
DDXPlus: A New Dataset For Automatic Medical Diagnosis

Arsene Fansi Tchango ^{*‡}, Rishab Goel ^{*†}, Zhi Wen [‡], Julien Martel [§], Joumana Ghosn [‡]

[‡] Mila-Quebec AI Institute

{arsene.fansi.tchango,zhi.wen,joumana.ghosn}@mila.quebec

[†] rgoel0112@gmail.com

[§] Dialogue Health Technologies Inc
julien.martel@dialogue.co

Abstract

There has been a rapidly growing interest in Automatic Symptom Detection (ASD) and Automatic Diagnosis (AD) systems in the machine learning research literature, aiming to assist doctors in telemedicine services. These systems are designed to interact with patients, collect evidence about their symptoms and relevant antecedents, and possibly make predictions about the underlying diseases. Doctors would review the interactions, including the evidence and the predictions, collect if necessary additional information from patients, before deciding on next steps. Despite recent progress in this area, an important piece of doctors' interactions with patients is missing in the design of these systems, namely the differential diagnosis. Its absence is largely due to the lack of datasets that include such information for models to train on. In this work, we present a large-scale synthetic dataset of roughly 1.3 million patients that includes a differential diagnosis, along with the ground truth pathology, symptoms and antecedents, for each patient. Unlike existing datasets which only contain binary symptoms and antecedents, this dataset also contains categorical and multi-choice symptoms and antecedents useful for efficient data collection. Moreover, some symptoms are organized in a hierarchy, making it possible to design systems able to interact with patients in a logical way. As a proof-of-concept, we extend two existing AD and ASD systems to incorporate the differential diagnosis, and provide empirical evidence that using differentials as training signals is essential for the efficiency of such systems. The dataset is available at https://figshare.com/articles/dataset/DDXPlus_Dataset/20043374.

1 Introduction

In a clinical conversation between a doctor and a patient, the patient usually initiates the discussion by specifying an initial set of symptoms they are experiencing. The doctor then iteratively inquires about additional symptoms and antecedents (describing the patient's history and potential risk factors), and considers, throughout the interaction, a differential diagnosis, i.e., a ranked list of plausible diseases the patient might be suffering from (Henderson et al., 2012; Guyatt et al., 2002; Rhoads et al., 2017), which is refined based on the patient's responses. During this multi-step process, the doctor tries to collect all relevant information to narrow down the differential diagnosis. Once the differential diagnosis is finalized, the doctor can ask the patient to undergo medical exams to eliminate most pathologies included in it and confirm the one(s) the patient is suffering from, or can decide to directly prescribe a treatment to the patient.

*Equal contribution

Aiming to aid doctors in such clinical interactions, there has been recently significant progress on building Automatic Symptom Detection (ASD) systems and Automatic Diagnosis (AD) systems, using machine learning and Reinforcement Learning (RL) techniques (Wei et al., 2018; Xu et al., 2019; Chen et al., 2021; Zhao et al., 2021; Guan and Baral, 2021; Yuan and Yu, 2021; Liu et al., 2022). These systems are meant to collect symptoms and antecedents relevant to the patient’s condition, while minimizing the length of the interaction to improve efficiency and avoid burdening patients with unnecessary questions. In addition, AD systems are tasked to predict the patient’s underlying disease to further aid doctors in deciding appropriate next steps. However, this setting differs from real patient-doctor interactions in an important way, namely the absence of the differential diagnosis. Based on the conversation alone, without further evidence such as medical exams, doctors tend to consider in their reasoning a differential diagnosis rather than a single pathology (Henderson et al., 2012). In addition to presenting a more comprehensive view of the doctor’s opinion on the patient’s underlying condition, the differential diagnosis also helps to account for the uncertainty in the diagnosis, since the symptoms experienced by a patient can point to multiple pathologies. Moreover, the differential diagnosis can help guide the doctor in determining which questions to ask the patient during their interaction. Thus, considering the differential diagnosis is especially important for better and more efficient evidence collection, for accounting for the uncertainty in the diagnosis, and for building systems that doctors can understand and trust. We believe that the absence of this key ingredient in recent AD/ASD systems is mainly due to the lack of datasets that include such information. The most commonly used public datasets, such as DX (Wei et al., 2018), Muzhi (Xu et al., 2019) and SymCAT (Peng et al., 2018), are all designed for predicting the pathology a patient is experiencing and do not contain differential diagnosis data.

To close this gap and encourage future research that focuses on the differential diagnosis, we introduce DDXPlus, a large-scale synthetic dataset for building AD and ASD systems. This dataset is similar in format to other public datasets such as DX (Wei et al., 2018) and Muzhi (Xu et al., 2019), but differs in several important ways. First, it makes a clear distinction between symptoms and antecedents which are not of equal importance from a doctor’s perspective when interacting with patients. Second, it is larger in scale, in terms of the number of patients, as well as the number of represented pathologies, symptoms and antecedents. Third, contrary to existing datasets which only include binary symptoms and antecedents, it also includes categorical and multi-choice symptoms and antecedents useful for efficient evidence collection. Moreover, some symptoms are organized in a hierarchy making it possible to design systems able to interact with patients in a logical way. Finally, each patient is characterized by a differential diagnosis as well as the pathology they are actually suffering from. To the best of our knowledge, this is the first large-scale dataset that includes both ground truth pathologies and differential diagnoses, as well as non-binary symptoms and antecedents. To summarize, we make the following contributions:

- We release a large-scale synthetic benchmark dataset of roughly 1.3 million patients covering 49 pathologies, 110 symptoms and 113 antecedents. The dataset is generated using a proprietary medical knowledge base and contains a mixture of multi-choice, categorical and binary symptoms and antecedents. Importantly, it also contains a differential diagnosis for each patient along with the patient’s underlying pathology.
- We extend two existing AD and ASD systems to incorporate the differential diagnosis and show that using this information as a training signal improves their performance.

2 Existing datasets and limitations

To build machine learning-based AD or ASD systems that medical doctors can trust, one needs to have access to related patient data, namely, the whole set of symptoms experienced by the patient, the relevant antecedents, the underlying pathology, and lastly the differential diagnosis associated with the experienced symptoms and antecedents. Unfortunately, there is no public dataset with all these characteristics. Existing public datasets of medical records, such as the MIMIC-III dataset (Johnson et al., 2016), often lack symptom-related data and are therefore inappropriate for training AD/ASD models. Other datasets, such as DX (Wei et al., 2018) or Muzhi (Xu et al., 2019), are of small scale, and don’t necessarily provide a holistic view of the symptoms and antecedents experienced by patients since they are derived from medical conversations. Indeed, if a symptom is not mentioned in a conversation, there is no way to determine if it was experienced by the underlying patient or not.

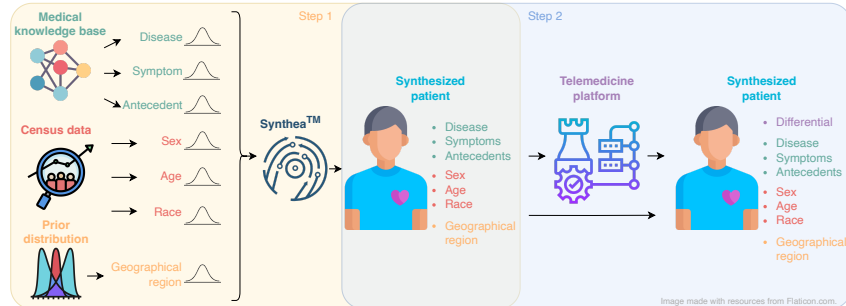


Figure 1: Overview of the data generation process of DDXPlus (Section 3).

To tackle these limitations, previous works (Peng et al., 2018; Kao et al., 2018) relied on the SymCAT database (AHEAD Research, 2011) for data synthesis. Unfortunately, SymCAT includes binary-only symptoms, which can lead to unnecessarily long interactions with patients, compared to categorical or multi-choice questions that collect the same information in fewer turns. For example, to ask about pain location, systems built using SymCAT need to ask several questions, such as *Do you have pain in your arm?*, *Do you have pain in your ankle?*, \dots , until every relevant location is inquired about. The same information can be obtained by asking a more generic question such as *Where is your pain located?* and obtain all the locations at once. Moreover, as noted by (Yuan and Yu, 2021), the symptom information in SymCAT is incomplete and, as a consequence, the synthetic patients generated using SymCAT are not sufficiently realistic for testing AD and ASD systems. Other datasets with more realistic synthetic patients have been recently proposed such as HPO (Köhler et al., 2021; Guan and Baral, 2021) and MedlinePlus (Yuan and Yu, 2021). However, like SymCAT, they contain only binary symptoms.

Most importantly, a critical shortcoming that is common to all above mentioned datasets is they do not contain differential diagnosis data. In what follows, we introduce the DDXPlus dataset and describe the undertaken steps towards its creation.

3 DDXPlus dataset

To generate the dataset, we proceed in two steps. For each patient, we first use a proprietary medical knowledge base (Section 3.1), public census data, and *Synthea™* (Walonoski et al., 2017) to synthesize their socio-demographic data, underlying disease, symptoms and antecedents (Section 3.2). Next, we generate the patient’s corresponding differential diagnosis using an existing commercial rule-based ASD system (Section 3.3). Figure 1 illustrates the data generation process.

3.1 Medical knowledge base

The dataset we propose in this work relies on a proprietary medical knowledge base (KB) which was constructed by compiling over 20 000 medical papers and extracting diseases’ incidence and prevalence, and symptoms and risk factors’ likelihood ratios across various age groups and geographies. This KB was used to design DXA, a rule-based ASD system that has been deployed in a commercial telemedicine platform. In total, the knowledge base covers a set of 440 pathologies and 802 symptoms and antecedents. The pathologies are grouped in overlapping subgroups based on common characteristics referred to as *chief complaints* (Aronsky et al., 2001; Thompson et al., 2006). In this work, we focus on pathologies belonging to the chief complaint related to cough, sore throat, or breathing issues. This group contains 49 pathologies covering 110 symptoms and 113 antecedents. Extending the dataset to all pathologies is left for future work.

Each disease d in the KB is characterized by either an incidence rate, a prevalence rate, or both values. Both rates are conditioned on the age, sex, and geographical region of the patient. The incidence rate measures the proportion of new occurrences of a disease in a population over a period of time while the prevalence rate captures the proportion of individuals in a population that have a disease at a particular time (Friis and Sellers, 2020). For each disease, a set of symptoms and antecedents describing the disease are provided together with their related probabilities. These probabilities

are conditioned on the age and sex of the patient. Thus, the probability values $p(s|d, age, sex)$ and $p(a|d, age, sex)$ are provided for each symptom s and each antecedent a . Unlike existing datasets mentioned above, evidences (i.e., symptoms and antecedents) within this KB are of several types. They can be binary (e.g., *Are you coughing?*), categorical (e.g., *What is the pain intensity on a scale of 0 to 10?*), or multi-choice (e.g., *Where is your pain located?*). In some cases, an evidence f_s (e.g., *pain location*) may be dependent on another evidence s (e.g., presence of *pain*), in which case the KB provides means to extract the corresponding conditional probability $p(f_s|s, d, age, sex)$. Finally, each disease is characterized by its level of severity ranging from 1 to 5 with the lowest values describing the most severe pathologies.

3.2 Patient generation

A synthetic patient is an entity characterized by an *age*, a *sex*, a geographical region *geo*, and who is suffering from a pathology d and is experiencing a set E of symptoms and antecedents. As a first step for modeling patients, we consider the following rule:

$$p(age, sex, geo, d, E) = p(age, sex, geo) \times p(d|age, sex, geo) \times p(E|d, age, sex).$$

This formulation relies on the fact that the set of evidence E experienced by a patient given a disease doesn't depend on the geographical region. In what follows, we present other rules and assumptions required to exploit the KB.

Assumptions on the socio-demographic data As we only have access to the marginal distributions, we assume that age, sex, and geographical region are independent. That is,

$$p(age, sex, geo) = p(age) \times p(sex) \times p(geo).$$

The distributions of age and sex can be obtained from census data. For this dataset, we used the 2010-2015 US Census data from the State of New York (US Government, 2015). For more details, see Section A.1. As a result of this choice, the default geographical region of a patient is "North America". Given that some pathologies in the KB can be contracted only if the patient is from a different geographical region, we embed the notion of recent travel when synthesizing a patient. Each synthesized patient is generated by simulating the fact that they recently travelled or not, and if they travelled, in which geographical region. We thus assume the availability of a prior distribution $p(travel)$ representing the proportion of the population travelling each month and we consider that the distribution regarding the geographical regions of destination is uniform². Based on these assumptions, we derive the following prior distribution $p(geo)$ over the geographical regions:

- Sample $u \sim \mathcal{U}(0, 1)$.
- If $u < p(travel)$, then randomly select a geographical region from the available set of geographical regions (see Section A.2). We used $p(travel) = 0.25$ for this dataset.
- If $u \geq p(travel)$, then set the geographical region to be "North America".

Assumptions on pathologies We use the disease's incidence rate, when available, as the disease prior distribution $p(d|age, sex, geo)$, and fall back on the disease's prevalence rate when the incidence rate is missing. This is one of the major limitations of the data generation process which needs to be addressed in future work. The incidence rate can be approximated by dividing the prevalence rate with a constant factor representing the average duration of the disease, and which may be different for each disease. Out of the 49 considered diseases, 8 are affected by this limitation (see Section A.3).

When the resulting rate is greater than 100% (e.g., an incidence rate of 200% means that an individual will likely develop the pathology twice a year on average), we cap it at 100% to avoid having a highly imbalanced dataset. Indeed, without this capping, the dataset would have been dominated by only a few pathologies, with more than half of the patients suffering from the three pathologies whose incidence rate is greater than 100% (i.e., URTI, Viral pharyngitis, and Anemia).

Finally, the KB also contains some diseases that have extremely low incidence rates, and therefore patients suffering from those pathologies are barely generated. To increase the chance of those pathologies to be represented within the dataset, we decide to cap the rate at a minimum of 10%. Thus, the rates used to generate the dataset lie between 10% and 100%. This alteration of the original rates leads to a more balanced dataset (see Figure 2).

²The dataset can be improved by using real travel and geographical destination statistics.

Assumptions on symptoms and antecedents Given a sampled disease d , the next step is to generate all the evidences (i.e., symptoms and antecedents) the synthesized patient will be experiencing. Given that the KB doesn't contain the joint distribution of all symptoms and antecedents conditioned on the disease, sex and age, a simplifying assumption is made according to which the evidences are independent of each other when conditioned on the disease, age and sex unless explicitly stated otherwise in the KB. In other words, we have:

$$p(E|d, age, sex) = \prod_{e \in E \setminus E_h} p(e|d, age, sex) \prod_{f_s \in E_h} p(f_s|s, d, age, sex),$$

where s is the evidence which f_s depends on, E is the set of evidences experienced by the patient, and E_h is the subset of evidences experienced by the patient which are dependent on other evidences also experienced by the patient.

Some evidences, such as pain intensity, are described as integer values on a scale from 0 to 10. However, the knowledge base only provides the average value of each such evidence given the disease, the age, and the sex of the patient. To inject some randomness in the patient generation process, the values of those evidences are uniformly sampled from the interval $[\max(0, v - 3), \min(10, v + 3)]$ where v is the average value present in the knowledge base.

Finally, in order to reflect the deployed ASD system that is based on the KB, we limit to 5 the maximum number of choices associated with multi-choice evidences such as *pain location*.

Tools As mentioned above, we use *Synthea*TM to synthesize patients. To generate the value associated with a categorical or multi-choice evidence, *Synthea*TM must be provided with a list of possible values together with their related conditional probabilities. *Synthea*TM then goes through that list incrementally and decides, for each possible value, if it is "on" or not based on its probability. The process stops as soon as 1 (resp. 5) possible value(s) is (are) "on" for categorical (resp. multi-choice) evidences or the provided list is fully processed. In this work, the possible values of an evidence are ordered in ascending order based on their conditional probability of occurrence to make sure rare values appear in the dataset.

3.3 Differential diagnosis generation

As mentioned above, the KB was used to build a rule-based ASD system which was deployed in a real-world telemedicine platform and which is capable of generating differential diagnoses. In the production environment, this platform expects to be given the age, sex and an initial symptom. Based on this information, it determines a set of chief complaints (and their associated pathologies) that are compatible with the provided information, and it engages in a question-answering session with the patient to collect information about the patient's symptoms and antecedents. It then generates a differential diagnosis of ranked pathologies. We leverage this platform to compute the differential diagnosis of each synthesized patient. In order to bypass the limitations of the platform and the errors it might make in the question-answering session by omitting to ask some relevant questions, we provide all evidences in one shot, as opposed to providing each evidence iteratively upon request of the platform. More specifically, we proceed according to the following high-level steps:

- We provide the age and the sex of the patient, the appropriate chief complaint, and we answer "yes" to the question "Are you consulting for a new problem?".
- We provide all the symptoms and antecedents experienced by the patient in one shot, at the beginning of the interaction. The motivation behind this is to make sure that the system is aware of all relevant information and doesn't miss on any of the patient's symptoms and antecedents.
- The platform may still inquire about additional questions. If that is the case, we answer "no" to those questions until we see a "QUIT" response from the platform or the maximum interaction length (30) is reached.
- When the maximum interaction length is reached, the platform does not produce a differential diagnosis. The corresponding patient is discarded from the dataset.
- When a "QUIT" response is provided by the platform, it contains a differential diagnosis. We further proceed by verifying if the underlying synthesized disease is part of the generated

differential diagnosis. If it is not the case (because the platform itself is not a perfect system or because the patient didn't have enough evidences for the rule-based system to include the simulated disease in the differential diagnosis), the patient is discarded from the dataset. Each pathology within the generated differential diagnosis has a score. Those scores are normalized to obtain a probability distribution.

Given that the platform treats the provided chief complaint as a clue instead of a hard constraint and given that the platform was built to consider all 440 pathologies found in the KB, it sometimes returns a differential diagnosis that contains pathologies which do not belong to the specified chief complaint. There are several options for handling this situation: (1) create an "other pathologies" category and assign it the cumulative mass of the corresponding pathologies, or (2) manually remove those pathologies from the differential diagnosis and re-normalize the distribution. We opt for the second option because we want to restrict the set of pathologies to the universe of 49 pathologies used to synthesize patients. On average, we removed $1.78 (\pm 1.68)$ pathologies from the generated differential diagnosis for an average cumulative probability mass of $0.10 (\pm 0.11)$. Statistics regarding the rank from which those pathologies are excluded are described in Section A.4.

3.4 Dataset characteristics

With the above assumptions and limitations, we generate, under the CC-BY licence, a dataset of roughly 1.3 million patients, where each patient is characterized by their age, sex, geographical region or recent travel history, pathology, symptoms, antecedents, as well as the related differential diagnosis. We divide the dataset into training, validation, and test subsets based on an 80%-10%-10% split and using stratified sampling on the simulated pathology.

Compared to existing datasets from the AD and ASD literature, our dataset has several advantages:

- To the best of our knowledge, this is the first large-scale dataset containing differential diagnoses. This information is important because doctors reason in terms of a differential and not a single pathology, as the evidence collected from a patient can sometimes point to multiple pathologies, some of which requiring additional medical exams before they can be safely ruled out.
- Unlike the *SymCAT* (AHEAD Research, 2011) and the Muzhi (Wei et al., 2018) datasets, which only contain binary evidences, our dataset also includes categorical and multi-choice evidences which can naturally match the kind of questions a doctor would ask a patient, and which can lead to more efficient evidence collection. Moreover, some related evidences are defined according to a hierarchy which can be used to design systems able to interact with patients in a logical way.
- Our dataset makes a clear distinction between antecedents and symptoms. Given that antecedents might be documented in a patient's medical records, one can decide to put less emphasis on retrieving them when designing an automated system.
- Each pathology in our dataset is characterized by a severity level. This information can be used to design solutions that properly handle severe pathologies.

3.5 Data analysis

We present summary statistics of the generated dataset. Those statistics are based on the entire set. Statistics on the train, validation, and test subsets are presented in Section A.5.

Types of evidences: The distribution of the types of the evidences related to the 49 pathologies selected from the KB is provided in Table 1. 6.7% of evidences are categorical and multi-choice.

Table 1: Distribution of the evidence types corresponding to the 49 pathologies selected from the KB.

	Binary	Categorical	Multi-choice	Total
Evidences	208	10	5	223
Symptoms	96	9	5	110
Antecedents	112	1	0	113

Number of evidences: Table 2 shows an overview of the synthesized patients in terms of the number of simulated evidences. On average, a patient has roughly 10 symptoms and 3 antecedents.

Table 2: Statistics describing the number of evidences of the synthesized patients.

	Avg	Std dev	Min	1st quartile	Median	3rd quartile	Max
Evidences	13.56	5.06	1	10	13	17	36
Symptoms	10.07	4.69	1	8	10	12	25
Antecedents	3.49	2.23	0	2	3	5	12

Pathology statistics: Figure 2 shows the histogram of the pathologies of patients in the generated dataset. Although there are three dominating pathologies (URTI, Viral pharyngitis, and Anemia), other pathologies are also well represented.




Figure 2: Histogram of the pathologies experienced by the synthesized patients.

Socio-demographic statistics: The statistics of the socio-demographic data of the synthesized patients are shown in Figure 3. As expected, these statistics are compliant with the 2015 US Census data of the state of New York used during the generation process (Section A.1).




Figure 3: The socio-demographic statistics of the synthesized patients.

Differential diagnosis statistics: The distribution of the length of the differential diagnosis of the synthesized patients is depicted in Figure 4 (left). As observed, the generated differential diagnosis can contain more than one pathology. It is also interesting to observe that the simulated pathology is ranked first for more than 70% of patients (see Figure 4 (right)).

3.6 Dataset usage

Based on the previous sections, it should be clear that the released dataset is meant for research purposes. Any model trained on this dataset should not be directly used in a real-world system prior to performing rigorous evaluations of the model performance and verifying that the system has proper coverage and representation of the population that it will interact with.

4 Experiments

4.1 Models

We consider two existing AD/ASD systems which were originally designed to predict the pathology a patient is suffering from, and adapt those systems to instead predict the differential diagnosis.




Figure 4: Statistics regarding (left) the length of the differential diagnosis, and (right) the rank of the patient’s simulated pathology within the differential diagnosis (y-axis on log scale).

AARLC AARLC (Yuan and Yu, 2021) is a model that has two branches, an evidence acquisition branch (trained using RL) whose goal is to determine the next evidence to inquire about, and a classifier branch trained in a supervised way to predict the patient’s disease. An adaptive approach is used to align the tasks performed by the two branches using the entropy of the distributions predicted by the classifier branch. We use the same settings as those described in Yuan and Yu (2021) but tune the ν and λ parameters together with the learning rate on the validation subset.

BASD This supervised learning approach builds on top of the ASD module proposed in Luo et al. (2020) (with the exception of the knowledge graph) and adds a policy network to predict the underlying patient diagnosis at the end of the acquisition process. The agent is made of an MLP with a number n_h of hidden layers of size 2048 which is tuned together with the learning rate on the validation subset. More details about this approach can be found in Appendix A.6.

To assess the impact of the differential diagnosis, we train two versions of each approach, one focused on predicting the patient’s disease and one trained to predict the differential diagnosis. After training, we use the probability distribution predicted by each version at the end of an interaction as the predicted differential diagnosis. The hyper-parameters of each version are tuned independently and the resulting optimal set is used to report model performance.

4.2 Experimental setup

Each patient provides an initial evidence to the model at the beginning of the interaction. The model then iteratively inquires about a symptom or an antecedent, until either all the relevant symptoms and antecedents have been collected or a maximum number of turns T is reached. At the end of the interaction, a differential diagnosis is predicted. We use $T = 30$ in all experiments. Training is performed using a NVIDIA V100 GPU.

4.3 Results

An AD system is typically tasked to collect (i) relevant evidences from a patient, (ii) make accurate predictions regarding the patient’s disease(s), and (iii) operate in a minimum number of turns. As such, we report on the interaction length (IL), and evaluate the evidence collection by measuring the recall (PER), precision (PEP), and F1 score (PEF1) of the inquired evidences. Additionally, to evaluate the differential diagnosis prediction, we calculate the recall (DDR@ k), precision (DDP@ k) and F1 score (DDF1@ k) when considering the top k entries of the predicted distributions. Finally, we report on the accuracy of the inclusion of the gold truth pathology (i.e., the pathology a patient was simulated from) in the predicted differential diagnosis (GTPA@ k). Definitions of these metrics and other details can be found in Appendix A.7. To compute these metrics, we post-process both the ground-truth differentials and the predicted ones to remove pathologies whose mass is less than or equal to 0.01. This threshold is selected to reduce the size of the differentials by removing highly unlikely pathologies.

Tables 3 and 4 show the results obtained for the two approaches. As expected, models trained to predict the patient disease exhibit better GTPA@1 scores than the ones trained to predict the differential diagnosis. This is explained by the fact that the patient disease is not always ranked at the first position in the corresponding differential diagnosis (see Figure 4 (right)). Still, models trained to predict the differential diagnosis manage to include the underlying patient pathology within the top-5 entries of their predictions at least 96% of the time (see GTPA@5).

Table 3: Interaction length, ground truth pathology accuracy, and evidence collection metrics of the trained agents as measured on the test set. Diff indicates whether the agent was trained to predict the differential diagnosis (✓) or not (×). Except IL, all values are expressed in %. Also, the higher the better for all metrics except IL.

Method	Diff	IL	GTPA@1	GTPA@3	GTPA@5	PER	PEP	PEF1
AARLC	✓	14.1	61.74	91.51	96.6	51.56	74.83	56.99
	×	7.27	98.86	99.79	99.93	35.91	79.38	45.51
BASD	✓	17.46	68.23	91.35	96.00	87.71	74.05	76.64
	×	17.85	96.65	98.32	98.34	88.15	73.26	76.17

Table 4: Evaluation of the differential diagnosis predicted by the trained agents as measured on the test set. Diff indicates whether the agent was trained to predict the differential diagnosis (✓) or not (×). All values are expressed in %. The higher the better for all metrics.

Method	Diff	DDR@1	DDP@1	DDF1@1	DDR@3	DDP@3	DDF1@3	DDR@5	DDP@5	DDF1@5
AARLC	✓	71.59	71.59	71.59	74.86	69.58	71.14	79.33	69.79	72.53
	×	72.73	72.73	72.73	37.31	91.78	51.25	28.56	95.97	41.05
BASD	✓	79.59	79.59	79.59	75.72	76.7	75.27	77.1	78.86	76.36
	×	71.35	71.35	71.35	37.19	92.05	51.25	28.42	96.54	41.05

Regarding the recovery of the differential diagnosis, it can be observed that models trained to predict the differential diagnosis do a better job than the ones trained to predict the patient disease (see DDF1@3 and DDF1@5) with an absolute improvement of more than 20%. This result suggests that in order to accurately predict the differential diagnosis the medical practitioners care about, AD systems must be trained using data containing information related to differential diagnoses.

Finally, as illustrated by the PEF1 scores, the usage of the differential diagnosis during the training stage contributes to the collection of relevant evidences when interacting with patients. This last observation together with the above results highlights the usefulness of the proposed dataset to the AD/ASD community.

5 Conclusion

We released a large-scale benchmark dataset of roughly 1.3 million patients suffering from pathologies that include cough, sore throat or breathing problems as symptoms. The dataset contains binary, categorical and multi-choice symptoms and antecedents. Each patient within the dataset is characterized by their age, sex, geographical region or recent travel history, pathology, symptoms, antecedents, as well as the related differential diagnosis. We extended two existing approaches from the AD/ASD literature (based on RL and non-RL settings) to leverage the differential diagnosis available in the proposed dataset. The obtained results provided empirical evidence that using the differential diagnosis as a training signal enhances the performance of AD/ASD systems. We hope that this dataset will encourage the research community to develop automatic diagnosis systems that can be trusted by medical practitioners as the latter operate with differential diagnoses when interacting with real patients. While extending the above mentioned approaches, we considered all pathologies as equally important. But in general, when establishing a differential diagnosis, medical practitioners likely ask questions to specifically explore and rule out severe pathologies. The proposed dataset has a severity flag associated with each pathology. This leaves room for exploring approaches that better handle severe pathologies. Finally, we would like to emphasize that this dataset should not be used to train and deploy a model prior to performing rigorous evaluations of the model performance and verifying that the system has proper coverage and representation of the population that it will interact with.

Acknowledgments and Disclosure of Funding

We would like to thank Dialogue Health Technologies Inc for providing us access to the knowledge base used throughout this work and their rule-based ASD system. We would also like to thank Quebec's Ministry of Economy and Innovation and Invest AI for their financial support.

References

- AHEAD Research. SymCAT: Symptom-based, Computer Assisted Triage. <http://www.symcat.com/>, 2011. Accessed: 2021-01-31.
- Dominik Aronsky, Diane Kendall, Kathleen Merkley, Brent James, and Peter Haug. A comprehensive set of coded chief complaints for the emergency department. *Academic emergency medicine : official journal of the Society for Academic Emergency Medicine*, 8:980–9, 11 2001. doi: 10.1111/j.1553-2712.2001.tb01098.x.
- Junying Chen, Dongfang Li, Qingcai Chen, Wenxiu Zhou, and Xin Liu. Diaformer: Automatic diagnosis via symptoms sequence generation. *arXiv preprint arXiv:2112.10433*, 2021.
- Robert H Friis and Thomas Sellers. *Epidemiology for public health practice*. Jones & Bartlett Learning, 2020.
- Hong Guan and Chitta Baral. A bayesian approach for medical inquiry and disease inference in automated differential diagnosis. *arXiv preprint arXiv:2110.08393*, 2021.
- Gordon Guyatt, Drummond Rennie, Maureen Meade, Deborah Cook, et al. *Users' guides to the medical literature: a manual for evidence-based clinical practice*, volume 706. AMA press Chicago, 2002.
- Mark Henderson, Lawrence M Tierney, and Gerald W Smetana. *The patient history: Evidence-based approach*. McGraw Hill Professional, 2012.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3:160035, 2016.
- Hao-Cheng Kao, Kai-Fu Tang, and Edward Chang. Context-aware symptom checking for disease diagnosis using hierarchical reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Sebastian Köhler, Michael Gargano, Nicolas Matentzoglou, Leigh C Carmody, David Lewis-Smith, Nicole A Vasilevsky, Daniel Danis, Ganna Balagura, Gareth Baynam, Amy M Brower, et al. The human phenotype ontology in 2021. *Nucleic acids research*, 49(D1):D1207–D1217, 2021.
- Wenge Liu, Yi Cheng, Hao Wang, Jianheng Tangi, Yafei Liu, Ruihui Zhao, Wenjie Li, Yefeng Zheng, and Xiaodan Liang. "my nose is running." "are you also coughing?": Building a medical diagnosis agent with interpretable inquiry logics. *arXiv preprint arXiv:2204.13953*, 2022.
- Hongyin Luo, Shang-Wen Li, and James Glass. Knowledge grounded conversational symptom detection with graph memory networks. In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 136–145, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.clinicalnlp-1.16. URL <https://www.aclweb.org/anthology/2020.clinicalnlp-1.16>.
- Yu-Shao Peng, Kai-Fu Tang, Hsuan-Tien Lin, and Edward Chang. Refuel: Exploring sparse features in deep reinforcement learning for fast disease diagnosis. *Advances in neural information processing systems*, 31:7322–7331, 2018.
- Jacqueline Rhoads, Julie C Penick, et al. *Formulating a Differential Diagnosis for the Advanced Practice Provider*. Springer Publishing Company, 2017.

- David Thompson, David Eitel, Christopher Fernandes, Jesse Pines, and James Amsterdam. Coded chief complaints-automated analysis of free-text complaints. *Academic emergency medicine : official journal of the Society for Academic Emergency Medicine*, 13:774–82, 08 2006. doi: 10.1197/j.aem.2006.02.013.
- US Government. 2010-2015 US Census Data. <https://www2.census.gov/programs-surveys/popest/datasets/2010-2015/>, 2015. Accessed: 2021-01-31.
- Jason Walonoski, Mark Kramer, Joseph Nichols, Andre Quina, Chris Moesel, Dylan Hall, Carlton Duffett, Kudakwashe Dube, Thomas Gallagher, and Scott McLachlan. Synthea: An approach, method, and software mechanism for generating synthetic patients and the synthetic electronic health care record. *Journal of the American Medical Informatics Association*, 25(3):230–238, 08 2017. ISSN 1527-974X. doi: 10.1093/jamia/ocx079. URL <https://doi.org/10.1093/jamia/ocx079>.
- Zhongyu Wei, Qianlong Liu, Baolin Peng, Huaixiao Tou, Ting Chen, Xuan-Jing Huang, Kam-Fai Wong, and Xiang Dai. Task-oriented dialogue system for automatic diagnosis. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 201–207, 2018.
- Lin Xu, Qixian Zhou, Ke Gong, Xiaodan Liang, Jianheng Tang, and Liang Lin. End-to-end knowledge-routed relational dialogue system for automatic diagnosis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7346–7353, 2019.
- Hongyi Yuan and Sheng Yu. Efficient symptom inquiring and diagnosis via adaptive alignment of reinforcement learning and classification. *arXiv preprint arXiv:2112.00733*, 2021.
- Xinyan Zhao, Liangwei Chen, and Huanhuan Chen. A weighted heterogeneous graph-based dialog system. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.

A Appendix

A.1 Demographic statistics from census data

To synthesize patients, one needs to have access to the prior distributions of age and sex of a population of interest. In this work, we rely on the 2010-2015 US census data US Government (2015) for the state of New York. Table 5 describes the corresponding statistics.

Table 5: The 2010-2015 census data of the state of New York.

	Category	Frequency
Sex	Male	0.4836
	Female	0.5164
Age	Less than 1-year	0.0154
	1-4-years	0.0461
	5-14-years	0.1146
	15-29-years	0.2132
	30-44-years	0.2025
	45-59-years	0.2042
	60-74-years	0.1399
	75-years and more	0.0641

A.2 Geographical regions

The following geographical regions are covered in the dataset: North Africa, West Africa, South Africa, Central America, North America, South America, Asia, South East Asia, The Caribbean, Europe, Oceania.

A.3 Missing incidence rates

The 8 pathologies whose incidence rates were not present for all the combinations of age, sex, and geographical regions are: Anemia, Inguinal hernia, Anaphylaxis, Allergic sinusitis, Chagas, Tuberculosis, Ebola, Chronic rhinosinusitis.

A.4 Differential diagnosis post-processing

As mentioned previously, pathologies that are part of the differential diagnosis generated by the rule-based ASD platform but that do not belong to the set of the 49 considered pathologies are excluded from the differential diagnosis. Table 6 describes, for each rank in the differential diagnosis, the proportion of patients for which the pathology at that rank is excluded from the differential diagnosis.

A.5 Subset data analysis

A.5.1 Evidence statistics

The statistics of the evidences experienced by the synthesized patients belonging to the train, validation, and test subsets are presented in Tables 7, 8, and 9 respectively. As illustrated, the evidences are similarly distributed across the three subsets.

A.5.2 Pathology statistics

Figures 5, 6, and 7 depict the histograms of the pathologies experienced by the synthesized patients in the train, validation, and test subsets. The pathologies are evenly distributed across the three subsets.

A.5.3 Socio-demographic statistics

The socio-demographic statistics of the synthesized patients from the train, validation and test subsets are respectively illustrated in Figures 8, 9, and 10.

Table 6: Proportion (%) of the patients for which pathologies are excluded at each rank from the differential diagnosis returned by the rule-based ASD platform.

Rank	Proportion
1	1.09
2	11.23
3	15.73
4	13.78
5	14.53
6	11.92
7	13.12
8	10.47
9	9.23
10	8.16
11	9.47
12	7.00
13	5.92
14	6.11
15	5.49
16	5.86
17	5.42
18	4.52
19	3.67
20	3.37
21	3.07
22	2.37
23	2.14
24	1.70
25	1.17
26	0.71
27	0.42
28	0.24
29	0.13
30	0.07
31	0.04
32	0.03
33	0.02
34	0.01
35	4.49e-03
36	1.76e-03
37	3.90e-04

Table 7: Statistics describing the number of evidences of the synthesized patients for the training set.

	Evidences	Symptoms	Antecedents
Avg	13.52	10.03	3.49
Std dev	5.08	4.71	2.23
Min	1	1	0
1st quartile	10	8	2
Median	13	10	3
3rd quartile	17	12	5
Max	36	25	12

Table 8: Statistics describing the number of evidences of the synthesized patients for the validation set.

	Evidences	Symptoms	Antecedents
Avg	13.76	10.27	3.49
Std dev	5.01	4.61	2.23
Min	1	1	0
1st quartile	10	8	2
Median	13	10	3
3rd quartile	17	12	5
Max	34	25	12

Table 9: Statistics describing the number of evidences of the synthesized patients for the test set.

	Evidences	Symptoms	Antecedents
Avg	13.72	10.23	3.50
Std dev	5.02	4.63	2.23
Min	1	1	0
1st quartile	10	8	2
Median	13	10	3
3rd quartile	17	12	5
Max	35	25	12




Figure 5: Histogram of the patient pathologies in the training set.




Figure 6: Histogram of the patient pathologies in the validation set.




Figure 7: Histogram of the patient pathologies in the test set.




Figure 8: The socio-demographic statistics of the synthesized patients from the training set.




Figure 9: The socio-demographic statistics of the synthesized patients from the validation set.




Figure 10: The socio-demographic statistics of the synthesized patients from the test set.

A.5.4 Differential diagnosis statistics

The histograms of the length of the differential diagnosis from the train, validation, and test subsets are respectively illustrated in Figures 11, 13, and 15. Similarly, the histograms of the rank of the simulated pathology within the differential diagnosis from the train, validation, and test subsets are respectively depicted in Figures 12, 14, and 16.




Figure 11: Statistics regarding the length of the differential diagnosis in the training set.

A.6 BASD approach

The BASD agent consists of an MLP network with 2 prediction branches:

- a policy branch whose role is to predict whether to stop or continue the interaction, and if the latter, what evidence to inquire about next;
- a classifier branch to make a prediction regarding the underlying patient disease.




Figure 12: Statistics regarding the rank of the simulated pathology within the differential diagnosis in the training set (y-axis on log scale).

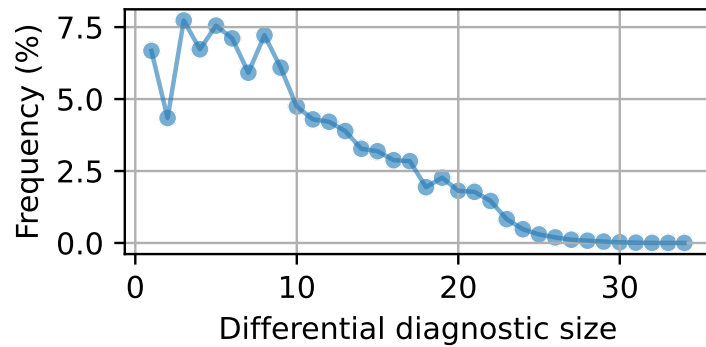


Figure 13: Statistics regarding the length of the differential diagnosis in the validation set.




Figure 14: Statistics regarding the rank of the simulated pathology within the differential diagnosis in the validation set (y-axis on log scale).

To train the network, we simulate dialogue states together with their target values. Assuming that a given patient has n evidences that they are experiencing, we simulate a dialogue state as follows:

1. Randomly select $p \in [1, n]$ representing the number of positive evidences already inquired about. Sample p evidences from the ones experienced by the patient and set them in the simulated dialog state.
2. Randomly select $q \in [0, T - p)$ representing the number of negative evidences already inquired, where T is the maximum number of allowed dialog turns. Sample q evidences from the ones not experienced by the patient and set them in the simulated dialog state.

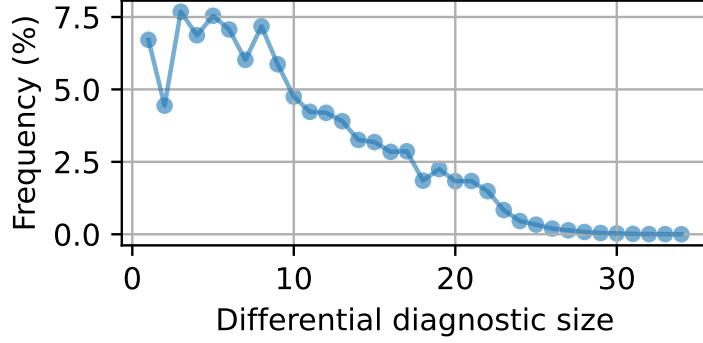


Figure 15: Statistics regarding the length of the differential diagnosis in the test set.

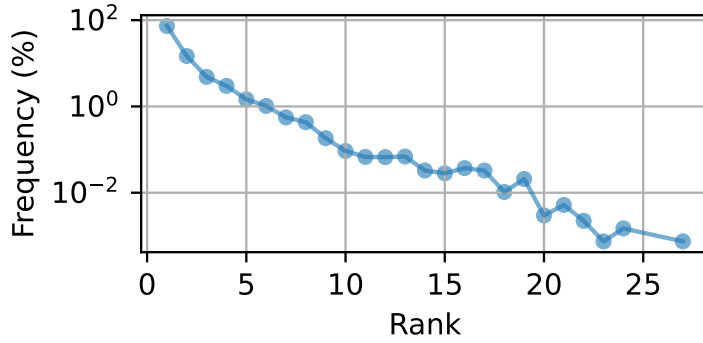


Figure 16: Statistics regarding the rank of the simulated pathology within the differential diagnosis in the test set (y-axis on log scale).

3. If $p = n$, set the target of the policy branch to "stop"; otherwise, set the target to be one of the experienced evidences that was not sampled at step 1.
4. Set the classifier branch target to be the ground-truth pathology.

Both branches are trained using the cross-entropy loss function and the classifier branch is only updated when the target of the policy branch is set to "stop".

A.7 Evaluation metrics

This section describes the metrics used to evaluate the performance of trained agents. The differentials generated by the trained models as well as the ground truth differentials are post-processed to remove pathologies whose mass is less than or equal to 0.01. This threshold, approved by our collaborating physician, is selected to reduce the size of the differentials by removing highly unlikely pathologies. Let D be the number of patients, T^i be the set of evidences collected by an agent from the i^{th} patient, Y^i be the ground truth differential, and \hat{Y}^i be the pathology distribution generated by the agent for that patient.

Interaction length (IL): The average interaction length is defined as:

$$IL = \frac{1}{|D|} \sum_{i=1}^{|D|} |T^i|. \quad (1)$$

Differential diagnosis recall for the top k entries (DDR@ k): At the end of an interaction, the agent produces a distribution \hat{Y} over the pathologies. We measure whether the top k pathologies \hat{Y}_k extracted from this distribution are present in the top k pathologies Y_k of the ground truth differential. If the ground truth differential has less than k pathologies, we set k to the size of the ground truth

differential. The metric is defined as:

$$DDR@k = \frac{1}{|D|} \sum_{i=1}^{|D|} \frac{|\hat{Y}_k^i \cap Y_k^i|}{|Y_k^i|}, \quad (2)$$

where \hat{Y}_k^i and Y_k^i are respectively the predicted and gold truth top k pathologies for the i^{th} patient.

Differential diagnosis precision for the top k entries (DDP@ k): This metric measures the precision of the differential diagnoses predicted by an agent when considering the top- k entries:

$$DDP@k = \frac{1}{|D|} \sum_{i=1}^{|D|} \frac{|\hat{Y}_k^i \cap Y_k^i|}{|\hat{Y}_k^i|}. \quad (3)$$

Differential diagnosis F1 for the top k entries (DDF1@ k): We combine the aforementioned recall and precision metrics to compute the F1 score of the differential diagnosis.

Gold truth pathology accuracy (GTPA@ k): This metric measures whether the differential diagnosis predicted by an agent contains the pathology p a patient was simulated from within its top- k entries:

$$GTPA@k = \frac{1}{|D|} \sum_{i=1}^{|D|} \mathbb{1}_{\hat{Y}_k^i}(p^i), \quad (4)$$

where $\mathbb{1}$ is the indicator function.

Positive evidence recall (PER): Let us suppose that the i^{th} patient is experiencing the set \mathcal{S}^i of symptoms and the set \mathcal{A}^i of antecedents. Also let us assume that the agent inquires the set $\hat{\mathcal{S}}^i$ (resp. $\hat{\mathcal{A}}^i$) of symptoms (resp. antecedents) from which $\hat{\mathcal{S}}_+^i \subseteq \mathcal{S}^i$ (resp. $\hat{\mathcal{A}}_+^i \subseteq \mathcal{A}^i$) is the set of symptoms (resp. antecedents) experienced by the i^{th} patient. Then, the recall for the positive evidences is calculated as:

$$PER = \frac{1}{|D|} \sum_{i=1}^{|D|} PER_i, \quad \text{where } PER_i = \frac{|\hat{\mathcal{S}}_+^i \cup \hat{\mathcal{A}}_+^i|}{|\mathcal{S}^i \cup \mathcal{A}^i|}. \quad (5)$$

Positive evidence precision (PEP): Similarly to the recall metric, we compute the precision of inquired evidences. We have:

$$PEP = \frac{1}{|D|} \sum_{i=1}^{|D|} PEP_i, \quad \text{with } PEP_i = \frac{|\hat{\mathcal{S}}_+^i \cup \hat{\mathcal{A}}_+^i|}{|T^i|}. \quad (6)$$

Positive evidence F1 (PEF1): We combine the aforementioned recall and precision metrics to compute the F1 score of the inquired evidences.

$$PEF1 = \frac{1}{|D|} \sum_{i=1}^{|D|} 2 \frac{PEP_i \times PER_i}{PEP_i + PER_i}. \quad (7)$$

A.8 Dataset release

We are releasing under the CC-BY licence a new large-scale dataset for Automatic Symptom Detection (ASD) and Automatic Diagnosis (AD) systems in the medical domain. The dataset contains patients synthesized using a proprietary medical knowledge base and a commercial rule-based ASD system. Patients in the dataset are characterized by their socio-demographic data, a pathology they are suffering from, a set of symptoms and antecedents related to this pathology, and a differential diagnosis. The symptoms and antecedents can be binary, categorical and multi-choice, with the potential of leading to more efficient and natural interactions between ASD/AD systems and patients. To the best of our knowledge, this is the first large-scale dataset that includes the differential diagnosis, and non-binary symptoms and antecedents. The dataset can be downloaded at https://figshare.com/articles/dataset/DDXPlus_Dataset/20043374.

A.8.1 Dataset documentation

In what follows, we use the term ‘‘evidence’’ as a general term to refer to a symptom or an antecedent. The dataset contains the following files:

- `release_evidences.json`: a JSON file describing all possible evidences considered in the dataset.
- `release_conditions.json`: a JSON file describing all pathologies considered in the dataset.
- `release_train_patients.zip`: a CSV file containing the patients of the training set.
- `release_validate_patients.zip`: a CSV file containing the patients of the validation set.
- `release_test_patients.zip`: a CSV file containing the patients of the test set.

Evidence description Each evidence in the `release_evidences.json` file is described using the following entries:

- `name`: name of the evidence.
- `code_question`: a code allowing to identify which evidences are related. Evidences having the same `code_question` form a group of related evidences. The value of the `code_question` refers to the evidence that needs to be simulated/activated for the other members of the group to be eventually simulated.
- `question_fr`: the query, in French, associated to the evidence.
- `question_en`: the query, in English, associated to the evidence.
- `is_antecedent`: a flag indicating whether the evidence is an antecedent or a symptom.
- `data_type`: the type of the evidence. We use “B” for binary, “C” for categorical, and “M” for multi-choice.
- `default_value`: the default value of the evidence. If this value is used to characterize the evidence, then it is as if the evidence was not synthesized.
- `possible-values`: the possible values for the evidence. Only valid for categorical and multi-choice evidences.
- `value_meaning`: The meaning, in French and English, of each code that is part of the `possible-values` field. Only valid for categorical and multi-choice evidences.

Pathology description The file `release_conditions.json` contains information about the pathologies patients in the datasets may suffer from. Each pathology has the following attributes:

- `condition_name`: name of the pathology.
- `cond-name-fr`: name of the pathology in French.
- `cond-name-eng`: name of the pathology in English.
- `icd10-id`: ICD-10 code of the pathology.
- `severity`: the severity associated with the pathology. The lower the more severe.
- `symptoms`: data structure describing the set of symptoms characterizing the pathology. Each symptom is represented by its corresponding name entry in the `release_evidences.json` file.
- `antecedents`: data structure describing the set of antecedents characterizing the pathology. Each antecedent is represented by its corresponding name entry in the `release_evidences.json` file.

Patient description Each patient in each of the 3 sets has the following attributes:

- `AGE`: the age of the synthesized patient.
- `SEX`: the sex of the synthesized patient.
- `PATHOLOGY`: name of the ground truth pathology (`condition_name` property in the `release_conditions.json` file) that the synthesized patient is suffering from.

- **EVIDENCES**: list of evidences experienced by the patient. An evidence can be either binary, categorical or multi-choice. A categorical or multi-choice evidence is represented in the format `[evidence-name]_@[evidence-value]` where `[evidence-name]` is the name of the evidence (name entry in the `release_evidences.json` file) and `[evidence-value]` is a value from the `possible-values` entry. Note that for a multi-choice evidence, it is possible to have several `[evidence-name]_@[evidence-value]` items in the evidence list, with each item being associated with a different evidence value. A binary evidence is simply represented as `[evidence-name]`.
- **INITIAL_EVIDENCE**: the evidence provided by the patient to kick-start an interaction with an ASD/AD system. This is useful during model evaluation for a fair comparison of ASD/AD systems as they will all begin an interaction with a given patient from the same starting point. The initial evidence is randomly selected from the evidence list mentioned above (i.e., **EVIDENCES**) and it is part of this list.
- **DIFFERENTIAL_DIAGNOSIS**: the ground truth differential diagnosis for the patient. It is represented as a list of pairs of the form `[[patho_1, proba_1], [patho_2, proba_2], ...]` where `patho_i` is the pathology name (`condition_name` entry in the `release_conditions.json` file) and `proba_i` is its related probability.

A.8.2 Responsibility statement

The authors declare that they bear all responsibility for violations of rights related to this dataset.