CSCI-461: Artificial Intelligence for Sustainable Development

Units: 4
Fall 2025 – Mon/Wed 4:00pm-5:50pm

Location: GFS 101

Instructor: Bistra Dilkina
Office: Ginsburg Hall (GCS) L2
Office Hours: 1 fixed-time office hour plus by appointment.

Contact Info: dilkina@usc.edu

Teaching Assistant: TBD
Office: TBD
Office Hours: 2 fixed-time office hours with TA
Contact Info: TBD

Course Catalogue Description

Hands-on AI for Sustainable Development: learn about data mining, ML, optimization, and fairness in the context of applications with environmental and societal benefits; team projects with real-world relevance.

Course Description

The course will focus on understanding how AI can be leveraged for social good and for making progress towards sustainable development. It will introduce AI concepts spanning *data mining, machine learning (classification, regression, ensembles, deep learning, feature importance), decision-making and optimization, and fairness in machine learning and algorithmic decision-making* in the context of informing applications in *environmental sustainability (biodiversity, climate, water, forests), disasters and climate change, poverty, homelessness, and health*. It will expose students to both core AI knowledge and cutting-edge research on the topic of using AI for social good applications. This course will introduce material through textbook reading, lectures, and academic research papers reading. In-class Python tutorials and individual Python assignments using real datasets will give students hands-on practice with the concepts and algorithms covered. In teams, students will experience *end-to-end the process of completing an applied AI project in the context of a societal or environmental domain*. The class will provide many opportunities for interactions with other students, the instructor and the TA.

Learning Objectives

Students will be able to:

1. Describe and apply AI methods covered in the course, including the basic concepts and the key algorithms
2. Describe pressing societal and environmental challenges, where AI has been successfully deployed to tackle them
3. Model societal challenges as mathematical problems that AI techniques can be applied to and recognize which AI techniques fit the problems
4. Prototype ML applications in Jupyter Notebook, including the full pipeline of a data-driven project

Prerequisite(s):

1. 1 from (CSCI 104 or CSCI 114)
2. CSCI 170

Non-CS undergraduate majors can fill in this D-clearance form, and will be considered as class space permits (after May 7): https://forms.gle/sb62ZvnYuTSrirW27

Recommended Preparation:

- **AI/ML exposure** is strongly recommended as this material will be reviewed and discussed in the context of sustainable development applications. Ideally, students would have taken at least one AI/ML or Data Science class, such as CSC I360, CSCI 467, DSCI 250, DSCI 352, ISE 225, or equivalent AI/ML experience.
- **Python** programming skills are strongly encouraged (classes that use Python depending on instructor: CSCI 353, CSCI 360, CSCI 445, EE 250, EE 364, also TAC 115 as introduction to Python). All coding assignments will be based on Python so students should either already have Python skills or be prepared to learn on their own in parallel with the course. Examples will be provided in the form of scaffolding for assignments.

Course Notes

The class is letter grade. Readings and slides will be posted to Blackboard. We will use Slack for Q&A and open discussions among students as well as Instructor and TA.

The class will provide many opportunities for interactions with other students, the instructor and TA. Please stop by my office during posted office hours. It can be to discuss course content, brainstorm ideas, ask a question about USC or careers, or just to stop by and say hello. If you cannot attend my posted office hours, please email me to arrange another time that is convenient for both of us.

Required Readings and Supplementary Materials

At this time, the growing area of AI applications in sustainable development and social good has not yet been formalized in a textbook. Therefore, most of the readings will be based on research papers and articles that will be provided to the class as PDF based handouts via Blackboard.

The following books provide useful background information about the techniques studied in class. However, they do not discuss the application to specific social good settings.

(PRML) Pattern Recognition and Machine Learning, Christopher Bishop, Springer; 2006

(ISLR) An Introduction to Statistical Learning with Applications in R, Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani, Springer, 2013. Available at http://faculty.marshall.usc.edu/gareth-james/ISL/ISLR%20Seventh%20Printing.pdf

(DL) Deep Learning, Goodfellow, I., Bengio, Y. and Courville A., 2016.

(MP) Applied Mathematical Programming. Bradley, Hax, and Magnanti (Addison-Wesley, 1977). Available at http://web.mit.edu/15.053/www/AMP.htm

Description and Assessment of Assignments

Paper reviews

The course will explore the course topics through a series of assigned readings in the form of research papers (and book chapters). Students will be assigned 1 research paper to read for a given week and submit a 1-2 page review for the assigned reading paper as homework. There will be 6 (+ 1 optional) such paper reviews assigned through the semester. Reviews will be assessed based on answering the following 5 questions (based on clarity and correctness):

1. What is the main problem/task addressed by the paper?
2. Why was AI needed and what AI approach was employed?
3. What does the paper's literature review suggest has been done previously, and how does the paper expand this previous work?
4. What data was used and were stakeholders / partner organizations engaged?
5. Discussion: what generalization to other settings/problems or what extension of the paper could be done?

Coding Assignments

Students will implement techniques studied in class by completing 4 Python notebook assignments. The assignments will be graded based on completion and correctness.

Class Presentation

Students will present individually or in small groups (depending on class size) an AI for Social Good research paper to the class. The paper will be selected from a list of papers selected by the instructor, or outside of the selected list with instructor permission. The presentations will be graded on clarity, completeness, and presentation style.

Peer presentation learning

Peer learning during student presentations will be assessed with a short questionnaire on each student presentation in terms of relevance to technical and social good course themes.

Semester Project

Students will work in small groups to carry out a class project. The focus of the class project will be to develop an innovate application of AI to address a social good problem. Students will leverage tools, concepts, and techniques presented in the class. The project involves identifying a problem related to social good, data sources available to inform the problem, and AI-based approaches to it. The project will involve: data cleaning and fusion,

data exploration/preporcessing/visualization, implementing at least 2 competing AI approaches, presenting results in terms of a table of multiple (at least 3) evaluation metrics and possible visualizations of outputs, and a discussion. Project topics will be suggested by the instructor, and also students will have the freedom to propose their own. The grade for the project will be based on the successful completion of the agreed upon project objectives. The deliverables include a project proposal (1-2 pages single space), premlinary paper draft (4-8 pages single space), final presentation (10-20 minutes) and a final report (10-15 pages single space). Examples of past projects (proposals, paper drafts, presentation and report) will be shared wit the class (with persmission from past students). They will be graded based on clarity, and completeness. The project is total 55% of final grade broken into deliverables throughout the semester as listed in Grading Breakdown.

Grading Breakdown

The grades for the students will be based on completion of the paper review assignments, presentation of a research paper, programming assignments, and a semester project. The breakdown for each of these categories is listed below. A more detailed explanation of the grading for each category is also provided.

Assessment Tool (assignments)	**% of Grade**
Class participation	2%
Paper Reviews	9%
Python assignment (4x4%)	16%
Class Presentation	15%
Peer Class Presentation Quizzes	3%
Project Proposal	5%
Project Preliminary Paper Draft	10%
Project Final Presentation	10%
Project Final Paper	30%
TOTAL	100%

Assignment Submission Policy
Assignments will be submitted via Brightspace.

Grading Timeline
Grades will be provided within 2 weeks of submission of the respective assignment.

Additional Policies
This is a discussion-based course, hence consistent attendance is expected. Missed classes with a valid excuse are allowed. Class participation will be scored based on engagements in course discussions. Class participation will be evaluated based on Q&A and engaging in discussions in class or on the class chat. At least 6 meaningful class interactions are needed to get full participation score. Meaningful interaction could be asking or answering questions during lecture, commenting on a paper being presented by the instructor or other students, or presenting updates on project progress, among others

Course Schedule: A Weekly Breakdown

	Topics/Daily Activities	Deliverables
Week 1	Introduction + Jupyter Notebooks Tutorial	
Week 2	Clustering + Clustering Tutorial / Clustering for Climate	Paper review 1
Week 3	Regression (Cross-Validation, Features Selection) / Applications	Python assignment 1: clustering
Week 4	Regression / Applications	Paper review 2 Python assignment 2: regression
Week 5	Classification (Trees, Ensembles, Feature Importance) / Applications	Project Proposal
Week 6	Deep Learning & Remote Sensing / Applications	Paper Review 3
Week 7	Student Paper Presentations	Paper Presentations + Reading (Quiz)
Week 8	Student Paper Presentations	Paper Presentations + Reading (Quiz) (Optional paper review)
Week 9	Student Presentations; Submodular Optimization / Applications	Paper Presentations + Reading (Quiz) Paper review 4
Week 10	Integer Programming Optimization / Applications	Python assignment 3: optimization
Week 11	Integer Programming Optimization / Applications	Project preliminary paper draft
Week 12	Multi-objective Optimization + Genetic Algorithms / Applications	Paper review 5
Week 13	AI & Fairness	Paper review 6
Week 14	AI & Fairness	Python assignment 4: fairness
Week 15	Project Presentations	Project Presentations
FINAL	Project Report	Due on University-scheduled date of the final exam

Papers covered in the class (subject to change):

Week 1: Introduction to **Sustainable Development + AI for Social Good Principles**

- "Artificial intelligence for social good: A survey." Shi, Zheyuan Ryan, Claire Wang, and Fei Fang. arXiv preprint arXiv:2001.01818 (2020).
- "Computational sustainability: Computing for a better world and a sustainable future." Gomes et al. Communications of the ACM 62.9 (2019)

- The role of artificial intelligence in achieving the Sustainable Development Goals. Vinuesa, R., Azizpour, H., Leite, I. et al. Nature Communications **11,** 233 (2020). https://doi.org/10.1038/s41467-019-14108-y

Week 2: Data Mining / **Environment**

- Book: Introduction to Data Mining, by Tan, Steinbach, Kumar. Chapter 8. Cluster Analysis: Basic Concepts and Algorithms [can be downloaded at: http://www-users.cs.umn.edu/~kumar/dmbook/ch8.pdf], ISLR Ch. 10
- "Representativeness-based Sampling Network Design for the State of Alaska." Hoffman, Forrest M., Jitendra Kumar, Richard T. Mills, and William W. Hargrove. 2013. Landscape Ecology

Week 3+4: Regression / **Energy + Climate Change**

- PRML Chapter 3, ISLR Ch. 3, 5 & 6
- A Large-Scale Study on Predicting and Contextualizing Building Energy Usage . J. Zico Kolter, Joseph Ferreira. AAAI 2011
- A Machine Learning Approach to Modeling Human Migration. Caleb Robinson and Bistra Dilkina. ACM SIGCAS Conference on Computing and Sustainable Societies 2018

Week 5: Classification **/ Illegal Wildlife Poaching**

- PRML Chapter 4 & 14, ISLR Ch. 4 & 8
- "Adversary models account for imperfect crime data: Forecasting and planning against real-world poachers." Gholami, Shahrzad, et al. AAMAS 2018.

Week 6: Deep Learning / **Satellite-based Land Cover + Poverty**

- PRML Chapter 5, DL Ch. 6 & 7
- "Combining satellite imagery and machine learning to predict poverty." Jean, Neal, et al. Science 2016.
- Large scale high-resolution land cover mapping with multi-resolution data. Robinson, C et al. In IEEE Conference on Computer Vision and Pattern Recognition 2019.

Week 7+8: Paper presentations on **Health | Energy | Environment | Public Sector** |

Week 9: Submodular Optimization / **Environmental Sensing + Public Health**

- Cost-effective outbreak detection in networks. Leskovec, J et al. In KDD 2007
- "Efficient sensor placement optimization for securing large water distribution networks." Krause, A. et al. Journal of Water Resources Planning and Management, 2008
- Using Social Networks to Aid Homeless Shelters: Dynamic Influence Maximization under Uncertainty. Amulya Yadav et al. AAMAS 2016

Week 10+11: Mathematical Programming / **Biodiversity Conservation**

- MP Chapters 1 & 9
- "Trade-offs and efficiencies in optimal budget-constrained multispecies corridor networks." Dilkina, Bistra, et al. Conservation Biology (2017)

Week 12: Multi-objective, Genetic Algorithms

- Deb, Kalyanmoy. "Multi-objective optimisation using evolutionary algorithms: an introduction." Multi-objective evolutionary optimisation for product design and manufacturing. Springer, London, 2011. Available at: https://link.springer.com/chapter/10.1007/978-0-85729-652-8_1

Week 13+14: AI and Fairness / **Crime**

- Tutorials: https://dssg.github.io/fairness_tutorial/, https://www.borealisai.com/en/blog/tutorial1-bias-and-fairness-ai/
- "Why Should I Trust You?": Explaining the Predictions of Any Classifier. Marco Tulio Ribeiro, Sameer Singh, Carlos Guestrin. In KDD, 2016
- A large-scale analysis of racial disparities in police stops across the United States. Pierson et al., https://arxiv.org/abs/1706.05678
- Julia Angwin, Jeff Larson, Surya Mattu, Lauren Kirchner, "Machine Bias". https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm

Academic Integrity

The University of Southern California is foremost a learning community committed to fostering successful scholars and researchers dedicated to the pursuit of knowledge and the transmission of ideas. Academic misconduct — which includes any act of dishonesty in the production or submission of academic work (either in draft or final form) — is in contrast to the university's mission to educate students through a broad array of academic, professional, and extracurricular programs.

This course will follow the expectations for academic integrity as stated in the USC Student Handbook. All students are expected to submit assignments that are their own original work and prepared specifically for this course and section in this academic term. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s). Students suspected of engaging in academic misconduct will be reported to the Office of Academic Integrity.

Other violations of academic misconduct include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

Academic dishonesty has a far-reaching impact and is considered a serious offense against the university. Violations will result in a grade penalty, such as a failing grade on the assignment or in the course, and disciplinary action from the university itself, such as suspension or even expulsion.

For more information about academic integrity see the student handbook or the Office of Academic Integrity's website, and university policies on Research and Scholarship Misconduct.

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment or what information requires citation and/or attribution.

Course Content Distribution and Synchronous Session Recordings Policies
USC has policies that prohibit recording and distribution of any synchronous and asynchronous course content outside of the learning environment.

Recording a university class without the express permission of the instructor and announcement to the class, or unless conducted pursuant to an Office of Student Accessibility Services (OSAS) accommodation. Recording can inhibit free discussion in the future, and thus infringe on the academic freedom of other students as well as the instructor. (Living our Unifying Values: The USC Student Handbook, page 13).

Distribution or use of notes, recordings, exams, or other intellectual property, based on university classes or lectures without the express permission of the instructor for purposes other than individual or group study. This includes but is not limited to providing materials for distribution by services publishing course materials. This restriction on unauthorized use also applies to all information, which had been distributed to students or in any way had been displayed for use in relation to the class, whether obtained in class, via email, on the internet, or via any other media. Distributing course material without the instructor's permission will be presumed to be an intentional act to facilitate or enable academic dishonestly and is strictly prohibited. (Living our Unifying Values: The USC Student Handbook, page 13).

Academic Integrity

The University of Southern California is foremost a learning community committed to fostering successful scholars and researchers dedicated to the pursuit of knowledge and the transmission of ideas. Academic misconduct — which includes any act of dishonesty in the production or submission of academic work (either in draft or final form) — is in contrast to the university's mission to educate students through a broad array of academic, professional, and extracurricular programs.

This course will follow the expectations for academic integrity as stated in the USC Student Handbook. All students are expected to submit assignments that are their own original work and prepared specifically for this course and section in this academic term. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s). Students suspected of engaging in academic misconduct will be reported to the Office of Academic Integrity.

Other violations of academic misconduct include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

Academic dishonesty has a far-reaching impact and is considered a serious offense against the university. Violations will result in a grade penalty, such as a failing grade on the assignment or in the course, and disciplinary action from the university itself, such as suspension or even expulsion.

For more information about academic integrity see the student handbook or the Office of Academic Integrity's website, and university policies on Research and Scholarship Misconduct.

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment or what information requires citation and/or attribution.

Course Content Distribution and Synchronous Session Recordings Policies

USC has policies that prohibit recording and distribution of any synchronous and asynchronous course content outside of the learning environment.

Recording a university class without the express permission of the instructor and announcement to the class, or unless conducted pursuant to an Office of Student Accessibility Services (OSAS) accommodation. Recording can inhibit free discussion in the future, and thus infringe on the academic freedom of other students as well as the instructor. (Living our Unifying Values: The USC Student Handbook, page 13).

Distribution or use of notes, recordings, exams, or other intellectual property, based on university classes or lectures without the express permission of the instructor for

purposes other than individual or group study. This includes but is not limited to providing materials for distribution by services publishing course materials. This restriction on unauthorized use also applies to all information, which had been distributed to students or in any way had been displayed for use in relation to the class, whether obtained in class, via email, on the internet, or via any other media. Distributing course material without the instructor's permission will be presumed to be an intentional act to facilitate or enable academic dishonestly and is strictly prohibited. (Living our Unifying Values: The USC Student Handbook, page 13).