

Vinum ex Machina: the winemaker gets forty tries, the machine gets a million

Nikita Khudov

PUBLISHED	EDITION	DOI	LICENCE	ONLINE
2026-08-13	v2026-09-29	10.5281/zenodo.23033982	CC BY-SA 4.0	vinumexmachina.com/keynote/

ABSTRACT

Irrigation, disease pressure and buyer behaviour are **already computed more reliably by a machine** than by the calendar and intuition, and every season it takes on more. What works today, what becomes normal within five years, and which part of the winemaker's work stays human once the algorithm has taken the rest.

Artificial intelligence in wine is written about much the way biodynamics was fifteen years ago: often, enthusiastically, and almost without figures. Robots between the rows, a neural network instead of a sommelier, a winery without people. So what actually works? At Vinum ex Machina we collected more than three hundred applications of AI in viticulture, winemaking and the wine trade and reopened the original source for every one of them, to look at the industry through its numbers rather than through its press releases. The numbers turned out to be encouraging: considerably more works than is generally assumed, and from what already works it is fairly easy to extrapolate what becomes normal within the next five years.

A winemaker who starts at twenty-five and retires at sixty-five will live through about forty vintages. Forty tries in a whole career, of which perhaps ten are genuinely difficult — the ones where the decision about a picking date or a fermentation temperature turned out to be a fork in the road. A model learning chess plays millions of games overnight and by morning knows more about openings than humanity accumulated in five hundred years.

It does not follow from that asymmetry that artificial intelligence does not work in wine. It follows that it works differently here than in advertising or logistics. Where the industry has learned to obtain observations quickly, AI settles in within a couple of seasons and already earns measurable money.

Where the answer arrives once a year, it is still catching up — but even that distance is closing faster than one might have expected five years ago: 146 of the 299 cases in our base, almost half, started in 2023 or later.

Testing a hypothesis like that needs a large body of evidence, so we built one: 307 documented applications of AI in wine, from peer-reviewed research to vendor press releases. We reopened the original source for every record, eight records had to be thrown out, and 299 remain, of which 287 are AI cases proper. As far as we know this is the largest audited catalogue of its kind, and every figure below is computed from it.

01 HOW A SOMMELIER LEARNS AND HOW A MODEL LEARNS

Anyone who has taken tasting courses already knows how machine learning works, just under different words. The WSET systematic approach requires a wine to be decomposed into a fixed set of parameters: colour intensity and hue, the character and intensity of the nose, acidity, tannin, alcohol, body, length of finish. Training a sommelier is first of all training them to turn a raw impression into a finite list of features, each of which can be named in a word and compared with another glass.

In machine learning the very same procedure is called feature engineering: from a shapeless “this wine is good” you have to obtain a row of a dozen or so numbers. The industry invented the technique decades before engineers needed it, and for the same reason — otherwise experience cannot be passed from one person to another.

Four terms are worth separating, because the trade press usually runs them together as synonyms.

Machine learning	A model that derives rules from examples instead of being given them by a person: show it ten thousand seasons of weather and yield and it will find a relationship nobody formulated.
Computer vision	The special case where the examples are images — a leaf, a bunch, a cork.
Language models	Trained to predict the next word in a text, which is why they reproduce a reference book beautifully and reason noticeably worse.
Agents	A layer over a language model that is permitted to take actions: search the web, write emails, open tickets.

The difference between them matters, because each runs into a data shortage of its own.

Three foundations

Wine has simultaneously far too much data and catastrophically little, depending on where exactly you look. A satellite photographs the vineyard every few days, a weather station writes a reading every hour, a sales platform records every transaction, and the French firm [Weenat](#) processes more than a billion data points a day and has deployed 25,000 sensors. That is an ocean. But an estate has exactly as much data on how one particular block behaved in a hot year as it has seen hot years, and as for how the wine turned out — one row per vessel per vintage.

The second problem is fragmentation. Sensor readings sit with the vendor, laboratory analyses in the oenologist's notebook, sales in the till system, tasting scores in somebody else's database, and none of those bodies of data answers the question the winemaker actually has.

Infrastructure is much the same. A vineyard is the textbook example of what the industry calls edge computing: the computation has to happen on site rather than in a data centre, because connectivity is either absent or unreliable. According to [Eurostat](#), only 43% of EU farms have internet access. The price of entry is concrete too: the French Bakus robot costs [€180,000 against €115,000 for a conventional tractor with a driver](#) and works out the same per hour at roughly 440 engine hours a year as the tractor does at 350. A data subscription costs quite different money: [Chouette](#) sells its own, cameras included, at [€200–400 per hectare per year](#).

Finally, the models themselves. Every model has a target variable — the thing it predicts — and features it predicts from. The choice of target variable decides what the model actually learns, and that is where the main trap is hidden: if the target variable is a critic's score, the model learns to predict one particular person's palate. The difference is not theoretical. A [study in the Journal of Wine Economics](#) found that the average correlation between Vivino's crowdsourced scores and critics' scores is 40%, while the critics agree with each other more closely: their average correlations with the others in the comparison run from 46% to 63%. Even the ground truth we train the model towards is itself noisy.

That is the ground-truth problem — the “correct answer” a model is checked against. In sensory work it barely exists. In a [study by the University of Adelaide and CSIRO](#) a model predicted 22 sensory descriptors of wine from grape data, with a panel of nine to eleven trained tasters as the reference; only five of the twenty-two came out convincingly, with R^2 above 0.7 — green flavour best of all.


Why wine turned out to be a hard case

Put three circumstances together. The industry's main production cycle yields one observation a year. The median estate is too small to justify a hundred and eighty thousand euro solution: the average EU wine holding is

1.4 hectares and 83.3% of holdings do not reach one. And the industry's end product is a subjective sensory impression that cannot be written down as a number without loss.

That trio explains the shape of the whole body of evidence rather well. In our base viticulture accounts for 177 cases, the wine business for 75 and winemaking and the laboratory for 47 — a vineyard-to-cellar ratio of 3.8 to 1. That does not make viticulture more important: it is simply that every satellite pass, every hour from a weather station and every vine under a camera turns into a row of a dataset today, while the cellar produces observations once a year. AI arrived where data accumulates faster than the question ages.

FIG. 1



02 A MAP OF THE EXPERIMENTS

How this was counted

Before arguing about how much artificial intelligence there is in wine, it is worth agreeing what counts as it, and we stated the frame explicitly and fairly broadly. AI here includes machine learning, computer vision, time-series forecasting, chemometrics with classifiers, language models, and autonomous navigation with perception. Outside the frame: measurement without prediction, cryptography and blockchain, and calculation from fixed coefficients.

The choice of frame changes the statistics more than one would think. Under a narrow definition, where only deep learning counts, **Amorim**'s optical cork sorting and **Pellenc**'s sorting cameras would drop out of the base; under ours they remain full cases. Conversely, 12 records fell outside it — no learning, no prediction — and they are flagged separately in the catalogue rather than counted as AI cases. More interesting: in about half of the disputed cases the word “AI” comes not from the company but from a journalist. On their own pages for the products in the catalogue, **Oculyze** and **Vinescapes** do not use “artificial intelligence” or “machine learning” once, though the first genuinely does use it. The press, in other words, errs in both directions.

The catalogue then went through a complete check: all 299 records were reopened against their sources by seventeen independent checkers whose brief was “find the error”, not “confirm it”. The checkers were AI agents, and the author reviewed their results and the corner cases. This is what it found.

INVENTED SOURCES	FULLY CONFIRMED	DISCREPANCY CORRECTED
0 no companies, no papers	181 61% of records	92 31% of records
RECORDS DELETED		
8 307 → 299		

Those four figures are the audit's own report rather than a field in the dataset: the published catalogue carries a source-confidence grade on every record, but not the history of its corrections. The most frequent defect was a specific figure simply not being on the page cited — about 55 cases out of 127. So 43% of the industry's claims turned out to disagree with their own source, though there is nobody here to blame. There is no independent body checking technology claims in this industry: the press retells the press release, the next publication retells the press, and by the third retelling the figure has started living a life of its own.

What is striking is that the louder a case is claimed to be working, the worse it checks out. Records resting on science and official reports were fully confirmed 66% of the time, those resting on the trade press 61%, and those resting on vendor marketing only 27% — and it is that last group that has the highest share of claims of commercial operation. I would add that the checkers found every fifth case in the base described with no quantitative result at all: the industry learned to talk about AI somewhat earlier than it learned to measure it.

What gets deployed and what stays in the papers

CLASSIFICATION

The tasks in this table are the audit's own working grouping, not a facet of the catalogue: the filter on the [register page](#) narrows the 299 cases by domain, technology, stage, country and source confidence.

TASK	CASES	OF WHICH RESEARCH	OF WHICH COMMERCIAL
Irrigation and water	49	8	22
Disease detection	41	7	19
Sensory and flavour	41	9	17
Consumer recommendation	37	1	18
Yield forecasting	30	6	13
Pruning	14	6	2

Pruning turned out to be the only task in the base with three times more research than deployments, and the explanation is fairly obvious. “How much water does the vine need” and “is there a spot on this leaf” are questions with a checkable answer that arrives within days, whereas “where to cut” is a question about the shape a vine should take in three years’ time. By domain the picture is the same: in winemaking 40% of cases remain research, in the business only 15%, while 71% of business cases reached operation. The economics of deployment in wine are set by how fast you can find out that the model was wrong.


A tour of the markets

FIG. 2

FIG. 2

Where the cases are, against where the wine is made

299 AUDITED CASES · 15 OF 42 COUNTRIES DRAWN



■ share of the cases ■ share of world wine production, 2021–2025 average

Each country is drawn twice on one scale: its share of the cases, and its share of the world's wine. Drawn: every country with at least 2% of either – 15 of the 42 the corpus names. A case run in several countries counts in each (multi-country cases: 33), so the case bars can add up to more than 100%. Platforms offered worldwide with no country named are not drawn (cases: 18). Production: OIV, wine excluding juices and musts, 2021–2025.

The cases are not where the wine is made. The United States accounts for 71 cases, or 23.7% of the base, against 9.5% of world wine production over 2021–2025, [by the OIV's count](#). Italy is the other way round: the largest producer, with 18.7% of the world's wine, and 9.0% of the cases. Spain makes 13.1% of the world's wine and holds 8.0% of the cases; Chile and Argentina make 4.5%

each and account for 8 and 4 cases. The United Kingdom, which makes almost no wine, accounts for 16 cases – more than Chile and Argentina together. The base follows the technology market – vendors, research groups and the trade press that covers them – more than the vineyard.

The United States also holds two of the most instructive stories. [A yield model by INRAE and Cornell University](#), built on 321 sampling points in a Cornell vineyard across 2018–2021, predicts yield with a 2–8% error and pruning weight with 15–20%. Beside it stands **Monarch Tractor**: an electric autonomous tractor, more than \$240m raised, 500-plus machines delivered, over 130,000 operating hours – then [layoffs in late 2025](#) and [a shutdown by April 2026](#), with the technology sold, reportedly to Caterpillar. The engineering is excellent in both; the only thing that separates them is who ends up paying for the hardware.

France, with 45 cases, gives the same contrast in pure form. **Naïo Technologies** is 350 machines in service, insolvency proceedings in June 2025, a €6.4m rescue from Mirova, Bpifrance and the Occitanie region, and a headcount cut from 80 to 21. And right beside it, **Weenat** with 30,000 users, 25,000 sensors and an €8.5m Series C in 2024. One company sold machines; the other sold data.

Portugal (16 cases) and Spain (24 cases) give two examples worth holding side by side. The EU-funded **VINBOT** project estimated yield from imagery across 27 plots and 17 varieties in Portuguese vineyards, and the project's own final report, published by the European Commission on CORDIS, states that the per-linear-metre estimates showed “a low or absent agreement and a very high error”. A publicly admitted failure is extremely rare in industry documentation, which makes the record more valuable rather than less. The second example, in Spain, is the [twenty-five-year collaboration between IRTA and Raventós Codorníu](#) in the Raimat vineyards, where a quarter of a century produced 20% less water and 65% more block uniformity – with none of it credited to artificial intelligence: Raimat's director calls AI a prospect rather than a current tool. It is probably the best inoculation against inflated expectations I know of, because any vendor would have sold that result as an achievement of machine learning.

Twenty-five years of work produced 20% less water and 65% more block uniformity, with none of it credited to artificial intelligence.

Germany, with 18 cases, has the longest-running AI case. **VitiMeteo**, which models the life cycle of downy and powdery mildew for precise spray timing, covers about 42,000 hectares through a hundred weather stations and has been in daily operation since 2003 – the longest-lived AI case in the catalogue. And the **FungiSens** project fitted the whole data-resolution problem into a single figure: conditions lethal to fungal spores were recorded inside the vine canopy on 9% of days, while a standard weather station [75](#)

[metres away, in the same vineyard](#), showed 1.4%. A near-sevenfold difference that comes from where the thermometer stands, not from the quality of the model.

New Zealand (15 cases) and Australia (24 cases) are the strongest in autonomous machinery that can count its own fuel. The **Prospr** sprayer from Robotics Plus burns 1.5 litres an hour against 9–10 for a diesel tractor — early results from a trial across three Marlborough vineyards show fuel use down by around 70% — and one operator runs several machines. The company was eventually bought by Yamaha. Chile (8 cases) gives the cleanest water figures, though [not strictly vineyard ones](#): the Argentine firm **Kilimo** reports cutting water use by 13% across 450 hectares of farmland and saving 1.5 million cubic metres over three years — in the Maipo basin, where agriculture takes about 68% of available water.

China (21 cases) stands apart, and cannot be skipped precisely because of its deployment model. While the West sells solutions one estate at a time, Ningxia is building a regional platform and connecting everyone to it: the “1+N+100” system serves [a hundred wineries](#). The results, though, are a single estate’s — GreatWall Terroir’s grape planting base, where drip irrigation with sensors replaced flood irrigation: water use down [from 700–800 to 220–260 cubic metres per mu](#), and five workers now covering more than 7,000 mu where they once covered 300. A caveat is obligatory here: the Chinese figures come from state media and cannot be independently verified. But the contrast between the two deployment models is real enough.

Two cases stand out in the cellar and the laboratory. The **Bionic Eye** system at the cork producer M.A. Silva inspects up to 40,000 corks an hour, and the figures the manufacturer gives are for consistency: human inspection holds at about 75%, machine inspection at about 100%. Those figures speak to stamina rather than to sharpness of sight — the machine does not tire by Thursday. And NMR spectroscopy, the basis of [Bruker’s wine profiling](#), which the company markets for detecting adulteration, entered the OIV Compendium of International Methods of Analysis of Wines and Musts in 2021 as a method for quantifying six wine parameters — a rare case of an industry technology reaching the status of an official method.

The wine business remains the only area with a genuine randomised experiment. The retailer **Wine Access** ran [three controlled trials](#) of email campaigns: copywriter-written against language-model-generated against hybrid, about 9,000 customers per cell. In two of the three trials the machine and hybrid emails matched or beat the human ones on profit, by as much as +9.36%. Meanwhile the copywriting staff cost \$375,000 a year, model licences \$1,000–1,200, and the hybrid option \$63,500–94,950. Beside it stands [Treasury Wine Estates](#), which cut the development cycle for a new wine from 12–18 months to under six, and [Naked Wines](#), where recommendations draw on 35.6 million reviews and a model forecasts five-year customer contribution at a projected 1.7x payback on new-customer spend. The other side of the same coin was [Winc](#): 77.5% revenue growth during the pandemic,

an IPO — and Chapter 11 just over a year later with \$36.75m of debt against \$50.3m of assets. Nothing suggests the wine-matching algorithm failed; the [customer-acquisition economics](#) did.

The copywriting staff cost \$375,000 a year, model licences \$1,000–1,200, and the hybrid option \$63,500–94,950.

The summary map

DOMAIN	CASES	REACHED OPERATION	WHAT ACTUALLY WORKS	WHO PAYS
Viticulture	177	41%	Data-driven irrigation, spray timing, disease detection	The estate, often with a subsidy
Wine business	75	71%	Churn prediction, personalisation, content generation	The company, payback in months
Winemaking and laboratory	47	45%	Fermentation control, authentication, sorting	A large producer or an equipment vendor

The overall economic picture came out harsher than the press releases sound. 11 records in the base stand at the “ended” stage, and almost all of them are machinery makers, while data services keep growing; add to them around a dozen acquisitions. **Bluewhite**, after being bought by Elbit, left agriculture for defence autonomy, and **Bloomfield**, after being bought by Kubota, [shifted its focus toward blueberries](#) — so the vineyard market, by all appearances, turned out not to be large enough even for those who had already entered it. At the same time the base documents about €26m in explicitly named European grants alone, and against a collapse in global agrifoodtech investment from \$51.7bn in 2021 to \$16.2bn in 2025, public budgets turned out to be a steadier source than venture capital. A significant part of AI in wine therefore exists not because it pays for itself but because it was funded.

03 POTENTIAL

What follows is about what is not in the base yet but is technically close. For each area it helps to walk the same chain: where the pain is, what data already exists, what can be extracted from it, what gets in the way, and what effect to expect.

The vineyard

WHERE THE PAIN IS

The main pain here is crop protection. In 1915 France [lost 70% of its grape crop to downy mildew](#), and in 2023 [several Tuscan wineries lost about 70% of their production](#), although the consortia of the main appellations put the fall

at 5–20%. Meanwhile many estates still spray by the calendar and by a forecast from a weather station thirty kilometres away.

WHAT DATA ALREADY EXISTS

The data for a different approach has already been collected: satellite imagery every few days, local weather stations, spore traps, tractor telemetry. The French **DeciTrait** system, which combines risk models with field observation, shows savings of 200–750 grams of copper per hectare and an average reduction of up to 35% in the pesticide-load index — and that is not the future but working practice.

WHAT CAN BE EXTRACTED FROM IT

Technically the closest thing is the move from calendar spraying to risk-based spraying, and to my mind it is the most underrated story in the whole subject. Next comes climate adaptation: harvest dates in most wine regions have shifted two to three weeks earlier over the last forty years, and models linking growing-season weather to ripeness deliver more here than any amount of robotics.

WHAT GETS IN THE WAY

What gets in the way is transferability. It is the central scientific limitation of the whole field and it recurs in almost every other paper. The Italian [De Nart](#) study on identifying variety from a leaf reports above 0.9 under cross-validation, and on an independent external set the authors state plainly that “no model can offer satisfactory performance” when only the single most probable class is considered. The German [Bendel](#) work on esca shows 88–95% accuracy at the symptomatic stage on manually labelled leaves (70–81% on the raw field data) and 62–86% pre-symptomatically — precisely the stage detection exists for; here the authors call their results promising but still in need of further evaluation. The DoctorP app, which diagnoses disease from a photograph of a leaf, is the main entry point to a JINR platform whose early model gave [above 95% on synthetic images and about 50% on real user photographs](#). Two systematic reviews, of 104 and 197 papers, independently list validating models across different vineyards and terroirs among the problems still open.

A model trained on one vineyard knows that vineyard, not viticulture. Terroir, which winemakers count as a virtue, is a generalisation problem for machine learning.

WHAT EFFECT TO EXPECT

As for the effect, it is realistic to expect a 15–30% reduction in sprays and a comparable saving in water on estates that already collect data. **Villa Sandi**, with two hundred connected hectares, reports spray reductions of up to 20% and water savings averaging 10%. That is hard to call a revolution; it is good agronomy amplified by measurement.

The cellar

WHERE THE PAIN IS

The cellar's pain is irreversibility. A mistake in the vineyard is corrected by the next operation; a mistake in fermentation is not corrected at all — and the winemaker takes decisions on three or four measurements a day and their own nose.

WHAT DATA ALREADY EXISTS

There is data here, but there is little of it and it is not joined up: density, temperature, volatile acidity, sometimes online tank monitoring. The **DTWINE** project, with a €1m budget, was building a digital twin of fermentation in 2021–2024 across four Spanish wine regions, testing it at an experimental winery of 30-litre tanks in Valencia, precisely so that realistic scenarios could be tried without risking large volumes of wine.

WHAT CAN BE EXTRACTED FROM IT

The most honest framing of the task is early deviation signals: not “make the wine” but “tell me that this tank is not behaving like the other forty”. Next comes blending as an optimisation problem — an assemblage has a finite number of components, measurable analytical parameters and an objective function set by the winemaker. A machine really can search that space faster than a person, provided the person sets the objective.

WHAT GETS IN THE WAY

What gets in the way is the same shortage of observations plus a lack of motivation: winemaking runs 40% research cases against 45% that reached operation, the longest distance between paper and cellar floor in the industry.

WHAT EFFECT TO EXPECT

The effect is modest in money and large in risk, which makes the right frame a winemaker's copilot rather than an autopilot. Cutting spoiled batches by a few per cent a year pays for the system at a large estate and pays for nothing at all at an estate of a hectare and a half.

The wine business

WHERE THE PAIN IS

The pain here is simple: wine sells worse than it is made. The industry produces more than it consumes, and the cost of acquiring a customer grows faster than the average order.

WHAT DATA ALREADY EXISTS

The data in this area is the best in the industry: it is already digital, nobody has to collect it again, the target variable is unambiguous (bought or did not buy), and the result is visible within days. Hence 71% of business cases reaching operation against 15% stuck in research.

WHAT CAN BE EXTRACTED FROM IT

What works is club churn prediction ([Commerce7](#) claims 74% accuracy), dynamic pricing, personalised recommendation and commercial content generation – with the caveat that the only clean experiment in this area showed a change in economics rather than a replacement of the copywriter: the hybrid is four to six times cheaper than the staff at comparable profit. A separate story is counterfeiting, which some insiders [feared in 2006](#) could reach 5% of the secondary market, and [as much as \\$550m worth of Rudy Kurniawan's fakes](#) may still be on the market, by one expert's 2016 estimate; isotope analysis coupled with NMR remains the only area where the technology has reached the status of an official method.

WHAT GETS IN THE WAY

What gets in the way here is not the technology but company size. Only 26 records in the base describe wine companies' own internal programmes; the overwhelming majority of the rest are vendors, universities and research consortia. [Gallo](#), [Pernod Ricard](#), [Concha y Toro](#) and [LVMH](#) have their own teams, and for everyone else AI is something you buy rather than something you build.

WHAT EFFECT TO EXPECT

The effect here is the highest of the three and the most boring: single-digit percentages on conversion and double-digit ones on the speed of operations. At [Gallo](#) a decision-making platform delivered [\\$890,000 of savings in the first year](#) with payback in about a year – a story not about the future of wine but about ordinary modern business.

04 THREE HORIZONS

Nearly inevitable by 2028: risk-based spraying everywhere there is a data subscription, simply because regulatory pressure on pesticides is growing faster than chemistry is getting cheaper. Language models will by then have taken the back office for good – wine descriptions, customer replies, translations, reporting. None of those applications will look like “AI in wine”; they will look like routine work disappearing.

Probable by 2032: at least one industry standard for measuring effect, because today every company counts its savings its own way, which is exactly why 43% of claims do not check out. In the same bracket, the consolidation of

data services around three to five platforms, as has already happened in row crops, and probably the first models that transfer between terroirs — not because someone invents a new architecture but because enough labelled data from different sites will have accumulated.

Pure speculation, for now: a model that predicts the sensory profile of a future wine from season data and must analytics, accurately enough for commercial decisions. The problem is technically well posed, and the University of Adelaide work showed that five descriptors out of twenty-two are already predicted convincingly. In practice it needs decades of carefully labelled vintages from many estates, and there is nobody to collect them: everyone has forty tries.

This is where it is worth returning to where we started. Everything artificial intelligence does well in wine, it does with the measurable: litres of water, grams of copper, hours of work, the probability of a customer leaving, a spot on a leaf. Everything it does badly lies where no correct answer exists: where to cut, how the wine will turn out, what to serve with this dish, whether it is worth the money.

It is tempting to call that a limitation of the technology, but look at the same boundary from the other side. An industry whose main product does not reduce to a number is an industry that cannot be optimised into a commodity. Wine is valued for precisely what resists measurement, and as long as that holds it has a protection no other agricultural product has. The machine will count the water, forecast the mildew, write the customer's email and free the winemaker from everything that can be counted at all. What wine is actually drunk for will stay in those forty tries — and by the look of it, for a long time yet.

Every figure in this text was checked against its original source as of September 2026; the wine production shares are the OIV's, as read in September 2026. The full catalogue of 299 records, with sources, confidence grades and audit results, is the [Casebook register](#), and it is open to checking: every figure taken from the catalogue is recomputed on this page at each build rather than typed in by hand.

*Cite this essay: Khudov, N. (2026). Vinum ex Machina: the winemaker gets forty tries, the machine gets a million [Report]. Zenodo.
<https://doi.org/10.5281/zenodo.23033982>*

*Cite this dataset: Khudov, N. (2026). Vinum ex Machina Casebook: an audited register of 299 AI applications in viticulture, winemaking and the wine business (Version v2026-09-27) [Data set]. Zenodo.
<https://doi.org/10.5281/zenodo.22982275>*

Text: [CC BY-SA 4.0](#), except where otherwise noted.

References

-
- 01 Armstrong, C. E. J., Niimi, J., Boss, P. K., Pagay, V., & Jeffery, D. W. (2023). Use of machine learning with fused spectral data for prediction of product sensory characteristics: The case of grape to wine. *Foods*, 12(4), Article 757. <https://doi.org/10.3390/foods12040757>
-
- 02 Banker, S. (2026, June 10). Gallo invests in agentic AI to improve supply chain capabilities. *Forbes*. <https://www.forbes.com/sites/stevebanker/2026/06/10/gallo-invests-in-agentic-ai-to-improve-supply-chain-capabilities/>
-
- 03 Bendel, N., Kicherer, A., Backhaus, A., Klück, H.-C., Seiffert, U., Fischer, M., Voegelé, R. T., & Töpfer, R. (2020). Evaluating the suitability of hyper- and multispectral imaging to detect foliar symptoms of the grapevine trunk disease Esca in vineyards. *Plant Methods*, 16(1), Article 142. <https://doi.org/10.1186/s13007-020-00685-3>
-
- 04 Bleyer, G., Kassemeyer, H.-H., Krause, R., Viret, O., & Siegfried, W. (2008). „VitiMeteo Plasmopara“ – Prognosemodell zur Bekämpfung von Plasmopara viticola (Rebenperonospora) im Weinbau [“VitiMeteo Plasmopara”: A forecasting model for controlling Plasmopara viticola (downy mildew) in viticulture]. *Gesunde Pflanzen*, 60(3), 91–100. <https://doi.org/10.1007/s10343-008-0187-1>
-
- 05 Bruker. (2021, May 10). International Organization of Vine and Wine (OIV) officially incorporates NMR method in Compendium of International Methods of Analysis of Wines and Musts [Press release]. <https://www.bruker.com/en/news-and-events/news/2021/international-organization-of-vine-and-wine-officially-incorporates-nmr-method-in-compendium-of-international-methods-of-analysis-of-wines-and-musts.html>
-
- 06 Chambres d'agriculture France. (n.d.). Decitrait. Retrieved September 28, 2026, from <https://chambres-agriculture.fr/etre-accompagne/nos-solutions-numeriques/decitrait>
-
- 07 Chambrin, A. (2021, September 20). Robot Bakus: plus intéressant économiquement qu'un enjambeur avec chauffeur? [Bakus robot: More economically attractive than a straddle tractor with a driver?]. *Entraid*. <https://www.entraid.com/articles/rentabilite-robot-bakus-vitibot-comparatif-enjambeur>
-
- 08 Commerce7. (2025, September 16). From churn prediction to ChatDTC: Commerce7 debuts AI features built for DTC growth [Press release]. *Wine Industry Advisor*. <https://wineindustryadvisor.com/2025/09/16/from-churn-prediction-to-chatdte-commerce7-debuts-ai-features-built-for-dtc-growth/>
-
- 09 Delattes, E. (2025, September 29). Agriculture : la pépée Weenat équipe 30 000 agriculteurs européens [Agriculture: Rising star Weenat equips 30,000 European farmers]. *L'Essentiel de l'Éco*. <https://lessentieldeleco.fr/3467-weenat/>
-
- 10 Delbecque, X. (2025, April 16). De nouveaux outils pour estimer le risque de maladies dans les vignes [New tools for estimating the risk of disease in vineyards]. *Réussir Vigne*. <https://www.reussir.fr/articles/de-nouveaux-outils-pour-estimer-le-risque-de-maladies-dans-les-vignes>
-
- 11 De Nart, D., Gardiman, M., Alba, V., Tarricone, L., Storchi, P., Roccotelli, S., Ammoniaci, M., Tosi, V., Perria, R., & Carraro, R. (2024). Vine variety identification through leaf image classification: A large-scale study on the robustness of five deep learning models. *The Journal of Agricultural Science*, 162(1), 19–32. <https://doi.org/10.1017/S0021859624000145>
-
- 12 Dubé, J.-P., & Xu, A. (2026). Large language models and creative content design: A case study of email marketing at Wine Access. *Quantitative Marketing and Economics*, 24(1), Article 1. <https://doi.org/10.1007/s11129-025-09303-9>
-
- 13 Eurostat. (2026, July 24). 43% of EU farms with internet access. <https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20260724-1>
-
- 14 Frank, M. (2006, December 20). Right bottle, wrong wine. *Wine Spectator*. <https://www.winespectator.com/articles/right-bottle-wrong-wine-3325>
-
- 15 Gessler, C., Pertot, I., & Perazzolli, M. (2011). Plasmopara viticola: A review of knowledge on downy mildew of grapevine and effective disease management. *Phytopathologia Mediterranea*, 50(1), 3–44. <https://oajournals.fupress.net/index.php/pm/article/view/5388>
-

- 16 Hekkert, G. (2026, April 15). Caterpillar acquires Monarch technology after electric tractor startup shuts down. Future Farming. <https://www.futurefarming.com/tech-in-focus/autonomous-semi-autosteering-systems/caterpillar-acquires-monarch-technology-after-electric-tractor-startup-shuts-down/>
- 17 Huband, M. (2026, July 22). How artificial intelligence is safeguarding a millennia-old practice [Partner content]. The Drinks Business. <https://www.thedrinksbusiness.com/2026/07/how-artificial-intelligence-is-safeguarding-a-millennia-old-practice/>
- 18 International Organisation of Vine and Wine. (2026). State of the world wine sector in 2025. https://www.oiv.int/sites/default/files/documents/OIV-State_of_the_World_Wine_Sector_in_2025_0.pdf
- 19 Joint Institute for Nuclear Research. (2023, April 18). OIYaI razvivayet platformu dlya raspoznavaniya bolezney rasteniy [JINR develops platform for plant disease detection] [Press release]. Nauchnaya Rossiya. <https://scientificrussia.ru/articles/oiai-razvivaet-platformu-dla-raspoznavania-boleznej-rastenij>
- 20 Kelly, C. (2025, Fall). Agriculture 4.0: Bloomfield Robotics. The Link: The Magazine of CMU's School of Computer Science. <https://magazine.cs.cmu.edu/bloomfield-robotics>
- 21 Kilimo. (2025, April 22). Microsoft and Kilimo: Artificial intelligence for smarter irrigation in Chile. <https://kilimo.com/en/microsoft-and-kilimo-artificial-intelligence-for-smarter-irrigation-in-chile/>
- 22 Kleb, M., Merkt, N., & Zörb, C. (2022). New aspects of in situ measurements for downy mildew forecasting. Plants, 11(14), Article 1807. <https://doi.org/10.3390/plants11141807>
- 23 Kopsacheilis, O., Analytis, P. P., Kaushik, K., Herzog, S. M., Bahrami, B., & Deroy, O. (2024). Crowdsourcing the assessment of wine quality: Vivino ratings, professional critics, and the weather. Journal of Wine Economics, 19(3), 285–304. <https://doi.org/10.1017/jwe.2024.20>
- 24 Li, P. (2024, December 2). Ningxia Helan Shan donglu jie "suan" fa "li" zoushang putaojiu zhihui fazhan zhi lu [Ningxia's eastern Helan Mountain foothills draw "power" from "computing" on the road to a smart wine industry]. China News Service. <https://www.chinanews.com.cn/cj/2024/12-02/10329190.shtml>
- 25 Marston, J. (2025, November 6). Ag robotics startup Naïo bounces back after a year of financial turmoil. AgFunderNews. <https://agfundernews.com/ag-robotics-startup-naio-bounces-back-after-a-year-of-financial-turmoil>
- 26 Martin-Graser, D. (2021, December 27). Moët & Chandon : quand l'IA révolutionne le Champagne [Moët & Chandon: When AI revolutionises Champagne]. CScience. <https://web.archive.org/web/20251119012708/https://www.cscience.ca/pratiques-viticoles-la-revolution-moet-hennessy/>
- 27 McGowan, B. (2017, July 18). E. & J. Gallo deploys data-driven app to harvest a better grape. CIO. <https://www.cio.com/article/230368/e-j-gallo-deploys-data-driven-app-to-harvest-a-better-grape.html>
- 28 Monarch Tractor shuts down: What went wrong with the \$500M electric tractor startup. (2026, April 8). Farm Progress. <https://web.archive.org/web/20260501183530/https://www.farmprogress.com/max-armstrong/monarch-tractor-shuts-down-what-went-wrong-with-the-500m-electric-tractor-startup>
- 29 Montes, J. (2025, December 9). Data-driven farming and AI are transforming water management in Chile. Source LATAM. <https://news.microsoft.com/source/latam/features/ai/data-driven-farming-chile/?lang=en>
- 30 Naked Wines plc. (2023). Annual report and accounts 2023. https://s28.q4cdn.com/964621086/files/doc_financials/2023/ar/NakedWines_AR23_FINAL.pdf
- 31 Oculyze Automation GmbH. (n.d.). Oculyze automated yeast cell counter: How it works. Retrieved September 29, 2026, from <https://www.oculyze.net/yeast-cell-counter/wine-fermentation/how-it-works/>
- 32 Olmos, R. (2023, May 29). Las nuevas líneas de trabajo del Centro de Investigación e Innovación de Viña Concha y Toro [The new lines of work at Viña Concha y Toro's Research and Innovation Centre]. Diario Financiero. <https://www.df.cl/df-lab/innovacion-y-startups/las-nuevas-lineas-de-trabajo-del-centro-de-investigacion-e-innovacion-de>

-
- 33 Parc Científic de la Universitat de València. (2023, August 7). IATA-CSIC presents an experimental winery that advances in the application of 'digital twins' in the wine industry. PCUV News. <https://news.pcuv.es/en/home-pcuv-institutos-pcuv-iata-iata-csic-presents-an-experimental-winery-that-advances-in-the-application-of-digital-twins-in-the-wine-industry>
-
- 34 Pernod Ricard. (2024, April 4). Pernod Ricard rated top European company in the 2024 Data & AI Human Capital Report: World rankings compiled by AlixPartners & Darwin X [Press release]. <https://www.pernod-ricard.com/en/media/pernod-ricard-rated-top-european-company-2024-data-ai-human-capital-report-world-rankings>
-
- 35 Poitiers, M. (2023, February 6). Chouette lève 5 millions d'euros pour aider les viticulteurs à surveiller et traiter leurs vignes [Chouette raises €5 million to help winegrowers monitor and treat their vines]. L'Usine Digitale. <https://www.usine-digitale.fr/article/chouette-l-agritech-qui-aide-les-viticulteurs-a-surveiller-et-a-traiter-leurs-vignes-leve-5-millions-d-euros.N2098361>
-
- 36 Redacción Interempresas. (2026, February 18). IRTA y Raventós Codorníu cumplen 25 años de colaboración vitivinícola [IRTA and Raventós Codorníu mark 25 years of wine-growing collaboration]. Interempresas. <https://www.interempresas.net/Vitivinicola/622921-IRTA-y-Raventos-Codorniu-cumplen-25-anos-de-colaboracion-vitivinicola.html>
-
- 37 Schmitt, P. (2016, December 8). \$550m worth of Kurniawan counterfeit wines being 'resold over and over again'. The Drinks Business. <https://web.archive.org/web/20260114064801/https://www.thedrinksbusiness.com/2016/12/550m-worth-of-kurniawan-counterfeit-wines-being-resold-over-and-over-again/>
-
- 38 "Several wineries in Tuscany have lost 70% of their production because of downy mildew." (2023, August 2). WineNews. https://winenews.it/en/several-wineries-in-tuscany-have-lost-70-of-their-production-because-of-downy-mildew_503118/
-
- 39 Skinner, T. (2026, February 17). Marlborough demo highlights role of AGVs in sustainable winegrowing. NZ Winegrower. <https://www.ruralnewsgroup.co.nz/wine-grower/wg-machinery/vineyard-robots-marlborough-two-robotics-plus-trial>
-
- 40 Smith, C. (2022, December 2). Online wine membership club Winc files for bankruptcy. The Drinks Business. <https://www.thedrinksbusiness.com/2022/12/online-wine-membership-club-winc-files-for-bankruptcy/>
-
- 41 Taylor, J. A., Bates, T. R., Jakubowski, R., & Jones, H. (2023). Machine-learning methods to identify key predictors of site-specific vineyard yield and vine size. American Journal of Enology and Viticulture, 74(1), Article 0740013. <https://doi.org/10.5344/ajev.2022.22050>
-
- 42 Thach, L. (2026, April 9). How Treasury Wine Estates leverages AI to create wines consumers love. Forbes. <https://www.forbes.com/sites/lizthach/2026/04/09/how-treasury-wine-estates-is-leveraging-ai-to-create-wines-consumers-love/>
-
- 43 Universität Hohenheim. (2023, July 27). Weinbau: Präzisere Pilz-Bekämpfung durch Digitalisierung [Viticulture: More precise fungus control through digitalisation] [Press release]. BioRegio STERN. <https://www.bioregio-stern.de/de/BranchenNews/weinbau-praezisere-pilz-bekaempfung-durch-digitalisierung>
-
- 44 Villa Sandi. (2026, July 20). Villa Sandi turns its vineyards into an intelligent ecosystem for precision viticulture [Press release]. PR Newswire. <https://www.prnewswire.com/news-releases/villa-sandi-turns-its-vineyards-into-an-intelligent-ecosystem-for-precision-viticulture-302829846.html>
-
- 45 VINBOT Consortium. (2018). Final report summary – VINBOT (Autonomous cloud-computing vineyard robot to optimise yield management and wine quality). CORDIS, European Commission. <https://cordis.europa.eu/project/id/605630/reporting>
-
- 46 Vinescapes Limited. (n.d.). VineMAP: Vineyard Mapping & Assessment Program. Retrieved September 29, 2026, from <https://www.vinemap.com/>
-
- 47 Watson, E. (2026, April 16). Exclusive: Monarch Tractor CEO—'We should have pivoted harder and faster'. AgFunderNews. <https://agfundernews.com/exclusive-monarch-tractor-ceo-we-should-have-pivoted-harder-and-faster>
-

48 Winc, Inc. (2022). Form 10-Q for the quarterly period ended September 30, 2022 [Quarterly report]. U.S. Securities and Exchange Commission. <https://www.sec.gov/Archives/edgar/data/1782627/000095017022025015/wbev-20220930.htm>

49 Xinhua. (2025, June 13). Smart tech fuels dynamic growth in northwest China wine industry. <https://english.news.cn/20250613/2d789a4410be4e2baadc698dc6add599/c.html>

Text: CC BY-SA 4.0, except where otherwise noted. Edition v2026-09-29 of <https://vinumexmachina.com/keynote/>; the living page may have changed since. Figures and counts are computed from the Vinum ex Machina Casebook, v2026-09-27, <https://doi.org/10.5281/zenodo.22982275>.