Automated Extraction of Substance Use and Co-occurring Disorders from Probation Records[1]

Karine Megerdoomian, The MITRE Corporation
Charles E. Horowitz, The MITRE Corporation
Amy B. Marsh, The MITRE Corporation

INDIVIDUALS IN THE criminal justice system who have a history of substance use disorder (SUD) have been shown to display a higher rate of rearrest and recidivism (Stahler et al., 2013; Fazel et al., 2016). Of particular concern are cases where the substance use disorder appears with co-occurring disorders[1] (CODs); these individuals are less likely to enter and successfully complete treatment and are at an even greater risk for criminal relapse (Monahan, 1992; Drake & Wallach, 1989). Comorbidity of mental illness with addiction is also associated with numerous negative health outcomes, increased risk of homelessness, loss of employment, or self-harm (SAMHSA, 2022). Given their multiple needs, these individuals often require specialized interventions, providing integrated mental health and substance use services, to facilitate their reintegration within society. Comprehensive research on the prevalence, trends, and correlates of SUDs and co-occurring mental illnesses is necessary to guide evidence-based, timely, and effective policies and programs aimed at increasing public safety and reducing recidivism (Fearn et al., 2017). Yet, there is a dearth of research on this topic within the justice system, and most studies on this population are limited to small-scale cohort studies.

The U.S. Probation and Pretrial Services Office (PPSO) produces billions of pages of information on individuals under supervision, including detailed social and psychological history on prior and current substance use, official diagnoses, and treatment information. This information is critical for probation officers and district chiefs to better assist the substance use population and to guide efficient intervention strategies. However, the data are predominantly stored in free text in multiple large documents rather than in structured format, making unassisted human review and analysis unfeasible, thus underscoring the need for automated knowledge discovery techniques. Automated extraction of structured meaning from narratives to find, interpret, and prioritize knowledge, with a focus on identifying social history information such as substance use and mental illness, can provide a reprieve from the time- and cost-prohibitive nature of human review, allowing probation officers and personnel to devote their time to higher priority tasks that require human cognitive skills.

This article describes the application of a Natural Language Processing (NLP) and Artificial Intelligence (AI) system to discover important information on substance use, mental illness diagnoses, and treatment history of individuals under supervision. We developed a system for the automatic detection of four main events in the social history of the individual under supervision within free-text probation documents: (i) any evidence of substance use (alcohol, prescription, and illegal drugs) also defined as "indefinite diagnoses" for SUD; (ii) official diagnoses for SUD; (iii) official diagnoses for COD; and (iv) history of contract or non-contract treatment for SUD. We also automatically identify and extract related information (e.g., temporal information, facilities, treatment type, treatment outcome) within the text. The results are combined with metadata information from the Probation and Pretrial Services Automated Case Tracking System (PACTS) on client demographics and supervision dates. The system applies analytics to large data sets (N=98,389) over multiple documents (254,585 total documents and over 14 million Chrono entries), fuses the extracted information in a structured form, and performs analytic reasoning to enhance results. The results show that about 93 percent of this population have had a substance use issue in their lives and about 15 percent have officially been diagnosed with SUD, while approximately 29 percent have also received a formal diagnosis for a co-occurring disorder. Top mental disorders that co-occur with addiction are depressive disorders, anxiety disorders, bipolar I disorder, ADHD, and posttraumatic stress disorder (PTSD). Nearly 58 percent of

[1] The Center for Substance Abuse Treatment defines co-occurring disorders as the presence of one or more mental disorders as well as one or more disorders relating to the use of alcohol and/or other drugs. A diagnosis of COD occurs when at least one disorder of each type can be established independently of the other and is not simply a cluster of symptoms resulting from the one disorder (SAMHSA, 2020, p. ix).

these individuals have undergone substance use treatment at some point in their lives, while 35 percent have been under treatment for SUD or intended to attend treatment while under supervision.

The advantage of automated knowledge discovery from probation documents is that the system can access information beyond what is listed in structured form in PACTS and can be re-applied to new data sets or scaled up to process larger data sets, at no additional cost. There are several projects that have applied AI technology to the identification of substance use or mental conditions in unstructured clinical text, especially within clinical notes fields in Electronic Health Records (EHR), but only a few studies have targeted comorbidity detection. These approaches can automatically extract a range of information on these conditions, allowing researchers to process larger data sets than a manual review would allow. To our knowledge, however, the current investigation is the first study to use NLP and AI methods on probation narrative text to automatically detect mental conditions, substance use, and comorbidity issues among individuals on supervision in the criminal justice system.[2] The system also generates more detailed information than in previous clinical works.

Background

Cohort Studies and Surveys

Each year, the National Survey on Drug Use and Health (NSDUH) collects data on a wide range of behavioral health issues from a representative sample of U.S. adults, including those under court-ordered supervision within the past year, specifically on probation or parole. Based on these self-reported responses by 201,400 individuals, the 2012 study estimated that among males 18-49, 40.3 percent of probationers and 38.3 percent of parolees had an alcohol or illicit drug use disorder in the previous year. With respect to substance use treatment, nearly half of male probationers and parolees needed treatment; however, only about a quarter of probationers and less than one third of parolees received some treatment in the previous year. About 10 percent of probationers reported that they were receiving treatment at the time of the survey, and about 3 to 7 percent had received treatment in prison or jail within the year (SAMHSA, 2014). For male individuals over 50 on probation or parole, Bryson et al. (2019) find that 21 percent of participants in the NSDUH survey reported a serious or moderate mental illness within the past year, and about 80 percent reported receiving some sort of mental health treatment.

These statistics show that the number of probationers and parolees with mental or substance use disorders whose treatment needs are not being met by community treatment and supportive services is significant. Yet, statistics for individuals under federal supervision who present with both SUD and COD are not readily available. Based on information in PACTS, Mangione (2019) finds that for post-conviction supervision, federal probation offices supervised 186,509 cases during fiscal year 2018. Of that number, 120,217 (64 percent) had substance abuse treatment conditions. During the same period, federal probation offices had 27,122 persons (14.5 percent) in substance use contract treatment.[3] The study adds that individuals with co-occurring disorders receive substance use and mental health services in an integrated fashion, but it does not provide statistics on these individuals.

A few studies have focused on the relationship between SUD with co-occurring mental illness and rates of recidivism on a smaller scale. Magee et al. (2021) conduct a retrospective cohort study of all individuals arrested in 2016 in Indianapolis, Indiana (N=22,939), by linking their arrest information with their clinical mental health and SUD diagnoses in the two years before the arrest. They found that 27.7 percent of the individuals in the study were formally diagnosed with SUD and 22.5 percent also had evidence of COD. The authors also found that individuals with SUD or co-occurring conditions in the preceding 2 years are at higher risk of repeat arrest, and they advocate interventions aimed at low-level offenders with behavioral health needs to prevent recidivism. Constantine et al. (2012) reach a similar conclusion after using a retrospective cohort design to study rearrest rate of inmates with serious mental illness diagnoses[4] in the Pinellas County, Florida, jail between July 1, 2003, and June 30, 2004, and their health and social service data from 2002 to 2006 (N=37,236). They find that 10.1 percent of the inmates from that period met the criteria. The authors argue that individuals with serious mental illness, especially with co-occurring SUD diagnosis, are at higher risk of felony rearrest, compared with other populations of inmates.

As this section demonstrates, previous studies of individuals with SUDs and CODs in the justice system have predominantly been limited to small-scale analyses in cohorts in specific counties, performed mostly manually, and relying almost exclusively on administrative data, from which results are then extrapolated to the general population. Analysis at the federal level typically depends on self-reported responses to the NSDUH survey questions, which does not emphasize cooccurrences of SUD and COD among the probationer and parolee population. Mangione (2019) is the most directly relevant study, yet it is limited to structured data provided in PACTS and does not have access to information on non-contact treatments or the social history information captured in free text form in probation documents.

Automated Approaches

Several efforts have focused on using automated NLP techniques, including machine learning, to extract smoking or substance use status (e.g., “Past smoker,” “Current smoker”). In these approaches, the authors typically use a pre-defined list of substance-related key phrases to identify text (paragraphs, sentences) containing potential substance use mentions from the notes, prior to processing the text through the NLP system. For instance, Uzuner (2008) describes several systems for classifying the smoking status of patients by using machine learning and rule-based algorithms, and reports F-scores ranging from 84 to 90.[5] Ni et al. (2021) developed an automated substance use detection system to identify substance use information in pediatric settings (N=3,890). Besides status (lifetime or current user), the system also detects substance

[2] Earlier versions of this study are described in Rowland et al (2018) and Megerdoomian et al (2019).

[3] “Contract” treatment refers to cases where Judiciary funds are used to pay facilities for treatment of individuals under supervision. PPSO will also frequently use treatment services that are available to the person under supervision in the community without cost to the federal Judiciary or through the individual’s own healthcare coverage. This is referred to as “noncontract” treatment. (Mangione, 2019).

[4] Serious mental illness includes schizophrenia, schizoaffective disorder, bipolar I disorder, major depressive disorder, other psychotic disorders, and other bipolar and mood disorders.

[5] F-measure or balanced F-score is a measure that combines precision and recall (harmonic mean). Precision (also called positive predictive value) is the fraction of relevant instances among the retrieved instances, while recall (also known as sensitivity) is the fraction of relevant instances that were retrieved.

categories (tobacco, marijuana, alcohol, opiates, any use) and agent (if family member was the user or the participant). The authors compare a knowledge-based model using logic rules and regular expressions and a deep-learning model trained on pre-annotated data and find that the deep-learning model performs better on most substance use categories and assertions, with sensitivity of 87.5 percent and specificity of 89 percent,[6] while their knowledge-based model outperforms the deep-learning model in detecting opiates use.

A couple of studies focused on identifying patients at increased risk of problem opioid use by applying NLP to electronic health records for patients receiving chronic opioid therapy (COT). Hylan et al. (2015) study chronic noncancer patients starting COT (N=2,752). Their algorithm addresses linguistic variation (different words with the same meaning), polysemy (single words with several meanings), negation (e.g., "reports no pain" vs. "reports severe pain"), ambiguity, and temporality. The algorithm resulted in a sensitivity of 60.1 percent and specificity of 71.6 percent. Carrell et al. (2015) also applied NLP to identify clinician entered descriptions of problem opioid use in the unstructured clinical notes of patients. They also capture terms that are negated, as well as terms qualified by uncertainty, historical reference, or reference to a person other than the patient. The false positive rate for patients identified by the NLP system was 41 percent. Authors conclude that human-assisted review of results is important for validation purposes.

Few studies have targeted both mental illness and substance use detection. Ridgeway et al. (2021) applied NLP to unstructured text sections of clinical notes in EHRs for HIV patients to detect mental illness and substance use among people living with HIV (N=778). The system performs keyword search using a list of pre-defined indicative words, negation terms, and regular expressions developed by subject matter experts. The study identified high rates of mental illness and substance use among patients in an urban HIV care clinic, nearly half of whom did not have a diagnosis code in the structured patient records, suggesting that relying on structured EHR fields alone to identify people with behavioral health disorders may miss a substantial number of patients. The NLP algorithm for detecting mental illness had a Positive Predictive Value (PPV)(Precision) of 98 percent and a Negative Predictive Value (NPV) of 98 percent. The NLP algorithm for detecting substance use had a PPV of 92 percent and an NPV of 98 percent.[7]

Wang et al. (2015) are one of the few studies that apply more in-depth analysis beyond regular expressions and logic rules and extract a wider range of attributes related to substance use. The system detects three main sub-categories of substance use (alcohol, drug, and nicotine use), but also extracts more fine-grained elements including amount, frequency, type (e.g., wine, alcohol, tobacco), status (current, past), method, and temporal information. The authors developed a knowledge-based NLP system that leverages substance use lexicons and annotated linguistic resources along with deep dependency parse relationships between tokens provided by the Stanford Dependency parser (Manning et al., 2014). The authors report the F-scores of 89.8, 84.6, and 89.4, respectively for alcohol, drug, and nicotine use statement detection. Performance on the extraction of attributes report average F-scores of 82.1 (amount), 90.3 (frequency), 80.8 (status), 88.7 (method), 96.6 (type), and 74.5 (temporal). The lower score on the temporal attribute is due to the variability in expression with temporal expressions.

As this review of the literature shows, NLP technologies have been used for extracting a range of information on substance use and mental illness from clinical notes. These approaches detect entities that correspond to Unified Medical Language System (UMLS) concepts (e.g., drugs, diseases, medications, or procedures). More recent approaches also detect features like negation and the experiencer or subject. Temporal analysis is typically limited to basic expressions that identify the status of use (current, past, lifetime), although more recent approaches have integrated analysis of more complex temporal expressions. The current study goes beyond previous works by building a system that can identify both substance use and mental conditions and related concepts, as well as mentions of treatments and diagnoses, including non-contract cases that are only discussed in narrative text. Furthermore, the system automatically detects all important dates and determines the status of treatments by leveraging textual information. The application of NLP allows the processing of larger data sets, combining structured PACTS information with knowledge discovered from relevant textual documents, delivering a comprehensive study of substance use and comorbidity detection among individuals in the criminal justice system.

Methods

System Overview

The Advanced Narrative Analytics System Infrastructure (ANAnSI) performs content extraction and detailed narrative analytics for knowledge discovery within a distributed high-performance system infrastructure. ANAnSI is a hybrid system that leverages linguistic resources including substance use and mental condition lexicons, and combines them with probabilistic algorithms as well as knowledge-based analytics to identify and extract rich event-based narrative analysis at the sentence level (i.e., *who did what to whom, where, and when* analysis). The system also uses linguistic knowledge in machine algorithms to perform reasoning tasks (e.g., temporal reasoning) and integrates machine learning-based components to make data-driven predictions (e.g., treatment outcome analysis). ANAnSI processes each sentence in the data collection and produces a detailed event-based analysis. Additional domain-specific analysis discovers properties relevant to substance use and mental illness. Table 1 shows the types of features automatically detected for each event, while Table 2 illustrates sample SUD treatment sentences with their corresponding analytic representation.

System analysis can also be viewed in terms of the relationships between each event and its participants, including the subject or agent of the action, the person affected by the action or the patient, and relations to temporal expressions. Figure 1 illustrates the analysis for the sentence in Table 2 where the Begun_by and Ended_by relations capture the start and end dates of the event. Shaded boxes represent discovered elements and arrows represent relationships between these elements.

[6] Sensitivity (True Positive Rate) refers to the proportion of those who received a positive result on this test out of those who actually have the condition, and Specificity (True Negative Rate) refers to the proportion of those who received a negative result on this test out of those who do not actually have the condition (when judged by the "Gold Standard").

[7] PPV and NPV in the clinical domain allow one to say how likely it is for a patient to have a specific disease. The positive predictive value is the probability that following a positive test result, that individual will truly have that specific disease. The negative predictive value is the probability that following a negative test result, that individual will truly not have that specific disease.

Data

The study focuses on 98,389 probation individuals under supervision as of October 2021. The information for this project was automatically obtained from the free text sections of Presentence Investigation Reports (PSIR), which represent investigations into the history of the person convicted of a crime before sentencing to determine if there are extenuating circumstances. More recent information was collected from psychological assessments and reports from treatment providers, and Chrono entries where probation officers record notes on office or home visits with the individual under supervision. All extracted events are associated with the offender's jurisdiction, criminal offense, and demographic information in the database for easy search and retrieval. A breakdown of the corpora used is shown in Table 3.

To prepare the data for processing, preprocessing steps are required.

1. Select corpus files: The PSIR documents and Chrono entries are explicitly tagged as such in PACTS and are easy to identify. However, psychological assessments may be classified under different document types. We therefore created a set of heuristic rules to automatically determine which documents should be treated as psychological evaluations by performing a keyword search on the Notes section of the PACTS metadata where users indicate additional information about the document type. For instance, if the Notes section contained terms such as "psy", "eval", "stable", "abel", "eval", "evl", "treatment", "trt", "ass" (for assessment) and did not include the terms "contract", "waiver", "receipt", or "no show", the documents were selected for analysis.
2. Extract text from PDF documents: The system applies generalized content extraction to the scanned and electronic PDF documents associated with the individuals under supervision. In addition, this component performs document structure analysis on the Presentence Investigation Reports to identify and parse out the different sections of the PDF documents and extracts the tabular profile and criminal information as well as all free text content per section. The following sections of the PSIR are predominantly used to extract relevant domain information: Mental and Emotional Health, Substance Abuse, Personal and Family Data, Juvenile Adjudications, Employment History, Education and Vocational History, Adult Criminal Convictions, Criminal History, and the PSIR cover page. This component further "cleans" the data by normalizing the textual content to maximize processing.

TABLE 1
Automatically Extracted Substance Use and Comorbidity Related Indicators

Feature	Value / Examples
Event type	Formal diagnosis, Substance use, Treatment
Status	Negated, non-negated
Mental condition	Substance Use Disorder, Anxiety, Schizoaffective disorder, etc.
Substance	Ethanol, Cocaine, Opioids, Promethazine, etc.
Usage method	Snort, drink, smoke, etc.
Temporal	Date, range ("since 2000"), duration ("for two months"), frequency ("twice a week")
Reporter	Self-report, medical professional, (medical) records, third party
Spatial	Facility, Organization, Location
Treatment type	Inpatient, outpatient
Procedure	Counseling, mental health treatment, rehabilitation, substance abuse treatment, etc.
Treatment outcome	Completed, discontinued, ongoing, participated, terminated, intended

TABLE 2
Representation of Automated Analysis for Sample Sentences

Sample Sentence	*Structured Representation of Analysis*
"The defendant reported he was unsuccessfully discharged from intensive outpatient substance abuse program at The Pyramid Rehabilitation Center, St. Louis, Missouri on December 10, 2012."	***EventType:*** Treatment ***Status:*** Non-negated ***MentalCondition:*** Substance Use Disorder ***Facility:*** The Pyramid Rehabilitation Center ***Location:*** St. Louis, Missouri ***Date:*** 12-10-2012 ***Reporter:*** self ***TreatmentType:*** Outpatient ***Procedure:*** Substance abuse program ***TreatmentOutcome:*** Terminated
"Information indicates that he was seen by Dr. Joseph Smith from January 30, 2011, until August 12, 2013."	***EventType:*** Treatment ***Status:*** Non-negated ***MentalCondition:*** N/A ***MedicalPerson:*** Dr. Joseph Smith ***StartDate:*** 1-30-2011 ***EndDate:*** 8-12-2013 ***Reporter:*** third party ***Procedure:*** Treatment

FIGURE 1
Event, Entity, and Relations Analysis

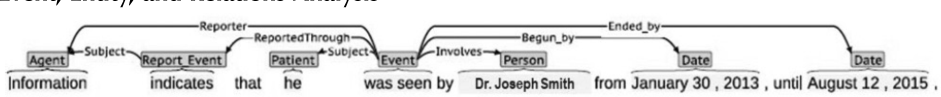

TABLE 3
Breakdown of Corpus for the Study

Documents Total	**9,721,569**
PSIR	121,518
Psychological Evaluation	133,067
Chronos	14,407,656
Sentences Total	**80,489,549**

N = *98,389*

Technical Approach

The ANAnSI architecture is provided in Figure 2. The *Information Extraction Component* takes the text extracted from PDF documents and Chrono notes as input and leverages open-source NLP tools for in-depth linguistic analysis and parsing. Stanford CoreNLP is a probabilistic system that performs entity recognition and sentence segmentation, detects the part-of-speech categories of each term, generates the dependency parse structure for each sentence, and detects temporal expressions (Manning et al., 2014). We used CoreNLP version 3.92 to generate dependency structures for all narrative statements. Apache cTAKES (clinical Text Analysis and Knowledge Extraction System) was developed specifically to extract and analyze clinical information from unstructured text (Savona et al., 2010).

The *Knowledge Discovery and Integration Component* builds on the results of the previous components to structure a complete analysis for events and their participants, to refine the recognition and classification of entities, and to infer the temporal relations between events in generating a timeline. This section also applies advanced linguistic analysis to improve argument and negation detection and improve precision of results. The complete list of entities employed by the system are person, role (e.g., doctor), location (city, state-or-province, country), facility, organization, temporal, money, salary. This component is able to detect and label the events in each sentence as well as related events such as a reporting event (*Medical records indicate that …*), an aspectual event that marks the end or beginning of the main event (*He began treatment in 1993*), or an event marking an intention (*Johnson stated that he would like to attend treatment for depression*). Finally, the system applies temporal reasoning techniques to infer complex temporal relations between events (e.g., *He became depressed after his infant brother died in 2000* infers that "became depressed" event began in or after the year 2000), compute temporal expressions with respect to the referred date (e.g., the diagnosis date can be computed based on the individual's date of birth in *He was diagnosed at age 20*, or the treatment date can be computed based on the document date in *She is currently undergoing treatment for anxiety* or *Medical records indicate that the defendant was terminated from the program three weeks ago*). In addition, this component links all temporal relations to obtain a complete temporal graph to capture the start and end dates of an event (e.g., *Jackson began treatment in March 2003 and was successfully discharged three months later* will be analyzed as having a start date of 2003-03-01 and an end date of 2003-06-01).

The resulting analyses are enhanced for the use case in the *Domain Analysis* component by identifying significant events and relations for the mental health and substance use domains. This component can detect paraphrases of mental condition mentions and substances at the appropriate level of granularity for PPSO. For example, terms such as *depression, chronic depression, depressive tendencies,* and *major depressive disorder* are all mapped to the more general term *depressive disorder.* In addition, it detects relevant events such as diagnoses, prescriptions, drug use, or treatments by identifying verbs commonly associated with these events (e.g., *attend, complete, hospitalize, undergo* are typically used to describe treatment events). This component also classifies the start or end of an event. For instance, verbs such as *enroll* or *enter* signal the beginning of a treatment event while *discharge* indicates the end of the treatment program. For each treatment event, the system also detects the treatment provider or facility mentioned in the sentence, the nature of treatment (e.g., inpatient or outpatient), and the procedure (e.g., anger management, drug rehabilitation). Any negated events are tagged as such.

As the system performs linguistic argument analysis to identify the participants in each event, it is possible to distinguish and ignore cases where a family member is mentioned rather than the individual under supervision (e.g., "*the defendant's mother suffered from Schizophrenia*"). The system also tags the source of the information (e.g., reported by a medical professional or self-reported). Note that a sentence like *The defendant denied smoking marijuana* is analyzed as a negated substance use event, but the reporter is tagged as "self-report." As descriptions of treatment events in probation documents can use very divergent wording to report the outcome, the system applies a trained Maximum Entropy machine learning algorithm to automatically classify the treatment outcome based on the categories shown in Table 4. The algorithm performs at 85 percent accuracy.

Extracted information from all documents associated with a particular individual is stored in the *Knowledge Model*, a Neo4j graph-based database management system, allowing all analyses to be compiled in a structured form with explicit links between related concepts and properties. The graph-based representation facilitates viewing all relevant information associated with a given individual, as well as obtaining an overview of all individuals under supervision in a specific district. The database provides the user with

FIGURE 2
Analytic Pipeline for ANAnSI

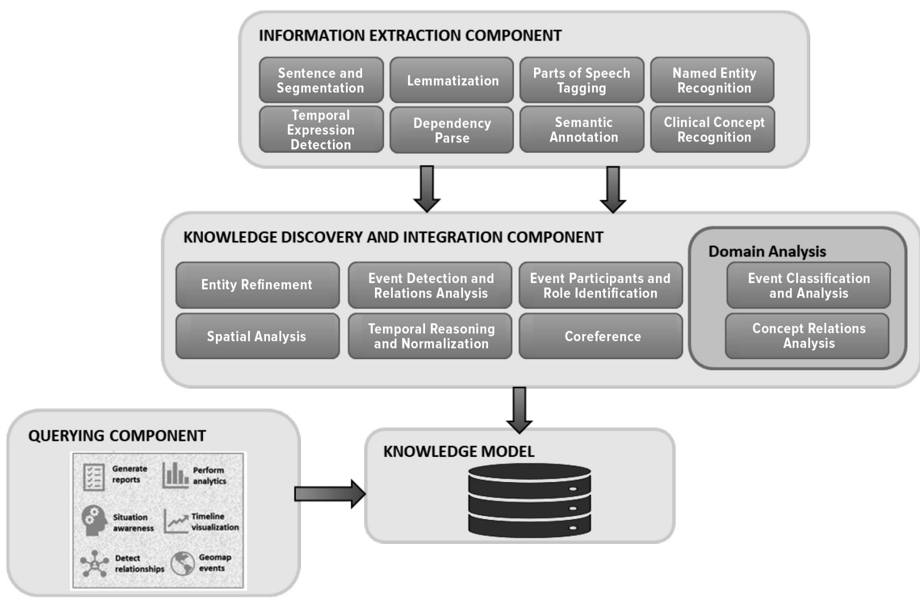

TABLE 4
Treatment Outcome Categories

Outcome Category	Description	Example
Participated	treatment occurred, but outcome not specified	"He received some substance abuse treatment while incarcerated as a teenager and as an adult."
Completed	treatment occurred and was successful	"He made 'good progress' in treatment and was successfully discharged on January 25, 2005."
Terminated	treatment occurred, but client was removed from the program before completion	"The defendant admitted he was discharged due to continued use while in treatment."
Discontinued	treatment occurred, but client left before completion	"He reported he was ordered into residential treatment but did not attend and absconded from supervision."
Ongoing	treatment is still taking place	"As a condition of his pretrial release, the defendant attends weekly drug treatment counseling sessions."
Intended	the client plans to enroll in treatment	"When asked, the defendant advised he would like to participate in the Residential Drug Abuse Program while incarcerated in the Bureau of Prisons."

a powerful query language to easily display answers to research questions, which can in turn be displayed in various formats, such as comma-separated value files, timeline view of events, or event relationship visualizations.

Results

Overview of Results

ANAnSI automatically processed over 62 million sentences from the documents associated with 98,389 individuals under supervision and identified events indicating formal SUD diagnoses, reports of substance use, presence of co-occurring disorders, and treatment participation information for the study group. Since recent events are of more interest to PPSO, the system also verifies if the detected event date or the document date falls within the supervision period for the individual. These results are shown in Figure 3 and details are presented in Table 5, contrasting all events in each class with the subset that was identified as occurring within the supervision period. We treat "successful" treatments as the ones that were not discontinued, terminated, or extended.

Results of the automated analysis show that about 93 percent of the individuals under supervision have had an issue with substance use at some point in their lives (see Diagnosed for SUD (formal & indefinite)). These include formal diagnoses for SUD as well as informal reports of substance use, either self-reported or reported by a medical professional or a third party. In contrast, only about 15 percent of the study group has received an official diagnosis at some point in their lives, according to the files analyzed. Note that these numbers reduce to 28 percent for all substance use reports and about 8 percent for formal diagnoses if limited to events within the supervision period. Furthermore, about 29 percent of the population under study has had at least one diagnosis for a co-occurring disorder with their substance use at some point in their life, and about 15 percent have received a COD diagnosis during their supervision period. Table 6 shows a comparison of the system results with the findings in previous research. As previously discussed, none of the previous study groups or research approaches correspond directly to the current analysis, yet a comparison of estimated values shows certain correlations.

The comparatively higher treatment scores found by ANAnSI reflect the treatment records during supervision as documented in the Chrono entries and include both contract and non-contract treatments. If we only take into account older information provided through PSIR or treatment reports, the percentage of individuals treated for SUD is lowered to 44 percent overall and 21 percent during their supervision period, which is closer to the rate discovered in previous studies. Similarly, the information on the individual's substance use history reflects data from the Chrono notes, while the information limited to PSIR and treatment documents finds that a smaller number (about 79 percent of the population under study overall and about 16 percent during the supervision period) has had a history with substance use.

Figure 4 illustrates the co-occurring disorders most often mentioned in the documents. These results correlate with findings in the literature where "mental disorders likely to co-occur with addiction include depressive disorders, bipolar I disorder, posttraumatic stress disorder (PTSD), personality disorders (PDs), anxiety disorders, schizophrenia and other psychotic disorders, ADHD, and eating and feeding disorders" (SAMHSA, 2020).

Meanwhile, Figure 5 shows the substances most often reported – about 74 percent of the individuals under supervision reported using alcohol (Ethanol), while 71 percent reported using marijuana (Marihuana) in their social histories or psychological evaluations. Other top substances include Cocaine, Methamphetamine, Ecstasy, and Heroin.[8]

FIGURE 3
Event Types Automatically Identified by ANAnSI

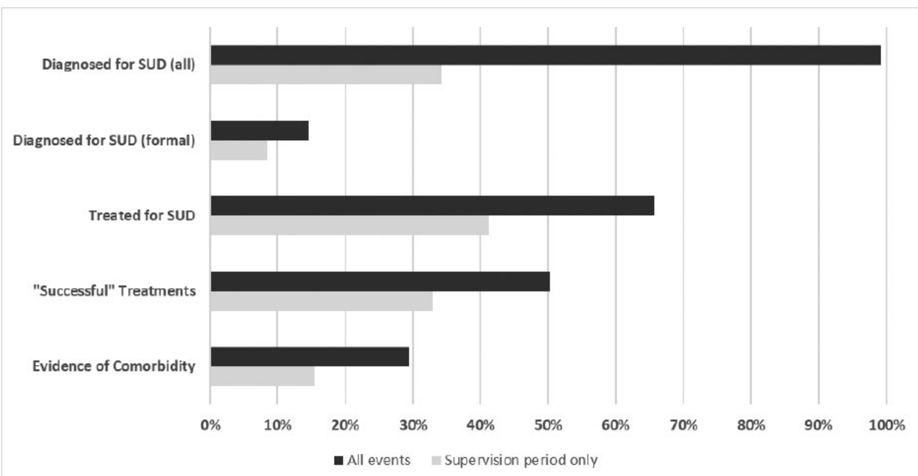

Percent values refer to the percentage out of the total number of clients (N=98,389).

TABLE 5
Results of Automated Event Detection for Substance Use and Co-occurring Mental Disorders

N=98,389 clients	All life events		During supervision period only	
	count	% of total	count	% of total
Diagnosed for SUD (formal)	14,321	14.56%	8,219	8.35%
Diagnosed for SUD (formal & indefinite)	91,077	92.57%	27,847	28.30%
Diagnosed for COD	28,288	28.75%	14,645	14.88%
Treated for SUD	57,507	58.45%	34,096	34.65%
Successful treatments	43,228	43.94%	26,581	27.02%

TABLE 6
Comparison of automated ANAnSI output with previous research findings

	SUD Diagnosis	Substance Use History	Treated for SUD	COD
ANAnSI (all)	14.56%	92.57%	58.45%	28.75%
ANAnSI (during supervision)	8.35%	28.30%	34.65%	14.88%
Magee et al., 2021	27.70%			22.50%
NSDUH, 2012	40.30%		25.00%	21.00%
Mangione, 2019		64.45%	14.54% (contract only)	
Sabol and Couture, 2008		60-80%		

[8] "Pharmaceutical Preparations" refer to general mention of the term "drugs" in the text.

FIGURE 4
Most Common Co-occurring Disorders Detected

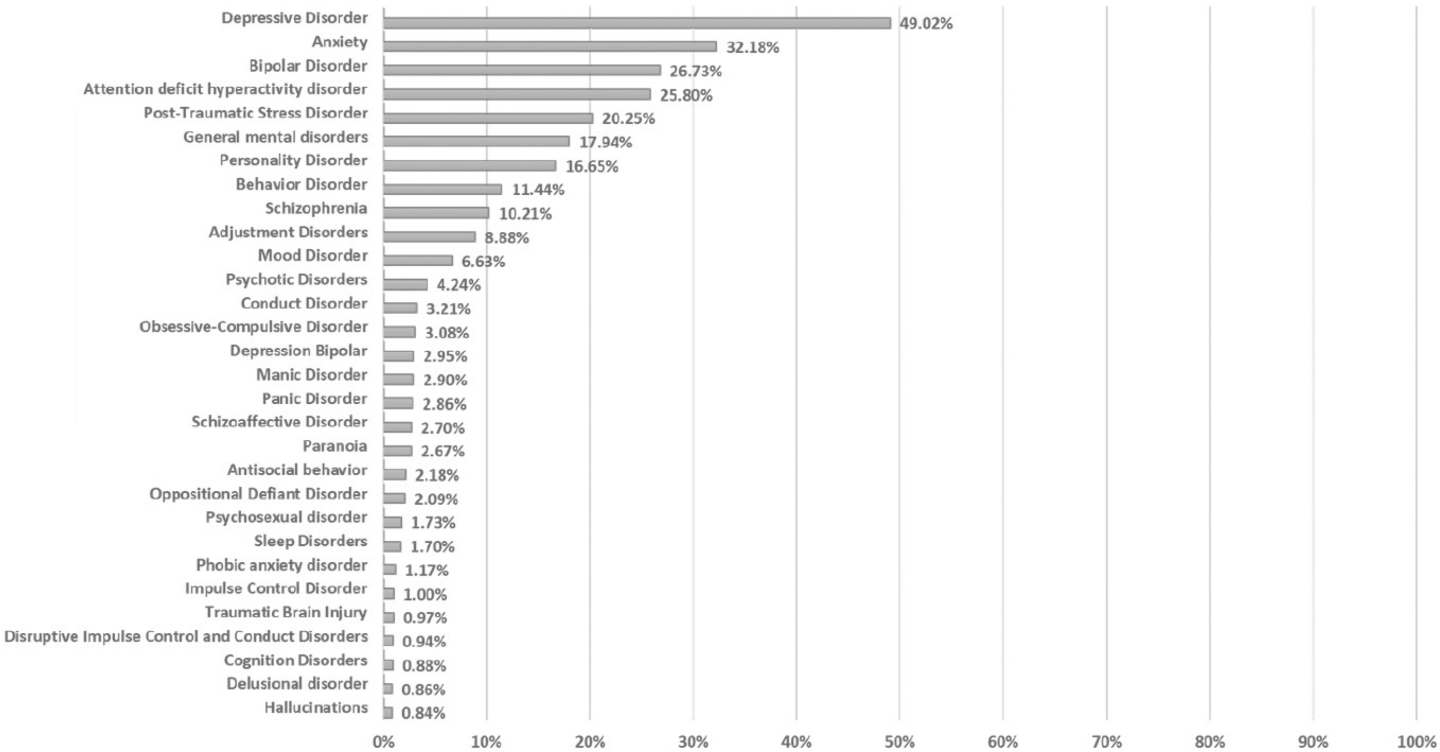

FIGURE 5
Percentage of Clients with SUD (all) that Used a Given Substance at Some Point in Their Life

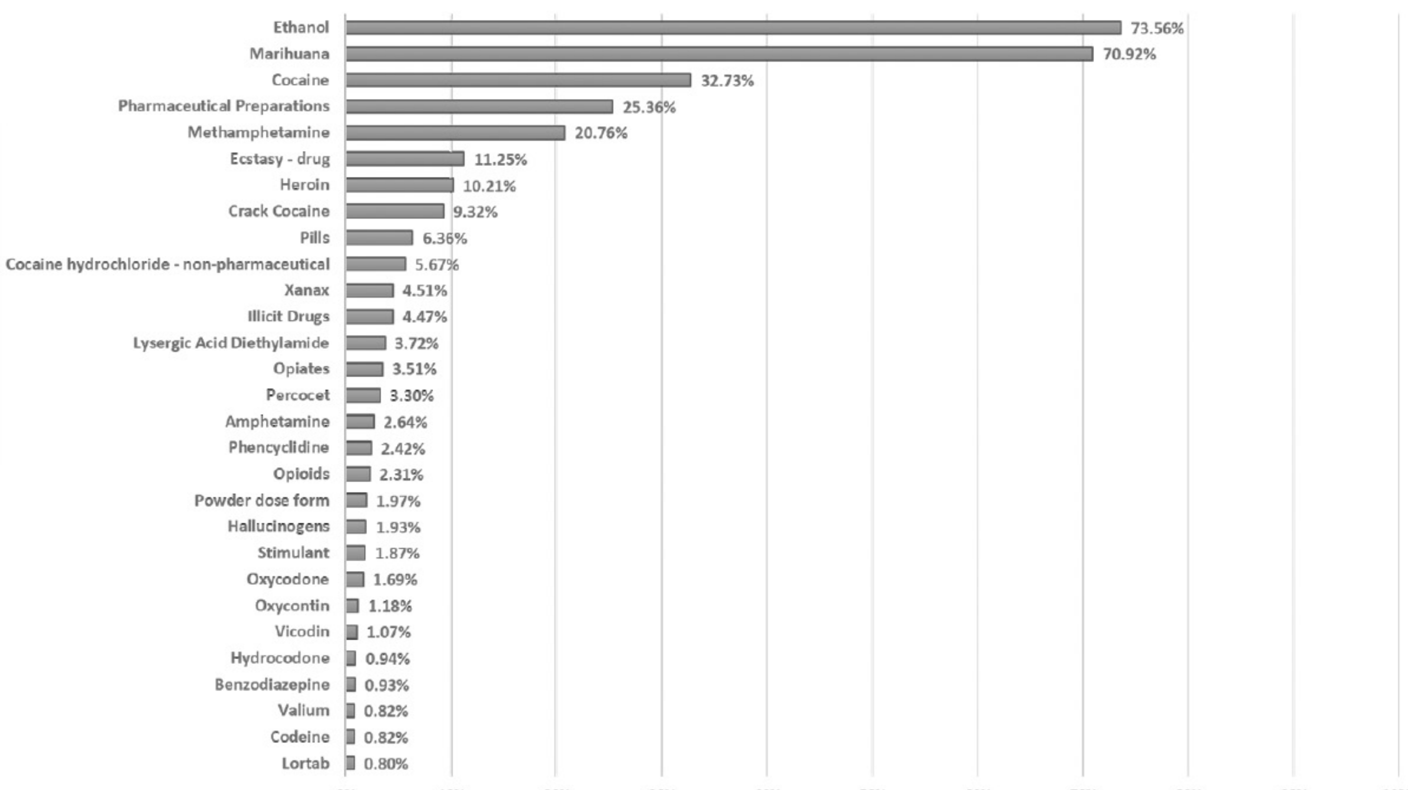

A study of the top treatment procedures results in an expected list where the majority of the population has participated in substance abuse treatment, substance abuse counseling, or rehabilitation programs, with a smaller group participating in a mental health treatment.

Performance Results

System performance was evaluated by creating a small reference sample of about 500 sentences from PSIR and treatment reports to measure the accuracy of the information extracted for each event type. The 500 sentences were manually annotated by team members for all event types and event attributes of interest. The language analytics results were then compared to the pre-annotated reference set to measure how many of the detected elements were accurate and to also calculate how many of the expected elements were not picked up by the system. System performance does not fare as highly when applied to the analysis of Chronos, however, given the more informal writing style and content, which often lacks full sentences and contains various shorthand as well as misspellings. (See Table 7.)

While this study makes an important contribution to advancing methods to extract substance use and mental condition information from text, there are limitations that could be addressed in future enhancements. The system does not consider same or co-referring events or substances repeated in distinct sentences. For instance, the sentences "*The defendant began substance use treatment in August 2010. He was successfully discharged in May 2012*" will be analyzed as depicting two distinct treatment events, instead of merging them as a single treatment with a begin date of 2010 and end date of 2012. Another challenge that was left unaddressed in the current version of ANAnSI is the distinction between events (e.g., diagnoses, treatments) that occurred in the past and those that are currently valid. This can be accomplished by leveraging the tense and aspect information that the system computes to enhance detection accuracy.

TABLE 7
Performance Results

Events & Attributes	Recall	Precision	F-measure
Diagnose	100	92.86	96.30
Usage	93.94	100	96.88
Treatment	90.91	100	95.24
Mental condition	94.12	100	96.97
Reporter	75	84.38	79.41
Date	76	73.1	74.52
Subject of event	100	98.28	99.13
Treatment outcome	100	70	82.35
Treatment procedure	81.82	100	90.00

Conclusion

This article describes a successful approach to the automatic extraction and analysis of probation narrative text in the mental health and substance use domain. The results provide evidence that the use of technology in identifying important information in free narrative text in administrative records is feasible and cost-effective, and any adaptations to new domains can be accelerated through probabilistic methods. These analytics can be further developed in various directions, depending on the mission needs of the organization.

Acknowledgments

We would like to acknowledge the guidance and support of the U.S. Probation and Pretrial Services Office throughout this project. The project was led and funded by the U.S. Probation and Pretrial Services Office and the Technology Solutions Office at the Administrative Office of the U.S. Courts as part of the Applied Technology Research and Development program and managed by the Judiciary & Legal Technology Modernization Department at the MITRE Corporation.

This (software/technical data) was produced for the U.S. Government under Contract Number USCA16D2019, and is subject to Judiciary Policy Clause 6-60, Rights in Data—General (JUN 2012).

No other use other than that granted to the U.S. Government, or to those acting on behalf of the U.S. Government under that Clause is authorized without the express written permission of The MITRE Corporation.

For further information, please contact The MITRE Corporation, Contracts Management Office, 7515 Colshire Drive, McLean, VA 22102-7539 (703) 983-6000.

Data Availability Statement

The data underlying this article are from the U.S. Probation and Pretrial Services Office–funded project to study substance use and co-occurring disorders among individuals under supervision. The data cannot be shared publicly due to protected health information of individuals that participated in the study.

References

Bryson, W. C., Cotton, B. P., Barry, L.C., Bruce, M. L., Piel, J., Thielke, S. M., and Williams, B. A. (2019). Mental health treatment among older adults with mental illness on parole or probation. *Health and Justice* 7(4).

Carrell, D. S., Cronkite, D., Palmer, R., Saunders, K., Gross, D. E., Masters, E. T., Hylan, T. R., and Von Korff, M. (2015). Using natural language processing to identify problem usage of prescription opioids. *International Journal of Medical Informatics*, 84(12):1057-1064.

Constantine, R. J., Robst, J., Andel, R., and Teague, G. (2012). The impact of mental health services on arrests of offenders with a serious mental illness. *Law and Human Behavior*, 36(3): 170-176.

Drake, R. E., and Wallach, M. A. (1989). Substance abuse among the chronic mentally ill. *Hospital and Community Psychiatry*, 40(10):1041-6.

Fazel, S., Hayes, A. J., Bartellas, K. Clerici, M., and Trestman, R. (2016). Mental health of prisoners: Prevalence, adverse outcomes, and interventions. *Lancet Psychiatry*, 3(9):871-81.

Fearn, N. E., Vaugh, M. G., Nelson, E. J., Salas-Wright, C. P., DeLisi, M,. and Qian, Z. (2016). Trends and correlates of substance use disorders among probationers and parolees in the United States, 2002–2014. *Drug and Alcohol Dependence*, 167:128-139.

Feucht, T. E., and Gfroer, J. (2011). Mental and substance use disorders among adult men on probation or parole: Some success against a persistent challenge.

Hylan, T. R., Von Korff, M., Saunders, K., Masters, El, Palmer, R. E., Carrell, D., Cronkite, D., Mardekian, J., and Gross, D. (2015). Automated prediction of risk for problem opioid use in a primary care setting. *The Journal of Pain*, 16(4):380-387.

Magee, L. A., Ranney, M. L., Fortenberry, J. D., Rosenman, M., Gharbi, S., and Wiehe, S. E. (2021). Identifying nonfatal firearm assault incidents through linking police data and clinical records: Cohort study in Indianapolis, Indiana, 2007–2016. *Preventive Medicine,* 149.

Mangione, C. (2019). Overview of substance use disorder occurrence and treatment in the Federal Judiciary. *Federal Probaation*

Journal, 83(2).

Manning, C. D., Surdeanu, M., Bauer, J., Finkel, J., Bethard, S. J., and McClosky, D. (2014). The Stanford CoreNLP Natural Language Processing Toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pp. 55-60.

Megerdoomian, K., Branting, K., Horowitz, C. E., Marsh, A. B., Modly, N., Petersen, S. J., Scott, E. O., and Wariyar, S. B. (2019). Automated narrative extraction from administrative records. *Proceedings of the Workshop on Artificial Intelligence and the Administrative State (AIAS 2019)*, Montreal, QC, Canada.

Monahan, J. (1992). Mental disorder and violent behavior: Perceptions and evidence. *American Psychologist*, 47 (4), 511–521.

Ni, Y., Bachtel, A., Nause, K., and Beal, S. (2021). Automated detection of substance use information from electronic health records for a pediatric population. *Journal of the American Medical Informatics Association*, 28(10):2116-2127.

Ridgeway, J. P., Uvin, A., Schmitt, J., Oliwa, T., Almirol, E., Devlin, S., and Schneider, J. (2021). Natural language processing of clinical notes to identify mental illness and substance use among people living with HIV: Retrospective cohort study. *JMIR Medical Informatics,* 9(3).

Rowland, M. G. (2018). Federal probation and pretrial services: What's going on and where are we going? (Presentation by the Chief of PPSO). Probation and Pretrial Services Office, Administrative Office of the U.S. Courts.

Sabol, W., and Couture, H. (2008). Prison Inmates at Midyear 2007. *Prison and Jail Inmates at Midyear.* Bureau of Justice Statistics Bulletin, June 2008, NCJ 221944.

SAMHSA (2014). The NSDUH report: Trends in substance use disorders among males aged 18 to 49 on probation or parole. March 6, 2014. Retrieved at https://www.samhsa.gov/data/sites/default/files/sr084-males-probation-parole/sr084-males-probation-parole/sr084-males-probation-parole.htm.

SAMHSA (2020). Substance use disorder treatment for people with co-occurring disorders, TIP42.

SAMHSA (2022). Key substance use and mental health indicators in the United States: Results from the 2022 National Survey on Drug Use and Health. Retrieved at https://www.samhsa.gov/data/sites/default/files/reports/rpt42731/2022-nsduh-nnr.pdf.

Savova, G. K., Masanz, J. J., Ogren, P. V., Zheng, J., Sohn, S., Kipper-Schuler, K. C., and Chute, C. G. (2010). Mayo clinical Text Analysis and Knowledge Extraction System (cTAKES): Architecture, component evaluation and applications. *Journal of American Medical Informatics Association* 17: 507–513.

Stahler, G. J., Mennis, J., Belenko, S., Welsh, W. N., Hiller, M. L., and Zajac, G. (2013). Predicting recidivism for released state prison offenders: Examining the influence of individual and neighbourhood characteristics and spatial contagion on the likelihood of reincarceration. *Criminal Justice and Behaviour,* 40(6):690-711.

Uzuner, Ö, Goldstein, I., Luo, Y., and Kohane, I. (2008). Identifying patient smoking status from medical discharge records. *Journal of the American Medical Informatics Association*, 15(1):14–24.

Wang, Y., Chen, E. S., Pakhomov, S., Arsoniadis, E., Carter, E. W., Lindemann, E, Sarkar, I. N., and Melton, G. B. (2015). Automated extraction of substance use information from clinical texts. *AMIA Annual Symposium Proceedings*, 2121-2130.