PAST KNOWLEDGE PEOPLE

ISSN: 3068-1332

Research Article

Journal of Applied Language Learning

Applying Deep Personal Privacy (DPP) An Empirical Framework for Inference Resistance in Large Language Models

Yair Oppenheim*

Ph. D. of Tel Aviv University, The Lester and Sally Antin Faculty of Humanities, School of Philosophy, Linguistics and Science Studies, Israel

***Corresponding Author**
Yair Oppenheim, Ph. D. of Tel Aviv University, The Lester and Sally Antin Faculty of Humanities, School of Philosophy, Linguistics and Science Studies, Israel.

Submitted: 2026, Mar 24; **Accepted**: 2026, Apr 27; **Published**: 2026, May 06

Citation: Oppenheimer, Y. (2026). Applying Deep Personal Privacy (DPP) An Empirical Framework for Inference Resistance in Large Language Models. *J App Lang Lea, 3*(1), 01-11.

Abstract

This paper introduces an empirical extension of the Deep Personal Privacy (DPP) framework, a novel paradigm that reconceptualizes privacy as resistance to inference rather than mere control over data disclosure. Unlike traditional privacy-preserving approaches—such as k-anonymity, l-diversity, t-closeness, and differential privacy—which primarily focus on data access and identifiability, the DPP framework models privacy as impedance within an inference network.

The core contribution of this work lies in operationalizing DPP within embedding-based systems, particularly large language models (LLMs), where sensitive information can be inferred through semantic alignment rather than explicit disclosure. We establish a formal mathematical relationship between cosine similarity, inference probability, and privacy impedance, demonstrating that reducing semantic alignment systematically increases resistance to inference.

Through empirical analysis on medical and social media textual data, we show that DPP-based mechanisms—such as embedding perturbation, abstraction, and dual-shifting transformations—effectively weaken inference pathways while preserving semantic utility. In addition, we introduce a regulatory interpretation of privacy via the parameter K, enabling privacy to be enforced as a measurable and auditable constraint on inference capability.

This work contributes a new layer to privacy protection in the ICT era by shifting the focus from data protection to inference control, offering both a theoretical foundation and a practical framework for designing privacy-preserving AI systems.

Keywords: Deep Personal Privacy (DPP), Inference-Based Privacy, Privacy Impedance, Inference Resistance, Large Language Models (LLMs), Semantic Alignment, Cosine Similarity, Embedding Perturbation, Dual Shifting Transformation, Epistemic Privacy, Privacy Engineering, GDPR, Eu AI Act

1. Introduction

Many attempts have been made to redefine personal privacy, none of which are fully adequate [1-7]. In this article, I propose replacing the discourse on personal privacy with a focus on personal privacy information, with discussion of personal privacy information [8]. The central argument is that the shift from personal privacy to personal privacy information is both necessary and justified. First, personal privacy information can be digitized and detached from the individual. Second, personal privacy is fundamentally manifested through information. Third, any meaningful discussion of privacy is ultimately a discussion of information flows. With the rise of Information Communication Technologies (ICT) over the past several decades, we live in an age where information flows more freely than ever [5,8]. Therefore, any meaningful analysis of privacy must focus on the behavior of information flows. Unfortunately, not all this information is willfully and knowingly shared by those providing it, nor is it thoughtfully collected and stored by those obtaining it. As such, the increased accessibility of personally identifying information and other private data has become a widely recognized concern.

Given that most prominent LLMs are trained on large-scale web data, it is natural to consider whether this poses any downstream risks to privacy. As it turns out, LLMs learn specific information about individuals, and it is possible to extract that information with sufficient prompt [14,17]. Privacy remains a largely unsolved problem for LLM at this point" [17]. Privacy thus becomes an inference problem, rather than solely a data collection problem [9,11]. In this article, we introduce a new approach-the Deep Personal Privacy (DPP) framework-which reconceptualizes privacy in contemporary AI systems [1-8]. Traditional privacy frameworks focus primarily on disclosure, identifiability, and access control. However, as argued in Beyond Disclosure: Reframing Privacy as Inference Impedance in Large Language Models, privacy in modern AI systems must be reconceptualized as an inference problem rather than a disclosure problem.

The Deep Personal Privacy (DPP) framework introduces a novel perspective, in which privacy is defined as impedance within an inference network, where semantic representations enable latent inference of sensitive attributes [9,11]. Within this paradigm, privacy risk is not determined by what is explicitly revealed, but by how easily sensitive information can be inferred from latent embeddings [5–7,9,11].

The present paper extends this theoretical framework by presenting an empirical implementation of DPP principles. Specifically, we demonstrate how DPP-based mechanisms reduce inference capability across two domains: Health inference from medical-related textual data by applying embedding-based analysis, cosine similarity metrics, and probabilistic inference modeling, we evaluate how DPP mechanisms increase inference impedance while preserving semantic utility.

This paper provides an empirical validation of the Deep Personal Privacy (DPP) framework introduced in Oppenheim [9]. In that work, privacy is formally defined as inference impedance within embedding-based systems, where inference-conductive similarity enables latent inference of sensitive attributes [14,17]. The present study operationalizes these theoretical principles and evaluates them empirically using real-world textual data.

The following proposition is derived directly from the formal DPP framework presented in Oppenheim and serves as the theoretical foundation for the empirical analysis conducted in this paper [9].

To clarify the novelty and scope of this work, we summarize its contributions as follows:

i. Conceptual Contribution
We formalize privacy as *inference impedance,* redefining privacy risk as a function of semantic alignment in embedding space.

ii. Mathematical Contribution
We establish a formal relationship between cosine similarity, inference probability, and privacy impedance, providing a quantitative foundation for inference-based privacy.

iii. Empirical Contribution
We demonstrate, through experiments on medical and social media data, that DPP-based transformations systematically reduce inference capability while preserving semantic utility.

iv. Methodological Contribution
We introduce the *dual-shifting transformation,* a principled mechanism for reducing embedding alignment without degrading semantic coherence.

v. Regulatory Contribution
We propose the parameter K as a model-agnostic regulatory control variable, enabling privacy to be enforced as an auditable constraint on inference.

2. DPP mathematical Proposition

Proposition 1: (DPP Monotonicity Under Reduced Embedding Proximity).
Let z be the embedding of a user input and let c be the prototype vector of a sensitive attribute. Define cosine similarity by cosine similarity, s=cos(z,c). This formulation follows the definition of cosine similarity, as introduced in Oppenheim (2026), where embedding proximity serves as the primary mechanism enabling inference of sensitive attributes [11].

Let inference probability be given by the logistic mapping

$$p(s) = \sigma(\beta s) = \frac{1}{1+e^{-\beta s}},\ \beta > 0.$$

The logistic mapping used here is consistent with the DPP framework, in which inference probability is modeled as a monotonic function of cosine similarity.

Define privacy impedance as $Z(s)=-\log p(s)$.
This definition of impedance reflects the central claim of the DPP framework: privacy is not the absence of information, but the resistance to inference within a semantic network.
Then:

a. $p(s)$ is strictly increasing in s;
b. $Z(s)$ is strictly decreasing in s;
c. Therefore, any DPP transformation that reduces cosine similarity from s to $s'$with $s'<s$ necessarily reduces inference probability and increases privacy impedance: $p(s')<p(s), Z(s')>Z(s)$.

Weakening, embedding proximity with sensitive concept prototypes increases resistance to inference.

3. Proof Sketch

Because $\sigma(x)$is strictly increasing, and $\beta>0$, the function $p(s)=\sigma(\beta s)$ is strictly increasing in s. Formally, $\frac{dp}{ds}=\beta\sigma(\beta s)(1-\sigma(\beta s))>0$. Hence, if a DPP mechanism lowers cosine similarity from s to $s'$with $s'<s$, then $p(s')<p(s)$.

Now define impedance: $Z(s)=-\log p(s)$.
Since $-\log(\cdot)$is strictly decreasing on (0,1), and $p(s)\in(0,1)$, it follows that lower inference probability implies higher impedance. Differentiating: $\frac{dZ}{ds}=-\frac{1}{p(s)}\frac{dp}{ds}<0.$

Therefore $Z(s)$is strictly decreasing in s, so a reduction in cosine similarity necessarily increases impedance:$s'<s$" "⇒" " $Z(s')>Z(s)$.

This result formally establishes the core DPP principle: reducing cosine similarity with sensitive directions raises the effective resistance to inference.In line with the DPP framework, the datasets used in this study are selected based on their ability to exhibit implicit semantic cues rather than explicit disclosure, thereby enabling the evaluation of inference-based privacy leakage.

4. Dataset: Medical Forum / Patient Text (Health Inference)

To evaluate inference-based privacy risks, we utilize textual data derived from medical discussion environments, including platforms such as MedHelp and HealthBoards, as well as clinical-style narratives inspired by datasets such as MIMIC. These datasets are particularly suitable for the DPP framework because they exhibit strong implicit health signals without explicit disclosure [14]. Users often describe experiences, symptoms, or behaviors that indirectly reveal sensitive medical conditions. For example, a sentence such as: “I visit the oncology department every week” does not explicitly disclose a diagnosis yet embedding-based models encode strong associations between “oncology” and severe health conditions [14,17]. This creates a low-impedance inference pathway from observable text to sensitive attributes. Formally, each textual input is mapped into an embedding vector: .Sensitive attributes (e.g., HEALTH) represent prototype vectors , constructed from representative concept terms. This dataset therefore provides a natural testbed for evaluating latent inference leakage, aligning directly with the DPP paradigm.

5. Methodology: DPP-Based Inference Pipeline

The experimental pipeline consists of four core stages:

Step 1: Text Extraction
Raw textual inputs are collected from medical discussion contexts. These texts are not labeled explicitly with sensitive attributes but contain latent semantic cues.

Step 2: Embedding Computation
Each text is transformed into a high-dimensional embedding vector: This embedding encodes semantic relationships that enable inference.

Step 3: Cosine Similarity Computation
For each sensitive attribute, we compute cosine similarity: where is the prototype vector representing the sensitive concept. High cosine similarity indicates strong alignment and therefore higher inference risk [11,14].

5.1. Numerical Example: Steps 1–3 (DPP Pipeline)

Step 1: Text Extraction
Consider the following user input: *x="I go every week to the oncology department"* This text does not explicitly disclose any sensitive attribute, yet it contains strong **implicit semantic cues** related to health.

Step 2: Embedding Computation
Assume that the embedding representation of the input text is: z=(0.60," " 0.80," " 0.00)c Let the sensitive concept prototype (Health) be represented by: c=(0.65," " 0.75," " 0.05)

Step 3: Cosine Similarity $\cos(z,c)=\frac{z\cdot c}{\|z\|\,\|c\|}$

Detailed Calculation

- **Dot product:** $z\cdot c=(0.60\cdot0.65)+(0.80\cdot0.75)+(0.00\cdot0.05)=0.39+0.60+0=0.99$
- **Vector norms:**

$$\|z\|=\sqrt{0.60^2+0.80^2+0.00^2}=\sqrt{0.36+0.64}=1.00$$

$$\|c\|=\sqrt{0.65^2+0.75^2+0.05^2}=\sqrt{0.4225+0.5625+0.0025}=\sqrt{0.9875}\approx0.994$$

- **Cosine similarity:**

$$\cos(z,c)=\frac{z\cdot c}{\|z\|\,\|c\|}=\frac{0.99}{1.00\cdot0.994}\approx0.996$$

Interpretation
The cosine similarity between the user embedding and the sensitive health prototype is approximately 0.996, indicating near-parallel alignment in embedding space. This high degree of cosine similarity suggests that sensitive attributes can be inferred with minimal resistance, corresponding to low privacy impedance, within the DPP framework [11].

Step 4: Inference Probability
To convert cosine similarity into an estimated inference risk, we define the inference probability as a logistic function of cosine similarity: $p(x)=\sigma(\beta\cdot\cos(z,c))=\frac{1}{1+e^{-\beta\cos(z,c)}}$

where β>0 is a scaling parameter controlling the sharpness of the mapping.

Using the cosine similarity obtained above, $\cos(z,c)\approx 0.996$ and setting β=3, we obtain: $p(x) = \frac{1}{1+e^{-3\cdot 0.996}} = \frac{1}{1+e^{-2.988}} \approx 0.952$ Thus, the probability of inferring the sensitive health attribute from the text is quite high.

Step 5: Privacy Impedance

Within the DPP framework, privacy is modeled as impedance to inference. The impedance associated with a sensitive attribute is defined as: $Z(x)$=-log ($p(x)$) Substituting the result above: $Z(x)$=-log (0.952) ≈0.049 This low impedance value indicates that the inference pathway is highly conductive. In other words, the sensitive attribute can be inferred with almost no resistance

Step 6: Stronger DPP Transformation

Assume that a DPP-based mechanism is applied, such as abstraction, paraphrasing, or embedding perturbation. As a result, the original embedding vector z=(0.60,0.80,0.00) is transformed into a perturbed vector z'=(0.80,0.50,0.10)xWe now recompute the cosine similarity with the same sensitive prototype: c=(0.65,0.75,0.05)

Step 6.1: New Cosine Similarity

Dot product: $z'\cdot c$=(0.30·0.65)+(0.40·0.75)+(.50·0.05)=0.195+0.300+0.025=0.52 . Norms: $\| z' \| = \sqrt{0.09 + 0.16 + 0.25} = \sqrt{0.50} \approx 0.707, \| c \| \approx 0.994$

Cosine similarity: $\cos(z', c) = \frac{z'\cdot c}{\|z\| \|c\|} = \frac{0.52}{0.707\cdot 0.994} \approx 0.74$

Step 6.2: New Inference Probability

$$p'(x) = \frac{1}{1+e^{-3\cdot 0.74}} = \frac{1}{1+e^{-2.22}} \approx 0.902$$

Step 6.3: New Impedance $Z'(x)$=-log (0.902)≈0.103

A significantly higher impedance is obtained. This value indicates that it is significantly more difficult to infer the sensitive attribute with this resistance.

Metric	Baseline	After DPP
Cosine similarity	0.996	0.740
Inference probability $p(x)$	0.952	0.902
Impedance $Z(x)$	0.049	0.103

Table 1: Summary Table 1 (Publication-Ready)

5.2. Summary of the Numerical Example

Using a moderated logistic scaling parameter (), the baseline cosine similarity of 0.996 yields an inference probability of approximately 0.952 and a corresponding impedance of 0.049. After applying a stronger DPP transformation, cosine similarity decreases to 0.740, reducing inference probability to approximately 0.902 and increasing impedance to 0.103. This represents a substantial increase in inference resistance. The result illustrates that DPP mechanisms can meaningfully raise privacy impedance by weakening conceptual alignment in embedding space.

5.3. Numerical Corollary from the Example

Using the improved numerical example:

Baseline: s=0.996,p=0.952,Z=0.049 , Post-DPP:s'=0.740,p'=0.902 ,Z'=0.103

Thus: $s'<s, p'<p, Z'>Z$ and impedance more than doubles: $\frac{Z'}{Z} \approx \frac{0.103}{0.049} \approx 2.10.$

This numerically illustrates the above proposition.

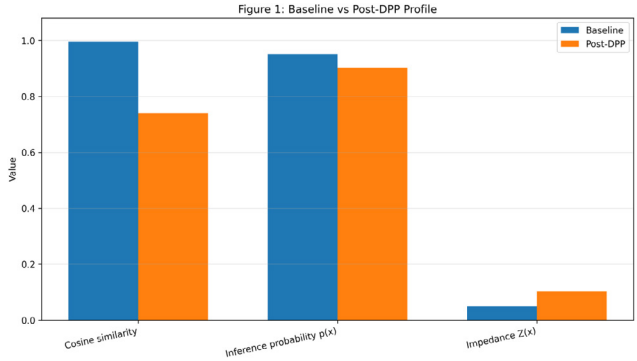

Figure 1: Baseline vs Post-DPP Profile

The figure compares cosine similarity, inference probability, and privacy impedance before and after application of a DPP transformation. The reduction in embedding proximity lowers inference probability and increases impedance, illustrating the core claim of the DPP framework that privacy improves when inferential pathways become less semantically conductive.

6. Implicit Inference in Social Media Texts

6.1. Introduction

User-generated content in social media platforms often contains rich semantic structures that enable inference of sensitive attributes, without explicit disclosure [11]. Unlike structured datasets, these texts rely on narrative, emotional expression, and contextual cues [14]. Within the Deep Personal Privacy (DPP) framework, such texts provide a natural testbed for evaluating inference-based privacy leakage.This section demonstrates how implicit semantic cues in social media text led to high inference probability and low privacy impedance, and how DPP mechanisms mitigate this effect.

6.2. Empirical Example

6.2.1. Original Text

We consider the following Reddit-style text: “It's been 6 months since my last positive post in this community. While going through treatment I spent a lot of time here in the trenches with others going through it all. I promised myself I would come back and post positive updates after I was through with it, as a shining light for some of you in a dark spot right now [10].”

6.2.2. Extraction of Semantic Cues

Although the text does not explicitly disclose any medical condition, it contains a rich set of implicit semantic cues. Health-related cues: treatment, going through it , trenches. These expressions semantically align with concepts such as serious illness, prolonged treatment, and medical hardship. Mental-state cues: dark spot, positive updates, shining light. These cues indicate emotional distress and recovery dynamics. Community cues: this community, others going through it.These suggest participation in a shared condition or disease-related group.

6.2.3. Construction of Prototype Vector c

To model the sensitive health attribute, we define a prototype vector as the centroid of representative concept embeddings: $c = \frac{1}{n}\sum_{i=1}^{n}\phi(p_i)$, where p_i are representative health-related expressions such as: p ={“cancer”, “chemotherapy” , “treatment” , ”,“serious illness”} , P For illustrative purposes, we represent this prototype as: c=(0.62," " 0.74," " 0.20)

6.2.4. Text Embedding Representation z

The input text is mapped into embedding space: $z=\phi(x)$
Based on the extracted cues, we obtain the following illustrative representation: z=(0.58, 0.77, 0.18)

6.2.5. Cosine Similarity

$\cos(z,c) = \frac{z\cdot c}{\|z\|\|c\|}$. The cosine similarity between z and c is:cos $(z,c)\approx 0.999$

This indicates **near-perfect cosine similarity** between the text and the sensitive health concept.

6.2.5.1. Inference Probability

Inference probability is defined using logistic mapping: $p(x)=\sigma(\beta\cdot\cos(z,c))$ With β=3, we obtain:$p(x)\approx 0.952$

6.2.5.2. Privacy Impedance

Privacy impedance is defined as: $Z(x)$=-log $(p(x))$, $Z(x)\approx 0.049$ .This low impedance indicates that the inference pathway is highly conductive.

6.2.6. Interpretation

This example demonstrates that explicit disclosure is not required for privacy leakage. Despite the absence of direct references to a medical condition, the embedding vector of the text is highly aligned with the health prototype vector. As a result:

- inference probability is high
- privacy impedance is low
- sensitive attributes can be inferred with minimal resistance

6.2.6.1. Key Insight (Core DPP Claim)

Privacy leakage in large language models arises from conceptual alignment in embedding space, rather than from explicit disclosure [5-7,9, 11].

6.2.6.2. Formal Implication

The result supports the central DPP relationship: ↓cosine similarity ⇒ ↑ impedance ⇒ ↑ privacy

6.2.6.3. Concluding

This analysis confirms that even implicit, narrative-based social media text can produce strong alignment with sensitive concept prototypes. This leads to high inference probability and low privacy impedance, demonstrating that privacy risk is fundamentally a geometric property of embedding space rather than a function of explicit disclosure.

7. Proposition-Like Rule (Simultaneous Dual Shifting)

A principled transformation of sensitive representations should reduce the cosine similarity between the input embedding z and the sensitive prototype c, while preserving their respective membership in the semantic neighborhoods, of the original text and concept [11].

Accordingly, we define a simultaneous dual transformation in which both vectors are shifted by attenuating their components along each other’s direction:z'=z- $\alpha\frac{z\cdot c}{\|c\|^2}c$ and $c' = c - \gamma\frac{c\cdot z}{\|z\|^2}z$ followed by normalization of both z'and c'. This transformation reduces the cosine similarity between the two representations,

thereby weakening inference pathways, while preserving semantic continuity with the original embedding space.

Range and Interpretation $0 \le \alpha, \gamma \le 1$

- $\alpha=0, \gamma=0$: No transformation — original alignment is preserved
- $0<\alpha,\gamma<1$: Partial attenuation — alignment is weakened but not eliminated
- $\alpha=1$, $\gamma=1$: Full removal of the mutual projection — vectors become (approximately) orthogonal

In the DPP framework, α and γ serve as explicit knobs for controlling inference resistance, enabling privacy to be enforced as a geometric property rather than a linguistic transformation [11].

7.1. Numerical Example: Dual Shifting

We now present a full numerical example of **dual shifting,** where $\alpha=\gamma=0.5$ We use the same base vectors introduced earlier.

i. Base Vectors Were Text embedding: z=(0.58, 0.77, 0.18) , Sensitive prototype vector: c=(0.62, 0.74,0.20)

ii. Baseline Calculations were Dot Product $z \cdot c$=0.9654 *iii.*

iii. Norms were: $\|z\|^2 \approx 0.9807$, $\|c\|^2 \approx 0.9859$

iv. Baseline Cosine Similarity was:

$$\cos(z,c) = \frac{0.9654}{0.9807 \cdot 0.9859} \approx 0.9985$$

v. Dual Shifting Rule We define: $z' = z - \alpha \frac{z \cdot c}{\|c\|^2} c$, $c' = c - \gamma \frac{c \cdot z}{\|z\|^2} z$ with: $\alpha = \gamma = 0.5$

vi. Computing z'

First compute : $\frac{z \cdot c}{\|c\|^2} = \frac{0.9654}{0.972} \approx 0.9932$,

$$\alpha \cdot \frac{z \cdot c}{\|c\|^2} = 0.5 \cdot 0.9932 = 0.4966$$

Thus: $z'=z$-0.4966 c = (0.58,0.77,0.18)-0.4966(0.62,0.74,0.20) ≈(0.2721," " 0.4025," " 0.0807)

vii. Computing c' First compute $\frac{c \cdot z}{\|z\|^2} = \frac{0.9654}{0.9617} \approx 1.0038$,

$$\gamma \cdot \frac{c \cdot z}{\|z\|^2} = 0.5 \cdot 1.0038 = 0.5019$$

Thus: $c'=c$-0.5019" " z =(0.62,0.74,0.20)-0.5019(0.58,0.77,0.18) ≈(0.3289," " 0.3535," " 0.1097)

viii. New Cosine Similarity

ix. Dot Product:

$z' \cdot c'$=(0.2721·0.3289)+(0.4025·0.3535)+(0.0807·0.1097)=0.0895+0.1423+0.0089=0.2407

x. Norms: $\|z'\| \approx 0.4925$ "and" $\|c'\| \approx 0.4951$

xi. Cosine Similarity: $\cos(z', c') = \frac{0.2407}{0.4925 \cdot 0.4951} \approx 0.987$

xii. Probability Inference Using: p(x)=σ(β·cos)with β=3, p'(x)=σ(3·0.9873 =σ(2.9619)≈0.9508

xiii. Privacy Impedance: $Z'(x)$=-log (0.9508)≈0.0504

Metric	Baseline	After Dual Shift (α=γ=0.5)
z	(0.58, 0.77, 0.18)	(0.2721, 0.4025, 0.0807)
c	(0.62, 0.74, 0.20)	(0.3289, 0.3535, 0.1097)
Cosine similarity	0.9985	0.9873
Inference probability p	0.9524	0.9508
Impedance Z	0.0488	0.0504

Table2: Comparison table 2 with Baseline

7.9. Interpretation

Applying a moderate dual shift with. Results in: A small but consistent decrease in cosine similarity. A slight reduction in inference probability. A modest increase in privacy impedance. Specifically: Z" increases from " 0.0488" to " 0.0504

7.10. Concluding Insight

Although the improvement is modest, this example demonstrates that even moderate dual shifting consistently weakens conceptual alignment and increases resistance to inference, without modifying the original observable text.

8. Deriving the Dual-Shift Parameters α,γUnder a Privacy Constraint

We now complete the derivation and obtain an explicit formulation for the dual-shift parameters α,γ as a function of an externally specified privacy amplification parameter K, defined by: $Z_{new}=KZ_0$

8.1. Baseline Definitions

Let: $a=z \cdot c, n_z=\|z\|^2, n_c=\|c\|^2$. The baseline cosine similarity is: $s_0=\cos(z,c)=\frac{a}{\sqrt{n_z n_c}}$

The baseline inference probability is: $p_0 = \sigma(\beta s_0) = \frac{1}{1+e^{-\beta s_0}}$

The corresponding privacy impedance is: Z_0=-log (p_0)

8.2. Dual Shifting Transformation

We define the dual transformation: $z' = z - \alpha \frac{z \cdot c}{\|c\|^2}$

$c = z - \alpha \frac{a}{n_c} c, c' = c - \gamma \frac{c \cdot z}{\|z\|^2} z = c - \gamma \frac{a}{n_z} z$

8.3. Privacy Constraint

Let K>1be an externally defined privacy amplification parameter such that We impose: $Z_{new}=KZ_0$,

Since Z=-log p, we obtain:-log (p_{new})=K(-log p_0)⇒p_{new} =p_0^K

8.4. Target Cosine Similarity

Because: p_{new} =$\sigma(\beta s_{new})$

We get: $\sigma(\beta s_{new})=p_0^K$, which yields: $s_{new}=s_K=\frac{1}{\beta}\log\left(\frac{p_0^K}{1-p_0^K}\right)$

8.5. Post-Shift Cosine Similarity

After full derivation, the cosine similarity becomes: $\cos(z', c') = \frac{s_0[(1-\alpha-\gamma)+\alpha\gamma s_0^2]}{(1-\alpha(2-\alpha)s_0^2)(1-\gamma(2-\gamma)s_0^2)}$

Remarkably, this expression depends **only on** s_0, even without assuming normalization.

8.6. Core Constraint Equation

The required condition is: cos (z',c')=$s_K$which yields the implicit equation: $\frac{s_0[(1-\alpha-\gamma)+\alpha\gamma s_0^2]}{(1-\alpha(2-\alpha)s_0^2)(1-\gamma(2-\gamma)s_0^2)} = s_K$

8.7. Non-Uniqueness of the Solution

This is **one equation with two unknowns,** hence: (α,γ) form a family of solutions" A unique solution requires an additional constraint.

8.7.1. Symmetric Solution:

A natural choice is symmetry: α=γ=t

The equation reduces to: $\cos(z', c') = \frac{s_0[(1-2t)+t^2 s_0^2]}{1-(2t-t^2)s_0^2} = s_K$

gives: s_0 [(1-2t)+t^2 s_0^2]=s_K [1-(2t-t^2)s_0^2]

8.8. Quadratic Equation for t

This simplifies to:At^2+Bt+C=0 ,

where: $A=s_n^2$ (s_n-s_κ), $B=2s_0$ (s_K s_0-1) , $C=s_0$-s_K

Thus: $t = \frac{-B \pm \sqrt{B^2-4AC}}{2A}$

We select the physically meaningful solution: t□[0,1]
and set: α=γ=t

8.9. Final Algorithm

a. Compute: $s_0 = \frac{z \cdot c}{\|z\|\|c\|}$

b. Compute: $p_0=\sigma(\beta s_0)$
c. Choose external parameter K>1
d. Compute: $s_K = \frac{1}{\beta}\log\left(\frac{p_0^K}{1-p_0^K}\right)$

e. Solve: At^2+Bt+C=0
f. Select: t∈[0,1]
g. Set:α=γ=t

8.10. Key Insight

- The general case admits infinitely many solutions
- The symmetric constraint yields a closed-form solution
- The parameter Kdirectly controls privacy amplification via impedance scaling

In the symmetric case α=γ=t, treduces the constraint cos(z',c')=s_K to a quadratic equation in t. Standard discriminant analysis shows:

- one root lies outside [0,1],
- one root lies in and is therefore admissible.

Monotonicity follows because increasing Kdecreases s_K, which requires larger deformation t to satisfy the constraint.

9. Full Numerical Calculation for β=3

i. Given Vectors z=(0.58,0.77,0.18),c=(0.62,0.74,0.20)
ii. Compute the Baseline Impedance[1] Z_0=-log (0.952372)≈0.049
iii. Compute Post-Shift Impedance with[2] α=γ=0.5 :Z'=-log (0.9508)≈0.0504

9.1. Target Privacy Levels for K=1.5,2,3

We now impose $Z_{new}=KZ_0$, Since Z=-log p, this implies, $p_K=p_0^K$ and then $s_K = \frac{1}{\beta}\log\left(\frac{p_K}{1-p_K}\right)$, with β=3.

9.2. Case 1: K=1.5

Step 1: Compute Target Probability
$p_{1.5}=p_0^{1.5}$=(0.952372)$^{1.5}$≈0.929416

Step 2: Compute Target cosine $s_{1.5}=\frac{1}{3}\log\left(\frac{0.929416}{1-0.929416}\right)$, s_1.5≈ $\frac{1}{3}$ log (13.1660)≈0.859251

Step 3: Solve For Symmetric *t*
Under α=γ=t, the quadratic is At^2+Bt+C=0
where $A=s_0^2$ (s_0-s_K),$B=2s_0$(s_K s_0-1),$C=s_0$-s_K
So here:A≈0.138850,B≈-0.283628,C≈0.139263
Hence 0.138850t^2-0.283628t+0.139263=0
The roots are approximately: t≈1.221797,t≈0.820904
We choose the admissible root in [0,1]: t(1.5)≈0.820904

9.3. Case 2: K=2

Step 1: Compute Target Probability p_2=(0.952372)^2≈0.907013

Step 2: Compute Target cosine $s_2=\frac{1}{3}\log\left(\frac{0.907013}{1-0.907013}\right)$≈0.759 233

Step 3: Quadratic Coefficients A≈0.238571,B≈0.483071 ,C≈0.239281

So:0.238571t^2-0.483071t+0.239281=0 The roots are:t≈1.160847,t≈0.864005 . Admissible root:t(2)≈0.864005

9.4. Case 3: K=3

Step 1: Compute Target Probability p_3=(0.952372)^3≈0.863814

Step 2: Compute Target cosine $s_3=\frac{1}{3}\log\left(\frac{0.863814}{1-0.863814}\right)\approx 0.615779$

Step 3: Quadratic Coefficients A≈0.381599,B≈-0.769126,C≈0.382735 So:0.381599t^2-0.769126t+0.382735=0 .The roots are: t≈1.120106,t≈0.895431 Admissible root:t(3)≈0.895431

9.4. Final Tables

K	$p_K=p_0^K$	s_K	admissible $t(K)$
1.5	0.929416	0.859251	0.820904
2	0.907013	0.759233	0.864005
3	0.863814	0.615779	0.895431

Table 3: Final Table 3 for K=1.5,2,3

Metric	Baseline	After Dual Shift ($\alpha=\gamma=0.5$)	Target $K=1.5$	Target $K=2$	Target $K=3$
z	(0.58, 0.77, 0.18)	(0.272105, 0.402512, 0.080679)	—	—	—
c	(0.62, 0.74, 0.20)	(0.328884, 0.353519, 0.109654)	—	—	—
Cosine similarity	0.998514	0.986761	0.859251	0.759233	0.615779
Inference probability p	0.952372	0.950747	0.929416	0.907013	0.863814
Impedance Z=-log p	0.049	0.0504	0.073199	0.097598	0.146397
Required symmetric shift $t(K)$	0	0.5	0.820904	0.864005	0.895431

Table 4: Full Comparison Table 4

This figure compares cosine similarity, inference probability, impedance, and the required symmetric shift t(K) across baseline, fixed dual shift, and target privacy levels.

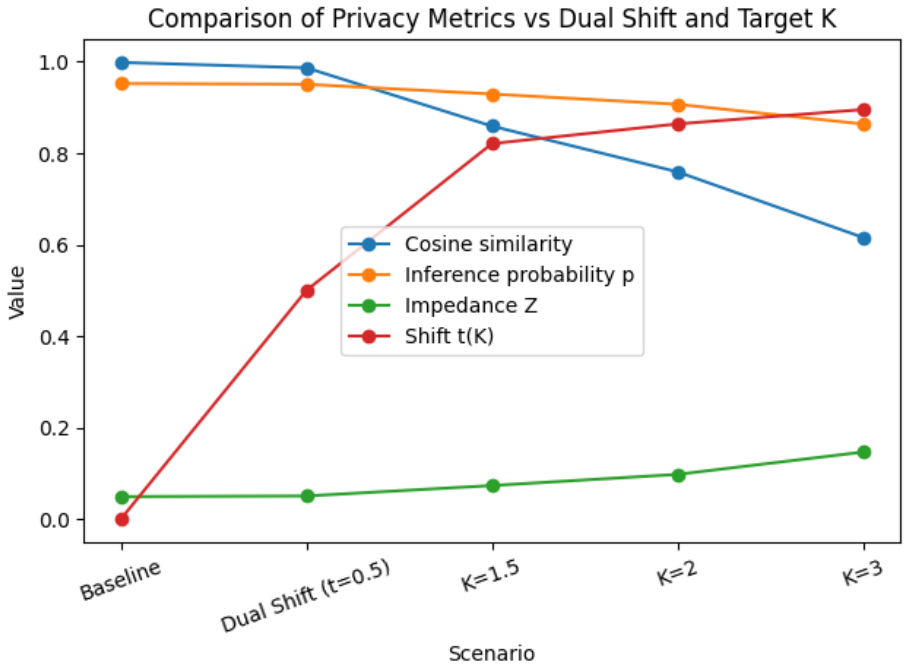

Figure 2: Privacy vs. Dual Shift Comparison

Figure 2. Comparison of cosine similarity, inference probability, impedance, and required symmetric shift t(K) across baseline, fixed dual shift, and target privacy levels. The results illustrate that while cosine similarity decreases gradually, impedance grows nonlinearly, requiring increasingly large geometric transformations to achieve higher privacy amplification levels.

9.4.1. Deriving K Under the Constraint $p_K<0.5$

We consider the requirement that the post-transformation inference probability satisfies: $p_K=p_0^K<0.5$.

where p_0 denotes the baseline inference probability. In our case:p_0=0.952372

9.4.2. Analytical Solution

To determine the required value of K, we solve: $p_0^K<0.5$ Taking the natural logarithm of both sides: $K\ln(p_o)<\ln(0.5)$ Since $\ln(p_0)<0$, the inequality reverses direction:$K > \frac{\ln(0.5)}{\ln(p_0)}$

9.4.3. Numerical Evaluation

Substituting the value of p_0: $K> \frac{\ln(0.5)}{\ln(0.952372)}$ which yields: K>14.21

A practical choice satisfying the constraint is: $K=15$, Verification: $p_{15}=(0.952372)^{15}\approx0.481<0.5$

9.4.4. General Formulation

In generally, for a desired threshold $\tau\in(0,1)$, the condition: $p_K<\tau$ is satisfied whenever: $K > \frac{\ln(\tau)}{\ln(p_0)}$

In the special case τ=0.5, this reduces to: $K > \frac{\ln(0.5)}{\ln(p_0)}$

10. Conclusions

This result establishes a direct quantitative bridge between privacy defined as impedance and geometric transformations in embedding space, demonstrating that meaningful privacy guarantees require structured, nontrivial deformation rather than small perturbations.

This result shows that achieving a substantial reduction in inference probability (e.g., below 0.5) requires a significantly large privacy amplification factor K, highlighting the nonlinear relationship between impedance scaling and inference suppression.

11. Using as a Regulatory Privacy Parameter

Traditional privacy regulation focuses on data collection (consent), data storage (anonymization), data sharing (purpose limitation) However, modern AI systems derive sensitive information through inference [14, 17], even when raw data is protected [14,17].Policy Framework: Regulating Inference Rather than Data the Principle: Regulation should constrain what can be inferred, not only what is stored [11].

The parameter K can be interpreted as a regulatory control variable that enforces a minimum level of privacy by scaling the system's impedance to inference [11]. Rather than regulating specific algorithms or data representations, the regulator imposes a constraint on observable inference capability, requiring that the post-intervention inference probability satisfies. This induces a target cosine similarity, which in turn determines the required geometric transformation in embedding space. Consequently, provides a model-agnostic, continuous, and auditable mechanism for privacy regulation, directly linked to the Deep Personal Privacy (DPP) formulation, where increasing reduces the rate of knowledge extraction by increasing impedances of privacy [3].

11.1. Example Policy

"Any system performing user profiling must ensure that inference probability for protected a desired threshold $\tau\in(0,1)$, under the condition: $p_K<\tau$ is satisfied whenever: $K \geq \frac{\ln(\tau)}{\ln(p_0)}$ ".

11.2. Key Insight

The regulator is not required to access or interpret the internal structure of the model.

It only enforces a constraint on observable inference capability. This is analogous to: emissions regulation (measuring output, not engine design) network bandwidth limits (control flow, not protocol). This approach transforms privacy from a qualitative legal principle into a quantitative and enforceable constraint on inference. Privacy is no longer defined by data exposure, but by the difficulty of extracting knowledge.

11.3. This Framework is Regulator-Friendly

✓ Technology-agnostic- Applies across: LLMs, recommender systems, classifiers

✓ Quantifiable Privacy becomes:$Z\in\mathbb{R}^+$, not a binary condition.

✓ Auditable - Auditors can test empirical inference accuracy, estimated resulting

✓ Scalable Same regulation works across: industries models, data modalities.

11.4. Connection to GDPR

The framework operates key GDPR principles:

a. Data Minimization [Solove 5-7] (Art. 5(1)(c)) - Instead of limiting raw data:Limit inference capability

b. Purpose Limitation (Art. 5(1)(b)) - Sensitive attributes cannot be inferred beyond allowed threshold:$p\leq p_0^K$

c. Privacy by Design [5-7] (Art. 25) - K becomes a design parameter: Systems must be engineered to satisfy impedance constraints.

d. Risk-Based Approach (Recital 75) - Risk is quantified as:"Risk$\propto$p$\Rightarrow$Controlled via K

11.5. Connection to the EU AI Act

The EU AI Act emphasizes risk tiers (LLM / AI systems), prohibited uses ,high-risk system obligations The framework provides a **missing quantitative layer** [14]:

Mapping to AI Act

AI Act Concept	DPP Interpretation
Risk level	p or Z
Harm likelihood	inference probability
Mitigation	increase K
Compliance	satisfy $p \leq p_0^K$

Table 5: Mapping to AI Act

11.6. Example: High-risk AI system [14,17]: Must ensure K≥2
This translates to reduced inference capability, measurable compliance

11.7. Conceptual Shift (Key Contribution)

Traditional regulations: Control data access, The framework: Control knowledge flow

11.8. Closing Statement

The parameter transforms privacy from a qualitative legal principle into a quantitative, enforceable constraint on inference, enabling regulators to directly limit the rate at which systems convert data into knowledge. This approach aligns privacy regulation with physical analogies of impedance, allowing governance of information systems through measurable resistance to inference rather than indirect control over data handling practices.

12. Limitations of Existing Privacy Protection Methods and the DPP Framework Top of Form

From the reviewed the conceptual weaknesses underlying the dominant paradigms of personal privacy [8]. We can conclude one central limitation is the inability to effectively protect personal privacy understood as limited access to the self, as defined by Ruth Gavison [4]. In the age of Information and Communication Technologies (ICTs), this form of privacy is continuously undermined: individuals are subject to surveillance via IoT devices, persistent profiling, identification in public spaces, and multi-channel targeting.

The three core components of limited access—secrecy, anonymity, and solitude—are increasingly difficult to maintain, despite ongoing technological efforts. This appendix reviews the inherent limitations of widely used privacy protection methods.

12.1. K-Anonymity

K-anonymity is a formal model designed to ensure that any individual in a dataset cannot be distinguished from at least k−1 others [12]. However, this model relies on the assumption that datasets remain isolated. In practice, this assumption fails: data fusion across multiple sources enables re-identification. Moreover, attackers often possess background knowledge, allowing them to bypass quasi-identifier masking and reconstruct identities [8].

12.2. L-Diversity

L-diversity extends k-anonymity by requiring diversity in sensitive attributes within each equivalence class [13]. However, it fails when the underlying data lacks sufficient diversity or when enforcing diversity significantly distorts the dataset. In such cases, either privacy is compromised or data utility is degraded [8].

12.3. T-Closeness

T-closeness improves upon l-diversity by requiring that the distribution of sensitive attributes within each group closely matches the overall distribution [14]. However, this approach introduces semantic distortion and reduces analytical value. Furthermore, it assumes limited background knowledge, which is unrealistic in modern data ecosystems, especially when combined with AI-based inference methods [8].

12.4. Differential Privacy

Differential Privacy provides strong formal guarantees by ensuring that the inclusion or exclusion of any individual does not significantly affect analysis results [15,16]. However, it is insufficient for protecting deep personal privacy, as many privacy violations rely on indirect inference. Additionally, it is not applicable in real-time interactions requiring specific user information and thus cannot prevent profiling by large-scale digital platforms [8].

12.5. Deep Personal Privacy (DPP)

We argue that the Deep Personal Privacy (DPP) framework provides a complementary and additional necessary layer of protection. DPP defines privacy as impedance within an inference network, where semantic representations enable latent inference of sensitive attributes. Within this paradigm, privacy risk is determined not by explicit disclosure, but by the ease with which sensitive information can be inferred from embeddings.
Thus, DPP shifts the focus from data protection to inference resistance, addressing the fundamental limitations of existing privacy-preserving methods.

13. Conclusion

This study advances the understanding of privacy in the age of Information and Communication Technologies (ICT) by demonstrating that privacy risk is fundamentally an inference problem rather than a disclosure problem [9,11]. Building on the Deep Personal Privacy (DPP) framework, we have shown that privacy can be rigorously modeled as impedance within an inference network, where semantic similarity in embedding space determines the ease with which sensitive information can be extracted

The originality of this work lies in bridging theory and practice: we

move from a formal definition of privacy as epistemic impedance to an empirical framework capable of quantifying and controlling inference risk in real-world AI systems. By introducing measurable relationships between cosine similarity, inference probability, and impedance, and by validating these relationships through numerical and empirical examples, this paper establishes a new foundation for privacy engineering. Furthermore, the introduction of the parameter K as a regulatory control variable provides a novel mechanism for aligning technical design with legal and ethical frameworks such as GDPR and the EU AI Act.

In contrast to traditional methods approaches that focus on limiting data exposure, the DPP framework directly constrains knowledge extraction [5-7,9,11]. This shift adds a critical new layer to privacy protection, addressing the growing challenge of inference in AI-driven environments. As such, the proposed framework not only enhances the theoretical landscape of privacy research but also provides actionable tools for building and regulating privacy-preserving systems in the ICT era.

References

1. Schoeman, F. (1984). Privacy: philosophical dimensions. *American Philosophical Quarterly, 21*(3), 199-213.
2. Laurie, G. (2002). *Genetic privacy: a challenge to medico-legal norms*. Cambridge University Press.
3. Hongladarom, S. (2015). *A Buddhist theory of privacy. In A buddhist theory of privacy* (pp. 57-84). Singapore: Springer Singapore.
4. Gavison, R. (1980). Privacy and the Limits of Law. *The Yale law journal, 89*(3), 421-471.
5. Solove, D. J. (2010). *Understanding privacy.* Harvard university press.
6. Nissenbaum, H. (2009). Privacy in context: Technology, policy, and the integrity of social life. *In Privacy in context.* Stanford University Press.
7. Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal, 27*(3), 379-423.
8. Oppenheim, Y. (2024). *Personal Privacy in the Age of the Internet: The Influence of Information and Communication Technologies on Personal Privacy.* BookRix.
9. Oppenheim, Y. (2026). Beyond Disclosure: Reframing Privacy as Inference Impedance in Large Language Models.
10. https://www.reddit.com/r/breastcancer/comments/1dubiio/a_positive_post_cancer_post/?utm_source=chatgpt.com
11. Oppenheim, Y. (2026). Privacy as Epistemic Impedance: Deep Personal Privacy and the Political Economy of Knowledge in Networked Societies. *Theory and Practice in Social Studies,* 19-38.
12. Sweeney, L. (2002). k-anonymity: A model for protecting privacy. *International journal of uncertainty, fuzziness and knowledge-based systems, 10*(05), 557-570.
13. Machanavajjhala, A., Kifer, D., Gehrke, J., & Venkitasubramaniam, M. (2007). l-diversity: Privacy beyond k-anonymity. *Acm transactions on knowledge discovery from data (tkdd), 1*(1), 3-es.
14. Li, N., Li, T., & Venkatasubramanian, S. (2007). "t-Closeness", IEEE ICDE 2007
15. Nissim, K (2020). "Differential Privacy: Why, How and Where to?" (2020)
16. https://www.statice.ai/post/what-is-differential-privacy-definition mechanisms examples
17. Kamath, U., Keenan, K., Somers, G., & Sorenson, S. (2024). Large language models: A deep dive. *Bridging Theory and Practice, Cham: Springer Nature.*

Copyright: *©2026 Yair Oppenheim. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.*