Comments on the EDPB Guidelines 02/2026 on Anonymisation

Reply of the OpenMined Foundation to the public consultation
Peter Ide-Kostic, Head of EU Policy
EU Transparency Register No 6393276102965-57
October 2026

1. Who we are

OpenMined is a non-profit foundation incorporated in the United States (a 501(c)(3) organisation) and a global open-source community of engineers and researchers, with members across the world, including contributors across at least eight EU Member States. Since 2017 we have built free, open-source software around a simple idea: organisations should be able to answer questions from data they cannot see. In the systems we build, data stays with its owner. A researcher sends in a question, the owner's systems compute the answer, and only the answer travels back. Universities, companies, and public institutions use this software to study sensitive data without copying or moving it.

We submit these comments from that practical experience. Our work is philanthropically funded; our principal donors are Coefficient Giving (formerly Open Philanthropy) and the Future of Life Institute. Our interest is not commercial: all of our software is free and open source.

2. Summary of our comments

That expertise exists, and the guidelines would be stronger for naming it. Differential privacy is the established mathematical method for exactly this problem. In plain terms: a way of releasing results is differentially private if what comes out would look essentially the same whether or not any one person's data had been included. The strength of that guarantee is a number. It can be chosen in advance, written into a policy, and checked by an auditor. Because the output barely depends on any single record, the three criteria are addressed at the root: a record that has almost no influence on the output cannot be isolated from it, linked through it, or inferred from it.

This is not a laboratory idea. The United States Census Bureau used differential privacy to publish the results of the 2020 census, and the method protects everyday usage statistics in widely used consumer software. The 2014 Opinion that these guidelines replace discussed it. If the successor text says nothing about it, readers may take the omission as a judgement, and assessments will remain qualitative where they could be quantitative.

Our suggestion. In sub-section 3.5.1, acknowledge differential privacy as the state of the art for bounding the risks that the three criteria test, and state that outputs released through an appropriately configured differentially private mechanism provide strong evidence that the criteria are met.

> **EXAMPLE IN PRACTICE**
>
> A statistics office wants to publish the average income of every neighbourhood in a city. Published naively, the figure for a small neighbourhood could give away an individual household's situation, much as the engineers' salaries do in Example 16 of the guidelines. Published with differential privacy, each figure is computed so that it would be almost unchanged whether or not any one household's data had been used. Readers still learn what neighbourhoods earn. But no figure meaningfully depends on any single household, so there is nothing to trace back, whatever other datasets a reader may hold. The protection is a property of the method itself, not an assumption about who might attack it.

5. Comment 3: computing without reading (paragraph 91)

Paragraph 91 discusses encryption as a way to limit access to data, and asks assessors to check whether anyone could decrypt it. That treatment assumes that data must at some point be decrypted to be used.

Several mature technologies break that assumption. Secure enclaves run computations inside a sealed compartment of a processor that even the machine's administrator cannot look into. Secure multi-party computation lets several organisations calculate a joint result while each one's data remains, at every moment, in fragments that are meaningless on their own. Homomorphic encryption allows calculations to be performed directly on encrypted data, producing an encrypted result. In each case, the party doing the computing never holds intelligible data.

This matters for the analysis of "means reasonably likely to be used". An entity whose only involvement is to run computations inside such a mechanism is in a different position from an entity that holds an encrypted file and might one day obtain the key.

Our suggestion. Extend paragraph 91 to state that technologies which compute on data without rendering it intelligible to the computing party (secure enclaves, secure multi-party computation,

and homomorphic encryption) should be taken into account when assessing which entities have access to the data and with what means.

> **EXAMPLE IN PRACTICE**
>
> Two hospitals want to know whether patients recover faster when two treatments are combined. Each hospital has treated some of the patients, so the answer lies in their two sets of records taken together. Neither hospital may show the other its records.
>
> With secure multi-party computation, the records never move. Each hospital's own computer takes part in a joint calculation, and what the two computers send each other are only scrambled numbers that mean nothing on their own: no names, no diagnoses, no pieces of any record. At the end, a single figure comes out for both hospitals, the recovery statistic across all patients.
>
> Both hospitals have learned the answer. Neither has seen the other's records, and someone eavesdropping on the connection would have seen only the scrambled numbers.

6. Comment 4: AI models that never absorb the data (paragraph 82)

Paragraph 82 observes that AI models and synthetic data represent correlations found in the original information, and that prompting a model can draw out information about a particular individual. We agree, and we think the final text can go one step further: this risk follows from a design choice, not from AI as such.

The risk arises when all training data is blended into a single set of model parameters. Once that has happened, personal details may be memorised, they can sometimes be extracted, and they cannot be selectively deleted. There are designs that avoid the blending. In retrieval-based and federated systems, source data stays where it lives, under its holder's access controls. The model consults a source at the moment it answers a question, the consultation obeys whatever rules the holder has set, and when the holder withdraws permission the consultations stop. The model itself holds no copy from which anything could be extracted.

Recognising this distinction would give the assessment a cleaner object. For a model that absorbed the data, the assessor faces an opaque set of parameters whose contents are hard to characterise, which is the difficulty paragraph 82 describes. For a system that consults data at answer time, the assessor can examine the released answers and the access rules, which is a much better defined task.

Our suggestion. Note, in paragraph 82 or in sub-section 3.5.1, that architectures which keep source data separate from the model and under the source's access control materially reduce the risks of memorisation and extraction, and that this design choice is relevant to the assessment of anonymity.

EXAMPLE IN PRACTICE

A company builds an assistant that answers doctors' medical questions. It can build it in two ways.

Built the first way, the assistant is trained directly on millions of patient records, which are blended into the model itself. Years later, an unusually phrased question may still draw out the details of a patient with a rare condition, and there is no reliable way to take that one patient back out of the model.

Built the second way, the assistant contains no patient records at all. When a doctor asks a question, the assistant sends a request to the participating hospitals, and each hospital's own system decides, under its own rules, what may be used for the answer. Every request is logged. If a hospital leaves the programme, its records simply stop being consulted from that day on.

7. Closing remark

Anonymisation done well is what allows data to serve science and public policy while leaving people protected. The draft guidelines describe the standard to meet; our suggestions aim to help them also describe how to meet it. We thank the EDPB for the opportunity to comment, and we offer our community's technical expertise, including practical demonstrations of the systems described above, in support of the scientifically grounded expertise that paragraph 84 calls for.

The OpenMined Foundation can be reached through the contact details provided with this submission.