Topic Modeling - Basic Handout

Machine Learning for Smarter Innovation

1 Topic Modeling - Basic Handout

Target Audience: Beginners with no technical background **Duration**: 30 minutes reading **Level**: Basic (no math, no code)

1.1 What is Topic Modeling?

Topic modeling is the automatic discovery of themes or subjects in a collection of documents. Given thousands of documents, topic modeling algorithms identify what those documents are about without anyone reading them all.

Think of sorting a massive pile of mail without opening each envelope. You might notice patterns - some envelopes have bank logos, others have charity appeals, some look like magazines. Topic modeling does something similar with text content, finding groups of words that tend to appear together and represent coherent themes.

A topic is essentially a cluster of related words. If many documents contain words like "battery," "charging," "power," and "drain" together, the algorithm identifies a "battery/power" topic. Documents can belong to multiple topics - a product review might discuss both battery life and design quality.

The key insight is that documents about the same subject use similar vocabulary. Articles about finance mention stocks, markets, and investments. Articles about health mention symptoms, treatment, and wellness. Topic modeling exploits these vocabulary patterns to organize large document collections automatically.

1.2 Why Does Topic Modeling Matter?

Organizations generate and collect far more text than humans can read. Customer feedback, support tickets, social media mentions, survey responses, research papers, internal documents - the volume overwhelms manual analysis. Topic modeling makes sense of this text at scale.

For customer insights, topic modeling reveals what customers actually discuss. Rather than assuming what matters, you discover the themes that emerge organically from thousands of conversations. A company might discover that customers talk about "ease of use" far more than expected, or that a specific feature generates disproportionate discussion.

For market research, topic modeling analyzes industry publications, news coverage, and competitive communications to identify emerging themes. What topics are gaining attention? What terminology is becoming standard? These insights inform strategy and positioning.

For knowledge management, topic modeling helps organize large document collections. A legal firm with millions of documents can automatically categorize and tag content, making relevant materials findable. A research organization can identify which topics their publications address and where gaps exist.

Topic modeling also supports content creation by revealing what audiences care about. Analyzing popular content in your domain shows which themes resonate, guiding editorial decisions and content strategy.

1.3 Key Concepts

1.3.1 Topics: Collections of Related Words

A topic is a group of words that frequently appear together across documents. The algorithm does not name topics - it discovers word clusters, and humans interpret what those clusters mean.

For example, analyzing product reviews might reveal: - Topic A: battery, charging, power, life, drain, hours (likely "Battery/Power") - Topic B: screen, display, bright, resolution, colors (likely "Display Quality") - Topic C: price, expensive, value, worth, cost (likely "Value/Price")

Topics can be broad or narrow depending on how the algorithm is configured. More topics create finer distinctions; fewer topics create broader categories. The right number depends on your purpose.

1.3.2 Documents: Mixtures of Topics

Each document contains a mixture of topics in various proportions. A product review might be 50% about battery life, 30% about display, and 20% about price. Another review might focus entirely on one topic or balance several equally.

This mixture model reflects reality - real documents rarely address only one subject. A news article about a company might discuss financial performance, new products, and leadership changes all in one piece. Topic modeling captures this complexity rather than forcing single-label categorization.

The topic mixture is called a document's topic distribution. This distribution becomes a compact representation of what the document is about, useful for search, recommendation, and organization.

1.3.3 Words: Building Blocks

Words are the fundamental units that define topics. The algorithm analyzes which words tend to co-occur across documents. Words that frequently appear in the same documents cluster into topics.

Importantly, the same word can contribute to multiple topics. "Apple" might appear in topics about technology and topics about nutrition. Context - the other words nearby - determines which topic the word represents in each document.

Common words like "the," "is," and "and" appear everywhere and do not help distinguish topics. These stopwords are typically removed before analysis, focusing on words that carry topical meaning.

1.3.4 Latent Themes: Hidden Structure

Topics are called "latent" because they are not explicitly marked in the documents - they are hidden patterns discovered by the algorithm. No one labeled which documents belong to which topics. The structure emerges from analyzing word patterns across the entire collection.

This unsupervised discovery is the power of topic modeling. You do not need to know in advance what topics exist. The algorithm finds whatever themes emerge from the data, potentially revealing patterns you would never have hypothesized.

1.4 How It Works (Plain English)

Topic modeling follows a logical process from raw documents to discovered themes.

Step 1: Prepare Documents

Text must be cleaned and standardized. Remove formatting, correct encoding issues, and handle special characters. Convert everything to consistent lowercase. Remove stopwords that appear everywhere but carry no topical meaning.

Step 2: Build a Vocabulary

Create a list of all unique words across all documents. Each document becomes a list of words from this vocabulary. Very rare words that appear in only one or two documents are often removed since they do not help identify broad themes.

Step 3: Discover Topics

The algorithm analyzes word co-occurrence patterns across all documents. Words that frequently appear together in the same documents cluster into topics. The algorithm finds the best way to explain the observed word patterns using the specified number of topics.

Step 4: Assign Topic Proportions

Each document receives a distribution over topics - the proportion of the document attributable to each topic. A document might be 60% Topic A, 25% Topic B, and 15% Topic C. These proportions sum to 100%.

Step 5: Interpret Results

Humans examine the top words in each topic and assign meaningful labels. Topic 1 with words "price, cost, value, expensive, affordable" might be labeled "Pricing Concerns." This interpretation step requires domain knowledge.

Step 6: Apply Insights

Use discovered topics to organize, search, monitor, or analyze documents. Track how topic prevalence changes over time. Identify which topics different customer segments discuss. Route documents based on topic content.

1.5 Real-World Applications

1.5.1 Customer Feedback Analysis

Companies analyze thousands of reviews, survey responses, and support tickets to understand customer concerns. Topic modeling reveals the main themes customers discuss and their relative prevalence. Product teams learn what features matter most; support teams identify common issues for proactive resolution.

1.5.2 Academic Research

Researchers analyze large literature collections to understand a field's evolution. Topic modeling identifies research themes, shows how they change over time, and reveals connections between areas. This supports literature reviews and identifies research gaps.

1.5.3 News and Media Monitoring

PR teams and analysts track news coverage to understand how topics evolve. Topic modeling processes thousands of articles to identify trending themes, emerging issues, and changes in discourse. This supports reputation management and strategic communications.

1.5.4 Internal Knowledge Organization

Large organizations struggle to organize accumulated knowledge. Topic modeling automatically tags and categorizes documents, making them searchable and discoverable. New employees can find relevant materials; teams can locate related work across the organization.

1.6 Common Misconceptions

1.6.1 "Topics Are Definitive Categories"

Topics are one valid interpretation of document themes, not the objectively correct categorization. Different algorithm settings produce different topics. Multiple valid topic structures usually exist for any document collection.

1.6.2 "The Algorithm Names Topics"

Algorithms discover word clusters, not topic names. The output shows groups of related words. Humans must examine these word groups and assign meaningful labels based on domain knowledge. "battery, power, charge, life" becomes "Battery Issues" through human interpretation.

1.6.3 "More Topics Are Better"

Too many topics fragment coherent themes into meaningless subdivisions. Too few topics merge distinct themes into overly broad categories. The right number balances granularity with interpretability and depends on your specific use case.

1.6.4 "Topic Modeling Works on Short Text"

Topic modeling works best on longer documents where word co-occurrence patterns can emerge. Short texts like tweets or single-sentence feedback may not contain enough words for reliable topic inference. Aggregating short texts into longer documents sometimes helps.

1.7 When to Use / When Not to Use

1.7.1 Use Topic Modeling When:

- You have a large collection of documents (hundreds to millions)
- Documents are long enough for word patterns to emerge
- You want to discover themes without predefined categories
- You need to organize, search, or monitor text collections
- Manual reading of all documents is impractical
- You want to track how themes change over time

1.7.2 Do Not Use Topic Modeling When:

- You have very few documents (under 100)
- Documents are very short (single sentences or phrases)
- You already know exactly what categories exist
- You need high-precision classification with labeled training data
- The text is highly technical with specialized vocabulary
- Real-time processing of individual documents is required

1.8 Getting Started Checklist

- ☐ Define what questions topic modeling should answer
- ☐ Gather a sufficient document collection (hundreds minimum)
- ☐ Ensure documents are long enough for meaningful analysis
- ☐ Clean text data and remove obvious errors
- ☐ Remove stopwords and very rare words
- ☐ Start with a reasonable number of topics (5-10)
- ☐ Examine top words in each discovered topic
- ☐ Assign meaningful labels to interpretable topics
- ☐ Validate topics make sense with domain experts
- ☐ Plan how topics will inform decisions or actions

1.9 Key Terms Glossary

() *	() *	
0.33333	0.33333	Definition
() *	() *	
0.33333	0.33333	A group of related words that frequently co-occur in documents
() *	() *	
0.33333	0.33333	Automatic discovery of themes in a document collection

()	()		
*	*		
0.33	0.33	0.33	Definition
()	()		
*	*		
0.33	0.33	0.33	Any u- text me- ne- ing an- a- lyzed (re- view, ar- ti- cle, email, etc.)
()	()		
*	*		
0.33	0.33	0.33	The dis- pro- tri- por- bu- tion tion of each topic present in a doc- u- ment
()	()		
*	*		
0.33	0.33	0.33	Hid- ten den struc- ture dis- cov- ered by the al- go- rithm, not ex- plic- itly la- beled

()	()		
*	*		
0.3333	0.3333	0.3333	Definition
()	()		
*	*		
0.3333	0.3333	0.3333	LDA - Latent Dirichlet Allocation, the most common topic modeling algorithm
()	()		
*	*		
0.3333	0.3333	0.3333	stopwords - words removed before analysis (the, is, and, etc.)
()	()		
*	*		
0.3333	0.3333	0.3333	The - entire collection of documents being analyzed

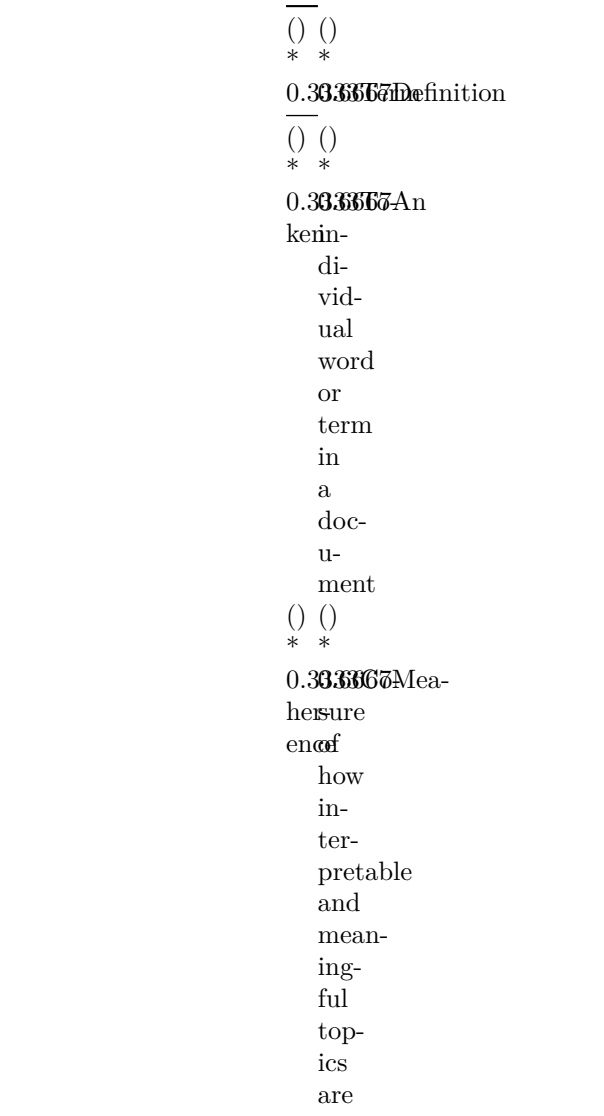

1.10 Next Steps

Ready to implement topic modeling? The intermediate handout covers Python implementation using Gensim and LDA, including text preprocessing, model training, topic interpretation, and visualization with pyLDAvis.

For immediate exploration without code, consider what document collections exist in your organization that might benefit from automatic theme discovery. Customer feedback archives, internal documents, industry publications, or research libraries all potentially contain hidden themes waiting to be revealed.

Start by manually examining a sample of documents and hypothesizing what topics might exist. This exercise builds intuition for what topic modeling discovers automatically at scale.

Topic modeling transforms overwhelming document collections into navigable themes. The algorithm finds patterns; you provide interpretation and purpose. Together, they reveal what thousands of documents are really about.