2026 ASEE ANNUAL CONFERENCE June 21–24, 2026 | Charlotte Convention Center, NC
Host Sponsor UNC Charlotte
Paper ID #53019

Evaluation of the Machine Learning and Artificial Intelligence Conceptual Framework (ML/AI-CF) in Engineering

Dr. Lauren Singelmann, Minnesota State University, Mankato

Lauren Singelmann is an assistant professor at Minnesota State University, Mankato with Iron Range Engineering. She earned her PhD from North Dakota State University in Electrical and Computer Engineering and STEM Education.

Evaluation of the Machine Learning and Artificial Intelligence Conceptual Framework (ML/AI-CF) in Engineering

Introduction

As machine learning (ML) and artificial intelligence (AI) become increasingly intertwined with engineering work, the need for engineers who are not only strong AI users, but also strong AI contributors is rapidly growing. Therefore, it is important to help engineering students build ML/AI literacy and expert-like thinking so they can contribute to projects and problems involving these emerging technologies. One key step towards developing expert-like thinking is using a conceptual framework to organize and retrieve technical information.

One conceptual framework for teaching ML/AI consists of four categories that can be used to categorize concepts and be applied to ML/AI applications. The categories include *data, task, algorithm,* and *assessment.* Preliminary work showed how the framework could be used to teach ML and AI concepts in engineering courses [1], and this paper expands on that existing work to demonstrate and evaluate how the framework is applied by students to actual ML/AI applications.

After students were introduced to the ML/AI Conceptual Framework (ML/AI-CF) during a 1-credit ML for Engineering course, they were asked to apply the framework to a journal paper about an engineering application of ML/AI. Students were given a general application area (e.g. self-driving cars, power systems, manufacturing), and they were allowed to choose the journal paper they reviewed. They then wrote a group report where they were asked to describe each of the journal papers using the ML/AI-CF. The instructor of the course also separately read and applied the framework to each journal paper. The instructor's and students' answers were then compared to assess the reliability of the framework, improve the framework based on student reasoning, and inform teaching practices in ML/AI.

Background

<u>Existing ML/AI Frameworks</u>

A variety of frameworks exist about ML/AI education and literacy. [2] defines AI literacy as understanding the techniques and concepts behind AI products and presents an AI Curriculum Framework with four stages for developing this literacy: building awareness, experimenting and familiarizing, fostering core AI topics, and becoming fluent. [3] created the Machine Learning Education Framework which consists of three categories: knowledge, skills, and attitudes, as well as sub-categories to build ML contributors. [4] reviewed 30 peer-reviewed articles to create and propose a four-part AI literacy TPACK Framework that includes knowing & understanding, using & applying, evaluating & creating, and AI ethics. [5] then extends this model by adding a fifth part of the framework: enabling AI. They also add a variety of student learning outcomes for each of the five components of the framework. [6] also reviewed conference papers, journal articles, and grey literature to develop an AI Literacy Framework that consists of 17 competencies (e.g. strengths and weaknesses of AI, ML steps, ethics, etc.) and 15 design considerations (e.g. acknowledging preconceptions, social interaction, reducing barrier to entry) for developing AI literacy. [7] interviewed professionals working in the intersection of AI and other domains to identify 52 core AI competencies for those outside of computer science. These competencies generally aligned with the ones presented in [6] but also listed competencies related to applying AI to specific disciplines outside of computer science.

These frameworks provide extremely valuable structure for ensuring students develop and interact with core competencies related to ML and AI. However, less published work has focused on identifying and structuring course concepts for developing ML/AI proficiency. Many of the existing frameworks mention the importance of "fostering core AI topics" [2] and "knowledge of ML methods" [3] as key AI literacy components, but these topics and methods are not explicitly or exhaustively named. In the interviews conducted in [7], the most popular competency discussed was "explain how ML techniques and their main algorithms work", but there was lack of agreement among participants about which main algorithms should be included. Therefore, maybe the question is not about *which* algorithms should be discussed, but rather *how* those algorithms should be discussed and framed.

The Motivation for a Conceptual Framework

In the foundational report *How People Learn* [8], one of the key findings was that competency is developed through "(a) hav[ing] a deep foundation of factual knowledge, (b) understand[ing] facts and ideas in the context of a conceptual framework, and (c) organiz[ing] knowledge in ways that facilitate retrieval and application." For ML/AI, learning must go beyond memorizing terms like classification, clustering, and neural networks; it should also involve understanding how these concepts relate to one another and forming a coherent structure that connects them. We define a conceptual framework using the framing provided in [8]; a conceptual framework is a way to organize facts and ideas that supports retrieval and application.

As defined in [2], AI literacy is the ability to understand the techniques and processes behind AI. By framing ML/AI topics in the context of a conceptual framework, learners are able to retrieve and apply that information to be able to make meaning out of the concepts.

The Machine Learning and AI Conceptual Framework (ML/AI-CF)

The Machine Learning and AI Conceptual Framework (ML/AI-CF) [1] is a framework that breaks machine learning concepts into four categories: *data, task, algorithm,* and *assessment*. These framework categories were identified by looking at a variety of machine learning textbooks, identifying common topics, and inductively coding the topics into groups [1]. *Data* describes the data being used and how it is being used and includes topics such as types of data (numerical, categorical, text) and other data vocabulary such as targets, classes, and features. *Task* describes the goal of the machine learning solution and includes topics such as supervised and unsupervised learning, classification, clustering, regression, etc. *Algorithm* describes the computational approach to completing the task and includes topics such as neural networks, support vector machines, DBSCAN clustering, etc. *Assessment* describes how the solution will be evaluated and includes topics like confusion matrices, accuracy, mean absolute error, etc. [1]

Many of the existing frameworks presented earlier are meant to be administrator- or educator-facing to create programs, curriculum, courses, and activities that meet the outcomes presented in these frameworks. However, the ML/AI-CF has a different audience and purpose; it is meant to be a student-facing tool to help students structure and discuss concepts. For example, when students learn about clustering as a *task*, they also have framing to understand that that specific task also has specific *data* needs, *algorithm* options, and *assessment* modes that may be different from classification *tasks*.

Preliminary published work presented the framework, how it can be implemented in a course, and student feedback about how the framework was useful when working in engineering co-ops [1]. In addition, it has been shown to be a helpful tool when helping assess and understand the strengths and limitations of large language models [9]. This paper extends this existing work by validating and improving the framework through an analysis of how students applied the framework to ML/AI applications in a variety of industries.

Methods

<u>*Course Context*</u>

The student data from this study comes from a 1-credit Machine Learning and Artificial Intelligence course for students in an Integrated Engineering program. The dataset comes from two iterations of the course: one in Spring 2024 and one in Spring 2025. The course runs for 7 weeks and the general topics are listed in Table 1. The ML/AI-CF is framed as four "fundamental principles" or "FPs", so each week discussed a different common task in ML/AI and the relevant considerations for each of the FPs when performing that task. For example, in Week 2 when students learned about the *task* of classification, they also learned about *algorithms* such as support vector machines and *assessment* metrics such as confusion matrices.

Assignments included watching videos covering the topics for the week, creating a concept map throughout the course, in-class activities, coding exercises, and working on a deep learning activity (DLA) – a final paper completed in groups.

Table 1
General Overview of Topics Covered in Course

Week	Topic
1	The 4 Fundamental Principles (FPs) of ML
2	Applying the 4 FPs of ML to Classification
3	Applying the 4 FPs of ML to Regression
4	Applying the 4 FPs of ML to Clustering
5	Applying the 4 FPs of ML to Association Analysis
6	Review and DLA Focus Time
7	Exams and Turning in Final Deliverables

Description of the DLA Assignment

Throughout the course, students worked in groups on a deep learning activity (DLA), which was a paper written in groups about how ML/AI is being applied in a specific industry. Students filled out a survey that asked them to choose an ML/AI application area of interest. Application areas included biomedical engineering, self-driving cars, robotics, manufacturing processes, smart electrical grids, environmental engineering, 3D printing, machine fault prediction, reliability and safety, and structural engineering. The instructor then grouped students into groups of 3-4 around topics of interest for each student. The DLA was split into multiple checkpoint assignments as listed in Table 2.

Table 2
DLA Checkpoint Assignments

Checkpoint	Due During Week	Description
1	2	Each group was assigned a review paper to read about ML/AI in their topic area. The instructor chose the review paper for each topic, and students were instructed to co-annotate the paper with their group in a shared document. Each student was instructed to add ~10 comments to the paper.
2	4	Each member of the team then chose two journal papers to read and annotate individually. Students were instructed to

		focus on making annotations about how the 4 FPs of ML were applied and any interesting takeaways. Each student was instructed to add ~10 comments to each of their journal papers.
3	6	Groups submitted a first draft of their DLA which included 1) an introduction and background describing their application area and how ML/AI is used, 2) an example case written by each team member that describes how the FPs of ML were applied, 3) limitations and conclusions about the state of ML/AI in that area.
4	7	Groups submitted a final draft of their DLA after receiving peer feedback during class in Week 6.

Dataset

The dataset was created by anonymizing each of the DLAs written by the student groups, extracting each of the journal papers cited in the bibliographies, and using a library database to secure a copy of each of the cited journal papers. The instructor of the course read each of the cited journal papers to determine if the paper should be included in the final dataset. 22 journal papers were part of the original dataset; 7 were excluded because they were review papers and did not describe a specific application of ML/AI, and 2 were excluded because they did not use ML/AI methods. It should be noted that many students were still able to use the framework to summarize these review papers, but these papers were still excluded for this study because there is no identifiable "correct" answer when many studies are being discussed rather than one specific study.

The instructor then applied the ML/AI-CF to the remaining 13 journal papers by documenting how each journal paper described their *data*, *task, algorithm,* and *assessment.* After this was completed, each of the student DLAs was read to see what the students wrote for each of the four categories. The instructor and student results were then compared, and each comparison was categorized using the agreement codes below:

Correct: The student and instructor had answers with similar major keywords and level of detail

Partially Correct: The student included correct information that was relevant to the framework category, but did not provide major details that the instructor identified

Miscategorized: The student included correct information, but it did not relate to the framework category being discussed

Incorrect: The student provided information that was not what was described in the initial journal paper

Not listed: The student did not provide any information for that framework category

Research Questions

Analysis was then conducted in three phases to answer the research questions (RQ) below:

> **RQ1:** What is the frequency and level of disagreement between student and instructor classifications within each framework category?
> **RQ2:** To what extent are these differences shared across students versus concentrated among a small number of individuals?
> **RQ3:** What qualitative features of student reasoning characterize instances of disagreement?

RQ1 and RQ2 were analyzed quantitatively; the agreement codes were tallied and organized by unit of analysis (framework category for RQ1 and student for RQ2). RQ3 was analyzed qualitatively; for each instance of disagreement between student and instructor, the disagreement was described and analyzed to identify what reasoning or interpretation challenges may have caused the disagreement. This was done by looking at the student and instructor responses, as well as the original publications that the students read and reviewed. These results were then used to inform updated definitions and guiding questions for an updated version of the ML/AI-CF.

Results and Findings

RQ1: Agreement within framework categories

The first research question aimed to determine which ML/AI-CF categories had the least and most disagreement between the student and instructor. For each of the four categories and thirteen students, the frequency of each agreement code was tallied. The results are shown in Table 3, and there were two main findings.

Table 3

Number of instances of each agreement code across each of the framework categories

Agreement Code	C	P	M	I	N
Data	11	1	0	0	1
Task	9	3	0	1	0
Algorithm	12	0	0	1	0
Assessment	8	3	0	1	1

C = Correct
P = Partially Correct
M = Misaligned
I = Incorrect
N = Not listed

Finding 1.1: All framework categories were applied correctly by at least 60% of students. The framework categories with the most instances of agreement were *Data* (11/13 students) and *Algorithm* (12/13 students), whereas the framework categories with the fewest instances of agreement were *Task* (9/13 students) and *Assessment* (8/13 students).

Finding 1.2: Most cases of disagreement (7/12 cases) were "partially correct", suggesting that students are generally using the ML/AI-CF correctly, but more attention could be placed on helping students understand what information is most important when identifying *tasks* and *assessment.* No items fell into the "miscategorized" agreement code, meaning no students gave a response that fell into a completely different category. This suggests that the framework categories are sufficiently different, and that these differences were made clear to students. However, in 3/12 cases of disagreement, students mentioned details that were not present in the journal paper, suggesting that students did not actually reference the original journal paper. These cases of disagreement will be further analyzed and discussed in RQ3.

RQ2: Agreement distribution across students

The second research question aimed to determine if the instances of disagreement were generally shared across students versus concentrated among a small number of individuals. For each student, the agreement codes for each category are shown in Table 4, and there were two main findings.

Table 4

Agreement codes for individual students

Student	1	2	3	4	5	6	7	8	9	10	11	12	13
Data	C	N	C	P	C	C	C	C	C	C	C	C	C
Task	I	P	C	C	C	C	C	P	C	C	C	C	P
Algorithm	C	I	C	C	C	C	C	C	C	C	C	C	C
Assessment	N	I	P	C	C	C	C	P	P	C	C	C	C

C = Correct
P = Partially Correct
M = Misaligned
I = Incorrect
N = Not listed

Finding 2.1: Instances of students providing "incorrect" information about the journal paper occurred only within a small number of students. This suggests only a few students did not refer to their original journal paper when applying the framework.

Finding 2.2: Instances of students providing "partially correct" information were more widespread across students. This suggests that small framework misinterpretations were more widespread, but that most students were still applying the framework correctly for most of the categories.

RQ3: Types of disagreement
The third research question aimed to qualitatively explore specific instances of disagreement. When comparing the student and instructor responses and referring back to the original journal papers that the students read, five main categories of disagreement emerged. These disagreement types are listed in Table 5, and there were two main takeaways.

Table 5
Disagreement types

Disagreement Type	**Number of Instances**	**Example Student Response**	**Example Instructor Response**
Providing incorrect information	3	Assessment: "Precision, recall, and F1 score"	Assessment: "Correlation coefficient (R) and mean squared error"
Missing information	2	Assessment: "RMSE compared to baseline"	N/A
Confounding *task* with the outcome/ goal	3	Task: "safely navigate a self-driving car down the road while driving in optimal conditions as well as suboptimal conditions"	Task: "regression to identify the center of a lane"
Confounding *assessment* with the outcome/goal	3	Assessment: "The prototype alerted the driver every time a dangerous situation approached the driver's vehicle."	Assessment: "'Prediction accuracy' was demonstrated by plotting ROC, as well as precision over time during dangerous driving conditions."
Confounding *data* with the outcome/goal	1	Data: "Library of defect features"	Data: "Images of batteries with labeled defects. Images were converted into specific parameters (e.g. gray max, length ratios, etc.)"

Finding 3.1: When analyzing instances of "partially correct" responses, most were tied back to students describing the overall outcome or goal of the application when applying the framework rather than the approach the authors took. For example, rather than describing that a self-driving car algorithm was set up to be a regression *task*, one student wrote that the main *task* was to

safely navigate the self-driving car. Although this was the intended goal of the study, it did not appropriately describe how the ML/AI model was framing the problem. This suggests that the definitions of *task, assessment*, and *data* can be refined to elicit responses that better discuss the ML/AI solution rather than the overall outcome of the study.

Finding 3.2: When analyzing instances of "incorrect" information in student responses, the specific keywords the students mentioned were not found anywhere in the original journal paper. Rather than the student simply making things up on their own, we hypothesized that the student may have used a generative AI tool to summarize the journal paper for them. To further explore this hypothesis, we gave ChatGPT the titles of a few of the journal papers and asked it to apply the framework. When looking at the responses, we saw that ChatGPT often responded using common, relatively non-specific methods (e.g. classification using a variety of algorithms including neural networks and decision trees with assessment metrics including accuracy and F1 score). In general, generative AI tools only have access to abstract-level information about most journal publications, so the tool is making a prediction of the methods used by using the information it has available.

Discussion

Overall, the ML/AI-CF was shown to be a helpful tool for students who are new to learning about ML/AI concepts. The framework allows students to organize these concepts in a way that allows for successful retrieval and application. The results of the paper suggest that the framework categories are helpful, but that the corresponding definitions could be refined to combat some of the disagreements between the students and instructors. The discussion provides updates to the ML/AI-CF that respond to student misconceptions, as well as some specific recommendations for instructors who wish to implement this specific assignment in a course.

Recommended updates to the ML/AI-CF

Four changes were made to address some of the disagreements that appeared to be from reasoning or interpretation challenges.

1) Refining the *task* category - The *task* category was intended to ask students which specific machine learning method (e.g. classification, regression, clustering) was being used in an application. However, some students interpreted this to be a more general goal/outcome. To better define this category, the refined definition and guiding questions was restated to focus on "problem formulation" rather than a "goal".
2) Refining the *assessment* category - The *assessment* category was intended to ask students how performance was measured or validated in an application. However, some students instead described the outcome of the study rather than how the researchers/developers knew that outcome was being met. To better define this category, a variety of assessment

categories were listed in the definition, including evaluation criteria, benchmarking, and other validation strategies. The focus on the word “criteria” is meant to bring the framing back to measurement and evaluation rather than describing the overall outcome.

3) Refining the *data* category - The previous definition of *data* in the ML/AI-CF used the very general description of “data and processes used”. This led to one student specifically describing a library of detect features that was the output of a study, rather than describing the images that were used to train and evaluate the model. An updated definition specifically calls out that *data* should describe specifically what was used for training and/or evaluation purposes.
4) Adding an *outcome* box - Finally, to help differentiate between the *outcome* of a study and the *task* or *assessment*, another version of the graphic was created that specifically calls out *outcome*. This version of the figure can be used when applying the ML/AI-CF to a specific application because it shows that the outcome is different than each of the four categories, but is ultimately a result of the framework categories being aligned and working together. The final outcome of an ML/AI application communicates the value and impact of its proposed solution, and it is reasonable that students would want to share that in their final DLA. By calling out *outcome* specifically, students are encouraged to describe that value and impact while differentiating between that *outcome* and the other categories of the ML/AI-CF.

These recommended changes led to updated wording for 1) the definitions of each framework category that were originally published in [1], and 2) guiding questions that the instructor used in course videos to help students apply the framework to a specific application. Table 6 shows how the recommendations changed the definitions of each framework category, and Table 7 shows how the recommendations changed the guiding questions for each framework category. Images that show the final definitions and guiding questions can be found in the Appendix.

Table 6

Recommended definition changes to the ML/AI-CF

Framework Category	Previous Definition from [1]	Refined Definition
Data	The characteristics of data and the processes used to utilize that data for machine learning processes	The characteristics **and source of the information used to train and evaluate the machine learning model**
Task	The goal of the machine learning solution	The **computational formulation** of the machine learning solution

Algorithm	The mathematical, statistical, and/or computational approach to completing the machine learning task	The mathematical, statistical, and/or computational approach to completing the machine learning task
Assessment	Metrics for quantitatively assessing the ability of the machine learning model to complete the task	**Evaluation criteria, benchmarks, and/or validation strategies** used to assess the ability of the machine learning model to complete the task

Table 7
Recommended "guiding question" changes to the ML/AI-CF

Framework Category	Previous Guiding Questions from Course Materials	Refined Guiding Questions
Data	What data do we care about, and what are the characteristics of this data?	**What kind of data will be used to train and test the model?**
Task	What goal are we trying to accomplish?	**How is the problem framed for the model?**
Algorithm	How are we going to computationally accomplish this goal?	How are we going to computationally accomplish this **task**?
Assessment	How are we going to show that we were successful?	**What criteria will be used to measure success?**

<u>*Future recommendations for implementing this assignment*</u>

For any instructor looking to implement a similar assignment in their own course, this work offers some suggestions for improving student learning during the assignment.

1) Encourage students to explicitly name each of the ML/AI-CF elements in their final DLA *and* the outcome. By naming each of the four elements *and* the outcome, students are encouraged to decouple the *task* and *assessment* elements from the overall goal/purpose of the application.
2) Provide students with both the definitions and guiding questions for the ML/AI-CF in the assignment documentation to help guide their reading of the journal paper. This will help

remind students of the definitions of each of the framework elements (rather than relying on their own interpretation of the framework category names).

3) Include two checkpoint assignments if possible: one where students submit the journal paper they want to review, and one where students submit an annotated copy of that journal paper with highlighted quotes showing each of the framework categories. This will help ensure that students are choosing journal papers with specific applications that use ML/AI methods and using the primary text to identify each of the framework elements. In addition, it offers an opportunity to provide formative feedback before groups submit their final DLAs.
4) Remind students that generative AI tools often cannot reliably give specific information about journal papers. These tools often only have access to abstract-level information about the journal paper, meaning that the tool is often just making a prediction about the methods based on the information available. Depending on the information available, these predictions may or may not be correct. This can be reinforced by showing examples in class of the tool incorrectly providing information about a study or having students test the use of the tool themselves.

Conclusion

As defined in [2], the goal of AI literacy is to be able to understand the techniques and concepts behind AI applications. This work demonstrates that the ML/AI-CF can help students organize ML/AI concepts in a way that allows them to build their AI literacy and describe a variety of ML/AI applications. The work also provided valuable insights about how the ML/AI-CF categories can be better defined and framed for students to help build expert-like thinking about ML/AI. The result is a refined version of the ML/AI-CF that can be used when teaching students about a wide variety of ML/AI concepts and applications. A variety of images that show the refined ML/AI-CF with definitions and guiding questions can be found in the Appendix.

References

[1] L. Singelmann and J. Covarrubias, “A Framework For Teaching Machine Learning For Engineers”, in *SEFI 2023 Annual Conference Proceedings*, 2023.

[2] M. Kandlhofer, G. Steinbauer, S. Hirschmugl-Gaisch, and P. Huber, “Artificial intelligence and computer science in education: From kindergarten to university”, in *2016 IEEE Frontiers in Education Conference (FIE)*, 2016, pp. 1–9.

[3] N. Lao, “Reorienting machine learning education towards tinkerers and ML-engaged citizens”, *Massachusetts Institute of Technology Cambridge*, 2020.

[4] D. T. K. Ng, J. K. L. Leung, S. K. W. Chu, and M. S. Qiao, “Conceptualizing AI literacy: An exploratory review”, *Computers and Education: Artificial Intelligence*, vol. 2, p. 100041, 2021.

[5] J. Southworth *et al.*, “Developing a model for AI Across the curriculum: Transforming the higher education landscape via innovation in AI literacy”, *Computers and Education: Artificial Intelligence*, vol. 4, p. 100127, 2023.

[6] D. Long and B. Magerko, “What is AI literacy? Competencies and design considerations”, in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–16.

[7] K. Tenório and R. Romeike, “AI Competencies for non-computer science students in undergraduate education: Towards a competency framework”, in *Proceedings of the 23rd Koli Calling International Conference on Computing Education Research*, 2023, pp. 1–12.

[8] J. D. Bransford, A. L. Brown, R. R. Cocking, and Others, *How people learn*, vol. 11. Washington, DC: National Academy Press, 2000.

[9] L. Singelmann, J. Elliott, Y. Wang, and J. J. Swanson, “Student, Faculty, and Industry Perceptions about Generative AI’s Capabilities to Perform Specific Tasks”, in *2025 ASEE Annual Conference & Exposition*, 2025.

Appendix

Figure A.1

The Machine Learning and Artificial Intelligence Conceptual Framework (ML/AI-CF) categories and definitions

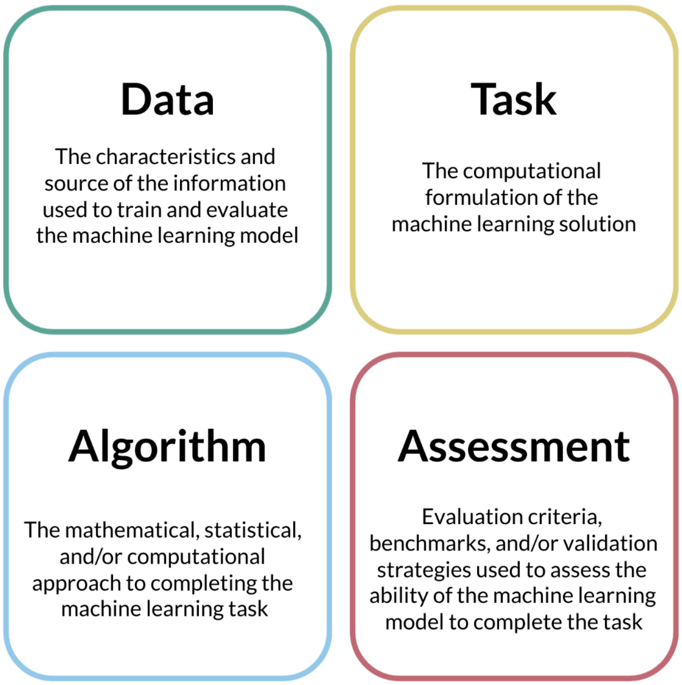

Figure A.2

The Machine Learning and Artificial Intelligence Conceptual Framework (ML/AI-CF) categories and guiding questions. These guiding questions can help students apply the ML/AI-CF to various applications

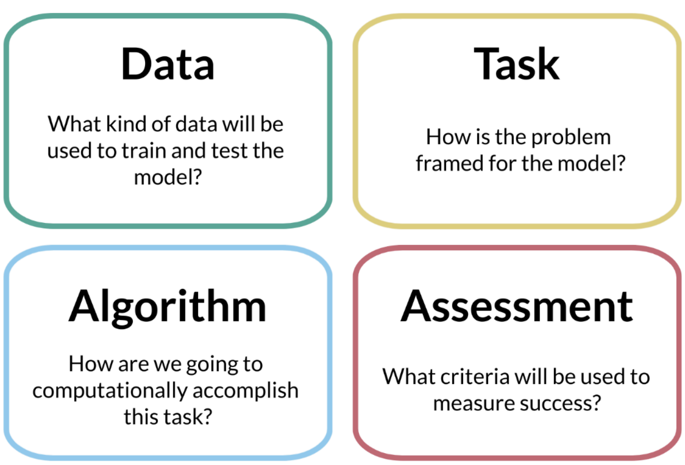

Figure A.3

The Machine Learning and Artificial Intelligence Conceptual Framework (ML/AI-CF) categories and definitions with the addition of the "Outcome" box. Although the "Outcome" is not part of the AI/ML-CF, the components of the framework work together to produce an outcome. When applying the ML/AI-CF, this image can help students differentiate between the way the model is structured (Data/Task/Algorithm/Assessment) and its overall goal/outcome.