

CHAI - Contextual Honor Codes for AI Literacy and Academic Integrity

Project Report



Authors

Manika Garg*, Sunčica Bruck*, Theo Bakker*

Department

Research Group Learning Technology & Analytics, Centre of Expertise Global & Inclusive Learning

Date

July 2026

Type project

Study research

Version

1.0

* Manika Garg: supervision, design, implementation and report. Sunčica Bruck: literature review, design and implementation. Theo Bakker: supervision.

Summary

CHAI, Contextual Honor Codes for AI Literacy and Academic Integrity, is a project of the Research Group Learning Technology & Analytics at The Hague University of Applied Sciences (THUAS). The project is funded by the AI Hub and has run since the fourth quarter of 2025. CHAI responds to a gap between broad institutional and legal guidance on generative AI (GenAI) and the concrete decisions that students and teachers face in individual courses. It closes that gap through a co-creation process in which students and teachers jointly agree on how GenAI may be used in a specific learning context, and record that agreement in a customizable Honor Code.

This report documents the first phase of CHAI. It sets out the literature and policy context that motivated the project, describes the methodology that was designed and then refined between two classroom pilots carried out in February and May 2026, and presents the resulting CHAI Implementation Toolkit for teachers. The report draws on voting data and on interviews with the two participating teachers to identify what worked well and what needed to change. It closes with recommendations for a next phase of the project, including a proposed split between a design phase and a teaching phase of CHAI.

Table of Contents

Summary	2
Table of Contents	3
1. Introduction.....	4
1.1 Background and problem context.....	4
1.2 Objectives and scope	4
1.3 Reading guide	4
2. Literature Review.....	5
2.1 Honor codes and academic integrity	5
2.2 Generative AI in higher education: opportunities and risks.....	5
2.3 Generative AI and academic integrity	6
2.4 Policy and regulatory context.....	6
2.5 Co-creation as a methodological approach	6
3. Methodology	6
3.1 Overall approach.....	6
3.2 The Honor Code template and AI involvement levels	7
3.3 The CHAI voting session.....	7
3.4 Dispute resolution	7
3.5 Students who miss the session	7
3.6 Revisiting the Honor Code.....	8
3.7 Data collection for this report.....	8
4. Pilots	8
4.1 Pilot I: Public Administration.....	8
4.2 Pilot II: International Business	9
5. Learnings from the Pilots	10
5.1 What worked.....	10
5.2 Challenges and critical feedback	10
5.3 Cross-pilot synthesis	11
6. Recommendations	11
6.1 Refining the methodology further	11
6.2 Scaling CHAI across THUAS	11
6.3 Directions for further research	11
7. References	12
Appendix: CHAI Honor Code Template.....	13

1. Introduction

1.1 Background and problem context

Generative AI tools such as ChatGPT have become part of daily academic life for students and teachers. A 2024 Digital Education Council survey, reported by Al Zaidy (2024), found that 86 percent of students use AI for academic tasks, while only 5 percent report being fully aware of their institution's guidelines on AI use.

The consequences of this gap are visible beyond the classroom. In January 2026, the Dutch newspaper NRC reported on a University of Amsterdam thesis fraud case that turned on a single question an examination board could not answer with confidence: was the work in front of them written by the student or by a chatbot (Klute, De Bruijn, & Van der Wees, 2026)? The case illustrates a wider difficulty. Institutional and legal frameworks, including the EU AI Act and THUAS's own policy on the use of generative AI (Dienst OKC, 2024), set boundaries at a general level. They leave many concrete questions about GenAI use in specific courses and assignments open to interpretation.

Two further issues compound this gap. Students and teachers frequently disagree, or are simply unaware of each other's expectations, about what counts as acceptable AI use for a given task. In addition, most institutional responses to GenAI treat regulation and AI literacy as separate questions, one about compliance and one about skills, rather than as a shared conversation between the people who teach and the people who learn.

1.2 Objectives and scope

CHAI was set up to close this gap by facilitating the co-creation of AI Honor Codes: agreements developed jointly by students and teachers that formalize a shared understanding of GenAI use within their own learning context. The project rests on two pillars. The first, AI regulation, translates broad institutional and legal guidelines into practical, context-specific agreements. The second, shared AI literacy, brings students and teachers to a common understanding of what GenAI tools can and cannot do, regardless of their starting level of experience.

The project has three main components: designing a generalizable Honor Code methodology and template, piloting that methodology in real courses to assess its impact, and developing recommendations, including an adjusted research proposal, for using the methodology more widely across THUAS. It is carried out by a multidisciplinary team. AI literacy is contributed by the Research Group Learning Technology & Analytics (Prof. Theo Bakker, Dr. Manika Garg, and Sunčica Bruck, MSc). Ethics and law are contributed by the Research Group Multilevel Regulation (Prof. Barbara Warwas and Dr. Kanan Dhru). Co-creation methodology is contributed by the Research Group Sustainable Talent Development (Prof. Ellen Sjoer and Dr. Miranda de Hei). The project is funded by the AI Hub.

1.3 Reading guide

Chapter 2 reviews the literature and policy context behind CHAI: honor codes, the effects of GenAI on learning, GenAI and academic integrity, and the relevant policy and regulatory context. Chapter 3 describes the CHAI methodology. Chapter 4 presents the two classroom pilots conducted in the Public Administration and International Business programmes, the refinements made between them, and the resulting Implementation Toolkit. Chapter 5 brings together the learnings from both pilots, based on voting data, interviews, and written feedback from the two participating teachers. Chapter 6 sets out recommendations for a next phase of the project. Chapter 7 lists the references cited in this report. The appendix reproduces the CHAI Honor Code template.

2. Literature Review

2.1 Honor codes and academic integrity

Honor codes have a long history in higher education as a mechanism for building shared commitment to academic integrity. Hall and Kuh (1998) describe how honor codes at American universities work best when students see them as owned by their own community rather than imposed by the institution. Tatum and Schwartz (2017) review evidence that honor codes reduce self-reported cheating, but only when students are actively involved in drafting and enforcing them. Hamann and Kerwin (2018) reach a similar conclusion in their account of an honor code used to nurture academic writing skills.

Beyond honor codes specifically, Garg and Goel (2022) reviewed the wider set of integrity strategies used in online and digital learning environments, spanning proctoring, authentication, and institutional policy alongside honor-code-style agreements. They find that no single strategy is sufficient on its own, and that effective integrity strategies combine technical, pedagogical, and policy measures rather than relying on any one of them in isolation. This finding supports CHAI's approach, which pairs a documented agreement, the Honor Code, with a pedagogical process of co-creation and dialogue, rather than treating either element as sufficient by itself.

Dix, Emery, and Le (2014) apply an investment model to explain why some honor codes succeed and others do not. Commitment to a code grows out of students' sense of investment, satisfaction, and the social norms they perceive around it, rather than from the text of the code itself. Outside education, Singh's (2011) study of corporate codes of ethics reaches a related conclusion: written codes only change behavior when they are embedded in a wider programme of communication and buy-in, not when they exist as a standalone document. Raman and Ramlogan (2020) extend this logic to a professional and clinical context in their study of an honour code in dental training. Codes are more effective, they find, when tied to specific and recognizable activities rather than to abstract principles.

Taken together, this literature suggests that an honor code's effectiveness depends less on its content and more on how it is created and maintained. This is the starting premise of CHAI. Rather than drafting a single institution-wide AI policy, the project asks students and teachers to co-create context-specific agreements, on the assumption that ownership and dialogue, and not the rules themselves, are what change behavior.

2.2 Generative AI in higher education: opportunities and risks

GenAI adoption in higher education has outpaced institutional response. Stone (2024) surveyed 733 undergraduates and found widespread use of GenAI alongside considerable uncertainty about what use is acceptable, a pattern also reported by Al Zaidy (2024). Both studies point to the same conclusion: students want clearer guidance and more opportunities to build AI literacy, and most institutions have not yet provided either.

At the same time, a growing body of research raises concerns about the cognitive effects of relying on GenAI for academic work. Kosmyrna et al. (2025) compared essay writing with and without AI assistance and found that participants who used an AI assistant showed weaker neural engagement and recalled less of their own writing. The authors describe this pattern as an accumulation of cognitive debt. This concern is directly relevant to the CHAI Honor Code, which separates No AI, GenAI-assisted, and Full GenAI into distinct categories rather than treating AI use as a single yes-or-no question. An earlier version of the categories, used in Pilot I, split GenAI-assisted use further into separate planning and editing categories. This was later consolidated into the single GenAI-assisted category described in Chapter 3, based on the experience of Pilot I (see Chapter 4). The underlying premise, that different degrees of AI involvement are likely to affect learning in different ways, holds across both versions.

2.3 Generative AI and academic integrity

The literature on GenAI and academic integrity generally agrees that detection-based enforcement is a losing strategy. Janssen (2025) argues that AI governance works best when organizations treat AI as part of a complex, adaptive system that co-evolves with policy and practice, rather than trying to control it through fixed rules alone. This view supports CHAI's emphasis on dialogue and on periodically revisiting agreements, instead of relying on a static, one-time policy statement.

The NRC investigation into the University of Amsterdam fraud case (Klute et al., 2026) shows what happens when this dialogue is absent. Examination board members described being unable to judge with confidence whether a piece of writing reflected a student's own understanding. Students, for their part, described genuine confusion about where the boundaries of acceptable AI use lay. This case surfaced during the design phase of CHAI and is one of the concrete examples that motivated the project's focus on making expectations explicit before a course or assignment begins, rather than adjudicating them after the fact.

2.4 Policy and regulatory context

CHAI operates within a layered policy context. At the European level, the AI Act sets requirements for transparency and risk management that apply to AI systems used in education, though it leaves specific classroom practice to individual institutions. At the national level, the Npuls AI-GO! Framework for AI literacy in education provides a shared vocabulary for what AI literacy means in Dutch higher education. At THUAS, the Beleidsregels voor het gebruik van (G)AI (Dienst OKC, 2024) and the accompanying Haagse Handreiking AI-gebruik (The Hague University of Applied Sciences, 2024) set out core principles and practical guidance for staff and students.

Both THUAS documents are deliberately broad, so that they remain applicable across many programmes and teaching contexts. This breadth is a strength for institutional coherence, but it also means that a student in one classroom and a student in another may work under different, unwritten expectations under the same underlying policy. CHAI is designed to sit directly below this layer of institutional policy and to translate its principles into agreements at the level of an individual course.

2.5 Co-creation as a methodological approach

Co-creation can take many forms, ranging from structured, multi-stage workshops to lighter formats built around a single shared decision point. The Research Group Sustainable Talent Development shared an overview of co-creation methods during CHAI's design phase. In practice, the session format used in both pilots centered on one method from that overview: presenting each activity to the class, voting on it anonymously, and discussing the result as a group, described in full in Chapter 3. This lighter design reflects a practical constraint as much as a methodological choice, since a full multi-stage workshop was not feasible within a single class period.

The approach still reflects a wider premise in the co-creation literature: participatory processes tend to produce more durable agreements than agreements imposed top-down, provided the process itself is well facilitated.

3. Methodology

3.1 Overall approach

CHAI's methodology has three components, developed in sequence: a generalized Honor Code template and set of AI involvement levels, a facilitated voting session in which a class agrees how those levels apply to its own activities, and a dispute resolution process for cases where the class does not reach agreement. The full methodology, including step-by-step guidance for teachers, is set out in the CHAI Implementation Toolkit (July 2026).

3.2 The Honor Code template and AI involvement levels

Every educational activity in a course or assignment can be assigned one of three levels of GenAI involvement, summarized in Table 1.

Level of AI use	Description
No AI	The activity is completed entirely without GenAI.
GenAI-assisted	GenAI may be used to plan or edit content that the student produces independently. The substance of the work must remain the student's own.
Full GenAI	GenAI may be used to complete any element of the task.

Table 1. The three CHAI levels of GenAI involvement.

These levels are recorded, activity by activity, in a customizable Honor Code template together with the reasoning behind each choice and, where relevant, how AI use should be acknowledged, for example by sharing prompts or citing GenAI-generated content. The full template is reproduced in the appendix.

3.3 The CHAI voting session

Within the session, the teacher and students vote anonymously, using a live polling tool such as Woodclap or Mentimeter, on the appropriate GenAI involvement level for each activity on a pre-agreed list. The teacher votes alongside the students, as an equal participant in the agreement. Results are displayed to the group in real time and discussed before a final level is agreed for each activity, with a revote where needed. For activities agreed at GenAI-assisted or Full GenAI, the group also agrees how AI use should be acknowledged, for example whether prompts must be shared or GenAI content cited. A facilitator, or the teacher directly, then drafts the resulting Honor Code and shares it with the class for confirmation.

3.4 Dispute resolution

Where a vote does not produce full consensus, a structured process, developed with the Research Group Multilevel Regulation and set out in full in the Implementation Toolkit, guides the group to a decision. For each contested activity, the group first names the learning objective at stake, then follows three steps.

1. **Vote.** The teacher and students vote on one of the three GenAI involvement levels.
2. **Surface the underlying values.** The group identifies what is driving each position, for example academic integrity, workload pressure, privacy, learning and development, or fairness between students.
3. **Resolve.** If the vote reaches full consensus, the agreed level is recorded and the group moves to the next activity. If the class remains divided, the guiding question is whether the disputed use affects assessment integrity. If it does, the teacher safeguards the decision and explains the reasoning to the class. If it does not, the class discusses a concrete example case, or votes again, until a majority position emerges.

This structure has been applied to cases such as a request to use AI for language support in an individually assessed essay, weighed against a learning objective of independent argumentation, and a request to use AI for slide design in a group presentation, weighed against a learning objective of oral presentation skill.

3.5 Students who miss the session

A student who is not present during the in-class voting session cannot vote with the rest of the group. The information package and the draft Honor Code are shared with absent students afterwards, so that they can review the agreements and raise concerns before the Honor Code is adopted. They then commit to the agreed code in the same way as the rest of the class.

3.6 Revisiting the Honor Code

The Honor Code is treated as a living agreement rather than a fixed document. Two reflection moments are planned in advance: a mid-semester checkpoint, typically halfway through the module, and an end-of-module reflection on what worked and what did not.

The Honor Code is also reopened when specific events occur, rather than waiting for the next planned moment:

- the university's AI policy changes;
- a student or teacher requests it;
- the code is not abided by a student or a group of students;
- many students ask for exceptions afterwards, for example when a large share of the class emails to opt out or request exceptions;
- AI tools evolve significantly mid-semester;
- a conflict takes place, either between teacher and student or between students.

3.7 Data collection for this report

This report draws on the materials generated during the design and pilot phases of CHAI: the project proposal and one-pager, the Honor Code documents produced in each pilot, anonymous voting data collected through Wooclap, semi-structured interviews conducted with both participating teachers after their pilot session, and a follow-up written reflection submitted by one of the teachers. Interviews were recorded, transcribed, and reviewed by the project team to identify recurring themes. Those themes are reported in Chapter 5.

4. Pilots

4.1 Pilot I: Public Administration

The first pilot was conducted on 25 February 2026 with a Bestuurskunde en Overheidsmanagement class, in the course Belangen behartigen (Representing Interests). The session was facilitated by Sunčica Bruck, with Dr. Manika Garg present as researcher. The class teacher, referred to as Teacher 1 in this report and identified in Table 2, is a member of THUAS's ethical committee. The class discussed five activities drawn from a debate-day assignment and voted on the appropriate GenAI level for each.

The resulting Honor Code recorded, for example, that describing one's own party's position should involve no GenAI. Students reasoned that the information is easy to find and that the exercise is meant to test their own understanding of it.

Field	Details
Programme	Bestuurskunde en Overheidsmanagement (Public Administration and Government Management)
Course	Belangen behartigen / Representing Interests (BO-DU1BBPI-21)
Date	25 February 2026
Teacher	Dr. Theo Zijderveld (referred to as Teacher 1 in this report)
Facilitator	Sunčica Bruck, MSc
Context	Per module (debate-day assignment)
Activities voted on	5

Table 2. Overview of Pilot I.

4.2 Pilot II: International Business

Between Pilot I and Pilot II, the methodology was refined based on Teacher 1's feedback. The four AI-involvement categories used in Pilot I, which separated planning and editing within GenAI-assisted use, were consolidated into the three levels described in Section 3.2. The signature requirement was also reconsidered: rather than collecting individual signatures, the agreed Honor Code was published on Brightspace for the whole class. This reduced a source of hesitation that Teacher 1 had flagged, though a signature may also create a stronger individual sense of commitment than a shared, published document. This trade-off has not yet been tested directly.

The second pilot was conducted on 11 May 2026 with an International Business Semester 2 class working on the Project Business Plan (Milestone 3). The class mentor, referred to as Teacher 2 in this report and identified in Table 3, worked with a course coordinator on the milestone deliverable. The session was facilitated by Dr. Manika Garg. Twenty-two students voted, using Wooclap, on the appropriate GenAI level for five learning activities selected from the milestone deliverable: writing the conclusion, writing the recommendations, rewriting and restructuring earlier milestone content, writing the executive summary, and writing short chapter summaries.

Voting results showed a consistent pattern. GenAI-assisted was the most common response across all five activities, but every activity also received votes for No AI and for Full GenAI, and several students changed their position between activities. For writing the conclusion, opinions ranged from students who argued AI should play no role, because students should be able to reach their own conclusions, to students who saw AI as useful for checking whether a conclusion followed from the evidence. For the more mechanical task of writing chapter summaries, sentiment shifted further toward Full GenAI. Several students explained this shift by noting that summarizing already-agreed content is less central to the learning objective than drafting original analysis.

The resulting Honor Code recorded GenAI-assisted as the agreed level for most activities, with an explicit note that the substance of the analysis, as distinct from its presentation, should remain student-owned.

Field	Details
Programme	International Business, Semester 2
Course	Project Business Plan (Milestone 3)
Date	11 May 2026
Mentor	Corina Tăbăcaru (referred to as Teacher 2 in this report)
Course coordinator	Mark Li Fo Sjoë
Facilitator	Dr. Manika Garg
Context	Milestone 3 assignment
Participants	22 students
Activities voted on	5

Table 3. Overview of Pilot II.

The step-by-step guidance developed across both pilots, covering how to prepare the activity list, run the voting session, and handle disagreement, was subsequently compiled into a standalone CHAI Implementation Toolkit for teachers (July 2026). The toolkit allows any teacher at THUAS to prepare and run a CHAI session independently, without depending on an external facilitator from the project team.

5. Learnings from the Pilots

5.1 What worked

Both teachers confirmed that CHAI addresses a real and current problem. Teacher 2 described GenAI as having entered education faster than the people designing and teaching courses have been able to respond. In her view, many courses still operate without a clear framework for AI use, and CHAI helps raise students' agency and awareness around responsible AI use. Teacher 1 similarly found value in testing whether students could articulate an honor code in their own terms, instead of simply receiving one from the institution.

The voting data from Pilot II showed that structured dialogue surfaces disagreement that would otherwise stay invisible. Students who wanted no AI involvement in a task sat alongside students who wanted full AI involvement in the very same task, and the session gave both groups a chance to explain their reasoning to the teacher and to each other. Explicit disagreement of this kind is difficult to generate through a written policy alone.

5.2 Challenges and critical feedback

Several concerns emerged in both pilots and were raised independently by both teachers.

Session length and depth. Both teachers felt that a single 45 to 60 minute session is a light intervention against what Teacher 2 called a deeply ingrained habit of outsourcing thinking to AI. Teacher 1 observed that keeping students engaged through repeated structured activities is harder in a real classroom than the research design assumed. Some students treated anonymous polling as an opportunity for jokes rather than as a genuine input.

Social desirability and anonymity. Teacher 2 raised the concern that voting in front of peers, the teacher, and an external researcher is not fully anonymous in practice, and that students may give answers they believe are expected of them rather than their honest view. Teacher 1 made a related point from the opposite direction: large-group discussion of ethical topics does not work well, and word clouds in particular produce low-quality input, because anonymity removes accountability without producing candor.

Whether AI use should be negotiable at all. This was the most significant disagreement to emerge from the pilots, and it was articulated most directly by Teacher 2. In her view, pedagogical decisions about AI use should already be settled during course design, not put to a vote during teaching. She proposed splitting CHAI into two distinct phases: a design phase, in which teachers and students, or student representatives, negotiate AI policy while a course is still being built, and a teaching phase, in which no further voting occurs and students instead receive a structured explanation of the reasoning behind decisions already made.

Teacher 1's feedback points in a related direction. He questioned whether the discrete levels of AI involvement reflect how AI is actually used in practice, since planning and editing frequently overlap. He suggested that the methodology should define types of AI use rather than levels of AI use. He also noted that No AI is unrealistic as an enforceable category outside a fully controlled environment, since students will use AI for idea generation regardless of what the Honor Code states. The more productive question, in his view, is how AI is used, not whether it is used at all.

Facilitation and ownership. Teacher 1 suggested that the methodology works better when the teacher leads the session directly, rather than an external facilitator, since students notice immediately when someone is unfamiliar with the class and classroom management skill affects how well the session runs. He also flagged that requiring a signature on the Honor Code may introduce an administrative and psychological barrier that is not necessary for the agreement itself to be meaningful. This concern was addressed in Pilot II, where the agreed Honor Code was published on Brightspace instead of requiring individual signatures (see Section 4.2).

5.3 Cross-pilot synthesis

Read together, the two pilots point to the same underlying tension from different angles. CHAI succeeds at making disagreement about AI use visible and giving it a structured outlet, which both teachers valued. It is less certain that a single classroom vote can settle what is, in practice, an ongoing and evolving question, or that it should try to. The pilots suggest that CHAI's greatest value may lie earlier in the process, while a course is still being designed and AI-related decisions are genuinely open. Its role during teaching is better understood as building shared understanding of decisions already made, rather than as negotiating those decisions in real time.

6. Recommendations

6.1 Refining the methodology further

Building on the refinements already made between Pilot I and Pilot II, described in Chapter 4, three further changes are recommended.

- Separate the design and teaching phases of CHAI explicitly, following the model proposed by Teacher 2. Use co-creation and voting during course design, when decisions are genuinely open, and reserve the teaching-phase session for explaining the reasoning behind agreed policy and for building AI literacy, without further voting.
- Test whether teacher-led delivery, now supported by the Implementation Toolkit, produces different classroom dynamics than the facilitator-led sessions used in both pilots to date. Teacher 1 suggested that students respond differently when the teacher leads the session directly.
- Explore whether a type-based framing of AI involvement, distinguishing kinds of AI use rather than degrees of AI use, would resolve the planning-and-editing overlap Teacher 1 raised in Pilot I more directly than the three-level scale already in place.

6.2 Scaling CHAI across THUAS

Before extending CHAI to additional programmes, the design-phase model proposed above should be piloted at least once, so that its practicality can be assessed against the same criteria used for the first two pilots. Any scaled version of CHAI should also address the anonymity concern raised in Pilot II, for example by separating the polling tool from any visible link to individual students and by considering asynchronous input alongside live voting. With the Implementation Toolkit now available, teachers across THUAS can prepare to run a CHAI session without relying on the original project team as facilitator. A next step is to track how these independently run sessions compare to the facilitator-led pilots reported here.

6.3 Directions for further research

An adjusted research proposal for the next phase of CHAI should examine whether design-phase co-creation produces Honor Codes that hold up better over the course of a semester than the single-session model piloted here. It should also include a follow-up measure of academic conduct or student-reported AI use after a Honor Code has been in place for several months, since neither pilot in this report includes that longer-term evidence. It would also be valuable to test a type-based framing of AI involvement directly against the current three-level scale, to establish whether it resolves Teacher 1's concern in practice and not only in principle.

7. References

- Al Zaidy, A. (2024). The impact of generative AI on student engagement and ethics in higher education. *Journal of Information Technology, Cybersecurity, and Artificial Intelligence*, 1(1), 30-38. <https://doi.org/10.70715/jitcai.2024.v1.i1.004>
- Dienst OKC. (2024). Beleidsregels voor het gebruik van (G)AI op De Haagse Hogeschool. The Hague University of Applied Sciences.
- Dix, E. L., Emery, L. F., & Le, B. (2014). Committed to the honor code: An investment model analysis of academic integrity. *Social Psychology of Education*, 17, 179-196. <https://doi.org/10.1007/s11218-013-9246-2>
- European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act).
- Garg, M., & Goel, A. (2022). A systematic literature review on online assessment security: Current challenges and integrity strategies. *Computers & Security*, 113, 102544. <https://doi.org/10.1016/j.cose.2021.102544>
- Hall, T. L., & Kuh, G. D. (1998). Honor among students: Academic integrity and honor codes at state-assisted universities. *Journal of Student Affairs Research and Practice*, 36(1), 2-18. <https://doi.org/10.2202/1949-6605.1069>
- Hamann, H., & Kerwin, M. (2018). Nurtured on Wikipedia, can an honor code foster better student writers? *New Directions for Community Colleges*, 2018(183). <https://doi.org/10.1002/cc.20315>
- Janssen, M. (2025). Responsible governance of generative AI: Conceptualizing GenAI as complex adaptive systems. *Policy and Society*, 44(1), 38-51. <https://doi.org/10.1093/polsoc/puae040>
- Klute, E., De Bruijn, M., & Van der Wees, J. (2026, January 11). What are theses worth now that AI is co-writing? NRC.
- Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task (arXiv:2506.08872). MIT Media Lab.
- Npuls. (2024). AI-GO! A framework for AI literacy in education.
- Raman, V., & Ramlogan, S. (2020). Academic integrity and the implementation of the honour code in the clinical training of undergraduate dental students. *International Journal for Educational Integrity*, 16, 9. <https://doi.org/10.1007/s40979-020-00058-2>
- Singh, J. B. (2011). Determinants of the effectiveness of corporate codes of ethics: An empirical study. *Journal of Business Ethics*, 101, 385-395. <https://doi.org/10.1007/s10551-010-0727-3>
- Stone, B. W. (2024). Generative AI in higher education: Uncertain students, ambiguous use cases, and mercenary perspectives. *Teaching of Psychology*, 52(3), 347-356. <https://doi.org/10.1177/00986283241305398>
- Tatum, H., & Schwartz, B. M. (2017). Honor codes: Evidence-based strategies for improving academic integrity. *Theory Into Practice*, 56(2), 129-135. <https://doi.org/10.1080/00405841.2017.1308175>
- The Hague University of Applied Sciences. (2024). Haagse Handreiking AI-gebruik: Handvatten om AI-tools te gebruiken in lijn met wet en waarden.

Appendix: CHAI Honor Code Template

Date: _____

Course ID: _____

Teacher: _____

Context (e.g. the entire course / individual assignment / module): _____

The goal of this Honor Code is to formalize the agreements reached during the CHAI session on the responsible use of GenAI in the specific learning context. The following rubric outlines the agreements made in this discussion:

GenAI Involvement Rubric

Level of AI use	Description of levels	Educational activities	Acknowledgment of AI use	Reasoning behind this choice (optional)
1. No AI	These tasks will be completed entirely without GenAI.	...	N/A	
2. GenAI-assisted	GenAI may be used to plan or edit individually produced content.	...	Do prompts or chat logs need to be provided? ☐ Yes ☐ No ☐ Only upon request from the instructor	
3. Full GenAI	GenAI may be used to complete any elements of the task.*	...	Does GenAI content have to be cited? ☐ Yes ☐ No ☐ Only upon request from the instructor	

** All external guidelines and limitations (e.g. the EU AI Act, GDPR, THUAS-wide AI policies) still apply.*

☐ I have read and understand the above agreements.

☐ The above rubric accurately reflects my own intentions with GenAI in this course.

☐ I agree to adopt this Honor Code as an affirmation of my commitment to upholding the values, principles and ethics of academic integrity in this course and THUAS as a whole.

Signature: _____

Name: _____