International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

AI-Powered Automation of NEPA Categorical Exclusions for Historic Districts and Federal Landmarks: An Integrated Multi-Agency Framework for Intelligent Permitting

Mohit Upadhyay

Abstract

NEPA categorical exclusion compliance for infrastructure projects in historic districts requires cross-referencing more than 30 federal agency CE catalogs, initiating NHPA Section 106 consultations with State Historic Preservation Offices, and generating legally defensible documentation, a workflow that currently consumes 15 to 30 days per site and costs telecommunications carriers $3,000 to $8,000 per application, based on the author's primary research projections in telecommunications program management. No integrated AI framework currently addresses this multi-agency determination challenge at scale. This paper proposes a multi-modal platform that resolves CE determination through retrieval-augmented generation (RAG) over a structured regulatory knowledge graph, transformer-based computer vision analysis of geospatial imagery and construction drawings, parallel federal API orchestration across 15 or more databases, and large language model (LLM)-driven generation of agency-specific compliance documentation. A tiered processing architecture targets Tier 1 CE determination in under two hours for standard cases, while a human-in-the-loop override design with immutable audit traceability maintains the legal defensibility standard required for operational deployment. Applied to U.S. 5G small cell programs, where NEPA/Section 106 compliance represents one of the documented bottlenecks on infrastructure deployment velocity, the framework addresses a compliance challenge that has resisted automation to date, with broader applicability to electric utilities, transportation agencies, and federal facility operators.

Keywords: *NEPA compliance, categorical exclusion, AI automation, retrieval-augmented generation, computer vision, multi-agency integration, intelligent permitting, historic preservation, 5G deployment, regulatory knowledge graph*

. Introduction

Enacted in 1969, the National Environmental Policy Act established the procedural foundation for federal environmental review in the United States, a mandate implemented through a framework whose operational complexity has grown substantially with the proliferation of agency-specific CE rules [1]. For routine, low-impact activities, agencies have established categorical exclusions: predefined classes of actions that do not require an Environmental Assessment or Environmental Impact Statement absent extraordinary circumstances [3]. CE determination sounds straightforward. In practice, when proposed projects intersect with historic districts, National Register-listed properties, or federally protected landmarks, it requires simultaneous resolution of overlapping CE frameworks from the FCC, USACE, DOT, DOI, DHS, DOE, and more than 25 additional federal entities, each with distinct condition thresholds, extraordinary circumstance triggers, and documentation requirements [4, 5]. The National Historic Preservation Act Section 106 adds a parallel consultation obligation with State Historic Preservation Offices [2], compounding both the analytical and administrative burden.

Tech Mahindra (Americas) Inc., USA

ORCID: 0009-0007-8072-0380

Telecommunications carriers deploying 5G small cell infrastructure in urban historic districts encounter this complexity as a direct operational constraint. FCC regulations at 47 CFR Section 1.1306(c) impose specific numeric thresholds, height increases, equipment cabinet additions, protrusion limits, ground disturbance extent, for antenna deployments on or near historic structures to qualify for CE treatment [4]. A single antenna modification on a historic building triggers FCC CE threshold analysis, NPS/NRHP historic property adjacency determination, and SHPO Section 106 consultation, each with distinct deliverables. Currently, licensed compliance analysts perform this work manually, database cross-referencing, construction drawing examination, site photograph review, boilerplate document production, at a cost of $3,000 to $8,000 per site and 15 to 30 days per application, based on the author's primary research projections. Against a national 5G densification requirement of approximately 800,000 small cell deployments [8], this per-site overhead represents a systemic constraint on deployment velocity.

Addressing this bottleneck requires integrating capabilities that existing systems have addressed only in isolation: machine understanding of regulatory text [15], RAG for factually grounded knowledge retrieval [6], graph-structured multi-hop regulatory reasoning [18], transformer-based visual analysis of site conditions [16], and LLM-driven generation of legally structured documents [17]. No prior work has assembled these capabilities into a unified architecture for the multi-agency CE determination problem in historic district contexts. This paper presents that architecture, which combines a machine-traversable CE knowledge graph, a federal API orchestration layer, a computer vision site analysis pipeline, and a human-in-the-loop document generation framework, to address the complete CE compliance workflow as an intelligent systems engineering problem.

2. Background and Related Work

2.1 The Compound Burden of NEPA and NHPA in Historic District Projects

Categorical exclusions were designed as an administrative efficiency mechanism, a recognition that not every federal action requires a full environmental review [1, 3]. Each agency's CE catalog reflects its specific program activities: the FCC's 47 CFR Section 1.1306(c) defines wireless-specific exclusions for antenna deployments [4]; the USACE's 33 CFR Section 230.9 covers minor maintenance dredging and routine infrastructure O&M; DOT Order 5610.1C addresses transportation facility modifications. Despite their operational diversity, all agency CE frameworks share one structural feature: extraordinary circumstances function as automatic disqualifiers, negating CE applicability regardless of project type when triggered. Adverse effects on National Register-listed or eligible properties represent the most universally applicable extraordinary circumstance trigger, appearing across every agency's CE framework without exception [3, 5].

Historic district projects activate this trigger precisely where CE determination would otherwise be most efficient. A small cell collocation that qualifies as a routine FCC Section 1.1306(c) CE in a modern commercial district becomes a multi-agency compliance event when the host structure is NRHP-listed. Section 106 of the NHPA obligates federal agencies and private applicants like telecommunications carriers acting under FCC licenses to consult SHPOs before undertaking actions that may affect historic properties [2, 5]. The compliance workflow components: FCC CE applicability analysis, historic property adjacency determination, Section 106 consultation initiation, and, in cases of adverse effect, ACHP coordination. The NCSHPO has documented that 40 to 60 percent of Section 106 submissions contain avoidable errors or omissions [19], creating downstream backlogs that add weeks to project timelines beyond the initial preparation burden.

2.2 AI Approaches to Regulatory Compliance Automation

The legal NLP literature has established the feasibility of machine understanding of regulatory texts. LEGAL-BERT [15] demonstrated that domain-adapted transformer models achieve strong performance on statutory classification and information extraction, capabilities that directly support parsing CE condition language and mapping project characteristics to applicable exclusion categories. Generative approaches alone introduce a critical risk for regulatory applications: LLM hallucination of plausible but factually unsupported regulatory provisions. RAG architectures address this challenge by grounding generation in retrieved authoritative source documents, ensuring that CE determination outputs cite specific regulatory provisions rather than producing synthetic text [6]. Gao et al. provide a comprehensive survey of RAG approaches applicable to knowledge-intensive legal and regulatory domains [11]. Efficient deployment of large-scale language models for such operational contexts is further supported by structural pruning techniques [14] that reduce inference cost without sacrificing domain accuracy.

Graph RAG extends standard retrieval by representing knowledge as a structured graph, enabling multi-hop reasoning across related nodes [18]. CE determination is fundamentally a multi-hop problem: resolving whether a given project qualifies for CE requires traversing from project type to applicable CE category to condition thresholds to extraordinary circumstance triggers to documentation requirements, a chain that spans agency boundaries and involves conditional dependencies. Chain-of-thought prompting [13] enables sequential regulatory reasoning for complex CE conditions. As Saren (2025) explains regarding AI-assisted NEPA reviews, the traceability of data sources and the assurance of human agency in final decisions are essential for the legal defensibility of AI-assisted CE determinations [7]. Foundation model capabilities [17, 20], including few-shot inference and transformer-based visual analysis [12], provide the generative and analytical capacity that unifies these components. Despite the maturity of these enabling technologies, no prior work has integrated them into a unified architecture for determining CE in historic districts.

3. Proposed Architecture

3.1 System Overview and Design Principles

Resolving multi-agency CE determination for historic district projects requires a system that can simultaneously traverse regulatory knowledge, analyze visual site conditions, retrieve authoritative government data, and generate agency-specific documentation. The proposed platform addresses this requirement through a seven-layer architecture (Figure 1) following a RAG + Computer Vision + Regulatory Logic Engine paradigm.

Proposed Platform Architecture — Seven Functional Layers

Layer	Function	Key Technologies
1. Data Ingestion and Normalization	Ingest permit forms, construction drawings, historic property data, satellite and Street View imagery	OCR, PDF parsing, DWG/IFC parser, Google Maps API, NPS IRMA API
2. Geospatial Intelligence	Identify historic district boundaries, landmark corridors, ROW zones, and ground disturbance areas	Google Street View API, PostGIS/ArcGIS, LiDAR, computer vision
3. Regulatory Knowledge Graph	Map project characteristics to CE categories, extraordinary circumstance triggers, and permit requirements	Neo4j, RAG over CE catalog, rule engine
4. AI Analysis Engine	Extract project dimensions, equipment counts, and visual impact indicators from drawings and imagery	Vision transformer, multi-modal LLM, OCR, 3D reconstruction
5. CE Determination Engine	Apply regulatory logic to produce CE determination, flag extraordinary circumstances, recommend EA/EIS	Deterministic rule engine + LLM chain-of-thought with citation enforcement
6. Document Generation	Auto-populate agency-specific forms, co-exclusion letters, SHPO consultation correspondence	Template engine, LLM text generation, e-signature integration
7. Audit and Defensibility	Log every data source, API call, AI inference, and human override for regulatory audit trail	Immutable time-series database, SHA-256 hashing, digital signature

Figure 1: Proposed Platform Architecture - Seven Functional Layers

Four principles govern the design. Human-in-the-Loop Override assures that every AI decision will be reviewed before issuance by a licensed analyst with a justification for the override recorded, and Source Traceability connects every output field back to the data source responsible for it, as required for the NEPA decision administrative record documentation [7]. Extraordinary Circumstance First defaults to flagging ambiguous cases rather than granting automated clearance. Agency-Specific Rendering produces distinct output documents for each participating agency from a shared underlying analysis, eliminating redundant multi-agency data gathering.

3.2 CE Knowledge Graph Design

At the core of the CE determination engine is a directed regulatory knowledge graph encoding the CE catalog of 30 or more federal agencies. Nodes represent CE category definitions, extraordinary circumstance trigger conditions, project characteristic thresholds, and permit form requirements. Graph RAG traversal [18] determines multi-agency CE applicability from a unified project description in a single inference pass, identifying applicable CE categories across agency boundaries simultaneously rather than executing sequential per-agency analysis. Historic district projects route through a specialized subgraph encoding nine CE trigger conditions that appear consistently across agency frameworks as extraordinary circumstance disqualifiers, summarized in Table 1.

Table 1. Historic District CE Trigger Conditions

CE Condition	Applicable Agencies	AI Analysis Required
No adverse effect on NRHP-listed or eligible properties	All agencies (extraordinary circumstance trigger)	Street View analysis + NPS NRHP database + SHPO lookup
Work within existing footprint/ROW only	FCC, FHWA, FTA, FAA, NPS, USACE	Drawing analysis + GIS boundary check
No ground disturbance beyond existing footprint	FCC Sec. 1.1306(c)(1)(D), FEMA, DOI, BLM	LiDAR + imagery change detection
No height increase >10% or 20 ft on existing structure	FCC Sec. 1.1306(c)(1)(A)	3D model analysis from drawings + Street View photogrammetry
No more than 4 new equipment cabinets	FCC Sec. 1.1306(c)(1)(B)	AI drawing review + equipment count extraction
No protrusion >20 ft beyond structure edge	FCC Sec. 1.1306(c)(1)(C)	Photogrammetric measurement from imagery
Existing ROW in active designated use	FCC Sec. 1.1306(c)(1)(i)-(ii)	GIS ROW database + land use classification
Routine O&M only, no significant change to historic character	NPS, USFWS, BLM, VA, NASA, TVA	Construction drawing scope classification
Section 106 consultation completed or no adverse effect finding	GSA, FBI, NPS, AFRH	SHPO API query + ACHP case log lookup

CE rule sets for standard project types are compiled into optimized decision trees at deployment, enabling sub-millisecond applicability screening for routine cases. Projects activating conditional extraordinary circumstance analysis route to full graph traversal with chain-of-thought LLM reasoning [13], preserving computational efficiency for the majority of cases while ensuring that complex regulatory logic receives thorough analytical treatment.

3.3 Multi-Source Federal Data Integration

Three source categories supply the platform's data inputs. Historic property and cultural resource systems, NPS IRMA [22], ACHP Section 106 case tracking, and the 50 state SHPO systems provide historic property adjacency status, prior Section 106 findings, and consultation contacts. Environmental and regulatory databases, FCC ULS, USACE ORM2, EPA EnviroFacts, and USFWS IPaC [21] supply regulatory permit history, environmental sensitivity flags, and extraordinary circumstance screening data. Geospatial and visual intelligence sources, Google Street View, USGS 3DEP, and commercial satellite imagery provide the site condition data that the computer vision analysis layer requires for CE threshold measurement. Table 2 presents the primary federal API inventory and its CE determination function.

Table 2. Key Federal APIs and CE Determination Relevance

API/Database	Agency	Data Provided	CE Function
NPS IRMA Data Store REST API [22]	NPS	NRHP property records, park boundaries, prior NEPA documents	Historic property adjacency; prior NEPA coverage lookup

API/Database	Agency	Data Provided	CE Function
ACHP Section 106 Case Tracking	ACHP	Active consultations, MOA/PA records	Consultation status; adverse effect history
FCC Universal Licensing System	FCC	Antenna registrations, structure heights, prior EA/CATEX records	FCC CE baseline; height increase calculation
USACE ORM2 / Nationwide Permit	USACE	Section 404/10 permit history, nationwide permit eligibility	USACE CE prior NEPA coverage screening
USFWS IPaC [21]	USFWS	ESA-listed species, critical habitat, migratory bird data	Extraordinary circumstance screening, all agencies
EPA EnviroFacts / ECHO	EPA	NPDES permits, Superfund sites, TRI facilities	Environmental sensitivity screening
USGS 3DEP	USGS	1-meter resolution elevation data, slope analysis	Ground disturbance extent; viewshed modeling
NPS Cultural Landscapes Inventory	NPS	Cultural landscape boundaries, character-defining features	Visual impact on historic character
NEPA Online (CEQ)	CEQ	Federal EIS/EA filings, agency NEPA contacts	Prior NEPA coverage tiering opportunities
State SHPO Systems (50 states)	50 state SHPOs	State nominations, prior reviews, consultation contacts	Section 106 consultation initiation; prior findings lookup

All API calls for a given project are launched in parallel at intake, reducing average data retrieval time from approximately 45 minutes under sequential processing to approximately 8 minutes, the author's primary research estimate from API latency analysis. Pre-cached monthly refreshes of historic district boundaries, NRHP records, and critical habitat maps eliminate real-time API calls for the most frequent lookups, concentrating API overhead on site-specific data that cannot be pre-computed.

3.4 Five-Stage Data Flow and Tiered Processing

Each project transits five pipeline stages. Stage 1 (Site Identification): submitted coordinates and project description trigger parallel database screening across NRHP, NPS, SHPO, and FCC systems to flag historic sensitivity and identify all applicable agencies. Stage 2 (Document Ingestion): vision transformer models [16] and OCR extract project dimensions, equipment counts, materials, and scope from construction drawings and photographs, producing a structured project characterization record. Stage 3 (Regulatory Mapping): knowledge graph traversal produces a CE applicability manifest and documentation requirements for all applicable agencies, identifying every CE condition that must be satisfied or flagged. Stage 4: (AI Analysis) The DeepLearning computer vision models assess the four key FCC Section 1.1306(c) CE conditions (height increase, equipment cabinet changes, protrusion changes, and soil disturbance) and historic character indicators constructed from the downloaded Street View data [4]. Stage 5 (Output Generation): agency-specific documentation is produced using template-driven generation with LLM-enhanced narrative text [17], then routed to the Analyst Workbench for mandatory human review before issuance.

Project complexity determines routing through a three-tier processing architecture (Table 3). This tiering mechanism is how the platform preserves human oversight for high-stakes determinations while enabling automated throughput for routine cases.

Table 3. Tiered Processing Architecture

Tier	Criteria	Processing Path	Target Turnaround
Green, Routine CE	No historic sensitivity, standard project type, existing ROW, prior CE coverage	Automated CE determination, AI-generated documents, analyst spot-check	< 2 hours
Yellow, Enhanced Review	Adjacent to historic district, minor construction, visual impact possible, new CE application	Full AI analysis, AI-drafted documents, mandatory analyst review before issuance	< 24 hours
Red, Complex or Extraordinary	Within historic district, ground disturbance present, potentially causing adverse effect, novel project type	Full AI analysis + extraordinary circumstance assessment, senior analyst review	3-5 business days

3.5 Human-in-the-Loop Override Architecture

No AI system should issue a regulatory determination without human review; such action is a legal requirement under the judicial defensibility framework applicable to NEPA decisions [7]. The platform implements this through the Analyst Workbench: every AI CE determination requires analyst approval before issuance. The extraordinary circumstance first-flag policy automatically escalates ambiguous cases rather than routing them to automated clearance, ensuring that the most consequential determinations receive expert attention regardless of model confidence level. Override decisions require structured justification, which is captured as a labeled training example and logged to the immutable audit trail alongside every data source queried, API call executed, and AI inference generated. SHA-256 hashing of the audit log provides the tamper-evident administrative record required for regulatory defensibility and FOIA response.

4. Performance Analysis

4.1 Projected SLA Improvements

Table 4 compares projected SLA performance across seven key stages of the CE determination workflow, contrasting the current manual baseline with the proposed AI platform. All figures represent the author's primary research projections based on architecture-level analysis of workflow stage durations and API response latencies.

Table 4. Projected SLA Improvements: Current Manual Baseline vs. AI Platform (Author's Primary Research Projections)

Process Step	Current Average	With AI Platform	Projected Improvement
Site sensitivity screening	1-3 days	< 5 minutes	~99%
CE applicability determination (Tier 1)	3-5 days analyst research	< 1 hour automated	~98%
FCC co-exclusion letter drafting	2-4 days	< 30 minutes automated	~97%
SHPO consultation letter preparation	1-2 weeks	< 4 hours (Tier 2)	~97%
Construction drawing review for CE thresholds	1-2 days per drawing set	< 2 hours AI analysis	~90%
Multi-agency coordination package (Tier 3)	4-8 weeks	3-5 business days	~85%
Post-construction compliance verification	2-3 weeks field inspection + report	AI imagery comparison + report	~80%

The most substantial gains occur where mechanical data retrieval dominates current practice. Site sensitivity screening, manual cross-referencing of NPS, SHPO, and FCC databases, drops from 1 to 3 days to under 5 minutes through parallel API orchestration. CE applicability determination drops from 3 to 5 days to under 1 hour for Green-tier projects through knowledge graph traversal. The Tier 3 multi-agency coordination package retains a 3 to 5 business day target by design; this constraint preserves the human quality gate for extraordinary circumstance determinations where AI acceleration should not substitute for professional regulatory judgment.

4.2 Telecommunications Sector Application

The FCC has estimated that approximately 800,000 small cell deployments are required to complete U.S. 5G coverage [8]. Current NEPA/Section 106 compliance timelines of 15 to 30 days per site represent one of the documented regulatory constraints on deployment velocity. For FCC Section 1.1306(c) CE determinations, the predominant compliance event in urban 5G programs, the Green-tier automated pathway reduces per-site analyst time substantially, generating per-site labor savings of $2,400 to $7,000 against a current cost of $3,000 to $8,000 per site, based on the author's primary research projections. Multi-agency projects, requiring simultaneous FCC, NPS, and SHPO review, benefit from cross-agency deduplication: a single unified site assessment replaces three independent analyses, projected to reduce document production time by 60 percent for this project category, the author's primary research projection from multi-agency workflow analysis.

The Army Corps of Engineers processes approximately 85,000 Section 404 permit applications annually [9], with an estimated 35 percent involving historic or culturally sensitive areas [9]. AI-assisted CE determination for this subset could substantially reduce USACE analyst workload for routine cases, freeing specialist capacity for complex, extraordinary circumstance determinations requiring genuine environmental expertise. The platform's multi-agency architecture is directly extensible to USACE CE workflows through the same knowledge graph and API integration infrastructure built for the telecommunications use case.

4.3 Continuous Learning and Regulatory Adaptation

Analyst overrides generate labeled training examples with mandatory structured justification, feeding monthly fine-tuning cycles. Every fix improves the system's accuracy on corner cases and project types it hasn't previously seen, effectively turning production into a continuous training signal, solving the key limitation of earlier rule-based compliance systems, which were incapable of evolving with changing regulations. SLA monitoring tracks per-agency statutory deadlines, such as 30 days for Section 106 responses from State Historic Preservation Offices (SHPO). Escalation notifications are automatic at 75 percent deadline expiry. Newly published agency CE catalog amendments automatically trigger knowledge graph update workflows, eliminating the need for manual curation of regulatory coverage currency for each CE.

5. Discussion

5.1 Cross-Sector Applicability

The architecture's modular agency connector design enables direct extension beyond telecommunications. Electric utilities executing FERC-regulated transmission line upgrades through historic district corridors, DOD installations managing construction on historically significant military properties, state DOTs executing FHWA-funded historic bridge rehabilitation, and GSA managing NEPA compliance for historic federal building modifications all face structurally identical multi-agency CE workflows. Each new agency CE catalog requires only a new connector and knowledge graph submodule, the core RAG, computer vision, and document generation layers require no modification. As federal agencies progressively adopt standardized REST API access consistent with NPS IRMA [22] and USFWS IPaC [21] design patterns, the platform's data coverage and processing efficiency will improve organically without architectural change.

5.2 Legal Defensibility Under Judicial Review

Under the judicial scrutiny framework for AI-assisted NEPA decisions [7], two requirements determine regulatory defensibility: source traceability to authoritative regulatory texts and preserved human agency decision-making authority. The proposed architecture satisfies both structurally, not as add-on safeguards but as core design constraints embedded from the outset. Source traceability is enforced architecturally; every CE determination output links to the specific regulatory provision, database record, and imagery analysis from which it was derived. Human authority is preserved through the mandatory analyst approval architecture: the extraordinary circumstance first-flag policy ensures that the most consequential determinations receive expert review regardless of AI confidence level, and the immutable audit log records every human decision with its justification for administrative record purposes [7].

5.3 Limitations and Implementation Prerequisites

Three structural prerequisites constrain immediate full-scale deployment. Federal API standardization across the 50 SHPO systems, which currently use heterogeneous proprietary formats, represents the most consequential barrier; NPS IRMA [22] and USFWS IPaC [21] demonstrate the technical feasibility of standardized REST access, but state-level adoption requires federal program coordination. Training data availability for novel project types limits initial accuracy for first-of-kind deployments in previously unsurveyed historic districts, a constraint that diminishes progressively as the continuous learning loop accumulates labeled determinations. The evolving legal framework for AI-assisted regulatory decisions [7] creates residual uncertainty about liability allocation when AI-assisted CE determinations are challenged on administrative record grounds, a question requiring resolution through regulatory guidance or judicial precedent before broad institutional adoption.

6. Conclusion

Scaling U.S. infrastructure deployment across the 5G telecommunications, electric grid modernization, and transportation rehabilitation sectors requires addressing a structural inefficiency that has persisted for decades: the per-site compliance burden of multi-agency NEPA categorical exclusion determination for projects in historic districts. This paper has presented an integrated AI platform that automates CE determination using graph RAG over a structured multi-agency CE knowledge graph, transformer-based computer vision for site condition analysis, parallel federal API orchestration, and LLM-driven agency-specific document generation. The framework reduces CE compliance timelines from weeks to hours for standard cases while maintaining the legal defensibility requirements established for AI-assisted NEPA decisions. Three contributions characterize the work. The multi-agency CE knowledge graph provides the first machine-traversable representation of the 30-plus federal agency CE catalog for the historic district determination context, enabling cross-agency CE resolution from a unified project description through a single graph traversal pass. The tiered processing model, Green/Yellow/Red, matches analytical depth to project complexity, preserving mandatory human oversight for extraordinary circumstance determinations while enabling automated throughput for routine cases. The human-in-the-loop override architecture with an immutable audit trail addresses the source traceability and independent agency judgment requirements that govern judicial review of AI-assisted NEPA determinations. Three directions for future work follow from these contributions: pilot deployment with a U.S. telecommunications carrier for FCC Section 1.1306(c) CE workflow validation and model performance measurement under real operational conditions; extension to USACE Section 404 permit automation using the knowledge graph and API integration architecture developed here; and engagement with CEQ, FCC, and ACHP to advance REST API standardization across the SHPO network, the data infrastructure prerequisite most critical to at-scale platform deployment.

References

[1] U.S. Government, "National Environmental Policy Act of 1969 (Public Law 91-190)," Enacted July 4, 2025. Available: https://www.govinfo.gov/content/pkg/COMPS-10352/pdf/COMPS-10352.pdf

[2] National Park Service, "National Historic Preservation Act of 1966." Available: https://www.nps.gov/subjects/archeology/national-historic-preservation-act.htm

[3] Office of NEPA Policy and Compliance, "40 CFR 1500-1508: CEQ Regulations for Implementing the Procedural Provisions of NEPA," 2005. Available: https://www.energy.gov/nepa/articles/40-cfr-1500-1508-ceq-regulations-implementing-procedural-provisions-nepa-ceq-1978

[4] Federal Communications Commission, "Environmental Rules, 47 CFR Sections 1.1301-1.1319," 2001. Available: https://docs.fcc.gov/public/attachments/FCC-01-319A1.pdf

[5] Advisory Council on Historic Preservation, "Section 106 Regulations, 36 CFR Part 800." Available: https://www.achp.gov/digital-library-section-106-landing/section-106-regulations

[6] P. Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," NeurIPS 33 (2020): 9459-9474. Available: https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html

[7] T. Saren, "Using AI in NEPA Review: Legal Challenges and Judicial Scrutiny," Environmental Law Reporter (2025). Available: https://www.elr.info/sites/default/files/files-general/55.UsingAIinNEPAReview_0.pdf

[8] Federal Communications Commission, "FCC Facilitates Wireless Infrastructure Deployment for 5G," 33 FCC Rcd 9088 (14), 2018. Available: https://www.fcc.gov/document/fcc-facilitates-wireless-infrastructure-deployment-5g

[9] Congressional Research Service, "U.S. Army Corps of Engineers: FY2024 Appropriations," 2023. Available: https://www.congress.gov/crs_external_products/IF/PDF/IF12370/IF12370.3.pdf

[10] H. Naveed et al., "A Comprehensive Overview of Large Language Models," ACM Trans. Intell. Syst. Technol. 16, no. 5 (2025): 1-72. Available: https://dl.acm.org/doi/full/10.1145/3744746

[11] Y. Gao et al., "Retrieval-Augmented Generation for Large Language Models: A Survey," arXiv:2312.10997 (2023). Available: https://arxiv.org/abs/2312.10997

[12] A. Vaswani et al., "Attention Is All You Need," NeurIPS 30 (2017). Available: https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

[13] J. Wei et al., "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models," NeurIPS 35 (2022): 24824-24837. Available: https://proceedings.neurips.cc/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html

[14] X. Ma, G. Fang, and X. Wang, "LLM-Pruner: On the Structural Pruning of Large Language Models," NeurIPS 36 (2023): 21702-21720. Available: https://proceedings.neurips.cc/paper_files/paper/2023/hash/44956951349095f74492a5471128a7e0-Abstract-Conference.html

[15] I. Chalkidis et al., "LEGAL-BERT: The Muppets Straight Out of Law School," Findings of EMNLP 2020, pp. 2898-2904. Available: https://aclanthology.org/2020.findings-emnlp.261/

[16] N. Carion et al., "End-to-End Object Detection with Transformers," ECCV 2020, pp. 213-229. Available: https://link.springer.com/chapter/10.1007/978-3-030-58452-8_13

[17] T. Brown et al., "Language Models Are Few-Shot Learners," NeurIPS 33 (2020): 1877-1901. Available: https://proceedings.neurips.cc/paper_files/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html

[18] D. Edge et al., "From Local to Global: A Graph RAG Approach to Query-Focused Summarization," arXiv:2404.16130 (2024). Available: https://arxiv.org/abs/2404.16130

[19] National Conference of State Historic Preservation Officers, "Section 106," NCSHPO, 1966. Available: https://ncshpo.org/resources/section-106/

[20] R. Bommasani et al., "On the Opportunities and Risks of Foundation Models," arXiv:2108.07258 (2021). Available: https://arxiv.org/abs/2108.07258

[21] U.S. Fish & Wildlife Service, "IPaC: Information for Planning and Consultation," Accessed: 2026. Available: https://www.fws.gov/service/information-planning-and-consultation

[22] National Park Service, "IRMA Data Store REST API Documentation," Accessed: Jun. 2026. Available: https://irmaservices.nps.gov/datastore/v8/documentation/datastore-api.html