ATENA-PRO: Generating Personalized Exploration Notebooks with Constrained Reinforcement Learning

Tavor Lipman
Tel Aviv University
Tel Aviv, Israel
tavorlipman@mail.tau.ac.il

Tova Milo
Tel Aviv University
Tel Aviv, Israel
milo@cs.tau.ac.il

Amit Somech
Bar-Ilan University
Ramat Gan, Israel
somecha@cs.biu.ac.il

ABSTRACT

We present ATENA-PRO, a framework for generating *personalized* data exploration notebooks, given an input dataset and user preferences. Via a dedicated wizard interface, users first specify their information needs from the desired exploration notebook. These specifications, alongside the input dataset, are fed to a *constrained* Deep Reinforcement Learning (CDRL) framework. Our CDRL framework is based on ATENA, a general-purpose DRL architecture for data exploration, augmenting it with a new *compliance reward scheme*, and a *specification-aware neural network architecture*, both crucial for the generation of personalized notebooks.

We demonstrate ATENA-PRO by inviting participants to explore several real-world datasets under various analytical tasks, pose their preferences, and examine the personalized exploration notebooks generated by our system.

CCS CONCEPTS

• Mathematics of computing → Exploratory data analysis; **• Information systems →** *Data analytics*; **• Theory of computation →** *Reinforcement learning*.

KEYWORDS

Automated Data exploration; AI for Data Analytics

ACM Reference Format:
Tavor Lipman, Tova Milo, and Amit Somech. 2023. ATENA-PRO: Generating Personalized Exploration Notebooks with Constrained Reinforcement Learning. In *Companion of the 2023 International Conference on Management of Data (SIGMOD-Companion '23), June 18–23, 2023, Seattle, WA, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3555041.3589727

1 INTRODUCTION

Data exploration is a challenging process performed by analysts and data scientists. Given an analytical task or a question over a dataset, the user sequentially employs exploratory queries on the data, such as filtering and grouping, in order to gain relevant insights for her task.

A popular, effective way to streamline this process is to examine existing data exploration notebooks created by other data scientists on the same dataset. Since, naturally, these notebooks may not always be available, we devised ATENA [1] – a system for automatically generating general-purpose exploration notebooks, given an input dataset.

However, general-purpose exploration notebooks (that do not consider the specific task at hand) are suboptimal, since in practice, even when working on the same dataset, users may have completely different information needs, derived from different analysis tasks.

To this end, we present ATENA-PRO, a new system for auto-generating personalized, task-relevant exploration notebooks. Using a notebook-builder interface, users first describe their information needs and preferences. Our interface allows for specifying desired properties of the exploratory notebook's structure, operations, and more importantly – its *continuity*, which ties together subsequent operations without specifying them explicitly. The interface is backed by LDX, a dedicated Specification *L*anguage for *D*ata E*X*ploration. The dataset, together with the LDX specifications derived from the user's information needs, are passed to a Constrained Deep Reinforcement Learning (CDRL) engine, a special type of DRL architecture designed to handle extra requirements, such as safety constraints [4]. Our CDRL engine is built on top of the generic exploration framework of [1] but uses a special reward scheme and neural network architecture designed to generate an interesting exploration notebook – that complies with the input specifications.

Example Use Case. Data Scientists Clarice and Dolly both work for a large media production company, and were assigned to analyze the Netflix Movies and TV Shows dataset [3]. Clarice's assignment is to discover countries with different, atypical viewing habits, compared to the rest of the world. Dolly, however, is assigned a different task on the same data – investigating the properties of "successful" TV shows that have more than one season. Clearly, these two tasks pose different exploration challenges.

Both analysts use the ATENA-PRO notebook builder UI, as depicted in Figure 1, to specify their preferences. As explained in more detail below, Clarice uses two *Comparison Components* to compare an (unknown yet) country to the rest of the world, while Dolly uses *Focus Components* in order to zoom in on interesting properties (to be discovered by ATENA-PRO) of TV shows with more than one season.

Figure 2a and 2b depicts (samples of) the exploration notebooks generated by ATENA-PRO for Clarice and Dolly.

As seen in Figure 2a, for the Atypical Country task the system chose to concentrate on *India*, and "compare" it to the rest of the world using a *count* aggregation over the attributes *rating* and *show type*. The exploratory steps show interesting and relevant insights, demonstrating how India is different from the rest of the world. For example, the *rating* comparison (See Cells 2 and 3), shows that

SIGMOD-Companion '23, June 18–23, 2023, Seattle, WA, USA

ACM ISBN 978-1-4503-9507-6/23/06.
https://doi.org/10.1145/3555041.3589727

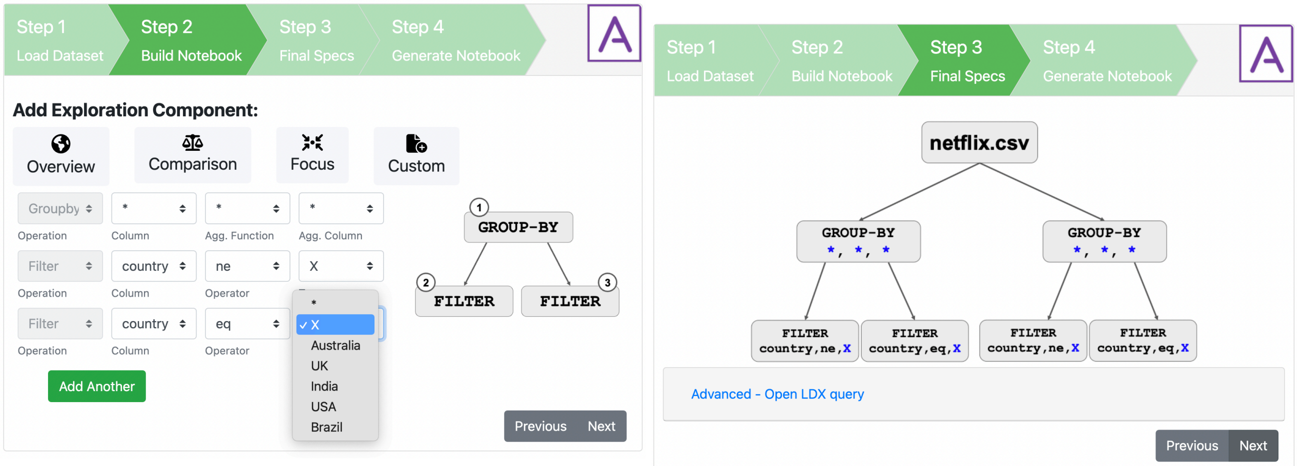

(a) Adding Exploration Components (b) Template Exploration Tree

Figure 1: ATENA-PRO GUI Notebook Builder. Specifications are described by adding ready-made exploration components (Figure 1a), or creating custom ones from scratch. Unfilled parameters (Marked in blue in Figure 1b) are then instantiated by ATENA-PRO CDRL Engine.

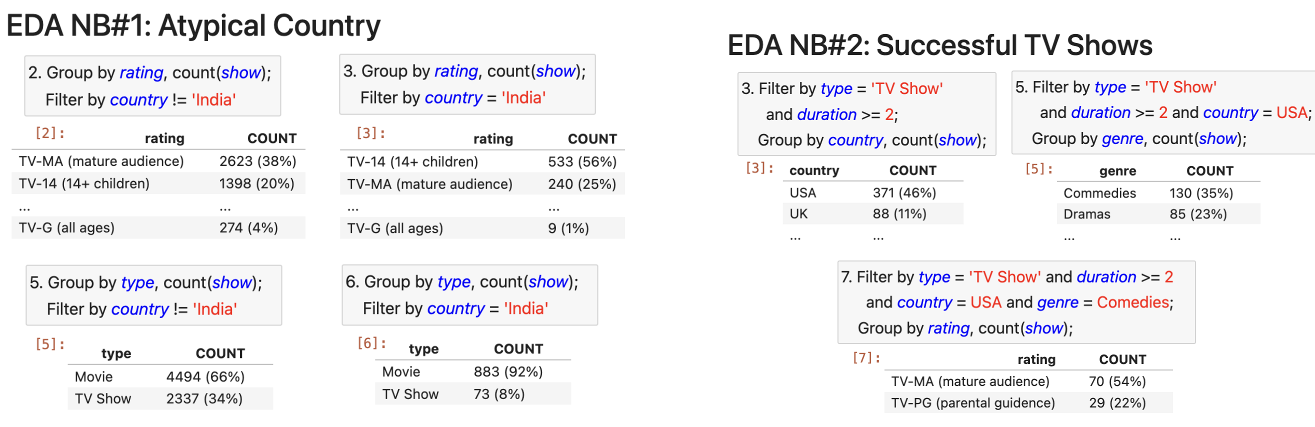

(a) "Atypical Country" Exploration Notebook (b) "Successful TV Shows" Exploration Notebook

Figure 2: Personalized Exploration Notebooks Generated by ATENA-PRO. The user specifications are translated to LDX, and fed to the ATENA-PRO CDRL engine, which then generates a compliant notebook. Each notebook cell depicts a query operation and its results.

whereas in the rest of the world the majority of titles are rated TV-MA (i.e., 17+), most titles in India are rated TV-14 (14+).

For Dolly's task, ATENA-PRO generates a completely different notebook (Figure 2b), exploring interesting properties of TV shows with two seasons or more. For instance, it first examines the *country* attribute of the successful shows (Cell 3), demonstrating that *most successful TV shows are made in the US.* Then, it focuses on such US-based successful shows, and examines their *genre* attribute (Cell 5), revealing that *successful American TV shows are mostly comedies.* □

Related Work. Numerous works seek to automate the process of data exploration using machine learning solutions (See [8] for an overview). In particular, systems such as [1, 9] use Deep Reinforcement Learning (DRL) to generate full exploratory sessions for a given input dataset. As mentioned above, such a generic output may often be unsuitable for the specific analytical task at hand.

Last, while our prototype system focuses on data manipulation operations, other works deal with recommendations of data visualizations. This is a parallel effort, as systems such as LUX [6] can pair adequate visualizations to the exploratory steps made by ATENA-PRO.

2 TECHNICAL OVERVIEW

2.1 Background & System Overview

Following [1], in the general DRL environment for data exploration, a neural-network–based agent *interacts* with an input dataset, with the goal of finding a *good* sequence of exploratory operations. This is done as follows:

Programmatic Exploratory operations. At each step in the exploratory session, the agent employs a query operation on either the raw data

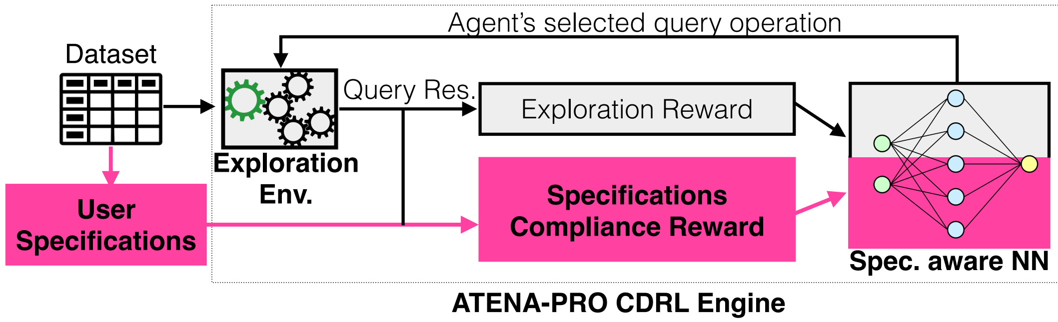

Figure 3: ATENA-PRO System Architecture

or intermediate results. Our prototype system currently supports the following, commonly used exploratory operations: (1) *Filter.* a filter operation is defined by `[F,attr,op,term]`, where `attr` is an attribute, `op` is a comparison operator (e.g. =, ≥, *contains*) and `term` is a numerical/textual term to filter by. (2) *Group-by and Aggregate.* This operation is defined by `[G,g_attr,agg_func,agg_attr]`, i.e., grouping on attribute `g_attr`, aggregating on column `agg_attr` with aggregation function `agg_func` (e.g., min, max, count). Last, an additional (3) *Back* command is supported, allowing the agent to go back to previous results in order to start a new exploration path.

Exploration Reward. The agent obtains a positive or a negative reward for its operations, encouraging exploratory queries that are: (i) interesting - by employing data-driven notions of *interestingness.* For example, the interestingness of a filter is measured using an *exceptionality* [10] measure which favors filter operations that cause a large deviation in value distributions. (ii) Diverse from one another - by computing the Euclidean distance of the vector representation of resulted views, and (iii) coherent, i.e., make sense to a human user - via a rule-based classifier (See [1] for more details.)

Learning Process. The goal of the agent is to learn how to obtain a maximal cumulative reward, by repeatedly interacting with the dataset (without requiring any labeled or training data). Since the space of possible exploration queries is significantly large, [1] suggests a neural network architecture with a "multi-softmax" layer, which generates separate probability distributions for the high-level operation types (filter, group & aggregate, back), then additional probabilities for the corresponding parameters' values.

We next describe the CDRL framework of ATENA-PRO used to generate *personalized* exploration notebooks.

ATENA-PRO system overview. Figure 3 depicts the architecture of ATENA-PRO. The pink elements in the figure mark the dedicated, novel components that are required to process users' specifications. Given an input dataset, the user first specifies her preferences for the desired exploration notebook via the Notebook Builder UI (See Section 2.2). Then, the input dataset and LDX specifications are passed to the CDRL engine of ATENA-PRO, containing (1) a *compliance* reward signal, designed to effectively guide the agent towards the narrow space of specifications-compliant notebooks (See Section 2.3); and (2) a specification-aware network architecture used to encourage the agent to select compliant exploratory steps with a higher probability.

In the CDRL learning process, the agent optimizes its policy via *both* the exploration reward and the compliance reward. Once done, the highest-scoring exploration process is presented in a notebook interface, as illustrated in Figure 2.

2.2 Exploration Specifications Interface

Following [1] we use a tree-based exploration model, s.t. each query operation is a node in the exploration tree, and is applied on the results of a parent operation. The "root" is the original dataset before any operation is applied.

Our exploration notebooks builder UI (See Figure 1) allows users to specify their requirements by gradually building a "template" exploration tree – adding nodes and partially (or fully) specifying the operations. ATENA-PRO will then complete the unspecified operations and parameters using the CDRL engine, as described below.

To further consider the *contextual* connection between query operations, ATENA-PRO also supports *continuity variables.* These variables are used to *match* between parameters of different operations, without explicitly stating them. For illustration, consider the following example.

Example 2.1. Figure 1 depicts the notebook construction process for the task of finding an atypical country. As explained in Section 3, users can also use a set of ready-made exploration components, to facilitate the tree construction. The user thus begins by creating a *comparison component*, containing a group-by followed by two subsequent filters, and specifying operations preferences (Figure 1a)): The group-by operation is left blank (to be instantiated by ATENA-PRO). Then, the two filters are specified to operate on the attribute *'country'*, where the *unspecified* filter term is captured by a continuity variable X. This ensures that ATENA-PRO will apply both filters on the same country, s.t. the first one filters *in* on that country, and the second filters *out.* The user then duplicates this tree component, review the final template tree (Figure 1b), and clicks "generate notebook".

The final, materialized notebook is generated by the CDRL engine of ATENA-PRO, as explained in Section 2.3. As depicted in Figure 2a, see that ATENA-PRO decides to instantiate the group-by with a *count* on the attribute *rating.* The two filter operations output is depicted in cells 2 and 3, where ATENA-PRO instantiated the continuity variable X to be the term 'India'. Thus comparing the *rating of Netflix titles in India to the rest of the world.* This comparison reveals that in India, most titles are rated TV-14, while in the rest of the world the majority of titles are rated TV-MA (17+). □

As mentioned above, the notebook builder is backed by LDX, a formal specification language for data exploration. LDX is based on Tregex [7], a query language for tree-structured data, but extends it in order to support the continuity variables as described above. The LDX specifications derived from the builder are then passed to our CDRL engine, together with the input dataset, which then generates compliant notebooks (as explained in Section 2.3).

2.3 CDRL Engine Design

Our goal is to ensure that the DRL agent generates exploratory notebooks that not only obtain a high exploration reward but are also fully compliant with the user specifications. Inspired by previous works in constrained reinforcement learning [2], our CDRL framework includes the following components:

(1) Compliance reward signal. We use both an end-of-session reward, which provides a positive reward if the output notebook is compliant with all specifications (and a negative reward otherwise); and an immediate reward, granted after each operation, that penalizes *specific* operations that violate the structural specifications. Our specification verification engine extends the one of TREGEX [5].

(2) Specification-aware neural network. Our neural network adapts its output layer w.r.t. the user specifications via built-in placeholders for operation "snippets" that are derived from the specifications. We therefore add *snippet* as a new high-level operation that can be activated by the agent. Snippets function as operation "shortcuts", eliminating the need for composing full, compliant operations from scratch. For example, using a snippet of `'F, Country, eq'` will only require the agent to choose a filter term.

The architecture derivation process is as follows (See the pink elements in Figure 4). First, given the user's specifications, we generate snippet neurons for the operational specification. All snippet neurons, as depicted in Figure 4, are connected to a dedicated multi-softmax segment, denoted σ_{snp}. The unassigned parameters of each snippet are then wired to the individual softmax segments of the parameters.

3 INTERACTIVE DEMONSTRATION

We implemented our notebook builder in React and Python 3, and the CDRL engine of ATENA-PRO using ChainerRL.

To further facilitate the exploration tree construction by the user, we incorporated several template exploration components (See Figure 1a): The *Comparison* component is the group-by with two subsequent filters as described in Example 2.1. The *Overview* component is a filter operation followed by 2 group-by ones, hence showing two different aggregation views for the *same* data subset. The *Focus Element* consists of filter and group-by operations, such that each filter focuses on a particular group found in the results of a previous grouping operation.

In our interactive demonstration, we invite the participants to compose their own, personalized exploration notebooks via the ATENA-PRO notebook builder UI.

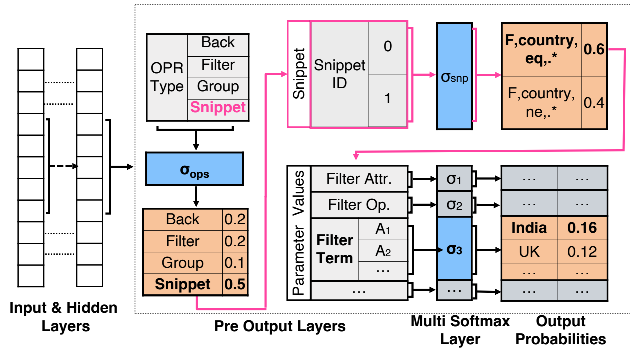

Figure 4: Specification-Aware Network Architecture

After a brief explanation of the UI, the audience is invited to upload their own dataset or select one from our collection, which contains numerous Kaggle datasets and analysis tasks such as *"find reasons for flight delays during the summer months"* using the Flight Delays Dataset, and *"discover properties of successful Google Mobile Apps, with more than one million installs"*, using the Google Play Store dataset.

When done specifying their requirements for the exploration notebook, users can then review the output notebook generated by ATENA-PRO, and examine the relevance of insights and discoveries it contains. For interested participants, we will provide a look under the hood of ATENA-PRO, showing the underlying LDX specifications of their composed notebook, the custom neural network generated by our CDRL framework, and its learning convergence plots.

ACKNOWLEDGMENTS

We thank Oz Zafar for his contribution to the implementation and presentation of our system. This work has been partially funded by the Israel Science Foundation grant number 2707/22 and the Binational US-Israel Science Foundation grant number 2018194.

REFERENCES

[1] Ori Bar El, Tova Milo, and Amit Somech. 2020. Automatically generating data exploration sessions using deep reinforcement learning. In *SIGMOD*.
[2] Gal Dalal, Krishnamurthy Dvijotham, Matej Vecerik, Todd Hester, Cosmin Paduraru, and Yuval Tassa. 2018. Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757* (2018).
[3] Netflix Dataset. 2022. www.kaggle.com/shivamb/netflix-shows.
[4] Javier Garcıa and Fernando Fernández. 2015. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research* 16, 1 (2015), 1437–1480.
[5] Tregex implementation. 2022. https://github.com/yandex/dep_tregex.
[6] Doris Jung-Lin Lee, Dixin Tang, Agarwal, et al. 2021. Lux: always-on visualization recommendations for exploratory dataframe workflows. *PVLDB* 15, 3 (2021), 727–738.
[7] Roger Levy and Galen Andrew. 2006. Tregex and Tsurgeon: Tools for querying and manipulating tree data structures.. In *LREC*.
[8] Tova Milo and Amit Somech. 2020. Automating exploratory data analysis via machine learning: An overview. In *SIGMOD*.
[9] Aurélien Personnaz, Sihem Amer-Yahia, Laure Berti-Equille, Maximilian Fabricius, and Srividya Subramanian. 2021. DORA THE EXPLORER: Exploring Very Large Data With Interactive Deep Reinforcement Learning. In *CIKM*.
[10] Matthijs van Leeuwen. 2010. Maximal exceptions with minimal descriptions. *Data Mining and Knowledge Discovery* 21, 2 (2010), 259–276.