

## Gold Standard

Dear 229 Project TAs,

Please accept our milestone report (attached as a separate pdf) and this brief summary:

<b>Members</b>	Dave Deriso, Bryan McCann, Vincent van Gogh		
<b>Title</b>	Classification of Van Gogh's Irises		
<b>Predicting</b>	We aim to classify the kind of irises in Van Gogh's garden		
<b>Data</b>	Data are publicly available <a href="#">here</a> (1 .csv file). There are (150) rows, each representing a <i>unique</i> example flower. There are (5) columns: Sepal Length (continuous), Sepal Width (continuous), Petal Length (continuous), Petal Width (continuous), and Species (class). Species is <i>labeled</i> and represent our ground-truth data with which we will train our model.		
<b>Features</b>	We have a 10 dimensional feature space consisting of: [the first (4) columns of the raw input data, 3 features based on PCA, 3 features based on gaussian kernels]. Our literature review has indicated that a lower dimensional representation of these data will improve the model more than the raw inputs alone (see report). An L1 shrinkage parameter is employed in the regression to discern whether or not a feature is informative (see results section in report).		
<b>Models</b>	We have continuous inputs and have a categorical output which go well with SVM, LDA, multinomial logistic regression, and naive bayes classifiers. We have implemented our models in Python.		
<b>Results</b>	<i>Training set = 75 random samples, Test Set = remaining 75 samples</i>		
	Model	Training MSE	Test MSE
	SVM	.9432	.7563
	LDA	...	...
	MNLR	...	...
	Naive Bayes	...	...
	Neural Network	...	...
<b>Future</b>	We are planning to continue adding new features, and we may even try some more advanced algorithms. We'd like to tap into some of the deep learning/neural network approaches Professor Ng mentioned, but we're not sure that we have enough data. If this doesn't go well, we are going to focus on improving existing models by training hyperparameters (such as the shrinkage term).		
<b>Specific Questions</b>	We have no specific questions, but will come to office hours for more help. If you must give us feedback, we'd like to know what other methods we should look at, but at this time there are no burning questions.		

Cheers,

Dave Deriso, Bryan McCann, Vincent van Gogh

## Details About Each Section

<b>Predicting</b>	What are you building? Specifically, what are the inputs and outputs. Don't explain the motivation for your topic. <b>(1-2 sentences max)</b>
<b>Data</b>	Exactly where did your data come from and what does your contain? (ie. What are in the rows and columns? Are examples labeled with ground truth? If you have images, are they color, normalized, etc?) <b>(2-3 sentences max)</b>
<b>Features</b>	How many features do you have and which features are the raw input data (ex. color, weight, location, etc) vs. features you have derived (ex. ICA, Gaussian Kernel)? Why they are appropriate for this task? <b>(3-4 sentences max)</b>
<b>Models</b>	Exactly which model(s) are you using? Why they are appropriate for this task? Why have you chosen to compare/contrast each model? <b>(3-4 sentences max)</b>
<b>Results</b>	Make a compact table of results. Each row should be a different model. The columns should be the training error and the test error. List how many samples are in each of the training and testing data sets. Obviously, these sets should be different. <b>(1-2 sentences max + 1 table max)</b>
<b>Future</b>	List plans for the future. <b>(2-3 sentences max)</b>
<b>Specific Questions</b>	In brief, if we could provide specific feedback about some aspect of your project, what could it be? Please make this succinct. <b>(1-2 sentences max)</b>