The Generative Portrait of Poland. Part I, “Poland for People”. Press-kit extract

Poland features in the answers of language models (LLMs) when it features in the question. When the question does not name Poland, it appears in seven in ten answers, but in questions about Europe in general, in fewer than one in two, and when it does appear, it is mentioned more often than recommended. It is the first map of its kind, measured rather than clicked through by hand: this is Poland as it emerges from 8,400 answers given by six models to questions from 12 countries.

What we studied. 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 we put 1,400 questions to them 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 were asked on behalf of personas, scenarios prepared for the study, and 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 makes 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. Part I reads these answers from the perspective of individuals, Part II from that of companies (700 questions, 4,200 answers). Because companies almost always named Poland in their questions, Part II also includes a separate shortlist measurement of 96 questions that do not name Poland, across eight decision areas (576 answers, 30 September). There are also two thematic measurements: Poland on request (the largest companies and industries) and Poland in the shadow of the war. The research programme as a whole comprises more than 14,500 answers from language models (LLMs). The subject of the report is the brand semantics of Poland: the entities (places, people, events, traits) that the models link with the name “Poland”, and the ways those entities connect to the situations that people and companies ask about.

Part I. Poland for People

1. Poland is in the answer when it is in the question. Poland appears in 95.1% of the 8,400 answers. What that figure shows is that the models stay on topic: when a question names Poland, they 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, Poland appears in 72.4%.

2. Visibility depends on the type of question, not on the model. Across 174 questions in which individuals are looking for a destination without naming 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, 45.1%; where it refers neither to a region nor to a trip, 20.0%. In the models’ “default Europe”, Poland appears in fewer than half of the answers, and without Ukraine, for which it is the obvious destination, in fewer than four in ten (38.5%). The differences between models are smaller than those between question types.

3. Poland competes with a list, not with a single country. In answers to questions that do not name Poland, the models list almost six countries on average. In questions about a region, Poland comes first (91.7%), ahead of Czechia (80.6%) and Hungary (73.2%). In questions about Europe in general, Germany heads the list (58.0%), while Poland (45.1%) sits in a cluster with Czechia, Portugal and Spain. One answer in five (21.1%) names the persona's home country.

4. 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 models differ more in recommendation rate than in mention rate. Asked directly about safety during the war in Ukraine (a separate measurement taken in October), in none of 792 answers do the models advise against Poland as a destination, but they almost always add a "but"; asked which countries in Europe as a whole are safe, they place it in the "with caution" group in between half and three quarters of answers.

5. Polish URLs come from the model itself more often than from search. Only one of the six models, Perplexity, searched the web live; the others supplied URLs themselves. Half of those URLs 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). So Polish institutions are present in the models' answers, but search rarely leads to them.

6. The canon appears on request, but is almost absent from practical answers. Asked directly about Poland, the models answer with the canon: Chopin, Wałęsa, Nicolaus Copernicus and John Paul II appear in every one of the 144 answers about the best-known Poles. In practical questions about travel, history and culture, the same canon barely features: from 12.0% of answers (Chopin) down to 0.8% (Skłodowska-Curie). We call this difference the activation gap. The canon enters practical answers only through the doors the question opens: a place to see (Chopin appears together with Warsaw or Żelazowa Wola in 97% of his mentions) or a chapter of history. The Polish Tourism Organisation (POT) campaign "Poland. More than you expected" builds on faces from the canon and on faces that are absent from the answers; the measurement was taken eight weeks after the campaign's launch and does not assess its effectiveness.

7. 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; the disagreement lies in the questions that do not name Poland. Where they do disagree is on how to talk about it: on the same question, the warmest and the coolest model differ by 1.8 points of tone on average, which is roughly twice what measurement variability alone would produce. 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.

8. Twelve portraits, no ranking. Chapter 8 reads the results market by market, using the same market card for each and without placing the markets in any order: Poland's visibility depends on what people in a given market ask about. Each portrait closes with a market-specific thread, that is, an association peculiar to one market (such as the Battle of Britain for the British or the Tatras

for Czechs), and a check on whether it makes its way into practical answers. Usually it does not; the Tatras and Gdańsk do, because they are places one can travel to. In all 12 markets, the answers also point to the authorities of the persona's home country. In the British market, gov.uk appears more often than gov.pl.

9. What this means for the promotion of Poland. The goal is not "for the models to recommend Poland" but presence in specific types of question. 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. A Polish source makes it into the answer when it addresses the situation of a person from a given market, in that person's language or in English. In questions about the cost of living in Poland, answers currently point to the numbeo.com website, while the tourism portal appears in 4.6% of answers, almost exclusively in English. 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.

What the report does not say. It measures neither what the consumer app shows, nor website traffic, nor 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, and the shortlist is not a check on the study. 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.