A Handshake Protocol for Cross-Model Embedding Interoperability *Patent Pending*

Vicktor Moberg

September 2025

We present a cryptographic handshake protocol to enable cross-model embedding interoperability. The protocol supports secure negotiation of schema metadata and a translation transform T_{AB} between heterogeneous embedding spaces. We formalize the mathematical setting, derive the closed-form solution for orthogonal Procrustes used to compute T_{AB}, and outline security properties (integrity, freshness, impersonation resistance) along with privacy-preserving variants using differential privacy and zero-knowledge proofs.

Novelty: The contribution is not the linear algebra itself but the *stateful, cryptographically attested protocol* that binds anchor selection, dimensional reconciliation, and Procrustes alignment to nonces, commitments, and signed transcripts with replay resistance and revocation.

Introduction

Embeddings serve as the working currency of retrieval-augmented generation (RAG), multi-agent systems, and search. Incompatibilities across embedding spaces currently force lock-in. We propose a handshake protocol that negotiates a mapping between spaces with cryptographic guarantees. Prior art demonstrates mathematical alignment; what is unique here is the combination of (i) anchor-policy negotiation under privacy constraints, (ii) transcript attestation and replay protection, and (iii) verifiable transform provenance tied to session transcripts.

Mathematical Formalization

Let $E_A \subset R^{d_A}$ and $E_B \subset R^{d_B}$ denote the embedding spaces for models A and B. Given an anchor set of m concepts $C = \{c_1, \ldots, c_m\}$, each model produces

$$\begin{aligned} X_A &\in R^{m \times d_A}, \text{rows } x_i^A = embed_A(c_i), \\ X_B &\in R^{m \times d_B}, \text{rows } x_i^B = embed_B(c_i). \end{aligned}$$

Goal: learn a mapping $T_{AB}: R^{d_A} \to R^{d_B}$ that preserves semantics. We focus on the orthogonal case where $T_{AB} \in O(d)$ with $d = min(d_A, d_B)$ after dimensional reconciliation.

Dimensional Reconciliation

If $d_A \neq d_B$, apply projection operators P_A and P_B (e.g., PCA top-d components or orthonormal linear layers) to obtain

$$Z_A = X_A P_A \in R^{m \times d}, Z_B = X_B P_B \in R^{m \times d}.$$

Protocol novelty: P_A and P_B are negotiated and attested inside the transcript to prevent silent dimension mismatches.

Orthogonal Procrustes: Closed-Form Solution

We seek $R \in O(d)$ minimizing $\| Z_A R - Z_B \|_F^2$. Expanding,

$$\| Z_A R - Z_B \|_F^2 \quad ¿\, Tr\left(\left(Z_A R - Z_B\right)^\top \left(Z_A R - Z_B\right)\right)$$

¿ ¿

Since $R^\top R = I$ and $Z_A^\top Z_A$ is constant, the minimization reduces to maximizing $Tr\left(R^\top M\right)$ where $M = Z_A^\top Z_B$. Let $SVD(M) = U \Sigma V^\top$. The optimum is $R^\star = U V^\top$.

Proof sketch.

By von Neumann's trace inequality, for any orthogonal R, $Tr\left(R^\top M\right) \leq \sum_i \sigma_i(M)$. Equality holds when R aligns singular vectors, i.e., $R = U V^\top$. Thus $R^\star$ minimizes the Frobenius error.

Protocol novelty.

The protocol *binds* $R^\star$ to an authenticated session via nonces and commitments; $R^\star$ is accepted only with a valid transcript.

Properties of $R^\star$

- **Stability:** small perturbations in M yield small changes in $R^\star$ when Σ has a spectral gap.
- **Isometry:** if Z_A and Z_B are isometric up to rotation, $R^\star$ recovers that rotation.
- **Complexity:** computing M and its SVD is $O\left(m d^2 + d^3\right)$ for m anchors and d dimensions.

Patent-relevant distinction: these properties are *bound* to transcript-level guarantees (freshness, replay resistance, revocation).

Security Model

Adversary controls the network but not private keys. Goals: (i) integrity of metadata; (ii) freshness (no replay); (iii) peer authentication; (iv) confidentiality of anchors and transforms when required.

6.1 Integrity & Freshness

Messages carry $HMAC = H(K, nonce \| header \| body)$. Nonce uniqueness per session prevents replay. **Novelty:** R^* is accepted only if bound to these commitments, ensuring provenance.

6.2 Authentication Without PKI

We support decentralized identifiers (DIDs) and a web-of-trust. Public keys are tied to DID documents; endorsements provide transitive trust. **Novelty:** peers not only align mathematically, but also cryptographically authenticate identity within the same transcript.

6.3 Privacy

Mitigations include dimension/tokenizer hints as ranges, DP noise on anchors, and ZK proofs for bounding norms without revealing exact values. **Novelty:** these privacy measures are enforced inside the handshake itself, not bolted on afterward.

Protocol Specification

7.1 Sequence Diagram (ASCII)

```
A -> B : ClientHello(nonce_A, protocol, dim_hint, tokenizer_hint)
B -> A : ServerHello(nonce_B, schema, anchor_policy, sig_B)
A <-> B : AnchorSetNegotiation(IDs | synthetic_seed)
A -> B : HMAC_A(metadata, anchors, nonce_A, nonce_B)
B -> A : HMAC_B(metadata, anchors, nonce_A, nonce_B)
A <-> B : Compute T_AB (Procrustes via SVD)
A <-> B : Optional: ZK proofs / DP stats

Novelty: each step contributes to a transcript; mismatches trigger
abort.
```

7.2 Message Schema (JSON Fragment)


```
{
  "protocol": "LLMHS-v0.3",
  "nonce": "...",
  "dim_hint": "512..1024",
  "tokenizer_hint": "bpe-range",
  "anchor_policy": "synthetic-seed:v1",
  "sig": "base64(ed25519)"
```

```
}

Novelty: schema values are not mere hints but cryptographically bound
commitments.
```

Pseudocode

```
function handshake(A, B):
  nonce_A <- rand()
  send A->B: ClientHello(nonce_A, hints)
  recv B->A: ServerHello(nonce_B, schema, anchor_policy, sig_B)
  anchors <- negotiate_anchors(anchor_policy)
  X_A <- embed_A(anchors); X_B <- embed_B(anchors)
  (P_A, P_B) <- reconcile_dims(X_A, X_B)
  Z_A <- X_A * P_A; Z_B <- X_B * P_B
  M <- Z_A^T * Z_B
  (U, Sigma, V) <- svd(M)
  R <- U * V^T
  send mutual HMACs over {hints, schema, anchors, R}
  return R
```

Efficiency & Caching

Handshake cost is dominated by SVD($d \times d$). Reuse R within session keys; cache by model-version and anchor policy. For streaming workloads, amortize anchor evaluation across batches. **Novelty:** cached transforms are usable only when the transcript ID, anchor policy hash, and model-version match.

Simulated Evaluation

Setup.

Anchors: 8000, policy = `real-phrases:v1`. Models: `text-embedding-3-large` ↔ `text-embedding-3-small`. Sweep across $d \in \{24, 32, 40\}$. Best performance: $d = 32$.

Cosine Similarity (Test Queries).

Before: -0.059. After: 0.670.

Table 1: Dimensional Sweep Results

d	Cosine Before	Cosine After	Time (s)	Notes
24	0.0220	0.6541	26.11	stable, decent
32	-0.0085	0.6597	25.18	**best overall**

d	Cosine Before	Cosine After	Time (s)	Notes
40	-0.0262	0.6310	25.26	degraded

Retrieval Examples.

- **“attention heads and residual streams”**
  Before: unrelated (Gregorian chant, whiskey notes, Ignatian examen).
  After: relevant (Azure pipelines, Transformers, neural nets).

- **“x64 windows tls client in assembly”**
  Before: scattered technical/math terms.
  After: precise matches (assembly sockets with Schannel, orthogonal Procrustes alignment).

- **“ignatian examen nightly reflection”**
  Before: poor (whiskey notes, generic math topics).
  After: close matches including Ignatian spirituality and examen.

- **“rtl-sdr telemetry capture”**
  Before: irrelevant (whiskey notes, generic ML).
  After: high-similarity matches with satellite telemetry and RTL-SDR topics.

Discussion & Limitations

Linear mappings may underfit when spaces differ nonlinearly; anchor bias can induce domain shift. DP noise may degrade alignment accuracy; ZK protocols add overhead. Future work: neural adapters regularized toward orthogonality; federated multi-party handshakes.

Conclusion

We gave a protocol, a closed-form alignment derivation, and a security analysis that together make cross-model embedding exchange practical. **Novelty (summary):** a cryptographic handshake that binds anchor policy, dimensional reconciliation, and Procrustes alignment to an attested, replay-resistant transcript with revocation and caching rules.