
A PERSONAL WINE ATLAS

About Codex Vini

project overview, rating scale, structure, and taxonomy

The project

Codex Vini — the book of wine — is a personal atlas of wines I have tasted, with my ratings attached. It is a reference tool, not a commercial site. The goal is to make the collection legible to myself: to see where I have tasted, what I have gravitated toward, and what remains unexplored.

The project is a sibling to *The Grand Cru Atlas*, which maps the world's classified vineyard sites, and to the other instruments in the *Vinotheca*: the *TasteRank Explorer* of grape-variety kinship, and *The Soul of Wine*, a study of regional identity. Where those projects are historical, taxonomic, and anthropological, *Codex Vini* is personal and cumulative — a record of what has actually crossed my palate.

Rating scale

Wines are scored on a 0–10 scale with one decimal of precision. The practical range in this codex runs from **8.9** to **10.0**, because the collection records only wines I considered worth noting. Lower-rated wines are omitted. Rough interpretive bands:

- 10.0** — transcendent; among the finest wines I have ever tasted
- 9.8–9.9** — extraordinary; consistent, memorable, benchmark examples
- 9.5–9.7** — excellent; distinctive, complete wines with few flaws
- 9.2–9.4** — very good; well-made wines with real character
- below 9.0** — noteworthy but with reservations

Ratings are subjective, contextual (bottle condition, food, mood), and drift over time as a palate develops. They are a snapshot of an impression, not an objective verdict.

Dynamic structure

Every wine lives as a single JSON entry in **wines.json**, stored in the GitHub repository that powers this site. Adding a new wine means committing one new object to that file; the site updates automatically within a minute.

This design choice — a flat JSON file in a git repository, instead of a database or spreadsheet — is deliberate. Git provides version history and backup. There is no server to run, no database to maintain, no authentication to manage. The atlas is self-contained, portable, and will outlive most web stacks.

Taxonomy

Wines are categorised by a single **grape/category** field. Grape varieties (Cabernet Sauvignon, Nebbiolo, Riesling) and named-wine categories (Barolo, Brunello di Montalcino, Champagne, Bordeaux Blend) sit as peers. A Châteauneuf-du-Pape is tagged *Châteauneuf-du-Pape Blend*, not decomposed into its lead grape Grenache. A varietal Grenache from another region is tagged *Grenache*. This matches how sommeliers and retail actually organise wine — by the identity on the label.

The reference lists draw on:

The **101 TasteRank grape varieties** — a personal distillation of the varieties worth knowing. See the companion grape-varieties document.

The **Soul of Wine region profiles** — 57 terroir profiles covering the world's significant wine regions. See the companion regions-reference document.

The rating-adjustment used in The Collection

Two charts in *The Collection* — average rating by country, and average rating by grape — use a statistical adjustment called **Bayesian shrinkage**. In short: it prevents a country with one lucky 10.0 from outranking a country with thirty consistent 9.5s. The full treatment, including the hierarchical-model derivation and a worked empirical-Bayes analysis of the current data, is in the companion shrinkage document.

Data

At time of this document, the codex holds **118 wines** across **14 countries**, **51 regions**, and **41 grape categories**, with a grand-mean rating of **9.52**. It grows by a few wines each week.

Every entry records: producer, wine name, vintage, country, region, grape/category, wine colour, rating, date tasted, and geographic coordinates for map placement.