UNIVERSITA' DEGLI STUDI DI PADOVA

DIPARTIMENTO DI SCIENZE ECONOMICHE ED AZIENDALI "M.FANNO"

CORSO DI LAUREA MAGISTRALE IN BUSINESS ADMINISTRATION

TESI DI LAUREA

"DIGITAL WINE: HOW PLATFORMS AND ALGORITHMS WILL RESHAPE THE WINE INDUSTRY"

RELATORE:

CH.MO PROF. Andrea Marella

LAUREANDO: Piercarlo Michelin

MATRICOLA N. 1184946

ANNO ACCADEMICO 2021 – 2022

Il candidato dichiara che il presente lavoro è originale e non è già stato sottoposto, in tutto o in parte, per il conseguimento di un titolo accademico in altre Università italiane o straniere.
Il candidato dichiara altresì che tutti i materiali utilizzati durante la preparazione dell'elaborato sono stati indicati nel testo e nella sezione "Riferimenti bibliografici" e che le eventuali citazioni testuali sono individuabili attraverso l'esplicito richiamo alla pubblicazione originale.

The candidate declares that the present work is original and has not already been submitted, totally or in part, for the purposes of attaining an academic degree in other Italian or foreign universities.
The candidate also declares that all the materials used during the preparation of the thesis have been explicitly indicated in the text and in the section "Bibliographical references" and that any textual citations can be identified through an explicit reference to the original publication.

Firma dello studente

TABLE OF CONTENTS

INTRODUCTION: WINE IN AI EARLY AGE

Wine is a complex product. From the vineyard to the glass, wine is the end result of a long chain of biological and chemical processes mediated by human intervention. The successful combination of nature, expertise, technology and accumulated traditions leads to an outstanding product, which is highly appreciated for the varied and complex flavor profiles one can experience during tasting, leading the consumer to an end state often defined as pleasure.

Wine journey starts in the vineyard where climate, soil characteristics and microbial communities affect grape chemical composition and therefore shape the potential sensory attributes of the final wine (van Leeuwen et al. 2020; Liu et al. 2019). In the winery, yeast-induced fermentation transforms grape juice or must into wine, and other processes such as clarification and ageing take place in order to obtain a product which expresses the desired flavor characteristics, while being safe to drink and appealing. The resulting wine chemistry is by far more complex than that of grape juice, and showcases thousands of volatile and non-volatile compounds which compose the flavor profile. Despite nature does most of the work, vineyard management practices and winemaking techniques are critical in ensuring that the final wine stands up to the expectations and possibly expresses traits of typicality or style that makes it unique and distinguished. Making the wrong predictions or failing to take adequate and timely actions may result in compromising the quality, or even worse in producing an unmarketable, faulty wine.

It is in the glass that the flavor profile finally unfolds. The nose, the mouth, the eyes are all stimulated during wine tasting and their sensations are then processed by the brain into a unique perception of flavor (Sheperd 2015; Smith 2019). Besides the sensory and motor systems, wine tasting activates central brain systems such as emotion, memory, language and reward, thus validating the claim that "…more brain systems are engaged in producing flavor perceptions than in any other human behavior" (Sheperd 2015, p. 2). In this way the tasters can perceive the wide range of flavor notes that the complex wine chemistry produces, such as red berries, tropical fruits, citrus, but also flowers, spices, wood, and others quite unexpected such as leather or boiled potato, which are not appealing by themselves but which contribute to create a balanced and harmonic flavor profile, that represents the quintessential definition of wine quality.

However, the experience of quality by consumers is necessarily linked to hedonic liking and pleasure, and liking the taste of wine has been established to be the main reason why people drinks wine by both literature (Charters & Pettigrew 2002) and market research (Thach & Chang 2015). However, despite consumers can have an immediate hedonic reaction when drinking, they have a significantly harder time in explaining why they like or not a wine. Experienced tasters have a superior capability to efficiently discern wine tasting notes, picking up the varietal, typicality and stylistic expressions of a wine and building up the ability to correctly discriminate among different wines, thus adding a psychological pleasure to their wine experience. This do not seem to come from a superior perceptual ability, but instead from experience and knowledge (Yang & Lee 2020).

The development of individual taste preferences is still far to be understood. Clearly, interpersonal differences may play a relevant role. Sensitivity to bitterness or the inability to detect specific odors are classic examples of physiological differences, as well as oral microbiota composition (Pérez-Jiménez et al. 2021). However, given the substantial involvement of the bran in generating the perception of flavor, psychological differences may play a bigger role, and represent an active field of research, as demonstrated by the growing body of literature on wine-evoked emotions (Niimi, Danner & Bastian 2019).

The inexplicability of individual taste and the difficulties of wine tasting represent a relevant hurdle for consumers as they often can't taste the wine prior purchasing, especially in the large-scale and online retail environments which represent a relevant share of market sales. This is why consumers typically refer to extrinsic cues of quality in order to be reassured of the fact that the wine they are going to purchase will meet their quality expectations. The most often cited extrinsic cues in the literature are price, grape variety, origin, brand and price (Goodman 2009; Pomarici et al. 2021). However, a big question mark can be raised on whether reliance on such extrinsic cues can represent an optimal strategy in wine choice, and on whether they can tell the inexperienced consumer something that goes beyond the mere expectation of quality (whatever quality means for the end consumer).

Another relevant phenomenon which further complicates wine choice is the proliferation of brands. It is exemplifying of this phenomenon the fact that the number of labels in the Italian wine market have doubled in the last 20 years (Mediobanca 2020). This often leads wine consumers to experience what Barry Schwartz defined as the "paradox of choice" (2004, cited in Ricci, Rokach, & Shapira 2015, p. 2), which indicates the phenomenon where the bewildering range of products offered to consumers is detrimental of their happiness, inducing them in a state of anxiety. Given that wine typically needs a relevant share of involvement and knowledge in order to elaborate the numerous wine cues, it can be reasonably argued that wine consumers experience more intensely this state compared to other categories. This is what empirical research is suggesting, as consumer purchasing decisions can often be traced back to risk-reduction strategies. In particular, recommendations and having previously tasted the wine rank among the top choice attributes in several consumer researches (Goodman 2009; Pomarici et al. 2021).

Given this context, it is not surprising that innovative startups and app developers considered wine choice a problem worth solving. Wine apps represent an early example of merging the physical and digital retail environments (Rooderkerk & Kök 2019) by delivering timely information directly to the smartphone upon taking a picture of a wine label, therefore behaving as *shopping assistants*. In this way consumers can access thousands of user ratings and reviews, or receive personalized recommendations, expressed as a percentage of match between the wine and the individual taste of the consumer. Wine apps that successfully gather a massive quantity of user-generated content can exploit a form of *wisdom of the crowd*, as users can benefit from the knowledge generated collectively by the whole users of the network.

In order to generate recommendations, collaborative filtering algorithms are fed up with users' ratings of the available items; then, they are capable of recommending new items to a single user by considering the ratings that similar users have given to such items, where similar users are the ones whose rating history overlaps with that of the considered user. The underlying assumption is that users with similar rating histories have similar tastes. However, such recommendation approach is often considered to work as a *black-box*, in the sense that it is solely based on numbers regardless of any other form of domain knowledge, and the explanations for a given recommendation are difficult to extract, except that other unknown users tend to like similarly the same items.

As AI increasingly permeates our everyday life, greater emphasis is being put in designing algorithms that are explainable (Calegari, Ciatto & Omicini 2020) and, at the same time, that instill trust in the end-user, the latter aspect especially relevant in the online shopping environment (Castellano et al. 2018). Furthermore, since the main objective of recommender systems is to improve decision-making, it would be straightforward to include some elements of the real-world wine choice problem in a digitized form in order to produce more explainable, tailored suggestions. This may be the reason why many wine apps started to introduce new features such as the *palate profiling*, thus promising to deliver truly personalized recommendations finally based on the unique (and not shared) taste of the user. What this approach does is simply creating a detailed user profile, where preferences towards specific features are collected, and then matching the user profile with the wines, which are also represented as a set of features. By effectively collecting accurate taste profiles and wine representations, tailored recommendations can be delivered, treating every consumer as a *segment of one* (Kotler, Kartajaya & Setiawan 2021). This approach, called content-based, appears particularly fitting with a complex problem such as the choice of wine.

The first and most evident benefit of such approach is that not only app developers get to know more about the users, but the users get also to know more about their taste preferences and the features they like or dislike about wines. In this way wine apps can serve as an instrument to develop users' wine knowledge based on wine intrinsic organoleptic properties and even to modulate their experience of wine, as it has been demonstrated that wine descriptors can improve the enjoyment of wine by setting the right expectations in the tasters (Danner et al. 2017). Such a user experience could give the consumer enough confidence to undertake the quality ladder towards wines of greater complexity and price (Smith 2019). But another potentially disruptive possibility is to perform sensory science at scale, by analyzing aggregated taste preferences and using them as a basis for experimentation and continuous learning. Considering that multiple studies seem to suggest that clear clusters of wine consumers can be identified based on their preferences towards some key flavors in wine (Francis & Williamson 2015), such an app could enable to perform a flavor-based segmentation, generating invaluable insights about consumer tastes trends and better advising how to market existing wines.

Moreover, considered the digital nature of wine apps, their potential is not limited to their core value delivered, but it extends to the set of connections they could establish along the supply chain. Taste trends data represent a valuable resource to optimize the logistics and distribution phases based on data. In the upstream of the supply chain, comprehensive chemical analysis, quantitative sensory data, and consumers' liking scores could be put in relation, and in this way relevant relationships between wine chemistry and consumer preferences could be discovered; such data-driven relationships could be translated into practices in the vineyard or the cellar by setting up an integrated data ecosystem, thus fostering the *precision viticulture* and *precision winemaking* approaches.

Wine platforms have the potential to fulfill the need for innovation in the wine digital customer experience while at the same time promoting the digitalization of the overall industry by providing a technological foundation which delivers critical services and offers connectivity. This happens in a scenario where the wine market is undergoing a period of deep transformation, facing the challenges of globalization, climate change, shifting consumers preferences and generational change, and where digitalization has been addressed by the EU to be the driver of change towards reaching sustainability and prosperity in the agricultural field (ISMEA 2020).

However, reaching a long-term industry-wide impact will depend on how wine apps companies design both their data ecosystem and their business model. The way in which wines and consumers are digitally represented (for example as a set of numbers or keywords) and the amount of knowledge that is put into the ecosystem (for example wine dictionaries or wine chemical data) will set a limit on the potential of the system to generate new knowledge. Similarly, the design of the user experience will determine how relevant data users will share with or generate within the app. Then, the definition of a successful business model which exploits the full potential of its digital nature is a necessary condition for striving in a hyper-competitive market. However, in the digital world the business model seldom represents a static concept, as continuous learning from data enables business model innovation and permits to adapt to a constantly evolving environment. Therefore, the objective of this thesis is twofold:

i. To examine the main approaches to digitally represent wine and consumer preferences;
ii. To assess how such digital representations are leveraged by digital business models in the wine industry.

An additional objective is to assess the impact of such business models on the overall wine industry, if any.

In the remainder of this chapter, a synthetic overview of the global wine market will be presented, while at the same time addressing current and future challenges and trends. Then, given their increasing relevance for the wine industry, the main AI applications in the wine supply chain and the foundations of digital business models will be examined.

In Chapter 1, the business models and the operating architecture of the AI-powered company will be addressed. Then, the basics of machine learning will be briefly presented coupled with an overview of recommender systems.

In Chapter 2, the approaches to represent wine digitally will be investigated by addressing the key sources of wine data, namely the online ecosystem and the chemical analysis of wine. The main data mining approaches will be exemplified through a case study on the Portuguese Vinho Verde wine by leveraging real-world data.

In Chapter 3, the logic and purpose behind a flavor-based recommender system will be explicated. Two business cases will be presented, the industry leader Vivino and the innovative startup Tastry. Both companies aim at matching wine with the unique taste of each wine consumer, but they do so through different data-driven approaches, which result in significantly different business models. Such diversity implies a different positioning along the wine supply chain and a different role on the wider digital wine ecosystem.

Finally, some conclusions will be drawn about the impact of the phenomena covered in this thesis on the future of the wine industry, in particular as regards the dynamics of competition and the mechanisms through which value is generated within a digital ecosystem.

I.1 Wine Industry Overview

The wine industry generated revenue for around $380 billion in 2019 (Statista 2021). While during the 1970-79 decade almost 60% of the wine in the world was produced and almost 50% was consumed within four European countries (Italy, France, Spain and Portugal) (Anderson, Pinilla & Holmes 2021), today wine represents a globalized market, where wine is exported all over the globe. Even though Italy, France, and Spain retained their leadership positions in terms of production, the globalization of the wine market profoundly reshaped the consumption patterns and the competitive environment. Such process, which accelerated since the 1990s, was exploited by New World (NW) wine producers such as Australia, Argentina, the US, and Chile, that produced highly appreciated wines and increased their competitiveness in the export market (Pomarici et al. 2021). Meanwhile, consumption consolidated in new countries around the world, with the US overcoming Italy and France as the first wine consumer market in the world.

By looking at the total world wine production and consumption over the last 20 years, as shown in *Figure 1*, it appears that the global wine market has somehow stabilized, as the globalization phenomenon reached maturity in this two decades. However, these curves are the result of a series of profound changes that are still ongoing.

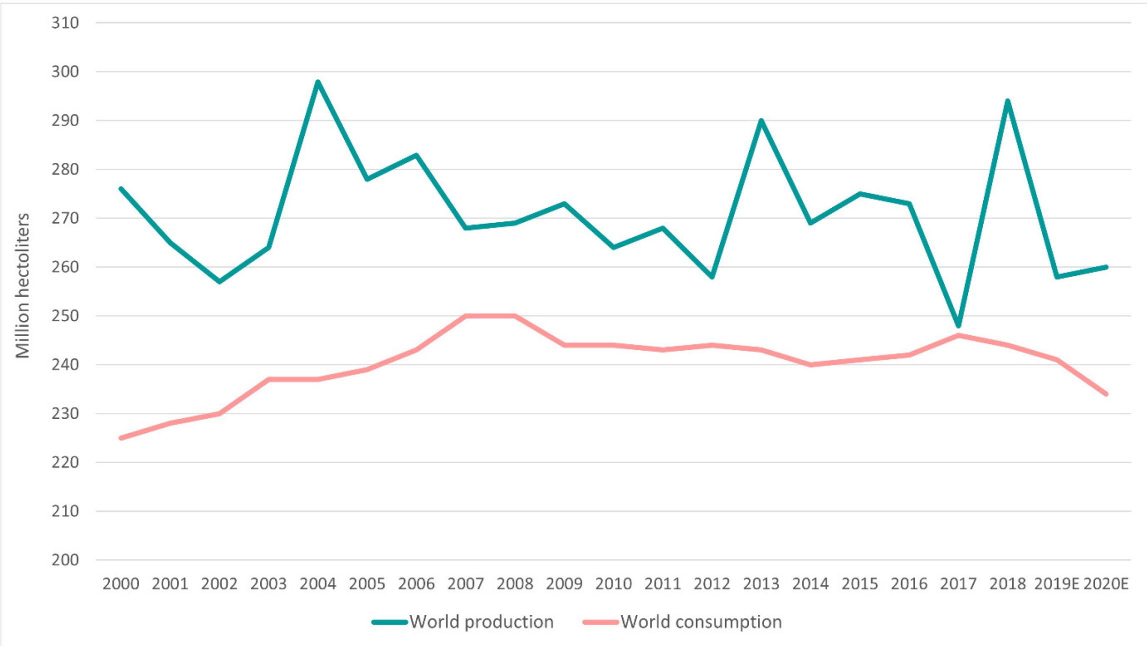

Figure 1. World wine production vs. consumption in the period 2000-2020, expressed in million hectoliters. Source: own elaboration on OIV (2018; 2021a) data.

On the production side, a first relevant feature to observe is the volatility of wine supply. This is due to the intrinsic nature of wine production, which can take place only once a year from fresh grapes and which is strongly influenced by weather conditions. As production decisions must be taken in advance by relying on forecasts about final demand and weather evolution, matching a volatile wine supply with a stagnating and fragmented demand is a critical and difficult task. Furthermore, an enduring excess of production over consumption can be observed, except for 2017 which represented an all-time low harvest. Part of this excess (which is on average around 30 mhl in the period 2000-2020) is due to the destination of a share of the wine produced to the distillation of spirits (such as brandy) or vinegar production. However, the remaining part is due to wine and grape oversupply. In order to mitigate this problem, a general tendency has been to reduce the vineyard surface area over the years, while increasing yield and quality (ISMEA 2020).

The composition of the world wine production in 2019 (*Figure 2*) highlights the enduring leadership of Italy, France, and Spain, which together still produce around 47% of the wine in the world. Italy consolidated as the world leader in wine production, and remains competitive across all price segments thanks to its vast heritage of grape varieties and its fragmented but efficient production system (Pomarici et al. 2021).

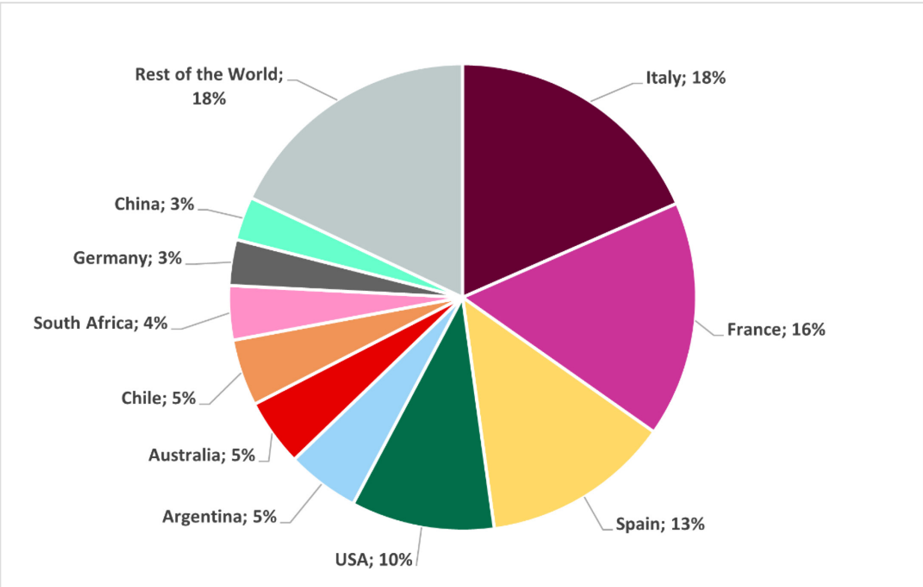

Figure 2. Composition of world wine production in 2019. Source: own elaboration on OIV (2021a) data.

Except the change in the lead between France and Italy, no significant changes occurred in the composition of world production during the period 2000-2020.

An evident change happened with regard to wine consumption, with the US market having a striking growth from around 22 mhl in 2000 (Anderson & Nelgen 2011) to 33 mhl in 2019 (OIV 2021a), making it the worlds' largest market for wine consumers. A similarly, outstanding growth happened in China, which passed from being a marginal market in the beginning of the millennium to represent today one of the most interesting markets for future growth, with around 15 mhl of wine consumed in 2019 (OIV 2021a). The current ranking in terms of world wine consumption is shown in *Figure 3*.

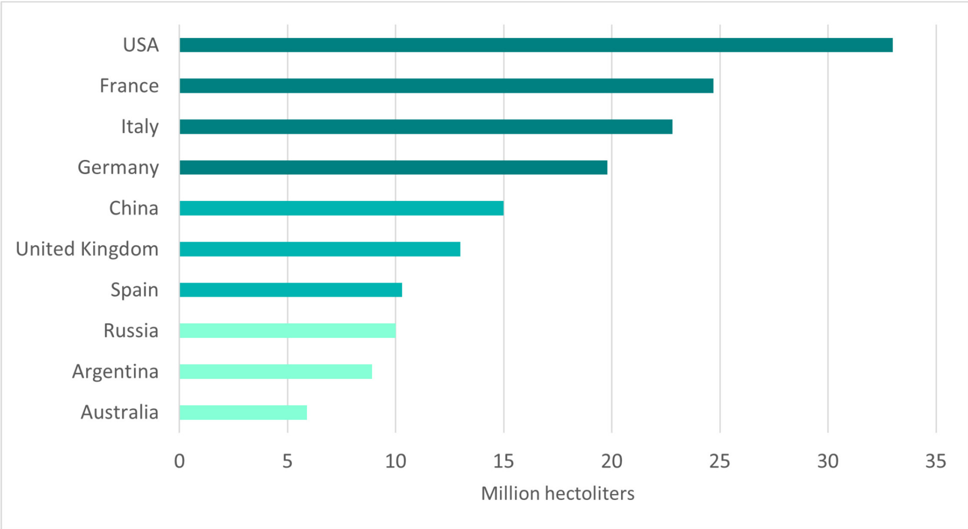

Figure 3. Ranking of world wine consumption in 2019. Source: own elaboration on OIV (2021a) data.

This ranking is the result of a longstanding tendency of decreasing consumption in the Old World countries, fostered by changing drinking patterns in the European population. While France, Italy and Spain national consumption decreased by 16%, 17% and 21% respectively between the five years periods 2005-09 and 2015-19, wine consumption increased in the US, China and Australia by 22%, 24,3%, and 21% in the same period (ISMEA 2020). The result of this dynamics can be appreciated in *Figure 4*, where it can be observed how the main New World countries' aggregated consumption reached and then overcame the core Old World countries' one, which suffered a steady decline in the first decade of the century.[1] This clearly shows how the New World countries were responsible for the mild overall growth of the world wine consumption in the last 20 years. It has to be notice however a relative stabilization in the last 5 years, thanks to an ongoing positive trend in Italy and the solidity of the German market.

[1] For this calculations, France, Italy, Spain, Portugal, and Germany were considered as core Old World of wine countries. On the other side, the New World countries were represented by: UK, Russia, Australia, Canada, US, Argentina, South Africa, and China.

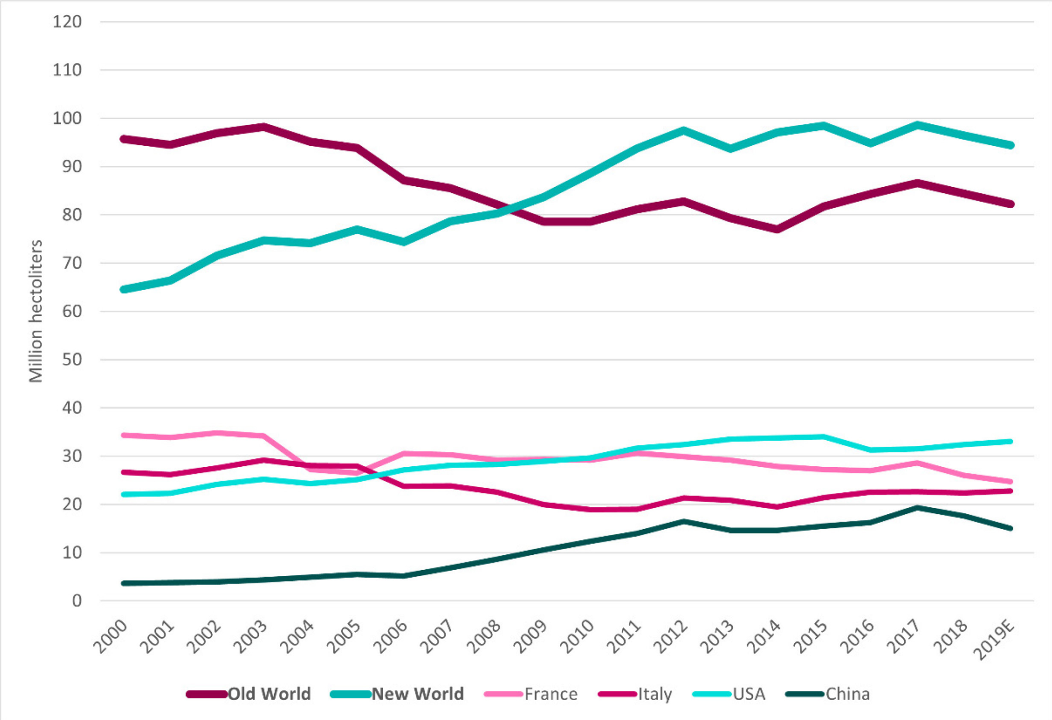

Figure 4. Dynamics of Old World vs. New World wine consumption. Source: own elaboration on Anderson & Elgen (2011), Anderson, Pinilla & Holmes (2021) and OIV (2021a) data.

The most defining dynamics of the last 20 years can be observed in the international trade of wine. In the period 2000-12, the share of world wine exports on world consumption increased from 27% to over 43%, and a similar increase occurred for the weight of exports on world wine production, passing from 23% to 41% (Anderson & Nelgen 2011; OIV 2018). In other words, today almost half of the wine in the world is consumed in a country different from the one in which it was produced (ISMEA 2020).

It appears clear then that the decrease of internal demand in Old World countries has been compensated with a sharp increase in exports, a hypothesis which is validated by the strong performance on the international markets of France, Italy and Spain which are the main exporters of wine both in terms of volume and value. However, they consolidated this position through significantly different strategies. In *Figure 5* the overall picture of world wine export in 2019 is presented by considering the volume of export in million hectoliters, the euro/liter price and the value of exports in million euros. France appears to hold a dominant position, with an outstanding value of exports of almost €10 billion. This figure is reached with a sensibly lower volume of exports compared with Italy and Spain, and is due to the remarkable price at which French wine is traded, on average almost €7 per liter. This price, almost double than the world average, reflects the positioning of France in the higher end of the wine market, which is fostered by the mythical aura around French wine and its well-known designations.

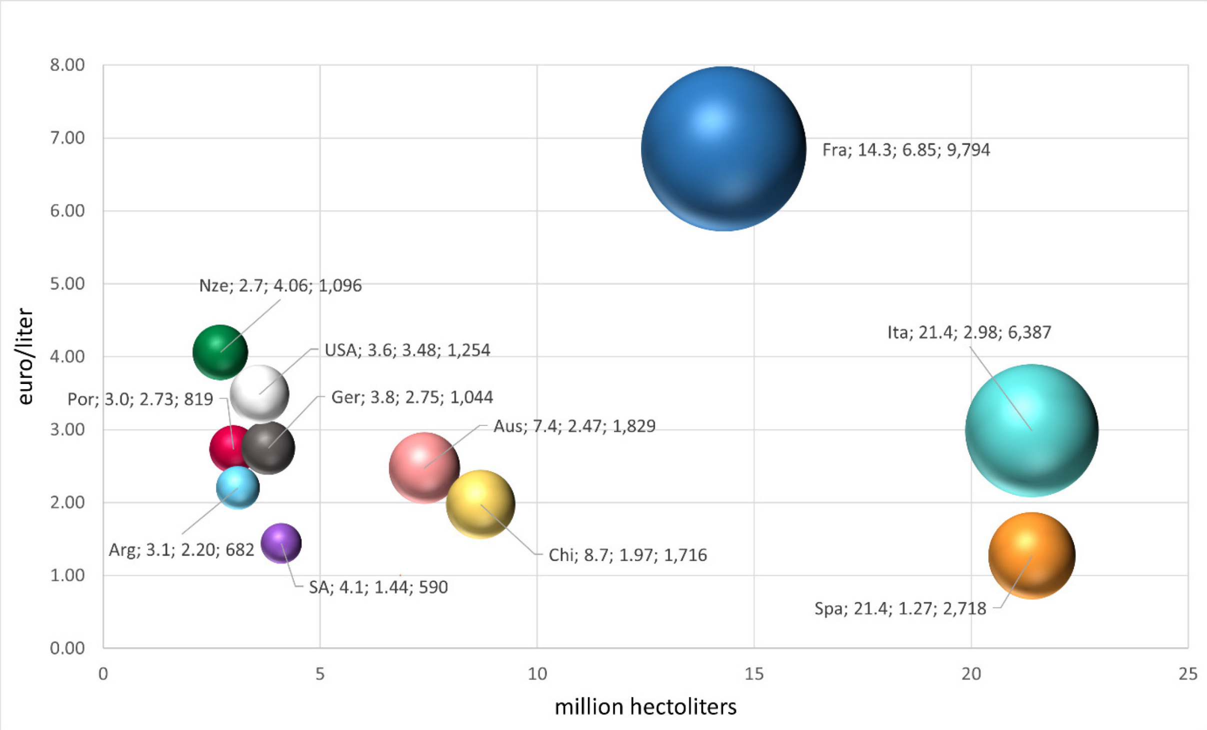

Figure 5. Main exporters of wine in 2019 (million hectoliters, euro/liter, million euros).[2] Source: own elaboration on OIV (2021a) data.

Italy and Spain share the first position in terms of volume at around 21 mhl; however, they perform substantially different in terms of value. Italy's vast and diversified wine production is capable to compete in all price segments, and the recent revaluation of its high-end wines enables Italy to reach more than €6 billion in terms of export value at an average price of €3 per liter. Spain exports instead are mainly directed towards the lower-end of the market, where big volumes of bulk wine and of wines without a designation of origin are traded at a low price. As for the New World countries, Australia and Chile show a solid performance, which is the result of a strong export orientation, where 62% and 73% of wine production is destined to the export respectively, and of the consolidated position of their brands in the global market. The performances of New Zealand (which exports 90% of its production) and the US are also worthy consideration, where the modest volume of wine are traded at an above average price, around €4 and €3.5 per liter respectively, enabling them to surpass the €1 billion mark in value. This reflects the international recognition of the quality of their wine production.

[2] On the labels are indicated in order the name of the country, the volume of exports in million hectoliters, the euro/liter price, and the value of exports in million euros. Additionally, the size of the bubbles represents the value of exports.

Over the 20-years span, the growth of wine exports in terms of volume has been accompanied by the growth in terms of value, which concentrated in the last decade. In fact, the value of world wine export has almost doubled in this period from €13.7 billion in 2000 to €31.7 billion in 2019, with the average price of exported wine shifting from €2.3/l to €3.1/l. While the volume of exports seems to have stabilized in the last 5 years at around 100 mhl, the growth in value seems more persistent, and can be interpreted within the general phenomenon of wine "premiumisation", as wine is experiencing a significant revaluation at the international level promoted by the well-established segment of *wine lovers*, who display a common knowledge and brand repertoire and who share the same appreciation towards fine wines (ISMEA 2020). Therefore, market analysts expect the wine market to continue to grow in value despite the stability of demand. However, it must be noted how the consumer segment which actually moves the market in many key countries is made of non-expert consumers who want to simply enjoy a good bottle of wine sometimes, possibly by purchasing it in the large-scale retail (ISMEA 2020). This segment is believed to be the key to understand the premiumization phenomenon, as a share of them is turning their wine choices towards *premium wines*, while at the same time another share is redirecting their choices towards entry-level *value wines*, which deliver great value at a low price. This enables to reframe the phenomenon into a "polarization" of the market, a tendency which has been observed across many other industry sectors and which has been exacerbated by the global pandemic (Kotler, Kartajaya & Setiawan 2021).

Such dynamics can be clearly observed by looking at the international exports by product categories. In 2019, the world wine exports by volume were composed by 54% of bottled wines (non-sparkling), 9% of sparkling wines, 33% of bulk wine and 4% of Bag-in-Box (ISMEA 2020).[3,4] During the period 2000-2019, the premiumization of bottled wines was accompanied by the considerable growth of sparkling wines, which are typically traded at a high average price and which were able to catch a positive momentum thanks to their association with celebrations and positive lifestyles, and by meeting the international taste due to their fresh and easy-drinking taste (ISMEA 2020). On the other side, a significant growth, bigger than that for bottled wines both in terms of volume and value, was experienced by bulk

[3] According to the World Customs Organization classification, the Bag-in-Box format consists of wine traded in packages between 2 liters and 10 liters, while bulk wine consists in wine traded in packages over 10 liters (ISMEA 2020).

[4] In the period 2000-2019, bottled wine grew by 4.2% in value and by 2.3% in volume; sparkling wine grew by 5.6% in value and by 5.8% in volume; bulk wine and BiB grew by 4.9% in value and by 3.4% in volume (ISMEA 2020). By considering the difference between the growth in value and volume, the growth of the average price can be obtained. In particular, bottled wine average price grew by 1.9%, thus confirming the premiumisation trend.

wine and Bag-in-Box. The role of this increasing supply of bulk wine is due to importing countries that buy bulk wine from producer countries and then bottle it for creating entry-level lines.

This last phenomenon appears clearer by looking at the overall wine supply chain (*Figure 6*). In the upstream and central part of the supply chain, exchanges of intermediate and finished products occur, thus promoting the smooth functioning of the market and enabling the creation of new entities which specialize at different stages of the chain and pursue different business models.

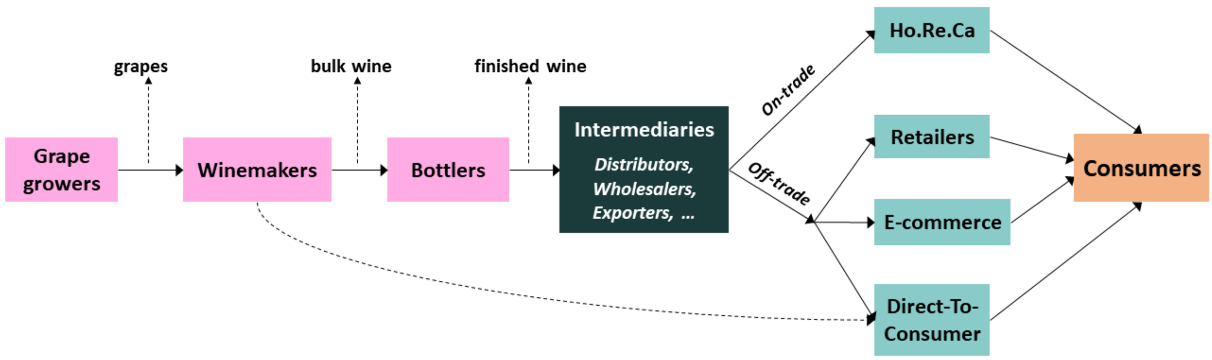

Figure 6. Wine supply chain overview. Source: own elaboration based on Pomarici, Boccia & Catapano (2012) and Pomarici et al. (2021).

A fundamental entity in many Old World countries are *cooperatives*, which gather grapes produced by a huge number of individual grape-growers, assuring them a fair remuneration for their work. Similarly, the exchanges of bulk wine enable the growth of *pure bottlers*, which buy large amounts of bulk wine in the market and then bottle it, sometimes even concluding the winemaking process by performing themselves ageing and blending (Pomarici et al. 2021). Often the competitive advantage of this business model relies on strong distribution networks and a preferential access to the destination market. This is the case of Germany, which has a minor internal production but which has a great commercial power and leverages it to re-export the imported wine (ISMEA 2020).

Another relevant topic with regard to the wine supply chain is integration. As it has been shown, matching wine supply and demand is challenging due to the intrinsic nature of wine, and competing in a globalized wine market requires access to new and geographically distant distribution chains. This is why a trend of downstream integrations has been observed, as the most established wine companies try to reach significant economies of scale and to strengthen their brand. A profoundly different strategy is that of many small family-owned wineries which perform autonomously all the stages of the wine chain and reach the final consumer by performing direct-selling, often exploiting online platforms. A niche production and a strong quality orientation has enabled them to be successful in the market and to perform well even during the pandemic (ISMEA 2020).

The final stage of the supply chain is typically split in two channels, depending on whether wine is consumed on site (*on-trade*) or at home (*off-trade*). The on-trade channel is largely represented by the Ho.Re.Ca. (which includes restaurants, cafés, catering, and restaurants), and increasingly encompasses wine bars and oenotourism, the latter representing a winning strategy for wineries to communicate the *sense of place*, which is a strong marketing element for wine. The off-trade channel can be broadly divided into large-scale retail, specialized retail (for example wine shops) direct-to-consumer selling (especially in the wine cellar), and e-commerce, which can represent both a standalone channel in the case of *pure players* (as the online marketplace Vivino) or which can be integrated into one of the traditional retail channels (for example in the case of a winery which performs direct-selling through a proprietary e-commerce website). While the relative size of the on-trade and off-trade channel is country-specific and therefore difficult to analyze at a global level (however with a general predominance of the latter), two major trends can be observed during the last two decades. Recently, the on-trade consumption has fueled the interest in wine of younger consumers in many countries both in the Old World and New World of wine, thanks to the social and ritual dimension of aperitivo and the growing interest in wine pairings for to achieve memorable culinary experiences (Wine Intelligence 2020). On the other side, large-scale distribution chains have now become the reference place where to buy wine for the majority of consumers, evolving into a space where consumers can make quality purchases by accessing to wines which were once exclusively sold through specialized retail shops, while at the same time being offered with a wide range of products for every price point (Pomarici, Boccia & Catapano 2012).

In order to conclude this industry overview, it is now the moment to make the structural differences between Old World and New World countries explicit, as they have been described by Pomarici et al. (2021). A first relevant difference regards the product, as Old World wine producing countries typically rely on their native grape varieties in order to offer a unique product in the market, and leverage protected designations of origins as an extrinsic cue of quality. On the contrary, New World countries mostly rely on well-known and appreciated international grape varieties by pursuing a more consumer-focused strategy, and make use of strong and recognizable brands in order to attract consumers. A second significant difference is the export orientation, as New World countries typically export more than 70% of their production, while in the Old World countries the strong domestic consumption absorbs a larger share of their production. A third fundamental difference, which is a reflex of the first two, is the different degree of concentration. New World countries typically show a higher concentration, with a focus in achieving economies of scale both in production and marketing which enabled them to perform product and process innovation and to build strong brands and distribution networks. Because of historical reasons and a different strategy focused on well-defined geographical designations, Old World countries show a lower level of concentration (with some exceptions), with thousands of wineries and specialized actors composing the wine chain. This difference can be easily appreciated by observing the ranking of the main international listed groups operating in the production and marketing of wine with a turnover of at least €150 million in 2018. In such list, the podium is occupied by Australia, Chile and China (Mediobanca 2020),[5] while considering the ten biggest wineries in the world by sales in 2019, the whole podium was taken by US wineries (ISMEA 2020).[6]

I.2 Present and Future Challenges and Opportunities

After having depicted the overall scenario of the wine industry through the last 20 years, it is the moment to address the current challenges that are shaping the future of the sector. Such challenges can be grouped into three main drivers of change: *globalization*, *digitalization*, and *sustainability*. This three drivers, which are strictly interrelated, must be analyzed through the lenses of the aftermath of the global pandemic.

While the globalization phenomenon has reached maturity in the last two decades, it will continue to be a primary driver of growth in the wine market, and fully exploiting its potential is still a critical challenge. Understanding consumer tastes in foreign markets, efficiently communicating wine values to different cultures, building an efficient distribution chain, and defining a business model with an optimal balance between volume and value, and between product and consumer focus, are all ongoing challenges without a unique identifiable solution (ISMEA 2020). Additionally, in an interconnected and digital world, consumer trends and lifestyles are becoming global and consumer preferences may swift rapidly. Wine consumers

[5] Respectively Treasury Wine Estates (€1.7 billion turnover), Viña Concha y Toro (€770 million) and Yantai Changyu Pioneer Wine (€617 million). This analysis excluded the US-based listed company Constellation Brands due to their repositioning in the beer market.
[6] Respectively EJ Gallo (3% of total wine sales by volume), Constellation Brands (1.7%) and The Wine Group (1.5%).

nowadays are drinking with moderation and their purchasing behaviors are strongly influenced by ethical issues, therefore they're open to try new low-alcohol wines and new sustainable packaging formats; they love to drink outside home and to share the experience with friends both physically and digitally and they like fresh and easy-drinking wines such as prosecco and rosé; their drinking choices shift from wine to beer and to other low-alcoholic beverages (Wine Intelligence 2020). The digital environments in which such consumers are increasingly moving, such as social network and wine apps, create new possibilities for studying consumers at scale, offering wine marketers new tools for understanding their motivations and tastes. Furthermore, the increasing level of digitalization in the wine sector opens to the opportunity of setting up smart supply chains globally, adapted for satisfying demand in a timely manner.

In the globalized wine market, the covid-19 pandemic started in 2020 had a severe impact, especially due to the sudden and protracted closure of the on-trade channels in most countries. By looking at the numbers, the decrease of global consumption in 2020 has been estimated around 3% compared to 2019, analogous to that of the 2008-2009 global financial crisis, while the international trade of wine slightly contracted in terms of volume (-1.7%), but experienced a much larger fall in terms of value (-6.7%) (OIV 2021a). However, these numbers cannot be fully attributed to the pandemic due to the combined effect of trade wars and geopolitical tensions which were already endangering the wine market. Furthermore, they are the results of an asymmetrical response by countries, with many of them capable of performing well in the overall year, as in the case of Italy where wine consumption reached the highest level of the decade at 24.5 mhl and which was successful in containing the negative effects on wine trade (OIV 2021a)[7]. Consequently, while it will take more time to quantify the enduring effects of the pandemic, it is useful to analyze in qualitative terms the changes that the pandemic fostered in the wine industry by identifying on the supply side the winning strategies of the companies which performed strongly during the year, and on the demand side the consumer trends which emerged (or intensified) in the same period.

The winning formula for wine producers during the pandemic was composed by e-commerce, home delivery, virtual tastings and oenotourism (during summertime, when most restrictions were lifted) (ISMEA 2020). Except for oenotourism, it appears evident that the common denominator of these activities is that they have all been performed mainly through digital means. Producers who had diversified their sales channels and had a strong online presence

[7] Italy's wine exports decreased by 2.4% in volume and by 2% in value, performing better compared to France (-4.9% and -10.8% respectively).

were more solid in facing the pandemic, especially during the months where a lockdown regime was imposed thus leading consumers to make their purchases through the internet. A significant share of these consumers bought wine online for the first time, and will continue to do so in the future. The value of beverage alcohol e-commerce grew by 42% in 2020, reaching US$24 billion (IWSR 2020).[8] This growth has been pushed by the significant investments that traditional large-scale retailers put in place to enhance their online offering towards an omnichannel strategy (IWSR 2020). Simultaneously, online pure players experienced triple-digit growth rates and significantly expanded their business through new offers. Wine delivery apps rapidly scaled their business, enabling wine producers to mitigate the losses from the disruption of the on-trade channel. Furthermore, the producers who were more ready to communicate with customers through the internet and with a solid customer base were also able to organize direct-selling from scratch by leveraging social media and messaging apps, or, if they were already doing so, they simply had to expand their network. Virtual wine tastings, with the participants connected via videoconference apps comfortably from home, were a successful vehicle for to reach existing customers and to attract new ones, and signaled a positive predisposition of wine consumers to be engaged through digital experiences. While the growth of e-commerce and digital players will be deepened in Chapter 2, *Section 2.1*, the significant benefits that the more digital-ready companies experienced during the pandemic, the significant investments that the wine and retail industry put in place, and the increased familiarity of consumers towards digital solutions in the wine world, they all seem to indicate that the wine industry is at a turning point and that the global pandemic has acted as a digital accelerator, similarly to other industries worldwide (Kotler, Kartajaya, Setiawan 2021).

From consumers' perspective, the pandemic strengthened their concern towards wellness and sustainability, which translate into a heightened emphasis on ingredients, authenticity, self-care, and environmental impact when making a purchasing decision (IWSR 2021c). While this trends could be clearly observed even in the previous years, the pandemic has pushed health-conscious and sustainable lifestyles to become aspirational. This could put greater pressure to wine producers in satisfying more rapidly consumer demand for *clean labels* (Asioli et al. 2017) which display information about nutritional aspects (such as calories intake per glass), ingredients list (exact sulfites content, free-from additives claims) and environmental impact (for example a carbon neutrality claim). Additionally, a number of

[8] This value was calculated across 10 core markets: Australia, Brazil, China, France, Germany, Italy, Japan, Spain, UK, and the US).

certifications may concur in signaling the *naturalness* of a wine, meaning that it is obtained from sustainable agricultural practices and is produced using minimum or no synthetic chemicals and oenological additives. Among them, organic wine is clearly the most well-known category and its certification is regulated by governments as in the case of the EU and the US. According to the IWSR (2021c), certified-organic wine volume consumption is increasing at a 9% on average a year (from 2014 to 2019, faster than traditional wine) and is expected to represent 4% of global wine volume consumption by 2024 (starting from the current 2.75%). This growth is particularly interesting considering that organic wine is typically traded at a higher than average price. However, despite representing one of the wine categories with the biggest potential for growth, industry insiders reveal a general confusion of consumers towards what organic wine means (Wine Intelligence 2020), and the recent surge in *natural wine* claims, including the privately-owned biodynamic certifications, just spurred more confusion. Therefore, the wine market is called to collaborate in conjunction with governing bodies in order to regulate the phenomenon and to effectively and transparently communicate with consumers about their progress in offering a healthier and more sustainable product. This effort includes promoting a healthy lifestyle which includes moderate and regular wine consumption as part of a balanced diet, as there is scientific evidence of the health benefits of such *Mediterranean way of drinking* in increasing longevity and reducing the risk of cardiovascular diseases (Giacosa et al. 2016). Additionally, from a business point of view, it is worth investigating how much consumers are willing to pay for such health- and sustainability-related product cues, especially with respect to traditional wine cues (see Chapter 2, *Section 2.2*). Finally, the demand for sustainability may promote the adoption of new packaging formats such as cans and Tetra Pak, that will inevitably affect the way in which wine is distributed and force the wine supply chain to adapt accordingly (Lockshin & Corsi 2020).

Sustainability and social values are deeply rooted in the younger generations, which consider the brands they are purchasing accountable for their broader impact on society. Furthermore, they typically research information online and they are inclined to be engaged through digital experiences, which they consider a natural extension of their world (Kotler, Kartajaya & Setiawan 2021). Since generational change is another critical challenge for the wine industry, setting up strategies and digital solutions to effectively engage with this generations is another critical task.

However, sustainability is not only a marketing value, but is an ongoing concern of the wine producing industry. Climate change is having a huge impact on viticulture (and consequently on wine production) through an increase in temperature and radiation, and by exposing large regions to water deficits and more frequent extreme weather events (van Leeuwen & Darriet 2016). The complex interaction of these elements affects grape yield and quality. While mild water deficits and radiation levels can lead to an improved quality of grapes in many regions, grape yield has been declining, therefore decreasing wine production; however, this delicate balance may last only until weather does not become too extreme in many traditional grape growing areas, especially in the Mediterranean, which may become unsuitable for the production of high quality wines. Furthermore, early-ripening is forcing grape-growers to change their established viticultural practices, and to consider shifting their wine production to non-local varieties. This poses a serious challenge for to not dissipate the typicality traits of wine production that compose the heritage of the Old World of wine countries and which largely represents their competitive advantage. Therefore, new tools and practices must be developed to help vine growers in their decision-making based on measurable parameters in the vineyard and on consolidated scientific knowledge.

As it will be shown throughout this thesis, digital solutions are a primary tool for addressing many of the aforementioned challenges, from understanding and engaging with the ever-changing wine consumer to optimize distribution and sales channels, and to support decisions in the vineyard and the winery, which are increasingly put under pressure both from consumers and the climate change. The digitalization process will eventually help to reach the objective of a sustainable development of the wine sector, which encompasses simultaneously the environmental, economic, and social dimensions (Mediobanca 2020).

To conclude this section, a look at future trends in the digital wine customer experience will be presented. As consumers in the market (including wine) increasingly polarize towards no-frills, good value for money offerings on one side and luxury premium offerings on the other, the new field of competition for the companies positioned on the higher-end of the market is represented by customer experience (Kotler, Kartajaya & Setiawan 2021).

Wine apps are among the first examples of blurring the lines between the physical and digital retail environment (Rooderkerk & Kök 2019). After taking a photo of a wine label or a wine list with their smartphone, an *image recognition* software is capable of matching such photo to the right wine bottle within a cloud-based database of wines. These *shopping-assistant* apps help users in choosing wine by accessing additional information (such as ratings and reviews from their communities, experts' opinions, tasting notes, food pairings) at the moment of purchasing. The most advanced apps also incorporate a recommendation system that displays a percentage score of how much a given wine matches the taste of the user and is capable of recommending other similar wines of potential interest. Following the suggestions coming from the field of Human-Computer Interaction (HCI) research, an effective way to enhance the user experience in these apps is to deliver synthetic and pleasant visual representations, such as (interactive) summaries showing the overall flavor profile of a wine and its main tasting notes.[9] Another common approach is to enable some degree of user control over the application functioning through a dialogue with the system, by requesting user's feedback and enabling him to adjust some parameters to his preference (Calero Valdez, Ziefle & Verbert 2016).

Recently, there has been a surge in the use of *augmented reality (AR)* experiences in marketing. Augmented reality is a form of human-computer interaction which overlays computer-generated information into the real world, thus "creating an illusion where both virtual and real objects coexist in the same place" (Crofton et al. 2019, p. 3). By exploiting this technology *living labels*[10] have been designed capable of telling stories, of displaying vineyards landscapes and winemaking phases, and of creating artistic expressions in real-time.

Furthermore, as consumers often perceive online shopping as a cold and impersonal experience, *virtual agents* are increasingly employed (both in a human form and not) in order to interact with customers in a warmer way and to create a social link with them (Castellano et al. 2017). Given the massive amount of data about users and wines that can be gathered into wine applications, it is only a matter of time before omniscient *virtual sommeliers*[11] will dialogue with the customers, suggesting the perfect wine, appropriate food pairings, and dynamically adapting their suggestions to customers' special requests or mood, both inside mobile apps and in the real world through AR experiences.

Another interesting field of innovation comes from neuroscience and its applications in the wine field. Flavor is created in the brain where multiple sensory stimuli are processed into a unique perception of flavor (Sheperd 2015). Such multisensory integration is not just simply the sum of the sensory stimuli, as they interact between each other through *cross-modal interactions* (Smith 2019). A classic example of these interactions is the smell-taste interaction where the vanilla aroma increases the perception of sweetness. Additionally, among the many brain systems activated for processing the flavor perception there are the central brain systems which are responsible for memory and emotion. Therefore, there has

[9] See https://www.vivino.com/IT/en/simone-capecci-picus-rosso-piceno-superiore/w/1128434
[10] See https://tactic.studio/living-wine-labels
[11] See https://dot-farm.net/tvs/en/

been an increasing interest in finding ways to enhance the wine tasting experience by leveraging cross-modal effects and emotions, leading to the emergence of a *multisensory experiential marketing* field in the wine industry (Spence 2019).

A number of peer-reviewed empirical studies have now evidenced the effects of music during wine tasting. The sonic properties of music can affect the perceived sweetness or acidity of wine, enhance fruity aromas, modulate mouthfeel, and ultimately increase the enjoyment of wine (Spence 2019). Therefore, carefully selecting wine-music pairings and delivering them to customers via mobile apps or social media can be an effective marketing tool.[12] Furthermore, soundscapes could be composed specifically for to enhance the flavor profile of a given wine. As vision is the dominant sense in humans, another relevant field of studies has focused on the effects of color, shapes and animations on taste and flavor perception (Huisman, Bruijnes & Heylen 2016).

The joint results of the research on cross-modal effects in wine tasting are now enabling to design new digital experiences through AR and VR technologies. *Virtual reality (VR)* is described as "an immersive human-computer interaction in which an individual can explore and interact with a three-dimensional computer-generated environment" (Crofton et al. 2019, p.2). This technology enables to perform virtual wine tours or to create appropriate contexts for virtual wine tastings, which may match with or enhance the emotional response towards a wine. However, virtual reality has also been employed to design more abstract digital environments which leverage colors, sounds, animations to augment the sensory perception of specific aromas and flavors in wine and beer.[13]

One may ask how far these digital experiences could go in the future. Will we be able to taste a digital wine, maybe as a free sample before purchasing it? Or will we be able to share a digital bottle with friends in a VR environment? As natural and digital world increasingly merge into one and as people is getting accustomed to buy and experience wine digitally, it is legit to ask whether flavor could be digitized in the same way that images and music have been in order to transmit flavor experiences via the Internet (Spence et al. 2017). However, as it will be thoroughly presented in Chapter 2, *Section 2.2.1*, flavor is in a strict sense the overall perception arising from olfaction, taste and chemesthetic sensations (such as the astringency of tannins). Therefore, in order deliver an authentic flavor experience it would be necessary to perform a digitally-controlled electrical and/or thermal stimulation of olfactory

[12] See https://www.krug.com/it/playlist/krug-echoes
[13] See https://www.rga.com/work/case-studies/a-sip-for-the-senses

receptors, taste buds and the trigeminal nerve. Furthermore, it has to be considered that around 75-90% of what we perceive as flavor in the mouth is actually delivered by the sense of smell via retronasal olfaction (Spence et al. 2017).[14] So, the considerable difficulty in effectively reaching and stimulating the odor receptors situated high-up inside the nose and their complex, little understood functioning pose serious limitations to the possibility of digitizing flavor. Therefore, a more straightforward approach could be to digitally-control the analogue delivery of scents, thus chemically stimulating the orthonasal olfaction (breathing in) and recreating only the aroma component of flavor. However, the complex chemistry of wine makes it difficult to obtain its aroma as a combination of basic olfactants,[15] therefore impeding the possibility of faithfully recreating a wine's aroma profile unless a real sample of the wine is vaporized. Therefore, as suggested by Spence et al. (2017), it can be concluded that the most promising opportunities in this field in the foreseeable future rely on the digital enhancement of flavor experiences by designing digital solutions that exploit the increasing knowledge about cross-modal effects.

I.3 What Artificial Intelligence Is

Artificial Intelligence (AI) has been defined in many ways. John McCarty, among the founders of AI as a discipline, gave the following definition: "[artificial intelligence] is the science and engineering of making intelligent machines, especially intelligent computer programs" (McCarthy 2007, p. 2). This definition of AI cannot leave aside what is meant by intelligence, which McCarthy described as "…the computational part of the ability to achieve goals in the world" (McCarthy 2007, p. 2).

To overcome the difficult task of defining intelligence, especially human intelligence, Somalvico (1992) proposed a more straightforward definition of AI as the discipline which studies the theoretical foundations, methodologies and techniques that allow the design of digital systems (hardware) and programs (software) capable of providing the electronic computer with performances that, to a common observer, would seem to belong exclusively to human intelligence.

[14] As the taste buds on the tongue are only responsible for the sensation of sweet, bitter, salty, sour and umami.
[15] Differently from colors where a small number of primaries can create a wide range of colors, there is less understanding on how to combine basic olfactants to obtain complex smells (Spence et al. 2017).

Some part of the research community believes that, eventually, machines will be able to display a human-like intelligence, what is commonly referred to as *strong AI*. Whether this goal will be reached or not, only the future will tell; however, today *weak* (or *narrow*) *AI* "…is already enough to transform the nature of firms and how they operate" (Iansiti & Lakhani 2020, p. 4). Today, the most prolific subfields of AI are represented by Machine Learning and Deep Learning (*Figure 7*).

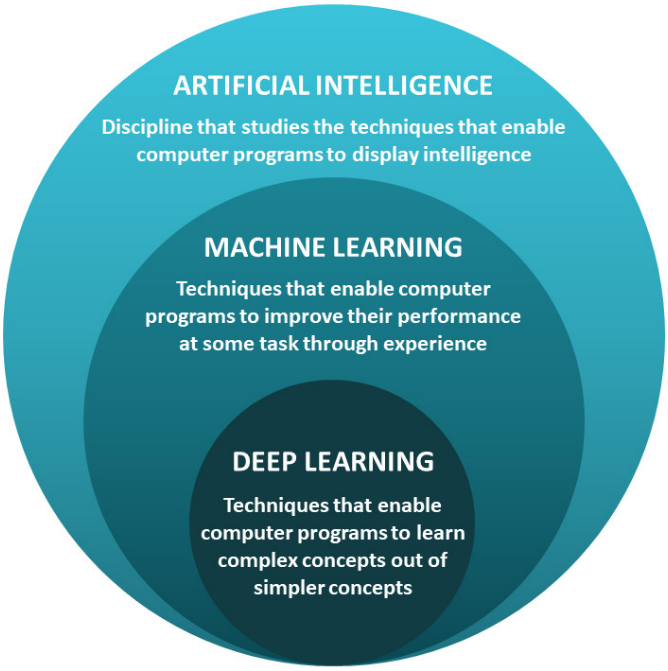

Figure 7. Venn diagram showing the relationship between Artificial Intelligence, Machine Learning, and Deep Learning. Source: own elaboration based on Somalvico (1992), Mitchell (1997) and Goodfellow, Bengio & Courville (2016).

Essentially, Machine Learning (ML) techniques enable computer programs to learn from experience and thus to improve their performance at some well-specified task (Mitchell 1997). Such experience is provided to them in the form of datasets, which consist of a collection of examples, which in turn are a set of features that typically represent a real-world object or phenomenon. Then, learning involves "…searching through a space of possible hypotheses to find the hypothesis that best fits the available training examples and other prior constraints or knowledge" (Mitchell 1997, p. 18). Problems easy to formalize and computationally difficult for humans (such as playing chess) are the typical target of machine learning algorithms.

Conversely, problems difficult to formalize but straightforward for humans, such as processing images or spoken words, have led to the emergence of Deep Learning (DL). Indeed, it is often difficult to design an accurate *representation* of a real-world phenomenon through a set of features, as several factors of variation usually influence the data; more abstract features are needed to be specified in order to enable the computer program to make sense of the variability in the data (Goodfellow, Bengio, & Courville 2016). The solution achieved by deep learning is to enable computers to build complex concepts by relating simpler concepts, thus representing the world in terms of a hierarchy of concepts (Goodfellow, Bengio, & Courville 2016).

The birth of this subfield of ML has been favored by the increasing availability of large and varied datasets, often defined as *big data*, and by the advancements in the computational power of computers. Deep leaning has enabled to perform tasks such as image and speech recognition, to develop self-driving cars, and to improve the algorithms' performance at many traditional machine learning tasks, in particular pattern recognition.

Two of the main outputs of machine learning models are accurate *predictions* based on data —assigning the inputs to the correct category (classification) or predicting a numerical value (regression)— and the discovery of relevant *patterns* in the data. The field of machine learning often overlaps with that of *data mining*, which is defined as "…the process of exploring data to generate new knowledge and insights that are useful for a particular domain" (Valente et al. 2018, p. 1). In this context, ML algorithms are useful tools to generate hypothesis and to improve decision-making based on data.

AI-based technologies have several applications in the wine industry along the supply chain. Such applications are enabled by the increasing availability of data that is generated at all stages thanks to the diffusion of sensor technology and the creation of digital touchpoints with customers.

At the vineyard, the combination of wireless sensors and satellite or drone images enable to monitor key parameters during grape growing regarding weather and soil conditions and to control for vine diseases (OIV 2021b). With this data, predictive models can be developed to assist grape-growers in performing vineyard treatments and in determining the optimal harvest time. Additionally, AI-powered robots are increasingly deployed in the vineyard to automate vineyard management practices and to perform monitoring.

Similarly at the winery the increasing availability of sensor-based solutions to chemically analyze wine at each stage of production[16] in a timely and cost-effective manner will allow for the control of key winemaking parameters and will enable the development of predictive

[16] See https://www.fossanalytics.com/en/products/winescan

models that will help achieve the target wine quality. Furthermore, AI enables to perform the monitoring of cellar and tank conditions by delivering real-time alerts and warnings (OIV 2021b).

AI and robots will also have a huge impact on the distribution phase by automating warehouses, optimizing logistic flows, and supporting inventory management. But it is at the very end of the supply chain where AI can shine. By engaging with consumers in digital environments such as e-commerce, wine apps and social media a vast amount of consumer data is generated. Such data can be combined with other internal or external sources to build a data ecosystem, which provides the raw material with which to perform data-driven marketing (Kotler, Kartajaya & Setiawan 2021). The aggregate data, in addition to represent a formidable source to analyze market trends, can be scaled down to individual consumers, therefore enabling to offer them personalized recommendations, pricing and marketing messages.

Shopping assistants apps today are relying heavily on machine (and deep) learning algorithms to deliver personalized wine recommendations, and they do so by leveraging different sources such as historical transactional data, ratings, reviews, and users' taste profiles. Given the growing convergence between physical and digital retail spaces, it's not hard to imagine that personalized wine recommendations could be provided in the future by a robot sommelier in the wine aisle or at the restaurant (Marr 2019).

Other AI applications can be leveraged in the at-home context and concern smart gadgets[17] such as smart decanters or smart wine cellars[18] capable of monitoring temperature, keeping track of wine aging, and suggesting food pairings. Additionally, the adoption of self-driving cars in the near future could lead to a significant increase in alcohol and wine consumption, as cars will turn into a space destined to entertainment (Marr 2019).

But AI potential does not limit to its applications within the wine supply chain. Computer science has become a fundamental component of conducting research, especially when dealing with complex biological systems (Emmott et al. 2006). A number of algorithms are now regularly employed in the data mining phase, with the hypotheses generated to be tested with consolidated analytical techniques.

[17] See https://www.winemag.com/wine-tech/
[18] See https://www.prestigeonline.com/th/pursuits/tech/winewall-ai-wine-celllar/

The chemistry of wine is composed of thousands of compounds that directly or indirectly influence its flavor, and the mechanisms through which these compounds interact to express the sensory attributes of a wine are complex and not yet fully understood. The most advanced techniques for the chemical analysis of wine are now capable of detecting thousands of volatile and non-volatile compounds within a wine sample, therefore offering a comprehensive picture of a wine (Pérez-Jiménez et al. 2021). Such datasets have been put in relation with quantitative data about sensory attributes of wines collected from sensory analysis performed by trained panels of experts or consumers in order to draw novel insights about the molecular basis of flavor.

By relating the datasets generated by the chemical and sensory analysis of wine with those containing consumer preferences (such as ratings collected from wine apps), deep learning algorithms could find interesting insights on which chemical compositions and sensory attributes drive consumer preference. Ideally, such datasets could also be used to train a computer program to develop its own ability to taste (and rate) wine. To conclude, relating these datasets could establish a link between the two extremes of the wine supply chain, opening up interesting scenarios to explore from a research and business point of view.

I.4 Digital Business Modelling

As the digital and physical world merge into one, especially thanks to the Internet of Things (IoT), products, services and business processes are increasingly digitized. This is not merely the result of the widespread adoption of digital technologies in business (and everyday life), but it also represents an objective to pursue for companies who want to exploit the opportunity offered by AI and analytics. By systematically *datafying* all activities and transactions which occur in the ongoing of any business, the pool of data generated can be leveraged to create additional value by optimizing existing processes or by delivering new value through other, new processes (Zeng 2018; Bagnoli et al. 2018). Creating a consistent *data pipeline* enables to extract highly valuable insights and to perform experimentation and continuous learning, while at the same time most digitized processes can be automated by leveraging machine learning algorithms. Ultimately, in a digitized business "…software instructions and algorithms make up the critical path in the way the firm delivers value" (Iansiti & Lakhani 2020, p. 4).

Iansiti & Lakhani (2020) claim that a new breed of firms, architected to exploit the full potential of data, algorithms and networks, is transforming the global economy. The digital foundation of their operating model is profoundly changing the way in which they create and deliver value to customers and how they capture it from them, leading to the definition of new business models and redefining the rules of competition.

A digital operating model essentially enables companies to outstandingly answer to three key challenges of traditional operating models (Iansiti & Lakhani 2020):

- *Scale.* A digital operating model is virtually infinitely scalable, as the marginal cost of serving an additional user is negligible. The additional storage and computing capacity needed can be bought on demand from cloud-based service providers. Additionally, as value is delivered mainly through digitized processes, the major growth constraint shifts from the organization of human labor to software development, which can be continuously improved internally or by outsourcing to external partners.
- *Scope.* Digitized processes are intrinsically modular and multisided. This means that they ease the creation of business connections, both by plugging into external networks of partners to provide complementary value or by connecting with existing internal processes to deliver additional value, thus adding a multiplicative factor to the value delivered by the company.
- *Learning.* As data is accumulated during the ongoing of the operations into a data pipeline, AI algorithms continuously learn and improve their performance based on that data and generate precious insights to promote innovation and the development of new products and services. As learning enables to increase scale and scope, more data is generated, thus improving the performance of AI algorithms even more and fostering a self-reinforcing loop.

Basically, the operating model defines how a company will deliver the promised value to its customers. The business model instead is a wider concept which encompasses "the overall value architecture and related mechanisms… [built around a] value proposition to generate value for target customers, place such value on the market and retain part of it to ensure economic and financial viability" (Ghezzi & Cavallo 2020, p. 519).

Among the business models which have been most successful in exploiting the potential of digital operating models there are undoubtedly *digital platforms*, which represent today an important share of the economy and which have had a disruptive impact on many traditional sectors (Evans & Gawer 2016). Essentially, platforms create value in two ways (Evans & Gawer 2016):

- by facilitating transactions or exchanges through the efficient *matching* of individuals and/or organizations that would otherwise have difficulty finding each other;
- by setting up an *innovation ecosystem* capable of attracting external innovators, which can deliver complementary value.

The most distinctive feature of platforms is the presence of *network effects*: the more users join the platform, the more attractive the platform will be for new users; furthermore, the more users join one side, the more value is generated for the users on the other side. Additionally, digital platforms have successfully leveraged Internet-connected mobile devices such as smartphones to access billions of users and to gather massive quantities of data to be analyzed and monetized (Evans & Gawer 2016).

It can be argued then that in successful platforms value is actually co-created through a collective effort which includes external partners and customers (Senyo et al. 2019). This means that platform leaders need "…a vision that extends beyond one's own firm… [aiming to] build and sustain an ecosystem of partners, where the platform leader has to be the equivalent of a captain" (Evans & Gawer 2016, p. 6). The successful build up of the platform's ecosystem could eventually lead to the creation of a *digital business ecosystem* defined as "…a socio-technical environment of individuals, organisations and digital technologies with collaborative and competitive relationships to co-create value through shared digital platforms" (Senyo et al. 2019, p. 53).

It must be acknowledged though that defining a successful business model seldomly means reaching a sustainable and long-term success in the market. In an era characterized by volatility, uncertainty, complexity, and ambiguity (*VUCA*), businesses need to continuously evolve and adapt in order to "match the speed of customer shifts and outpace the competition at the same time… Agility is the new name of the game" (Kotler, Kartajaya & Setiawan 2021, p. 183). This is why besides innovating products, services and processes, a number of scholars and practitioners are stressing the importance for companies to engage in innovation processes that involve their entire business model (Ghezzi & Cavallo 2020).

Business Model Innovation (BMI) deals with designing “novel, non-trivial changes to the key elements of a firm’s business model and/or the architecture linking these elements” (Foss & Saebi 2018, p. 201, cited in Ghezzi & Cavallo 2020, p. 520).

A number of disciplines have become popular to guide the design and innovation of digital business models. *Design Thinking* is an approach to innovation and problem solving based on design principles. It is typically an iterative process which consists in an initial exploration of the problem space, an ideation stage where ideas are generated by leveraging creative and divergent thinking, and concludes with delivering a working solution to be tested; after evaluation, the cycle restarts based on the feedback from the experiment, possibly leading to reframe the initial problem (Micheli et al. 2019). The pillars of design thinking are represented by a *human-centered approach*, which entails empathizing with customers, a continuous *reframing* of the initial problem by embracing *ambiguity* and *diversity* (for example through interdisciplinary collaboration), and the reliance on *visualization techniques* to explore abstract concepts (Dell’Era et al. 2019). In a BMI process, the main aim of the Design Thinking approach is to gather qualitative insights from customers, understanding their needs and motivations and the problem to be solved for them.

While being developed for the design of new business models by innovative startups, the *Lean Startup* approach adapts perfectly also for the BMI process by validating the hypothesis underlying a business model through an iterative process based on quantitative data. The Lean Startup approach "consists of a scientific, hypothesis-driven approach to entrepreneurship, where entrepreneurs translate their vision – i.e. business idea – into falsifiable hypothesis which are embedded in a first version of a business model” (Ghezzi & Cavallo 2020, p. 521). Such hypotheses are tested by delivering *Minimum Viable Products (MVPs)* to evangelists[19], which will generate a feedback. MVPs are then adjusted according to those feedbacks and tests are iterated, thus creating a *build-measure-learn* loop. The loop goes on until all key hypotheses are validated, which means that a *product-market fit* has been achieved (Ghezzi & Cavallo 2020).

As digital operating models are based on data and software, software development represents an integral part of the BMI process. *Agile development* “refers to a number of agility-enabling practices for software development that value the centrality of individuals and interactions, the incremental delivery of working software, collaboration with customers and response to

[19] Minimum Viable Products (MVPs) are “the smallest set of activites needed to disprove a hypothesis” (Ghezzi & Cavallo 2020, p. 521), while evangelists “are expert prospects who can provide informed and useful feedback” (Ghezzi & Cavallo 2020, p. 521).

change" (Ghezzi & Cavallo 2020, p. 521). Agile methods therefore aim at building elegant and ready-to-use software solutions in a rapid, iterative way, while at the same time enabling evolution and adaptation to change through continuous improvement (Schneider 2017).

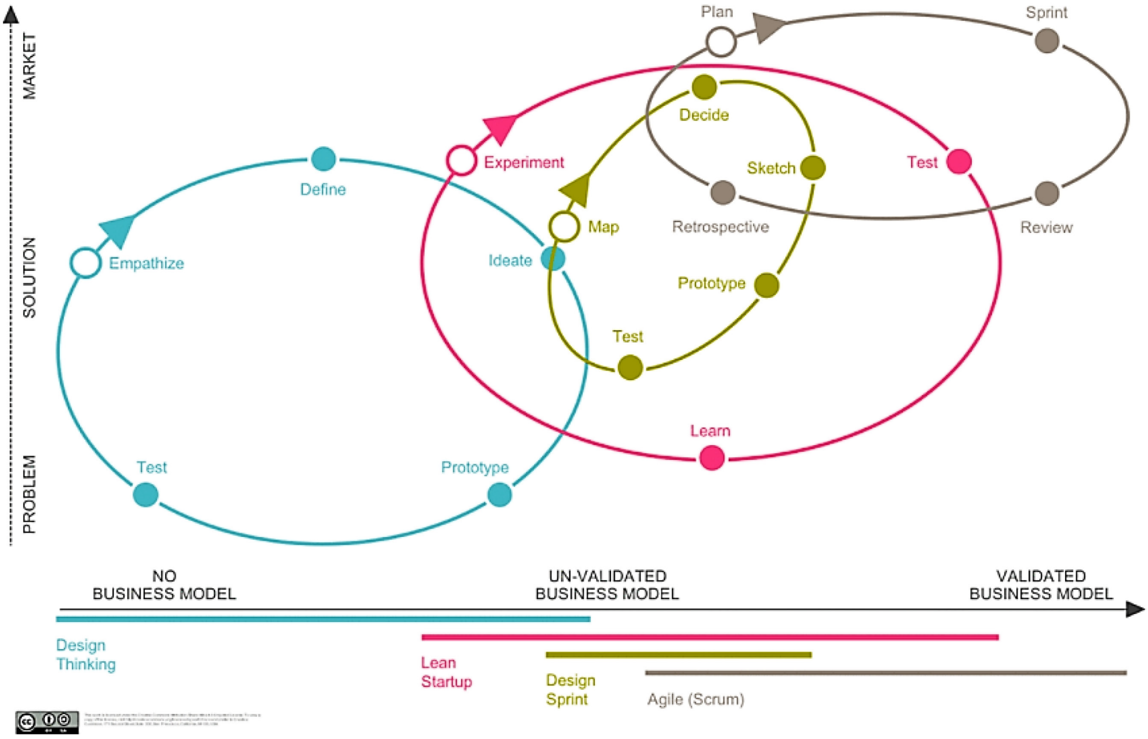

Figure 8. Design Thinking, Lean Startup, Design Sprint[20] and Agile Scrum in the innovation spectrum. Source: Problem2Value (2017), CC-BY-SA, https://creativecommons.org/licenses/by-sa/4.0/legalcode.

The iterative application of these practices enables to move within the innovation spectrum from a problem space to a market-ready solution, and from an unvalidated to a validated business model through qualitative and quantitative data (see *Figure 8*).

A useful tool to support these practices is represented by the Business Model Canvas,[21] developed by Alexander Osterwalder. It consists in a visual tool which represents a business model in a synthetic but comprehensive way. It is composed of nine building blocks which represent a firm's:

- *offering* (Value Propositions),
- *infrastructure* (Key Partners, Key Activities, Key Resources),
- *customers* (Customer Relationships, Customer Segments, Channels), and

[20] A *Design Sprint* consists in sharing insights, ideating, prototyping and testing a concept all in a 5-day sprint (Problem2Value 2017).
[21] See https://www.strategyzer.com/canvas/business-model-canvas

- *finances* (Cost Structure, Revenue Streams).

Such chart is useful not only for the initial design of a business model, but also during the BMI process in order to visualize and communicate the building blocks of the business model that will be affected (for a practical application, see Chapter 3, *Sections 3.2.1 and 3.2.2*).

There are quite a number of digital business models which have established in the wine industry. Since 2010, *shopping assistants* apps have proliferated, aiming at matching consumers and wines both in the physical retail space and in the online environment.

Oftentimes such apps have built their own *online marketplace*, by assembling their proprietary inventory or by relying on an ecosystem of external distributors and wineries. Furthermore, these apps have been constantly improving and introducing new features by leveraging the data generated by their users and the transaction history.

Additionally, digital wine businesses have engaged with the communities of wine lovers by creating wine-specific *social networks* where they could create their own profiles, rate and review wines, share pictures and comment others'.

The industry leader Vivino represents a virtuous example of an integrated platform which has successfully developed all the aforementioned business models into a unique platform (shopping assistant, marketplace and social network), thus delivering a comprehensive digital wine experience.

Other businesses have focused on *home delivery* by leveraging a more locally clustered network of distributors, guaranteeing the timely delivery of wine within a restricted time frame.

Digital business models have also addressed the needs of the most demanding wine consumers by offering services such as *cellar management*,[22] which enables them to organize and keep track of their wines and to record the tasting notes experienced during tasting. Apps have also been developed to facilitate *fine wine trading*.[23]

A kind of digital business model that could possibly become relevant for the wine market in the upcoming years is represented by *wine analytics*[24] delivered as a software-as-a-service (SaaS). The increasing amount of data generated within online marketplaces and the equally

[22] See https://apps.vinocell.com/
[23] See https://vindome.net/
[24] See https://www.enolytics.com/

vast amount of data that is being generated through IoT sensors across the wine supply chain will enable to generate highly valuable insights which can be delivered via software on the cloud on a subscription-based model.

CHAPTER 1: THE AI-POWERED COMPANY

1.1 The AI Factory

When business models are being continuously innovated, experimenting new ways to create and capture value from customers, the operating model becomes crucial in delivering the actual value promised to the customers. And when products, services and processes are digitized, the actual bottleneck of operating models shifts from human labour organization to the software architecture of the company.

And as "software, analytics, and AI are reshaping the operational backbone of the firm" (Iansiti & Lakhani 2020, p. 20), a new breed of firm emerges, characterized by digital scale, scope, and learning.

Zeng (2018) and Iansiti, Lakhani (2020) recognize that traditional IT systems are inadequate to provide a digital operating model and harness the full potential of analytics and AI. In particular, although improving the performance of many operating processes, traditional enterprise IT systems typically mirror the siloed and specialized architecture of the organizations in which they are deployed (Iansiti, Lakhani 2020).

Oftentimes this translates into *inconsistent data collection*, undermining the possibility of aggregating and integrating datasets and therefore limiting the potential for analytics and AI applications. Furthermore, the adoption of highly specialized and oftentimes incompatible systems may foster *architectural inertia*, thus preventing work to be organized in new ways.

The combined effect of these two aspects significantly constrains the scale, scope and learning potential of traditional firms.

According to Iansiti and Lakhani (2020), companies in the age of AI need to be rearchitected on a new integrated, highly modular digital foundation. The authors propose to industrialize data gathering, analytics, and decision making to reinvent the core of the modern firm into an "AI factory" (*Figure 9*).

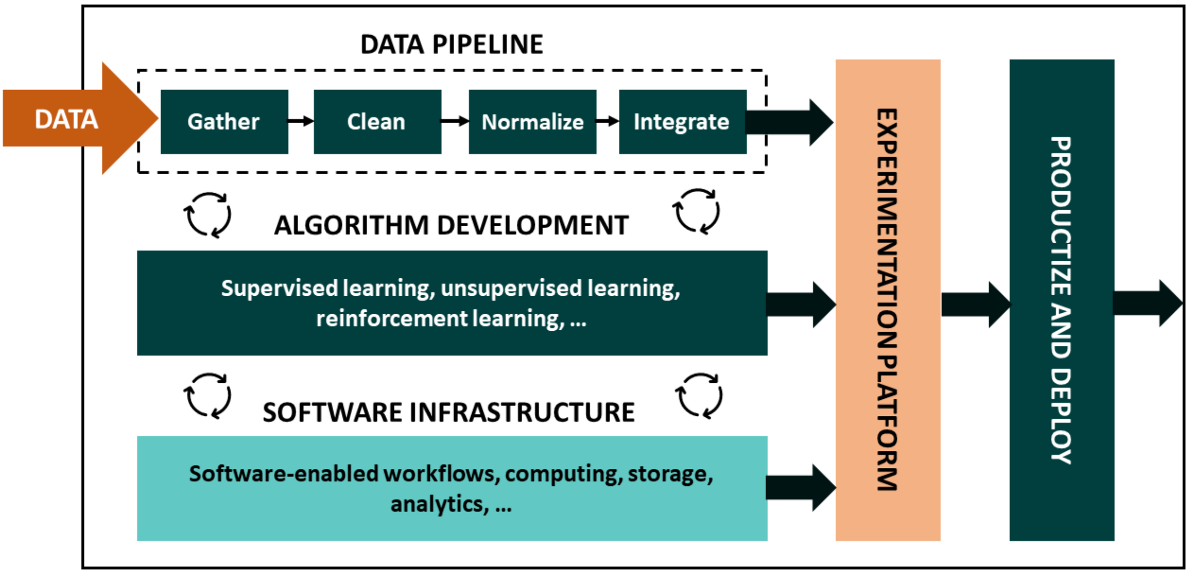

Figure 9. The AI Factory architecture. Source: redrawn from Iansiti & Lakhani (2020), p. 58.

Data is the fuel of the AI factory, and the *data pipeline* represents its first critical process as it gathers, cleans, integrates, and processes data in a consistent way. The data is fed into the pipeline through a process called *datafication,* which enables to extract data systematically from activities and transactions that are naturally ongoing in the business (Iansiti & Lakhani 2020). In fully digital operating models, companies can gather massive amounts of live data about every exchange and communication with customers during the process of delivering value to them (Zeng 2018). Such data can be integrated with other internal or external data sources which will then undergo the same processing steps, thus resulting into a single stream of consistent data.

The second fundamental element of the AI factory consists in *algorithms*, which represent its beating heart. In fact, in a digital operating model algorithms drive some of the most critical processes in the delivery of the promised value to customers. They typically do so by leveraging a consistent data pipeline to generate accurate predictions about future states or actions of the business (Iansiti & Lakhani 2020). Such predictions then translate into informed human decisions or into automated responses. Furthermore, algorithms typically experience significant feedback loops, thus improving their own accuracy over time. Given their operational relevance, it must be stressed how algorithms must be carefully designed to express the underlying business logic that the company is trying to optimize (Zeng 2018).

A third and innovative element of the AI factory is an *experimentation platform*. This mechanism very well represents the business model innovation focus underlying digital business models. In fact, new hypothesis generated through data and algorithms are continuously tested in order to ensure that the suggested decisions are having the intended effect on the business (Iansiti & Lakhani 2020). Such experimentation platforms leverage the digital nature of the operating model to perform A/B testing at scale, where a random sample of users is exposed to the target change (the treatment) and another random sample of users has the same customer experience as usual (the control). Then, only if the difference between the two outcomes is statistically significant the algorithmic output has a causal effect on the business outcome. In this way, not only BMI is promoted by generating new knowledge and new ways of incrementing the value delivered, but it becomes also possible to overcome the loss of causality which is traditionally addressed to AI applications (Iansiti & Lakhani 2020).

Finally, the *software infrastructure* should be designed in order to fully empower the AI factory functioning. Firstly, "softwaring" as many business activities as possible, by breaking down business decisions into simpler elements and replicating them through software, may eventually enable to automate entire processes and to perform data collection naturally as part of the business processes (Zeng 2018). Secondly, the flowing of data should be fostered by building a secure and accessible platform. APIs (application programming interfaces) are the mechanisms through which clean and consistent data is made available to different players within the platform while ensuring control on who can access and edit the data (Zeng 2018). This enables software developers to rapidly sample the data they need and build, deploy, and execute AI applications (Iansiti & Lakhani 2020). Clearly, the more data flows across the network, the stronger the feedback loops, the more experimentation can be done, and the more applications can be implemented, ultimately resulting in more value to be generated.

Thanks to such communication standards, the resulting software infrastructure and its underlying code is inherently modular. Digital organizations should mirror this modular design by letting agile teams to work independently and simultaneously on clearly defined business- and customer-focused goals (Iansiti, Lakhani 2020), thus reaching innovation at the module level.

Nowadays, software infrastructures are increasingly on the cloud and scalable on demand. The definition of a new operating architecture through software is fostering a new organization of work and will ultimately lead to new digital operating models, capable of exploiting the full potential of AI and analytics.

1.2 Platform Business Models

In the era of digitalization of products and services, digital operating models are business models that offer new ways for creating and capturing customer value. The rules of competition are changing as industries' boundaries are getting more and more blurred, and these aspects are constantly reshaping many sectors of the economy. In this new competitive arena, emerging digital companies come very often into collision with the traditional companies. This occur whenever these digital companies are able to create new effective operating models that help them in meeting their customer needs in new ways, through the digitization of the most critical tasks in the delivery of value. (Iansiti & Lakhani 2020).

Even if it is clear that every business would benefit from having an accessible and consistent data pipeline, it is also true that some business models in particular have a natural advantage in exploiting data analysis and AI, and are more predisposed to create strong feedback loops, which are at the base of the competitive advantage of the next generation of companies. This occur because these business models themselves can enhance the collection of an enormous amount of data, allowing those interactions that will then be digitized by the operative model and processed by the AI analysis to gather valuable information and to promote continuous learning and improvements. Conversely, if a company is missing a profitable and effective business model, this will impede any chance for generating a valuable data pipeline and for getting a real and sustainable competitive advantage, exposing it to present and future competition.

For Napier et al. (2020) modern business models (MBMs) are a combination of three elements:

1. A *software-as-a-service* (SaaS) product that generates critical value for its users and promotes repeated, daily interactions;
2. A *multisided platform* structure, where an ecosystem of third-party sellers and partners can match with customers, generating network effects;
3. A *data lake* powered by AI that enables personalization of the offer, generates valuable insights, and drives matching between sellers and customers.

According to Zeng (2018), smart businesses rise when an ecosystem of different players that pursue the same goal is coordinated through an online network, exploiting data to identify and satisfy customers' needs in a quicker and more efficient manner compared to traditional companies.

Both these models have a common foundation in network effects, learning from data and an ecosystemic approach. All these elements will be deepened in the next paragraphs.

1.2.1 Multisided platforms

Multisided platforms coordinate the needs of different groups of customers, who need each other to engage in value-increasing exchanges, and internalize the outcoming network externalities (Evans 2003a) (Evans 2016).

Platforms move from traditional pipeline businesses, which operate through the classic value-chain model, and put in place three key shifts (Van Alstyne et al. 2016):

1. *From resource control to resource orchestration*. In platforms the crucial assets are external to the company and consist in its communities and the networks they create through their interactions and the resources they contribute.
2. *From internal optimization to external interaction*. In platforms, "value" is created by facilitating interactions between its participants, typically external producers and consumers. The focus of strategy becomes to persuade participants and to govern (or at best influence) the overall ecosystem.
3. *From a focus on customer value to a focus on an ecosystem value*. Platforms seek to maximize the total value of an expanding ecosystem in a circular, iterative, feedback-driven process.

Nowadays, some of the most successful business models are indeed based on multisided platforms. Platforms are in fact occupying stable positions in the top 10 companies in the world by market capitalization – as of June 2020 (PwC), from the second position downwards we find in order Apple, Microsoft, Amazon, Alphabet, Facebook, Tencent, Alibaba, Visa. Moreover, in recent years platforms started to permeate more and more our everyday life, mostly thanks to the worldwide diffusion of mobile devices and mobile broadband networks, which enabled businesses to reach billions of people, and, together with the development of the Internet of Things (IoT), lead to a deeper integration of online and physical worlds (Evans 2016). But this is only the last step of a process that started in the mid-1990s with the rise of the web-economy and the dawn of global online platforms.

Even though platforms were already popular since the '80s as a tool for developing software and products through modularity and incremental innovation (Gawer & Cusumano 2013) and platforms as a business model have ancient economic connection with matchmaking (Evans 2016), in the early 2000s Rochet and Tirole began to investigate the rise of multisided platforms as new, increasingly important business entities. What they observed was that many markets with network externalities were two (or multi) sided markets, where value was created by the interaction between two (or more) distinct groups of customers through a platform. Since firms based on a platform were operating in two-sided markets, defining the business model was a crucial aspect for "getting both sides on board". Therefore, instead of simply setting a price level for their service, the fundamental aspect was to set up a price structure which was able to attract users on both sides; this would lead to often treat one side as a profit center and the other as a loss leader, or, at best, financially neutral (Rochet & Tirole 2003). This aspect has important implications in terms of competition. For example, when users on one side of the market tend to connect to several platforms for benefiting of the same service (they multihome), price competition will increase on the other side.

Evans (2003a) highligts the necessary conditions for the rise of a platform business:

1. There are two or more distinct groups of customers.
2. There are externalities associated with customers A and B becoming connected or coordinated in some fashion.
3. An intermediary is necessary to internalize the externalities created by one group for the other group.

Indirect network effects are the main reason behind the birth of platforms: one group of customers attributes value to the presence of another group, and through the platform they benefit from an increased probability of finding a suitable match; conversely, these network externalities represent an opportunity for entrepreneurs, who must find pricing, product and investment strategies to coordinate the interdependent demand and to exploit those profit opportunities (Evans 2003a). What makes a platform a critical aspect for its members are information and transaction costs, as well as free-riding, which will prevent customers to internalize the externalities by themselves (Evans 2003a).

Evans (2003a) identifies three traditional types of multi-sided platforms:

1. *Market-makers* enable members of distinct groups to transact with each other.
2. *Audience-makers* match advertisers to audiences.
3. *Demand-coordinators* make goods and services that generate indirect network effects across two or more groups.

The traditional competition dynamics are deeply reshaped by the presence of indirect network effects. A crucial aspect for firms is to get to a critical mass, what in literature is defined as the “chicken and egg” issue: without “both sides on board”, little or no value would be delivered, so the problem is to attract at least a critical mass of customers on one side to attract the other one thanks to indirect network effects. This aspect is the reason for the common practice of giving away for free (or even to subsidize) the access to the platform to one side (Evans 2003b). Evans (2003a) (2003b) observed that despite the wide literature on the first-mover advantage and the established conception of the importance of growing market share rapidly, many successful multisided platforms evolved slowly, carefully finding their optimal price structure, investing in operating infrastructure and technology, and growing customers on all sides. Furthermore, many early-entrants lost their leadership position to late entrants, and some even exited the market. This aspect was further analyzed by Zhu and Iansiti (2012), who developed a dynamic model of competition between platforms and empirically observed how strength of network effects, quality, and customer expectations are key elements in determining the success of an entrant directly competing with an incumbent player for the same market. They came to the conclusion that even in markets with significant network effects, when indirect network effects and consumers’ discount factor of future applications are lower than a certain threshold, an entrant player with higher quality can gain market share over time, and eventually overtake the incumbent player (Zhu & Iansiti 2012). In these quality-driven markets, even a large installed-base is not a long-lasting sustainable advantage. Furthermore, competition is increasingly based on incremental innovation, that companies pursue in order to subtract customers from competitors (Evans 2016).

The ongoing shift from a PC-browser to mobile-app centric model offers new opportunities and threats. The capital cost of starting a platform has declined thanks to technological change, and this, together with the fact that consumers can switch between different platforms at zero cost, exacerbated competition. On the other side, given the time spent by consumers on their smartphones, the new critical asset for which platforms are competing is consumers’ attention: by offering engaging and captivating content, platforms can gain consumers’ attention and generate profit by the selling it to other companies (i.e. other sellers or advertisers) (Evans 2016). This is increasingly crucial today when mobile apps are becoming part of the shopping experience.

Evans and Gawer (2016) conducted a global survey on platforms business models, identifying the presence of a global trend. They identified 176 platform companies all over the world with at least 1 billion of market cap or valuation, for a total market value of more than $4.3 trillion.

Asia holds the highest number of platforms (82) North America accounts for 72% of the value (with 64 platforms). Europe has a marginal share of total value with only 27 platforms who passed the $ 1 billion cut (4% of the total). With regards to the employment aspect, publicly traded platforms employ more than 1.3 million direct employees, but since these business models rely on an ecosystem of third-parties, the indirect effect on employment is also extremely significant.

After a comprehensive literature review they also defined four fundamental types of platform business models (Evans & Gawer 2016):

1. *Transaction platforms*. A transaction platform is a technology, product or service that acts as an intermediary that facilitates transactions between different users, buyers, or suppliers.
2. *Innovation platforms*. An innovation platform is a technology, product or service that serves as a foundation on top of which other firms develop complementary technologies, products or services.
3. *Integrated platforms*. An integrated platform is a technology, product or service that is both a transaction platform and an innovation platform. An example of this category is Apple, which has both matching platforms like the App Store and a large third-party developer ecosystem that supports content creation on the platform.
4. *Investment platforms*. Investment platforms are firms that have developed a platform portfolio strategy and act as a holding company, active platform investor or both.

The survey highlighted several ongoing trends (Evans & Gawer 2016):

- Integrated platforms, while small in number, are becoming predominant. In addition, both transaction and innovation platforms are evolving in the direction of integration. The ability to combine efficient matching with an ecosystem of developers is becoming the competitive advantage that companies try to pursue.
- There is a significant inhomogeneity between regions. The main regional factor behind the rise of platforms is the access to a large demand, as experienced by the US, India and China. The latter, also benefited from local regulation who limited foreign competition. At the same time, in smaller countries where a large demand was missing, it was also possible to launch some successful platforms thanks to high technological capabilities and to a consolidated start-up culture.

- Competition between platforms is increasing, and so are incentives to consolidate different platforms through mergers and acquisitions. Moreover, traditional firms are getting interested in platform business models and are starting to develop strategies to adapt to the mutated competitive environment, growing platforms organically or via acquisitions.

While platforms are becoming a global trend, expanding the playfield of competition and threatening traditional businesses, managers must search for new strategies to innovate their products and services and to internalize some of the benefits associated with this business model. Hagiu and Altman (2017) identified four possible strategies to transforms a company's products and services into a multisided platform:

1. *Opening the door to third parties*. It means connecting your customer base to third party sellers, who can offer additional or complementary products. In this strategy, having a large, desirable customer base, which generates frequent interactions, and whose heterogeneous needs could be served by third parties without the risk of cannibalizing your offer, is the precondition for success.
2. *Connecting customers*. It consists in selling your product or service to various customer segments that interact or transact with each other outside your offering. Therefore, by expanding your offer it is possible to capture at least part of those interactions.
3. *Connecting products to connect customers*. In this strategy you are selling two products or services, each one to a different customer base, but the two different customer bases interact outside your offering. You can modify your offerings in a way that at least part of that transactions occurs through one or both of your products or services.
4. *Supplying to a multisided platform*. In this strategy you create an offering for your customers' customers that increase the value of the product or service they buy from your customers.

Between all the challenges associated with these strategies, the most crucial one is to accept the transition from a world where companies have a total control over the processes of value delivering to customers, to a world where managers can only have an influence on an ecosystem who creates the actual value (Hagiu & Altman 2017).

1.2.2 Strategic networks analysis

As a consequence to the rise of platform business models, strategy is shifting from managing internal resources to the art of managing the firm's networks and leveraging the data that flows through them (Iansiti & Lakhani 2020). Companies who can act as network hubs, connect various businesses through digital nodes, aggregate massive amounts of data and extract value through analytics and AI, will reach critical competitive advantage.

For this reason, Iansiti and Lakhani (2020) propose to implement a systematic network analysis in firms' strategic analysis. The idea behind this approach is that through the creation of connections between firms and between networks, and with the aggregation of the data resulting from these interactions, eventually two kinds of effects would emerge (Iansiti & Lakhani 2020):

- *network effects*, which consists in the value added by increasing the number of connections within and across networks.
- *learning effects*, which capture the value added by increasing the amount of data flowing through the same networks.

The objective of the analysis is to evaluate the potential opportunity for value creation and value capture generated by the connections to networks, and by leveraging network and learning effects. The impact on value of these two effects is particularly sensitive to scale. Only little value will be generated in the beginning, with few connections and modest amounts of data, but, with an increase in the scale, the value curve will increase more sharply. The stronger the network and learning effects, the sharper the value curve.

A wide network analysis should start with the mapping all the major networks that the business is potentially connectable to, realising where the valuable data is and what are the value creation opportunities that exist. The digital operating model enables the company to easily connect with a multitude of other complementary networks. Firstly, it is necessary to analyze each network on an individual level, due to its specific properties, opportunities, willingness to pay, and level of competition, and then after that, to analyze the potential synergies across the networks.

Next is to consider which is the potential of each network in the business when it comes to value creation and value capture. There are plenty of dynamics which shape value creation and value capture in a network structure. The factors that can impact on value creation are:

1. *Network Effects*. The underlying value or utility of a product or service is exponentially proportional to the number of users utilizing the service. Network effects consist in direct network effects whenever the users value the presence of other users, and indirect network effects when users on one side of the network attribute value to the presence of users on another side of the network. Moreover, indirect network effects can be either two-sided (reciprocal between the two groups of users) or one-sided. Another aspect to consider is the strength of such effects. While in general it is true that the larger the network, the greater the value, firms which rely on weak network effects will not get any concrete competitive advantage and will be exposed to challenge by competitors. On the other hand, solid network effects can lead to a more long-lasting competitive advantage and to an increase in market concentration.
2. *Learning Effects*. Learning effects can either add value to existing network effects or generate value on their own. Generally, the more data is used to train and optimize an algorithm, the more accurate the algorithm's output will be, and the more complex the problem that the algorithm can be applied to solve. As operating models grow to embody a multitude of algorithms, learning effects will increase the impact of scale and scope on the company's value creation. As the user base increases and the scale becomes larger, the more data that is available, and the greater the value. At the beginning, the accuracy of most algorithms increases with the square root of the number of data points, at least in the first stage, and then it levels off as the algorithm is fully trained. But when a business is driven by more than one algorithm, the combined value of their learning effects can increase.
3. *Clusters*. Network's structure also has an important impact on how the value of a network increases with its size. Clustering affects the sustainability of network-based business models. Many digital business models are operating globally (i.e., people travelling do not care much about the number of Airbnb hosts in their home cities; instead, they care about the number of hosts in the cities they wish to visit). But, generally speaking, global networks are more concentrated around a small number of critical hubs. Typically, barriers to competition are high, and for the dominant player it is relatively easy for the dominant player. On the contrary, when a network is highly clustered, grouped around individual urban locations, the global scale of a company will not be relevant for customers in a given cluster (i.e., Uber drivers in a Boston neighbourhood will care only about the number of riders available in that same neighbourhood, and the same is true for riders). Therefore, the more a network is

fragmented into local clusters, the smaller the impact of scale and network effects, making it easier for challengers to enter. Typically, clustered networks are highly competitive, as any competitor with local scale can reach similar efficiencies, and the impact of scale is effectively limited to the level required to serve the local cluster. Clustered network with such structure makes it possible for a competitor with less scale to reach critical mass in a local network and also to take off through a lower price or a differentiated offer.

Finally, the strength and structure of network and learning effects change over time. Changes can either strengthen or weaken the value creation curves, and consequently markets can become much more or less competitive. When the strength of a network effect decreases, affected markets become less concentrated.

Regarding the value capture, the ability to capture value in a platform business model depends on many competitive dynamics. The factors that can impact on the value capture are:

a. *Multihoming*. Multihoming refers to the feasibility of competitive alternatives, particularly in situations where users or service providers in a network can create links with multiple platforms or hub companies ("homes") at the same time. If a network hub suffers competition from another hub connecting to a network in a similar way, the ability of the first network hub to gain value from the network will be threatened, especially if the switching costs are low enough to enable users to easily use the other hub.
b. *Disintermediation*. Wherein nodes in a network can easily bypass the firm to connect directly, disintermediation can also be a significant problem for capturing value. This issue is a very frequent, especially in marketplaces that provide only a connection between network participants. In these circumstances most (if not all) the value created is delivered after the first connection is made. Hubs have put in place various mechanisms to deter disintermediation, some of them include requiring terms of service that demand users to conduct all transactions on the platform or blocking users from exchanging contact information, at least before payment is confirmed. These types of strategies, however, are often not effective. A more honourable way to discourage disintermediation is to increase the value for users of conducting business through the hub. Hubs in fact may facilitate transactions by providing insurance, payment escrow, or communication tools, resolving disputes, or monitoring transactions. However, these services might become less valuable to users after they

build strong trust among themselves. Another effective way to reduce disintermediation is the reduction of transaction fees and make up the revenue on different market sides. If disintermediation is a threat, providing complementary services could be more effective than charging transaction fees.

c. *Network bridging*. Network bridging can improve and even save a company's business model by creating new connections through previously separate economic networks, making use of more favourable competitive dynamics and different willingness to pay. Network participants can enhance their ability to create and capture value when they connect to multiple networks, connecting with each other to create important synergies. It is worth pointing out that data-driven resources are almost inevitably useful in many scenarios and on multiple sides of the network. Companies that manage to build a critical mass of users can use this resource to acquire value on new and different networks. As the most successful network hubs connect between markets, they can be increasingly effective in driving connections between previously disconnected sectors.

Ultimately, with a strategic approach to networks, companies can express the intrinsic value of their multisided business models.

1.3 Machine Learning

In computer science, an *algorithm* is a set of rules or instructions that enable computer programs to perform tasks or to solve a given problem.[25] Such instructions need to be written in detail by programmers, and typically enable the computer program to process some inputs into a desired output. However, besides representing a costly and time-consuming activity, it is often difficult to write accurate instructions, especially when dealing with real-world problems, where a number of factors have to be taken into account and which are often solved by humans thanks to their prior knowledge about the world. In 1959 Arthur Samuel, among the pioneers of AI, wrote in an article where he was applying machine learning techniques to the game of checkers that "programming computers to learn from experience should eventually eliminate the need for much of this detailed programming effort" (Samuel 1959, p. 211). In fact, since artificial intelligence was theorized, great emphasis has been placed on the

[25] See the Cambridge Dictionary's definition: https://dictionary.cambridge.org/dictionary/english/algorithm

possibility of building intelligent computer programs, capable of automatically learning from data the best strategy to solve a problem or to perform a task (Mitchell 1997).

Machine learning (ML) enables to overcome the difficulty of writing accurate instructions by feeding a computer program with experience of the real world in the form of datasets and then by training a *learning algorithm* to learn a model which performs well not only on the data used for training, but also on new, never seen before data relating to the same problem. In other words, the iteration of the inputs-algorithm-outputs cycle enables the computer program to improve its performance over time and to learn a general model which performs well on a specific real-world task. The inherent shift promoted by this technique consists in passing from designing accurate algorithms, which transform inputs into outputs, to providing a set of well-defined inputs and outputs and then letting the computer program learn a model that transforms inputs into outputs through learning algorithms. Learning in this context is defined in a broad way as the ability to improve automatically from data (Mitchell 1997) and is usually framed into an optimization problem where the model parameters are learned by minimizing a cost function.

Machine learning is particularly useful when:

- it is possible to collect huge amounts of data about a real-world phenomenon and to extract from them a set of features which represent it well;
- the solution to be learned can be defined as a list of formal (mathematical) rules, but it is too computationally difficult for humans to find it.

ML techniques have enabled computers to efficiently perform a wide variety of useful real-world tasks, both in business and in the everyday life, such as automating entire activities in the manufacturing of products and in the delivery of services, recommending products to online customers, identifying customer segments, communicating with customers in human language through chatbots, making medical diagnoses, assessing credit risk, driving autonomous vehicles, and supporting at various stages scientific research, among others.

Since the objective of machine learning is to learn a model that generally works good with the data it is presented in probabilistic terms, but not to derive general rules which are true for every single piece of data (inductive approach), the output of such models is typically a *prediction* (Goodfellow, Bengio & Courville 2016). The model's predictions are then evaluated through a performance measure chosen by a human supervisor, and the model is trained until its performances are satisfactory. Despite ML techniques derive only probabilistic rules (which can lead, however, to generate hypotheses to be tested through more robust approaches), prediction is a fundamental component of most human activities, which can be described as a sequence of data, prediction, judgement, action, and outcomes (Agrawal, Gans & Goldfarb 2016); the superior capability of computers to store and analyze huge amounts of data to find patterns enables them to make more accurate predictions than a human could ever do, when appropriately designed. This highlights the fact that the real shift behind machine learning is the increased availability of accurate predictions based on data to perform tasks.

Additionally, the increasing availability and the decrease in the cost of computing power means not only that many more businesses can implement ML techniques to enhance their predictions, but also that many tasks where prediction was not typically a component could be reframed into prediction problems (Agrawal, Gans & Goldfarb 2016). In general, predictions may be related both to the present and the future. This point is clearer by looking at the main functions of a machine learning system, as described by Malone, Rus & Laubacher (2020):

- *descriptive*, when the system explains what happened by analyzing the data;
- *predictive*, when the system predicts what will happen based on the data;
- *prescriptive*, when the system makes suggestions on what actions to take based on the data.

It is now useful to formally examine the main components of a machine learning system. The components to be specified in order to obtain a well-defined learning problem have been described in the definition of machine learning given by Mitchell (1997, p. 2):

> "A computer program is said to learn from experience *E* with respect to some class of tasks *T* and performance measure *P*, if its performance at tasks in *T*, as measured by *P*, improves with experience *E*".

Defining the task (*T*) essentially means specifying what the output of the machine learning model should be. The two fundamental ML tasks are (Goodfellow, Bengio & Courville 2016):

- *classification,* where the computer program has to specify which category each input belongs to;
- *regression,* where the computer program has to predict a numerical value given the inputs.

These two tasks are equivalent except for the fact that in the first case the output is a label which represents a category, while in the second it is a numerical value.

The performance measure (P) displays how well the model learned by the learning algorithm performs on the data, and must be chosen according to the specific task and by taking into consideration the overall goal of designing the machine learning system. Typical performance measures are *accuracy* for classification tasks and *mean squared error* for regression.

Accuracy simply represents the proportion of inputs for which the model predicts the correct output (Goodfellow, Bengio & Courville 2016) and can be calculated as the number of correct predictions divided by the number of total predictions. Accuracy, however, is not always a completely satisfactory indicator of the performance of the model. As an example, if the task to be performed by the system is to recommend a good wine to an online customer based on his preferences,[26] the main goal could be to maximize the *precision* of the system, that is the number of correct good wines predictions (true positives) over the total good wines predictions (the sum of true positives and false positives). In fact, having a high precision is more vital for the system than having a high overall accuracy: the goal is not necessarily to classify all the examples correctly, but instead to suggests some good wines to customers with confidence, at the cost of misclassifying some good wines as bad wines (more on this point in *Section 2.3*).

The Experience (E) is arguably the most critical component of a machine learning system. The quantity and quality of data used for learning inevitably impacts the quality of the model, and any bias in the data will result in a bias the final model (Malone, Rus & Laubacher 2020). Experience of the real world is fed into the machine learning system through *datasets*, which are composed as collections of *examples* that represent an object or phenomenon. In turn, each example is typically a vector containing a number of *features* that can be quantitatively measured. Therefore, the final design of a dataset is a matrix where each row corresponds to an example. Measuring and designing the right set of features to extract for the target task is key in solving most straightforward machine learning applications (Goodfellow, Bengio & Courville 2016).

The data can be collected once for all (*batch learning*) or by interacting with the environment over time (*online learning*). The nature of the data to be collected is strictly related to the

[26] This problem can be thought as a binary classification problem, where the categories to be predicted are "good wines" and "bad wines".

learning paradigm chosen to train the model, which can be typically one of three: supervised, unsupervised and reinforcement learning.

In supervised learning, which represents the most widespread paradigm, the ML system experiences a dataset where each example is associated with a label, provided by a human supervisor (Goodfellow, Bengio & Courville 2016). Such labels enable to perform a classification task: the system is presented with many examples that represent a specific class, and has to learn a model that correctly classifies each of them in the right class. This task can be adapted to address many real world problems, from simply labelling e-mails as spam, to recommend products to customers or to perform health diagnoses.

In *unsupervised learning*, the dataset experienced by the system without labels: it is up to the ML system to discover regularities and patterns in the data. A classic example of a task performed by unsupervised learning algorithms is *clustering*, where the dataset is divided into clusters of similar examples (Goodfellow, Bengio & Courville 2016). This enables for example to perform customer segmentation, or to analyze large bodies of text for to model the topics contained as clusters of similar words. Another useful task is *association rule mining*, where correlations in the examples are leveraged to discover interesting rules, for example for the purpose of market basket analysis. Additionally, *anomaly detection* analyzes patterns in the data to enable predictive maintenance of machines or to detect faults in the production processes.

Finally, in *reinforcement learning* an agent interacts with an environment by performing an action and receiving a reward (which can be positive, neutral, or negative). The goal of the agent is to learn from the feedback received the strategy to maximize the total sum of rewards. The agent learns from trial and error and, when its performance is satisfactory, can be deployed in the real world. The typical trade-off faced by such ML systems is deciding whether it is more appropriate to train the model over a long period (*exploration* phase) or whether it's already performing well enough and it's more convenient to employ it and let it learn in the real world (*exploitation*) (Iansiti & Lakhani 2020). Besides enabling to build autonomous robots, reinforcement learning is useful for a number of other tasks which require decision making in dynamic contexts. Some examples are dynamically setting product prices, selecting digital ads to present to online users or creating an optimal portfolio in finance.

The data collected through these learning paradigms constitutes the *training data* for the model. Basically, the goal of an ML system is to learn the model, among all the possible ones it could learn, that best maps the inputs into the desired outputs. In a more formal way, the learning algorithm must find within its *Hypothesis Space* (which consists in all the hypotheses that such algorithm could possibly implement) the hypothesis that best fits the available data and other additional knowledge or constraint encoded by a human supervisor (Mitchell 1997). Therefore, learning typically involves finding, among the set of functions that can be learned by the algorithm, the function that best approximates the target function (Goodfellow, Bengio & Courville 2016).

Once enough training data is collected and an appropriate learning algorithm has been chosen, the ML system is trained until its performance on the training data is satisfactory, which means that the error on the training set is minimized. However, reaching a small *training error* is not completely satisfactory. In fact, when designing ML systems, we are interested in *generalization,* that is the ability of the learned model to perform well on new data pertaining to the same task. In other words, we want to minimize the *generalization error* on unseen data as well. The generalization error can be defined as "the expected value of the error on a new input… drawn from the distribution of inputs we expect the system to encounter in practice" (Goodfellow, Bengio & Courville 2016, p. 108), and represents the true error of the model. This is why the performance of the model is usually tested on test sets containing new unseen data coming from the same initial distribution.

When our model has a training error lower than the generalization error, it is said to *overfit*. This can be interpreted as if the learning algorithm has learned the training data too well in such a way that it perfectly explains all the variability in the data, but at the same time it cannot generalize on new data as the learned model is too complex and specialized on the training data. More formally, a "hypothesis overfits the training examples if some other hypothesis that fits the training examples less well actually performs better over the entire distribution of instances (i.e., including instances beyond the training set)" (Mitchell 1997, p. 67). This is where a popular principle becomes useful, the *Occam's Razor*, which poses a preference for simple hypotheses. This is due to the fact that simple hypotheses are more likely to generalize (Goodfellow, Bengio & Courville 2016). However, the learning algorithm should still be able to learn an enough complex hypothesis to achieve a good training error.

This leads to conclude that to find the optimal balance between training error and generalization error it is necessary to choose a learning algorithm whose ability to find complex hypotheses is adequate for the task to be performed, and then providing it with an adequate amount of data (Goodfellow, Bengio & Courville 2016).

An additional remark on this point comes from the *No Free Lunch Theorem*, which says that in general there does not exist a better machine learning algorithm than another over all the possible data distributions (Goodfellow, Bengio & Courville 2016). This connects to the initial consideration that machine learning models have to work well in the real world. Some learning algorithms perform better than others on some data distributions. Given that the data collected in the real world naturally display some regularities and patterns, and thus have specific distributions, it is critical to design/choose learning algorithms that perform well on the data distributions they will encounter when performing their specific tasks in the real world. Furthermore, it may be useful to put some preferences towards some solutions over others into the learning algorithm to help reduce the generalization error (Goodfellow, Bengio & Courville 2016). Such process is called *regularization*.

In order to help businesses to design and implement machine learning systems, Agrawal, Gans, & Goldfarb (2016) developed the *AI Canvas (Figure 10)*. Such visual representation enables to keep track of all the critical decisions to be made when employing prediction to perform a task or to solve a problem. In particular, the top row displays the single components for to complete the task, that are the prediction based on data, the judgement of the outputs of the predictive model, the actions associated to those outputs and the final evaluation of the overall outcome of the task. The bottom row instead represents three key considerations for to improve the performance of the model, that consist in determining the data needed for training, the data needed for running the algorithm in its real-world application and how to exploit feedback loops.

PREDICTION	JUDGMENT	ACTION	OUTCOME
What do you need to know to make the decision?	How do you value different outcomes and errors?	What are you trying to do?	What are your metrics for task success?

INPUT	TRAINING	FEEDBACK
What data do you need to run the predictive algorithm?	What data do you need to train the predictive algorithm?	How can you use the outcomes to improve the algorithm?

Figure 10. The AI Canvas. Source: redrawn from Agrawal, Gans & Goldfarb (2018).

As previously stated, the power of machine learning algorithms is that they are well suited for many real-world problems, while at the same time they do not need a strong programming effort as they are capable of automatically learning from data. However, the choice of the right learning algorithm is critical in achieving a good performance and generalization.

1.3.1 Decision Trees & Random Forests

We can now go through an example to exemplify the method to build a machine learning system. A classic example of a learning problem is provided by Mitchell (1997): to decide whether to go or not to play tennis, based on the available weather information. The dataset representing the problem, as shown in *Table 1*, is a matrix where each row represents an example (a day), which is described by a set of attributes (outlook, humidity, and wind) and to which is attached a label (PlayTennis). Having a label attached to each example expressing a positive or negative class means that we are in a supervised learning paradigm and this enables us to frame the problem into a binary classification task (with two classes Yes/No). Therefore, we want to train a machine learning algorithm to learn a model which exploits some rules about the three attributes to correctly predict whether to play tennis or not, when facing new data in future days.

Example	*Attributes*			*Goal*
Day	**Outlook**	**Humidity**	**Wind**	**Play Tennis**
D1	Sunny	High	Weak	No
D2	Sunny	High	Strong	No
D3	Cloudy	High	Weak	Yes
D4	Rainy	High	Weak	Yes
D5	Rainy	Normal	Weak	Yes
D6	Rainy	Normal	Strong	No
D7	Cloudy	Normal	Strong	Yes
D8	Sunny	High	Weak	No
D9	Sunny	Normal	Weak	Yes
D10	Rainy	Normal	Weak	Yes
D11	Sunny	Normal	Strong	Yes
D12	Cloudy	High	Strong	Yes
D13	Cloudy	Normal	Weak	Yes
D14	Rainy	High	Strong	No

Table 1. Dataset for the "go play tennis" problem. Source: adapted from Mitchell (1997), p. 59.

A powerful and widespread learning algorithm for classification tasks is *Decision Tree*. Such algorithm represents the target function as a decision tree where, starting from a *root node* to some *leaf nodes*, a tree is grown downwards by testing for some attribute at each node in order to sort the examples into the correct class (Mitchell 1997). Each *branch* of the tree then represents the possible values that the tested attribute could take and represent the decision rule that enables to partition the dataset and to move the examples downwards until all of them are sorted into the correct class, and for this reason they are called decision branches.

In general, Decision Tree learning is appropriate when (Mitchell 1997):

- each example is described by a fixed set of attributes which can take a small number of possible discrete values, or even real values when the learning algorithm is capable of setting some intervals that partition such continuous attributes values into newly created discrete-valued attributes.
- The target function to be learned has discrete output values, such as a boolean classification (yes or no), or multiple discrete-valued classes.
- Disjunctive expressions may be needed; in fact, an additional feature of Decision Tree is that the learned tree can be re-represented as a sequence of if-then rules.
- The training data may contain missing values or errors, as decision trees are typically robust to errors.

Such algorithms have been successfully employed to tackle a varied range of tasks such as making health diagnoses, assessing credit risk, classify equipment malfunctions by their cause, and recommending products to online users.

Considering all the characteristics listed above, Decision Tree learning is suitable for to solve the “go play tennis” problem. However, it remains to understand how the attributes will be tested in each node. In fact, the algorithm grows the tree from the root downwards by testing one attribute in every node. Given that the goal of the algorithm is to correctly classify all the examples, the root node is created by performing a statistical test on all attributes to determine which one is the best predictor of the target class, meaning that it correctly classifies the most examples, and then choosing such attribute. Then, each possible value of the root attribute creates a branch where the examples will descend towards the correct leaf node. In each leaf node the test will be repeated again by choosing the best attribute that can correctly classify the most examples in that new node, then creating new branches and continuing to partition the training set of examples. The process goes on until all the examples are correctly classified.

The *splitting criterion* that enables to choose the attributes among the candidates through which create the nodes is a statistical property called *information gain*. Such measure is defined as the "expected reduction in entropy caused by partitioning the examples according to… [an] attribute" (Mitchell 1997, p. 57). Entropy measures the homogeneity (purity) of the examples in each node; it suffices to say here that its value ranges between 0 and 1, and it is equal to zero when all the examples have the same label (they belong to the same class). It is straightforward than that the attribute which is expected to reduce entropy the most will be the best attribute for classifying the data in that node.

We can finally implement our algorithm and observe the tree that has grown (Figure 11). The algorithm chose *Outlook* as the attribute for the root node, as it has been considered the best predictor for the target classes. For every possible value of Outlook, that are *Sunny*, *Cloudy*, and *Rain*, a branch is created, partitioning the initial training dataset according to those values and carrying the examples to the leaf nodes, where subtrees will grow and the whole process will be repeated, creating many levels in the decision tree, until the learned tree is capable of perfectly classifying all the training examples.

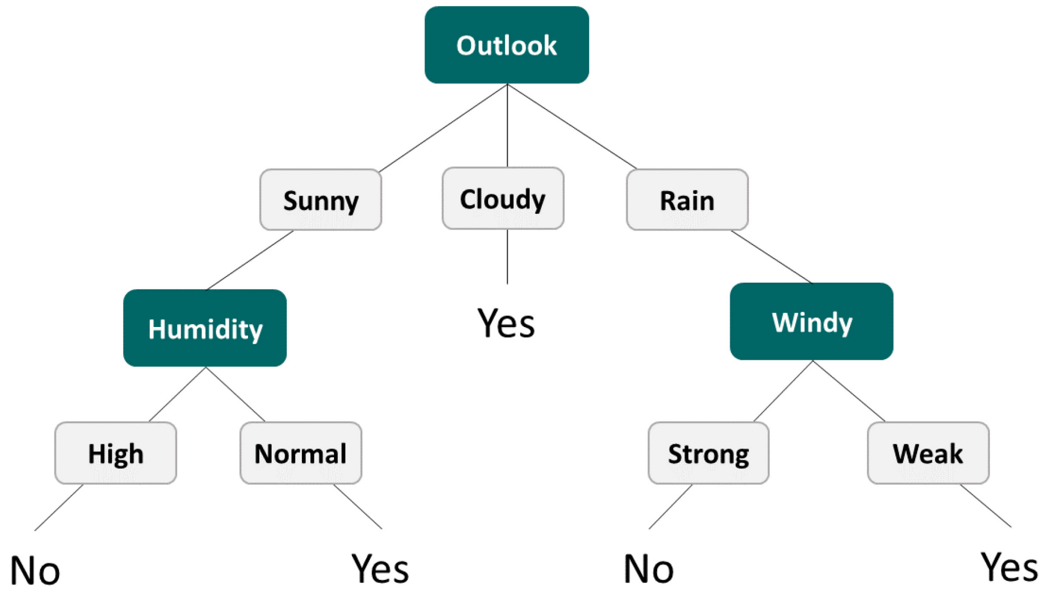

Figure 11. Decision Tree for deciding whether to play tennis or not. Source: redrawn from Mitchell (1997), p. 53.

The main drawback of decision trees is that sometimes they can grow so deep and so complex that they will perfectly fit all the data in the training set, but the resulting model will not perform better than any simpler model on new unseen data. A primary solution to overfitting consists in *pruning* the tree at a certain level of depth by removing some subtrees.

The inherent strength of decision trees relies in the fact that they are said to be *white-boxes*, or to have a transparent design. In fact, they can also “be re-represented as sets of if-then rules to improve human readability” (Mitchell 1997, p. 52). This means that they exploit a symbolic approach which can be easily understood by humans. This is key aspect when machine learning algorithms are employed to enhance human decision-making.

Despite their success and wide diffusion, decision trees have some limitations. In particular, many real world problems may need hundreds or thousands of attributes to be described, with each attribute carrying only a small amount of information (Breiman 2001). This can be the case when analyzing large bodies of text or when making complex medical diagnoses. In such cases, a single decision tree will likely lack the capacity needed for to learn a general model. To overcome this issue, the *Random Forest* algorithm has been proposed. In this method, a large number of decision trees are grown (this is why they are called forests), each of them learned on a random subset of the training examples (rows) and attributes (columns) (this is why random). Then, each tree in the forest will express a vote for the class to which assign each example. The most popular class is usually defined through a majority rule, and so each example is classified.

The injection of randomness in the selection of inputs and features typically enables random forests decrease the generalization error and to outperform decision trees in terms of accuracy. However, this comes at the expense of explainability, as the large number of trees and attributes does not allow humans to appreciate the inner workings of the learned model, contrarily to what happens in decision trees. In fact, random forests are treated as *black-boxes*, where knowledge is distributed across a wide number of trees and nodes and not explicitly represented (Calegari, Ciatto & Omicini 2020). Still, it is possible to extract some of that knowledge in the form of the attributes importance for the classification task.

Machine learning systems have been successful in leveraging the digital representation of the world in the form of data to perform some well-specified tasks with a performance beyond human capabilities. This has been enabled by the *datafication* process of all interactions between people and software and with the diffusion of sensors to gather data from the real world. However, the problems tackled by ML algorithms are usually very narrow, and their performance is highly dependent on the data representations they are given. This means that traditional ML algorithms need to process some sets of features which are carefully designed and that clearly represent the underlying problem or phenomenon to be studied. However, it is clear that many real-world phenomena can only be described with a countless number of attributes, interconnected through complex relationships. It is here that humans leverage their accumulated knowledge of the world to learn complex concepts which help them to make sense of the phenomena around them, then isolating the relevant features of a problem and eventually solving it.

1.3.2 Artificial Neural Networks

Deep Learning (DL) is a subset of machine learning which entails the techniques that enable computer programs to learn complex concepts by relating simpler concepts, thus representing the world as a hierarchy of concepts (Goodfellow, Bengio & Courville 2016). This translates into a superior capability of learning complex models (functions), by searching into a very large Hypotheses Space.

Deep learning enabled computers to reach unprecedented performances on complex real world tasks like computer vision, natural language processing and speech recognition, robotics, bioinformatics, search engines, and online recommendations (Goodfellow, Bengio & Courville 2016). In fact, DL techniques are suitable for tasks where the training data is noisy and complex, such as sensor data from cameras and microphones, online browsing history or the composition of biological samples (Mitchell 1997).

The development of these methods became possible in the first decade of the 2000s thanks to the exponential increase in the size and variety of datasets and the increase in computational power of computer machines. Besides these enabling conditions, a fundamental step was the possibility to combine faster CPUs and new general-purpose GPUs (graphics processing units) to perform *parallel computing*. This has been a key aspect in the development of deep learning techniques and highlights the biological inspiration that has been taken from the human brain. In fact, our brain is made of billions of densely interconnected neurons, where each neuron receives an input in the form of an electrical impulse, processes that input, and then transmits its output to numerous other neurons. Despite this intricate architecture, our brain enables us to take some decisions or actions extremely quickly, so fast that the steps that the information could have made by being transmitted through neurons are a very limited number. This has led to the understanding that the human brain heavily leverages parallel processes, and that the learned representations are therefore distributed over many neurons (Mitchell 1997).

In order to replicate and exploit the powerful network architecture of the human brain in a computer program, the first step is to create a mathematical model of its basic unit, the neuron. The first attempt to replicate the functioning of a neuron in artificial intelligence is the *perceptron*, introduced in the 1950s. Analogously to a neuron, the perceptron receives some inputs into an input layer and then transforms them into the output layer through weighted connections. Formally, a perceptron receives a vector containing real-valued inputs and maps them into an output value by learning a weight vector (Mitchell 1997). Typically, the output can take only the values +1/-1, making it suitable for binary classification tasks. Since the function learned is just a linear combination of the inputs, the perceptron is a linear classifier and works only for linearly separable data.

However, to solve complex problems with real-world data we often need computer programs capable of learning nonlinear, complex functions. To fulfill this need and to leverage the power of the brain architecture (and thus of parallel computing), *artificial neural networks* have been developed, composed of many interconnected units (artificial neurons).

A *Deep Feedforward Neural Network* represents the quintessential example of an artificial neural network. Such networks simply map some set of input values to output values by approximating a function, which is composed by many simpler, nonlinear functions connected in a chain structure (Goodfellow, Bengio & Courville 2016).

The architecture of feedforward networks is composed by many units that act in parallel, grouped into layers. Each unit works as a neuron as it receives inputs from many other connected units and produces its own output, typically through a nonlinear function (Goodfellow, Bengio & Courville 2016). Additionally, the chain structure of layers means that each layers is a function of the layer that preceded it (Goodfellow, Bengio & Courville 2016).

A typical network is structured into an input layer, a certain number of hidden layers, and an output layer (*Figure 12*). Essentially, the input layer contains labelled training examples, and the output layer must learn an approximated function that returns a value for each example close to their label. The fundamental characteristic of feedforward networks is that the behavior of the hidden layers is not directly specified by the training data, but the learning algorithms can use them to process the inputs in order to automatically discover new useful features of the training examples, which will help them in finding the best approximation of the target function (Mitchell 1997; Goodfellow, Bengio & Courville 2016).

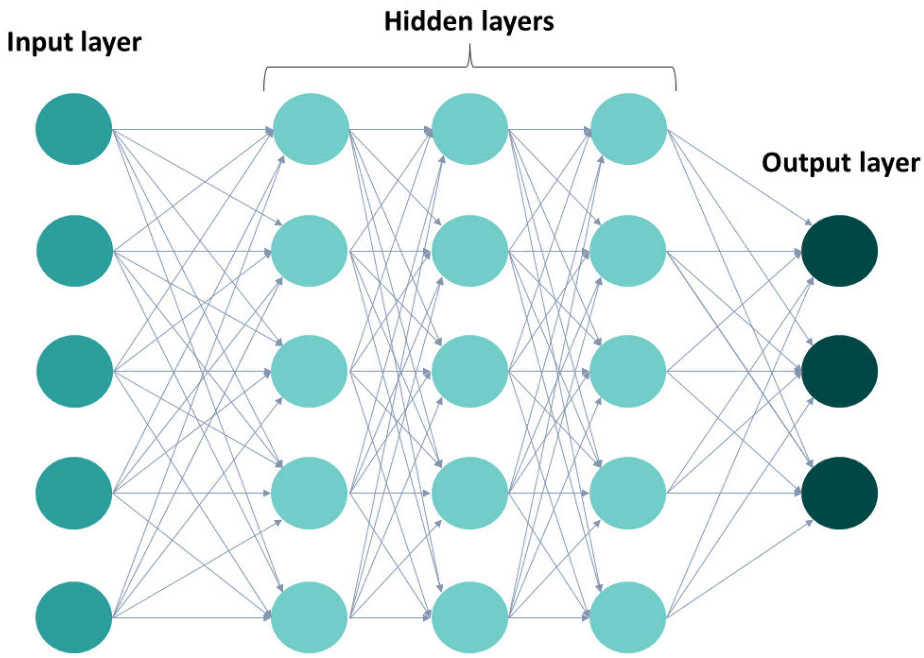

Figure 12. Deep Neural Network. Source: own elaboration.

A network with many hidden layers is said to be deep, while the number of units in each layer determines its width.

The learning algorithm in artificial neural networks is typically *Backpropagation*. The learning problem to solve here is to search into the large hypothesis space defined by all possible weights for all the units in the network and to learn the optimal weight values which minimize the error (Mitchell 1997). Essentially, the algorithm initializes all the network weights to a small, close to zero random value; then, each training example is processed through the network iteratively. At every iteration, an error for the output on that example is calculated at the output node; then, the gradient with respect to that error is computed, which provides information on the direction towards which to adjust each weight of each unit through small steps at every iteration. This procedure goes on until the network has learned a set of weights that enable the network to achieve a satisfactory performance.

An important result of deep neural networks is the *Universal Approximation Theorem*, which states that, regardless of what function we want to learn, there exists a large neural network capable of approximating such function. In this sense, "feedforward networks provide a universal system for representing functions" (Goodfellow, Bengio & Courville 2016, p. 195).

However, finding the network capable of learning the desired function representation remains an open problem, and is up to the ability of the designer of the DL system.

Clearly, with deep learning the problem of the interpretability and explainability of the model learned becomes particularly severe, as deep learning algorithms combine a purely numerical (s*ub-symbolic*) approach with an extremely complex structure that goes beyond human cognitive capabilities. This is why, despite their capability to solve complex real-world problems and the striking levels of accuracy achieved in many tasks, there is a growing concern over their widespread adoption in the economy and everyday life as biased AI behaviors can be harmful for people, businesses, or the whole society. This is why *explainable AI (XAI)* is emerging as a new research field, aiming at studying the design of intelligent systems capable of showing or explaining to humans the reasons behind their behaviors (Calegari, Ciatto & Omicini 2020). Some approaches have been proposed in order to generate *hybrid AI approaches* by means of integration or combination of sub-symbolic and symbolic approaches, therefore combining the performances of the former with the inherent explainability of the latter. An example of this approach are the AI systems where, starting from a sub-symbolic classifier, such as a deep neural network, a decision tree is then extracted.

Additionally, today major software companies are promoting responsible AI principles,[27] where, along with some essential objectives of *reliability* and *safety, privacy* and *security*, AI systems must also ensure *fairness* and *inclusiveness*, they should be *transparent* and therefore understandable and they have to be *accountable*, meaning that people have to take responsibility for the behavior of AI systems.

Explainable and fair AI is a key challenge in a time where the increasing sensorization and datafication of the real world and the drop in the cost of prediction are leading to an economic shift where many more tasks, from manufacturing products to facing customers in retail environments, are being performed by humans and computers working side-by-side and augmenting each other – the latter crunching large amounts of data to find patterns and to generate accurate prediction, and the former performing judgement (Malone, Rus & Laubacher 2020; Agrawal, Gans & Goldfarb 2016).

[27] See https://www.microsoft.com/en-us/ai/responsible-ai

1.4 Recommender Systems

Among the many real world applications of machine learning, recommender systems are certainly among the ones that reached a widespread adoption. From recommending movies and news to helping choose the perfect bottle of wine, such systems are the fundamental engine behind the success of many digital platform business models, where a large number of items and users need to be efficiently matched.

In a formal definition, *Recommender Systems (RSs or RecSys)* “are software tools and techniques that provide suggestions for items that are most likely of interest to a particular user” (Ricci, Rokach & Shapira 2015, p. 1). Operationally, such systems leverage data about items, users, and their interactions and apply machine learning algorithms to present new relevant items to a given user, according to his preferences or needs.

The need for recommender systems arose when the newborn e-commerce websites in the late 90s and early 2000s started to offer online customers with unprecedented assortments of items[28] in terms of quantity, quality and variety. This meant that customers had to deal with an overwhelming number of available alternatives when making their shopping decisions, which, instead of incrementing their well-being, ended up causing an information overload. This led consumers to make inconsistent purchasing decisions and to experience a state of anxiety in the online environment, falling into the “paradox of choice” described by Barry Schwartz (2004, cited in Kotler, Kartajaya & Setiawan 2021, p. 60).

Schwartz’s argument goes further to state that because humans have a selective attention, eliminating choices will ultimately improve happiness. Indeed, RecSys development had been based on the observation that individuals often lack enough knowledge to make an autonomous decision and therefore rely on recommendations provided by others (Ricci, Rokach & Shapira 2015). It follows that such recommendations may be delivered by AI systems when combining machine learning techniques with the huge amount of data about users and items available in online environments, thus building intelligent systems capable of matching specific customer needs with the right items (Kotler, Kartajaya & Setiawan 2021). Indeed, RecSys represent a striking example of “large scale usage of machine learning and data mining algorithms in commercial practice” (Ricci, Rokach & Shapira 2015 p. 17) and a prominent solution to the *information overload* risk.

[28] In this context, items can represent both products and services.

Today, RecSys play a fundamental role for the User Experience of the most prominent platforms such as Amazon, YouTube, Netflix, Spotify, LinkedIn, Facebook, Tripadvisor, IMDb, but are present to various degrees in a multitude of smaller ones, thus becoming a standard for the development of an online platform.

The reason behind the success of these systems is that they are beneficial for both the service providers and the users at the same time. In the case of service providers, recommendations produced by the system are valuable for generating additional revenue at relatively low risk (Calero Valdez, Ziefle & Verbert 2016). The benefits of implementing a recommender system for this category can be summarized as following (Ricci, Rokach & Shapira 2015):

- Increasing the number of items sold;
- Selling more diverse items;
- Increasing the user satisfaction;
- Increasing user fidelity;
- Better understanding of the user preferences.

The latter point is worth to be highlighted: as users continuously interact with the system, leaving feedbacks or executing transactions, the data collected can be leveraged to generate a data-driven understanding about users preferences at an unprecedented scale and speed. This knowledge can be then employed for many business purposes such as products and services development, better focusing the marketing effort, and making logistics and distribution more efficient.

From users' perspectives, recommender systems improve their user experience (UX) in various ways. They are a form of human-computer interaction (HCI) which leverages at the same time *visual representations* that facilitate human understanding and a constant *dialogue with the system* via feedbacks, ratings and other forms of interaction (Calero Valdez, Ziefle & Verbert 2016). This makes them an effective *information retrieval* tool, thus helping online customers to navigate through a multitude of items and to find the ones they may be interested in. However, the development of such systems enabled them to become today proper *decision making* tools (Calero Valdez, Ziefle & Verbert 2016). In fact, while early RecSys' outputs were simply ranked lists of the most popular items, today recommendations are typically *personalized*, meaning that "different users or user groups benefit from diverse, tailored suggestions" (Ricci, Rokach & Shapira 2015, p. 1). To do so, RecSys reframe the recommendation into a prediction problem, trying to predict how much a user will like/need an item by leveraging the recorded user preferences, the constraints set, and the available information about the items. Furthermore, the value of personalized recommendations is strengthened by the fact that many RecSys are *item-specific*, meaning that every component of the system is optimized in order to give an effective recommendation for the specific kind of items considered, including the graphical interface, the algorithms employed, and the data sources leveraged (Ricci, Rokach & Shapira 2015). The characteristics discussed above, combined with the increasing availability of data and the level of sophistication many RecSys have reached, have made them today powerful assistants in supporting customers to make better decisions in online shopping environments. It follows that the potential benefits of RecSys for users are the following (Ricci, Rokach & Shapira 2015):

- Finding some good items;
- Finding all good items related to a particular user need;
- Getting recommended a sequence of items that is pleasing as a whole, or a bundle of items that make sense to be bought together;
- Getting an overall better browsing experience, searching between items that are more likely to fall within the scope of the users' interests;
- The possibility to provide feedback to the system on what the users like or dislike;
- The possibility of helping others and of influencing them with their feedbacks;
- The possibility of expressing themselves within a community.

Ultimately, RecSys are a primary example of an AI system allowing companies to perform real-time market research, which then empowers them to deliver the personalized online experiences increasingly demanded by users in a quick and consistent manner and at scale, as described by Kotler, Kartajaya, & Setiawan (2021).

In general terms, RecSys' tasks consist in predicting the utility of a set of items, comparing such items based on their utilities, and then suggesting the most useful ones to a user. Its architecture is determined by how it processes the data referred to three kinds of fundamental objects, as defined by Ricci, Rokach & Shapira (2015):

a) *Items.* Items are the recommended objects and are characterized in terms of their complexity and value/utility. RecSys can represent items in various ways, from simple ID codes to rich sets of attributes or other more sophisticated representation approaches. In the latter cases, they can leverage items' properties and features to various degrees in order to learn how the utility of those items depends on such features.

b) *Users*. Users are the end users of a RecSys, as a consequence their goals and characteristics must be taken into account in order to deliver a personalized experience. Therefore, every user is represented by building a user model that incorporates its preferences and needs. Data that can be effectively leveraged for this purpose may include the list of the user ratings, its sociodemographic attributes, its behavior pattern data, or even its relations with other users.

c) *Transactions*. Transactions refer to every recorded interaction between users and the system, which are stored to be leveraged by the recommendation algorithm. The quintessential form of transaction data is represented by the explicit rating the user gives to an item, which may take various forms (i.e., numerical, ordinal, binary). However, today implicit ratings such as browsing history, clicks and view times, are increasingly leveraged by state-of-the-art RecSys.

The main design choice the developer of an RecSys must face regards the algorithmic approach to be undertaken; this decision will have an impact on users' and items' representations, on the type of transaction data leveraged, on the external data sources to be integrated, and ultimately on the effectiveness of the system in the specific domain for which it is developed. A profound understanding of the domain where the system is going to be implemented is therefore essential to assure correspondence between the recommendation algorithm and the underlying logic of decision-making in such domain.

Despite the many possible approaches to recommendation, there are two main paradigms for recommendation that have gained widespread adoption and have led the development of the field, namely the collaborative filtering (CF) and content-based (CB) approaches.

Collaborative filtering approaches leverage the ratings produced by a community of users to deliver personalized recommendations to an active user, thus incorporating the principle of *word-of-mouth* without the need for integrating any external data about either items or users (Ricci et al. 2011). This can be achieved with a few different methods. *Neighbourhood-based* approaches, the first and simplest form of collaborative filtering, directly employ the user-item ratings stored in the system to predict ratings for new items, in two alternative ways (Ricci et al. 2011; Ricci, Rokach & Shapira 2015):

- *User-based approach*: the system recommends new items that other similar users had liked in the past, where similar users are the ones whose rating patterns overlap with those of the active user.
- *Item-based approach*: the system recommends new items similar to those the active user liked in the past, where similar items are those which have been rated similarly by many users.

User-based recommenders may generate more original and engaging recommendations, however, in large e-commerce platforms where the number of users outweighs the number of items, item-based recommenders are preferred for being accurate and more computationally efficient (Ricci, Rokach & Shapira 2015). Summing up, neighborhood-based approaches capture local associations in the stored user-item ratings data to make predictions and deliver personalized recommendations in a simple yet efficient way, making them the most popular recommendation approach.

More recently, *model-based* approaches emerged as a more sophisticated alternative. This family of collaborative filtering approaches aims at learning a predictive model out of the stored ratings by uncovering latent factors[29] that lie in the patterns of user-item interactions data. In this way, the recommender may learn key attributes of items (like the movie genre) and users (like the preference for science-fiction movies) that explain the observed ratings without the need to code them. A high correspondence between item and user factors then leads to a recommendation (Ricci et al. 2011). However, the latent factors may represent both easy interpretable and completely uninterpretable dimensions, and the purely numerical nature of these systems makes them typical *black-boxes*. Still, such approaches deliver a higher prediction accuracy than other collaborative filtering techniques and easily integrate implicit user feedbacks and temporal dynamics, thus capturing the evolution of users' taste over time (Ricci et al. 2011).

Content-based approaches, on the other hand, take a completely different route to recommendation. These systems leverage the available data about items' features, either manually encoded into the system or integrated from external knowledge sources, to recommend new items similar to those the active user has liked in the past (Ricci et al. 2011). The recommendation process is performed by matching up the user profile, where his preferences towards attributes are stored, with the attributes of a set of items. The success of

[29] Latent variables are variables that are not directly observed but are rather inferred from other variables that are observed.

these systems is critically linked to items representation. This may be obtained by simply extracting keywords from items' descriptions and user generated content, or via more sophisticated semantic indexing techniques that introduce concepts by integrating external knowledge sources (like ontologies and encyclopedic knowledge) or other natural language processing techniques (more on items representation in *Section 2.1*) (Ricci, Rokach & Shapira 2015). By storing the user ratings of well-represented items (rich in meaningful attributes for the specific domain considered), the system can effectively build a user profile that represents in a structured way the user preferences underlying those ratings, thus creating a formidable resource for refining recommendations and for other business and marketing purposes.

By confronting content-based and collaborative filtering approaches, the former ones have several advantages compared to the latter ones (Ricci et al. 2011):

- ✓ CB recommenders are *user-independent*, meaning that in order to start providing accurate recommendations they need only the active user ratings.
- ✓ CB recommenders are inherently more *transparent*, as explanations on why a recommendation has been presented could be provided easily as a list of the item's features that matched the users' profile.
- ✓ CB recommenders are capable of handling *newly-added items* not yet rated by the community of users, as the recommendation process only needs the item's content.

However, the advantages of CB approaches are coupled with some drawbacks, where CF approaches outperform (Ricci et al. 2011):

- ✓ CF recommenders are capable of recommending items for which *attributes are not available* or too difficult to extract, thanks to the ratings of the community of users.
- ✓ CF recommenders are capable of recommending items to *new users* who did not rate any item yet, thanks to the ratings of the community of users.
- ✓ CF recommenders are based on word-of-mouth and user ratings may be a better *indicator of quality* of items than the set of attributes used to describe the items.
- ✓ CF recommenders can make more original and surprising recommendations, as long as other users' rating patterns go in some unpredicted ways.

By comparing the two approaches, it looks clear how a combination of them could enable to overcome the limitations of both. Indeed, modern recommender systems are a combination of two or more recommendation techniques, and are typically referred to as *hybrid recommender systems*. In this way, state-of-the-art RecSys are capable today of integrating numerous and diverse sources of data to deliver highly accurate personalized recommendations and satisfying user experiences, while at the same time gathering massive amounts of data about users and items.

Looking into the future directions of development in this field, there is a clear trend towards integrating contextual information into the recommendation process. This does not include only easy to access information such as the knowledge level of the user or the time of the day, but may potentially expand in the near future to personality traits, moods and emotions of the user by capturing behavioral data (such as physiological data, facial expressions, speech and so on) through mobile and wearable devices and in-store sensors (Ricci, Rokach & Shapira 2015; Calero Valdez, Ziefle & Verbert 2016).

1.4.1 Content-based recommender systems architecture

Essentially, content-based recommender systems need to effectively represent items, to learn the user profile, and to put in place some strategies for comparing the two. Ricci et al. (2011) described the high level architecture of a CB recommender (as depicted in Figure 13), which typically comprises three fundamental components (that are responsible for the three key stages of the recommendation):

- *Content Analyzer*. This component must pre-process the unstructured data that is fed into the system in order to represent the content of items in a structured way. The data leveraged is typically in text form, retrieved from product descriptions, documents, web pages, news, or other sources; feature extraction techniques are usually performed, for example by representing a text document into keyword vectors.
- *Profile Learner*. This component collects data representative of the user preferences, for example the historical user feedback/ratings and the item representations related to those ratings (retrieved from the relative repositories), which will serve as positive and negative training examples; then, it employs supervised machine learning techniques to learn a model of user preferences (the user profile).
- *Filtering Component*. Finally, this component leverages the learned profile and generates a recommendation for the active user by matching a new item representation against the user preferences representation (stored in the user profile). This may result in a binary or continuous relevance judgement computed using some similarity metrics.

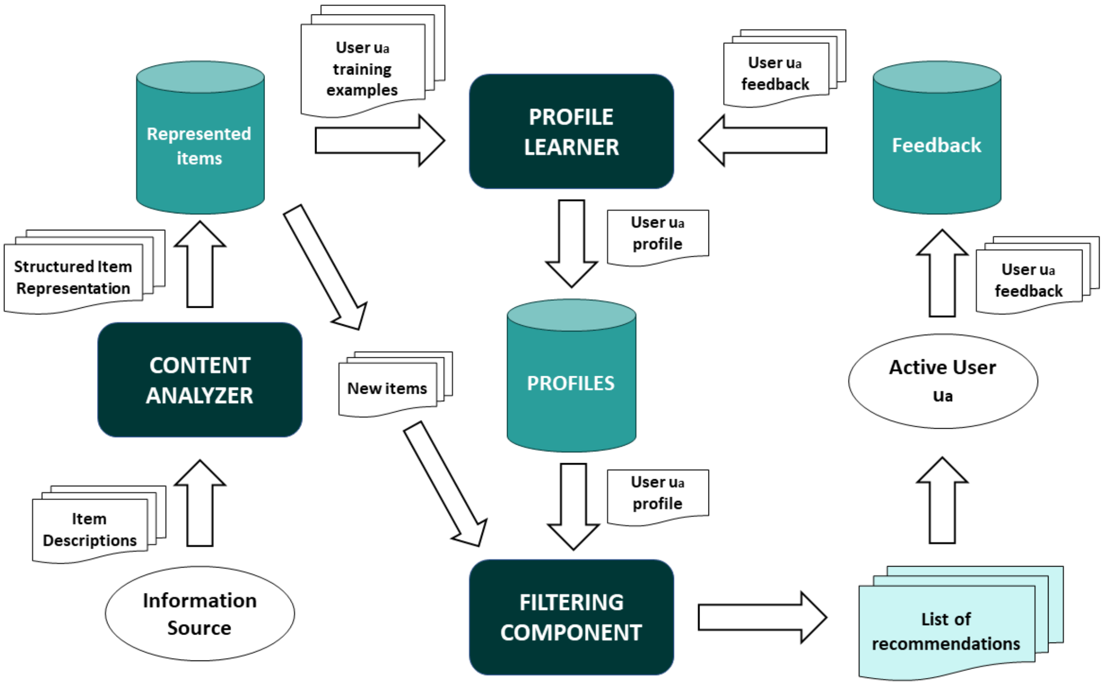

Figure 13. Content-based recommender system architecture. Source: redrawn from Ricci et al. (2011), p. 76.

It must be noticed how the list of recommendations that is generated by the system and presented to the active user is then subject to further feedback by him. In this way, the Profile Learner performs the learning process iteratively on the newly updated training set, thus automatically updating the user profile. This enables the system to both take into account the dynamics of user preferences over time (Ricci et al. 2011) and to put in motion a self-reinforcing loop which allow for the continuous refinement of recommendations.

1.4.2 Recommender systems evaluation

As anticipated in the previous sections, RecSys are information retrieval and decision-making tools of vital importance in online environments. In order to be really useful for users, they must be carefully designed in order to suit both the underlying logics of the domain in which they are employed and the purchasing decision processes of target customers. A well-designed RecSys will ultimately increase the number of transactions on the platform and revenues as well, while helping customers to make better decisions and making them enjoy an improved user experience (UX). Given its strategic relevance and the multitude of objectives it can pursue, it appears clear that the evaluation of a RecSys' performance must take into account several dimensions beyond the mere accuracy of the recommendation algorithm.

Indeed, it has been demonstrated in the literature that a high recommendation accuracy does not necessarily always correlate with user satisfaction, while other factors such as the knowledge level of the user may be relevant (Calero Valdez, Ziefle & Verbert 2016). Additionally, the credibility of the RecSys has been addressed as another vital dimension for increasing user trust in recommendations (Ricci, Rokach & Shapira 2015). These considerations highlight the need for implementing proper evaluation metrics, and for generating explanations of how the system works that may increase the acceptance of RecSys by users.

While accuracy still holds as one of the critical evaluation criteria for a recommender system, especially when comparing different algorithmic approaches, several other metrics have been proposed in the literature (Ricci et al. 2011):

- *Coverage*, that is the proportion of the total items available that the recommendation system can actually recommend;
- *Confidence*, which represents the probability that the value predicted by the system is indeed true (or alternatively the confidence interval of the recommendation);
- *Trust*: which represents user's trust in the system recommendations as calculated by online tests or the number of repeated users;
- *Novelty*, which measures how many of the successful recommendations are made for items that the user did not know about before;
- *Serendipity*, which measures how surprising and unexpected the successful recommendations are;
- *Diversity*, which measures the diversity, calculated as the opposite of similarity, across the recommended items.

While implementing a mix of these metrics may give a comprehensive view of the performance of the system to service providers and developers, in order to ultimately offer an overall better experience for the customer, the latter may still struggle to trust such systems. This is due by the fact that many recommender systems are treated as *black boxes*, offering no clue to users on why a recommendation was done. However, except in the case of purely numerical and computationally complex algorithms such as state-of-the-art collaborative filtering, which present an objective difficulty in interpreting their results, explanations for recommender system are generally possible and may be presented to users while pursuing some of the following objectives (Ricci et al. 2011):

- *Transparency*, meaning to provide an honest explanation of how the system works;
- *Trust*, which means trying to increase users' confidence in the recommendations by improving the transparency and accuracy of the system;
- *Persuasiveness*, which means trying to increase user acceptance of the recommendations and so to buy more;
- *Effectiveness*, meaning to help users make good decisions according to their purchasing criteria, for example by helping them to evaluate how well the suggested items match their preferences;
- *Satisfaction***,** meaning to increase users' enjoyment of interacting with the system and ultimately to provide a better user experience.

As the algorithmic approach chosen determines the architecture of the recommendation process, thoroughly transparent explanation should give an intuition of the underlying algorithm functioning. For example, collaborative-based approaches may generate explanations in terms of "Customers Who Bought This Item Also Bought…", while content-based approaches may list the item's features that match user preferences and which have been chosen to generate the recommendation.

CHAPTER 2: THE DIGITAL REPRESENTATION OF WINE

The digital representation of products and services is what enables companies to digitize the critical activities in the delivery of value, thus transforming and enhancing the customer experience in a variety of ways while gathering massive amounts of data about user preferences, which represent an increasingly strategic asset.

Wine is an intrinsically data-rich product, due to its complexity from a chemical point of view and to its equally complex sensory properties in terms of taste, aroma, and palate sensations. Therefore, wine can be effectively represented both in terms of text, for example through a description of its sensory characteristics, and in terms of numbers, for example through a dataset where each wine is represented as a set of physicochemical attributes. The digital representation of wine has important implications for the customer experience (CX) and for the wine value and supply chain, being the foundation upon which to build a digital ecosystem.

In this chapter the role of data in the wine industry will be deepened by focusing on the customer experience, which is increasingly online-based, and on the scientific research, which is a primary source of innovation in the wine industry. Two increasingly important sources of wine data will be deepened, namely the *user-generated content* in online wine platforms and the *physicochemical analysis* of wine. The implications of exploiting the data generated through these approaches will be addressed.

Finally, a data mining approach to wine will be exemplified through a case study by leveraging two digital representations of the same wine (one in terms of physicochemical analysis, and the other in terms of expert reviews) using real-world data from publicly available datasets and by constructing simple machine learning models.

2.1 Online Wine Customer Experience

2.1.1 Consumer trends and e-commerce growth

Since before the Covid-19 pandemic, the global wine market is undergoing a period of deep transformation, facing several critical challenges of various nature, as highlighted in Section

I.2. The globalization of the wine market, while increasing the growth opportunities for many wines businesses, has also highlighted the difficulty of matching a volatile supply with a fragmented and ever-changing demand (Ismea 2020). Not only are their tastes evolving, but consumers are also more health-conscious and more demanding on sustainability issues, while expecting the same level of digital customer experience as other online services they use every day.

Prior to the pandemic, Wine Intelligence (2020) provided a snapshot of global trends in the wine market from the consumer side:

1) Wine involvement is increasing worldwide but, interestingly, a declining level of wine knowledge is also recorded;
2) Visual cues, relating to both label and packaging design, are increasingly strong determinants of purchase;
3) Wine consumer population is maturing as the largest cohort of wine drinkers is aged 55+, however younger cohorts over-index in terms of volume of wine consumed and total spend on wine;
4) Both the frequency of wine drinking and the average spend on wine have been increasing in the on-premise channel, as consumers seek memorable experiences when dining out; putting it in perspective, this trend is part of a wider wine *premiumization* phenomenon, which sees an overall decline in the frequency of wine consumption coupled with an increased spend per bottle, as demonstrated by the growing shares of premium and ultra-premium[30] wines in total wine sales and exports (See Section I.1);
5) Not only consumers are diversifying their varietal repertoires (especially towards niche red varietals) and putting greater emphasis on the country of origin of wine, but they are also frequently switching out of wine in favor of other beverages (alcoholic or not), making the demand side more and more fragmented;
6) The rise of health-conscious lifestyles and ethical consumerism are driving moderation in alcohol intake and a greater attention towards labelled information, opening the door to organic wines, low-alcohol wines, and to new and more sustainable packaging formats.

[30] According to the classification proposed by Rabobank, price per bottle for Premium wines ranges between 5€ and 7,99€, for Super-premium between 8€ and 13,99€, and for Ultra-premium between 14€ and 49,99€ (Mediobanca 2022).

After the global pandemic, and as a result of it, many of these trends have evolved considerably. First of all, premiumization is likely to split into three paths (IWSR 2021b):

- consumers in developed markets who financially benefited from lockdowns will continue to increase their spending on wine and drive the premiumization trend;
- some consumers will stick to their *tried and trusted* wines and avoid experimenting;
- others will decrease their spending and look for the best *value-for-money* alternatives;

These paths will inevitably increase competition at all price levels and fragment the market even more. Additionally, the pandemic has led to a renowned awareness towards sustainability issues and the importance of well-being, thus promoting the strong performances of organic, low-intervention wines and low-alcohol wines (IWSR 2021a).

The lockdown regime meant a prolonged closure of the on-trade channel and a limited access to the off-trade. This has led to a rediscovery of *at-home* wine consumption and made consumers more willing to try new packaging formats such as *bag-in-box* (larger and convenient for home consumption) and *cans* (smaller and suitable for home delivery) (IWSR 2021a). But the most evident and profound change the pandemic has brought to the wine market (and possibly the most relevant in the long term) is the exponential growth of the *e-commerce* channel and the recourse to digital tools for engaging with consumers (2021a), at a time when traditional shopping patterns were inevitably disrupted.

Overall, the total beverage alcohol ecommerce value reached US$24 billion in 2020 across 16 key global markets,[31] increasing by almost +43% in 2020 (versus +12% in 2019) (IWSR 2021e). Furthermore, the IWSR (2021e) forecasted a growth of total beverage alcohol ecommerce sales of +66% over the next five years to reach US$42 billion, therefore projecting ecommerce to represent a share of 6% of all off-trade volumes in 2025 (compared to less than 2% in 2018). Their research also observed how today one-quarter of total alcohol drinkers across the world reportedly buy alcohol online, with one third of them having made their first purchase during the pandemic.

These numbers become even more relevant considering that wine heavily over indexes online; in fact, while retaining only a 14% value share of the total drinks market, wine retains a 40%

[31] Markets examined represent over 90% of total alcohol e-commerce value and include: Australia, Brazil, Canada, China, Colombia, France, Germany, Italy, Japan, Mexico, Netherlands, Nigeria, South Africa, Spain, the United Kingdom and the United States.

share online, potentially signaling the great value that wine consumers attribute to the possibility of discovering, comparing, reviewing, and rating wines (IWSR 2021a).

In the US, the biggest wine consumer market in the world, the value of alcohol e-commerce grew of over 80% in the period 2019-2020, projecting the US to overtake China as the largest alcohol e-commerce market in the world by the end of 2021 (IWSR 2020). This growth is significant considering that, according to the IWSR (2020), e-commerce is poised to account for 7% of total off-trade beverage alcohol volume in the US by 2024 (versus only 1% in 2019) (IWSR 2020). While still lagging behind with respect to other edible grocery categories online, relaxed legislation during covid times favored online alcohol sales and home deliveries, pushing online alcohol sales up 131% since 2019 (Rabobank 2021). Furthermore, the strategic channels of online grocery and marketplaces grow by 271% over just two years and are now four times larger than they were in 2019 (Rabobank 2021), with the 44% of American alcohol e-shoppers who only started buying alcohol online in 2020 (compared to 19% in 2019) (IWSR 2020).

In Italy, the first wine producer country and the third wine consumer market in the world, 8 million people bought wine online in 2020, representing the 27% of total wine consumers (versus the 17% in 2018) (Nomisma 2021). According to the Mediobanca (2022) survey on the wine sector, the growth of the online channel has been fostered by specialized online players (pure players), which intercept the 90% of the wine e-commerce of main wine producers and which experienced a striking growth of +132,8% in sales over 2019. This growth still kept traction in 2021, with online wine sales growing by +30,4%. However, it must be noted how online pure players sales represent still only 1,3% of total turnover of the Italian wine market in 2021, showing a lower penetration of the channel compared to the US.

Overall, online alcohols sales growth in the US can be split into four main channels (Rabobank 2021):

- *Omnichannel grocery retailers*, which experienced a peak in demand during lockdowns and therefore rapidly increased the number of store locations offering alcohol online;
- *Marketplaces*, which were joined by thousands of retailers seeking to rapidly access the online channel;

- *Licensed specialty retailers*, which historically used to ship hard-to-find wines to elite consumers but which proved to be resilient in serving local markets during covid times;
- *Direct-To-Consumer*, that represents the traditional channel through which wineries sell wine online through mailing lists or via their own website.

After the growth of the online alcohol space and the subsequent consolidation of the channel, the IWSR (2021e; 2022) identified two relevant streams:

- ✓ A traditional e-commerce accessed via *websites*, leveraged by omnichannel or online specialists players, which is typically enjoyed by older consumers seeking value for money and a wide product offering.
- ✓ A more modern e-commerce accessed via *smartphone apps*, which typically consist in marketplaces or on-demand services such as fast home delivery. This stream is favored by younger consumers who seek for premium brands and new interesting wines, and which are willing to pay an extra price for receiving their wine within specified delivery times.

Both streams are significantly relevant across all markets surveyed. Despite some country-specific variations, a general equilibrium can be observed, with a slight preference for websites (favored by 55% of surveyed consumers) (IWSR 2022).

Putting it in perspective, the global pandemic has transformed wine e-commerce from a niche targeting wine enthusiasts to a mainstream phenomenon, with a user base consisting of the 40-50% of the population who like wine and buys it regularly (Wine Intelligence 2021a). Furthermore, the wine e-commerce growth has been successful in recruiting the younger, high spending, and curious consumers which represent the future profitability of the industry.

2.1.2 Online wine platforms

Be they app-based or web-based, online wine platforms such as marketplaces connect consumers with a myriad of wines offered by retailers and wine producers. They make revenues by collecting fees from third-party sellers, by selling data and insights to suppliers, and by promoting brands on the marketplace (Rabobank 2021).

These platforms aim at *matching* users with the perfect wine for them by crafting a user experience aimed at optimizing their decision-making process. For example, the user interface can be dynamically customized for to display the information most relevant to a single user. Increasingly, more and more sophisticated AI algorithms are leveraged to deliver personalized wine recommendations based on user preferences and relevant patterns emerging from the data. Furthermore, users can access thousands of ratings and reviews from other users, thus benefiting from the so-called *wisdom of the crowd*.

Another crucial aspect of online wine platforms which emerged in the past 10 years has been the focus on the creation of wine *communities*, where likeminded users could exchange wine knowledge and support shared values (Wine Intelligence 2021b). This aspect have been promoted in many platforms as users can build their own profile, leave reviews and comments, and discuss in forums within the platform. Finally, ecommerce platforms in China pioneered the *fast delivery* of alcoholic beverages via apps, a phenomenon which is becoming global thanks to the traction gained during the pandemic.

However, these platforms are first and foremost a formidable vehicle for *mobile marketing*, granting the possibility to wine brands to reach pre-qualified consumers at the precise time and location the purchase decision will be made, in both on-premise and off-premise environments (Lotus Growth 2015). Furthermore, in-app marketing campaigns have become extremely efficient and cost-effective thanks to the possibility of collecting data along each point of the conversion funnel and continuously refining consumer segmentation (Lotus Growth 2015).

The game-changing fact about online wine platforms is their mobile nature. Typically, customer experience in an online setting substantially differs from traditional brick-and-mortar shopping for the possibility of getting instant access to a broader range of information, to apply personalized research filters, to compare prices across channels and other platforms as well, and, importantly, to benefit from tailored recommendations and user-generated ratings and reviews. The coexistence of physical and digital sales channels results in the so-called *research shopping*, where consumers gather information about the product in the physical store and then order it online (*showrooming)*, or where consumers research about a product online and then visit a physical store to finalize the purchase (*webrooming*) (Rooderkerk, Kök 2019). The advent of mobile wine applications contributed to blur the lines between showrooming and webrooming, enabling consumers to seamlessly integrate online touchpoints into their shopping experience in the brick-and-mortar store, in what can be defined as "*omnirooming*" (Rooderkerk, Kök 2019). The possibility of delivering additional information at the exact place and time of purchase creates today great opportunities for platforms to engage with consumers and to help them in a time constrained setting where the purchasing experience can be overwhelming, ultimately increasing wine brands' sales.

Lotus Growth (2015) conducted a quantitative research on the sales impact of consumer engagement and wine promotions within the Hello Vino smartphone app. They highlighted how the app can influence purchase behavior during each stage of the buyer lifecycle, in particular (Lotus Growth 2015):

- *Pre-purchase recommendation*. By leveraging users' personal taste preferences, predefined food matching options and contextual information such as specific wine-buying occasions, location and timing, the app can deliver relevant wine recommendations.
- *Engagement & purchase intent*. Upon recommendation, engagement is promoted by offering options such as "Save to Shopping List", "Find Nearby Stores", or directly submit an order request through Hello Vino.
- *Post-purchase feedback*. After purchase, consumer engagement with the recommended brands continues thanks to the possibility of sharing reviews and keeping personal notes about the wine tasted within the app.
- *Brand loyalty & advocacy*. The online app environment enables to deliver post-purchase surveys to gather feedback about consumers' repurchase intention and about whether consumers would recommend the wine purchased to others.

By promoting three wines on the platform and then presenting consumers with a post-campaign survey, the research obtained some interesting results (Lotus Growth 2015):

- 29% of consumers confirmed the purchase of the promoted wine;
- 88% of consumers will purchase the same brand again within the next 6 months;
- 87% of consumers will recommend the promoted brand to others, each to 3.9 people on average;
- 63% of consumers were not aware of the promoted brand prior to engagement in the Hello Vino app;
- 45% of consumers were previously not aware of the promoted brand, but did confirm purchase of the wine.

In order to appreciate the variety of online platforms that populate the market, it is useful to illustrate some examples:

- *Vivino*,[32] with over 50 million users, is the world leading wine platform offering a complete digital wine experience: it is a marketplace, selling wine through third-parties and proprietary inventory, it is a shopping assistant, delivering tailored recommendations and additional information in the brick-and-mortar store, and it is a social network, enabling users to create their own profile, rate and review wines, share wine pictures and leave comments, interacting between each other as a proper community. Vivino enables consumers to take a picture of a wine label and then matches it with its cloud-based database of over 13 million wine labels through image recognition algorithms, thus enabling consumers to access additional information and to compare different wines simultaneously. By accessing the wine page, users get information about average user rating and online price, overall taste profile and specific tasting notes, and get some highlights about convenience and quality (ranking, awards), coupled with synthetic wine "facts", such as region, regional style, grapes, alcoholic content and producer. Additionally, it is possible to take a picture of a wine list to access information about all wines in a restaurant cart.
- *Naked Wines*[33] is a UK based online platform which supports independent winemakers by providing a marketplace for Direct-To-Consumer selling. By granting micro-wineries access to its proprietary distribution chain, Naked Wines can offer a wide range of high quality and novel, extravagant wines. Furthermore, by cutting out the "middle-man" from the distribution chain, they also claim to offer wines at a considerably lower price. The platform promotes user engagement thanks to the possibility of directly interacting with winemakers and reviewing wines.
- *Drizly*[34] is a US based online platform that relies on a capillary network of local stores to deliver alcohol beverages to customers at home within 60 minutes. Customers only need to enter their address and see the available stores in their area, choose some items within the local assortment of products, and place an order.

While Vivino behaves as shopping assistants, by providing customers with a digitally-enhanced tailored experience, Drizly deploys a locally-deployed strategy by connecting local stores with customers by digital means, and Naked Wines targets a quality-oriented and

[32] See https://www.vivino.com/
[33] See https://www.nakedwines.com/
[34] See https://drizly.com/

sustainability-sensitive segment by making the smaller, independent producers benefit from the shortened distances of the digital world.

All these platforms digitally-enhance the wine customer experience in their own way, and experienced terrific growth in 2020:

- Vivino experienced a global year on year sales growth of 103%, reaching US$265m (Thach 2021);
- Naked Wines grew its sales by 68% to reach more than US$300m dollar of revenues (Sweney 2021);
- Drizly reported a 350% growth in sales (Furnari 2021).

Furthermore, a number of important acquisitions took place within the wine e-commerce space: Drizly was acquired by Uber, *UVINUM* and *Bodeboca* have been acquired by the Pernord Ricard group, and *Tannico* in Italy have been acquired by the Campari group (ISMEA 2020). Additionally, Vivino raised US$155m of funding for to fuel its expansion and to build its new AI recommendation engine (Thach 2021). All these M&A and investment activity altogether signal that the wine platforms are here to stay and are poised to grow further in the next years, possibly setting the standard for the new customer experience in wine omnichannel shopping.

2.1.3 Extrinsic wine attributes and purchasing behavior

In order to improve wine consumer decision making, driving sales and customer satisfaction, it is necessary to gain a deep understanding of how wine attributes influence purchasing behavior. Indeed, wine is deemed as a complex product not only due to its complex sensory profile, but also for the numerous attributes the consumer needs to process upon purchase.

Since in most brick-and-mortar shopping occasions wine cannot be tasted before purchase, wine consumers typically elaborate on extrinsic product attributes to infer wine quality (Sàenz-Navajas et al. 2014), with most of these quality cues being carried by the label and the packaging. In fact, contrarily than intrinsic cues which refer to the organoleptic properties of the wine and the relative sensory perceptions that are experienced upon tasting, extrinsic cues refer typically to attributes such as vintage, brand, country-of-origin, region, appellation, price, awards, label design, bottle weight, and others.

Among them, the attributes related with the place of origin, including the presence of a *protected designation of origin (PDO)*, have always received particular attention by both the industry and the consumers. In fact, wine is a traditional product whose history has deep roots in the oldest wine producing countries and regions. This not only leads the expectation that a wine from a certain region will carry specific and unique sensory properties, but also evokes in consumers the feeling that by tasting the wine they are experiencing the values and the history of a place. This explains why consumers may use PDOs as an indication of quality and why many wine producing countries have leveraged them as a mean for *regional branding*.

A particularly relevant role in conveying wine information to the consumer is played by the *back-label,* which may display statements such as flavor descriptions, winery history, production method description, and food matching suggestions (Mueller et al. 2010). Together with the *front label*, which typically displays the legally required information about wine (winery's name, grape variety, grape origin, vintage year and alcohol content), they represent the most immediate and cost-effective marketing vehicle for wine producers to communicate with the final consumer at the point of purchase and to influence their wine choices (Mueller et al. 2010).

However, other factors such as having previously tasted the wine or the wine being recommended by someone trusted can overshadow all the previously mentioned cues, as it has been suggested in the literature (Goodman 2009). Furthermore, the way wine consumers elaborate wine attributes depends on nationality and sociodemographic factors, as well as their wine involvement and wine knowledge level (Sàenz-Navajas et al. 2014).

It has been of great interest for both researchers and the industry to determine which product cues have the greatest impact on wine consumers' purchasing decisions, and a large body of literature on this topic has been produced over the last 30 years.

Notable examples of quantitative techniques applied to wine consumer research are *Discrete Choice Experiment* (DCE) and *Best-Worst Scaling* (BWS) (Lockshin & Corsi 2019). A DCE combines product features into sets of competing alternatives (for example wine labels displaying different wine attributes). Then, the respondents have to state the preference by choosing the preferred alternative (or none of them). At the end of the experiment, the utility of each feature is estimated and thus its contribution to the choice decision. BWS works in a similar fashion, but the respondents are asked to indicate which feature influenced their choice the most (the “best”) and which one influenced it the less (the “worst”), for every choice set. The effectiveness of this techniques lies in resembling a real-world purchasing decision, with consumers forced to trade-off between different product attributes. The results of some worth mentioning wine consumer studies are reported below.

Lockshin et al. (2006) measured how key extrinsic cues such as brand, region, price and awards influence consumers’ purchasing decision. The results showed how low involvement consumers tend to rely on price and awards, with a gold medal being the attribute increasing the probability of choice the most at the lower and middle price points. Additionally, a well-known region was shown to increase the desirability of small brands. Mueller et al. (2010) focused instead on the influence of wine back label information on wine choice, finding that a combination of winery history, elaborate taste descriptions and food pairing had the most positive impact, while chemical ingredients list caused a largely negative impact.

Goodman (2009) gave a monumental contribution to wine consumer research by carrying out an international comparison of retail consumer wine choice through a BWS experiment. He found that previous trial and recommendation were the key attributes influencing wine choice across most markets, followed by grape variety and origin of the wine. However, some interesting country-specific influencers emerged, such as brand in Brazil and China, food matching in Italy and France, origin of the wine in France and grape variety in Austria.

Although wine consumer research through *Stated Preferences* data (such as DCE and BWS techniques) is now a relatively mature field, novel questions arise as consumer trends evolve. Capitello et al. (2021) investigated how environmental and terroir cues are processed by young consumers, as both ethical consumerism and the demand for authentic products are on the rise. In particular, they conducted a DCE for analyzing the effect of a carbon claim in conjunction with terroir cues. While PDO emerged as young consumers’ preferred attribute in wine choice, they observed that the carbon claim had a positive impact on utility for the 46% of the sample. Furthermore, they identified some strategic patterns through which wineries can leverage environmental and terroir attributes, such as (Capitello et al. 2021):

- enhancing the modernity of a wine;
- connecting the origin of a wine with typical production specificities;
- echoing product quality for less-known PDOs.

Despite this field of research gave a great contribution to understanding wine consumers’ purchasing behavior, Lockshin & Corsi (2012) criticized the predominance of convenience-based samples, which are not statistically representative of the population of wine drinkers, thus undermining the possibility of deriving generalizable results.

Another field of wine research leverages *Revealed Preference* data, such as *panel data*. In this case it is possible not only to observe the actual consumer choices, instead of hypothetical ones, but it is possible also to link each purchase to a specific individual (Lockshin & Corsi 2019). Panel data enables to conduct insightful analyses with regards to price, brand performance, consumer segmentation, promotions effectiveness and so on. Typically, this kind of data has been the prerogative of large wine retailers and big market research firms, making it difficult to access and very expensive to purchase.

Today, with the growth of online wine platforms, such data is automatically collected upon every transaction at an unprecedented speed and scale. As these platforms continuously gather data about wines and consumers, not only they can profit by selling them along the wine supply chain, but they also represent an ideal platform for rapid experimentation, generating precious insights that are relevant both for the optimization of the platform’s performance and for the industry overall.

Furthermore, the flourishing of online wine content enabled the proliferation of new sources of wine data, and opened up to new ways of investigating wine consumers’ preferences.

2.1.4 Expert reviews and user-generated content

Before wine platforms and apps emerged, wine content was already proliferating online through online magazines, websites, forums, blogs, etc. In an omnichannel retail world and within the context of “research shopping”, all these information sources are readily available to consumers along their customer journeys and may play a role in shaping their opinion of the product and ultimately influencing their purchasing decision (Nosi, Mattiacci & Sfodera 2019).

Since these information sources are typically owned and managed by third parties, Nosi, Mattiacci & Sfodera (2019) suggest that wine producers and wine brands should pay increasing attention to the *digital narrative* of their wines. To explain the point, they conducted a research to investigate how a specific grape variety (Sangiovese) is narrated online by non-winery-owned websites across four countries (including Australia, Canada, the UK and the USA). They observed that the majority of online content relates to technical aspects of the varietal and adopts a professional slant, with only marginal references to culture, history, and folklore, which are elements that could be more effective for engaging with common wine drinkers (Nosi, Mattiacci & Sfodera 2019). Furthermore, they observed that how the digital wine ecosystem in highly concentrated, with few sources acting as gatekeepers of wine information.

As wine magazines and wine guides turned to the electronic format, they contributed to generate a large volume of online publicly available information about wine, focusing in particular on the sensory description of wines. This represents an invaluable source of unstructured data that can be processed to extract useful knowledge about wines. This *text mining* process usually involves a preliminary phase of text pre-processing and standardization, and a subsequent phase where concepts and relationships within the text are structured by applying proper algorithms; furthermore, visualization techniques may be employed to improve human interpretation of the data (Valente et al. 2018). The results of these techniques are useful to test hypotheses or to generate novel ones, and to find hidden relationships in the data which can generate insights relevant to a particular domain, where a certain level of expertise is needed.

Valente et al. (2018) analyzed the sensory characteristics of 7,000 Chenin blanc and Sauvignon blanc wines by leveraging the publicly available data from the *John Platter Wine Guide to South African Wines*. This guide contains over 15,000 wine sensory reviews of South African wines redacted by appointed experts. Valente et al. (2018) applied a data visualization technique known as *formal concept lattices* to observe the *sensory space* of wines, successfully identifying aroma sensory attributes that are unique to or strongly associated with each varietal (with *Gunpowder*, *Capsicum* and *Gooseberry* associated with Sauvignon blanc, and *Biscuit* and *Vanilla* with Chenin blanc). Furthermore, they identified *Complexity* as an important driver of style for Chenin blanc wines.

In a seminal paper, Chen et al. (2014) coined the term "*Wine Informatics*" to refer to the growing area of wine research which leverages wine sensory reviews as the domain knowledge. In their work they developed the "*Computational Wine Wheel"*, a method for automatically retrieving sensory notes in a structured way from unstructured wine reviews, closely inspired by the famous Ann C. Noble's *Wine Aroma Wheel*[35], a famous tool for helping wine practitioners detect wine sensory notes during tastings. In particular, they

[35] See: https://www.winearomawheel.com/

leveraged online reviews from *Wine Spectator*,[36] as their compact and precise review style, fully-focused on tasting notes, provides consistent data for research at scale (with over 400,000 wines reviewed). Dong et al. (2020) successfully applied this method to study a dataset of over 14,000 Bordeaux wines from 2000 to 2016. In particular, they used classification algorithms to build a predictive model for classifying outstanding wines (scored 90+) based on wine reviews, and successfully identified the distinct tasting notes for this category (*Long*, *Black Currant*, *Apple*, *Fig*).

The most relevant and fastest growing online source of wine data is clearly represented today by user-generated content in the form of ratings and reviews on online wine platforms. Vivino alone has gathered almost 90 million reviews and over 250 million ratings as of October 2022.[37] Kotonya, De Cristofaro & De Cristofaro (2018) conducted a thorough analysis of the Vivino platform, getting some precious results:

1) Vivino's user-generated ratings and reviews display the same rich knowledge of wines as professional wine reviews, while not being biased by wine prices.
2) Vivino's users tend to prefer and rate local wines, with some strong geographical similarities in how wines from adjacent countries are rated;
3) Vivino's user-generated data can be successfully leveraged to build models to predict wine quality and provide personalized wine recommendations.

These results are extremely relevant in two ways. Firstly, point 2) suggests that, by aggregating massive amounts of data, emerging local clusters may be efficiently targeted and serviced. Secondly, user-generated content is a formidable resource to enrich the information available to recommender systems to generate personalized recommendations (for example by enriching items representations and user profiles, in the case of a content-based recommender, or by leveraging patterns in the ratings in the case of a collaborative filtering recommender). This aspect becomes even more relevant considering that there is no actual difference between perceptual capacities of wine experts and novices (Smith 2019). What actually differs is the level of wine expertise and the consistency of their vocabulary. However, by employing appropriate text mining techniques at scale, structured knowledge can still be generated through user reviews.[38]

[36] See https://www.winespectator.com/
[37] See https://www.vivino.com/about
[38] In fact, Vivino proved the consistency of its users' judgement by analyzing over 100,000 expert ratings and establishing a relation between Vivino ratings and the most influential expert ratings, finding that a 4.0 Vivino

2.2 Wine's Chemical Fingerprint

2.2.1 Wine chemistry and flavor

From a chemical point of view, wines are a complex beverage with plenty of compounds that significantly contribute to sensory properties and at the same time impact stability and affect product safety (Waterhouse et al. 2016). Even though the quality of wine is typically framed in terms of prestige, provenance, style, vintage, age, varietal origin, or other attributes, its legitimate quality lies in the sensory characteristics that are derived from its chemical composition (Jackson 2017).

Unlike commodities where producers seek for homogeneity, in specialty products such as wine variation from standards is an appreciated and celebrated aspect. Consumers expect that wines with different labels have a different smell, a different taste, and a different look; from a chemist's point of view, wines are expected to have different chemical compositions (Waterhouse et al. 2016). The study of wine chemistry focuses on these differences – to explain how there can be hundreds of thousands of different wine compositions, and how the myriad of choices that a winemaker faces can lead to these differences (Waterhouse et al. 2016). The number of chemical compounds in wine is basically uncountable. The aim of a wine chemist is not to enumerate every compound, but rather to identify compounds or classes of compounds, that will directly or indirectly determine key quality aspects of the wine, such as organoleptic properties (aroma, appearance, flavor,), stability, and safety (Waterhouse et al. 2016). In fact, the identification of compounds that contribute to wine aroma and taste is an extremely important aspect in understanding the overall quality and consumer acceptance (Pérez-Jiménez et al. 2021).

Generally speaking, a dry table wine is a mildly acidic (pH3–4) hydroalcoholic solution where water and ethanol are the two major components, typically accounting for about 97% on a weight-for-weight (w/w) basis. The remaining compounds, responsible for most of the flavor and color of wine, are typically present at <10g/L, and many key odorants are found at part-per-trillion (ng/L) concentrations (Waterhouse et al. 2016). For example in Sauvignon white wines, varietal thiols (3-mercaptohexanol and 3-mercaptohexyl acetate) that contribute

rating correlates with a 90 point expert rating. See https://www.vivino.com/wine-news/vivino-5-star-rating-system towards passionfruit and grapefruit aroma respectively have a perception threshold of < 70 ng/L (Pérez-Jiménez et al. 2021).

Wine is produced by the fermentation of grape juice or must (juice and solids), a chemical process where the complete or partial transformation of grape sugars to ethanol and CO_2 take place. However, during the winemaking process and wine storage many other chemical changes occur. This is readily exemplified by the volatile composition of a wine, which is far more complex than that of grape juice.

These volatile components determine the aroma of wine and are often classified based on when they are formed (Waterhouse et al. 2016):

- *Primary Compounds*, which are present in the grape and persist unchanged into wine.
- *Secondary Compounds*, which are formed during alcoholic or malolactic fermentation due to either:
    a) Normal metabolism of sugars, amino acids, etc.
    b) Transformation of grape-specific precursors.
- *Tertiary Compounds* formed during wine storage, for example, as a result of
    a) Extraction from oak.
    b) Microbial spoilage or chemical tainting.
    c) Abiotic transformation of precursor compounds in wine.

Wine production includes the techniques and technologies used in the transformation of grapes into wines of a target style (Waterhouse et al. 2016).

During the process of *vinification*, where grape sugars are converted into alcohol by yeast, at the same time, a variety of other important chemical changes occur due to a multitude of other grape components. Many red wines also undergo a secondary fermentation which is called malolactic fermentation (MLF), where malic acid is converted into lactic acid. Production techniques differ depending on which grapes (red or white) are used and on which type of wine is being produced. For example, when producing white wines the grape solids (like skin and seeds) are removed soon after, while on the contrary, red wines undergo the maceration and fermentation steps in contact with grape solids. Through these steps, many polyphenols are extracted, conferring to red wine its distinctive color and organoleptic properties. In addition, red wine styles typically go through a maturation process in tank or barrels, where color is stabilized and mouthfeel properties are refined. Other common winemaking processes aim at stabilizing, fining and filtering the wine in order to ensure that the product has the desired sensory profile and it is stable from a microbiological point of view.

The complexity of wine chemistry is reflected in the complex sensory perceptions that humans can perceive by drinking a glass of wine. Hundreds, if not thousands, of chemical compounds are responsible for aroma (ortho and retronasal), taste, aftertaste and mouthfeel sensations (Pérez-Jiménez et al. 2021).

Flavor is defined as the perception resulting from the combined stimulation of taste buds, olfactory organs, and chemesthetic receptors within the oral cavity.

In particular (Waterhouse et al. 2016):

- *Olfaction*, or smell, consists in the detection of odorants by olfactory receptors (ORs) which are located within the nasal cavity. Their interactions with wine's volatile compounds induce signals interpreted by the human brain as *aromas* (Pérez-Jiménez et al. 2021). Odorant volatiles reach the nasal cavity through two routes:
    - *Orthonasal olfaction* consists in the detection of odorants without tasting, for example, by smelling the headspace of the wine.
    - *Retronasal olfaction* consists in the detection of odorants that travel from the oral cavity to the nasal cavity. For example, this occurs after swallowing.
- *Taste* involves the detection of small molecules by taste receptors situated in the taste buds. Taste receptors have been grouped in five classes – "sweet," "sour," "bitter," "salty," and "umami" –, but only the first three appear to be routinely experienced in wine.
- *Chemesthesis* involves the chemical activation of those receptors that are responsible for sensations of pain, temperature, and touch. Typically related to wine is astringency, that is the perceived loss of lubrication in the mouth which can be triggered by condensed tannins and other phenolic compounds.

Those perceptions are elicited through the complex interactions occurring within the so-called *wine matrix*, that is composed by two fractions (Pérez-Jiménez et al. 2021):

- *Volatile aroma compounds* such as thiols, methoxypyrazines, some terpenes and volatile phenols typically attribute very specific aromas to wines. A lot of other volatile compounds, such as esters, higher alcohols and other terpenes, have been demonstrated to contribute to the complexity of wine aroma rather than attributing specific, readily identifiable aromas.

- *Non-volatile or semi-volatile compounds* contribute mainly to mouthfeel sensations and tastes, as well as wine color. Sugars, sugar alcohols and tannins have been linked to wine mouthfeel while organic acids have been shown to impart a sour taste. In addition, alcohols can influence the perception of aromas and tastes, and are responsible for the hot sensation in mouth during wine tasting. As with wine volatile metabolites, non-volatile compounds in wine can come from the grapes, yeast, and reactions during winemaking and ageing.

Flavor *perception* is not only affected by a few active compounds, instead, there are many others that are not flavor active but are able to enhance or suppress the perception of flavor and flavor active compounds. Typically, *activity values* for a particular compound are calculated as the ratio of a compound's concentration to its sensory threshold in an appropriate matrix (Waterhouse et al. 2016):

Activity value = Concentration / Detection threshold

However, despite useful for an initial screening, an activity value >1 is insufficient to determine if the compound is important in the overall perception of flavor.

In fact, some relevant effects occur between matrix compounds (Waterhouse et al. 2016):

- *Masking*. The perceived intensity of a flavor can be decreased by the presence of other flavor compounds.
- *Additive or synergistic effects*. Groups of homogeneous compounds, that is, series of alkyl esters or ketones, can reach sensory threshold through additive effects even if all compounds individually have activity values <1. Another possible effect is called "synergism", that is an increase in stimulus intensity beyond what is predicted for simple additive effects.

Besides the physicochemical interactions occurring between wine matrix compounds, the human aspect of flavor perception also needs to be considered to truly understand the molecular basis of wine flavor (Pérez-Jiménez et al. 2021). Flavor is created through the interaction of stimuli with multiple sensory, motor and central behavioral systems by the brain. Therefore, it is a complex task to understand the relationships between stimuli and flavor perception as psychological constructs are different from individual to individual.

Besides a few innate tendencies, such as a dislike of bitter substances and a liking for sweetness, human sensory responses seem to be primarily based on experience, not reflex. Therefore, flavor preferences are potentially malleable and primarily culturally based. Knowledge and experience appear to be the predominant factors. For example, familiarity abets odor discrimination, and often enhances perceived intensity and pleasantness, and repeated exposure often modifies preferences, promoting acceptance (Jackson 2017).

In addition, physiological factors including breathing, saliva, and oral microbiota contribute to the ultimate flavor perception. This explains why the same compound or group of compounds that induces a pleasant aroma for some may be perceived as very unpleasant by others, or may not be detected at all. Careful consideration of both wine chemistry and human perception is paramount in understanding the complex concept that is wine flavor (Pérez-Jiménez et al. 2021).

The full spectrum of interactions above described gives rise to the concept of quality. Wine quality can be viewed from many angles, but fundamentally, quality depends on the wine's physicochemistry, and how it is detected and processed sensorially (Jackson 2017).

However, intrinsic attributes are mediated by the role of extrinsic attributes that consumers value, such as authenticity and sustainability, monetary value and the socio-cultural context of consumption.

A wine's quality is also intimately linked with flavors that may donate distinctive stylistic, varietal, and possibly regional attributes. (Jackson 2017). Intrinsic quality is driven by grape cultivars and grape quality, production style. But the most critical factor in the development of a wine's quality is the winemaker, as it is ultimately on their decisions that a wine will evolve, and the attributes it will eventually possess (Jackson 2017).

Finally, it has be considered that grapegrowing techniques and winemaking practices are not static, and typical values may change dramatically with changes in fashion, technology, or climate. A modern consumer's expectation of a high- quality wine is the result of both improved technical capacities as well as accumulated traditions (Waterhouse et al. 2016).

2.2.2 Fingerprinting and Flavoromics

The influence of variety, soil and climatic conditions on the taste of wine is referred to as the "*terroir effect*". While several organoleptic characteristics of wine are influenced, wine typicity in relation to terroir is largely shaped by volatile compounds (van Leeuwen et al. 2020).

The existence of a sensory space that can be unique to individual wine types, such as wines from specific varieties or geographical locations, has been largely discussed in wine research, confirming the existence of chemical signatures associated with the area of production of the wines (Slaghenaufi et al. 2021).

Van Leeuwen et al. (2020), by means of a literature review, deconstructed the effect of measurable soil and climate parameters on grape and wine aroma compounds, in order to explicit the terroir effect on wine typicity. In fact, they individuated 4 key factors for the expression of aromatic typicity:

- *Air temperature*
- *Radiation*
- *Vine nitrogen status*
- *Vine water status*

However, the overall terroir effect is expressed through the grapevine variety (Van Leeuwen et al. 2020). As an example, Sauvignon blanc aroma in cool climates such as France and New Zealand is shaped by a balance of green and peppery aromas and fruity aromas, such as grapefruit and passion fruit. In very cool climates instead, aromas of asparagus and bell pepper may dominate. Warm climate Sauvignon blanc, such as in many regions of USA or Australia, is dominated by passion fruit aroma. Severe water deficits and low vine nitrogen status exert a negative impact on the aromatic typicity of Sauvignon blanc wines, while high radiation has a positive effect, favoring its fruity aromas.

In a seminal work, King et al. (2014) compared the chemical compositions and sensory profiles of Malbec wines from Mendoza, Argentina and California, USA. They found that sensory profiles significantly differed between the two countries, with the Argentinian Malbec displaying more ripe fruit aromas and a sweeter taste and the Californian Malbec showing artificial fruit and citrus aromas, with a bitter taste. However, they correlated these differences in the sensory and chemical composition of wine mainly with the different altitude at which the vines were grown.

Recently, Slaghenaufi et al. (2021) investigated the chemical composition and sensory profile of Lugana and Verdicchio, two Italian white wines from the same grape variety but produced in two different regions. They successfully found some chemical markers on the volatile composition of wines that helped in discriminate between two wine clusters, the first composed only of Lugana wines and described by “fruity” and “minty” notes, and the second composed mainly of Lugana wines described by “fermentative” and “spicy” notes.

These studies, among others, confirm the primary importance of geographical origin on the volatile composition of wines (Slaghenaufi et al. 2021).

During its transformation from the vineyard to the glass, the wine develops its complex *metabolome*, that is the set of small-molecules chemicals found within a biological sample (Cozzolino 2016). Such molecules, and the various interactions that occur between them, are potentially responsible for shaping the sensory profile of a wine, however only a limited amount of them has been identified (Cozzolino 2016).

It became of great interest for wine researchers and the wine industry to develop new methods for analyzing the wine metabolome to uncover the complexity of wine chemistry, and how the latter shapes its sensory profile.

The major drawback of traditional wine chemical analysis was that it was conducted in a targeted manner, aiming at identifying only a limited number of compounds with accuracy and sensitivity, but allowing to detect only metabolites that were present in high concentration (Pinu 2018).

Today, *untargeted metabolomics* methods enable researchers to detect thousands of metabolites in a single run, and to discover new compounds that were not expected to be in the sample. While this comes at the cost of precision, employing both methods in the analysis of the same wine enables to get a comprehensive picture of its chemistry.

In a volatile untargeted metabolomic experiment, the components present in the sample are separated in gas phase through *Gas chromatography (GC)*, and the ionized molecules are detected by a *Mass spectrometer (MS)* which delivers the resulting spectra of the wine sample and enables to perform quantitative and qualitative analysis (Waterhouse et al. 2016).

Such spectra of a wine sample represent a *chemical fingerprint*, which is the metabolome as it has been shaped by vine growing and winemaking factors (Amargianitaki & Spyros 2017), telling the whole history of a wine.

Pérez-Jiménez et al. (2021) thoroughly presented wine *flavoromics* as an emerging field of research that leverages volatile untargeted approaches and data-driven approaches to establish relationships between the volatile metabolome and the sensory properties of a wine.

In particular, these experiments are conducted by combining untargeted volatile profiling with human descriptive sensory analysis carried out by trained experts, capable of quantifying aromas and tastes in a consistent manner. Then, analytical approaches are used in order to establish relationships between compounds and specific sensory attributes in the wines.

By shedding a light into the complex processes that shape flavor at a molecular level, flavoromics studies represent a fundamental step towards precision enology, supporting wine producers in ensuring wine quality, and towards the emerging field called *personalized wine,* where trends in consumer preferences from different markets are considered while producing wine (Pérez-Jiménez et al. 2021).

2.3 Case Study: Data Mining from Vinho Verde Wine

Classification and regression tasks are the fundamental tasks of machine learning. Many real-world problems can be reframed in terms of prediction by assigning an object to a class or by predicting a continuous value (such as a rating). Even recommender systems, which are employed to model customer preferences and to suggest them items that they may like, are just a more complex system based on simple classification and regression tasks, that is whether the customer will like the suggested item or not, and how much he will like it. In this brief case study, different machine learning models will be employed to predict the quality of wine from physicochemical analysis and from expert reviews by analyzing publicly-available datasets.

The wine considered is *Vinho Verde*, a popular Portuguese wine from the Minho region in the northwest of Portugal. The term "verde" stands for green as it is consumed young, and it is appreciated for its freshness, especially in the summer. Characterized by fruity aromas, it may sometimes present a subtle fizziness. It is produced in white, red and even rosé variants, being the white the most common style and the most exported one.

2.3.1 Data mining from physicochemical analysis

Cortez et al. (2009) proposed first a data mining approach to predict human wine taste preference by relying on standard analytical tests which are typically employed at the certification step. The dataset, freely accessible through the UCI Machine Learning Repository, is composed by 4898 samples of white wines and 11 physicochemical attributes and one quality attribute. Each sampled wine was protected by a PDO and was tested by the official certification entity.

After the physicochemical analysis, a sensory test was conducted by three experts through blind tasting, thus assigning a score from 0 (very bad) to 10 (excellent), and the final score was calculated as the median of the three. This last attribute, referred to as “quality”, will be the target class for the prediction task, since the objective of the model is to predict the quality of wine from its physicochemical properties.

By first looking at the distribution of the scores in Fig. 14 by plotting a bar chart, it is possible to observe a normal shape distribution, with the highest number of wines belonging to the class associated with a 6 score. Such a distribution is typical of real-world data and it was expected in the contest of quality evaluation, where wine is tested not only for to assure safety and quality but also with the goal to stratify prices.

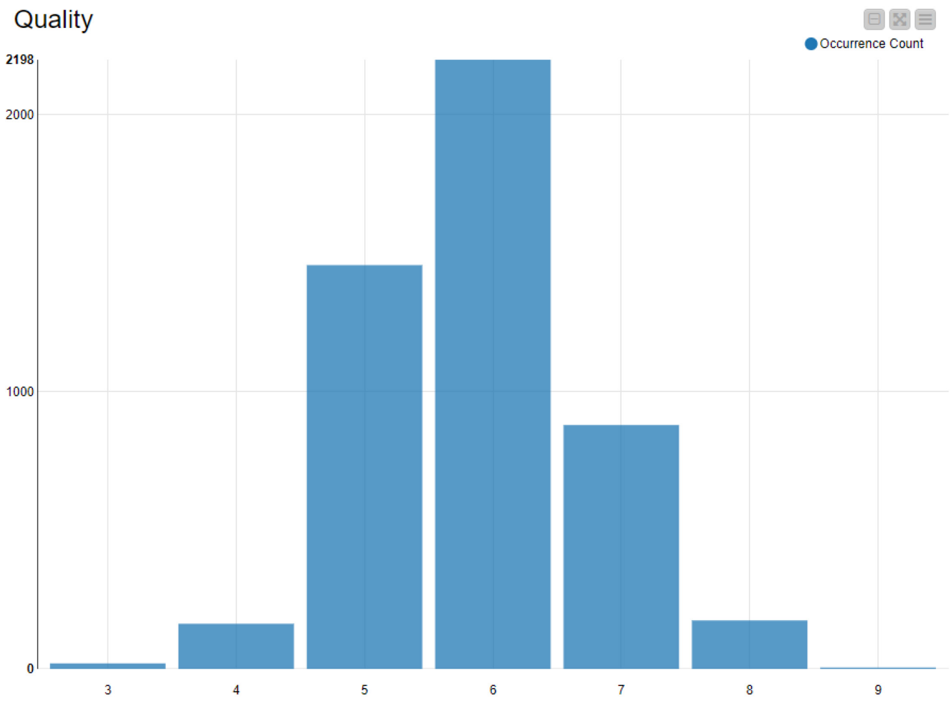

Figure 14. Bar chart of quality scores. Source: own elaboration.

As a first step, it could be useful to look at the correlation matrix between the attributes (Fig. 15), in order to see which attributes may show a correlation with the target attribute quality.

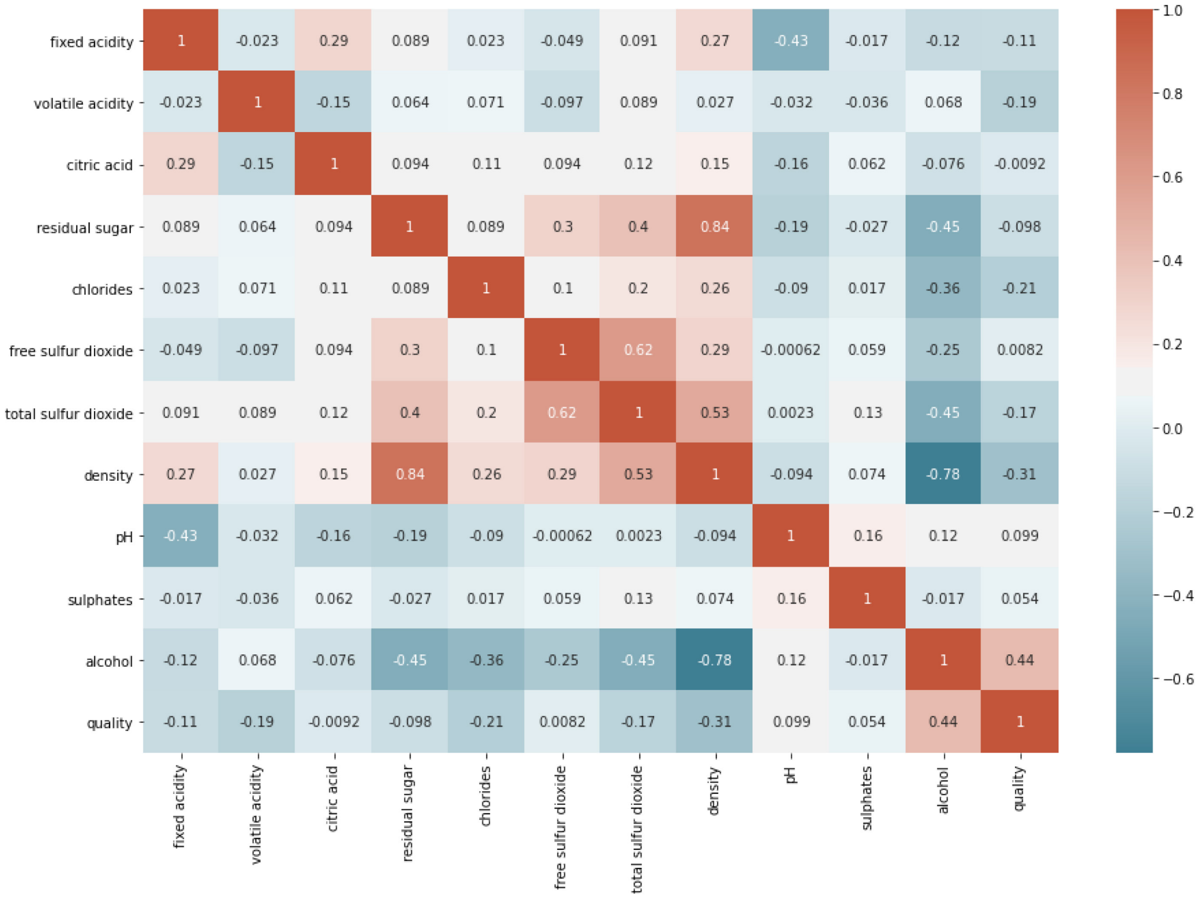

Figure 15. Attributes correlation matrix. Source: own elaboration.

The most evident positive correlation with quality is alcohol, which is consistent with the long-standing assumption that wine with higher alcoholic degree is more likely to get higher scores from the critics (citation needed). Density instead shows a negative correlation; density is primarily correlated with residual sugar, that is the sugar that is left after the fermentation, and with total sulfur dioxide, which is the sulfur dioxide both naturally generated during wine fermentation and added as an oenological additive for to protect wine during the fermentation from oxidation and to inhibit microbial activity. In fact, these two attributes are also slightly negatively correlated with quality. Finally, volatile acidity and chlorides are also negatively correlated with quality, the first being associated with a vinegary odor when in high quantities, and the second conferring a salty taste to wine when reaching a certain threshold.

In order to try to predict wine quality from these physicochemical attributes, some simple machine learning models will be employed. For this purpose, KNIME open-source analytics platforms will be used. The software enables to automate machine learning models through workflows composed by interconnected nodes.

Given the unbalanced dataset, it is useful as a first attempt to convert the problem of predicting quality into a binary classification task. Therefore, two categories will be created, the first one labeled as "good" (1.060 wines) which comprises wines with a score higher than 6, and the "bad" (3.838 wines) class which comprises the remaining wines.

The first model employed is the *Decision Tree learning*, which is a classic supervised learning algorithm particularly suitable for binary classification tasks. The algorithm partitions the dataset into subsets until each partition is "pure", that means it is composed of only samples of the same class. The partitioning is made at each node of the tree by a rule that separates the data based on the value of a given feature. The workflow is shown in *Figure 16*.

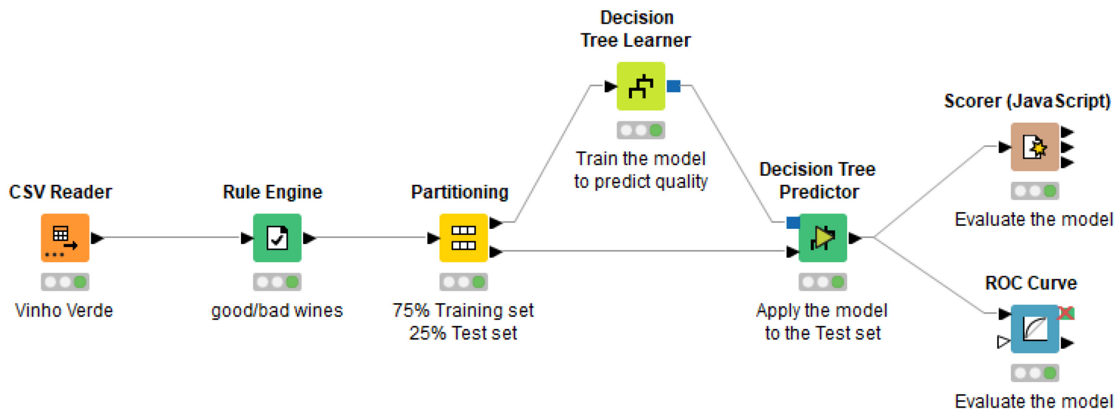

Figure 16. Workflow for a Decision Tree classification task. Source: own elaboration.

The "Partitioning" node is used to split the dataset into a training set (75% of the samples) and a test set (25% of the sample). The training set is sent to the Decision Tree learner node which learns the model and the applies it to the test set. This procedures aims at ensuring that the model learned from the algorithm in the training set can generalize on new, unseen data.

Given the white-box nature of the Decision Tree algorithm, it is possible to extract the model as a graphical representation (Figure 17).

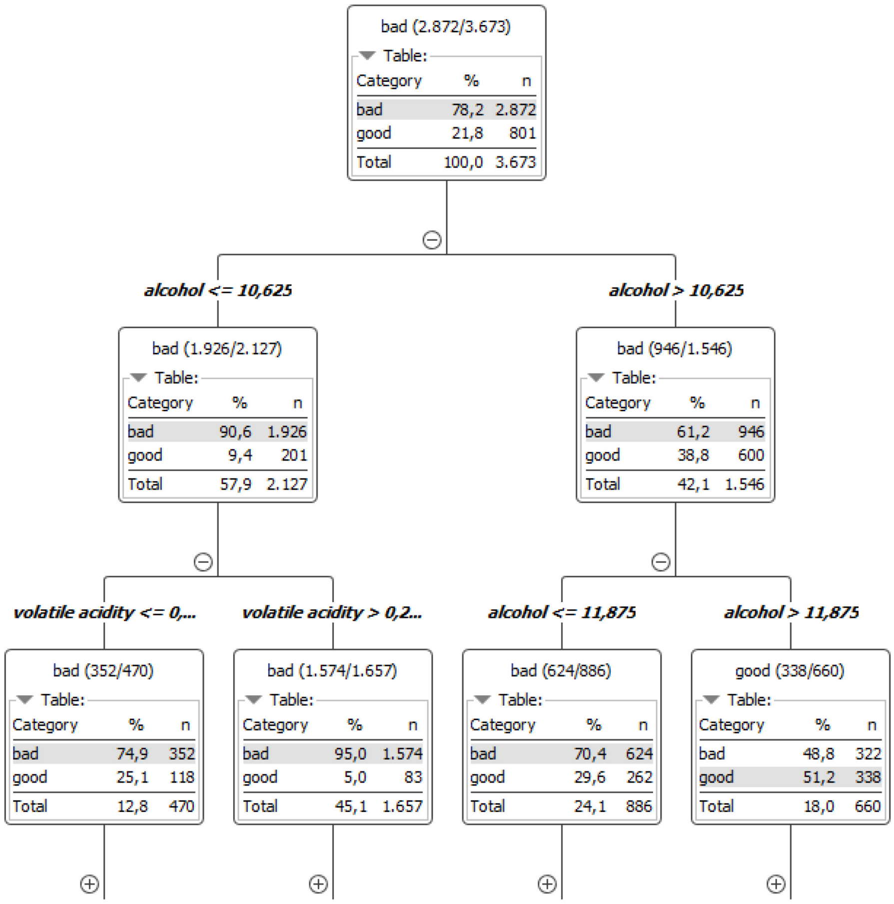

Figure 17. Decision Tree graphical representation. Source: own elaboration.

As it can be observed, the split criterion on the top nodes are based on alcohol and volatile acidity values. In particular, 90,6% of wines with an alcohol degree lower than 10,625 are classified as “bad wines”, and the 95% of such wines are classified as “bad” when reaching a volatile acidity higher than 0,205. This result is consistent with oenological theory and tells that the model identified the majority of the bad wines (1.574 out of 2.872 bad wines) through a lower alcohol content and a higher volatile acidity.

Once the model has run and produced a prediction output on the test set, the model can be evaluated on the test (Fig. 18).

Decision Tree - Scorer View

Confusion Matrix

	good (Predicted)	bad (Predicted)	
good (Actual)	154	105	59.46%
bad (Actual)	92	874	90.48%
	62.60%	89.27%	

Class Statistics

Class	True Positives	False Positives	True Negatives	False Negatives	Recall	Precision	Sensitivity	Specificity	F-measure
good	154	92	874	105	59.46%	62.60%	59.46%	90.48%	60.99%
bad	874	105	154	92	90.48%	89.27%	90.48%	59.46%	89.87%

Overall Statistics

Overall Accuracy	Overall Error	Cohen's kappa (κ)	Correctly Classified	Incorrectly Classified
83.92%	16.08%	0.509	1028	197

Figure 18. Decision Tree scorer. Source: own elaboration.

The overall accuracy simply gives the proportion of correct classifications over the total, which amounts at 83.92%. But given the imbalanced dataset, this value is pushed by the high number of bad wines which the model correctly classified. In fact, by taking the "good" as the reference class and looking at its statistics a large number of False Positives (92) and False Negatives (105) occur with respect to True Positives (154).

Since the major goal of the model was to correctly classify the good wines, other metrics are more relevant for these purpose:

- *Precision*, which is the proportion of True Positives correctly categorized out of all the True Positives and False Positives;
- *Recall*, which is the proportion of True Positives correctly categorized out of all the True Positives and False Negatives.

There is a tradeoff between Precision and Recall; the higher the Precision, the less False Positives the model classifies, but at the same time some True Positives may be overlooked. The higher the Recall instead, the more probability there is that the model will correctly classify all the True Positives, but in doing so it may capture some False Positives.

Since both metrics are unsatisfactory (Recall = 59.46%, precision = 62.60% for the "good" class), another algorithm will be employed to try to improve the results.

The *Random Forest* algorithm employs an ensemble of Decision Trees, each one trained in random subsets of the training data, and merges them into a unique output through a majority vote rule. By aggregating many Decision Trees the Random Forest typically outperforms the former. Therefore, a workflow analogous to the first one was built in order to train the model. The results are shown in *Figure 19*.

Random Forest - Scorer View

Confusion Matrix

	good (Predicted)	bad (Predicted)	
good (Actual)	145	114	55.98%
bad (Actual)	30	936	96.89%
	82.86%	89.14%	

Class Statistics

Class	True Positives	False Positives	True Negatives	False Negatives	Recall	Precision	Sensitivity	Specificity	F-measure
good	145	30	936	114	55.98%	82.86%	55.98%	96.89%	66.82%
bad	936	114	145	30	96.89%	89.14%	96.89%	55.98%	92.86%

Overall Statistics

Overall Accuracy	Overall Error	Cohen's kappa (κ)	Correctly Classified	Incorrectly Classified
88.24%	11.76%	0.600	1081	144

Figure 19. Random Forest scorer. Source: own elaboration.

While the overall accuracy has increased to 88.24% as expected, it is interesting to look at how the model performed for the "good" class. In fact, while the Recall decreased to 55.98%, at the same time the Precision increased to 82.86%. In other words, it sensibly decreased the probability of a False Positive, while accepting some more False Negatives. That is, in the fictional task of suggesting a wine to a consumer, the model is more confident in predicting a wine that is actually "good", and on the other side is less likely to predict as "good" a wine that is "bad". Given that a consumer buying wine has made an investment and has a certain expectation towards its consumption, which may be particularly relevant for a social occasion (such as a dinner with friends), the impact of a wrong recommendation may be worse than overlooking some good wines, incorrectly labelling them as "bad". Precision instead may be considered a more relevant measure when recommending good wines, and may be a fundamental driver of trust. This is a typical trade-off which has to be taken into account when developing a User-centered application.

The Random Forest algorithm behave as a black-box, that means it is difficult if not impossible for a human to fully comprehend the internal functioning of the model, and all the steps that led to the output. In this case, this is due to the fact that Random Forests are made of hundreds of Decision Trees. Despite this shortcoming, it is still possible to observe the attributes which were employed the most by the algorithm in making the splits. For example, the density attribute was chosen 31 times for the top split criterion (level 0 of the tree) while the most popular attribute at level 1 and 2 was alcohol (employed 44 and 74 times respectively). This gives an idea about which were the most relevant attributes the model considered in developing its prediction.

Finally, the last classifier employed is the Multilayer Perceptron (see *Section 1.3.3*). Also called feedforward neural network, its goal is to approximate some function f∗ by learning the value of the parameters that result in the best function approximation. These models are called feedforward because information flows through the function being evaluated from x, through the intermediate computations used to define f, and finally to the output y (Goodfellow et al. 2016). Given their non-linear learning capabilities, MLPs are powerful classifiers. The set up of the model entails setting the hyperparameters, which must be chosen arbitrarily. Given that multiple hidden layers did not improve the overall accuracy, one single layer and 10 neurons per each layer were set as hyperparameters. The results are shown in Figure 20.

Multilayer Perceptron - Scorer View

Confusion Matrix

	good (Predicted)	bad (Predicted)	
good (Actual)	112	147	43.24%
bad (Actual)	58	908	94.00%
	65.88%	86.07%	

Class Statistics

Class	True Positives	False Positives	True Negatives	False Negatives	Recall	Precision	Sensitivity	Specificity	F-measure
good	112	58	908	147	43.24%	65.88%	43.24%	94.00%	52.21%
bad	908	147	112	58	94.00%	86.07%	94.00%	43.24%	89.86%

Overall Statistics

Overall Accuracy	Overall Error	Cohen's kappa (κ)	Correctly Classified	Incorrectly Classified
83.27%	16.73%	0.426	1020	205

Figure 20. Multilayer Perceptron scorer. Source: own elaboration.

The model underperformed with respect to both Random Forest and Decision Tree in terms of overall accuracy, but slightly outperformed the Decision Tree model in terms of Precision.

The intrinsic sub-symbolic nature of neural networks and its black-box design does not let us have a glimpse of what attributes were considered as more relevant in the classification decision.

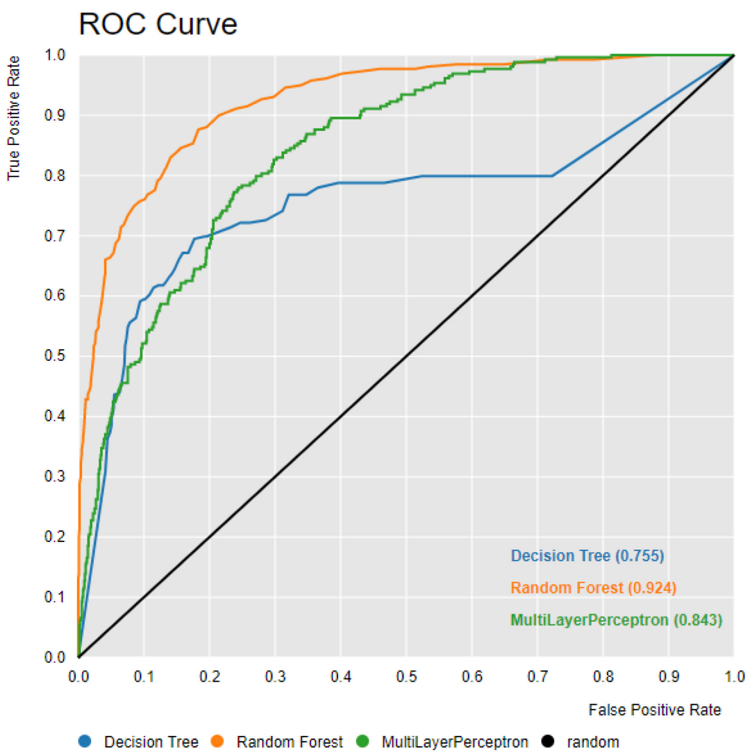

Figure 21. ROC curve comparing different classification models. Source: own elaboration.

Finally, it is possible to compare all the classifiers in a synthetic and intuitive way through a Receiving Operator Curve (ROC). Such curve plots the True Positive Rate against the False Positive Rate for the positive class ("good" in this case) across different discrimination thresholds, that are the probability at which the positive class is chosen over the negative class.

The black line represents the performance of a random classifier, that is the worst possible performance a classifier could achieve, while a perfect classifier will be a single point lying in the upper-left angle. Therefore, the higher the curve on the left side of the graph, and the closer to the asymptotes, the better the performance on predicting the positive class.

It is also possible to confront the Area Under Curve value, that literally represents the area under the plotted curve and is 1 for a perfect classifier and 0.5 for a random classifier. Again, the Random Forest classifier with and AUC = 0.924 performed best.

2.3.2 Text mining from Wine Enthusiast reviews

Another relevant source of wine information are expert reviews. Chen et al. (2014) proposed a new field of research that leverages wine sensory reviews in order to extract knowledge. In particularly, they proposed to analyze expert reviews due to the thousands of on-line reviews produced every year and their consistency. In fact, expert reviews are typically short and essential, consisting only in two or three sentences, and employ a precise and consistent vocabulary, which has a clear link with wine's intrinsic characteristics.

The credibility of the reviews is enhanced by the blind tastings procedure, which gives to the reviewers just general information about vintage, variety or appellation of origin. Wine ratings are based on a 100 points scale, but only wines receiving a score of 80 or higher are reviewed. Therefore, ratings span from acceptable (80-82 points) to classic (98-100 points).

For to analyze Vinho Verde expert reviews, a large dataset comprising more than 130.000 reviews from Wine Enthusiast was downloaded from Kaggle, an online data science community. The dataset is composed by the full text reviews and relevant attributes such as country, designation, points, price, province, and others.

From this datasets, only the reviews from the Vinho Verde province were extracted. This filter led to 410 expert reviews. The reviews were then divided into two classes according to their rating, namely "good" and "excellent", the former comprising wines rated less than 90, the former including the wines rating more or equal to 90. Additionally, given the intrinsic differences between red, white and rosé wine styles, and the consequent difference in the tasting notes employed to describe them, only white styles were chosen, that means Alvarinho, Loureiro, Arinto, and Avesso varieties and the remaining wines labelled as "Portuguese White".

The final datasets consisted of 333 reviews, consisting of 79 "excellent" wines and 254 "good" wines.

The goal of the analysis this time was to predict the class of each wine based on tasting notes derived from the sensory reviews. To do so, it was necessary to create a consistent vocabulary, in order to tag only the terms directly related to taste. Therefore, a first step consisted in analyzing the reviews through the Natural Language Processing toolkit available in KNIME, with the objective of highlighting through a Word Cloud the 200 most popular terms across all reviews. Then, the most relevant words in terms of sensory characteristics of wine were chosen and wrote into a vocabulary comprising a total of 52 attributes directly related to tastes, aromas and perceptions of wine (See the Appendix).

After the dictionary was created, a new workflow was started in order to apply the dictionary to the reviews and to perform a classification task. The workflow is represented in Figure 22.

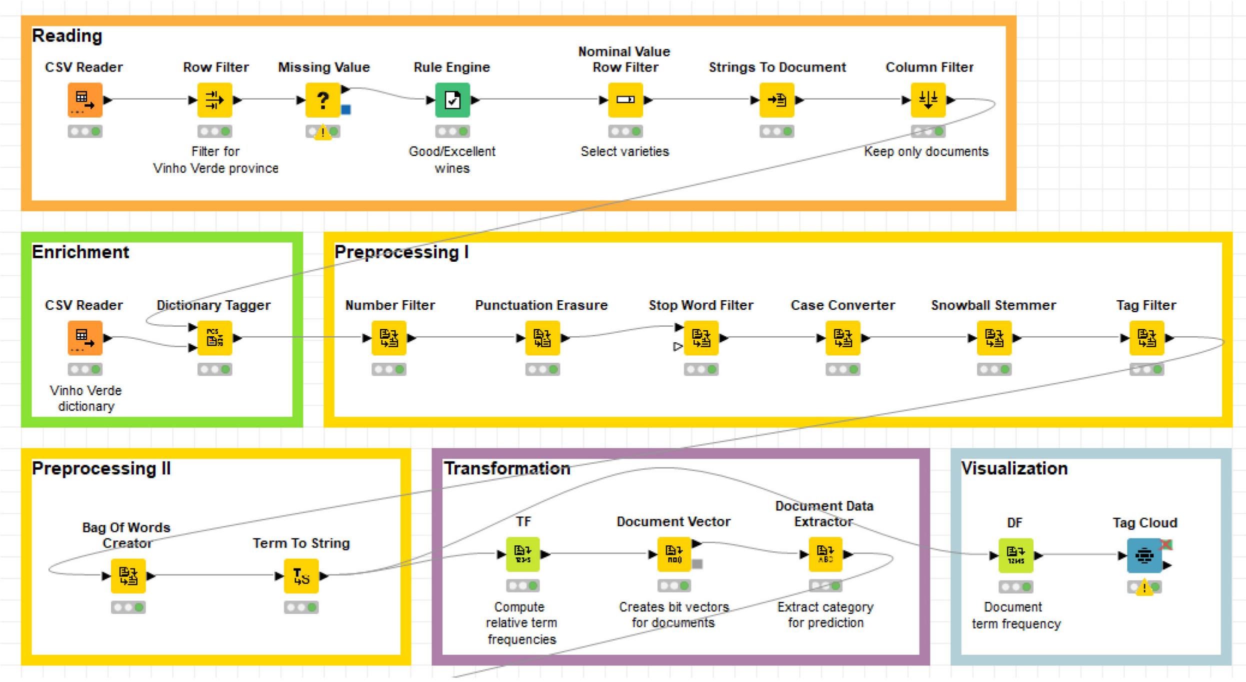

Figure 22. Workflow for text mining of Vinho Verde reviews. Source: own elaboration.

After creating a document for each review, comprising only the text of the review, the Dictionary Tagger node tagged all the terms indicated in each reviews. By navigating inside the documents, it was now possible to highlight the tagged terms, as shown in Figure 23. Additionally, by allowing for case sensitive tagging it was possible to catch many declinations of the same words.

> The name is not a Portuguese greeting, but a blending of Alvarinho and Loureiro, the grapes in the wine. It is a rich, juicy wine full of apricot, apple and pear flavors. It has just the right amount of tangy texture and refreshing acidity. The wine is fruity, perfumed and bright.

Figure 23. Example of tagged review. Source: own elaboration.

The documents than went through two Preprocessing phases. In the first phase, all the text contained in the documents was standardized by removing numbers, punctuation, stopwords, and upper cases, and then proceeding through a stemming process. Finally, a filter selected only tagged terms to keep, in order to get rid of unnecessary terms. In the second Preprocessing phase, the Bag of Words node created a column with all the terms contained in all the documents. This step is used for to represent each document as a collection of attributes.

Finally in the Transformation phase, the Document Vector node transformed each document into a bit vector, where each document is represented by a collection of attributes and the absence or presence of an attribute is indicated by a Boolean (0,1). After extracting the categories "excellent" and "good" from each document, this transformation process allowed to finally perform a classification task.

At the same time, by computing the document frequency, that is the number of documents in which each term appeared, a new and cleaner version of the Word Cloud was computed, showing the most popular 50 words associated with Vinho Verde, as shown in Figure 24.

Figure 24. Word Cloud representing the 50 most popular words across all reviews. Source: own elaboration.

Finally, the Random Forest algorithm was applied in order to predict the quality of the wine from the bit vectors. The results are shown in Figure 25.

By looking at the splitting attributes for the Random Forests, the most popular at level 0 were *ripe*, *rich*, *light* and *pear*.

Scorer View

Confusion Matrix

	excellent (Predicted)	good (Predicted)	
excellent (Actual)	10	17	37.04%
good (Actual)	2	71	97.26%
	83.33%	80.68%	

Overall Statistics

Overall Accuracy	Overall Error	Cohen's kappa (κ)	Correctly Classified	Incorrectly Classified
81.00%	19.00%	0.416	81	19

Figure 25. Random forest scorer from the text mining workflow. Source: own elaboration.

2.3.3 Conclusions

By deploying simple machine learning based workflows it was possible to predict the quality of wine, both from physicochemical characteristics and from expert sensory descriptions.

Additionally, it was possible to get some insights about what attributes could drive wine quality by looking at the most relevant attributes for the models. In fact, the results suggests that a quality Vinho Verde wine has a moderately high alcoholic degree and the most valuable tasting notes could be peach, creamy, and rich. At the same time, the Word Cloud gives a clear picture of the range of tasting notes related to Vinho Verde wines, which is characterized by its acidity, its crispiness, and its fruity notes such as citrus and apple. On the contrary, the physicochemical data only gave some insights about hygiene factors in wine production, such as the level of volatile acidity and density which usually are bad signals of the quality of production. This attributes could be further investigated to obtain insights for winemakers and marketers.

Overall, the models performed better in classifying the bad wines, than to identify the good ones. In fact, few of the attributes could explain the actual taste of wine, and few of them were correlated with quality. Additionally, the unbalanced datasets hampered the performance of the models, even though it is a common problem when working with real-world data.

The massive generation of data, even in the wine industry, offers the opportunity to employ more and more AI-based models for to complete a variety of tasks. On one side, by developing adequate analysis standards which could relate the physicochemical characteristics of wine with its sensory attributes and quality perception in a clear and effective way, it could be possible to produce models useful to set production parameters, for example for targeting valuable markets as suggested by Cortez et al. (2009).

On the other side, the text analysis of reviews may generate precious insights into the human sensory perceptions of wines. Both expert reviews and textual sensory descriptions in large datasets could be employed to generate high level insights, given their trained ability to detect tasting notes and the advancements in Natural Language Processing. Consumer reviews instead could be employed for clustering consumers tastes and to build better and better predictive models of user preferences through continuous learning. Recommender systems based on preferred tasting notes are already a reality. Given the complexity of the wine quality perception and the fluctuations over time of taste preferences, consumers and experts sensory descriptions represent an unprecedented opportunity to investigate the relations between sensory perceptions and quality.

The two digital representations of wine, both in terms of chemical data and sensory reviews, are complementary to each other. By defining a clear link between them it could be hypothetically possible to build a comprehensive model of digital wine, where user preferences could be analyzed automatically, extracting production parameters and insights for target marketing, thus enabling to better meet consumer demand. By offering to customers a compelling and satisfactory experience, wine platforms could foster a higher engagement with wine, therefore promoting their use and the generation of more user data in a self-reinforcing cycle.

Platforms are not only advanced e-commerce businesses and/or social networks. In their broader definition, they represent a *keystone*, that is a central hub of a network that resembles an ecosystem, acting as a technological and architectural foundation for the development and innovation of the whole ecosystem (Gawer & Cusumano 2013). In the next chapter, once deepened the results of this chapter into an integrated view of digital wine taste preferences, two different business models with two different approaches to wine digital representation will be compared to assess their potential as *keystone* companies, at the center of a digital wine value ecosystem.

CHAPTER 3: DIGITAL WINE BUSINESS MODELS

The proliferation of wine apps and e-commerce platforms has brought the wine industry into the digital age. However, in a world where digital, personalized experiences are becoming a hygienic factor more than a competitive advantage, the dynamics of platform-based competition may erode profitability. Platform businesses who want to strive for success are expected to deliver user-centric, compelling experiences, while at the same time generating significant value for all the networks participants they engage with.

Among the many successful and potential digital wine business models, two companies, Vivino and Tastry, have been chosen for being analyzed due to their potential impact on the wine industry. In particular, their way of leveraging consumer and wine data for building accurate models of wines' attributes and consumers' preferences opens up to the opportunity of creating value at all levels of the supply chain. Additionally, their digital, modular, and data-driven nature enforced by the consistent data pipelines they created, enables them to become an innovation platform for the whole wine sector.

The possibility of creating a modern business ecosystem for wine by leveraging data and user-centric platforms could be of great interest for the wine industry, which is facing deep transformations and challenges due to increasing competition in a globalized market, shifts in consumer preferences and the impact of climate change on wine production. Both small and big, product-focused or marketed-oriented producers could benefit from the insights and services generated in such system.

In this chapter, first it will be addressed the study of consumer preferences by means of sensory science, then, different approaches in building a wine recommender system based on flavor will be explicated. Afterwards, two companies, an industry leader (Vivino) and an innovative startup (Tastry), will be analyzed in order to highlight how the design choices in creating wines' and consumer's models have an impact on the way data is leveraged to create value in their business models. Finally, the potential impact at the industry level of such platforms is explored, and major industry challenges will be addressed.

3.1 Flavor-Based Wine Recommender Systems

3.1.1 Wine consumer segmentation criteria and sensory science

Wine is a complex beverage from a chemical and, consequently, sensory point of view. Wines offer complex flavor profiles, with taste, olfactory, visual and chemesthetic perceptions combined by our brains into a unique perception. The wide variety of flavors and aromas experienced during wine consumption elicit in consumers hedonic liking and pleasure.

However, given its experiential nature, consumers cannot assess the quality of a wine before consumption. Therefore, they tend to infer the quality of a wine from its extrinsic attributes, such as brand, origin, price, labelling and packaging (as discussed in Section 2.1.1). Such attributes have been demonstrated to exert a high influence on the purchasing decision and represent an integral part of the experience of wine quality (Charters & Pettigrew 2007).

The information-intensive nature of wine results in wine marketing being inherently difficult. Additionally, competition in the global market and shifts in consumer preferences have contributed to an even more fragmented market, therefore making the definition of effective segmentation criteria critical for targeting the modern wine consumer.

Traditional wine consumer segmentation criteria are based on socio-demographics, psychographics (in particular, wine involvement and knowledge), lifestyle and behavior (for example, frequency of consumption and consumption occasion) (Pomarici et al. 2017).

Given that wine is a product rich in attributes which may exert an influence on consumer behavior and purchasing decision, an alternative way of segmenting wine consumers is by measuring their preferences towards specific wine attributes.

Pomarici et al. (2017) proposed to use a combination of different segmentation criteria in order to study the preferences of a sample of 504 wine consumers from the New York metropolitan area. In particular, they collected consumers' preferences towards a set of wine attributes through a Best-Worst scaling method, in addition to consumers' psychographic characteristics. The latter included the level of wine involvement, subjective knowledge, innovativeness (willingness to try new wines), and loyalty proneness. Based on such combination of segmentation criteria, they were able to identify four relevant consumer segments (Pomarici et al. 2017):

- *Experientials*. They represent the largest segment of the sample (34%) and they value most previous experience followed by recommendations. Additionally, they are the most innovative segment of all, indicating openness to try new wines. This segment displays a high degree of confidence deriving from previous experience and a propensity to acquire information for future purchases. Given their self-confidence, recommendations are preferred via impersonal sources.

- *Connoisseurs*. The second largest segment (29% of the sample), they value most in order grape variety, wine maker, brand name, and country of origin. They are the most involved in the wine category, the most loyal to their choices and show the highest subjective knowledge of wine. Given their commitment with wine, this segment shows a more complex purchasing behavior that involves the analysis of several attributes, and a strong loyalty to brands that meet their expectations in terms of such attributes. They display a keen interest in learning more about wine, but again they tend to avoid expert recommendations, preferring instead to build their own opinions.
- *Risk minimizers*. They are the segment (18% of the sample) which values the most recommendations, while relying also heavily on previous experience, and considering important other cues such as price and brand name. Given that the main risk perceived by these consumers is to choose the wrong wine, this segment relies heavily on recommendations and previous information in order to mitigate such risk and to simplify their choice.
- *Price sensitive*. This segment (18% of the sample) values price and store promotions more than any other segment. They show the lowest involvement and knowledge in the wine category, and are among the least innovative consumers. Given their focus in price, this segment shows a preference for buying wine mainly in promotion. For such consumers, who do not have established wine preferences, price is capable of driving both purchasing decision and brand loyalty.

The above identified segments, while lacking in generalizability due to the design of the experiment, are consistent with previous results in the wine marketing literature and such segmentation criteria may be a valuable base upon which setting the marketing strategy, mainly through modulating the marketing mix with respect to the peculiar preferences for wine attributes of each identified segment. Additionally, an interesting result of this study is the consumer preference towards the recommendations and previous experience attributes across all segments. This could be explained with the fact that wine complexity makes for a difficult purchasing decision, therefore forcing consumers to often rely on risk-reduction strategies irrespective of the involvement level. Furthermore, more involved and knowledgeable consumers appear to be reluctant to critics' or friends' recommendations, preferring to build their own personal preferences upon impersonal information sources.

While modulating the marketing activity based on consumer preferences on specific wine attributes is a valuable practice itself, it must be considered how most of the attributes which could be controlled in such a way are extrinsic, with limited or no relation to the actual intrinsic properties of wine. While the marketing mix may be effective in driving consumer purchasing decision at the point of sale, where typically consumers cannot taste wine prior their purchasing decision and so they rely on extrinsic attributes in order to set an expectation of quality, this may not guarantee market success as consumers will test such expectations upon consumption. Repurchasing decision and brand loyalty may significantly depend upon meeting consumers' expectations in terms of flavor.

The issue of flavor is particularly relevant for wine, as it is a beverage capable of delivering complex flavor profiles, with a huge degree of variation across different wines. Consumers seek for a wine which they will enjoy drinking, or which is fitting for a particular social occasion or for to be paired with a particular food. More experienced consumers find an additional source of pleasure in appreciating the peculiar differences within flavor profiles, and by discriminating wines according to such characteristics, often relating them with attributes such as variety, origin, vintage, and winemaking practices.

Charters and Pettigrew (2007) investigated through qualitative interviews how consumers perceive wine quality, with the goal of identifying the dimensions which interact in defining an overall quality assessment for consumers. They concluded that although quality is a multidimensional construct, ultimately intrinsic quality attributes are the most relevant for wine consumers when assessing quality upon consumption. Pleasure is the end state connected with the experience of quality, with taste being the main catalytic dimension for reaching such state.

Reinforcing this view, Thach and Olsen (2006) surveyed a sample of US Millennials in order to investigate how to target such consumers and highlighted that the major reason why Millennials do drink wine is that they like the taste, while the opposite is true for the Millennials who don't drink wine.

The primary role of taste in the experience of wine quality questions the effectiveness of market segmentation based solely on extrinsic cues of quality. However, as discussed in Section 2.2, not only wine has an extremely complex chemical composition, which results in complex interactions between compounds that generate a difficult to predict overall perception of flavor, but also consumers differ significantly in their perceptions. In fact, as a unified flavor perception is processed by the human brain by integrating the multisensory stimulus coming from tasting the wine (Smith 2019), individual physiological and psychological differences have a role in shaping individual perceptions and, ultimately, flavor preferences. This subjectivity of flavor further increases the complexity of wine choice, and demands more insights into consumer flavor preferences in order to develop effective targeting strategies.

Sensory evaluation, defined as a "scientific discipline used to evoke, measure, analyze, and interpret reactions to stimuli perceived through the senses" (Lesschaeve 2007), has been addressed by the wine industry and research community as the reference methodology in order to gather quantitative, unbiased, and actionable information about consumers sensory preferences (Francis & Williamson 2015).

Common goals of sensory science are (Lesschaeve 2007) (Yang & Lee 2020):

- characterize sensory properties and identify which characteristics influence the preference for wine by consumers;
- test whether consumers can differentiate between specific attributes of wine;
- assist winery operations by understanding consumer preferences in order to design better wine styles to meet their expectations in the market;
- assess the sensory impacts of viticultural or enological treatments on finished wines.

Additionally, sensory science methods have been employed to assess whether it was possible to identify flavor and aroma profiles linked to regional typicity by comparing wines from different geographical regions or denominations of origins (see *Section 2.2.2*), and to investigate if consumers are able to perceive such differences and discriminate wines according to them (Souza Gonzaga et al. 2020).

The reliance on sensory science is rooted in the observation that wine experts cannot predict consumer liking or market success (Lesschaeve 2007). In fact, no relationship between quality judgement of experts and liking of consumers has generally been observed (Francis & Williamson 2015). Despite that, wine industry has typically relied on internal or external wine experts, while implementing a sensory program is still perceived as an expensive investment (Lesschaeve 2007).

The fundamental method employed in sensory science is the *descriptive sensory analysis* (DA), in which a trained panel characterizes the sensory properties of a sample of wines by identifying and quantifying the intensity of specific attributes. Sensory data collected from such panels must be accurate, sensitive, repeatable, and reproducible (Lesschaeve 2007).

A consumer-based trained panel has been advocated as best practice for external validity (Francis & Williamson 2015) and for their use of a consumer-relevant and non-technical-based language (Yang, Lee 2020).

Once a consistent list of attributes has been generated through descriptive analysis and each sample has been characterized in terms of intensity of such attributes, a consumer hedonic test may be carried under blind tasting conditions. In such way, the particular appearance, aroma and flavor attributes and their intensity can be related to consumer preference or liking (Francis, Williamson 2015). Such approach relies on the principles of psychophysics: while consumers find it difficult to verbalize which tastes, aromas, colors and chemesthetic sensations they perceive and are driving their liking, they can still react to such sensory stimuli and deliver an overall assessment of liking (Lesschaeve 2007).

In *preference mapping*, one of the most popular sensory evaluation methods, consumer liking responses are overlaid on a two (or more) dimensional plot with the sensory characteristics of a set of wines (Francis & Williamson 2015).

Francis and Williamson (2015) gathered the results of several preference mapping studies over the last 15 years in order to assess whether there are consistent wine sensory attributes related to consumer liking. A common result of such studies is that consumer preferences are not homogeneous, with clear clusters of consumers identified based on preferences, and some studies have shown that consumers can be surprisingly responsive in their liking judgements to small sensory differences for some key attributes (Francis & Williamson 2015).

Study details	Consumer group	Attributes related to liking	
		Positively related	Negatively related
Ten 'inexpensive' California Chardonnay wines, 126 consumers (Yegge and Noble 2001)†	Cluster 1 (25%)	Floral aroma	Oak
	Cluster 2 (14%)	Oak, peach/apricot	Floral
	Cluster 3 (25%)	Moderate oak, moderate apple or peach/apricot	Floral, astringency
12 international Chardonnay wines, <$US15, 361 consumers (Lésschaeve et al. 2002)	Cluster 1 (moderate proportion)	Alcohol, spicy/smoky oak, lingering aftertaste, sour, dry, bitter	Sweet, fruity, berry, smooth
	Cluster 2 (low proportion)	High overall flavour, sweet, vanilla, toasted oak	Sour, bitter, dry
	Cluster 3 (moderate proportion)	Moderate sweetness, sourness, low-moderate oak flavour	High flavour
	Cluster 4 (high proportion)	Sweet, smooth, fruit	Acid, bitter, dry
12 white wine blends, $CAN6–8, 115 consumers (Lésschaeve and Findlay 2005)	Cluster 1 (18%)	Alcohol burn, banana aroma	Earthy, musty, oak
	Cluster 2 (18%)	Sour, bitter, earthy, vanilla, oak	Smooth, sweet, melon, banana
	Cluster 3 (33%)	Smooth, sweet, melon, banana	Acid, bitter, earthy, vanilla, oak
	Cluster 4 (31%)	Alcohol, apple, tropical fruit and pear aroma	–
14 white wines: Riesling, Chardonnay, Pinot Gris, $A8–20, low sugar, 203 consumers (Francis et al. 2010)	Cluster 1 (41%)	Oak, alcohol heat, butter, viscosity	Citrus
	Cluster 2 (22%)	Citrus, acidity	Oak
	Cluster 3 (36%)	–	Bitterness, sourness, astringency
Ten Sauvignon Blanc wines, $US6–20, 109 consumers (Lund et al. 2009)	Cluster 1 (77%)	Fruity (stone fruit, passionfruit, box wood), capsicum, asparagus	–
	Cluster 2 (23%)	Flinty/mineral, bourbon	–
Neutral white wine with added Sauvignon Blanc aroma compounds, seven wines, 150 consumers (King et al. 2011)	Cluster 1 (26%)	–	Tropical fruit, cat urine, sweaty
	Cluster 2 (43%)	Fresh green, green vegetal	–
	Cluster 3 (31%)	Overall fruit aroma, tropical fruit, cat urine, sweaty, confectionary	Fresh green
12 Riesling, Chardonnay and Sauvignon Blanc wines, $US15–20, 120 consumers (Lésschaeve et al. 2012)	Cluster 1 (77%)	Tropical fruit, tree fruit, slight sweetness	Oak, vegetal, heat, dry,
	Cluster 2 (23%)	Oak, vegetal, heat, dry,	Tropical fruit, tree fruit, slight sweetness

†The liking scores of two clusters (36% of the consumers tested) could not be modelled well. –, not identified or inconclusive.

Figure 26. Sensory attributes of white wines found to be related to liking for clusters of wine consumers from preference mapping studies. Source: Francis & Williamson (2015)

These studies proved that despite the inherently variation across consumer preferences, consistent clusters which show clear preference towards key flavor attributes do exist. More meaningful conclusions are generally achieved when consumers are grouped into clusters with similar preferences as opposed to just reporting the average liking across the entire customer group, which would masks the underlying strong differences in responses, with the most liked samples from the group average often being the ones that are not disliked strongly by any of the consumers comprising the clusters, but are not necessarily the most liked by any cluster (Francis & Williamson 2015).

The differences in sensory preferences in a consumer sample can then be related to demographic data, such as gender and age; wine involvement and experience; or more complex measures such as psychological profile (also called psychographics); food 'neophobia' (preference for avoiding new foods or beverages); or genetically based physiological differences (Francis & Williamson 2015).

A powerful tool for researchers and the industry consists in putting in relation the quantitative descriptive sensory analysis and the elicited consumer sensory preferences with the chemical analysis of wine in order to understand the complex relationship between chemical compounds and sensory properties (Yang & Lee 2020). To do so, multivariate statistical methods such as Principal Component Analysis are usually conducted (Yang & Lee 2020). By correlating objective analytical measurements with consumer liking scores, the objective parameters that drive consumer likes and/or dislikes can be identified and, eventually, the optimal product formulation for a particular consumer segment can be determined (Lesschaeve 2007).

An interesting study by Mora et al. (2021) tried to uncork the definition of a trendy red wine for Spanish young consumers and to identify the sensory and chemical drivers of liking by means of a two-staged experiment. In order to reconnect young consumers with red wine, the study proposed to "reverse-engineer" the intrinsic properties of a perfect red wine for such target, starting from the conceptualization of an "ideal" red wine by means of an online survey.

In the second stage of the experiment, emotional and hedonic response of consumers has been recorded with regards to 12 commercial red wines, and the 5 wines which elicited the most different responses were further chemically analyzed in terms of oenological parameters and volatile compounds. Consumer emotional response in the sensory field have recently seen a surge in interest as emotions have been advocated as being more effective than the sole liking in discriminating between samples, as samples scoring similar for liking may elicit different emotions in consumers. Additionally, they have been reputed useful in finding relationships between wine extrinsic of intrinsic aspects and preferences/emotions (Mora et al. 2021).

The results of the experiment indicated that for young Spanish wine consumers flavor is the most important attribute of wine, followed by aroma, price and Protected Designation of Origin, and that a trendy wine is associated to terms such as *sparkling*, *soft*, *fresh*, *fruity*, *sweet*, *light*, and *balanced*. Furthermore, by performing a Principal Component Analysis (PCA), oenological parameters and volatile compounds families could be related with the emotions elicited, identifying in particular which physico-chemical characteristics were responsible for a positive or negative emotional response. In general, liking and positive emotions were positively related to volatile compounds from the organic acids group, while the presence of benzenoids and lactones, and the lower presence of terpenoids and norisoprenoids were associated with negative emotions. Furthermore, some individual molecules potentially responsible for the differences in emotional responses have been addressed by looking at their concentrations across the wine samples and considering their

Odor Activity Values (OAV). For example, isovaleric acid (acids), responsible for the "cheesy" aroma, and β-Phenylethanol (benzenoids), responsible for floral notes, have been found to have the highest concentrations in wines 1 and 6, which are the wines that elicited the most positive hedonic and emotional responses. The relative importance that the cheesy aroma may have in driving consumer liking highlights once again the complexity of the interactions between chemistry and flavor occurring within the wine matrix.

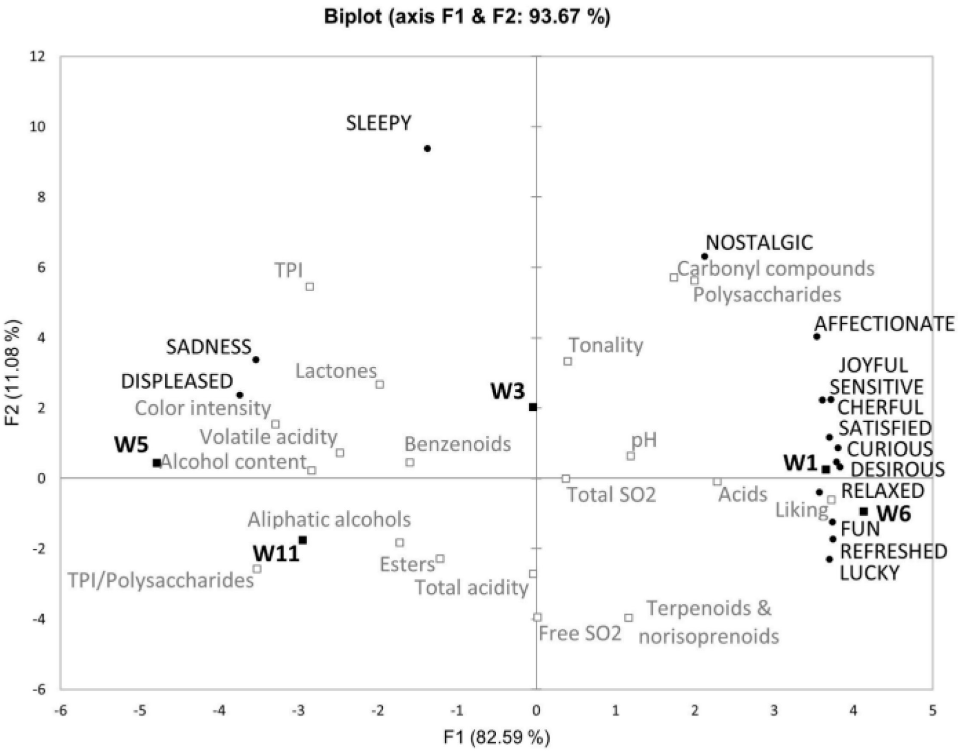

Figure 27. PCA of emotion categories elicited by five selected red wines with liking and physico-chemical characteristics as supplementary variables. Source: Mora et al. (2021)

The results suggest that soft and smooth wines with presence of floral and fruity aromas were more liked and elicited more positive emotions, while acidic, astringent, and alcoholic wines with presence of clove, coconut, leather and spice aromas elicited more negative emotions (Mora et al. 2021). Interestingly, the identified sensory attributes driving positive emotional response are consistent with the attributes that emerged during the previous conceptualization of a trendy wine by consumers, thus proving the consistency of the experiment.

Sensory analysis is also at the basis of formulating accurate wine descriptions. Such descriptions, which are typically available in the back-labels and now also in online platforms interfaces, represent a formidable way of communicating the taste of wine to consumers, and so to mitigate the perceived risk of purchasing wine. However, the expectations of taste that have been set while reading the wine description will be tested by consumers during consumption. Then, a mismatch between expectations and sensory properties may have two outcomes:

- when the mismatch is small, liking will tend to move towards expectations (assimilation);
- when the mismatch is large, liking will tend to move away from expectations (contrast).

The effect of wine descriptions on expectations, liking, emotions and willingness to pay for consumers has been studied by Danner et al. (2017) in relation to a sample of Australian consumers, which were asked to taste three commercially available white wines in 2 separate sessions. The first session consisted in a blind tasting, while the second session comprised an informed tasting with two different information level, namely basic (only sensory descriptions of wines) and elaborate (accurate and emotional sensory description plus highlights on quality and winery history). At each sessions consumers reported their liking, their wine-evoked emotions, their willingness-to-pay for such wines, and, before tasting wines in session 2, their expected liking.

Their results showed that the information level had a significant effect on all investigated variables as the elaborate information level evoked higher expectations before tasting the wines, plus resulted in higher liking ratings, elicitation of more intense positive (e.g. contented, happy and warm-hearted) and less intense negative emotions (e.g. embarrassed and unfulfilled), and a substantial increase in willingness to pay after tasting the wines compared to the blind condition, with the basic condition ranging in-between (Danner et al. 2017).

Another interesting result, consistent with the aforementioned assimilation/contrast theory, is that if the liking rating after tasting the wines matched the expected liking or exceeded the expectations by 1 point on a 9-point hedonic scale, participants felt the most intense positive emotions and the least intense negative emotions, whereas if the expectations were not met or the actual liking exceeded the expectations by > 2 points, participants felt less intense positive and more intense negative emotions (Danner et al. 2017). This study highlighted not only the importance of well written and accurate wine descriptions, but also demonstrated how information can influence consumers' wine drinking experience and possibly behaviour (Danner et al. 2017).

Sensory science has built a bridge between sensory properties, consumer preferences and chemistry yet to be fully explored. The potential in guiding winery operations towards designing wine styles that better meet market expectations, the ability to identify sensory properties which can discriminate between different clusters of consumers, and again the possibility to test winemaking practices and consumer perception of regionality traits, are all critical aspects in enhancing value creation and competitive performance across wine producers of all size.

On the other side, sensory experiments are deemed as expensive, and their results are difficult to translate to the market directly as consumer tastes show a high degree of inter-variability and are constantly evolving, therefore challenging the generalizability of their results.

A cost-effective, large-scale infrastructure for sensory science, which allows for flexibility and continuous learning, would be a desirable and powerful tool.

3.1.2 Flavor-based approaches to wine recommendations

Wine-apps are typical multisided platforms: they aim at *matching* users and wines from third-party wine producers, and by attracting more and more participants to their network increasing value is created for both sides.

While the e-commerce spread regarded all sectors of retail, it is interesting to observe the number and variety of platforms that have been launched specifically addressing the wine industry. Besides the many reasons that may explain this phenomenon, it may be argued that the success of virtual-assistant apps and online platforms relies on the fact that such platforms are particularly suitable for a complex purchasing decision such as with regards to wine.

First of all, through their matching function, platforms are specialized in matching a fragmented offer with a wide and varied demand. Secondly, most platforms nowadays employ to some extent recommendation systems, which automatically show or suggests items to users that may be of their interest (*see Section 1.5*).

In fact, such systems have been developed to deal with the overwhelming offer to which users were exposed to in the e-commerce platforms, leading consumers to experience the "paradox of choice" in which consumers instead of benefiting from the immense offer of products that is proposed to them, they instead experience anxiety which diminishes their well-being, paralyzing their ability of making choices.

The aforementioned complexity of wine in terms of intrinsic and extrinsic attributes, and the level of individual knowledge often required to make a purchase decision with confidence coupled with the vast and fragmented offer in the retail and spaces inevitably leads consumers to experience such paradox even in the traditional brick-and-mortar setting. Therefore, not only recommender systems are a critical function for a wine e-commerce, but with the advent of mobile apps and virtual shopping-assistants that enable consumers to access personalized information directly in the point of sales, they potentially represent a dramatic shift in the way consumers experience wine choice.

The common motif of the recent advancements in digital technologies is undoubtedly the *user-centrism*. Now that such technologies have reached widespread adoption and satisfying levels of performance, the focus is gradually shifting on enhancing the user experience by making it more streamlined and compelling and by digitally addressing consumer frustrations across the overall customer journey (Kotler et al. 2021). To do so, state-of-the-art recommender systems play a central role in learning a model of user preferences which can be exploited for to deliver a personalized experience.

Recommender systems are one of the most widespread business applications of artificial intelligence. They rely on machine learning algorithms (and, in more sophisticated declinations deep learning) in order to build a predictive model which is capable of learning or approximate consumer preferences and to suggest relevant items to them by filtering across the whole set (or a subset of) the available items. As discussed in Section 1.5, there are two main paradigms of recommendation (Ricci et al. 2011):

- Collaborative filtering, which basically makes recommendations to the active user based on items that other users with similar tastes liked in the past;
- Content-based, which matches up the attributes of a user profile in which preferences and interests are stored, with the attributes of an item, in order to recommend to the user new, interesting items based on what he has liked in the past.

However, a combination of the two is generally deemed as the best practice.

Collaborative filtering approaches have been the most popular ones thanks to their easy deployment and their satisfactory performance in terms of accuracy. However, as intelligent systems move towards more user-centric experiences, recommender systems' goals extended beyond accuracy of the algorithms towards providing helpful, enjoyable, and personalized experience that ultimately leads to user retention and satisfaction (Ricci et al. 2011). Such renowned focus has led to emphasize several additional evaluation metrics such as coverage, confidence, trust, novelty, risk, and serendipity.

Bringing these topics into the world of wine, an interesting question is whether a content-based approach may overall perform better than a collaborative filtering one.

One first argumentation that could be posed is that the subjectivity of taste, which is rooted not only in demographics and other easy-to-measure criteria, but also in physiological and psychological differences, poses an insurmountable problem in terms of suggesting a wine the user will certainly like only based on the average liking of its “nearest neighbors”, such as in the collaborative filtering paradigm. However, collaborative filtering algorithms have been successfully employed in fields such as movies and music where individual taste plays a predominant role, and they have proved to be not only accurate but also to be capable of delivering recommendations, thus recommending niche, unexpected but pleasant products to users.

A more compelling argument to propend for a content-based recommendation paradigm for wine ties back to the ultimate reason for which recommender systems have been conceived. In fact, recommender systems are designed to help consumers make better decisions (Ricci et al. 2011). In a competitive market like wine, where consumers naturally experience the paradox of choice when purchasing, while reportedly their levels of individual knowledge about wine continue to decrease (Wine Intelligence 2020), the possibility for the industry to influence the way in which customers make purchasing decisions is of paramount importance and may be a driver to increase engagement with wine.

Smith (2019) argues that AI collaborative filtering recommendations are mostly based on the “liking of the crowd” and do not give the underlying explanation of the reasons for such liking. In fact, collaborative systems are black boxes since the only explanation for an item recommendation is that unknown users with similar tastes liked that item (Ricci et al. 2011).

Consumers should instead be able to understand the reasons behind their choices and their liking in order to build up their knowledge and to be able to formalize their preferences. Given the reported relevance of individual knowledge in wine purchasing decision, refined understanding of wine and own taste preferences will eventually lead consumers to scale up on the quality and price ladder by experiencing and enjoying wines of greater complexity and interest (Smith 2019). This highlights the importance of transparency, that means that explanations on how the recommender system works can be provided by explicitly listing content features or descriptions that caused an item to occur in the list of recommendations (Ricci et al. 2011).

Directly connected to this need for a transparent design, another relevant argument could be made about what constitutes a good wine choice from a consumer perspective. As previously highlighted, the main concern for wine consumers is actually flavor, and the impossibility to taste wine prior consumption forces them to evaluate extrinsic cues of quality and to perform risk-reduction strategies. This means that, for a virtual assistant to significantly improve consumers' decision making, flavor-related knowledge must be taken into account. Further on this point, consumer-based trained panels performance has shown to be comparable to that of experts panels in sensory characterization and discrimination of wines (Yang & Lee 2020), thus suggesting that generally with training most tasters can come to make fine discriminations and with that improve the satisfaction and reward they get from tasting good quality wines (Smith 2019). The relevance of wine descriptions and its influence on liking and willingness-to-pay as highlighted by Danner et al. (2017) again stresses the importance of accurate and effective flavor descriptions and their role in increasing wine enjoyment, with potential benefits for both consumers and producers. Ultimately, it is the interplay between tasting and knowing that leads to refined discrimination and a better understanding of the wine and to go beyond mere liking (Smith 2019).

By building up knowledge about wine flavor, consumers may start to focus only on extrinsic cues which may directly relate to the organoleptic properties such as variety, origin, winemaking practices and overlook attributes such as price, promotion and brand. Again, the formalization of intrinsic properties through consistent flavor descriptions will enable flavor profiles to be a more effective source of differentiation in consumers' minds and to be a fundamental element in the purchasing decision, thus pushing towards a flavor-based competition and better aligning consumer flavor expectations with producers' offer.

Ultimately, the role of a flavor-based recommender systems will consist in learning the user flavor preferences and lead him to make good and satisfying choices based on the intrinsic properties of wine. Building flavor profiles of wines and flavor preferences of consumers is a tremendous tool for both consumers and producers willing to engage into a wine platform, as consumers could be offered wines whose flavor profile matches their taste, while producers may target consumers' preference profiles.

This is a significant step from collaborative filtering recommendation, where consumers tastes are just calculated as a probability of liking based on the average liking from unknown people whose tastes overlap in the data. A content-based recommender system with a flavor focus will ensure customers that recommendations are based on what it is ultimately important to them, that is flavor. Additionally, it would play a role in improving customers' understanding and enjoyment of wine, with potential benefits for their involvement and loyalty. In this way, the increasingly relevant objectives of explainability and trust would be reached.

The architecture of a content-based recommender system relies on two fundamental components: items representations and the users' profiles. In practice, the system applies supervised learning algorithms in order to generate a predictive model that is based on users' feedback. Given a new item, a filtering component is capable of predicting whether it is going to be of interest to the user by comparing the features in the item representation to those in the representation of user preferences (the user profile) (Ricci et al. 2011).

In order to build the user model, the system will employ users' feedbacks (may them be implicit, derived from users activity, or explicit, such as ratings and reviews) and the attributes of the items the user has liked or disliked in the past. The quality and quantity of feedbacks clearly depends on the level of engagement users have with the platform and on the user experience design. In order to deliver compelling and satisfying recommendations, the learning component of such system is vital, as it can take into account the dynamic nature of user preferences and deal with or even promote such dynamism.

However, given the primary role of items' features in setting user preferences and, eventually, in promoting interaction in the platform, items representation and the accurate extraction of relevant features is at the heart of such system.

Items representations may be derived from internal and external data sources, and then a content analyzer will extract the features of interest and shift the item representation from the original feature space to the target one (Ricci et al. 2011). When each item is described by the same set of attributes, and there is a known set of values the attributes may take, the item is represented by means of structured data (Ricci et al. 2011). It is important to notice however that no content-based recommendation system can provide suitable suggestions if the analyzed content does not contain enough information to discriminate items the user likes from items the user does not like (Ricci et al. 2011). The case of wine is exemplifying of this:

if the systems aims at recommending wines the user will like, items representation must reflect in some fashion the taste of wine.

As showed in Chapter 2, wine is a complex and information-laden product, which makes it well suited for a content-based approach. Information about wine style, grape variety, vintage, winery, and even flavor descriptions are typically available and well-displayed in wine online platforms. However, such richness in data may be underexploited by classic non-personalized recommender systems or by basic forms of collaborative filtering ones. Additionally, information such as flavor descriptions may be episodic, unstructured and inconsistent.

In this sense, it is not sufficient to employ an effective strategy to exploit the value that potentially lies in the data, but the system should also be designed in order to be able to gather consistent data and to extract it in a suitable form upon necessity.

A number of strategies may be employed in order to enrich the item representation with features that are relevant to learning user preferences. Ontologies and dictionaries may be used to enrich the system representation of complex concepts, as could be the case of wine. It is in this in this field that sensory science could play an increasingly relevant role; with its capacity to translate wine organoleptic properties into consistent and accurate descriptors, descriptive analysis is a powerful technique to generate actionable data for an intelligent system. Another relevant strategy stems from recognizing that many platforms have a treasure in user-generated content in the forms of text, such as rating and comments, which may be rich in item descriptors and which is suitable to be processed in order to extract tags.

Today, several wine apps claim to have developed accurate wine recommendations, and, despite the pretty small volume of wine e-commerce pre-pandemic, a survey conducted in 2015 by Wine Business.com highlighted how 25% of the sampled US wine consumers have wine apps and use them for their purchasing decisions. There is so a clear interest of consumers in wine apps, which is destined to grow as the business increasingly moves to digital after the pandemic. The claim of "helping consumers to choose the perfect wine with artificial intelligence" has naturally evolved during the years, and today sophisticated, truly wine-specific recommender systems are being deployed in the market, and are successfully tackling the problem of digitizing flavor preferences.

Vivino, the market leader in wine apps, has recently introduced in March 2021 its new recommendation system. With the impressive number of over 50 million active users, who wrote over 70 million reviews, Vivino realized that the real competitive advantage of its platform was its community and decided to revolutionize the user experience by leveraging the content its users created everyday.

At Vivino they decided to leverage user reviews in enriching wine information. In particular, by using Natural Language Processing algorithms, they extracted keywords related to wines' taste and flavor from user reviews and aggregated them in order to create synthetic and graphically pleasant representations of a wine flavor profiles. Users can also select one flavor and browse all the reviews mentioning all keywords related to that flavor through a sort of real-time text analysis. Other dashboards related to taste characteristics such as body, acidity, or dryness/wetness are created in a similar fashion.

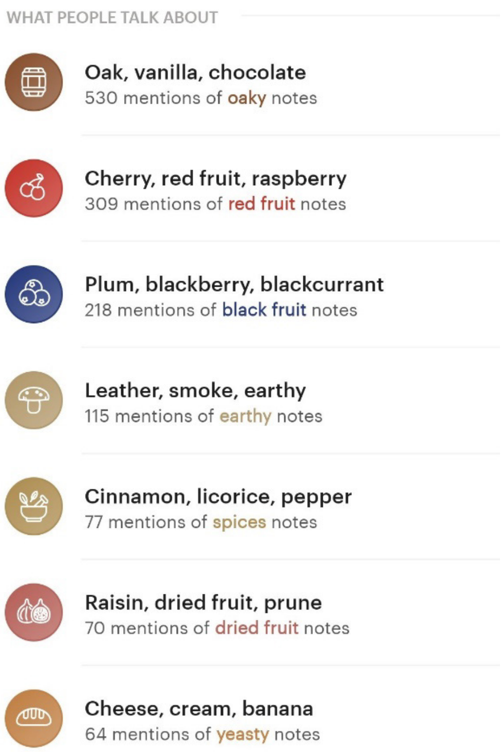

Figure 28. Flavor profile of an Italian red wine based on Vivino users' reviews. Source: Vivino mobile app.

The aggregation method in wine flavors and characteristics gives a structure to the massive amount of user-generated content, and even pushes the users to write reviews consistent with this framework by letting them access such keywords with one tap upon writing a review, thus educating them in delivering more and more accurate sensory descriptions. Interestingly, the way in which tasting notes are structured resembles the famous "Wine Aroma Wheel" (Fig. 29), created by Professor Ann C. Noble at the University of California Davis as a guiding instrument in tasting sessions.

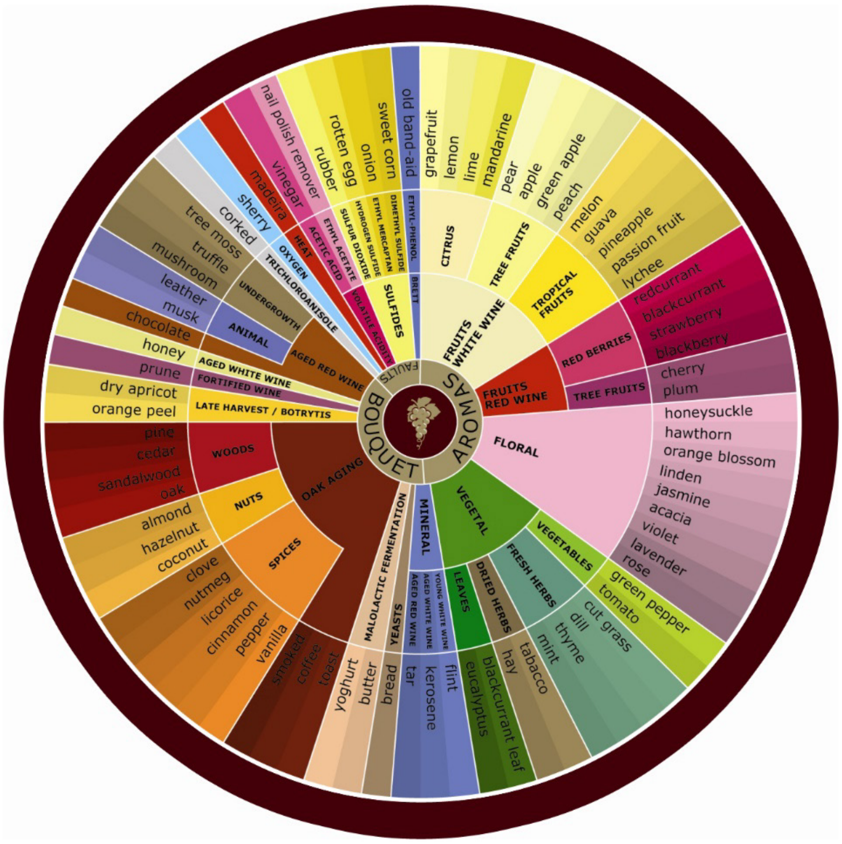

Figure 29. Wine Aroma Wheel. Source: Aromaster (2016), CC-BY-SA 3.0
https://creativecommons.org/licenses/by-sa/3.0/deed.en

With this changes, Vivino revolutionize the user interface with effective dashboards that describe wine taste in an immediate, consumer-driven way, fostering clarity and trust. It took the immense amount of unstructured data that users are generating and structured into a flavor-based system which is suitable to be employed to enrich wine representations and, consequently, user taste profiles.

Tasting notes are then employed into the construction of a detailed user profile, which keeps track of the users preferred tasting notes, along with the wine styles, regions, and grapes the user has tried. After rating at least 5 wines, the recommendation engines starts to calculate a matching score of how much a new wine will fit the taste of the user.

Such user-generated wine representations represent also a non-valuable resource per se, representing a collective consumer sensory perception of wine organoleptic properties. The integration of a sensory science approach in structuring the data implies that such data could eventually be employed for an experimentation platform or for external research.

While Vivino took existing AI approaches and brought them to the state-of-the-art, other companies are experimenting with radical approaches by integrating AI capabilities with a science-based approach.

Tastry, a California-based startup, developed a patent-pending AI software that provides wine recommendations based on chemistry. To do so, wines are comprehensively tested for volatile and non-volatile compounds in an untargeted manner. The untargeted approach is a precise choice in order to get a non-reductionist picture of the wine matrix. By having such rich dataset, AI is expected to be able to learn by itself the complex interactions between matrix compounds that lead to flavor expression. In this way, they claimed to have developed a new analytical chemistry approach which, by leveraging AI and sensory science methods, have decoded the flavor matrix.

The other fundamental part of their system architecture consists in consumers palate data, where consumer palate is profiled upon the completion of a survey.

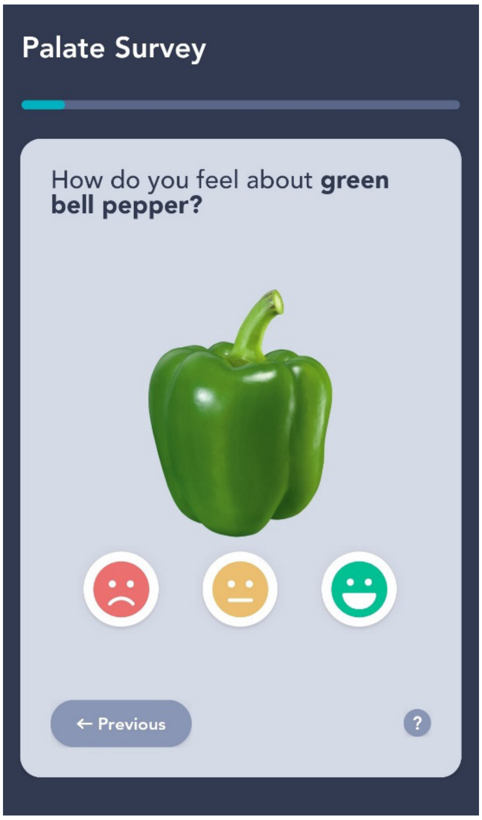

Figure 30. Bottlebird palate profiling survey. Source: Bottlebird mobile app

The quiz, which can be accessed via their *Bottlebird* mobile app or in in-store kiosks, lasts around a minute and asks to the user simply to rate how they feel about some food ingredients, which clearly are related to relevant wine flavor and aromas, or whose liking or not is instrumental to characterize the palate of a consumer. It is interesting to notice that the user is questioned about how does it feel about and not about how much does he like. This could be a precise choice in order to elicit an emotional response which is regarded as more informative of actual preferences, in accordance with previously mentioned studies from Mora et al. (2021) and Danner et al. (2017).

The resulting datasets coming from chemical analysis and palate surveys are employed conjointly to generate wine recommendations to users based on chemistry and on their unique taste preferences. The proprietary AI technology is capable of matching wine flavor profiles with consumer palates by translating them into the same feature space. This results in recommendations which are tailored to every single user and based in anything else but flavor.

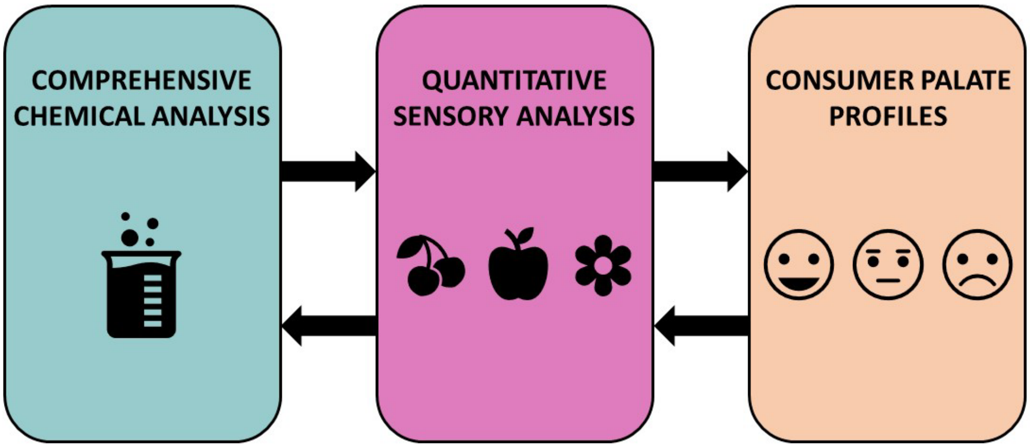

Figure 31. Tastry's datasets for wine recommendations. Source: own elaboration.

The overall output is again represented by a matching score that represents how much a wine will meet the individual taste of the user. Similarly to Vivino, users can then explore their palate preferences through dashboards and even modify them, thus finely tuning their user profile.

The overall architecture of the AI system built by Tastry is represented in Figure 31. It is interesting to notice how the data can flow in both directions, from comprehensive chemical analysis to consumer palate profiles and vice-versa, thus signaling how both wine producers and wine consumers can benefit from the insights generated by such system.

By using two substantially different approaches, both companies have successfully obtained a high-level digital wine representation, and leveraged it to create value for their consumers with flavor-related wine recommendations, potentially improving the way in which they make their wine choices.

Both approaches achieve an unprecedented level of personalization, focusing on the individual palate preference and by matching it to the bewildering range of wine offered in the market. Additionally, if adequately supported with learning components, such systems may dynamically adapt to the change and the development of taste preferences. Such personalization achieved at scale is called segments of one, and can only be achieved through intelligent systems capable of leveraging the massive amounts of data gathered by platforms about items and users.

Finally, this recommender systems may only be a primordial form of virtual sommeliers, which could leverage state-of-the-heart bots and the same knowledge and data in order to deliver tailored recommendations through warmer, human-like interactions.

However, these two approaches to wine representations show also significant differences, which reflect on the value and scope of their datasets, and also reflect in their business models and their potential industry-wide impact. Such differences will be explored in the next sections.

3.2 Case Studies: Vivino and Tastry

As digital platforms businesses have their foundation in data and networks, the way in which they leverage the data in order to deliver valuable and compelling services to all the users in the network is at the base of their competitive advantage and service differentiation.

By naturally gathering huge amounts of data during the ongoing of their operating activities, to structure and recombine data from different sources enables them to ideate and offer new services to existing and new target customers, increasing the scope of their business by simply adding to their modular and flexible digital foundation.

While the digital nature allows for this flexibility, the business model definition is still a critical step in defining how the company will operate and will set the boundaries to its potential scope and growth. It will define how company resources will be employed to create unique value for the participants to the platform, and how the company will tackle competition in a hypercompetitive market such as the one of platforms, where profitability is drained by multihoming. It will ultimately define what is the quality advantage in the user experience offered, which is the fundamental driver of competition, especially for the new entrants to the market.

In this section, two companies with a similar mission, helping consumers to find wine that they like, will be analyzed. However, significantly business models have been built around this idea, the main difference being the way in which data is collected and leveraged. The consequences of such difference go beyond the single business level, but extend to the way such companies may impact the wider wine industry.

3.2.1 Vivino's business model

Vivino is world's largest online wine marketplace and most downloaded wine app. Based in San Francisco, California, it was founded in 2010 by Heini Zachariassen and Theis Søndergaard. Its value proposition consists in delivering a "unique wine shopping experience that uses community data to suggest personalized wine recommendations, making wine discovery and purchase fun, accessible, and effortless for wine drinkers of every level".

Its mission is to help consumers choosing the right wine by matching them with an impressive offer 14 million wines offered in its marketplace coming from third-party sellers and proprietary inventory. With over 50 million users as of 2021, Vivino is by far the largest online wine community in the world.

Due to the pandemic boost on wine consumptions, wine sales on Vivino's platform more than doubled in 2020 to $265 million. Additionally, the company reportedly raised $155 million in funding in 2021, which will foster its expansion in terms of staff (up to 300 people) and countries served (actually 17), and will fund the development of their new recommendation engine. While company valuation remains undisclosed, it may reportedly be around $600 to $800 million (Thach 2021). However, the company still didn't reach profitability.

Like many successful digital applications, Vivino was born with the idea to solve a user-centered problem, that is to choose wine from the bewildering offer available at any brick-and-mortar store. The first iteration of Vivino was essentially a platforms where users shared ratings and reviews of the wines they tasted. The technological solutions which enabled the success of the platform was the introduction of the possibility of scanning the label of a wine to get immediately on your mobile phone information about users ratings and reviews of that wine. Such solution enabled consumers to seamlessly shift from physical to digital space and substantially improved their shopping experience. As consumers started to scan wines and share reviews, its community and wine databases grow more and more. As of today, more than 1,5 billion of labels have been scanned, and users produced more than 75 million reviews and 213 million ratings. With such an impressive amount of user generated content and such an easy to deploy mobile technology, Vivino is capable of employing the *wisdom of the crowd* to deliver trustable recommendations in every purchasing occasion.

Vivino's impressive growth has permitted to put in place the necessary investments in order to become a marketplace and build an Amazon-like supply chain, attracting over 700 wine retailers which partner with Vivino in order to offer nearly 13 million wines from around the world. The integration of a marketplace in its shopping assistant app produced a profound change in the business model in terms of investments and future growth, and enabled Vivino to offer a thorough experience and to further blur the line between physical and digital shopping.

Additionally, its increasing data pipeline was enriched by transactional data, setting the basis for market intelligence capabilities, which may be exploited in the near future in the form of a service, as suggested by the words of founder Heine Zachariassen which acknowledged how Vivino has gradually become a data company.

Wineries selling on Vivino marketplace have the opportunity to enrich their displayed information, even visually with live-motion covers, and may get more visibility on the platforms through the promotions service offered by Vivino.

With the shopping experience gradually shifting back from brick-and-mortar to online, Vivino needed to deliver an engaging shopping experience also in the digital space. The solution was, again, to rely on the immense user-generated content, which was employed in order to extract flavor descriptors from wines and so enriching wine representations. The enhanced wine representations are then feed into the learning algorithms to create a user's unique taste profile, which is then matched with new, unseen wines. Users can then browse their own preferred styles, countries, and grape varieties, eventually increasing their knowledge and enjoyment of wine. The introduction of such a structured system opens the opportunity to get higher-level insights into users preferences, which could in turn employed for targeting or for to develop new functions.

However, Vivino did not stop with creating an efficient recommender system, but enriched the user experience with dashboards and compelling graphics again exploiting the rich information gathered through its users and partners. This interactive graphic style reflects the increasing introduction of game design elements (*gamification*) into various aspects of the user experience.

A clear example is the introduction of rewards, which can be achieved for example by scanning certain wines according to a certain rule, gaining the possibility to earn promotions and badges. Another relevant feature consists in taste adventures, in which the consumers that decides to embark on will get some notions about specific wine styles and will be invited to try them. This features are conceived in order to promote wine involvement and to develop consumers knowledge through the app, eventually resulting in consumer loyalty and profitability.

Such forms of personalization may potentially generate a lock-in effect in which consumers feel engaged and compelled to loyalty to the platform.

A last but equally relevant addition to the user experience regards the introduction and enhancement of social network functions inside the Vivino app. Not only users can set their profile picture and their biography, but recently a whole new section has been opened where consumers engage like in a typical social network settings, sharing posts and pictures, and reacting to and commenting other users' content. This feature opens up new possibilities to monitor users activity and to get precious insights by analyzing the *word of mouth* and the consumption occasions through, respectively, natural language processing and image recognition.

The core business of Vivino gradually shifted from a shopping-assistant to a state-of-the-art online marketplace, powered by accurate recommendations that are at the core of a system that promotes user engagement and trust, and which stimulates him to get more involved with wine by trying new wines and developing his knowledge. At the same time, value creation opportunities increase as more rich data is acquired about users taste profiles, which could be potentially exploited by offering to third-party vendors and wineries precious analytics services.

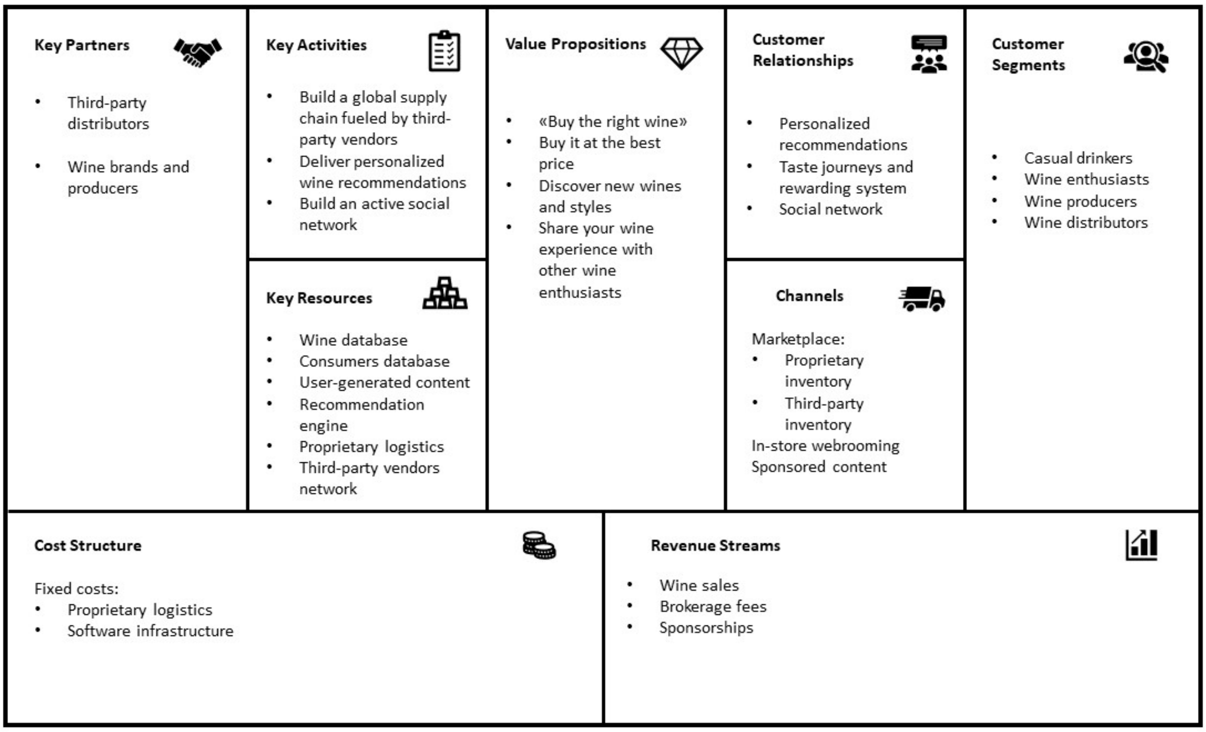

Figure 32. Vivino's business model canvas. Source: own elaboration.

3.2.2 Tastry's business model

Tastry is Software-as-a-Service and AI company based in San Luis Obispo, California. It was founded in 2014 by Katerina Axelsson. Its value proposition consists in "matching sensory products to people using chemistry, machine learning and every person's individual palate".

Tastry has raised over $6 million in funding and is now currently in the Later Stage Venture Capital. The company has 18 employees and it is currently making profits.

The company was founded on the idea that every palate is unique, and so recommendations should be based on an objective method capable of taking into account such uniqueness. A patent-pending AI-software has been developed by merging an analytical chemistry dataset with a consumers' palate dataset, thus projecting the user preferences into the same feature space as the wine flavor matrix and enabling the learning AI to recommend wines that match with consumers' palate.

In order to do so, Tastry chemically analyzes wines in an untargeted manner in order to get a snapshot of the complex wine matrix. This in-house data generation represents one of the most critical steps and the bottleneck for their operations. However, they claim to be already able with their actual infrastructure to analyze one third of all wines in the US every year.

On the other side, palate profiles data is generated through a quick survey which users can perform both on the Bottlebird app, which is powered by Tastry AI, and at in-store kiosk situated in partnering grocery stores.

The shopping experience offered by Tastry is based on two pillars. First, recommendations are unbiased and science-based, and are calculated on the unique palate of the consumer. Second, recommendations are calculated based on the inventory of the store to which the consumer interfaces. This overcomes a critical aspect of other virtual assistant apps, which asked the consumers an active effort in scanning several wines, to be picked up from the huge offer in the shelves.

This software architecture is at the basis of the business connections that Tastry is capable of create. By generating data starting from wine chemistry to users' palate, and to store inventories, and by relating such data through an integrated system, value can be created across the whole supply chain, and traditional industry silos can be broken down.

From a consumer perspective, the user experience is enhanced by unique, tailor-made recommendations, which are calculated as a matching score as soon as the consumer interfaces with a store or online inventory, thus simplifying the purchasing process to an unprecedented level. Additionally, its palate profile can be explored through interactive dashboards, which also allow for a finer tuning of preferences. In this way, not only the consumer can experience complete transparency into the recommendations system, but has also some degree of independency which lets him get more involved with the system.

However, wine consumers are just one of the target customers to which Tastry aims to offer a compelling, data-driven experience.

In fact, dashboards and insights are generated for both wine producers and retailers. Consumers data generated by Tastry are extremely valuable for the reason that they explicitly represent the unique and real taste preferences of consumers. Such palate profiles can be exploited not only at an individual level for successfully matching users and products, but can also be aggregated in order to get heatmaps of consumer tastes. Such maps are useful not only to individuate major taste trends, but also to identify niches in which producer could propose their products. The established relation between human palate and chemistry enables even to predict market success of a wine based on its chemical analysis, or to extract objective parameters that could be useful for wine producers to craft their wine styles. Currently, wine producers are granted a free access to Tastry analytics upon submitting their wine for chemical analysis.

As for retail stores, analytics about wines performances in the store could be combined with analytics about relevant taste trends in that geographical area in order to optimize the inventory assortment.

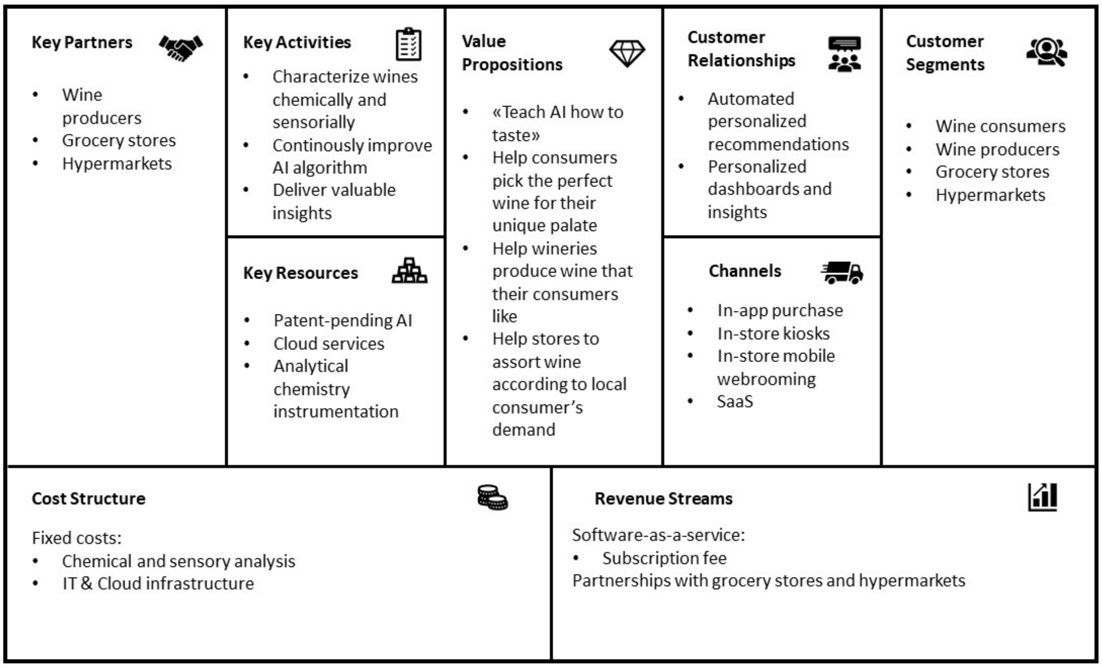

Figure 33. Tastry's business model canvas. Source: own elaboration.

The strength of Tastry business model relies in the fact that despite the initial investment necessary to chemically analyze thousands, or even million of wines, the huge AI capabilities that have been developed make their data extremely attractive at all levels of the supply chain, multiplying the profit opportunities. Furthermore, their innovative patented method of analyzing sensory products could also be transferred in other adjacent industries, significantly increasing the scope of their business.

3.2.3 Discussion and conclusions

Both companies, Vivino and Tastry, started their business for the same objective – helping consumers to choose wines they will like. However, significant differences occurred in their business models, and could be the key to interpret the potential dynamics of competition.

The first and foremost difference regards in which kind of data the two companies leveraged to build their digital business model. Vivino relied since the beginning on the wisdom of the crowd, that is, ratings and reviews generated by its users, in order to give advice to consumer at the moment of purchasing. With the growth of its community, such data became so valuable and consistent that it has been leveraged in a new, creative way to enhance wine representations and set up a new generation recommender system based on flavor. Tastry instead preferred a science-based approach and generated in-house the relevant data by analyzing chemically the wines and by letting the AI extract the relevant representation.

The data generated by the two companies, besides employed for the same task of recommendations, have different value and scope. Vivino users taste profile combined with the impressive size of its marketplace make such data a formidable source of market insights, which could be employed for to make more efficient the supply chain and for to timely analyze market trends. Such trends may be analyzed for flavor at an unprecedented scale.

On the other side, Tastry data is more expensive to generate, but the relations created across data at the different levels of the supply chain make such data extremely appealing and valuable to different actors. In particular, the most interesting and innovative aspect is perhaps the connection between chemistry and consumers' preferences, which opens up to a further integration of AI technologies with wine production by providing objective and measurable parameters. The possibility for consumers to use the same data to find the perfect wine from the stores' shelves and for retailers to optimize store inventory adds additional value to the data.

Besides both data generated are inherently valuable for the wine industry, there are differences in their scale and scope. Vivino data benefit from a scale which is unprecedented, and may generate high level insights that could benefit the whole industry. Tastry data instead has a broader scope, thanks to the software infrastructure that seamlessly relates the supply chain most critical levels in order to deliver to the consumer the perfect wine.

It is clear by analyzing their business models that the two companies are positioned at different levels of the supply chain. Vivino is fully embracing its marketplace nature, focusing on creating state-of-the-art logistics and a compelling user experience that promotes engagement and loyalty, therefore positioning itself at the end of the industry supply chain. Tastry, on the contrary, tends to vertically integrate in the supply chain by establishing innovative relationships in the data. However, it must be noted that Tastry does not directly sell wine and does not aim to do so, being its core product its SaaS. This significantly limits the competition between the two.

Another relevant difference with regards to competition is that while Vivino operates in a global scale, and generates its competitive advantage in operating as such, Tastry sets up a locally-deployed strategy, based on the proximity between the user and the product. This aspect is relevant as the ability to effectively serve locally some markets could be the key for success for platforms directly competing with global-scale competitors. Having a full range of stores connected with Tastry software and equipped with kiosks in a well-defined geographical area may push consumers in that area to find it more convenient to prefer that platform.

Perhaps the most interesting aspect when comparing these two companies is the potential industry impact of their business models. As Vivino is currently reshaping the industry towards online and DTC sales, and is changing the way in which consumers purchase wine, Tastry enables unprecedented connections between wine production and wine platforms, promoting the potential integrations of intelligent systems in the winery.

As a concluding remark, it must be noticed how the scale of Vivino and its digital foundation enable them to continuously enhance their platform by adding additional features and eventually further vertically integrating in the industry.

3.3 Towards a Digital Wine Ecosystem

The companies analyzed in the previous sections are clear examples of successful digitalization in the wine industry. By relying on different approaches on wine digital representation, objectives such as offering a more satisfying shopping experience have been achieved, while at the same time opening new scenarios for the industry.

The power of such platforms relies in the user-centered approach, that enables to identify and address customers frustration points. By breaking down wine representation to its fundamental, intrinsic properties, better advice can be tailored to consumers. With such refined model, user preferences can be learned with extreme accuracy, therefore allowing for a personalization at scale. The user profiles set in this way are the basis for an individual targeting strategy, that is, a segmentation of one. The matching function then is able to connect consumers with the wine that its unique taste likes. Continuous learning and algorithm improvement can then adapt to shifts in preferences, or can even promote the user exploration of adjacent wines.

The consequences of such framework go well beyond the operating model of such companies. First of all, the enhancement of the user experience may have a broader positive effect in reconnecting people with wine, especially by engaging younger generations which may enjoy a digital, gamified approach to wine. The segments of one enable extreme personalization by optimizing the user interface and by modulating the marketing mix according to the user profile. Being digital and modular by nature, platforms could continuously introduce new features and integrate with new technologies to make the shopping experience even more compelling, as in the case of the upcoming trend of contextual marketing. AR, VR, and mixed reality are still underdeveloped opportunities to communicate with the consumer a rich in attributes product such as wine. The rise of virtual sommeliers in the form of chatbots is also easily predictable and deployable in existing platforms.

Secondly, the insights that are generated in state-of-the-art platforms are of an unprecedented scale and scope compared to traditional research. In fact, these platforms are becoming real experimentation platforms, where valuable insights may be derived which could be employed by other actors in the supply chain. By integrating sensory science and chemical analysis into the functioning of such platforms, as in the case of Tastry, data integration across the whole supply chain may be promoted, enabling for the adoption of data-driven approaches at all levels. As the adoption of Industry 4.0 technologies in agriculture rises, the potential to implement in the productive process the user preferences is huge. Similarly, the potential to track user preferences and leverage them to optimize inventories and logistics is huge.

The function across which to integrate all these frameworks is *matching*. In fact, most of the marketing effort and data extraction are performed starting from this relatively simple operation. As granular representations of wines effectively represent digitally the characteristics of a wine, and as people tries and rates wines, uncovering their flavor preferences, the more effectively the machine learning algorithm can learn a preference profile that reproduces faithfully the taste of a consumer. When such profile is known, much of the segmentation criteria which are usually used to appraise such taste profile become accessories.

Matching wines and consumers poses a double choice: on one side, existing wines can successfully be matched with consumer tastes, while on the other side, by leveraging such data, new wines styles could be studied and tailored on consumers preferences. This dichotomy is open to the wine producer's interpretation, as their sensibility will determine to which approach to propend. In both cases consumer satisfaction may only increase. As New World countries are more inclined to design new wine styles, they may propend to leverage consumer preference data in order to reverse-engineer wine recipes, while in contrast with Old World countries, which are more focused on protecting typicity and regionality, will benefit of the matching function in order to find a consumer segment for their product.

Another relevant effect of platforms is the change in the traditional distribution chain, as online players gradually increase their share and as Direct-to-Consumer selling becomes an affordable alternative. Platforms such as NakedWines explicitly targeted at small independent producers enable them to gain visibility and to benefit from the infrastructure necessary for online DTC selling. More in general, platform marketplaces may disrupt traditional distributors role as they are guided by different principles and are cutting them out from the distribution chain. Considering that Italian producers still rank among their major concerns the access to international markets (Mediobanca 2020), this change is welcomed and may exert a positive effect in the internationalization of the wine market.

By reasoning on peculiar, consumer-focused principles, wine platforms may be a dramatic driver in promoting fairer competition between brands solely based on flavor.

The resulting framework of such considerations is a new representation of the wine industry in terms of a value ecosystem, where wine platforms occupy a central hub role.

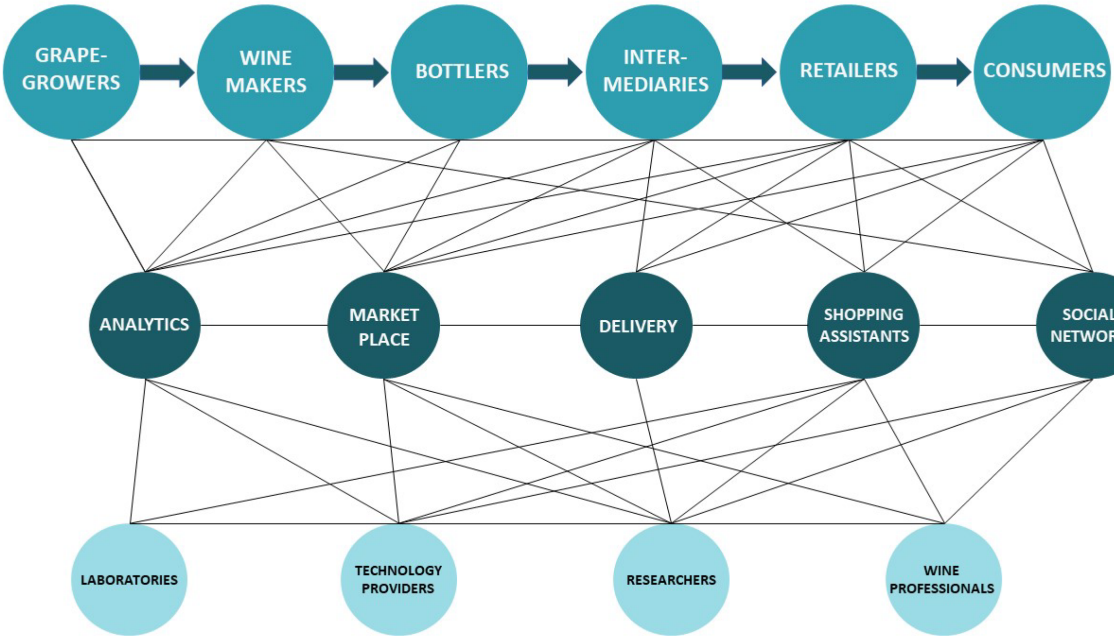

Figure 34. Representation of the digital wine ecosystem. Source: own elaboration.

Given their modularity and scope, wine platforms may position themselves along different point in the value ecosystem, in any case trying to connect with the wider number of agents possible.

The seamless flowing of data that is promoted by wine apps will result on a further integration of the supply chain, especially at the left hand, by integrating with production practices, and at the right end, by integrating with the shopping experience.

In this context, the scientific community and wine professionals are seen as fertilizers. The first ones may have a role in developing intelligent systems according to science-based approaches, while in the other direction researchers may use wine platforms data as a basis for research. Wine professionals will instead be asked to enrich the knowledge of platforms with ontologies and dictionaries, and take part in structuring the data and training the intelligent systems. On the opposite, they may use the platform output as a basis for their work.

The resulting role of platforms in focusing and accelerating wine research and innovation is consistent with the business ecosystem theory, and will eventually lead the wine industry to overcome critical challenges that await the industry.

CONCLUSIONS

The main objectives of this thesis, namely *(i) to examine the main approaches to digitally represent wine and consumer preferences* and *(ii) to assess how such digital representations are leveraged by digital business models in the wine industry*, have been thoroughly analyzed through Chapters 2 and 3.

It has been demonstrated how wine can be digitally represented through information in the form of text coming from the online environment, such as experts' or consumers' ratings and reviews. The growth in terms of users in online platforms has made it possible to generate extremely large unstructured databases of wine sensory notes to be leveraged to communicate wine, to analyze trends and in general to discover new knowledge about wine consumers' preferences. Similarly, effective techniques are now available to chemically analyze the entire wine metabolome, thus having a complete picture of its complex chemistry. These data coupled with sensory analysis techniques are enabling researchers to better understand the relationships between wine chemical composition and human sensory perceptions, including their hedonic responses and emotions. In this context, despite there's still a long way to go for to validate this hypothesis, it is now a concrete possibility to correlate consumers' preferences with specific patterns in wine chemistry, therefore better informing wine producers during the production process. On the other side, consumers could be offered wines they may like by leveraging data about their preferences on a scientific basis.

These sources of data have been leveraged by many startups which attempt to match individual consumer preferences (*palates*) to the bewildering range of wines available online. As discussed in Chapter 3, the industry leader Vivino leverages millions of data generated by its users, in the form of ratings and reviews, to generate an algorithmically efficient recommender system and a satisfying user experience. At the same time, other players like the American startup Tastry have adopted a science-driven approach, by analyzing the wines available in the market from a chemical and sensory point of view and then searching for relevant relationships with consumers' preferences.

While Vivino's approach to recommendation is suitable for its Amazon-like architecture, serving millions of consumers and millions of wineries and retailers, Tastry's approach is tailored to the single consumer and is ideal to be deployed locally. In fact, after their system will reach a critical mass by gathering enough data, Tastry's approach is by-design efficient at a granular level, as a new wine added to the platform does not need thousands of reviews to be effectively recommended and a consumer does not need to rate dozens of wines for the system to learn his preferences. The established relationships between wine chemistry and consumer preferences enable wine consumers to be confident in finding wines they will like, help wine producers to better target their audience, and support wine retailers in making better inventory assortments to satisfy the local demand. While both companies display an integrative approach along the supply chain, Tastry makes a step further by establishing a direct connection between both ends of the chain, enabling the data to flow seamlessly.

An additional objective of the thesis was to assess the impact at the industry-level of such digital business models. By looking at the overall picture of the wine industry two strategic themes clearly emerged, namely the role of *precision matching* and the emergence of a *digital wine value ecosystem*.

The precision matching between wines and consumers enabled by online wine platforms reshapes the wine supply chain from the ground up. Serving every customer as a *segment of one*, online wine platforms deliver the level of personalization and user experience that is expected in the digital age. By focusing primarily on consumer tastes and on how to drive consumer satisfaction, such platforms are initiating a flavor-based competition which represent a huge opportunity for shaking up the rules in market typically led by heritage, brands, and institutions. This represents an opportunity for both New World Wine producers and niche micro-wineries to be in the spotlight, as well as for natural wineries to serve the most curious wine enthusiasts. This does not mean, however, that Old World producers need to blindly follow flavor trends at the expense of their wine typicity. On the contrary, the improved segmentation enabled by wine platforms ideally permits to match every wine to its audience. In this scenario, Old World producers can thrive following two paths: by finding the right balance between typicity and market-driven tastes, or by exalting some differentiating flavor properties of their wines to target some identified niches of wine consumers.

We are going towards the direction of *flavor maps*, capable of characterizing the tastes of the wine drinkers at a local level and opening interesting scenarios for all the players willing to attack market niches. This approach of micro-segmenting the market is particularly fitting for small wine producers, as they can exploit differentiation and overcome the constraints of their small scale. Bigger producers on the other hand can increase the penetration of their diversified portfolio of brands with better market fit, guided by such data-driven segmentation of the market.

Big online wine platforms are increasingly occupying a central role in the wine supply chain, acting as gatekeepers and moving the wine market. By building impressive Amazon-like structures through an ecosystem of third parties, such platforms generate everyday massive amounts of data about wines and consumer preferences which can be leveraged by the whole industry. Due to their digital nature, they are an ideal platform for performing rapid experimentation and analyzing trends. This knowledge is then monetized by delivering services along the entire supply chain.

Ultimately, the digitalization in the wine industry has not only shortened the distance between producers and consumers, but also between every other actor in the supply chain. As consumers benefit from better user experience, wineries and wine brands optimize their market positioning, retailers maximize their sales, and new knowledge about the industry is continuously generated, the resulting landscape is that of a *digital wine value ecosystem* in the process of consolidation. By enabling novel business connections through data, online wine platforms had taken a hub role for promoting the co-creation of a value that is higher than the sum of the parts. The perspectives on the advancements of wine chemistry and AI applications and the deployment of intelligent technologies in the vineyard and the winery seem to suggest that precision matching will soon be followed by *precision viticulture and winemaking*, thus realizing an even stronger integration along the digital wine supply chain.

As a concluding remark, it can be said that even though the wine industry had never been perceived as being at the forefront of digital innovation, wine is by its intrinsic nature well-poised to experience the benefits of digitalization. In fact, there are today solid technologies available and urgent motivations for to embrace the digital transformation of the wine sector. The complex challenges of market fragmentation, increasing brand competition, sustainability, and generational change impose the industry to rethink itself, and digital has already proved to be an effective answer to manage complexity.

Winemaking has always represented a striking example of human ingenuity across centuries, as technique and technology have adapted to changes in culture, tastes, and environmental conditions.

Tradition and digital technology may be an effective mix for addressing current industry challenges, while assuring everyone is enjoying a good glass of wine.

BIBLIOGRAPHY

AGRAWAL, A., GANS, J., & GOLDFARB, A., 2016. The Simple Economics of Machine Intelligence. *Harvard Business Review* [online]. Available at <https://hbr.org/2016/11/the-simple-economics-of-machine-intelligence> [Accessed: 31/03/2022].

AGRAWAL, A., GANS, J., & GOLDFARB, A., 2018. A Simple Tool to Start Making Decisions with the Help of AI. *Harvard Business Review* [online]. Available at <https://hbr.org/2018/04/a-simple-tool-to-start-making-decisions-with-the-help-of-ai> [Accessed: 23/04/2022].

AMARGIANITAKY, M., & SPYROS, A., 2017. NMR-based metabolomics in wine quality control and authentication. *Chemical and Biological Technologies in Agriculture*. Vol. 4, (Article n. 9, 2017).

ARAPITSAS, P., et al., 2020. Use of Untargeted Liquid Chromatography–Mass Spectrometry Metabolome To Discriminate Italian Monovarietal Red Wines, Produced in Their Different Terroirs. *Journal of Agricultural and Food Chemistry*. Vol. 2020, 68, 47, pp. 13353-13366.

ANDERSON, K. & S. NELGEN, S., 2011. *Database of Global Wine Markets, 1961 to 2009* [online]. Available at <https://economics.adelaide.edu.au/wine-economics/databases#global-wine-markets-1961-to-2009-a-statistical-compendium> [Accessed: 23/04/2022].

ANDERSON, K., PINILLA V., & HOLMES, A.J., 2021. *Annual Database of Global Wine Markets, 1835 to 2019* [online]. Available at <https://economics.adelaide.edu.au/wine-economics/databases#annual-database-of-global-wine-markets-1835-to-2019> [Accessed: 03/06/2022].

ASIOLI, D., et al., 2017. Making sense of the "clean label" trends: A review of consumer food choice behavior and discussion of industry implications. *Food Research International*, 99 (1), pp. 58-71.

BAGNOLI, C., et al., 2018. *Business Model 4.0: I modelli di business vincenti per le imprese italiane nella quarta rivoluzione industriale* [online]. Available at <https://edizionicafoscari.unive.it/it/edizioni4/libri/978-88-696-287-1/> [Accessed: 06/07/2022].

BREIMAN, L., 2001. Random Forests. *Machine Learning*, 45 (1), pp. 5-32.

CALEGARI, R., CIATTO, G., & OMICINI, A., 2020. On the integration of symbolic and sub-symbolic techniques for XAI: A survey. *Intelligenza Artificiale*, 14 (1), pp. 7–32.

CALERO VALDEZ, A., ZIEFLE, M., & VERBERT, K., 2016. HCI for Recommender Systems: The Past, the Present and the Future. In: ASSOCIATION FOR COMPUTING MACHINERY, 2016. RecSys '16: Proceedings of the 10th ACM Conference on Recommender Systems. Boston, Massachusetts, 15-19/09/2016. New York: Association for Computing Machinery, pp. 123-126.

CAPITELLO, R., AGNOLI, L., CHARTERS, S. & BEGALLI, D., 2021. *Labelling environmental and terroir attributes: young Italian consumers' wine preferences* [online]. Available at: <https://www.sciencedirect.com/science/article/pii/S0959652621012105> [accessed: 23/02/2022].

CASTELLANO, S., et al., 2018. Uncovering the role of virtual agents in co-creation contexts: An application to the online wine business. *Management Decision*, 56 (6), pp. 1232-1246.

CHARTERS, S., & PETTIGREW, S., 2007. The dimensions of wine quality. *Food Quality and Preference*, Vol. 18 (7), pp. 997-1007.

CHEN, B., RHODES, C., CRAWFORD, A., HAMBUCHEN, L., 2014. *Wine Informatics: Applying Data Mining on Wine Sensory Reviews Processed by the Computational Wine Wheel* [online]. Available at: <https://ieeexplore.ieee.org/document/7022591> [accessed: 25/02/2022].

CORTEZ, P., et al., 2009. Modeling wine preferences by data mining from physicochemical properties. *Decision Support System.* Vol. 47 (2009), pp. 533-547.

COZZOLINO, D., 2016. Metabolomics in Grape and Wine: Definition, Current Status and Future Prospects. *Food Anal. Methods.* Vol. 9 (2016), pp. 2986-2997.

CROFTON, E. C., et al., 2019. Potential applications for virtual and augmented reality technologies in sensory science. *Innovative Food Science & Emerging Technologies*, 56 (102178), pp. 1-7.

DANNER, L., et al., 2017. "I like the sound of that!": Wine descriptions influence consumers' expectations, liking, emotions and willingness to pay for Australian white wines. *Food Research International*, 99 (1), pp. 263-274.

DELL'ERA, C., et al., 2020. Four kinds of design thinking: From ideating to making, engaging, and criticizing. *Creativity and Innovation Management*, 29 (2), pp. 324-344.

DONG, Z., XIAOWAN, G., SYAMALA, R., & BERNARD. C., 2020. Understanding 21st Century Bordeaux Wines from Wine Reviews Using Naïve Bayes Classifier. *Beverages,* 2020,6 (1),5.

DRESSLER, M., & PAUNOVIC, I., 2020. Converging and diverging business model innovation in regional intersectoral cooperation–exploring wine industry 4.0. *European Journal of Innovation Management*, 24 (5), pp. 1625-1652.

EMMOTT, S., et al., eds., 2006. *Towards 2020 Science*. Cambridge: Microsoft Research Ltd.

EVANS, D. S., 2003a. The Antitrust Economics of Multi-sided Platform Markets. *Yale Journal of regulations,* Vol. 20, pp. 325.

EVANS, D. S., 2003b. *Some Empirical Aspects of Multi-sided Platform Industries* [online]. Cambridge: NERA Economic Consulting. Available at <https://www.rnejournal.com/articles/evans_final_sept03.pdf> [accessed: 12/01/2022].

EVANS, P. C., & GAWER, A., 2016. *The Rise of the Platform Enterprise: A Global Survey* [online]. Available at <https://www.thecge.net/app/uploads/2016/01/PDF-WEB-Platform-Survey_01_12.pdf> [Accessed: 05/07/2022].

FINE DINING LOVERS, 2020. *Meet the Woman Behind the AI that Tastes Wine* [online]. Available at: <https://www.finedininglovers.com/article/tastry-ai-tastes-wine> [Accessed: 03/04/2022].

FRANCIS, I. L., & WILLIAMSON, P. O., 2015. Application of consumer sensory science in wine research. *Australian Journal of Grape and Wine Research*, 21 (S1), pp. 554-567.

FURNARI, C., 2021. *Uber To Acquire Drizly In $1.1 Billion Deal*. [online]. Available at <https://www.forbes.com/sites/chrisfurnari/2021/02/02/uber-to-acquire-drizly-in-11-billion-deal/> [Accessed: 10/04/2022].

GAWER, A., & CUSUMANO, M. A., 2013. *Industry Platforms and Ecosystem Innovation* [online]. Available at <https://onlinelibrary.wiley.com/doi/abs/10.1111/jpim.12105> [accessed: 16/06/2022].

GHEZZI, A., & CAVALLO, A., 2020. Agile Business Model Innovation in Digital Entrepreneurship: Lean Startup Approaches. *Journal of Business Research*, 110, pp. 519-537.

GIACOSA, A. et al., 2016. Mediterranean Way of Drinking and Longevity. *Critical Reviews in Food Science and Nutrition*, 56 (4), pp. 635-640.

GOODMAN, 2009. An international comparison of retail consumer wine choice. *International Journal of Wine Business Research*, 21 (1), pp. 41-49.

GOODFELLOW, I., BENGIO, Y., & COURVILLE, A., 2016. *Deep Learning* [online]. Available at <https://www.deeplearningbook.org/> [Accessed: 08/08/2022].

HAGIU, A., & ALTMAN, E., 2017. *Finding the Platform in Your Product* [online]. Harvard Business Review. Available at <https://hbr.org/2017/07/finding-the-platform-in-your-product> [accessed: 18/07/2022].

HUISMAN, G., BRUIJNES, M., HEYLEN, D. K. J., 2016. A Moving Feast: Effects of Color, Shape and Animation on Taste Associations and Taste Perceptions. In: ASSOCIATION FOR COMPUTING MACHINERY, 2016. *ACE '16: Proceedings of the 13th International Conference on Advances in Computer Entertainment Technology*. Osaka, Japan, 9-12/11/2016. New York: Association for Computing Machinery, article no. 13, pp. 1-12.

IANSITI, M., & LAKHANI, K. R., 2020. *Competing in the Age of AI: Strategy and Leadership when Algorithms and Networks Run the World*. Boston, Massachusetts: Harvard Business Review Press.

ISMEA, 2020. *Il mercato del vino In Italia e nel Mondo prima e dopo il Covid-19: Situazione, tendenze e sfide per il vino a livello mondiale* [online]. Available at <https://www.reterurale.it/flex/cm/pages/ServeBLOB.php/L/IT/IDPagina/22350> [Accessed: 08/08/2022].

IWSR, 2020. *Beverage alcohol ecommerce value grows by 42% in 2020, to reach US$24 billion* [online]. Available at <https://www.theiwsr.com/beverage-alcohol-ecommerce-value-grows-by-42-in-2020-to-reach-us24-billion/> [Accessed: 04/08/2022].

IWSR, 2021a. *Global wine trends to watch in 2021* [online]. Available at < https://www.theiwsr.com/global-wine-trends-to-watch-in-2021/ > [Accessed: 03/03/2022].

IWSR, 2021b. *5 key trends that will shape the global beverage alcohol market in 2021* [online]. Available at <https://www.theiwsr.com/5-key-trends-that-will-shape-the-global-beverage-alcohol-market-in-2021/> [Accessed: 03/03/2022].

IWSR, 2021c. *Health and sustainability cues boost the organic wine movement* [online]. Available at <https://www.theiwsr.com/health-and-sustainability-cues-boost-the-organic-wine-movement/> [Accessed: 03/03/2022].

IWSR, 2021d. *Increasing demand for nostalgic flavours in drinks* [online]. Available at <https://www.theiwsr.com/increasing-demand-for-nostalgic-flavours-in-drinks/> [Accessed: 29/04/2022].

IWSR, 2021e. *Led by the US, beverage alcohol ecommerce value expected to grow +66% across key markets 2020-2025* [online]. Available at <https://www.theiwsr.com/led-by-the-us-beverage-alcohol-ecommerce-value-expected-to-grow-66-across-key-markets-2020-2025/> [Accessed: 09/03/2022].

IWSR, 2022. *Ecommerce channel evolves into two overlapping worlds* [online]. Available at <https://www.theiwsr.com/ecommerce-channel-evolves-into-two-overlapping-worlds/> [Accessed: 09/03/2022].

JACKSON, R. S., 2014. *Wine Science: Principles and Applications*. 4th ed. London: Academic Press, pp. 1-17.

JACKSON, R. S., 2017. *Wine Tasting: A Professional Handbook*. 3rd ed. London: Academic Press, pp.1-16.

KENSTENBAUM, R., 2019. Direct-To-Consumer Selling Is Coming To Retail Categories You Couldn't Have Even Imagined. *Forbes* [online]. Available at: <https://www.forbes.com/sites/richardkestenbaum/2019/09/25/direct-to-consumer-selling-is-coming-to-retail-categories-you-couldnt-have-even-imagined/> [Accessed: 05/06/2022].

KETTMAN, M. 2019. Katerina Axelsson of Tastry on Teaching Computers to Taste Wine. *Wine Enthusiast* [online]. Available at <https://www.winemag.com/2019/04/04/katerina-axelsson-tastry-on-teaching-computers-to-taste-wine/ > [Accessed: 05/07/2022].

KING, E., et al., 2014. Regional sensory and chemical characteristics of Malbec wines from Mendoza and California. *Food Chemistry,* Vol. 143 (2014), pp. 256-267.

KOTLER, P., KARTAJAYA, H., & SETIAWAN, I., 2021. *Marketing 5.0: Technology for Humanity*. Hoboken, New Jersey: John Wiley & Sons Inc.

KOTONYA, N., DE CRISTOFARO, P., DE CRISTOFARO, E., 2018. *Of Wines and Reviews: Measuring and Modeling the Vivino Wine Social Network.* University College London.

LESSCHAEVE, I., 2007. Sensory Evaluation of Wine and Commercial Realities: Review of Current Practices and Perspectives. *Am. J. Enol. Vitic, 58:2 (2007),* pp. 252-257.

LIU, D., et al., 2019. From the Vineyard to the Winery: How Microbial Ecology Drives Regional Distinctiveness of Wine. *Frontiers in Microbiology*, 10 (2679), pp. 1-13.

LOCKSHIN, L., & CORSI, A.M., 2012. Consumer behaviour for wine 2.0: A review since 2003 and future directions. *Wine Economics and Policy,* Vol. 1, pp. 2-23.

LOCKSHIN, L., & CORSI, A. M., 2020. Key research topics likely to generate Australian and other wine producer countries' support during the period 2020-2030. *International Journal of Wine Business Research*, Vol. 32 (4), pp. 493-502.

LOCKSHIN, L., et al., 2006. Using simulations from discrete choice experiments to measure consumer sensitivity to brand, region, price, and awards in wine choice. *Food Quality and Preference,* Vol. 17 (2006), pp. 166-178.

LOCKSHIN, L., CORSI, A.M., 2019. *Consumer Research for Wine.* University of South Australia, Ehrenberg Bass Institute for Marketing Science.

LOTUS GROWTH, 2015. *The Influence of Mobile Apps on Wine Purchases* [online]. Available at <https://www.hellovino.com/research> [Accessed: 23/06/2022].

MALONE, T. W., RUS, D., & LAUBACHER, R., 2020. *Artificial Intelligence and the Future of Work* [online]. Available at <https://workofthefuture.mit.edu/wp-content/uploads/2020/12/2020-Research-Brief-Malone-Rus-Laubacher2.pdf> [Accessed: 24/04/2022].

MARR, B., 2019. The Incredible Ways The 4th Industrial Revolution And AI Are Changing Winemaking. *Forbes* [online]. Available at <https://www.forbes.com/sites/bernardmarr/2019/07/03/the-incredible-ways-the-4th-industrial-revolution-and-ai-are-changing-winemaking/> [Accessed: 05/06/2022].

MARR, B., 2021. Artificial Intelligence Can Now Taste – Transforming Winemaking With Tastry. *Forbes* [online]. Available at: <https://www.forbes.com/sites/bernardmarr/2021/06/28/artificial-intelligence-can-now-tastetransforming-winemaking-with-tastry/> [Accessed: 05/06/2022].

MARR, B., 2021. Vivino: Choose Your Next Great Wine With Big Data and Artificial Intelligence. *Forbes* [online]. Available at: <https://www.forbes.com/sites/bernardmarr/2021/05/07/vivino-choose-your-next-great-wine-with-big-data-and-artificial-intelligence/> [Accessed: 05/06/2022].

MASTROBERARDINO, P., CALABRESE, G., CORTESE, F., & PETRACCA, M., 2019. Sustainability in the wine sector. An empirical analysis of the level of awareness and perception among the Italian consumers. *British Food Journal,* Vol. 122 No. 8, 2020 pp. 2497-2511.

MCCARTHY, J., 2007. *What Is Artificial Intelligence?* [online]. Available at <http://jmc.stanford.edu/articles/whatisai/whatisai.pdf> [Accessed: 09/07/2022].

MEDIOBANCA, 2020. *Indagine sul settore vinicolo 2020* [online]. Available at < https://www.areastudimediobanca.com/it/product/indagine-sul-settore-vinicolo-2020> [Accessed: 23/07/2022].

MEDIOBANCA, 2022. *Il settore vinicolo in Italia* [online]. Available at <https://www.areastudimediobanca.com/it/product/il-settore-vinicolo-italia-ed-2022> [Accessed: 12/09/2022].

MICHELI, P., et al., 2019. Doing Design Thinking: Conceptual Review, Synthesis, and Research Agenda. *Journal of Product Innovation Management*, 36 (2), pp. 124-148.

MITCHELL, T., 1997. *Machine Learning*. New York: McGraw-Hill, pp.1-126.

MORA, M., et al., 2021. Exploring young consumers' attitudes and emotions to sensory and physicochemical properties of different red wines. *Food Research International,* Vol. 143 (2021) 110303.

NAPIER, L., et al., 2020. *Modern Business Models Will Drive the Post-Pandemic World* [online]. MIT Sloan Management Review. Available at

<https://sloanreview.mit.edu/article/modern-business-models-will-drive-the-post-pandemic-world/> [accessed: 21/06/2022]

NIIMI, J., DANNER, L., & BASTIAN, S. E. P., 2019. Wine leads us by our heart not our head: emotions and the wine consumer. *Current Opinion in Food Science*, 27, pp. 23–28.

NOMISMA, 2021. *L'e-commerce nel vino prima e dopo l'emergenza coronavirus: un'analisi dell'Osservatorio Nomisma Wine Monitor* [online]. Available at <https://www.nomisma.it/ecommerce-vino-dati-da-nomisma-wine-monitor/> [Accessed: 04/07/2022].

NOSI, C., MATTIACCI, A., & SFODERA, F., 2019. Online wine ecosystem: the digital narrative of Sangiovese. *British Food Journal*, Vol. 121 No. 11, 2019, pp. 2683-2695.

OIV, 2018. *State of the vitiviniculture world market* [online]. Available at <https://www.oiv.int/public/medias/5958/oiv-state-of-the-vitiviniculture-world-market-april-2018.pdf> [Accessed: 08/07/2022].

OIV, 2021a. *State of the world vitivinicultural sector in 2020* [online]. Available at <https://www.oiv.int/en/technical-standards-and-documents/statistical-analysis/state-of-vitiviniculture> [Accessed: 08/07/2022].

OIV, 2021b. *Digital trends applied to the vine and wine sector: A comprehensive study on the digitalisation of the sector* [online]. Available at <https://www.oiv.int/en/oiv-life/digitalisation-the-future-of-the-vine-and-wine-sector> [Accessed: 08/07/2022].

PÉREZ-JIMÉNEZ, M., et al., 2021. Application of untargeted volatile profiling and data driven approaches in wine flavoromics research. *Food Research International*, 145 (110392), pp. 1-15.

PINU, F. R., 2018. Grape and Wine Metabolomics to Develop New Insights Using Untargeted and Targeted Approaches. *Fermentation*, 2018, 4.

POMARICI, E., BOCCIA, F., & CATAPANO, D., 2012. The wine distribution systems over the world: an explorative survey. *New Medit,* 11 (4), pp. 23-32.

POMARICI, et al., 2017. One size does (obviously not) fit all: Using product attributes for wine market segmentation. *Wine Economics and Policy,* Vol. 6 (issue 2), pp. 98-106.

POMARICI, E. et al., 2021. The Italian Wine Sector: Evolution, Structure, Competitiveness and Future Challenges of an Enduring Leader. *Italian Economic Journal*, 7, pp. 259-295.

PROBLEM2VALUE, 2017. *When, which ... Design Thinking, Lean, Design Sprint, Agile?* [online]. Available at <https://www.problem2value.com/blog/when-which-design-thinking-lean-design-sprint-agile.html> [Accessed: 08/07/2022].

PWC, 2020. *Global Top 100 companies by market capitalization. Update to 30 June 2020* [online]. Available at <https://www.pwc.com/gx/en/audit-services/publications/assets/global-top-100-companies-june-2020-update.pdf> [accessed: 24/03/2022].

RABOBANK, 2021. *2022 Alcohol E-commerce Playbook: The Size of the Prize* [online]. Available at <https://research.rabobank.com/far/en/sectors/beverages/2022-alcohol-e-commerce-playbook.html> [Accessed: 20/08/2022].

RICCI, F., et al., eds., 2011. *Recommender Systems Handbook*. New York: Springer Science+Business Media, pp.1-186, 257-297; 479-510.

RICCI, F., ROKACH, L., & SHAPIRA, B., 2015. Recommender Systems: Introduction and Challenges. In: RICCI, F., et al., eds., 2015. *Recommender Systems Handbook*. 2nd ed. New York: Springer Science+Business Media, pp. 1-34.

RISTIC, R., et al., 2018. Wine-related aromas for different seasons and occasions: Hedonic and emotional responses of wine consumers from Australia, UK and USA. *Food Quality and Preference*, Vol. 71, pp. 250-260.

ROCHET, J. C., & TIROLE, J., 2003. Platform Competition in Two-Sided Markets. *Journal of the European Economic Association,* Vol.1, issue 4, pp. 990-1029.

ROODERKERK, R. P., & KÖK, A. G., 2019. Omni-Channel Assortment Planning. In: GALLINO, S., & MORENO, A., eds., 2019. *Operations in an Omnichannel World*. Cham: Springer Nature Switzerland AG, pp. 51-86.

SAENZ-NAVAJAS, M. P., BALLESTER, J., PEYRON, D., VALENTIN, D., 2014. Extrinsic attributes responsible for red wine quality perception: A cross-cultural study between France and Spain. *Food Quality and Preference,* Vol. 35 (2014), pp. 70-85.

SAMUEL, A. L., 1959. Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development*, 3 (3), pp. 210-229.

SCHNEIDER, J., 2017. *Understanding Design Thinking, Lean, and Agile*. [online]. Available at <https://www.oreilly.com/library/view/understanding-design-thinking/9781491998410/> [Accessed: 18/06/2022].

SENYO, P. K., LIU, K., & EFFAH, J., 2019. Digital business ecosystem: Literature review and a framework for future research. *International Journal of Information Management*, Vol. 47, pp. 52-64.

SHEPERD, G. M., 2015. Neuroenology: how the brain creates the taste of wine. *Flavour*, 4 (19), pp. 1-5.

SLAGHENAUFI, D., et al., 2021. Two Sides to One Story - Aroma Chemical and Sensory Signature of Lugana and Verdicchio Wines. *Molecules*, Vol. 2021, 26, 2127.

SMITH, B., 2019. Getting More Out of Wine: wine experts, wine apps and sensory science. *Current Opinion in Food Science*, 27, pp. 123-129.

SOMALVICO, M., 1992. *Intelligenza Artificiale* [online]. Available at <https://www.treccani.it/enciclopedia/intelligenza-artificiale_%28Enciclopedia-Italiana%29/> [Accessed: 08/09/2022].

SOUZA GONZAGA, L., et al., 2020. Defining wine typicity: sensory characterisation and consumer perspectives. *Australian Journal of Grape and Wine Research. Vol* 27 (2021), pp. 246–256.

SPENCE, C., et al., 2017. Digitizing the chemical senses: Possibilities & pitfalls. *International Journal of Human-Computer Studies*, 107, pp. 62-74.

SPENCE, C., 2019. Multisensory experiential wine marketing. *Food Quality and Preference*, Vol. 71, pp. 106-116.

SPIELMANN, N., CHARTERS, S., 2013. The dimensions of authenticity in terroir products. *International Journal of Wine Business Research*, Vol. 25 No. 4, pp. 310-324.

SPIELMANN, N., GELINAS-CHEBAT, C., 2012. Terroir? That's not how I would describe it. *International Journal of Wine Business Research*, Vol. 24 No. 4, pp. 254-270.

STATISTA, 2021. *Wine market revenue worldwide from 2012 to 2025* [online]. Available at <https://www.statista.com/statistics/922403/global-wine-market-size/> [Accessed: 28/07/2022].

SWENEY, M., 2021. *Naked Wines benefits from Covid online boom as sales soar* [online]. Available at <https://www.theguardian.com/business/2021/jun/11/naked-wines-benefits-from-covid-online-boom-as-sales-soar> [Accessed: 21/07/2022].

THACH, L., 2021. Vivino Continues Watershed Year In Online Wine Sales With $155 Million In Funding And CEO Appointment. *Forbes* [online]. Available at: <https://www.forbes.com/sites/lizthach/2021/06/03/vivino-continues-watershed-year-in-online-wine-sales-with-155-million-in-funding-and-ceo-appointment/> [Accessed: 05/06/2022].

THACH, L., & CHANG, K., 2015. 2015 Survey of American Wine Consumer Preferences. *Wine Business* [online]. Available at <https://www.winebusiness.com/news/article/160722/> [Accessed: 28/05/2022].

THACH, E. C., OLSEN, J.E., 2006. Market Segment Analysis to Target Young Adult Wine Drinkers. *Agribusiness*, Vol. 22 (3), pp. 307–322.

VALENTE, C. C., et al., 2018. Modelling the sensory space of varietal wines: Mining of large, unstructured text data and visualisation of style patterns. *Scientific Reports*, 8 (4987), pp. 1-13.

VAN ALSTYNE, M. W., PARKER, G. G., & CHOUDARY, S. P., 2016. *Pipelines, Platforms, and the New Rules of Strategy* [online]. Harvard Business Review. Available at <https://hbr.org/2016/04/pipelines-platforms-and-the-new-rules-of-strategy> [accessed: 24/03/2022].

VAN LEEUWEN, C., & DARRIET, P., 2016. The Impact of Climate Change on Viticulture and Wine Quality. *Journal of Wine Economics*, 11 (1), pp. 150-167.

VAN LEEUWEN, C., et al., 2020. Recent advancements in understanding the terroir effect on aromas in grapes and wines. *OENO One*, 54 (4), pp. 985-1006.

WATERHOUSE, A. L., SACKS, G. L., & JEFFERY, D. W., 2016. *Understanding Wine Chemistry*. Chichester, West Sussex: John Wiley & Sons, Inc., Introduction, Chapter 19, Chapter 32.

WINE INTELLIGENCE, 2020. *Global Trends in Wine 2020* [online]. Available at <https://www.wineintelligence.com/downloads/global-trends-in-wine-2020-report-updated-and-released-as-open-source/> [Accessed: 08/07/2022].

WINE INTELLIGENCE, 2021a. *Wine e-commerce's huge opportunity in 2021* [online]. Available at <https://www.wineintelligence.com/wine-e-commerces-huge-opportunity-in-2021/> [Accessed: 07/05/2022].

WINE INTELLIGENCE, 2021b. *How will consumer needs prompt wine ecommerce models to evolve post covid-19?* [online]. Available at <https://www.wineintelligence.com/how-will-consumer-needs-prompt-wine-ecommerce-models-to-evolve-post-covid-19/> [Accessed: 07/05/2022].

WINE INTELLIGENCE, 2022. *How wine can keep its lead in e-commerce* [online]. Available at <https://www.wineintelligence.com/how-wine-can-keep-its-lead-in-e-commerce/> [Accessed: 07/05/2022].

YANG, J., & LEE, J., 2020. Current Research Related to Wine Sensory Perception Since 2010. *Beverages*, 6 (47), pp. 1-18.

ZENG, M., 2018. Alibaba and the Future of Business. *Harvard Business Review* [online]. Available at <https://hbr.org/2018/09/alibaba-and-the-future-of-business> [Accessed: 18/01/2022].

ZHU, F., IANSITI, M., 2012. Entry Into Platform-based Markets. *Strategic Management Journal,* Vol. 33, issue 1, pp. 88-106.

WEBOGRAPHY

https://www.vivino.com/IT/en/simone-capecci-picus-rosso-piceno-superiore/w/1128434

https://dot-farm.net/tvs/en/

https://tactic.studio/living-wine-labels

https://www.rga.com/work/case-studies/a-sip-for-the-senses

https://www.krug.com/it/playlist/krug-echoes https://www.fossanalytics.com/en/products/winescan

https://www.winemag.com/wine-tech/

https://www.prestigeonline.com/th/pursuits/tech/winewall-ai-wine-celllar/

https://www.strategyzer.com/canvas/business-model-canvas

https://apps.vinocell.com/

https://vindome.net/

https://www.enolytics.com/

https://dictionary.cambridge.org/dictionary/english/algorithm

https://www.microsoft.com/en-us/ai/responsible-ai

https://www.vivino.com/

https://www.nakedwines.com/

https://drizly.com/

https://www.winearomawheel.com/

https://www.winespectator.com/

https://www.vivino.com/about

https://www.vivino.com/wine-news/vivino-5-star-rating-system

https://agriculture.ec.europa.eu/farming/geographical-indications-and-quality-schemes/geographical-indications-and-quality-schemes-explained_en#geographicalindications

https://www.kaggle.com/

https://pair.withgoogle.com/

APPENDIX: Vinho Verde Dictionary

peach	caramel	bright	edgy
melon	candy	light	smooth
apple	ripe	dry	yellow
green apple	almond	rich	balanced
citrus	floral	creamy	juicy
tropical	fruity	sweet	round
lemon	yeasty	dense	intensity
lime	aromatic	mineral	perfect
grapefruit	zest	tangy	attractive
pineapple	spice	off-dry	delicious
pear	acidity	soft	
apricot	crisp	clean	
orange	prickle	summer	
plum	fresh	fine	

Source: own elaboration.