

JERTE - TIRANA · 15.05.2026

Artificial Intelligence in the service of the State

Challenges and perspectives for public administration

XXXIst European Days of Territorial Representatives

Boris Cuanoud - Prefect of the District of Morges - Canton of Vaud - Switzerland

Presentation outline

01 Generative AI: how it works and its limits

Understanding it in order to use it well

02 AI in administration

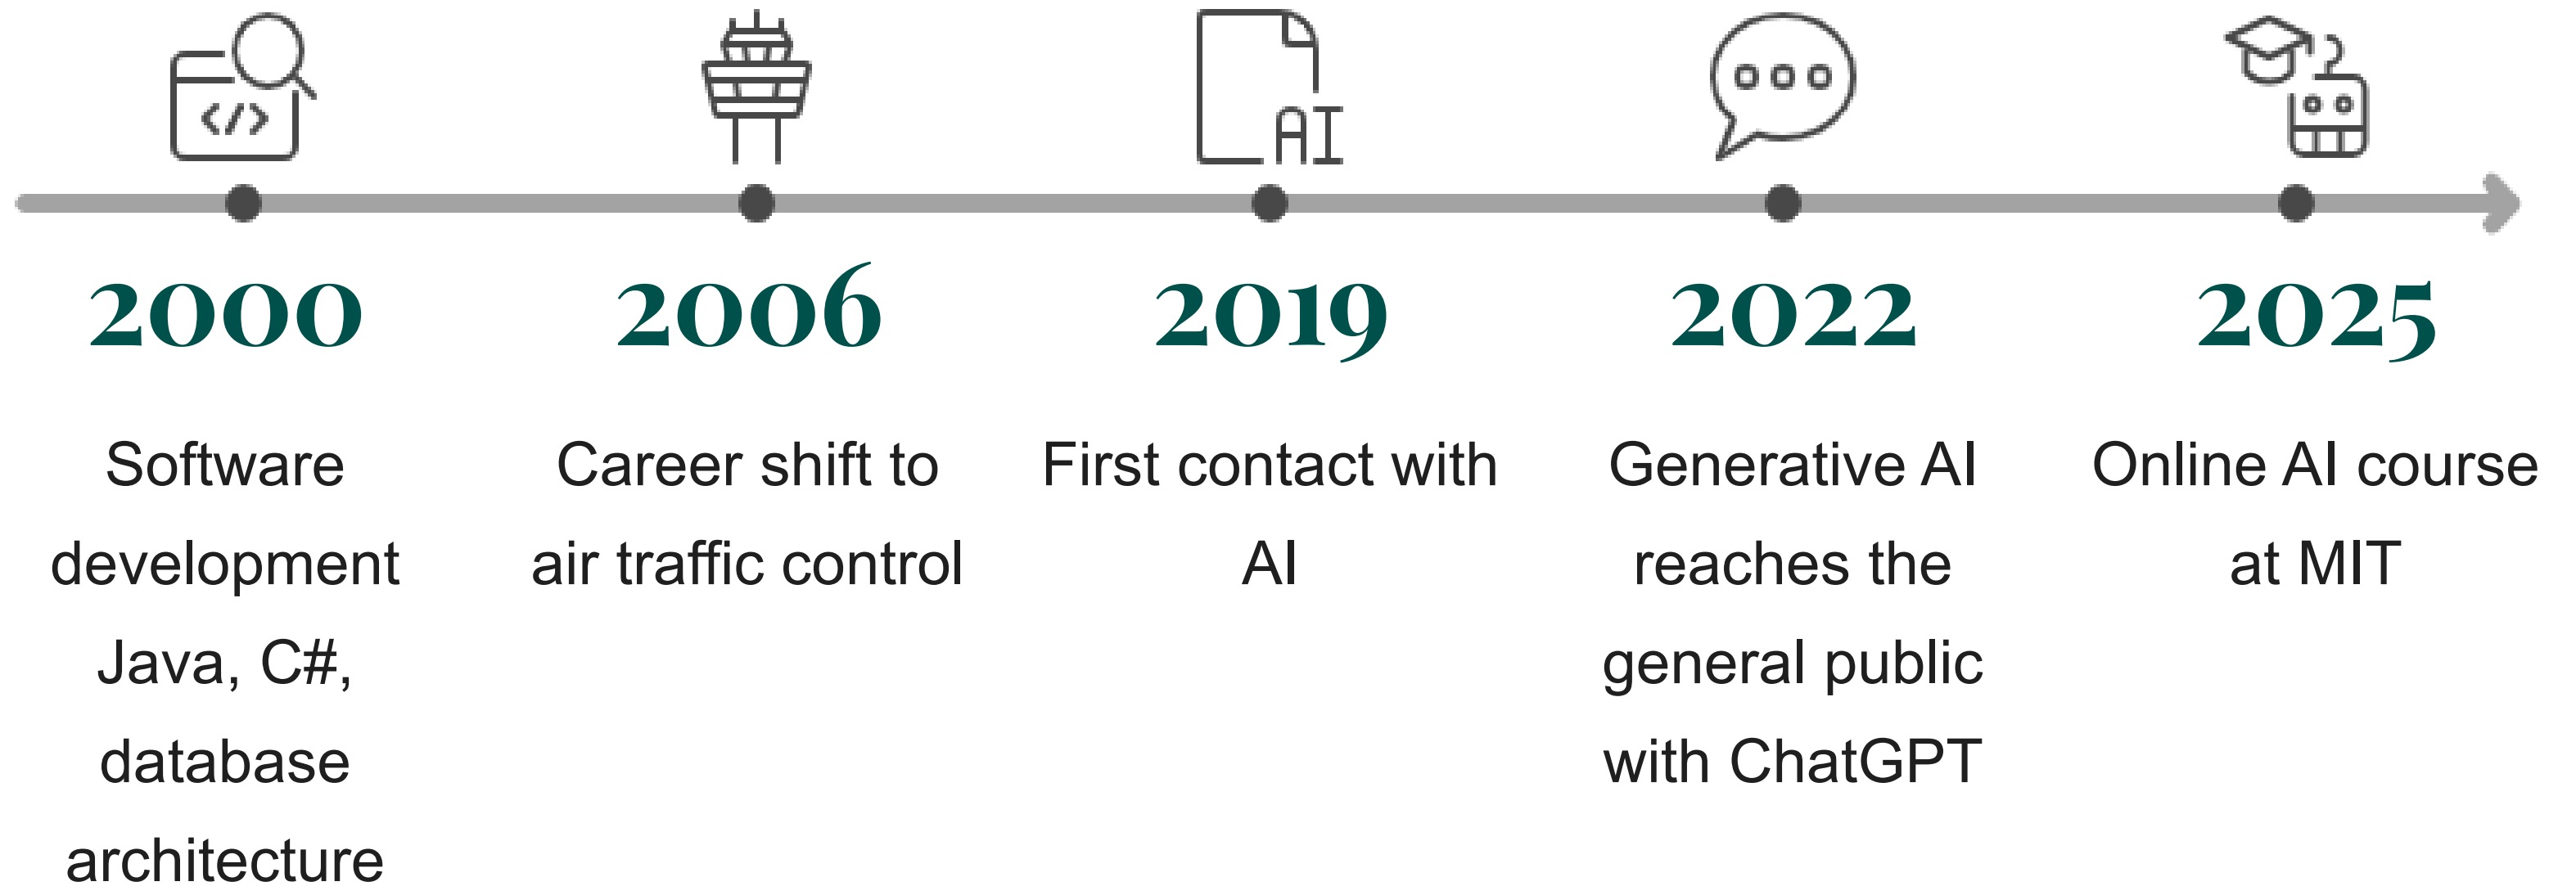
From individual use to its applications in an institutional setting

03 AI in society and its impacts

Ethical and philosophical issues

A prefect's view of AI

My perspective on AI, shaped by earlier experience as a software developer after my university studies



01

Generative AI: How It Works and Its Limitations

Understanding it is the key to using it well.

How does a Large Language Model (LLM) work?

An LLM **predicts the most likely next word** after what you have written. It captures only statistical regularities — enough to produce coherent text, but not to guarantee its truthfulness.

LLMs learn to predict one word at a time
LLMs learn to predict one word at a time
LLMs learn to predict one word at a time
LLMs learn to predict one word at a time

SEARCH ENGINE



Retrieves an **existing word** from a set of indexed documents. The same query returns the same result.

LANGUAGE MODEL



Creates a **new response** each time it is called. No guarantee of reproducibility or accuracy.

Key concepts behind how an LLM works

Tokens

The basic units the model manipulates
→ fragments of words or punctuation.

Dimensions and semantic vector spaces

How the model represents meaning → each token (word) becomes a point in a space with thousands of dimensions.

Analogy: GPS coordinates.

Temperature

A setting that controls how variable the answers are → from deterministic to highly inventive.

Parameters

The model's internal knowledge → values learned during training. A single model can contain billions of them.

Context window

The amount of information the model can draw on during a chat.

Hallucinations and zone of competence

A hallucination is factually incorrect content produced by exactly the same mechanism as correct content. Nothing in the process distinguishes the true from the plausible.

Subtle trap: a mix of true and false in the same sentence, stated with the same confidence.

The model is trained to **always produce an answer**, even without reliable information.

Practical consequence: never accept factual information without verification, especially on important matters.

**GOLDEN
RULE**

**What AI produces is not the end product. It's where our work starts.
Critical thinking remains essential.**

Beyond chat: RAG and MCP

An AI cannot access data it has not been connected to. Without access, **it hallucinates rather than admit it does not know.**

RAG

Retrieval-Augmented Generation

Inject the relevant passages from a targeted document base, for example a legal text repository, into the query.

MCP

Model Context Protocol

A bridge that connects the AI to external tools and databases.

RAG: the AI answers from the documents you've given it, rather than relying only on its general knowledge.

02

AI in administration

From individual use to its applications in an institutional setting

Master your craft before the tool

A 3D modeling tool, however powerful, will only produce a nice drawing of an unliveable building in the hands of someone who does not know how to design a building.

THE ARCHITECT ANALOGY

An **expert in their field** who is new to AI will get better results than an AI expert who does not know the subject.

Good news: if you are good at your job, you already have the most important skill for working with AI.

A word of caution: AI can create an illusion of competence in the person using it.

Best practices in use

ARTICLE 4 — EU AI ACT

Providers and deployers of AI systems must ensure a **sufficient level of AI literacy** among their staff: an understanding of the opportunities, the risks, and the impacts.

- **The GPS analogy:** don't tell the taxi "take me somewhere nice" → aim for precision, not exhaustiveness.
- **Deliberately test the limits** of AI in your area of expertise to sharpen your judgement.
- Don't get hung up on the **type of technology** (they are constantly evolving). The point is to master the fundamentals of generative AI.
- **Iterate, refine, reject** : value comes from the dialogue, not from the first answer.

Use cases for administration

01 RAG on a legal corpus

An MCP that makes all legal and regulatory databases searchable in natural language.

03 Writing assistance

Notes, letters and meeting minutes
→ always under human supervision.

02 Summarising and rewording

Long administrative documents, reports, briefs and meeting minutes.

04 Cross-source search

Across scattered document repositories, where traditional tools fail to bridge the gap.

How should administration approach AI?

- **Keep humans at the heart of decisions:** AI augments judgement, it does not replace it.
- **AI proposes, humans decide:** an AI answer marks the start of our work, not the end → critical re-reading and validation are essential.
- **Set clear limits on use:** what can safely be shared with an AI tool? Pay close attention to data protection and sovereignty.
- **Verify accuracy:** AI can produce inaccurate or fabricated content. Always check the sources.

03

AI in society and its impacts

Beyond our public administrations: upheavals, complementarity and ethical issues.

Economic and social upheavals

JERTE 2026 — "Resilience in the age of vulnerabilities of populations".

- **Labour market:** some tech firms are laying off staff to fund the *data centres* that will be essential by 2030.
- **From mechanisation to AI:** digging a trench for a pipe once required many hands. Mechanisation reduced it to a single operator, multiplied by the machine. AI extends this same logic into administrative and intellectual work. Human capability is augmented by AI.
- **Skills transfer:** professional reconversion — towards which jobs?
- **Pace mismatch:** technological acceleration is exponential, while human adaptation remains linear. Resilience means designing systems that respect the biological tempo without giving up the benefits of innovation.

Human / AI complementarity

WHAT HUMANS BRING

Critical thinking and judgement
Creativity and intuition
Knowledge of real-world context
Final decision

WHAT AI BRINGS

Processing speed
Handling large volumes
Pattern recognition
Constant availability

AI does not replace people, but the person who uses AI may replace the one who does not.

Ethical and philosophical issues

- **Disillusionment:** *deepfakes*, synthetic voices, blurring of truth and falsehood in the public sphere.
- **Humanity will adapt**, but at what cost?
- Human **value** has not vanished; it has *shifted*.

AI will become what we decide to make of it.

IN CLOSING

AI: the question is no longer *whether*, but *how*.

Boris Cuanoud — Prefect of the District of Morges · Canton of Vaud · Switzerland