vinetur

Study Finds AI Tools Alter Conclusions on Wine Innovation

An analysis of 9,439 patents in 53 countries found some systems pointed to divergence, others to convergence.

Tuesday, August 18, 2026



A new study in the Journal of Wine Economics says the way researchers measure technological similarity in the wine business can sharply change the picture of how innovation has evolved across countries and product areas.

The paper, by Francesca Chiaradia of the University of St Andrews, examined 9,439 wine patent abstracts filed in 53 countries from 1970 to 2023. Using artificial intelligence tools that turn text into numerical vectors, the study compared five language models to see how closely wine inventions resembled one another over time.

The central finding was not just that wine technology has changed, but that different models often point to different trends.

That matters because technological similarity is used to track how ideas spread, where producers cluster around common methods, and where new invention is breaking away from older standards. In the wine trade, that can affect how companies benchmark competitors, spot emerging techniques, and assess markets across regions.

Chiaradia built the dataset from Lens.org, a public patent database, and focused on patent abstracts rather than claims. The paper argues that abstracts give a clearer description of an invention and contain less legal language. The study then processed the text and used five systems to generate sentence embeddings, a method that maps written language into a mathematical form so patents can be compared through cosine similarity scores ranging from 0 to 1.

The five models were GTE, S-BERT, USE, PaECTER, and Doc2Vec. According to the study, the choice among those models led to noticeably different readings of the wine patent record. GTE, S-BERT, and USE showed a decline in average similarity from the 1970s until roughly the late 1990s, followed by a flatter pattern later on. PaECTER and Doc2Vec, by contrast, suggested an increase in convergence.

The paper says those differences are large enough to affect any broad claim about whether wine innovation is becoming more standardized or more diverse. In other words, the answer depends in part on the model being used.

Across the full sample, the average semantic similarity also varied by model. GTE and PaECTER produced the highest overall averages, from 38.7% to 41.9%. S-BERT and USE were lower, at 33.2% and 28.9%. The study found that U.S. and Asian patents tended to show greater convergence than those from other regions, with scores of 46.5% to 46.6% in some comparisons.

The data also showed that patenting activity in wine technology is concentrated. The United States led individual countries with 1,516 patents, followed by South Korea with 1,453 and France with 903. Europe accounted for 4,196 patents and Asia for 3,032. Within Europe, France, Germany, and Spain held 63.1% of the region’s wine innovations

in the sample. In North America and Asia, a small number of countries also dominated filings. The United States, South Korea, and Japan accounted for 87.2% and 75.7% of the shares in their respective continents, the paper said.

The study grouped the patents into five broad technology areas and found that food chemistry and chemical engineering dominated, followed by machinery and mechanical elements. Europe showed relatively stronger specialization in wine measurement technology. North America stood out more in digital technologies and consumer products. In Africa and Oceania, a decline in chemistry-related patents was offset by greater focus on handling and packaging technologies such as bottles, corks, and storage.

One of the paper’s strongest regional findings is that similarity within countries was often higher than similarity across countries. The United States showed a long decline in the similarity of wine patent language, which became more stable in the early 2000s. Europe moved in the opposite direction later in the period. Similarity there rose after 2010 and spread across fields, especially in mechanical engineering and instruments.

The paper does not claim a direct cause, but it points to several possible drivers for Europe’s pattern, including lower trade barriers inside the European Union and the gradual harmonization of agricultural rules. The broader implication is that wine innovation remains closely tied to the economic and institutional setting in which an invention is created, even in a global industry with extensive trade.

The study also found sharp differences across technical areas. Technologies linked to biotechnology, organic chemistry, and chemical engineering showed a stronger drop in similarity from 1970 to 1990 in the models trained on larger and more refined datasets. After that, the results suggested more stability around the turn of the century, especially in standardized equipment such as engines, pumps, and other mechanical components. Consumer goods patents, which the paper describes as less technically complex, showed higher average similarity and more stable patterns after the mid-1990s.

To test whether the language models were producing credible results, the paper used three validation methods. First, it compared model predictions to International Patent Classification codes, which are formal labels assigned to patents. The results were mixed. Accuracy ranged from 55.4% to 58.1% for most models in the overall comparisons, and precision varied by field. GTE improved discrimination, while USE recorded stronger accuracy and precision in some cases. Even so, the paper says all models performed poorly when judged against standardized labels, which suggests that current tools still struggle to capture similarity beyond a basic level.

Second, the study looked at patent drawings. It downloaded images from patent filings and measured visual similarity using the Structural Similarity Index Measure, or SSIM. The average image similarity was 39.69%, a figure that aligned closely with the scores generated by GTE. Chiaradia says that match supports the view that more advanced, better-tuned models can provide a more reliable reading of technological convergence.

Third, the paper used a retrieval-augmented generation system tied to a large language model to test whether a dynamic system would trace similarity in a way that matched the static models. According to the study, the RAG-based results closely followed the patterns found by GTE, S-BERT, and USE, giving the author more confidence in the trends those models produced.

The paper places the findings in a wider debate over how economists and industry analysts use machine learning to read innovation data. Patent records are increasingly mined to identify industrial change, but the study warns against treating AI outputs as neutral or interchangeable. The same set of wine patents can tell different stories depending on the system used to analyze them.

The author also argues that the findings say something about the structure of the wine business itself. Higher similarity in some New World regions, especially in areas tied to software, packaging, online advertising, labeling, and quality control, may point to stronger standardization or tighter innovation clusters. At the same time, more sophisticated technologies in chemistry and mechanical engineering appear to remain more closely tied to the local conditions of production and research.

The study is descriptive and does not try to prove why those differences emerged. It also covers wine only, not the broader beverage market. Chiaradia says the results are limited by the current generation of language models and could change as newer systems are developed. The paper suggests that future research could extend the same approach to spirits and other alcoholic drinks, where patent activity and competition have also been growing.