



City Research Online

City, University of London Institutional Repository

Citation: Butcher, B., Zilka, M., Hron, J., Cook, D. & Weller, A. (2024). Optimising Human-Machine Collaboration for Efficient High-Precision Information Extraction from Text Documents. *ACM Journal on Responsible Computing*, 1(2), pp. 1-27. doi: 10.1145/3652591

This is the published version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/33255/>

Link to published version: <https://doi.org/10.1145/3652591>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk



Optimising Human–Machine Collaboration for Efficient High-Precision Information Extraction from Text Documents

BRADLEY BUTCHER, University of Sussex and The Alan Turing Institute, Brighton, UK

MIRI ZILKA, University of Cambridge, Cambridge, UK

JIRI HRON, Google DeepMind, London, UK

DARREN COOK, Imperial College London, London, UK

ADRIAN WELLER, University of Cambridge and The Alan Turing Institute, Cambridge, UK

From science to law enforcement, many research questions are answerable only by poring over a large amount of unstructured text documents. While people can extract information from such documents with high accuracy, this is often too time-consuming to be practical. On the other hand, automated approaches produce nearly-immediate results, but are not reliable enough for applications where near-perfect precision is essential. Motivated by two use cases from criminal justice, we consider the benefits and drawbacks of various human-only, human–machine, and machine-only approaches. Finding no tool well suited for our use cases, we develop a human-in-the-loop method for fast but accurate extraction of structured data from unstructured text. The tool is based on automated extraction followed by human validation, and is particularly useful in cases where purely manual extraction is not practical. Testing on three criminal justice datasets, we find that the combination of the computer speed and human understanding yields precision comparable to manual annotation while requiring only a fraction of time, and significantly outperforms the precision of all fully automated baselines.

CCS Concepts: • **Information systems** → **Information extraction**; • **Human-centered computing** → **Interactive systems and tools**; **HCI design and evaluation methods**;

Additional Key Words and Phrases: Human-computer collaboration, human-in-the-loop, information extraction, weak supervision

ACM Reference Format:

Bradley Butcher, Miri Zilka, Jiri Hron, Darren Cook, and Adrian Weller. 2024. Optimising Human–Machine Collaboration for Efficient High-Precision Information Extraction from Text Documents. *ACM J. Responsib. Comput.* 1, 2, Article 16 (June 2024), 27 pages. <https://doi.org/10.1145/3652591>

B. Butcher and M. Zilka are contributed equally to this research.

J. Hron are part of the work done while a PhD student at the University of Cambridge, UK..

BB acknowledges support from the European Research Council (ERC), grant agreement No 851538. MZ acknowledges support from EPSRC grant EP/V025279/1, The Alan Turing Institute, and the Leverhulme Trust grant ECF-2021-429. AW acknowledges support from a Turing AI fellowship under grant EP/V025279/1 and the Leverhulme Trust via the Centre for the Future of Intelligence.

Authors' addresses: B. Butcher, University of Sussex and The Alan Turing Institute, Brighton, UK; e-mail: b.butcher@sussex.ac.uk; M. Zilka, University of Cambridge, Cambridge, UK; e-mail: mz477@cam.ac.uk; J. Hron, Google DeepMind, London, UK; e-mail: jirihron@google.com; D. Cook, University of Liverpool, Liverpool, UK; e-mail: hldcook@liverpool.ac.uk; A. Weller, University of Cambridge and The Alan Turing Institute, Cambridge, UK; e-mail: aw665@cam.ac.uk.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2024 Copyright held by the owner/author(s).

ACM 2832-0565/2024/06-ART16

<https://doi.org/10.1145/3652591>

1 INTRODUCTION

High precision extraction of structured data from unstructured text has valuable applications in medicine [72], finance [22], science [50, 54, 73], and law enforcement [12, 35, 71]. In this article, we explore **information extraction (IE)** use cases in criminal justice. This area is of particular interest because the existing data sources—whether publicly available [3, 15, 55, 65] or restricted [39]—are usually derived from routinely collected administrative data, which often omits key information like victim details, or mitigating and aggravating circumstances [81]. This information, and much more, can however be found in long unstructured documents like trial transcripts. Making such information easily accessible would, among else, enable much needed research into bias in the judicial system.

There are several ways to achieve this. While humans can often extract information from unstructured documents with near-perfect accuracy, their work is time-intensive, costly, and possibly traumatizing (e.g., when the content is violent or sexual) [19]. On the other hand, even state-of-the-art automated methods still cannot extract information from text with human-level accuracy [30, 79]. The main focus of our article are therefore human–computer collaborative approaches to IE with emphasis on the *tradeoff between processing speed and extraction accuracy*.

Our study is centered around three distinct datasets and two IE tasks. Two of the datasets contain information about criminal cases from legal documents and news articles. For both, the task is to extract structured (tabular) information about the victims, defendants, and situational circumstances. The third dataset contains online communications between convicted child sexual predators, and police officers posing as minors; the task is to identify specific predatory behavior patterns [18]. In all cases, the text documents are too long for *manual* extraction at scale [37], and the high precision necessary for the extracted data to be useful to researchers rules out fully *automated* approaches. Unlike in most of the existing literature on human–machine collaboration [e.g., 5, 62], all our tasks are examples of the under-explored scenario where computers achieve lower accuracy, but can significantly speed up processing.

In Section 2, we discuss how this type of human–computer complementarity affects design choices within collaborative IE, and identify a gap among existing solutions (see Table 1). In Section 3, we fill this gap by developing ELICIT, a new human-validated IE tool which leverages the speed of modern machine learning models, and the near-perfect accuracy of humans. ELICIT is built on weak supervision approaches [63, 80], utilizing multiple independent algorithms to provide suggestions from which the user selects the final answer. In Section 4, we perform a case-study comparing manual extraction and ELICIT, with two of this paper’s authors serving as the human validators. ELICIT achieves precision on-par with manual annotation, significantly outperforming the state-of-the-art among *automated* IE tools [63, 80], while using orders of magnitude less processing time (Section 4.2).¹ In Section 4.3, we further explore: (i) increasing the number of computer-generated suggestions to improve recall; (ii) deferral as a way to tradeoff performance for time; and (iii) ranking and fine-tuning as auxiliary ways to use the validated data to attain even better performance.

Our main contributions are summarized below:

- (1) We analyze the relative strengths of various human–computer approaches to IE from long unstructured texts, in the context of two criminal justice based use cases.
- (2) We design and implement ELICIT, an interactive IE tool which combines weak supervision and human validation to enable IE that is faster yet almost as accurate as manual annotation.

¹While having authors of this article perform the human validation confounds the results, the speed up is larger than could be made up for even much faster reading manual annotators.

Table 1. Overview of Human-Computer Configurations Applicable to IE (Section 2.2), and the Degree to which they Meet the Requirements for Our use cases as Identified in Section 2.1.3

	Flexibility	Precision	Recall	Reproducibility	Time-efficiency	Shielding
Manual only (Section 2.2.1)	high	high	high	depends	low	low
Manual with search (Section 2.2.1)	medium	high	—	tradeoff —	↓	↓
Assisted annotation (Section 2.2.2)	medium	high	medium	depends	↓	↓
Human validation (Section 2.2.3)	depends	high	medium	medium	↓	↓
Deferring to a human (Section 2.2.4)	medium	↑	medium	depends	↓	↓
Validation in training (Section 2.2.5)	depends	↑	medium	medium	↓	↓
Full automation (Section 2.2.6)	medium	medium	medium	high	high	high

Arrows show progressive increase (e.g., the “low→high” in the “Time-efficiency” column means the average time needed to complete annotation is decreasing with every row).

- (3) Creating three new datasets, we demonstrate ELICIT significantly speeds up the annotation process, while maintaining near-perfect precision. Accuracy can be further traded-off for processing time via a deferral setting.

While we focus on criminal justice use cases, we believe ELICIT can be useful in other areas where the imperfect accuracy of automated methods, and the slowness and other issues of manual annotation, pose a problem (e.g., legal texts, medical records, meta-analysis of literature).

2 REQUIREMENTS OF HIGH-PRECISION INFORMATION EXTRACTION

IE and data labeling are used in a variety of settings, each with a distinct set of challenges. Not all settings are similarly suitable for human-computer collaboration [37], and practitioners often resist solutions which over focus on automation instead of providing assistance [26]. To ground our investigation of human-computer collaboration in IE, we start by exploring two use cases related to the criminal justice domain (Section 2.1). Centering the discussion around use cases allows us to examine how attributes of different tasks affect system desiderata. This enables us to evaluate and compare available solutions, and identify existing gaps in the literature and tool landscape (Section 2.2).

2.1 Criminal Justice Use Cases

2.1.1 Task 1: Criminal Trial Information. Research in criminal justice is severely hindered by lack of high-quality publicly available datasets [23, 43, 45, 48]. A great deal of useful information is contained in unstructured documents such as court transcripts, sentencing remarks, judgments, and press articles. Extracting this information into a structured format can enable research into critical issues like racial bias in criminal justice [16, 41]. The documents are however too numerous and lengthy for manual extraction. This creates a need for a *time-efficient* yet *accurate* IE tool. To enable high-quality extraction across a wide variety of settings, the system should also be *flexible*, and able to take advantage of efficient IE methods for a given task when they exist.

To ensure *reproducibility* of any analysis facilitated by the tool, it should be feasible to document the dataset creation steps in a way that allows independent reconstruction. This is a particular concern under ambiguity. To illustrate, consider extracting whether a victim or suspect was legally considered vulnerable from sentencing remarks.² Such information is seldom mentioned explicitly—unless directly relevant to the ruling—but can sometimes be inferred. While humans are skilled at making such inferences, they often need to read a large portion of the text for context, and may disagree with each other’s judgments. These disagreements, and a lack of explicit

²Sentencing remarks summarize the judge’s ruling in a criminal trial. They typically describe the crime, and any mitigating or aggravating circumstances.

rules for resolving them, are a core problem in reproducibility [53]. When adding any human element, it is thus key to establish guidelines around how potentially ambiguous data should be recorded.

2.1.2 Task 2: Online Child Sexual Exploitation Discourse. Online sexual grooming of minors is an increasing problem in the digital age [28]. Many abusers seek to establish physical contact offline [67]. To assist the law enforcement, academics have been trying to in advance identifying predators who steer the relationship toward physical encounters [10, 57, 74, 75]. To date, this important work has relied on time-consuming manual annotation of conversations involving child sex offenders. While full automation is possible, it currently comes at a cost of low accuracy [18].

Our second task is identifying signs of offline contact solicitation in online chats. The goal is to annotate conversations in accordance with the “self-regulation” theory of child grooming [24], which has already been employed in Reference [18]. In practice, this means detecting whether each conversational instance (i.e., a continuous message exchange without more than a one-hour break) contains any of the following *offender behaviors*: (1) rapport building, (2) control, (3) challenges, (4) negotiation, (5) use of emotions, (6) testing boundaries, (7) use of sexual topics, (8) mitigation, (9) encouragement, and (10) risk management. See Reference [18] for a qualitative description of these behaviors, and further discussion.

When done manually, an expert annotator scans the offender’s messages one-by-one, deciding for each whether one or more of the above ten behaviors is present. For reference, performing manual annotation on 24 chats took a forensic psychologist over 600 hours [18]. Designing a solution with high *time-efficiency* and *accuracy* is therefore crucial. Reducing the user’s exposure to this difficult content is a significant additional benefit [68]. For optimal results, the system should again be *flexible* enough to enable incorporation of already existing IE solutions. Finally, due to the more subjective nature, *reproducibility*, or the ability easily compare disagreement in annotations is even more pertinent here than for Task 1 (Section 2.1.1). In particular, since annotators can substantially disagree, we want the data creation process to be reproducible at the level of these differences. We can then monitor and see if there is significant inter-annotator disagreement. Since we require the tool to be time-efficient, using multiple human annotators is significantly more feasible than with fully manual annotation.

2.1.3 Comparing the Tasks. General design considerations of human-in-the-loop IE tools based on existing workflows have been studied in Reference [60]. However, the specificity of our use cases (Sections 2.1.1 and 2.1.2) requires that we explore the design space in context of our goals [46]. Inspecting the previous sections, we identified several **common requirements**:

- **Time-efficiency**, i.e., orders of magnitude faster than manual extraction.
- **Accuracy**, i.e., the results contain little incorrect information (*precision*), and are as complete as possible (*recall*).
- **Flexibility**, i.e., can extract variety of user-specified information, and incorporate existing IE tools.
- **Reproducibility**, i.e., the information extraction process should be replicable with appropriate documentation, and inter-annotator disagreement should be trackable.
- **Shielding**, i.e., reducing the user’s exposure to the text by only highlighting relevant sections.

Time-efficiency is essential due to the large quantities of text. Moreover, in high-stakes domains such as criminal justice, accuracy of the extracted information is paramount, and must not be sacrificed for processing speed.

2.2 Existing Tools and their Drawbacks

We now discuss various approaches for addressing the two tasks from Sections 2.1.1 and 2.1.2 with focus on their key requirements (see Table 1 for an overview).

2.2.1 Manual Extraction and Search. Humans with appropriate background are the gold standard for both tasks (given sufficient time and motivation). However, ensuring high recall requires the reader to scan through the whole text, which is prohibitively slow for our use cases [37]. Besides fatigue, both our tasks (Section 2.1) expose the reader to potentially disturbing content which may impact their well-being [4]. Humans can also take advantage of contextual knowledge, and make flexible on-the-fly decisions; while often advantageous, this flexibility may hinder reproducibility when no fixed protocol is followed, or when the task requires a level of subjective judgment (as in Task 2).

Human labeling can be sped up by making the source documents searchable. Beyond digitisation, this may take the form of keyword search, or regular expression matching. An example of a relevant tool is OpenSearch [66], an open-source fork of the more well-known Elasticsearch [52]. While this preserves flexibility and accuracy, the boost to time-efficiency is often small, especially when assigning a correct label requires understanding large portions of the text (as in Task 2). Reproducibility concerns also remain, unless a very specific set of rules is detailed and followed. If such rules exist, an approach with a greater level of automation may produce similar performance more efficiently.

2.2.2 Assisted Annotation. Assisted annotation tools are commonly used across many domains, including law, medicine, and political science [32, 51, 69, 78]. Their aim is to produce annotated text, with labels assigned to each annotation. The tools are most often used in an iterative process: the algorithm proposes annotations, the user makes modifications to correct mistakes, the tool uses this feedback to learn better recommendations, and so on. For example, [20] uses a semi-supervised model to predict the most probable labels for each instance. They show this increases speed and accuracy of data labeling. Some other recent examples of these tools are: *prodi.gy*, *lighttag*, and *CLIEL* [27, 49, 58]. Since the user has full control over the final annotation, accuracy is ensured provided there is sufficient time and motivation. The algorithmic proposals can provide a speed-up, but it is limited in cases where the user is still required to read or scan through large chunks of the document. For the same reason, the reduction of user fatigue, and exposure to disturbing material, is rather limited. Analogously to Section 2.2.1, reproducibility can be difficult to achieve.

2.2.3 Human Validation. By human validation, we mean methods where automated algorithms make all label predictions, *each* of which is then reviewed by a human. When combined with *passage retrieval* [e.g., 40], this setup exploits the complementary human and computer abilities (accuracy and speed). Beyond time-efficiency and accuracy, human validation also improves reproducibility, at least at the level of the machine predictions. The level inter-annotator agreement can be measured by using multiple annotators and improved, if needed, by providing clear guidelines. The only open-source tool from this category we found is *Rubrix* [17]. *Rubrix* satisfies many of our desiderata. However, its main purpose is labeling many short texts, rather than extracting a set of interdependent variables from larger documents.

2.2.4 Deferring to a Human. Deferral is an *extension* of human validation (Section 2.2.3) where the human reviews *only some* of the predictions. This allows trading-off accuracy for time-efficiency, and alleviates user fatigue and exposure to harmful material. The cases to defer are typically chosen using an estimate of prediction confidence. When these estimates are well-calibrated, significant time savings may be attained with little performance loss. Deferral can be particularly

useful for large amounts of relatively low-stake decisions, e.g., moderating social media comments [42]. However, calibrating confidence estimates remains a challenge, especially in deep learning [31]. Without calibration, the potential for improvement is often marginal. Deferral style solutions can be obtained by combining any Human validation (Section 2.2.3) and automated labeling (Sections 2.2.5, 2.2.6) algorithm, provided the latter outputs confidence estimates.

2.2.5 Validation in Training. Validation using human experts can be costly and time-consuming. An alternative to deferring to humans during labeling is to only use their feedback for model training. We specifically refer to a form of active learning where the user is only asked to label several strategically selected examples at the beginning, after which the fine-tuned model labels the rest of the dataset in a fully automatic mode. An example of a tool from this class is IWS [9]. Compared with human validation (Section 2.2.3) and deferral (Section 2.2.4), this method enables further time-savings at the price of performance reduction. While the impact on reproducibility and user fatigue and well-being is positive, the drop in performance is often too large to satisfy the requirement of near-perfect precision (Section 2.1.3).

2.2.6 Full Automation. Automated IE is an area of active research [1, 30]. Algorithms in this category do not defer any of their predictions to humans. This provides the best time-efficiency and reproducibility. However, even state-of-the-art algorithms [30, 63, 64, 76, 79] cannot achieve the level of performance our use cases require.

3 ELICIT: A SYSTEM FOR USER-VALIDATED INFORMATION EXTRACTION FROM TEXT

Comparing existing tools (Section 2.2) to our requirements (Section 2.1.3) reveals a gap. Manual and assisted annotation methods (Sections 2.2.1 and 2.2.2) are time-intensive and provide little shielding, while approaches that do not validate the final extractions (Sections 2.2.5 and 2.2.6) tend to compromise accuracy. Human-in-the-loop systems (Sections 2.2.3 and 2.2.4) offer a balance, but the only existing open-source tool, Rubrix [17], is not optimized for use-cases which necessitate extracting multiple interdependent variables from lengthy documents for extraction of interdependent variables from long documents common to our use-cases (Section 2.1). Therefore, a new system is required.

3.1 System Design

Based on Section 2.1.3 requirements, we designed and implemented ELICIT, a flexible, accurate, and time-efficient tool for IE from text. Its core functionality falls into the *human validation* category (Section 2.2.3), with specialization on extraction of interdependent variables from long complex texts. ELICIT consists of two high-level stages (Figure 1): (1) automated passage retrieval and label suggestion, and (2) human label validation. The prediction *accuracy* relies on the user’s ability to infer the correct label from the retrieved passages, while the *time-efficiency* gains come mainly from the automated passage retrieval. Since the user only interacts with text excerpts, they are *shielded* from having to engage with the entirety of the potentially disturbing text. The automated passage retrieval and label suggestions also allows for step (1) to be *reproducible*, and step (2) easily comparable between annotators.

To ensure *flexibility*, the automated first step utilizes weak supervision—a method for combining multiple “weak labels” generated by *labeling functions* [63]. For intuition, weak supervision is analogous to assigning labels by combining answers from multiple annotators, where each annotator provides a label and a corresponding confidence level. The final label is then a weighted average of the predicted labels. The weights are proportional to (re)calibrated confidence, where the calibration is used to adjust for miscalibrated annotators. In ELICIT, we use different

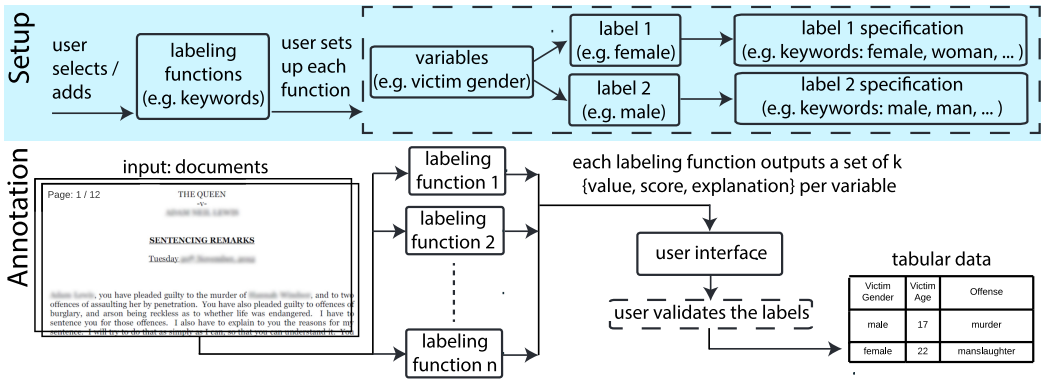


Fig. 1. A high-level diagram of the ELICIT framework. *Setup*: User chooses the labeling functions, and the information (variables) to be extracted; the labeling functions (1) identify relevant sections of the text; and (2) suggest a label for the user to validate. *Annotation*: For each variable, labeling functions provide possible values and accompanying explanations. The user validates these, creating a tabular dataset where each row corresponds to a document, and each column to an extracted variable.

automated labeling methods as “annotators”. Instead of a single final label, we produce a list ranked by average calibrated confidence, for the human to validate (see Appendix G for details). Beyond its flexibility, a key advantage of weak supervision for our use cases is that it can enable high overall accuracy even if the labeling functions are not individually performant.

3.2 Core Features and User Interface

Leveraging large language models. In ELICIT, the user defines a set of questions (e.g., “Was the victim vulnerable?”), and an ensemble of *labeling functions* (ranging from keyword lookup to neural nets; see Setup in Figure 1). Each labeling function is tasked with: (a) identifying a part of the text relevant to the question; and (b) suggesting an answer label for the user to validate (e.g., “Victim was vulnerable”). The most successful labeling functions we tested utilize large language models [11, 21, 33] to achieve one or both of the above tasks. Large language models can achieve impressive results on passage retrieval and information extraction tasks [2, 34], although not without limitations and biases [7].

Providing explanations. To enable human verification, label predictions must be accompanied by *explanations*. For the tasks we consider (Section 2.1), explanations are equivalent to a relevant snippet of the original text. If the provided snippet is insufficient, the user can open a pop-up window containing a larger section of the text. Presenting only relevant snippets is not only faster than manual reading, but also shelters the user from sensitive, graphic, and otherwise problematic content. This reduces both harm to the user, and their mental fatigue.

User interface. The user validates candidate labels—e.g., “Victim is female.”—within the user interface (Figure 2). Too many candidates may overwhelm the user. To streamline use and reduce fatigue, we *merge predictions* of the same label if their explanations (snippets) significantly overlap. In the Figure 2 example, the user is asked to validate the victim sex. Rows correspond to suggested label values. The extended snippets for “female” are unrolled. You can see that the explanation with the highest confidence was highlighted by 3-out-of-5 labeling functions.

Note that each label value (e.g., male, female) can be supported by multiple explanations (e.g., “She was described ...”). These are ordered in the UI by a *ranking* function, with the highest confidence explanations placed on the left. Our ranking model is similar to Reference [64], except we




Fig. 2. User interface screenshot. In this example, the user is asked to validate the victim’s sex. The user is presented with two snippets of the text as explanations. The reported level of agreement corresponds to the maximal LF agreement for a single explanation for each label. The full explanations are presented to the user in a pop-up window, following a mouse click on the box.

compute scores on a per-explanation rather than per-document level (see Figure E1 and Appendix G.2 for details). See supplementary materials for a video demonstration of ELICIT’s user interface.

Top- K validation. To improve recall, labeling functions can nominate up to K candidate tuples—(label value, confidence score, explanation)—for each label, instead of one per document (see Annotation in Figure 1). We refer to this as *top- K validation*. Increasing K improves recall, but can increase burden on the user.

3.3 Advanced Features

Continual adaptation of candidate ranking. As the user validates the extracted information, a new tabular dataset is created. We can use this newly created data to calibrate the confidence output by the labeling functions, and continually *adapt the ranking function to user needs* as in [8] (see Appendix G.2 for details). This will improve the chance that the correct answer will be presented to the user first, making the validation even faster.

Fine-tuning labeling functions. Beyond adapting the ranking, the data provided by user feedback can be used to *continually improve (fine-tune) the labeling functions*. This can be especially useful for improving recall among the automatically generated candidates. If a labeling function is updated and a new candidate is found for an already labeled document, ELICIT’s user interface alerts the user to this fact.

Deferral. ELICIT can be combined with *deferral* (Section 2.2.4), where only candidates with low assigned confidence are validated by the user. This allows further tradeoff between time and accuracy, and will work well only if the confidence scores are well calibrated. In Section 4, we assume the automated tool comes with its own confidence score, which is used to decide whether to defer to the user via ELICIT. We note that the automated extraction algorithm need *not* be one of the ELICIT labeling functions. This allows specialization, i.e., fine-tuning one system for automated extraction (e.g., SNORKEL, [63]), and ELICIT for assisting the user (calibrated scores and explanations, high recall, and so on).

4 CASE STUDY: APPLYING ELICIT TO CRIMINAL JUSTICE USE CASES

We describe the datasets (Section 4.1.1), the information to be extracted (Section 4.1.2), and ELICIT’s labeling functions (Section 4.1.3). Evaluation of ELICIT’s core features is presented in

Section 4.2, and of its advanced features in Section 4.3. ELICIT and the evaluation code will both be made available on GitHub.

4.1 Setup

4.1.1 Data Sources. For Task 1 (Section 2.1.1), we use two self-compiled datasets. Both contain information on criminal trials for the offence of murder in the United Kingdom. For Task 2 (Section 2.1.2), we utilize an existing dataset.

Crown Court Sentencing Remarks. *Sentencing remarks* are a transcript of judge’s remarks delivered when announcing a sentence. They usually include a summary of the offence, and sentence justification (incl. mitigating and aggravating circumstances). The United Kingdom Judiciary publishes sentencing remarks for cases within the realm of public interest [38]. There are 343 published sentencing remarks, covering a range of offences, including murder, manslaughter, and misconduct in a public office. The vast majority are murder cases; we filter out other offence in the dataset. To facilitate evaluation, we manually extracted 18 variables from 20 sentencing remarks, ranging from 500 to 14,000 words in length. In Section 4.2, we only report results for variables which occurred in the 20 remarks 5 or more times.

News articles. *Law Pages* is a legal resource website which allows the general public to search for sentencing information, filtering for certain types of offences [25]. The website also links to news articles relating to each case. Using the filtering tool, we constructed a dataset of sentencing metadata for murder cases tried in the Crown Court, linked to unstructured text from corresponding news articles. We only include articles from the five most common outlets (see Figure B1 in the appendix). For validation, we manually labeled 20 cases using the relevant subset of variables from the sentencing remarks (10 out of the original 18 variables; see Table B1 in the appendix).

Perverted Justice. For Task 2 (Section 2.1.2), we use data from the Perverted Justice website [59], which contains real-world chat-based conversations between adults later convicted of grooming offences, and adult decoys posing as children. We use the 24 annotated chats with 10 variables mentioned in Section 2.1.2, which originate from Reference [18]. 5 of the 24 chats are used for training and fine-tuning, leaving 19 for evaluation in Section 4.2.

4.1.2 Information Targeted for Extraction. For each task (Section 2.1), we extract a different set of variables.

Task 1 (Section 2.1.1): Information we extract about the *victim*: race, sex, religion, sexual-orientation, employment status, pregnancy, disability, and whether the victim was considered vulnerable, or suffered physical, mental or domestic abuse at the hands of the defendant. Information we extract about the *defendant*:³ prior convictions, relationship to the victim, sexual or racial motivation, offence premeditation, remorse, and whether age was a mitigating factor.

Task 2 (Section 2.1.2): For each conversation, we extract the ten indicator variables described in Section 2.1.2.

4.1.3 ELICIT’s Labeling Functions. We employ distinct LFs for each of our use cases (Section 2.1). The *same* LFs are used in the SNORKEL baseline [64], which employs an automated algorithm instead of a human for label validation.⁴

³The defendant’s demographics are not extracted as these are usually collected as admin data.

⁴Unlike in **Named-entity recognition (NER)**, we are not just assigning labels to words or phrases. This is best understood in the context of the Perverted Justice dataset, where we look for evidence of specific predatory behaviors in the conversation. These depend as much on tone and the context of the conversation, as on particular words or phrases being used. This is the main reason why we choose Q&A algorithms instead of NER.

Task 1 (Section 2.1.1): A short description of the five labeling functions we use is below; see Appendix F for details.

- LF1 **Transformer Q&A → Zero-shot Sequence Classifier.** We use RoBERTa fine-tuned for question-answering on the Squad2 dataset [44, 61]. The question-answering capability is then used to retrieve relevant text excerpts by associating each variable of interest (e.g., victim sex) with one or more questions (e.g., “What sex was the victim?”). The model outputs excerpts with associated probabilities $P_{QA}(\text{excerpt} \mid \text{question})$. We then use a RoBERTa **Natural Language Inference (NLI)** model [77] to assign labels $P_{NLI}(\text{label} \mid \text{excerpt})$. The score we use to select the top- K candidates is the product $P_{NLI}(\text{label} \mid \text{excerpt}) \cdot P_{QA}(\text{excerpt} \mid \text{question})$.
- LF2 **Transformer Q&A → Similarity.** The relevant sections are extracted as in the first step above. In the second step, we replace the NLI model by an alternative scoring rule: the cosine similarity S_{cos} between RoBERTa embeddings of the excerpts and the labels. The top- K then maximize $S_{\text{cos}}(\text{excerpt}, \text{label}) \cdot P_{QA}(\text{excerpt} \mid \text{question})$.
- LF3 **Sentence-level similarity.** We use a transformer to embed every sentence and label. Cosine distance is then again used to identify semantic similarities. Sentences with similarity above a user-defined threshold are extracted. The top- K sentences with similarity above a user-defined threshold are taken as candidates.
- LF4 **Keyword Search.** Each label is associated with a set of user-specified keywords. For example, for “victim sex”, the label value “male” could be associated with keywords like “man”, “male”, “Mr.”, and so on. In this case, top- K does not apply due to no clear way of assigning scores to different candidates.
- LF5 **Visual Q&A → Zero-shot Sequence Classifier.** Same as LF1, except with RoBERTa replaced by LayoutLM [76], a visual question-answering model fine-tuned on the Squad2 and DocVQA datasets [47, 61]. LayoutLM operates directly on PDFs, making use of layout and visual information. The top- K candidates are obtained as in LF1.

Task 2 (Section 2.1.2): We modify LF1, LF2, and LF3 from above, and add one new LF (referred to as LF6 for coherence). For LF1 and LF2, the key difference are the questions for the Q&A transformer. Here, we derive these from the original coding dictionary [18]. For example, for *rapport building* (one of the ten indicators to be extracted; see Sections 2.1 and 4.1.1), the questions we use start with *Is the offender ...?* with the ellipsis substituted by: (i) *giving a compliment*; (ii) *accepting a compliment*; (iii) *building a special bond*; (iv) *being romantic*; (v) *showing interest*; (vi) *talking about personal similarities*. The questions used for the other indicators can be found in Appendix F. For LF3 the only modification is setting the categories to the ten indicators of interest (Section 2.1.2).

- LF6 **Pre-trained NLI sequence classifier.** We use a fine-tuned version of RoBERTa-large for NLI obtained from [18]. For all offender’s messages, the classifier assigns ten scores, each specifying how model’s confidence that given behaviour of interest (Section 2.1.2) is present. Predictions with confidence of at least 0.4 can be flagged for user validation. If there are more than K such predictions (see Section 3), the K highest confidence messages per each of the ten labels are returned. If no messages exceeded the threshold, “no evidence” is returned.

Examples of labeling functions outputs for the Perverted Justice dataset are included in Table B3.

4.1.4 ELICIT’s Human Validators. The evaluations presented in this article were conducted by the paper authors and their collaborators. Specifically, for Task 1, manual extraction for 25 cases was done by two authors, one of whom also completed the ELICIT extraction for 100 cases. For Task 2, manual extraction from 66 conversations was done by two forensic psychologists as part of their MSc dissertations. Validation and ELICIT extraction were conducted by an author who




Fig. 3. Annotation time vs. word count on Sentencing remarks. For top-1 ELICIT (blue), time is constant, whereas manual annotation (purple) scales roughly linearly with word count. Increasing the allowed answers per LF from 1 to 3, doubles the annotation time.

is also a forensic psychologist. Ethics clearance was received, separately, for both tasks, via the participants’ departmental ethics board. Due to the potentially traumatizing nature of the data, no external participants were recruited.

4.2 Results: Core Features

4.2.1 Time-Efficiency. The difference in annotation times between manual and semi-automated extraction depends on length of the documents. Figure 3 shows the annotation times for top-1 and top-3 ELICIT vs. manual annotation, on the Sentencing remarks dataset (Section 4.1.1). Top-1 and top-3 ELICIT respectively took 102 and 227 seconds on average, *independently of the document length*. Compared with manual annotation, top-1 ELICIT achieved an order-of-magnitude speed-up for documents containing 4,000 or more words. For Task 2, annotators only timed the overall time. On average, each conversation took 2.4 minutes to annotate, achieving a $\sim 20x$ time reduction compared with manual annotation. Top-1 and top-3 achieved similar timing, as the annotator was more familiar with the system when performing the top-3 validation. While validators are co-authors of this article (Section 4.1.4), the improvements observed here show a qualitative difference which is very unlikely to disappear with other users due to limits of human reading speed.

4.2.2 Performance. We compare the precision and recall of ELICIT to an automated baseline called SNORKEL [63] (see Appendix E for a majority rule baseline). In Figure 4, we show the comparison of top- K validation in ELICIT, using $K = 1, 3$. As expected, for both tasks, the precision of ELICIT is comparable to manual extraction, irrespective of K .⁵

For Task 1, the user incorrectly validated only five (out of 260) of the instances. The failures were due to contradictions within the retrieved snippets. For example, the user selected “not premeditated” after seeing the quote “I cannot find that there was a lack of premeditation such as to amount to a mitigating factor”. However, reading the full text, the judge later added that “Whilst I do not find there was a significant degree of planning or premeditation as an aggravating factor, equally I cannot find that there was a lack of premeditation such as to amount to a mitigating factor.”

For full automation, precision varied among the extracted variables. It fell below 0.5 for 4/13 (Task 1), resp. 3/10 (Task 2), of the variables. Precision of both models was better on Task 2, likely

⁵Top-1 ELICIT did not retrieve any “Challenge” label candidates, which is reported as zero precision.

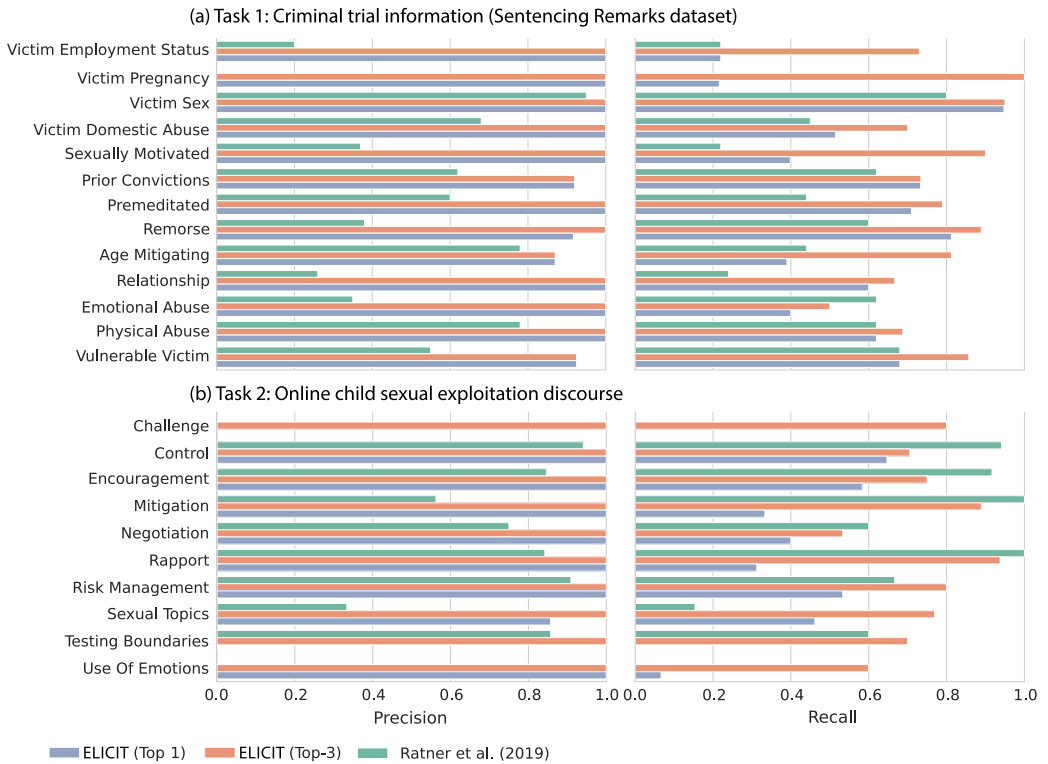


Fig. 4. Weighted precision and recall performance. Top-1 (blue) and Top-3 (orange) ELICIT are compared with fully automated extraction from Reference [64]. Precision and recall are weighted by per-class support as to not misrepresent the performance due to class imbalance. Precision and recall can be affected by variable prevalence. While prevalence is the same for each variable in Task 2 (a True/False label is assigned for occurrence of each behavior in every conversation), the same is not true for Task 1 (Table B2).

due to addition of LF6. ELICIT’s overall superior precision is evidence that the retrieved text snippets provide the user with sufficient context for label validation.

Increasing the number of candidates from top-1 to top-3 had little impact on precision but it significantly improved in *recall*. For Task 1, top-1 performed slightly better than the automated baseline with 0.56 and 0.46 mean recall, while top-3 outperforms both with 0.78 mean recall. The K -related recall improvement is even more pronounced for Task 2, where the automated baseline outperformed top-1 ELICIT with 0.58 vs. 0.33. However, top-3 ELICIT increased recall to 0.75, surpassing both. In ELICIT, the high accuracy of the user implies that low recall only occurs when relevant information is not retrieved by the LFs (modulo user fatigue). The large recall improvement when moving from top-1 to top-3 thus suggests that the LF confidence is not sufficiently well-calibrated. Similarly to other weak supervision methods, recall can be further improved by adding and improving the LFs (e.g., by adding keywords or paraphrased questions). Figure 4 only reports performance only on the Sentencing Remarks dataset. We report the precision and recall for the press articles dataset (Task 1) in Figure B2 in the appendix.

4.3 Results: Advanced Features

4.3.1 Validation vs. Deferral. When time efficiency is critical, it may be worth considering a more efficient solution than top-1 validation. We use the Sentencing Remarks data (Section 4.1.1)

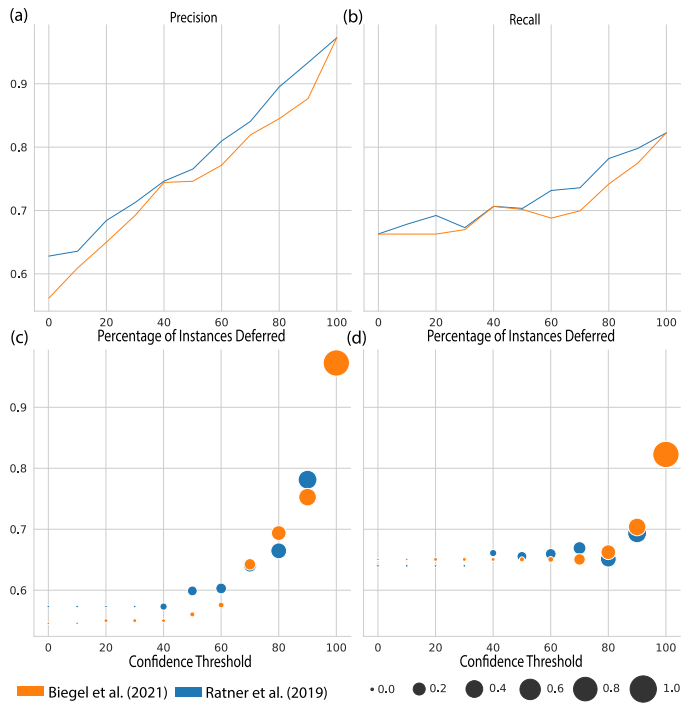


Fig. 5. Performance under different deferral schemes. (a, b) Precision and recall as a function of percentage of instances deferred to ELICIT. The lowest confidence instances are deferred first. (c, d) Precision and recall when instances below a given confidence threshold are deferred from to ELICIT. Point sizes are scaled by the number of instances under a given confidence threshold.

to demonstrate how performance changes for two types of deferral: (1) fixed *budget* deferral, where only a certain percentage of the cases can be deferred to a human; and (2) fixed *threshold* deferral, where every prediction below a certain confidence is deferred to a human. Figure 5 shows the tradeoff in precision and recall for fixed budget (a and b) and threshold (c and d). The colours distinguish between ranking methods with Ratner et al. [64] in blue, and Biegel et al. [8] in orange. Both achieve similar results in our setting.

Precision improves linearly with percentage of deferred cases. Recall, however, improves faster when more cases are human validated. In this setting, a saving of 40% in human time results in a loss of ~20% precision and ~10% recall, highlighting the value of adding a human-in-the-loop element. Deferral may be the preferred option in settings where either (1) the difference in performance between the automation and human validation is small, (2) the confidence produced by the automated methods is well-calibrated, or (3) reducing time and human-effort is a higher priority.

4.3.2 Improving Recall Via Training. With every human validation, we obtain an additional label beyond the original dataset. This can be used to improve recall by fine-tuning the labeling functions. Beyond improved performance on new instances, ELICIT also alerts the user if fine-tuning resulted in new validation opportunity in already reviewed data.

Figure 6 shows the improvement in recall when fine-tuning on 20 validated cases. The average improvement in recall was 0.075, with a maximum improvement of 0.15 (for domestic abuse, 0.7 increased to 0.85). These auxiliary uses of the validated data—refining the ranking, and fine-tuning the LFs—may improve performance enough to make deferral or even fully automated extraction feasible in the longer term in certain settings.




Fig. 6. The recall delta before and after fine-tuning with ELICIT on the Sentencing remarks dataset.

4.3.3 Reflections from the User’s Perspective. One of the goals behind developing ELICIT was to improve the well-being of human annotators. Although we did not measure well-being quantitatively, we documented reflections from the two authors who annotated *both* manually and using ELICIT (Section 4.1.4). For Task 1, the author who extracted information about murder cases tried in the UK, said: “Reviewing large amounts of court materials in such detail is mentally draining. I found it could limit the quantity which I could label in a single sitting. Using ELICIT limits the exposure to unnecessary material, allowing more labeling to be completed before a break is required.” The validator for Task 2, who is an experienced annotator, found the user interface improved their own ability to perform the task, stating that “Manual message-by-message labeling is difficult because it does not really reflect how humans think naturally, while chunking the messages together and asking me “Is this rapport?” seems more familiar as an everyday task.” This validator also stated that while the content was equally unpleasant to read, it was a relief to spend less time with it.

While promising, the above reflections are insufficient to draw conclusions about the usability of the user interface. This would require a more methodical approach and many more participants, making it a subject of future work.

5 ETHICS AND SOCIAL IMPACT

Data. All the data used in this study are publicly available. For Task 1, although the original sentencing remarks and news articles contained full names, we removed these when creating the datasets. These could be re-identified by finding the original publication; however, as the data is in the public domain, this is unlikely to result in further harm to the individual’s privacy. For Task 2, the data is publicly available via the Perverted Justice website [59]. The website contains conversations between adults and adult decoys pretending to be minors. No actual minors were involved in the conversations. The website only publishes chats that have led to convictions. We note that it is highly unlikely that the offenders consent to publish these conversations. While contestable, this data is the only publicly available source in the domain, and has been widely used in the relevant literature.

Potential impact of this research. The goal of this work is to facilitate creation of new tabular datasets, which would enable new avenues of quantitative and mixed-methods research. This work opens new routes to fill critical data gaps, which is much needed to improve our understanding of the criminal justice system [81]. However, we acknowledge our research can lead to creation of low-quality, misleading, or cherry-picked datasets, if used irresponsibly. We stress the necessity of accompanying any created dataset by detailed extraction methodology documentation, including how ambiguity is dealt with, true/false positive/negative rates from sample testing, and disclosure of any other known biases and errors.

6 CONCLUSION

We present a framework for human-validated information extraction from text. We centered our investigation around two use cases from the criminal justice domain. Based on their commonalities, we identified several key functionality requirements: accuracy, time-efficiency, flexibility, reproducibility, and reducing user’s exposure to harmful content. We reviewed suitability of different collaborative settings and tools with respect to the identified requirements. Since none satisfied all our requirements, we developed ELICIT which we release as open-source.

ELICIT is a flexible tool useful for a variety of information extraction tasks. Its design is inspired by weak supervision approaches: we use a set of algorithmic annotators (labeling functions) to identify relevant pieces of text, which are then validated by the user. Compared with manual annotation, ELICIT can attain comparable accuracy at fraction of time. This is achieved by leveraging the complementary strengths of humans and machines in our use cases: speed and accuracy.

We perform a case-study, evaluating ELICIT on three extraction tasks based on our criminal justice use cases. In each case, we achieve accuracy close to manual annotation with *orders of magnitude lower time investment*. ELICIT significantly outperforms the *automated* extraction on both precision and recall. We demonstrate that recall can be further improved by using the already validated data for fine-tuning. We further quantify the tradeoff between human effort and performance within a deferral setup. Finally, based on our own experience extracting information manually and of using ELICIT, we felt using ELICIT required less emotional strain compared with manual annotation.

Our framework can be particularly effective for extraction of factual information from very long documents, when high precision is an essential requirement. Beyond helping human annotators do their work more effectively, we believe the continual learning component of our system (Section 4.3.2) is a promising direction for improving machine performance via human feedback. Learning from human feedback is a topic of growing importance in machine learning [e.g., 13, 29, 36], where fine-tuning language models on human feedback can yield significant gains [6, 14, 56, 70]. These approaches largely rely on human annotators hired specifically for the purpose of ranking model outputs based on their quality. In contrast, ELICIT uses the interactions with users themselves to become more performant, demonstrating a complementary venue to effectively collecting and learning from human feedback. We hope ELICIT’s value-led design—combined with engineering choices informed by task-specific requirements—inspires further research in this direction.

APPENDICES

A F1 SCORES




Fig. A1. F1 scores all tasks: (a) sentencing remarks, (b) perverted justice, and (c) press articles. Top-1 (blue) and Top-3 (orange) ELICIT is compared with fully automated extraction from Reference [64].

B PRESS ARTICLES DATASET




Fig. B1. Distribution of press site coverage for murder cases from the Law Pages website. Color indicates whether we use articles from the corresponding sites: (green) use, (red) do not use. The five highest frequency sites are used.




Fig. B2. Weighted precision and recall performance on the Law Pages dataset. ELICIT (orange) Top-3 is compared with fully automated extraction (green) from Reference [64]. Precision and recall are weighted by per-class support to account for class imbalance.

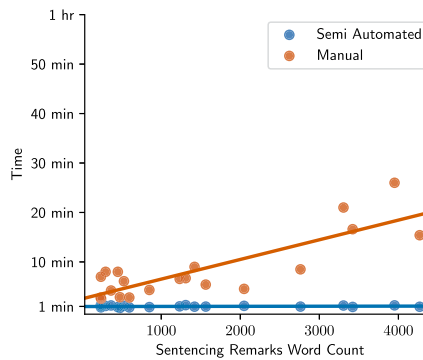


Fig. B3. Annotation time against Law Pages word count. For ELICIT semi-automated annotation (blue), time is constant, whereas manual annotation (orange) scales roughly linearly with word count.

C VARIABLES

Table B1. List of Variables for Sentencing Remarks and News Articles

Variable	Sentencing Remarks	News Articles
Victim Sex	✓	✓
Victim Domestic Abuse	✓	✓
Victim Vulnerable	✓	✓
Victim Pregnancy	✓	✗
Victim Employment Status	✓	✓
Victim Religion	✗	✗
Victim Race	✗	✗
Victim Disability	✗	✗
Victim Sexual Orientation	✗	✗
Physical Abuse	✓	✗
Mental Abuse	✓	✗
Remorse	✓	✓
Prior Convictions	✓	✓
Sexually Motivated	✓	✓
Racially Motivated	✗	✗
Age Mitigating	✓	✗
Premeditation	✓	✓
Relationship	✓	✓

Check marks and crosses indicate whether a variable was extracted for the corresponding dataset. Variables with no or only few relevant examples were removed from annotation process.

Table B2. Number of Non-Abstain Data Points of Variables for Sentencing Remarks

Variable	number of non-abstains
Victim Sex	20
Victim Domestic Abuse	14
Victim Vulnerable	15
Victim Pregnancy	16
Victim Employment Status	13
Victim Religion*	0
Victim Race*	0
Victim Disability*	2
Victim Sexual Orientation*	2
Physical Abuse	17
Mental Abuse	10
Remorse	10
Prior Convictions	17
Sexually Motivated	12
Racially Motivated*	0
Age Mitigating	17
Premeditation	20
Relationship	19

The reported (weighted) precision and recall can only be evaluated on the non-abstain instances. * denotes variables which are not included in evaluation (see Table B1).

D LABELING FUNCTION OUTPUT EXAMPLES

Table B3. Examples of LF1 Outputs for the Perverted Justice Dataset

Variable	Value	Confidence	Explanation
Rapport	True	92%	<i>PRED</i> : how do you dress <i>PRED</i> : what type of food do you like <i>VICT</i> : um wel i lik tanks <i>VICT</i> : i love pizza <i>PRED</i> : what scares you
Control	True	100%	<i>PRED</i> : doesnt matter <i>PRED</i> : what ever <i>VICT</i> : u tell me <i>PRED</i> : what would you like to me to do <i>VICT</i> : watever u wana do relly <i>PRED</i> : you just want to mess around dont you <i>VICT</i> : do u? <i>PRED</i> : maybe <i>VICT</i> : lol <i>VICT</i> : well i wana kno wat
Use of Emotions	True	83%	<i>PRED</i> : i swaer i dont want you in trouble because that trouble for me <i>PRED</i> : you know what i mean <i>VICT</i> : ya <i>VICT</i> : i just wana b carful <i>PRED</i> : it turned off until i pay it friday

E MAJORITY RULES




Fig. C1. Weighted precision and recall performance for Sentencing Remarks. Majority Rules (blue) and Top-3 (orange) ELICIT is compared with fully automated extraction from Reference [64]. Precision and recall are weighted by per-class support as to not misrepresent the performance due to class imbalance.

F SCHEMA EXAMPLES

We provide a small example of the schemas used for the variables *victim sex* and *prior convictions*. The schemas—which are intended to be designed by a domain expert (here by the authors)—differ between datasets to reflect the relevant context. The employed labeling functions (Section 4.1.3)

require three schemas: category, question, and keyword. The category and question schemas are used by the labeling functions based on Q&A algorithms, and the keyword schema is used by the keyword search. We represent the schemas here as lists; in reality, these are defined in YAML files.

F.1 Crown Court Sentencing Remarks dataset

Category Schema.

- Victim sex
 - Male
 - Female
- Prior Convictions
 - Prior Convictions
 - No Prior Convictions

Question Schema.

- Victim sex
 - What sex was the victim?
 - Was the victim male?
 - Was the victim female?
- Prior Convictions
 - Prior convictions?
 - Did the defendant have prior convictions?
 - Previous crimes?

Keyword Schema.

- Victim sex
 - male:
 - * male
 - * man
 - * boy
 - female:
 - * female
 - * woman
 - * girl
- Prior Convictions
 - Prior Convictions:
 - * prior convictions
 - * previous convictions
 - * criminal record
 - No Prior Convictions:
 - * No prior convictions
 - * Previous good character

F.2 Perverted Justice Dataset

Category Schema.

- Rapport
 - Rapport
 - No Rapport
- Control
 - Control

- No Control
- Negotiation
 - Negotiation
 - No Negotiation
- Challenge
 - Challenge
 - No Challenge
- Use of emotions
 - Use of emotions
 - No use of emotions
- Mitigation
 - Mitigation
 - No mitigation
- Encouragement
 - Encouragement
 - No encouragement
- Risk Management
 - Risk Management
 - No Risk Management
- Sexual Topics
 - Sexual Topics
 - No Sexual Topics
- Testing Boundaries
 - Testing Boundaries
 - No Testing Boundaries

Question Schema.

- Rapport
 - is the offender giving a compliment?
 - is the offender accepting a complement?
 - is the offender building a special bond?
 - is the offender being romantic?
 - is the offender showing interest?
 - is the offender talking about personality?
 - is the offender talking about personal similarities?
- Control
 - is the offender being persistent?
 - is the offender talking about consent?
 - is the offender trying to please the victim?
 - is the offender complying with requests?
 - is the offender jealous?
 - is the offender being compliant?
 - is the offender being assertive?
 - is the offender asking a rhetorical question?
 - is the offender being patronising?
 - is the offender asking for permission?
 - is the offender checking for engagement?
- Negotiation

- is the offender offering incentives?
- is the offender making plans to meet?
- is the offender persuading the victim?
- is the offender defensive?
- is the offender talking about alcohol?
- is the offender talking about drugs?
- is the offender arranging plans?
- Challenge
 - is the offender mocking the victim?
 - is the offender insulting the victim?
 - is the offender confronting the victim?
 - is the offender rejecting the victim?
 - does the victim trust the offender?
- Use of emotions
 - is the offender showing concern?
 - is the offender looking for validation?
 - is the offender shocked?
 - is the offender angry?
 - is the offender sad?
 - is the offender confused?
 - is the offender embarrassed?
 - is the offender happy?
 - does the offender reassure the victim?
 - does the offender ask for reassurance?
- Mitigation
 - does the offender implicate themselves in a crime?
 - does the offender have a sexual preference for children?
- Encouragement
 - does the offender express willingness to engage?
 - does the offender encourage the victim?
 - does the offender comply with the victim?
 - does the offender flirt with the victim?
 - does the offender request a picture of the victim?
- Risk management
 - does the offender ask if the victim is real?
 - does the offender ask if the victim is a cop?
 - does the offender ask about the victim's mom?
 - does the offender ask about the victim's dad?
 - does the offender ask about the victim's family?
 - does the offender talk about the dangers on the internet?
 - does the offender ask about meeting the victim?
- Sexual Topics
 - is the offender talking about sexual topics?
 - is the offender talking about fantasies?
 - is the offender talking about sexual preferences?
 - is the offender talking about pornography?
 - is the offender talking about sexual acts?
 - is the offender talking about relationships?

- is the offender talking about age differences?
- Testing Boundaries
 - does the offender set boundaries?
 - does the offender check the victim’s willingness to engage?
 - does the offender talk about sex?
 - does the offender talk about relationships?
 - does the offender talk about sharing pictures?
 - does the offender talk about meeting offline?
 - does the offender talk about fantasies?
 - does the offender talk about sharing pictures?
 - is the offender being secretive?
 - is the offender bored?

G RANKING

G.1 Differences Between ELICIT and Ratner et al. [64]




Fig. E1. An overview of differences between our method and [64]. In [64], the labeling functions predict one label value per document, which is then encoded in a vector λ_i (eventually a matrix λ over all documents). This λ is used to learn a generative model $p(Y_i|\lambda_i)$. In our case, the labeling functions can produce multiple answers, each with an associated explanation. Common explanations are grouped into a λ_i vector, and corresponding probability $P(Y_i|\lambda_i)$ is produced on a per-explanation, rather than per-document, basis.

G.2 Ratner et al. [64]

We use the weak supervision method presented in Reference [64] in order to get a confidence for each explanation. We use this confidence to rank the explanations to minimize the number of explanations a user must interact with before reaching a valid answer.

A weak label matrix, λ , is formed as a $e \times m$ matrix, where m is the number of labeling functions, and e is the total number of explanations over n documents: $e = \sum_{i \in n} |e_i|$ where e_i is the set of explanations for a document i . Let Σ be the $m \times m$ covariance matrix of λ , then the parameter \hat{z} can be estimated by solving the following matrix completion problem:

$$\hat{z} = \underset{z}{\operatorname{argmin}} \left\| \left(\Sigma^{-1} + z z^T \right) \odot \Omega \right\|_F$$

where $\|\cdot\|_F$ is the Frobenius norm, and $\Omega \in \mathbb{R}^{m \times m}$ is a positive semidefinite matrix which encodes conditional independence structure amongst the labeling functions. If Ω is correctly specified, \hat{z}_i

will represent how often the labeling function $f_i(\cdot)$ independently reaches the same conclusion as the other labeling functions, and with a sufficient number of labeling functions, will serve as a proxy for labeling function accuracy. The generative model is then a function that transforms \hat{z} , and values of the labeling functions for an explanation $\lambda_{i,e=j}$ into a probabilistic label:

$$\hat{p}(y \mid \lambda_{i,e=j}) = f(\hat{z}, \lambda_{i,e=j})$$

G.3 Biegel et al. [8]

As validation progresses (explanation, value, λ_i , valid) tuples are collected. We use the human-disagreement penalty proposed in Reference [8] to penalize the generative model when there is disagreement between the model and the validated data.

In their article, they add a penalty to the matrix completion when optimizing for \hat{z} . This penalty is a simple quadratic difference between the human and the probabilistic label: $\sum_{i \in D} (f(\hat{z}, \lambda_{i,e=j}) - y_i)^2$. In effect, the optimization is penalized for disagreeing with the human. The penalty is scaled by a hyper-parameter α :

$$\hat{p}(y \mid \lambda_{i,e=j}) = f(\hat{z}, \lambda_{i,e=j}) + \alpha Pe(\mathbf{z})$$

We use $\alpha = 100$ as the authors do in their experiments.

ACKNOWLEDGMENTS

The authors thank Nikolaos Aletras, and Roi Reichart for valuable discussions.

REFERENCES

- [1] Kiran Adnan and Rehan Akbar. 2019. Limitations of information extraction methods and techniques for heterogeneous unstructured big data. *International Journal of Engineering Business Management* 11 (2019).
- [2] Monica Agrawal, Stefan Hegselmann, Hunter Lang, Yoon Kim, and David Sontag. 2022. Large language models are few-shot clinical information extractors. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics, 1998–2022.
- [3] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. *ProPublica* (2016). Retrieved from <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [4] Andrew Arshnt and Daniel Etcovitch. 2018. The human cost of online content moderation. *Harvard Journal of Law and Technology* (2018).
- [5] Yannis Assael, Thea Sommerschild, Brendan Shillingford, Mahyar Bordbar, John Pavlopoulos, Marita Chatzipanagiotou, Ion Androutsopoulos, Jonathan Prag, and Nando de Freitas. 2022. Restoring and attributing ancient texts using deep neural networks. *Nature* 603, 7900 (2022), 280–283. DOI: <https://doi.org/10.1038/s41586-022-04448-z>
- [6] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom B. Brown, Jack Clark, Sam McCandlish, Chris Olah, Benjamin Mann, and Jared Kaplan. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *CoRR* abs/2204.05862 (2022). <https://arxiv.org/abs/2204.05862>
- [7] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [8] Samantha Biegel, Rafah El-Khatib, Luiz Otávio Vilas Boas Oliveira, Max Baak, and Nanne Aben. 2021. Active WeaSuL: Improving weak supervision with active learning. *CoRR* abs/2104.14847 (2021). <https://arxiv.org/abs/2104.14847>
- [9] Benedikt Boecking, Willie Neiswanger, Eric Xing, and Artur Dubrawski. 2021. Interactive weak supervision: Learning useful heuristics for data labeling. In *International Conference on Learning Representations*.
- [10] Peter Briggs, Walter T. Simon, and Stacy Simonsen. 2011. An exploratory study of internet-initiated sexual offenses and the chat room sex offender: Has the internet enabled a new typology of sex offender? *Sexual Abuse* 23, 1 (2011), 72–91. DOI: <https://doi.org/10.1177/1079063210384275>
- [11] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric

- Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*. H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901.
- [12] Gonçalo Carnaz, Vitor Beires Nogueira, Mário Antunes, and Nuno Ferreira. 2018. An automated system for criminal police reports analysis. In *International Conference on Soft Computing and Pattern Recognition*. Springer, 360–369. DOI : https://doi.org/10.1007/978-3-030-17065-3_36
- [13] Paul F. Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in Neural Information Processing Systems* 30 (2017).
- [14] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. arXiv:2210.11416. Retrieved from <https://arxiv.org/abs/2210.11416>
- [15] Maria-Veronica Ciocanel, Chad M. Topaz, Rebecca Santorella, Shilad Sen, Christian Michael Smith, and Adam Hufstetler. 2020. JUSTFAIR: Judicial system transparency through federal archive inferred records. *PLOS ONE* 15, 10 (2020), 1–20. DOI : <https://doi.org/10.1371/journal.pone.0241381>
- [16] John Tyler Clemons. 2014. Blind injustice: The Supreme Court, implicit racial bias, and the racial disparity in the criminal justice system. *Am. Crim. L. Rev.* 51 (2014), 689.
- [17] Rubrix community. 2021. Rubrix. Retrieved from <https://github.com/recognai/rubrix>
- [18] Darren Cook, Miri Zilka, Heidi DeSandre, Susan Giles, Adrian Weller, and Simon Maskell. 2022. Can we automate the analysis of online child sexual exploitation discourse? *CoRR* abs/2209.12320 (2022). <https://arxiv.org/abs/2209.12320>
- [19] Cristina Criddle. 2021. Facebook Moderator: “Every Day was a Nightmare”. Retrieved from <https://www.bbc.co.uk/news/technology-57088382>
- [20] Michael Desmond, Michael Muller, Zahra Ashktorab, Casey Dugan, Evelyn Duesterwald, Kristina Brimijoin, Catherine Finegan-Dollak, Michelle Brachman, Aabhas Sharma, Narendra Nath Joshi, and Qian Pan. 2021. Increasing the speed and accuracy of data labeling through an AI assisted interface. In *26th International Conference on Intelligent User Interfaces* (College Station, TX, USA) (IUI '21). Association for Computing Machinery, New York, NY, 392–401. DOI : <https://doi.org/10.1145/3397481.3450698>
- [21] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Minneapolis, Minnesota. Association for Computational Linguistics. 4171–4186.
- [22] Xiao Ding, Yue Zhang, Ting Liu, and Junwen Duan. 2015. Deep learning for event-driven stock prediction. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- [23] Thomas Douglas, Jonathan Pugh, Ilina Singh, Julian Savulescu, and Seena Fazel. 2017. Risk assessment tools in criminal justice and forensic psychiatry: The need for better data. *European Psychiatry* 42 (2017), 134–137.
- [24] Ian A. Elliott. 2017. A self-regulation model of sexual grooming. *Trauma, Violence, & Abuse* 18, 1 (2017), 83–97. DOI : <https://doi.org/10.1177/1524838015591573>
- [25] David Ferguson. 2010. The Law Pages. Retrieved from <https://www.thelawpages.com/>
- [26] Jessica L. Feuston and Jed R. Brubaker. 2021. Putting tools in their place: The role of time and perspective in human-AI collaboration for qualitative analysis. 5, CSCW2, Article 469 (oct 2021), 25 pages. DOI : <https://doi.org/10.1145/3479856>
- [27] Matías García-Constantino, Katie Atkinson, Danushka Bollegala, Karl Chapman, Frans Coenen, Claire Roberts, and Katy Robson. 2017. CLIEL: Context-based information extraction from commercial law documents. In *International Conference on Artificial Intelligence and Law*.
- [28] Emily A. Greene-Colozzi, Georgia M. Winters, Brandy Blasko, and Elizabeth L. Jeglic. 2020. Experiences and perceptions of online sexual solicitation and grooming of minors: A retrospective report. *Journal of Child Sexual Abuse* 29, 7 (2020), 836–854. DOI : <https://doi.org/10.1080/10538712.2020.1801938>
- [29] Shane Griffith, Kaushik Subramanian, Jonathan Scholz, Charles L. Isbell, and Andrea L. Thomaz. 2013. Policy shaping: Integrating human feedback with reinforcement learning. *Advances in Neural Information Processing Systems* 26 (2013).
- [30] Ralph Grishman. 2019. Twenty-five years of information extraction. *Natural Language Engineering* 25, 6 (2019), 677–692. DOI : <https://doi.org/10.1017/S1351324919000512>
- [31] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *International Conference on Machine Learning*. PMLR, 1321–1330.
- [32] Shohreh Haddadan, Elena Cabrio, and Serena Villata. 2019. Yes, we can! mining arguments in 50 years of US presidential campaign debates. In *ACL 2019-57th Annual Meeting of the Association for Computational Linguistics*. 4684–4690.
- [33] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Thomas Hennigan, Eric Noland, Katherine Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karén Simonyan, Erich Elsen, Oriol Vinyals,

- Jack Rae, and Laurent Sifre. 2022. An empirical analysis of compute-optimal large language model training. *Advances in Neural Information Processing Systems*. 6, 35 (2022), 16–30. https://proceedings.neurips.cc/paper_files/paper/2022/hash/c1e2faff6f588870935f114ebe04a3e5-Abstract-Conference.html
- [34] Allen H. Huang, Hui Wang, and Yi Yang. 2022. FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research* 40, 2 (2022), 806–841.
- [35] Benjamin W. K. Hung, Shashika R. Muramudalige, Anura P. Jayasumana, Jytte Klausen, Rosanne Libretti, Evan Moloney, and Priyanka Renugopalakrishnan. 2019. Recognizing radicalization indicators in text documents using human-in-the-loop information extraction and NLP techniques. In *2019 IEEE International Symposium on Technologies for Homeland Security (HST)*. IEEE, 1–7. DOI : <https://doi.org/10.1109/HST47167.2019.9032956>
- [36] Hong Jun Jeon, Smitha Milli, and Anca Dragan. 2020. Reward-rational (implicit) choice: A unifying formalism for reward learning. *Advances in Neural Information Processing Systems* 33 (2020), 4415–4426.
- [37] Jialun Aaron Jiang, Kandrea Wade, Casey Fiesler, and Jed R. Brubaker. 2021. Supporting serendipity: Opportunities and challenges for human-AI collaboration in qualitative analysis. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 94 (apr 2021), 23 pages. DOI : <https://doi.org/10.1145/3449168>
- [38] Judiciary. 2022. Courts and Tribunals Judiciary: Judgements. Retrieved from <https://www.judiciary.uk/judgments/>
- [39] Ministry of Justice. 2022. Data First: Criminal Courts Linked Data. Retrieved from <https://www.gov.uk/government/publications/data-first-criminal-courts-linked-data>
- [40] Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP'20)*. Online. Association for Computational Linguistics. 6769–6781.
- [41] Margaret Bull Kovera. 2019. Racial disparities in the criminal justice system: Prevalence, causes, and a search for solutions. *Journal of Social Issues* 75, 4 (2019), 1139–1164.
- [42] Vivian Lai, Samuel Carton, Rajat Bhatnagar, Q. Vera Liao, Yunfeng Zhang, and Chenhao Tan. 2022. Human-AI collaboration via conditional delegation: A case study of content moderation. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22*). Association for Computing Machinery, New York, NY, USA, Article 54, 18 pages. DOI : <https://doi.org/10.1145/3491102.3501999>
- [43] David Lammy. 2017. The lammy review: An independent review into the treatment of, and outcomes for, black, asian and minority ethnic individuals in the criminal justice system. *London: Lammy Review* (2017).
- [44] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. arXiv:1907.11692. Retrieved from <https://arxiv.org/abs/1907.11692>
- [45] Wayne A. Logan and Andrew Guthrie Ferguson. 2016. Policing criminal justice data. *Minn. L. Rev.* 101 (2016), 541.
- [46] Maximilian Mackeprang, Claudia Müller-Birn, and Maximilian Timo Stauss. 2019. Discovering the sweet spot of human-computer configurations: A case study in information extraction. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–30.
- [47] Minesh Mathew, Dimosthenis Karatzas, R. Manmatha, and C. V. Jawahar. 2021. DocVQA: A dataset for VQA on document images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2200–2209.
- [48] Paul Millar and Akwasi Owusu-Bempah. 2011. Whitewashing criminal justice in Canada: Preventing research through data suppression. *Canadian Journal of Law and Society/La Revue Canadienne Droit et Société* 26, 3 (2011), 653–661.
- [49] Ines Montani and Matthew Honnibal. 2018. Prodigy: An annotation tool for AI, machine learning & NLP. Available online: <https://prodi.gy> (accessed on 14 April 2024).
- [50] Zara Nasar, Syed Waqar Jaffry, and Muhammad Kamran Malik. 2018. Information extraction from scientific articles: A survey. *Scientometrics* 117, 3 (2018), 1931–1990. DOI : <https://doi.org/10.1007/s11192-018-2921-5>
- [51] Mariana Neves and Ulf Leser. 2014. A survey on annotation tools for the biomedical literature. *Briefings in bioinformatics* 15, 2 (2014), 327–340.
- [52] Elastic NV. 2010. Elasticsearch. Retrieved from www.elastic.co
- [53] Pepijn Obels, Daniel Lakens, Nicholas A. Coles, Jaroslav Gottfried, and Seth A. Green. 2020. Analysis of open data and computational reproducibility in registered reports in psychology. *Advances in Methods and Practices in Psychological Science* 3, 2 (2020), 229–237.
- [54] Elsa A. Olivetti, Jacqueline M. Cole, Edward Kim, Olga Kononova, Gerbrand Ceder, Thomas Yong-Jin Han, and Anna M. Hiszpanski. 2020. Data-driven materials research enabled by natural language processing and information extraction. *Applied Physics Reviews* 7, 4 (2020), 041317. DOI : <https://doi.org/10.1063/5.0021106>
- [55] Pablo A. Ormachea, Gabe Haarsma, Sasha Davenport, and David M. Egleman. 2015. A new criminal records database for large-scale analysis of policy and behavior. *Journal of Science and Law|jscilaw.org* September 1, 1 (2015), 2.
- [56] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda

- Askill, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*. 6, 35 (2022), 30–44. https://proceedings.neurips.cc/paper_files/paper/2022/hash/b1efde53be364a73914f58805a001731-Abstract-Conference.html
- [57] Rachel O’Connell. 2003. A Typology of Child Cyberexploitation and Online Grooming Practices. Retrieved from <http://image.guardian.co.uk/sys-files/Society/documents/2003/07/17/Groomingreport.pdf>
- [58] Tal Perry. 2021. Lighttag: Text annotation platform. arXiv:2109.02320. Retrieved from <https://arxiv.org/abs/2109.02320>
- [59] Perverted justice: A dataset. Available online: <http://perverted-justice.com/> (accessed on 14 April 2024).
- [60] Sajjadur Rahman and Eser Kandogan. 2022. Characterizing practices, limitations, and opportunities related to text information extraction workflows: A human-in-the-loop perspective. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI ’22). Association for Computing Machinery, New York, NY, Article 628, 15 pages. DOI : <https://doi.org/10.1145/3491102.3502068>
- [61] Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Melbourne, Australia. Association for Computational Linguistics, 784–789.
- [62] Charvi Rastogi, Liu Leqi, Kenneth Holstein, and Hoda Heidari. 2023. A taxonomy of human and ML strengths in decision-making to investigate human-ML complementarity. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 11, 1 (2023), 127–139.
- [63] Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. 2017. Snorkel: Rapid training data creation with weak supervision. In *Proceedings of the VLDB Endowment. International Conference on Very Large Data Bases*, Vol. 11. NIH Public Access, 269. DOI : <https://doi.org/10.1007/s00778-019-00552-1>
- [64] Alexander Ratner, Braden Hancock, Jared Dunnmon, Frederic Sala, Shreyash Pandey, and Christopher Ré. 2019. Training complex models with multi-task weak supervision. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 4763–4771. DOI : <https://doi.org/10.1609/aaai.v33i01.33014763>
- [65] Cynthia Rudin, Caroline Wang, and Beau Coker. 2020. The age of secrecy and unfairness in recidivism prediction. *Harvard Data Science Review* 2, 1 (March 2020). Retrieved from <https://hdsr.mitpress.mit.edu/pub/7z10o269>
- [66] Amazon Web Services. 2021. OpenSearch. Retrieved from www.opensearch.org
- [67] Joy Shelton, Jennifer Eakin, Tia Hoffer, Yvonne Muirhead, and Jessica Owens. 2016. Online child sexual exploitation: An investigative analysis of offender characteristics and offending behavior. *Aggression and Violent Behavior* 30 (2016), 15–23. DOI : <https://doi.org/10.1016/j.avb.2016.07.002>
- [68] Miriah Steiger, Timir J. Bharucha, Sukrit Venkatagiri, Martin J. Riedl, and Matthew Lease. 2021. The psychological well-being of content moderators: the emotional labor of commercial moderation and avenues for improving support. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [69] K. Stoykov and S. Chelebieva. 2019. Legal data extraction and possible applications. In *IOP Conference Series: Materials Science and Engineering*, Vol. 618. IOP Publishing, 012037.
- [70] Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. LaMDA: Language models for dialog applications. arXiv:2201.08239. Retrieved from <https://arxiv.org/abs/2201.08239>
- [71] Bernhard Walzl, Georg Bonczek, and Florian Matthes. 2018. Rule-based information extraction: Advantages, limitations, and perspectives. *Jusletter IT (02 2018)* (2018).
- [72] Yanshan Wang, Liwei Wang, Majid Rastegar-Mojarad, Sungrim Moon, Feichen Shen, Naveed Afzal, Sijia Liu, Yuqun Zeng, Saeed Mehrabi, Sunghwan Sohn, and Hongfang Liu. 2018. Clinical information extraction applications: A literature review. *Journal of Biomedical Informatics* 77 (2018), 34–49. DOI : <https://doi.org/10.1016/j.jbi.2017.11.011>
- [73] Leigh Weston, Vahe Tshitoyan, John Dagdelen, Olga Kononova, Amalie Trewartha, Kristin A. Persson, Gerbrand Ceder, and Anubhav Jain. 2019. Named entity recognition and normalization applied to large-scale information extraction from the materials science literature. *Journal of Chemical Information and Modeling* 59, 9 (2019), 3692–3702. DOI : <https://doi.org/10.1021/acs.jcim.9b00470>
- [74] Rebecca Williams, Ian A. Elliott, and Anthony R. Beech. 2013. Identifying sexual grooming themes used by internet sex offenders. *Deviant Behavior* 34, 2 (2013), 135–152. DOI : <https://doi.org/10.1080/01639625.2012.707550>
- [75] Georgia M. Winters, Leah E. Kaylor, and Elizabeth L. Jeglic. 2017. Sexual offenders contacting children online: an examination of transcripts of sexual grooming. *Journal of Sexual Aggression* 23, 1 (2017), 62–76. DOI : <https://doi.org/10.1080/13552600.2016.1271146>
- [76] Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. 2020. LayoutLM: Pre-training of text and layout for document image understanding. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD’20)*. Association for Computing Machinery, New York, NY, 1192–1200. <https://doi.org/10.1145/3394486.3403172>
- [77] Wengpeng Yin, Jamaal Hay, and Dan Roth. 2019. Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the*

9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Hong Kong, China. Association for Computational Linguistics. 3914–3923.

- [78] Ashwini V. Zadgaonkar and Avinash J. Agrawal. 2021. An overview of information extraction techniques for legal document analysis and processing. *International Journal of Electrical & Computer Engineering (2088-8708)* 11, 6 (2021).
- [79] Gohar Zaman, Hairulnizam Mahdin, Khalid Hussain, and A. Rahman. 2020. Information extraction from semi and unstructured data sources: A systematic literature review. *ICIC Express Lett.* 14, 6 (2020), 593–603. DOI: <https://doi.org/10.24507/icicel.14.06.593>
- [80] Jieyu Zhang, Cheng-Yu Hsieh, Yue Yu, Chao Zhang, and Alexander Ratner. 2022. A survey on programmatic weak supervision. arXiv:2202.05433. Retrieved from <https://arxiv.org/abs/2202.05433>
- [81] Miri Zilka, Bradley Butcher, and Adrian Weller. 2022. A survey and datasheet repository of publicly available US criminal justice datasets. In *Thirty-Sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Received 16 October 2023; revised 7 February 2024; accepted 15 February 2024