Development of a Generative AI-Based Model for Guiding Grape Variety Selection Under Contemporary Climate Dynamics

Joel Z. Harms

Department of Bioresource Engineering, McGill University, Montreal
Submitted July 2024

English Abstract. Utilizing the ideal grape varieties in the environment and terroir where they stand to produce the highest quality is common in long-established viticultural regions. It has been formalized by demarcating appellations and growths where the best grapes are grown. Changes in the climate, however, may alter which varieties are better grown and may have the highest potential in a certain region. In addition, grape cultivation is still expanding to new areas that do not have an extensive history of grapes being grown for wine. Utilizing the available grape variety diversity estimated between about 1300-1600 commercially used grape varieties would allow growers to adapt to changing climates and to establish new quality regions more quickly. However, guidance is needed since the number of choices and climates, as well as market conditions and consumer demand, make this a complex problem. A generative AI architecture for climate-based variety recommendation is developed, tested, and applied to current and future climate conditions to give data-based guidance for growers in a changing climate and new wine development. This novel varietal recommendation system ranks 1300 commercial grape varieties for a given climate, generating predicted popularity and suitability scores. With this tool, trials can be directed to varieties that are already likely to perform well, reducing the risks associated with choosing new grape varieties to plant. This is critical, as new plantings require significant capital and time investments. Furthermore, it is possible to extract climate indices from future climate projections to make predictions about the popularity and suitability of varietals globally, allowing for current guidance and long-term planning. The novelty of this tool is being the first crop recommender system that ranks varieties and assigns suitability scores. It can do this for 1300 grape varieties simultaneously, whereas the few examples of previous crop recommender systems only recommended one crop and only for a couple of dozen crops. This tool is furthermore not limited to a region, as it can predict popularity and suitability globally for present and future climates.

French Abstract.

L’utilisation des cépages idéaux dans l’environnement et le terroir où ils se trouvent pour produire la meilleure qualité est courante dans les régions viticoles établies de longue date. Elle a été formalisée en délimitant les appellations et les crus où sont cultivés les meilleurs raisins. Les changements climatiques peuvent cependant modifier les variétés les mieux cultivées et susceptibles d’avoir le plus grand potentiel dans une certaine région. En outre, la culture de la vigne continue de s’étendre à de nouvelles zones qui n’ont pas une longue histoire de culture de la vigne pour le vin. L’utilisation de la diversité des cépages disponibles, estimée entre environ 1 300 et 1 600 cépages utilisés commercialement, permettrait aux producteurs de s’adapter aux changements climatiques et d’établir plus rapidement de nouvelles régions de qualité. Cependant, des conseils sont nécessaires car le nombre de choix et de climats, ainsi que les conditions du marché et la demande des consommateurs, rendent ce problème complexe. Une architecture d’IA générative pour la recommandation de variétés basée sur le climat est développée, testée et appliquée aux conditions climatiques actuelles et futures pour fournir des conseils basés sur des données aux producteurs dans un climat changeant et le développement de nouveaux vins. Ce nouveau système de recommandation variétale classe 1300 cépages commerciaux pour un climat donné, générant des scores de popularité et d'adéquation prévus. Grâce à cet outil, les essais peuvent être orientés vers des variétés qui sont déjà susceptibles de bien se comporter, réduisant ainsi les risques associés au choix de nouveaux cépages à planter. Cela est essentiel, car les nouvelles plantations nécessitent des investissements importants en capital et en temps. En outre, il est possible d'extraire des indices climatiques à partir de projections climatiques futures pour faire des prévisions sur la popularité et l'adéquation des cépages à l'échelle mondiale, ce qui permet de disposer de conseils actuels et de planifier à long terme. La nouveauté de cet outil est d'être le premier système de recommandation de cultures qui classe les variétés et attribue des scores d'adéquation. Il peut le faire pour 1300 cépages simultanément, alors que les quelques exemples de systèmes de

recommandation de cultures précédents ne recommandaient qu'une seule culture et seulement pour quelques dizaines de cultures. Cet outil ne se limite pas non plus à une région, car il peut prédire la popularité et l'adéquation à l'échelle mondiale pour les climats actuels et futurs.

Acknowledgments

"The race doesn't go to the swift,
Nor the battle to the strong,
Nor food to the wise,
Nor wealth to the smart,
Nor recognition of the skilled.
Instead, timing and circumstances meet them all."
Ecclesiastes 9;11

In the theme of finding the right climatic circumstances for specific grapevines I must of course also acknowledge the timing, circumstances out of my control, and the helping hands that have allowed me to complete this degree. Firstly, I must acknowledge the provision and leadership of my God and Christ Jesus, who holds my life and therefore also my academic career in his hands. Secondly, I must acknowledge the love, patience, and support I have received from my mother and my younger sister throughout my academic journey. Without my younger sister's guidance in writing and presenting this Thesis would still read like a high-school assignment. Thirdly, of course, I need to thank my community in the form of the Arabic Evangelical Baptist Church of Montreal for their continuous prayers and encouraging words.

Then of course the guidance and help of my colleagues and supervisors as well as committee members must also be mentioned. I want to thank Dr. Israel (Jacob) Liberty, Dr. Christopher Nzediegwu, Prof. Shiv Prasher, and Prof. Amit Gross for initiating me into research and allowing me to learn from them so much. I want to thank Dr. Julien J. Malard-Adam for his guidance and trust, allowing me to work alongside him and guiding me to publish my first article. I want to thank Prof. Jan Adamowski for his continued trust and for allowing me to not only pursue an M.Sc. and soon Ph.D. degree with him but also for giving me the chance to expand my experience in lecturing among others. I want to thank Dr. Nathaniel Newlands and Prof. Castellarin through whose support I was able to work in the field of Viticulture though it is outside of the usual research areas of McGill University. I also want to thank my co-supervisor Prof. Viacheslav Adamchuk who always offers timely advice.

I also want to acknowledge the lovely colleagues that get to interact with and who take time to meet and talk with me, who make every workday brighter, these are Birkhoff, Theo, Matt, Erik, Pierre-Luc, Ali, Vincent, Derrick, Raphael, Anthony, Qi, Arusha, and Ehsan.

Contributions

The supervision of this Thesis was conducted by Prof. Adamowski, and co-supervised by Prof. Adamchuk, the committee consisted of Dr. Newlands and Prof. Castellarin who were also involved in the work assisting with conceptualization and planning of the work together with the supervisors. Prof. Sun has also held a position on the committee to assist in tracking progress. I conceptualized, programmed, and developed the models as well as wrote this Thesis. Some editorial assistance was provided by Jemima Harms and Prof. Adamchuk. The conceptualization and initial trial stage of this research was supported by a Canada Graduate Scholarship – Master's Program from the Natural Sciences and Engineering Research Council of Canada.

Contents

List of Tables

List of Figures

List of Abbreviations

GST – Growing Season Temperature
AnnP – Annual Precipitation
GSP – Growing Season Precipitation
HMP – Harvest Month Precipitation
AnnT – Annual Average Temperature
RPT – Ripening Period Temperature
MJT – Mean January/July Temperature
GDD – Growing Degree Days
HI – Huglin Index
GSDTR – Growing Season Diurnal Temperature Range
RPDTR – Ripening Period Diurnal Temperature Range
CNI – Cold Night Index
VPD_GS – Average Vapour Pressure Deficit over the Growing Season
VPD_SU – Average Vapour Pressure Deficit over the Summer
SRAD_GS – Average Downward Shortwave Radiation over the Growing Season
SRAD_SU – Average Downward Shortwave Radiation over the Summer
ML – Machine Learning
DL – Deep Learning
AE – Auto-Encoder
ECG – Electrocardiogram
MRI – Magnetic Resonance Images
CAE – Coupled-Auto-Encoder
ANN – Artificial Neural Network
XGBoost – Extreme Gradient Boosting Machine
TPE - Tree-structured Parzen Estimator
POWER – Prediction of Worldwide Energy Resources
GDEM – Global Digital Elevation Model
ASTER – Advanced Spaceborne Thermal Emission and Reflection Radiometer
MSE – Mean Squared Error
NDCG – Net-Discounted-Cumulative-Gain
RMSE – Root Mean Squared Error
auPRC – area under Precision-Recall Curve
IDW – Inverse Distance Weighted Interpolation
SSP126 – Shared Socioeconomic Pathway 1: “Taking the Green Road”
SSP585 – Shared Socioeconomic Pathway 5: “Taking the Highway”

1 Introduction and Literature Review

1.1 The Effect of Climate on Grape Cultivation

The climate is part of the so-called ‘terroir’ concept (van Leeuwen, 2022). Along with soil properties and management practices the local climate helps to elicit specific and valued ‘local’ qualities from the grapes grown and therefore in the wine made from them (Jones, 2018; van Leeuwen, 2022; van Leeuwen et al., 2012, 2018). A changing climate means that some of these attributes will continue to shift as they are already shifting (van Leeuwen & Darriet, 2016). In major viticultural regions, this can largely be measured by the change in harvest dates to be generally earlier than they have been in the past (van Leeuwen & Darriet, 2016). The advancing of harvest days is linked to the temperature characteristics during the growing season. Precipitation changes are also likely to cause differences in irrigation regimes in the future (Prada et al., 2024). Where vineyards are not already irrigated precipitation shifts may enhance stress levels in the plants through drought, alternatively where precipitation is increasing during the growing season disease pressure may increase (Prada et al., 2024). This however, does not necessitate the future unsuitability of current wine-growing regions, in general, as grape growers have many tools to adapt to changing conditions, such as through management, by adjusting irrigation regimes, pruning and shading among others (Reshef et al., 2017; van Leeuwen et al., 2013). They may also be able to select varieties more suitable to their current and future conditions that originate from other regions or have been bred for resistance to certain conditions (Wolkovich et al., 2018). The selection of varieties that are available to grape growers is exceptional, Robinson et al., (2013) list 1368 grape varieties that are cultivated for wine-making globally; the database developed by Anderson & Nelgen, (2020) names 1595 grape varieties. The majority of the world grape growing regions, however, are dominated by a few dozen of the so-called ‘international varieties’, these are most often traditional French varietals from well-known regions such as Burgundy and Bordeaux which are known for their quality and are hence emulated globally (Johnson & Robinson, 2001). These varietals can be found in some proportion in most of the world's grape-growing regions, especially in the New World since it does not have native 'noble grapes', which is a title given only to varieties of *Vitis vinifera* (Johnson & Robinson, 2001; Robinson et al., 2013). The majority of the world's wine production is through these grapes, creating a lack of diversity and adaptability of wine regions to changing conditions.

Grape diversity is enormous as mentioned before and even continues to expand. Since roughly 200 years ago, to adapt grape growing to the novel disease pressure that appeared due to imported plant material from the New World, interspecific hybrids were developed that aim to maintain the quality of 'noble varieties' while incorporating the disease resistance of native North-American or East-Asian varieties (Johnson & Robinson, 2001; Robinson et al., 2013). Many of the native New World and hybrid varieties now find their main use as rootstocks for wine production and have 'noble varieties' grafted to them so that the benefit of the noble fruit and resistant roots is gained (Johnson & Robinson, 2001; Robinson et al., 2013). Varieties are also developed for cold resistance to be grown in cold climates like Quebec or North Dakota (Robinson et al., 2013). Some wine regions use hybrids to produce wine, such as Quebec or Nova Scotia, though most wine regions do not use these varieties at all since they are usually regarded as producing wine of lesser quality compared to noble grapes (Johnson & Robinson, 2001; Robinson et al., 2013).

The other origin of grape diversity is the countless regionally important noble grape varieties that have been continuously used for many decades and adapted to their regional climatic conditions. Specialized varieties of this type generally come from the Caucasus, North Africa, Europe, or Western Asia, which is roughly the native distribution of the *Vitis vinifera* species (Robinson et al., 2013). These well-tested varieties are of particular interest for quality wine production as climate change shifts wine-growing regions (Wolkovich et al., 2018). However, since the specialized noble varieties from 'lesser known' viticultural regions are adapted to specific conditions, it will be important to match the correct varieties to the correct climatic conditions to guide selection. There is also the possibility that for specific wine-growing regions hybrids may be the best way forward, either way, guidance is required to make the most of the available diversity and to adapt winemaking for the future.

1.1.1 Guiding Variety Selection

The selection of the ideal variety for a specific terroir, even if only considering the climatic component, is no small feat. Still, it may allow for a synergistic relationship between the vine and the place which only occurs in very specific conditions (Jones, 2018). Not only are the climate requirements highly specific, but these requirements are not known for all varieties. The selection is complicated by the variability in sensitivity between varieties such that some are very climate-dependent while others are much more stable in the quality of produced wine (Davis et al., 2019). For these reasons, studies are usually limited in the number of grapes they consider. Fraga et al., (2016) for example evaluate 44 Portuguese varieties, and Mavromatis et al., (2020) evaluate 29 Greek varieties. An approach is needed that can truly help growers guide which of the 1300-1600 cultivated grape varieties they may want to make use of for their climatic conditions. This requires changes to the past approaches in this field, namely that many grape-growing regions and the majority of the grape varieties used today need to be considered in a standardized way, in other words, a global approach is required. Localized studies will by design be very limited in the number of grapes that can be assessed and the range of climate conditions that projections are made for. To develop a global tool, a global approach is also necessary, fortunately, data relating to which regions grow which amounts of which grape variety is provided by Anderson & Nelgen, (2020). Furthermore, a global approach increases the data that can be used to more specifically narrow down the optimal conditions for each grape variety, so that most grape varieties will be able to be somewhat understood and can thus be contextualized among the large body of grapevine diversity. What this requires is a good selection of climatic indices to allow for the characterization of the grape varieties by the climatic conditions in which they are grown.

Skahill et al., (2022, 2023) focus on using a linear approach and classify the suitability by only one climatic variable the growing season temperature (GST). The GST is a good indicator of quality and can give general information about the potential of a grape in a certain location. However, the climatic variables that affect wine quality are not the same for each variety, even the precise timing of climatic conditions affects different varieties differently (Davis et al., 2019). Therefore, grape variety selection should include many climatic factors if it is to be useful to rank all or most cultivated grape varieties. This suggests that a high-dimensional climate space needs to be considered as was done by Puga et al., (2022) who used 16 climate variables and derived indicators to classify the world's wine regions. A subset of these variables is used by Hall & Jones, (2010) to classify Australia's wine regions. Puga et al., (2022) use the 16 indicators to classify 813 viticultural regions from the dataset by Anderson & Nelgen, (2020) which is as close to the global scale as possible, all that is missing is to connect this meaningfully with the grape varieties grown in these regions. The 16 climatic indicators that were considered as well as their significance and ranges are shown in the table below which is based on (Hall & Jones, 2010; Hewer & Gough, 2021; Puga et al., 2022).

Table 1: The climatic indices useful for global climate classification may also be sufficient to classify the suitability of the same regions' grape varieties. This table is mainly derived from Hall & Jones, (2010) as well as (Puga et al., 2022) and (Hewer & Gough, 2021).

Variable	Abbr.	Required Variables	Significance
Annual Precipitation	AnnP	Precipitation at all time scales	The amount of water naturally available to the vine during a year.
Growing Season Precipitation	GSP	Monthly or daily precipitation	The amount of water naturally available to the vine during the period it is most needed but also

			when it is susceptible to disease because of precipitation.
Harvest Month Precipitation	HMP	Monthly or daily precipitation	The amount of water that can disrupt harvest, can cause harvest loss and diluted flavors.
Annual Average Temperature	AnnT	Mean temperature at any time scale	The prevailing temperature conditions across the year.
Growing Season Average Temp.	GST	Monthly or daily mean temperature	The temperature during the time of major vine growth and fruit building is a predictor of ripening ability (Jones et al., 2005).
Ripening Period Average Temp.	RPT	Monthly or daily mean temperature	The temperature during the final stages of ripening.
Mean January/July Temp.	MJT	Monthly or daily mean temperature	The temperature is during the hottest month and usually during the beginning of the ripening process.
Growing Degree Days	GDD	Daily max and min temperature	The amount of heat available for the entire ripening process, this index is also known as the Winkler index (Winkler, 1974).
Huglin Index	HI	Daily max and mean temperature as well as latitude	The amount of heat available for ripening corrected for the latitude and therefore day length. It also only considers a shorter growing season and ends in September/March instead of October/April as GDD does.
Growing Season Diurnal Temperature Range	GSDTR	Daily max and min temperature	The variability between day and night temperatures during the growing season has a major impact on the flavor development of some varieties (Davis et al., 2019).
Ripening Period Diurnal Temperature Range	RPDTR	Daily max and min temperature	Similar to the above but limited to the final ripening stages.

Cool Night Index	CNI	Monthly or daily minimum temperature	The lowest temperature in the assumed harvest month hints at the potential of the grapes to develop secondary metabolites important for color and aroma (Tonietto & Carbonneau, 2004).
Growing Season Vapor Pressure Deficit	VPD_GS	Monthly or daily temperature and relative humidity	A proxy to potential water stress conditions that may lead to the development of additional flavors (Kovalenko et al., 2021), or cause stress over the entire growing season.
Summer Vapor Pressure Deficit	VPD_SU	Monthly or daily temperature and relative humidity	Similar to above but only over the hottest months (June to August or December to February).
Growing Season Average Downward Surface Shortwave Radiation	SRAD_GS	Monthly or daily all-sky downward surface shortwave radiation	Another dimension of plant stress is related to plant stomatal conductance, vapor pressure deficit, and temperature (Gowdy et al., 2022). Radiation also affects metabolic processes and flavors (Reshef et al., 2017).
Summer Average Downward Surface Shortwave Radiation	SRAD_SU	Monthly or daily all-sky downward surface shortwave radiation	Similar to the above but only considering the hottest months of the year.

Alternative climate indices are also presented by Hewer & Gough, (2021), many covering similar realms in terms of their effect on grapes and grape quality. However, one different realm not covered in Table 1 above is included through variables such as extreme cold days and potential frost days. Cold damage is a major concern and dictates the winter survivability of various grape varieties, and is one of the major limitations of where grapes can be cultivated in the first place which tends to be between the 30th and the 50th parallels (Aney, 1974; Johnson & Robinson, 2001; Robinson et al., 2013). Of particular interest to whether grape cultivation is even a possibility is the 'extreme cold days' index (Aney, 1974). Grape varieties will generally die in temperatures below -20 degrees Celsius (Hewer & Gough, 2021). Specific cold-hardiness is of course also variety dependent (Robinson et al., 2013). In contrast to the other climate variables, it does not necessarily relate to the quality of grapes produced but generally, it is a pre-requisite of whether grapes can grow at all. Learning cold hardiness from the global varietal datasets is a bit difficult as many methods exist to protect grapes from cold damage, allowing them to be grown in many regions that could be considered unsuitable when looking at the climate data without the adaption measures. Other than quality data, the cold hardiness is easily established as it can be directly measured objectively, whereas quality and the terroir characteristics of wine are much more high dimensional problems that could not be solved by a simple threshold. For this reason in this research, the focus will be on the terroir characteristics and not the general viticultural potential as adaption measures exist and since this type of suitability can be easily established by means other than artificial intelligence or other advanced data-science methods such as through thresholds (Hewer & Gough, 2021). Growers should be aware of the general viticultural potential and the cold hardiness of their grapes, but this information is more available and where not, easily established by laboratory or field trials (Howell, 2001). For these reasons cold-hardiness is not considered in this study.

1.2 Crop Recommendations Systems

Crop recommendation systems are not completely novel in general they exist in some form for a few dozen crops. However, current approaches are very limited not only in scope but also in the utility of the approach to the end-user. Recent examples of crop recommendation systems come from a Kaggle competition, based on a dataset from the Department of Agriculture of India (Garanayak et al., 2021; Gopi & Karthikeyan, 2023; Islam et al., 2023). This dataset has also been supplemented by additional local datasets (Musanase et al., 2023). Musanase et al., (2023) is also the study that first utilized artificial neural networks for crop recommendation whereas previous papers focus mainly on tree-based approaches. The previous research papers all utilize a classification criterion that assumes that there is one crop that can be grown on any piece of land. So, they may therefore be thought of as crop classification models rather than recommender systems, which is an important difference. This absolute perspective is not useful for choosing grape varieties. Which variety should be grown is dependent not only on suitability but also on marketing potential, the regional legislations regarding grape variety selection, as well as the preference of the winemaker or vineyard owner, or the general regional demand. Furthermore, the model should be able to return the suitability of all grape varieties to be useful for enhancing diversity effectively.

For grape varieties, a recommendation system should be developed that ranks the varieties by suitability and helps the growers to narrow down the grape vine diversity usefully. This is likely a better approach for crop recommendation systems in general, as models should not dictate optimum solutions, as they have no guarantee of this from the training data. As this is a relatively new field, approaches will still change and hopefully soon adapt. Additionally, the errors of the machine learning (ML) algorithm need to be accounted for and communicated to the user. Making an absolute prediction leads to the user expecting a degree of certainty that the model cannot possess. Ranking all possible varieties may, on the other hand, imply that the model is suggesting or recommending, which in practice will be more useful and desirable for the growers as it does not take their place but rather supports them.

1.2.1 New Approaches

Regression between the climate as the independent and grape varieties as the output variables, which is the approach used by previous crop recommender systems (Garanayak et al., 2021; Gopi & Karthikeyan, 2023; Islam et al., 2023; Musanase et al., 2023), may not be the best approach to make grape varietal recommendations. Developing a model that directly works from climate or other variables to crops is promising for classification and finding one crop as has been shown in the literature (Garanayak et al., 2021; Gopi & Karthikeyan, 2023;

Musanase et al., 2023). However, this also limits the usability of the model for other tasks, the moment the inputs change training needs to be conducted from scratch. More flexible models may be useful for allowing for transfer-learning, so that models can be re-used for similar but not identical tasks, or fine-tuning, so that models may be improved for a specific case study, or imputation, so that models may be used to fill-in incomplete data. This aligns more closely with the modern machine and deep learning (ML/DL) paradigms (Howard & Gugger, 2020), and is practiced by the leading models (Chen et al., 2020; Devlin et al., 2019; Pan & Yang, 2010; Radford et al., 2021). The discipline of learning useful representations has gained momentum and learning transferable representations is often considered even more important than the final prediction because of its wide-ranging usefulness (Chen et al., 2020; LeCun & Misra, 2021; Radford et al., 2021).

It needs to be considered that such 'multi-purpose' models require great amounts of computational time and very large datasets of millions of data points to train (Radford et al., 2021), which may not be available in the field of viticulture. This to a large part depends on the size of the models used, with billions of parameters that need tuning and hence a lot of very diverse data, but more efficient alternatives exist that can provide similar usefulness. Generally, to be a multi-purpose model the model is generally trained to take in the same type of data that it produces or to complete it in some way, through imputation for example (Devlin et al., 2019). This means it is usually a generative model as it doesn't reduce the input but produces data 'equivalent' to the input. A more simple architecture that achieves that effect is the Auto Encoder (AE) (Bank et al., 2020). AEs are traditionally used to remove noise from images or to reduce their size as they compress the input information and then recreate it. Derived versions of this architecture have been used to create multipurpose models and useful representations (Caciularu & Goldberger, 2023; Cohen Kalafut et al., 2023; Iatrou et al., 2022; Radhakrishnan et al., 2023).

Of these, I will highlight Radhakrishnan et al., (2023) which uses a very simple version of the AE approach but with great design. Two sets of AEs with a shared latent space, or shared representation at the bottleneck of the AE are used, for a schematic of one AE please see Figure 1.

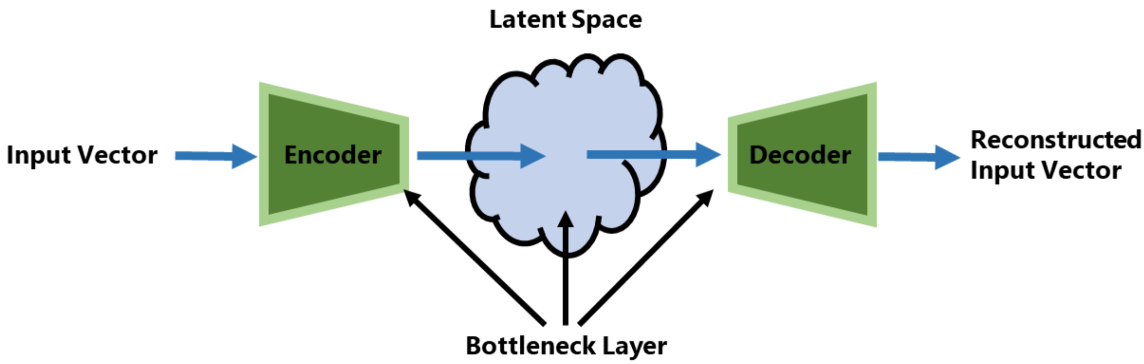

Figure 1: A simple diagram of the Auto-Encoder architecture. Both the encoder and decoder are usually ANN-based and the smallest layer is generally in the middle, being the last layer of the encoder which is shared by the decoder and often referred to as the latent space.

Using one AE for electrocardiogram (ECG) data and another for cardiac magnetic resonance images (MRI) the model was able to learn a shared representation through the shared latent space or bottleneck for both images and ECG data (Radhakrishnan et al., 2023). This allows the model not only to predict the ECG for any MRI or vice-versa, when some of the information is missing, but it also allows for further study by allowing researchers to find patterns in cardiovascular health in its latent space (Radhakrishnan et al., 2023). The representations can further be decoded by different decoders to predict other phenotypes of interest (Radhakrishnan et al., 2023), providing an extremely flexible DL network based on quite a simple foundation.

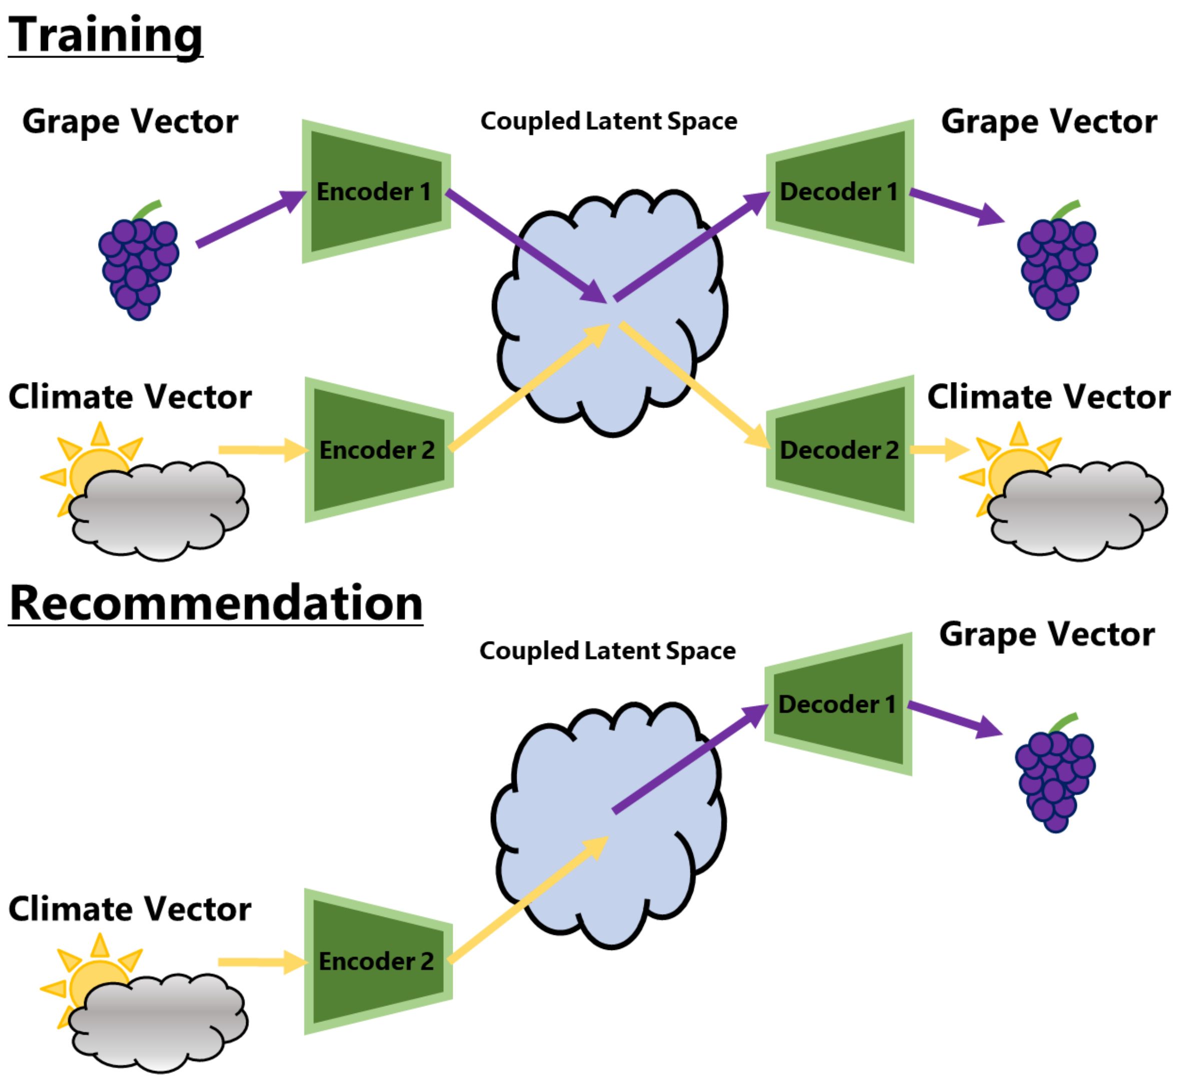

Figure 2: Shows the proposed methodology for creating the Coupled-Auto-Encoder Network (CAE) for grape variety recommendations.

The possibility of developing a similarly useful system for grape growers and the wine industry is exciting. Theoretically creating a 'terroir' embedding is possible. It may be a combined set of AEs of which one is dedicated to the grape growing areas of all the potential varieties and the others, with a shared latent space, would be devoted to the various factors of terroir. In this thesis I present the development of such a model architecture with two coupled AEs (CAE) that combine the proportions of which grape varieties are grown with the local climate, its utility for predicting and helping recommend suitable grape varieties is evaluated for past and future climates. However, since it is designed using principles from representation learning it remains adaptable for other grape-related prediction problems and can be taught to relate other terroir factors in the same way or be applied to enhance breeding programs since grape varieties can be clustered in the latent space by their climate attributes. Such an architecture is shown in Figure 2 and will be developed and tested in this Thesis.

2 Methodology

2.1 Model Definition

The proposed terroir CAE architecture is shown in Figure 2 above. Each encoder and each decoder have 3 fully connected ANN layers with 1024 neurons and bias terms. The encoder additionally includes a linear projection head to produce the latent space representation or act as the bottleneck that reduces the representation from 1024 to 256 neurons. The specific hyper-parameters for the CAE such as the bottleneck size are determined using a Tree-structured Parzen Estimator (TPE) based hyper-parameter optimization over 1600 runs for 5 hyper-parameters (Bergstra et al., n.d.):

1. Learning rate from e^{-10} to 1
2. Batch size from 256 to 768
3. Mixing rate which determines the level of interconnection between latent spaces from 0.5 to 0.12 (discussed more later)
4. Size of bottleneck from 128 to 256
5. Activation function which is a choice between GELU (Hendrycks & Gimpel, 2023), LeakyReLU, Sigmoid, Softsign

The optimal hyper-parameters found that were used for model training are the following:

1. Learning rate = 0.0004835
2. Batch size = 256
3. Mixing rate = .33
4. Bottleneck = 256
5. Activation function = Sigmoid

The model is implemented using the PyTorch library in Python (Paszke et al., 2019). For evaluation and other metrics we generally utilize the Scikit-learn functions (Pedregosa et al., 2011).

2.2 Baseline Models

As baseline models, we utilize the ANN approach from Musanase et al., (2023) as well as a tree-based model based on the XGBoost architecture which is similar to the approaches used by previous crop recommender systems (Garanayak et al., 2021; Gopi & Karthikeyan, 2023). ANNs are a layered collection of multiple-non-linear-regression-like models which allow it to theoretically approximate any function (Hornik et al., 1989). The output is produced by the last layer and each preceding layer takes as its input the output of the previous. ANNs usually refer to the shallow version (<3 layers) of this architecture whereas deeper versions are usually called deep neural networks, these are the basis for most DL architectures. XGBoost or extreme gradient boosting machines are a type of decision tree ensemble, where predictions are averaged over all trees in the ensemble. Decision trees are essentially mathematical flow-charts that at every layer linearly divide the dataset based on one of the input features. In particular, XGBoost was used since it is a very common architecture and well-known so the Scikit-learn implementation can be used (Pedregosa et al., 2011).

The baseline models' hyper-parameters were also tuned using the TPE approach (Bergstra et al., n.d.). The number of evaluations is scaled to account for the fewer parameters/options so the ANN model used 200 runs and the XGBoost model used 125 runs. The parameters found are listed in the Tables 2 & 3 below:

Table 2: ANN hyper parameters found after TPE optimization

Parameter	**Value**
Learning Rate	4.7727e-05
Batch Size	256
Activation Function	Softsign

525

Table 3: XGBoost model hyper parameters found after TPE optimization

Parameter	**Value**
Data Per Tree	100%
Max Depth	2
Number of Estimators	50

2.3 Grape Data Preprocessing

The input data comes from the database (Anderson & Nelgen, 2020). To normalize the acreage values for each variety in each region we define the popularity score according to Eq.
530 (1):

$$p\_i = \frac{\log(a_{\mathrm{i}}+1)}{\sum_A \log(a+1)+10^{-8}} \tag{1}$$

Where p_i is the popularity of variety i in a specific region and A is the set of the areas of all grape varieties whereas a_i is the area of variety i. This score ensures that the value is between 0 and 1 for best compatibility with the neural networks and ML in general. Additionally, due
535 to its definition it also ensures that 'small' varieties have a non-negligible value which would have been the case if scaling by the total area in a region for example (for a grape with 1% share in growing area the score is 0.15 instead of 0.01). Still, we remove grapes from the 1595 varieties if they constitute less than 1% of the area in all the regions that we have data for, this

leaves 1300 grape varieties. This is done to ensure that all varieties that the model can make predictions for it also it has sufficient data for.

Since a few varieties are very popular globally, namely the international varieties (Johnson & Robinson, 2001), and it is desirable to normalize this popularity to determine the suitability of a variety, the suitability score is also introduced in Eq (2) below:

$$s\_i = \frac{\mathrm{p\_i} * \frac{\max(P_{\mathrm{i}})}{1300}}{1300} \tag{2}$$

Where s_i is the suitability of variety i, it depends on p_i which is the popularity of variety i in a specific region, and P_i is the set of the popularities of the same variety over the training set. Varieties that are never very popular therefore will receive a relative advantage against varieties that generally make up most regions.

2.4 Climate Data Preprocessing

Climate data for the model originates from the NASA Prediction of Worldwide Energy Resources (POWER) Project: NASA/POWER CERES/MERRA2 Native Resolution Daily Data v2.5.5 accessed on 2024/01/29. This is supplemented by data obtained from TerraClimate (Wang et al., 2016). TerraClimate has a resolution of about 4x4 km and provides data on the monthly temporal scale while the data from the POWER project can provide daily data but at a resolution of about 50x50 km. The climate data was obtained in one location per region in the dataset, the approach from Puga et al., (2022) was followed and the same reference locations were used as well except for British Columbia (BC). Here the location was moved so that the calculated GDD more closely matched the true GDD of the initial location. This was considered important as one of the subsequent case studies presented here will be conducted in BC. The POWER project data is the source of all temperature-related data of the 16 indicators listed in Table 1. To increase the accuracy of the data the temperature is corrected by elevation since the 50x50 km resolution can greatly over or underestimate the elevation especially since many wine-growing regions from the training set are in mountainous regions. The elevation information comes from the Global Digital Elevation Model (GDEM) from the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER), which has a 30m resolution. The temperature used is calculated by Eq. (3):

$$t = t_{raw} + 0.0061 * (e_{POWER} - e_{ASTER}), \tag{3}$$

Where t is the temperature that will be used later, t_{raw} and e_{POWER} are the temperature and elevation derived from the POWER dataset respectively while e_{ASTER} is the elevation that is closer to the real elevation from the ASTER GDEM. The factor to convert the change in elevation to temperature (0.0061 C/m) is derived from Wallace & Hobbs, (2006). To calculate the climate indicators that are defined by the growing seasons we adjust the growing season definitions to assure equal length of growing season (for when daily calculations are required) for the north and south hemispheres and for leap and non-leap years. This is done in the following way: for locations in the southern hemisphere where the previous year was not a leap year the growing season ranges from the 274th day of the year of the previous year to the 122nd day of the current year or the 121st day in case the previous was a leap year. Northern hemisphere growing seasons are defined to range from the 90th to the 274th day of the year. This corresponds closely to the monthly definitions given by Puga et al., (2022). This is required especially for the training data, since individual years will be used for training to maximize the number of data points, and since all sources of bias during training should be removed. Later the model will likely not require using these same definitions as we anticipate long-term averages will be the input and the model should then be robust to small differences.

TerraClimate is used where only monthly data is required, which are the indices related to precipitation (AnnP, GSP, and HMP) as well as vapor pressure deficit (VPD_GS and VPD_SU).

The equations for calculating GDD and HI are given by (Hall & Jones, 2010).

For model training the climate variables are normalized using min/max normalization such that as for the grape data every input variable ranges from 0 to 1. The normalization is shown in Eq. (4):

$$v_{norm,i} = \frac{v_{raw,i} - \min(v_{raw})}{\max(v_{raw}) - \min(v_{raw})}, \tag{4}$$

Where $v_{norm,i}$ is the normalized variable of region i whereas $v_{raw,i}$ is the corresponding raw variable for the same region. Regions for which "unusual" climate characteristics were identified, i.e. regions and years which according to Hall & Jones, (2010) would have to be classified as not suitable for wine production, are excluded from the training dataset. This is done as a measure to control for faulty data and to not teach the model to rely on more certain 'high quality' regions and to not learn misleading trends. Specifically, these are regions for which the GDD is less than 850 or greater than 2700, where the HI is less than 1200 or greater than 3000 and the GST is less than 13 C or more than 24 C (Hall & Jones, 2010). The excluded regions will however be considered in the test set to confirm whether the model is able to make sensible predictions for 'odd' regions, or regions with faulty data.

2.5 Training

Climate indices are calculated for every region in the dataset of Anderson & Nelgen, (2020), for every year between and including 1996 and 2020 accounting for a total of 25 years. Individual years are used for training to (a) prevent over-fitting by allowing for more variability, or a sort of natural augmentation of the dataset, in the climate indices and (b) increase the overall amount of training data both critical for well-performing DL models (Howard & Gugger, 2020). The testing set consists of the average of the 10 years not used for training. The testing set includes all regions (903) so that the robustness of the model can be properly assessed as well as its ability to predict beyond the training range and for unseen grape combinations. Both of these actions are taken to ensure that the model will perform under likely usage scenarios where 'unlikely' regions may be more common and also long-term averages will be used by the growers instead of yearly data. Since the grape acreage data has three distinct time points 2000, 2010, and 2016 so the training years selected include these years and also 2 years prior and after these 3 distinct years for a total of 15 years of training data for each region. Since regions do not keep the same proportions of the varieties each training year uses the closest of the three grape acreage time-points as its label. Regions which do not exist at all time points only are used in the training set for the 5 years where they are considered available. The testing set uses the average over all available years as the label. An overview of the training and test set years are shown in Table 4 below:

Table 4: The years used for training (blue) versus testing (violet) sets are shown.

Year	'96	'97	'98	'99	'00	'01	'02	'03	'04	'05	'06	'07	'08	'09
'10	'11	'12	'13	'14	'15	'16	'17	'18	'19	'20				

From the training set 20 % is used as a validation set on each run, and cross-validation is run to align with previous work in crop recommender systems and also the ANN and XGBoost modelling practices (Islam et al., 2023; Musanase et al., 2023). Further, it allows us to gauge the robustness of the approaches somewhat. We use 5-fold cross-validation meaning that every model is trained with 80% of the training set, and 20% of the training set is used to implement early stopping and track validation loss. Early stopping was chosen so that all models, which generally have different required epochs for training are trained to the same degree, here 50 consecutive epochs of non-improvement of recommendation performance are used to stop training. The size of the training dataset is 7230 points (5784 training and 1446 for validation) and the test dataset contains 903 data points. This means that 1446 grape growing area vectors are contained in the training set, each with 5 years of climate indices while the test set utilizes one set of 10-year averages for 903 regions and averaged grape growing areas over the years for these regions. In the next section, the losses used for training and validation are discussed.

2.5.1 Losses and Metrics

The baseline models are trained by taking the climate data as input and predicting the popularity score of all the 1300 grapes considered, as would be the traditional approach (Musanase et al., 2023). The loss used to train the models is the mean squared error loss (MSE). This was done to create a regression rather than a classification problem such that the models learn for any climate input to give a likely popularity value to all of the considered grape varieties. For the CAE model, the loss also has an MSE component but also a joining loss component. The joining loss is necessary for connecting the two component models (Radford et al., 2021; Radhakrishnan et al., 2023). Here a similarity loss is chosen similar to the loss used by Radford et al., (2021) and the loss used by Radhakrishnan et al., (2023), the objective of this loss is to maximize the cosine similarity between the bottlenecks or the latent spaces of the models. It is quite convenient to calculate since it essentially is the cross-entropy of the cosine similarity of the two latent spaces (Radford et al., 2021). Cross-entropy is the most commonly used loss for classification and therefore, also widely used in the previous crop recommendation systems (Musanase et al., 2023). The cross-entropy component of the joining-loss measures whether for each dimension of the latent space, the corresponding dimension of the other models' latent space is the most in agreement with. Thereby the two latent spaces are shifted until they align. In addition to that, to ensure an even better connection between the latent spaces during training time some of the neuron activations are switched between the models as was done by (Radhakrishnan et al., 2023). In this way, not only the joining loss is encouraging the models to have a similar embedding but the switched activations will carry any misalignment into the final prediction such that the MSE component also adds to adjusting the model's latent spaces. The rate at which these activations are switched during training time is given by the mixing rate hyperparameter which was set to 0.33. This means that each latent space neuron had a probability of 0.33 to be switched with a neuron from the other model (this is calculated only for one model so that the chance for the other model remains the same). The AdamW optimizer was used to approximate the methodology from (Radford et al., 2021).

To practice data augmentation we also set a dropout rate of around 15% where input grape varieties or individual climate indices are set to 0. This is done both for baseline and CAE models. This allows the model to see configurations of grape varieties that are possible but do not exist in the dataset and also prevents it from giving too high of weight on any individual climate index. As mentioned above data augmentation is one of the most important tools in DL model development to create models that are robust to small input changes. It benefits the models similarly to additional data without requiring additional data (though if available additional data would be preferable) (Howard & Gugger, 2020).

Validation is performed by using the net-discounted-cumulative-gain (NDCG) metric. For this, the CAE model is put into recommendation mode, so that only climate is given and the latent space of the climate encoder is passed completely to the grape-variety decoder. NDCG measures whether the ordering of results is close to the ideal order of results, it is a common metric for recommendation systems and hence used here (Wu et al., 2019). As explained above if the validation loss on the validation set does not decrease for 50 consecutive epochs training is stopped and the model is evaluated on the test set.

Other metrics used for evaluation are the top 1 accuracy; whether the top-ranked variety matches the real top-ranked variety (the main metric used by previous crop recommender systems), and additionally the top 5 accuracy which measures whether the actual most popular variety is within the top 5 most highly ranked varieties. Both of these measure the utility of the model to characterize the main few varieties of these regions. The Root Mean Squared Error

(RMSE) is used as a proxy to how well the models reduce the training loss (MSE). The root of the MSE is used as it makes the scale of the error more easily visible since all the values are between 0 and 1. Having similar RMSE is expected from all models after training since it indicates that the training was successful and the training setup used was fair, or if one of the models is not at all suitable for solving the problem. Next, the area under the Precision-Recall Curve (auPRC) is used. Precision measures the proportion of true positives divided by the total positives in the model. The Recall also known as sensitivity measures how many of the positive cases are identified (Musanase et al., 2023). In this case, each grape variety is treated as one sample (so there are 1300 samples per data-point in the test set). Each variety that exists in the region is treated as a true positive and then all the predicted values above the threshold are treated as predicted positive. The Precision-Recall curve then plots the Recall and Precision for various thresholds for each data-point. Recall would be maximized where all varieties found are predicted with a certainty of 1 and Precision where all varieties not grown are given a score of 0. The average area under the Precision-Recall curve (auPRC) over all data-points therefore evaluates how sensitive the model is to including the varieties grown in the region as positive versus including other varieties. It is an overall measure of how selective the model is towards the relevant varieties. It may be seen that the auPRC score is a sort of combined measure of the other metrics and is therefore best used to evaluate the overall performance of the model. In the past Precision and also Recall have been used as metrics to evaluate recommender systems separately from one another (Darban & Valipour, 2022), here both are combined to give a more holistic view of the model performance.

3 Discussion & Results

The results of the trained model evaluated on the test set alongside the baseline models are shown in Figure 3 below. Starting from the right to left the results will be discussed. Firstly, the d1RMSE is shown which is the difference between one and the RMSE score. This representation is chosen for consistency such that the higher value is better as with the other scores. This score is a proxy of the MSE score which is used as the training loss or a component of the training loss of the models. As can be seen in Figure 3, all models can minimize the training loss roughly equally well. The paired t-test between the 5 ANN and CAE models performances for each cross-validation fold indicated that there is no significant difference between the ANN and CAE models in this metric. Because all models can minimize the training loss similarly it can be supposed that the difference in architecture does not affect the ability of the models to minimize the training loss but the difference lies in whether they minimize it in a way that is useful for the task that the models are intended for.

This is now seen with the NDCG loss which scores how close the ranking of the varieties matches the ranking that the test regions have. Firstly, statistically with $p<0.05$ the NDCG scores of the CAE models are significantly better than those of the ANN models. Additionally, it is observed that the ANN models have a much larger variance in their scoring efficiency. This would then suggest that not only does the CAE model outperform the ANN on average but it consistently outperforms the ANN approach and gives more reliable results which can be seen across all metrics shown in Figure 3. The CAE model performs more consistently than the ANN models, not just better but with more reliability. The CAE may through the bottleneck that it possesses remove some of the noise that can lead to erratic behaviour of neural network models. The same principle by which the bottleneck may prevent overfitting as it forces selection of only the significant inputs from the previous layer (Howard & Gugger, 2020). Additionally, the latent space of the CAEs is informed by both the encoded climate and the encoded grape inputs which it has to align creating further restraint on the bottleneck layer to further prevent overfitting, the ANN models have neither of these mechanisms.

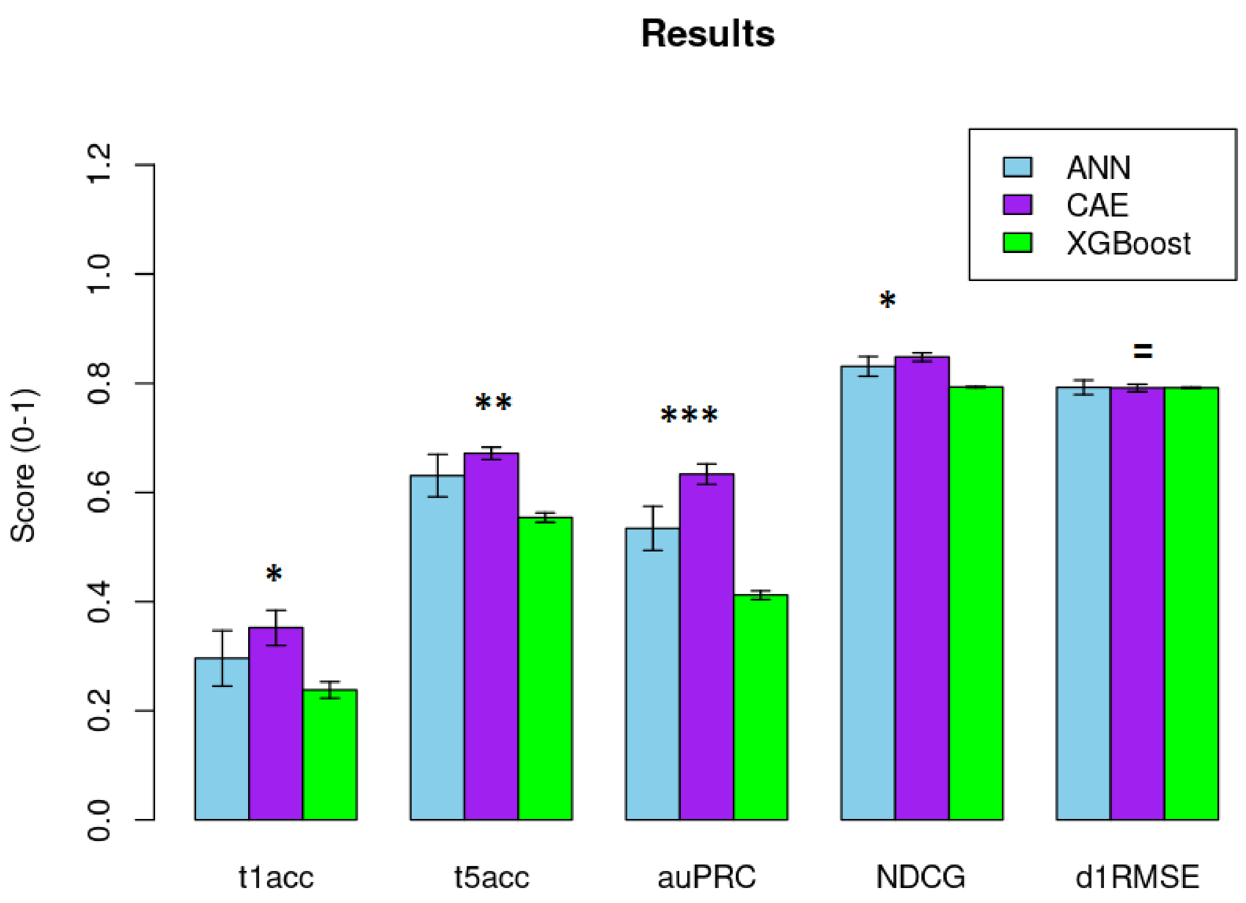

Figure 3: Results of the trained models on the test set. Significances are shown for a one-sided paired t-test between the CAE model and the ANN model. (* for p<0.05, ** for p<0.01, * for p<0.001, = for p>0.05)**

Robustness is very important when evaluating an approach, especially a DL approach for which certain robustness cannot be guaranteed in general. DL models can erratically react to unseen data which makes it a critical issue to have robust DL approaches for users to trust the model's output, or to be able to gauge the approximate error (Howard & Gugger, 2020).

Seeing then in Figure 3 the consistency of the CAE results across the metrics is encouraging and may suggest that the models are worthy of trust, which will surely also advance the technique in the eyes of the end-users.

The accuracy scores also clearly show the superiority of the proposed CAE approach over both the ANN and XGBoost approaches, top 1 accuracy shows significantly better performance from CAE and high significance ($p<0.01$) between CAE and ANN performances for the top 5 accuracy. Past models of crop recommendations systems research perform very highly in terms of accuracy since they are trained only for this objective (Musanase et al., 2023), but the CAE approach outperforms these approaches when all are trained for recommendation rather than classification.

Lastly, moving to the score that combines some aspects of the above-mentioned scores the auPRC score. Here the clearest difference in performance can be seen between the models, with $p<0.001$ the CAE model outperforms the ANN model clearly, not to mention the XGBoost model. Again, with high consistency between the various runs. In conclusion the initial model results on the test set, which includes regions not seen and years not seen before in the training set, are very promising. The sets are very different but the performance of the CAE model is consistent and consistently better when it comes to the objectives of recommending suitable grape varieties. This is remarkable considering the CAE model is the only DL method tested which should otherwise be the type of model reacting more erratically (Howard & Gugger, 2020).

To further evaluate the CAE model's ability to predict suitable varieties for other regions the wine-growing regions of British Columbia are used. Climate indices are calculated according to the training and testing set, the average value over the years 1996 to 2020 is used as the input. The locations are this time hand-picked from the wine-growing regions, to select a location that is representative of the prevalent growing conditions. Data of the most common grapes in the sub-regions of BC is available for the top 4 varieties only but for 16 sub-regions (of which 3 are regions that combine multiple towns) (Withler & Geldart, 2023). The results of the prediction can be seen in Table 5 below.

The combined regions from the dataset are separated here for a total of 20 regions. We assume that the true variety distribution remains the same in each split-off region for the sake of simplicity. The varieties are colour-coded to be easily visually distinguished by the reader. When the varieties are also predicted to be in the top 4 they are marked in green, if they are predicted in the top 10 they are marked in yellow, and 10-25th are marked in orange with red past that. At first look, it can be seen that most of the varieties in the top 4 in reality are also ranked among the top 4 by the model, roughly 62.5% of varieties. Another 26.25% were ranked within the top 10 varieties for the climate. As a reminder to the reader, at this point, only the popularity is shown not the suitability which will be discussed later. Overall, 88.75% of the top 4 varieties from the British Columbia wine regions are predicted to be at least in the top 10 based on popularity. The worst of the predictions are Maréchal Foch (59th) for Vancouver Island and La Crescent (64th) for the Thompson Valley. Considering the model has the option of choosing any of 1300 grape varieties the worst prediction is within the top 95.5% of grapes. This is still quite a good considering that the global acreage of Maréchal Foch is about 465 ha and of La Crescent 114 ha compared to 21067 ha for Gewürztraminer or 171940 ha for Pinot Noir. However, this error in ranking points out a weakness of using only the popularity score, and the necessity of normalizing for overall popularity. In Table 6 the top ranked varieties by suitability score are shown which accounts for this problem. For suitability or potential there unfortunately is no validation data, but it is nonetheless interesting to observe.

Table 5: Results of popularity predictions for 20 selected locations in British Columbia versus the true popularity of the same varieties according to the Grape Growers BC report (Withler & Geldart, 2023).

Region	Popularity	Predicted Rank	Top 4 Predicted Grapes
Oliver	Merlot	3	Pinot Noir
	Cabernet Sauvignon	5	Chardonnay
	Chardonnay	2	Merlot
	Cabernet Franc	7	Syrah
Penticton*	Pinot Noir	1	Pinot Noir
	Pinot Gris	3	Chardonnay
	Merlot	8	Pinot Gris
	Chardonnay	2	Gewürztraminer
Naramata*	Pinot Noir	1	Pinot Noir
	Pinot Gris	3	Chardonnay
	Merlot	7	Pinot Gris
	Chardonnay	2	Syrah
Kaleden*	Pinot Noir	1	Pinot Noir
	Pinot Gris	3	Chardonnay
	Merlot	8	Pinot Gris
	Chardonnay	2	Gewürztraminer
Osoyoos	Merlot	3	Pinot Noir
	Cabernet Franc	6	Chardonnay
	Cabernet Sauvignon	5	Merlot
	Syrah	4	Syrah
Kelowna	Pinot Noir	1	Pinot Noir
	Chardonnay	2	Chardonnay
	Pinot Gris	3	Pinot Gris
	Riesling	10	Syrah
Similkameen Valley	Merlot	6	Pinot Noir
	Cabernet Sauvignon	3	Chardonnay
	Cabernet Franc	13	Cabernet Sauvignon
	Chardonnay	2	Riesling
Okanagan Falls	Pinot Noir	1	Pinot Noir
	Chardonnay	4	Pinot Gris
	Pinot Gris	2	Syrah
	Gewürztraminer	5	Chardonnay
Summerland**	Pinot Noir	1	Pinot Noir
	Gewürztraminer	4	Chardonnay
	Chardonnay	2	Pinot Gris
	Pinot Gris	3	Gewürztraminer
Peachland**	Pinot Noir	1	Pinot Noir
	Gewürztraminer	7	Chardonnay
	Chardonnay	2	Cabernet Sauvignon
	Pinot Gris	5	Merlot

Region	Popularity	Predicted Rank	Top 4 Predicted Grapes
West Kelowna	Pinot Noir	1	Pinot Noir
	Chardonnay	2	Chardonnay
	Pinot Gris	3	Pinot Gris
	Riesling	10	Syrah
Lake Country* **	Pinot Noir	1	Pinot Noir
	Pinot Gris	4	Chardonnay
	Riesling	10	Merlot
	Chardonnay	2	Pinot Gris
Vernon* **	Pinot Noir	1	Pinot Noir
	Pinot Gris	6	Chardonnay
	Riesling	5	Cabernet Sauvignon
	Chardonnay	2	Merlot
Vancouver Island	Pinot Noir	1	Pinot Noir
	Pinot Gris	3	Chardonnay
	Ortega	11	Pinot Gris
	Maréchal Foch	59	Seyval Blanc
Fraser Valley	Pinot Noir	1	Pinot Noir
	Bacchus	20	Pinot Gris
	Siegerrebe	38	Chardonnay
	Pinot Gris	2	Chasselas
Shuswap	Maréchal Foch	10	Pinot Noir
	Ortega	26	Chardonnay
	Siegerrebe	22	Aligoté
	Pinot Noir	1	Pinot Blanc
Thompson Valley	Pinot Noir	1	Pinot Noir
	La Crescent	64	Chardonnay
	Chardonnay	2	Gewürztraminer
	Riesling	8	Cabernet Sauvignon
Kootenays	Pinot Noir	1	Pinot Noir
	Gewürztraminer	11	Chardonnay
	Chardonnay	2	Cabernet Sauvignon
	Pinot Gris	6	Pinot Blanc
Gulf Islands	Pinot Noir	1	Pinot Noir
	Pinot Gris	3	Chardonnay
	Chardonnay	2	Pinot Gris
	Gewürztraminer	29	Bacchus
Lillooet	Cabernet Franc	10	Pinot Noir
	Merlot	7	Chardonnay
	Riesling	4	Pinot Gris
	Pinot Noir	1	Riesling

*** These were considered the same region by Withler & Geldart, (2023) and so are assumed to share the true popularity data. Predictions were generated for separate locations.**
**** Similar to the above**
***** Similar to above**

Table 6: Results of suitability predictions for 20 selected locations in British Columbia versus the true popularity of the same varieties according to the Grape Growers BC report (Withler & Geldart, 2023). In brackets under the Region name the mean suitability is listed.

Region	Popularity	Predicted Rank	Suitability Score	Top 4 Predicted Suitable Grapes
Oliver	Merlot	7	0.41	Pinot Noir
	Cabernet Sauvignon	12	0.37	Muscat Swenson
(0.014)	Chardonnay	4	0.45	Castel
	Cabernet Franc	10	0.39	Chardonnay
Penticton*	Pinot Noir	1	0.66	Pinot Noir
	Pinot Gris	3	0.52	Blattner Reds
(0.012)	Merlot	12	0.35	Pinot Gris
	Chardonnay	5	0.49	Gewürztraminer
Naramata*	Pinot Noir	1	0.66	Pinot Noir
	Pinot Gris	2	0.59	Pinot Gris
(0.015)	Merlot	7	0.42	Blattner Reds
	Chardonnay	5	0.54	Gewürztraminer
Kaleden*	Pinot Noir	1	0.68	Pinot Noir
	Pinot Gris	2	0.56	Pinot Gris
(0.011)	Merlot	9	0.35	Gewürztraminer
	Chardonnay	4	0.51	Chardonnay
Osoyoos	Merlot	7	0.55	Muscat Swenson
	Cabernet Franc	8	0.53	Pinot Noir
(0.017)	Cabernet Sauvignon	10	0.50	Castel
	Syrah	6	0.55	Chardonnay
Kelowna	Pinot Noir	4	0.64	Muscat Swenson
	Chardonnay	9	0.58	Blattner Reds
(0.019)	Pinot Gris	3	0.65	Pinot Gris
	Riesling	22	0.41	Pinot Noir
Similkameen	Merlot	6	0.16	Pinot Noir
Valley (0.004)	Cabernet Sauvignon	3	0.25	Chardonnay
	Cabernet Franc	13	0.05	Cabernet Sauvignon
	Chardonnay	2	0.28	Riesling
Okanagan Falls	Pinot Noir	1	0.67	Pinot Noir
	Chardonnay	5	0.53	Pinot Gris
(0.013)	Pinot Gris	2	0.63	Gewürztraminer
	Gewürztraminer	3	0.59	Syrah
Summerland**	Pinot Noir	1	0.68	Pinot Noir
	Gewürztraminer	3	0.53	Pinot Gris
(0.010)	Chardonnay	4	0.49	Gewürztraminer
	Pinot Gris	2	0.55	Chardonnay
Peachland**	Pinot Noir	1	0.53	Pinot Noir
	Gewürztraminer	7	0.16	Chardonnay
(0.005)	Chardonnay	2	0.31	Cabernet Sauvignon
	Pinot Gris	4	0.21	Pinot Gris

805

Region	Popularity	Predicted Rank	Suitability Score	Top 4 Predicted Suitable Grapes
West Kelowna (0.018)	Pinot Noir	3	0.65	Blattner Reds
	Chardonnay	7	0.59	Pinot Gris
	Pinot Gris	2	0.67	Pinot Noir
	Riesling	22	0.43	Muscat Swenson
Lake Country* ** (0.021)	Pinot Noir	7	0.65	Muscat Swenson
	Pinot Gris	6	0.66	Blattner Reds
	Riesling	23	0.46	Ehrenfelser
	Chardonnay	12	0.62	Sovereign Opal
Vernon* ** (0.013)	Pinot Noir	1	0.56	Pinot Noir
	Pinot Gris	5	0.39	Chardonnay
	Riesling	14	0.34	Pinot Blanc
	Chardonnay	2	0.42	Maréchal Foch
Vancouver Island (0.005)	Pinot Noir	1	0.77	Pinot Noir
	Pinot Gris	3	0.36	Chardonnay
	Ortega	8	0.17	Pinot Gris
	Maréchal Foch	58	0.0088	Madeleine×Angevine7672
Fraser Valley (0.004)	Pinot Noir	1	0.66	Pinot Noir
	Bacchus	27	0.030	Pinot Gris
	Siegerrebe	32	0.027	Chardonnay
	Pinot Gris	2	0.41	Chasselas
Shuswap (0.004)	Maréchal Foch	8	0.14	Pinot Noir
	Ortega	29	0.062	Chardonnay
	Siegerrebe	11	0.11	Pinot Blanc
	Pinot Noir	1	0.39	Ruby
Thompson Valley (0.013)	Pinot Noir	1	0.58	Pinot Noir
	La Crescent	60	0.04	Blattner Reds
	Chardonnay	3	0.43	Chardonnay
	Riesling	18	0.30	Gewürztraminer
Kootenays (0.003)	Pinot Noir	1	0.34	Pinot Noir
	Gewürztraminer	11	0.05	Chardonnay
	Chardonnay	2	0.22	Pinot Blanc
	Pinot Gris	5	0.081	Cabernet Sauvignon
Gulf Islands (0.008)	Pinot Noir	1	0.76	Pinot Noir
	Pinot Gris	7	0.42	Madeleine×Angevine7672
	Chardonnay	2	0.56	Chardonnay
	Gewürztraminer	36	0.037	Bacchus
Lillooet (0.003)	Cabernet Franc	10	0.042	Pinot Noir
	Merlot	7	0.065	Chardonnay
	Riesling	4	0.13	Pinot Gris
	Pinot Noir	1	0.40	Riesling

*** These were considered the same region by Withler & Geldart, (2023) and so are assumed to share the true popularity data. Predictions were generated for separate locations.**
**** Similar to the above**
***** Similar to above**

For interpretability of the tables, the color coding for the suitability scores are as follows: dark green is assigned to a score above e^{-2} so about 0.135 is considered excellent suitability, between 0.135 and e^{-5} is colored in light green which is considered above average suitability (about 0.0067). Defining which score is suitable objectively is difficult so here the histogram from Figure 4 is used to define categories by the distribution of suitability scores. The averages of the regions are not quite the same but it should be noted that the number of grape varieties ranked highly is very small for each region, and therefore the average is largely dependent on the score of the first few varieties, which are much larger than the subsequent scores. This is also seen in the histogram of the ratings for Kelowna in Figure 4 below.

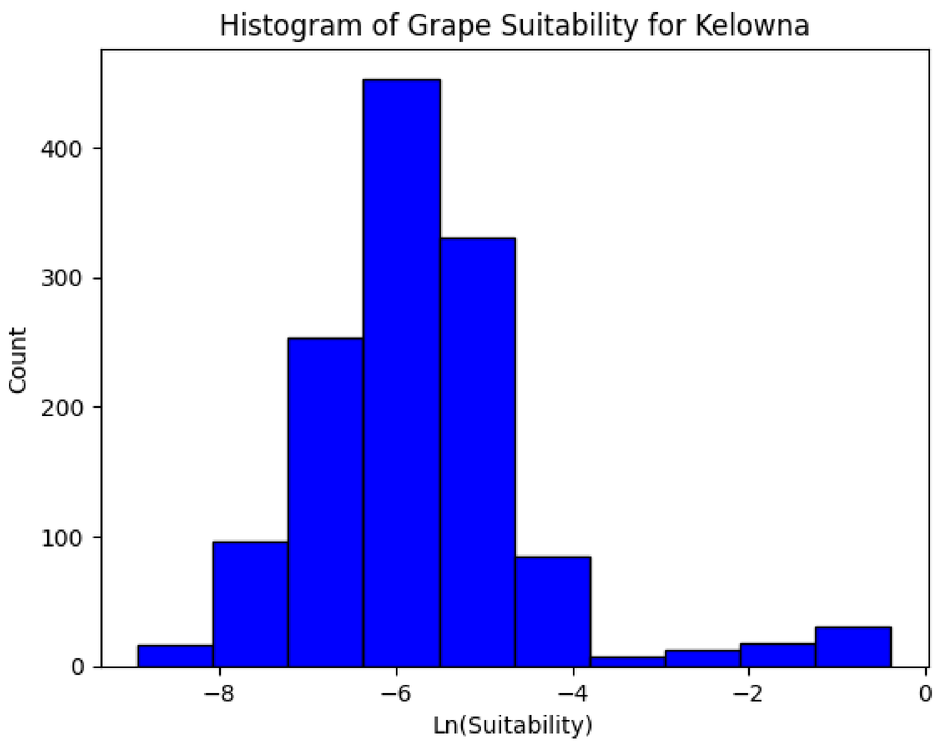

Figure 4: The log-Histogram over all grape varieties' suitability scores for the Kelowna region of British Columbia.

Regarding the suitability scores, it is clearly visible that the suitability rankings does not disqualify any of the varieties grown in the BC wine regions but assigns them at least above average suitability. This is a promising result for using the suitability score in the future. What is needed however, is a clear definition of which threshold means what, when is something very suitable versus just possible to be grown? Such definition will likely need to be developed iteratively and may need to utilize empirical experience from applying the model in practice and gauging the response and advice of a multitude of experts. Since such field trials are outside of the scope of the thesis the matter of the suitability score will be left here. Results are promising, suitable varieties in the regions tested were always ranked above average for the region. In the future, clear definitions of high versus low versus medium suitability will be required for the purposes of easier usability, for now quantiles or above/below averages may be used to gauge suitability.

The next case study is conducted on the future suitability of Pinot Noir using climate models. This is to test the utility of the model for planning future wine regions and adapting grape varieties to climate change scenarios as may be required in the future (Wolkovich et al., 2018). Since there is no validation data necessarily in this area, anecdotal evidence will need to be used. Mainly the ease with which the method can be applied to climate projections is assessed. For this case study the source of the climate data is the CNRM-CM6-1-HR model from CNRM-CERFACS (CNRM-CERFACS, n.d.). It is part of the CMIP6 exercise. This model was selected because it provides relatively high-resolution (50x50 km) data of the required climate variables while also providing some in daily time-step which is required to calculate the indices used in this study from Puga et al., (2022). The calculations of the climate indices are the 20 year averages for the historical suitability prediction from the years 1995 to 2015 and for the future scenarios (SSP126: low emission scenario and SSP585: high emission scenario), the years from 2025 to 2045 are used. The average of 20 years is used for modelling and the growing seasons are defined as by Puga et al., (2022). It is assumed that small differences that may arise due to differences in one or two days in the growing season will be accounted for by the averaging. Additionally, it is assumed that end-users will likely not take the time to adapt our definition from the training set when using the model so this also simulates the likely usage scenario of the model as mentioned in the methodology section. The indices are calculated largely the same but since vapour pressure deficit is not available from the climate model directly the relative humidity is used to calculate this value according the handbook of the FAO Chapter 3 equations (11), (12) and (19) (FAO, 1998).

First, the historical Pinot Noir suitability from this new dataset is shown in Figure 5 which serves as a baseline to interpret the future predictions from.

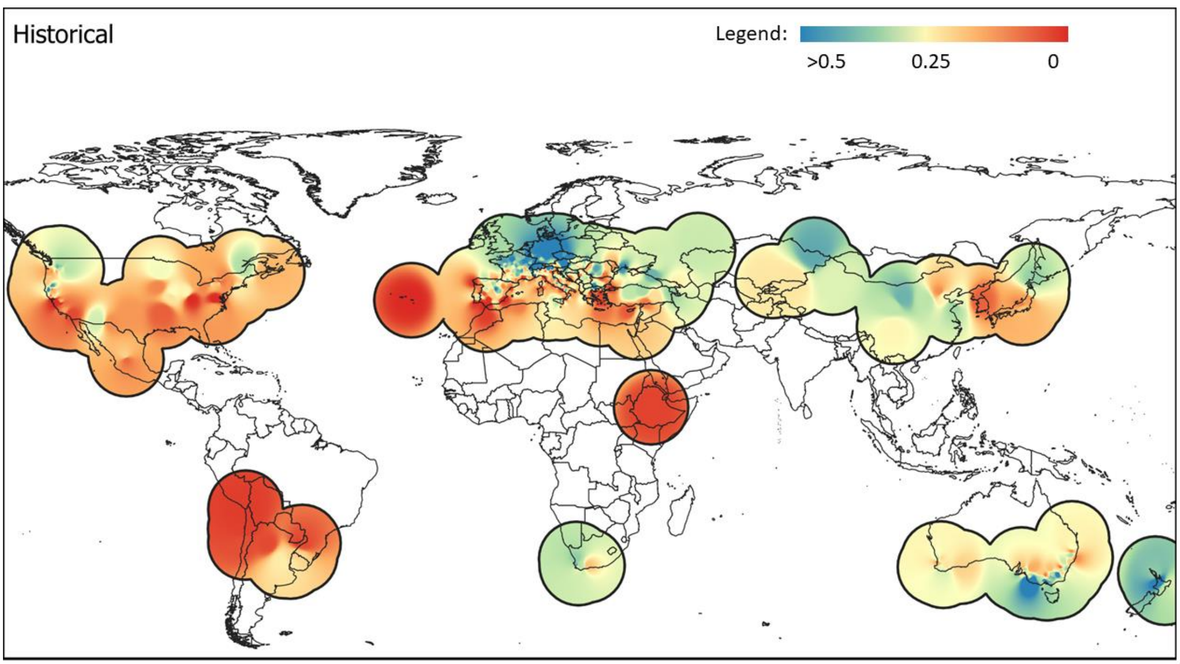

Figure 5: Historical Pinot Noir Suitability Prediction, Inverse Distance Weighted interpolation is shown with p=5.

The maps contain public sector information licensed under the Open Government Licence v3.0.

The Pinot Noir suitability is predicted at 617 of the locations previously defined (these were filtered as previously the regions for the training set), then Inverse Distance Weighted (IDW) interpolation was run (p=5). Interpolated data is shown 10 degrees latitude around the sampling locations, for better visualization. The sampling locations used are shown in the Appendix in Figure A1. The interpolation approach is used to better visualize trends across regions while evaluating every 50 km was unfeasible within the timeframe of this degree and available computational resources. The colouring of the map does not match the colours with which suitability was defined before. This is done to more clearly show the differences in suitability since Pinot Noir has a very high average suitability score over the evaluated regions so using the previous colouring scheme would lead to a map lacking contrast.

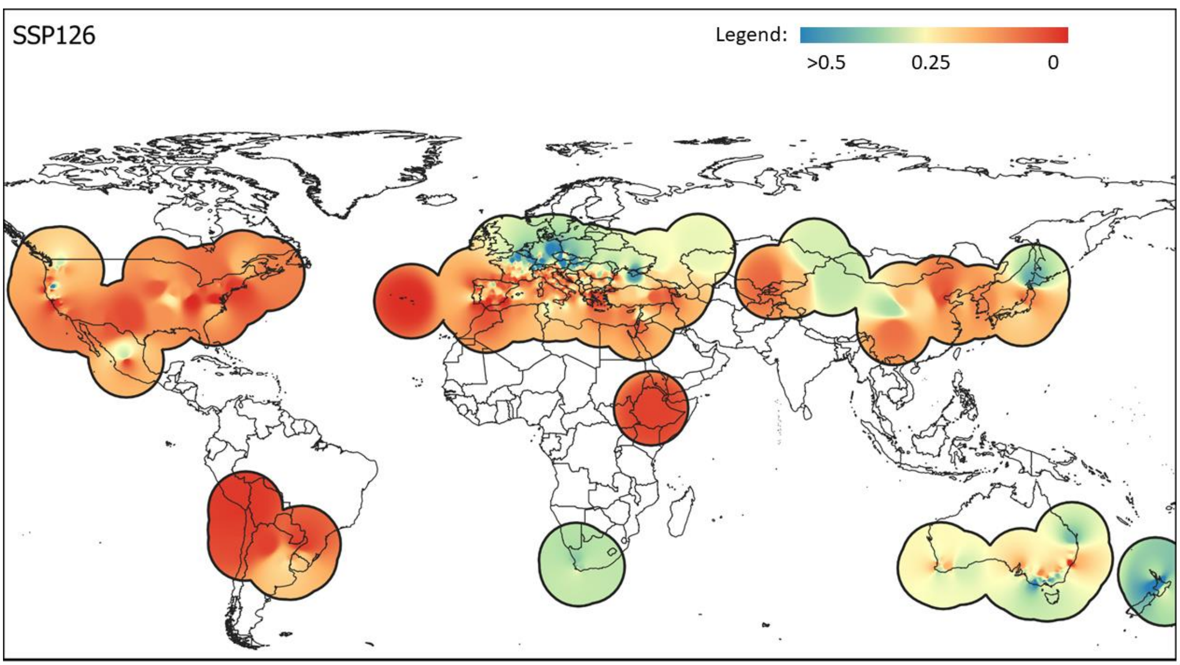

Figure 6: Pinot Noir Suitability Prediction under the SSP 1 scenario which assumes a temperature increase of 2.6 degrees Celsius, Inverse Distance Weighted interpolation is shown with p=5.

The colouring here depends on the average suitability score of Pinot Noir which for the evaluated regions lies roughly at 0.25. The colouring scheme is maintained for the following maps of the climate projections with SSP126 shown in Figure 6 above and SSP585 shown in Figure 7 below.

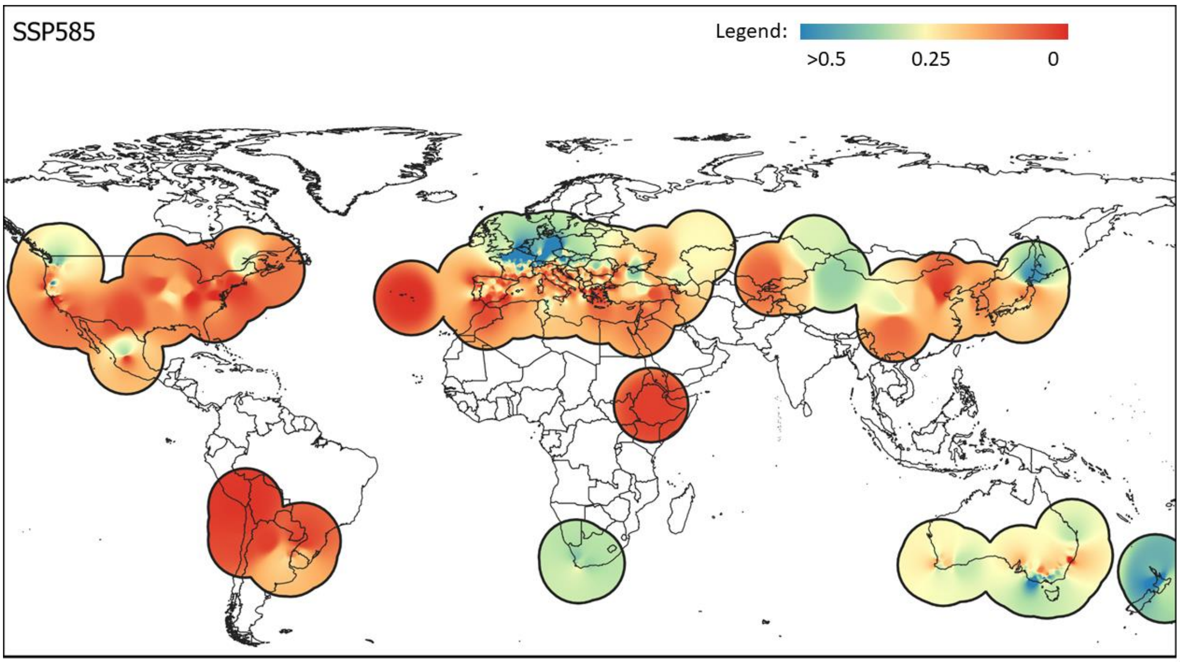

Figure 7: Pinot Noir Suitability Prediction under the SSP 5 scenario which assumes a temperature increase of 8.5 degrees Celsius, Inverse Distance Weighted interpolation is shown with p=5.

The maps contain public sector information licensed under the Open Government Licence v3.0.

Firstly, it is demonstrated that quite simply a global suitability map can be created using the CAE models presented. This, since it can be applied to past and future climates, is a quite useful tool in the hands of the wine industry and grape growers to guide long-term planning. Here only the results for Pinot Noir are shown, but with the same predictions 1300 grape varieties are ranked at the same time. Due to its flexibility, and even expandability, for which this approach was chosen in the first place means that it also hasn’t reached its potential yet.

At the moment it is discouraged to use the figures or models to make investment decisions. While the model was shown here to be reliable it is now applied to a new dataset and more aspects of the model may want to be tested before too much trust is given to it. The historical prediction of Pinot Noir suitability aligns almost exactly with the expected regions for Pinot Noir which are in the northern European wine-growing regions, South Africa, and regions like Oregon which is quite a clear fit according to the map generated in the figures above (Johnson & Robinson, 2001). The suitability of some regions in Australia and not others nearby is also to be expected as Pinot Noir is grown in Australia but only where altitude or sea breezes create favorable conditions (Johnson & Robinson, 2001). The detail of the model's suitability predictions is further confirmed by the prediction for southern California which is based on the wine region of Santa Barbara which is an outlier geographically, but as the model shows and as is the case in reality, is known to produce Pinot Noir (Johnson & Robinson, 2001).

Interestingly, climate projections do not predict large deviations from the geographic regions in which Pinot Noir is suitable. This is to be expected as the time frame used in this study is quite small in the scale of climate. Small changes are visible however, under both scenarios Pinot Noir has a lower average suitability, which is expected as Pinot Noir thrives in rather cooler regions. Past research has also shown that likely Pinot Noir's suitability will decrease due to increasing temperature which the variety is rather sensitive too (Skahill et al., 2022, 2023). This included regions where Pinot Noir is currently very popular such as the Willamette Valley of Oregon. The models shown here predict a decrease in suitability which matches the results from Skahill et al., (2022).

Some other interesting developments are that the models predict that Arkansas may become a better region in the future to produce Pinot Noir. For such discoveries the model is most useful as it can help growers and industry anticipate climate trends and potential before otherwise possible.

4 Conclusion

The world of viticulture is facing uncertainty in terms of the suitability of grape varieties given our changing climate. To assist in guiding through this uncertainty, a generative deep learning model is developed in this research that aims to assists grape growers with guidance on the popularity and suitability of grape varieties for their climate. With ease this model can also be applied to predict future suitability at the global level. To the best of my knowledge, this is the first time such a tool is developed, that can rank 1300 grape varieties by suitability for any climate. A total of 16 indices covering major aspects of climatic conditions known to impact wine quality were used in the definition of climate. The model has been shown here to be more reliable and more consistent than previously used approaches for similar tasks such as artificial neural networks and XGBoost. Through its architecture the model seems to be less susceptible to irrational predictions, even compared with the simpler models, an important factor for its application in practice. Additionally, due to the underlying design, the model is useable through transfer learning and may be adapted to solve other similar problems relating to the climate-dependent attributes of grape varieties. By providing not only a popularity but also a suitability score, smaller grape varieties that may have been overlooked are rated among the more famous international varieties to give the best possible utility to end users, the wine industry and grape growers. However, future research will need to define specific thresholds of suitability and ideally confirm these with empirical evidence from grape growers or viticultural experts for the suitability score to be directly useful for decision making. Additionally, more data and more fine-resolution data would also aid in further increasing the utility of the grape variety recommendation models. No investments should be made based on this research at the moment but the approach has shown its utility and will hopefully help to provide guidance to industry in the near future.

References

Anderson, K., & Nelgen, S. (2020). *Database of Regional, National and Global Winegrape Bearing Areas by Variety, 1960 to 2016* (Version slightly revised May 2021) [Dataset]. Wine Economics Research Centre.

Aney, W. W. (1974). Oregon climates exhibiting adaptation potential for vinifera. *American Journal of Enology and Viticulture*, *25*(4), 212–218.

Bank, D., Koenigstein, N., & Giryes, R. (2020, March 12). *Autoencoders*. arXiv.Org. https://arxiv.org/abs/2003.05991v2

Bergstra, J. S., Bardenet, R., Bengio, Y., & Kégl, B. (n.d.). *Algorithms for Hyper-Parameter Optimization*.

Caciularu, A., & Goldberger, J. (2023). An entangled mixture of variational autoencoders approach to deep clustering. *Neurocomputing*, *529*, 182–189. https://doi.org/10.1016/j.neucom.2023.01.069

Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). *A Simple Framework for Contrastive Learning of Visual Representations* (arXiv:2002.05709). arXiv. https://doi.org/10.48550/arXiv.2002.05709

CNRM-CERFACS. (n.d.). *Models and Contribution to CMIP6—CNRM-CERFACS contribution to CMIP6*. Académie de Versailles. Retrieved July 5, 2024, from http://www.umr-cnrm.fr/cmip6/spip.php?rubrique8

Cohen Kalafut, N., Huang, X., & Wang, D. (2023). Joint variational autoencoders for multimodal imputation and embedding. *Nature Machine Intelligence*, *5*(6), Article 6. https://doi.org/10.1038/s42256-023-00663-z

Darban, Z. Z., & Valipour, M. H. (2022). GHRS: Graph-based Hybrid Recommendation System with Application to Movie Recommendation. *Expert Systems with Applications*, *200*, 116850. https://doi.org/10.1016/j.eswa.2022.116850

Davis, R. E., Dimon, R. A., Jones, G. V., & Bois, B. (2019). The effect of climate on Burgundy vintage quality rankings. *OENO One*, *53*(1), Article 1. https://doi.org/10.20870/oeno-one.2019.53.1.2359

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding* (arXiv:1810.04805). arXiv. https://doi.org/10.48550/arXiv.1810.04805

FAO. (1998). *Chapter 3—Meteorological data*. https://www.fao.org/4/x0490e/x0490e07.htm

Fraga, H., Santos, J. A., Malheiro, A. C., Oliveira, A. A., Moutinho-Pereira, J., & Jones, G. V. (2016). Climatic suitability of Portuguese grapevine varieties and climate change adaptation. *International Journal of Climatology*, *36*(1), 1–12. https://doi.org/10.1002/joc.4325

Garanayak, M., Sahu, G., Mohanty, S. N., & Jagadev, A. K. (2021). Agricultural Recommendation System for Crops Using Different Machine Learning Regression Methods. *International Journal of Agricultural and Environmental Information Systems (IJAEIS)*, *12*(1), 1–20. https://doi.org/10.4018/IJAEIS.20210101.oa1

Gopi, S. R., & Karthikeyan, M. (2023). Effectiveness of Crop Recommendation and Yield Prediction using Hybrid Moth Flame Optimization with Machine Learning. *Engineering, Technology & Applied Science Research*, *13*(4), Article 4. https://doi.org/10.48084/etasr.6092

Gowdy, M., Pieri, P., Suter, B., Marguerit, E., Destrac-Irvine, A., Gambetta, G., & van Leeuwen, C. (2022). Estimating Bulk Stomatal Conductance in Grapevine Canopies. *Frontiers in Plant Science*, *13*. https://doi.org/10.3389/fpls.2022.839378

Hall, A., & Jones, G. v. (2010). Spatial analysis of climate in winegrape-growing regions in Australia. *Australian Journal of Grape and Wine Research*, *16*(3), 389–404. https://doi.org/10.1111/j.1755-0238.2010.00100.x

Hendrycks, D., & Gimpel, K. (2023). *Gaussian Error Linear Units (GELUs)* (arXiv:1606.08415). arXiv. https://doi.org/10.48550/arXiv.1606.08415

Hewer, M. J., & Gough, W. A. (2021). Climate change impact assessment on grape growth and wine production in the Okanagan Valley (Canada). *Climate Risk Management*, *33*, 100343. https://doi.org/10.1016/j.crm.2021.100343

Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, *2*(5), 359–366. https://doi.org/10.1016/0893-6080(89)90020-8

Howard, J., & Gugger, S. (2020). *Deep Learning for Coders with Fastai and Pytorch: AI Applications Without a PhD*. O'Reilly Media, Incorporated. https://books.google.no/books?id=xd6LxgEACAAJ

Howell, G. S. (2001). Grapevine cold hardiness: Mechanisms of cold acclimation, mid-winter hardiness maintenance, and spring deacclimation. *Proceedings of the ASEV 50th Anniversary Annual Meeting, Seattle, Washington, June 19-23, 2000*, 35–48.

Iatrou, M., Karydas, C., Tseni, X., & Mourelatos, S. (2022). Representation Learning with a Variational Autoencoder for Predicting Nitrogen Requirement in Rice. *Remote Sensing*, *14*(23), Article 23. https://doi.org/10.3390/rs14235978

Islam, M. R., Oliullah, K., Kabir, M. M., Alom, M., & Mridha, M. F. (2023). Machine learning enabled IoT system for soil nutrients monitoring and crop recommendation. *Journal of Agriculture and Food Research*, *14*. https://doi.org/10.1016/j.jafr.2023.100880

Johnson, H., & Robinson, J. (2001). *The world atlas of wine* (5th ed.). Octopus Publishing Group.

Jones, G. V. (2018). The Climate Component of Terroir. *Elements*, *14*(3), 167–172.

Jones, G. V., White, M. A., Cooper, O. R., & Storchmann, K. (2005). Climate Change and Global Wine Quality. *Climatic Change*, *73*(3), 319–343. https://doi.org/10.1007/s10584-005-4704-2

Kovalenko, Y., Tindjau, R., Madilao, L. L., & Castellarin, S. D. (2021). Regulated deficit irrigation strategies affect the terpene accumulation in Gewürztraminer (Vitis vinifera L.) grapes grown in the Okanagan Valley. *Food Chemistry*, *341*, 128172. https://doi.org/10.1016/j.foodchem.2020.128172

LeCun, Y., & Misra, I. (2021). *Self-supervised learning: The dark matter of intelligence*. Meta AI. https://ai.facebook.com/blog/self-supervised-learning-the-dark-matter-of-intelligence/

Mavromatis, T., Koufos, G. C., Koundouras, S., & Jones, G. V. (2020). Adaptive capacity of winegrape varieties cultivated in Greece to climate change: Current trends and future projections. *OENO One*, *54*(4), 1201–1219. https://doi.org/10.20870/oeno-one.2020.54.4.3129

Musanase, C., Vodacek, A., Hanyurwimfura, D., Uwitonze, A., & Kabandana, I. (2023). Data-Driven Analysis and Machine Learning-Based Crop and Fertilizer Recommendation System for Revolutionizing Farming Practices. *Agriculture*, *13*(11), Article 11. https://doi.org/10.3390/agriculture13112141

Pan, S. J., & Yang, Q. (2010). A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, *22*(10), 1345–1359. https://doi.org/10.1109/TKDE.2009.191

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., … Chintala, S. (2019). *PyTorch: An Imperative Style, High-Performance Deep Learning Library* (arXiv:1912.01703). arXiv. https://doi.org/10.48550/arXiv.1912.01703

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.

Prada, J., Dinis, L.-T., Soriato, E., Vandelle, E., Soletkin, O., Uysal, Ş., Dihazi, A., Santos, C., & Santos, J. A. (2024). Climate change impact on Mediterranean viticultural regions and site-specific climate risk-reduction strategies. *Mitigation and Adaptation Strategies for Global Change*, *29*(6), 52. https://doi.org/10.1007/s11027-024-10146-0

Puga, G., Anderson, K., Jones, G., Tchatoka, F. D., & Umberger, W. (2022). *A climatic classification of the world's wine regions*. *56*(2). https://doi.org/10.20870/oeno-one.2022.56.2.4627

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., & Sutskever, I. (2021). *Learning Transferable Visual Models From Natural Language Supervision* (arXiv:2103.00020). arXiv. https://doi.org/10.48550/arXiv.2103.00020

Radhakrishnan, A., Friedman, S. F., Khurshid, S., Ng, K., Batra, P., Lubitz, S. A., Philippakis, A. A., & Uhler, C. (2023). Cross-modal autoencoder framework learns holistic representations of cardiovascular state. *Nature Communications*, *14*(1). https://doi.org/10.1038/s41467-023-38125-0

Reshef, N., Walbaum, N., Agam, N., & Fait, A. (2017). Sunlight Modulates Fruit Metabolic Profile and Shapes the Spatial Pattern of Compound Accumulation within the Grape Cluster. *Frontiers in Plant Science*, *8*. https://doi.org/10.3389/fpls.2017.00070

Robinson, J., Harding, J., & Vouillamoz, J. (2013). *Wine grapes: A complete guide to 1,368 vine varieties, including their origins and flavours*. Penguin UK.

Skahill, B., Berenguer, B., & Stoll, M. (2022). Temperature-based Climate Projections of Pinot noir Suitability in the Willamette Valley American Viticultural Area. *OENO One*, *56*(1), Article 1. https://doi.org/10.20870/oeno-one.2022.56.1.4889

Skahill, B., Berenguer, B., & Stoll, M. (2023). Climate Projections for Pinot Noir Ripening Potential in the Fort Ross-Seaview, Los Carneros, Petaluma Gap, and Russian River Valley American Viticultural Areas. *Agronomy*, *13*(3), Article 3. https://doi.org/10.3390/agronomy13030696

Tonietto, J., & Carbonneau, A. (2004). A multicriteria climatic classification system for grape-growing regions worldwide. *Agricultural and Forest Meteorology*, *124*(1), 81–97. https://doi.org/10.1016/j.agrformet.2003.06.001

van Leeuwen, C. (2022). 9 - Terroir: The effect of the physical environment on vine growth, grape ripening, and wine sensory attributes. In A. G. Reynolds (Ed.), *Managing Wine Quality (Second Edition)* (pp. 341–393). Woodhead Publishing. https://doi.org/10.1016/B978-0-08-102067-8.00005-1

van Leeuwen, C., Bois, B., Roby, J.-P., & Resseuguier, L. (2012). *Towards a unified terroir zoning methodology in viticulture.*

van Leeuwen, C., & Darriet, P. (2016). The Impact of Climate Change on Viticulture and Wine Quality. *Journal of Wine Economics*, *11*(1), 150–167. https://doi.org/10.1017/jwe.2015.21

van Leeuwen, C., Roby, J. P., & Resseguier, L. de. (2018). Soil-related terroir factors: A review. *OENO One*, *52*(2), 173. https://doi.org/10.20870/oeno-one.2018.52.2.2208

van Leeuwen, C., Schultz, H. R., Garcia de Cortazar-Atauri, I., Duchêne, E., Ollat, N., Pieri, P., Bois, B., Goutouly, J.-P., Quénol, H., Touzard, J.-M., Malheiro, A. C., Bavaresco, L., & Delrot, S. (2013). Why climate change will not dramatically decrease viticultural suitability in main wine-producing areas by 2050. *Proceedings of the National Academy of Sciences*, *110*(33), E3051–E3052. https://doi.org/10.1073/pnas.1307927110

Wallace, J. M., & Hobbs, P. V. (2006). 10—Climate Dynamics. In J. M. Wallace & P. V. Hobbs (Eds.), *Atmospheric Science (Second Edition)* (pp. 419–465). Academic Press. https://doi.org/10.1016/B978-0-12-732951-2.50015-6

Wang, T., Hamann, A., Spittlehouse, D., & Carroll, C. (2016). Locally Downscaled and Spatially Customizable Climate Data for Historical and Future Periods for North America. *PLOS ONE*, *11*(6), e0156720. https://doi.org/10.1371/journal.pone.0156720

Winkler, A. J. (1974). *General viticulture*. Univ of California Press.

Withler, C., & Geldart, G. (2023). *2022 B.C. Wine Grape Acreage Report*. Grape Growers BC. https://www.grapegrowers.bc.ca/sites/default/files/resource/files/2022%20BC%20Wine%20Grape%20Acreage%20Report_FINAL.pdf

Wolkovich, E. M., García de Cortázar-Atauri, I., Morales-Castilla, I., Nicholas, K. A., & Lacombe, T. (2018). From Pinot to Xinomavro in the world's future wine-growing regions. *Nature Climate Change*, *8*(1), 29–37. https://doi.org/10.1038/s41558-017-0016-6

Wu, L., Li, S., Hsieh, C.-J., & Sharpnack, J. (2019). *SSE-PT: Sequential Recommendation Via Personalized Transformer*. https://openreview.net/forum?id=HkeuD34KPH

A1. Appendix

Figure A1: Locations at which the Pinot Noir Potential was evaluated.