

## ATLANTA LEGAL AID SOCIETY: AI TAGGING TOOL Q&A

Topic	Question	Answer
Proposal Submission	How do I submit a proposal?	Email proposals to Virginia O'Reilly (vmoreilly@atlantalegalaid.org) by May 29th.
Budget, Contractual & Commercial	What is the allocated budget, budget range, or ceiling for this project?	We have allocated \$24,000 for this portion of the project, but we have some flexibility within the overall project budget and may be able to increase that amount based on scope and proposed approach.
Budget, Contractual & Commercial	What is the preferred pricing model, milestone payment structure, and is there a holdback percentage prior to final acceptance?	We anticipate paying 10% on contract signature, 40% after the delivery of the first tool iteration (model and UI), 40% after the achievement of the proposed F1 score and the delivery of the remaining materials, and 10% after the certification that training and documentation are complete and replicable by ALAS.
Budget, Contractual & Commercial	Is there an incumbent vendor? If so, why are you looking elsewhere, and are they bidding?	There is no incumbent vendor.
Budget, Contractual & Commercial	Will proposals for both RFPs be evaluated independently, or is there a preference/discount expected for a combined award?	Proposals for all RFPs will be evaluated over the same time period. We have not yet decided if we prefer to choose the same vendor for two or more projects.
Budget, Contractual & Commercial	Are there staffing location expectations (e.g., US/India blended team) or restrictions on offshore access, VDI, and VPNs?	There are no expectations for vendor location; however all transactions must be made in USD between ALAS and a US-based entity. There are no restrictions on VPN, VDI, or data access because all code will be hosted in public repositories and all data shared will be publicly available.
Budget, Contractual & Commercial	Will ALAS provide a standard contract/IP assignment form regarding ownership and AGPLv3 redistribution?	The contract terms will include provisions that recognize the "work for hire" nature of all the deliverables, and the licensing requirements for any included works in the deliverables.
Budget, Contractual & Commercial	Is professional liability (E&O) insurance required, and at what minimum coverage level?	Please disclose your current limits for Professional Liability (E&O) and Cyber Incident Response insurance. While there is no mandatory minimum required to submit a proposal, insurance adequacy commensurate with the risk profile may be a factor in our evaluation
Budget, Contractual & Commercial	If training data differs materially from assumptions, will that be handled via a change order, or should contingencies be baked into the fixed price?	Because there are a number of answers to the size, format, and scope of the training data, any contingencies should be baked into the fixed price.
Training Data & Document Ecosystem	What is the estimated size, format, and structure of the currently available training dataset and document-to-authority mappings?	The datasets will be raw data in the form of { document, list-of-authorities } where "document" will be, generally, in MS-Word (docx), OpenDocument (ODT), HTML, or "searchable" PDF format, or a URL to one of the above, and "list-of-authorities" will be an array of "authority", where "authority" is { citation , last-retrieved-text-of-authority}. All "authorities" will be either Federal Code (USC), Federal Regulation (CFR), or a rule from the Federal Rules of Civil Procedure (FRCP) or Rules of Evidence (FRE).
Training Data & Document Ecosystem	Will ALAS provide the data (Georgia-specific, Federal, Partner datasets), or should vendors supplement with public sources? Are the partner datasets (Texas, Minnesota, etc.) accessible now and fully validated?	ALAS will provide the data as it becomes available. The datasets have not yet been compiled.

## ATLANTA LEGAL AID SOCIETY: AI TAGGING TOOL Q&A

Topic	Question	Answer
Training Data & Document Ecosystem	Are the datasets currently clean and labeled, or is vendor preprocessing and annotation required?	The datasets will be raw data in the form of { document, list-of-authorities } where "document" will be, generally, in MS-Word (docx), OpenDocument (ODT), HTML, or "searchable" PDF format, or a URL to one of the above, and "list-of-authorities" will be an array of "authority", where "authority" is { citation , last-retrieved-text-of-authority}. All "authorities" will be either Federal Code (USC), Federal Regulation (CFR), or a rule from the Federal Rules of Civil Procedure (FRCP) or Rules of Evidence (FRE).
Training Data & Document Ecosystem	What is the distribution of document types (HTML, PDF, etc.), and are there known formatting inconsistencies, encoding issues, or expected length variations?	The documents will generally be MS-Word (docx), OpenDocument (ODT), HTML, and "searchable" PDF.
Training Data & Document Ecosystem	Can the model be trained on this data considering privacy constraints, and do historical documents contain sensitive client data?	All materials provided will be publicly available data. No confidential information nor personally identifiable information will be included.
Training Data & Document Ecosystem	Where do the canonical Georgia self-help materials live (since the provided URL 404s), and is data ingestion in scope?	Georgia self-help materials can be reached at <a href="https://www.georgialegalaid.org">https://www.georgialegalaid.org</a>
Training Data & Document Ecosystem	Are there any limitations on utilizing synthetic data or augmentation techniques?	There are no limitations on using synthetic data or augmentation techniques as long as the proposed F1 score is achieved.
Technical Architecture & Infrastructure	What is the target deployment environment (Cloud, On-premise, Hybrid) and hosting setup (managed Drupal hosting vs. dedicated VPS)?	ENVIRONMENT: Managed Drupal (Pantheon) + Private Cloud AI. SETUP: Vendor must provision the AI engine to integrate with the core managed host, optimizing for performance between tools while ensuring the architecture remains provider-agnostic.
Drupal Integration, Workflow & Retraining	What compute environment will ALAS have available post-handoff for staff to run the autonomous retraining workflow?	Open to environment based on solution. We don't have a current environment used for other tools.
Technical Architecture & Infrastructure	What server specifications (CPU/GPU, RAM) will ALAS provision for model training and inference? Is GPU inference in budget scope, and who provides the initial training compute?	Vendor will perform the initial model training. Inference, and retraining based on user feedback, will need to take place on ALAS infrastructure. Vendor will strive to develop a solution with ongoing costs as low as possible.
Technical Architecture & Infrastructure	Should the solution be fully embedded as a native Drupal module or allowed as an external microservice with API integration?	A native Drupal module is preferred. If the solution includes an external microservice with API integration then ALAS must be able to deploy and host such a microservice, and to that end, the source code to the microservice must be provided as a deliverable.
Technical Architecture & Infrastructure	Are you open to an architecture that uses commercially available LLMs (via BYO keys/Intelligent Model Routing), or does that violate the proprietary dependency evaluation criterion?	Yes. Commercial LLMs (via BYO keys) are acceptable. To satisfy the Sustainability/Open-Source criteria, the tool must use the Drupal AI abstraction layer to remain provider-agnostic and swappable. All data must reside in a Private Instance and never be used to train third-party models.
Technical Architecture & Infrastructure	What is the current configuration of the Georgia LACI platform (Drupal version, existing AI modules, provider, notification systems in place)?	The Georgia LACI platform will be deployed as a deliverable under a related RFP, see (link here)
Technical Architecture & Infrastructure	Is a staging/dev environment available for vendor access during the project, or must the vendor stand up their own instance?	Vendor must provide their own environment for development and staging.

## ATLANTA LEGAL AID SOCIETY: AI TAGGING TOOL Q&A

Topic	Question	Answer
Model Strategy, Performance & Evaluation	What is the minimum acceptable F1 score or target performance threshold for initial delivery versus over time?	<b>TARGET:</b> Vendors must propose the highest achievable F1 Score within the budget/timeline; this proposed target will serve as the benchmark for deliverable acceptance. <b>GROWTH:</b> We prioritize a human-in-the-loop (HITL) architecture and a defined retraining workflow to ensure performance iteratively improves beyond the initial delivery.
Model Strategy, Performance & Evaluation	How will the evaluation dataset be managed (vendor-created vs. ALAS gold-standard holdout)?	<b>DATASET:</b> Vendor is responsible for creating the initial training and evaluation datasets using ALAS content. <b>EVALUATION:</b> ALAS "reviewers" will validate a specific subset to serve as the gold-standard holdout, which will be used to benchmark the vendor's proposed F1 score during deliverable acceptance.
Model Strategy, Performance & Evaluation	Will extraction and prediction be evaluated with separate F1 scores or combined, and will vendors see the methodology before final testing?	<b>SEPARATE.</b> Extraction and Prediction will be evaluated with individual F1 scores. <b>THRESHOLD:</b> ALAS expects a higher performance baseline for Extraction (identifying the correct text/citations) than for Prediction (topic classification/tagging). <b>METHODOLOGY:</b> The evaluation framework and metrics will be shared with vendors, but the "gold-standard" holdout dataset will remain private until final testing.
Model Strategy, Performance & Evaluation	Should the solution prioritize explicit extraction, predictive tagging, or both equally? Should the model suggest only known authorities or generalize to unseen ones?	<b>BOTH EQUALLY.</b> While both are core requirements, Extraction prioritizes Precision for citation reliability. <b>GENERALIZE:</b> The model should generalize to the broader target legal corpus or dataset. It is expected to suggest relevant authorities from the entire available source material, using the initial training subset as a benchmark. The human-in-the-loop (HITL) workflow will be used to validate these broader suggestions, creating the ground-truth data needed to refine the model in future retraining cycles.
Model Strategy, Performance & Evaluation	Beyond hitting the baseline F1 score, will formal acceptance depend on a qualitative manual review by ALAS staff, and how is success measured overall?	Yes. Formal acceptance depends on both the quantitative F1 benchmark and a qualitative internal human audit. <b>SUCCESS MEASURES:</b> 1) Meeting the vendor-proposed F1 benchmark on the holdout set, and 2) Passing a manual qualitative spot-check by ALAS staff to ensure extracted text and predicted tags meet professional legal standards for accuracy and contextual relevance.
Drupal Integration, Workflow & Retraining	What are the expectations for the reviewer interface? Are there wireframes, brand guidelines, or should it just be the described accept/reject workflow?	<b>INTERFACE:</b> No pre-existing wireframes; the vendor is expected to design these during the Discovery Phase. <b>UI:</b> Functional accept/reject/edit workflow integrated into Drupal as described in the Scope of Work. <b>BRANDING:</b> The core LACI branding should serve as the starting point for the interface design. However, to facilitate reuse, the implementation must remain brand-agnostic (white-labeled). The vendor should utilize a clean, professional theme and standard CSS/theming practices that allow future organizations to easily apply their own branding without requiring core code changes.

## ATLANTA LEGAL AID SOCIETY: AI TAGGING TOOL Q&A

Topic	Question	Answer
Drupal Integration, Workflow & Retraining	Should the interface include advanced analytics (e.g., error trends, retraining dashboards)?	Core functionality and the correction/retraining workflow are the top priorities. Essential logs (e.g., accuracy, error counts, and correction history) are required to support retraining; however, advanced visual dashboards should be proposed as a cost-effective "stretch" goal to ensure the limited budget is focused on the inference engine first.
Drupal Integration, Workflow & Retraining	How does a new review begin, what does the human-in-the-loop workflow look like, how many concurrent reviewers will use the tool, and what is the expected daily document volume?	TRIGGER: A review can be initiated through multiple channels: automated "cron" checks for document changes in external authorities, manual uploads of new content, or bulk ingestion during the onboarding of a new organization. HITL WORKFLOW: The system generates suggested authorities and categories, flagging them for human oversight. Through the reviewer interface, a "Monitor" audits these suggestions to Accept, Reject, or Edit. The system must log every modification made by the reviewer to establish a "ground-truth" dataset, which will be used to track performance and drive future retraining cycles. USER SCALE: The tool is designed for a small, specialized team, typically fewer than 10 concurrent users. DOCUMENT VOLUME: While daily document volume fluctuates based on legislative and regulatory update cycles, the architecture must support a burst capacity of approximately 1,000 documents per day to accommodate the onboarding of new organizations or the migration of historical datasets.
Drupal Integration, Workflow & Retraining	How frequently does ALAS anticipate running retraining cycles (monthly, weekly, on-demand), and will there be a transition period for assisted retraining?	FREQ: Primarily on-demand based on correction volume. TRANSITION: Vendor must provide assisted retraining during the Finalization Phase to ensure independent ALAS operation.
Scope Boundaries, Process & Timeline	What is the overall project schedule, and when exactly will the selected vendor receive access to training data and staging environments?	See "Project Phase" in the RFP for the timeline. We expect integration into, e.g., staging to happen in Phase 3. Training data will be provided before the start of Phase 2.
Scope Boundaries, Process & Timeline	Will the final module be published to Drupal.org or Git repositories, and should documentation align with community contribution standards?	The final deliverables will be published at <a href="https://codeberg.org/LACI">https://codeberg.org/LACI</a> and code must conform to the standard at <a href="https://codeberg.org/LACI/laci-core/wiki/CodingStandards">https://codeberg.org/LACI/laci-core/wiki/CodingStandards</a> to be accepted.
Scope Boundaries, Process & Timeline	Should the system support all authority types equally at first, and are there expectations to support additional jurisdictions beyond Georgia/Federal at launch?	The system should support all Federal authority types equally at first. There is not currently an expectation to support Georgia authorities at launch. However, the deliverables (code, data, documentation) must be sufficient for ALAS to train the solution to support Georgia authorities.
Scope Boundaries, Process & Timeline	May extraction and prediction capabilities be delivered and validated in separate project phases?	Yes. A phased delivery and validation approach is acceptable. PAYMENT: While ALAS is open to milestone deliverables, the final project disbursement will remain contingent upon the successful validation of both components as a unified, integrated solution.
Scope Boundaries, Process & Timeline	Does the portability documentation require adapting the actual extraction logic for differing state citation formats, or just outlining the process?	The portability documentation needs to outline the process for extracting citations in different state formats.

**ATLANTA LEGAL AID SOCIETY: AI TAGGING TOOL Q&A**

<b>Topic</b>	<b>Question</b>	<b>Answer</b>
Scope Boundaries, Process & Timeline	Does the requirement to deliver trained model weights apply exclusively to vendor-trained components (SLMs), excluding commercial-API tiers?	The requirement to deliver model weights is included so that partners can implement their own infrastructure to host this tool. They will need all the data and documentation necessary to do so.
Scope Boundaries, Process & Timeline	Are there ongoing initiatives, parallel projects, or current challenges that may affect project timelines?	No.