

Data Science Applications

Assignment Semester 2 2025 – Marking Rubric





Marking Rubric

The following rubric outlines the specific standards required for each question in the assignment.

Criteria in bold at the top of each rubric box provide a holistic view on what differentiates answers under each rubric grade. Other criteria in each rubric box provide more specific guidance about components of an answer that are usually required to meet each holistic criterion.

Criteria designated as 'hurdle' in the left most column are those that must be met to gain a certain grade (usually either a grade of 2 or 3).

Criteria designated as 'differentiator' in the left most column are those that help to distinguish between assignment responses at each rubric grade.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
1a	5%	Demonstrates a very strong ability to apply text vectorisation.	Demonstrates a good ability to apply text vectorisation.	Demonstrates some ability to apply text vectorisation.	Demonstrates a limited ability to apply text vectorisation.	Does not demonstrate an ability to apply text vectorisation.
Hurdle (for a 2)		Accurately vectorises the 'written_report' column, using word embeddings and TF-IDF.		Vectorises the 'written_report' column with some errors in the application of the vectorisation methods or only uses one of the two methods.		
Differentiator		Accurately applies and checks at least five different types of appropriate text cleaning steps.	Accurately applies and checks four different types of appropriate text cleaning steps.	Accurately applies and checks three different types of appropriate text cleaning steps.	Accurately applies and checks two different types of appropriate text cleaning steps.	Accurately applies and checks no more than one type of text cleaning step, or no accurate text cleaning applied.
Differentiator		Correctly interprets the outcomes of at least two types of appropriate checks on the vectorisations completed, in a thorough and insightful way.	Correctly interprets the outcomes of one appropriate check on the vectorisations completed, in a thorough and insightful way.	Correctly interprets the outcomes of at least one appropriate check on the vectorisations completed.	Applies at least one appropriate check on the vectorisations completed but does not interpret its outcome or does not interpret its outcome correctly.	Applies at least one check on the vectorisations completed but the check is not appropriate, and the outcome of the check is not interpreted or not interpreted correctly.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
1b	5%	Demonstrates a very strong ability to apply a clustering algorithm.	Demonstrates a good ability to apply a clustering algorithm.	Demonstrates some ability to apply a clustering algorithm.	Demonstrates a limited ability to apply a clustering algorithm.	Does not demonstrate an ability to apply a clustering algorithm.
Hurdle (for a 2)		Accurately applies a clustering algorithm to the dataset using subsets of both the embedding and TF-IDF features.				Applies a clustering algorithm to the dataset with errors in the way the algorithm has been applied or does not use subsets of both the embedding and TF-IDF features.
Hurdle (for a 3)		Examines the clustering outputs using internal validation, including a correct interpretation of the internal validation output.			Does not examine the clustering outputs or does not interpret the validation output or does not interpret the validation output correctly.	
Differentiator		Justifies the features and type of clustering algorithm used and the selection of all required hyperparameters (distance method, linkage method, and number of clusters) with comparison to alternative choices, and with reference to the business context.	Discusses the features and type of clustering algorithm used and the selection of all required hyperparameters (distance method, linkage method, and number of clusters), and the selections seem reasonable.	States the features and type of clustering algorithm used and the selection of some required hyperparameters (distance method, linkage method, and number of clusters), and the selections seem reasonable.	States the features and type of clustering algorithm used and the selection of some required hyperparameters (distance method, linkage method, and number of clusters), but the selections do not seem reasonable.	Does not select features or hyperparameters such as the distance method, linkage method, or number of clusters.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
1c	5%	Demonstrates a very strong ability to build and interpret a discriminator model.	Demonstrates a good ability to build and interpret a discriminator model.	Demonstrates some ability to build and interpret a discriminator model.	Demonstrates a limited ability to build and interpret a discriminator model.	Does not demonstrate an ability to build and interpret a discriminator model.
Hurdle (for a 2)		Accurately applies a random forest algorithm to the cluster analysis.				Applies a random forest algorithm to the cluster analysis, with errors in how the algorithm has been applied.
Hurdle (for a 3)		Accurately calculates variable importance.			Calculates the variable importance, with errors in how the algorithm has been applied.	
Differentiator		Examines the top 10 keywords and provides insights on the cluster analysis' validity and individual clusters' interpretation.	Discusses the top 10 keywords and their relationship to different clusters.	States the top 10 keywords.	States the calculated variable importances.	Does not comment on the calculated variable importances.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
1d	10%	Demonstrates a very strong ability to examine a clustering algorithm's outputs.	Demonstrates a good ability to examine a clustering algorithm's outputs.	Demonstrates some ability to examine a clustering algorithm's outputs.	Demonstrates a limited ability to examine a clustering algorithm's outputs.	Does not demonstrate an ability to examine a clustering algorithm's outputs.
Hurdle (for a 2)		Examines the clustering output using manual validation, including a correct interpretation of the output of the manual validation.				Does not examine the clustering output, does interpret the validation output, or does not interpret the validation output correctly.
Differentiator		Uses at least five unique manual validation analyses, including variables/features/responses beyond the pathology commentary text.	Uses four unique manual validation analyses, including variables/features/responses beyond the pathology commentary text.	Uses three unique manual validation analyses.	Uses two unique manual validation analyses.	Uses only the discriminator model output from Q1c.
Differentiator		Describes the key characteristics of each cluster with a strong link to the problem context. Intuitively labels each cluster.	Describes the key characteristics of each cluster with some link to the problem context.	Outlines the key characteristics of each cluster with some link to the problem context.	Outlines some characteristics of the clusters, but these are not key characteristics, or there is no link to the problem context.	Some characteristics of the clusters are outlined, but these are not based on the output of the clustering algorithm.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
1e	5%	Demonstrates a very strong ability to apply clustering outcomes to a business problem.	Demonstrates a good ability to apply clustering outcomes to a business problem.	Demonstrates some ability to apply clustering outcomes to a business problem.	Demonstrates a limited ability to apply clustering outcomes to a business problem.	Does not demonstrate an ability to apply clustering outcomes to a business problem.
Hurdle (for a 2)		Uses correct spelling and grammar.				Contains spelling and/or grammar mistakes.
Differentiator		Suggests actionable insights related to the usefulness of information provided by pathologists, based on the examination of the clustering outputs.	Describes how the clusters relate to the usefulness of information provided by pathologists, based on the examination of the clustering outputs.	Outlines how the clusters relate to the usefulness of information provided by pathologists, based on the examination of the clustering outputs.	Outlines the cluster attributes, but these are not related to the usefulness of information provided by pathologists.	Outlines the cluster attributes, but these are not related to any business relevant issues.
Differentiator		Communicates in a highly suitable way for the management team of Betahelp, with no edits required before use.	Communicates in a way that is suitable for the management team of Betahelp, with only minor edits required before use.	Communicates in a way that is mostly suitable for the management team of Betahelp.	Communicates in a way that is mostly not suitable for the management team of Betahelp.	Communicates in a way that is not suitable for the management team of Betahelp.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
2a	5%	Demonstrates a very strong ability to clean the dataset.	Demonstrates a good ability to clean the dataset.	Demonstrates some ability to clean the dataset.	Demonstrates a limited ability to clean the dataset.	Does not demonstrate an ability to clean the dataset.
Hurdle (for a 3)		Identifies privacy issues and effectively anonymises the data.			Does not identify privacy issues or does not anonymise the data.	
Hurdle (for a 4)		Splits the data into training, validation, and test sets at an appropriate stage to avoid leakage in the model to be built later in Question 2.		Does not split the data into training, validation, and test sets, or does not split the data at an appropriate stage to avoid leakage in the model to be built later in Question 2.		
Differentiator		Accurately applies and checks at least five different types of appropriate cleaning steps.	Accurately applies and checks four different types of appropriate cleaning steps.	Accurately applies and checks three different types of appropriate cleaning steps.	Applies at least three different appropriate cleaning steps, although the output of the cleaning is not checked, or there are errors in the steps undertaken.	Applies some cleaning steps, but these steps are not appropriate for the problem context.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
2b	5%	Demonstrates a very strong ability to propose a unit of analysis that is suitable for the given problem.	Demonstrates a good ability to construct a unit of analysis that is suitable for the given problem.	Demonstrates some ability to construct a unit of analysis that is suitable for the given problem.	Demonstrates a limited ability to construct a unit of analysis that is suitable for the given problem.	Does not demonstrate an ability to construct a unit of analysis that is suitable for the given problem.
Hurdle (for a 3)		States the entity and the basis of the timestamps.			Does not state the entity and/or does not state the basis of the timestamps.	
Hurdle (for a 3)		Constructs a pandas table containing the correct entity and timestamp values.			Does not construct a pandas table containing the correct entity and timestamp values.	
Differentiator		Proposes a suitable unit of analysis with strong reference to the problem context.	Proposes a suitable unit of analysis with some reference to the problem context.	Describes a suitable unit of analysis with reference to the problem context.	Describes a suitable unit of analysis but does not reference the problem context.	Describes a unit of analysis but it is not suitable for this problem context.
Differentiator		Justifies a suitable selection of timestamps (regular/event/both) with strong reference to the problem context.	Justifies a suitable selection of timestamps (regular/event/both) with some reference to the problem context.	Describes a suitable selection of timestamps (regular/event/both) with some reference to the problem context.	Describes a suitable selection of timestamps (regular/event/both) but does not reference the problem context.	Does not describe the selection of timestamps.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
2c	5%	Demonstrates a very strong ability to construct a response variable that is suitable for the given problem.	Demonstrates a good ability to construct a response variable that is suitable for the given problem.	Demonstrates some ability to construct a response variable that is suitable for the given problem.	Demonstrates a limited ability to construct a response variable that is suitable for the given problem.	Does not demonstrate an ability to construct a response variable that is suitable for the given problem.
Hurdle (for a 3)		Constructs a suitable response variable to use in the model.			Constructs a response variable but it is not suitable for use in the model.	
Hurdle (for a 3)		Calculates the correct response variable values and adds them to the pandas table constructed in Q2b.			Does not calculate the correct response variable values and add them to the pandas table constructed in Q2b.	
Differentiator		Justifies the response variable constructed with strong reference to the problem context.	Justifies the response variable constructed with some reference to the problem context.	Describes the response variable constructed with reference to the problem context.	Describes the response variable constructed but does not reference the problem context.	Does not describe the response variable constructed.
Differentiator		Correctly interprets the outcomes of at least two types of appropriate checks on the response variable constructed, in a thorough and insightful way.	Correctly interprets the outcomes of one appropriate check on the response variable constructed in a thorough and insightful way.	Correctly interprets the outcomes of at least one appropriate check on the response variable constructed.	Applies at least one appropriate check on the response variable constructed but does not interpret its outcome or does not interpret its outcome correctly.	Applies at least one check on the response variable constructed but the check is not appropriate and the outcome of the check is not interpreted or is not interpreted correctly.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
2d	5%	Demonstrates a very strong ability to exercise critical thinking to plan how a model will be evaluated.	Demonstrates a good ability to exercise critical thinking to plan how a model will be evaluated.	Demonstrates some ability to exercise critical thinking to plan how a model will be evaluated.	Demonstrates a limited ability to exercise critical thinking to plan how a model will be evaluated.	Does not demonstrate an ability to exercise critical thinking to plan how a model will be evaluated.
Differentiator		Suggests four evaluation metrics to use that are highly suitable for the business problem to be solved with classification.	Describes four suitable evaluation metrics, with strong reference to the business problem to be solved with classification.	Describes four suitable evaluation metrics, with some reference to the business problem to be solved with classification.	Describes four suitable evaluation metrics but does not link this description to the business problem to be solved with classification.	Outlines four evaluation metrics, but these do not seem sensible given the business problem that needs to be solved with classification.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
2e	20%	Demonstrates a very strong ability to construct a classifier.	Demonstrates a good ability to construct a classifier.	Demonstrates some ability to construct a classifier.	Demonstrates a limited ability to construct a classifier.	Does not demonstrate an ability to construct a classifier.
Hurdle (for a 2)		Correctly constructs a neural network to classify acute diagnosis as a binary outcome.				Constructs a neural network but does not classify acute diagnosis as a binary outcome.
Hurdle (for a 2)		Fine-tunes model architecture, hyperparameters, feature selection, regularisation, and optimisation algorithms.				Does not fine-tune all five options.
Differentiator		Improves predictions between the initial and final iterations, using a range of appropriate metrics and sensible inspection of each model to justify subsequent iterations.	Attempts to improve predictions between the initial and final iterations, using a range of appropriate metrics and sensible inspection of each model to justify subsequent iterations.	Takes an iterative approach, but more appropriate metrics could have been used, including an inspection of each model.	Does not take an iterative approach but takes steps to prevent the classifier from over or under-fitting to the training data.	Does not take steps to prevent the classifier from over or under-fitting to the training data.
Differentiator		Constructs a wide range of useful, explainable, and intuitive features from every table in the dataset.	Constructs useful, explainable, and intuitive features from most tables in the dataset.	Constructs features from most tables in the dataset, but some could be more useful, explainable or intuitive.	Constructs few features, and they lack usefulness, explainability or intuitive value.	Uses only existing data columns without further feature engineering or incorrectly constructs features.
Differentiator		The code is structured and very easy to read.	The code is structured and easy to read.	The code is structured and mostly easy to read.	The code lacks structure and is mostly difficult to read.	The code lacks structure and is difficult to read.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
2f	5%	Demonstrates a very strong ability to interpret relevant measures of success for a classifier.	Demonstrates a good ability to interpret relevant measures of success for a classifier.	Demonstrates some ability to interpret relevant measures of success for a classifier.	Demonstrates a limited ability to interpret relevant measures of success for a classifier.	Does not demonstrate an ability to interpret relevant measures of success for a classifier.
Hurdle (for a 1)		Calculates how good the final classification model's predictions are using the evaluation metrics suggested in Question 2d.				
Differentiator		Compares the final classification model's outcomes to those from at least two suitable benchmark models and intuitively explains the validity of each benchmark.	Compares the final classification model's outcomes to those from at least two suitable benchmark models.	Compares the final classification model's outcomes to one suitable benchmark model.	Compares the final classification model's outcomes to a benchmark model, but the benchmark model is not wholly suitable.	Compares the final classification model's outcomes to a benchmark model, but the benchmark model is not suitable.
Differentiator		Correctly interprets the evaluation metrics calculated, with strong reference to the business problem to be solved and with language that is highly suitable for the management team at Betahelp.	Correctly interprets the evaluation metrics calculated, with strong reference to the business problem to be solved and with language that is mostly suitable for the management team at Betahelp.	Correctly interprets the evaluation metrics calculated, with some reference to the business problem to be solved and with language that is mostly suitable for the management team at Betahelp.	Correctly interprets the evaluation metrics calculated, but with no reference to the business problem to be solved or using language that is mostly not suitable for the management team at Betahelp.	Does not correctly interpret the evaluation metrics calculated. Does not reference the business problem to be solved. Does not use language suitable for the management team at Betahelp.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
2g	5%	Demonstrates a very strong ability to interpret the behaviours of a classifier.	Demonstrates a good ability to interpret the behaviours of a classifier.	Demonstrates some ability to interpret the behaviours of a classifier.	Demonstrates a limited ability to interpret the behaviours of a classifier.	Does not demonstrate an ability to interpret the behaviours of a classifier.
Hurdle (for a 2)		Calculates feature importance correctly.				Does not calculate feature importance correctly.
Hurdle (for a 3)		Calculates partial dependence and SHAP feature explanations correctly.			Does not calculate partial dependence and SHAP feature explanations correctly.	
Hurdle (for a 4)		Evaluates potential unfair socioeconomic, racial, or gender biases in the model.			Does not assess fairness of bias.	
Differentiator		Correctly interprets the model explanation outputs, prioritising which details to calculate and share, with strong reference to the business problem to be solved, and with language highly suitable for the management team at Betahelp.	Correctly interprets the model explanation outputs, prioritising which details to calculate and share, with strong reference to the business problem to be solved, and with language suitable for the management team at Betahelp.	Correctly interprets the model explanation outputs, prioritising which details to calculate and share, with some reference to the business problem to be solved, and with language mostly suitable for the management team at Betahelp.	Correctly interprets the model explanation outputs, prioritising which details to calculate and share, with no reference to the business problem to be solved, or with language not suitable for the management team at Betahelp.	Does not correctly interpret model explanation outputs. Does not reference the business problem to be solved. Does not use language suitable for the management team at Betahelp.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
3a	5%	Demonstrates a very strong ability to construct and apply an LLM prompt.	Demonstrates a good ability to construct and apply an LLM prompt.	Demonstrates some ability to construct and apply an LLM prompt.	Demonstrates a limited ability to construct and apply an LLM prompt.	Does not demonstrate an ability to construct and apply an LLM prompt.
Hurdle (for a 2)		Correctly constructs a Python prototype that outputs a reusable LLM prompt which, when applied to each of the five random lab reports, produces structured, analytical commentary similar to the existing human-written pathology reports.				Constructs an LLM prompt but it is not reusable or does not result in suitable output.
Differentiator		The prompt contains all five components of a strong prompt.	The prompt contains four components of a strong prompt.	The prompt contains three components of a strong prompt.	The prompt contains two components of a strong prompt.	The prompt contains one component of a strong prompt.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
3b	5%	Demonstrates a very strong ability to apply generative AI to a business problem.	Demonstrates a good ability to apply generative AI to a business problem.	Demonstrates some ability to apply generative AI to a business problem.	Demonstrates a limited ability to apply generative AI to a business problem.	Does not demonstrate an ability to apply generative AI to a business problem.
Differentiator		Suggests actionable insights from how the LLM outputs compare to the human pathologists' advice.	Describes how the LLM outputs compare to the human pathologists' advice.	Outlines how the LLM outputs compare to the human pathologists' advice.	Outlines the LLM outputs, but these are not compared to the human pathologists' advice.	States the LLM outputs.
Differentiator		Provides a detailed analysis and assessment of the strengths and weaknesses of using the GenAI commentaries on lab results, with strong reference to the needs of all healthcare stakeholders.	Provides an analysis and assessment of the strengths and weaknesses of using the GenAI commentaries on lab results, with some reference to the needs of all healthcare stakeholders.	Outlines many of the strengths and weaknesses of using the GenAI commentaries on lab results, with some reference to the needs of most healthcare stakeholders.	Outlines only one or two strengths or weaknesses of using the GenAI commentaries on lab results, or does not reference the needs of healthcare stakeholders.	Outlines only one or two strengths or weaknesses of using the GenAI commentaries on lab results and does not reference the needs of healthcare stakeholders.
Differentiator		Communicates in a way that is highly suitable for the management team of Betahelf.	Communicates in a way that is suitable for the management team of Betahelf.	Communicates in a way that is mostly suitable for the management team of Betahelf.	Communicates in a way that is mostly not suitable for management team of Betahelf.	Communicates in a way that is not suitable for the management team of Betahelf.



Question	Weight	Significantly above pass level	Above pass level	Pass level	Below pass level	Significantly below pass level
		5	4	3	2	1
4	10%	Demonstrates a very strong ability to communicate model outcomes to a non-technical audience.	Demonstrates a good ability to communicate model outcomes to a non-technical audience.	Demonstrates some ability to communicate model outcomes to a non-technical audience.	Demonstrates a limited ability to communicate model outcomes to a non-technical audience.	Does not demonstrate an ability to communicate model outcomes to a non-technical audience.
Hurdle (for a 1)		Prepares a presentation to summarise findings in a video format.				
Hurdle (for a 4)		Does not appear to be reading word for word from a script.			Appears to be reading word for word from a script.	
Differentiator		The video has a clear start, middle, and end. There are clear transitions between all sections.	The video has a clear start, middle, and end. There are clear transitions between most sections.	The video has some structure. There are transitions between some sections.	The video has some structure. There are no transitions between sections.	The video lacks structure.
Differentiator		Communicates clearly and concisely in a way that is highly suitable for the management team at Betahelp.	Communicates clearly and concisely in a way that is suitable for the management team at Betahelp.	Communicates clearly and concisely in a way that is mostly suitable for the management team at Betahelp.	Communicates in a way that is mostly not suitable for the management team at Betahelp.	Communicates in a way that is not suitable for the management team at Betahelp.
Differentiator		Provides well considered answers to all three questions asked by the management team.	Provides answers to all three questions asked by the management team, two are well considered.	Provides answers to all three questions asked by the management team, one is well considered.	Provides answers to all three questions asked by the management team, but none are well considered.	Provides answers to one or two of the questions asked by the management team.
Total	100%					



Data Science Applications

Assignment Semester 2 2025 – Marking Rubric

Students are advised that a mark of zero (0) will be allocated for any question (or sub-question) where there is no attempt made or the marker finds the attempt is completely unsatisfactory.