

# NLP for Emotional Context

From Words to Understanding What Users Really Feel

Week 3: Machine Learning for Smarter Innovation

Transform 50,000 Reviews into Actionable Design Insights

## Four Stages of Understanding

1. **The Challenge** - Why understanding emotion in text is impossibly hard
2. **First Solution & Its Limits** - Traditional NLP works... until it doesn't
3. **The Breakthrough** - How Transformers changed everything
4. **Design Synthesis** - From ML insights to user empathy

**Core Question:** How can we understand the emotions of 50,000 users without reading 5 million words?

---

Large-scale emotional analysis enables mass personalization - automated sentiment understanding processes millions of user expressions beyond manual capacity

## Which Reviews Reveal Why Users Really Quit?

### The Scenario You Face:

- Your product has 50,000 user reviews
- Average review: 100 words
- Total: 5 million words to process
- Hidden inside: The 10 reviews that explain 80% of churn

### The Human Limit:

- Reading speed: 200 words/minute
- Time to read all: 417 hours (10 weeks!)
- Reviews arrive: 500 new ones daily
- **You're already behind before you start**

### Real Review Examples:

- “Love it but...” (quit in 2 weeks)
- “Not bad for the price” (renewed 3 years)
- “Just what I expected!” (1-star rating)
- “Finally someone gets it” (5-star champion)

charts/review\_volume\_growth.pdf

### Context Changes Everything

#### Same Word, Different Meanings:

##### “Cold”

- Restaurant: “The soup was cold” → Negative
- Service: “Support was cold” → Negative
- Beer: “Nice and cold” → Positive
- Logic: “Cold hard facts” → Neutral

##### “Fast”

- Delivery: “Super fast!” → Positive
- Battery: “Dies too fast” → Negative
- Customer service: “Too fast, felt rushed” → Negative

#### The Computer's Problem:

charts/context\_dependency.pdf

### What People Say What They Feel

#### Layer 1: Literal

- “Great product”
- Clear positive
- Easy to detect
- 15% of reviews

#### Example:

“This is excellent. I love every feature. Will buy again.”

→ Genuinely positive

#### Layer 2: Sarcastic

- “Just wonderful”
- Opposite meaning
- Context reveals truth
- 25% of reviews

#### Example:

“Oh great, another update that breaks everything”

→ Actually negative

#### Layer 3: Cultural

- “It’s okay”
- Varies by culture
- UK: Negative
- US: Neutral

#### Example:

“It does what it says”

UK: Disappointed

US: Satisfied

**The Complexity:** 60% of emotional meaning is NOT in the words themselves but in how they're used together

Interpretive complexity limits consensus - emotional content carries inherent ambiguity even among expert human evaluators

## Why This Problem Explodes in Complexity

charts/complexity\_explosion.pdf

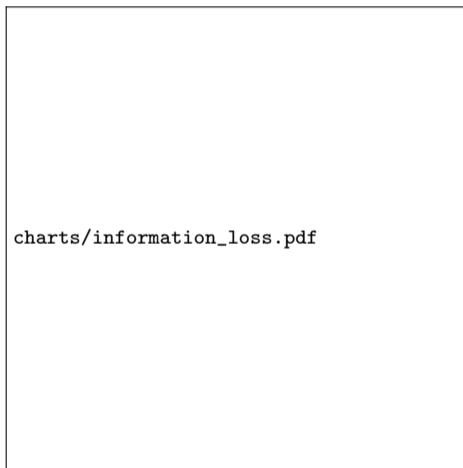
## The Mathematical Impossibility

### Information Content Analysis:

Component	Bits	Information
One word	14 bits	Which of 10,000 words
+ Position	7 bits	Where in sentence
+ Emotion	3 bits	8 emotion types
+ Context	10 bits	Surrounding words
+ Sarcasm	1 bit	Yes/No
<b>Total per word</b>	<b>35 bits</b>	
<b>Per review (100 words)</b>	<b>3,500 bits</b>	
<b>50,000 reviews</b>	<b>175M bits</b>	<b>22 MB</b>

### Traditional Compression:

- Keyword counts: 10,000 → 100 words
- Compression ratio: 100:1
- Information lost: **99%**



## How Would YOU Quickly Scan 1000 Reviews?

### Human Strategy:

1. Look for emotion words: “love”, “hate”, “terrible”
2. Count positive vs negative
3. More positive → Happy customer
4. More negative → Unhappy customer

### Computer Implementation:

- Build word lists
- Positive: [great, excellent, love, amazing...]
- Negative: [bad, terrible, hate, awful...]
- Count occurrences
- Calculate:  $\text{Positive\%} - \text{Negative\%}$

### Example Analysis:

“I love the design but hate waiting. The quality is excellent though shipping was terrible.”

**Count:** 2 positive, 2 negative

**Score:**  $50\% - 50\% = \text{Neutral (0)}$

Seems reasonable!

## The Classic Approach in Action

### Step-by-Step Process:

#### 1. Original Review:

"The product quality is excellent excellent excellent but customer service terrible terrible"

#### 2. Tokenize (split into words):

The, product, quality, is, excellent, excellent, excellent, but, customer, service, terrible, terrible

#### 3. Count Each Word:

Word	Count	Word	Count
excellent	3	terrible	2
product	1	service	1
quality	1	customer	1

#### 4. Convert to Vector:

1, 1, 1, 3, 0, 0, 0, 1, 1, 2, 0, ...  
(10,000 dimensions, mostly zeros)

charts/bow\_visualization.pdf

## Not All Words Are Equal

### The Problem with Raw Counts:

- “The” appears 1000 times → Important?
- “Revolutionary” appears once → Not important?
- Common words dominate
- Rare but meaningful words ignored

### TF-IDF Solution:

$$\text{TF-IDF} = \underbrace{\frac{\text{count in doc}}{\text{total words}}}_{\text{Term Frequency}} \times \log \underbrace{\frac{\text{total docs}}{\text{docs with word}}}_{\text{Inverse Document Freq}}$$

- TF: How often in THIS review
- IDF: How rare across ALL reviews
- Product: Important AND distinctive

### Example Calculation:

Word	TF	IDF	TF-IDF
“the”	0.15	0.01	0.0015
“product”	0.05	0.69	0.0345
“excellent”	0.10	1.38	0.138
“revolutionary”	0.02	3.40	0.068

**Result:** Meaningful words now weighted higher than common words!

Rarity signals importance - terms appearing selectively carry more discriminative power than ubiquitous vocabulary

## When Simple Reviews Get Perfect Scores

Review Text	Human	BoW+TFIDF	Match
"This product is absolutely fantastic! Best purchase ever!"	Positive	Positive (95%)	
"Terrible quality. Waste of money. Very disappointed."	Negative	Negative (92%)	
"Good product, fair price, happy with purchase"	Positive	Positive (88%)	
"Broken on arrival. Customer service un- helpful. Never again."	Negative	Negative (94%)	
"Excellent! Exceeded all expectations! Highly recommend!"	Positive	Positive (97%)	

**Average Accuracy: 93.2%**

This is why Bag of Words dominated for 40 years!

Explicit signals enable simple methods - literal expression reduces analytical complexity compared to implicit communication

## Performance Degradation with Complexity

Review Type	Example	Human	BoW	Accuracy
Simple & Direct	"Great product!"	Pos	Pos	95%
Mixed Sentiment	"Good but overpriced"	Neg	Pos	67%
Sarcasm	"Oh great, it broke. Just perfect!"	Neg	Pos	23%
Context Dependent	"Not bad for the price"	Pos	Neg	31%
Subtle Emotion	"It's fine, I guess"	Neg	Neu	44%
Negation	"Not the worst I've seen"	Neu	Neg	28%

### The Pattern:

- Simple: 95% → Works great!
- Real-world: 44% → Worse than coin flip
- Sarcasm: 23% → Actively wrong

**Average: 51% (Random guessing: 50%)**

## Tracing the Failure Through Real Examples

### Example: "Not bad for the price"

#### What BoW Sees:

- Words: [not, bad, for, the, price]
- "bad" → Negative word (-1)
- "not" → Negation word (ignored)
- Score: Negative

#### What Got Lost:

- "not bad" = Actually positive
- "for the price" = Qualified satisfaction
- Relationship between words
- Overall: Mild positive sentiment

**Root Cause:** Word order contains meaning!

### Information Loss Calculation:

Information Type	Bits Lost
Word positions	420 bits
Word relationships	1,200 bits
Grammar structure	350 bits
Contextual meaning	890 bits
<b>Total Lost</b>	<b>2,860 bits</b>
<b>Total Kept</b>	<b>640 bits</b>
<b>Kept Percentage</b>	<b>18%</b>

charts/information\_preserved.pdf

## Let's Be Honest About Our Mental Process

### Trace Your Thinking:

**Sentence:** “I went to the bank. The water was nice.”

1. Read “I went to the bank”
2. Your brain: Could be financial or river...
3. Read “The water was nice”
4. Your brain: Ah! River bank, not financial
5. Re-interpret: “bank” = riverbank
6. Understand: Person enjoyed the riverside

### What You Actually Did:

- Held multiple meanings open
- Used later words to resolve earlier ones
- **Selectively focused on “water” to understand “bank”**

### The Key Observation:

You DON'T compress the sentence into a summary.

You KEEP all words available and SELECTIVELY ATTEND to relevant ones when needed.

**This is the breakthrough insight!**

## A Fundamentally Different Approach

### Old Way (Bag of Words):

- Take all words
- Compress to counts
- Lose relationships
- Try to reconstruct meaning

### Analogy:

Like taking a book, counting word frequencies, throwing away the book, then trying to understand the story from counts.

**Result:** Lost 82% of information  
Can't recover context

### New Way (Attention):

- Keep all words
- Keep all positions
- For each word, look at all others
- Decide how much to focus on each

### Analogy:

Like keeping the entire book and using a smart highlighter that knows which sentences help understand other sentences.

**Result:** Keep 100% of information  
Use what's relevant when needed

Instead of asking “What to keep?” we ask “What to focus on?”

Paradigm shifts replace constraints with capabilities - architectural innovations enable previously impossible approaches

## No Math Yet - Just Percentages

**Example:** Understanding “bank” in: “The bank by the river is peaceful”

**Step 1: For word “bank”, look at all other words**

Word	Relevance to “bank”
The	Low
by	Medium
the	Low
river	<b>Very High</b>
is	Low
peaceful	Medium

**Step 2: Convert to percentages (must sum to 100%)**

Word	Focus %
The	5%
by	15%
the	5%
river	<b>55%</b>
is	5%
peaceful	15%

**Step 3: Use these percentages to understand “bank”**

“bank” meaning =  
 $5\% \times \text{meaning}(\text{“The”}) +$   
 $15\% \times \text{meaning}(\text{“by”}) +$   
 $5\% \times \text{meaning}(\text{“the”}) +$   
 $55\% \times \text{meaning}(\text{“river”}) +$   
 $5\% \times \text{meaning}(\text{“is”}) +$   
 $15\% \times \text{meaning}(\text{“peaceful”})$   
= Mostly “river” context  
= Riverbank, not financial bank!

**These percentages ARE the attention weights!**

Normalized weighting enables comparison - probability distributions force explicit trade-offs among competing information sources

## Words as Directions in Space

Start Simple: 2D Space

charts/word\_vectors\_2d.pdf

Scale to Real NLP: 768D Space

**Same principle, more dimensions:**

- 2D:  $[x, y]$  coordinates
- 768D:  $[x, x, \dots, x]$  coordinates

More dimensions = More nuanced meaning

Can distinguish "bank (river)" from "bank (money)" from "bank (slope)"

**The Attention Calculation:**

1. Compute alignment (dot product) with all words
2. Higher alignment = Pay more attention
3. Convert to percentages
4. Use percentages to blend meanings

**Geometry gives us semantic similarity!**

## Each Step Has a Purpose

### Step 1: SCORE - Find Relevance

- Action: For each word, compute alignment with all others
- Why: Identify which words help understand this one
- Math:  $\text{Score}_{ij} = Q_i \cdot K_j$
- Plain: How relevant is word  $j$  to understanding word  $i$ ?

### Step 2: NORMALIZE - Create Percentages

- Action: Convert scores to probabilities (0-100%)
- Why: Need weights that sum to 1.0 for averaging
- Math:  $\alpha_{ij} = \frac{e^{\text{Score}_{ij}}}{\sum_k e^{\text{Score}_{ik}}}$
- Plain: What percentage of attention should word  $j$  get?

### Step 3: AGGREGATE - Blend Information

- Action: Weighted sum of all word meanings
- Why: Combine context based on attention weights
- Math:  $\text{Output}_i = \sum_j \alpha_{ij} \cdot V_j$
- Plain: New understanding using focused context

charts/attention\_three\_steps.pdf

## Real Numbers, Real Calculation

**Task:** Computing attention for “nice” in: “The service was nice”

**Given: Word Vectors (Simplified to 2D)**

The = [1, 0], service = [2, 3], was = [0, 1], nice = [3, 2]

**Step 1: Calculate Scores**

Query (nice) = [3, 2]

Word	Key	Q·K	Score
The	[1, 0]	$3 \times 1 + 2 \times 0$	3
service	[2, 3]	$3 \times 2 + 2 \times 3$	12
was	[0, 1]	$3 \times 0 + 2 \times 1$	2
nice	[3, 2]	$3 \times 3 + 2 \times 2$	13

**Step 2: Convert to Percentages (Softmax)**

Word	$e^{\text{Score}}$	Attention %
The	20	0.7%
service	162,755	<b>63.1%</b>
was	7	0.0%
nice	442,413	36.2%

**Step 3: Aggregate**

Values: The=[1,1], service=[4,5], was=[1,2], nice=[5,4]

Output for “nice” =

$$\begin{aligned}
 &0.007 \times [1,1] + \\
 &\mathbf{0.631 \times [4,5]} + \\
 &0.000 \times [1,2] + \\
 &0.362 \times [5,4] \\
 &= [0.01, 0.01] + \mathbf{[2.52, 3.16]} + [0, 0] + [1.81, 1.45] \\
 &= [4.34, 4.62]
 \end{aligned}$$

**Result:** 63% attention on “service”

The word “nice” is understood primarily in context of “service” → Customer service sentiment!

**Low-dimensional principles generalize to high dimensions - mathematical operations remain identical regardless of vector size**

## The Power of Looking Forward AND Backward

### Traditional (Left-to-Right Only):

Sentence: "The bank charges are too high"

Reading "bank":

Can see: "The"; Cannot see: "charges are too high"; Guess: Could be river or financial...; **50% chance of error**

### BERT (Bidirectional):

Same sentence: "The bank charges are too high"

Reading "bank":

Can see: "The" AND "charges are too high"; "charges" → financial context; Certain: Financial bank; **99% accurate**

`charts/bert_bidirectional.pdf`

### Standing on the Shoulders of Human Knowledge

#### **BERT's Pre-training Data:**

Wikipedia (2.5B words) + BookCorpus (800M words) = 3.3B total


Equivalent: 16,500 books, 104 years reading time

#### **What BERT Learned:**

Language patterns, world knowledge, cultural contexts, emotional expressions, sarcasm patterns, domain vocabularies

#### **Training Cost:**

4 days on 64 TPUs (\$50-100K) - but reusable for everyone!



charts/pretraining\_scale.pdf

## From General Knowledge to Your Reviews

### Your Fine-tuning Process:

#### 1. Start with Pre-trained BERT

Has general language understanding; knows emotions, sarcasm, context; but doesn't know YOUR products

#### 2. Prepare Your Data

1,000 labeled reviews; Positive/Negative/Neutral labels; your product-specific language

#### 3. Fine-tune (Transfer Learning)

Show BERT your reviews; it adjusts parameters slightly; learns domain specifics while keeping general knowledge intact

#### 4. Time & Cost

2-3 hours on GPU; \$10-50 cloud compute; as few as 500 examples needed

### Example: Learning Company Slang

#### Before Fine-tuning:

"This app is sick!" → Negative (illness)

#### Your Training Data:

"Sick UI design!" → Positive

"Sick features!" → Positive

"Sick performance!" → Positive

#### After Fine-tuning:

"This app is sick!" → Positive (cool)

## Transformers Solve the Context Problem

`charts/performance_comparison_final.pdf`

### What BERT Found in 50,000 Reviews

charts/emotion\_taxonomy.pdf

#### 6 Core Emotion Clusters:

1. **Frustration (28%)** - Wait times, complex UI, missing features
2. **Delight (22%)** - Unexpected features, beautiful design, fast
3. **Confusion (18%)** - Unclear instructions, hidden functions
4. **Trust (15%)** - Data security, reliable service, transparent
5. **Disappointment (12%)** - Unmet expectations, quality issues
6. **Satisfaction (5%)** - Met expectations, good value

## When and Why Emotions Shift

charts/user\_journey\_emotions.pdf

## Which Emotions Predict User Loss?

charts/churn\_priority\_matrix.pdf

### High Impact on Churn:

**1. Frustration + Confusion (72% quit)**

“Can’t figure it out and support doesn’t help”

**Design Action:** Better onboarding + in-app help

**2. Disappointment + Distrust (64% quit)**

“Not what was promised, feels sketchy”

**Design Action:** Align marketing with reality

### Low Impact on Churn:

**3. Minor Frustrations (8% quit)**

“Annoying but I deal with it”

**Design Action:** Fix in regular updates

## Real Users, Not Imagined Ones

### The Enthusiast

15% of users

"Love it!", "Game changer"

Delight 78%, Anticipation 22%

**Need:** Advanced features

### The Struggler

35% of users

"Trying to...", "Can't find..."

Confusion 61%, Frustration 39%

**Need:** Better guidance

### The Pragmatist

30% of users

"It works", "Does the job"

Satisfaction 82%, Neutral 18%

**Need:** Reliability

### The Critic

20% of users

"Should have...", "Missing..."

Disappointment 54%, Frustration 46%

**Need:** Feature parity

**These personas emerged from BERT clustering - not designer assumptions**

---

Language-based segmentation reveals latent groups - behavioral patterns emerge from expression style rather than demographics

## Personalized Understanding, Automated Delivery

### The Scale Challenge:

1M users  $\times$  100 reviews = 100M opinions (impossible to read manually)

### BERT's Solution:

Process all in 24hrs, understand individually, group by emotional need, generate targeted responses

### Personalization:

Frustrated  $\rightarrow$  Support offer; Confused  $\rightarrow$  Tutorial; Delighted  $\rightarrow$  Beta invite; Disappointed  $\rightarrow$  Feedback survey

`charts/empathy_scale_pyramid.pdf`

### Emotion-Driven Engagement at Scale

#### The Challenge:

400M users; make each feel special; drive social sharing; increase engagement

#### NLP Analysis Revealed:

Users want validation; nostalgia drives sharing; uniqueness matters; discovery excites

#### Design Response:

Emotion words (“Rebellious”); unique stats (“Top 0.5%”); nostalgia (“Played 47x in March”); social proof

`charts/spotify_wrapped_emotions.pdf`

**Results:**

## Apply These Techniques to Real Data

### Workshop Exercise (45 minutes):

#### Dataset:

5,000 app store reviews (your choice of category), mix of 1-5 star ratings, real user language

#### Your Tasks:

1. Load BERT → 2. Fine-tune on 500 → 3. Analyze 4,500 → 4. Discover clusters → 5. Identify pain points → 6. Recommend designs

#### Tools:

Jupyter notebook, pre-processed data, loaded BERT, visualization functions

### Expected Outputs:

1. Emotion distribution chart (percentages)
2. Pain point priority list (by rating impact)
3. Data-driven personas (not assumptions)
4. Actionable design recommendations

### Learning:

Use BERT for emotions, interpret attention, convert ML to design, experience scale

**5,000 reviews → 5 key insights in 45 min!**

---

Practice bridges theory and application - implementing techniques reveals operational challenges theory abstracts away

`charts/nlp_method_decision.pdf`

## Three Levels of Hands-on Learning

### Exercise 1: Basic Sentiment Analysis

**Time:** 20 minutes

**Difficulty:** Beginner

**Task:**

Use pre-trained model; analyze 100 reviews; classify positive/negative; calculate accuracy

**Learning Goal:**

Understand basic NLP pipeline

**Tools:**

TextBlob or Hugging Face pipeline

### Exercise 2: Intermediate Emotion Detection

**Time:** 45 minutes

**Difficulty:** Medium

**Task:**

Fine-tune BERT; 6-class emotions; analyze attention weights; visualize results

**Learning Goal:**

Work with transformers

**Tools:**

Transformers library; pre-labeled dataset

### Exercise 3: Advanced Full Pipeline

**Time:** 90 minutes

**Difficulty:** Challenging

**Task:**

Scrape real reviews; preprocess text; fine-tune model; generate personas; create dashboard

**Learning Goal:**

End-to-end implementation

**Deliverable:**

Emotion insights report

**Resources:** Notebooks at [github.com/ml-design-course/week3-nlp](https://github.com/ml-design-course/week3-nlp)

Progressive complexity builds competence - sequential skill development enables advanced capability through mastered fundamentals

### From Words to Design Insights

#### Technical Understanding:

Why context matters in language; how Bag of Words loses information; what attention mechanisms do; how BERT reads bidirectionally; why pre-training + fine-tuning works

#### Practical Skills:

Identify emotion in text at scale; use pre-trained models effectively; fine-tune for specific domains; interpret attention visualizations; convert ML outputs to insights

#### Design Applications:

Data-driven persona creation; emotion-based user journeys; priority matrices from ML analysis; scalable empathy systems; personalization strategies

#### Remember:

**The Goal:** Not to read faster, but to understand deeper

**The Method:** Selective attention to what matters

**The Result:** True user understanding at scale

## Next Week: Classification & Problem Definition

---

Technical capability transforms analytical scope - methods enable understanding previously inaccessible through manual approaches

## Key Takeaway

Understanding emotion at scale is not about reading faster—it's about attending smarter

**Next Week:** Classification & Problem Definition  
From emotions to actionable categories