

GenAI Bootcamp 2026–2027

Course Title: Becoming AI-First: Generative AI in Business Practice

Instructor: Prof. Nikos Mylonopoulos

Contact Details: nmylonopoulos@alba.acg.edu

OVERVIEW

Raising productivity is the objective that SEV has set for Greek enterprises, in the SEV and IOBE study of May 2026, “Productivity: National Objective, Collective Responsibility”, and it is the objective of this bootcamp. Generative AI is a general-purpose technology that can raise the productivity of individuals, teams, and entire processes. For the employees of SEV member companies, the gains come in two forms. Augmentation means that the same people produce more and better work in less time when they direct AI well. Redesign means rethinking processes, roles, and services with AI as the starting point, so the gains reach the whole organization, not just the individual. The risks are equally clear. A company that misunderstands what AI can and cannot do wastes effort or acts on wrong output. A company that adopts AI superficially, using it to do old things a little faster, sees no lasting productivity gain and fails to adapt on time.

The central idea is **AI+, not +AI**: not adding AI on top of existing ways of working, but systematically asking how AI can augment every task, decision and interaction, and learning its limits through disciplined practice.

The 2026–2027 cycle strengthens the bootcamp's practical, applied character. About three-quarters of the contact time is spent at the keyboard. Every exercise starts from a real business use case or a realistic case scenario, drawn from finance, sales and customer service, marketing and communications, human resources, operations and supply chain, legal and compliance, and data and reporting, and is then repeated by each participant. The final session is a demo fair in which participants, individually or in pairs, build and demonstrate a working solution to a problem from their own work.

The bootcamp is designed for corporate employees who are expected to lead or support the adoption of generative AI in their companies. It is a working seminar, not a lecture series. Every session pairs one theme with extended hands-on practice, and every participant leaves with prompts, assistants, and working methods.

CONTENT & FLOW

The five sessions form a sequence. The first four pair one theme with prompting practice and business use cases. In each of Sessions 1 to 4, a Discussion stream covers one theme in about 30 minutes, plus Q&A. During the Practice stream, participants engage with techniques, tools, and use cases on the keyboard. All sessions are delivered in Greek with materials in English.

The four themes of the Discussion stream are:

1. **The technology of AI, the properties of LLMs, and why it is different this time.** Large language models (LLMs) are the engine of generative AI: they predict the next word from everything they have read. This is why they can draft, summarize, analyze and code, and also why they sometimes state falsehoods with confidence.
2. **The economics of the AI revolution and how technology changes the competition and the economy.** Generative AI makes intelligence available at near-zero marginal cost, so that tasks which once needed skilled hours now cost minutes. Enterprise computing and the internet show that productivity gains arrive only after firms reorganize their work around the new technology.
3. **The transformation of organizations with an agentic operating model and the skills for the AI era.** An AI agent is software that plans multi-step tasks, uses tools such as code, other systems, and the web, and completes them with limited supervision. In an agentic operating model, agents run entire processes while people set goals, check results, and handle exceptions. The session examines how roles and skills change under this model and where its productivity gains come from: fewer hand-offs and faster cycles.
4. **Governance for agentic AI, risk and compliance.** When AI acts rather than advises, an error can reach a customer, a supplier or a regulator before anyone has read it. Governance means the rules that decide what an agent may do on its own, what a person must approve, and how every action is recorded and traced. The session covers these rules, the main risks (confidential data, wrong output, security) and the obligations of the EU AI Act that matter to a manager.

The fifth session is a demo fair, where participants, individually or in pairs, build a working solution for a problem from their own work and demonstrate it to the cohort.

The Practice stream follows three rules.

- **Every exercise emphasizes particular skills.** Seven transferable mechanics recur across the program: document interrogation; structured extraction; synthetic data; adversarial review; roleplay and rehearsal; decomposition and chaining; verification. Participants learn what capability they are acquiring, not only what task they are completing.
- **Business example first, own work second.** Each mechanic is introduced through a concrete example from a business function (e.g. a contract, a product recall, a supply shortage, a set of customer calls, a financial statement, a job opening) and is then repeated by each participant.
- **Compose, do not paste.** Ready-made prompts are used to demonstrate structure. Every session contains a composition exercise in which participants write a prompt of their own from scratch and evaluate it.

Exercises are delivered through the Exercise Hub, a web library of some 150 tested exercises with copy-ready prompts, sample files, and reference material. Each cohort has its own link, updated live during the session, which stays accessible after the bootcamp. Each practice block ends with a structured debrief, and we distribute the slides after each session.

Participant requirements. The practice depends on each participant having working access to AI tools. The lab provides computers; participants log into their own accounts. Own laptops are welcome.

- A single paid subscription (ChatGPT Plus, Claude Pro, Google AI Pro, or employer-provided Microsoft 365 Copilot) is sufficient.
- Participants relying on free accounts need several of them, across ChatGPT, Gemini, Claude, and one or two others. Free-tier limits are too low to sustain the exercises, and file upload and analysis are restricted on free tiers.
- Accounts must be created and tested before Session 1. Participants without working access cannot take part in the practice.
- No confidential company or personal data is entered into public AI tools. All exercises use sample, public or anonymised material.

LEARNING OUTCOMES

Upon completing the bootcamp, participants will be able to:

- Understand the fundamentals of AI, generative AI and large language models: what they can do, how they fail, and how to judge their output.
- Compose effective prompts from scratch, using context, personas, goals, reasoning steps, constraints, and format, and use AI itself to improve them (metaprompting).
- Apply generative AI to concrete use cases in their own function: documents and contracts, research and verification, data and reporting, customer and market insight, recruitment and people processes, operations and supply-chain problems.
- Build and maintain a reusable assistant (custom GPT, Gem, Project or Copilot agent) for a recurring task.
- Understand the economic, competitive and organizational implications of AI, including the transformation of work, adoption patterns, governance and the EU AI Act.
- Begin building an AI-first way of working, moving from ad hoc use to systematic practice that raises the productivity of their own work and their team's, and help colleagues do the same.

INDICATIVE USE CASES BY BUSINESS FUNCTION

The list draws on the exercises most used across the program's past cohorts and on recent company seminars. Use cases are chosen to demonstrate the unlimited variety of ways in which imaginative AI uses can remove hours of routine work or improve the quality of a decision. The list is indicative: a selection is used in each cohort, adapted to the mix of functions in the room. Case scenarios are fictional; documents, filings and datasets are public or synthetic.

Business function	Use cases practised
Strategy and general management	A company profile built by metaprompting and deep research · Competitive threat analysis · A strategic decision stress-tested by a council of AI experts · A reusable weekly brief on one's sector: regulation, competitors and market signals, each item with a source and a date · Discovering AI opportunities in the business with AI role cards · Selecting and prioritising AI use cases for one's own unit
Finance and control	Analysis of a listed company's published financial statements · One set of results explained to five audiences: investors, media, employees, the regulator and the bank · Data analysis and dashboards from a spreadsheet · PDF and text to Excel · An acquisition valuation from public filings to a board presentation · A simulated earnings call between a CEO agent and an analyst agent
Sales, business development and customer service	From a client brief to a proposal: what is asked, the questions to send back, the structure and three budget scenarios · A pitch stress-tested by a synthetic client panel that is told not to flatter · Fee and scope negotiation with an AI counterpart · A sales call and a difficult customer call, rehearsed · A preparation pack and a recap for a client meeting, with actions, owners and dates · Customer-service call transcripts analysed for sentiment and recurring complaints
Marketing and communications	Synthetic customer personas and market research · A campaign pre-tested against a synthetic audience that includes people who would reject it · A press release or article from a brief, with every fact checked against the brief and AI-writing markers removed · One approved text turned into a LinkedIn post, an internal note and three social posts, checked for consistency · A brand-accurate slide deck from one prompt · A social-media crisis played over six rounds · A hostile journalist interview, rehearsed
Human resources	Interview practice from the interviewer's or the candidate's seat, on a scenario the participant builds · A job description, synthetic applicants and their evaluation against the criteria · Role architecture and a KPI framework · An AI performance and leadership coach · Preparing a difficult conversation, from a vague request to a reusable coaching assistant · A difficult message rewritten in neutral, constructive and empathetic registers

Business function	Use cases practised
Operations, production and supply chain	A product fault in a new vehicle line: recall exposure reviewed by a panel of expert personas · A semiconductor shortage decomposed into a problem tree with options and trade-offs · Maintenance logs, inspection reports or shift notes extracted into a database · A maintenance-request triage assistant · Checking a product, plant or process against regulatory requirements · Meeting notes turned into a summary with actions, owners and dates
Legal, compliance and governance	Review of a complex contract or terms of service for hidden risks · The company's exposure assessed against a named regulation, with sources and gaps · A contract drafted and then audited by an agent against the EU AI Act and GDPR · Which company material may enter which AI tool, and under what conditions · A company AI acceptable-use policy, researched and drafted · A three-year archive of board minutes read for patterns and unresolved items · The diligence statement: what the person did, what the model did and where the model erred
Data, reporting and small tools	Conversational analysis of a sample dataset · A dashboard from a spreadsheet · A set of reports loaded into a Project and turned into a one-page summary of insights with citations · AI inside Excel and PowerPoint, demonstrated · A no-code mini-tool, such as an expense tracker from receipts or a document converter · A small game coded in one prompt, as a first contact with code
Personal productivity, reusable prompts and assistants	The ten-step anatomy of a prompt, from a vague request to a reusable assistant · Reverse-engineering the prompt behind a finished piece · From a vague request to an executive-ready email · Explaining a complex topic in a one-minute pitch · A prompting coach that questions the prompt instead of answering it · One's own reusable assistant (custom GPT, Gem, Project or Copilot agent) for a recurring task · A personal library of five tested prompts, with a note on when to use each and what to check

ABOUT THE INSTRUCTOR

Dr. Nikos Mylonopoulos is Professor of Digital Business at ALBA Graduate Business School, The American College of Greece, and Director of the SEV Center of Excellence in Creative Leadership. His expertise spans artificial intelligence, digital business, and enterprise systems. He has taught at Warwick Business School and other European universities. He designs and leads executive seminars and workshops for senior business leaders. His research is published in journals including Management Science and Psychology & Marketing. He successfully managed Alba's academic programs during Greece's economic crisis and co-authored the foundational study on digital skills in Greece.