BA2: Digital Korea

Week 9: Sentiment Analysis

Steven Denney

Korean Studies
Leiden University

April 13, 2026

Today's Agenda

1. Review & check-in: from meaning to feeling
2. Sentiment arithmetic: word + word + word = ?
3. What is sentiment analysis?
4. How dictionary methods work (3 steps + preprocessing)
5. Korean sentiment dictionaries (KNU)
6. Orange under the hood
7. Application: Moon Jae-in's tweets
8. Minimum Orange workflow
9. Looking ahead

Review & Check-In

From Meaning to Feeling

Week 8 asked what words *mean*; Week 9 asks how they *feel*

Last week (embeddings): words as vectors capture **meaning**.

This week (sentiment): words carry **feeling**, and the researcher decides in advance which feelings. Embeddings: *what does this word mean in context?*

Sentiment: *does this word feel positive or negative, yes or no?*

Sentiment Arithmetic

Sentiment Arithmetic: Word + Word + Word = ?

Positive words add, negative words subtract, neutral words count for nothing

Each positive word adds $+1$, each negative adds -1, words not in the dictionary are ignored. The sum is the tweet's raw sentiment score. Negative examples work the same way, just with the signs flipped.

Your Turn: Guess the Sentiment

What is the sentiment of each combination?

1. 행복 (happy) + 사랑 (love) + 희망 (hope) = **?**
2. 공포 (fear) + 고통 (pain) + 슬픔 (sadness) = **?**
3. 감사 (thanks) + 어려움 (difficulty) + 희망 (hope) = **?**
4. 기쁨 (joy) + 위기 (crisis) + 신뢰 (trust) = **?**

Think about it

Cases 3 and 4 are **mixed**. This is where dictionary methods get interesting — and where their limitations start to show.

Sentiment Arithmetic: Answers

#	Words	Hits	Result
1	행복 + 사랑 + 희망	3 pos, 0 neg	**Positive (+3)**
2	공포 + 고통 + 슬픔	0 pos, 3 neg	**Negative (−3)**
3	감사 + 어려움 + 희망	2 pos, 1 neg	**Positive (+1)**
4	기쁨 + 위기 + 신뢰	2 pos, 1 neg	**Positive (+1)**

Key insight

Dictionary methods treat every word independently. They don't understand *context*: "overcome a crisis" and "cause a crisis" would score the same for 위기.

What Is Sentiment Analysis?

What Is Sentiment Analysis?

Sentiment analysis = determining whether a text expresses a positive, negative, or neutral attitude.

Examples in everyday life

- Product reviews (5 stars?)
- Social media (happy? angry?)
- News coverage (favorable? critical?)

Why it matters for research

- Track public opinion over time
- Compare tone across leaders
- Detect emotional framing

Two Main Approaches

Dictionary-based (this week)

- Humans write the word lists
- Computer looks words up, counts hits
- Transparent, no training data needed

Machine learning-based

- Train a model on labeled examples
- Captures context and nuance
- Less transparent ("black box")

This week: dictionary methods — the researcher sets the rules; the computer applies them at scale.

Step 1: How a Word Gets a Value

A sentiment dictionary is two word lists

positive.txt

감사 행복 희망
사랑 축하 발전
… 4,868 words

negative.txt

위기 걱정 파괴
고통 어려움 가난
… 9,824 words

Researchers decide, word by word, which carry positive feeling and which negative. Some dictionaries are hand-annotated; others are built computationally (e.g. by propagating labels across a multilingual knowledge graph). Either way, the output is two lists.

A word is **positive** if it appears in `positive.txt`, **negative** if in `negative.txt`, otherwise it doesn't count.

Step 2: Matching Words in a Document

For every content word, ask two yes/no questions

Input: 국민 여러분께 감사드리며 희망을 함께 만들어 갑시다
(People of the nation, thank you — let's build hope together)

Kiwi keeps 5 content words: 국민 감사 드리 희망 만들

Word	In a list?	Result
감사 (thanks)	positive.txt	**positive hit**
희망 (hope)	positive.txt	**positive hit**
국민, 드리, 만들	neither	no match

Totals: **2 positive**, **0 negative**, across 5 content words.

Step 3: Turning Matches Into a Score

One simple formula — the only one at play in Week 9

$$\text{score} = 100 \times \frac{\text{positive hits} - \text{negative hits}}{\text{number of content words}}$$

Applied to our 5-word example: $100 \times \frac{2 - 0}{5} = +40.00$

Why divide by length?

Long tweets would always look more extreme than short ones — more words, more chances to match. Dividing by length turns the score into a **rate**: share matched positive minus share matched negative, $\times 100$.

Preprocessing: What Counts as a "Word"

For sentiment we keep verbs and adjectives too — not just nouns

Korean is agglutinative: 감사드리며 packs several morphemes together. Kiwi splits it and tags each with a part of speech. We keep four tag types:

Tag	Meaning	Example	English
NNG	Common noun	감사	gratitude
NNP	Proper noun	한국	Korea
VV	Verb stem	이기	to win
VA	Adjective	좋	good

Plus two filters: length ≥ 2 characters; drop $\sim$20 very common stopwords (하다, 있다, 것, 수, …).

Why keep verbs and adjectives?

Feeling lives in action and description as much as in things. 행복하 (happy), 사랑하 (love), 걱정 (worry) would be invisible if we kept only nouns. For topic modeling we often keep nouns only — sentiment needs the fuller picture.

Korean Sentiment Dictionaries

Korean Sentiment Dictionaries: The Landscape

A few Korean lexicons are publicly available

Dictionary	Notes
KNU (Park et al. 2018)	**Course dictionary.** 14,854 words, polarity lists we load into Orange
KOSAC (Jang et al. 2013)	SNU, manually annotated on news, intensity + class
Chen & Skiena (2014)	Built into Orange's Multilingual option (Korean included)

Other options exist (NRC Emotion Lexicon, KSenticNet) — useful if you keep working with Korean sentiment.

The KNU Korean Sentiment Lexicon

Published by Park et al. (2018) at Kunsan National University

KNU annotators read Korean dictionary entries, emoticons, neologisms, and other lexicons, and decided word by word which carry positive feeling and which carry negative. We use the two plain word lists they published:

File	Words	Contents
`positive.txt`	4,868	Words KNU labelled positive
`negative.txt`	9,824	Words KNU labelled negative

These are KNU's own polarity lists with the header stripped — no intensity scores, no modifications by us.

What's in the KNU Word Lists?

A peek inside — these are the words the widget will look for

`positive.txt` (samples)

- 감사 (gratitude)
- 행복 (happiness)
- 축하 (celebration)
- 사랑 (love)
- 희망 (hope)
- 발전 (development)

`negative.txt` (samples)

- 가난 (poverty)
- 파괴 (destruction)
- 위기 (crisis)
- 걱정 (worry)
- 어려움 (hardship)
- 고통 (suffering)

Plain text, one word per line, no header. You can open these files in any text editor.

Orange Under the Hood

Orange's Sentiment Analysis Widget: Two Options

Same scoring formula, different word lists

Option 1: Multilingual (built-in) ← we use this

- Method = Multilingual, Language = Korean
- Orange ships with its own Korean word lists (Chen & Skiena 2014)
- Zero setup, zero files to download

Strengths: one click; no file management; works immediately.

Weaknesses: built by propagating English sentiment labels across a multilingual graph — not Korean-native; smaller Korean vocabulary; may miss domain-specific words.

Option 2: Custom Dictionary (KNU)

- Method = Custom Dictionary
- You provide `positive.txt` and `negative.txt`
- KNU: 14,854 Korean words, human-annotated

Strengths: Korean-native; larger coverage; pairs well with Kiwi morphological preprocessing; better for research use.

Weaknesses: requires downloading files and a Python Script step; more moving parts.

Under the Hood: What the Widget Actually Does

No black box — the same 3 steps as before, now inside Orange

When you connect a Corpus to Sentiment Analysis (Custom Dictionary + `positive.txt` / `negative.txt`), here's what happens to **every document**:

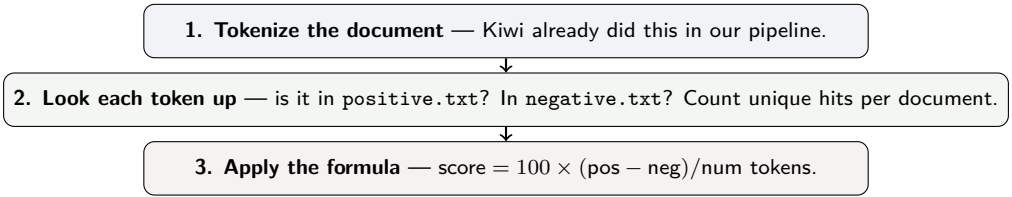

Output: a new `sentiment` column, one number per tweet. Connect it to a Box Plot, Distributions, or Line Plot.

One Method, Two Places to See It

Same pipeline in Orange and in the interactive — numbers match

In Orange (you build this)

- 5 widgets + 1 Python Script
- Custom Dictionary + KNU
- Ends at a Box Plot by period

In the interactive (pre-computed)

- 6 panels to click through
- Same Kiwi + KNU pipeline
- You can read individual tweets

Orange gives you the tool. The interactive opens it up so you can see what the tool is doing to your tweets.

Application: Moon Jae-in's Tweets

Case Study: Moon Jae-in on Twitter

3,148 tweets from @moonriver365 (2012–2020), scored with KNU

The corpus

- 3,148 tweets, 2012–2020
- Most active in the 2012 campaign
- Tweeting dropped after he took office (2017)
- Topics: inter-Korean relations, COVID, Japan, democracy

Research questions

1. Is the tone positive or negative overall?
2. Does sentiment change across periods?
3. Do events leave visible marks?

Three Political Periods: My Editorial Split

The `period3` column is my grouping, not an official classification

Period	Dates	Tweets	What's happening
Pre-presidency	2012-01 to 2016-11	1,973	Opposition leader — 2012 campaign, DP chairmanship, legislative politics
Transition	2016-12 to 2017-05	393	Park Geun-hye impeachment crisis through Moon's campaign and inauguration
Presidency	2017-05 to 2020-06	782	In office — inter-Korean summits (2018), Japan trade dispute (2019), COVID-19 (2020)

A research choice, not a fact

I drew the boundaries at the impeachment vote and the inauguration. A different researcher might cut the timeline differently — by election cycle, Twitter-volume phase, topic. **The categories are part of your analysis, not a given.**

Sentiment Arithmetic with Real Tweets

A purely positive tweet: celebrating a birthday

Tweet (2018-01-24, presidency):
생일 축하, 고맙습니다. 생일을 챙기지 않는 삶을 살아왔는데 ⋯ 진심으로 감사드립니다.
(Thank you for the birthday wishes. I've lived without celebrating birthdays … I am truly grateful.)

Matched word	In dictionary
축하 (celebration)	**positive**
감사 (gratitude)	**positive**

2 positive matches, 0 negative. Orange's formula: $100 \times (2 - 0)/14$ tokens $= +14.3$.

Sentiment Arithmetic with Real Tweets

A mixed-sentiment tweet: grief, solidarity, hope

Tweet (2020-03-30, presidency):
어려움 속에서도 서로를 격려해가며 신뢰와 협력으로 재난을 이겨가고 있는 국민들께 한없는 존경과 감사를 드립니다.
(Even amid hardship, I offer deepest respect and gratitude to the people enduring this disaster with trust and cooperation.)

Matched word	In dictionary
신뢰 (trust)	positive
존경 (respect)	positive
감사 (gratitude)	positive
어려움 (hardship)	negative
재난 (disaster)	negative

3 positive, 2 negative in 20 tokens: $100 \times (3-2)/20 = +5.0$. The tweet leans positive because gratitude and respect outweigh the hardship words.

Moment in Time: The Impeachment Vote

Tweet (2016-12-09, transition, 4,445 favorites):

국민이 이겼습니다. 대통령 탄핵은 끝이 아니라 시작입니다.

(The people have won. The presidential impeachment is not the end but the beginning.)

Dictionary scoring:

- Matched words: **0**
- 국민, 이기, 대통령, 탄핵, 시작 — none of these are in KNU
- Score: **0** (neutral)

Context:

- Posted day after the impeachment vote
- Tone is triumphant
- Dictionary says "neutral" — the method misses the political significance entirely

Critical Reading: What the Dictionary Misses

A negative-scored tweet that is rhetorically assertive

Tweet (2019-08-02, presidency, 19,110 favorites):

우리는 다시는 일본에게 지지 않을 것입니다. 수많은 역경을 이겨내고 ··· 어려움을 극복할 역량이 있습니다.

(We will never again lose to Japan. We have overcome many adversities ... we have the capability to overcome difficulties.)

Dictionary scoring:

- 어려움 (difficulty) → **negative**
- 역경 (adversity) → **negative**
- 극복, 역량, 일본 — not in KNU
- Score: -16.67 (2 neg hits / 12 tokens)

Human reading:

- Tone is defiant and assertive
- Negative *toward Japan* specifically
- Positive *self-framing* (national strength)
- The dictionary sees just "difficulty"

Lesson

Dictionary-based sentiment measures **word-level valence**. It doesn't capture the **target** (negative toward whom?) or the **rhetorical purpose** (rallying, not despairing).

Orange Demo

Interpreting Your Results

Once you have sentiment scores, ask:

1. **Distribution.** What's the overall shape? Skewed positive, neutral, negative?
2. **Outliers.** Read the most extreme tweets. Do they make sense?
3. **Patterns.** Does sentiment track time, topic, or event?

Always validate

Grimmer, Roberts & Stewart: "All quantitative models of language are wrong — but some are useful." Read 10–20 tweets across the score range. Does the algorithm agree with you?

Looking Ahead

For Next Week

Week 10: Topic Modeling (LDA)

- Latent Dirichlet Allocation (LDA)
- Discovering hidden thematic structure in a corpus
- Choosing the number of topics
- Interpreting and labeling topics

Recommended reading:

- Grimmer, Roberts & Stewart — Chapter 13: Topic Models

For Next Week

Interactive, assignment, preparation

Explore. Click through all 6 panels of the *Sentiment Analysis: Moon Jae-in's Tweets* interactive on the course website.

Build. The minimum Orange workflow on `moon_twitter.csv` — ending in a Box Plot by `period3`.

Write. 2–3 sentences: which period is most positive, and which matched words from the interactive explain why?

Download from the Data & Scripts page

`moon_twitter.csv` — that's all you need for the Multilingual option.

KNU dictionaries (`positive.txt`, `negative.txt`) and the preprocessing script are also there if you want to compare.