

# MANOS ATHANASSOULIS – TEACHING EXPERIENCE AND PHILOSOPHY

*Teaching is to inspire students to become the best possible scientists, practitioners, and ultimately citizens.* Toward this path, the teaching process guides students to build their technical skills, self-confidence, creativity, and passion for technology and science. Today computer science and data management are two subjects that make our world spin, offering tremendous opportunities for creativity, innovation, and research.

**Teaching Philosophy.** Exposing the students to a real practical problem is key to motivate them regarding the intrinsic and often laborious details of a subject. When I taught a Database Systems class, I designed from scratch a class that demanded throughout the semester increasingly more engaged system building. The goal was to gradually expose the class to the complexity of a database system and motivate them by the tangible result of their labor. Covering theoretical topics on database design and data modeling naturally helped them to understand how to proceed in the class projects, and build a “complete” system.

Solving real problems and building complete systems is no easy task. Thorough preparation of lectures, diligent design of homework questions, projects, and exams is paramount to enable the students to fulfil their potential. Additionally, making them feel comfortable asking any question and seeking support is a key effort toward ensuring that the students *actively participate* in a class. I particularly encourage my students to ask questions and to question things I present to crystallize their understanding. I also *ask them questions* to uncover more of their questions and observations. Finally, as a faculty member, I aim to provide an environment that will offer equal opportunity to all students who are interested in learning or in conducting research. I consider removing barriers from knowledge a duty both as an academic and as a citizen.

**Teaching Assistant.** As a postdoc at Harvard University, I have been the lead teaching fellow in two classes: the undergraduate class Data Systems (fall 2015, fall 2016) and the graduate class Big Data Systems (fall 2014, spring 2016), both offered by Prof. Stratos Idreos.

The goal of Data Systems is to teach the concepts of modern data systems and to build a complete data system from scratch. As the lead teaching fellow, I was responsible for preparing and presenting supplementary material, for guiding the students through their project designs and implementations, and for building automatic testing infrastructure for the project. The class was also offered as part of the Harvard Extension School, for which I offered similar support to remote students through online collaboration platforms. In retrospect, it was a challenging yet rewarding teaching experience.

Big Data Systems is an advanced class offered at both Harvard College and Harvard Extension School. It covers state-of-the-art research papers from the data management community. In addition to studying the cutting edge advancements of the community, a key element of this class is introducing students to research through small-scale research projects. One of these projects resulted in publishing a workshop paper with two students in the area of data management on modern hardware (in DAMON, a workshop collocated with SIGMOD, the premiere conference of data management).

For each instance of these classes, I have received awards as a teaching fellow. In particular, I received the “Harvard University Certificate of Teaching Excellence” in fall 2014 and the “Harvard University Certificate of Distinction in Teaching” in fall 2015, spring 2016, and fall 2016.

During my doctoral studies at EPFL, I was the lead teaching assistant in the undergraduate class Introduction to Database Systems (named Databases for some years) for four semesters (spring 2010 through spring 2013), offered by Prof. Anastasia Ailamaki. As a teaching assistant at EPFL, I received the 2013 “EPFL Outstanding Teaching Assistant Award”.

**Instructor.** During spring 2017, I taught the class *Database Systems* at Tufts University, offered at both the undergraduate and the introductory graduate level. I designed the lectures, created homework assignments, projects and exams, prepared slides and notes, and regularly met with students. My role was challenging: I had to make the class suitable for a group of students with varied backgrounds and academic experience. Several positive comments I received from students show that the class was successful in engaging them. The class website is archived at <http://manos.athanassoulis.net/classes/Comp115-Spring2017>.

**Mentorship.** Both as a PhD student and as a postdoc, I had the privilege to work closely with a number of great students. As a postdoctoral researcher at Harvard, I work closely with Michael Kester, a PhD student supervised by Prof. Stratos Idreos. Our research results on access path selection in main-memory data analytics query engines were recently presented at the top-tier data management conference (SIGMOD). I further worked with Zheng Yan, an undergraduate research intern. Our work proposing a new data structure that offers efficient updates on bitmap indexes was also presented at SIGMOD and received the award of “Most Reproducible Paper”. Earlier, as a senior PhD student at EPFL, I worked closely with Iraklis Psaroudakis, a PhD student supervised by Prof. Anastasia Ailamaki. Our work on building query processing systems that employ work sharing was presented at the top-tier conferences on data management (VLDB and SIGMOD). I also worked with Andreas Schädli, a master’s student at EPFL, with whom we found and corrected a bug in the popular performance benchmark in data management TPC-H.

**Tutorials.** In spring 2016, I delivered a tutorial at the ACM SIGMOD international conference on data management. Offering a tutorial was a rewarding task: teaching a subject I am passionate about to an interested and demanding audience. I presented my research results as well as my broader view of the area of designing access methods. The preparation and delivery of a tutorial is a dense teaching experience. It requires preparing aesthetically pleasing slides, abstracting out the common principles of multiple different approaches, and delivering an engaging presentation. The material of this tutorial is publicly available at <http://manos.athanassoulis.net/publication.html#DesignAMTutorial>.

**Teaching.** My general background permits me to teach various introductory-level courses in computer science, computer systems, and software engineering. My research and teaching experience makes me well-suited for teaching subjects related to: (i) databases, (ii) data management systems, (iii) big data systems, (iv) computer systems design, principles, and methodology, and (v) cloud infrastructure.

I am also interested in teaching a class in data system design and access method design. Over the last ten years, increasingly more applications rely on analyzing data to make key decisions (anything from scientific exploration, to business intelligence and government). The workflow of *generating data/monitoring*, then *curating*, and finally *storing, maintaining, and analyzing* data, faces tremendous challenges because of the sheer size of the data maintained and the data access patterns needed to be efficiently supported. The goal of the class is to cover in detail the common system-level design principles (physical data layout, query language, query processing with code generation, and work sharing) and the performance tradeoffs of hardware (multi-threaded execution and GPUs as well as modern storage and non-volatile memories) and of access method designs (read-optimized vs. write-optimized data structures). The ultimate goal of this class is to expose students to real-life data management challenges, guiding them to build the necessary systems skills and intuition needed for both academic research and industrial research and development.