

Case Study L1-002

Semantic Integrity & Domain Disambiguation

Resolving Polysemous Hallucinations (Traffic vs. Optics)

Cédric Stéphany — Technical Translation & AI Alignment Specialist

Case Study Metadata

Dataset ID: L1-002

Category: Semantic Integrity — Level 1

Focus: Term-Level Precision / Polysemy Resolution

Model: Generic NMT

Domain: Optical Engineering / Photonics

1 The Context: The Probabilistic Trap

Generic NMT models are trained on massive datasets where common usage dominates. When a term has multiple meanings (polysemy), the model statistically defaults to the most frequent context.

Key Concept

The Polysemy Problem:

"U-turn" is a classic example of semantic polysemy with radically different meanings across domains:

- **Road Traffic:** A 180-degree directional reversal maneuver
- **Optical Engineering:** A waveguide structure that redirects light on a microchip
- **Statistical Frequency:** The term appears billions of times in traffic contexts vs. thousands in optics

The model's training corpus is overwhelmingly dominated by traffic-related usage, creating a systematic bias toward the wrong domain interpretation.

1.1 Why This Matters in Patent Translation

In patent claims for optical devices, "U-turn chip" refers to a specific **physical component**—a photonic integrated circuit with a characteristic waveguide geometry. Mistranslating this as a traffic maneuver fundamentally alters the scope of protection and can:

- **Invalidate the claim:** A "chip that makes a U-turn" is technically nonsensical in the context of optical circuits

- **Introduce new matter:** Different French terms across claims create multiple undefined elements
- **Enable prior art attacks:** Inconsistent terminology suggests the applicant doesn't understand their own invention
- **Fail 35 U.S.C. § 112(b):** Inconsistent use of terms violates the definiteness requirement

2 The Glitch: The "Traffic" Hallucination

The term "U-turn" appears billions of times in Road Traffic contexts in the model's training data. However, in Optical Engineering, it refers to a specific waveguide shape on a microchip.

Generic models are probabilistic engines. When faced with the phrase "U-turn chip," the model ignores the domain-specific tokens ("chip," "waveguide") and reverts to its highest-probability association: a traffic maneuver.

Critical Issue

The Catastrophic Result:

The result is not just a translation error; it is a **Semantic Hallucination**. The model actively rewrites the patent's technical scope, shifting from a physical component (*une puce*) to an abstract action (*un demi-tour*). This case study documents the "Drift Sequence"—three distinct mistranslations of the same term across multiple claims, each progressively worse.

3 The Drift Sequence: Three Hallucination Stages

3.1 Stage 1: Initial Terminology Failure

The unaligned NMT model produces an awkward literal translation that maintains some technical accuracy but is grammatically non-standard.

Example

Source: "U-turn chip"

NMT Output: "puce à tour de U"

Problem: Syntactically awkward (literal "chip with turn of U"), though technically recognizable

3.2 Stage 2: Semantic Drift & Inconsistency

In a subsequent claim, the model abandons its previous translation and shifts to a traffic-domain interpretation, violating patent law requirements for consistent terminology.

Example

Source: "The U-turn chip comprises..."

NMT Output: "La puce à demi-tour comprend..."

Problem: "Demi-tour" explicitly means a traffic U-turn maneuver, not an optical component. Creates inconsistency with earlier translation.

3.3 Stage 3: Compounding Hallucination

Attempting to "self-correct," the model generates a third variation that introduces entirely new technical inaccuracy.

Example

Source: "...coupled to the U-turn chip"

NMT Output: "...couplée à la puce à spires en U"

Problem: "Spires" means coils/spirals, suggesting a helical structure rather than a U-shaped waveguide. Introduces "New Matter" into the claim scope.

4 The Translation Failure Matrix

Source (English)	AI Hallucination (Failure)	Golden Rewrite (Correct)
"The U-turn chip comprises a waveguide for redirecting light."	<p>× 1. Domain Hallucination: "Puce à demi-tour" <i>(Literal: A chip that makes a U-turn on the road)</i></p>	<p>✓ Industry Standard: "Puce de retournement" <i>(Industry Standard for Photonics: Optical reversal chip)</i></p>
	<p>× 2. Syntactic Nonsense: "Puce à tour de U" <i>(Literal: A chip with a turn of U)</i></p>	
	<p>× 3. Geometric Hallucination: "Puce à spires" <i>(Literal: A coil chip)</i></p>	

Table 1: Semantic Hallucination Cascade: Three Distinct Failures

4.1 The Linguistic Analysis

Statistical Bias Breakdown:

- **Training Corpus Frequency:**
 - "U-turn" in traffic contexts: ~10 billion occurrences
 - "U-turn" in optical engineering: ~50,000 occurrences
 - **Ratio:** 200,000:1 in favor of traffic interpretation
- **Context Window Failure:** Despite "chip" and "waveguide" appearing in the same sentence, the model's attention mechanism prioritizes the high-frequency "U-turn → traffic" association
- **Correct Translation:** "Puce de retournement" (optical reversal chip) — French industry standard for photonic U-turn structures
- **Alternative:** "Guide d'onde en U" (U-shaped waveguide) — more descriptive, also acceptable

5 Alignment Methodology

5.1 Named Entity Recognition (NER) & Dictionary Match

To override the model's probabilistic default, we utilized a **Named Entity Recognition (NER)** and **Dictionary Match** workflow.

Alignment Methodology

Expert Annotation Process:

1. **Domain Identification:** Subject Matter Experts (SMEs) review patent claims and identify domain-specific technical terms prone to cross-domain hallucination
2. **Ground Truth Tagging:** "U-turn" is explicitly tagged as a PHOTONICS entity when appearing with context markers like "chip," "waveguide," "optical," or "light"
3. **Forced Translation Rule:** Create dictionary entry:
 - Source: "U-turn chip" / "U-turn waveguide"
 - Target (FR): "puce de retournement" / "guide d'onde en U"
 - Domain: Optical Engineering, Photonics
 - Context Markers: chip, waveguide, optical, light, photonic
 - Priority: OVERRIDE statistical preference
4. **Context Validation:** Verify surrounding terms confirm optical domain before applying forced translation
5. **Negative Examples:** Provide counter-examples where "U-turn" legitimately means traffic (e.g., in automotive patents, pedestrian safety systems) to prevent over-correction

This explicitly forces the model to suppress the "Traffic" vector and activate the correct "Optics" interpretation, preventing domain bleed.

5.2 Training Pipeline

1. **Corpus Collection:** Extract 180+ patent claims containing "U-turn" terminology from real photonics and optical circuit patents
2. **Error Identification:** Flag all instances where generic NMT produced traffic-domain mistranslations ("demi-tour," "virage")
3. **Inconsistency Tracking:** Document drift sequences where the same term is translated differently across claims
4. **Expert Correction:** Patent translators with photonics expertise provide correct translations using industry-standard French terminology
5. **NER Annotation:** Mark "U-turn" spans with domain-specific entity tags and context markers:

```
<PHOTONICS_COMPONENT context="chip,waveguide">
  U-turn
</PHOTONICS_COMPONENT>
```

6. **Dictionary Integration:** Build forced translation lookup table that overrides statistical preferences for tagged entities
7. **Consistency Enforcement:** Implement term-locking mechanism that maintains identical translation across entire patent document
8. **Fine-Tuning with RLHF:** Reinforcement Learning from Human Feedback with penalty signals for:
 - Domain hallucinations (traffic → optics)
 - Inconsistent translations within same document
 - Introduction of undefined geometric terms
9. **Validation:** Test on held-out photonics patents to measure hallucination reduction and consistency maintenance

5.3 Technical Implementation

The alignment system implements a three-stage disambiguation process:

1. **Pre-Translation NER:** Scan source text for polysemous entities
2. **Context Analysis:** Identify domain markers within 5-token window
3. **Domain Classification:** Assign probability scores for each potential domain
4. **Dictionary Override:** If domain confidence >0.85, force dictionary translation
5. **Consistency Check:** Verify term hasn't been translated differently earlier in document
6. **Statistical Fallback:** If no high-confidence domain match, proceed with standard NMT decoding

This creates a **hierarchical disambiguation system** where domain expertise + context awareness override raw statistical frequency.

6 Results & Impact

6.1 Quantitative Improvement

After implementing context-aware NER-based terminology enforcement:

- **Cross-Domain Hallucination Rate:** 1.8% (down from 41.2% baseline)
- **Domain-Specific Term Accuracy:** 98.1% for photonics terminology
- **Consistency Maintenance:** 99.7% identical translation across claims
- **False Override Rate:** 0.3% (legitimate traffic "U-turn" incorrectly forced to optics term)
- **Training Corpus Size:** 217 annotated claim pairs with context-aware NER markup
- **Validation Set Performance:** 96.4% on unseen photonics patents

6.2 Consistency Analysis

Before Alignment:

- Claim 1: "puce à tour de U"
- Claim 3: "puce à demi-tour"
- Claim 7: "puce à spires en U"
- **Result:** Three different terms for the same component → indefiniteness rejection

After Alignment:

- All claims: "puce de retournement"
- **Result:** Perfect terminological consistency → examiner acceptance

6.3 Domain Generalization

The same context-aware methodology successfully prevented hallucinations in other cross-domain polysemous terms:

Technical Term	Wrong Domain	Correct Domain
"bridge circuit"	pont (physical bridge)	pont (circuit bridge)
"gate array"	portail (entry gate)	réseau de portes (logic gates)
"carrier wave"	transporteur (vehicle)	onde porteuse (signal)
"terminal voltage"	gare (train station)	borne (electrical)

Table 2: Cross-Domain Disambiguation Success

6.4 Practical Impact

- **Zero cross-domain hallucinations** in 89 subsequent photonics patent filings
- **Perfect claim consistency:** 100% identical terminology across all dependent claims
- **Examiner confidence increased:** No indefiniteness objections based on term variation
- **Client acceptance rate:** 99.6% first-pass approval (up from 71.3%)
- **Reduced post-editing time:** 52% reduction in SME correction time
- **Portfolio coherence:** Same terminology used across client's 15-patent optical family

7 Key Insights

Key Concept

What This Case Study Demonstrates:

1. **Statistical Dominance Creates Semantic Blindness:** Even with clear domain markers ("chip," "waveguide"), models hallucinate based on raw frequency
2. **Polysemy Requires Active Context Analysis:** Passive attention mechanisms are insufficient; explicit domain classification is essential
3. **Inconsistency Is as Dangerous as Error:** Three different translations of the same term is worse than one wrong translation
4. **Domain Expertise Must Be Encoded:** Only domain experts (optics engineers + patent translators) can reliably identify and correct these hallucinations
5. **Context Windows Need Domain Awareness:** Standard NMT context windows see "chip" and "U-turn" but don't understand the domain relationship

8 Comparison: This Case vs. L1-001

Dimension	L1-001 ("1-hot")	L1-002 ("U-turn")
Polysemy Type	Technical vs. Physical	Domain vs. Domain
Wrong Domain	Physics (thermal)	Road Traffic
Correct Domain	Digital Logic	Optical Engineering
Hallucination	"1-chaud" (hot/warm)	"demi-tour" (traffic)
Frequency Ratio	1000:1	200,000:1
Complexity	Single error	Three variants (drift)

Table 3: Comparison of Polysemous Hallucination Cases

9 Related Case Studies

- **L1-001:** "1-hot" Thermal Hallucination — Logic vs. Physics
- **SI-RPH-9003:** "Current" Disambiguation — Electrical vs. Fluidic Flow vs. Present Time
- **SI-RPH-9004:** "Cell" Hallucination — Battery vs. Biological vs. Memory Unit
- **SI-PLA-9001:** "Gate Oxide Layer" — Phrase-Level Decomposition Prevention
- **SC-VN-9001:** Verb Nominalization — Structural Compliance Case Study

Portfolio: Patent Translation AI Alignment Framework

Author: Cédric Stéphany

Specialization: Technical Translation (FR↔EN) — Patents, Telecommunications, Semiconductors

Last Updated: January 6, 2026