

PROGRAM PROPOSAL, UNIVERSIDAD DE BUENOS AIRES

DRAFT

Licenciatura en Sistemas Complejos

Bachelor of Systems and Decision Science

UNIVERSIDAD DE BUENOS AIRES

FACULTAD DE CIENCIAS ECONÓMICAS · FACULTAD DE INGENIERÍA

Duration: 3.5 years

Cohort: 30–40 students

8-week blocks

VIEW: AUTHORITIES

Preface

This is a proposal, a draft document, not an approved program. It presents the Licenciatura en Sistemas Complejos, Bachelor of Systems and Decision Science, for consideration as a new undergraduate degree at the Universidad de Buenos Aires.

It is written for three audiences simultaneously. For the academic authorities of UBA, it presents the intellectual foundations, pedagogical model, curriculum, and operational structure of the program in sufficient detail to evaluate its rigor and its coherence. For the faculty who would teach in the program, it presents the disciplinary architecture and the pedagogical commitments. For the students who would inhabit it, it presents the argument, the founding conviction that the program is a response to, and the bet it is making about what the world needs.

The document is bilingual: Spanish by default, English available via the toggle in the top right. Course names and key institutional terms are given in Spanish.

If approved, this document will become the full account of what was approved and why. Until then, it is a proposal under active development. Feedback is welcome.

Executive Summary

3.5 YEARS	39 COURSES	6 STUDIOS
30–40 STUDENTS PER COHORT	8 WEEKS PER BLOCK	0 WRITTEN FINAL EXAMS

The Licenciatura en Sistemas Complejos, Bachelor of Systems and Decision Science, is a new undergraduate degree at the Universidad de Buenos Aires designed to form a generation of builders capable of operating at the intersection of technology, institutions, and power in a world of abundant artificial intelligence.

The program's objective is formation, the development of a kind of person, not the transmission of a body of content or the credentialing of a professional skill. Specifically, it forms people who possess the two things that AI cannot yet replace: judgment under genuine uncertainty and the capacity to construct legitimacy. People who can decide what to build, not just how to build it. People who can take responsibility for outcomes that affect others at scale.

The curriculum spans mathematics, physics, biology, computer science, machine learning, economics, game theory, institutional design, political economy, philosophy of science, and cultural analysis. It is completed by six Studio projects of increasing ambition, culminating in the creation of a real company during the

final Studio. The program runs in eight-week blocks, a structure borrowed from the École Polytechnique, with three courses per block and a Studio running continuously across two consecutive blocks.

The program admits 30 to 40 students per cohort through a competitive three-stage selection process. There are no written final exams. All primary evaluation is in person, in real time: written exams, oral defenses, live coding, or public defense depending on the course. AI tools are mandatory throughout.

The world has enough people who know how to write a prompt. It needs people who can decide what to build, why it matters, and how to make it last.

One useful historical reference is Oxford's PPE, Philosophy, Politics and Economics, a degree built to form people capable of reasoning across institutions, policy, and public life. The analogy is not curricular equivalence but institutional ambition: this program asks what a similarly integrative undergraduate formation would require today, in a world shaped by AI, complex technical systems, and institutional fragility. Its answer combines philosophy, economics, and political judgment with mathematics, physics, computer science, machine learning, and a sustained builder orientation.

It was designed by studying reference programs, MIT, Stanford, Caltech, École Polytechnique, Oxford, Cambridge, Harvard, Carnegie Mellon, the Santa Fe Institute, and asking which elements of those traditions could be combined into a coherent formation for this moment. The answer became the curriculum.

Context and Motivation

1.1 The Moment

We are living through a transition comparable in scope to the Industrial Revolution. Artificial intelligence is not a new tool in a stable landscape. It is a force that is reorganizing what kinds of work are possible, what kinds of organizations are viable, what kinds of knowledge are scarce, and what kinds of people are needed to navigate the transition.

The most visible consequence is the automation of cognitive labor. Tasks that required years of training, writing, coding, legal research, financial analysis, medical diagnosis, are being performed by AI systems at a fraction of the cost and time. This is not a future prediction. It is happening now, and the rate of change is accelerating.

The less visible consequence is more important: the automation of cognitive labor raises the value of what cannot be automated. Judgment. Legitimacy. Responsibility. The capacity to decide what to build, not just how to build it. The capacity to operate in institutions, with their politics, their cultures, their histories, their competing interests, and make things happen despite them. The capacity to take responsibility when systems fail and other people bear the cost.

These capacities are not the primary focus of traditional university education. They require a different kind of formation, one that combines disciplinary rigor, cross-domain integration, and sustained practice under real constraints.

1.2 The Limits of Traditional Disciplines

The modern university is organized around disciplines that were designed for a different world. Each does its work with excellence: computer science forms people who build sophisticated software; economics forms people who model markets rigorously; political science forms people who analyze power in depth; business schools form people who manage existing organizations effectively. That disciplinary depth is a real achievement, and this program respects it and builds on it.

What a single discipline is unlikely to produce by itself is fluency across all of them, the capacity to move between technology, economics, law, culture, and institutions, where many of the most complex problems live. Reference programs such as computer science at MIT, Oxford's PPE, or economics at Harvard show the strength of disciplinary depth. Disciplinary integration is a distinct, complementary task that requires its own curricular architecture.

The result is a systematic gap between the complexity of the problems that need to be solved and the formation of the people who are supposed to solve them. Regulators who do not understand the technology they regulate. Technologists who do not understand the institutions their technology disrupts. Economists who do not understand the cultural and political constraints on the policies they recommend. Leaders who do not understand the systems they are responsible for.

1.3 The Gap

This gap is not new. But AI is making it catastrophically more visible. As AI systems become more capable, the people directing them need to be more capable, not less. The risk is not that AI replaces human judgment. The risk is that human judgment atrophies at exactly the moment when it matters most.

The gap has a specific shape. On one side: an unprecedented supply of technical capability, analytical power, and information processing. On the other side: a shortage of people capable of deciding what to do with it, of setting objectives,

navigating institutional constraints, earning the trust of the people affected, and taking responsibility for the outcomes.

Filling that gap is the purpose of this program.

1.4 Why Argentina and Why UBA

Argentina is not an obvious place to launch a program of this ambition. It is, however, a necessary one.

Latin America is systematically underrepresented in the global conversation about artificial intelligence, institutional design, and the future of technology. The frameworks being developed to govern AI, to regulate platforms, to design the institutions of the post-AI world, these are being developed primarily in the United States, Europe, and China. Latin America is largely absent from the design process and will largely inherit the results.

This is not inevitable. Argentina has one of the most educated populations in Latin America, a strong tradition of mathematics and science at the university level, a culture of intellectual seriousness, and a generation of young people who are as connected to global technological and intellectual currents as any in the world.

The Universidad de Buenos Aires is the natural home for this program. UBA is the largest and most prestigious university in Argentina, with a tradition of academic rigor and intellectual independence that is rare in the region. It has produced Nobel laureates, Fields Medal winners, and generations of scientists, lawyers, doctors, and economists who have shaped Argentine public life.

What UBA has not produced, and what Argentina needs, is a generation of people capable of building the institutions, companies, and systems that the post-AI world requires. The Licenciatura en Sistemas Complejos is designed to produce that generation.

1.4.1 Why UBA Should Host This Program

The choice of UBA as host is neither accidental nor ceremonial. Argentina has several universities of high quality, and this program could be imagined at more than one of them. The combination that UBA offers, however, scale, a tradition of rigor, intellectual independence, free admission, and regional prestige, is not easily replicated by any other institution in the country.

UBA is the heir to the scientific tradition that produced Bernardo Houssay, Luis Federico Leloir, and César Milstein, three Nobel laureates trained in its classrooms, in different disciplines, across three generations. That tradition is not merely a historical fact: it is a way of working, an expectation of seriousness, an insistence that university formation be a formation in the craft of thinking carefully. The founding of the Facultad de Ciencias Exactas y Naturales, the creation of the Carrera de Sociología, and more recently the Licenciatura en Ciencia de Datos, show that UBA is capable of generating new disciplines when intellectual conditions demand it. This program is situated in that same lineage.

To this is added a decisive feature: UBA is a free public university. This means the program can admit the most capable students without filtering them by family income. The selection the program requires is pedagogical, the cohort is small because the formation it offers demands closeness with faculty and mentors, but it is not economic. In a program designed to form people who will serve the public interest, this is not an administrative detail: it is a condition of possibility. A formation of leaders that is filtered by family wealth produces leaders filtered by family wealth. UBA, by contrast, allows the program to seek talent wherever talent happens to be.

1.4.2 Why Economic Sciences and Why Engineering

The program is jointly hosted by the Facultad de Ciencias Económicas and the Facultad de Ingeniería. This architecture is not an administrative compromise: it is the spine of the program. Its intellectual core rests on the integration between mathematical and computational rigor, the tradition of Ingeniería, and economic,

political, and institutional reasoning, the tradition of Ciencias Económicas. Neither faculty alone covers the disciplinary range the program requires; together they cover it naturally.

This kind of initiative, sustained by two large faculties, each with its own tradition and its own teaching body, is precisely what a large public university can offer and what private institutions, generally organized around a single faculty or a single disciplinary profile, are rarely positioned to mount. UBA has a rich history of cross-faculty collaboration, from the historic cátedras paralelas to contemporary inter-faculty institutes. The program inscribes itself in that institutional practice.

The dual home also gives the program a double credibility. A formation of this kind needs to be taken seriously both by the tradition of engineering and the exact sciences and by the tradition of economic analysis and public policy. If the program were solely an Ingeniería initiative, its graduates would be read as technologists with an interest in institutions; if it were solely a Ciencias Económicas initiative, they would be read as economists with an interest in technology. The program requires both readings at once, and therefore requires both homes at once.

1.4.3 The Argentine and Latin American Problem This Program Addresses

Argentina has historically produced individual scientists and technicians of the first rank. What has been harder, and this is a structural observation, not a critique of any particular period of government, is translating that individual capacity into durable technical and scientific institutions that outlast their founders and accumulate over time. The country has the human raw material; what is frequently missing is the institutional architecture that this raw material needs in order to become collective capacity.

The post-AI transition will be sharper in countries that import frameworks instead of designing them. AI governance, data infrastructure, regulatory regimes, and technical standards are being written in a small number of places in the world, and Latin America largely receives them pre-written. The consequence is predictable: the decisions that most affect the region are taken in other languages,

in other institutions, with other priorities. This program is explicitly a step toward forming the people capable of designing those frameworks locally, not as an act of rhetorical sovereignty, but as a concrete technical capacity.

Argentina also faces a specific set of institutional challenges, monetary instability, regulatory turnover, technology adoption gaps across sectors and across regions, that require people who simultaneously understand technical systems and institutional design. This profile is exactly the one the program seeks to form. And there is an additional reason, painful but pertinent: a significant share of Argentine graduates in engineering and computer science emigrate. Building a program that increases local leverage, the set of opportunities, networks, and projects that make it worthwhile to stay and build here, is also a modest contribution to that problem.

1.4.4 The Public University Mission

UBA's foundational mission, free, high-quality, accessible education in the service of national development, is not a slogan: it is the condition that makes this program possible. A formation that aspires to produce people capable of serving the public interest justifies itself as a public investment, not as a private product. That is why admission is independent of family income; that is why the small cohort size reflects a pedagogical decision, not a rationing mechanism by price; that is why the proposal is honest when it presents itself to a public university.

The program's outputs are aligned with that mission in a direct way. Its graduates are intended to enter public service, institutional design, technical leadership in Argentine firms, research, and the founding of organizations that the country needs. The knowledge produced along the way, the institutional designs, public-service prototypes, and research outputs generated in the Studios and integrative projects, is a public good and is made available as such.

Ultimately, the graduates' capacity to build durable Argentine institutions is itself a contribution to *la cosa pública*. A public university that forms people capable of strengthening the country from within is fulfilling its mission in the most

fundamental sense. This program, in its design and in its aspirations, understands itself as a contemporary expression of that historic UBA vocation.

1.5 Why Now

Every technological transition produces a window of opportunity for new institutions. The people and organizations that shape the transition, who define its norms, build its infrastructure, design its governance, are disproportionately formed in the early years of the transition, not after it has stabilized.

We are in that window now. The institutions, companies, and norms that will govern AI are being built in the next five to ten years. The people building them are being formed now. A program that begins admitting students today will produce its first graduates in three and a half years, precisely when the transition is at its most critical and most open to influence.

This is not a program for the world that exists. It is a program for the world that is being built.

The Contemporary Intellectual Context

2.1 The Conversation This Program Is Entering

The questions this program addresses, how to form people capable of governing complex systems, what institutions need to survive technological transitions, what remains irreducibly human when intelligence is abundant, are not only academic questions. They are being debated right now in the most serious publications operating at the intersection of ideas and power.

This program is a contribution to that conversation. Not through analysis but through formation. The full list of publications that constitute that intellectual context, *Foreign Affairs*, *Palladium*, *Le Grand Continent*, *Noema*, *American Affairs*, 421, is included in **Appendix D**, together with a short note on each one's relevance to the program.

2.2 What This Program Adds

These publications describe, analyze, and argue. They produce knowledge about the problems of the post-AI world. What they do not produce is people. This program's contribution to the conversation is not another analysis of what forming strategic capacity requires or another argument for why institutional capacity matters. Its contribution is the graduates themselves, people formed to exercise the judgment, construct the legitimacy, and bear the responsibility that the moment demands.

The Founding Argument

3.1 The Objective of the Program

The objective of the Licenciatura en Sistemas Complejos operates at three levels.

The immediate objective is to form people capable of building systems, technical, institutional, economic, that survive contact with reality in a world where artificial intelligence is abundant and cheap.

The deeper objective is to produce people who possess the two things that AI cannot yet replace: judgment under genuine uncertainty and the capacity to construct legitimacy. People who can decide what to build, not just how to build it. People who can take responsibility for outcomes that affect others.

The strategic objective is to begin forming the capacity the country needs to participate in, and shape, the institutions, companies, and governance frameworks of the post-AI world. Not inheriting the frameworks designed elsewhere but contributing to their design from Argentina.

In one sentence: to form builders with the mathematical rigor, scientific literacy, computational depth, institutional understanding, and philosophical grounding to create systems that matter, and the judgment and legitimacy to be trusted with them.

3.2 What AI Changes

Artificial intelligence changes the economics of knowledge. Tasks that required a human being with years of training can now be performed by a system trained on vast amounts of data. Previous waves of automation, mechanical, electrical,

digital, automated physical and routine cognitive labor. AI automates non-routine cognitive labor: the kind of work that previously required judgment, expertise, and flexibility.

The consequences for education are direct. If AI can write code, what should a computer science education teach? If AI can analyze legal documents, what should a law school teach? If AI can generate business plans, what should a business school teach? These questions do not have easy answers. But they have one common implication: the value of education cannot lie primarily in the transmission of content or the development of routine skills. It must lie in the development of capabilities that AI does not yet have and may not have for a long time.

3.3 What AI Does Not Change

Three things remain irreducibly human even as AI becomes more capable.

Judgment under genuine uncertainty. AI systems optimize for specified objectives. They are extraordinarily good at this. What they cannot do, yet, is decide what the objective should be when objectives are in conflict, when the consequences are irreversible, when the information is genuinely incomplete, and when other people's lives are affected by the outcome. That is judgment. It requires not just intelligence but wisdom, experience, and a capacity for moral reasoning that current AI systems do not possess.

Legitimacy. Institutions, organizations, and systems require the trust and cooperation of the people they govern. That trust is not granted automatically, it is earned through the quality of relationships, the consistency of behavior, the fairness of processes, and the credibility of the people involved. An AI system can optimize a governance mechanism. It cannot, by itself, make that mechanism legitimate. Legitimacy is a human achievement, constructed through culture, history, and the quality of human relationships.

Responsibility. When systems fail and other people bear the cost, someone must be accountable. Accountability requires human beings who can be held responsible, who can explain their decisions, bear the consequences, and make amends. AI systems cannot be held responsible in the morally relevant sense. As AI systems become more capable and more consequential, the humans who direct them must become more capable of bearing responsibility, not less.

3.4 The Formation Imperative

These three capacities, judgment, legitimacy, and responsibility, cannot be taught in the conventional sense. They cannot be transmitted through lectures or assessed through written exams. They are formed through practice, pressure, and accountability. Through building things and watching them fail. Through defending decisions under adversarial questioning. Through working in teams where coordination is hard and the stakes are real. Through being responsible for outcomes that affect other people.

This insight is the foundation of the program's pedagogical model. The program does not primarily teach. It forms. The difference is not semantic. Teaching transmits content. Formation develops character, judgment, and capability. The Studio system, six projects of increasing ambition, each involving real stakes and public accountability, is the mechanism through which formation happens.

The courses are not secondary. Mathematical rigor, physical intuition, computational depth, economic reasoning, philosophical clarity, these are necessary conditions for good judgment. A person who cannot reason formally cannot evaluate the assumptions in a model. A person who does not understand physics cannot reason about the physical constraints on a system. A person who has not thought seriously about epistemology cannot understand the limits of what they know. The courses build the intellectual foundation. The Studios build the judgment on top of it.

3.5 Formation for Public Responsibility

This program is designed as a public elite formation for high-impact responsibilities. That phrase requires precision.

Elite, in this context, does not mean private status, social privilege, or personal superiority over graduates of other programs. It means a selective, rigorous, public formation for a specific capacity: bearing the weight of decisions that affect many people. Every complex society needs people with that capacity, people who can take part in the design of its institutions, the running of its organizations, the construction of its infrastructure, and the navigation of crises that no institution fully anticipates. That capacity is not produced automatically. It is formed, through education, through experience, through the sustained pressure of high expectations.

Forming technical and institutional leaders is one of the oldest functions of universities, public and private. Oxford's PPE, France's grandes écoles, the great American research universities, all of them were designed, in part, to produce people capable of sustaining high-stakes decisions in complex societies. What none of them was designed to produce, and what the post-AI world needs, are people who combine technical depth with institutional literacy, scientific rigor with cultural understanding, and the capacity to build with the capacity to deliberate and account for the result. This program does not compete with the existing degrees at UBA: it fills a space that no single discipline addresses on its own.

The program admits a small cohort, between thirty and forty students, through a competitive selection process. That is a fact about the pedagogical structure: a public elite formation, built around Studios with public accountability, requires proximity, intensity, and a small-cohort environment. It is not a claim about the relative worth of its students compared to those of other programs. Its graduates will not all become officials or company directors. Some will build companies, others will design institutions, others will do research, others will work in public

administration, others will contribute to strengthening Argentine institutions from positions that do not yet exist. All of them will be prepared to take responsibility for systems that affect many people.

Program Identity and Positioning

4.1 What This Program Is

The Licenciatura en Sistemas Complejos is a degree for builders. Not builders in the narrow sense of software engineers or product managers, builders in the full sense of people who take a problem that exists in the world and convert it into a system that solves it. The system might be a company, an institution, a piece of software, a policy, an organization, or some combination of all of these. What matters is that it works, that it survives contact with reality, that it earns the trust of the people it serves, and that someone is accountable for it.

The program trains builders by integrating four dimensions that rarely appear together in one undergraduate degree: mathematical and scientific rigor, computational depth, institutional literacy, and philosophical grounding. These are not four separate tracks. They are four dimensions of a single formation. A graduate who works across them is prepared to approach problems that do not fit comfortably inside a single discipline.

4.2 What This Program Is Not

It is not a computer science degree. Computer science programs produce excellent programmers and researchers. They do not produce people who understand the institutional, cultural, and political dimensions of the systems they build. A computer science graduate can write the code for a platform that disrupts an industry. They are rarely equipped to understand the regulatory environment it will face, the cultural norms it will violate, the institutional resistances it will encounter, or the responsibility they bear for the consequences.

It is not an economics degree. Economics programs produce people who can model markets and analyze policy. They do not produce people who can build things. The gap between economic analysis and institutional practice is enormous, and most economics graduates never cross it.

It is not a political science degree. Political science programs produce analysts and, sometimes, practitioners of politics. They do not produce people with the technical depth to understand the systems, computational, financial, biological, that modern governance must regulate and manage.

It is not a business school. Business schools train managers of existing organizations. This program aims at a different task: preparing people who can create new institutions, products, and systems, reason about the foundations within which they operate, and take responsibility for outcomes that affect many people.

It is something new: a degree that learns from those traditions, adapts their useful mechanisms, and integrates them into a coherent formation for the post-AI world.

4.3 The PPE Comparison

One useful historical reference for what this program attempts is Oxford's PPE, Philosophy, Politics and Economics. PPE was created in 1920 as an interdisciplinary formation for people who would work in politics, public administration, diplomacy, journalism, and intellectual life. Its later influence shows the force of a simple idea: some public and institutional responsibilities require an education that crosses disciplines from the beginning.

PPE's insight was that governing complex societies requires fluency across disciplines, that a person who understands only politics without economics, or only economics without philosophy, is systematically unequipped for the decisions they will face. The integration of three disciplines into a single coherent formation was PPE's great contribution.

This program takes that insight into a different moment. The post-AI world requires fluency not only across the humanities and social sciences, but also across technical systems, science, economics, and institutions. It also requires a more explicit practical orientation: the capacity to build and sustain systems, not only to analyze them.

It is worth being precise about the scope of the comparison. What this program borrows from PPE is the **integrative ambition**: the conviction that some responsibilities require formation that crosses disciplines from the start, not early specialization. What this program does **not** borrow from PPE is its pedagogical mechanism. PPE works through Oxford's tutorial system, weekly one-on-one sessions where a student defends an essay to a tutor, a practice that produces analytical maturity by accumulation over three years. This program adopts a different mechanism, the Studio system and public defense, fitted to a different institutional context and to a world where building matters as much as arguing. The comparison is of ambition, not of mechanism.

If PPE was one twentieth-century answer to the formation of public judgment, this program asks what an analogous answer would require in the twenty-first century.

4.4 The Builder Orientation

The program's builder orientation is not a preference or a stylistic choice. It is a pedagogical and philosophical commitment grounded in a specific view of what knowledge is for.

Knowledge that cannot be applied to the construction of something real is incomplete. Not wrong, incomplete. The test of whether you understand a system is not whether you can describe it but whether you can build one, fix one, or identify precisely why it is failing. This is the insight behind the Studio system, behind the requirement that every evaluation take the form of a defense of something functional, and behind the program's insistence that AI tools be used aggressively rather than avoided.

In a world where AI can generate descriptions and analyses of almost unlimited sophistication, the ability to describe and analyze is no longer sufficient evidence of understanding. The ability to build, to make decisions about what to build, to take responsibility for the outcome, and to learn from the failure, is what distinguishes a person who genuinely understands from a person who merely appears to.

4.5 The Name

The degree is called **Licenciatura en Sistemas Complejos** in Spanish and **Bachelor of Systems and Decision Science** in English. The difference is deliberate.

"Sistemas Complejos" makes a claim about the world, that what matters now are systems that resist simple optimization, that behave differently at scale than their parts suggest, that fail in ways nobody predicted. The name is a bet on what kind of thinking the next decades will demand.

In English the translation is not literal. "Complex systems" in the Anglophone academic tradition carries a specific technical connotation, chaos theory, nonlinear dynamics, the Santa Fe Institute in its most mathematical register. That is part of what this program does but it is not all of it. "Systems and Decision Science" communicates the full scope: building systems, making decisions, and understanding the structures, technical, economic, institutional, cultural, within which both happen.

The two names point at the same program from different angles. Neither is the marketing version. Both are true.

Academic References and Inspiration

5.1 The Design Process

This program was not designed by starting from a blank page. It was designed by studying reference undergraduate and graduate programs and asking three questions about each: What does this program do especially well? What does it leave out by design? What can be learned from it when building a different formation?

The result is a program that lives in the space between its references: from PPE, the integration of philosophy, politics, and economics; from engineering and computer science, technical rigor; from MS&E, attention to technological, economic, and social systems; from Caltech and Polytechnique, scientific seriousness; and from studio and capstone models, the orientation toward building under real constraints.

5.2 Structural Inspiration

École Polytechnique, Cycle Ingénieur

A central structural reference is the Cycle Ingénieur of the École Polytechnique. Polytechnique's tronc commun runs in intensive blocks shorter than traditional semesters, with fewer courses taken concurrently but more courses completed in total. This program adapts that logic of concentration into an eight-week block structure. Each block is a complete intellectual unit. Students take three courses per block with the Studio project running as a continuous thread across two consecutive blocks.

Colorado College, The Block Plan

Colorado College's Block Plan offers another reference for the intensive block principle at undergraduate level. Students take one course at a time for 3.5 weeks, giving complete attention to a single subject before moving to the next. This program adapts the principle without copying the format: students take three courses per block rather than one, but the shared pedagogical idea is that concentration allows greater depth of work.

MBA Programs, Harvard, INSEAD, London Business School

Professional programs such as Harvard Business School, INSEAD, and London Business School have long used seven to eight week modules rather than traditional semesters. This structure allows for intensive immersion in a subject followed by a clean break before the next. For this program, those models serve as institutional evidence that a modular organization can sustain intensity without relying on the classic semester calendar.

5.3 Disciplinary Inspiration

Oxford PPE, Philosophy, Politics and Economics

As discussed in Section 4, PPE is an important historical reference for interdisciplinary formation at undergraduate level. This program takes its core insight, that governing complex societies requires fluency across disciplines, and reformulates it for an environment where technical systems, AI, and institutional design are inseparable. The comparison is not an attempt to copy PPE or surpass it, but to clarify the kind of integration this program is trying to build.

Stanford MS&E, Management Science and Engineering

MS&E is one of the disciplinary references closest to this program in content. Its declared mission is to advance the design, management, operation, and interaction of technological, economic, and social systems. Its curriculum combines optimization, probability, statistics, organizational theory, economics,

ethics, and computer science, all within an engineering school. This program learns from that combination and adapts it to a different purpose: an undergraduate degree with greater emphasis on AI, institutional design, philosophical foundations, and Studio projects as a central formative structure.

MIT 6-4, Artificial Intelligence and Decision Making

MIT's 6-4 program offers a reference for connecting artificial intelligence and decision theory within a mathematically rigorous formation oriented toward real systems. This program takes that proximity between AI, statistics, optimization, and decision as one basis for its mathematical and computational spine.

MIT 6-14, Computer Science, Economics, and Data Science

MIT's 6-14 program is a reference for combining computer science and economics through game theory, mechanism design, optimization, and statistical reasoning. This program adapts that connection in a sequence that runs from microeconomics through game theory and mechanism design to institutional design.

Stanford Symbolic Systems Program

Stanford's SymSys program is a reference for studying computation, logic, philosophy, psychology, and linguistics as facets of a common problem: the nature of mind and intelligence. This program's courses in formal reasoning, epistemology, and philosophy of AI adapt that sensibility to a formation more oriented toward systems and decisions.

Caltech, Core Curriculum and IDS

Caltech's insistence on deep mathematical and scientific foundations for all students, regardless of specialization, informs this program's requirements in physics, biology, mathematics, and probability. The difference is one of purpose: here those foundations serve a generalist formation oriented toward building and governing systems.

Santa Fe Institute, Complexity Science

The Santa Fe Institute is a central reference for complexity, emergence, networks, and adaptive systems, and for the idea that certain patterns appear across domains that traditional disciplines treat as separate. The program's course in *Sistemas Complejos*, and the broader conviction that graduates should recognize complexity patterns across domains, draw on that intellectual contribution.

UCL Integrated Engineering Programme

UCL's IEP offers a reference for project-based, interdisciplinary learning centered on teams as the primary unit of formation. The Studio system in this program adapts that idea: sustained project work appears not as a supplement, but as a structural part of the education.

Minerva University

Minerva is a reference for treating critical thinking frameworks, decision-making tools, and interdisciplinary application as an explicit part of the curriculum, not as a hoped-for byproduct of traditional courses. This program adapts that idea through core requirements in formal reasoning, decision theory, and epistemology.

5.4 Reference Table

Institution	Program	Primary Contribution to This Design
École Polytechnique	Cycle Ingénieur	8-week block structure, tronc commun, cohorte, intensity
Colorado College	Block Plan	Intensive block pedagogy, depth over breadth
Oxford	PPE	Interdisciplinary leadership formation, historical precedent
Stanford	MS&E	Technology + economics + institutions as unified discipline

MIT	6-4, 6-14	AI + decision theory; CS + economics as unified disciplines
Stanford	Symbolic Systems	Computation + philosophy + cognitive science integration
Caltech	Core + IDS	Mathematical and scientific foundations for all students
Santa Fe Institute	Complexity Science	Complexity as cross-domain unified discipline
UCL	IEP	Project-based formation, Studios as structural center
Minerva	Full program	Explicit thinking curriculum, decision frameworks as core
Harvard Business School	MBA	Case method, real decisions, 8-week module structure
INSEAD / LBS	MBA	8-week intensive module validation

Pedagogical Model

6.1 Formation, Not Teaching

The central pedagogical commitment of this program is the distinction between formation and teaching. Teaching transmits content. Formation develops judgment, character, and capability. Both are necessary. Neither is sufficient alone.

Traditional university education is optimized for teaching. Lectures, problem sets, written exams, these are efficient mechanisms for transmitting content and verifying that it has been retained. They are not effective mechanisms for developing the capacity to make good decisions under uncertainty, to build things that work, to lead teams under pressure, or to take responsibility for outcomes that affect other people.

This program is optimized for formation. The courses are necessary but not sufficient. The Studios are central. The oral defenses are central. The adversarial reviews are central. The requirement to build functional systems and defend them publicly is central. These are the mechanisms through which judgment is developed, not by being told what good judgment looks like, but by exercising it under conditions where the consequences are real and the accountability is inescapable.

6.2 The Block System

The eight-week block system is the structural expression of the formation philosophy. By concentrating attention on fewer subjects for shorter, more intensive periods, the block system creates the conditions for the kind of deep engagement that formation requires.

In a traditional semester system, a student taking five courses simultaneously must divide their attention five ways for sixteen weeks. The result is breadth without depth. The block system inverts this: three courses for eight weeks, with the Studio running continuously across two blocks. The load is informed by intensive programs such as MIT or Polytechnique, but the emphasis is on concentration and block-level accountability.

A typical week combines three courses, two long Studio sessions, a weekly progress review, lab or build time, and a decision log where each team records assumptions, AI use, tests performed, failures found, and next commitments. At the end of each block, students submit a defensible body of work; at the end of the double block, the Studio is defended before a panel.

The concrete load, in hours, is laid out below. Each course occupies between fourteen and twenty hours per week depending on character, technical or humanities, split across class, lab or discussion, and independent work. Each Studio occupies fourteen to sixteen hours per week across sixteen weeks. Three courses simultaneously plus the active Studio amount to roughly seventy hours per week, comparable to the intensity of programs such as Polytechnique or Caltech.

Component	Class	Lab/Discussion	Independent work	Weekly	Block/Studio total
Technical course (Calculus, Physics, ML, systems)	6	2	10–12	18–20	144–160
Humanities course (History, Ethics, Culture)	4	2	8–10	14–16	112–128
Studio	6	—	8–10	14–16	224–256

The block totals (8 weeks) and Studio totals (16 weeks) assume independent work is real work, not nominal self-study: technical reading, hard problems, debugging, defensible writing. The load is high and stated as such: the program is not compatible with a parallel full-time job.

The block boundaries also create natural checkpoints. Every eight weeks, a body of work must be completed and defended. There is no way to coast through a block and catch up at the end of the semester. The rhythm of production and accountability is sustained, and it is exactly this rhythm that builds the capacity to work under pressure, meet real deadlines, and produce something defensible on a fixed schedule.

6.3 The Studio System

The six Studios are the spine of the program. Each Studio runs for sixteen weeks, two consecutive blocks, and culminates in a public defense before a panel that includes people from outside the program. The Studios are not courses. They are projects with real deliverables, real teams, real constraints, and real accountability.

STUDIO I

Escritura y Oratoria, Writing and Rhetoric

Trains the fundamental communication skills that everything else in the program requires. Students cannot defend their work if they cannot write clearly and speak convincingly. Runs in Year 1, before the technical content has fully accumulated, because the communication skills must be developed early enough to be practiced throughout the rest of the program.

STUDIO II

Construir y Romper, Build and Break

Forces students to build something with a physical component, something that interacts with the world beyond the screen, and then attack what another team has built. The physical component is essential: it confronts students with the constraints that the physical world imposes on software and systems.

STUDIO III

Fiabilidad Bajo Presión, Reliability Under Pressure

Simulates operating a real system in production, defining service level objectives, managing incidents, writing postmortems, executing rollbacks. Develops operational and organizational skills: how to coordinate under pressure, how to communicate during a crisis, how to learn from failure.

STUDIO IV

Diseño de Sistemas Institucionales, Institutional Systems Design

Requires designing a system that crosses technical and institutional boundaries, a marketplace, a compliance platform, a public service. Students apply both technical depth and institutional literacy simultaneously for the first time at full scale.

STUDIO V

Sistema Crítico bajo Restricciones, Critical System Under Constraints

The most technically demanding project before the final Studio. Students build something that operates under real constraints, regulatory, financial, safety, or ethical. The constraint is the point. Building under constraints is what distinguishes a professional from a hobbyist.

STUDIO VI

Venture from Zero

The culminating Studio. Teams create and operate a real company over sixteen weeks, under genuine legal, financial, organizational, and regulatory constraints. They incorporate a company or advance as far as legally viable, build an MVP, seek real customers, validate demand, measure, iterate, and account for the decisions they made. The pedagogical justification is not startup culture or product survival: it is the integrated test of operating under real constraints, where technical design, market demand, finance, compliance, legitimacy, and leadership meet.

6.4 The Cohort as Formation Environment

The program admits 30 to 40 students per cohort. This is not a scaling constraint, it is a pedagogical requirement. Formation happens in relationship with other people. The pressure of working in a team with high expectations, the experience of defending your work before peers who understand it deeply enough to challenge it, the culture that develops when a small group of people share three and a half years of sustained intellectual pressure, these are the primary products of the small cohort, not its byproducts.

6.5 AI as Mandatory Tool

AI tools are mandatory throughout the program. Students are expected to use them constantly and fluently, for code generation, analysis, drafting, simulation, research, and anything else that accelerates their work. There is one rule, stated once and applied everywhere: you must be able to explain, defend, and take responsibility for everything you submit.

This approach reflects the program's foundational conviction: in a world where AI can produce sophisticated outputs in almost any domain, the relevant human capability is not the ability to produce outputs but the ability to evaluate, improve, and take responsibility for them.

6.6 No Electives Until the Final Semester

The program has no electives until the final semester, where one guided elective is offered. This is a design decision, not an omission. The program optimizes for cognitive coherence and a shared basis for judgment. Electives fragment that basis, they produce specialists who speak different languages too early.

The diversity of the cohort is not produced by different course choices. It is produced by subjecting people with different backgrounds and perspectives to the same rigorous common framework and observing how they make decisions under the same constraints. That is the diversity that matters for building systems.

Evaluation Framework

7.1 The AI Problem in Evaluation

If students can use AI freely, and in this program they must, then any written work done outside the classroom is essentially unverifiable. You cannot know how much of a take-home essay or problem set is the student's. This forces a fundamental rethink of what evaluation is for and how it should be conducted.

The answer is not to ban AI. The answer is to ensure that the primary evaluation happens in person, in real time, without preparation time for that specific problem. Everything done at home with AI is preparation for the real evaluation, not the evaluation itself.

7.2 The Principle

The only evaluations that are AI-proof are those that happen in person, in real time, on material the student could not have prepared specifically for. All primary evaluation in this program takes this form. Take-home work done with AI counts toward the grade but is preparation and evidence of engagement, not the primary test of understanding.

An important clarification. The program does not eliminate written exams. Mathematics, physics, probability, algorithms, and other theory-heavy courses are primarily evaluated by in-person written exam, open book, on unseen problems designed so that having the book open is irrelevant. What the program eliminates is the *comprehensive end-of-program written final*: there is no terminal written exam at graduation. Each block closes with its own primary evaluation, in person and in real time, in the modality appropriate to that course (written exam, oral defense, live coding, or public defense).

7.3 Open Book, Open Notes

All in-person evaluations are open book and open notes. No memorization is tested. The problems are designed so that having the book open is irrelevant, what matters is knowing how to think, which tool to apply, and why.

This is not a concession to students. It is a more honest and demanding form of assessment. A student who does not understand linear algebra will not be saved by having Strang on their desk. A student who does not understand thermodynamics will not be saved by having Feynman open in front of them. The book helps someone who already understands, it cannot substitute for understanding. This is the model used by the grandes écoles and by Caltech, where open-book exams with extremely hard problems are the standard.

7.4 Evaluation Types

Code	Type	Description
WE	Written Exam	In person, on paper, open book, open notes, no AI, no internet. Hard unseen problems requiring genuine reasoning.
TP	Take-home Work	Take-home work with full AI assistance. Counts toward grade. Primarily preparation for in-person evaluation.
LC	Live Coding	In-person programming session. AI permitted but evaluator questions every decision in real time.
OD	Oral Defense	Student defends written work under adversarial questioning with no preparation for specific questions.
PD	Public Defense	Defense before a panel including external reviewers. Used for Studios and major projects.
WP	Written Portfolio	Accumulated written work with full revision history. The writing itself is the evaluation.

7.5 The AI Rule

AI tools are mandatory. Students are expected to use them constantly and fluently for every take-home component of every course. There is one rule, applied without exception:

You must be able to explain, defend, and take responsibility for everything you submit.

The in-person evaluation is where this rule is enforced. A student who used AI without understanding will be exposed immediately under adversarial questioning. A student who used AI as a genuine tool for thinking will have no difficulty defending the work.

7.6 Evaluation by Course Type

Course Type	Primary Evaluation	Secondary Evaluation
Mathematics and Physics	Written Exam (WE), open book, hard unseen problems	Weekly problem sets (TP)
Programming and Systems	Live Coding (LC), AI permitted, explain every decision	Projects and portfolios (TP)
Quantitative Social Science	Written Exam (WE) or Oral Defense (OD)	Empirical studies (TP)
Humanities and Social Analysis	Oral Defense (OD), adversarial questioning	Essays with revision history (TP)
Design and Product	Public Defense (PD), real users or external panel	Prototype with documented research (TP)
Studios	Public Defense (PD), external panel	Full project built with AI
Writing	Written Portfolio (WP), revision history	Ongoing drafts and rewrites

Curriculum Overview

8.1 The Three-and-a-Half Year Arc

Año 1, Fundamentos. Year 1 builds the mathematical, scientific, and computational foundations on which everything else rests. By the end of Year 1, students can reason formally, compute, handle probability and statistics with rigor, build and deploy ML systems, program at the systems level, reason about the physical constraints on the systems they build, study living systems as complex adaptive systems, and have completed two team projects.

Año 2, Sistemas. Year 2 builds the systems, technical, economic, and institutional, within which builders must operate. By the end of Year 2, students understand algorithms and their limits, the mathematics of strategic interaction, the physics of risk and fat tails, the foundations of distributed computation, the architecture of financial systems, the political economy of regulation, historical modes of institutional failure, and the epistemological limits of AI.

Año 3, Síntesis. Year 3 integrates everything. Strategy, product, control, ethics, institutional design, negotiation, complex systems, decision theory, organizational behavior, information theory, product design, the synthesis courses where the mathematics of Year 1 and the systems of Year 2 are applied to the full complexity of building things that matter.

Año 3.5, Salida. The final semester is the closing and the opening simultaneously. Students complete Venture from Zero, take Culture, Legitimacy and Norms as a final synthesis on social acceptance, meaning, and authority, take their one guided elective, and participate in a founding seminar on AI, judgment, and responsibility, the philosophical argument of the program, encountered after three years of living it.

Why three and a half years

The 3.5-year length is intentional. The program treats undergraduate education as foundation, not specialization. Mastery of any single domain at depth, research-level economics, doctoral-grade computer science, terminal depth in any discipline, is the natural work of graduate study. Graduates of this program are explicitly prepared to compete for top international masters and doctoral programs in the disciplines they want to pursue further. The structure compresses what most programs spread across four to five years into three and a half, freeing graduates to specialize abroad while still in their early twenties. This positions the undergraduate degree as the start of a longer trajectory, not as a terminal qualification.

The Depth Model: Mastery, Literacy, and Integration

The obvious objection deserves a direct answer: no three-and-a-half-year program can produce simultaneous mastery in mathematics, physics, biology, computer science, machine learning, economics, law, politics, philosophy, product design, organizational behavior, and venture creation. The breadth is deliberate; the depth, by contrast, is calibrated honestly. The plan distinguishes explicitly between three levels of expectation, and that distinction is what makes the curricular load coherent.

Mastery. There is a bounded core of domains in which students reach genuine working depth, sufficient to build: mathematics (linear algebra, probability, optimization, abstract math), computer science (programming, systems, algorithms, distributed systems), machine learning and AI engineering, and mechanism and product design. In these domains the graduate operates as a practitioner, not as a spectator.

Literacy. A second set of domains is studied with reading depth and analytical rigor rather than with a specialist's vocation: economics and political economy, philosophy of science, cultural and institutional analysis, law and regulation, organizational behavior. The graduate can engage with the primary literature critically, evaluate arguments, and operate in interdisciplinary teams without

needing a specialist at hand for every decision. **Integration.** The Studios serve the specific function of forcing students to combine what they have mastered with what they can read, under realistic constraints. That is where breadth stops being a list of courses and becomes an effective capacity.

The model has clear precedents: École Polytechnique, Caltech, and Oxford's PPE programs each show, in their own register, how an undergraduate degree can cultivate mathematical maturity, scientific literacy, or integrative capacity without promising total specialization in every domain it touches. The program's honest claim is analogous. An MIT MSc graduate in computer science will go deeper in CS than a graduate of this program; an LSE MSc in economics will go deeper in economics. **This program's wager is different: to form someone with enough technical mastery to build, enough disciplinary fluency to integrate, and enough Studio-trained judgment to know which mode to apply when.** Those who wish to specialize further continue into graduate study (§12.2, pathway 5); that transition is not a flaw of the design but the natural trajectory of a generalist undergraduate degree oriented toward cross-domain integration.

8.2 The Full Curriculum

Año 1, Fundamentos

Bloque A	Bloque B	Bloque C	Bloque D
Álgebra Lineal y Computación I	Álgebra Lineal y Computación II	Probabilidad y Estadística	Machine Learning e Ingeniería de Modelos
Razonamiento Formal y Demostración	Microeconomía y Lógica de Decisión	Programación de Sistemas	Biología de Sistemas, Evolución y Ecología
Cálculo y Análisis	Fundamentos de Programación y Datos	Física I: Mecánica, Energía y Termodinámica	Física II: Electromagnetismo, Señales y Computación Física
<i>Studio I, Escritura y Oratoria (A+B)</i>		<i>Studio II, Construir y Romper (C+D)</i>	

Año 2, Sistemas

Bloque A	Bloque B	Bloque C	Bloque D
Algoritmos, Complejidad y Optimización	Teoría de Juegos y Diseño de Mecanismos	Redes y Sistemas Distribuidos	Economía Política, Derecho y Regulación
Econometría e Inferencia Causal	Procesos Estocásticos y Riesgo	Sistemas de Datos	Historia de las Instituciones, la Tecnología y el Poder
Deep Learning y Modelos Fundacionales	Métodos Formales y Verificación	Seguridad, Criptografía y Confianza	Epistemología, Ciencia y los Límites de la IA
<i>Studio III, Fiabilidad Bajo Presión (A+B)</i>		<i>Studio IV, Diseño de Sistemas Institucionales (C+D)</i>	

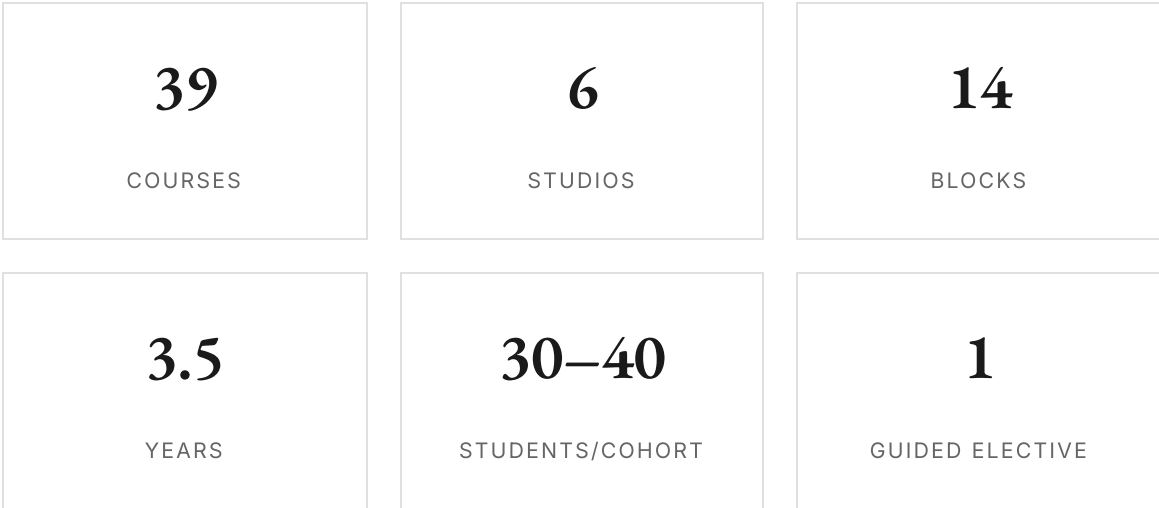
Año 3, Síntesis

Bloque A	Bloque B	Bloque C	Bloque D
Diseño y Estrategia de Producto	Laboratorio de Diseño Institucional	Sistemas Complejos	Métodos de Investigación, Medición y Diseño Experimental
Control, Observabilidad y Medición	Liderazgo, Negociación y Coordinación	Teoría de la Decisión	Finanzas, Dinero y Banca
Ética y Responsabilidad bajo Incertidumbre	Sistemas Embebidos e Interacción Física	Escritura Avanzada en Español	Optimización Aplicada y Métodos Numéricos
<i>Studio V, Sistema Crítico bajo Restricciones (A+B)</i>		<i>Studio VI, Venture from Zero (C+D)</i>	

Año 3.5, Salida

Bloque A	Bloque B
Seminario Fundacional: IA, Juicio y Responsabilidad	Venture from Zero, continuación y cierre
Cultura, Legitimidad y Normas	Seminario de Egreso: El Constructor en el Mundo
Electiva Guiada	Defensa Pública Final

8.3 Program Statistics



Course Descriptions

Año 1, Fundamentos

Álgebra Lineal y Computación I

Linear Algebra and Computation I

OBJETIVO / OBJECTIVE

To build the mathematical backbone of the program. Linear algebra is the language of data, machine learning, optimization, and quantitative reasoning. Students who do not master it at depth will hit a ceiling in every quantitative course that follows.

CONTENIDO / CONTENT

Vector spaces, linear transformations, column and null spaces, matrix factorizations (LU, QR), eigenvalues and eigenvectors, singular value decomposition. Foundations of probability through linear algebra: covariance matrices, PCA, expected values.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Strang, Gilbert. *Linear Algebra and Learning from Data*. Wellesley-Cambridge Press, 2019.
- Strang, Gilbert. *Introduction to Linear Algebra*. 6th ed. Wellesley-Cambridge Press, 2023.
- Axler, Sheldon. *Linear Algebra Done Right*. 4th ed. Springer, 2024.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 18.06 Linear Algebra](#)
- [MIT 18.065 Matrix Methods in Data Analysis](#)
- [Imperial College Mathematics for Machine Learning, Coursera](#)
- [3Blue1Brown: Essence of Linear Algebra, YouTube](#)

Álgebra Lineal y Computación II

Linear Algebra and Computation II

OBJETIVO / OBJECTIVE

From linear algebra to optimization and learning. The block where mathematics connects directly with machine learning.

CONTENIDO / CONTENT

Optimization: least squares, gradient descent, saddle points, constraints, basic convexity. Regularization and learning from data. Neural network structure as matrix operations. Backpropagation as the chain rule applied to computational graphs.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Strang, Gilbert. *Linear Algebra and Learning from Data*. Parts II–IV. Wellesley-Cambridge Press, 2019.
- Boyd, Stephen, and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 18.065 Matrix Methods in Data Analysis](#)
- [Stanford EE364A Convex Optimization](#)

EVALUACIÓN / EVALUATION

Razonamiento Formal y Demostración

Formal Reasoning and Proof

OBJETIVO / OBJECTIVE

To train precise reasoning and the discipline of falsification. To build the capacity to work rigorously in unfamiliar territory.

CONTENIDO / CONTENT

Propositional and predicate logic, proof techniques (direct, contradiction, induction), counterexample construction, basic set theory, relations.

Combinatorics and counting arguments (generating functions, pigeonhole principle, probabilistic method) treated as intensive applications of the proof techniques. AI may propose proofs, students find the gaps.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Velleman, Daniel J. *How to Prove It*. 3rd ed. Cambridge University Press, 2019.
- Sipser, Michael. *Introduction to the Theory of Computation*. 3rd ed. Cengage, 2012.
- Enderton, Herbert B. *A Mathematical Introduction to Logic*. 2nd ed. Academic Press, 2001.
- Graham, Ronald L., Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics*. 2nd ed. Addison-Wesley, 1994.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 6.1200 Mathematics for Computer Science](#)
- [Stanford Phil 151 Metalogic](#)
- [Stanford CS 103 Mathematical Foundations of Computing](#)

EVALUACIÓN / EVALUATION

WE, Written proof exam

OD, Oral defense of proof portfolio

Cálculo y Análisis

Calculus and Analysis

OBJETIVO / OBJECTIVE

To build the calculus and mathematical analysis that the rest of the program silently assumes. From limits and continuity to multivariable calculus and elements of real analysis, with enough rigor that probability, optimization,

physics, machine learning, and stochastic processes can rely on it rather than work around it.

CONTENIDO / CONTENT

Limits, continuity, derivatives, and Taylor approximation. Riemann integration. Sequences and series, convergence criteria. Multivariable calculus: gradients, Jacobians, Hessians, multiple integrals. Constrained optimization, Lagrange multipliers. Elementary ordinary differential equations. Closing unit on real analysis and topology: metric spaces, compactness, uniform convergence, continuity as structure preservation, fixed-point theorems.

The course is deliberately dense for a single block: eight weeks for material many programs spread over a full year. The compression is possible because it runs in parallel with Razonamiento Formal y Demostración, where the proof discipline rigorous calculus requires is built, and with Álgebra Lineal y Computación I, which provides the vector language. The goal is not to exhaust Tao, it is for Probabilidad, Procesos Estocásticos, Física, Machine Learning, and the optimization courses to be able to assume calculus rather than work around it.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Spivak, Michael. *Calculus*. 4th ed. Publish or Perish, 2008. *Single-variable calculus treated rigorously, written as an introduction to real analysis.*
- Tao, Terence. *Analysis I*. Hindustan Book Agency, 2016. *Real analysis built from the axioms of the reals, self-contained and rigorous.*
- Hubbard, John H., and Barbara Burke Hubbard. *Vector Calculus, Linear Algebra, and Differential Forms*. 5th ed. Matrix Editions, 2015. *Multivariable calculus with linear algebra and differential forms, the natural bridge into ML, optimization, and physics.*
- Strogatz, Steven H. *Nonlinear Dynamics and Chaos*. 2nd ed. CRC Press, 2014. *Selected chapters on ODEs and qualitative behavior, an early bridge to Sistemas Complejos.*
- Munkres, James R. *Topology*. 2nd ed. Pearson, 2000. *Selected chapters of point-set topology for the closing unit: metric spaces, compactness, continuity.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 18.01 Single Variable Calculus](#)
- [MIT 18.02 Multivariable Calculus](#)
- [MIT 18.100A Real Analysis](#)
- [Caltech Ma 1abc Mathematics](#)
- [École Polytechnique tronc commun \(modules de mathématiques\)](#)

EVALUACIÓN / EVALUATION

WE, Written exam with unseen problems and short proofs

TP, Weekly problem sets and proof exercises

Microeconomía y Lógica de Decisión

Microeconomics and Decision Logic

OBJETIVO / OBJECTIVE

To reason about scarcity, incentives, and tradeoffs before touching a business model. To develop economic intuition before the program demands formal modeling.

CONTENIDO / CONTENT

Utility, marginal reasoning, price formation, game theory fundamentals, market structure, market failures.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Varian, Hal R. *Intermediate Microeconomics*. 9th ed. W. W. Norton, 2014.
- Kahneman, Daniel. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.
- Schelling, Thomas C. *The Strategy of Conflict*. Harvard University Press, 1960.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 14.01 Principles of Microeconomics](#)
- [Yale ECON 159 Game Theory](#)
- [Stanford Econ 1](#)

EVALUACIÓN / EVALUATION

Fundamentos de Programación y Datos

Foundations of Programming and Data

OBJETIVO / OBJECTIVE

To write software that does what it says it does and handle basic data with judgment. Not coding speed, conceptual precision: understanding what code does, why it fails, how to verify it does what you believe, and how to model the data it uses.

CONTENIDO / CONTENT

Data structures, testing, specification-driven development. SQL, data modeling, basic normalization, joins, simple indexes, and small migrations. Shell, git, debugging, profiling, SSH, development environments. Specification first, implementation second, tests always. This course gives early contact with data; Data Systems treats them rigorously in Year 2.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Abelson, Harold, and Gerald Jay Sussman. *Structure and Interpretation of Computer Programs*. 2nd ed. MIT Press, 1996.
- Sedgewick, Robert, and Kevin Wayne. *Algorithms*. 4th ed. Addison-Wesley, 2011.
- McConnell, Steve. *Code Complete*. 2nd ed. Microsoft Press, 2004.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 6.100A Introduction to CS Programming in Python](#)
- [MIT Missing Semester](#)
- [Stanford CS 106B Programming Abstractions](#)

EVALUACIÓN / EVALUATION

Probabilidad y Estadística

Probability and Statistics

OBJETIVO / OBJECTIVE

To reason rigorously about uncertainty, evidence, and inference. The quantitative foundation on which ML, econometrics, stochastic processes, and every decision under uncertainty rests.

CONTENIDO / CONTENT

Probability spaces, random variables, distributions, expectation and variance, central limit theorem, Bayesian inference, estimation, hypothesis testing, confidence intervals. Entropy, mutual information, and KL divergence as measures of uncertainty, evidence, and updating.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Blitzstein, Joseph K., and Jessica Hwang. *Introduction to Probability*. 2nd ed. CRC Press, 2019.
- Wasserman, Larry. *All of Statistics*. Springer, 2004.
- Jaynes, Edwin T. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 6.3700 Introduction to Probability](#)
- [MIT 18.650 Fundamentals of Statistics](#)
- [Stanford CS 109 Probability for Computer Scientists](#)
- [Caltech Ma 3](#)

EVALUACIÓN / EVALUATION

WE, Written exam

TP, Weekly problem sets

Programación de Sistemas

Systems Programming

OBJETIVO / OBJECTIVE

To understand computation as a physical and temporal process. The focus is on the mental model, how computation actually happens in hardware and operating systems.

CONTENIDO / CONTENT

Memory hierarchy, concurrency, I/O, processes, scheduling, isolation, resource control, debugging under load and resource contention.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Bryant, Randal E., and David R. O'Hallaron. *Computer Systems: A Programmer's Perspective*. 3rd ed. Pearson, 2015.
- Arpaci-Dusseau, Remzi H., and Andrea C. Arpaci-Dusseau. *Operating Systems: Three Easy Pieces*. ostep.org.
- Kerrisk, Michael. *The Linux Programming Interface*. No Starch Press, 2010.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 6.1800 Computer Systems Engineering](#)
- [MIT 6.1810 Operating Systems Engineering](#)
- [Stanford CS 110 Principles of Computer Systems](#)

EVALUACIÓN / EVALUATION

LC, Live systems analysis

TP, Debugging portfolio and systems labs

Física I: Mecánica, Energía y Termodinámica

Physics I: Mechanics, Energy and Thermodynamics

OBJETIVO / OBJECTIVE

Physics for builders, not for physicists. To understand the physical world well enough to build things that interact with it, and to recognize when abstractions ignore real constraints.

CONTENIDO / CONTENT

Classical mechanics: forces, energy, momentum, rotation. Thermodynamics and entropy. Order-of-magnitude reasoning. Physical estimation: how much energy does it need? How much does it weigh? At what scale does it stop working?

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Feynman, Richard P., Robert B. Leighton, and Matthew Sands. *The Feynman Lectures on Physics*. Vol. I. Caltech, 2013. feynmanlectures.caltech.edu.
- Mahajan, Sanjoy. *The Art of Insight in Science and Engineering*. MIT Press, 2014.
- Halliday, David, Robert Resnick, and Jearl Walker. *Fundamentals of Physics*. 11th ed. Wiley, 2018.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 8.01 Classical Mechanics \(Walter Lewin\)](#)
- [Caltech Ph 1a Physics](#)
- [École Polytechnique tronc commun Physique](#)

EVALUACIÓN / EVALUATION

WE, Written exam with unseen problems

TP, Problem sets and estimation exercises

Física II: Electromagnetismo, Señales y Computación Física

Physics II: Electromagnetism, Signals and Physical Computation

OBJETIVO / OBJECTIVE

Electromagnetism and signal physics for builders. Enough quantum mechanics to understand semiconductors and the physical limits of computation.

CONTENIDO / CONTENT

Electromagnetism: fields, potential, circuits, waves. Signal physics: noise, frequency, spectrum. Basic quantum mechanics: semiconductors, physical limits of computation.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Feynman, Richard P., Robert B. Leighton, and Matthew Sands. *The Feynman Lectures on Physics*. Vol. II. Caltech, 2013.
- Griffiths, David J. *Introduction to Electrodynamics*. 4th ed. Cambridge University Press, 2017.

- Griffiths, David J. *Introduction to Quantum Mechanics*. 3rd ed. Cambridge University Press, 2018.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 8.02 Electricity and Magnetism](#)
- [MIT 8.04 Quantum Physics I](#)
- [Caltech Ph 1b](#)

EVALUACIÓN / EVALUATION

WE, Written exam

TP, Problem sets and physical system analysis

Machine Learning e Ingeniería de Modelos

Machine Learning and Model Engineering

OBJETIVO / OBJECTIVE

To build, train, evaluate, deploy, and monitor ML models, and to know when not to use them. The mathematical foundations were established in earlier blocks. This block teaches the engineering.

CONTENIDO / CONTENT

Data pipelines, feature engineering, training pipelines, evaluation methodology, overfitting, regularization, deployment, and monitoring. Data and model versioning, evaluation pipelines, offline/online validation, drift detection, rollback, privacy, security, fairness checks, model cards, incident response for model failures, and criteria for deciding when not to use ML.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Bishop, Christopher M. *Pattern Recognition and Machine Learning*. Springer, 2006.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. 2nd ed. Springer, 2009.
- Géron, Aurélien. *Hands-On Machine Learning*. 3rd ed. O'Reilly, 2022.
- Huyen, Chip. *Designing Machine Learning Systems*. O'Reilly, 2022.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 6.3900 Introduction to Machine Learning](#)
- [Stanford CS 229 Machine Learning](#)
- [Stanford CS 329S Machine Learning Systems Design](#)
- [Caltech CMS 155 Machine Learning and Data Mining](#)
- [Andrej Karpathy: Neural Networks Zero to Hero, YouTube](#)

EVALUACIÓN / EVALUATION

LC, Live ML pipeline session

TP, End-to-end ML project

Biología de Sistemas, Evolución y Ecología

Systems Biology, Evolution and Ecology

OBJETIVO / OBJECTIVE

To study living systems as complex adaptive systems: evolution, ecology, regulatory networks, immunity, epidemics, robustness, and fragility. This is not a generic introduction to biology but a lens for understanding adaptation, selection, feedback, coevolution, ecological collapse, and optimization under constraints.

CONTENIDO / CONTENT

Evolution by selection and adaptive dynamics. Genetics and cell biology only as needed to understand regulation, inheritance, and variation. Biological networks, feedback, and homeostasis. Immunity as a system of detection, memory, and response. Ecology, stability, resilience, trophic cascades, and collapse. Epidemiology and diffusion. Evolutionary game theory, cooperation, competition, and coevolution. Robustness, fragility, redundancy, tradeoffs, and path dependence. Careful use of biological analogies in institutions, AI, and systems design.

The goal is for students to internalize a domain of complexity not designed by humans. Physics teaches laws, constraints, and conservation; biology teaches adaptation, failure, selection, and survival under changing conditions. The course

requires models, estimates, and causal diagrams, not taxonomic memorization or a pre-medical orientation.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Carroll, Sean B. *The Serengeti Rules*. Princeton University Press, 2016.
- Maynard Smith, John. *Evolution and the Theory of Games*. Cambridge University Press, 1982.
- Nowak, Martin A. *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap Press, 2006.
- Levin, Simon A. *Fragile Dominion: Complexity and the Commons*. Perseus Books, 1999.
- Alberts, Bruce, et al. *Essential Cell Biology*. 5th ed. W. W. Norton, 2019.
Reference for cellular mechanisms, not the backbone of the course.
- Dawkins, Richard. *The Selfish Gene*. 40th Anniversary ed. Oxford University Press, 2016.

CURSOS DE REFERENCIA / REFERENCE COURSES

- MIT 7.013 Introductory Biology
- MIT 7.91J Foundations of Computational and Systems Biology
- Caltech Bi 1 Principles of Biology
- École Polytechnique tronc commun Biologie

EVALUACIÓN / EVALUATION

OD, Defense of analysis of a biological system

TP, Quantitative analysis with estimates and models

Año 2, Sistemas

Historia de las Instituciones, la Tecnología y el Poder

History of Institutions, Technology and Power

OBJETIVO / OBJECTIVE

To understand how organizations, states, markets, and transformative technologies actually work, and how they fail. Historical cases as stress tests of governance models, together with the study of how each major technological transition reorganized power, economy, and forms of life.

CONTENIDO / CONTENT

First half, institutions: institutional failure modes and their structural causes. Why empires collapse, why regulatory regimes calcify, why effective organizations become pathological. Second half, technology as engine of change: the printing press, steam engine, electricity, computer, internet, AI. How each technology changed what kinds of organizations are possible, who holds power, and what forms of life become imaginable.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Acemoglu, Daron, and James A. Robinson. *Why Nations Fail*. Crown, 2012.
- North, Douglass C. *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, 1990.
- Scott, James C. *Seeing Like a State*. Yale University Press, 1998.
- Hirschman, Albert O. *Exit, Voice, and Loyalty*. Harvard University Press, 1970.
- Olson, Mancur. *The Logic of Collective Action*. Harvard University Press, 1965.
- Dalio, Ray. *Principles for Dealing with the Changing World Order*. Simon & Schuster, 2021.
- Hobsbawm, Eric. *The Age of Revolution, The Age of Capital, The Age of Empire*. Vintage, 1962–1987. *The 1789–1914 trilogy, historical scaffolding for understanding how the modern world was formed by the dual revolution*.
- Mokyr, Joel. *The Lever of Riches*. Oxford University Press, 1990. *Technological creativity and economic progress as historical forces*.
- Chandler, Alfred D. *The Visible Hand*. Harvard University Press, 1977. *The managerial revolution and how organizations became the dominant form of coordination*.
- Perez, Carlota. *Technological Revolutions and Financial Capital*. Edward Elgar, 2002. *Long-cycle theory of technology, capital, and institutional adjustment*.

- Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age*. W. W. Norton, 2014.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 17.S914 Political Development](#)
- [MIT 15.137 Entrepreneurship and the Evolution of American Capitalism](#)
- [Oxford PPE Politics component](#)
- [Oxford History of Science and Technology Program](#)
- [Stanford HSST Program](#)
- [Harvard Kennedy School Governance and Institutions](#)

EVALUACIÓN / EVALUATION

OD, Adversarial oral defense of comparative essay

TP, Essay with revision history

Algoritmos, Complejidad y Optimización

Algorithms, Complexity and Optimization

OBJETIVO / OBJECTIVE

Structural reasoning about search, constraints, and the limits of computation. The central question: which problems are solvable, at what cost, and what does it mean for a problem to be hard.

CONTENIDO / CONTENT

Graphs, dynamic programming, complexity classes, computability limits, linear programming, convex optimization and duality. Applied optimization: integer programming, network flows, matching, scheduling, routing, queues, resource allocation, logistics, relaxations, heuristics, sensitivity, and formulation of real operational problems.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Cormen, Thomas H., et al. *Introduction to Algorithms*. 4th ed. MIT Press, 2022.
- Boyd, Stephen, and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

- Sipser, Michael. *Introduction to the Theory of Computation*. 3rd ed. Cengage, 2012.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 6.1210 Introduction to Algorithms](#)
- [MIT 6.C571 Optimization Methods](#)
- [MIT 15.053 Optimization Methods in Management Science](#)
- [Stanford CS 161 Design and Analysis of Algorithms](#)
- [Stanford MS&E 211 Introduction to Optimization](#)

EVALUACIÓN / EVALUATION

WE, Written exam with algorithm design problems

TP, Algorithm portfolio with formal proofs

Econometría e Inferencia Causal

Econometrics and Causal Inference

OBJETIVO / OBJECTIVE

To move from correlation to causality. The central question: under what conditions can one claim that X causes Y, and what identification strategy guarantees it.

CONTENIDO / CONTENT

Regression, instrumental variables, differences in differences, regression discontinuity, experimental design. A/B testing, statistical power, sample size, multiple hypotheses, measurement bias, external validity, surveys, instrument design, qualitative research as evidence, ethics of experiments, and critical reading of empirical results.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Angrist, Joshua D., and Jörn-Steffen Pischke. *Mostly Harmless Econometrics*. Princeton University Press, 2009.
- Cunningham, Scott. *Causal Inference: The Mixtape*. Yale University Press, 2021. mixtape.scunning.com.
- Pearl, Judea, and Dana Mackenzie. *The Book of Why*. Basic Books, 2018.

- [MIT 14.32 Econometric Data Science](#)
- [Stanford MS&E 226 Causal Inference](#)
- [Caltech IDS 126](#)

OD, Defense of empirical study

TP, Empirical study with real data

Deep Learning y Modelos Fundacionales

Deep Learning and Foundation Models

To understand in depth the architectures, mechanisms, and scaling laws of the AI systems that define this era, not as users, but as engineers and critics.

Transformer architectures, attention mechanisms, diffusion models, large language models, scaling laws, emergent capabilities, in-context learning, RLHF, fine-tuning, distillation, multimodality.

- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. deeplearningbook.org.
- Prince, Simon J.D. *Understanding Deep Learning*. MIT Press, 2023. udlbook.github.io.
- Vaswani et al. “Attention Is All You Need.” NeurIPS, 2017.
- Kaplan et al. “Scaling Laws for Neural Language Models.” arXiv, 2020.

- [Stanford CS 224N NLP with Deep Learning](#)
- [Stanford CS 236 Deep Generative Models](#)
- [MIT 6.S898 Deep Learning](#)
- [fast.ai Practical Deep Learning](#)

- [Andrej Karpathy: Neural Networks Zero to Hero, YouTube](#)

EVALUACIÓN / EVALUATION

LC, Live architecture session

OD, Defense of deep learning replication

TP, Replication at reduced scale

Teoría de Juegos y Diseño de Mecanismos

Game Theory and Mechanism Design

OBJETIVO / OBJECTIVE

From individual optimization to strategic interaction. To design systems where rational agents interact and incentives produce desired outcomes.

CONTENIDO / CONTENT

Game theory, Nash equilibria, auction theory, matching markets, incentive compatibility, social choice, computational complexity of mechanisms.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Nisan, Noam, et al. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- Roth, Alvin E. *Who Gets What, and Why*. Houghton Mifflin Harcourt, 2015.
- Mas-Colell, Andreu, et al. *Microeconomic Theory*. Oxford University Press, 1995. (game theory chapters)

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Stanford MS&E 232 Introduction to Game Theory](#)
- [MIT 14.126 Game Theory](#)
- [Yale ECON 159 Game Theory](#)
- [Harvard CS 1360 Economics and Computation](#)

EVALUACIÓN / EVALUATION

WE, Written exam

OD, Defense of mechanism design project

Procesos Estocásticos y Riesgo

Stochastic Processes and Risk

OBJETIVO / OBJECTIVE

To model path-dependent processes, detect non-ergodicity, and identify ruin scenarios. To understand why conventional risk metrics fail catastrophically in fat-tailed domains.

CONTENIDO / CONTENT

Markov chains, random walks, Poisson processes, Brownian motion. Fat-tailed distributions, the ergodicity problem, Kelly criterion, Black Swan dynamics. The second half treats fat tails as the central argument, not an appendix.

The integration of fat tails is not an appendix, it is the central argument of the second half of the course. Connects directly with Biology (ecological collapse is a fat-tailed event), Physics (thermodynamics of far-from-equilibrium systems), and Epistemology (the problem of induction is Taleb's central philosophical concern).

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Ross, Sheldon M. *Introduction to Probability Models*. 12th ed. Academic Press, 2019.
- Taleb, Nassim Nicholas. *Statistical Consequences of Fat Tails*. STEM Academic Press, 2020.
- Norris, J. R. *Markov Chains*. Cambridge University Press, 1997.
- Peters, Ole, and London Mathematical Laboratory papers on ergodicity. *The distinction between ensemble averages and time averages, if a process is non-ergodic, the average across the group does not describe the experience of the individual, which fundamentally changes how risk, inequality, and decisions should be analyzed.*
- Cirillo, Pasquale, and Nassim Nicholas Taleb. *The Logic of Risk* and joint papers on fat tails, wars, and pandemics. *Rigorous statistical formalization of the fat-tails framework with direct implications for policy and risk management.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Stanford MS&E 221 Stochastic Modeling](#)
- [MIT 6.7700 Discrete Stochastic Processes](#)

- [London Mathematical Laboratory: Ergodicity Economics](#)
- [Nassim Taleb's Real World Risk Institute](#)

EVALUACIÓN / EVALUATION

WE, Written exam

OD, Defense of risk analysis

Métodos Formales y Verificación

Formal Methods and Verification

OBJETIVO / OBJECTIVE

To prove mathematically that a system does what it says it does. If AI generates code, formal verification becomes more important, not less.

CONTENIDO / CONTENT

Formal specification, Lean 4, program verification, invariants, pre/postconditions, dependent types. Students verify properties of real programs.

*Connects directly with Razonamiento Formal (which provides the logic), Cálculo y Análisis (which provides mathematical maturity), and Programación de Sistemas (which provides the context of real software). With AI generating code, this discipline becomes **more important, not less.***

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Avigad, Jeremy. *Mathematics in Lean*. leanprover-community.github.io.
- Pierce, Benjamin C., et al. *Software Foundations*. softwarefoundations.cis.upenn.edu.
- Chlipala, Adam. *Certified Programming with Dependent Types*. MIT Press, 2013.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Carnegie Mellon 15-815 Interactive Theorem Proving](#)
- [MIT 6.820 Foundations of Program Analysis](#)
- [Lean 4 / Mathlib community curriculum](#)
- [Software Foundations](#)

Redes y Sistemas Distribuidos

Networks and Distributed Systems

OBJETIVO / OBJECTIVE

To design protocols that work under loss, delay, and adversarial conditions. To build systems that survive partial failures.

CONTENIDO / CONTENT

Network architecture, protocol design, distributed consensus, replication, consistency tradeoffs, fault tolerance, CAP theorem. Data pipelines and distributed storage.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Kurose, James F., and Keith W. Ross. *Computer Networking*. 8th ed. Pearson, 2021.
- Kleppmann, Martin. *Designing Data-Intensive Applications*. O'Reilly, 2017.
- Van Steen, Maarten, and Andrew S. Tanenbaum. *Distributed Systems*. 4th ed. distributed-systems.net.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 6.5820 Computer Networks](#)
- [MIT 6.5840 Distributed Systems](#)
- [Stanford CS 144 Introduction to Computer Networking](#)
- [Cambridge Distributed Systems](#)

EVALUACIÓN / EVALUATION

Seguridad, Criptografía y Confianza

Security, Cryptography and Trust

OBJETIVO / OBJECTIVE

To understand how trust is constructed technically and where it breaks.

CONTENIDO / CONTENT

Algebraic structures used operationally in cryptography: groups, rings, fields, modular arithmetic, enough number theory to understand RSA, Diffie-Hellman, and elliptic curves. Cryptographic foundations: symmetric/asymmetric, hashing, signatures, protocols. Threat modeling, system security, attack surfaces, trust architectures.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Katz, Jonathan, and Yehuda Lindell. *Introduction to Modern Cryptography*. 3rd ed. CRC Press, 2020.
- Anderson, Ross J. *Security Engineering*. 3rd ed. Wiley, 2020.
- Boneh, Dan, and Victor Shoup. *A Graduate Course in Applied Cryptography*. cryptobook.us.
- Artin, Michael. *Algebra*. 2nd ed. Pearson, 2010.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Stanford CS 255 Introduction to Cryptography](#)
- [MIT 6.5620 Computer Systems Security](#)
- [Coursera: Cryptography I & II \(Dan Boneh\)](#)

EVALUACIÓN / EVALUATION

LC, Live threat modeling session

TP, Threat model and security architecture

Economía Política, Derecho y Regulación

Political Economy, Law and Regulation

OBJETIVO / OBJECTIVE

To understand how regulation, taxation, legal liability, and public policy determine what can be built.

Regulatory frameworks, liability regimes, compliance strategy, institutional incentives, the interaction between law and technical architecture. Operational law for technological systems: data protection, privacy, intellectual property, open-source licensing, contracts, procurement, product liability, consumer protection, platform regulation, AI governance, audit, reporting, and compliance by design. Students work with conflicting frameworks, Hayek alongside Polanyi, Land alongside Srnicek and Williams.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Acemoglu, Daron, and James A. Robinson. *The Narrow Corridor*. Penguin Press, 2019.
- Lessig, Lawrence. *Code: And Other Laws of Cyberspace*. Basic Books, 1999.
- Tirole, Jean. *Economics for the Common Good*. Princeton University Press, 2017.
- Polanyi, Karl. *The Great Transformation*. Beacon Press, 1944. *The market is not natural, it was politically constructed. Essential counterpoint to Hayek.*
- Zuboff, Shoshana. *The Age of Surveillance Capitalism*. PublicAffairs, 2019.
- Land, Nick. Selections from *Fanged Noumena*. Urbanomic, 2011. *Accelerationism, intelligence as a process exceeding the human, capital as AI avant la lettre, taught as object of analysis, not doctrine.*
- Srnicek, Nick, and Alex Williams. *Inventing the Future*. Verso, 2015. *Take Land's premises and turn them leftward, the contrast is pedagogically potent.*
- Brown, Wendy. *Undoing the Demos*. Zone Books, 2015. *Neoliberalism as a rationality that converts everything into market logic, including politics and subjectivity.*
- Davidson, James Dale, and William Rees-Mogg. *The Sovereign Individual*. Touchstone, 1997. *Predicted in 1997 how technology would shift power between individuals and states, the framework remains remarkably useful for analyzing current dynamics.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 14.42 Environmental Policy and Economics](#)

- [Stanford Law and Technology](#)
- [Oxford PPE Politics and Philosophy components](#)

EVALUACIÓN / EVALUATION

OD, Adversarial defense of policy analysis

TP, Policy analysis with stakeholder map

Epistemología, Ciencia y los Límites de la IA

Epistemology, Science and the Limits of AI

OBJETIVO / OBJECTIVE

To understand what counts as knowledge, what makes a model valid, and what machines can and cannot do. The course that examines the program's founding claim: that AI cannot yet be trusted with judgment.

CONTENIDO / CONTENT

The course is organized around three questions: (1) what counts as valid scientific knowledge, and how does that question change over time? (2) what is intelligence, and what distinguishes it, if anything distinguishes it, between humans and machines? (3) what does a model hide when it explains something, and when does idealization stop being useful? The plan combines philosophy of science (falsifiability, paradigms, model validity) with philosophy of mind and AI (the Chinese Room, the symbol grounding problem, the alignment problem).

Philosophy of science is not abstract, it is operational for those who build systems. The course anchors on four core readings: Kuhn on how paradigms change, the Dreyfus–Dennett debate on what intelligence is, Cartwright on idealization in scientific models, and Russell on the alignment problem. Those four are the structural spine of the course; other authors appear as extensions or references, not as parallel reading. Connects directly with Métodos Formales (Lakatos), Machine Learning (Cartwright on models), and Sistemas Complejos (Polanyi on tacit knowledge).

SECONDARY READINGS

Selected excerpts, not full reading:

- Lakatos, *Proofs and Refutations*: how mathematics works as human activity; connects with Métodos Formales.
- Polanyi, *Personal Knowledge*: tacit knowledge and what an expert knows but cannot formalize.
- Christian, *The Alignment Problem*: recent history of the alignment problem.
- Feyerabend, *Against Method*: epistemological anarchism, critical complement to Kuhn.

FURTHER READING

Available for individual depth or Studio projects:

- Hacking, *Representing and Intervening*
- Latour, *Science in Action*
- Hui, *The Question Concerning Technology in China*
- Simondon, *On the Mode of Existence of Technical Objects*
- Haraway, *Cyborg Manifesto*
- Bostrom, *Superintelligence*

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Kuhn, Thomas S. *The Structure of Scientific Revolutions*. 4th ed. University of Chicago Press, 2012.
- Dreyfus, Hubert L. *What Computers Can't Do*. Harper & Row, 1972. *Core reading. The most powerful phenomenological critique of AI, was right about GOF AI, his arguments about embodiment remain unresolved by LLMs.*
- Dennett, Daniel C. *Consciousness Explained*. Little, Brown, 1991. *Core reading. The functionalist / deflationary position, direct counterpoint to Dreyfus; the Dreyfus–Dennett debate is the central axis of philosophy of AI.*
- Cartwright, Nancy. *How the Laws of Physics Lie*. Oxford University Press, 1983. *Core reading. Models idealize, they don't describe, crucial for ML builders.*
- Russell, Stuart. *Human Compatible*. Viking, 2019. *Core reading. The alignment problem from a founder of the field, why machines optimizing*

misspecified objectives are dangerous.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Stanford Symbolic Systems Program](#)
- [Stanford Phil 80 Mind, Matter, and Meaning](#)
- [MIT 24.09 Minds and Machines](#)
- [Oxford Philosophy of Mind](#)
- [Cambridge HPS](#)

EVALUACIÓN / EVALUATION

OD, Defense of philosophical analysis of a real AI system

TP, Written analysis

Sistemas de Datos

Data Systems

OBJETIVO / OBJECTIVE

To design systems where data survives real use: storage, querying, transactions, recovery, replication, migrations, governance, and operations. The course that makes explicit the substrate every venture, public service, ML system, and monitoring dashboard needs.

CONTENIDO / CONTENT

Relational model, advanced SQL, data modeling, normalization, and schema design. Indexes, query planning, storage, logs, transactions, isolation, concurrency, recovery, and backups. Replication, partitioning, consistency, analytical vs operational systems, data warehouses, lakehouses, pipelines, streaming, data quality, lineage, migrations, observability, access control, privacy, retention, governance, and data operations under incidents.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Kleppmann, Martin. *Designing Data-Intensive Applications*. O'Reilly, 2017.
- Ramakrishnan, Raghu, and Johannes Gehrke. *Database Management Systems*. 3rd ed. McGraw-Hill, 2003.

- Silberschatz, Abraham, Henry F. Korth, and S. Sudarshan. *Database System Concepts*. 7th ed. McGraw-Hill, 2019.
- Hellerstein, Joseph M., and Michael Stonebraker, eds. *Readings in Database Systems*. 5th ed. 2015. *Selections on architecture, transactions, distributed systems, and the evolution of the field*.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [CMU 15-445/645 Database Systems](#)
- [MIT 6.5830 Database Systems](#)
- [Stanford CS 145 Data Management and Data Systems](#)
- [Berkeley CS 186 Introduction to Database Systems](#)

EVALUACIÓN / EVALUATION

LC, Live data-system design and debugging session

TP, Schema, migration, and reliability project

Año 3, Síntesis

Finanzas, Dinero y Banca

Finance, Money and Banking

OBJETIVO / OBJECTIVE

To understand how things get funded, how money works, how capital is structured. First half: corporate finance. Second half: monetary systems, central banking, credit creation.

CONTENIDO / CONTENT

Financial statements, valuation, capital structure, unit economics, fundraising. Monetary systems, central banking, payments infrastructure. Bitcoin as a case study of distributed systems applied to money.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Brealey, Richard A., Stewart C. Myers, and Franklin Allen. *Principles of Corporate Finance*. 14th ed. McGraw-Hill, 2022.

- Mehrling, Perry. *The New Lombard Street*. Princeton University Press, 2011.
- Ammous, Saifedean. *The Bitcoin Standard*. Wiley, 2018.
- Alden, Lyn. *Broken Money*. Timestamp Press, 2023.
- Minsky, Hyman P. *Stabilizing an Unstable Economy*. McGraw-Hill, 2008.
Financial instability is endogenous to capitalism, not an external shock, the hedge / speculative / Ponzi phases explain every financial crisis since publication.
- Hayek, Friedrich A. “The Use of Knowledge in Society” and *Denationalisation of Money. The spontaneous order of the price system as a coordination mechanism for dispersed knowledge, and the radical proposal of private monetary competition.*
- Hobart, Byrne, and Tobias Huber. *Bubbles and the End of Stagnation*. Stripe Press, 2024. *Bubbles as engines of progress, not only destructive forces, how speculative energy drives technological adoption.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 15.401 Finance Theory](#)
- [Yale ECON 252 Financial Markets \(Robert Shiller\)](#)
- [Stanford MS&E 246 Financial Risk Analytics](#)

EVALUACIÓN / EVALUATION

OD, Defense of financial model and investment memo

TP, Financial model built with AI

Diseño y Estrategia de Producto

Product Strategy and Design

OBJETIVO / OBJECTIVE

To decide what to build and for whom, and to design it so people actually use it. The course covers strategy (market, positioning, distribution, advantage) and design (user research, prototyping, interaction, experimentation) as a single discipline, because in practice they are: strategy without design produces products no one uses, design without strategy produces products that do not scale.

CONTENIDO / CONTENT

Competitive analysis, segmentation, market sizing, beachhead market, positioning, differentiation, network effects, switching costs, economies of scale and learning, unit economics, pricing, channels, distribution, and launch strategy. User research, interviews, contextual observation, jobs-to-be-done, journey maps, personas as hypotheses, problem definition, and success criteria. Information architecture, flows, prototyping, usability, accessibility, onboarding, error handling, design for failure, and interface design with AI. Product analytics, experimentation, A/B testing, and iterative learning. Narrative for investors, users, and partners.

The fusion is deliberate: the usual split between strategy and design is an artifact of how business schools and design schools are organized, not how real products are built. The team that decides what to build is the team that decides how to build it. This course trains both decisions together. Connects directly with Studio VI.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Rumelt, Richard P. *Good Strategy Bad Strategy*. Crown Business, 2011.
- Christensen, Clayton M. *The Innovator's Dilemma*. Harvard Business Review Press, 1997.
- Ries, Eric. *The Lean Startup*. Crown Business, 2011.
- Norman, Donald A. *The Design of Everyday Things*. Basic Books, 2013.
- Cooper, Alan, Robert Reimann, and David Cronin. *About Face: The Essentials of Interaction Design*. 4th ed. Wiley, 2014.
- Bryar, Colin, and Bill Carr. *Working Backwards*. St. Martin's Press, 2021. *How Amazon unifies strategy, writing, operating mechanisms, and product design as one practice.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Stanford CS 183 Startup](#)
- [Stanford d.school Design Thinking](#)
- [MIT 15.835 Entrepreneurial Marketing](#)
- [MIT Integrated Design and Management](#)
- [Carnegie Mellon Human-Computer Interaction](#)

Control, Observabilidad y Medición

Control, Observability and Measurement

OBJETIVO / OBJECTIVE

To understand feedback, stability, and observability, and what measurement does to the things it measures. The course that fuses systems control with the epistemology and politics of metrics.

CONTENIDO / CONTENT

Control theory, stability analysis, controller design. Observability: what can and cannot be measured. Goodhart's law and Campbell's law, when a metric stops measuring what it claims to. The politics of metrics: perverse incentives, gaming, capture. The difference between measuring and managing. Epistemology of indicators: what a measurement constructs about the system measured. Resilience: graceful degradation, recovery, blast radius.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Åström, Karl J., and Richard M. Murray. *Feedback Systems*. Princeton University Press, 2008.
- Leveson, Nancy G. *Engineering a Safer World*. MIT Press, 2011.
- Meadows, Donella H. *Thinking in Systems*. Chelsea Green, 2008.
- Porter, Theodore M. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton University Press, 1995. *The historical construction of quantification as authority*.
- Muller, Jerry Z. *The Tyranny of Metrics*. Princeton University Press, 2018. *Why measurement-driven management systematically distorts what it claims to optimize*.

- Espeland, Wendy Nelson, and Mitchell L. Stevens. “Commensuration as a Social Process.” *Annual Review of Sociology*, 1998. *How making things commensurable is itself a political act.*
- Goodhart, Charles. “Problems of Monetary Management” (1975) and Strathern’s later restatement: “*When a measure becomes a target, it ceases to be a good measure.*”

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Stanford EE 263 Linear Dynamical Systems](#)
- [MIT 6.3100 Dynamical Systems and Control](#)
- [Caltech CDS 110 Introduction to Control](#)

EVALUACIÓN / EVALUATION

- | | |
|------------------|--|
| WE, Written exam | OD, Defense of system stability analysis |
|------------------|--|

Ética y Responsabilidad bajo Incertidumbre

Ethics and Responsibility under Uncertainty

OBJETIVO / OBJECTIVE

To work systematically through hard cases. A program that forms people to make decisions affecting others at scale needs a dedicated space for practical ethics.

CONTENIDO / CONTENT

Jonas, MacIntyre, Williams. Tragic tradeoffs are irreducible. Every important decision has an ethical dimension that clean frameworks cannot resolve.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Jonas, Hans. *The Imperative of Responsibility*. University of Chicago Press, 1984.
- MacIntyre, Alasdair. *After Virtue*. 3rd ed. University of Notre Dame Press, 2007.
- Williams, Bernard. *Ethics and the Limits of Philosophy*. Harvard University Press, 1985.
- Parfit, Derek. *Reasons and Persons*. Oxford University Press, 1984.

- [Oxford Practical Ethics](#)
- [Harvard Kennedy School Ethics and Public Policy](#)
- [Oxford PPE Philosophy component](#)

OD, Defense of ethical analysis of a hard case

TP, Written case analysis

Laboratorio de Diseño Institucional

Institutional Design Laboratory

OBJETIVO / OBJECTIVE

To design institutions that survive contact with real people. The course where students synthesize mechanism design, organizational theory, political economy, and legal reasoning into institutional architecture.

CONTENIDO / CONTENT

Institutional diagnosis: public or private problem, stakeholders, mandates, resources, legal constraints, existing capacities, and failure modes. Design of rules, incentives, decision mechanisms, access rights, accountability, and enforcement. Polycentric governance, commons, market design, procurement, regulation, legitimacy, and capture. Monitoring design, metrics, audit, appeal procedures, institutional maintenance, and transition plans. Each student produces an institutional architecture with explicit assumptions, abuse analysis, implementation strategy, and defense before practitioners.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Ostrom, Elinor. *Governing the Commons*. Cambridge University Press, 1990.
- Roth, Alvin E. *Who Gets What, and Why*. Houghton Mifflin Harcourt, 2015.
- Williamson, Oliver E. *The Economic Institutions of Capitalism*. Free Press, 1985.
- Simon, Herbert A. *The Sciences of the Artificial*. 3rd ed. MIT Press, 1996.
Bounded rationality, satisficing, design as discipline, Simon is the bridge between

AI, economics, and organizational theory.

- Ostrom, Vincent. *The Intellectual Crisis in American Public Administration*. University of Alabama Press, 1973. *Polycentricity, institutional design isn't choosing between hierarchy and market; there's a third path of polycentric governance.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Elinor Ostrom Workshop, Indiana University](#)
- [Santa Fe Institute Complexity and Governance](#)
- [Oxford Blavatnik School of Government](#)

EVALUACIÓN / EVALUATION

PD, Public defense before external panel with practitioners

TP, Institutional design document

Liderazgo, Negociación y Coordinación

Leadership, Negotiation and Coordination

OBJETIVO / OBJECTIVE

To get things done through other people when you have no authority and information is imperfect, and to coordinate organizations when incentives do not align on their own. The course combines interpersonal skill in negotiation with a structural view of how teams and organizations work.

CONTENIDO / CONTENT

Distributive and integrative negotiation, BATNA, ZOPA, anchors, concessions, packages, credible commitments, and focal points. Negotiation under incomplete information, power asymmetries, coalitions, mediation, and multi-party conflict. Written and oral persuasion: framing, audience, evidence, authority, trust, and timing. Leadership without formal authority, running difficult meetings, productive disagreement, and crisis communication. The organization as a coordination system: principal-agent, coordination costs, delegation, hiring, performance management, operating cadences, postmortems, and team dynamics

(trust, conflict, accountability, burnout). Cases of organizations that scale, bureaucratize, fragment, or fail under pressure. Simulations and diagnostics on real or simulated teams, with debrief.

The fusion is natural: you cannot lead what you do not structurally understand, nor negotiate well without knowing how counterparties work as organizations.

*Principal-agent explains **why** your team or your counterpart is not doing what you expect; negotiation and leadership techniques say **what to do about it**. The course connects directly with Studios IV, V, and VI, where real counterparties and organizational dynamics under pressure appear.*

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Fisher, Roger, William Ury, and Bruce Patton. *Getting to Yes*. Penguin Books, 2011.
- Voss, Chris. *Never Split the Difference*. Harper Business, 2016. *Negotiation from a former FBI hostage negotiator, immediately practical.*
- Schelling, Thomas C. *The Strategy of Conflict*. Harvard University Press, 1960. *Focal points and credible commitments, bridge between game theory and real negotiation.*
- Grove, Andrew S. *High Output Management*. Random House, 1983.
- Horowitz, Ben. *The Hard Thing About Hard Things*. Harper Business, 2014.
- MacIntyre, Alasdair. *After Virtue*. 3rd ed. University of Notre Dame Press, 2007. *Why leadership cannot rely on universal ethical frameworks and needs situated judgment.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Harvard Program on Negotiation](#)
- [MIT Sloan 15.665 Power and Negotiation](#)
- [Stanford MS&E 280 Organizational Behavior](#)
- [INSEAD Organisational Behaviour](#)

EVALUACIÓN / EVALUATION

Sistemas Embebidos e Interacción Física

Embedded Systems and Physical Interaction

OBJETIVO / OBJECTIVE

Build systems that interact with the physical world. The course that closes the gap between software and hardware: sensors, actuators, microcontrollers, real-time communication, and the physical constraints an embedded system imposes on design.

CONTENIDO / CONTENT

Microcontrollers and embedded platforms (Arduino, ESP32, Raspberry Pi). Sensors and instrumentation, actuators and control. Communication: I2C, SPI, UART, wireless networks. Real-time systems, latency, and determinism. Physical constraints: power consumption, thermal dissipation, electromagnetic robustness. Edge ML for systems with bounded compute budget. Connects with Física I and the physical component of Studio II.

The program is jointly hosted by Ciencias Económicas and Ingeniería; technical depth in physical systems is part of what the Engineering dimension commits to deliver. The course connects the physical foundations from Year 1 with the practice of building systems that matter in the real world, not only in the browser.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Lee, Edward A., and Sanjit A. Seshia. *Introduction to Embedded Systems: A Cyber-Physical Systems Approach*. 2nd ed. MIT Press, 2017. *The reference text for cyber-physical systems at university level; combines formal rigor with design practice.*
- Horowitz, Paul, and Winfield Hill. *The Art of Electronics*. 3rd ed. Cambridge University Press, 2015. *The canonical electronics text for engineers who need to design real circuits, not only analyze them.*

- Warren, Henry S. *Hacker's Delight*. 2nd ed. Addison-Wesley, 2012. *Bit-level operations, low-level optimization; useful for programming microcontrollers with limited resources.*
- Banzi, Massimo, and Michael Shiloh. *Getting Started with Arduino*. 3rd ed. Maker Media, 2014.
- Warden, Pete, and Daniel Situnayake. *TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers*. O'Reilly, 2019. *ML inference on microcontrollers; connects Deep Learning with real embedded systems.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 6.S08 Interconnected Embedded Systems](#)
- [Berkeley EECS 149 Introduction to Embedded Systems](#)
- [Stanford CS 149 Parallel Computing](#)
- [ETH Zurich Embedded Systems](#)

EVALUACIÓN / EVALUATION

PD, Public defense of a working embedded system with documented design decisions

TP, Functional prototype with hardware + firmware + measurement plan

Sistemas Complejos

Complex Systems

OBJETIVO / OBJECTIVE

To understand how order, chaos, and structure emerge from the interaction of simple agents, and why complex systems resist prediction and centralized control.

CONTENIDO / CONTENT

Emergent behavior, complex networks, agent-based models, phase transitions, self-organization, power laws, coevolution. Complexity patterns cross domains.

Connects everything: Biología (Año 1) with Economía (Año 1–2) with Redes (Año 2) with Instituciones (Año 2) with Riesgo y fat tails (Año 2). Complexity patterns cross domains, the cascade failure in a power grid, a bank run, an epidemic, and an institutional collapse are the same phenomenon viewed from different disciplines.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Mitchell, Melanie. *Complexity: A Guided Tour*. Oxford University Press, 2009.
- Barabási, Albert-László. *Network Science*. Cambridge University Press, 2016.
- Miller, John H., and Scott E. Page. *Complex Adaptive Systems*. Princeton University Press, 2007.
- West, Geoffrey. *Scale*. Penguin Press, 2017. *Universal scaling laws in biology, cities, and economies, the same mathematical patterns appear in radically different systems, from metabolic rates to urban innovation.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Santa Fe Institute Complexity Explorer](#)
- [MIT 18.354 Nonlinear Dynamics](#)
- [Oxford Complex Systems](#)
- [Barabási Network Science](#)

EVALUACIÓN / EVALUATION

OD, Defense of agent-based model

TP, Agent-based model of a real complex system

Teoría de la Decisión

Decision Theory

OBJETIVO / OBJECTIVE

To unify the formal machinery of decision-making under uncertainty. The course that formalizes judgment. AI can optimize a given objective function. The human question is: what is the right objective function?

CONTENIDO / CONTENT

Expected utility theory and its limits, Bayesian updating, value of information, sequential decisions, multi-criteria decision making, Knightian uncertainty, robust decision making.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Raiffa, Howard. *Decision Analysis*. Addison-Wesley, 1968.
- Savage, Leonard J. *The Foundations of Statistics*. Wiley, 1954. *The foundational axioms of decision under uncertainty*.
- Kahneman, Daniel. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.
- Gigerenzer, Gerd. *Gut Feelings: The Intelligence of the Unconscious*. Viking, 2007. *Fast heuristics that outperform formal models in uncertain environments, empirical counterpoint to expected utility theory*.
- Tetlock, Philip E., and Dan Gardner. *Superforecasting*. Crown, 2015. *The disciplined practice of calibrated prediction, connects directly with Bayesian updating*.
- Simon, Herbert A. *Models of Bounded Rationality*. MIT Press, 1982. *Why real agents satisfice rather than optimize*.
- Williams, Bernard. *Ethics and the Limits of Philosophy*. Harvard University Press, 1985. *Tragic tradeoffs are irreducible, every important decision has an ethical dimension that clean frameworks cannot resolve*.
- Jonas, Hans. *The Imperative of Responsibility*. University of Chicago Press, 1984. *Ethics for technological societies, the principle of responsibility toward future generations as a constraint on present decisions*.

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Stanford MS&E 250 Decision Analysis](#)
- [Harvard Kennedy School Decision Science](#)
- [Oxford PPE Philosophy component](#)

EVALUACIÓN / EVALUATION

Escritura Avanzada en Español

Advanced Writing in Spanish

OBJETIVO / OBJECTIVE

To master analytical and argumentative writing in Spanish. For people who will operate in Latin America, the ability to write precisely in Spanish is not secondary, it is central.

CONTENIDO / CONTENT

Analytical essay, executive memo, technical note, institutional report, short speech, public column, and explanation for non-specialist audiences. Argument structure, thesis, evidence, rhythm, lexical precision, conceptual translation between technical jargon and public Spanish, line editing, and deep revision. AI use for drafts, counterarguments, and editing, with full authorial responsibility. The portfolio requires multiple versions of the same text, an editorial decision log, and oral defense of style, structure, and evidence.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Borges, Jorge Luis. Selected essays on precision and style.
- Walsh, Rodolfo. Selected political and journalistic writing.
- Martínez Estrada, Ezequiel. *Radiografía de la Pampa*. Losada, 1933.
- Cortázar, Julio. Selected essays on structure and form.

EVALUACIÓN / EVALUATION

WP, Written portfolio with full revision history

Métodos de Investigación, Medición y Diseño Experimental

Research Methods, Measurement and Experimental Design

OBJETIVO / OBJECTIVE

To build, measure, and defend claims with evidence. The course that operationalizes the question 'how do I know this is true?' Essential when AI can generate convincing text but cannot generate evidence.

CONTENIDO / CONTENT

Formulating hypotheses and tractable research questions. Classical experimental design: control, randomization, replication, balance, power, and sample size. Quasi-experimental designs: regression discontinuity, difference-in-differences, instrumental variables at the project level. Internal, external, construct, and measurement validity. Operationalization: what is being measured, and what is actually being measured? Reproducibility, pre-registration, open science, open data, open code, transparent reporting. Critical reading of papers, identifying weak studies, meta-analysis. Scientific communication: paper structure, peer review, response to reviewers. Research ethics: consent, IRB, integrity, conflicts of interest. Cases: A/B tests in product work, policy evaluation, clinical trials, replication of findings in ML.

The course is operational, not philosophical: Epistemología (Year 2) raises the limits of knowledge, this course builds the tools. Paired with Control, Observabilidad y Medición (Year 3 A): that one looks at measurement from the controlled system; this one looks at measurement from the standpoint of an evidence claim. Connects with every Studio: defensible hypotheses, measurements that count, replicability of results, honest postmortems.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Shadish, William R., Thomas D. Cook, and Donald T. Campbell. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Houghton Mifflin, 2002.
- Montgomery, Douglas C. *Design and Analysis of Experiments*. 10th ed. Wiley, 2019.
- Booth, Wayne C., Gregory G. Colomb, Joseph M. Williams, Joseph Bizup, and William T. FitzGerald. *The Craft of Research*. 4th ed. University of Chicago Press, 2016.

- Bergstrom, Carl T., and Jevin D. West. *Calling Bullshit: The Art of Skepticism in a Data-Driven World*. Random House, 2020. *Critical reading in the age of AI and mass data: how to distinguish signal from confident bullshit.*
- Ritchie, Stuart. *Science Fictions: How Fraud, Bias, Negligence, and Hype Undermine the Search for Truth*. Metropolitan Books, 2020. *The contemporary replication crisis, context for why methods matter operationally.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Stanford GSB Methodology of Empirical Research](#)
- [MIT 14.310x Data Analysis for Social Scientists](#)
- [Berkeley D-Lab Computational Social Science](#)
- [Center for Open Science training](#)

EVALUACIÓN / EVALUATION

OD, Defense of a research design proposal with measurement plan and pre-analysis

TP, Replication audit of a published study

Optimización Aplicada y Métodos Numéricos

Applied Optimization and Numerical Methods

OBJETIVO / OBJECTIVE

To model real problems as optimization problems and solve them numerically with responsibility: knowing which method applies, what assumptions it demands, what errors it introduces, and when the result is reliable. The course that provides the technical backbone that ML, finance, control, operations, and research assume.

CONTENIDO / CONTENT

Numerical linear algebra: stability, conditioning, decompositions (LU, QR, SVD, Cholesky), iterative methods. Convex and nonconvex optimization: gradient, Newton, quasi-Newton, interior-point methods, constraints, duality, Lagrange multipliers. Linear and integer programming, simplex, branch and bound, relaxations. Network flows, scheduling, routing, assignment, transportation

problems. Stochastic methods: Monte Carlo, MCMC, discrete-event simulation, stochastic optimization. Sensitivity analysis, robustness, numerical stability, and computational pitfalls. Cases: resource planning, decisions under real constraints, model calibration, simulation of operational and physical systems.

*The course is deliberately practical: the goal is not to become a professional optimizer, but for the student to be able, when an operational problem appears (route planning, resource allocation, model calibration, system simulation), to model it, choose the method, run it numerically without fooling themselves, and judge the solution. Connects with *Algoritmos, Complejidad y Optimización* (which provides the theoretical base), *Procesos Estocásticos y Riesgo* (the probabilistic dimension), *Finanzas* (portfolio and risk models), *Machine Learning* (training as optimization), and the Studios where real physical and operational constraints appear.*

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Nocedal, Jorge, and Stephen J. Wright. *Numerical Optimization*. 2nd ed. Springer, 2006.
- Trefethen, Lloyd N., and David Bau III. *Numerical Linear Algebra*. SIAM, 1997.
- Bertsimas, Dimitris, and John N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, 1997.
- Hillier, Frederick S., and Gerald J. Lieberman. *Introduction to Operations Research*. 11th ed. McGraw-Hill, 2021.
- Boyd, Stephen, and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. *Text shared with *Algoritmos, Complejidad y Optimización*; here the emphasis is on numerical methods and application.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [MIT 15.053 Optimization Methods in Management Science](#)
- [Stanford EE364B Convex Optimization II](#)
- [CMU 10-725 Convex Optimization](#)
- [Caltech ACM 113 Mathematical Optimization](#)

Año 3.5, Salida

Cultura, Legitimidad y Normas

Culture, Legitimacy and Norms

OBJETIVO / OBJECTIVE

To understand the informal structures that sustain systems, or tear them apart. Popular culture is where a society processes its anxieties about the future.

CONTENIDO / CONTENT

Analysis of how cultural norms, legitimacy claims, and social expectations shape institutional behavior. Includes anime and science fiction canon treated as primary analytical material: Neon Genesis Evangelion, Ghost in the Shell, Foundation, Neuromancer.

These works are not entertainment: they are cultural artifacts that articulate the tensions between technology, institutions, identity, and power in ways academic papers cannot. AI writes convincing-but-shallow cultural analyses, students expose missing assumptions and connect cultural narratives to systemic stability.

PRIMARY CULTURAL CANON

Anime:

- *Neon Genesis Evangelion*, technology, dysfunctional institutions (NERV), individual responsibility under institutional pressure, the limits of human control over systems we don't understand.
- *Serial Experiments Lain*, identity, reality, and the dissolution of the boundary between the digital and the human; philosophy of mind disguised as anime.
- *Ghost in the Shell*, consciousness in a world of artificial bodies; connects directly with Epistemología and the Limits of AI.

- *Ergo Proxy*, post-collapse societies, authoritarian governance, what it means to be human when machines are conscious.
- *Akira*, technology, state power, and destruction.

Science fiction:

- Asimov, *Foundation*, psychohistory: can the behavior of complex systems be predicted and designed at societal scale? Connects directly with *Sistemas Complejos*.
- Asimov, "The Last Question", entropy, information, and the limits of computation.
- Gibson, *Neuromancer*, invented the concept of cyberspace; defines the aesthetic of the human-machine relationship.
- Dick, *Do Androids Dream of Electric Sheep?* / *Blade Runner*, empathy as a test of humanity.
- *Black Mirror* (selected episodes), unintended consequences of technology.
- *The Matrix*, simulation, reality, and the Chinese Room turned into film.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Douglas, Mary. *How Institutions Think*. Syracuse University Press, 1986.
- Bourdieu, Pierre. *The Logic of Practice*. Stanford University Press, 1990.
- Henrich, Joseph. *The WEIRDest People in the World*. Farrar, Straus and Giroux, 2020.
- Fisher, Mark. *Ghosts of My Life*. Zero Books, 2014. *Hauntology, culture as the repetition of cancelled futures; the theoretical frame for analyzing why we keep processing the same technological anxieties.*
- Han, Byung-Chul. *The Burnout Society and Psychopolitics*. Stanford University Press, 2015 / Verso, 2017. *Control is no longer disciplinary but operates through self-exploitation and performance, students will live this.*
- Girard, René. *Violence and the Sacred*. Johns Hopkins University Press, 1977. *Mimesis, scapegoats, the founding violence of every culture, explains social media dynamics and consensus formation better than any contemporary sociologist.*

- Illich, Ivan. *Tools for Conviviality*. Harper & Row, 1973. *Institutions become counterproductive past a certain threshold, meta-relevant: students are in an educational program that should ask itself this question.*
- Debord, Guy. *The Society of the Spectacle*. Zone Books, 1994. *The total mediation of experience by images, written in 1967, describes 2025 better than any contemporary analysis.*
- Weber, Max. *The Protestant Ethic and the Spirit of Capitalism*. 1905. *How religious ideas shaped economic behavior, foundational for understanding the cultural bases of modern capitalism.*
- Siedentop, Larry. *Inventing the Individual*. Belknap Press, 2014. *Genealogy of Western individualism from early Christianity, essential context for debates on consciousness, subjectivity, and autonomy.*
- O’Neil, Cathy. *Weapons of Math Destruction*. Crown, 2016. *How algorithms encode and amplify inequality, the culture of algorithmic opacity.*
- Zuboff, Shoshana. *The Age of Surveillance Capitalism*. PublicAffairs, 2019. *The extraction of behavioral data as a new form of capital accumulation, the culture of normalized surveillance.*

CURSOS DE REFERENCIA / REFERENCE COURSES

- [Stanford Anthropology and Cultural Analysis](#)
- [MIT Comparative Media Studies](#)
- [École Polytechnique HSS, Humanités et Sciences Sociales](#)
- [Cambridge Social Anthropology](#)

EVALUACIÓN / EVALUATION

OD, Defense of critical essay connecting cultural artifact to institutional problem

TP, Essay with full revision history

Seminario Fundacional: IA, Juicio y Responsabilidad

Founding Seminar: AI, Judgment and Responsibility

OBJETIVO / OBJECTIVE

The founding argument of the program, given explicitly after three years of living it. What AI changes, what it does not, what judgment is, why legitimacy matters.

CONTENIDO / CONTENT

Final synthesis of the program. What AI automates and what it displaces in the knowledge economy. Judgment: decision under uncertainty, misspecified objectives, responsibility for consequences, and the capacity to say no. Legitimacy: trust, authority, public explanation, institutions, and consent of affected people. Alignment, specification failures, delegation to opaque systems, automation bias, and accountability. Case discussions: automated public systems, foundation models, medical or financial decisions, critical infrastructure, and ventures built in Studio VI. The seminar culminates in a publicly defended position on what must remain a human responsibility.

TEXTOS PRINCIPALES / PRIMARY TEXTS

- Russell, Stuart. *Human Compatible*. Viking, 2019.
- Dreyfus, Hubert L. *What Computers Can't Do*. Harper & Row, 1972.
- Jonas, Hans. *The Imperative of Responsibility*. University of Chicago Press, 1984.
- Christian, Brian. *The Alignment Problem*. W. W. Norton, 2020.

EVALUACIÓN / EVALUATION

OD, Oral defense of position paper

TP, Position paper on what AI changes and what it does not

Electiva Guiada

Guided Elective

OBJETIVO / OBJECTIVE

An elective course in an area of deepening chosen by the student in consultation with the faculty. Format, texts, and evaluation defined by the supervising faculty member.

EVALUACIÓN / EVALUATION

Defined by supervising faculty

Studio Descriptions

Studios are the spine of the program. They are not courses. They are sixteen-week projects with real deliverables, real teams, real constraints, and real accountability. All Studios culminate in a public defense before a panel that includes external reviewers.

STUDIO I, AÑO 1, BLOQUES A+B

Escritura y Oratoria / Writing and Rhetoric

OBJETIVO / OBJECTIVE

To formally train the written and oral communication skills on which everything else in the program rests.

Concrete example. A 4000-word essay arguing a contested position on AI governance, with full revision history (drafts, peer feedback, rewrites) and a recorded oral defense before a panel that includes at least one external reviewer.

ESTRUCTURA / STRUCTURE

Writing emphasis in the first eight weeks. Oral emphasis in the second eight weeks. Both practiced throughout. Material comes from courses being taken simultaneously: students write and present about real content, not abstract exercises.

ENTREGABLES / DELIVERABLES

- Portfolio of rewritten essays with complete revision history: original drafts, feedback received, rewrites
- Recorded oral presentations with self-evaluation

- Accountability memo on claims, evidence used, and limits of what the team can defend
- Final oral defense without notes

EVALUACIÓN / EVALUATION

WP, Written portfolio with revision history

Recorded presentations with self-evaluation

CRITERIOS / CRITERIA

- Clarity of argument and traceability of evidence
- Ability to recognize uncertainty and limits
- Responsibility for claims that affect others

PROGRAMAS DE REFERENCIA / REFERENCE PROGRAMS

Stanford PWR Program in Writing and Rhetoric · MIT CI-H
Communication Intensive · Harvard Bok Center Writing Program

STUDIO II, AÑO 1, BLOQUES C+D

Construir y Romper / Build and Break

OBJETIVO / OBJECTIVE

Build a system with a mandatory physical component. Another team attacks it. Roles rotate.

Concrete example. A working sensor + actuator system that responds to a physical input, for example a moisture-sensing irrigation controller for a public park, accompanied by a threat model documenting at least three failure modes and a postmortem written by the team that attacked it.

ESTRUCTURA / STRUCTURE

Teams of 4-5 students. One team builds, another attacks. Roles rotate at the midpoint. The physical component is mandatory: the system must measure, move, or interact with something tangible. AI assists both sides.

ENTREGABLES / DELIVERABLES

- Specification document
- Test plan
- Functional prototype with physical component
- Measurement log
- Threat model
- Incident report
- Failure report
- Postmortem and team retrospective

EVALUACIÓN / EVALUATION

PD, Public defense of functional system

CRITERIOS / CRITERIA

- Match between specification, tests, and observed behavior
- Quality of measurement and instrumentation
- Explicit treatment of failure modes, affected users, and physical risks

PROGRAMAS DE REFERENCIA / REFERENCE PROGRAMS

MIT 6.S08 Interconnected Embedded Systems · Caltech ME/EE 75
Projects in Experimental Engineering · UCL Integrated Engineering
Programme

STUDIO III, AÑO 2, BLOQUES A+B

Fiabilidad Bajo Presión / Reliability Under Pressure

OBJETIVO / OBJECTIVE

Operate a system under simulated production conditions. Define SLOs, build monitoring, manage incidents, execute rollbacks, write postmortems.

Concrete example. A monitored web service running for 8 weeks under simulated production load, with explicit SLOs, an on-call schedule, at least one staged incident with a written postmortem, and a rollback procedure documented and rehearsed.

ENTREGABLES / DELIVERABLES

- Service Level Objectives (SLOs)
- Reliability test plan
- Monitoring dashboard
- Incident logs
- Documented and rehearsed rollback procedure
- Postmortems
- Team coordination retrospective

EVALUACIÓN / EVALUATION

PD, Defense simulating a post-incident review

CRITERIOS / CRITERIA

- Realism of SLOs and alerts
- Quality of incident response and communication with affected users
- Technical honesty of the postmortem and preventive measures

PROGRAMAS DE REFERENCIA / REFERENCE PROGRAMS

Google Site Reliability Engineering methodology · MIT distributed systems labs

STUDIO IV, AÑO 2, BLOQUES C+D

Diseño de Sistemas Institucionales / Institutional Systems Design

OBJETIVO / OBJECTIVE

Design a system that crosses technical and institutional boundaries. AI simulates the behavior of regulators, attackers, and users.

Concrete example. A complete institutional design document for a real-world problem, for example a digital identity system for a hypothetical municipal government, including incentive analysis, stakeholder map, regulatory feasibility assessment, and a red-team review by another team.

ENTREGABLES / DELIVERABLES

- Requirements document
- Incentive analysis
- Stakeholder and affected-user map
- Partial formal specification
- Monitoring plan
- Abuse and institutional-capture analysis
- Financial viability assessment
- Public defense

EVALUACIÓN / EVALUATION

PD, Public defense before panel including practitioners from the relevant institutional domain

CRITERIOS / CRITERIA

- Coherence between technical design, incentives, and institutional constraints
- Treatment of legitimacy, affected stakeholders, and possible abuses
- Operational, financial, and regulatory viability

PROGRAMAS DE REFERENCIA / REFERENCE PROGRAMS

Oxford Blavatnik School of Government capstone projects · Harvard Kennedy School policy exercises

STUDIO V, AÑO 3, BLOQUES A+B

Sistema Crítico bajo Restricciones / Critical System Under Constraints

OBJETIVO / OBJECTIVE

Design and partially build a system that operates under real constraints, regulatory, financial, safety, or ethical. The constraint is the point.

Concrete example. A partially-implemented critical system, for example a vote-tallying protocol or a financial settlement mechanism, with formal verification of at least one property, a hazard analysis, a budget plan, and a red-team report that the team responds to.

ENTREGABLES / DELIVERABLES

- Hazard analysis
- Verification plan
- Test matrix
- Safety margins
- Monitoring design
- Red team report
- Budget and resource plan
- Revised design responding to red team findings
- Team retrospective

EVALUACIÓN / EVALUATION

PD, Public defense with red team findings as adversarial input

CRITERIOS / CRITERIA

- Rigor of verification and test evidence
- Explicit treatment of failure consequences and affected users
- Technical response to adversarial findings and safety limits

PROGRAMAS DE REFERENCIA / REFERENCE PROGRAMS

MIT capstone engineering projects · Caltech senior thesis projects · UCL IEP final projects

Venture from Zero / Venture from Zero

OBJETIVO / OBJECTIVE

Create a real company over 16 weeks, under genuine legal, financial, organizational, and regulatory constraints. This is not a simulation. Teams legally incorporate a company, build an MVP, find real customers, sell or demonstrate verifiable demand, get paid where appropriate, measure, iterate, and account for the decisions they made. The pedagogical justification is not startup culture but the integrated test of operating under real constraints.

Concrete example. A legally incorporated company with a functional MVP, interviews and evidence of real demand, first customers or pilots, financial model, initial cap table, written decision log, and a pitch defended before a panel including investors, entrepreneurs, and technical reviewers.

ESTRUCTURA / STRUCTURE

The first weeks cover the operational mechanics of company creation: incorporation, cap table, fundraising, first customers, basic accounting, compliance, taxes. These are not taught in the abstract: they are executed immediately as part of launching the venture. The Studio integrates everything learned in three years: technical skills, design, institutions, finance, communication, and leadership.

ENTREGABLES / DELIVERABLES

- Incorporated company (or in process)
- Functional MVP
- Evidence of real demand: paying customers, pilots, letters of intent, or documented usage
- Updated financial model, initial cap table, and documented decision log
- Operational checklist: accounting, compliance, taxes, support, and responsibilities

- Legal, financial, technical, and reputational risk register
- AI usage report
- Red team report by peers
- Complete postmortem
- Final public defense before technical, institutional, entrepreneurial, and investor reviewers

EVALUACIÓN / EVALUATION

PD, Final public defense before investors, entrepreneurs, and technical or institutional reviewers

CRITERIOS / CRITERIA

- Verifiable evidence of demand and learning
- Legal, financial, operational, and AI-use accountability
- Legitimacy of the venture before customers, affected users, and external reviewers

PROGRAMAS DE REFERENCIA / REFERENCE PROGRAMS

Y Combinator Startup School · Stanford CS 183 Startup · MIT delta v accelerator · Entrepreneur First

Admissions

11.1 Who This Program Is For

This degree is not for everyone. It is for a specific kind of person: someone who finds traditional disciplinary boundaries frustrating rather than comforting. Someone who has already built something, a piece of software, a community, a business, an argument, and wants to understand more deeply why it worked or failed. Someone who reads across fields without being asked to. Someone who is honest about what they don't know and precise about what they do. Someone who wants to operate in the world, not just describe it.

The cohort is deliberately small, 30 to 40 students per entering class. This is not a scaling constraint. It is a pedagogical requirement. Formation happens in relationship with other people. The culture that develops when a small group of people share three and a half years of sustained intellectual pressure cannot be replicated at scale.

11.2 The Three-Stage Selection Process

Stage 1, Written Application

Two essays. The first describes a system the applicant has observed, built, or broken, anything from a neighborhood to a piece of software to a family business. The evaluators are looking for the ability to see structure, identify failure modes, and reason about causality. The second addresses what they believe AI changes and what it does not. This filters for people who have actually thought about the founding premise of the program rather than just wanting a prestigious degree. Academic records are reviewed as context, not as the primary filter.

Stage 2, Quantitative and Logical Examination

A purpose-built three-hour exam. Four sections:

- **Mathematical reasoning and estimation**, order-of-magnitude problems, logical deduction, combinatorial thinking, basic probability. Tests how someone thinks, not what they have memorized.
- **Formal reasoning**, logical puzzles, argument analysis, identifying invalid inferences.
- **Computational thinking**, no syntax required. Pseudocode, algorithmic thinking, debugging a described process.
- **Analytical reading**, a short dense text, scientific or philosophical, followed by questions requiring extraction of the argument, identification of assumptions, and evaluation of evidence.

The exam does not assume that all applicants have had reliable high-school calculus. It does require fluency with elementary algebra, functions, symbolic manipulation, estimation, proportional reasoning, graph reading, and mathematical reasoning. The program owns calculus from the first block; admissions verifies that the student can enter that pace without turning the first year into remediation.

Stage 3, Oral Interview

Thirty minutes. Two interviewers minimum.

- **Ten minutes defending the written essay**, interviewers push back, propose counterexamples, ask what was missed. Tests intellectual honesty and comfort with not knowing.
- **Ten minutes on a live problem**, a scenario with incomplete information, reasoning out loud. Could be a Fermi estimation, a system design question, an ethical dilemma. The answer is not the point. The thinking is.
- **Ten minutes of genuine conversation**, what they are reading, what problem obsesses them, what they have built.

11.3 What the Process Selects For

The process is designed to find people who reason precisely under uncertainty, are honest about what they don't know, think about systems naturally, and have some evidence of making or building things. Secondary school grades are one signal among many. Some of the best applicants will have uneven academic records because they were spending their time building things. The process is designed to find them.

11.4 No CBC

Students admitted to this program do not pass through the CBC. The program's selective admissions process, with its three stages and its explicit requirements for mathematical reasoning and computational thinking, serves the same function as a foundational year, ensuring that admitted students arrive with the prerequisites for the program's demanding first year. Calculus is not delegated to secondary school or to the CBC: it is taught inside the degree from Year 1, Block A, but on a base of reasoning and algebra that admissions must verify.

Graduate Profile

12.1 What Graduates Can Do

A graduate of this program can:

- Identify a real problem and frame it with precision
- Build a functional solution with code, mathematics, and institutional design
- Build, deploy, and evaluate ML and AI systems
- Design products that people actually use
- Read a balance sheet, model a business, and create a company
- Design incentive structures and institutional architectures
- Reason about physical and biological constraints on systems
- Recognize complexity patterns across domains
- Question the epistemic foundations of their own models
- Negotiate, persuade, and build coalitions without formal authority
- Lead teams and make unpopular decisions
- Write with clarity and speak with conviction in Spanish and English
- Use AI to move fast without losing judgment
- Navigate regulatory and political environments
- Reason about ethical obligations and assign responsibility
- Launch things, measure results, and correct course
- Take responsibility when it matters

12.2 Where Graduates Go

The program does not produce a single career type; it produces a formation that finds expression in several ways. The five pathways below are not curricular tracks, all students take the same program, but the occupational profiles that the formation naturally prepares for. Most graduates will move between two or more over the course of their careers.

One of the most common and deliberately prepared pathways is continuing into graduate study abroad. The 3.5-year structure is designed to free graduates into top international masters or doctoral programs while still in their early twenties, where they can deepen at research level in the discipline they choose. That deepening is not a side outcome of the program's design, it is the trajectory the design anticipates.

1. Founders and operators of technical companies

Creating companies that combine technical depth with organizational discipline. Studio VI, Venture from Zero, is the direct preparation: legal incorporation, MVP, first customers, fundraising. The strategy, product, organizational behavior, negotiation, finance, product design, and data systems courses provide the operational base; Algorithms, ML, Deep Learning, Systems Programming, and Distributed Systems provide the technical depth for building what the company requires. Typical destinations: startup founders, early technical hires at early-stage companies, operators scaling products.

2. Technical and product leadership in AI-heavy organizations

Roles where deciding what to build and how to build it is as critical as the implementation itself. The Deep Learning and Foundation Models, Epistemology and Limits of AI, Decision Theory, Product Design, and Data Systems courses provide the technical and conceptual vocabulary to lead AI teams without losing judgment about what those systems can and cannot do. Typical destinations: technical leads at software companies, product managers on AI platforms, ML team leads, architects of critical systems.

3. Institutional design and public-sector modernization

Working in or with the state to design mechanisms, regulations, and services that work. Studio IV, Diseño de Sistemas Institucionales, is the direct preparation. The Political Economy and Regulation, Institutional Design Laboratory, Game Theory and Mechanism Design, Culture and Legitimacy, and History of Institutions courses provide the analytical frame; the technical courses provide the capacity to reason about the systems the state has to regulate or build. Typical destinations: state digital modernization teams, regulatory agencies with a technological dimension, international organizations, technically-grounded policy consultancies, foundations that design mechanisms.

4. Financial, infrastructure, and risk systems

Building and operating systems where failure costs are high and fat-tailed distributions matter. The Stochastic Processes and Risk, Finance Money and Banking, Control Observability and Measurement, Data Systems, Security and Cryptography, Formal Methods and Verification, and Networks and Distributed Systems courses form the base. Studio III, Reliability Under Pressure, trains operation under real production conditions; Studio V, Critical System Under Constraints, trains design under regulatory or safety constraints. Typical destinations: financial infrastructure engineering, quantitative analysts with fat-tail discipline, reliability and critical-infrastructure teams, risk management at central banks or regulators, fintechs with an infrastructure component.

5. Research and graduate study

Continuing into graduate programs or entering research groups in complex systems, machine learning, decision science, computational political economy, technically-grounded public policy, or philosophy of AI. The Year 1 mathematical and scientific depth, the Studios as applied research practice, and the program's interdisciplinary orientation prepare graduates particularly well for master's and doctoral programs at the intersection of several disciplines, the places where

narrow disciplinary training hits its limits. Typical destinations: master's and doctoral programs in CS, ML, economics, public policy, complex systems, decision science, and philosophy of science, in Argentina and abroad.

What these pathways share is not a sector or a title. It is the capacity to operate at the level of complexity that the post-AI world demands and the willingness to take responsibility for the outcomes.

12.3 The Long-Term Bet

This program is a bet on a specific future: that the most important work of the next fifty years will be done at the intersection of technology, institutions, and power, and that the people capable of doing that work need a formation that does not currently exist.

Latin America is systematically underrepresented in the global conversation about artificial intelligence, institutional design, and the governance of technological change. The frameworks being developed to govern AI are being designed primarily elsewhere. This program is, among other things, a contribution to changing that, not by producing people who analyze the problem, but by producing people who build the solutions.

This is not a degree for people who want to study systems. It is a degree for people who want to build them.

References by Course and Institution

Course	Reference Institution	Reference Course / Program
Álgebra Lineal I y II	MIT	18.06, 18.065 Linear Algebra
Álgebra Lineal I y II	Stanford	EE364A Convex Optimization
Álgebra Lineal I y II	Imperial College London	Mathematics for Machine Learning (Coursera)
Razonamiento Formal	MIT	6.1200 Mathematics for Computer Science
Razonamiento Formal	Stanford	Phil 151 Metalogic, CS 103
Cálculo y Análisis	MIT	18.01 + 18.02 Calculus, 18.100A Real Analysis
Cálculo y Análisis	Caltech	Ma 1abc Mathematics
Cálculo y Análisis	École Polytechnique	Tronc commun (modules de mathématiques)
Cálculo y Análisis	Cambridge	Mathematical Tripos, Part IA Analysis
Historia de las Instituciones, la Tecnología y el Poder	MIT	17.S914 Political Development; 15.137
Historia de las Instituciones, la Tecnología y el Poder	Oxford	PPE Politics; History of Science and Technology
Historia de las Instituciones, la Tecnología y el Poder	Harvard	Kennedy School Governance and Institutions
Historia de las Instituciones, la Tecnología y el Poder	Stanford	HSST Program
Microeconomía	MIT	14.01 Principles of Microeconomics

Microeconomía	Yale	ECON 159 Game Theory
Fundamentos de Programación y Datos	MIT	6.100A, Missing Semester
Fundamentos de Programación y Datos	Stanford	CS 106B
Probabilidad y Estadística	MIT	6.3700, 18.650
Probabilidad y Estadística	Stanford	CS 109
Probabilidad y Estadística	Caltech	Ma 3
Programación de Sistemas	MIT	6.1800, 6.1810
Programación de Sistemas	Stanford	CS 110
Física I/II	MIT	8.01, 8.02, 8.04
Física I/II	Caltech	Ph 1a, Ph 1b
Física I/II	École Polytechnique	Tronc commun Physique
Machine Learning	MIT	6.3900
Machine Learning	Stanford	CS 229
Machine Learning	Caltech	CMS 155
Algoritmos y Optimización	MIT	6.1210, 6.C571
Algoritmos y Optimización	Stanford	CS 161, MS&E 211
Econometría	MIT	14.32
Econometría	Stanford	MS&E 226
Econometría	Caltech	IDS 126
Deep Learning	Stanford	CS 224N, CS 236
Deep Learning	MIT	6.S898
Deep Learning	UC Berkeley	CS 285
Teoría de Juegos	Stanford	MS&E 232
Teoría de Juegos	MIT	14.126

Teoría de Juegos	Yale	ECON 159
Teoría de Juegos	Harvard	CS 1360
Procesos Estocásticos	Stanford	MS&E 221
Procesos Estocásticos	MIT	6.7700
Procesos Estocásticos	London Mathematical Laboratory	Ergodicity Economics
Métodos Formales	Carnegie Mellon	15-815
Métodos Formales	MIT	6.820
Redes y Distribuidos	MIT	6.5820, 6.5840
Redes y Distribuidos	Stanford	CS 144, CS 244B
Redes y Distribuidos	Cambridge	Distributed Systems
Finanzas, Dinero y Banca	MIT	15.401
Finanzas, Dinero y Banca	Yale	ECON 252
Finanzas, Dinero y Banca	Stanford	MS&E 246
Seguridad y Criptografía	Stanford	CS 255
Seguridad y Criptografía	MIT	6.5620
Economía Política	MIT	14.42
Economía Política	Stanford	Law and Technology
Economía Política	Oxford	PPE Politics and Philosophy
Cultura y Legitimidad	Stanford	Anthropology
Cultura y Legitimidad	MIT	Comparative Media Studies
Cultura y Legitimidad	École Polytechnique	HSS
Epistemología	Stanford	Symbolic Systems, Phil 80
Epistemología	MIT	24.09 Minds and Machines
Epistemología	Oxford	Philosophy of Mind
Epistemología	Cambridge	Philosophy of Science

Diseño y Estrategia de Producto	Stanford	CS 183, d.school Design Thinking
Diseño y Estrategia de Producto	MIT	15.835 Entrepreneurial Marketing, Integrated Design and Management
Diseño y Estrategia de Producto	Carnegie Mellon	Human-Computer Interaction
Control y Resiliencia	Stanford	EE 263
Control y Resiliencia	MIT	6.3100
Control y Resiliencia	Caltech	CDS 110
Ética y Responsabilidad	Oxford	Practical Ethics, PPE Philosophy
Ética y Responsabilidad	Harvard Kennedy School	Ethics and Public Policy
Lab. Diseño Institucional	Indiana University	Ostrom Workshop
Lab. Diseño Institucional	Santa Fe Institute	Complexity and Governance
Lab. Diseño Institucional	Oxford	Blavatnik School of Government
Liderazgo, Negociación y Coordinación	Harvard	Program on Negotiation
Liderazgo, Negociación y Coordinación	MIT Sloan	15.665 Power and Negotiation
Liderazgo, Negociación y Coordinación	Stanford	MS&E 280 Organizational Behavior
Liderazgo, Negociación y Coordinación	INSEAD	Organisational Behaviour
Sistemas Complejos	Santa Fe Institute	Complexity Explorer
Sistemas Complejos	MIT	18.354
Sistemas Complejos	Oxford	Complex Systems
Teoría de la Decisión	Stanford	MS&E 250
Teoría de la Decisión	Harvard Kennedy School	Decision Science
Métodos de Investigación, Medición y Diseño Experimental	Stanford	GSB Methodology of Empirical Research

Métodos de Investigación, Medición y Diseño Experimental	MIT	14.310x Data Analysis for Social Scientists
Métodos de Investigación, Medición y Diseño Experimental	Berkeley	D-Lab Computational Social Science
Métodos de Investigación, Medición y Diseño Experimental	Center for Open Science	Open Science training
Sistemas de Datos	CMU	15-445/645 Database Systems
Sistemas de Datos	MIT	6.830 Database Systems
Sistemas de Datos	Stanford	CS 145 Data Management and Data Systems
Sistemas de Datos	Berkeley	CS 186 Introduction to Database Systems
Optimización Aplicada y Métodos Numéricos	MIT	15.053 Optimization Methods in Management Science
Optimización Aplicada y Métodos Numéricos	Stanford	EE364B Convex Optimization II
Optimización Aplicada y Métodos Numéricos	Carnegie Mellon	10-725 Convex Optimization
Optimización Aplicada y Métodos Numéricos	Caltech	ACM 113 Mathematical Optimization
Biología de Sistemas, Evolución y Ecología	MIT	7.013; 7.91J
Biología de Sistemas, Evolución y Ecología	Caltech	Bi 1
Biología de Sistemas, Evolución y Ecología	École Polytechnique	Tronc commun Biologie

Evaluation Framework by Course

Course	Primary Evaluation	Secondary Evaluation
Álgebra Lineal y Computación I	WE, Written Exam	TP, Weekly problem sets
Álgebra Lineal y Computación II	WE, Written Exam	TP, Computational project
Razonamiento Formal y Demostración	WE, Written proof exam	OD, Oral defense of proof portfolio
Cálculo y Análisis	WE, Written exam with unseen problems and short proofs	TP, Weekly problem sets and proof exercises
Microeconomía y Lógica de Decisión	WE, Written Exam	OD, Defense of market analysis project
Fundamentos de Programación y Datos	LC, Live coding session	TP, Weekly coding assignments
Probabilidad y Estadística	WE, Written Exam	TP, Weekly problem sets
Programación de Sistemas	LC, Live systems analysis	TP, Debugging portfolio and labs
Física I: Mecánica, Energía y Termodinámica	WE, Written Exam with unseen problems	TP, Problem sets and estimation exercises
Física II: Electromagnetismo, Señales y Computación Física	WE, Written Exam	TP, Problem sets and system analysis
Machine Learning e Ingeniería de Modelos	LC, Live ML pipeline session	TP, End-to-end ML project
Biología de Sistemas, Evolución y Ecología	OD, Defense of analysis of a biological system	TP, Quantitative analysis with estimates and models
Algoritmos, Complejidad y Optimización	WE, Written Exam with algorithm design problems	TP, Algorithm portfolio with formal proofs

Econometría e Inferencia Causal	OD, Defense of empirical study	TP, Empirical study with real data
Deep Learning y Modelos Fundacionales	LC, Live architecture session	OD + TP, Defense of replication project
Teoría de Juegos y Diseño de Mecanismos	WE, Written Exam	OD, Defense of mechanism design project
Procesos Estocásticos y Riesgo	WE, Written Exam	OD, Defense of risk analysis
Métodos Formales y Verificación	TP, Lean 4 proof portfolio	OD, Defense of formal verification project
Redes y Sistemas Distribuidos	LC, Live architecture and failure analysis	TP, Architecture document and protocol spec
Finanzas, Dinero y Banca	OD, Defense of financial model and memo	TP, Financial model built with AI
Seguridad, Criptografía y Confianza	LC, Live threat modeling session	TP, Threat model and security architecture
Economía Política, Derecho y Regulación	OD, Adversarial defense of policy analysis	TP, Policy analysis with stakeholder map
Historia de las Instituciones, la Tecnología y el Poder	OD, Adversarial defense of comparative essay	TP, Essay with revision history
Epistemología, Ciencia y los Límites de la IA	OD, Defense of philosophical analysis of real AI system	TP, Written analysis
Diseño y Estrategia de Producto	PD, Public defense of strategy and product with usability evidence	TP, Strategic and design portfolio built with AI
Control, Observabilidad y Medición	WE, Written Exam	OD, Defense of system stability and measurement critique
Ética y Responsabilidad bajo Incertidumbre	OD, Defense of ethical analysis of a hard case	TP, Written case analysis
Laboratorio de Diseño Institucional	PD, Public defense before external panel with practitioners	TP, Institutional design document
Liderazgo, Negociación y Coordinación	LC, Live negotiation exercise with real stakes	OD, Defense of organizational diagnosis

Sistemas Embebidos e Interacción Física	PD, Public defense of a working embedded system with documented design decisions	TP, Functional prototype with hardware, firmware, and measurement plan
Sistemas Complejos	OD, Defense of agent-based model	TP, Agent-based model of real complex system
Teoría de la Decisión	OD, Defense of decision analysis of multi-criteria problem	TP, Decision analysis built with AI
Escritura Avanzada en Español	WP, Written portfolio with full revision history	
Métodos de Investigación, Medición y Diseño Experimental	OD, Defense of a research design proposal with measurement plan and pre-analysis	TP, Replication audit of a published study
Sistemas de Datos	LC, Live data-system design and debugging session	TP, Schema, migration, and reliability project
Optimización Aplicada y Métodos Numéricos	WE, Written exam with modeling problems and numerical analysis	LC, Live design and debugging session on a real optimization problem
Seminario Fundacional: IA, Juicio y Responsabilidad	OD, Oral defense of position paper	TP, Position paper on what AI changes
Cultura, Legitimidad y Normas	OD, Defense of critical essay	TP, Essay with full revision history
Electiva Guiada	Defined by supervising faculty	
Seminario de Egreso	OD, Final oral reflection (ungraded)	
Studio I, Escritura y Oratoria	WP, Written portfolio with revision history	Recorded presentations with self-evaluation
Studio II, Construir y Romper	PD, Public defense of functional system	
Studio III, Fiabilidad Bajo Presión	PD, Defense simulating post-incident review	
Studio IV, Diseño de Sistemas Institucionales	PD, Public defense before practitioners	

Studio V, Sistema Crítico bajo
Restricciones

PD, Public defense with red team
findings

Studio VI, Venture from Zero

PD, Final public defense before
investors

Full Bibliography

Abelson, Harold, and Gerald Jay Sussman. *Structure and Interpretation of Computer Programs*. 2nd ed. MIT Press, 1996.

↗

Acemoglu, Daron, and James A. Robinson. *The Narrow Corridor: States, Societies, and the Fate of Liberty*. Penguin Press, 2019.

↗

Acemoglu, Daron, and James A. Robinson. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown, 2012.

↗

Aleksandrov, A. D., A. N. Kolmogorov, and M. A. Lavrent'ev. *Mathematics: Its Content, Methods and Meaning*. Dover, 1999.

↗

Alberts, Bruce, et al. *Essential Cell Biology*. 5th ed. W. W. Norton, 2019.

↗

Alden, Lyn. *Broken Money: Why Our Financial System Is Failing Us and How We Can Make It Better*. Timestamp Press, 2023.

↗

Ammous, Saifedean. *The Bitcoin Standard: The Decentralized Alternative to Central Banking*. Wiley, 2018.

↗

Anderson, Ross J. *Security Engineering: A Guide to Building Dependable Distributed Systems*. 3rd ed. Wiley, 2020.

↗

Artin, Michael. *Algebra*. 2nd ed. Pearson, 2010.

↗

Angrist, Joshua D., and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, 2009.

↗

Arpaci-Dusseau, Remzi H., and Andrea C. Arpaci-Dusseau. *Operating Systems: Three Easy Pieces*. Arpaci-Dusseau Books, 2018. ostep.org.

↗

Åström, Karl J., and Richard M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, 2008.

↗

Avigad, Jeremy. *Mathematics in Lean*. Lean Prover Community, 2023. leanprover-community.github.io.

↗

Axler, Sheldon. *Linear Algebra Done Right*. 4th ed. Springer, 2024.

↗

Barabási, Albert-László. *Network Science*. Cambridge University Press, 2016.

↗

Bishop, Christopher M. *Pattern Recognition and Machine Learning*. Springer, 2006. ↗

Blitzstein, Joseph K., and Jessica Hwang. *Introduction to Probability*. 2nd ed. CRC Press, 2019. probabilitybook.net. ↗

Boneh, Dan, and Victor Shoup. *A Graduate Course in Applied Cryptography*. cryptobook.us, 2023. ↗

Banzi, Massimo, and Michael Shiloh. *Getting Started with Arduino*. 3rd ed. Maker Media, 2014. ↗

Bergstrom, Carl T., and Jevin D. West. *Calling Bullshit: The Art of Skepticism in a Data-Driven World*. Random House, 2020. ↗

Bertsimas, Dimitris, and John N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, 1997. ↗

Booth, Wayne C., Gregory G. Colomb, Joseph M. Williams, Joseph Bizup, and William T. FitzGerald. *The Craft of Research*. 4th ed. University of Chicago Press, 2016. ↗

Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, 2014. ↗

Bourdieu, Pierre. *The Logic of Practice*. Stanford University Press, 1990. ↗

Boyd, Stephen, and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. stanford.edu/~boyd/cvxbook. ↗

Brealey, Richard A., Stewart C. Myers, and Franklin Allen. *Principles of Corporate Finance*. 14th ed. McGraw-Hill, 2022. ↗

Bryant, Randal E., and David R. O'Hallaron. *Computer Systems: A Programmer's Perspective*. 3rd ed. Pearson, 2015. ↗

Bryar, Colin, and Bill Carr. *Working Backwards: Insights, Stories, and Secrets from Inside Amazon*. St. Martin's Press, 2021. ↗

Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age*. W. W. Norton, 2014. ↗

Carroll, Sean B. *The Serengeti Rules*. Princeton University Press, 2016. ↗

Cartwright, Nancy. *How the Laws of Physics Lie*. Oxford University Press, 1983. ↗

Chandler, Alfred D. *The Visible Hand: The Managerial Revolution in American Business*. Harvard University Press, 1977. ↗

Chlipala, Adam. *Certified Programming with Dependent Types*. MIT Press, 2013. adam.chlipala.net/cpdt. ↗

Christensen, Clayton M. *The Innovator's Dilemma*. Harvard Business Review Press, 1997. ↗

Christian, Brian. *The Alignment Problem: Machine Learning and Human Values*. W. W. Norton, 2020. ↗

Cialdini, Robert B. *Influence: The Psychology of Persuasion*. Harper Business, 2006. ↗

Cooper, Alan, Robert Reimann, and David Cronin. *About Face: The Essentials of Interaction Design*. 4th ed. Wiley, 2014. ↗

Cormen, Thomas H., Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. 4th ed. MIT Press, 2022. ↗

Cover, Thomas M., and Joy A. Thomas. *Elements of Information Theory*. 2nd ed. Wiley, 2006. ↗

Cunningham, Scott. *Causal Inference: The Mixtape*. Yale University Press, 2021. mixtape.scunning.com. ↗

Dalio, Ray. *Principles for Dealing with the Changing World Order: Why Nations Succeed and Fail*. Simon & Schuster, 2021. ↗

Dawkins, Richard. *The Selfish Gene*. 40th Anniversary ed. Oxford University Press, 2016. ↗

Debord, Guy. *The Society of the Spectacle*. Zone Books, 1994. ↗

Dennett, Daniel C. *Consciousness Explained*. Little, Brown, 1991. ↗

Douglas, Mary. *How Institutions Think*. Syracuse University Press, 1986. ↗

Dreyfus, Hubert L. *What Computers Can't Do: A Critique of Artificial Reason*. Harper & Row, 1972. ↗

Enderton, Herbert B. *A Mathematical Introduction to Logic*. 2nd ed. Academic Press, 2001. ↗

Feynman, Richard P., Robert B. Leighton, and Matthew Sands. *The Feynman Lectures on Physics*. Vols. I–II. Caltech, 2013. feynmanlectures.caltech.edu. ↗

Feyerabend, Paul. *Against Method*. 4th ed. Verso, 2010. ↗

Fisher, Mark. *Ghosts of My Life: Writings on Depression, Hauntology and Lost Futures*. Zero Books, 2014. ↗

Fisher, Roger, William Ury, and Bruce Patton. *Getting to Yes: Negotiating Agreement Without Giving In*. Penguin Books, 2011. ↗

Géron, Aurélien. *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*. 3rd ed. O'Reilly, 2022. ↗

Gigerenzer, Gerd. *Gut Feelings: The Intelligence of the Unconscious*. Viking, 2007. ↗

Girard, René. *Violence and the Sacred*. Johns Hopkins University Press, 1977. ↗

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. deeplearningbook.org. ↗

Graham, Ronald L., Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. 2nd ed. Addison-Wesley, 1994. ↗

Griffiths, David J. *Introduction to Electrodynamics*. 4th ed. Cambridge University Press, 2017. ↗

Griffiths, David J. *Introduction to Quantum Mechanics*. 3rd ed. Cambridge University Press, 2018. ↗

Grove, Andrew S. *High Output Management*. Random House, 1983. ↗

Halliday, David, Robert Resnick, and Jearl Walker. *Fundamentals of Physics*. 11th ed. Wiley, 2018. ↗

Han, Byung-Chul. *The Burnout Society*. Stanford University Press, 2015. ↗

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. 2nd ed. Springer, 2009. hastie.su.domains/ElemStatLearn. ↗

Henrich, Joseph. *The WEIRDest People in the World*. Farrar, Straus and Giroux, 2020. ↗

Hillier, Frederick S., and Gerald J. Lieberman. *Introduction to Operations Research*. 11th ed. McGraw-Hill, 2021. ↗

Hirschman, Albert O. *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Harvard University Press, 1970. ↗

Hobart, Byrne, and Tobias Huber. *Bubbles and the End of Stagnation*. Stripe Press, 2024. ↗

Hobsbawm, Eric. *The Age of Revolution: 1789–1848*. Vintage, 1962. *The Age of Capital: 1848–1875*. Vintage, 1975. *The Age of Empire: 1875–1914*. Vintage, 1987. ↗

Horowitz, Paul, and Winfield Hill. *The Art of Electronics*. 3rd ed. Cambridge University Press, 2015. ↗

Hubbard, John H., and Barbara Burke Hubbard. *Vector Calculus, Linear Algebra, and Differential Forms: A Unified Approach*. 5th ed. Matrix Editions, 2015. ↗

Horowitz, Ben. *The Hard Thing About Hard Things: Building a Business When There Are No Easy Answers*. Harper Business, 2014. ↗

Imbens, Guido W., and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge University Press, 2015. ↗

Jaynes, Edwin T. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003. ↗

Jonas, Hans. *The Imperative of Responsibility: In Search of an Ethics for the Technological Age*. University of Chicago Press, 1984. ↗

Jurafsky, Daniel, and James H. Martin. *Speech and Language Processing*. 3rd ed. draft. web.stanford.edu/~jurafsky/slp3. ↗

Kahneman, Daniel. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011. ↗

Kaplan, Jared, et al. "Scaling Laws for Neural Language Models." arXiv:2001.08361, 2020. ↗

Katz, Jonathan, and Yehuda Lindell. *Introduction to Modern Cryptography*. 3rd ed. CRC Press, 2020. ↗

Kerrisk, Michael. *The Linux Programming Interface: A Linux and UNIX System Programming Handbook*. No Starch Press, 2010. ↗

Kleppmann, Martin. *Designing Data-Intensive Applications*. O'Reilly, 2017. ↗

Hellerstein, Joseph M., and Michael Stonebraker, eds. *Readings in Database Systems*. 5th ed., 2015. ↗

Huyen, Chip. *Designing Machine Learning Systems*. O'Reilly, 2022. ↗

Krug, Steve. *Don't Make Me Think, Revisited*. 3rd ed. New Riders, 2014. ↗

Kuhn, Thomas S. *The Structure of Scientific Revolutions*. 4th ed. University of Chicago Press, 2012. ↗

Kurose, James F., and Keith W. Ross. *Computer Networking: A Top-Down Approach*. 8th ed. Pearson, 2021. ↗

Lencioni, Patrick. *The Five Dysfunctions of a Team*. Jossey-Bass, 2002. ↗

Lessig, Lawrence. *Code: And Other Laws of Cyberspace*. Basic Books, 1999. ↗

Leveson, Nancy G. *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press, 2011. ↗

MacIntyre, Alasdair. *After Virtue: A Study in Moral Theory*. 3rd ed. University of Notre Dame Press, 2007. ↗

MacKay, David J. C. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003. inference.org.uk. ↗

Mahajan, Sanjoy. *The Art of Insight in Science and Engineering: Mastering Complexity*. MIT Press, 2014. ↗

March, James G. *A Primer on Decision Making: How Decisions Happen*. Free Press, 1994. ↗

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995. ↗

Maynard Smith, John. *Evolution and the Theory of Games*. Cambridge University Press, 1982. ↗

Nowak, Martin A. *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap Press, 2006. ↗

Levin, Simon A. *Fragile Dominion: Complexity and the Commons*. Perseus Books, 1999. ↗

McConnell, Steve. *Code Complete: A Practical Handbook of Software Construction*. 2nd ed. Microsoft Press, 2004. ↗

Meadows, Donella H. *Thinking in Systems: A Primer*. Chelsea Green, 2008. ↗

Mehrling, Perry. *The New Lombard Street: How the Fed Became the Dealer of Last Resort*. Princeton University Press, 2011. ↗

Milgrom, Paul, and John Roberts. *Economics, Organization and Management*. Prentice Hall, 1992. ↗

Miller, John H., and Scott E. Page. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press, 2007. ↗

Minsky, Hyman P. *Stabilizing an Unstable Economy*. McGraw-Hill, 2008. ↗

Mitchell, Melanie. *Complexity: A Guided Tour*. Oxford University Press, 2009. ↗

Mokyr, Joel. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press, 1990. ↗

Montgomery, Douglas C. *Design and Analysis of Experiments*. 10th ed. Wiley, 2019. ↗

Muller, Jerry Z. *The Tyranny of Metrics*. Princeton University Press, 2018. ↗

Munkres, James R. *Topology*. 2nd ed. Pearson, 2000. ↗

Ramakrishnan, Raghu, and Johannes Gehrke. *Database Management Systems*. 3rd ed. McGraw-Hill, 2003. ↗

Nisan, Noam, et al. *Algorithmic Game Theory*. Cambridge University Press, 2007. ↗

Nocedal, Jorge, and Stephen J. Wright. *Numerical Optimization*. 2nd ed. Springer, 2006. ↗

Norman, Donald A. *The Design of Everyday Things*. Revised and Expanded ed. Basic Books, 2013. ↗

Norris, J. R. *Markov Chains*. Cambridge University Press, 1997. ↗

North, Douglass C. *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, 1990. ↗

Olson, Mancur. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press, 1965. ↗

Ostrom, Elinor. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990. ↗

Ouyang, Long, et al. "Training Language Models to Follow Instructions with Human Feedback." NeurIPS, 2022. ↗

Parfit, Derek. *Reasons and Persons*. Oxford University Press, 1984. ↗

Pearl, Judea, and Dana Mackenzie. *The Book of Why: The New Science of Cause and Effect*. Basic Books, 2018. ↗

Perez, Carlota. *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*. Edward Elgar, 2002. ↗

Peters, Ole. "The Ergodicity Problem in Economics." *Nature Physics* 15 (2019): 1216–1221. ↗

Pierce, Benjamin C., et al. *Software Foundations*. Electronic textbook, 2023. softwarefoundations.cis.upenn.edu. ↗

Polanyi, Karl. *The Great Transformation: The Political and Economic Origins of Our Time*. Beacon Press, 1944. ↗

Polanyi, Michael. *Personal Knowledge: Towards a Post-Critical Philosophy*. University of Chicago Press, 1958. ↗

Porter, Theodore M. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton University Press, 1995. ↗

Prince, Simon J.D. *Understanding Deep Learning*. MIT Press, 2023. udlbook.github.io/udlbook. ↗

Raiffa, Howard. *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*. Addison-Wesley, 1968. ↗

Ries, Eric. *The Lean Startup*. Crown Business, 2011. ↗

Ritchie, Stuart. *Science Fictions: How Fraud, Bias, Negligence, and Hype Undermine the Search for Truth*. Metropolitan Books, 2020. ↗

Ross, Sheldon M. *Introduction to Probability Models*. 12th ed. Academic Press, 2019. ↗

Roth, Alvin E. *Who Gets What, and Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt, 2015. ↗

Rumelt, Richard P. *Good Strategy Bad Strategy: The Difference and Why It Matters*. Crown Business, 2011. ↗

Russell, Stuart. *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking, 2019. ↗

Schelling, Thomas C. *The Strategy of Conflict*. Harvard University Press, 1960. ↗

Shadish, William R., Thomas D. Cook, and Donald T. Campbell. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Houghton Mifflin, 2002. ↗

Scott, James C. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press, 1998. ↗

Sedgewick, Robert, and Kevin Wayne. *Algorithms*. 4th ed. Addison-Wesley, 2011. ↗

Simon, Herbert A. *Models of Bounded Rationality*. MIT Press, 1982. ↗

Simon, Herbert A. *The Sciences of the Artificial*. 3rd ed. MIT Press, 1996. ↗

Silberschatz, Abraham, Henry F. Korth, and S. Sudarshan. *Database System Concepts*. 7th ed. McGraw-Hill, 2019. ↗

Sipser, Michael. *Introduction to the Theory of Computation*. 3rd ed. Cengage, 2012. ↗

Spivak, Michael. *Calculus*. 4th ed. Publish or Perish, 2008. ↗

Strang, Gilbert. *Introduction to Linear Algebra*. 6th ed. Wellesley-Cambridge Press, 2023. ↗

Strang, Gilbert. *Linear Algebra and Learning from Data*. Wellesley-Cambridge Press, 2019. ↗

Strogatz, Steven H. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. 2nd ed. CRC Press, 2014. ↗

- Taleb, Nassim Nicholas. *Statistical Consequences of Fat Tails*. STEM Academic Press, 2020. ↗
- Tao, Terence. *Analysis I and II*. Hindustan Book Agency, 2016. ↗
- Tetlock, Philip E., and Dan Gardner. *Superforecasting: The Art and Science of Prediction*. Crown, 2015. ↗
- Tirole, Jean. *Economics for the Common Good*. Princeton University Press, 2017. ↗
- Trefethen, Lloyd N., and David Bau III. *Numerical Linear Algebra*. SIAM, 1997. ↗
- Van Steen, Maarten, and Andrew S. Tanenbaum. *Distributed Systems*. 4th ed. distributed-systems.net, 2023. ↗
- Varian, Hal R. *Intermediate Microeconomics: A Modern Approach*. 9th ed. W. W. Norton, 2014. ↗
- Vaswani, Ashish, et al. "Attention Is All You Need." NeurIPS, 2017. ↗
- Velleman, Daniel J. *How to Prove It: A Structured Approach*. 3rd ed. Cambridge University Press, 2019. ↗
- Voss, Chris. *Never Split the Difference: Negotiating As If Your Life Depended On It*. Harper Business, 2016. ↗
- Wasserman, Larry. *All of Statistics: A Concise Course in Statistical Inference*. Springer, 2004. ↗
- West, Geoffrey. *Scale: The Universal Laws of Growth, Innovation, Sustainability, and the Pace of Life in Organisms, Cities, Economies, and Companies*. Penguin Press, 2017. ↗
- Williams, Bernard. *Ethics and the Limits of Philosophy*. Harvard University Press, 1985. ↗
- Williamson, Oliver E. *The Economic Institutions of Capitalism*. Free Press, 1985. ↗
- Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs, 2019. ↗
- Brown, Wendy. *Undoing the Demos: Neoliberalism's Stealth Revolution*. Zone Books, 2015. ↗
- Cirillo, Pasquale, and Nassim Nicholas Taleb. *The Logic of Risk Taking*. arXiv, 2015. ↗
- Davidson, James Dale, and William Rees-Mogg. *The Sovereign Individual: Mastering the Transition to the Information Age*. Touchstone, 1997. ↗
- Epictetus. *The Discourses and Enchiridion*. Loeb Classical Library, Harvard University Press. ↗
- Espeland, Wendy Nelson, and Mitchell L. Stevens. "Commensuration as a Social Process." *Annual Review of Sociology* 24 (1998): 313-343. ↗
- Goodhart, Charles. "Problems of Monetary Management: The U.K. Experience." *Papers in Monetary Economics*, Reserve Bank of Australia, 1975. ↗

Hacking, Ian. *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge University Press, 1983. ↗

Haraway, Donna. "A Cyborg Manifesto." *Socialist Review*, 1985. ↗

Hayek, Friedrich A. *Denationalisation of Money: The Argument Refined*. Institute of Economic Affairs, 1976. ↗

Hayek, Friedrich A. "The Use of Knowledge in Society." *American Economic Review* 35, no. 4 (1945): 519-530. ↗

Hui, Yuk. *The Question Concerning Technology in China: An Essay in Cosmotechnics*. Urbanomic, 2016. ↗

Illich, Ivan. *Tools for Conviviality*. Harper & Row, 1973. ↗

Lakatos, Imre. *Proofs and Refutations: The Logic of Mathematical Discovery*. Cambridge University Press, 1976. ↗

Land, Nick. *Fanged Noumena: Collected Writings 1987-2007*. Urbanomic, 2011. ↗

Latour, Bruno. *Science in Action: How to Follow Scientists and Engineers Through Society*. Harvard University Press, 1987. ↗

Lee, Edward A., and Sanjit A. Seshia. *Introduction to Embedded Systems: A Cyber-Physical Systems Approach*. 2nd ed. MIT Press, 2017. ↗

Martínez Estrada, Ezequiel. *Radiografía de la Pampa*. Babel, 1933. ↗

O'Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown, 2016. ↗

Ostrom, Vincent. *The Intellectual Crisis in American Public Administration*. University of Alabama Press, 1973. ↗

Savage, Leonard J. *The Foundations of Statistics*. Wiley, 1954. ↗

Siedentop, Larry. *Inventing the Individual: The Origins of Western Liberalism*. Belknap Press, 2014. ↗

Simondon, Gilbert. *On the Mode of Existence of Technical Objects*. Univocal, 2017. ↗

Srnicek, Nick, and Alex Williams. *Inventing the Future: Postcapitalism and a World Without Work*. Verso, 2015. ↗

Weber, Max. *The Protestant Ethic and the Spirit of Capitalism*. 1905. ↗

Weil, Simone. *La gravedad y la gracia / Gravity and Grace*. 1947. ↗

Warden, Pete, and Daniel Situnayake. *TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers*. O'Reilly, 2019. ↗

Warren, Henry S. *Hacker's Delight*. 2nd ed. Addison-Wesley, 2012. ↗

Intellectual Conversation, Reference Publications

The publications below represent the intellectual context within which the program situates itself, the conversation it is entering and, through its graduates, intending to shape. None is required reading in the curriculum; they serve as reference for anyone wishing to situate the program within its contemporary context.

Foreign Affairs

The flagship journal of serious geopolitical and international relations thinking. Unmatched in its access to practitioners, former heads of state, central bank governors, senior diplomats, writing alongside leading academics. Foreign Affairs sets the terms of the serious conversation about global order, institutional design, and the governance of technological change. Its relevance to this program is direct: the graduates of this program will operate in the world that Foreign Affairs describes and analyzes.

Palladium Magazine

An important contemporary publication on the formation of high-responsibility public cadres, state capacity, and institutional futures. Palladium's on-the-ground reporting on Singapore's civil service, the Gulf states' technocratic cadres, and China's cadre system constitutes a body of case studies in exactly the question this program is trying to answer: how do you deliberately form people capable of governing complex systems? Samo Burja's great founder theory, the argument that

functional institutions depend on a small number of people with genuine understanding of how they work, is one of the intellectual foundations of this program's approach to forming technical and institutional leaders.

Le Grand Continent

The most serious European publication on institutional futures, geopolitical analysis, and strategic thinking. Based in Paris and affiliated with the École Normale Supérieure, Le Grand Continent publishes in multiple languages with contributors drawn from the highest levels of European policy and academia. Its focus on European institutional design, strategic autonomy, and the governance of technological change provides a non-Anglophone perspective that is essential counterweight to the US-centric framing of most AI governance debates.

Noema Magazine

Published by the Berggruen Institute, Noema operates at the intersection of technology, philosophy, geopolitics, and global-scale questions. Its genuinely global editorial vision, drawing on Chinese scholars, European philosophers, Latin American thinkers, and Anglophone voices in roughly equal measure, makes it a reference aligned with the intellectual aspiration of this program. Noema asks what it means to govern well when the technological substrate of society is changing faster than the institutions designed to govern it.

American Affairs

The leading journal of post-neoliberal political economy. American Affairs provides the most rigorous intellectual framework for understanding the failures of the current economic order and the institutional requirements of what might replace it. Its focus on industrial policy, the political economy of technology, and the relationship between economic structure and political legitimacy is directly relevant to the institutional design and political economy courses in this program.

421 News

Born in Buenos Aires in 2024, 421 is a publication about culture, technology, and philosophy built around the conviction that digital autonomy, cognitive sovereignty, and independent thinking are the essential capacities of the post-AI individual. 421 is both an intellectual reference for this program and evidence that the conversation this program is entering is not exclusively foreign, it is being built here, in this city, by people who share the program's foundational commitments.
