

**ECON 4444 Artificial Intelligence and the Future of Work**  
**Anton Korinek, UVA Economics and Darden School of Business**

**Spring 2020 (Preliminary Syllabus)**

**(Preliminary Version – will be updated soon)**

*Personal note: I am offering this course because I believe that artificial intelligence will change our world more profoundly than we can imagine in coming decades. This course invites you to join me on a journey from our evolutionary origins to present technological disruptions and their effects on the economy towards the future of life in our universe.*

**Course webpage:** UVA Collab

**Lecture:** Mon and Wed 2 – 3:15pm in Gilmer Hall 130

**Office hours:** tbd

**Teaching Assistant:** Miguel Mascarua

**Discussion sessions:** tbd

**TA office hours:** tbd

**Course description:**

Advances in artificial intelligence and automation have proceeded rapidly in recent years and have reached an inflection point that will have profound implications for the future of humanity.

How does rapid progress in automation and artificial intelligence affect our economy and society? This course analyzes the short- and medium-run implications for employment, economic growth, and inequality. It also studies the long-run implications of artificial intelligence rivaling human intelligence.

**Learning goals:** At the conclusion of this course, you will be able to:

- assess the implications of the rise of AI for humanity
- evaluate how different sectors of the economy will be affected by automation
- demonstrate your understanding in an in-depth course paper on a specific topic in the field
- devise a plan for how to effectively manage your career and life in the age of AI

**Grade composition:**

24 % Class preparation and participation

20 % Problem sets

20 % Mid-term quizzes

18 % Course paper

18 % Course presentations

No opportunities for extra credit.

If you cannot deliver an assignment on time, please contact the TA as soon as possible to discuss how to manage the situation. Little can be done after an unsatisfactory grade has been assigned.

**Policy on Use of Electronics:**

The use of laptops and other electronics in the classroom is forbidden because research shows that they interfere with effective learning (see e.g. <https://nyti.ms/2hVMrVo>).

If you need to use your smartphone, please step outside the classroom – you are welcome to do so for urgent matters and do not need to excuse yourself.

**Class Preparation, Attendance, and Participation:**

Class attendance and participation are an essential part of your learning experience. Each week, I will assign preparatory readings that we will discuss in the following week. Preparing these materials will help you to actively and effectively participate in classroom discussions, including in your response to cold calls. This will count for a significant part of your grade.

Regardless of your comfort level with the materials, there are many effective ways to participate in the classroom, including asking good questions or articulating ideas and insights that help others understand better. Here are some suggestions for effective classroom participation:

- The best contributions are well-formulated analyses or questions that relate directly to the preparatory readings. Look for ways to discuss the questions at hand that help others understand the material and the fundamental concepts behind it.
- If your own understanding of the concepts is a bit murky, you can still participate by presenting some of the key facts in the preparatory readings.
- If engaging in discussion is natural for you, think about ways to make your comments precise, on topic, and helpful for fellow students.
- If you are normally reluctant to engage in classroom discussion, try to push yourself a bit, or discuss with me how to best include you.

**Course Presentations and Papers:**

Part of the course will feature student team presentations, accompanied by papers, on how different sectors and/or areas of the economy – for example, the area in which you are planning to spend your career – are being revolutionized by AI. We can discuss topics in my office hours.

Please start forming teams of 4 or larger (maximum of 8) with classmates who have similar interests. The following deadlines will guide you along the process of preparing your presentation. You have a choice between two date ranges for your presentation (early or late) – please let me know by Jan 31<sup>st</sup>:

Due date for selection of <b>team members</b> and <b>topic</b>	<b>tbd</b>	
Due date for draft list of <b>source materials</b>		
<b>Choice</b> between two date ranges		
Due date for <b>draft slides</b>		
<b>Office hours</b> to discuss slides		
Date of <b>presentations</b> (5 min per speaker)		
Due date for <b>course paper</b>		

**Components of presentation grade:**

- 8 pt: Depth of preparation (explanation of mechanisms, multitude of sources)
- 4 pt: Quality of slides and preparedness
- 3 pt: Coordination among team members
- 3 pt: Professionalism and Timeliness of Materials and Presentation (stay within time limit)

**Course Paper:**

Seven days after your presentation, each team is expected to submit a paper on the topic presented. The paper should be about 4000 words long for a team of four (about 8 pages using standard formatting), and proportionately longer for larger teams. I will particularly value original and actionable proposals for how to prepare workers in your sector of analysis for a successful career in the age of AI.

Any late submission will lose 25% of the points if submitted within the first 24 hrs after the deadline, 50% if submitted within 48 hrs, 75% if submitted within 72 hours, and will receive a zero grade if submitted later than that.

**Problem Sets:**

There will be five short problem sets, which will help you to digest the materials covered in class and will be useful in preparing for the quizzes.

You are allowed to collaborate in teams (up to 4 students). If you do so, then please list at the beginning of the problem set who the other team members were. Nonetheless, you are responsible for handing in your own assignment and will be graded solely on what you hand in yourself.

**Quizzes:**

There will be two quizzes on [tbd] and [tbd]. You will be responsible for all the materials covered in class, including the preparatory readings. The materials will be cumulative.

**Medical Excuses and Religious Holidays:**

If you will have to miss a class for an excused reason as per university policy (religious holiday or illness), please email our TA (and CC me) as soon as you find out about the circumstance that will prevent you from attending in order to make sure that you do not lose participation points.

**Outline of course topics (may be adjusted over the course of the semester):**

Note: specific reading assignments will be distributed in the week prior to each class

Introduction (Week 1):

- Syllabus and course requirements
- Determining the course topics
- Asking the big questions:
  - Will AI/robots take our jobs?
  - Will they replace humanity?
  - What can we do about it?

Part 1: Intelligence (Weeks 2 & 3):

- What is intelligence?
  - agents and their goals, take 1
  - substrates
- Evolution of biological intelligence
  - history
  - benefits and costs
  - emotions
  - empathy and cooperation
  - language and human culture
- (\*) Evolution of artificial intelligence:
  - history of AI: intelligent design of intelligence
  - a non-technical primer on AI
  - recent developments
  - co-evolution of man and AI
  - Scenarios for the near future
- Substrate independence
  - computation and consciousness
- Comparing humans and machines
- Potential scenarios for the far future: singularity and intelligence explosion

Part 2: Labor Markets and Public Policy (Weeks 4, 5 and 6):

- Mankind, skills, specialization and labor
- Pre-industrial revolution:
  - subsistence and the reign of Malthus
- Industrial revolution:
  - technological progress and automation
  - job destruction, job creation
- (\*) Technological change and labor
  - factor bias
  - unemployment or falling wages?
  - worker-replacing progress and economic singularity

- Recent labor market developments:
  - decline in labor share
  - job polarization
  - labor vs human capital
- Digitization and superstars in the information economy:
  - increasing returns
  - value of data
- (\*) Inequality:
  - technological forces
  - political feedback loops
  - subsistence and the return of Malthus?
- (\*) Public policy:
  - UBI and other benefits
  - redistribution vs predistribution
  - what shall we tax? labor vs capital; robots vs scarce factors?
  - steering technological progress
- Social attitudes:
  - deriving meaning & dignity from work
  - are concerns about inequality outdated?

#### Part 3: Beyond Human Biology (Weeks 7 & 8):

- Humanism:
  - anthropocentrism and our most cherished beliefs
- Engineering better humans:
  - history of enhancing intelligence
  - genetics
  - transhumanism
  - homo deus versus the useless?
- The future of life
- (\*) Scenarios for superintelligence and role for humanity
  - enhanced humans
  - whole brain simulation
  - brain-computer interfaces
  - artificial superintelligence
- (\*) Potential outcomes: existential catastrophe or eternal life?
  - the great filter
  - friendly AI?

#### Part 4: Agency & Goals (Weeks 9, 10 and 11):

- Agents and their goals, take 2:
  - modeling the world & the construct of agency
  - orthogonality hypothesis
  - agency problems

- Final goals and instrumental goals:
  - instrumental convergence hypothesis
  - representation of goals
  - natural selection of goals
- Thermodynamics and goals
- Human goals:
  - from biological evolution: single- and multi-cellular perspectives
  - from cultural evolution
  - agency problems
  - spirituality and the meaning of life
  - What are your goals?
- AI goals:
  - programmed goals vs AI drives
  - representation problems
  - the control program: an evolutionary perspective
  - the control problem: an agency perspective
- Artificially Intelligent Agents (AIAs):
  - evolution of AIAs
  - distribution of resources between humans and AIAs
  - signs of AIAs in our world
- What can we do about it?

Part 5: Student group presentations (Weeks 12 & 13):

- see detailed description above

Conclusion: Preparing for the Age of AI (Week 14):

- Why this is the most exciting time to be alive in history
- Reaping short-term gains from AI
- The bigger picture and the future of life
- What can you do?

**Books:**

Barrat, James (2013), *Our Final Invention: Artificial Intelligence and the End of the Human Era*, Thomas Dunne Books.

(\*)Bostrom, Nick (2014), *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press.

(\*)Brynjolfsson, Erik and Andrew McAfee (2015), *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W.W. Norton.

Domingos, Pedro (2015), *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*, Basic Books.

Gordon, Robert (2016), *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*, Princeton University Press.

Ford, Martin (2015), *Rise of the Robots: Technology and the Threat of a Jobless Future*, Basic Books.

Harari, Yuval Noah (2017), *Homo Deus: A Brief History of Tomorrow*, Harper.

Kaplan, Jerry (2015), *Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence*, Yale University Press.

Kurzweil, Ray (2005), *The Singularity Is Near: When Humans Transcend Biology*, Viking.

Lee, Kai-Fu (2018), *AI superpowers: China, Silicon Valley, and the new world order*, Houghton Mifflin.

Moravec, Hans (1988), *Mind Children: The Future of Robot and Human Intelligence*, Harvard University Press.

(\*)Tegmark, Max (2017), *Life 3.0: Being Human in the Age of Artificial Intelligence*, Knopf.

**Academic Articles:**

Autor, David (2014), "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives* 29(3), pp. 3-30.

Berg, Andrew, Edward F. Buffie, and Luis-Felipe Zanna (2017), "Robots, Growth, and Inequality: Should We Fear the Robot Revolution? (The Correct Answer is Yes)," forthcoming, *IMF Working Paper*.

Freeman, Richard B. (2015), "Who Owns the Robots Rules the World," *IZA World of Labor*.

(\*)Frey, Carl Benedikt, and Michael A. Osborne (2013), "The Future of Employment: How Susceptible Are Jobs to Computerisation?" Oxford University paper.

Keynes, John Maynard (1930), "Economic Possibilities for our Grandchildren," in *Essays in Persuasion*, Harcourt Brace, pp. 358-373.

(\*)Korinek, Anton and Joseph Stiglitz (2019), "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," in Agrawal et al.: *The Economics of Artificial Intelligence*, NBER and University of Chicago Press, May 2019.

(\*)Alan M. Turing (1950), "Computing Machinery and Intelligence," *Mind* 59(236), pp. 433-460.