On-campus Students:
Time and Room: See UAccess

Distance Ed Students:
Access lecture video recordings through Content page of D2L

Dr. Matthias Poloczek
poloczek@email.arizona.edu

Office Hours: **See Announcement in Class**

Course Description:

We consider optimization problems whose objective functions are unknown and hence have to be learned from data. Such problems are pervasive in science and industry, e.g., when

- designing prototypes in engineering,
- automated tuning of machine learning algorithms, e.g., in deep learning,
- optimizing control policies in robotics,
- developing pharmaceutical drugs, and many more.

Bayesian optimization methods are popular in the machine learning community due to their high sample-efficiency and have become a key technique in the area of "automatic machine learning". We introduce a general framework in which to understand and formulate such **optimal learning** problems, and provide a survey of problems, methods, and theoretical results.

Course Prerequisite(s):

This course is intended for a wide audience interested in machine learning and optimization. Learners from Industrial Engineering, Systems Engineering, Engineering Management, other fields of engineering (e.g., ECE, AME, BME), applied math and computer science are very welcome.

The success of machine learning relies on mathematical models fitted to data and the development of algorithmic techniques. The instructor will assume solid knowledge of elementary probability theory, linear algebra, and calculus. In particular, students should be comfortable with random variables, conditional probability and (conditional) expectation, common probability distributions and their properties (e.g., binomial, geometric, exponential, Poisson, Gaussian).

Students should also be comfortable with coding small programs for activities, and will be expected to complete a larger program for the project. While Python (version 3, for example with Anaconda, GPy, GPyOpt, and matplotlib) is recommended, the choice of the programming language is yours. The instructor cannot provide individual assistance ("debugging"). There will be an online forum to facilitate solving coding issues collaboratively.

Note that the algorithmic techniques covered in this course boost the performance of **deep learning** by orders of magnitudes. Thus, this course provides substantial synergies with "*Bayesian Machine Learning II*", although it is not a formal prerequisite for the latter. Both courses complement each other.

Please e-mail the instructor if you are concerned about having appropriate prerequisites.

Instructor and Contact Information:

Prof. Dr. Matthias Poloczek poloczek@email.arizona.edu
Office Hours: TBA, or by appointment

Teaching Assistant: TBA

Course Format and Teaching Methods:

This course is based on lectures, homework sets, and group projects. This class will use web-based D2L (Desire to Learn) as the only means of distributing class materials including assignments. All assignments and submissions will be uploaded by the student into Assignment Submission Folders on D2L before the appropriate due dates. For each set you will be given a *grace period* of 24 hours, unless stated otherwise. Later submissions will not be graded. Exceptions from this rule will be made only if there is a convincing, properly documented reason. Your grades for this course will also be available on D2L. You will need a UANet ID to access D2L at the following site: http://d2l.arizona.edu/. You are expected to check D2L frequently for class information.

For distance learning students, you will find your team mates remotely and meet online for your projects (Instructor will provide you detail instruction about that in the recorded lecture). If you are not able to form team remotely, you are allowed to carry out the project individually based on recorded lectures accessible through D2L.

Course Topics:

- Overview of applications and models of Optimal Learning.
- Foundations in probability theory and Bayesian statistics.
- Common statistical models and applications.
- The value of information and useful value functions.
- Popular acquisition functions in Bayesian optimization.
- The Knowledge Gradient.
- Best practices.
- Applications in simulation optimization, image classification, etc.
- Theoretical performance guarantees for selected acquisition functions.
- Advanced methods for hyperparameter treatment.

This list is tentative. The instructor reserves the right to make modifications if appropriate.

Course Objectives:

During this course students will

- Develop a framework to understand and formulate problems in optimal learning.

- Recognize the exploration versus exploitation trade-off and the curse of dimensionality.
- Gain an overview of algorithmic techniques in Bayesian machine learning.
- Implement a Bayesian optimization algorithm to solve problems in optimal learning.
- Analyze the performance of a state-of-the-art optimization method.
- Discuss challenges and future directions of Bayesian machine learning.

Expected Learning Outcomes:
Upon completion of this course, students will be able to

- Identify and categorize problems in optimal learning.
- Describe uncertainties arising in such problems using a variety of statistical models.
- Apply a variety of algorithmic techniques to solve sequential optimization problems efficiently in the presence of uncertainty.
- Understand proof strategies to provide robustness guarantees for optimization algorithms.

Absences and Class Participation Policy:

The UA's policy concerning Class Attendance, Participation, and Administrative Drops is available at: http://catalog.arizona.edu/policy/class-attendance-participation-and-administrative-drop

The UA policy regarding absences for any sincerely held religious belief, observance or practice will be accommodated where reasonable, http://policy.arizona.edu/human-resources/religious-accommodation-policy.

Absences pre-approved by the UA Dean of Students (or Dean Designee) will be honored. See: https://deanofstudents.arizona.edu/absences

Participating in the course and attending lectures and other course events are vital to the learning process. Students are responsible for all materials covered during class. As such, attendance is strongly recommended. Occasionally, attendance may be required for special events such as Project Presentations. Students who miss class due to illness or emergency are required to bring documentation from their health-care provider or other relevant, professional third parties. Failure to submit third-party documentation will result in unexcused absences.

Class Guidelines:
All students:

- Check D2L regularly.
- Turn-in assignments by due date/time (allow for D2L "glitches").
- Treat instructors, speakers and peers with respect.
- Always behave in an ethical manner.
- Students are encouraged to use laptops for note taking and other class activities.
- Make sure that you are subscribed to the respective forum at D2L: after clicking on 'Discussions' in D2L, click the arrow next to the forum name to subscribe to the whole forum.

On-campus students:

- Arrive on-time, turn off cell phones, beepers, social networks, etc.
- Attend class regularly and participate in class discussions and activities.

Distance Ed students:

- View lectures in a timely manner, preferably within 48 hours of the lecture date.
- In some cases, Content will be developed "For Distance Ed Students Only". Pay careful attention to all communications.

Course Communication:

D2L will be used for normal course communication. Both D2L and UA e-mail will be used for any critical items.

To assert that every students has access to the same information, such questions will only be answered in class or in our D2L forum. Do not inquire about projects and assignments *via email*. This policy does not apply to personal or sensitive matters, e.g., special needs.

Required Text:

Students will not be required to purchase a text book. Recommended books include

- W.B. Powell and I.O. Ryzhov, *Optimal Learning*, ISBN-13: 978-0470596692
- C.E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, ISBN-13: 978-0262182539, available at http://www.gaussianprocess.org/gpml/chapters/RW.pdf
- C.M. Bishop, *Pattern Recognition and Machine Learning*, ISBN-13: 978-0387310732
- E. Çınlar, *Probability and Stochastics*, ISBN-13: 978-0387878584

Further references will be provided in class and on the course website.

Required Materials:

Computer with internet access and your choice of programming tools.

Required Extracurricular Activities:

None

Assignments and Examinations:

The Final Grade will be based on the *project* (50%), *readings* (5%), the *UnLectures* (15%, approximately equally weighted), *participation* (5%), the *presentation* (15%), and the *peer reviews* (10%).

Project: There will be a project in the second half of the semester that allows students to apply lessons learned in a research project of their choosing. The project involves

- *Formulating an original research proposal* based on the content of the course,
- *Solving it*, where "solving" could mean computing the optimal solution, evaluating different algorithms experimentally, or proving structural results,
- *Writing a project report*, and giving a *presentation* of the work.

Students are encouraged to pursue a research topic that they find interesting, e.g., related to their thesis research or private interests. Please feel free to consult with the instructor about the proposal early.

For the project report, an interactive document (e.g., Jupyter, Matlab, Mathematica) is suggested which states the problem and results, and has interactive demonstrations of the work. (A static document is also accepted without penalty.)

Readings: From time to time, students will be pointed to reading materials that may provide the foundation for topics presented in class or a deeper treatment. Students are required to write a concise summary (at most one page) for each such assignment.

UnLectures: UnLectures are activities designed to (1) explore and deepen concepts presented in this course, (2) discover new techniques in a collaborative fashion, and (3) promote critical thinking and reflection on key topics. Each UnLecture will involve an in-class session and pre- and post-session write-ups. More details will be provided one week prior to an UnLecture.

Peer reviews: Students will occasionally be asked to review other students' submissions according to criteria specified in class. These reviews will be "single blind", i.e. reviewers will know the identity of the author, whereas the reviewers' identities are only known to the instructor.

Presentation: Students will present selected scientific works in class. The instructor will provide a list of suitable topics and manuscripts that for example apply techniques discussed in the course to problems in engineering, health care, transportation, machine learning etc. In addition, students may propose works, e.g., related to their research project. In this case, students are to reach out to the instructor early.

Participation: Participations points are given for active participation in class and for assisting other students in the online forum (but not for raising questions). Active participation in class is not applicable for remote students.

For group assignments, students only need to hand in a single submission per group that lists the names and UANet IDs of all students involved in the preparation and their individual contributions.

Grading Scale and Policies:

Final Grades for this course will be computed as follows:

>90%	A
>80%	B
>70%	C
>60%	D
<60%	E

Requests for incomplete (I) or withdrawal (W) must be made in accordance with University policies, which are available at http://catalog.arizona.edu/policy/grades-and-grading-system#incomplete and http://catalog.arizona.edu/policy/grades-and-grading-system#Withdrawal respectively.

Honors Credit:

Honors Credit will not be available at this time.

Scheduled Topics/Activities:

All assignments will be announced on D2L with their respective due dates.

All homework/projects/presentations and papers are to be submitted by the due date/time to the D2L Assignment Submission Folder unless otherwise specified. All D2L activities, including Discussions, Surveys and Quizzes, must be complete by the due date/time. No late work is accepted except where a grace period is given. No extra credit is available.

Classroom Behavior Policy:

To foster a positive learning environment, students and instructors have a shared responsibility. We want a safe, welcoming, and inclusive environment where all of us feel comfortable with each other and where we can challenge ourselves to succeed. To that end, our focus is on the tasks at hand and not on extraneous activities (e.g., texting, chatting, reading a newspaper, making phone calls, web surfing, etc.).

Threatening Behavior Policy

The UA Threatening Behavior by Students Policy prohibits threats of physical harm to any member of the University community, including to oneself. See http://policy.arizona.edu/education-and-student-affairs/threatening-behavior-students.

Accessibility and Accommodations:

At the University of Arizona we strive to make learning experiences as accessible as possible. If you anticipate or experience physical or academic barriers based on disability or pregnancy, you are welcome to let me know so that we can discuss options. You are also encouraged to contact Disability Resources (520-621-3268) to explore reasonable accommodation.

If our class meets at a campus location: Please be aware that the accessible table and chairs in this room should remain available for students who find that standard classroom seating is not usable.

Code of Academic Integrity

Students are encouraged to share intellectual views and discuss freely the principles and applications of course materials. However, graded work/exercises must be the product of independent effort unless otherwise instructed. Students are expected to adhere to the UA Code of Academic Integrity as described in the UA General Catalog. See: http://deanofstudents.arizona.edu/academic-integrity/students/academic-integrity.

The University Libraries have some excellent tips for avoiding plagiarism, available at http://www.library.arizona.edu/help/tutorials/plagiarism/index.html.

Selling class notes and/or other course materials to other students or to a third party for resale is not permitted without the instructor's express written consent. Violations to this and other course rules are subject to the Code of Academic Integrity and may result in course sanctions. Additionally, students who use D2L or UA e-mail to sell or buy these copyrighted materials are subject to Code of Conduct Violations for misuse of student e-mail addresses. This conduct may also constitute copyright infringement.

UA Nondiscrimination and Anti-Harassment Policy

The University is committed to creating and maintaining an environment free of discrimination; see http://policy.arizona.edu/human-resources/nondiscrimination-and-anti-harassment-policy

Our classroom is a place where everyone is encouraged to express well-formed opinions and their reasons for those opinions. We also want to create a tolerant and open environment where such opinions can be expressed without resorting to bullying or discrimination of others.

Additional Resources for Students:

UA Academic policies and procedures are available at http://catalog.arizona.edu/policies

Student Assistance and Advocacy information is available at http://deanofstudents.arizona.edu/student-assistance/students/student-assistance

Confidentiality of Student Records

All student records are held in strict confidence. Additional information can be found at http://www.registrar.arizona.edu/personal-information/family-educational-rights-and-privacy-act-1974-ferpa?topic=ferpa

Subject to Change Statement:

Information contained in the course syllabus, other than the grade and absence policy, may be subject to change with advance notice, as deemed appropriate by the instructor.

This syllabus is tentative and the instructor reserves the right to make modifications if appropriate.

M. Poloczek, May 9, 2018