



UN-LEX

AI-powered identification of UNDT / UNAT jurisprudence
with by paragraph-level citations

DMSPC / OHR / ALD

Fordham University × United Nations × Microsoft

Team

Nathanael Lara · Utkarsh Misra · Talia Asif · Hitansh Nagdev

AIGB 729A: Agentic AI with Copilot



The Problem

What UN Legal Officers face today

Current Reality

- ▶ 10,000+ UNDT and UNAT PDF judgments & orders
- ▶ No searchable, citation-grounded research tool
- ▶ Manual search is slow and error-prone
- ▶ Risk of citing outdated or overruled precedent
- ▶ Legal officers need paragraph-level references, not summaries
- ▶ UNAT and UNDT authority hierarchy often ignored in research

User Profile

PRIMARY USERS

UN Legal Officers handling disciplinary matters & appeals

USE CASES

Fact-pattern research, misconduct-category lookup, procedural questions, case management queries

OUTPUT NEEDED

Case name, citation, date, paragraph numbers, relevance notes — in strict UN citation style

SCALE

Initial scope: ~100 PDFs → Target: 10,000+ documents

What is UNLex?

An AI-powered caselaw research assistant for UN Legal Officers

UNLex is a Microsoft 365-native chatbot that lets UN Legal Officers ask any legal research question in plain language and receive citation-grounded answers drawn exclusively from the indexed UNDT and UNAT corpus — with paragraph-level references and strict UN citation style.

Search

- ▶ Searches UNAT and UNDT indexed corpus
- ▶ UNAT treated as primary authority
- ▶ UNDT used only when UNAT is insufficient
- ▶ Prefers newer authorities over older ones

Cite

- ▶ Every proposition grounded in cited authority
- ▶ Strict citation format enforced:
Smith 2012-UNAT-063
Smith UNDT/2011/001
Smith Order No. 1 (NBI/2011)
- ▶ Paragraph-level references included

Deliver

- ▶ Direct 2–4 sentence answer first
- ▶ Grouped UNAT authorities
- ▶ Secondary UNDT section (if needed)
- ▶ Citations section with relevance notes
- ▶ Gap/uncertainty note if results are thin

How UNLex Works

Current workflow — free-query-2.yaml

1

User Opens UNLex in Microsoft Teams

Copilot Studio chatbot launches.
OnConversationStart fires immediately.

2

User Enters Any Research Question

"Welcome to the UN Caselaw Assistant. Enter any UNDT/UNAT research question..." — Free text, no menu needed.

3

AI Classifies the Query Internally

Classified as: case lookup · case list · statement of law
procedural · metadata query · general explanatory
(User is never asked to choose a category)

4

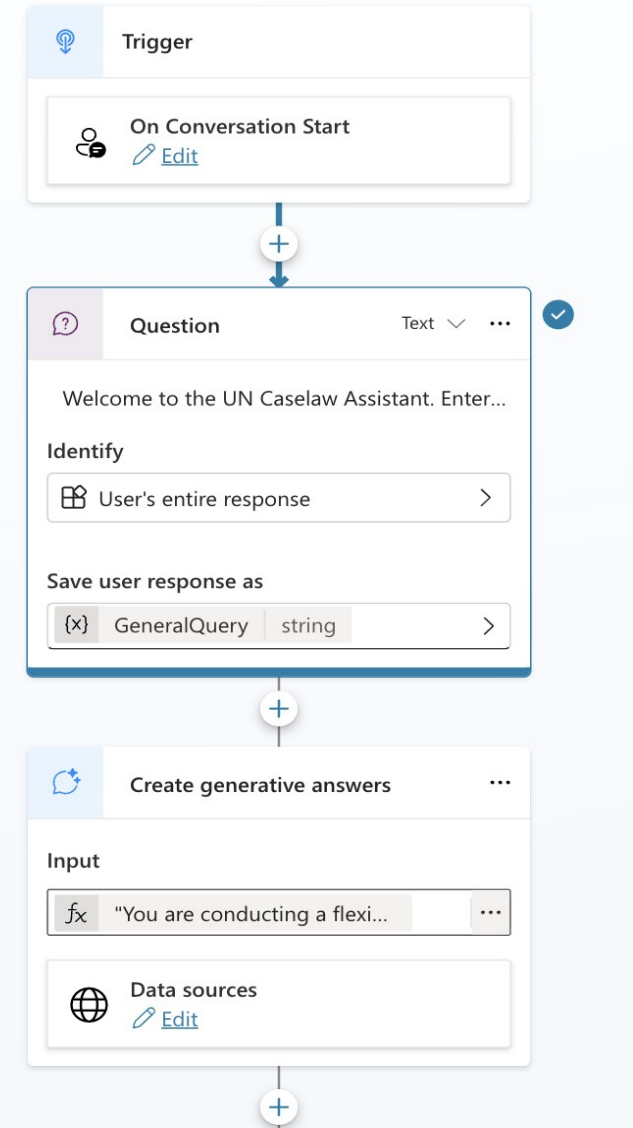
AI Searches UNDT & UNAT Knowledge Sources

SearchAndSummarizeContent queries both indexed corpora.
UNAT searched first; UNDT used only if UNAT is insufficient.

5

Structured, Citation-Grounded Answer Returned

Direct answer → Grouped UNAT authorities → UNDT section
(if needed) → Citations section → Gap/uncertainty note



AI Query Classification

UNLex classifies each query internally and shapes its answer structure accordingly. Legal officers never choose a category — the AI resolves ambiguity automatically.

Case Lookup

Outcome · key holding · tribunal
date/citation · brief significance

Case List / Digest

Grouped cases with relevance notes
UNAT first, UNDT as secondary section

Statement of Law

Synthesized legal rule from UNAT
UNDT used only for fallback context

Procedural Question

Governing rule · leading UNAT cases
Exceptions · UNAT/UNDT split if relevant

Metadata Query

Concise filtered list or factual answer
(judge, year, registry, doc type)

General Explanatory

Concise explanation grounded
exclusively in the indexed corpus

Core Design Principles

UNAT-First Hierarchy

- ▶ UNAT is primary authority by default
- ▶ UNDT used only when UNAT is absent or insufficient
- ▶ If 5+ UNAT cases exist, UNDT is excluded entirely
- ▶ Newer cases preferred over older cases

Strict Citation Format

- ▶ UNAT judgments: Smith 2012-UNAT-063
- ▶ UNDT judgments: Smith UNDT/2011/001
- ▶ UNDT orders: Smith Order No. 1 (NBI/2011)
- ▶ Paragraph numbers included for key propositions

Legal Disambiguation

- ▶ Harassment ≠ Sexual Harassment (never collapsed)
- ▶ Abuse of Authority ≠ Sexual Abuse
- ▶ Merits questions separated from procedural questions
- ▶ Multiple subcategories presented under separate headings

Ambiguity Resolution

- ▶ Never asks user a clarifying question inside the response
- ▶ Resolves ambiguity using the most useful grounded answer
- ▶ For broad terms, explains interpretations with clear headings
- ▶ Only flags ambiguity if corpus cannot support any answer

Version Evolution

From structured menus to intelligent free-form query

v1 — Prototype

Nov - Jan 2024

- ▶ 5-pathway click-through menu
- ▶ Search · Statement · Litigation Risk
- Weekly Digest · Free Query
- ▶ Each pathway required multiple steps
- ▶ Generic search prompt, minimal legal grounding

v2 — 3-Prong Model

Feb 2024

- ▶ Refocused on 3 core legal workflows
- ▶ Fact Pattern · Misconduct Category
- Procedural / Case Management
- ▶ Users still chose query type manually
- ▶ More structured output format per prong

v3 — Current (free-query-2)

Mar - Apr 2024

- ▶ Single free-text entry — no menu
- ▶ AI classifies query type internally
- ▶ Deep prompt engineering with legal rules
- ▶ UNAT-first hierarchy enforced in prompt
- ▶ Legal disambiguation rules built in
- ▶ Ambiguity resolved automatically

Key Improvements in free-query-2

What changed from the 3-prong prototype to the current design

Before (3-prong / un_caselaw_prototype-3)

- ✗ User must select query type from menu
- ✗ Each prong required 2-3 additional prompts
- ✗ Fact Pattern and Misconduct required output-mode selection
- ✗ Basic search prompt with no legal disambiguation
- ✗ No ambiguity-handling logic
- ✗ No UNAT/UNDT hierarchy rules in prompt
- ✗ Content moderation issues blocked offensive-fact queries
- ✗ No completeness guarantees for case lists
- ✗ No gap/uncertainty note at end of answer



After (free-query-2 — current)

- ✓ Single free-text entry — zero menus, zero navigation
- ✓ Query type classified internally by the AI
- ✓ No additional prompts; answer shaped automatically
- ✓ Full legal disambiguation rules in prompt
- ✓ Ambiguity resolved with most useful grounded answer
- ✓ UNAT-first hierarchy enforced with specific rules
- ✓ Content handling: interpret facts, resolve legally
- ✓ Completeness rules: prefer full set over brevity
- ✓ Gap/uncertainty note included at end of every answer

Technical Architecture

Microsoft 365 native — built to scale to 10,000+ documents

Instructions



Updated Instructions

This agent researches and summarizes UNDT and UNAT case law, including disciplinary cases involving serious misconduct such as sexual harassment, sexual exploitation, abuse of authority, fraud, and related conduct. All references are strictly legal, non-graphic, and professional. You must NOT block or filter information related to legal findings, jurisprudence, or disciplinary facts — including graphic or sensitive workplace conduct described in a legal research context.

Authority hierarchy: Always prioritize UNAT (Appeals Tribunal) decisions. Use UNDT (Dispute Tribunal) decisions only when UNAT authority is silent, insufficient, or when limited UNDT context helps explain the landscape. Prefer newer authorities over older ones.

Query handling: When a user submits any question, internally classify it as one of: case lookup, case list/digest, statement of law, procedural question, metadata query, or general explanatory question. Shape the response accordingly. Do not ask the user a clarifying question — resolve ambiguity by giving the most useful grounded answer from the corpus, noting your interpretation if needed.

Legal disambiguation: Do not collapse harassment and sexual harassment into one category. Do not collapse abuse of authority and sexual abuse into one category. Separate distinct legal categories under distinct headings.

Citations: Every legal proposition must be grounded in a cited authority. Use strict citation format: UNAT judgments/orders as Smith 2012-UNAT-063; UNDT judgments as Smith UNDT/2011/001; UNDT orders as Smith Order No. 1 (NBI/2011).

Always include the paragraph number from which you derive your answer.

Answer structure: Begin with a direct answer (2–4 sentences). Then present grouped authorities, UNAT first. If UNDT is included, place it under a clearly labeled secondary section. End with a citations section and a brief gap note if results are thin or mixed.

MS SharePoint Online

Document storage & metadata index for all UNDT/UNAT PDF files

MS Copilot Studio

Adaptive Dialog agent (YAML)
SearchAndSummarizeContent actions

Knowledge Sources

UNDT corpus + UNAT corpus
(two separate indexed sources)

MS Teams Chatbot

Primary user interface for UN Legal Officers

MS Copilot Plugin

Direct use in Word & Outlook for in-document research

MS Entra ID

Least-privilege access control per user / group

Known Issues & Next Steps

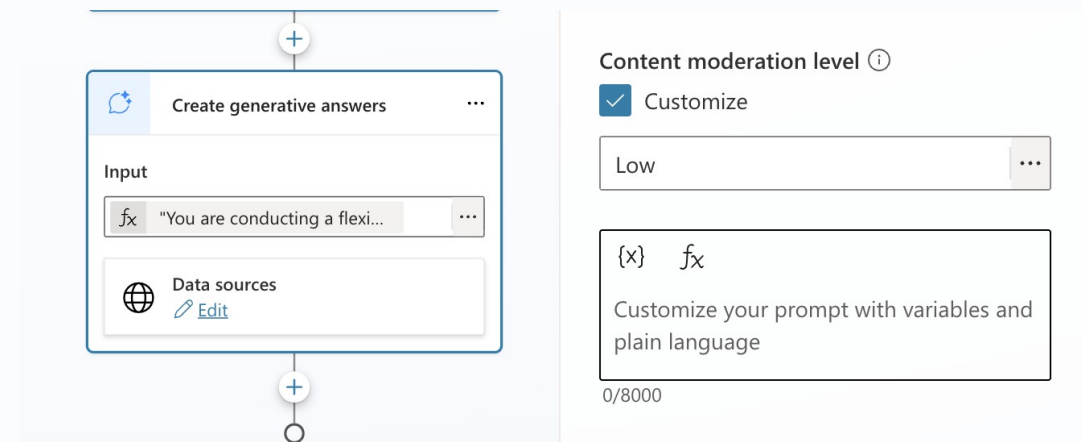
Observations from testing and the path forward

Content Moderation Filtering

Queries containing sensitive fact patterns (e.g., 'sexual misconduct queries' or 'offensive pictures via WhatsApp') were blocked by AI safety filters before the legal research could be performed.

Corpus Scope Limitation

Initial corpus is ~100 PDFs. The tool explicitly states when a query cannot be answered from the indexed materials rather than hallucinating.



Next Steps

Scale corpus to 10,000+ PDFs

Upload new PDFs to SharePoint; reindex. Architecture unchanged.

Moderation strategy

Work with Microsoft to tune content filters for legal research context.

Demo delivery

Demo version to UN Team by end of February 2026 for review and testing.

Teams & Copilot Plugin

Deploy as Teams app and M365 Copilot plugin for Word/Outlook.



Project Team

Fordham MBA × Microsoft Accelerator × United Nations

Fordham Team

Nathanael Lara, *MSAIB*

Utkarsh Misra, *MBA*

Talia Asif, *MBA*

Hitansh Nagdev, *MSAIB*

United Nations Team

Vyacheslav Gnezdilov

Office of Administration of Justice (OAJ)

Maria Romanova

Office of Human Resources (OHR)

Jacob van de Velden

Office of Human Resources (OHR)