

# Variational Autoencoders for Generating Hyperspectral Imaging Honey Adulteration Data

Tessa Phillips

University of Auckland  
Auckland, New Zealand

tphi709@aucklanduni.ac.nz

Waleed Abdulla

University of Auckland  
Auckland, New Zealand

w.abdulla@auckland.ac.nz

## Abstract

*Honey fraud and adulteration are an increasing concern globally. Hyperspectral imaging and machine learning can detect adulterated honey within a known set of honey, where we have captured data at different sugar concentrations. Previous work in this area has used a minimal number of honey types, as sample preparation and data capture is a time-consuming process. This paper develops a new approach using variational autoencoders (VAEs) for generating adulterated honey data for unseen honey types. The results show that the binary adulteration detector can achieve on average 81.3% accuracy on unseen honey types by adding the generated data to the existing training data. Without including the generated data while training, the classifier can only achieve 44% on unseen honey types.*

## 1. Introduction

There is an increasing risk of honey fraud [1], as honey is a large and growing industry globally and a high-value export in New Zealand (NZ) [8]. Quality assurance is essential to protecting high-quality honey, as honey is the third most adulterated food product globally [3, 8, 20]. Honey fraud typically involves the dilution of honey with cheap sugar syrups, processing honey unnaturally to change its properties, incorrect labelling of geographical or botanical origin, or feeding bees sugars rather than botanical nectars [3].

There are existing quality assurance measures for premium honey types such as NZ Manuka honey [4]. Hyperspectral imaging can be used as a tool for detecting the botanical origins and quality of pure honey [7, 12–16]. NZ honey is a premium product and sells for 20% more than other honey sold in the USA [8]. These honey types should not be adulterated with other ingredients such as sugar syrup.

Honey adulteration can also be detected using spectroscopy and hyperspectral imaging. Most of the exist-

ing work on honey adulteration detection has focussed on a minimal set of different honey types, or in many cases, one honey type. Several different types of sugar syrup were added to Manuka honey and captured using near-infrared spectroscopy, and aquaphotomics [24]. This work shows promise in using spectral approaches to classify adulteration with many different adulterants. The experiment used one type of Manuka honey, so it is uncertain if this approach is suitable for other types of honey. Several papers have investigated the use of Fourier transform infrared (FTIR) spectroscopy to detect the adulteration of honey with cane sugar [6, 23]. The honey was adulterated between 0.5% and 25% cane sugar concentration. Only one type of honey was used in [6] which predicted the sugar concentration with an accuracy of 93.75%. When using three different honey types to classify adulteration [23], the classification accuracy was below 80%.

Previous research has been done on hyperspectral imaging to detect adulteration. A small dataset of 56 samples was used to detect adulteration with 95% accuracy [21]. This approach uses a neural network to calculate a percentage of pixels in the image to be either sugar or honey. Recent work has captured a much larger database of 8525 training and testing examples from 341 honey samples covering 11 different honey botanical origin types from seven unique brands. The honey was adulterated at 5%, 10%, 25%, and 50% for each honey type [18].

While a larger data set of adulterated honey with 8525 examples covering 11 different honey types have been developed using hyperspectral imaging [18], it still covers far fewer kinds of honey than have been captured as pure honey samples [19]. Sample preparation and hyperspectral image capture are time-consuming. Hence, the datasets used for adulteration detection have typically been tiny. A larger dataset of adulterated honey is required to detect adulteration of many different honey types.

Recent advances in machine learning have improved data generation techniques. Generative adversarial networks (GANs) and variational autoencoders (VAEs) can generate

very realistic fake data using deep neural networks trained on the actual data. GANs work by using two networks, a generator and an adversary. The generator network learns how to generate realistic fake data, and the adversary learns how to determine if the data is real or fake [5]. These two networks are trained in parallel; thus, the other also has to improve as each one improves. This technique can generate very realistic fake data, including images that are difficult for humans to detect as fake.

Variational autoencoders extend a traditional autoencoder structure, where they learn a distribution in the latent (feature) space for every training example. [22]. VAEs create new data by reconstructing the sampled latent space with the decoder network. The latent space created by VAEs satisfies certain conditions, which make it easy to manipulate to create variations on the data. The latent space must be continuous and map similar data close together in an organised way. Using regularisation on the latent space enforces these conditions [2]. This regularisation uses the Kullback-Leibler (KL) divergence between the returned latent distribution and a standard Gaussian [2, 22]. By including this regularisation term in the loss function for the VAE, we can now manipulate the latent space [2].

The variational class embodiment autoencoder (VCEAE) was developed to have higher generalisation performance on unseen data [15]. This technique is a combination of a VAE and a class embodiment autoencoder (CEAE) [14], which is a semi-supervised autoencoder that utilises weighted class labels to train the latent space. We consider this technique for data generation, as the latent space has the properties for data transformation similar to a VAE. The generated data is trained using class labels and input data for the VCEAE technique.

This work aims to generate fake honey adulteration data with a VAE and VCEAE by utilising the small adulteration dataset. The data generator can then create a generated dataset covering a more comprehensive set of honey types for adulteration detection.

Specifically, the objectives are:

- Construct a VAE to generate fake data from new honey types trained on the existing honey adulteration dataset.
- Evaluate the data generated using a leave one out evaluation strategy.
- Apply the data generation to the existing honey botanical origins dataset.

The organisation of the paper is as follows: Material and Methods describes the data preparation, the design of the VAEs and the evaluation strategy used. Results and discussion shows the generated data compared to the real data,

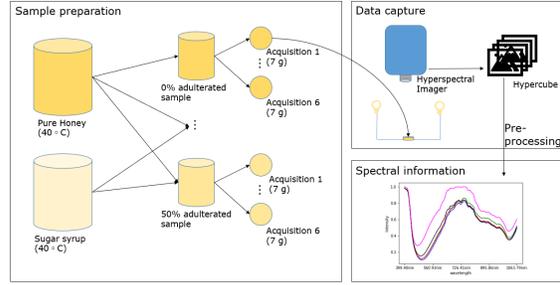


Figure 1. Showing the sample preparation process for one honey type in our database.

along with tabulated results and a short discussion. The paper is then concluded, and future work presented.

## 2. Material and Methods

This section details the methods used to generate fake adulteration data for honey using variational autoencoders and an existing dataset. This section discusses the dataset collected and used for honey adulteration and the subset we apply for this work. We also discuss variational autoencoder architectures used to generate the new data and the evaluation techniques.

### 2.1. Honey adulteration dataset

This section provides an overview of the methodologies we use to capture hyperspectral images of adulterated honey. The dataset is available online [18]. Figure 1 illustrates the procedure followed for sample preparation and capturing of this dataset.

#### 2.1.1 Data preparation

Homogenous samples from a wide range of honey types were prepared to the desired sugar concentrations of 5%, 10%, 25%, 50% [10]. Existing work has developed and used a dome system with halogen bulbs which ensures an even broadband light source over the entire sample [10]. The images of honey have been captured with a hyperspectral imager SOC-710 from Surface Optic. The hypercubes captured are between 400 – 1000nm in wavelength with a 5nm increment and 520x696 spatial resolution. This wide range of wavelengths has been used in previous work in this application [9, 10].

Because a hyperspectral imaging system measures the intensity of light reflected into the camera, the calibration process is crucial to obtaining consistent images. Calibration is performed using a dynamic white reference technique [11].

Segmentation is a key preprocessing step in our work. This technique is a point of difference between the two

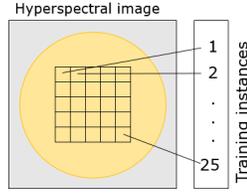


Figure 2. Segmentation process for our dataset, showing how the image is split into a five by five grid to obtain 25 samples.

groups working in this area [10, 21]. The approach we take is to use segmentation rather than directly using the entire hyperspectral images, as this enables training of machine learning techniques with more training examples. Figure 2 clearly illustrates this process.

Normalisation is applied to the final dataset before using machine learning techniques to speed up the training process. The normalisation technique is a standard scalar, which scales the standard deviation and means to one and zero, respectively [12].

### 2.1.2 Data makeup

The dataset [18] comprises 12 different honey products from seven different brands with 11 different botanical origins labels. Half of the samples are from Manuka honey, a premium NZ honey type, and the other half are from other NZ honey. Table 1 shows the makeup of the dataset from these different kinds of honey.

Table 1. Overall make-up of the adulterated honey dataset from each brand and botanical origins label of honey.

Brand	Class	Adulteration Concentration					Sum
		0%	5%	10%	25%	50%	
C1	Clover	150	150	300	300	300	1200
	MultiFloral	150	150			150	450
C10	ManukaUMF5		150	150	150	150	600
	ManukaUMF15		150	150	150	150	600
	ManukaUMF20		150	150	150	150	600
C4	ManukaUMF10		150	150	150	125	575
C5	ManukaBlend		150		150	150	450
C7	BorageField		150	150	150	150	750
	Kamahi		150	150	150	150	750
	Rewarewa		150	150	150		450
C8	ManukaBlend		150	150	150	150	750
C9	Manuka		150	300	300	300	1350

It is only feasible to use data from classes with complete information for this work on generating new data. The dataset was limited to only include data from the following brand and honey combinations: C4 Multifloral, C7 Borage-Field, C7 Kamahi, and C8 ManukaBlend.

### 2.2. Variational Autoencoder structures

The variational autoencoder structure we used has been explained in [15]. We consider both the traditional variational autoencoder and the variational class embodiment autoencoder (VCEAE) for this work. The benefit of the VCEAE over the traditional variational autoencoder is that

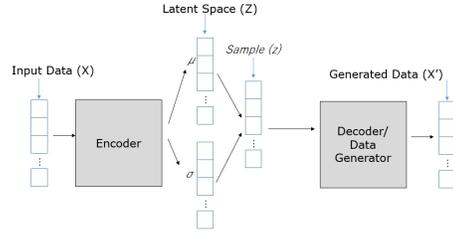


Figure 3. Variational autoencoder structure

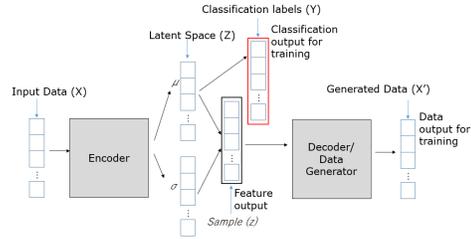


Figure 4. VCEAE architecture

the feature space categorizes the class labels in a more accurate way. However, a potential downside is that the overall impact of regularisation may be less; thus, the latent space may not be as trivial to manipulate when generating data.

The variational autoencoder structure used is given in figure 3. The overall structure of the decoder and encoder is two fully connected layers of size 128 and 74, respectively. The latent space or encoding layer is of size 20. A dropout rate of 0.001 is used to prevent over-fitting. The network is trained with a batch size of 32 and a learning rate of 0.001 for 100 epochs.

The VCEAE network parameters are very similar to the variational autoencoder, using the same encoding and decoding networks. However it also applies a classification weight of 0.4 for training the classification output. This parameter value has been used successfully in previous work [15]. The VCEAE architecture used is shown in detail in figure 4.

### 2.3. Evaluation Strategy

Because we have selected only a small number of honey types to be used in this work, it is critical to use a cross-validation strategy to ensure fair evaluation of this work on a range of data. The strategy used is a leave-one-out cross-validation strategy, where for all the four kinds of honey, each honey is put aside for validation, with the other three used for training. The VAE training and the latent space transforms only use the training honey, and this is then applied and evaluated on the validation honey.

Because there are only four kinds of honey used in this strategy, visualisation of the results will be shown on all the

validation honey individually, and some results are averaged over the cross-validation performance.

### 2.3.1 Using regression to interpolate between concentrations

The benefit of a variational autoencoder is that it is possible to transform the latent space to generate new data between data points or related to the existing data in some way. We utilise this in our work by transforming un-adulterated honey into adulterated honey and evaluating the performance compared to the ground truth at 5%, 10%, 25%, and 50%.

The methodology used here is to find the transform vectors for each honey type and concentration, the vector from the mean of un-adulterated honey to the mean of each concentration of that same honey. These vectors for all honey types in the training set are used in a linear regression problem to find the transformation between un-adulterated honey and honey of any concentration. The latent space transform is linear, as adding complexity to this model would likely over-fit with such a small dataset.

## 2.4. Applying the models to the broader honey dataset

Although evaluation of the technique is performed using cross-validation, we use the entire dataset to train the autoencoder models for application of this technique in a broader dataset of honey found at [19]. This dataset has no ground-truth data, so it cannot be used for evaluation but will provide a valuable dataset of generated data for determining adulteration in honey.

## 3. Results and Discussion

This section discusses the results of the variational autoencoders for data generation. The results are compared to the ground truth work using the cross-validation strategy discussed in section 2.3.

We first use t-distributed stochastic neighbour embedding (t-sne) to evaluate how spaced out our data is from each concentration to another. Figure 5 shows a t-sne representation of our entire data space (training plus validation data). This figure shows that our data space is generally made up of groups; these are the individual honey types at each concentration. The concentrations are not all aligned to one cluster, and the transformation between each concentration is non-trivial.

### 3.1. Cross-Validation

We now utilise each honey type as a validation set to compare our results to ground truth data. Each honey in the set (Manuka, ManukaBlend, Kamahi, and BorageField)

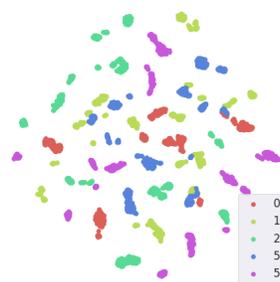


Figure 5. T-sne representation of training and validation data.

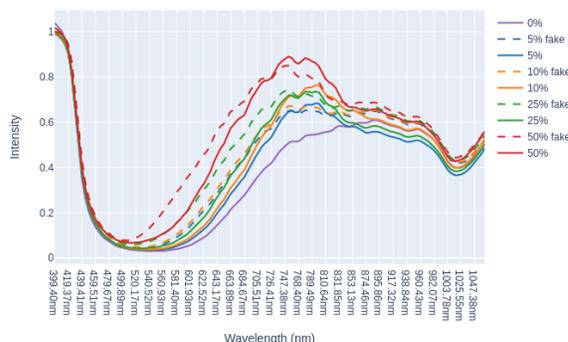


Figure 6. Average spectrum of Manuka honey, showing the mean spectrum of the ground truth data vs the VAE generated data.

will be used as validation honey, using the rest for training. This evaluation considers both the VAE and the VCEAE.

### 3.1.1 Validation - Manuka

Figure 6 shows the mean ground truth compared to the mean generated spectrum for the validation honey type, Manuka using the VAE. The ground truth values are labelled 0%, 5%, 10%, 25%, and 50% whereas the generated data is labelled as 5% fake, 10% fake, 25% fake, and 50% fake. This data shows that the generator did well creating adulterated honey data, particularly in the lower wavelengths. The 10% adulterated honey was not well matched to the ground truth; however, the spectrum looks to be within a reasonable approximation for the other concentrations. For the 50% concentration, the generated data does not peak to the same value as the ground truth data. This discrepancy could be a consequence of the linear approximation we have made when transforming the latent space, as the 50% adulterated honey has a much higher peak in its spectral response than the other concentrations.

Figure 7 shows the mean ground truth compared to the mean generated spectrum for the validation honey type, Manuka using the VCEAE. These results show that overall the data generator struggled on this honey type. There are discrepancies between the real and generated honey at each concentration, particularly in the higher wavelengths.

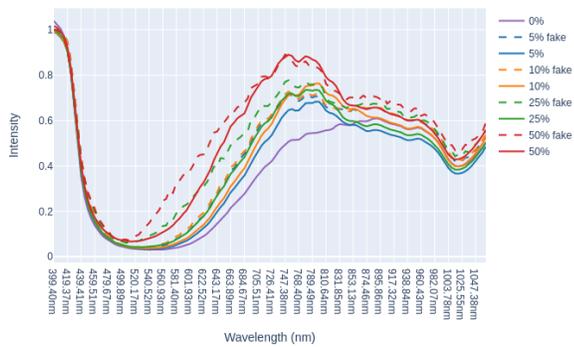


Figure 7. Average spectrum of Manuka honey, showing the mean spectrum of the ground truth data vs the VCEAE generated data.

The 50% concentration has the closest comparison between generated data and ground truth, however, the approximation is not entirely accurate around the peak of the spectral response. The spectrum represents a reasonable approximation; however, a classifier might struggle to correctly classify the ground truth data when trained on the generated data.

Compared to the VAE, the results from using the VCEAE were not as accurate, except for 50% adulterated honey. The VAE had a much better fit for the 10% and 5% concentrations and was still primarily accurate for the 50% honey. There was minimal difference between the two responses for 25% adulterated honey.

### 3.1.2 Validation - ManukaBlend

Figure 8 shows the mean ground truth compared to the mean generated spectrums for the validation honey type, ManukaBlend using the VAE. The spectral response matches very well for the 50% concentration, however it does not match well for 25%. For 5% and 10% concentrations, the spectral response generally matched well; however, it would not accurately distinguish these two concentrations from each other. This error is likely because the honey types used in training the VAE are struggling to generalise to this data type.

Figure 8 shows the mean ground truth compared to the mean generated spectrums for the validation honey type, ManukaBlend using the VCEAE. The response differed from the ground truth data for 25% and 50% adulterated honey. The response for 5% and 10% adulterated honey was reasonable; however, the generated 5% response matched better with the 10% ground truth than 5% ground truth. Similarly, the 10% generated data better matched with the 5% ground truth.

Comparing the VAE and VCEAE for ManukaBlend honey, we can see that the VAE had an overall better comparison to the ground truth data. The 50% concentration

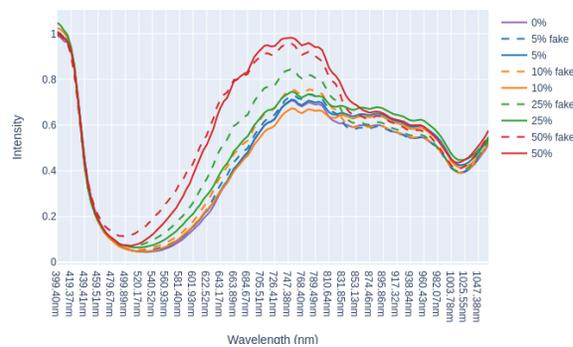


Figure 8. Average spectrum of ManukaBlend honey, showing the mean spectrum of the ground truth data vs the VAE generated data.

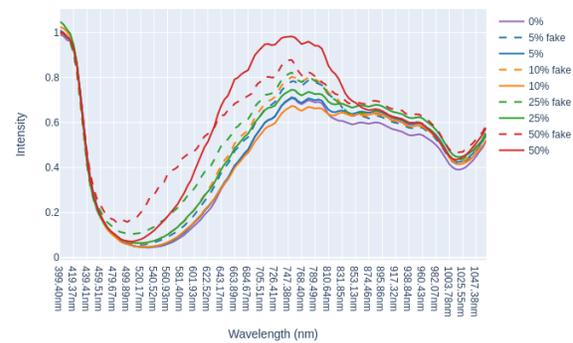


Figure 9. Average spectrum of ManukaBlend honey, showing the mean spectrum of the ground truth data vs the VCEAE generated data.

matched much better, and the other concentrations of adulteration were comparable between the VAE and VCEAE.

### 3.1.3 Validation - Kamahi

Figure 10 shows the mean ground truth compared to the mean generated spectrums for the validation honey type, Kamahi using the VAE. Overall the results show that the data generator can represent this honey type well; however, it struggles on the 10% and 25% concentrations. Overall, the generator creates data that should represent the adulteration; however, a classifier trained on the generated data would struggle to accurately classify the ground truth data for a multi-class classification problem.

Figure 10 shows the mean ground truth compared to the mean generated spectrum for the validation honey type, Kamahi using the VCEAE. These results show that the response with the VCEAE is more accurate around the peak for this honey type; however, it is not very accurate around lower and higher wavelengths, especially for the 50% adulterated honey. Compared to the VAE on this honey type, the response is better around the middle wavelength range but worse at the lower and upper wavelengths.

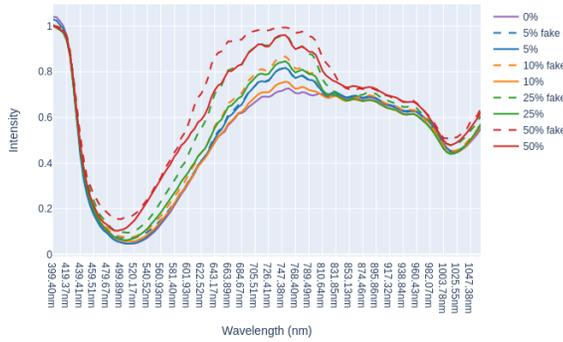


Figure 10. Average spectrum of Kamahi honey, showing the mean spectrum of the ground truth data vs the VAE generated data.

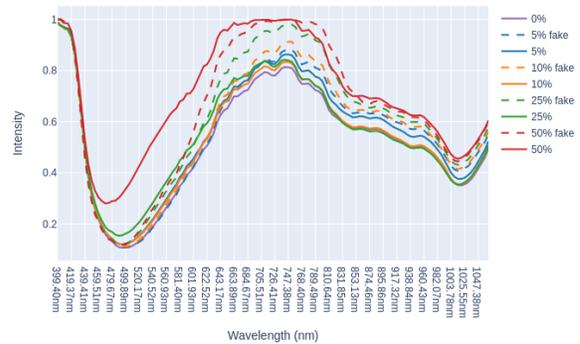


Figure 12. Average spectrum of BorageField honey, showing the mean spectrum of the ground truth data vs the VAE generated data.

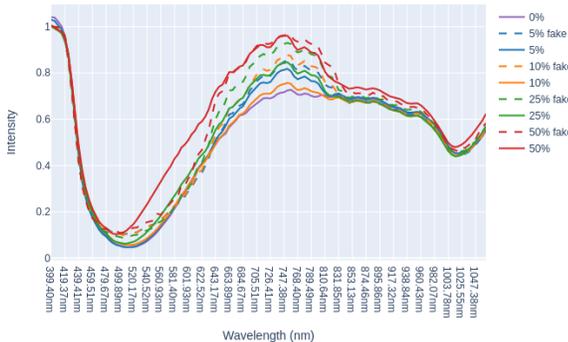


Figure 11. Average spectrum of Kamahi honey, showing the mean spectrum of the ground truth data vs the VCEAE generated data.

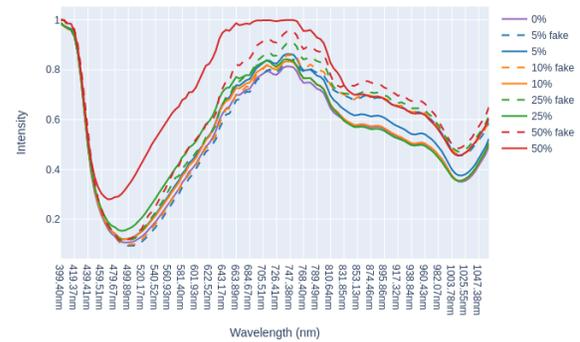


Figure 13. Average spectrum of BorageField honey, showing the mean spectrum of the ground truth data vs the VCEAE generated data.

### 3.1.4 Validation - BorageField

Figure 12 shows the mean ground truth compared to the mean generated spectrums for the validation honey type, BorageField using the VAE. These results show that overall the data generator struggled on this honey type. For all concentrations, the average spectrum does not line up. For the 50% adulterated spectrum, this was in the lower wavelengths, and for the rest, the mid wavelength range did not line up with the ground truth data. The shape of the spectrums was similar; however, the average values are not accurate for any concentration. For 5% and 10% honey, the generated spectrum is close enough to the ground truth to detect adulteration accurately, but a multi-class adulteration concentration classifier would struggle to classify the ground truth accurately.

Figure 12 shows the mean ground truth compared to the mean generated spectrums for the validation honey type, BorageField using the VCEAE. These results show that the VCEAE did not perform well on this data type. Similarly to the VAE, the overall shape was reasonably accurate for the 50% adulterated honey; however, the average values did not match the ground truth. The generated values were very different from the ground truth for the lower concentrations.

Overall, the VAE and VCEAE did a reasonable job approximating the data for most honey types. The VAE is typically more accurate across more concentrations, but the VCEAE tended to fit the 50% adulterated data better.

### 3.2. Classification Results

We used the generated data alongside the training data from other honey types to train classifiers for multi-class and binary adulteration classification. The results are for each validation fold and an average result calculated using a macro average. We use a traditional KNN classification technique with  $K = 5$ , which has performed well on honey classification previously [14, 15].

Although feature reduction with autoencoders and SVMs have commonly been used to classify honey quality [14–16], a KNN classifier will be suitable to show the impact of the generated data on the generalisation performance. We compare our techniques to an equivalent classifier trained on only the training data without the generated data from the same honey type.

Table 2 shows the results of this for the binary classification case, where the task is detecting adulteration in honey. These results show that by adding the generated

Table 2. Binary classification results for adulteration detection on an unseen honey type using leave one out strategy comparing different data generation methods.

Generator	Honey Type	Precision	Recall	F1	Avg Acc
VAE	BorageField	0.720	0.650	0.540	0.717
	Kamahi	0.870	0.740	0.750	0.868
	Manuka	0.750	0.600	0.500	0.753
	ManukaBlend	0.920	0.800	0.830	0.915
	Average	0.815	0.698	0.655	0.813
VCEAE	BorageField	0.620	0.620	0.400	0.625
	Kamahi	0.800	0.690	0.660	0.804
	Manuka	0.790	0.610	0.550	0.789
	ManukaBlend	0.930	0.820	0.880	0.928
	Average	0.785	0.685	0.623	0.787
No Data	BorageField	0.500	0.400	0.440	0.500
	Kamahi	0.500	0.400	0.440	0.500
	Manuka	0.260	0.400	0.320	0.260
	ManukaBlend	0.500	0.400	0.440	0.500
	Average	0.440	0.400	0.410	0.440

data, the classification accuracy can improve with the generated adulteration data. Both the VAE and VCEAE have improved over the benchmark case with no data generator. The highest accuracy on the unseen honey type is 81.3% using the VAE. This accuracy is a considerable improvement over the average accuracy of 44% without using generated data. Data capture for honey adulteration with a hyperspectral camera is a very time-consuming process, so using this generated data to supplement the actual adulteration data helps to improve the generalisation performance of an adulteration detector.

Table 3 shows the unseen honey type results for the multi-class adulteration concentration classification problem. These results show that the generated data is not suitable for the multi-class problem in contrast to the binary classification problem. There is still an improvement over using no generated data; however, the accuracy is far from acceptable in the real world. Interestingly the Manuka honey type struggles to get accurate classifications for all three data generation strategies with lower performance than the other honey types. This poor performance might be due to the data being quite different to the other honey types. The ManukaBlend honey type had the best performance, especially with the VCEAE data generation strategy. It is reasonable to assume that the data generation strategies could be helpful for this problem in future, and perhaps with more comprehensive data collected, we could generate more realistic fake data.

These results are a positive step towards an overall adulteration detection system. With an accuracy of 81.3% on unseen honey types for binary adulteration detection, we can use this as part of an overall honey quality detection system combined with botanical origins classification. By utilising generated data, we can extend our adulteration dataset to the broader set of honey that has been captured previously [19]. Figure 14 shows the average spectrum for each concentration on our new generated adulteration dataset. We can see that the general trends that we saw in the real adulteration data is followed in the generated data.

This new dataset is available online for the research com-

Table 3. Multi-class classification results for adulteration classification on an unseen honey type using leave one out strategy comparing different data generation methods.

Generator	Honey Type	Precision	Recall	F1	Avg Acc
VAE	BorageField	0.440	0.460	0.360	0.443
	Kamahi	0.330	0.330	0.270	0.333
	Manuka	0.230	0.250	0.120	0.149
	ManukaBlend	0.430	0.260	0.330	0.433
	Average	0.358	0.325	0.270	0.339
VCEAE	BorageField	0.400	0.250	0.280	0.400
	Kamahi	0.460	0.310	0.350	0.463
	Manuka	0.330	0.450	0.250	0.261
	ManukaBlend	0.430	0.380	0.390	0.431
	Average	0.405	0.348	0.318	0.389
No Data	BorageField	0.230	0.080	0.120	0.233
	Kamahi	0.210	0.100	0.140	0.211
	Manuka	0.050	0.240	0.070	0.056
	ManukaBlend	0.220	0.160	0.180	0.224
	Average	0.178	0.145	0.128	0.181

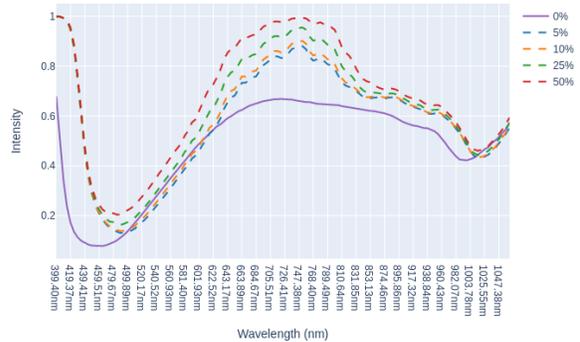


Figure 14. Average spectrum of generated adulteration data applied to the broader honey botanical origins dataset.

munity [17]. This now provides generated adulteration data for 21 unique botanical origins of honey from 11 brands. This dataset includes samples of premium Manuka honey, as well as non-Manuka honey types.

## 4. Conclusion

We have successfully used a VAE and VCEAE to generate realistic new data for adulterated honey hyperspectral images. The data can improve the generalisation of different honey types with no adulterated honey images. These techniques were evaluated with a leave-one-out strategy for each honey type and achieved an average of 81.3% on unseen honey types using a VAE to generate supplemental data. This performance is an improvement from 44% without generated data on unseen honey types. We have applied this data to the much more comprehensive honey botanical origins dataset, which has created a new adulteration dataset covering many honey types and brands.

## References

- [1] Canadian Food Inspection Agency. Report: Enhanced honey authenticity surveillance (2018 to 2019). Technical report, Canadian Food Inspection Agency, 2019.

- [2] Jaehoon Cha, Kyeong Soo Kim, and Sanghyuk Lee. On the transformation of latent space in autoencoders. *arXiv preprint arXiv:1901.08479*, 2019. 2
- [3] Norberto L García. The current situation on the international honey market. *Bee World*, 95(3):89–94, 2018. 1
- [4] Alodia Girma, Wonjae Seo, and Rosemary C She. Antibacterial activity of varying umf-graded manuka honeys. *PloS one*, 14(10):e0224495, 2019. 1
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [6] J Irudayaraj, R Xu, and J Tewari. Rapid determination of invert cane sugar adulteration in honey using ftir spectroscopy and multivariate analysis. *Journal of food science*, 68(6):2040–2045, 2003. 1
- [7] Saeid Minaei, Sahameh Shafiee, Gerrit Polder, Nasrolah Moghadam-Charkari, Saskia van Ruth, Mohsen Barzegar, Javad Zahiri, Martin Alewijn, and Piotr M Kuś. Vis/nir imaging application for honey floral origin determination. *Infrared Physics & Technology*, 86:218–225, 2017. 1
- [8] New Zealand Consulate-General Los Angeles. *The New Zealand honey phenomenon in the USA*. New Zealand Consulate-General Los Angeles, 2015. 1
- [9] Ary Noviyanto. *Honey Botanical Origin Classification using Hyperspectral Imaging and Machine Learning*. PhD thesis, The University of Auckland, 2018. 2
- [10] Ary Noviyanto and Waleed H Abdulla. Honey dataset standard using hyperspectral imaging for machine learning problems. In *2017 25th European Signal Processing Conference (EUSIPCO)*, pages 473–477. IEEE, 2017. 2, 3
- [11] Ary Noviyanto and Waleed H Abdulla. Segmentation and calibration of hyperspectral imaging for honey analysis. *Computers and Electronics in Agriculture*, 159:129–139, 2019. 2
- [12] Ary Noviyanto and Waleed H Abdulla. Honey botanical origin classification using hyperspectral imaging and machine learning. *Journal of Food Engineering*, 265:109684, 2020. 1, 3
- [13] Ary Noviyanto and Waleed H. Abdulla. Signifying the information carrying bands of hyperspectral imaging for honey botanical origin classification. *Journal of Food Engineering*, 292:110281, 2021. 1
- [14] Tessa Phillips and Waleed Abdulla. Class embodiment autoencoder (CEAE) for classifying the botanical origins of honey. In *2019 International Conference on Image and Vision Computing New Zealand (IVCNZ)*. IEEE, dec 2019. 1, 2, 6
- [15] Tessa Phillips and Waleed Abdulla. Generalisation techniques using a variational ceae for classifying manuka honey quality. In *2020 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 1631–1640. IEEE, 2020. 1, 2, 3, 6
- [16] Tessa Phillips and Waleed Abdulla. Developing a new ensemble approach with multi-class svms for manuka honey quality classification. *Applied Soft Computing*, page 107710, 2021. 1, 6
- [17] Tessa Phillips and Waleed Abdulla. Hyperspectral imaging generated adulterated honey dataset, Feb 2022. Available at [https://auckland.figshare.com/articles/dataset/Hyperspectral\\_imaging\\_generated\\_adulterated\\_honey\\_dataset/19134698](https://auckland.figshare.com/articles/dataset/Hyperspectral_imaging_generated_adulterated_honey_dataset/19134698). 7
- [18] Tessa Phillips, Bradley Coleman, Shunji Takano, and Waleed Abdulla. Hyperspectral imaging adulterated honey dataset, Aug 2021. Available at [https://auckland.figshare.com/articles/dataset/Hyperspectral\\_Imaging\\_adulterated\\_honey\\_dataset/16441686/1](https://auckland.figshare.com/articles/dataset/Hyperspectral_Imaging_adulterated_honey_dataset/16441686/1). 1, 2, 3
- [19] Tessa Phillips, Ary Noviyanto, and Waleed Abdulla. Hyperspectral imaging honey database. 4 2020. Available at <https://figshare.com/s/25afe30ff531b8f1e65f>. 1, 4, 7
- [20] Ron Phipps. *International Honey Market*. CPNA International, Ltd., 2020. 1
- [21] Sahameh Shafiee, Gerrit Polder, Saeid Minaei, Nasrolah Moghadam-Charkari, Saskia van Ruth, and Piotr M. Kuś. Detection of honey adulteration using hyperspectral imaging. *IFAC-PapersOnLine*, 49(16):311 – 314, 2016. 5th IFAC Conference on Sensing, Control and Automation Technologies for Agriculture AGRICONTROL 2016. 1, 3
- [22] Irhum Shafkat. Intuitively understanding variational autoencoders. Online. 2
- [23] S Sivakesava and J Irudayaraj. Prediction of inverted cane sugar adulteration of honey by fourier transform infrared spectroscopy. *Journal of food science*, 66(7):972–978, 2001. 1
- [24] Xinhao Yang, Peiwen Guang, Guoze Xu, Siqi Zhu, Zhenqiang Chen, and Furong Huang. Manuka honey adulteration detection based on near-infrared spectroscopy combined with aquaphotomics. *LWT*, 132:109837, 2020. 1