

AI Agents & Context-Aware Automation

Controlled AI agents that interpret information, use tools and context, and automate bounded workflows while keeping critical decisions under human supervision.

What we offer

EBRE AI develops AI agents for information-intensive and operational processes. Agents can work with documents, structured data, APIs and approved tools to search, interpret, classify, prepare outputs and support the next step in a workflow.

Context can be built around the current user, case, journey or external situation, allowing the same architecture to support monitoring, personalised recommendations, internal assistance, research workflows and repetitive operational tasks.

Capabilities, tools, approvals and audit are defined for each deployment so the agent works within clear operational boundaries.

Where it creates value

Information work

- Search, extract and classify information
- Monitor sources and detect relevant changes
- Summarise cases, documents or opportunities

Context and decisions

- Combine current context with available data
- Generate recommendations and next steps
- Adapt outputs to the user, case or situation

Workflow automation

- Use approved APIs and local tools
- Prepare drafts, records and actions
- Apply human approval and keep an audit trail

Working foundations

Evolution and research engine: Opportunity discovery, source grounding, assessment, pilot design, human review and evidence-based iteration.

Reusable agent base: FastAPI, MongoDB, Keycloak, files, typed tools, approvals, audit and configurable agent behaviour.

Operational agents: Existing assistants for municipal services, incidents, scheduling and internal workflows using constrained tools.

Document and data components: OCR, structured extraction, APIs and data services that can be exposed safely to agent workflows.

How a pilot can start

1. Choose one manual workflow, information task or decision where context and repetition create measurable effort.
2. Connect the minimum sources and tools the agent needs and define what it can read, propose and execute.
3. Run real cases with the agreed supervision and compare results with the current process.
4. Extend the validated agent, tools or workflow to additional users and processes when the results justify it.

What a pilot should prove

Output quality

Does the agent interpret the available information correctly and produce useful, relevant outputs?

Operational value

Does it reduce manual effort, response time or repetitive work in the selected process?

Control and reliability

Can the workflow operate consistently with clear permissions, approvals, traceability and failure handling?