

Article

Not peer-reviewed version

---

# Conceptual Model and Working Prototype of a Personal Legal and Social AI Assistant for Decision Support in Bulgarian Social Services

---

[Penyo Georgiev](#)\*

Posted Date: 22 May 2026

doi: 10.20944/preprints202605.1550.v1

Keywords: artificial intelligence; social services; legal informatics; intelligent assistant; retrieval augmented generation; semantic retrieval; decision support; large language models; child protection; Bulgarian legislation



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Conceptual Model and Working Prototype of a Personal Legal and Social AI Assistant for Decision Support in Bulgarian Social Services

Penyo Georgiev

Faculty of Economic and Social Sciences, University of Plovdiv "Paisii Hilendarski", 24 Tsar Asen Str.,  
4000 Plovdiv, Bulgaria; penyo@uni-plovdiv.bg

## Abstract

Social service professionals operate in legally sensitive, administratively intensive, and context-dependent environments in which decision-making requires the simultaneous interpretation of regulatory norms, institutional procedures, and individual case circumstances. This paper proposes a conceptual model of a Personal Legal and Social Artificial Intelligence (AI) Assistant intended to support professional decision-making in social services, and demonstrates its functionality through a working prototype. The model is formulated as a domain-specific retrieval-augmented generation (RAG) framework in which a controlled legal and social document corpus is processed through text extraction, chunking, semantic indexing via SentenceTransformer embeddings, top-k retrieval through cosine similarity, and bounded large-language-model reasoning to produce grounded and explainable responses. The proposed framework is informed by three successive prototype versions and by observed sensitivity to corpus scope, document prioritization, and prompt constraints. The current prototype version operates on a prioritized corpus of sixteen Bulgarian normative acts complemented by three supplementary resources, comprising 883 indexed fragments, and uses DeepSeek as the reasoning model accessed through the OpenRouter API. The functionality of the model is validated through a representative use case concerning child protection, in which the prototype identifies the applicable legal provisions, exposes the retrieved documentary evidence, and generates a four-part structured analysis comprising legal qualification, applicable provisions, legal consequences, and recommendations for action. The main contribution lies in the formalization and prototype-level demonstration of a domain-specific AI assistant that combines legal grounding, social-context awareness, and bounded language-model reasoning for trustworthy decision support in regulated social-service practice.

**Keywords:** artificial intelligence; social services; legal informatics; intelligent assistant; retrieval-augmented generation; semantic retrieval; decision support; large language models; child protection; Bulgarian legislation

---

## 1. Introduction

The rapid digitalization of public services increasingly demands intelligent systems capable of supporting knowledge-based activities [1,2]. This need is especially pronounced in social services, where specialists operate in environments shaped by legal regulation, institutional accountability, vulnerable target groups, and substantial documentary workload [3,4]. Professional decisions in this domain rarely rely on a single source of information; rather, they emerge from the simultaneous interpretation of laws, by-laws, methodological guidelines, case-specific information, and service provision rules [5].

Although contemporary AI systems provide advanced natural language processing capabilities [6], general-purpose conversational models remain insufficient for legally constrained and socially sensitive environments. In such contexts, unsupported or weakly grounded responses may lead to

procedural inconsistencies, incomplete regulatory orientation, or inappropriate recommendations [7]. Recent investigations into the behavior of large language models in regulated domains have repeatedly demonstrated that linguistic plausibility does not guarantee normative adequacy, and that unconstrained generation can produce statements that are stylistically convincing but legally incorrect or untraceable to authoritative sources. The effective use of AI in social services therefore requires domain-specific assistants that operate within explicit documentary boundaries, rely on controlled knowledge sources, and generate traceable outputs under professional supervision [8,9].

In the international literature, research on AI-supported decision-making in social services has developed along several complementary directions. Early work on the application of computers in social work established the value of structured knowledge representation for supporting professional activity [10]. More recent research on electronic case management systems emphasizes that digitalization in social services requires integration of documentary traceability, role-based access, and procedural control [11]. The emergence of large language models (LLMs) has opened new possibilities for legal question answering and document-grounded reasoning; at the same time, research on specialized legal LLMs confirms that domain specificity, citation accuracy, and explainability require careful architectural design beyond standard query-response interaction [12,13]. Retrieval-augmented generation has emerged as a particularly suitable paradigm for regulated domains because it externalizes factual grounding into a curated corpus, which in turn allows responses to remain anchored to verifiable source passages rather than to the parametric knowledge of the model [9]. In parallel, the research group at the University of Plovdiv has developed a series of personal assistant models across different domains: virtual educational environments [14–16], ontology-based knowledge representation [17], and domain-specific diagnostic support [18]—the last of which represents a direct architectural predecessor of the present work.

Despite this body of research, integrated assistants designed specifically for Bulgarian social-service practice remain absent from the published literature. The Bulgarian normative corpus governing social services is extensive, fragmented across several primary acts and a substantial body of implementing regulations, and is updated frequently; at the same time, the language is Bulgarian and the procedural logic is country-specific, which limits the applicability of internationally trained legal LLMs. This combination of regulatory density, linguistic specificity, and procedural particularity defines a clear gap in the current research landscape.

The aim of this paper is to present and demonstrate a conceptual model of a Personal Legal and Social AI Assistant for decision support in social services, accompanied by a working prototype. The proposed model is not formulated as an autonomous legal expert, but as a specialized component for professional support that combines normative documents, semantic retrieval, language-model reasoning, and human validation. The assistant is conceptualized as a RAG-based [9] and professionally supervised AI system, designed to improve regulatory orientation, case-related analysis, and the generation of structured responses. The scientific novelty of the proposed model lies in three directions: (a) it is, to the best knowledge of the author, the first documented domain-specific RAG prototype calibrated for Bulgarian social legislation with a focus on child protection; (b) the model is validated through three successive prototype versions that empirically demonstrate that a precisely selected smaller corpus of case-determining normative acts yields better results than a broader but more peripheral corpus; and (c) the demonstration includes explicit display of retrieved evidence alongside the generated analysis, which addresses the explainability requirement that is central for professional accountability in regulated environments.

## 2. Materials and Methods

### 2.1. Domain-Specific Foundations

The conceptual framework is grounded in the assumption that decision support in social services must be both legally informed and socially contextualized. The assistant therefore operates at the intersection of legal informatics, social-service casework, and intelligent information

processing. Its design is informed by the functional requirements identified during the development of the system, by the structure of the available normative corpus, and by observations from successive prototype stages. The knowledge base is built from a controlled documentary corpus that may include laws, implementing regulations, methodological instructions, service descriptions, and organization-specific procedural documents. This controlled corpus is essential because the assistant is expected to work with verifiable sources rather than with unrestricted external knowledge. The conceptual framework presented in Figure 1 formalizes the generalized model established already in Version 1 of the prototype; the specific technologies used in each development phase evolve across versions, but the controlled documentary corpus and the human-in-the-loop principle remain unchanged.

The proposed model follows the principle of bounded assistance. The assistant supports professionals in identifying relevant legal provisions, synthesizing documentary evidence, structuring case-oriented information, and generating draft analytical outputs. At the same time, it does not replace the social worker, case manager, or institutional specialist. Final interpretation, proportionality assessment, and case-specific decision-making remain the responsibility of the human expert. This human-in-the-loop principle is central to the trustworthiness of the model [7,17] and is consistent with established positions on the role of AI in social work, where automation is treated as a complement to professional judgement rather than as a substitute for it.

## 2.2. Retrieval-Augmented Architecture

From an informatics perspective, the assistant is modeled as a retrieval-augmented generation architecture based on a sequence of text extraction, segmentation, vector indexing, semantic retrieval, and structured response generation. In the first stage, legal and social documents are acquired, transformed into machine-processable text, segmented into semantically meaningful chunks, and represented through embeddings. The resulting vectors are stored in an indexed search structure that supports semantic retrieval. In the second stage, a user query or case-oriented prompt activates the retrieval process, which identifies the most relevant documentary fragments. In the third stage, the retrieved context is assembled into a bounded prompt and submitted to the large language model, which generates a structured answer tied to the selected sources.

The reasoning layer is responsible for transforming the retrieved legal-social context into a professional response. Its role is not to provide unrestricted knowledge, but to operate strictly within the documentary set supplied by the retrieval layer. Prompt constraints are used to limit the generation space, require documentary grounding, and structure the output into interpretable sections such as legal qualification, applicable provisions, legal consequences, and recommended actions. The architecture therefore combines semantic retrieval with bounded language-model reasoning, which is more suitable for regulated domains than direct prompt-only interaction [9,12].

To ensure the documentary grounding of the architecture, the prototype incorporates an explicit prompt template that supplies the language model with the relevant normative context, ordered by importance, and explicitly prohibits citation of laws, articles, or provisions outside the provided documentary set. The output structure is fixed in four sections: (1) legal qualification, (2) applicable provisions, (3) legal consequences, and (4) recommendations for action. This combination of semantic retrieval, bounded prompting, and structured output concretizes the conceptual model and shows why it is more suitable for regulated domains than freely generative conversational models [6,9,12].

## 2.3. System Implementation

The current prototype version of the assistant is built as a Django-based web application following a clear separation between an offline indexing stage and an online analysis stage. The offline stage is triggered by an explicit administrative action and is responsible for transforming the static legal corpus into a queryable semantic structure. PDF documents stored in a controlled directory are parsed by the server, segmented into chunks of fixed length, and submitted to a SentenceTransformer model that produces dense vector representations. Each fragment is persisted

together with its source metadata and the corresponding embedding vector in a serialized index file (embeddings.pkl), which transforms the static collection of documents into a high-dimensional space in which semantic proximity can be calculated mathematically.

The online stage is responsible for analyzing user-submitted situations. When the user submits a case description through the web interface, the server loads the precomputed index, encodes the query into a vector through the same SentenceTransformer model, and computes cosine similarity between the query vector and all indexed fragments. The top-seven most relevant chunks are then assembled into an enriched prompt that explicitly enumerates the retrieved legal context, instructs the model to ground its analysis exclusively in these passages, and prescribes the four-section output structure. The constructed prompt is dispatched through the OpenRouter API to the DeepSeek reasoning model, whose structured textual response is returned asynchronously to the client and stored within the temporary session.

Although the initial pilot version of the system (Version 1) used the open-weight Llama 3 model accessed through a local LlamaIndex/ChromaDB pipeline, subsequent iterations of the prototype revealed that response quality, multilingual handling of Bulgarian legal terminology, and robustness were more reliably achieved by accessing a hosted reasoning model through a controlled API. The current prototype therefore uses DeepSeek as the reasoning layer, while preserving the architectural commitment that the model never generates content outside the retrieved documentary context.

#### 2.4. Document Corpus

The document corpus is the architectural foundation of the assistant. In the current prototype, the corpus comprises sixteen Bulgarian normative acts directly relevant to social-service practice, supplemented by three thematic resources containing practical examples and frequently asked questions. The normative acts are explicitly ordered by importance for child-protection casework, with three acts marked as of paramount importance: the Social Services Act, the Regulations for the implementation of the Social Services Act, and the Child Protection Act. Subsequent acts in the prioritized order include the Regulations for the implementation of the Child Protection Act, the Social Assistance Act, the Regulations for the implementation of the Social Assistance Act, the Law on Protection from Domestic Violence, the Regulations for the implementation of the law on protection from domestic violence, and the Law on Family Benefits for Children, among others.

The three supplementary resources extend the legal corpus with practical knowledge that is not directly contained in the normative texts: a case study and intervention guide on AI-supported counseling for children at risk of school dropout, a question-and-answer document concerning children at risk of dropping out of the education system, and a question-and-answer document concerning the social service “Information and Counseling for Children without Disabilities”. These resources do not replace the normative basis of the analysis; rather, they enrich the retrieval space with realistic procedural context and intervention vocabulary, which improves the relevance of the retrieved chunks for case-oriented queries.

The current corpus is segmented into 883 fragments, which represents a deliberate reduction from earlier prototype versions. In Version 1, the corpus contained thirty-five laws and 1,049 chunks; in Version 2, the corpus was reorganized into thirty-four documents with 2,097 chunks; in the current Version 3, the corpus was reduced and reprioritized to sixteen laws, three resources, and 883 fragments. This reduction was a direct response to the observation that broader corpora can degrade retrieval precision, because peripheral or only weakly related documents introduce competing chunks that displace more relevant passages in the top-k results. The transition from breadth to prioritized depth is therefore not a simplification, but an architectural design decision informed by empirical observation across prototype iterations.

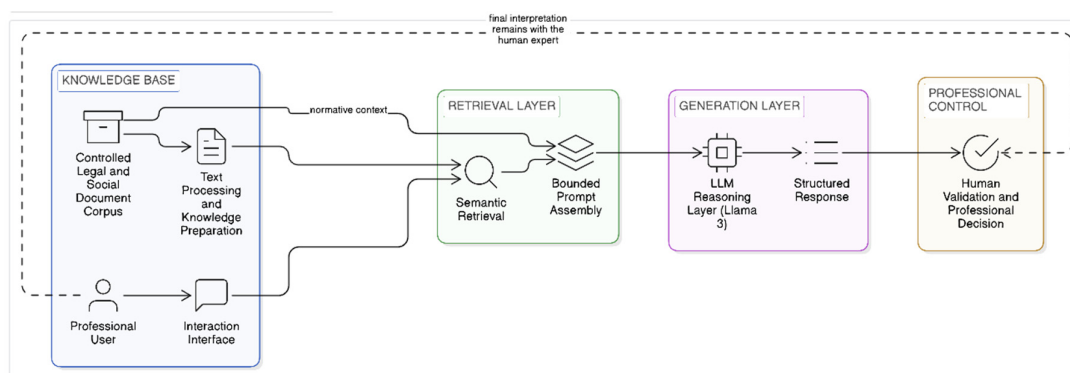
### 3. Results

This section presents the principal outcomes of the prototype-based investigation. It first summarizes the evolution of the assistant across three successive implementation versions, then

presents the conceptual architecture and the interaction workflow of the current prototype version, and finally demonstrates the operational behavior of the system through a representative case in the domain of child protection. The demonstration includes the user-facing interface, the retrieved documentary evidence, and the generated four-section legal analysis.

### 3.1. Conceptual Architecture

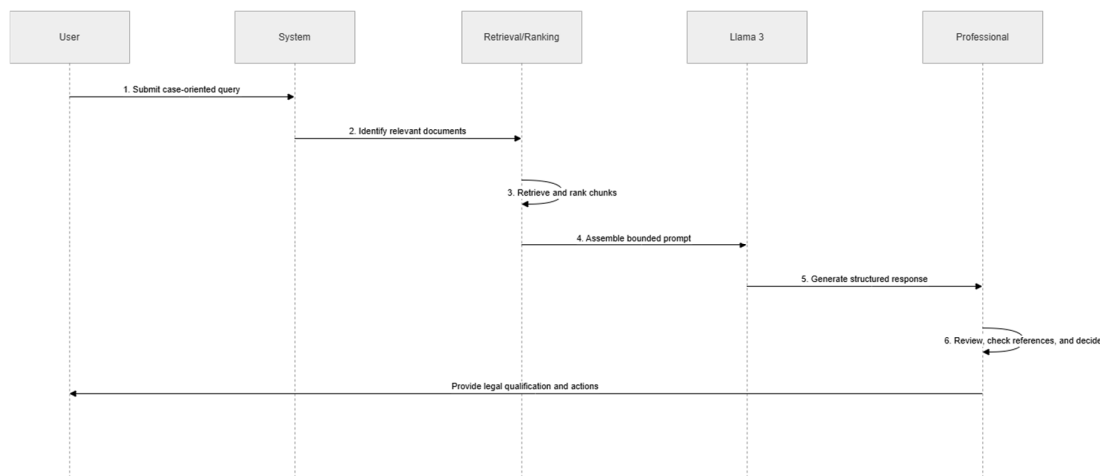
The conceptual architecture of the assistant is shown in Figure 1. The architecture is organized into four layers—a knowledge base containing the controlled legal and social document corpus together with the corresponding text-processing and knowledge-preparation routines; a retrieval layer responsible for semantic retrieval and bounded prompt assembly; a generation layer in which the reasoning model produces a structured response; and a professional-control layer in which the social worker validates the response and makes the final case decision. The diagram explicitly indicates that final interpretation remains with the human expert, reflecting the bounded-assistance principle of the model.



**Figure 1.** Conceptual architecture of the Personal Legal and Social AI Assistant.

### 3.2. Interaction Workflow

The interaction workflow is summarized in Figure 2. The workflow begins with a case-oriented input formulated by the professional user. The input may take the form of a question, a concise problem statement, or a description of a client situation requiring legal-social interpretation. The assistant then performs semantic retrieval over the controlled documentary corpus by matching the vector representation of the query against the pre-indexed fragments through cosine similarity and assembles a structured bounded prompt from the most relevant text fragments. This retrieved context is submitted to the reasoning model, which produces a structured answer tied to the selected documentary basis. The generated result is returned asynchronously to the user interface, and the final stage is professional validation, in which the social worker reviews the response, checks the cited evidence, and decides whether and how the output should be applied in practice.



**Figure 2.** Interaction workflow of the assistant.

### 3.3. Prototype Evolution

The development materials indicate a successive transition through three prototype versions. Version 1 represents a pilot prototype operating with thirty-five laws and 1,049 chunks, in which chunking, ChromaDB, LlamaIndex, and Llama 3 are used for retrieval-based legal analysis. Version 2 moves the system into a web environment and expands the documentary corpus to thirty-four documents and 2,097 chunks, while in an intermediate branch it reaches thirty-five documents and 1,798 chunks. Version 3, which represents the current prototype version, operates with a prioritized set of sixteen laws, three resources, 883 fragments, and a cumulative record of forty-three case analyses performed during development testing. In this version, the normative corpus is explicitly ordered by importance, the reasoning model has been changed to DeepSeek accessed through OpenRouter, and the prompt template enforces a four-section structured output: legal qualification, applicable provisions, legal consequences, and recommendations for action.

The reduction of the corpus from the broader early versions to the current set of sixteen documents reflects the practical observation that the quality of the analysis depends not only on the size of the corpus, but also on its composition, prioritization, and retrieval quality. When the corpus is too narrow, the generated support becomes incomplete; when it becomes broader, the precision of retrieval and document ranking becomes decisive for the quality of the response. This confirms that the effectiveness of the assistant depends not only on the language model itself, but also on corpus preparation, chunk granularity, retrieval quality, and prompt design. Consequently, the conceptual model treats corpus scope and retrieval control as architectural factors rather than as secondary prototype details.

### 3.4. Indexing and Retrieval Processes

Figures 3 and 4 illustrate the two principal processes of the prototype. Figure 3 depicts the offline preprocessing of the legal corpus, which is a prerequisite for the functioning of semantic retrieval. The process is initiated by an administrative action; the server extracts unstructured text from all PDF documents, segments it into semantically meaningful passages of fixed length, and generates vector representations through the SentenceTransformer model. The resulting data structure, containing the original text, source metadata, and the corresponding vector, is serialized into a file. This approach transforms the collection of static documents into a multidimensional space in which semantic proximity can be calculated mathematically.

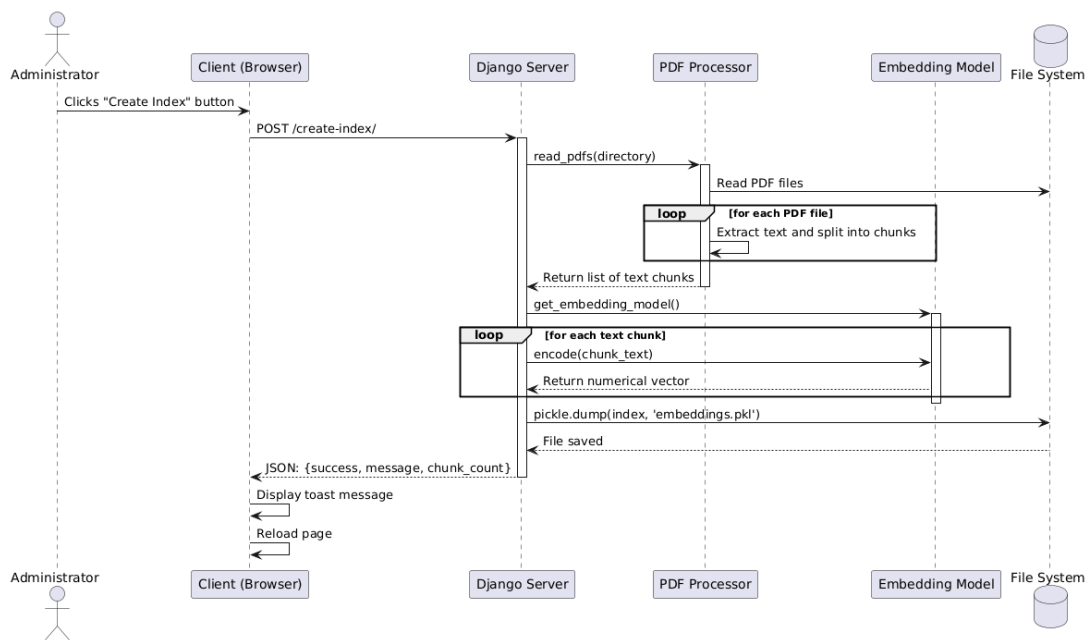


Figure 3. Sequence diagram of the offline process for creating the semantic index of the legal corpus.

Figure 4 illustrates the synergy between classical information retrieval and generative artificial intelligence in the online analysis stage. After a query is submitted, the server first performs semantic retrieval by transforming the user text into a vector and comparing it with all indexed fragments through cosine similarity. The identified top-seven relevant passages are injected as context into a prompt for the reasoning model. This approach, known as Retrieval-Augmented Generation, ensures that the analysis generated by the model is grounded in specific normative sources rather than relying on the model’s parametric knowledge. The model response is returned asynchronously and presented to the user without requiring a page reload, providing a smooth user experience even in the presence of potentially longer AI processing.

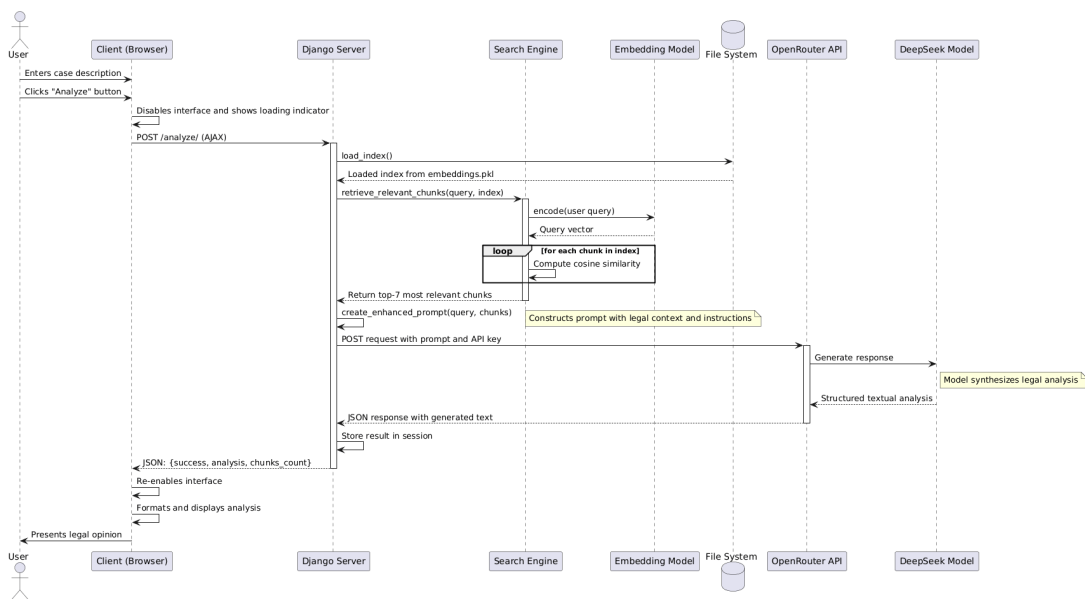


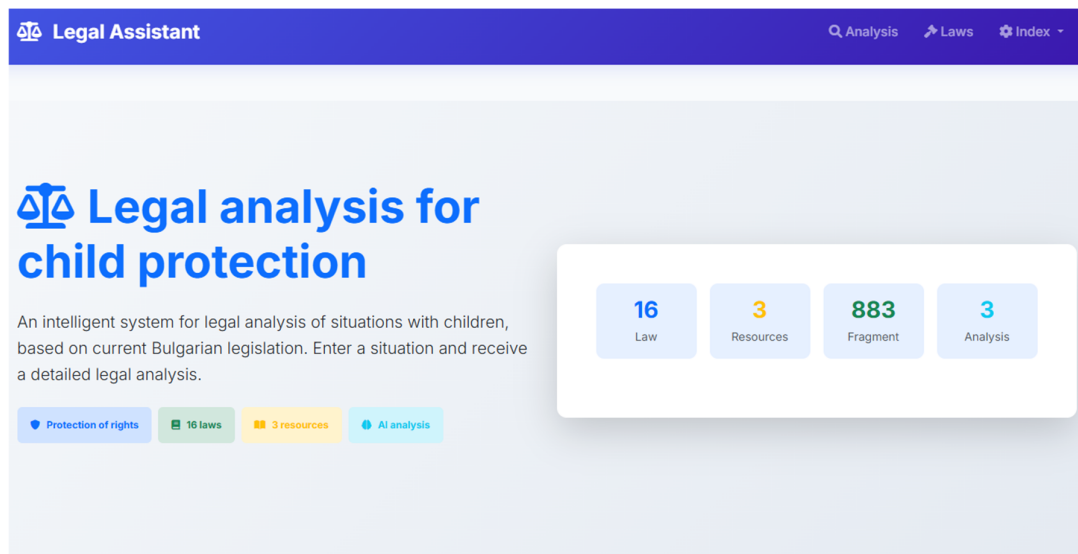
Figure 4. Sequence diagram of the online process for analyzing a user-submitted situation.

### 3.5. Use Case: Child Protection Scenario

To demonstrate the experimental behavior of the prototype, this subsection walks through a representative case in the domain of child protection. The case has been selected because child protection is one of the most regulation-intensive areas of Bulgarian social-service practice, requiring the simultaneous interpretation of the Child Protection Act, the Regulations for its implementation, the Law on Protection from Domestic Violence, and several adjacent acts. The walkthrough covers the main interface of the assistant, the case-description input, and the four-section analysis generated by the reasoning model.

#### 3.5.1. Main Interface

The main interface of the assistant is shown in Figure 5. The interface presents the title of the application, a short description of its purpose, and a summary of the current state of the knowledge base, including the number of indexed laws (sixteen), supplementary resources (three), text fragments (883), and previously generated analyses (three at the moment of the screenshot). The interface is intentionally minimal, oriented towards professional users who require an immediate overview of the system state and direct access to the analysis function.



**Figure 5.** Main interface of the assistant, showing the title, purpose, and a summary of the indexed corpus (16 laws, 3 resources, 883 fragments).

#### 3.5.2. Case Input and Retrieval

Figure 6 shows the case-input form. The user describes the situation in free text and triggers the analysis. For the illustrative example, the entered description is: “Mother testifies that father beats child. Parents are divorced. Child lives with mother, but has visits with father.” Upon submission, the system encodes the description into a vector representation through the SentenceTransformer model and computes cosine similarity against all 883 indexed fragments. The top-seven most relevant chunks are then retrieved and exposed to the user before being injected into the prompt for the reasoning model. This explicit exposure of retrieved evidence is a deliberate design choice: it allows the social worker to verify the documentary basis of the subsequent analysis and reinforces the explainability of the response.

## Q Legal analysis of a situation

Describe the situation with the child and receive a detailed legal analysis

The screenshot displays the 'Describe the situation' input form. The form has a blue header with a pencil icon and the text 'Describe the situation'. Below the header, there is a section titled 'Description of the situation:' with a list icon. The text area contains the following text: 'Mother testifies that father beats child. Parents are divorced. Child lives with mother, but has visits with father.' Below the text area, there is a lightbulb icon and the text 'Describe the situation in detail for a more accurate analysis'. At the bottom of the form, there are four buttons: 'Analyze the situation' (purple), 'Use index' (blue), 'Create index' (grey), and 'Clear analysis' (red). Below the form, there is a green header for 'Analysis results' with a document icon, a 'Latest analysis' button, a 'Copy' button, and a 'Situation+Analysis' button. Below the green header, there is a blue header for 'Analyzed situation:' with a question mark icon.

**Figure 6.** Case-description input form and main controls of the analysis interface, illustrated with the child-protection example used in the demonstration.

### 3.5.3. Structured Legal Analysis

Figure 7 shows the generated four-section analysis for the entered case across three vertically scrolled views (Figure 7a, 7b, 7c). The reasoning model, prompted with the retrieved chunks and the enforced output structure, identifies the Child Protection Act as the primary applicable law, citing Article 1, item 11 to qualify the situation under the legal definition of a “child at risk” as a child who is a victim of abuse, violence, exploitation or any other inhuman or degrading treatment or punishment within or outside his family. The model then complements the qualification with the definition of “prevention” from the same act (Article 1, item 13) and with a cross-reference to the Social Services Act (Article 1, item 8). In the situation-analysis section, the model interprets the post-divorce visitation arrangement as a recurring exposure to risk; in the practical-aspects section, it identifies the Child Protection Authority of the relevant municipality (Article 1, item 15) as the competent body for assessment and intervention; and in the recommendations section, it lists four sequential steps: reporting the case, performing a case-specific risk assessment, providing social services to the child and the mother, and coordinating the action across the Child Protection Authority, the court, the police, and social-service providers. The conclusion of the response explicitly states that the described situation qualifies the child as a “child at risk” under the Child Protection Act and requires immediate legal and social intervention.

📄 Analysis results Latest analysis

📄 Copy
🔍 Situation+Analysis

🔍 Analyzed situation:

“ Mother testifies that father beats child. Parents are divorced. Child lives with mother, but has visits with father. ”

🔍 Search statistics 7 relevant snippets

☰ View found fragments ▼

🤖 Analysis by DeepSeek AI model

### Legal analysis of the situation

**\*\*1. Legal basis\*\***

The main law applicable to the situation described is the **\*\*Child Protection Act\*\***. According to Art. 1, item 11 of this Act, a "Child at Risk" is defined as a child who is "a victim of abuse, violence, exploitation or any other inhuman or degrading treatment or punishment within or outside his family" (Document 3: Child Protection Act.pdf). This definition directly covers the case where the child has been beaten by his father.

Additionally, the Child Protection Act defines "prevention" as "protection of the child through information, assistance, support and services" (Art. 1, item 13 of the same document). Services within the meaning of the Act are the social services specified in the Social Services Act (Art. 1, item 8 of Document 3).

**\*\*2. Situation analysis\*\***

According to the information provided, the child falls into the **\*\*\*child at risk\*\*\*** category due to the violence exercised by the father. This is grounds for intervention by the child protection authorities. The fact that the parents are divorced and the child has visits with the father does not eliminate the risk, but makes it specific and recurring during the contacts.

The legal implications are that the case requires immediate assessment and action to protect the child. The aim is to ensure the child's physical and psychological safety. This may include temporarily restricting or suspending visitation with the father until the risk is assessed, as well as providing support and services for the child and mother.

**\*\*3. Practical aspects\*\***

Based on the legal definition of a "child at risk", the practical approach should be focused on **\*\*prevention and protection\*\***. This means activating the child protection system. The Child Protection Authority (CPA) in the municipality where the child has his/her current address (Art. 1, item 15 of Document 3) should be notified. The CPA has the obligation to assess the situation, protect the child and offer appropriate services.

Although the additional resources provided (R1, R2, R3) focus on the risk of school dropout, the logic of the intervention is similar: identifying a child at risk, assessing the situation and providing targeted services and support. In the case of physical violence, the priority is immediate safety.

**\*\*4. Recommendations for action\*\***

1. **\*\*Reporting the case\*\***: The mother or anyone else aware of the situation must immediately report the case to the **\*\*Child Protection Authority\*\*** in the municipality or to the **\*\*Police\*\***. This is the first and most important step in activating the protective mechanisms.
2. **\*\*Assessment and protection\*\***: After reporting, the CPD is obliged to carry out an assessment of the risk to the child (Document 3. definition of a child at risk). Based on this assessment,

**Figure 7. a.** Structured legal analysis generated by the reasoning model for the child-protection case (a): analyzed situation and legal basis. **b.** Structured legal analysis generated by the reasoning model for the child-protection case (b): situation analysis and practical aspects.

1. **Reporting the case**: The mother or anyone else aware of the situation must immediately report the case to the **Child Protection Authority** in the municipality or to the **Police**. This is the first and most important step in activating the protective mechanisms.

2. **Assessment and protection**: After reporting, the CPD is obliged to carry out an assessment of the risk to the child (Document 3, definition of a child at risk). Based on this assessment, immediate protection measures can be taken, including temporary arrangements regarding contact with the father.

3. **Provision of services**: The child and the mother must be offered **social services** within the protection system. This may include counselling, psychological support, social assistance and legal advice (Document 3, Art. 1, p. 8 and 13).

4. **Coordination**: Actions must be coordinated between the CPD, the court (on matters of exercising parental rights and visitation), the police (in the event of a crime report) and social service providers.

**Conclusion**: The described situation clearly qualifies the child as a "child at risk" under the Child Protection Act. This requires immediate legal and social intervention to end the violence and ensure the child's well-being. All actions must be based on the legal definitions and procedures provided for in the law.

Analysis Date: Now # Number of analyses: 3

**Figure 7. c.** Structured legal analysis generated by the reasoning model for the child-protection case (c): recommendations for action and conclusion.

### 3.6. Document Corpus and Retrieved Evidence

In addition to the analysis interface, the system exposes the underlying document corpus and the retrieved evidence to the professional user. Figure 8 presents the corpus-management view, in which the sixteen normative acts are displayed in their prioritized order, each annotated with its level of importance. Figure 9 presents the corresponding view of the three supplementary resources, with brief descriptions of their content. This transparency of the corpus is essential in regulated practice, where the legitimacy of any generated analysis depends on the user's ability to verify the documentary basis on which the analysis was produced.

Figures 10a and 10b illustrate the retrieved-evidence view for the child-protection case. After the analysis is generated, the user can expand the search-statistics panel to inspect each of the seven relevant snippets that informed the response, together with the source document from which each snippet was retrieved. The displayed snippets in the example originate primarily from the Child Protection Act, the Regulations for its implementation, the Social Services Act, and the Law on Family Benefits for Children, which corresponds directly to the laws cited in the generated analysis. This direct correspondence between retrieved evidence and generated response operationalizes the bounded-assistance principle at the level of the user interface: the analysis is not a free generation, but a structured synthesis of explicitly retrieved and verifiable passages.

**Uploaded laws and documents**  
Comprehensive legal framework for child and family protection

Search among the laws... All categories

- 1. Social Services Act.pdf** #1  
★ Of paramount importance  
Overview
- 2. Regulations for the implementation of the Social Services Act.pdf** #2  
★ Of paramount importance  
Overview
- 3. Child Protection Act.pdf** #3  
★ Of paramount importance  
Overview
- 4. Regulations for the implementation of the Child Protection Act.pdf** #4  
★ Of paramount importance  
Overview
- 5. Social Assistance Act.pdf** #5  
★ Secondary importance  
Overview
- 6. Regulations for the implementation of the Social Assistance Act.pdf** #6  
★ Secondary importance  
Overview
- 7. Law on Protection from** #7
- 8. Regulations for the** #8
- 9. Law on Family Benefits for** #9

**Figure 8.** Corpus-management view of the assistant, displaying the sixteen Bulgarian normative acts in their prioritized order with importance annotations.

**Resources and practical examples**  
Additional materials with questions, answers and case studies

- EXAMPLES Case-Interventions - AI - school dropout counseling.pdf** Resource #1  
Case study and interventions  
Detailed case study with analysis, interventions and recommendations.  
Overview Info
- Questions and answers - children at risk of dropping out of the education system.pdf** Resource #2  
Questions and answers about dropping out of the education system  
Frequently asked questions and answers about children at risk of dropping out.  
Overview Info
- Questions and answers - related to the social service information and counseling for children without disabilities.pdf** Resource #3  
Questions and answers about social services  
Questions and answers about the social service "Information and Consulting".  
Overview Info

**About resources**

- Educational purpose**  
Provide practical examples and explanations
- Practical application**  
Decision-making assistance and interventions
- Supplement to the laws**  
Enrich legal analysis with real cases

**Figure 9.** Supplementary resources view of the assistant, displaying the three thematic resources that complement the normative corpus.

## Analyzed situation:

“ Mother testifies that father beats child. Parents are divorced. Child lives with mother, but has visits with father. ”

## Search statistics

7 relevant snippets

### View found fragments

#### 4. Regulations for the implementation of the Child Protection Act.pdf Law

the regulations; b) (amended - SG, issue 57 of 2009, in force from 24.07.2009, issue 34 of 2017) monthly allowances for raising a child with a permanent disability under Art. 8e of the Family Allowances for Children Act; c) (repealed, previous item "d" - SG, issue 24 of 2004) additional...

#### 4. Regulations for the implementation of the Child Protection Act.pdf Law

text of the reasons for the requested assistance) Note. The requested type of assistance should be circled - item 1 or item 2. I DECLARE: I. Marital status: Personal data Declarant Husband/wife Name \_\_\_\_\_

#### 9. Law on Family Benefits for Children.pdf Law

of the assistance under para. 1 shall be determined in accordance with the procedure of art. 4a, para. 3 as the total amount of assistance for the family depending on the number of children for whom it is received. (5) (In force from 1.01.2016 - SG, issue 57 of 2015, amended, issue 99 of 2017, in force from 1.01.2018...

#### 9. Law on Family Benefits for Children.pdf Law

#### 9. Law on Family Benefits for Children.pdf Law

benefits that families receive for the child under the Child Protection Act. Art. 8f. (New - SG, issue 88 of 2018, in force from 1.01.2019) (1) Monthly benefits for a child without the right to a survivor's pension from a deceased parent shall be provided to the other ...

#### 1. Social Services Act.pdf Law

In para. 2, a comma shall be placed after the word "patronage" and the words "financed by the municipalities" shall be added. § 22. In the Service Provision Activities Act (promulgated, State Gazette, No. 15 of 2010; amended, No. 83 of 2013, No. 14 of 2015, No. 97 of 2016, No. 21 of 201...

#### 3. Child Protection Act.pdf Law

, in force from 1.07.2020 - amended, issue 101 of 2019); e) for whom there is a risk of dropping out of school or who has dropped out of school. 12. (Previous item 6 - SG, issue 36 of 2003, previous item 7, issue 14 of 2009) "A child with outstanding talents" is a child who ...

#### 3. Child Protection Act.pdf Law

the guardians and the persons with whom the child lives by virtue of another legal basis, to guarantee his rights and protect his interests. 7. (Previous item 1 - SG, issue 36 of 2003, previous item 2, issue 14 of 2009) "Family environment" is biological...

**Figure 10. a.** Retrieved-evidence view of the assistant (a): the top relevant snippets identified through semantic retrieval, with their source documents. **b.** Retrieved-evidence view of the assistant (b): continuation of the retrieved snippets, showing the documentary basis of the generated analysis.

#### 4. Ethical Aspects and Data Protection

The system is designed with implicit data-protection mechanisms and ethical constraints that function as design-level safeguards. Personal data are not collected, since the application does not require registration or the input of identifying attributes. User case descriptions are processed transitively and stored only within a temporary server session with a duration of twenty-four hours, after which they are automatically destroyed. A manual deletion option is also available. Communication with the external language model is limited to what is necessary for response generation, and the data submitted for analysis are not stored in the application's own database.

Regarding the ethical framework, the system is instructed to base its conclusions solely on the provided corpus of Bulgarian normative acts, thereby minimizing the risk of generating speculative or unsupported claims. The application functions as an auxiliary analytical tool rather than as a substitute for professional legal advice; this is reinforced both by the bounded prompting strategy and by the explicit exposure of retrieved evidence, which allows the social worker to verify the documentary basis of every generated response. Despite these built-in safeguards, the absence of an explicit privacy notice and terms-of-use declaration in the front-end remains a limitation that should be addressed in a subsequent development phase to achieve full transparency regarding data processing and the intended purpose of the system. A second area for future development concerns the assessment of potential biases in the retrieved evidence, since the prioritization of the corpus shapes which passages are most likely to be surfaced for any given query; this is a recognized characteristic of retrieval-augmented systems and is addressed in the present work through corpus curation and explicit ordering of importance, but should be subject to systematic empirical study in future iterations.

#### 5. Discussion

The proposed framework contributes to the design of AI systems for regulated social domains in several important ways. First, it formalizes the assistant as a domain-specific support system based on retrieval and bounded generation rather than as a generic conversational interface. Second, it combines legal grounding, interpretation of social context, and structured response generation within a single architecture operating over a controlled corpus of normative documents. Third, it places human validation at the end of the processing chain, which is essential for professional accountability in social services. Fourth, the prototype demonstrates that explicit exposure of retrieved evidence in the user interface is feasible and useful: the social worker can verify the documentary basis of each analysis at the point of use, which strengthens both trustworthiness and explainability without compromising the speed of the analysis.

The three developed prototype versions also reveal an important informatics insight. The effectiveness of the assistant depends not only on the language model itself, but on the coordination between corpus preparation, text segmentation, indexing, retrieval quality, and prompt design. The transition from broader early corpora to a more compact and prioritized documentary set shows that increasing documentary volume does not automatically improve the quality of the analysis. On the contrary, when the corpus becomes broader, retrieval precision and document ranking become decisive for answer quality, because peripheral or only weakly related documents introduce competing chunks that displace more relevant passages in the top-k results. The evolution of the reasoning model from a locally hosted Llama 3 instance to an API-accessed DeepSeek model further illustrates that pragmatic considerations—robustness, multilingual handling of Bulgarian legal terminology, and the quality of structured generation under explicit prompt constraints—are integral to the maturity of such prototypes and may justify departures from the initial technological setup.

The added value of the present work lies not only in the conceptual description of the model, but also in the documented prototype evolution, the prototype workflow, and the experimental case study with explicit retrieved evidence. In this sense, the paper presents not only an abstract architectural idea, but a prototype-informed model of a legal-social AI assistant operating over a concrete set of Bulgarian normative sources. This makes the framework suitable both as a standalone contribution and as a basis for future integration into broader digital platforms for the management and support of social services. In particular, the present assistant is being developed as a controlled support component within a larger web-based platform for social-service case management within the SocialITAI research project of the DUEkos programme at the University of Plovdiv “Paisii Hilendarski”, which provides the broader research context for its later embedded use.

The present study has several limitations that should be acknowledged. The contribution is primarily architectural and prototype-oriented; it does not yet include a formal empirical evaluation with end users, a large-scale benchmark of retrieval precision against a labeled ground-truth set, or a controlled comparative experiment against alternative reasoning models. The illustrative use case demonstrates that the system produces a coherent and well-structured response for a typical child-protection scenario, but does not constitute a statistical evaluation of the system’s behavior across diverse case profiles. Future work should therefore include scenario-based validation with practising social workers, retrieval evaluation against curated query sets, comparative analysis of alternative embedding models and reasoning models, and a more formal investigation of explainability, bias, and privacy in social-service contexts. A further line of development concerns the embedding of the assistant inside the broader case-management platform, where it can be invoked directly from the case context and operate over a combined corpus of normative documents and case-specific data.

## 6. Conclusions

This paper proposed and demonstrated a conceptual model of a Personal Legal and Social AI Assistant for decision support in social services. The model combines a controlled normative corpus, semantic retrieval through SentenceTransformer embeddings and cosine similarity, bounded response generation through a hosted reasoning model, structured output organized into four interpretable sections, and human validation in a single framework intended for professional support in a legally sensitive and socially specific environment. In the current prototype version, the system operates on a prioritized corpus of sixteen Bulgarian normative acts and three supplementary resources, comprising 883 indexed fragments, and uses DeepSeek as the reasoning model accessed through the OpenRouter API.

The main scientific contribution of the study lies in the formalization and prototype-based demonstration of a domain-specific AI assistant that combines legal grounding, social context, and controlled language-model reasoning. The added value of the work lies not only in the conceptual description of the architecture, but also in the traceable evolution of the prototype across three versions, the prototype workflow, the corpus-management interface, the retrieved-evidence view, and the illustrative case in the domain of child protection. In this form, the proposed model can be regarded both as a standalone contribution and as a basis for future integration into broader digital environments for the management and support of social services. Future work will focus on scenario-based validation with practising specialists, on retrieval evaluation against curated query sets, on comparative analysis of alternative reasoning models, on the introduction of an explicit privacy notice in the front-end, and on the embedding of the assistant inside the case-management platform developed within the SocialITAI research project (D24-PF-016).

**Funding:** This research is related to the work on project No. D24-PF-016 “Formation and Development of Digital Competencies and Use of Artificial Intelligence to Improve the Quality of Social Work (SocialITAI)” as part of the Strategic Research and Innovative Development Program—Digital Sustainable Ecosystems—Technological Solutions and Social Models for Ecosystem Sustainability (DUEkos) of the University of Plovdiv “Paisii Hilendarski”. The article processing charge was funded by the same project.

**Institutional Review Board Statement:** Not applicable. The study did not involve research on human subjects; the prototype assistant processes only user-submitted free-text case descriptions in a transient session-based manner without collecting personal data.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The Bulgarian normative acts used as the legal corpus of the assistant are publicly available through the official State Gazette and the consolidated legal databases maintained by the Republic of Bulgaria. The supplementary resources, prompt templates, and indexed embeddings used in the current prototype version are available from the author upon reasonable request, subject to the data-handling constraints described in Section 4.

**Acknowledgments:** The author thanks Prof. Dr. Lyudmila Vekova, principal investigator of the SocialITAI project, and the colleagues from the project research team for the supportive scientific environment in which this work was carried out, as well as the participating social-service organizations for their feedback during prototype testing. During the preparation of this manuscript, the author used a large language model assistant for purposes of language editing and consistency checking; the author has reviewed and edited the output and takes full responsibility for the content of this publication.

**Conflicts of Interest:** The author declares no conflicts of interest.

## References

1. Eurofound. Social Services in Europe: Adapting to a New Reality; Publications Office of the European Union: Luxembourg, 2023; ISBN 978-92-897-2339-8.
2. European Commission. Digital Economy and Society Index (DESI) 2022: Bulgaria; European Commission: Brussels, Belgium, 2022.
3. Vekova, L.; Vazova, T.; Georgiev, P.; Uzhikanova-Kovacheva, E. An Innovative Model for Developing Digital Competences of Social Workers. *Strategies for Policy in Science and Education* 2025, 33, 1s, 60–77. <https://doi.org/10.53656/str2025-1s-3-inn>.
4. Council of Ministers of the Republic of Bulgaria. National Program “Digital Bulgaria 2025”; Council of Ministers of the Republic of Bulgaria: Sofia, Bulgaria, 2019.
5. Vekova, L.; Uzhikanova-Kovacheva, E.; Georgiev, P.; Vazova, T. Information Systems and Technologies with the Implementation of Artificial Intelligence as an Interdisciplinary Approach for the Development of Digital Competencies and Training of Social Workers. In *Proceedings of the Jubilee International Scientific Conference “The Interdisciplinary Approach in Economic and Social Sciences” (ICISS 2025)*; ARPHA Proceedings 2025, 10, 545–557. <https://doi.org/10.3897/ap.10.e0545>.
6. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; et al. On the Opportunities and Risks of Foundation Models. *arXiv* 2021, arXiv:2108.07258.
7. Eubanks, V. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*; St. Martin’s Press: New York, NY, USA, 2018.
8. Saxena, D.; Kaur, H. Artificial Intelligence in Social Work: A Systematic Review. *British Journal of Social Work* 2022, 52, 1–20. <https://doi.org/10.1093/bjsw/bcab159>.
9. Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-T.; Rocktäschel, T.; et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *Advances in Neural Information Processing Systems* 2020, 33, 9459–9474.
10. Schoech, D. Using Computers in Social Work Practice. *Social Work* 1982, 27, 178–182.
11. Gillingham, P. Electronic Case Management Systems: The Impact on Social Work Practice. *British Journal of Social Work* 2015, 45, 1859–1875. <https://doi.org/10.1093/bjsw/bcu018>.
12. Chalkidis, I.; Fergadiotis, M.; Malakasiotis, P.; Aletras, N.; Androutsopoulos, I. LEGAL-BERT: The Muppets Straight Out of Law School. In *Findings of EMNLP 2020*; Association for Computational Linguistics: Online, 2020; pp. 2898–2904.

13. Niklaus, J.; Chalkidis, I.; Stürmer, M. Swiss-Judgment-Prediction: A Multilingual Legal Judgment Prediction Benchmark. In Proceedings of the Natural Legal Language Processing Workshop; Punta Cana, Dominican Republic, 2021; pp. 1–8.
14. Stoyanov, S.; Orozova, D. Intelligent Personal Assistants in a Virtual Learning Space. In Proceedings of the Annual Conference of the Technical University of Sofia, Plovdiv Branch; Plovdiv, Bulgaria, 2014.
15. Todorov, J.; Stoyanov, S.; Valkanov, V.; Daskalov, B.; Popchev, I. Learning Intelligent System for Student Assistance—LISSA. In 2016 IEEE 8th International Conference on Intelligent Systems (IS 2016); 2016; pp. 753–757. <https://doi.org/10.1109/IS.2016.7737397>.
16. Todorov, J.; Valkanov, V.; Stoyanov, S.; Daskalov, B.; Popchev, I.; Orozova, D. Personal Assistants in a Virtual Education Space. In Practical Issues of Intelligent Innovations; Studies in Systems, Decision and Control, Vol. 140; Springer: Cham, Switzerland, 2018; pp. 131–153. [https://doi.org/10.1007/978-3-319-78437-3\\_6](https://doi.org/10.1007/978-3-319-78437-3_6).
17. Valkanov, V.; Todorov, J.; Daskalov, B.; Stancheva, N.; Stoyanov, S. Personal Assistant with Ontology Based Knowledge. In Proceedings of the International Conference Big Data, Knowledge and Control Systems Engineering; Sofia, Bulgaria, 2017.
18. Tabakova-Komsalova, V.; Stoyanov, S.; Doukovska, L.; Stoyanov, I.; Cherecharov, S. Personal Assistant Supporting Diagnosis of Livestock Poisoning. In 2022 International Conference Automatics and Informatics (ICAI); 2022; pp. 189–192. <https://doi.org/10.1109/ICAI55857.2022.9960100>.
19. Reimers, N.; Gurevych, I. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing; Association for Computational Linguistics: Hong Kong, China, 2019; pp. 3982–3992. <https://doi.org/10.18653/v1/D19-1410>.
20. DeepSeek-AI. DeepSeek-V3 Technical Report. arXiv 2024, arXiv:2412.19437.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.