

RPLT

Law. Tax. Solutions.

21 November 2026, 11.30 – 11.55

W - Amsterdam

Chiara Agostini

AI AGENTIC ON TRIAL

From Character. AI Decision to the governance of agentic systems

INPLP® International Network of Privacy Law Professionals

About RPLT

RPLT. It is an **independent, multidisciplinary firm that integrates legal and tax expertise, established in the market since 1949.**



+200 professionals, including lawyers and chartered accountants



+30 practice areas



7 offices in Italy



5 International Desk

Three questions, one decision

What happened	The Character.AI case: facts, procedure, findings;
What the Italian DPA decided	Violations, corrective orders, sanction;
What it means for agentic AI	Principles transferred by analogy to autonomous agents.

Key rule of this deck

The decision is about a conversational chatbot; agentic AI is our extrapolation.

The case: a companion chatbot, globally deployed



Character Technologies, Inc., Menlo Park (CA), U.S.A.

Data Controller	Character Technologies, Inc., Menlo Park (CA), U.S.A.
Offered Service	Character.AI: **a generative artificial intelligence platform enabling users to create and hold conversations with virtual "Characters", based on proprietary LLMs**.
Essential timeline	• 16 September 2022: web beta — 23 May 2023: app beta; • 8 April 2024: **final version on www.character.ai**, **also offered in Italian**; • November 2024: version dedicated to minors; • 31 May 2025: appointment of the EU representative (VeraSafe Ireland Ltd.); • 1 July 2026: new privacy policy.

Why the italian DPA could decide

Targeting Criterion

Article 3(2)(a) GDPR:
offering services to data subjects in the Union, as the service has been available in Italy and in Italian since 8 April 2024 and the privacy policy also addressed data subjects in the EEA.

No one-stop-shop

Article 56 GDPR requires an establishment (single or main) in the EU:
without one, each supervisory authority independently exercises the powers under Article 58 GDPR.

Operational consequence

A non-EU provider localising its service (language, currency, campaigns) accepts the jurisdiction of each national DPA:
designation of the representative under Article 27 is not a formality, but the condition for engagement with the authority.

The five alleged infringements

#	Issue	Provisions alleged to be infringed
1	Incomplete and unintelligible information (2023 and 2025 privacy policies)	Articles 12.1, 13.1–2, 14.1–2
2	No information on LLM pre-training	Articles 14.1–2, 21.1 and 4
3	Inadequate measures by default for minors	Articles 24.1, 25.2
4	DPIA adopted late	Articles 5.2, 35.1
5	EU representative appointed late and non-functioning contact link	Article 27.1 and 4

Transparency: the defects

Language	English-only information for a service offered in Italian: translation is not a nice-to-have.
Semantics	“Inferences” classified as information provided directly by the data subject; “grievance” instead of “complaint”; US concepts of “sensitive” data and de-identification.
Multiple legal bases	Consent and legitimate interest (or contract and legitimate interest) cannot be invoked for the same purpose: there is a single legal basis, which determines the rights that may be exercised.
Retention	Generic wording tied to the purpose, qualifications (“*unless the law provides for a shorter period*”), probabilistic verbs (“*may be deleted*”).
Transfers	No indication of third countries, adequacy decisions or Chapter V safeguards.
Misrepresentation of processing	Information leading EEA users to believe that post-training was taking place when it had not yet begun.

Training data: no transparency

The Company's defence

Third-party open-source datasets, "incidental" presence of personal data and "unintentional" processing; disproportionate-effort exception under Article 14(5)(b); provision of information through blogs, the Help Center, Reddit and Discord.

The DPA's response

Checking for personal data in datasets is the controller's responsibility under accountability; third-party forums are unsuitable information tools; the Article 14(5)(b) exception requires a documented assessment and published compensatory measures.

Opt-out

No infringement was found, not because the right to opt out was available, but because the obligation to put in place specific mechanisms for unconditional and prior opt-out in relation to pre-training arises from the evolving interpretation provided by the EDPB in Opinion No 28/2024 (adopted on 17 December 2024), which cannot be applied retroactively to Character's pre-training operations, which had ceased prior to that date

Inadequate measures for minors

The initial illicit setup

At the time of its launch, **Character.ai verified the age of users solely on the basis of a self-declaration**, and the service allowed that **even minors** under the age of 15 could have a **public profile by default**: measures that clearly breached the GDPR and the EDPB's interpretation.

The updated compliant framework

Later, during the proceedings, Character Inc. launched a separate LLM for minors, parental insight and blocking of open chats; it introduced, moreover, a multi-tiered age verification system based on a proprietary behavioural age estimation model (which analyses interaction signals), combined, in the event of a user dispute, with biometric or documentary verification via the third-party provider Persona Identities, Inc. (selfies and, where applicable, identity documents, retained for 7 days).

Risk assessment, representative, sanction

DPIA	Presumed high risk: innovative technology, large scale, vulnerable data subjects. The DPIA should have preceded launch whilw Character elaborated the DPIA only on 14 November 2024 and this DPIA was also deficient on minors and pre-training.
Representative	EU representative appointed late
Sanction	€158,000.00
Orders (120 days)	update the privacy policy; determine the fate of pre-training data; ensure the age gate.

From chatbot to agent — what changes

The chatbot

Generates text in response to a prompt; processes user data and, in pre-training, third-party data.

The agent

It follows an objective, plans, maintains persistent memory, calls external tools and APIs, acts on third-party environments and data, and learns from its own execution trajectories.

Why Decision No. 487/2026 remains relevant

A chatbot and an agent are different cases but an agent AI takes to the extreme exactly the issues analised and sanctioned in the decision 487/2026. That’s why its principles are applicable by analogy to an AI agent.

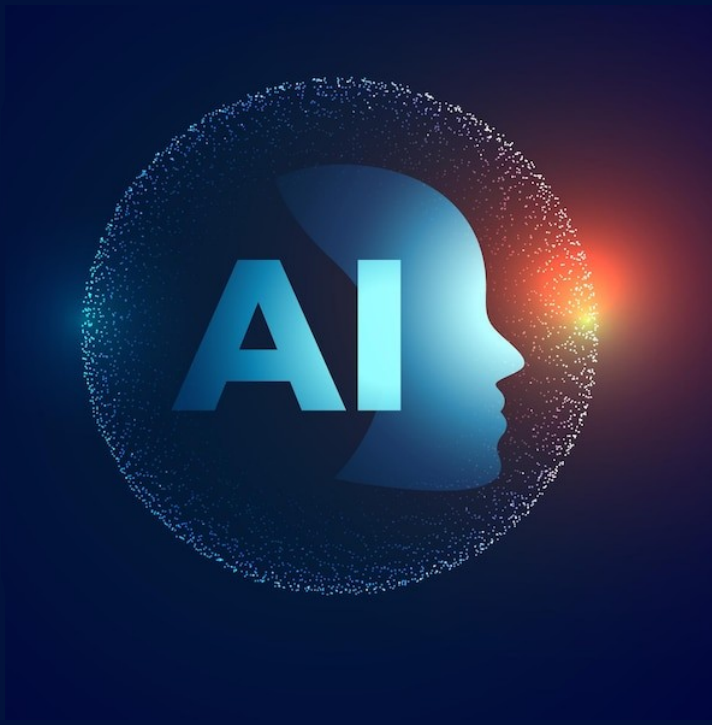

Transparency first

WHY CHARACTER.AI'S DECISION MATTERS FOR AGENTIC AI

1. **Transparency comes before autonomy** The first compliance question is simple: **can the user clearly understand what the system actually does?** The Italian DPA treated unclear, incomplete and misleading privacy disclosures as a central violation, not a secondary issue.

2. **"AI agent" is not a sufficient description**
   If the system remembers past interactions, builds context, calls tools, contacts third parties, or uses data to improve future performance, each of these functions must be described clearly and specifically.

3. **No vague or conditional wording**
   Statements such as "*may use*", "*may transfer*", or "*may improve the service*" are not enough where the processing is real and structured. The notice must describe actual processing activities, not abstract possibilities.

4. **Language matters**
   Information must be accessible, understandable and in the user's language. In the Character.AI decision, even translation quality and misleading terminology were treated as legally relevant.

TAKEAWAY: for AI agents, transparency is not just about saying that AI is used. It is about explaining **how the agent works, what data it uses, and what it is allowed to do**.

One purpose, One legal basis

LEGAL BASIS: NO “MULTI-BASIS” SHORTCUTS

1. **The Italian DPA rejected the “multi-basis” approach**
   The same processing activity cannot be justified at the same time by consent, contract and legitimate interest. For each purpose, the controller must identify **one specific legal basis**, clearly stated.

2. **Each function needs separate legal analysis**
   Agentic systems do not perform just one activity. They may execute tasks, retain memory, generate profiles, call external tools, share outputs, or use interaction data to refine performance. These are **different processing purposes**, **not one single operation**: task execution may rely on contract; long-term memory may require a different basis; profiling, training, or external sharing may require another. The legal basis must follow the actual function, not the label of the product.
   Pursuant to the Italian DPA, it is important to determine each legal basis as it determines what rights the individual can exercise and how the data controller must structure the processing.

TAKEAWAY: For AI agents, compliance starts by breaking the system into functions and asking: what is this processing for, and what is its single legal basis?

Govern the FULL DATA CHAIN



Govern the FULL DATA CHAIN

MEMORY, INFERENCES, THIRD-PARTY DATA AND TRANSFERS

1. **Inferences are real personal data issues**
   The DPA treated inferential outputs as a substantive privacy issue. If the system draws conclusions about preferences or behaviour, those outputs must be explained, limited and justified.

2. **Memory is not neutral storage**
   If an AI agent stores prior prompts, preferences, profiles, tool outputs or behavioural patterns, that memory becomes part of the processing architecture and must be governed accordingly.

3. **Third-party data remain the controller's responsibility**
   If the agent reads emails, queries databases, browses external sources or contacts other persons, the use of third-party data cannot be treated as incidental. Accountability remains with the controller.

4. **Extra-EEA transfers must be mapped**
   If the agent relies on external providers, APIs, cloud services or tools outside the EEA, the controller must know where the data goes and on what transfer mechanism that flow relies.

The italian takeaway - what the character.Ai decision suggests — by analogy — for ai agents

1. **SAY CLEARLY WHAT THE AGENT DOES**
   A **clear, complete, non-misleading privacy notice**, in the user's language: what the agent remembers, which tools it calls, who it contacts, what it learns from interactions. No vague or conditional wording.

2. **ONE LEGAL BASIS PER PURPOSE**
   No "*multi-basis*" shortcuts: each function — task execution, memory, profiling, training, tool use, data sharing — needs **its own single legal basis**, clearly stated.

3. **GOVERN MEMORY AND INFERENCES**
   Profiles and conclusions drawn by the agent are personal and inferred data: they must be explained, limited and justified.

4. **OWN THE FULL DATA CHAIN**
   Third parties' data, retention periods and extra-EEA transfers across the whole action chain — prompts, logs, memory, tool outputs, providers — remain the controller's responsibility.

5. **RESTRICTIVE DEFAULTS**
   The agent need to start with **minimum necessary permissions**: no access to tools, external systems or public outputs unless specifically enabled for a defined purpose.

6. **HUMAN APPROVAL FOR HIGH-RISK ACTIONS**
   Sending data, contacting third parties, transactions, sensitive data, irreversible steps: autonomy stops, **human authorization begins**.

YOUR CASES, PLEASE

Have you handled — or seen — a decision on agentic AI in your jurisdiction or a case where the principles are applicable by analogy to agentic AI?

Thank you!

RPLT



Chiara Agostini
Partner TMT Privacy

chiara.agostini@rplt.it

+39 02 87313335