6.S976 Engineering Leadership in the Age of AI, Fall 2019
Monday and Wednesday, 12:30-2:00 p.m., Room: 4-231

Mr. David Martinez, Associate Division Head, MIT Lincoln Laboratory
Contact: 781-981-7505, dmartinez@ll.mit.edu
Twitter handle for course: @ELAAI16
Office and hours: 35-433C, Office Hours: Monday 10 am-12 pm and Wednesday 2-3 pm or by appointment

Dr. David Niño, Senior Lecturer, Gordon-MIT Engineering Leadership Program
Contact: 617-324-4677, dnino@mit.edu

Overview

Artificial intelligence (AI) will continue to revolutionize many industries, for example, driverless cars, finance, national security, medicine, e-commerce, to name a few. President Rafael Reif stated in a recent article, "To prepare society for the demands of the future, institutions must equip tomorrow's leaders to be 'AI bilingual'. Students in every field will need to be fluent in AI strategies to advance their own work. And technologists will need equal fluency in the cultural values and ethical principles that should ground and govern the use of these tools." [1] This course will equip MIT graduate engineering students to lead, develop, and deploy AI systems in ways that augment human's capabilities, while providing positive impact to society. It will teach at the intersection of engineering leadership and artificial intelligence, and deliver learning experiences in ways that are informed by research but applied to engineering products and/or services.

The course will begin with a brief AI history, including highlights of representative successes in the application of AI. The course will be project-based requiring the students to formulate a strategic roadmap of an innovative AI application. Several key engineering leadership principles will serve as the foundation for developing an AI application roadmap. These principles will include establishing a strategic vision, identifying candidate customers, project execution, and the building of teams with complementary strengths. The students will be required to present the strategic vision and roadmap proposal to a selected industry/academic panel with a broad range of expertise. The final presentation will require a balance between technical depth and breath, as well as emphasis on strategic vision and uniqueness of the AI application. The course will also include invited speakers from industry that are practitioners in the development and deployment of AI applications. At the completion of this course, the students will have the necessary skills to lead AI teams. This course is offered through the Gordon-MIT Engineering Leadership Program.

Course Units: 12 credits, 3-0-9

[1] L. Rafael Reif, February 10, 2019, "Prepare students for a future of artificial intelligence," *Financial Times*.

Recommended Readings *(A subset will be required and identified during class)*

1. David Martinez and David Nino, "Engineering Leadership in the Age of AI," Lecture Notes.
2. David Martinez, et.al., "High Performance Embedded Computing Handbook: A Systems Perspective," CRC Press, 2008, selected chapters.
3. Erik Brynjolfsson and Andrew McAfee (2016), "The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies," W.W. Norton & Company.
4. Paul Daugherty and James Wilson (2018), "Human + Machine: Reimagining Work in the Age of AI," Harvard Business Press.
5. John Doerr (2018), "Measure What Matters: Objectives and Key Results," Portfolio Penguin.
6. Key papers on AI system architecture subcomponents will be recommended throughout the course.
7. Key papers on Engineering Leadership will be recommended throughout the course.
8. Readings will also be posted to Stellar.

Learning Objectives

At the completion of this course, students will be able to lead AI teams based on four core competencies:

- Understanding an end-to-end AI architecture at the system engineering level
- Applying engineering leadership principles
- Creating a strategic vision and development plan focused on a product or service
- Developing an execution strategy staffed by a diverse and multidisciplinary team
- Demonstrating an AI conceptual design based on an industry challenge problem using the Raspberry Pi

Exams

There will be three exams consisting of multiple-choice and true/false questions. These will be administered in class and will be based on assigned readings, lectures, and class discussions. These will be open-book, open-note exams. **Final grade will be based on best 2 exams out of three.**

Make-up exams will be administered only under circumstances involving serious and documented illnesses, or other excused absences. If such circumstances occur, students must inform us immediately in order to coordinate a make-up exam. Documentation for excused absences will be required.

Team Projects and Presentations

The purpose of these projects is to help students develop a practical understanding of how to apply course concepts to real, client-based problems. Small groups will be formed early in the semester to focus on specific AI application problems from two companies (one entrepreneurial company and one large company). We will provide some time in class for students to work on together via case studies but we also expect students to work together outside of class. Teams will prepare two types of presentations; one to class peers for feedback purposes, and a final presentation to a panel of industry and academic experts.

There will be three deliverables for the class project, at the end of the semester, as part of the experiential learning element of the course:

1. An AI strategic roadmap addressing one of the two companies challenge problems. This strategic roadmap will stress all elements of engineering leadership including project management. Content of the strategic roadmap will be highlighted in class.

2. An AI conceptual design demo showing key capabilities of your team solution to the company challenge problem. This part will address project development suitable for a 1 semester class.

3. A document that will be handed over to one of the two companies at the conclusion of the class containing: a) hard copies of slides for both the strategic roadmap and the AI conceptual design demo, b) a jupyter notebook file containing the procedure you use to execute the demo.

Class Participation

Students will be expected to actively contribute to class discussions and to honor the group project obligations. Class participation grades are based on participation, attendance, and results of a final peer and industry/academic panel presentations. Additional information about grading will be provided in class and final grades for class participation will be determined after the last day of class.

Areas Covered in the Course are Based on the Following Framework:

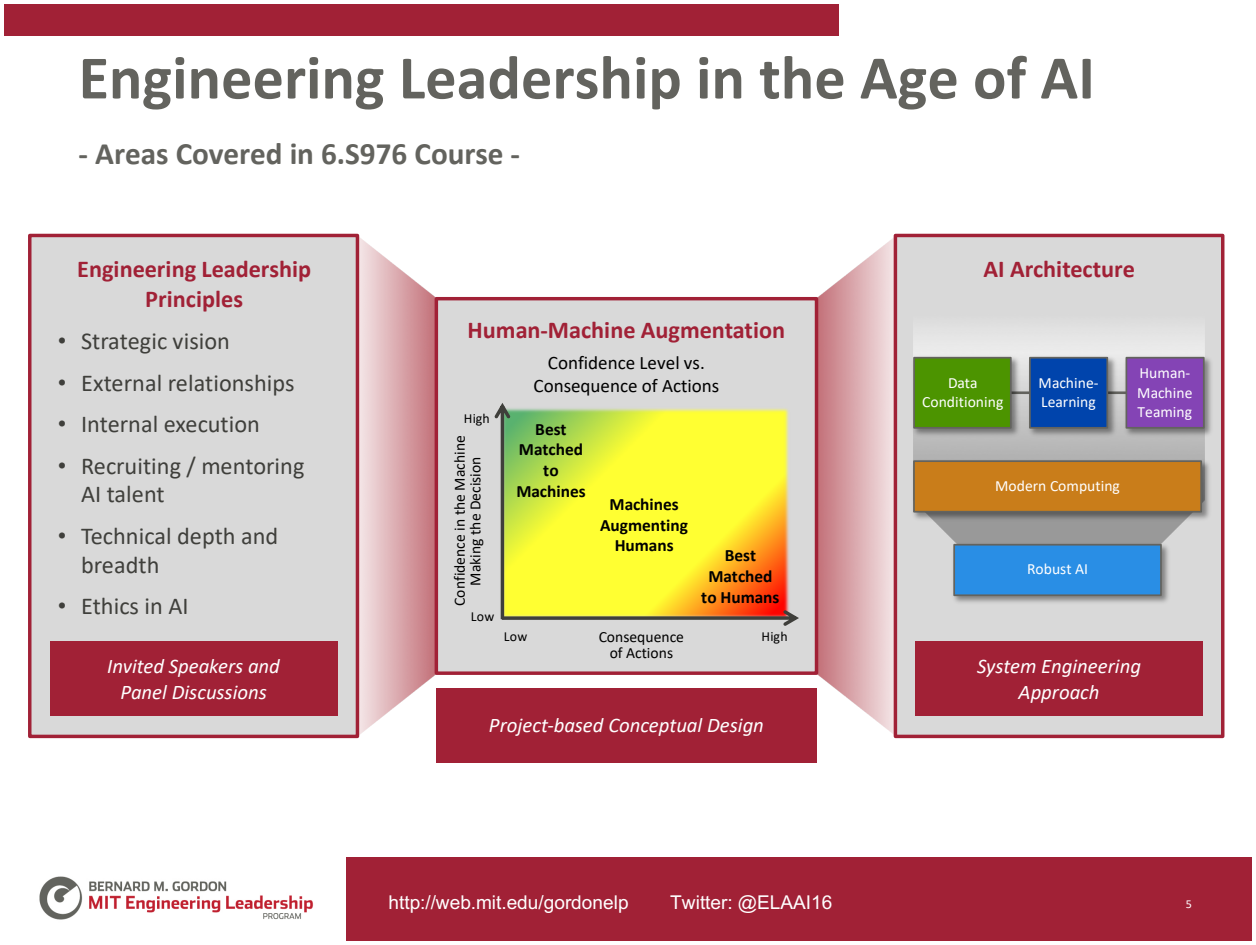

The AI Canonical Architecture is Central to the Course Content:

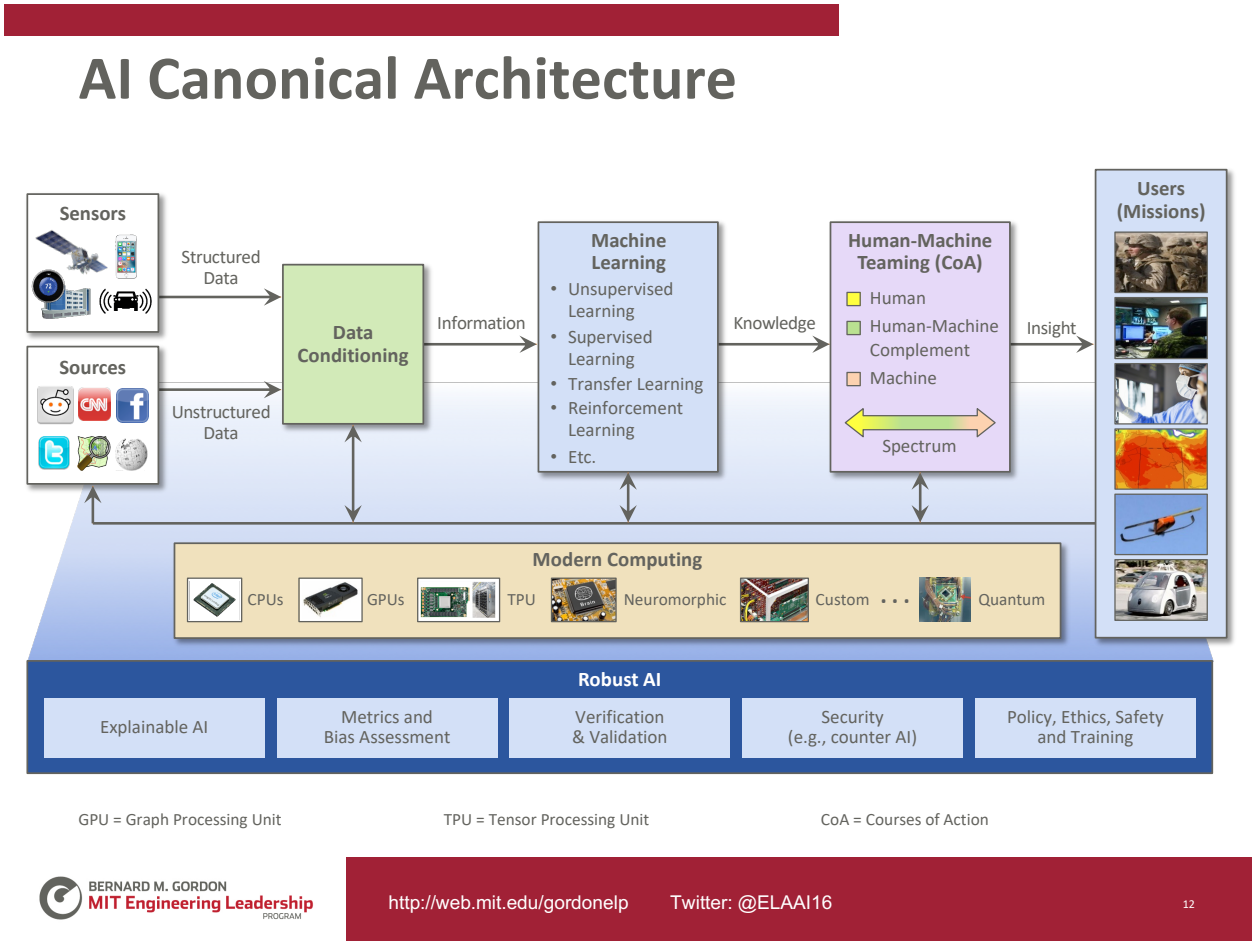

Final Grading Breakdown:

Individual contributions based on three exams (best 2 out of 3)	40%
Class participation	20%
Final team project and presentations	40%
Total	100%

Major Assignments, Exams, and Dates

1st Exam focused on fundamentals of AI architecture		September 23rd
2nd Exam focused on Lectures 6-10 (see Lecture #s below)		October 16th
3rd Exam focused on Lectures 11-15 (see Lecture #s below)		November 25th
Note: best 2 out of 3 exams will count towards grade	40%	
Class participation	20%	Presentation of strategic roadmap and AI conceptual design for 1 of the 2 challenge problems to peers
Final presentation of AI project	40%	December 2nd and December 4th presentations to industry/academic panel
Total	100%	

Lecture #	Class Date	Topic	Description	Major Events Dates
1	4 Sept	**Introduction and engineering principles**	**Overview course content and present key AI accomplishments plus introduction to engineering principles**	
2	9 Sept	**AI canonical architecture**	**System engineering approach to developing an AI product or service**	
3	11 Sept	**Data conditioning**	**Unstructured and structured data**	**Misty Robotics invited speaker**
4	16 Sept	**Machine learning algorithm taxonomy**	**Classes of ML algorithms and examples of successful implementations**	**Start identifying teams to work on class project**
5	18 Sept	**Modern computing as enabling technology**	**Types of computing engines and key computational drivers. Review of key concepts prior to exam**	
	23 Sept	**Exam focused on subsystem components of AI architecture**	**Exam addressing data conditioning, ML algorithm taxonomy and modern computing**	**1st Exam**
6	25 Sept	**Robust AI**	**Key issues in making AI systems robust. Confidence in a machine making a decision vs. consequence of action**	
7	30 Sept	**Human-machine teaming**	**Human-machine teaming as an enabler to augmenting human capabilities**	**Bose Wellness Division invited speaker**
8	2 Oct	**Overview of engineering leadership principles**	**Strategic development model and development plan for a product or service Part 1**	
9	7 Oct	**Tools and techniques in formulating a strategic vision**	**Internal and external factors in formulating a strategic vision Part 2**	**Finalize class teams**
10	9 Oct	**External relationships**	**Forming teams and understanding customer needs. Review of key concepts prior to exam**	
	16 Oct	**Exam focused on additional subsystem components of AI architecture and engineering leadership principles**	**Exam addressing additional AI architecture subcomponents and material covered so far on engineering leadership principles**	**2nd Exam on Lectures 6-10**
11	21 Oct	**Internal execution and managing-by-results**	**Measuring what matters and fostering a culture of innovation**	**RPi demo by Christos Samolis**

12	23 Oct	**Tools and techniques in formulating a strategic development plan**	**Show examples of strategic development plans integrating engineering leadership principles**	**RPi demo by Bruke Kifle**
13	28 Oct	**Leading from a position of strength in technical depth and breath**	**Discussion on balance between technical depth and breath**	**Invited AI practitioners as panel members**
14	30 Oct	**Effective recruiting and mentoring AI talent**	**Importance of diverse and multidisciplinary AI teams. Personal board of advisors**	
15	4 Nov	**Ethics in AI**	**Discussion on ethics in AI**	
	6 Nov	**Peer presentation by the teams on strategic roadmap**	**Misty Robotics teams present strategic roadmap <u>except</u> for conceptual design demo**	**Presentation by the teams (25 min/team)**
	13 Nov	**Peer presentation by the teams on strategic roadmap**	**Bose teams present strategic roadmap <u>except</u> for conceptual design demo**	**Presentation by the teams (25 min/team)**
	18 Nov	**Peer presentation of AI conceptual design for Bose challenge problems**	**Presentation by Bose teams of AI conceptual design and receive feedback**	**Presentation by Bose teams (20 min/team)**
	20 Nov	**Peer presentation of AI conceptual design for Misty Robotics challenge problems**	**Presentation by Misty Robotics of AI conceptual design and receive feedback. Also review of key concepts prior to exam**	**Presentation by Misty Robotics teams (20 min/team)**
	25 Nov	**Exam focused on last set of lectures 11-15**	**Applying engineering leadership principles to the development of an AI products or service**	**3rd Exam on Lectures 11-15**
	27 Nov	**No class Lecture. Teams preparation for final presentations**	**Teams use this time to prepare final presentations**	
	2 Dec	**Final presentation of AI conceptual design**	**Presentation of AI conceptual design by Bose teams to industry/academic panel (~20 min/team)**	**Invite industry/academic panel**
	4 Dec	**Final presentation of AI conceptual design**	**Presentation of AI conceptual design by Misty Robotics teams to industry/academic panel (~20 min/team)**	**Invite industry/academic panel**
16	9 Dec	**Review of core competencies (Part I)**	**Review course takeaways: end-to-end AI architecture, engineering leadership principles, and team projects (Part I)**	

17	11 Dec	**Review of core competencies (Part II)**	**Review course takeaways: strategic vision, development plan, execution strategy, and team projects (Part II)**	

Supplemental Readings on AI Canonical Architecture (required reading will be identified during class)

1. J. Sterman, "System Dynamics at Sixty: The Path Forward," 2018 System Dynamics Society. *https://onlinelibrary.wiley.com/doi/10.1002/sdr.1601*
2. N. J. Nilsson, "The Quest for Artificial Intelligence," Cambridge University Press, 2009.
3. W. Moore, et.al., "The AI Stack: a Blueprint for Developing and Deploying Artificial Intelligence," 2018.
4. George Heilmeier's catechism: Successful Technical Proposals, Cecilia M. Elliott, University of Illinois at Urban-Champaign, 2013.
5. Moll, Zalewski, Pi8llai, Madden, Stonebraker, Gadepally, "Exploring big volume sensor data with vroom", VLDB'17.
6. *Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): A public-access intensive care unit database*, Saeed, et al., Crit Care Med. 2011.
7. Pedro Domingos, "The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World," Basic Books, February 2018.
8. Deep Learning, Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, Nature, vol. 521, 28 May 2015.
9. Gradient-Based Learning Applied to Document Recognition, Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, Proc. Of the IEEE, November 1998.
10. A. Krizhevzky, I. Sutskever, and G. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," NeuroIPS 25, pp. 1106-1114, 2012.
11. B. Borgstrom, M. Brandstein, and R. Dunn, "Improving Statistical Model-based Speech Enhancement with Deep Neural Networks", 2018 IWAENC.
12. Williams, A. Waterman, and D. Patterson, "Roofline: An Insightful Visual Performance Model for Multicore Architectures", Communications of the ACM vol. 52 no. 4 (Apr 2009).
13. N. Jouppi, C. Young, N. Patil, and D. Patterson, "A Domain-Specific Architecture for Deep Neural Networks", Communications of the ACM vol. 61 no. 9 (Sep 2018).
14. V. Sze, Y-H Chen, T-J Yang, J. Emer, „ Efficient Processing of Deep Neural Networks: A Tutorial and Survey," Proceedings of the IEEE, vol. 105, no. 12, 2017.
15. K. Simonyan, A. Vedaldi, A. Zisserman, "Deep Inside Convolutional Networks: Visualizing Image Classification Models and Saliency Maps," arXiv: 1312.6034v2, April 2014.
16. J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson, "Understanding Neural Networks Through Deep Visualization," Deep Learning Workshop, ICML, 2015.
17. Evtimov, et.al., "Robust Physical-World Attacks on Deep Learning Models," arXiv:1707.08945, 2017.
18. Goodfellow, P. McDaniel, and N. Papernot, "Making Machine Learning Robust Against Adversarial Inputs," ACM, 2018.
19. G. Kasparov, "Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins," 2017, PublicAffairs.

Supplemental Readings on Engineering Leadership (required reading will be identified during class)

20. T. Chamorro-Premuzic, M. Wade, J., Jordan, "As AI Makes More Decisions, the Nature of Leadership Will Change," HBP, 2018.
21. David Nadler and Mike Tushman, "A Model for Diagnosing Organizational Behavior," Organizational

Dynamics, Autumn 1980.
22. P. Lencioni, “The Five Dysfunctions of a Team: A Leadership Fable,” Jossey-Bass, 2002.
23. P. Schwartz, “The Art of the Long View: Planning for the Future in an Uncertain World,” Crown Business, 1996.
24. Y. Shen, R. Cotton, and K. Kram, “Assembling Your Personal Board of Advisors,” MIT Sloan Management Review, 2015.
25. N. Craig and S. Nook, “From Purpose to Impact,” HBR, 2014.
26. D. Ancona, “In Praise of the Incomplete Leader,” HBR, 2007.
27. C.K. Prahalad and G. Hamel, “The Core Competence of the Corporation,” HBR, 1990.
28. R. Kaplan, “What to Ask of the Person in the Mirror,” HBR, 2007.
29. J. Collins and J. Porras, “Building Your Company Vision,” HBR, 1996.
30. P. Williams, “Coach Wooden: The 7 Principles That Shaped His Life and Will Change Yours,” Revell, 2011.
31. J. Rao and J. Weintraub, “How Innovative is Your Company Culture,” MIT Sloan Management Review, 2013.
32. P. Drucker, “Management and the World’s Work,” HBR, 1989.
33. J. Zenger, J. Folkman, and S. Edinger, “Making Yourself Indispensable,” HBR, 2011.
34. D. Garvin, “How Google Sold Its Engineers on Management,” HBR, 2013.
35. W. Deresiewicz, “The American Scholar: Solitude and Leadership,” Lectyre at U.S. Military Academy at West Point Delivered in October 2009.
36. HBR Guide to Project Management. Harvard Business School Publishing, 2013.
37. D. Ancona and H. Gregersen, “Problem-led Leadership: An MIT Style of Leadership”, MIT Leadership Center White Paper, 2017.
38. B.J. Avolio, F.O. Walumbwa, T.J. Weber, “Leadership: Current theories, research, and future directions”, Annual review of Psychology, 2009.
39. R. Kramer, “Trust and distrust in organizations: Emerging perspectives, enduring questions”, Annual Review of Psychology, Vol. 50, Pages 569-598, February 1999.
40. J. Beyer and D. Niño, “Ethics and Cultures in International Business”, Journal of Management Inquiry, Vol. 8, pp. 287-297, 1999.
41. M.E. Brown and L.K. Treviño, “Ethical leadership: A review and future directions”, The leadership quarterly, 2006.
42. L.K. Treviño and M.E. Brown, “Managing to be ethical: Debunking five business ethics myths”, Academy of Management Perspectives, 2004.
43. D. Vaughan, “The Dark Side of Organizations: Mistake, Misconduct, and Disaster”, Annual Review of Sociology, Vol. 25., pp. 271-305, 1999.
44. W.T. Lynch and R. Kline, “Engineering practice and engineering ethics”, Science, technology, & human values, 2000.

Academic Integrity

MIT's Academic Integrity policy reads, in part: "Fundamental to the academic work you do at MIT is an expectation that you will make choices that reflect integrity and responsible behavior. MIT will ask much of you. Occasionally, you may feel overwhelmed by the amount of work you need to accomplish. You may be short of time, working on several assignments due the same day, or preparing for qualifying exams or your thesis presentation. The pressure can be intense. On the "Working Under Pressure" page, we suggest resources to help you manage your workload and prevent yourself from becoming overwhelmed. However, no matter what level of stress you may find yourself under, MIT expects you to approach your work with honesty and integrity." (see more information at http://integrity.mit.edu/). We expect to uphold these standards in our class and this is essential to your personal learning and to our ability to assess learning. Violating the Academic Integrity policy in any way (e.g., plagiarism, unauthorized collaboration, cheating, etc.) will result in official Institute sanction and please let us know if you have any questions about these issues. (Adapted from guidelines recommended by the MIT Teaching and Learning Lab at http://tll.mit.edu/).

Students with Disabilities

If you need disability-related accommodations, please let us know early in the semester. If you have not yet been approved for accommodations, contact the Student Disability Services at sds-all@mit.edu for more information. We look forward to working with you to assist with approved accommodations.

Inclusive Learning

MIT and GEL values an inclusive learning environment. We aim to create a sense of community in our class; a place where everyone will be treated with respect and where everyone feels safe to share their ideas and perspectives. We welcome individuals of all backgrounds, beliefs, ethnicities, national origins, gender identities, sexual orientations, religious and political affiliations – and other visible and nonvisible differences. Participants are expected to contribute to an open and inclusive environment for every other participant. If you feel this standard is violated or moving in this negative direction, please let us know (Adapted from guidelines recommended by the MIT Teaching and Learning Lab at http://tll.mit.edu/).