Syllabus · 2026

tripleten

AI Agents & Automation

4–5 months | Online | Part-time

Table of contents

What is AI Agents & Automation 03

Learn to build automations that live between a company’s systems and do not break 05

Program structure 06

Week 0: Entry assessment 07

Sprint 1: How AI works and where it pays off 07

Sprint 2: Building and evaluating an AI solution 09

Sprint 3: The agent in code, at work 12

Sprint 4: A process between teams 14

Sprint 5: Capstone defense 15

Career preparation 16

Active job search 17

What is AI Agents & Automation?

AI Agents & Automation is a 4–5 month online program for professionals with 2+ years in one business function who have run a recurring process and the systems it runs on (CRM, ERP, ticketing, spreadsheets): set them up, fixed what broke, explained them to others. Revenue and sales operations, business and process analysts, implementation and onboarding in SaaS, support and service operations.

It takes that experience, adds the automation and AI engineering skills that job postings ask for, and proves them on four projects in a virtual company environment built on real business cases. You study 10–12 hours a week without quitting your job. Where AI Transformation Management is about carrying a solution through the company so that people use it, AI Agents & Automation is about building the automation yourself, step by step, and keeping it running.

What you’ll learn

- You start with how models and agents actually work and where AI pays off in a business process.
- You learn to set tasks for AI and give it the right context, build automations and agents without code, connect them to the systems a company runs on, and check what AI produces before anyone relies on it. You learn to count the economics of a solution and write the rules it runs under.
- In the second half you go where no-code stops: you write the code that links the steps of an automation and own what an AI coding assistant wrote, integrate systems through APIs, webhooks and MCP, build agents on a corporate no-code platform and in code through a model API, and put the automation into production with monitoring, failure handling and a runbook for the people who support it.

Learn to build automations that live between a company's systems and do not break

From the first week you work as an AI Agents & Automation Specialist in a virtual company, assigned to the department closest to your own experience: operations, customer support, sales, marketing or finance. The role is the same for everyone in the program; the department you start in and the process you work on are yours. The environment is built for the program; the cases, data, systems and approvals in it come from real business practice: 18 months of operational records, process documents, system sandboxes, and the stakeholders you have to convince.

Four projects, each a standalone case

Each has its own artifact and its own before-and-after number. You start with your own department, and project 4 always takes the work to a second team. Together they add up to the capstone you defend at the end in front of a practitioner and a manager who hires for similar roles.

1

Where the team loses time

2

The first AI agent

3

The agent in code, at work

4

A process between teams

Lessons are not tied to one tool. You work in the Microsoft stack and can pick another tool for your project if you can justify the choice; the work is graded by the same criteria.

Program structure

21 weeks

Week 0 is the entry assessment. Weeks 1–9 are the shared core, the same for all programs: seven modules on how AI works, how to build with it and how to judge it. Weeks 10–18 are the program part: three modules on code, agents in production and integrations between systems, and in weeks 17 and 18 you bring the projects together and pre-defend them. Weeks 19 and 20 are the capstone defense. The time is split evenly between the core and the program part.

5 sprints

A sprint is one stage of your capstone and ends with a checkpoint on your project. The modules inside a sprint give the lessons, with hands-on labs on training tasks, that the stage needs.

10–12 h/week

We recommend 10–12 hours a week. You can pause the program once for up to 4 weeks.

Certification

After the defense you receive a certificate that lists the role requirements you closed and links to your final project page with its measured results.

Week 0

Entry assessment

Sprint 0

Where your experience already counts

Before the first module you take an assessment against the requirements of real postings for your target role: what your experience already covers, what the program adds, and what is left to close. The result sets your pace and the support you get; the path is the same for everyone. You also choose the department you start in and the process you will work on, and receive your case brief at the end of the week.

Project kickoff: Project 1 starts in week 1.

Sprint 1 · Weeks 1–3

How AI works and where it pays off

Module C1

How models and agents work

This sprint lays the foundation. You learn how large language models produce an answer, how models differ and how to pick one for a task, and what turns a model into an agent: tools, skills, connectors, MCP, subagents. You learn where models get it wrong or biased and which data must never be uploaded.

By the end you can choose a model and a toolset for a task and name the data that stays out of it.

Microsoft 365 Copilot · Azure AI Foundry · Copilot Studio · ChatGPT · Claude · Gemini

Module C2

Finding the opportunity: baseline and priorities

Where AI gives an effect in a business process and where it does not. You collect candidate processes, interview the process owner, map the process as is, fix the baseline before any change (time, money, errors), compare the options (drop the work, simplify it, plain automation, AI, buy ready-made) and rank them by effect, difficulty and risk.

By the end you have a first business case for the process you chose.

Excel · Microsoft Forms · Visio · Microsoft Whiteboard

Project 1

Where the team loses time

You get 18 months of your department’s data and notes on how the team’s day goes. You map where the time goes and, for each spot, weigh the options: an AI agent, a simple automation without AI, a ready-made tool, or removing the step. You choose what to automate first and how. Result: a report on where the time goes, what the options are and where to start.

Checkpoint 1: business case and baseline

Building and evaluating an AI solution

Module C3

Setting tasks for AI and working with context

You turn a work task into a brief for AI: what is needed, for whom, in what format, how to tell the result is good. You choose the context, documents, data and examples, and what to keep out. You get structured output, build a knowledge base without code, brief an agent (goal, steps, tools, where it stops and asks a person) and rework the brief until the result is stable across different inputs.

By the end you have a first version of an assistant for your process and a set of test inputs.

Microsoft 365 Copilot · SharePoint · Copilot Studio

Module C4

Automations and agents without code

You break the process into steps and decide what AI does, what a plain rule does and what a person does. You build an automation with an AI step: trigger, steps, conditions, a human check. You build an agent with tools and set what it may do, what data it sees and where it hands the work back to a person. You learn when robotic process automation is the right answer and when integration is, and how to debug from run logs.

By the end you have a prototype of your solution.

Power Automate · Copilot Studio

Module C5

Connecting systems: APIs, webhooks, MCP

How systems talk to each other: API requests and keys, webhooks and HTTP steps, custom connectors, and MCP to connect an agent to a working system. Where code is unavoidable, you write a simple script with an AI assistant.

You finish with an action plan for your solution.

Power Automate · Copilot Studio · Azure AI Foundry · Python · VS Code · GitHub Copilot

Module C6

Checking AI output and evaluating agents

You build a test set from real cases: ordinary, edge and critical, with reference answers. You set quality criteria and the price of an error, find the typical failures (invented facts, wrong calculations, an agent doing too much), use a model as a judge where it helps and a person where it must, compare versions so that a fix does not break what worked, and set up guardrails, acceptance testing and monitoring after launch.

By the end you have an evaluation report with a version comparison and one failure analyzed in full.

Copilot Studio · Azure AI Foundry · Excel

Module C7

Rules for working with AI, risk and the economics of a solution

The risks of your use case: data, compliance, model errors, reputation. Which data may go into which tool, who needs which access, and how to check it in practice. Rules built into the solution itself: limits for the agent, a human check, an action log. What regulation requires (GDPR, HIPAA, SOX, the EU AI Act, NIST AI RMF). The full cost of the solution: licenses, model calls, people’s time on checking, support. Build or buy, and how to assess a vendor. The business case and the plan to measure the effect after launch.

By the end you can defend an action plan: rules, risks, total cost, go or no-go.

Excel · Copilot Studio · Power Automate

Project 2

The first AI agent

You build an AI agent that automates one of the team’s processes: for example, it reads a new request, sorts it and drafts a reply. You run 200 past records through it and check the result against what really happened, work out what it costs and how many hours it saves the team, and defend the plan to the head of the department, played by a mentor. Result: a working agent, a report on how often it gets it right, and a payback estimate.

Checkpoint 2: defense of the plan

The agent in code, at work

The core taught you to build and judge a solution; the program part teaches how to make it live between systems and keep it running.

Module TC1

Code for automations

You write Python scripts with an AI agent that link the steps of an automation: pull the data, process it, pass it on. You work with an AI coding assistant and own what it wrote: read it, understand it, fix it.
You read and edit JavaScript and TypeScript with AI where the platform runs on it. You use Git the right way: branches, pull requests, review, rollback.
You write automated tests and evals that run before every change, use SQL to pull and check the data for the automation, and find the cause of an error from the logs.

Azure AI Foundry · Python · VS Code · GitHub · GitHub Actions · GitHub Copilot · Promptfoo

Module TC2

Agents in code and production

You build an agent through a model API: tools, steps, function calls. You manage prompts and context in code (templates, versions) and connect the agent to the knowledge base (RAG). You set limits for the agent, a human check and an action log.
You put it into production: environments, schedule, error handling, retries, API limits; logs, monitoring, alerts and failure investigation; service accounts, secrets and least privilege; and a support guide for the people who will run it.

Azure AI Foundry · Azure Functions · Key Vault · Application Insights · LangChain · LangGraph · Langfuse · GitHub Actions

Sprint 3 · Weeks 10–13

Project 3

The agent in code, at work

You rebuild the agent from project 2 in code with an AI coding assistant and connect it to your department’s working system, so it takes the action on its own. You write checks that run before every change and track what the agent costs and where it makes mistakes. Result: an agent that handles every new record by itself, the code in a repository, and a guide for IT.

Checkpoint 3: the agent at work

A process between teams

Module TC3

System integrations and no-code agents

You connect systems through REST APIs and webhooks: documentation, keys, OAuth, data formats. You learn how integration platforms work and how to choose one, automate work in the CRM (objects, events, permissions) and link the CRM, the ticketing system and the ERP into one end-to-end process. You build an agent on a corporate no-code platform (knowledge sources, actions, publishing to the team channel, permissions), connect an MCP connector to a company system, and decide where a no-code agent is enough and where code is needed.

Power Automate · Copilot Studio · Teams · Dataverse · HubSpot · Salesforce · ServiceNow · Workato · n8n

Project 4

A process between teams

Project 3 was one agent in code; here you connect several systems so work moves from team to team without anyone copying data by hand. For example: deal closed → CRM updated → client gets an email → accounting gets what it needs for the invoice. You build it without code and explain why code isn't needed here. Result: an automation between two teams and the number of manual steps removed.

Checkpoint 4: pre-defense with a working demo, weeks 17–18

Capstone defense

Defense

An automation system for the company

You bring the four projects together into one case: the agent in your department and the automation between two teams, how they fit together, how they are tested and monitored, and what to automate next. Projects 1–4 are the proof that the system works.

Passing earns the certificate with the list of closed role requirements and a link to your project page.

30 minutes online, in front of a practitioner from your field and a manager who hires for similar roles

10 min

Presentation and demo

15 min

Questions and objections from the seats of security, the CFO and the team lead

5 min

Feedback

They ask what a real stakeholder asks: how you know the agent is reliable and what your evaluation set is, what happens when it fails and how a customer reaches a human, what it costs in production and what happens if volume doubles, how IT takes it over and what breaks if you leave.

Career preparation

Prepare for your target role before, during, and after the program

Before the program

- Your target role with its salary range from real postings, which parts of your experience already count, and a target list of 10–15 openings paying no less than you earn today

During the program

- After every project: raw material for your résumé and LinkedIn, the artifact and the number
- Your match with target openings, recalculated as modules close requirements
- Group review by function with people from similar roles, and a peer pair through the program

After the defense

- A résumé built from the four projects, with numbers
- Applications to specific openings through our platform
- AI interview practice for each opening
- A live mock interview with a coach before a real interview
- Offer negotiation preparation

Target role on exit: AI Agents & Automation Specialist

Other titles in postings: AI Automation Specialist · AI Implementation Specialist · CRM and RevOps Automation roles · Automation Specialist

Active job search

For 6 months after graduation you keep the career support and access to the program materials, including those updated to new job posting requirements.

Guidance from the career team on your target openings

Applications, networking and interviews managed through our platform

AI interview practice for every opening you apply to

Live career workshops

A live mock interview with a coach when you land an interview

Move into a role AI won’t replace, without starting from scratch

4–5 months | Online | Part-time

Book a call to talk to a career mentor

tripleten