease of use. Each tool has unique features, functionalities, and pricing options. When choosing a sentiment analysis tool, consider factors such as your specific project requirements, ease of integration, scalability, and the level of technical expertise needed.

Exploring these tools will empower project managers to effectively harness the power of sentiment analysis and gain valuable insights into public perceptions. Let us explore some of the top tools available today. The first is IBM Watson Natural Language Understanding. IBM Watson NLU is a comprehensive AI tool that offers sentiment analysis as one of its core features. It utilizes advanced NLP techniques and machine learning algorithms to analyze text and extract sentiments.

With its robust capabilities, Watson NLU provides accurate sentiment classification, entity recognition, and keyword extraction. The next is Google Cloud Natural Language API. Google Cloud Natural Language API provides a wide range of language processing functionalities, including sentiment analysis. Leveraging Google's powerful machine learning models, this tool can analyze text and return sentiment scores and labels for each document. It also offers entity recognition, content classification, and syntax analysis.

Microsoft Azure Text Analytics is a comprehensive text analytics service that offers sentiment analysis as part of its features. suite of features. Using cutting-edge AI models, this tool can detect sentiments, extract key phrases, and perform language detection. It supports multiple languages and provides rich insights into sentiment strengths and scores. The next is Hugging Face Transformers.

Hugging Face Transformers is an open-source library that provides state-of-the-art machine learning models for various language processing tasks, including sentiment analysis. With a wide selection of pre-trained models, developers and data scientists can easily implement sentiment analysis in their products. This library supports popular deep

learning frameworks such as TensorFlow and PyTorch. The next is RapidMiner. RapidMiner is an end-to-end data science platform that offers a range of analytical capabilities, including sentiment analysis.

With its intuitive interface, RapidMiner allows users to build sentiment analysis models using a visual workflow or programming. It supports text processing, feature engineering, and sentiment classification. Now, let us look at the use cases for AI-enabled sentiment analysis across industry verticals. The first is AI-driven sentiment analysis in retail.

The retail sector can harness the power of AI-driven sentiment analysis to gain a deeper understanding of customer perception about their brand and pinpoint upcoming market trends. Here is an outline of its use case in retail. First is understanding customer sentiments. AI sentiment analysis tools can evaluate customer reviews, social media posts, and other user-generated content to understand customer feelings about a brand or its products.

These insights can help retailers identify areas of success or areas that need improvement. The next is personalized marketing. By understanding individual customer sentiment, retailers can tailor their marketing and advertising strategies to better resonate with their target audience. This personalized approach can increase engagement and improve overall marketing effectiveness. Customer service improvement.

AI for sentiment analysis can identify negative sentiments in real time. Allowing customer service teams to address complaints or issues promptly. This proactive approach can help improve customer satisfaction and brand perception. Next is product development. Customer sentiment analysis can reveal

What features customers like or dislike about a product. This valuable feedback can guide product development teams to modify products to better meet customer expectations. Competitive analysis. By applying sentiment analysis to public opinion about competitors, retailers can identify strengths and weaknesses in their rivals' offerings. Such insights can help them position themselves.

Their product more effectively in the market. Trend forecasting. AI for sentiment analysis can identify consumer attitudes and preference shifts over time. Providing early warnings of changing market trends, retailers can use this information to stay ahead of the curve, adjust their strategies, and cater to evolving customer needs. The second is AI-based sentiment analysis in tourism and hospitality.

The hospitality sector can greatly benefit from using AI for sentiment analysis in many ways. Here is a snapshot of its use case. First is understanding guest feedback. Hotels, restaurants, and other hospitality businesses receive massive amounts of feedback through online reviews, social media posts, and direct customer interactions. AI-based sentiment analysis can help.

These businesses efficiently parse through these feedback data to understand customer sentiments, uncovering valuable insights about their services. Enhancing guest experience. Sentiment analysis can reveal specific aspects of the guest experience that are especially delightful or problematic. This information can help hospitality businesses tailor their services to meet guest expectations better. Leading to increased customer satisfaction and loyalty.

The next is real-time service recovery. AI for sentiment analysis can identify negative sentiments in real-time, allowing the hospitality businesses to interact. Hospitality businesses to intervene immediately and resolve the issue before it escalates. This can significantly improve the overall guest experience and prevent potential reputational damage. The next is the strategic decision-making by gauging public sentiments

About different aspects of their offerings, businesses in the hospitality sector can make data-driven decisions regarding their services. Whether revamping a restaurant's menu or redesigning a hotel's room, sentiment analysis can provide actionable insights that help enhance business strategy. Competitor analysis. By applying sentiment analysis to reviews of competitor hotels and restaurant businesses can identify their own relative strengths. and weaknesses this information can help them position their services more effectively in the market the next comes trends trend identification ai for sentiment analysis can help hospitality businesses identify emerging trends or changing guest preferences this can be particularly useful in such a dynamic industry where staying on top of trends can provide a significant competitive advantage the third example is the use of ai for sentiment analysis and telecommunication

AI driven sentiment analysis plays a crucial a critical role in the telecom industry, providing a range of use cases to improve customer experience, business operations and strategic decision making. Here is how it can be employed. The first is for customer experience management. The telecom sector often has to analyze massive volumes of customer interaction across various channels, including call centers, social media, email and more.

So, AI-based sentiment analysis can process this data to understand customer sentiment and identify pain points. Enabling proactive customer service and improving customer experience. The next is churn prediction and prevention. Telecom providers can identify dissatisfied customers and predict potential churn by analyzing customer sentiments over time. This information can be leveraged to implement targeted strategies and offers to retain at-risk customers.

Service improvement sentiment analysis can reveal underlying issues with service quality, network coverage, or pricing causing customer dissatisfaction. Telecom providers can use these insights to improve these areas and enhance customer satisfaction. The next is competitive analysis. Sentiment analysis can be applied to public discussions and reviews about competitors to understand their strengths and weaknesses from the customer's perspective.

This can help telecom providers position their services more effectively and identify opportunities for differentiation. Next comes product development. By understanding how customers feel about different features and services, telecom providers can align their product development efforts with customer needs and preferences. Next comes marketing and sales.

Sentiment analysis can provide insights into how marketing and sales messages resonate with the audience. This can inform the development of more effective marketing strategies and sales pitches. Real-time decision-making. In an industry as dynamic as telecom, real-time sentiment analysis can help providers respond quickly to emerging issues or opportunities, leading to better decision-making and agility. The next example is AI-driven sentiment analysis in healthcare.

In the healthcare sector, AI sentiment analysis can be harnessed in various ways to improve patient care, service delivery, and operational efficiency. Here are some potential use cases. One is patient feedback analysis. AI-driven sentiment analysis can interpret patient feedback from various sources, such as social media, online reviews, and patient surveys. By identifying positive and negative sentiments, healthcare providers can better understand patients' experiences, address concerns, and enhance service quality.

Next is improving patient care. Sentiment analysis can aid in interpreting patients' feelings and emotions conveyed through consultations or written in patient records or communications. Healthcare professionals can better understand a patient's mental and emotional state, potentially leading to more personalized and effective care.

Pharmacovigilance: Sentiment analysis can aid in drug safety surveillance, known as pharmacovigilance. Potential adverse drug reactions can be identified easily by analyzing sentiments in patient forums.

Social media posts and other online platforms. Clinical trial monitoring. AI sentiment analysis can track participants' experiences in clinical trials. Analyzing sentiments in patient reports can provide early indications of side effects, efficacy, and adherence issues. The next is healthcare marketing.

Using sentiment analysis, healthcare organizations can understand public perceptions and sentiments about their brands, services, or specific marketing campaigns. This information can guide marketing strategies and communication. Policymaking: government and regulatory bodies can use sentiment analysis to gauge public sentiments toward health policies or public health issues. Aiding in policy formulation and public health intervention.

Mental health analysis. Sentiment analysis could potentially help monitor patients' mental health by analyzing their written or spoken communication for signs of negative sentiments or distress. The next example we will consider is AI-powered sentiment analysis in banking. AI sentiment analysis plays a significant role in banking by enhancing customer experience, risk management, and brand perception. Here are some

Potential use cases. The first is customer service improvement. AI in sentiment analysis can analyze customer feedback from various sources such as social media, online reviews, and customer surveys. By identifying and understanding the sentiment, banks can improve their services and address customer complaints more effectively. Risk management sentiment analysis can assist in early warning signal detection for credit risk management.

For example, negative sentiments from a business customer's can indicate potential financial distress, which can be factored into the bank's risk assessment process. The next is product and service development. By understanding customer sentiments toward various banking products and services, banks can gain valuable insights into what customers appreciate and what they do not. This can guide the development of new products and refinement of existing services.

The next is brand perception. So, sentiment analysis can help banks understand the overall perception of their brand in the marketplace. They can identify positive and

negative sentiments, understand trends over time, and compare their sentiment scores with competitors. Marketing strategy: By analyzing customer sentiment toward various marketing campaigns, banks can understand which messages resonate with customers and why. This can help guide future marketing strategies. Customer segmentation: Sentiment analysis can categorize customers based on the sentiments. This segmentation can tailor communication and offers to different customer segments, enhancing personalization and customer satisfaction. The next is fraud detection.

While not a direct application, sentiment analysis can indirectly support fraud detection. Unusual negative sentiments or sudden sentiment changes can indicate potential fraudulent activity, triggering further investigation. So, to conclude this module, we have discussed what customer sentiment analysis is and its various types. We also understood the role of AI in customer sentiment analysis and how it actually works. We gave an overview of the various AI tools that can be used for sentiment analysis.

Finally, we discussed various use cases of AI-enabled sentiment analysis across industry verticals. These are some of the sources from which material for this module was taken. Thank you.