

Gamma Version Revision Report – Artificial Intelligence Knowledge Area

Feedback on the gamma version was solicited from the AAI and EAI communities via the general membership email lists in October 2023. We discuss below the feedback received and how those changes were incorporated into the delta version. 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:

This looks very archaic and is not appropriate for training today's AI students. Just about everything is deep learning nowadays, and there is a huge demand for jobs in this space. A student who followed this curriculum could not get a job in AI nor could they get into a PhD program. While I believe one course covering older AI techniques, search, etc. should be offered, everything in computer vision and NLP and even much of robotics, depends on having a strong understanding of deep learning. Because deep learning is such an empirical science, a firm foundation in experimental design, hypothesis testing, and statistics is essential for modern machine learning and deep learning work. Those topics should not be in data science, but instead should be a core course. As AI is now used in many production systems, I'd also argue content about MLOps, detecting bias in models, data-centric AI, etc., is critical. The document seems like it is from 20 years ago and it is very non-aligned with what students need to know today, having been a professor for a decade and worked in AI at companies for 5 years.

How incorporated:

Why not incorporated:

The CS2023 CS and KA cores represent minimal sets of topics, and indeed it does contain the “older AI techniques, search, etc.” you mention while also including required coverage of working with data, ML evaluation, and deep learning. Instructors are welcome to expand coverage to other topics, if they can support it. To cover the in-depth knowledge that a student would need to pursue a career or graduate study in AI, it is expected that they would take additional elective courses beyond the small CS and KA core, which would go much more in-depth into deep learning, etc.

Feedback comment:

Depth-first search and breadth first search are usually covered in algorithms classes, nowadays, so I don't know if they make much sense for AI courses.

How incorporated:

Why not incorporated:

Search is a foundational idea for AI, including how modern deep learning algorithms work. DFS and BFS are cross-listed topics between AI and Algorithms, and the overlapped hours (3 on uninformed search) are counted under Algorithms, not AI hours, as noted in the footnote under the core hours table.

Feedback comment:

This is very good, but needs a bit more emphasis on the ethical issues because they're so important.

How incorporated:

Although study of ethical issues are recommended to be incorporated throughout and tied with applications and case studies, we agree that additional time would be useful for it, and so have increased both the CS and KA core hours by 1 each (from 2 hrs to 3 hrs each) in the Applications and Societal Impact to give additional depth to generative models (per another comment) and ethics.

Why not incorporated:

Feedback comment:

AI-Introduction -> KA Core #9: lumping "ethics" at the end of a 1-line list suggests that it's a relatively minor issue. It would be much better to add a sub-bullet for Ethical behaviour after #9.

How incorporated:

Why not incorporated:

Ethics is definitely not a minor issue, and neither are the economic and societal impacts of AI. But this is just the introduction. There is only one additional KA hour allotted to cover all of these topics, and so while students should be introduced to them, proper in-depth study of these topics would fall under the other KUs.

Feedback comment:

AI-Introduction: change "4. Enumerate the characteristics of a specific problem." Under Illustrative Learning Outcomes to "...specific problem, including the ethical issues"

How incorporated:

Why not incorporated:

The problem characteristics defined above in this KU do not include ethics, which is more closely tied to the application than how to characterize a problem. The same abstract AI problem (as introduced here in this KU) could be manifested in a video game or a real-life application, with very different ethical

implications. Consequently, we lump studying ethics (and economic and societal implications) in with studying the application of AI, which is examined under the “Applications...” KU.

Feedback comment:

AI-ML: CS Core #5: Doesn't make it clear that there are (usually) ethical issues around data. Add “f. Ethical issues around data (incl Bias, IP and Privacy)”

How incorporated:

Why not incorporated:

We definitely agree that there are ethical issues around data, which is why these ethical issues are included as AI-ML: CS Core #9 and KA Core #16. The curriculum doesn't cover how these topics should be taught (i.e., their grouping or order), just what should be covered, so instructors are free to cover these topics alongside when they are first talking about data.

Feedback comment:

AI-ML core: add c. Privacy, Intellectual Property & oversimplification

How incorporated:

Added intellectual property, but we're not sure what is meant precisely by oversimplification (perhaps removing context in a prediction or preventing drill-down investigation of the reasons for a prediction?). Instead, we're going to add that as explainability.

Why not incorporated:

Feedback comment:

AI-ML: NONE of the 16 Learning Outcomes listed have any reference to Ethics

How incorporated:

Good point. Added an ethics-focused learning outcome.

Why not incorporated:

Feedback comment:

AI-SEP: It looks as though Ethics is just one of the learning outcomes, and that it irrelevant to (1) and (3).

How incorporated:**Why not incorporated:**

Definitely not. The study of these issues should inform student's answer to (1) and (3), even if ethics is not explicitly mentioned in the prompt. The hope would be that students identify ethical, societal, and economic implications in formulating their AI solution (which is #1) and analyzing the failure modes (#3) even without being explicitly prompted to do so.

Feedback comment:

The entire "Agent" module makes no mention of ethical issues although these are of considerable importance, especially around "Believable agents (synthetic characters, modeling emotions in agents)" and Human Agent interactions.

How incorporated:

Good point. Added.

Why not incorporated:

Feedback comment:

AI-Agents: Illustrative learning outcome #2: should add "and discuss the Ethical implications of these"

How incorporated:

Done.

Why not incorporated:

Feedback comment:

AI-Robotics: There are HUGE ethical issues around robotics, which are ignored here. Suggest adding e. Ethical issues.

How incorporated:

Done, but we put it under a separate heading of applications.

Why not incorporated:

Feedback comment:

General issue, from multiple people: Constraint satisfaction should be a required KA core topic.

Specific feedback:

Person #1 I understand the wish to emphasize on recent trends in AI but it should not be done at the detriment of areas that have been central to the success of AI. Constraint satisfaction has long been (from the mid-90s up to the late 2000s) the main provider of industrial applications of AI (going from nurse rostering to configuration of products, assembly lines, etc). As such, it should belong to Core CS. Its fundamental aspects (backtrack search, constraint propagation, global constraints) are ubiquitous as they allow us, for instance, to understand the strengths and weaknesses of SAT solvers or and the more recent explainability issue.

Person #2 Constraint satisfaction is included under "AI Search: Search" as "Non-Core". The CS Core focuses on graph search, e.g. A* Search. The KA Core contains bi-directional search, beam search, two-player adversarial games, and implementation of A* search. Constraint satisfaction is categorized as "Non-Core", along with topics such as variations on A* (IDA*, SMA*, RBFS). Surely constraint satisfaction at least belongs in the KA Core. Is it really more important to include beam search, for example, than anything about constraint satisfaction? Constraint Satisfaction Problems have their own chapter in Russell and Norvig and in Poole and Mackworth. There is a Constraints journal and a Constraint Programming conference and an Association for Constraint Programming. There are many applications. At the same time the basic idea of constraint satisfaction should be easy to introduce to undergraduates: Sudoku and scheduling university classes are constraint satisfaction problems, and the basic concept of backtrack search is not overly complex. I venture that students who proceed with further AI study or application are likely to benefit more from having been introduced to constraint satisfaction than beam search. Using Science Direct, I found in the Artificial Intelligence journal a single paper with "beam search" in the title versus 53 with "constraint satisfaction" in the title. For the Engineering Applications of Artificial Intelligence journal I found again a single paper with "beam search" in the title versus 12 with "constraint satisfaction" in the title. Over all the publications covered by the Science Direct search I confess to having found a surprising 71 results with "beam search" in the title, but there were 310 with "constraint satisfaction" in the title. And the Constraints journal is a Springer journal, not covered by Science Direct. :-) Ideally, constraint satisfaction would even be in the CS Core. I realize that there is a desire to keep the cores small. But the proposed CS Core here is very narrowly focused on graph search. Backtrack search to satisfy constraints might be substituted for some of the graph search material. It would be good if constraint propagation (consistency processing) could be included as well. If there is not room for both in the CS Core then perhaps one in the CS Core and one in the KA Core. The Appendix: Core Topics and Skill Levels can be modified accordingly. An Illustrative Learning Outcome might be: Model a logic puzzle or Sudoku puzzle as a constraint satisfaction problem. Solve with backtrack search. Determine how much arc consistency can reduce the search space.

How incorporated:

Done. Constraint satisfaction is now a KA core topic under AI-Search. KA core hours increased from 4 to 6 accordingly. Added corresponding learning outcome.

Why not incorporated:

Feedback comment:

There needs to be more coverage of how deep generative models work (especially including fundamental ideas like attention), which is where most of the latest advances are coming from. It should also include the idea of foundational models, how to use them, and aspects on their training from big data.

How incorporated:

Done. This has been incorporated into “Applications and Societal Impact”, with the hours correspondingly increased from 2 to 3 for both the CS and KA cores to accommodate it.

Why not incorporated:
