PRESS RELEASE · for publication from 10 November 2026, 8:30 CET

Poland is in the answers of language models when it is in the question. “The Generative Portrait of Poland” is the first measurement of how six models talk about Poland to individuals and companies from 12 countries

Warsaw, 10 November 2026. When a question does not name Poland, language models bring it in of their own accord in seven answers out of ten, but in questions about Europe “in general”, in fewer than one in two, and when they do bring it in, they mention it more often than they recommend it. To companies they talk about Poland almost always, even when nobody asks about it: it comes up in 98 answers out of 100. Its role, however, depends on the decision area, and in logistics Poland loses on geography, not on cost. These are the main findings of the report “The Generative Portrait of Poland”.

The first semantic map of Poland, measured rather than clicked through by hand. The study covers six generative AI models queried through their official APIs: ChatGPT, Gemini, Claude, Perplexity, Grok and DeepSeek. Between 25 August and 28 September 2026 the models received 1,400 questions in 11 languages, 700 on behalf of individuals and 700 on behalf of companies from 12 markets: China, Czechia, France, Germany, Italy, Norway, Slovakia, Spain, Sweden, Ukraine, the United Kingdom and the United States. The questions concern everyday decisions: where to go for a weekend, where to study, where to move, where to build a plant, whom to contract production to, how to vet a business partner. That produced 8,400 answers; each was assessed by the same model judge, and the results were checked through the five control layers described in the annex. The research programme as a whole comprises more than 14,500 answers from language models (LLMs). The report has two parts: “Poland for People” and “Poland for Business”.

Part I. Poland for People

Poland is in the answer when it is in the question. Poland appears in 95.1% of the 8,400 answers. When the question names it, the models talk about it in 99.9% of answers. The real test lies in the 1,452 answers to questions that do not name Poland: there it appears in 72.4%.

Visibility depends on the type of question, not on the model. Across 174 questions from individuals that do not name Poland, the models mention it in 71.3% of answers. Where the question points to a region that includes Poland, the figure is 91.7%; where it concerns Europe in general, it falls to 45.1%, and without Ukraine, for which Poland is the obvious destination, to 38.5%. In the models’ “default Europe”, Germany heads the list (58.0%), while Poland (45.1%) sits in a cluster with Czechia, Portugal and Spain. The differences between models are smaller than those between question types.

Mentioned more often than recommended. When a model brings Poland into an answer of its own accord, it explicitly recommends it in 55.9% of cases; it almost never advises against it (1.7%), but in four cases out of ten it mentions Poland without recommending it. Here the question type works the opposite way from visibility: in questions about a region, Poland appears almost every time, yet is recommended in only 47.0%; in questions about Europe in general, it appears less often, but when it does, it is recommended in 62.3%. The model brings it in when it has an argument for it. The models' tone about Poland is moderately positive (6.95 on a 1–10 scale) and almost identical across all markets.

The canon appears on request, but is almost absent from practical answers. Asked directly about the best-known Poles, the models answer with the canon: Chopin, Wałęsa, Nicolaus Copernicus and John Paul II appear in every one of the 144 answers. In practical questions about travel, history and culture, the same canon barely features (from 12.0% of answers for Chopin down to 0.8% for Skłodowska-Curie). The report calls this difference the activation gap: the canon enters a practical answer only through the doors the question opens (a place to see or a chapter of history).

Polish URLs come from the model itself more often than from search. Only one of the six models searched the web live. Of the URLs the models supply themselves, half lead to Polish sites (49.5%) and one in three to official ones (33.5%). Sources from search account for 57% of all citations, but only 22.4% of them lead to Polish sites (6.5% to official ones). Polish institutions are present in the models' answers, but search rarely leads to them. In questions about the cost of living in Poland, the answers point to the numbeo.com website, and in questions from the United Kingdom to gov.uk more often than to gov.pl.

Asked about the war, the models do not advise against Poland, but almost always add a "but". In a separate thematic measurement taken in October, none of the 792 answers to individuals advises against Poland as a destination because of the war in Ukraine; 15.7% recommend it without any caveat. In a question about safe countries across Europe as a whole, Poland is placed in the "with caution" group in between half and three quarters of answers. NATO appears in 78.1% of answers, Schengen in 5.3%.

Part II. Poland for Business

The models put Poland forward to companies of their own accord. When a company asks without naming Poland, the models mention it in 98.4% of the 576 shortlist answers and explicitly recommend it in 82.5%. What matters is the decision area: in questions about a sales market, Poland is recommended in 93.1% of answers; about engineers, in 93.0%; about contract manufacturing, in 69.4%; about a warehouse for Europe, in 62.3%.

Poland loses on geography, not on cost. In questions about sales markets, incentives and service centres, its competitors are its regional neighbours: Czechia, Romania, Hungary and Slovakia. In questions about suppliers, manufacturing and logistics, Germany, Italy, Austria and the Netherlands join the list. In answers about a warehouse for the whole of Europe, the models know Polish locations (53 of 72 answers), but the hubs are the Netherlands and Germany, because from there goods reach the west in two to three days; cost is Poland's argument in every one of these answers. To suppliers, the models describe Poland as a base that is growing, and Germany as reliable.

The corporate canon appears on request, but in use the registers win. Asked about the largest Polish companies, the models name KGHM in every one of the 144 answers. In the 4,200 answers to companies' questions, KGHM appears in 21 and PZU in one; the leaders are Krajowy Rejestr Długów, a Polish debt register (143 answers), BIG InfoMonitor, a credit-information bureau (97), and PSE, the transmission-system operator (95), because they address the company's situation. A company that asks for names in an industry gets a Polish company in two answers out of three (99 of 150), but when it asks for certified subcontractors, only in 5 of 30. Aviation, electronics and pharmaceuticals appear rarely even on request (in one in four, one in six and one in fourteen answers about the strongest industries), and the study asked no questions about companies in these sectors.

Polish institutions are in the models' answers, but search in the market's language rarely leads to them. Of the URLs the models give companies themselves, 54.5% lead to Polish sites (35.0% to official ones); of the sources from live search, 25.0% (6.9%). The Polish Investment and Trade Agency (PAIH) is among the three most-cited domains in each of the 12 markets, and in questions about incentives it appears in 26 of 27 answers with URLs given in the answer and in 3 of 12 with sources from search. A separate measurement from October shows the same thing from the other side. The same model answered each question twice: once without the tool, that is, without going to the internet at the moment of answering, and once in search mode, that is, after searching the web. With search, Poland does not appear in the answer more often, and Polish sites are just as rare among the sources as among the URLs given without the tool. For some of the models, search mode runs with a different instruction, so the report reads this difference with caution (Part II, Chapter 6).

The war: Poland recommended with conditions. Asked directly about investment, nearshoring, energy or insurance in Poland, the models did not advise against Poland as a choice because of the war in any of 864 answers, but only 3.8% of answers to questions about investment risk recommend it without any caveat, and some answers point boards that are sensitive to risk to a country further west, most often Czechia. As a base for the reconstruction of Ukraine, Poland is recommended in every answer. Companies hear different arguments from individuals: the European Union (86.9% of answers) and energy (75.0%) instead of Belarus and ministry warnings.

The models disagree, and none of them is "the best for Poland"

The six models agree on whether to talk about Poland in 93.4% of the 1,400 questions; where they disagree is on how: on the same question, the warmest and the coolest model differ by 1.8 points of tone on average. How the models line up depends on what we count: five of the six come first or last on at least one measure. Newer versions of the models shifted the share of explicit recommendations of Poland by a dozen or so points within seven weeks, in both directions. Manual verification confirmed 14 of the 19 factual errors flagged by the judge in 8,400 answers; errors the judge did not flag were not estimated.

What this means for the promotion of Poland

The goal is not "for the models to recommend Poland" (to companies they already do) but presence in the specific types of question and decision areas in which Poland is to be a candidate rather than a backdrop. Where Poland appears rarely, it is recommended more often, so promotion should supply the models with arguments rather than presence. The canon is capital that can be activated immediately, but only through doors: linking an entity with a place, a route, a decision

situation. In industrial areas, Poland needs evidence of maturity, not of growth: certifications, references, the names of subcontractors in the languages of the markets. Poland's image in the answers changes faster than campaigns do; this semantic drift has to be measured at regular intervals, with the same set of questions, because a one-off audit shows only how things stood.

> **Grzegorz Miłkowski, author of the report:** "The models do not speak ill of Poland. To individuals they often simply say nothing about it, and when they do speak of it, they mention it rather than recommend it; to companies they almost always speak of it, but in a role determined by the question. This is the first map of how Poland looks in the models' answers, measured on 8,400 answers, not clicked through in a chat window."

What the report does not say

The report measures the models' answers returned through the API within a dated window (main study 25 August – 28 September 2026; control and additional measurements 30 September – 2 November 2026), not what the consumer app shows, not website traffic and not the effectiveness of any campaign. It does not forecast investment decisions and does not assess the work of any institution. It does not rank markets or models. It does not combine figures from the study with those from the shortlist, because these are two sets with separate denominators. It does not compare individual models' recommendation rates in the main text, because four of them were subject to a length cap on answers in the study (annex). It is a dated snapshot, a map of the representation of Poland in the answers of six models to questions from individuals and companies in 12 markets, and it should be read as such.

About the report

"The Generative Portrait of Poland" is a semantic map of Poland: a record of the places, names, companies, arguments and sources that six language models associate with the country when they answer individuals and companies from 12 countries. The report is a single document in two parts (Part I, "Poland for People", and Part II, "Poland for Business") with a methodological annex and open data (personas and question scenarios). The report is published in Polish ("Polski portret generatywny") and in English, in PDF format. For editors, there are two extracts (Part I and Part II), a "Key findings and quotes" document for each part, and 12 local notes, one for each market.

Knowledge partner: AI Business Center Foundation (aibusiness.center).

Under the patronage of Business Centre Club.

About the author and the publisher

The author of the report is Grzegorz Miłkowski, founder and CEO of Brand Semantics P.S.A., an advisory and technology firm that works on brand semantics and its measurement in language models. Grzegorz Miłkowski is also President of the Management Board of the AI Business Center Foundation, the report's knowledge partner. The measurement was carried out in Semantio, the company's own research application, which queries the models through their official APIs and assesses the answers with the same judge within a dated window.

Report, press kit and open data:
https://semantio.pro/en/reports/generative-portrait-of-poland/

Interactive semantic map of Poland:
https://semantio.pro/en/reports/generative-portrait-of-poland/map

---

Media contact: Michał Grzebyk, COO, Brand Semantics, m.grzebyk@brandsemantics.eu, +48 782 010 316.