

# Cover Sheet: Request 14290

## PHI 3XXX: Ethics, Data, and Technology

### Info

Process	Course New Ugrad/Pro
Status	Pending at PV - University Curriculum Committee (UCC)
Submitter	John Palmer palmerj@ufl.edu
Created	9/27/2019 2:04:22 PM
Updated	10/14/2019 7:13:10 PM
Description of request	This course exposes students to ethical issues related to data science, algorithmic decision-making, and artificial intelligence. Students will grapple with foundational concepts in ethics, economics, and policy-making. Beginning with a brief introduction to philosophical ethics, the then course pairs theoretical discussions of ethics, economics and policy-making with concrete issues in emerging technologies. Discussion topics include: cost-benefit analysis, risk, markets and market failures, economic valuations of technology, justice and fairness, and property rights. We will apply these concepts in assessing emerging technologies like autonomous weapons, big data policing algorithms, facial recognition cameras, and "creative" machine learning algorithms.

### Actions

Step	Status	Group	User	Comment	Updated
Department	Approved	CLAS - Philosophy 011615000	John Palmer		9/27/2019
No document changes					
College	Conditionally Approved	CLAS - College of Liberal Arts and Sciences	Joseph Spillane	The College Curriculum Committee conditionally approves this request, with the following changes required: 1) start the course description with "Addresses"; 2) please clarify the prerequisites; 3) remove "speaking" from the last bullet point in grading under the short paper	10/14/2019
3K E&T Consult Form.pdf					
Department	Approved	CLAS - Philosophy 011615000	John Palmer	Changes requested by CLAS Approver have all been made within the form	10/14/2019
No document changes					
College	Approved	CLAS - College of Liberal Arts and Sciences	Joseph Spillane		10/14/2019
No document changes					
University Curriculum Committee	Pending	PV - University Curriculum Committee (UCC)			10/14/2019
No document changes					
Statewide Course Numbering System					
No document changes					
Office of the Registrar					
No document changes					

Step	Status	Group	User	Comment	Updated
Student Academic Support System					
No document changes					
Catalog					
No document changes					
College Notified					
No document changes					

## Course|New for request 14290

### Info

**Request:** PHI 3XXX: Ethics, Data, and Technology

**Description of request:** This course exposes students to ethical issues related to data science, algorithmic decision-making, and artificial intelligence. Students will grapple with foundational concepts in ethics, economics, and policy-making. Beginning with a brief introduction to philosophical ethics, the then course pairs theoretical discussions of ethics, economics and policy-making with concrete issues in emerging technologies. Discussion topics include: cost-benefit analysis, risk, markets and market failures, economic valuations of technology, justice and fairness, and property rights. We will apply these concepts in assessing emerging technologies like autonomous weapons, big data policing algorithms, facial recognition cameras, and “creative” machine learning algorithms.

**Submitter:** John Palmer palmerj@ufl.edu

**Created:** 10/14/2019 1:51:35 PM

**Form version:** 3

### Responses

**Recommended Prefix** PHI

**Course Level** 3

**Course Number** XXX

**Category of Instruction** Intermediate

**Lab Code** None

**Course Title** Ethics, Data, and Technology

**Transcript Title** Ethics, Data & Tech

**Degree Type** Baccalaureate

**Delivery Method(s)** On-Campus

**Co-Listing** No

**Effective Term** Fall

**Effective Year** 2020

**Rotating Topic?** No

**Repeatable Credit?** No

**Amount of Credit** 3

**S/U Only?** No

**Contact Type** Regularly Scheduled

**Weekly Contact Hours** 3

**Course Description** Addresses ethical issues related to data science, algorithmic decision-making, and artificial intelligence. Pairs theoretical discussions of ethics, economics, and policy-making with concrete issues in emerging technologies.

**Prerequisites** Sophomore standing or higher standing or (PHI 2010(C) or PHI 2100(C) or PHI 2630(C)) or (philosophy major or philosophy minor) or data science major.

**Co-requisites** N/A

**Rationale and Placement in Curriculum** This course is designed in the first instance to serve as the required ethics course for the new undergraduate major in Data Science to be offered in CLAS. The Departments of Mathematics and Statistics are currently taking the lead in developing this new major, with the support of the UF Provost and the CLAS Dean. The Chairs of those departments have reviewed the syllabus of the proposed course and have indicated their approval. The course is also designed to be an appealing elective for Philosophy majors and minors. It should also be an attractive course for a range of students from other majors, both inside and outside of CLAS. More generally, the respond to the societal need for an informed citizenry able to grapple with the ethical issues arising in the present era of rapidly developing technology.

**Course Objectives** 1. Identify and explain the philosophical dimensions of foundational concepts in economics and public policy as they pertain to data science and its applications.

2. Develop a basic vocabulary for discussing the ethical dimensions of data science and its applications.
3. Analyze issues and policies concerning emerging "big data" technologies through the application of ethical concepts.
4. Critique public policies, social practices, and political-economic institutions that shape, and are shaped by, scientific discovery and technological design.
5. Discern the structure of arguments, representing them fairly and clearly, and evaluating them for cogency.
6. Formulate original arguments, anticipate objections, and respond in a conscientious fashion.
7. Read and discuss complex philosophical texts from both historical sources and contemporary works.
8. Speak and write clearly and persuasively about abstract and conceptually elusive matters.

**Course Textbook(s) and/or Other Assigned Reading** Since there is no altogether suitable textbook for this course, the readings indicated on the Weekly Schedule of Topics that follows will be made available to students through a Canvas course site.

**Weekly Schedule of Topics** COURSE SCHEDULE

Week 1

Introductions/Syllabus

Topic: Technology and Society

Reading: Winner, Langdon (1997). "Technologies as forms of life." In Kristin Shrader-Frechette & Laura Westra (eds.), *Technology and Values*. Rowman & Littlefield. pp. 55--69.

Topic: Technology: What is it?

Reading: Stanford Encyclopedia of Philosophy, "Philosophy of Technology" (Section 2.5)

Week 2

Topic: Technology and Society

Reading: Mesthene, Emmanuel G. (1997). "The role of technology in society." In Kristin Shrader-Frechette & Laura Westra (eds.), *Technology and Values*. Rowman & Littlefield. pp. 71--85.

Topic: Philosophy, Metaethics, Normative Ethics, and Applied Ethics

Reading: Shelley Kagan. *Normative Ethics*. Westview Press. 1997. (pp. 1 – 22)

Week 3

Topic Normative Ethics - Utilitarianism

Readings: Eggleston, Ben (ed.). *Utilitarianism (by John Stuart Mill): With Related Remarks from Mill's Other Writings*. Hackett Publishing Company, (2017), chapter 2 & Feldman, Fred. "Utilitarianism: Pro and Con" in *Introductory Ethics*. Prentice Hall (1978)

Topic: Technology and moral side constraints: The Case of Nuclear Weapons

Readings: Robert Nozick, "Moral Constraints and Moral Goals." In Darwall, Stephen L. (ed.) (2003). *Deontology*. Wiley-Blackwell. (2003), pp. 83-89 &

Anscombe, Elizabeth, "Mr. Truman's Degree," In *The Collected Philosophical Papers of G. E. M. Anscombe*, vol. III, *Ethics, Religion and Politics*. Blackwell (1981), 62-71.

Week 4

Topic: Property Rights

Readings: John Locke, "Of Property," from *Two Treatises of Government*, Book II & David Hume, "Of Justice and Property," from *Enquiries Concerning the Human Understanding and Concerning the Principles of Morals*.

Topic: Technology and Intellectual Property

Readings: Hettinger, Edwin C.. "Justifying intellectual property." In *Philosophy and Public Affairs* 18 (1989):31-52. & Adam Moore and Ken Himma, "Intellectual Property" (*Stanford Encyclopedia of Philosophy*)

Week 5

Case Study: Technology and Intellectual Property Rights

Readings: "These Stunning AI Tools are about to Change the Art World" (*Slate*) & "AI System 'should be recognized as inventor'" (*BBC News*) & "The AI-Art Gold Rush Is Here" (*The Atlantic*)

Topic: Transparent machine learning

Readings: Kilian Weinberger, Cornell: "Interpretable Machine Learning: What are the limits and is it necessary?" (<https://www.youtube.com/watch?v=w0KFH8aSfG4>) & "Making Intelligence Intelligible" (Interview)

Week 6

Topic: Transparent machine learning

Reading: Vredenburg, "Is there a right to explanation?" Manuscript.

Case Study: Intellectual property rights v. the right to explanation

Reading: Burrell, J. (2016). "How the machine 'thinks': Understanding opacity in machine learning algorithms." *Big Data & Society*. <https://doi.org/10.1177/2053951715622512>

Week 7

Topic: Public goods and market failures

Readings: Tyler Cowen, "Public Goods" (<https://www.econlib.org/library/Enc/PublicGoods.html>) & Jonathan Anomaly, "Public Goods and Government Action," *Politics, Philosophy, and Economics*. Sage Publishing, 2015

Topic: Technology, public goods, and government intervention

Readings: McGinn, Robert E. "Technology, Demography, and the Anachronism of Traditional Rights." *Journal of Applied Philosophy* 11 (1994):57-70.

Week 8

Case study: Facebook, democracy, and market failures

Readings: "Russian content on Facebook, Google and Twitter reached far more users than companies first disclosed" (WashingtonPost) & "Why Fake News Spreads like Wildfire on Facebook" (Chicago Tribune)

Topic: Risk - Small probabilities of great harm

Readings: Shrader-Frechette, Kristin. "Technological risk and small probabilities." *Journal of Business Ethics* 4 (1985):431 - 445. & "Managing Risks in Public Policy: Impact vs. Probability?" (Huffington Post)

Week 9

Topic: Risk - Cost-benefit analysis v. the precautionary principle

Readings: Sunstein, Cass R. "Cost-benefit analysis and the environment." *Ethics* 115 (2005):351-385. (pp. 351 – 354) & Gardiner, Stephen M. "A core precautionary principle." *Journal of Political Philosophy* 14 (2006): 33–60

Risk Case study: Superintelligence and existential risk

Reading: Nick Bostrom, *Superintelligence* (Oxford, 2014) pp. 140 – 149.

Week 10

Topic: Technology and Work

Readings: "Humans Need not Apply" (<https://www.youtube.com/watch?v=7Pq-S557XQU>) & Marx, *Capital*, ch. 15, sections 5 - 6 & Lenman, James. "On becoming redundant or what computers shouldn't do." *Journal of Applied Philosophy* 18 (2001):1–11.

Topic: Distributive Justice - Egalitarianism

Reading: Excerpts from John Rawls. *A Theory of Justice*. (Harvard Univ. Pr., 1999).

Week 11

Topic: Distributive Justice and the Right to Work

Readings: Kavka, Gregory S. "Disability and the Right to Work": *Social Philosophy and Policy* 9 (1992):262-290. & Nozick, Robert. "Distributive justice." *Philosophy and Public Affairs* 3 (1973): 45 – 60

Case Study: Artificial Intelligence and the Worker

Readings: "Robots could replace 1.7 million American truckers in the next decade" (LA Times) & Vardi, "What the industrial revolution really tells us about the future of automation and work" (The

Conversation) & Vardi, "Humans, Machines, and Work: The Future is Now" (<https://www.youtube.com/watch?v=5ThiCIGEBes&feature=youtu.be>)

#### Week 12

Topic: Killing by Machine - The Ethics of Autonomous Weapons

Readings: Strawser, Bradley Jay. "Moral Predators: The Duty to Employ Uninhabited Aerial Vehicles." *Journal of Military Ethics* 9 (2010):342-368. & "Attack of the Killer Robots" (Buzzfeed) & Purves, Duncan, Jenkins, Ryan & Strawser, Bradley J. "Autonomous Machines, Moral Judgment, and Acting for the Right Reasons." *Ethical Theory and Moral Practice* 18 (2015):851-872.

#### Week 13

Topic: Big data, surveillance and privacy

Reading: Screening of "Citizenfour"

Topic: Surveillance and big data policing

Reading: Doyle, Tony. "Privacy and perfect voyeurism." *Ethics and Information Technology* 11 (2009):181-189. & Rachels, James (1975). "Why privacy is important." *Philosophy and Public Affairs* 4 (4):323-333. & "House passes NSA spying bill after Trump tweets cause confusion" (Reuters)

#### Week 14

Topic: Why does Privacy Matter?

Readings: Calo, Ryan, "The Boundaries of Privacy Harm." *Indiana Law Journal*, Vol. 86, No. 3, 2011. Available at SSRN: <https://ssrn.com/abstract=1641487> & Moxie Marlinspike, "Why 'I've got nothing to hide' is the Wrong Way to Think about Surveillance (wired.com)

Case Study: Surreptitious Surveillance on the Internet

#### Week 15

Topic: Fair Machine Learning

Reading: Binns, R. "Fairness in Machine Learning: Lessons from Political Philosophy" *Proceedings of Machine Learning Research* 81 (2018):1–11. & Ferguson, A. (2017). *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. New York: NYU Press. Retrieved from <http://www.jstor.org/stable/j.ctt1pwtb27> (86 – 101 and 62 – 83).

#### Week 16

Topic: Fair Machine Learning

Reading: Kleinberg lecture: <https://www.youtube.com/watch?v=4X3Z7FPwA8> & Kleinberg, J., Mullainathan, S., Raghavan, M. (2016). "Inherent trade-offs in the fair determination of risk scores." (arXiv preprint arXiv:1609.05807) & "AI programs exhibit racial and gender biases, research reveals" (The Guardian) & Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, Noah A. Smith. (2019). *The Risk of Racial Bias in Hate Speech Detection*. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678

**Grading Scheme** Emerging technology case studies (500 pts): We will have six or seven in-class discussions of case studies concerning emerging technologies. Students will be assigned to discussion groups. Each group member will provide a written analysis of the case study and will use this analysis to inform the discussion. Each group will submit the "best" written analysis at the end of the discussion. Grades for the case studies will be determined by (a) participation in the case study discussion and (b) the quality of the group's submitted written analysis. See the Course Schedule for the dates of the case studies.

Short Paper (200 pts): (1200 – 1500 words): This short paper involves a critical assessment of a topic that we have discussed in class. I will give you a choice of two topics for this paper; I will also provide the structure that the paper must follow. One of the topics will be theoretical, addressing an argument for a particular position in ethical theory. The other topic will involve applying ethical concepts in order to assess a specific technology.

This paper assesses whether students have acquired:

- skill in discerning the structure of arguments, representing them fairly and clearly, and evaluating them for cogency.
- skill in formulating original arguments, anticipating objections, and responding in a

conscientious fashion.

- skill in reading and discussing complex philosophical texts from both historical sources and contemporary works
- and skill in writing clearly and persuasively about abstract and conceptually elusive matters.

Short reading quizzes (200 pts): There will be irregular reading quizzes, designed to make sure that you are doing and understanding the readings, and to double-check that you are coming to class.

Class attendance and participation (100 pts): A good philosophy course requires good discussion. Good discussion requires regular attendance and healthy conduct. Therefore, attendance and conduct round out the remaining 10% of the course grade. Every unexcused, recorded absence beyond the first results in a 10% deduction from the participation/conduct grade (So, that's 10% of 10% of your course grade). Students may be removed from the class roster for frequent absences. A 10% deduction may also occur for transgressions such as disruptive behavior, falling asleep, sending text messages, or surfing the internet.

#### Grade Scale

Grade Value

100-93=A

A=4.0

92-90=A-

A-=3.67

89-86=B+

B+=3.33

85-82=B

B=3.00

81-79=B-

B-=2.67

78-76=C+

C+=2.33

75-72=C

C=2.00

71-69=C-

C-=1.67

68-66=D+

D+=1.33

65-62=D

D=1.00

61-60=D-

D-=0.67

59-0=E

E=0.00

**Instructor(s)** Dr. Duncan Purves (Philosophy)

Dr. Amber Ross (Philosophy)

**Attendance & Make-up** Yes

**Accommodations** Yes

**UF Grading Policies for assigning Grade Points** Yes

**Course Evaluation Policy** Yes

External Consultation Results (departments with potential overlap or interest in proposed course, if any)

Department	Name and Title
_____	_____
Phone Number	E-mail
_____	_____
Comments	

Department	Name and Title
_____	_____
Phone Number	E-mail
_____	_____
Comments	

Department	Name and Title
_____	_____
Phone Number	E-mail
_____	_____
Comments	