

Social sentiment analysis: A guide to LLM automation

Social media is buzzing 24/7, and keeping track of how people feel about your or your competitor's brand can feel overwhelming. Many social media pros already have sentiment analysis in their toolkit to save time and understand their audience at scale.

Some use off-the-shelf tools, while others combine multiple third-party solutions. With the rise of large language models, however, setting up your own data pipeline has become easier than ever. But how do you do it? Can you build a reliable social media data extraction and processing system without solid data analysis skills?

Turns out you can. In this article, we'll explore important considerations, essential tools, and a step-by-step guide to setting them up.

What is sentiment analysis?

At its heart, sentiment analysis tries to determine whether a piece of text is expressing a positive, neutral, or negative attitude. It can be used on everything from short social media posts and product reviews to long-form articles and internal company documents.

The point of automating it, though, is that you don't have to **manually sift through thousands of comments, reviews, or messages** to understand the general picture. Instead, sentiment analysis tools can process and categorize vast amounts of text in real-time.

As a social media marketer, you can use it for short- and long-term goals: **monitoring brand perception, tracking how audiences react to campaigns, responding to customer feedback.** It also helps with identifying potential PR disasters early, understanding customer sentiment over time, and changing content strategies based on audience reactions. And of course, needless to say, all of this can be applied to both your brand and your competitors' so you can be on top of everything.

LLMs for sentiment analysis: yay or nay?

So now that we know what social sentiment analysis is, how do we actually implement it?

Traditional sentiment analysis relies on controllable machine learning or rules-based approaches. You collect the data and define all the rules yourself — which is good if you're looking for transparency and control. However, this typically requires a team of data experts to set up and maintain. They would use various libraries like VADER, TextBlob, or spaCy, and the process would involve training models, fine-tuning parameters, and continuously updating the system.

Access to large language models (LLMs) have become a shortcut to this complexity. Ever since ChatGPT entered the scene, people realized that models like it can be used for sentiment analysis without extensive manual rule-setting. Arguably, LLMs can also interpret context at a deeper level, detect sarcasm, subtle shifts in tone, and domain-specific phrases that traditional models might miss. They can also work with multiple languages.

Here are a few solid **reasons why people use LLMs for sentiment analysis**:

- **Faster deployment and integration with other tools.** Unlike traditional sentiment models that require extensive training, dataset curation, and fine-tuning, LLMs are ready to use almost instantly through APIs. This drastically reduces setup time.
- **Great vocabulary.** Since they're trained on vast and diverse datasets, LLMs can recognize slang, idioms, and niche industry jargon that traditional models might miss.
- **Multilingual capabilities.** Many LLMs support multiple languages out of the box, often without requiring specific translations or retraining, which makes them useful for global-minded brands and businesses.
- **Context sensitivity.** LLMs can analyze longer text passages and maintain awareness of how meaning shifts across multiple sentences, which makes them better at handling complex discussions or opinion pieces.
- **Adaptability.** Whether analyzing product reviews, financial reports, or social media comments, LLMs can adjust to different contexts with minimal customization, making them useful across industries.

That said, relying solely on LLMs for sentiment analysis isn't always ideal. Because LLMs operate on pre-trained knowledge and predefined rules set by their creators, there are certain pitfalls to watch out for — such as biases in training data, lack of explainability, potential overgeneralization. Stuff you need to watch out for when making decisions.

In the next section, we'll learn how to set up a practical pipeline

How to set up an LLM automation for sentiment analysis

1. Choose your data sources

First, you have to **decide which data sources matter most**. Are you analyzing customer reviews? Yours or competitors'? Is it going to be industry articles? Forums? Social media posts? Which social media platform is it going to be? Knowing your end goal helps you pick the right scope and format.

2. Collect social media data

Sentiment analysis is only as good as the text it processes, so a reliable set of tools to collect relevant content will make or break your sentiment analysis pipeline. This could mean pulling website data from publicly available APIs, using or even creating data extraction tools — or web scrapers.

Why scrapers, you might ask? One of their biggest advantages is that they get data in real time — which is what you need for sentiment analysis. However, the web scraping landscape is vast, with solutions ranging from simple browser extensions to complex, custom-built automations — depending on the complexity of the scraped website. The key factor for you will be convenience — having **all the necessary scrapers in one place with an accessible API**.

One of the easiest ways to collect comments and posts at scale is with [Apify social media scrapers](#). These tools automate data extraction from major platforms and produce high-quality input for your sentiment analysis.

With high-quality data in place, you're ready to start processing sentiment and extracting valuable insights.

3. Clean your data

As a sidequest — with some APIs — you will have to remove duplicates, deal with missing information, or filter out anything that doesn't fit your project. You might also have to standardize letter casing, strip special characters or HTML tags, and refine your data so it's easier to process. Sometimes you would have to use a combination of automated labeling, third-party annotation services, or in-house tagging to build a reliable dataset.

4. Choose the LLM

Select the AI model or tool that fits your needs. You can choose from a variety of LLMs, from OpenAI's ChatGPT to Anthropic's Claude and Google Gemini.

Which LLM is best for sentiment analysis? There's no single "best" model, but some stand out:

- **GPT-3/GPT-3.5** – Handles complex language with minimal training but may require more resources.

- **Claude (Anthropic)** – Strong at contextual understanding and safety-focused, making it a good option for nuanced sentiment analysis.
- **Google Gemini** – Multimodal capabilities with strong language comprehension, useful for diverse sentiment tasks.
- **Llama, Falcon, Bloom** – Open-source, customizable, but needs setup.
- **BERT-based models (RoBERTa, DistilBERT)** – Efficient for short reviews with fine-tuning.

For quick sentiment analysis, a fine-tuned BERT model works well. For multilingual or nuanced sentiment detection, GPT-3.5, Claude, or an open-source alternative may be better. You'll need to consider accuracy, F1 scores, and inference speed to find the right fit.

5. Automate the workflow

Now, you need to automate the workflow. Ideally, the **data should flow directly from the scrapers to LLM**. Most LLMs have an API key, which you can use to integrate the data, and even push it further down the pipeline with different tools like Make, Zapier, or n8n. Some platforms such as Apify can unite the whole process for you.

The screenshot shows the Apify web interface for a workflow. At the top, there's a header for 'Tiktok Comments Scraper' with a price of '\$15.00/month + usage' and a 'Start' button. Below this, there's a description: 'Effortlessly scrape Tiktok data: comments, users, and all the other insights. Overcome API limitations, and ensure real-time monitoring for seamless extraction.' A navigation bar includes 'Input', 'Information', 'Runs 0', 'Builds 77', 'Integrations 1', 'Monitoring', 'Issues 0', 'Saved tasks 0', and 'Admin settings'. A sub-header states: 'The integrations set up here will trigger for all runs of this Actor, including its saved tasks.' Below this is a '+ Add integration' button. The main area displays a workflow diagram with two nodes: 'Tiktok Comments...' (epctex/tiktok-comment-...) and 'LLM Dataset Processor...' (dusan.vystrcil/llm-...). A blue arrow points from the first node to the second, indicating the data flow.

6. Analyze and visualize results

Once you get the analysis from your chosen LLM, it's ready for you to take a look and use the insights in your strategy. Sometimes, you might want to **visualize the results in graphs and tables for easier interpretation and presentation**. For this, use tools that you already know — Tableau, Power BI, or even Google Sheets are powerful enough.

Pro tip: Include the visualization tools in your workflow to make sure the data flows directly from web scrapers to LLM and from LLM to these tools.

Want to know how to automate sentiment analysis step-by-step? You can read the full setup with Apify's scrapers and LLM Dataset Processor here:

<https://blog.apify.com/sentiment-analysis-tools/>

If we're to leave you with just one insight from this guide, it's **that automating social sentiment analysis with LLMs is more accessible than ever**. Whether you're monitoring brand perception, analyzing competitors, or identifying trends, a well-structured pipeline will help you stay ahead in an ever-changing web landscape.

So if you're ready to set up a sentiment analysis pipeline — start by choosing the data overview you want to achieve, then collect that with data extraction tools, and plug the data into a fitting LLM. The right combination of tools can turn overwhelming amounts of social media chatter into actionable insights. Good luck!

