KMC Reflection Paper

The application of AI, ML, Robotics and Automation in the Wine Industry

Participant: Mimi GIRAUD

Section: GEMBA Asia

Name of professor: Professor Phil Parker

Confidential (NDA signed): No

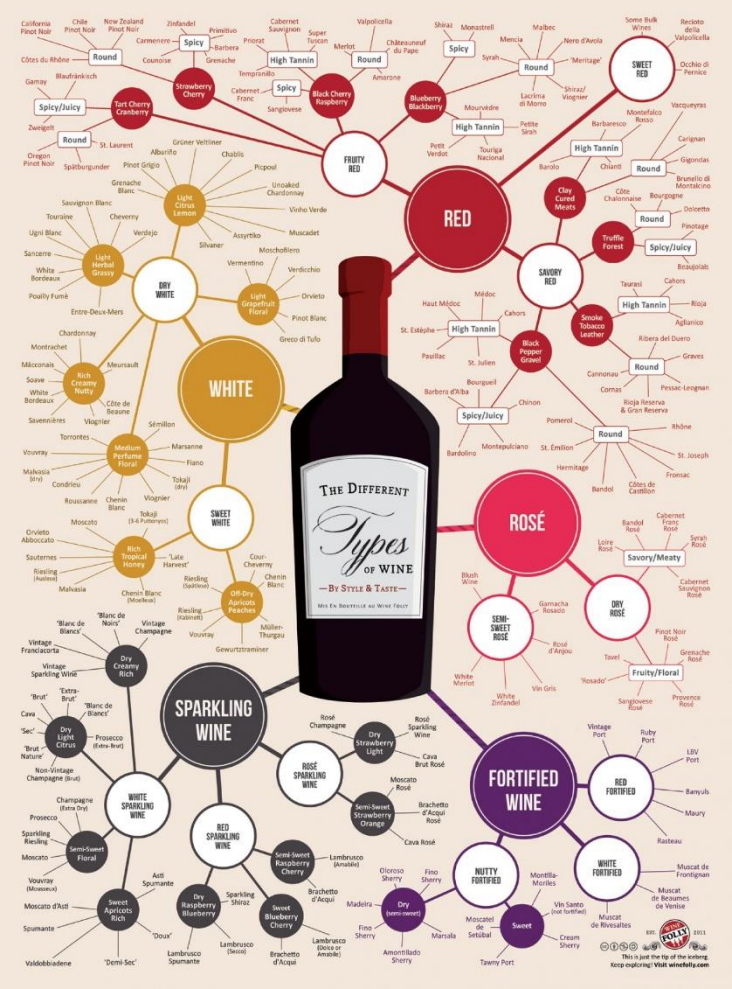

Summary

Here below is a summary of the current and future applications of Artificial Intelligence, Machine Learning, Robotics and Automation in the Wine and Wine Tasting industry. Surprisingly to my initial thoughts, wine is an increasingly crowded space in the tech scene, with thousands of start-ups, web platforms and apps being released to aid wine drinkers in their pursuit of an easy wine drinking experience. Following the trend of wine and wine education penetration in a lot of new countries, particularly in the New world and Asia, we can expect not only new areas growing wine, but new consumers and consumers habits coming to rise. Over the next few decades, the wine ecosystem will undergo a sea change, right from the way consumers buy wine to involvement of machines in the production, and AI coming to the aid of wine buyers and enthusiasts (when it comes to choosing the right one, pairing or not with their food and willingness to pay). From innovations to secure trademarks and track the whole supply chain, to assessing wine quality and building wine lover profiles, the use of AI systems is also likely to involve computerized forecasts, such as early-season vintage forecasts in a vineyard, or sales and price forecasts in a shop. In these cases, the forecasts are expected to improve through time, as more and more data are gathered, and the AI system continually adjusts itself based on newly found patterns in the data.

Company Mini-Case #1 - BrentaPack

Using corks imprinted with individual codes, BrentaPack, of the **Labrenta Group** has created the "IDCORK" system, which allows users to access details of a wine's history via an app. It's based on the conviction that while labels can be copied, corks cannot. This invention is designed to "guarantee the authenticity of a wine." In a few seconds and with a few clicks, you will be able to ensure that the contents of the bottle are not an imitation.

At the company's factory, the ID corks are printed with a personal number and then photographed from all angles. This information is scanned into a database, along with details of where and when the cork was harvested and made, and its individual characteristics. From there, wineries that have purchased the corks go into the database and insert all the information they want to include. For example, where the wine comes from, when the cork was inserted, the brand, the region, the vintage, the blend, and a picture of the bottle. After the bottle has been opened, consumers take an image of their cork and the app uses "visual recognition" to compare the on-screen shot and the cork that's just been pulled.

The second step in the process is to enter the cork's printed code, which reveals all the information stored in the database, including a picture of what the bottle should look like and a description of how the wine should taste.

In the future, the corks will also include GPS trackers, meaning that if a consumer accesses the IDCORK app, wineries will be able to see where an individual bottle has been opened.

- *Source: Corks with "Fingerprints" to Combat Wine Fraud, Diana Goodman | Posted Wednesday, 13 November 2013, https://www.wine-searcher.com/m/2013/11/corks-with-fingerprints--to-combat-wine-fraud*

Company Mini-Case #2 - Vivino

Vivino is an online wine community, database and mobile application where users can buy, rate and review wines. Vivino was founded in 2010 by Heini Zachariassen and Theis Sondergaard. As of 2018, Vivino had a wine database containing over 9 million different wines, and had 31 million users. It is now the world's most popular wine community and most downloaded mobile wine app. Vivino's million users contribute ratings for millions of wines from around the globe, and collectively, this database makes up the largest wine library in the world. Vivino uses label recognition technology to help guide wine purchases. With Vivino application, the user simply takes a photo of the wine label they are considering and is instantly provided the wine's rating, average price and review from the community of users. Vivino is also leveraging an AI solution from app-growth developer Shortcut Media for its newsletters: it serve its subscribers mobile-enabled newsletters via email, which congregate a series of reviews, ratings and wine discussions from its community, based on data from each individual reader. The email is enabled with deep links for users to seamlessly continue to its mobile app, which is significantly driving opens. The online wine destination is leveraging an artificial intelligence-powered algorithm that taps into behavioral data to serve targeted newsletters to users. Vivino users are getting personalized emails that give them an updated look at what is going on in the wine community.

- *Source: Vivino website, https://www.vivino.com/press*
- *Source: Vivino sees app open increase through AI-developed newsletters, Brielle Jaekel, https://www.mobilemarketer.com/ex/mobilemarketer/cms/news/email/22982.html*

Company Mini-Case #3 – Wine Ring

Wine Ring, headquartered in Syracuse, New York and founded in 2010, offers one of the most personal wine selection experiences available. Wine Ring is an AI-driven, B2B personalization software for the wine industry, focused on individual consumer preference and is used in over 150 countries. It has two applications, one for consumers, the other wine for merchants. The Wine Ring software has eight patents, with other patents pending. For consumers, unlike other apps that offer wine suggestions based on pairing suggestions or expert ratings, Wine Ring bases suggestions on individual preferences. This app uses advanced algorithms to develop a personal profile based on consumers' rating of wines and then recommends bottles based on their taste profile. The more wine you drink and rate, the better the AI and the better the wine recommendations. For merchants, Wine ring software analyzes consumer sales and ratings, and then makes inventory-based recommendations for individual consumers and households. It works in digital and physical settings, front and back of house. So retailers, restaurants, importers, distributors, producers, wine clubs, tasting rooms, trade groups - any company focused on consumer wine preference, can market and sell to preference, while optimizing inventory.

- *Source: Wine and AI: A Perfect Pairing of Technology and Tradition, Amy Cravens, 6 March 2017, https://thespoon.tech/wine-and-ai-a-perfect-pairing-of-technology-and-tradition/*
- *Source: Wine Ring website. https://www.winering.com/faq/*

Company Mini-Case #4 – Wine-Searcher

Wine-Searcher is a wine price comparison website originated in New Zealand. The Wine-Searcher database and search engine bring together wines and prices from merchants around the world. It was created in 1999, and is now used by millions to locate, compare and purchase wines. The database grows daily and is constantly monitored for quality. Manual and automated procedures are run daily to remove lists that are out-of-date or incorrect in any way. If you are a merchant or winery, it's free to add your price list. The search engine offers two levels of service: the free version prioritizes the results of Wine-Searcher's sponsors. The pro version shows all results for the searched-for wine. Wine-Searcher uses artificial intelligence to classify wines, linking the hundreds of thousands of products and tens of thousands of retailers to produce wine suggestions and pricing based on inputted search terms. Wine Searcher is also integrating label recognition technology and developing a chatbot to improve user interaction with the site. Wine-Searcher's website is indeed piloting Casey, an AI-engineered chatbot that strives to make wine purchase suggestions. Casey starts with a question as to whether the user is looking for a wine to go with food or just for easy drinking, helping less-knowledgeable drinkers to engage and start a conversation. Aiming to be an AI somm, it uses an interface like a messaging app, and opens in a dedicated browser window. You can ask it questions about the kind of wine you're looking for, and it can make a recommendation, including about what food to pair it with and a nearby store where it's available to buy. Casey is able to understand what a user is saying, taking that language and turning it into a database query, all while keeping the interaction conversational, allowing users to effectively search our database without requiring an exact wine's name or any wine-specific vocabulary. the real marker of the AI is that it understands wine and wine language – it knows appellations, grapes, vintages, wine styles, food and where the wine is sold. Casey's recommendations are made by comparing the average price of a wine with its average critic and user scores. It then gives the wines a new value, called the QPR, or quality price ratio, and recommends you the wines with the highest QPR that are available near you and match what you've asked it. You can also manually change your location and currency and check out wines available in other regions if you'd like.

- *Source: Wine and AI: A Perfect Pairing of Technology and Tradition, Amy Cravens, 6 March 2017, https://thespoon.tech/wine-and-ai-a-perfect-pairing-of-technology-and-tradition/*
- *Source: Wine-searcher website, https://www.wine-searcher.com/m/2017/10/ai-know-there-s-an-answer*

Company Mini-Case #5 - Lincoln Agritech Ltd

Lincoln Agritech Ltd, a New Zealand research and development company owned by Lincoln University is developing a computerised system to make early-season predictions on the grape yield a vineyard is likely to harvest. Lincoln Agritech is working on creating a more convenient system that uses electronic sensors to accurately count grapes. The sensors will capture and analyse grape bunches within individual rows, and assess the number, sizes and distribution of grape bunches. The data will then be fed into computer algorithms, which have been designed by the University of Canterbury, to predict grape yield at harvest time. New data will be added to the system each year, leading to continuous improvements in the model's accuracy, with the system's predictive power improving over time as more data is gathered under different conditions. Profitable wine production depends on early knowledge of the grape yield that is likely to be harvested each season. Estimating the yield as soon as possible allows marketers to know how much wine will end up being produced. The project is funded by the Ministry of Business, Innovation and Employment (MBIE) and NZ Winegrowers. Collaborating partners include Plant and Food Research, Lincoln University, the University of Canterbury, CSIRO (Adelaide), NZ Winegrowers and local

winegrowers in the Marlborough region. This 5 years research programme aims to develop a tool to deliver much more accurate yield estimations and will benefit the industry by supporting better crop management, smoother processing and market forecasting based on capacity to supply. It will also reduce manual process. This will make the job easier and the results more accurate. It will indeed require less workers as people that would normally be doing the manual sampling will now be instead capturing images and feeding the data through an algorithm and getting the results that way.

- *Source: Artificial Intelligence to help wine profits flow, 29 September 2017, https://www.mscnewswire.co.nz/keeping-in-touch/item/5873-artificial-intelligence-to-help-wine-profits-flow*
- *Source: Lincoln Agritech website, https://www.lincolnagritech.co.nz/capabilities/capabilities-and-projects/grape-yield-scanner/*

Tool Case #1 – Visual and Facial recognition

Visual recognition is applied to build up databases and pull information from databases through API.

For instance, the customer can either scan the bottle label or open the bottle, photograph the cork and use the apps "visual recognition" feature to compare the on-screen shot and the cork that's just been pulled. Both these actions generate all the information stored in the database, including a picture of what the bottle should look like and a description of how the wine should taste.

Vivino enables users to take a photo of any wine label or wine list with their mobile device, and it instantly identifies the wine delivering ratings, reviews and average pricing for every bottle. To deliver this experience, Vivino required reliable image recognition on a massive scale. First, recognizing the precise bottle in your hand is a significant challenge. Not only does the app need to recognize one brand from the next, but it must also discern the vintages within a brand. To accomplish this, Vivino relies on Vuforia's advanced computer vision technology to identify the exact wine label captured by the app. Second, to handle the massive number of unique wines, Vivino uses the Vuforia Cloud Recognition Service to store each wine label image. Vuforia then matches the incoming scans from the app with the appropriate bottle. By hosting the images in the cloud, Vivino can update this database dynamically, and manage millions of image targets for a single app without updating the app itself. The Vuforia Cloud Recognition Service is an enterprise class Image Recognition solution that enables developers to host and manage Image Targets online. Cloud Recognition is available with the Development, Cloud, Pro, and Enterprise licenses. Usage is determined by the total number of image recognitions, or "recos", per month that your app performs and is counted when a target is matched.

Pushing the visual recognition further, we can apply it to the consumer's facial recognition level. Face recognition, which is already in use on some computers and smart phones, will play an even greater part in the new sales technology for wineries. At wineries or retailers, potential customers will be able to order online without having to login or enter credit card information. Potential buyers will also be able to talk to AI bots – those sales queries that pop up on computer screens – both by voice and through the written word – once they are online at a retailer or winery's site. Consumer choosing and buying behaviors will be recorded and analysed to feed the marketing mix.

- *Source: Artificial Intelligence Changes the Wine World, Liza B. Zimmerman, 30 January 2018 https://www.wine-searcher.com/m/2018/01/artificial-intelligence-changes-the-wine-world*
- *Source: vuforia website, https://vuforia.com/case-studies/vivino.html*

Tool Case #2 – Pattern recognition and Neural networks

Vivino uses a combination of pattern recognition and OCR (Optical Characters Recognition). The engine used for the pattern recognition comes from www.kooaba.com but is powered by Vivino's own "bucket" of more than 50 million wine labels, all sent by the users. While this pattern recognition figures out which existing photo is the closest match, Vivino run an OCR analysis in parallel, and use this to fine-tune the vintage (year) and also often the grape information of the wine. The whole process takes less than a second on the servers. The app then tracks which wines the user scan and rate.

More largely, WineEngine combines image recognition algorithms with neural networks to provide fast and reliable recognition of wine, beer and spirit labels. Image recognition is used to provide a better user experience and more reliable results than approaches based on barcodes, Optical Character Recognition or QR codes. The process of setting up the API begins with adding the reference label images into the WineEngine collection. The API automatically generates a unique fingerprint for each label image as there are added. These fingerprints are stored in a private database per account. New label images can be added to the collection as needs grow and they will be immediately searchable. To search for a label, a photo of the bottle is sent to the API. WineEngine then generates a fingerprint for the bottle label and instantly compares it against the fingerprints of the entire collection of labels.

To improve the quality of the matches provided, a neural network that automates the process of carefully cropping the photos is used. The neural network locates the label in the photo and masks the rest of the image, removing irrelevant content from the photo to enable more robust label matching. Search results are returned quickly and provide filenames for all matching images from your collection. Wine label searches also return the vintage year detected from the query image.

- *Source: https://www.quora.com/What-are-the-basic-computer-vision-and-recognition-matching-techniques-used-by-the-Vivino-wine-scanner*
- *Source : TinEye website, https://services.tineye.com/WineEngine*

Tool Case #3 – Robots

Still very new and on the back of self-driving cars, the Wine robots wander the vineyards gathering information about the state of the vines (such as vegetative development, water status, production, and grape composition), just like vineyard managers used to do.

VinBot is an all-terrain autonomous mobile robot (Summit XL) with a set of sensors capable of capturing and analysing vineyard images and 3D data by means of cloud computing applications, to determine the yield of vineyards and to share information with the winegrowers. It can work 8 hours per day and climb slopes up to 45°. Developers say the robot will allow winegrowers to become more competitive. VinBot responds to a need to boost the quality of European wines by implementing precision viticulture (PV) to estimate the yield (amount of fruit per square metre of vine area: kg/m2). Winegrowers need to be able to estimate yield accurately to perform yearly canopy management techniques, and harvest vineyard areas sequentially, according to the optimal ripeness of the grape in each area, which improves wine quality.

Wall-Ye V.I.N. robot is self-propelled and carries out the labor-intensive vineyard tasks of pruning and de-suckering (removing unproductive young shoots) to recording important data on the soil, fruit and vine stocks. Wall-Ye draws on tracking technology, artificial intelligence and mapping to move from vine to vine, recognise plant features, capture and record data, memorise each vine, synchronise six cameras and guide its arms to wield tools.

TED roams around on wheels to cultivate soil and uproot weeds. This robot neatly weeds between the vineyard rows (for those people who don't use sheep to keep their weeds under control). Perhaps the world's most prestigious wine maker, Château Mouton Rothschild, is testing TED at its Château Clerc Milon estate. In June, the company teamed up with Naio Technologies in an experiment. Philippe Dhalluin, managing director of the châteaux, is convinced robots are part of a "green" future.

Another example of wine robot. is the VineScout, which will be used in Portugal at Symington Family Estates' vineyards during the 2018 season to check vine vitals such as leaf temperatures.

It will be possible to automate a vineyard completely within 5 to10 years. Robots will do the whole thing. We already are getting drones to do the flying over the vineyards, which can tell you which vines need water or treating.

- *Source: https://www.science20.com/news_articles/wine_production_now_with_more_robots-152703*
- *Source: Artificial intelligence in the wine industry? Not yet, please! David Morisson, 2 April 2018, http://winegourd.blogspot.com/2018/04/artificial-intelligence-in-wine.html*
- *Source: Wine Robot 'Wall-Ye' Performs Complex Vineyards Tasks For $32,000, Rachel Tepper, The Huffington Post, 7 December 2017, https://www.huffingtonpost.in/entry/wine-robot-wall-ye_n_1923808*
- *Source: Robots in vineyards, giant bottles, and bubbly: Here's what the wine world will look like this year, Bloomberg, 5 Jan 2018, https://economictimes.indiatimes.com/magazines/panache/robots-in-vineyards-giant-bottles-and-bubbly-heres-what-the-wine-world-will-look-like-this-year/articleshow/62376308.cms Source: Vine scout website, http://vinescout.eu/web/newsgallery-2*
- *Source: Robotnik website,* https://www.robotnik.eu/vinbot/

Modeling the complex human taste is an important focus in wine industries. The main purpose of most studies using **classifiers tools** was to **predict wine quality based on physicochemical data**. Here below is a description of the various tools I could find and related studies. As a conclusion, alcohol content and volatile acidity levels seem to be the most influential factors, followed by sulphates, citric acid and fixed acidity levels to define the quality of a wine.

- *Source: How to Use Machine Learning to Predict the Quality of Wines, Ashwin Hariharan , 7 February 2018 https://medium.freecodecamp.org/using-machine-learning-to-predict-the-quality-of-wines-9e2e13d7480d*

Tool case #4 - Naïve Bayesian / Bayesian classifier and Decision Tree

The objective of this classification experiment was to investigate physicochemistry properties in wine that influence the taste, hence the quality of a wine. The dataset is a wine quality dataset that is publicly available for research purposes. It consists of two separate datasets, red wine and white wine. Both dataset contains 1,599 instances with 11 attributes for red wine and 4, 989 instances and the same 11 attributes for white wine. Each instance is classified into quality attribute that ranges between 0 (very bad) and 10 (excellent). The attributes are physicochemical data such as alcohol, PH and sulfates. Two classification algorithms, Decision tree and Naïve Bayes are applied on the dataset and the performance of these two algorithms is compared. Both datasets are separated into training and testing set by using 10-fold cross-validation method. The training data is randomly portioned into 10 sets of equal size and the algorithms are executed 10 times.

Decision tree-: The first algorithm chosen is decision tree using ID3 method. The ID3 method gives the highest accuracy among other decision tree's methods such as random tree, FT (Functional Tree), NBTree, and simple cart in WEKA. This method builds the tree from top down, with no backtracking. Furthermore, by using this method, the tree is built faster and the leaves are shorter when compared to other methods.

Naïve Bayes: The second algorithm is Naïve Bayes algorithm that is able to find two highly correlated attributes that is similar to attribute selection purpose. By using Naïve Bayes algorithm, the model is built faster and it is highly scalable as compared to other Bayes' algorithms in WEKA.

Results showed that Decision tree (ID3) outperformed Naïve Bayesian techniques particularly in red wine, which is the most common type. The study also showed that two attributes, alcohol and volatile-acidity contribute highly to wine quality. White wine is also more sensitive to changes in physicochemistry as opposed to red wine, hence higher level of handling care is necessary. This research concludes that classification approach will give rooms for corrective measure to be taken in effort to increase the quality of wine during production.

- *Source: Classification-based Data Mining Approach for Quality Control in Wine Production, P. Appalasamy, A. Mustapha, N.D. Rizal, F. Johari and A.F. Mansor, 2012. Journal of Applied Sciences https://scialert.net/fulltext/?doi=jas.2012.598.601*

Tool Case #5 – Logistic regression, Decision Tree and Random Forest

Ordinal logistic regression, decision tree, and random forest can also be used as methods to create models to predict wine quality based on wine chemical properties. In these models, chemical properties are used as the independent variables. Wines are divided into three groups based on wine quality. Wines with quality score less than 6 are grouped as 1, wines with quality score 6 are grouped as 2, and wines with quality score better than 6 are grouped as 3. Also, data is split into train data and test data to check the performance of the models. Random forest model seems to be the best model to predict wine quality according to wine chemical properties, as the accuracy rates of both red wines and white wines exceed 0.7. However, in this study, the problem is simplified by sorting wine scores into three groups. To make more accurate prediction, such as predicting the exact score, we will need better models and the accuracy rates may decrease.

- *Source: Wine Data Analysis Part III 3 September 2016* https://rstudio-pubs-static.s3.amazonaws.com/206313_6bb9329710904ab29387b303688099d5.html

Tool Case #6 – K Nearest Neighbors

Another example applies K-Nearest Neighbors to wines Quality Classification (in R) and to predict wine quality given a dataset of red wine samples and their quality, e.g. low, medium, high. If we take a test sample with an unknown quality and the task is to correctly classify the wine using a set of physicochemical features, e.g. acidity, density, alcohol, pH, k-Nearest Neighbors identifies the k number of observations that are most proximate to the test sample, as defined by some distance metric, e.g. Euclidean. From this set of k-neighbors, majority rule is used to predict the class. If k=3 and the nearest neighbor wines' qualities are {*low, low, medium*}, then we would classify the test sample as a low-quality wine. The same approach is extended to subsequent test samples.

Source: Introduction to K-Nearest Neighbors with Red Wines Quality in R , Chansoo Song, 19 February, 2018, https://medium.com/nyu-a3sr-data-science-team/k-nn-with-red-wines-quality-in-r-bd55dcba4fd7