

C Detailed Task-specific Static Prompts

Prompt Strategies	Reddit-Impacts
Basic Prompt	<p>[Task Description]: You are a medical AI trained to identify and classify tokens into three categories: Clinical Impacts, Social Impacts, and Outside ('O'). Your task is to extract and classify the clinical and social impacts from this dataset, considering your knowledge of the lifestyle of this population and the potential clinical and social impacts they might experience.</p>
	<p>[Entity Types with Definitions]: 'Clinical Impacts' refer to tokens describing the effects, consequences, or impacts of substance use on individual health or well-being, as defined in UMLS. 'Social Impacts' describe the societal, interpersonal, or community-level effects, also based on UMLS definitions. Any token not falling into these categories should be labeled as 'O'.</p>
	<p>[Format Specification]: For example, the sentence 'I was a codeine addict.' is tokenized and labeled as follows: ['I', 'was', 'a', 'codeine', 'addict', '.'] with labels ['O', 'O', 'O', 'Clinical Impacts', 'Clinical Impacts', 'O']. Your task is to predict and return the label for each provided token, ensuring the number of output labels matches the number of input tokens exactly. The output format should be tokens with their labels: ['I-O', 'was-O', 'a-O', 'codeine-Clinical Impacts', 'addict-Clinical Impacts', '.-O'].</p>
Description of datasets	The data you are working with has been collected from 14 forums on Reddit (subreddits) that focused on prescription and illicit opioids, and medications for opioid use disorder. This dataset represents a social media context, coming from individuals who may use prescription and illicit opioids and stimulants.
High-frequency instances	In this dataset, high-frequency clinical impacts include 'withdrawal', 'rehab', 'addicted', 'detox', 'overdosed', and 'rehab'. High-frequency social impacts include 'lost', 'homeless', 'charged', 'streets', 'jail', and 'disorderly'.
UMLS knowledge	The Unified Medical Language System (UMLS) is developed by the U.S. National Library of Medicine (NLM) to integrate and standardize diverse medical terminologies and coding systems. It consists of three main components: the Metathesaurus, Semantic Network, and SPECIALIST Lexicon, supporting medical information retrieval and semantic analysis. You understand medical terminology and concepts from UMLS.
Error analysis	Possible analysis of prediction errors: If a sentence describes the background information of an event, facility, or project, then even if it mentions keywords related to social impact like 'at jail', it still cannot be determined as describing a patient being in jail. It is essential to clearly determine whether the sentence is describing the patient's condition. Second, if the sentence is about the usage, operation, or introduction of a drug or medicine, it does not belong to the patient's clinical impacts, even if it mentions some symptoms. Pay attention to whether the sentence contains words like 'if' that indicate conditions.

Table 8. Specific static prompts for each component we used for the REDDIT-IMPACTS dataset.

Prompt Strategies	BC5CDR
	<p>[Task Description]: You are a medical AI trained to identify and classify tokens into three categories: 'Disease', 'Chemical' and Outside ('O'). Your task is to extract and classify the Disease and Chemical related concepts from this dataset.</p> <p>[Entity Types with Definitions]: 'Disease' is a particular abnormal condition that adversely affects the structure or function of all or part of an organism and is not immediately due to any external injury. Diseases are often known to be medical conditions that are associated with specific signs and symptoms. A disease may be caused by external factors such as pathogens or by internal dysfunctions. For example, internal dysfunctions of the immune system can produce a variety of different diseases, including various forms of immunodeficiency, hypersensitivity, allergies, and autoimmune disorders. 'Chemical' in this context refers to substances or compounds with specific chemical properties and structures. These can include drugs, neurotransmitters, elements or ions, vitamins, and other medically relevant chemicals. Any token not falling into Disease categories should be labeled as 'O'.</p> <p>[Format Specification]: For example, the sentence 'The hypotensive effect of 100 mg / kg alpha-methyldopa was also partially reversed by naloxone.' is tokenized and labeled as follows: ['The', 'hypotensive', 'effect', 'of', '100', 'mg', '/', 'kg', 'alpha-methyldopa', 'was', 'also', 'partially', 'reversed', 'by', 'naloxone', '.'] with labels ['O', 'Disease', 'O', 'O', 'O', 'O', 'O', 'O', 'Chemical', 'O', 'O', 'O', 'O', 'O', 'Chemical', 'O']. Your task is to predict and return the label for each provided token, ensuring the number of output labels matches the number of input tokens exactly. The output format should include tokens with their labels: ['The-O', 'hypotensive-Disease', 'effect-O', 'of-O', '100-O', 'mg-O', '/-O', 'kg-O', 'alpha-methyldopa-Chemical', 'was-O', 'also-O', 'partially-O', 'reversed-O', 'by-O', 'naloxone-Chemical', '.-O'].</p>
Basic Prompt	
Description of datasets	<p>The data you are working with is BC5CDR dataset, a benchmark dataset for biomedical natural language processing, created from PubMed abstracts. It includes annotations for two entity types—chemicals and diseases—and their relationships, specifically chemical-induced disease interactions. The dataset is widely used for tasks such as named entity recognition and relation extraction, supporting research in biomedical text mining and information extraction.</p>
High-frequency instances	<p>In this dataset, high frequency of 'Disease' include 'pain', 'toxicity', 'renal', 'failure', 'disease', 'hypotension'; high frequency of 'Chemical' include 'cocaine', 'acid', 'dopamine', 'nicotine', 'morphine', 'lithium'.</p>
UMLS knowledge	<p>The Unified Medical Language System (UMLS) is developed by the U.S. National Library of Medicine (NLM) to integrate and standardize diverse medical terminologies and coding systems. It consists of three main components: the Metathesaurus, Semantic Network, and SPECIALIST Lexicon, supporting medical information retrieval and semantic analysis. You understand medical terminology and concepts from UMLS.</p>
Error analysis	<p>Possible analysis of prediction errors: The prediction errors mainly stem from challenges in distinguishing between entity boundaries and contextual usage. For instance, multi-token entities were partially labeled, causing boundary mismatches. Additionally, certain terms such as "receptor" or "antagonist" were incorrectly labeled as 'O', despite being part of chemical-related entities. Misclassification also occurred in sentences with conditional phrases or background information, where the relation between entities was not accurately captured. Furthermore, entities mentioned in descriptive or abstract contexts, were sometimes overlooked. These errors highlight difficulties in handling complex sentence structures, context-specific classification, and multi-token entity recognition.</p>

Table 9. Specific static prompts for each component we used for the BC5CDR dataset.

Prompt Strategies	MIMIC III
Basic Prompt	<p>[Task Description]: You are a medical AI trained to identify and classify tokens into 13 categories: 'CONDITION/SYMPOM', 'DRUG', 'AMOUNT', 'TIME', 'MEASUREMENT', 'LOCATION', 'EVENT', 'FREQUENCY', 'ORGANIZATION', 'DATE', 'AGE', 'GENDER' and Outside ('O'). Your task is to extract and classify the concepts from this dataset.</p> <p>[Entity Types with Definitions]: 'ORGANIZATION' refers to entities or groups associated with healthcare or emergency medical services. These could be specific departments, teams, or services within a medical or emergency response organization. 'DATE' in this context refers to specific calendar dates. These dates are typically used to mark particular events, appointments, or deadlines. 'AGE' in this context refers to the length of time that a person has lived or the number of years since their birth. It can be expressed in various formats, including numerical values, abbreviated forms, or written out in words. 'GENDER' in this context refers to the socially constructed roles, behaviors, activities, and attributes that a given society considers appropriate for men and women. It encompasses the identities of 'male' and 'female,' which are often associated with biological sex but are also shaped by cultural and social factors. 'FREQUENCY' in this context refers to the rate or regularity at which an event or phenomenon occurs. It can describe how often something happens, ranging from sporadic or irregular occurrences to more regular or constant patterns. 'EVENT' in this context refers to specific occurrences or actions that take place, particularly in a medical or clinical setting. These can include procedures, assessments, or other significant incidents. 'LOCATION' in this context refers to specific places or areas, particularly within a healthcare or medical setting. These can include types of facilities, specific locations within a facility, or other relevant places. 'MEASUREMENT' in this context refers to quantitative assessments or values used to evaluate specific physiological or medical parameters. These can include vital signs, laboratory test results, numerical values, or other metrics related to patient health. 'TIME' in this context refers to specific points or periods in the temporal continuum, particularly as they relate to healthcare or medical events. These can include general time references, specific durations, or events tied to time. 'AMOUNT' in this context refers to specific quantities or dosages, particularly in a medical or pharmaceutical setting. These can include measurements of medication, frequency or number of administrations, and methods of delivery. 'DRUG' in this context refers to specific medications or pharmaceutical substances used in the treatment, prevention, or diagnosis of diseases. These can include brand names, generic names, or forms of administration. 'CONDITION/SYMPOM' in this context refers to physical or subjective signs that indicate a medical condition or disease. These can include sensations of discomfort, specific types of pain or discomfort, respiratory issues, or gastrointestinal symptoms. Any token not falling into categories above should be labeled as 'O'.</p> <p>[Format Specification]: For example, the sentence "The patient was readmitted to the hospital on 2195-6-6 due to fevers to 103 at the rehabilitation facility despite being on intravenous antibiotics HISTORY OF PRESENT ILLNESS 55 year-old female presents with 2/5 week history of non-bloody diarrhea" is tokenized and labeled as follows: ['The', 'patient', 'was', 'readmitted', 'to', 'the', 'hospital', 'on', '2195-6-6', 'due', 'to', 'fevers', 'to', '103', 'at', 'the', 'rehabilitation', 'facility', 'despite', 'being', 'on', 'intravenous', 'antibiotics', 'HISTORY', 'OF', 'PRESENT', 'ILLNESS', '55', 'year-old', 'female', 'presents', 'with', '2/5', 'week', 'history', 'of', 'non-bloody', 'diarrhea'] with labels ['O', 'O', 'O', 'EVENT', 'O', 'LOCATION', 'LOCATION', 'O', 'O', 'O', 'O', 'MEASUREMENT', 'MEASUREMENT', 'MEASUREMENT', 'O', 'LOCATION', 'LOCATION', 'LOCATION', 'O', 'O', 'O', 'DRUG', 'DRUG', 'O', 'O', 'O', 'O', 'AGE', 'O', 'GENDER', 'O', 'O', 'AMOUNT', 'AMOUNT', 'O', 'O', 'CONDITION/SYMPOM', 'CONDITION/SYMPOM']. Your task is to predict and return the label for each provided token, ensuring the number of output labels matches the number of input tokens exactly. The output format should include tokens with their labels: ['The-O', 'patient-O', 'was-O', 'readmitted-EVENT', 'to-O', 'the-LOCATION', 'hospital-LOCATION', 'on-O', '2195-6-6-O', 'due-O', 'to-O', 'fevers-MEASUREMENT', 'to-MEASUREMENT', '103-MEASUREMENT', 'at-O', 'the-LOCATION', 'rehabilitation-LOCATION', 'facility-LOCATION', 'despite-O', 'being-O', 'on-O', 'intravenous-DRUG', 'antibiotics-DRUG', 'HISTORY-O', 'OF-O', 'PRESENT-O', 'ILLNESS-O', '55-AGE', 'year-old-O', 'female-GENDER', 'presents-O', 'with-O', '2/5-AMOUNT', 'week-AMOUNT', 'history-O', 'of-O', 'non-bloody-CONDITION/SYMPOM', 'diarrhea-CONDITION/SYMPOM'].</p>
	<p>The data you are working with is MIMIC-III (Medical Information Mart for Intensive Care) dataset, a large, publicly available database containing de-identified health data from critical care patients at the Beth Israel Deaconess Medical Center. It includes structured data, such as demographics, lab results, and vital signs, as well as unstructured data, such as clinical notes and discharge summaries. The dataset is widely used for research in machine learning, natural language processing, and clinical decision support to improve healthcare outcomes.</p>
	<p>In this dataset, high frequency of 'CONDITION/SYMPOM' include 'pain', 'chest', 'cough', 'breath', 'nausea', 'abdominal'; high frequency of 'DRUG' include 'iv', 'lasix', 'ceftriaxone', 'oxygen', 'ns', 'coumadin'; high frequency of 'AMOUNT' include 'iv', '2', '1', 'mg', 'days', 'one'; high frequency of 'TIME' include 'day', 'admission', 'prior', 'last', 'ago', 'morning'; high frequency of 'MEASUREMENT' include 'bp', 'hr', 'pressure', 'blood', 'rr', 'rate', 'heart'; high frequency of 'LOCATION' include 'hospital', 'right', 'home', 'floor', 'emergency', 'micu'; high frequency of 'EVENT' include 'ct', 'placed', 'cxr', 'intubated', 'exam', 'review'; high frequency of 'FREQUENCY' include 'chronic', 'intermittent', 'daily', 'occasionally', 'frequent', 'intermittently'; high frequency of 'ORGANIZATION' include 'ems', 'service', 'surgery', 'pcp', 'emergency', 'neuro', 'medicine'; high frequency of 'DATE' include '2171114', '21491117'; high frequency of 'AGE' include '60', '80yo', '78', '61', 'seventyeightyearold', '69'; high frequency of 'GENDER' include 'man', 'woman', 'f', 'male', 'female', 'm'.</p>
	<p>The Unified Medical Language System (UMLS) is developed by the U.S. National Library of Medicine (NLM) to integrate and standardize diverse medical terminologies and coding systems. It consists of three main components: the Metathesaurus, Semantic Network, and SPECIALIST Lexicon, supporting medical information retrieval and semantic analysis. You understand medical terminology and concepts from UMLS.</p>
	<p>The prediction errors stem from several factors. Entity boundary recognition issues were common, particularly with multi-token entities like "shortness of breath" or "paroxysmal nocturnal dyspnea," where some tokens were missed or incorrectly labeled as 'O.' Additionally, the model struggled with entity type confusion, such as distinguishing "pain" as a symptom versus its contextual use related to location. Context-dependent misinterpretations also contributed to errors, especially in handling negations like "denies chest pain" or temporal references such as "last few months." Overlapping entities posed further challenges, where closely related terms (e.g., "MI" and "CABG") interfered with accurate classification. Finally, rare or unseen entities in the training data led to occasional misclassifications, highlighting gaps in the model's ability to generalize.</p>

Table 10. Specific static prompts for each component we used for the MIMIC III dataset.

Prompt Strategies	NCBI
	<p>[Task Description]: You are a medical AI trained to identify and classify tokens into five categories: DiseaseClass, SpecificDisease, Modifier, CompositeMention and Outside ('O'). Your task is to extract and classify the DiseaseClass, SpecificDisease, Modifier and CompositeMention from this dataset.</p> <p>[Entity Types with Definitions]: 'DiseaseClass' refers to a classification system or category used to group various medical conditions or diseases based on certain characteristics, such as their nature, affected biological systems, or underlying causes. 'SpecificDisease' appears to describe particular diseases that are identified and classified based on their specific clinical features, genetic origins, or biochemical abnormalities. 'Modifier' refers to specific attributes or variations or conditions that can modify or influence the presentation, progression, or characteristics of a disease, alter the manifestation or course of a disease, potentially affecting its diagnosis, treatment, and prognosis. 'CompositeMention' describes medical conditions or characteristics that are composed of several elements or features, often involving multiple tissues, organs, or systems. Any token not falling into these categories should be labeled as 'O'.</p> <p>[Format Specification]: For example, the sentence 'Histidinemia. Classical and atypical form in siblings.' is tokenized and labeled as follows: ['Histidinemia.', 'Classical', 'and', 'atypical', 'form', 'in', 'siblings.']. with labels ['SpecificDisease', 'O', 'O', 'O', 'O', 'O', 'O']. Your task is to predict and return the label for each provided token, ensuring the number of output labels matches the number of input tokens exactly. The output format should include tokens with their labels: ['Histidinemia.-SpecificDisease', 'Classical-O', 'and-O', 'atypical-O', 'form-O', 'in-O', 'siblings.-O'].</p>
Basic Prompt	
Description of datasets	<p>The data you are working with is NCBI disease corpus, a collection of 793 PubMed abstracts fully annotated at the mention and concept level to serve as a research resource for the biomedical natural language processing community. Each PubMed abstract was manually annotated by two annotators with disease mentions and their corresponding concepts in Medical Subject Headings (MeSH) or Online Mendelian Inheritance in Man (OMIM). The public release of the NCBI disease corpus contains 6892 disease mentions, which are mapped to 790 unique disease concepts. Of these, 88 percent link to a MeSH identifier, while the rest contain an OMIM identifier. We were able to link 91 percent of the mentions to a single disease concept, while the rest are described as a combination of concepts.</p>
High-frequency instances	<p>In this dataset, high-frequency 'DiseaseClass' include 'disorder', 'abnormalities', 'tumors', 'mental', 'disorders', 'retardation'. High-frequency 'SpecificDisease' include 'deficiency', 'syndrome', 'dystrophy', 'familial', 'myotonic', 'colorectal'. High-frequency 'Modifier' include 'tumor', 'tumour', 'APC', 'choroideremia', 'DM', 'DMD'. High-frequency 'CompositeMention' include 'breast', 'ovarian', 'cancer', 'muscular', 'andor', 'becker'.</p>
UMLS knowledge	<p>The Unified Medical Language System (UMLS) is developed by the U.S. National Library of Medicine (NLM) to integrate and standardize diverse medical terminologies and coding systems. It consists of three main components: the Metathesaurus, Semantic Network, and SPECIALIST Lexicon, supporting medical information retrieval and semantic analysis. You understand medical terminology and concepts from UMLS.</p>
Error analysis	<p>The prediction errors in the NCBI dataset primarily stem from challenges in distinguishing composite mentions and modifiers within complex biomedical contexts. For instance, entities like "BRCA1 gene" were incorrectly segmented, with "BRCA1" labeled as a modifier instead of being part of the composite mention. Additionally, multi-token composite mentions such as "breast and ovarian cancer" were not consistently labeled, with individual tokens occasionally missed or misclassified. Contextual ambiguity, such as distinguishing between mentions of general biological terms (e.g., "tumor") and their specific functional roles (e.g., "tumor suppressor"), also contributed to errors.</p>

Table 11. Specific static prompts for each component we used for the NCBI dataset.

Prompt Strategies	Med-Mentions
	<p>[Task Description]: You are a medical AI trained to identify and classify tokens into two categories: Disease and Outside ('O'). Your task is to extract and classify the Disease related concepts from this dataset.</p> <p>[Entity Types with Definitions]: 'Disease' is a particular abnormal condition that adversely affects the structure or function of all or part of an organism and is not immediately due to any external injury. Diseases are often known to be medical conditions that are associated with specific signs and symptoms. A disease may be caused by external factors such as pathogens or by internal dysfunctions. For example, internal dysfunctions of the immune system can produce a variety of different diseases, including various forms of immunodeficiency, hypersensitivity, allergies, and autoimmune disorders. Any token not falling into Disease categories should be labeled as 'O'.</p> <p>[Format Specification]: For example, the sentence 'A total of 200 children and adolescents with type 1 diabetes, ages 9-18 years, completed the DEPS-R Turkish version.' is tokenized and labeled as follows: ['A', 'total', 'of', '200', 'children', 'and', 'adolescents', 'with', 'type', '1', 'diabetes', 'ages', '9-18', 'years', 'completed', 'the', 'DEPS-R', 'Turkish', 'version.']. with labels ['O', 'O', 'O', 'O', 'Disease', 'O', 'Disease', 'O', 'Disease', 'Disease', 'Disease', 'Disease', 'O', 'Disease', 'O', 'O', 'Disease', 'Disease', 'Disease']. The output format should include tokens with their labels: ['A-O', 'total-O', 'of-O', '200-O', 'children-Disease', 'and-O', 'adolescents-Disease', 'with-O', 'type-Disease', '1-Disease', 'diabetes.-Disease', 'ages-Disease', '9-18-O', 'years.-Disease', 'completed-O', 'the-O', 'DEPS-R-Disease', 'Turkish-Disease', 'version.-Disease'].</p>
Basic Prompt	
Description of datasets	<p>The data you are working with is Med-Mentions, a new manually annotated resource for the recognition of biomedical concepts. What distinguishes Med-Mentions from other annotated biomedical corpora is its size (over 4,000 abstracts and over 350,000 linked mentions), as well as the size of the concept ontology (over 3 million concepts from UMLS 2017) and its broad coverage of biomedical disciplines.</p>
High-frequency instances	<p>In this dataset, high frequency 'Disease' related entities include 'patients', 'cells', 'treatment', 'cancer', 'analysis', 'disease', 'clinical'.</p>
UMLS knowledge	<p>The Unified Medical Language System (UMLS) is developed by the U.S. National Library of Medicine (NLM) to integrate and standardize diverse medical terminologies and coding systems. It consists of three main components: the Metathesaurus, Semantic Network, and SPECIALIST Lexicon, supporting medical information retrieval and semantic analysis. You understand medical terminology and concepts from UMLS.</p>
Error analysis	<p>The prediction errors in the Med-Mentions dataset are primarily due to challenges in identifying complex and overlapping disease mentions, as well as distinguishing between general biomedical terms and specific disease entities. Multi-token entities such as "renal pedicle occlusion" or "intention-to-treat analyses" were often partially labeled, with some tokens being misclassified or excluded. Additionally, the presence of nested or overlapping mentions, such as "prostate cancer" and its relationship to broader contexts like "treatment disparities," led to inconsistent labeling. The model also struggled with domain-specific terminology, misclassifying general terms like "maternal genotype" or "outcome" as disease mentions. These errors highlight limitations in handling nuanced biomedical language, especially when entities span multiple tokens or overlap with related terms.</p>

Table 12. Specific static prompts for each component we used for the Med-Mentions dataset.