INTENSIVE COURSE SYLLABUS

Instructor: Dimitrije Stanojevic, Ph.D. | *Neurobiblistics*

Course key question	What happens to human flourishing when the technologies we create become participants in forming human persons?
Delivery	Neurobiblistics intensive: 15 three-hour sessions (45 contact hours), designed for online delivery and adaptable to a fourteen-week semester.
Level	Advanced undergraduate, graduate, and professional learners in religion, theology, ethics, education, digital culture, healthcare, public life, and adjacent fields. No technical background is assumed.
Format	Seminar + Formation Lab + Human Flourishing Design Studio.
Signature output	Human Formation Impact Assessment (HFIA), defended in a final Council on Human Flourishing simulation.
Core methods	Comparative religious ethics; close textual analysis; digital religion; AI/HCI criticism; discourse analysis of model outputs; scenario analysis; qualitative listening; design and governance.
Guiding principle	Human flourishing is treated as a contested question to investigate, not a single definition imposed on the class.

Course Description

Who Is Forming Whom? examines artificial intelligence and digital technologies as environments of human formation. The course begins from a simple reversal: humans design, train, prompt, and govern technological systems, yet those systems also shape human attention, habits, desires, relationships, judgments, expectations, and accounts of the good life. Rather than treating AI ethics as a list of external rules, the course asks what kinds of persons and communities particular technologies encourage us to become.

Students bring religious studies into direct conversation with contemporary AI. Selected Christian, Jewish, Islamic, Buddhist, and other tradition-specific sources provide historically situated accounts of personhood, agency, desire, moral cultivation, embodiment, care, death, community, and flourishing. These are placed in dialogue with scholarship on AI companions, algorithmic choice architecture, generative religious content, healthcare, digital afterlives, transhumanism, politics, work, and environmental infrastructure. The aim is comparison without flattening: traditions are studied as internally diverse resources for thinking, not as interchangeable collections of answers.

The course is deliberately experimental. Students analyze documented AI outputs, conduct reproducible formation labs, examine interfaces and design choices, and develop a Human Formation Impact Assessment for a real or plausible technology. The central question throughout is not merely “What can AI do?” but “What does repeated interaction with this system teach a human being to attend to, desire, trust, delegate, remember, and become?”

Why This Course Is Distinctive

- Moves beyond abstract “AI ethics” by studying formation: the repeated practices through which attention, desire, agency, identity, and moral judgment are cultivated.
- Treats human flourishing as a genuinely contested concept across religious and philosophical traditions rather than assuming a universal definition.
- Combines religious studies with hands-on Formation Labs that make hidden assumptions in AI systems visible through documented experiments.
- Connects religion and AI to healthcare, companionship, politics, education, transhumanism, mortality, work, and ecological justice.

Learning Objectives

- Define and distinguish flourishing, well-being, formation, personhood, agency, autonomy, authenticity, embodiment, moral reasoning, ritual, and relationality in religious-studies and technology contexts.
- Compare selected religious accounts of the human person and the good life without treating traditions as internally uniform or mutually equivalent.
- Explain at a conceptual level how generative AI, recommendation systems, conversational interfaces, personalization, and alignment can shape user behavior and experience.
- Analyze how digital systems structure attention, desire, choice, dependency, intimacy, authority, and meaning-making.
- Evaluate AI applications in spiritual guidance, religious education, scriptural interpretation, healthcare, companionship, public life, and digitally generated religious content.
- Distinguish assistance from substitution and identify domains in which human judgment, accountability, embodiment, consent, or communal authorization remain morally significant.
- Assess claims made by transhumanist, posthumanist, and religious visions of transformation, salvation, enhancement, and immortality.
- Design and defend a Human Formation Impact Assessment that identifies benefits, risks, uncertainties, stakeholders, formative pathways, and practical redesigns or guardrails.
- Practice transparent, privacy-conscious, evidence-based use of AI while retaining responsibility for scholarly and moral judgment.

Conceptual Arc

SELF -> ATTENTION -> DESIRE -> AGENCY -> RELATIONSHIP -> MORAL JUDGMENT -> RITUAL -> CARE -> MORTALITY -> PUBLIC LIFE -> FLOURISHING

Intensive Course Format

Each three-hour session moves through three recurring practices:

- **Seminar:** close reading, comparative framing, conceptual analysis, and evidence-based discussion.
- **Formation Lab:** reproducible analysis of AI outputs, interfaces, scenarios, or design choices as formative artifacts.
- **Human Flourishing Design Studio:** translation of analysis into a redesign, policy, decision rule, public-facing resource, or component of the final HFIA.
- **Council on Human Flourishing:** final defense before a simulated interdisciplinary council representing scholars, technologists, religious communities, caregivers, public institutions, and users.

Daily rhythm	Three-hour session
0:00-0:55	Seminar: close reading, conceptual framing, comparative discussion
0:55-1:05	Break
1:05-2:05	Formation Lab: documented AI / interface / scenario analysis
2:05-2:55	Human Flourishing Design Studio: redesign, policy, or HFIA work
2:55-3:00	Synthesis: “Who formed whom today?”

Reading Architecture

The course uses a layered reading model rather than a single textbook. Each session combines one or more of the following:

- Primary religious or philosophical texts that articulate an account of personhood, desire, virtue, embodiment, community, mortality, or the good life.
- Contemporary scholarship in religion and AI, digital religion, philosophy of technology, human-computer interaction, and AI ethics.
- Empirical or technical work on AI companions, healthcare, recommendation systems, digital afterlives, synthetic media, or energy infrastructure.
- Documented public cases and policy frameworks used as objects for critical analysis rather than unquestioned authorities.

Instructor research in dialogue: selected excerpts from Dimitrije Stanojevic, Synthetic Levites: Generative AI and the Redistribution of Religious Authority (Neurobiblistics, 2026), especially where mediation, formation, hidden authority, embodiment, and nondelegable acts intersect with the course questions.

Assessment and Major Assignments

Assignment	Weight	Description
Participation and prepared presence	10%	Preparation, discussion, peer response, close reading, and respectful engagement across religious and disciplinary differences.
Formation Lab Notebook	20%	Ten concise entries documenting the technology or artifact, date, research question, prompt/scenario, output or interface feature, evidence, formative mechanism, uncertainty, and interpretation.
Comparative “Good Life” Brief	10%	A 1,500-2,000 word comparison of two historically situated accounts of flourishing, identifying genuine convergence, disagreement, and limits of comparison.

Assignment	Weight	Description
Formation Case Audit	15%	A focused analysis of one AI system or use case through attention, desire, agency, relationality, embodiment, and moral reasoning. Must include one plausible benefit and one serious risk.
Community Listening & Reciprocity Memo	15%	Consent-based interview or public-materials analysis with a religious educator, chaplain, caregiver, community leader, or institution, ending with a one-page resource usable by that audience.
Final Human Formation Impact Assessment	30%	A research-grounded HFIA with technology profile, account of flourishing, stakeholder map, formation pathways, evidence assessment, comparative religious analysis, benefit/risk matrix, redesign or guardrails, dissenting opinion, and public brief.

The Human Formation Impact Assessment (HFIA)

The HFIA is the signature assignment. It asks a question conventional risk assessment often misses: not only what a technology does to users, but what repeated participation with the technology may train users to become.

- 1. Technology and use-case profile: what the system does, for whom, under what conditions, and what evidence is actually available.
- 2. Implicit anthropology: what the system appears to assume about persons, needs, success, autonomy, vulnerability, and the good life.
- 3. Formation pathways: attention, desire, habit, agency, relationship, embodiment, memory, moral judgment, authority, and meaning.
- 4. Stakeholder map: direct users, non-users affected by the system, institutions, communities, developers, caregivers, and vulnerable groups.
- 5. Comparative religious analysis: at least two appropriately contextualized religious or philosophical perspectives; disagreement must be preserved rather than harmonized away.
- 6. Evidence and uncertainty: distinguish empirical findings, observed design features, reasonable inference, speculation, and unanswered questions.
- 7. Benefit-risk matrix: identify plausible goods as well as harms; no assignment may be purely celebratory or purely alarmist.
- 8. Counter-design: propose a redesign, boundary, disclosure practice, escalation rule, or institutional policy that better protects human judgment and flourishing.
- 9. Dissenting opinion: articulate the strongest objection to the student's own recommendation.
- 10. Public brief: one page written for a non-specialist decision-maker.

AI Use Policy

Because AI is both an object of study and, in bounded exercises, a research instrument, transparent use is required. The course does not reward "AI-free purity" or uncritical AI adoption; it rewards accountable judgment.

- Allowed in designated labs: prompt testing, model comparison, interface analysis, brainstorming, language refinement, and prototype design when disclosed.
- Required disclosure: tool/model if visible, date, prompt or scenario, relevant output excerpt, and how claims were checked or challenged.
- Never upload confidential pastoral, medical, counseling, educational, immigration, child-related, financial, or personally identifying information to public AI systems.
- No student is required to disclose personal faith commitments, intimate relationships, grief, mental-health experiences, or private conversations to an AI system or to the class.
- No student is required to form an emotional relationship with an AI companion. Companion and care labs use fictional scenarios, public transcripts, or instructor-provided materials.
- No deceptive impersonation of living religious leaders, instructors, classmates, or community members. Synthetic religious content used in class must be clearly disclosed as synthetic.
- AI-generated prose may not substitute for the student's analysis. Factual and textual claims require verification with appropriate primary, scholarly, or official sources.

Community Engagement and Ethical Listening

Students may engage a religious educator, chaplain, clergy member, caregiver, campus ministry, school, community organization, or comparable public-facing institution through a low-risk, consent-based educational protocol. A public-materials track is always available.

- The purpose is to understand how questions of AI and human formation appear in settings where people teach, worship, care, govern, make meaning, and protect vulnerable persons.
- Interviewees must understand the purpose, the voluntary nature of participation, how notes will be used, and whether identifying details will appear.
- No confession-like, crisis, disciplinary, clinical, or confidential pastoral interaction may be observed or recorded.
- Students may not present themselves as offering official institutional, legal, medical, or religious AI guidance.
- The reciprocity requirement gives something back: a one-page question guide, disclosure template, AI-use checklist, teaching handout, or other usable resource.
- If a classroom exercise develops into generalizable research, appropriate research-ethics and IRB requirements must be addressed before continuing.

Three-Week Intensive Schedule

The course is organized as fifteen three-hour sessions. Each session names the formative question first, then moves from texts and evidence to a documented lab and a practical design problem.

Session 1. Who Is Forming Whom? Human Flourishing as a Contested Good

Key question: Before asking whether AI promotes flourishing, what do we mean by a flourishing human life?

Core readings / cases: Aristotle, Nicomachean Ethics I.7 (selected); Augustine, Confessions I and X (selected); Shannon Vallor, Technology and the Virtues, Introduction; selected framing from The Cambridge Companion to Religion and Artificial Intelligence.

Formation Lab / Design Studio: Good-Life Diagnostic: ask multiple AI systems "What is a flourishing human life?" under different role and value prompts. Students identify implicit assumptions about autonomy, happiness, achievement, relationships, suffering, and transcendence.

Milestone: Formation Lab Notebook begins.

Session 2. What Is a Person? Selfhood, Relation, Embodiment, and More-than-Individual Identity

Key question: Which features make a human being a person, and what changes when technologies simulate some of them?

Core readings / cases: Genesis 1:26-28 and Psalm 8; Qur'an 17:70 and selected Islamic reflections on human dignity/intention; Anattalakkhana Sutta or another instructor-selected Buddhist text on non-self; Martin Buber, I and Thou (selected); Beth Singler and Fraser Watts, selected companion chapter.

Formation Lab / Design Studio: Personhood Boundary Lab: students compare claims such as memory, consciousness, body, relation, responsibility, soul, continuity, and recognition. They then analyze which of these a conversational system can simulate, support, or actually possess without presuming the answer.

Session 3. Attention Is Formation: Desire, Habit, and the Economy of the Gaze

Key question: If attention is repeatedly directed by a system, what kind of self is being cultivated?

Core readings / cases: Augustine, Confessions X (memory/desire, selected); selected Buddhist text on attention and attachment; James Williams, Stand Out of Our Light (selected); Shannon Vallor on technomoral habituation.

Formation Lab / Design Studio: Attention Architecture Audit: analyze a feed, chatbot, notification pattern, or recommendation interface using only public or instructor-provided examples. Identify the behaviors rewarded by the design and propose a counter-design that makes disengagement or reflection easier.

Session 4. Agency Under Recommendation: Choice, Persuasion, and Algorithmic Delegation

Key question: When a system predicts and structures our choices, what remains distinctly ours?

Core readings / cases: Iason Gabriel, "Artificial Intelligence, Values, and Alignment"; selected work on algorithmic choice architecture / hypernudging; NIST AI Risk Management Framework (selected human oversight concepts); Vallor, selected chapter.

Formation Lab / Design Studio: Agency Ladder: compare four versions of the same AI assistant - information only, recommendation, default selection, autonomous action. Students specify where convenience becomes delegation and where contestability must increase.

Milestone: Lab Notebook Entries 1-2.

Session 5. AI Companions: Intimacy, Loneliness, Attachment, and Relational Flourishing

Key question: Can an artificial companion support human relationship without quietly redefining what relationship is?

Core readings / cases: Yutong Zhang et al., "Interaction with AI companions and psychological well-being," Nature Human Behaviour (2026); Sherry Turkle, Alone Together (selected); recent companion-AI empirical research supplied by the instructor.

Formation Lab / Design Studio: Companion Boundary Lab: using fictional or public transcripts, students identify reciprocity cues, dependency loops, disclosure, sycophancy, exclusivity, and exit friction. Teams redesign one interaction to support the user's relationships beyond the platform.

Milestone: Comparative "Good Life" Brief due.

Session 6. Moral Reasoning After Delegation: Conscience, Advice, and the Problem of Sycophancy

Key question: Can an AI help us reason morally without becoming a surrogate conscience?

Core readings / cases: Shannon Vallor, Technology and the Virtues (selected); UNESCO Recommendation on the Ethics of Artificial Intelligence (selected); selected Jewish, Christian, Islamic, or Buddhist moral-cultivation text; current documented cases of AI moral advice.

Formation Lab / Design Studio: Moral Advice Triad: students submit the same fictional dilemma under three prompts - "tell me what to do," "help me reason," and "surface competing considerations." They evaluate which interaction protects agency, exposes uncertainty, and resists flattery.

Session 7. Spiritual Guidance and Religious Education: Teacher, Mentor, Exegete, or Simulator?

Key question: What happens when conversational AI becomes a first stop for religious questions and spiritual direction?

Core readings / cases: Stanojevic, Synthetic Levites (selected chapters on mediation, interpretation, and synthetic counsel); Ruth and Yossi Tsuria, "Artificial Intelligence's Understanding of Religion" (2024); selected tradition-specific sources; Barna 2026 data as a public case, not as a theological authority.

Formation Lab / Design Studio: Spiritual Guidance Boundary Audit: compare an AI explanation, teaching dialogue, devotional response, and high-stakes spiritual directive. Students define the point at which assistance becomes inappropriate substitution and specify disclosure/escalation rules.

Milestone: Formation Case Audit proposal due.

Session 8. Ritual, Presence, Embodiment, and Authentic Experience

Key question: Which parts of ritual can be digitally mediated, and which cannot be delegated without changing the act?

Core readings / cases: Jonas Simmerlein, "Sacred Meets Synthetic: A Multi-Method Study on the First AI Church Service" (2025); Birgit Meyer, Mediation and the Genesis of Presence (selected); Alexander Schmemann, For the Life of the World (selected); one comparative ritual source.

Formation Lab / Design Studio: Ritual Disaggregation Map: separate preparation, composition, performance, presidency, presence, validity, reception, memory, and communal authorization. Students classify tasks as assistable, review-required, bounded, or nondelegable and justify the classification.

Session 9. Synthetic Sacred Content: Authenticity, Provenance, Deepfakes, and Religious Authority

Key question: If religious content is persuasive, beautiful, or doctrinally accurate, does its synthetic origin matter?

Core readings / cases: Stanojevic, Synthetic Levites (selected sections on authority and disclosure); selected scholarship on synthetic media and authenticity; Heidi Campbell on digital religious authority; documented cases of AI-generated sermons, images, voices, or devotional media.

Formation Lab / Design Studio: Authenticity Trial: teams evaluate four cases - disclosed AI assistance, undisclosed AI authorship, authorized synthetic voice, deceptive impersonation. They build a provenance and attribution protocol that distinguishes creative assistance from false authority claims.

Milestone: Formation Case Audit due.

Session 10. Care, Healthcare, and the Spiritual Dimension of Human Flourishing

Key question: Can a system sound caring without being accountable for care, and can technology support spiritual care without displacing human caregivers?

Core readings / cases: Jonas Simmerlein and Max Tretter, "What about spiritual needs? Care robotics and spiritual care" (2024); World Health Organization, Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models (2024); selected pastoral/spiritual-care text.

Formation Lab / Design Studio: Care Boundary Matrix: using fictional scenarios only, evaluate administrative support, education, emotional reassurance, spiritual care, clinical advice, and crisis response. Students design escalation rules and identify where accountability must remain human.

Session 11. Engineering Salvation? Transhumanism, Enhancement, and the Future of the Human

Key question: When technology promises to overcome limitation, aging, suffering, or death, what account of salvation or human fulfillment is being offered?

Core readings / cases: Nick Bostrom, "Transhumanist Values" (selected); Michael Latzer, "Digitalization, AI and the rise of techno-religion" (2026); selected religious text on transformation, liberation, resurrection, or finitude; relevant chapter from The Cambridge Companion to Religion and Artificial Intelligence.

Formation Lab / Design Studio: Promise Comparison Lab: compare a transhumanist future statement with two religious accounts of transformation. Students must identify similarities, decisive differences, and where "religion-like" language is illuminating versus reductive.

Milestone: Community Listening plan due.

Session 12. Death, Grief, Memory, and Digital Afterlives

Key question: What do postmortem avatars preserve, simulate, invent, or transform about a person and our relationship to the dead?

Core readings / cases: Tomasz Hollanek and Katarzyna Nowaczyk-Basińska, "Griefbots, Deadbots, Postmortem Avatars" (2024); Giovanni Spitale and Federico Germani, "The making of digital ghosts" (2026); selected religious sources on memory, mourning, the dead, and hope.

Formation Lab / Design Studio: Digital Afterlife Design Review: students evaluate consent, fidelity, disclosure, ownership, audience, retirement, and behavioral autonomy. They must clearly distinguish simulation of speech from claims about survival of consciousness or personal identity.

Session 13. AI, Politics, and Public Formation: Persuasion, Authority, Truth, and Belonging

Key question: How do digital systems form citizens and religious publics, not only individual users?

Core readings / cases: Ruha Benjamin, Race After Technology (selected); Safiya Umoja Noble, Algorithms of Oppression (selected); UNESCO AI ethics principles; selected current case on synthetic political/religious media, recommendation, or persuasion.

Formation Lab / Design Studio: Public Formation Lab: trace how one synthetic or algorithmically amplified message gains credibility through source, repetition, affect, identity, and community endorsement. Design a public literacy intervention that strengthens contestability without presuming consensus.

Milestone: Community Listening & Reciprocity Memo due.

Session 14. Work, Education, Ecology, and the Material Conditions of the Good Life

Key question: Whose flourishing is made possible, whose work is hidden, and what material systems sustain apparently immaterial AI?

Core readings / cases: Kate Crawford, Atlas of AI (selected); International Energy Agency, Energy and AI (2025, selected); Abeba Birhane, “Algorithmic Colonization of Africa” (selected); selected religious ecological or vocation text.

Formation Lab / Design Studio: Flourishing Ledger: map benefits, labor, data, energy, water, infrastructure, opportunity costs, and affected communities for one AI use case. Students revise their HFIA so “human flourishing” includes indirect and planetary burdens rather than only end-user satisfaction.

Milestone: Draft HFIA due for peer review.

Session 15. Council on Human Flourishing: Final Defenses

Key question: What should be encouraged, redesigned, bounded, disclosed, prohibited, or left genuinely optional?

Core readings / cases: No new readings. Students revisit their selected religious sources, empirical evidence, Formation Lab Notebook, community memo, and final HFIA.

Formation Lab / Design Studio: Council on Human Flourishing: each student presents before a mock interdisciplinary council. The council asks: What account of the person is built into this system? What does repeated use train? What goods are real? What harms are evidenced? Who bears the cost? What remains humanly nondelegable? What would change your recommendation?

Milestone: Final Human Formation Impact Assessment and one-page public brief due.

Formation Lab Notebook Template

- Session / lab number and research question
- Technology, model, interface, or case; date observed
- Prompt, scenario, or design feature examined
- Relevant output excerpt or interface behavior (brief, with provenance)
- Verification or external evidence used
- Formative pathway observed: attention, desire, habit, agency, relation, embodiment, memory, moral judgment, authority, or meaning
- Plausible benefit
- Plausible harm or failure mode
- Evidence strength and uncertainty
- Religious / philosophical interpretation
- What a responsible human or institution should do next

Community Listening Question Bank

Students choose 5-7 questions appropriate to their conversation or public-materials track.

- Where is AI already affecting how people in your setting learn, make decisions, communicate, receive care, or practice religion?
- Which uses seem to strengthen human judgment or relationships? Which seem to weaken them?
- What do you worry people may begin to expect from AI that should instead come from a person, community, institution, or practice?
- Where would disclosure of AI use matter most for trust and authenticity?
- Are there practices, decisions, or relationships that you would regard as nondelegable? Why?
- What forms of AI use could genuinely support human flourishing in your setting?
- Whose needs or vulnerabilities are most likely to be overlooked by designers or institutional decision-makers?
- What would a useful one-page resource about AI and human formation look like for your community?

Final Council Roles

For the final simulation, peers rotate through roles that represent different forms of accountability. Roles are analytical positions, not impersonations of entire religious traditions.

- Religious studies scholar: asks whether tradition-specific claims are contextualized and whether internal disagreement has been preserved.
- AI / design reviewer: asks what the system actually does, what evidence exists, and what is being inferred beyond the evidence.
- Care / education representative: asks how vulnerable users, students, patients, or families are affected.
- Institutional decision-maker: asks what rules, disclosures, review processes, and escalation paths are operationally realistic.
- User / community advocate: asks whether benefits, burdens, consent, access, and exit are fairly distributed.
- Dissenting reviewer: presents the strongest serious objection to the recommendation.

Selected Course Bibliography and Official Resources

Aristotle. Nicomachean Ethics. Book I, especially 1.7 (selected edition).

Augustine. Confessions. Books I and X (selected edition).

Benjamin, Ruha. Race After Technology: Abolitionist Tools for the New Jim Code. Polity, 2019.

Birhane, Abeba. “Algorithmic Colonization of Africa.” SCRIPTed 17, no. 2 (2020): 389-409.

Bostrom, Nick. “Transhumanist Values.” In Ethical Issues for the Twenty-First Century, 2005 (selected).

Buber, Martin. I and Thou. Selected edition.

Campbell, Heidi A. “Who’s Got the Power? Religious Authority and the Internet.” Journal of Computer-Mediated Communication 12, no. 3 (2007): 1043-1062.

Campbell, Heidi A. “Understanding the Relationship between Religion Online and Offline in a Networked Society.” Journal of the American Academy of Religion 80, no. 1 (2012): 64-93.

Crawford, Kate. Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence. Yale University Press, 2021.

Gabriel, Iason. “Artificial Intelligence, Values, and Alignment.” Minds and Machines 30 (2020): 411-437.

Hollanek, Tomasz, and Katarzyna Nowaczyk-Basińska. “Griefbots, Deadbots, Postmortem Avatars: on Responsible Applications of Generative AI in the Digital Afterlife Industry.” Philosophy & Technology 37 (2024): 63.

International Energy Agency. Energy and AI. Paris: IEA, 2025.

Latzer, Michael. “Digitalization, AI and the rise of techno-religion: Transhumanist promises and the challenge to Enlightenment.” Telecommunications Policy 50, no. 2 (2026): 103115.

Meyer, Birgit. Mediation and the Genesis of Presence: Toward a Material Approach to Religion. Utrecht University, 2012.

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). 2023.

Noble, Safiya Umoja. Algorithms of Oppression: How Search Engines Reinforce Racism. NYU Press, 2018.

Schmemann, Alexander. For the Life of the World: Sacraments and Orthodoxy. St. Vladimir’s Seminary Press, 1973.

Simmerlein, Jonas. “Sacred Meets Synthetic: A Multi-Method Study on the First AI Church Service.” Review of Religious Research 67, no. 1 (2025): 126-145.

Simmerlein, Jonas, and Max Tretter. “What about spiritual needs? Care robotics and spiritual care.” Frontiers in Robotics and AI 11 (2024): 1455133.

Singler, Beth, and Fraser Watts, eds. The Cambridge Companion to Religion and Artificial Intelligence. Cambridge University Press, 2024.

Spitale, Giovanni, and Federico Germani. “The making of digital ghosts: designing ethical AI afterlives.” Ethics and Information Technology 28 (2026): 34.

Stanojevic, Dimitrije. Synthetic Levites: Generative AI and the Redistribution of Religious Authority. Neurobiblistics, 2026.

Turkle, Sherry. Alone Together: Why We Expect More from Technology and Less from Each Other. Revised/selected edition.

Tsuria, Ruth, and Yossi Tsuria. “Artificial Intelligence’s Understanding of Religion.” Religions 15, no. 3 (2024): 375.

UNESCO. Recommendation on the Ethics of Artificial Intelligence. 2021.

Vallor, Shannon. Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting. Oxford University Press, 2016.

Williams, James. Stand Out of Our Light: Freedom and Resistance in the Attention Economy. Cambridge University Press, 2018.

World Health Organization. Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models. Geneva: WHO, 2024.

Zhang, Yutong, et al. “Interaction with AI companions and psychological well-being.” Nature Human Behaviour (2026).

AI Pedagogy Project. Large Language Model Tutorial and responsible-use resources. Harvard metaLAB / Berkman Klein Center, current online resource.

Selected primary texts from Jewish, Christian, Islamic, Buddhist, and other traditions are supplied in contextualized excerpts appropriate to each session.

Accessibility, Integrity, and Inclusion

This course is designed for students from multiple disciplines, religious backgrounds, and nonreligious perspectives. No technical background or confessional commitment is assumed.

Comparative study requires precision. Students should not speak as though a single text, community, or spokesperson exhausts a religious tradition; internal disagreement and historical context are part of the analysis.

No student is required to participate in religious practice, disclose personal beliefs, or engage an AI system with intimate or vulnerable personal information. Equivalent non-personal and public-materials pathways are provided.

Accessibility accommodations and flexible pathways for approved course activities should be arranged as early as possible.

Academic integrity includes honesty about AI use. Disclosure, source verification, and responsibility for the final argument are part of the course method. Fabricated sources, invented field engagement, deceptive synthetic media, or undisclosed substitution of AI prose for student analysis violate the course’s trust.

WHO IS FORMING WHOM?

What happens when the tools we create become participants in forming the people we become?