

Version Beta Revision Report – Artificial Intelligence Knowledge Area

Feedback was solicited from specific reviewers on the beta version; the notes below indicate how this feedback was incorporated into subsequent versions of the AI KA. This log focuses on the most major comments and issues; other minor edits were taken directly, and other edits were made independently by the AI subcommittee.

Feedback comment:

Reviewers suggested various minor edits for clarity and additional examples (e.g., generative AI), especially in the preamble, but throughout the document.

How incorporated:

Edits were combined with edits made by the AI subcommittee to create cohesive text.

Why not incorporated:

Feedback comment:

AI/ML: NB would not be my first choice for an algorithm for CS core. I think linear regression or k-NN is more suitable for the general CS student.

How incorporated:

Agreed. Changed NB to linear regression.

Why not incorporated:

Feedback comment:

AI/ML: Linear regression and other topics are listed under electives. These seem to be in the KA core already. Should we take them out here?

How incorporated:

Why not incorporated:

These will remain as subtopics of supervised learning under electives since 1) they may not have been covered under core and 2) there are certainly additional details that instructors may wish to cover on these topics.

Feedback comment:

AI/ML: Learning Outcomes: Various edits for ordering, deduplication, and content were suggested.

How incorporated:

These edits were incorporated into the document and merged with edits made by the AI subcommittee.

Why not incorporated:

Feedback comment:

Add the following Knowledge Units: Deep Learning, Cognitive Systems, AI Ethics and Societal Impact; Human-Centered AI, Conversational AI, Recommendation Systems.. Include societal and ethical issues in all knowledge units.

How incorporated:

We've added cognitive systems as an explicit part of AI-Agents.

Why not incorporated:

Most of these already fall under existing knowledge units. We've added cognitive systems as an explicit part of AI-Agents. Ethics and applications is suggested to be incorporated throughout all topics. Wherever a KU has specific ethical aspects (e.g., data privacy and intellectual property in those KUs involving data), ethics is listed explicitly with suggested topics relating to that KU. Other topics without explicit ethics topics would instead depend on the ethics topics covered in AI-SEP.

Feedback comment:

Reduce the number of units in search and add units in Ethical AI

How incorporated:

Why not incorporated:

The subcommittee feels that the number of hours in search is necessary to cover the specified topics. Ethics is part of AI-SEP.

Feedback comment:

AI/Fundamental Issues: Learning outcomes need to be expanded so there is a learning outcome for each topic. For example, there should be a learning outcome for the ethical and philosophical topics. General Comment: the lists of Learning Outcomes sound more like assessment tasks rather than a statement of what the student has learned.

How incorporated:

Why not incorporated:

The learning outcomes may cut across topics, helping to synthesize the various topics together

Feedback comment:

AI/Search: There are more learning outcomes here than there are topics. This areas has been taught for the longest period of time. There is a need to be more strategic instead of teaching everything here that we have ever taught. We can now teach search in terms of the applications that are using search in our everyday use of AI rather than teaching search in terms of games.

How incorporated:

Why not incorporated:

The learning outcomes also cover the subtopics, which is why there are so many.

Agreed that search can be taught in conjunctions with applications, but that is the instructor's choice. "Two-player adversarial games" implies games in its most general form, not specifically board or video games. This could certainly be taught using economic games as the application.

Feedback comment:

AI/ML: The definition of types of ML tasks should be expanded, at least to include neural networks in unsupervised and supervised learning

How incorporated:

Why not incorporated:

There are certainly other forms of ML tasks, but these are the main ones that are covered across most or all ML courses. Neural nets are an approach, not a fundamental ML problem formulation like supervised, semisupervised, etc. learning.

Feedback comment:

AI/ML: “no free lunch” This phrase is too informal compared to other topics

How incorporated:**Why not incorporated:**

That name is correct – it is called the “No Free Lunch” theorem

Feedback comment:

AI/ML - Electives: It is not clear how these topics expand on the CS core topics. For example, this includes an overview of reinforcement learning, but RL is a topic in the core.

How incorporated:**Why not incorporated:**

The specific way to expand on the CS or KA core topics would be up to the instructor; the RL under the KA core specifies it should be an overview, with the elective topics delving much deeper into specific aspects, per the subtopics.

Feedback comment:

Learning Outcomes should be mapped to topics

How incorporated:**Why not incorporated:**

The template for the AI KA follows that of the general CS 2023 format, which does not include this mapping.

Feedback comment:

AI-SEP: The topics should be lists of application areas: medicine, health, sustainability, social media, etc. and then the subtopics should be specific to each application area, and include a list of the major ethical issues in that area, human-centered AI should feature in each application area in a similar way to ethical issues: specific human-centered issues for each application area.

How incorporated:**Why not incorporated:**

Many different applications are listed, but there are certainly others that could be covered. The curriculum is intended to give the most flexibility possible to the instructor, and so the topics and issues are framed in general, with the responsibility placed on the instructor to identify the specific aspects of these issues for their chosen applications.

Feedback comment:

AI/Robotics: Add a topic on human-robot Interaction

How incorporated:

Done.

Why not incorporated:

Feedback comment:

increase the emphasis on machine learning and specifically deep learning since they are dominating the field.

How incorporated:

We have expanded coverage of both of those topics, with corresponding hours increases.

Why not incorporated:

Feedback comment:

AI education should focus on those aspects of AI that are distinguishable from knowledge areas in Computer Science Education. My review recommends removing areas such as search from an AI curriculum. Search is a class of algorithms that is common to CS and AI education. Assuming that selected core CS courses are part of an AI curriculum, the Knowledge Areas for AI should include topics such as search techniques in the context of their use in specific Knowledge Areas. Search should be considered a pre-requisite knowledge similar to programming and mathematics.

How incorporated:**Why not incorporated:**

Search is foundational to AI, and the basis of many modern methods. Those aspects which are studied commonly in other KAs, e.g., uninformed search, are cross-listed.

Feedback comment:

I recommend the following table of AI Knowledge Areas in which highlighted areas are in the core and other areas to be selected as electives to reach a total expected number of units for a degree in AI.

AI Knowledge Areas Units

AI Foundations 2

Symbolic AI Systems: Knowledge Representation and Reasoning 2

Overview of Machine Learning 4

Deep Learning 4

Applications and Societal Impact 2

Ethical Implications of AI 2

Human-Centered AI 2

Reinforcement Learning

Integrating Symbolic and Deep Learning Systems

Advanced Deep Learning

Probabilistic Representation and Reasoning

Planning and Design

Recommender Systems

Agents

Natural Language Processing

Robotics

Computer Vision

AI in Games

Evolutionary Systems

Conversational AI

History of AI

Use-inspired AI electives: Health, Climate Change, Energy Sustainability, Social Media, Creative AI, Games Design, Generative Writing, and others

How incorporated:

Why not incorporated:

Most of these new suggested KUs are already included under other KUs, and separate out topics in a way that does not follow the CS 2023 format (e.g., splitting ML into ML, deep learning, advanced deep learning). Importantly, the ethical issues are incorporated with Applications under AI-SEP.
