



Common challenges of deep reinforcement learning applications development: an empirical study

Mohammad Mehdi Morovati¹ · Florian Tambon¹ · Mina Taraghi¹ · Amin Nikanjam¹ · Foutse Khomh¹

Accepted: 17 May 2024 / Published online: 14 June 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

Machine Learning (ML) is increasingly being adopted in different industries. Deep Reinforcement Learning (DRL) is a subdomain of ML used to produce intelligent agents. Despite recent developments in DRL technology, the main challenges that developers face in the development of DRL applications are still unknown. To fill this gap, in this paper, we conduct a large-scale empirical study of **927** DRL-related posts extracted from Stack Overflow, the most popular Q&A platform in the software community. Through the process of labeling and categorizing extracted posts, we created a taxonomy of common challenges encountered in the development of DRL applications, along with their corresponding popularity levels. This taxonomy has been validated through a survey involving 65 DRL developers. Results show that at least 45% of developers experienced 18 of the 21 challenges identified in the taxonomy. The most frequent source of difficulty during the development of DRL applications are *Comprehension*, *API usage*, and *Design problems*, while *Parallel processing*, and *DRL libraries/frameworks* are classified as the most difficult challenges to address, with respect to the time required to receive an accepted answer. We hope that the research community will leverage this taxonomy to develop efficient strategies to address the identified challenges and improve the quality of DRL applications

Keywords Deep reinforcement learning · Machine learning · Deep learning · Stack overflow · Programming issues · Software reliability · Empirical study

Communicated by: Shaukat Ali

✉ Mohammad Mehdi Morovati
mehdi.morovati@polymtl.ca

Florian Tambon
florian-2.tambon@polymtl.ca

Mina Taraghi
mina.taraghi@polymtl.ca

Amin Nikanjam
amin.nikanjam@polymtl.ca

Foutse Khomh
foutse.khomh@polymtl.ca

¹ SWAT Laboratory, Department of Computer and Software Engineering, Polytechnique Montréal, 2500 Chem. de Polytechnique, Montréal, QC H3T 1J4, Canada

1 Introduction

Reinforcement Learning (RL) has begun making its mark across a range of industrial sectors, from autonomous vehicles (Aradi 2020) and traffic engineering (Xiao et al. 2021) to healthcare systems (Yu et al. 2021). Recently we have been also witnessing an increasing adoption of RL to solve different software engineering tasks, from automatic code improvement (Wan et al. 2018), to test case prioritization (Bagherzadeh et al. 2021), and program debloating (Heo et al. 2018). Reinforcement Learning differs significantly from other subcategories of Machine Learning (ML) such as supervised and unsupervised learning, as it includes an agent that interacts with an environment to learn how to perform a sequence of actions leading to the best cumulative final rewards (Nikanjam et al. 2022). In other words, in RL, an agent learns to act in a way that modifies its behavior gradually to achieve the best final result; which makes traditional software quality assurance techniques inadequate for RL.

Deep Reinforcement Learning (DRL) is an integration of Deep Learning (DL) and RL, also known as Deep RL, to address challenges, such as high-dimension input data (Arulkumaran et al. 2017). Combining DL and RL enables DRL to discover compact low-dimensional representations of high-dimensional data automatically (Arulkumaran et al. 2017).

Although there exist studies on testing and debugging RL programs (Zolfagharian et al. 2022; Tambon et al. 2023), the main challenges and obstacles that developers face while developing RL applications are still unclear (Zhang et al. 2020). Moreover, because of basic differences between the paradigm of traditional software applications and ML applications (Morovati et al. 2023; Islam et al. 2020), it is expected that developers of ML applications face different types of challenges in the implementation process of such applications. Thus, DRL developers may face different challenges from other types of software systems (including traditional software systems as well as other subcategories of ML applications) (Nguyen et al. 2020; Du and Ding 2021; Dulac-Arnold et al. 2021).

As an example, Listing.1 shows a SO post (70562317) related to a DRL application, representing a challenge in implementing the method to choose an optimal action which is specific to DRL development and differs from ML and DL development challenges.

Although there exist some studies regarding challenges in the development of DL (Zhang et al. 2019; Rao and Frtunikj 2018), ML applications (Lwakatere et al. 2019; de Souza Nascimento et al. 2019), to the best of our knowledge there is no study on the challenges that developers face when developing DRL applications. The study by Yahmed et al. (2023) is the most closely related work to this research. It examines the challenges that developers face during the deployment process of DRL systems but does not consider the challenges occurring in the early development phases prior to deployment. In this study, we examine the following research questions:

RQ1. *What are the common challenges of DRL application development?*

RQ2. *How are the identified challenges perceived by DRL practitioners?*

RQ3. *Are DRL application development challenges language- and/or framework-specific?*

To answer these research questions, we manually examined and categorized 927 Stack Overflow (SO) posts that are related to DRL development. We report our results as a taxonomy of challenges in DRL application development. Besides, we conducted a survey of DRL developers/practitioners to validate our findings. Moreover, we investigated the dependency of the identified challenges on programming languages and libraries/frame-

```
1 def act(self, some_input, state):
2     mu, var, state_value = self.model(some_input, state)
3     # mu contains info required for gradient
4     mu = mu.data.cpu().numpy()
5     # mu is detached and now has forgot all the
6     # operations performed in "self.action_head"
7     sigma = torch.sqrt(var).data.cpu().numpy()
8     action = np.random.normal(mu, sigma)
9     action = np.clip(action, 0, 1)
10    action = torch.from_numpy(action/1000)
11    return action, state_value
```

Listing 1 SO post (70562317) showing a challenge in the development of the RL action

works used for DRL development. The contributions of this study are summarized as follows.

- We provide the first large-scale empirical study of the challenges in the development of DRL applications,
- We categorize challenges in DRL application development and propose a taxonomy,
- We conduct a survey with DRL practitioners to validate the identified common challenges of DRL application development,
- We examine the relationship between the identified challenges and the programming languages and libraries/frameworks used to develop DRL applications.

The Rest of the Paper is as Follows We describe the methodology of our study in Section 2. In Section 3, we report our findings including the taxonomy of DRL development challenges. Section 4 discusses the implications of the highlighted findings. Afterward, we review related works in Section 5. Threats to the validity of our research, and conclusion/future works are discussed in Sections 6 and 7, respectively.

2 Methodology

This section describes the methodology we follow in this study. This methodology is illustrated in Fig. 1.

2.1 Extracting Posts from Stack Overflow (SO)

We rely on Stack Overflow (SO) as the main source of information in this study; similar to several previous studies which utilized data exclusively obtained from SO for analyses (Zhang et al. 2019; Alshangiti et al. 2019; Hamidi et al. 2021). SO is known as the largest




Fig. 1 High-level view of the used methodology

technical question and answer (Q&A) website creating a public knowledge base in various areas (Zhu et al. 2022), with 23.4 million questions and 19.6 million users as of December 2022 (StackExchange 2022). In the software development community, SO provides a platform for developers to exchange about coding issues; improving their coding knowledge. To extract SO posts related to DRL, we use Stack Exchange Data Explorer¹ which provides access to up-to-date SO data. Overall, we use a list of DRL-related tags and keywords to collect DRL-related SO posts. Listing 2 presents an example of a query used to collect SO posts containing ‘*deep-learning*’ and ‘*reinforcement-learning*’ tags at the same time.

```
1 SELECT * FROM Posts
2 WHERE Tags LIKE '%<deep-learning>%'
3 AND Tags LIKE '%<reinforcement-learning>%'
```

Listing 2 Sample query to extract SO posts that have ‘*deep-learning*’ and ‘*reinforcement-learning*’ tags simultaneously

To gather the list of DRL-related tags and keywords, we follow previous study (Alshangiti et al. 2019) using a snowballing approach in which we start with posts that have ‘*deep-learning*’ and ‘*reinforcement-learning*’ tags, simultaneously. In the next step, we collect all tags assigned to the SO posts gathered in the previous step. Then, we include DRL-related tags to our list of tags (e.g., ‘*dqn*’). We continue this process and expand the list of DRL-related tags until we are unable to add any new tags to our list. Besides, we create a list of DRL-related keywords based on the list of collected DRL-related tags. Firstly, we include all DRL-related tags (such as ‘*reinforcement learning*’) to the list of DRL-related keywords. Moreover, we add expanded forms of DRL-related tags which are acronyms. As an example, we add ‘*Deep Q-Learning*’ which is the expanded form of ‘*dqn*’. The complete list of tags and keywords used to extract SO posts are available in our replication package (Morovati et al. 2023). In summary, we collected SO posts that meet at least one of the following criteria:

- Posts having one of the identified tags (e.g., ‘*drl*’, ‘*dqn*’, etc.)
- Posts with a combination of identified tags (e.g., combination of ‘*deep-learning*’ and ‘*reinforcement-learning*’)
- Posts with a combination of identified tags and keywords in their title or body (e.g., ‘*reinforcement-learning*’ tag and ‘*deep*’ in the post title)
- Posts including identified keywords in their title or body (e.g., ‘*drl*’)

After extracting all posts and removing duplicates, we obtained 3,083 posts. Then, we filtered out the posts without an accepted answer which leaves us with 927 posts. We chose to remove posts without an accepted answer, similar to previous studies (Nikanjam et al. 2022; He et al. 2023) because the correctness of any of such responses would not be inferable, which could potentially bias our results. We also collected information about the time taken by each post to receive an accepted answer and used it as an indicator of the level of difficulty of the question, similar to the approach employed in previous works (Haque et al. 2020; Zahedi et al. 2020).

2.2 Manual Inspection

During this step, a team of four raters (three Ph.D. candidates and a senior research staff who all are practitioners of DRL development) is responsible for labeling the collected SO posts. Following a methodology similar to prior works (Humbatova et al. 2020; Islam et al.

¹ <https://data.stackexchange.com/stackoverflow/query>

2019), we split the collected SO posts into 10 parts, each of which is inspected in a dedicated labeling round. All the discussions and referenced source codes in each post are thoroughly reviewed. The raters use an open coding method (Lune and Berg 2017) to label SO posts and categorize them. Each SO post is reviewed by two raters. We use the “Google Sheets” platform (GoogleSheet 2020) to save all extracted labels in an online environment. That is, all raters put generated labels in the shared document, but they do not have access to the labels that other raters assigned to each post. After labeling SO posts in each round, the raters meet to discuss disagreements and resolve conflicts. In the case they fail to resolve a disagreement, a third rater reviews the SO post and makes a decision about its label, acting as a tie-breaker. Besides, the raters review the generated labels in each meeting to ensure their comprehensiveness and granularity (combining similar labels generated by different raters or dividing a label into separate ones). The justification for dividing SO posts into 10 labeling rounds, with approximately 10% of the total collected posts in each round, emerges from the necessity for raters’ discussion in iterative meetings. These discussions are essential for resolving labeling conflicts and reaching an agreement, reviewing generated labels, and creating finer or coarser-grained labels. In case of changing existing labels, raters re-review the previously labeled posts to ensure that assigned labels are in synchronization with newly generated labels. We made this decision to allow for continuous improvement of the labeling process. This way raters have the opportunity to resolve their conflicts at the end of each round, similar to the technique used in previous studies (Humbatova et al. 2020; Islam et al. 2019). Besides, in the case that any of the raters suggests generating a new label, all raters meet to discuss and reach an agreement on that new label. After completing labeling all 927 posts, all raters meet to finalize the generated labels, categorize them, and create the taxonomy. Then, the first two authors review all of the labeled posts again to ensure that their assigned labels are in sync with the final generated taxonomy. Regarding the posts in which the questioner asks more than one question belonging to multiple categories of challenges, we repeat that record in our dataset and assign a different label to each record. For instance, post #45382763 has been identified to belong to two categories, comprehension and design problems. Although we could not report inter-rater agreement level due to the lack of prior defined categories, we calculate inter-rater agreement between the pair of raters who investigate each SO post after finalizing the labels using Cohen’s Kappa (McDonald et al. 2019) and obtained an 86% agreement level. Table 1 presents detailed information on the labeling procedure.

During the manual inspection of SO posts, we filtered out 57 posts that were not related to DRL development. Generally, some questioners may add DRL-related tags to their posts by mistake or because of unfamiliarity with DRL and its real capabilities to solve their considered problems (such as #60958362). We also filtered out posts that are too general and which could not be considered as reporting about a challenge in DRL development (e.g., #3972812).

It is worth mentioning that 70% of the DRL-related questions in SO still remain without any accepted answer. This is consistent with previous findings by Alshangiti et al. (2019) that 61% of ML-related SO posts remain without any accepted answer. Multiple factors could explain this finding. In some cases, the person who asks the question responds to the question after a while, but she does not assign the accepted answer badge to the post (e.g., #45364837). Some users also ask basic questions irrelevant to DRL but assign DRL-related tags to them. These questions receive negative scores and remain without any accepted answer (e.g., #50544568). We also observe posts where the person asking the question forgot to assign the accepted answer badge to an answer, based on upvotes to the response, comments of the person who asked the question, or other people with the same problem under the response

Table 1 Detailed information on the manual labeling process

Round	Analyzed SO posts	Conflicts	Relevant to DRL	New categories
1	100	11	95	15
2	100	7	93	1
3	72	10	67	–
4	100	16	96	4
5	100	14	90	1
6	100	9	96	–
7	100	5	95	–
8	100	7	93	1
9	100	5	94	–
10	55	3	51	–
Total	927	87	870	22

The boldface cells mean the total of all values on the column

(e.g., #63250935). About the posts with accepted answers, it should be mentioned that 16% of them have been answered by the user who published the question. This usually happens when a user asks a very specific question that remains unanswered for a long time, and then the same user finds the response elsewhere and adds it to his original post (e.g., #2723999).

2.3 Taxonomy Construction and Validation

Similar to previous studies (Vijayaraghavan and Kaner 2003; Humbatova et al. 2020), we use a bottom-up methodology to create the taxonomy. In fact, after completing each labeling round, we categorize all generated labels belonging to a similar theme into a group. Next, we build up parent nodes in a way that makes sure that categories and their subcategories adhere to a ‘*is a*’ relationship. Concerning that the raters may provide new labels during each labeling round, we need to update the taxonomy which means adding a new category, a new subcategory, or combining two categories/subcategories. After any update on the taxonomy which leads to a new version of it, all the authors have a debate on the newly generated version of the taxonomy in a group meeting. By completing the final labeling round and integrating all updates on the taxonomy, all paper’s authors carry out a careful inspection of the produced taxonomy (including all categories and subcategories) over a meeting and finalize it.

Interviews and surveys stand as the two popular methods to validate the results of qualitative studies (Hove and Anda 2005; Aldhaen 2020). Considering the advantages of conducting surveys, including cost-effectiveness, generalizability, reliability, and versatility (DeCarlo 2018; Nekkanti and Reddy 2016), we assessed the comprehensiveness and representativeness of the obtained taxonomy using a survey with DRL developers/practitioners who are not involved in the construction of the taxonomy. Nevertheless, it is noteworthy that several preceding studies have presented their findings without undergoing any validation process (Zhang et al. 2019; Islam et al. 2019).

While we build our taxonomy based on SO posts, we use Github to identify potential respondents to our survey. We collect a list of survey participants from collaborators of GitHub repositories related to DRL. Specifically, we extract GitHub repositories mentioning ‘*deep reinforcement learning*’ in their description using GitHub’s search API V3; a rest API that receives a query and returns a list of repositories that satisfy conditions stated in the

query. In other words, we use ‘*deep reinforcement learning*’ as the keyword to generate the search query of GitHub search API V3². Given that GitHub search API limits access to only the first 1,000 results, we follow the methodology used in Morovati et al. (2023) and run several different queries to achieve less than 1,000 repositories for each query. That is, we divide the duration of the search for repository creation date between Jan 1, 2010, and Jan 31, 2023 (the date of running queries) into snapshots of 1 month. Thus, we execute 157 GitHub search requests to collect 7,244 repositories. Subsequently, we filter out forked and disabled repositories. The complete list of repositories and a sample search query to extract DRL-related repositories are accessible via the replication package (Morovati et al. 2023). In the next step, we check the repositories’ contributors and collect the contributors mentioning their email addresses, obtaining 2531 unique email addresses of developers.

We use Qualtrics (Qualtrics 2023), an online survey tool for designing and conducting surveys, to create survey forms. Table 2 presents the structure of our survey questionnaires. The survey starts with general questions regarding the participant’s current role and experience in DRL development (Section 1 of Table 2). Subsequently, we delve into specific questions regarding each DRL development challenge which is pointed out in the finalized taxonomy (Section 2 of Table 2). Concerning the potential complexity and difficulty of comprehending the whole taxonomy as a single figure within the survey, we present challenges (subcategories) in groups based on their respective main categories. Besides, to ensure clarity, a detailed description is given with each challenge offering participants a thorough understanding of the asked challenges. For each identified challenge, we ask three questions including 1) a ‘*yes/no*’ question identifying whether the answerer has faced the identified challenge, 2) the severity of the challenge, and 3) the amount of effort required to address the challenge. If a participant answers ‘*yes*’ to the first question, s/he will have the next two Likert-scale questions which are related to the severity of challenges and the required effort to address them. We also provide a free-text question in the final part of the survey asking the participants about any challenges in developing DRL applications not listed in our provided taxonomy (Section 3 of Table 2). These free-text questions allow us to collect possible challenges that we may miss in the taxonomy. The full survey questionnaire is available in our replication package (Morovati et al. 2023).

We conduct a comprehensive analysis of top DRL-related GitHub repositories to ensure the completeness of the generated taxonomy. Besides, investigating the challenges faced by developers engaged in the development of real-world DRL applications may further enhance the generalizability of the generated taxonomy. Our methodology for selecting DRL-related repositories aligns with established approaches documented in similar prior studies (Morovati et al. 2023, 2024; Humatova et al. 2020). Initially, we extracted 7,244 repositories to identify contributors of DRL-related repositories. Then, we select the top 100 repositories based on the highest number of stars. Subsequently, we conduct searches for ‘*challenge*’, ‘*difficult*’, and ‘*complex*’ keywords within repositories’ documentation, commit messages, and closed issues to identify potential challenges encountered by developers during the development of these repositories. Our search yielded 746 occurrences of the aforementioned keywords within the top 100 DRL-related repositories. Using a confidence level of 95% and a confidence interval of 5%, we conduct sampling, resulting in the selection of 254 instances. In the next step, two raters independently examine 254 randomly selected commits and issues to find out whether they pertained to real DRL-related development challenges. The raters carefully review commit messages, issue discussions, and investigate any changes made in relation to the specified keywords.

² <https://docs.github.com/en/rest>

Table 2 Survey structure

Sec.1	General questions				
	<i>Email address (optional)</i>	<input type="text"/>			
	<i>Years of experience on DRL</i>	<input type="text"/>			
...					
Sec.2	DRL development challenges review				
	DRL Challenges				
	<i>Comprehension</i>		Facing this challenge	Challenge severity	Needed effort to address
	<i>Reward</i>		<input type="radio"/> Yes <input type="radio"/> No	<input type="radio"/> Minor <input type="radio"/> Major	<input type="radio"/> Low <input type="radio"/> Medium <input type="radio"/> High
...			<input type="radio"/> Yes <input type="radio"/> No	<input type="radio"/> Minor <input type="radio"/> Major	<input type="radio"/> Low <input type="radio"/> Medium <input type="radio"/> High
Sec.3	Free-text question				
	<i>Any challenge that we have missed</i>	<input type="text"/>			
	<i>Any suggestion or comment</i>	<input type="text"/>			

3 Results

This section presents and discusses the results of our study. All the materials used in this study, including the collected data, are publicly available online in our replication package at Morovati et al. (2023).

RQ1: What are the common challenges of DRL application development?

The final taxonomy of challenges of DRL application development is arranged into a tree structure with five high-level categories in which leaves (subcategories) refer to the challenges. Figure 2 shows the taxonomy of common challenges in the development of DRL applications. The number in parentheses presented next to each category/subcategory is the absolute number of identified SO posts that are categorized into that category (the absolute frequency of each challenge). To give a better understanding of identified challenges, a brief description of each category/subcategory is provided in the following.

DRL Issues This category focuses on the challenges that developers may face while developing the DRL part of their applications. Challenges that belong to this category are specific to DRL application development. That is, compared to challenges classified under other categories that can be shared among other kinds of ML-related applications (e.g. supervised learning), challenges of the *DRL issues* category are only faced during DRL application development.

- a. *Design Problem*: Instances where the user asks advice for designing a solution and implementing a DRL application for their specific problem or scenario. For example, developers asked for recommendations to implement different parts of DRL applications for Curve Fever or mini-golf games. Another such challenge in this category is related to designing the properties of each object in a tank game. Although we tried to make finer-grained subcategories under this subcategory, we found out that it is impractical to split this subcategory further, as the majority of this subcategory’s posts primarily comprised users seeking guidance on high-level conceptual queries about defining their problem within a DRL context.

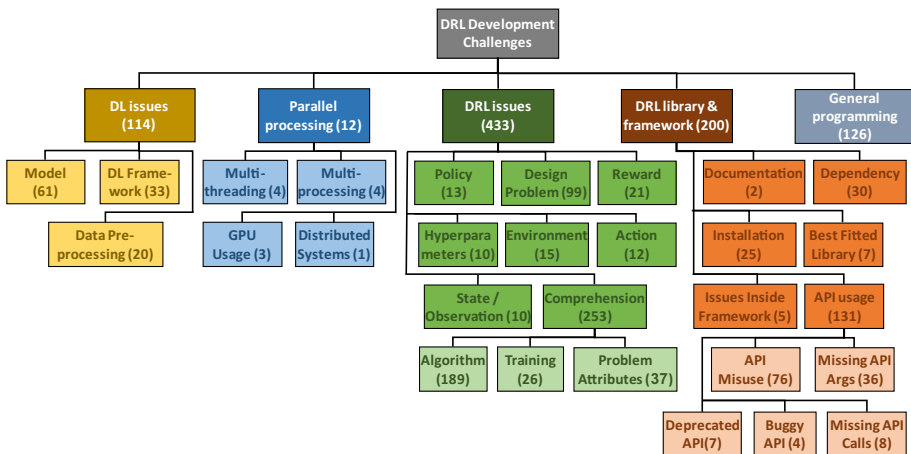


Fig. 2 Taxonomy of common challenges in DRL development. Numbers represented with each category/subcategory indicate the number of SO posts categorized into that category/subcategory

- b. *Comprehension*: Challenges about the meaning or details of theoretical concepts in DRL, i.e., misunderstanding about basic formulas of different DRL algorithms. For instance, a developer mentioned “*I’m trying to make a learning football game from scratch using Deep Q-learning algorithm (without convolutional network though). I just couldn’t figure out what does Φ stand for in this algorithm.*”, or one which is related to the difference between SARSA and Q-learning algorithms in terms of collecting the next policy value.
- b.1. *Training*: This subcategory comprises inquiries concerning the theoretical concepts of the DL model and its significance in DRL applications. The DL model plays an essential role in the training and decision-making process of the DRL agents. Indeed, DL models in DRL applications serve to represent policy of the agent, estimate the value of state-action, and learn mappings from states to values.
- b.2. *Problem attributes*: SO posts containing theoretical queries about attributes of DRL problems (such as reward, action, state, etc.) are categorized in this subcategory. SO posts in this category are related to the conceptual aspects of reward, action, and state, not their implementation details. For example, questions on how to formulate states, actions, and the reward signal for a particular DRL problem, or why a particular definition of the states is not suitable for a problem.
- b.3. *Algorithm*: Questions related to fundamental concepts of various DRL algorithms, such as actor-critic, Q-learning, SARSA, etc., are categorized as algorithm subcategories. The main duty of DRL algorithms is updating parameters of the policy and value function according to the observed states and achieved rewards. Our investigation into these algorithms reveals that Q-learning (28%) and DQN (15.3%) are the most common DRL-related algorithms posing challenges for developers. Conversely, Actor-Critic (3.1%), Proximal Policy Optimization (2.6%), and Trust Region Policy Optimization (0.5%) are the least queried DRL-related algorithms.
- c. *Policy’s Loss*: This category refers to challenges about DRL learning policy’s loss. As an example, questions regarding implementing a customized loss function or any problems in loss calculation methods are categorized in this group.
- d. *Reward*: Challenges in the implementation of reward, e.g., not using the negative reward to penalize each added time step are categorized in this subcategory. An example of this category is a post asking “*I am implementing the basic RL algorithm to play the game Flappy Bird. I want to be able to process the screen and recognize whether a point has been scored or the bird has died. Processing the screen returns a stacked numpy array. The reward function then needs to assign a reward to the provided array, but I have no idea how to go about this*”.
- e. *Action*: Questions/Challenges related to the action(s). e.g., possible actions in a specific game or implementing a ‘chooseAction’ method for a PacMan bot.
- f. *State/Observation*: Questions/Problems regarding the state(s) or the agent observation(s), e.g., handling large state spaces. An example of such instance is a user asking “*I implemented a 3x3 OX game by q-learning (it works perfectly in AI v.s AI and AI v.s Human), but I can’t go one step further to 4x4 OX game since it will eat up all my PC memory and crash...Since I need to calculate each Q value (for each state, each action), I need such a large number of array, is it expected? any way to avoid it?*” and the accepted answer offered suggestions on reducing their state space size by considering symmetries and other tricks.
- g. *Environment*: Questions/Problems pertinent to the environment, e.g., designing a custom environment. For example a user asked “*I’m very new to Ray RLlib and have an issue with using a custom simulator my team made. We’re trying to integrate a custom Python-based*

simulator into Ray RLlib to do a single-agent DQN training. However, I'm uncertain about how to integrate the simulator into RLlib as an environment”.

- h. *Hyperparameters*: Questions/Challenges related to the hyperparameters of the RL algorithm, e.g., setting the discount factor too high. A good example of this group is demonstrated in a post where the user shared the code for the learning algorithm and reported that the loss keeps increasing and the model is not learning. The accepted answer states that *“The main problem I think is the discount factor, gamma. You are setting it to 1.0, which means that you are giving the same weight to the future rewards as the current one”.*

DRL Libraries/Frameworks This category refers to the challenges that developers face when they are trying to use DRL-specific libraries/frameworks (e.g., KerasRL (Plappert 2016), RLlib (Liang et al. 2018), Tensorforce (Schaarschmidt et al. 2018), etc.). Challenges that software developers face when using libraries/frameworks have been extensively studied for traditional software systems development (Decan et al. 2019; Nguyen et al. 2010) and also for DL applications development (Arpteg et al. 2018). However, despite the large number of SO posts related to the usage of DRL libraries/frameworks (i.e., 200 SO posts), these challenges are yet to be examined for DRL application development.

- a. *Installation*: Questions/problems regarding installing/uninstalling DRL-related libraries/frameworks or missing libraries. Issues categorized in this subcategory can often stem from an incompatibility between DRL-related frameworks/libraries and other libraries. For example, a user described her issue as *“when I try to install gym[box2d] I get the following error: I tried: pip install gym[box2d]. on anaconda prompt I installed swig and gym[box2d] but I code in python3.9 env and it still not working (my text editor is pycharm) gym is already installed”.*
- b. *Dependency*: This subcategory includes questions/challenges about the mismatch between versions of installed libraries/frameworks and problems in installed versions of libraries, e.g., when the version of the installed *OpenAI Gym* is not compatible with *Python*. An instance of this subcategory is a user reporting getting an error while installing *OpenAIGym* and the answer pointed out that *“the error means that the package has dependency requirements that conflict with one another”.*
- c. *API usage*: This subcategory includes questions about the usage of arguments, attributes, methods, etc. of an API. It also includes questions about the default values, implemented method, existence of attributes, or methods in an API. An example from this group of issues is a user reporting not knowing how to get the weights of the network using the correct API methods: *“I'm using RLlib to train a reinforcement learning policy (PPO algorithm). I want to see the weights in the neural network underlying the policy. After digging through RLlib's PPO object, I found the TensorFlow Graph object. I thought that I would find the weights of the neural network there. But I can't find them”.* This subcategory is subdivided into five subcategories to delineate more detailed challenges.
- c.1. *API misuse*: This subcategory covers SO posts that mention misunderstanding of API usage. In other words, API misuse occurs when DRL developers try to utilize an API in a manner that is not aligned with its intended purpose.
- c.2. *Missing API call*: Questions related to the absence of necessary API calls within a code snippet are classified in this subcategory.
- c.3. *Missing API args*: When SO posts discuss challenges that DRL developers face due to the absence of one or more essential arguments in an API call, we classify them under this subcategory.

- c.4. *Buggy API*: This subcategory includes SO posts that inquire about calling APIs resulting from the bugs within the implementation of APIs. It is worth mentioning that this challenge is distinct from issues related to the implementation of DRL applications themselves; rather, they pertain to the implementation of the API.
- c.5. *Deprecated API*: In this subcategory, we cover questions about calling a deprecated API which has been altered or removed from the library/framework.
- d. *Documentation (using newly added features)*: This subcategory of issues occurs when a developer wants to use a feature of a DRL library/framework, but there is no documentation for it. For example, a user who could not find the required documentation for the *Neural Network Approximator* in *ReinforcementLearning.jl*: “*I have decided to use a Neural Network Approximator. But the docs do not discuss much about it, nor are there any examples where a neural network approximator is used. I am stuck on how to figure out how to use such an approximator*”.
- e. *Best fitted library for a special task (library suitability)*: Instances where the user asks about the best DRL libraries/frameworks for customizing agents, based on the requirements of the problem. An instance of this group was observed in a post where a user had a customized state space and was looking for a library that supports it: “*I’ve had some luck training an agent using keras-rl, specifically the DQNAgent, however, keras-rl is under-supported and very poorly documented. Any recommendations for RL packages that can handle this type of observation space? It doesn’t appear that openai baselines, nor stable-baselines can handle it at present*”.
- f. *Problems inside DRL frameworks*: Including issues that are encountered because of internal faults, i.e., bugs in the DRL frameworks. For example, a user kept getting a *numpy* error when calling the *model.learn()* function and it was found to be an official bug in the used library *Stable Baselines3*.

DL Issues This category represents the challenges that arise specifically from the DL part of DRL applications. As the challenges belonging to this category are shared by both DL and DRL applications, we use the high-level categories of the taxonomy provided by Humbatova et al. (2020) for DL applications.

- a. *Model*: Questions regarding the DL model including model layer, activation function, load/save model, etc. For instance, a user who was implementing a DRL model asked for advice on back propagation: “*I am struggling with the implementation of the back propagation. Since the rewards are so big, the error values are huge, which creates huge weights. After a few training rounds, the weights to the hidden layer are so big, my nodes in the hidden layer are only creating the values -1 or 1*”.
- b. *Data Preprocessing*: Questions about preparing data to be fed into the DL model, e.g., the shape of the input matrix. As an example, there was a user who asked “*I am learning how to use Gym environments to train deep learning models built with TFLearn. At the moment my array of observations has the following shape: (210, 160, 3). Any recommendations on what is the best way to reshape this array so it can be used in a TensorFlow classification model?*”
- c. *DL framework*: Questions about the usage of DL frameworks (e.g., Keras, TensorFlow, PyTorch, etc.) in development of DL part in DRL applications. For instance, a user who wanted to use *Huber loss* in their model which was written using *Keras*, but this loss function is not readily available in *Keras*, and the accepted answer implements the *Huber loss* function.

Parallel Processing & Multi-threading This category focuses on the challenges associated with running DRL applications as parallel or distributed applications. While running different types of DL applications in a parallel style is a common practice (e.g., using GPU or multi-core CPU), it is important to note that the architecture of DRL applications differs from other types of DL applications. As a result, running a DRL application as a parallel or distributed application introduces new challenges that may not occur in DL application development.

- a. *GPU usage*: Questions/Problems regarding utilizing GPU for running DRL applications. For example, a user was facing performance issues when training a DQN on GPU: “I am try to train a DQN model with the following code. The GPU (cuda) usage is always lower than 25 percent. I know the tensorflow backend is consulting the GPU resources, but the usage is low. Is there any way I can improve the utilization of the GPU (When I train a CNN network, the GPU (cude) utilization is around 70 percent)?”
- b. *Distributed processing*: Questions/Challenges about running DRL applications as a distributed software. For example when a user wanted to implement the *Asynchronous Advantage Actor Critic (A3C)* model for reinforcement learning in the local machine, she posed a question about the possibility of implementing it in a distributed manner: “Would it be easier/faster/better to implement this using the distributed TensorFlow API? In the documentation and talks, they always make explicit mention of using it in multi-device environments. I don’t know if it’s an overkill to use it in a local async algorithm”.
- c. *Multi-threading*: Questions/Problems about running DRL applications as a multi-threaded software. An example of this category was observed in a post where the user asked about the possibility of reducing the DRL application’s training time by running it on multiple threads concurrently: “My friend and I are training a DDQN for learning 2D soccer. I trained the model about 40.000 episodes but it tooks 6 days. Is there a way for training this model concurrently? For example, I have 4 core and 4 thread and each thread trains the model 10.000 times concurrently. Therefore, time to training 40.000 episodes are reduced 6 days to 1,5 days like parallelism of for loop”.
- d. *Multi-processing*: This subcategory refers to challenges stemming from running DRL applications on multiple processing units (e.g., multiple CPUs). As an example, a user asked a question about using Ray³ in a multi-processing style.

General Programming Issues This category contains programming and coding mistakes occurring when developing DRL applications. The challenges in this category could not be classified in any of the other categories described above. For example, a user had a question about how to slice a 3D *numpy* array in an RL finite MDP application (#67089715).

Overall, the majority of the analyzed SO posts have been assigned to categories *Comprehension*, *Design problem*, *Model*, and *API Usage*. Aside from *Design problem* related questions, which are often quite specific (i.e., related to particular implementations of DRL), the majority of questions asked by DRL developers concern issues that apply to DRL applications in general.

³ <https://www.ray.io/>

Finding 1: The taxonomy of challenges in DRL development is structured into five main categories including *DRL issues*, *DL issues*, *DRL libraries/frameworks*, *parallel processing & Multi-threading*, and *general programming issues*, where *DRL issues* (48.9%) and *DRL libraries/frameworks* (22.6%) categories are the most popular. Among the challenges, *comprehension* (28.6%), *API usage* (14.8%), and *design problem* (11.2%) are the most prevalent DRL development challenges. *Comprehension* (58.4%). This can however be explained by the fact that survey respondents being more experienced in DRL which would limit such questions.

Table 3 shows the number and percentage of the SO posts within our dataset having different tags and/or keywords that we used in Section 2.1. Figure 3 presents the distribution of SO posts related to DRL application development for a period of 13 years, from 2009 to

Table 3 Detailed information regarding tags and keywords used to extract DRL-related SO posts

	Tag/ Keyword	PostsDRL (#)	DRL (%)	DRL LF (%)	DL (%)	PP (%)	Gen (%)
Tags	Reinforcement-learning	710	96	58	86.8	83.3	54.8
	Openai-gym	208	10	59	10.5	25	25.4
	Tensorflow	159	11.3	14	47.4	16.7	20.6
	Q-learning	137	25.2	2	13.2	–	7.1
	Deep-learning	127	13.4	6.5	32.5	–	15.1
	Keras	78	5.1	6	29.8	–	7.9
	Pytorch	65	6.5	3	18.4	–	7.9
	Dqn	39	4.4	2	7	8.3	5.6
	Stable-baselines	24	–	5.5	1.8	–	5.6
	Keras-rl	19	–	3.5	5.3	–	2.4
	Tensorflow-agents	12	0.5	2.5	2.6	–	1.6
	rllib	9	0.5	4	–	–	–
	Tensorboard	8	–	2	–	–	2.4
Starcraftgym	5	–	2	–	–	0.8	
Keywords	Reinforcement learning	313	44.8	24	36	41.7	19.8
	Deep	136	16.4	10.5	23.7	16.7	11.9
	ppo	117	11.5	18	11.4	8.3	13.5
	Deep reinforcement learning	37	4.6	2.5	5.3	8.3	4
	a3c	17	1.6	0.5	3.5	25	1.6
	Deep q-learning	16	2.5	1	0.9	–	1.6
	Actor critic	11	1	–	4.4	8.3	0.8
	ddqn	9	–	1	1.8	8.3	–
	drl	5	0.7	0.5	–	–	0.8
	dql	5	0.7	0.5	0.9	–	–
	Deep rl	2	0.5	–	–	–	–
	trpo	2	0.5	–	–	–	–
	Hierarchical	1	0.2	–	–	–	–
	Trust region	1	0.2	–	–	–	–
	Policy optimization						

Columns show the number (#) and percentage (%) of SO posts categorized in the main categories having mentioned tags/keywords ('DRL', 'DRL LF', 'DL', 'PP', and 'Gen' refer to 'DRL issues', 'DRL libraries/frameworks', 'DL issues', 'Parallel processing', and 'General programming', respectively)

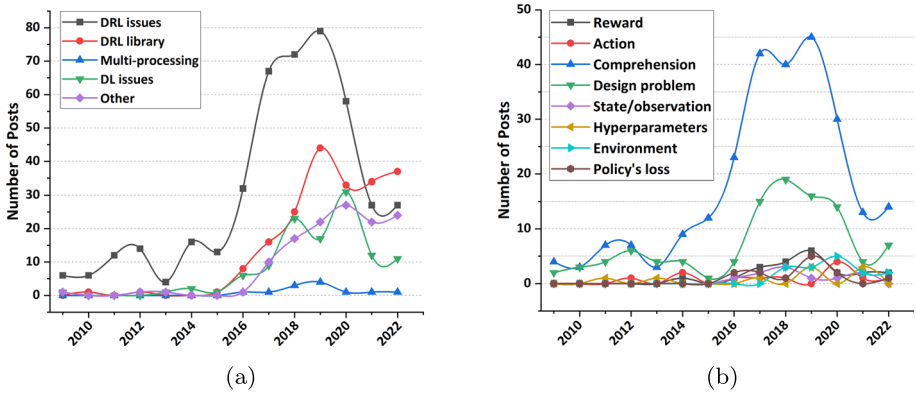


Fig. 3 Distribution of DRL related SO posts per years (a) high-level categories of the provided taxonomy, (b) subcategories of 'DRL issues' category

2022. From Fig. 3(a), we observe a substantial surge of inquiries about DRL development in 2016; reaching a peak in 2019 and 2020. Additionally, as can be seen on Fig. 3(b), *comprehension* and *design problem* questions dominated posts about DRL application development challenges. It is also noticeable that *API usage*, the second most common DRL application development challenge, was at its peak in 2018.

Figure 4 depicts the distribution of time taken by SO posts from different categories to receive an accepted answer. This duration is an indicator of the difficulty level of the questions mentioned in the SO posts in the development of both traditional (Haque et al. 2020; Zahedi et al. 2020) and ML software (Alshangiti et al. 2019; Chen et al. 2020). *Parallel Processing* is the category with the highest average time taken before receiving an accepted answer.






Fig. 4 Duration to receive an accepted answer (hours)

This can be explained by the fact that using multi-processing or distributed processing in DRL is not necessarily widespread and also requires particular knowledge and expertise. The remaining categories need a nearly similar average time frame to receive an accepted answer, with the *general programming issues* category having the shortest average duration. We attribute the shortest average time of the *general programming issues* category to the fact that this category contains generic challenges that do not require expertise in DL or DRL.


Although we note a high proportion of outliers in Fig. 4, we have a relatively low median, first, and third quartiles, for all categories compared to the average. Meaning that, although the majority of posts received an accepted answer in a relatively short period of time (in general less than 10 days), a sizable number of posts required a longer time (more than 20 days). Therefore, even within categories, there is quite some discrepancies inside the subcategories (challenges) themselves. As can be seen on Fig. 5(a), SO posts categorized as *Model* have a higher average waiting time to receive an accepted answer than the other subcategories of the *DL Issues* category. Among subcategories under *DRL Issues*, *hyperparameters* includes SO posts with the highest average required time for receiving an accepted answer (Fig. 5(b)). About *DRL libraries/frameworks*, SO posts belonging to *API usage*, *Dependency* and *Installation* subcategories require the longest average time before receiving an accepted answer (Fig. 5(c)). Regarding subcategories within *parallel Processing & multi-threading*, it is notable that the average duration required to receive an accepted answer for SO posts




(a) DL issues



(b) DRL issues



(c) DRL libraries/frameworks



(d) Parallel processing & multi-threading

Fig. 5 SO posts' required time to receive an accepted answer

classified under *multi-processing* subcategory exceeds one year. It should be also taken into account that the small number of SO posts in this subcategory might bias the results, with respect to the fact that small data may not represent the distribution of classes in the population adequately (Bruer et al. 2015).

Although our methodology to measure the difficulty level of addressing challenges aligns with prior studies analyzing SO posts (Alshangiti et al. 2019; Decan et al. 2019), it is worth noting that some studies have utilized the number of posts within each category as a metric to illustrate the difficulty associated with addressing the related challenges (Bangash et al. 2019; Hamidi et al. 2021). In the case of considering the number of posts as an indicator of the difficulty level for tackling challenges (the number mentioned along with each category in Fig. 2), the results would closely mirror the difficulty levels observed for the challenges categorized as *DL issues* and *parallel processing* categories depicted in Fig. 5. For instance, *design problem* is the second most difficult challenge in the *DRL issues* category when we use ‘number of posts’ or ‘duration to receive an accepted answer’ to measure the difficulty level. Conversely, the difficulty levels for challenges falling under *DRL issues* and *DRL libraries/frameworks* differ by using these two metrics. For example, as illustrated in Fig. 5(b), *hyperparameters* emerges as the most challenging subcategory within the *DRL issues* category, even though *comprehension* has the highest number of posts. This discrepancy may be attributed to the nature of inquiring about *hyperparameters*, which may necessitate various implementations and a longer time to respond compared to other challenges. Furthermore, addressing SO posts categorized as *comprehension* mainly seeks users’ background knowledge, which does not necessarily require practical implementation or application execution.

Finding 2: The proportion of SO posts containing questions related to DRL application development increased significantly during the period (2016 - 2021), reaching a peak between 2019 and 2020. In terms of average time before the reception of an accepted answer, *Parallel processing* is the category that required the longest period of time, in comparison to the other categories. Overall, we observe a high variance in the time to answer of the posts from the different categories of challenges identified in our taxonomy, suggesting that their impact on the development process of DRL applications is not equal.

RQ2: How are the identified challenges perceived by DRL practitioners?

We cross-check the taxonomy generated based on SO posts using a validation survey. 65 ML practitioners participated in our survey to assess our identified challenges, including 55% researchers (Master’s and Ph.D. students, research assistants, and professors), 27% ML/SE engineers, 11% developers, and 7% data scientists. Among our respondents, 86% have at least 1 year of DRL development experience and 46% have more than 3 years of experience. Table 4 summarizes the responses of participants for each DRL development challenge contained in the taxonomy. For each challenge, we provide the percentage of developers who reported having experienced that challenge (based on the answers to the ‘yes’ or ‘no’ question). Besides, we asked the participants about the severity of each challenge and the effort required to address them. Results show that all challenges presented in our taxonomy were encountered by the survey respondents. Moreover, no additional challenges were proposed by the survey participants through the open-text questions. This indicates that our provided taxonomy is representative of challenges faced by developers during the DRL application development.

It is worth noting that the survey questions incorporate the challenges depicted in the third level of the taxonomy (Fig. 2), ensuring the survey's conciseness.

According to the survey results, the majority of our respondents have been confronted with challenges classified as *DRL issues* (68.9% average over all subcategories in this category). This observation is aligned with the proportion of SO posts categorized as *DRL issues* (see Fig. 2). Conversely, challenges belong to *Parallel Processing & Multi-threading* category have been experienced the least with only 45.25% of respondents (which is the lowest proportion among all categories) reporting having faced challenges leveraging parallel processing. This finding is reflected by the results of our quantitative analysis of SO posts, which show that only 1.5% of posts contained questions related to the *Parallel Processing & Multi-threading* category. It should be also taken into consideration that previous research showed that there is a growing trend toward the number of studies on DRL (Panzer and Bender 2022; Kiran et al. 2021). Among the challenges identified in our taxonomy, developers who participated in our survey specify *reward* (86%), *environment* (83%), *hyperparameters* (80%), and *design problem* (75%) as the most common challenges in DRL development. Although only 14.8% of SO posts contained questions about *API usage*, 38% of survey respondents identified it as a challenging issue in DRL development. Given that *reward*, *environment*, *hyperparameters*, and *design problem* are fundamental components of an RL application (Lorenz 2022), it is expected that survey participants reported them as the most encountered challenges. For instance, defining the environment is known as a crucial step in an RL application development process that affects the convergence of an agent's behavior significantly (Reda et al. 2020).

This however contrasts with the number of SO posts identified as *DRL environment* (i.e., 1.6% of SO posts). We also note that *comprehension*, the most frequent challenge in terms of the number of SO posts in our taxonomy (28.6% of SO posts), has been reported by 52% of survey participants as a non-challenging issue. The explanation for this variance lies in the experience level of the survey respondents. Indeed, it is indicated that 84% of the survey participants have at least 1 year of experience in DRL development. In other words, experienced practitioners are less likely to seek help for understanding the fundamental DRL concepts because they have already mastered these basics. Moreover, it can be interpreted as the fact that the most challenging steps in the development of DRL applications for DRL practitioners are related to providing an optimized solution for various DRL-related problems, not just addressing a DRL-related problem. Therefore, to fulfill the specific requirements of various DRL developers with different experience levels, it is important to acknowledge that DRL developers have unique needs at different stages of the DRL development journey. Moreover, it should be taken into consideration that the survey was conducted in 2023; nearly a decade after DRL started to become mainstream (Mnih et al. 2015; Li 2017).

To enhance the completeness of our provided taxonomy, we scrutinize 254 sampled commits and issues extracted from real-world DRL-related repositories (Section 2.3). Upon thorough examination of the sampled commits, we did not identify any instances of challenges being mentioned in relation to the DRL. Furthermore, an analysis of sampled closed issues revealed a consistent pattern wherein users primarily seek support on the utilization of DRL applications offered by the repositories.

According to our 13-year analysis of the distribution of DRL-related SO posts, there has been a drop in the number of SO posts categorized as *DRL issues* after its peak in 2019 (Fig. 3(a) and (b)). This phenomenon may be attributed to various factors. Initially, it suggests a potential progression in the mastery of fundamental DRL concepts among developers over these years leading to a reduction in challenges encountered and thereby a decrease in the number of DRL-related posts on SO. This can stem from the fact that

the growth of DRL popularity in the community results in increased accessibility of DRL tutorial resources including books, tutorials, videos, and papers. These resources aid DRL developers in enhancing their understanding of foundational DRL concepts. Besides, these resources mostly address various DRL problems, including repositories of practical DRL examples that facilitate comprehension of DRL concepts. Moreover, as time has passed, the accumulation of SO posts regarding DRL development has delivered a rich source of DRL development challenges. As a result, many DRL developers can potentially find answers to their questions among the existing SO posts. It should be also taken into consideration that previous research showed that there is a growing trend toward the number of studies on DRL (Rao and Frtunikj 2018; Li 2017), so the drop in DRL-related SO posts does not imply a drop in the popularity of research on this topic.

Finding 3: Survey respondents encountered all the challenges included in our taxonomy. They identified *reward*, *design problem*, *environment*, and *hyperparameters* as the most common challenges in the development of DRL applications. However, this observation highlights a notable contrast with the findings from the analysis of SO posts, wherein *comprehension* with 28.6% of SO posts is the most frequent challenge. The variability observed can be attributed to the fact that DRL practitioners mostly face challenges when they are trying to provide solutions for different DRL problems.

We also ask survey respondents about the severity and needed effort to address the challenges identified in the provided taxonomy (Table 4). In general, the majority (exceeding 57%) of the survey respondents indicated that the most frequent challenges from their viewpoint (i.e., *Reward*, *Environment*, *Design problem* and *Hyperparameters*) are major or critical. Moreover, at least 63% of the survey participants considered the level of effort required to address these challenges, to be “Medium” or “High”. The majority (more than 52%) of the survey participants consider the other identified challenges to be of *Minor* severity level, and to require a *Low* level of effort. In general, the participants consider that *Installation* and *API usage* challenges require a low level of effort, which might signal that DRL libraries/frameworks have good documentation and usability in general (Mojica-Hanke et al. 2023).

We compared the time-to-answer of the posts (from different challenges categories) with the effort reported by our survey participants for the different challenges categories and made the following observations:

- *Hyperparameters*, and *design problems* are the subcategories of *DRL issues* that took longer time before receiving an accepted answer. Survey respondents also reported them as severe and requiring a high effort from ML developers.
- The average time required to receive an accepted answer for *State/observation* and *comprehension* SO questions are comparable (even though some SO posts within *comprehension* subcategory took a bit longer to receive an accepted answer than posts belonging to *State/observation*). The survey participants also assessed these two subcategories of challenges (i.e., *comprehension* and *State/observation*) as easy to resolve in general.
- *API usage* and *Dependency* challenges are the groups of challenges that took the longest time before their questions received an answer. This result is in contrast with the survey participants’ estimation of the effort required to fix them.

The severity and required effort reported by the survey participants for each challenge are strongly correlated (using kendall’s tau (Gibbons 1993)) in a positive direction (Frost 2019)

Table 4 Result of the survey of DRL practitioners

Challenges		Response		Severity			Required effort		
		Yes	No	Minor	Major	Critical	Low	Medium	High
DRL Issues	Comprehension	48%	52%	69%	23%	8%	60%	23%	17%
	Reward	86%	14%	32%	37%	31%	28%	28%	44%
	Action	49%	51%	74%	11%	15%	66%	20%	14%
	Environment	83%	17%	32%	34%	34%	31%	29%	40%
	Hyperparameters	80%	20%	43%	32%	25%	37%	26%	37%
	Design problem	75%	25%	38%	37%	25%	34%	31%	35%
	Policy	62%	38%	54%	37%	9%	54%	32%	14%
	State/Observation	68%	32%	52%	32%	16%	49%	37%	14%
DRL libraries/ frameworks	Installation	43%	57%	85%	12%	3%	78%	15%	7%
	Dependency	54%	46%	78%	14%	8%	77%	18%	5%
	API usage	38%	62%	86%	9%	5%	78%	14%	8%
	Documentation	51%	49%	69%	23%	8%	71%	20%	9%
	Bugs inside framework	48%	52%	75%	14%	11%	70%	15%	15%
	Best fitted library	51%	49%	70%	25%	5%	65%	25%	10%
Parallel processing	GPU usage	54%	46%	72%	20%	8%	69%	20%	11%
	Multi-threading	45%	55%	68%	24%	8%	65%	18%	17%
	Multi-processing	49%	51%	57%	31%	12%	51%	23%	26%
	Distributed systems	43%	57%	58%	28%	14%	55%	14%	31%
DL Issues	Model	48%	52%	71%	23%	6%	66%	26%	8%
	Data preprocessing	48%	52%	57%	25%	18%	63%	25%	12%
	DL framework	46%	54%	66%	25%	9%	64%	27%	9%

Highlighted cells indicate values that are important for the analysis and need additional attention

and statistically significant ($P - value < 0.05$). Hence, more severe challenges necessitate more effort from developers.

Finding 4: Survey respondents highlighted *Reward*, *Environment*, *Hyperparameters*, and *Design problem* as the most severe and intricate-to-address challenges in DRL development.

RQ3: Are DRL application development challenges language- and/or framework-specific?

We extract information about the programming languages used to develop DRL applications from the collected SO posts. It should be mentioned that the posts have been collected without any distinction on the programming language and frameworks used. Figure 6 presents the proportion of posts using *Python* programming language for the different identified categories of challenges.

As can be seen, *Python* is by far the dominant programming language for all categories of challenges. However, the proportion of posts mentioning other programming languages and containing *DRL issues* is non-negligible (i.e., 20.2%). This high ratio is mostly attributable to *Java* (5.2% of all posts), *C++* (4.7% of all posts), and *R* (4.7% of all posts) programming


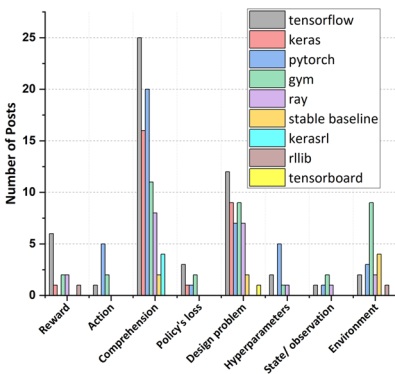


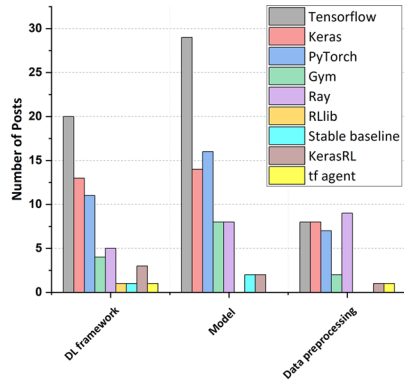
Fig. 6 Programming languages mentioned in the posts belong to various categories of challenges

languages. It is also noteworthy that investigating the relationship between used programming languages and the challenges (subcategories) within each category reveals that *Python* stands out as the predominant programming language across all DRL development challenges. Based on these results, we conclude that there is no relationship between DRL development challenges and used programming languages. This finding is in accordance with prior research (Morovati et al. 2023; Humatova et al. 2020) which reported that *Python* is the most popular programming language for ML-enabled applications. These results about programming languages used in DRL applications development are also supported by our validation survey where all participants mentioned *Python* as the programming language they use for developing DRL applications. Besides, 20% of participants reported *C/C++*, and 12% mentioned other programming languages in addition to *Python* (e.g., *C#* 4%, and *Java* 3%).


We also examine the mentioned frameworks and libraries in posts related to different challenge categories. Figure 7(a) shows the number of times that each library/framework has been mentioned in posts belonging to various subcategories of the *DRL issues* category. With an exception for the *environment* and *state/observation* subcategories, *Keras*, *Tensorflow*, and *PyTorch* are the most used libraries/frameworks by SO users, and they are the most popular libraries/frameworks for developing ML and DL (Morovati et al. 2024). Considering that there are several libraries/frameworks specifically designed to ease DRL application development (e.g. *KerasRL*, *RLlib*, etc), Fig. 7(a) exposes that SO users usually prefer to use popular ML libraries/frameworks which can be leveraged to implement DRL applications as a subdomain of ML. On the other hand, *gym* (Brockman et al. 2016) is the most popular library/framework in the *environment* subcategory which is reasonable as it is the most popular library for implementing various RL environments and provides a standard benchmark containing a large number of well-known RL environments (Panerati et al. 2021). This observation can be related to the inherent nature of SO posts, which may not necessarily provide details about the libraries/frameworks employed by the user posting questions. It is noteworthy that 46% of 885 examined SO posts lack any reference to the used DRL-related libraries/frameworks.



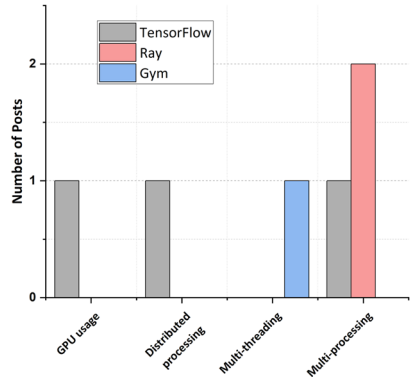
(a) DRL issues



(b) DL issues



(c) DRL libraries/frameworks



(d) Parallel processing & multi-threading

Fig. 7 Common libraries/frameworks used for developing DRL applications

As an example, post #56312962 serves as an illustrative case where no information has been mentioned regarding the utilized libraries-frameworks to implement the DRL application.

Figure 7(b) presents the libraries/frameworks mentioned in the posts classified as *DL issues*. Results show that *TensorFlow*, *Keras*, and *PyTorch* are the most popular libraries/frameworks in the *DL framework* and *model* subcategories. Given that the challenges within the *DL issues* category pertain to the DL parts of DRL applications, it is not surprising to see *TensorFlow*, *Keras*, and *PyTorch* are frequently mentioned since they are the most used libraries in the development of DL-enabled applications (Morovati et al. 2023; Humbatova et al. 2020). Moreover, the generality of *Ray* in the *Data preprocessing* subcategory, compared to other libraries/frameworks can be attributed to the fact that *Ray* encompasses not only DRL-related libraries (e.g., *RLLib* (Liang et al. 2018)) but also various other libraries for a wide range of ML-related tasks at the same time, including scalable datasets, model training, and hyperparameter tuning (Pumperla et al. 2023) which may ease developing DRL applications.

About the most frequently referenced libraries/frameworks in the *DRL libraries/frameworks* category, as illustrated in Fig. 7(c), *gym* received the largest number of questions, especially on the topics of *API usage*, *installation*, and *dependency* challenges. Although

gym is the most popular library for implementing RL/DRL environments (Panerati et al. 2021), comparing Fig. 7(a) and (b) may indicate that *gym* has less matured documentation and tutorials, in comparison to *Keras*, *Tensorflow*, and *PyTorch*. It is also worth mentioning that *Keras*, *Tensorflow*, and *PyTorch* are general-purpose ML-related libraries/frameworks which are more developed, compared to *gym* which is implemented specifically for RL development (Brockman et al. 2016).

Regarding parallelization and multi-threading, as can be seen in Fig. 7(d), the majority of issues are reported against *TensorFlow*, particularly regarding *GPU usage* and *distributed processing*. It can be related to the fact that *Tensorflow* is considered the most popular ML framework (Openja et al. 2022). On the other hand, most of the *multi-processing* challenges relate to the *ray* framework, which could be attributed to the fact that *ray* supports multiprocessing and RL at the same time (Moritz et al. 2018).

The respondents of our survey corroborated these findings; with 86% of them reporting *PyTorch* as their preferred framework, followed by *TensorFlow* (50%), and *Keras* (39%). Some participants also mentioned *KerasRL* (6%) and *JAX* (6%). We note that *PyTorch* was cited more than *Tensorflow* and *Keras* by participants of our validation survey compared to SO posts. This can be explained by the fact that SO posts do not necessarily include information about the frameworks used by users asking questions. It is also worth mentioning that we collected SO posts over a period of 13 years (from 2009 to 2022), while *PyTorch* was introduced only in 2016 (in comparison to *Tensorflow* and *Keras* released in 2015).

Finding 5: *Python* is by far the dominant programming language used for DRL development. While *TensorFlow*, *Keras*, and *PyTorch* are frequently mentioned in posts reporting challenges faced by DRL developers, *gym* is the most challenging libraries to install (see *installation*, and *dependency* challenges) and use (see *API usage* challenges). Our results also show that *ray* is the most challenging library when dealing with *multi-processing* and *data preprocessing* operations.

4 Discussion

Based on the findings of our study, in this section, we discuss the state of DRL application development and highlight some research avenues for researchers and practitioners.

Through this study, we gained a thorough understanding of frequently asked questions regarding DRL development to enable the community to explore potential approaches for mitigating these challenges, minimizing errors, and enhancing the reliability of DRL applications. Based on our provided taxonomy, one can see that some challenges faced in the development of DRL applications are common to all types of DL applications. For example, managing dependencies when using DL libraries/frameworks is a prevalent challenge in DL applications. However, dependency management can be more complex in DRL, because of the need for synchronization among a larger number of libraries in the development of DRL applications (e.g., aligning the Python version with the DRL libraries/frameworks and the library that manages RL environment). Similarly to what was suggested by Huang et al. (2022) to tackle dependency management challenges in DL applications, DRL researchers can provide a dependency knowledge graph for DRL libraries/frameworks to mitigate this challenge.

Regarding the provided taxonomy, it is worth mentioning that all of the categories and subcategories of challenges directly relate to the DRL application development. However, some of the challenges may be observed in other ML-related applications. For instance, all of the challenges belonging to *DRL libraries/frameworks* (e.g. *API usage*, *Installation*, *Dependency*, etc.) have been faced by all developers who use ML/DL libraries/frameworks. But it should be also taken into consideration that all of the investigated SO posts in this study have been achieved after a comprehensive filtering process making sure all of the extracted SO posts are about challenges in DRL application development. On the other hand, *Reward*, *Environment*, *Action*, *State/Observation*, and *Policy* challenges are Specific to DRL applications. It should be also taken into consideration that what makes this taxonomy valuable in the DRL community is the fact that their frequency, importance, and severity would be different in DRL application development compared to other ML-related applications. For example, 14.8% of challenges in DRL application development are related to *API usage*, whereas it is only 5.3% in DL-related applications (Humbatova et al. 2020).

Finding 2 revealed that 27.3% of DRL development challenges categorized as *comprehension* are related to the lack of sound understanding of basic DRL concepts. In other words, 58.4% of posts belonging to *DRL issues* category (DRL-specific category) are about *comprehension* challenge. This finding highlights the need for documentation and tutorials to help DRL developers who are not experts in DRL, in the development of DRL applications. A roadmap for the development of DRL applications would also help developers navigate through the implementation of DRL applications with fewer misunderstandings of DRL concepts. The need for such material is emphasized by a post⁴ asking questions about the difference between RL and DRL. By providing a roadmap that systematically expands DRL developers' understanding, developers will be supported in overcoming the most common challenge in DRL application development. An illustration of such guidance is the work conducted by Garg et al. (2019) on creating a roadmap for DL development. The need for good documentation and guidance is also emphasized by the survey participants who mentioned that '*although there are a number of tutorials to start working on DRL, a few issues are shared between many of them*'. The participants also noted that many of the DRL tutorial documents cover only a specific domain of DRL. Participants also lamented the poor usability of DRL-related tutorials, claiming that they often contain a lot of unnecessary materials.

Leveraging our Findings 1 and 2, researchers can develop debugging tools to help developers identify the issues early on during DRL application development. Debugging tools can significantly reduce DRL development and maintenance costs. For instance, considering the limited documentation available for most of the DRL libraries/frameworks in comparison with DL libraries for example (e.g., *TensorFlow*), a helpful approach would be proposing techniques and tools to assist DRL developers when using different DRL APIs. This could help mitigate DRL API issues. An example of such techniques focusing on the challenges of software API usage, is the work by Xie et al. (2022) which proposes an approach to automatically extract the API parameter constraints of DL libraries/frameworks.

Finding 2 of this study regarding challenges associated with installation and dependency management of DRL libraries/frameworks is aligned with previous studies on dependency management in software development (Cao et al. 2022) in general, as well as in DL applications (Han et al. 2020). Considering the complex nature of DRL application development (due to the communication of several libraries), challenges regarding libraries/frameworks installation and their dependency management become more intricate, compared to other types of DL applications. This highlights the need for tools (e.g., package manager) to support

⁴ <https://stackoverflow.com/questions/37973108>

dependency management. As an example, researchers can provide a tool (such as Maven⁵ for Java) that automates the identification and installation of DRL libraries/frameworks that are best suited for a specific Operating System (OS) and a specific version of *Python*. Additionally, such tools can assist DRL developers in synchronizing the installed DRL-related libraries/frameworks when they need to update some of them.

Our results in Finding 2 also stress the need for supporting tools and documentation, for parallel and distributed DRL application developments. Questions related to this topic took a long time before receiving an accepted answer on Stack Overflow. DRL experts could consider developing pre-configured packages to support parallel/distributed DRL application developments. As an example of the same task in the ML-related development, Openja et al. (2022) examined ML application deployment practices on Docker and reported that a significant number of ML developers use Docker to manage dependencies, environment, and the execution of ML applications.

5 Related Works

We now report and discuss the related literature.

5.1 SO Posts Analysis

Beyer et al. (2020) investigated the automatic classification of SO posts. They manually labeled 1000 posts, and identified 7 categories of questions: 1) API changes, 2) API usage, 3) Conceptual, 4) Discrepancy, 5) Learning, 6) Errors, and 7) Review. Leveraging the labeled dataset, they developed two approaches for the automatic classification of SO posts. In the first approach, they used the labeled dataset to extract some regular expression patterns and used these patterns to predict the category of other posts; achieving a performance of 0.91 for both precision and recall. In the second approach, they trained Random Forest and Support Vector Machine (SVM) classifiers using the labeled dataset. The best results were obtained using the Random Forest classifier, i.e., a precision of 0.88 and a recall of 0.87.

Alshangiti et al. (2019) investigated SO posts related to ML development. They used a tag-based snowball sampling approach to extract SO posts related to ML, starting with the ‘machine-learning’ tag. Their results revealed that a higher number of ML-related posts remain without any accepted answer (61%), in comparison with general domain questions (48%). They also reported that ML-related questions need 10 times longer to receive an answer, compared to general domain questions. Next, they compared the ratio of expert users in ML and web development (the most popular domain of programming in SO), showing that the number of ML experts is significantly less than that in web development. Afterward, they reviewed the most challenging ML development phases revealing that data preprocessing and manipulation, and model deployment and environment setup are the two most error-prone phases.

Bangash et al. (2019) conducted an empirical study of SO posts related to ML. They used the ‘machine_learning’ tag to extract 28,010 posts published between 2008 and 2018. Next, they used Latent Dirichlet allocation (LDA) (Jelodar et al. 2019) to categorize the extracted posts into 44 detailed topics. Then, they showed that *code errors*, *Algorithms*, and *Labeling* are the most discussed topics in ML. They then combined the 44 topics identified using LDA into 4 main groups including frameworks, implementation, sub-domain (RE), and algorithms.

⁵ <https://maven.apache.org/>

They reported that nearly 51% of all ML-related SO posts belong to the implementation group. Afterward, two of the researchers manually examined 230 sampled posts and reported that most of the questions stem from the fact that ML novice developers try to use ML in their software systems. They extracted information about the number of questions with accepted answers and concluded that ML-related questions are harder to answer than general domain SO questions. They also observed that only 65.6% of ML-related SO posts have appropriate tags; which might suggest that many users are not knowledgeable enough to assign the proper ML-related tags to their posts.

Hamidi et al. (2021) examined the challenges that developers may face in the development of ML systems, based on their discussion in SO. They studied 43,950 ML-related SO posts submitted between 2008 and 2020. First, they showed that *Python* is the most popular programming language for ML development, and *C#* and *C/C++* are the least popular programming languages for ML development. Then, they reported that model building and model evaluation are the two most challenging steps in ML development while model monitoring is the least questioned phase. They also report that questions regarding model requirements, data collection/processing, and model-building steps receive less accepted answers than others. This may stem from the fact that questions about these steps are more difficult to answer or the lack of active knowledgeable developers on SO to answer questions related to these steps.

Although these previous works investigated SO questions related to ML/DL application development, to the best of our knowledge, none of them examined the challenges of DRL development specifically.

5.2 RL and DRL Quality Assurance

In this section, we report on studies about the quality assurance of RL and DRL applications.

Zhang et al. (2021) proposed strategies to help DL and DRL developers detect and resolve quality issues in their applications. Nikanjam et al. (2021) proposed a methodology for automatically detecting faults in DL applications, using graph transformations. development of DL models. These studies primarily targeted issues occurring in the training program of DL models. The issues considered in these aforementioned study fall within the '*DL issues*' category of our provided taxonomy, which constitutes only a small portion of DRL issues.

Nikanjam et al. (2022) also investigated challenges categorized as *DRL issues* in our proposed taxonomy. They examined questions/discussions about four popular DRL frameworks (including *gym* (Brockman et al. 2016), *Tensorforce* (Kuhnle et al. 2017), *Dopamine* (Castro et al. 2018), and *Keras-rl* (Plappert 2016) on GitHub and SO) and extracted 329 SO posts about DRL. They categorized these posts into six groups: basic concepts, without acknowledgment, implementation issues, answered by the owner, relative questions, and others. They reported that 'without acknowledgment' and 'implementation' questions are the most common DRL-related questions in SO, accounting for 32% and 27%, respectively. They also showed that in 2% of their studied SO posts, the answer has been posted by the questioner. They report that DRL-related SO posts take an average and median time of 2.07 days and 13 hours, respectively, before receiving an accepted answer. This period is longer than the time taken by DL-related SO posts to receive an accepted answer; which is 5 hours on average. This finding implies that DRL-related questions might be more difficult to answer than DL-related questions. Nikanjam et al. (2022) also proposed a taxonomy of faults in DRL models. with 11 different issue types. Our study differs from what Nikanjam et al. (2022) carried out in several aspects. Firstly, they concentrated on four specific libraries/frameworks

designed for DRL development, ignoring other SO posts containing source codes using other Python-based libraries/frameworks or not mentioning any script. For instance, a number of SO posts inquire about DRL concepts (e.g. #52838439) without mentioning any scripts, a scope not covered by Nikanjam et al. (2022). Moreover, their research was just on the DRL model, whereas our study delves into the challenges developers may face throughout the development of entire DRL applications, without any limitation to a specific section of DRL applications. Last but not least, Nikanjam et al. (2022) examined SO posts reporting program faults during the development of DRL applications. It is imperative to distinguish between challenges and program faults and note that challenges do not necessarily equate to program faults. A fault denotes a defect or error leading to a discrepancy between the expected and achieved results or observed behavior (Morovati et al. 2023). Software development challenges, in contrast, encompass any difficulty or complexity encountered in completing a development task. These challenges may arise from various factors, including technical complexities, resource constraints, or lack of expertise. Nikanjam et al. (2022) also investigated challenges categorized as ‘DRL issues’ in our proposed taxonomy. They examined questions/discussions about four popular DRL frameworks (including *gym* (Brockman et al. 2016), *Tensorforce* (Kuhle et al. 2017), *Dopamine* (Castro et al. 2018), and *Keras-rl* (Plappert 2016) on GitHub and SO) and extracted 329 SO posts about DRL. Since Nikanjam et al. (2022) focused on software faults, they excluded 305 SO posts from their dataset, generating their taxonomy based on 24 such posts that included faults explicitly. Notably, all SO posts in their dataset are included in our study. Furthermore, our research incorporates SO posts that do not pertain to program faults but instead pose queries related to the comprehension of DRL concepts. For example, the comprehension subcategory includes 253 SO posts, none of which reports program faults. They categorized these posts into six groups: basic concepts, without acknowledgment, implementation issues, answered by the owner, relative questions, and others. They reported that ‘without acknowledgment’ and ‘implementation’ questions are the most common DRL-related questions in SO, accounting for 32% and 27%, respectively. They also showed that in 2% of their studied SO posts, the answer has been posted by the questioner. They report that DRL-related SO posts take an average and median time of 2.07 days and 13 hours, respectively, before receiving an accepted answer. This time period is longer than the time taken by DL-related SO posts to receive an accepted answer; which is 5 hours on average. This finding implies that DRL-related questions might be more difficult to answer than DL-related questions. Nikanjam et al. (2022) also proposed a taxonomy of faults in DRL models. The most significant difference between their study and ours is their focus on the DRL model only and the fact that they only collected data about faults mentioned in SO; disregarding all the other types of questions.

Yahmed et al. (2023) conducted a study on the challenges of deploying DRL systems, based on the questions that developers ask on SO. In the first step, they extracted 357 SO posts related to the deployment of DRL systems. Next, they categorized collected SO posts into 4 categories with respect to their deployment platform, including ‘server/cloud’, ‘mobile/embedded system’, ‘browser’, and ‘game engine’. Their results showed that the number of SO posts regarding deployment has grown over the 7 studied years. They manually examined the extracted SO posts and identified 31 challenges related to DRL deployment. They grouped these challenges into 11 categories (and proposed a taxonomy): general questions, deployment infrastructure, data preprocessing, RL environment, communication, agent load/save, performance, environment rendering, agent export, request handling, and continuous learning. The proposed taxonomy has been evaluated via a survey with practitioners. Their results show that DRL developers struggle the most with deployment infrastructures and RL environments. Their results also show that communication-related challenges (procedure,

connection loss, configuration of remote setting, and model convergence) are the most difficult challenges to address, in terms of the time to an accepted answer. The main difference between the work of Yahmed et al. and our study is the focus of the study. Yahmed et al. examined DRL deployment challenges while we examined challenges faced by DRL developers during the application development phase, i.e., prior to deployment. In other words, they ignored all DRL development steps before deployment and focused only on the deployment phase of the DRL applications, which occurs after complete implementation of them. At the opposite extreme, we investigated the SO posts regarding the whole pipeline of DRL applications development. Besides, comparing the number of studied SO posts in these two studies (357 vs 927 SO posts) demonstrates that Yahmed et al filtered out the whole dataset of SO posts on DRL development to achieve posts specifically talking about the deployment step of the DRL application.

6 Threats to Validity

We now discuss threats to the validity of our study.

Construction Validity. Our methodology and labeling process can be a potential validity threat. We have thoroughly described our process and the tags used to collect the posts. As no previous taxonomy on this subject exists, we used an open coding approach with multiple rounds and cross-checking to ensure continuous improvement and consistency of the labeling. We further validate our results via a survey with 65 DRL practitioners.

Internal Validity. As users do not necessarily provide suitable tags for their questions, our search might have missed some DRL challenges. For example, post [37973108](#) is related to DRL, but it does not have any specific tag mentioning DRL. Nonetheless, tag usage was necessary for a consistent methodology and we believe that the number of posts gathered and analyzed (927) is sizable enough to provide a good representation of the challenges faced by DRL developers. Moreover, we used a snowballing approach to expand our basic set of DRL-related tags used to extract DRL-related SO posts, similar to previous studies (Ayman et al. 2019). In addition to extracting posts based on the DRL-related tags, we used a set of keywords to extract DRL-related SO posts without any DRL-related tags to address this issue, as other researchers followed a similar methodology (Peruma et al. 2022). Another source of threat to the validity of this study arises from the potential overlap between users posing questions in SO posts and the developers associated with the DRL-related GitHub repositories that are used for our survey. To address this concern, we provided a detailed description for each category and subcategory of the taxonomy within the survey, intentionally avoiding any specific information references to any SO post as an example of challenges. Despite these precautions, the prospect of overlap between these two groups of DRL developers remains a possibility. Although utilizing the duration to receive an accepted answer has been employed in several studies (Alshangiti et al. 2019; Decan et al. 2019) as a metric to gauge the difficulty level of SO posts, it may present a potential threat to the internal validity of this research. Reproducibility of issues in ML-enabled systems poses significant challenges and sometimes needs more time compared to traditional software systems (Shah et al. 2024). In other words, addressing certain questions in ML-enabled systems (such as those related to parallel processing) requires specific configurations or environments, so the question may be quite simple, but it requires a long time to receive an accepted answer. To address this concern, we also considered the number of posts, which has been used as an indicator to measure the difficulty level of SO posts in a few studies (Bangash et al. 2019;

Hamidi et al. 2021). However, as discussed in Section 4, this metric yielded similar results to the methodology we employed (i.e., the time to receive an accepted answer).

External Validity. While there might exist other DRL challenges that practitioners are facing, we conducted our study using SO which is the largest technical Q&A platform in the software development community. Moreover, all the challenges identified in our provided taxonomy have been validated by our survey participants. Also, respondents of the survey did not report any other challenge that is not included in our provided taxonomy; which is a good result, with respect to completeness. Another external factor that could pose a threat to the validity of our results is the possibility that users who raised SO posts may mostly fall into the category of less experienced users. To mitigate this concern, we examined the top 100 DRL-related GitHub repositories to extract the challenges that GitHub developers mention in their development process. However, it appears that experienced developers generally refrain from detailing their challenges in GitHub commit messages or issues. It is noteworthy that our provided taxonomy may represent high-level categories of challenges, covering almost all aspects of DRL application development. Therefore, refining the taxonomy to offer more detailed categorization could be explored as a potential avenue for future research in this study.

Reliability Validity. We described our methodology in detail and provided a replication package (Morovati et al. 2023) to allow others to replicate our results and expand our study.

7 Conclusion and Future Works

In this study, we conducted a large-scale empirical study of 927 DRL-related posts extracted from SO. We examined all posts manually to identify the challenges that developers face when developing DRL applications. We found that *Python* is by far the most popular programming language and *TensorFlow*, *Keras*, *PyTorch*, and *OpenAI Gym* are the most frequently used libraries/frameworks for developing DRL applications. We categorized DRL development challenges into five groups including *DRL issues*, *parallel processing & multi-threading*, *DRL libraries/frameworks*, *DL issues*, and *general programming*. An analysis of the received response by the investigated SO posts shows that DRL comprehension, DRL libraries/frameworks API usage, and designing a problem using DRL algorithms are the most challenging parts of DRL applications' development. Furthermore, parallel processing/multi-threading and DRL libraries/frameworks challenges required a longer time to receive an accepted answer. We proposed a taxonomy of challenges and validated it using a survey of 65 DRL developers. The developers confirmed the frequency, severity, and required effort to address identified challenges. We hope that the reported results in this paper will stimulate the development of DRL quality assurance tools and guide the research community toward solving the identified challenges.

Acknowledgements This work was supported by: Fonds de Recherche du Québec (FRQ), the Canadian Institute for Advanced Research (CIFAR) as well as the DEEL project CRDPJ 537462-18 funded by the National Science and Engineering Research Council of Canada (NSERC) and the Consortium for Research and Innovation in Aerospace in Québec (CRIAQ), together with its industrial partners Thales Canada inc, Bell Textron Canada Limited, CAE inc and Bombardier inc.

Funding The authors have no financial or proprietary interests in any material discussed in this article. All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.

Data Availability Statement The dataset generated during the current study is available in the replication package, which is accessible via Morovati et al. (2023).

Declarations

Conflict of interest The authors have no conflict of interest.

The authors have no competing interests to declare that are relevant to the content of this article.

References

- Aldhaen M (2020) Interview versus questionnaire from the perspective of cbe members. *Int J Educ, Learn Develop* 8(2):21–41
- Alshangiti M, Sapkota H, Murukannaiah PK, Liu X, Yu Q (2019) Why is developing machine learning applications challenging? a study on Stack Overflow posts. In: 2019 acm/ieee international symposium on empirical software engineering and measurement (esem). IEEE, pp 1–11
- Aradi S (2020) Survey of deep reinforcement learning for motion planning of autonomous vehicles. *IEEE Trans Intell Transp Syst* 23(2):740–759
- Arpteg A, Brinne B, Crnkovic-Friis L, Bosch J (2018) Software engineering challenges of deep learning. In: 2018 44th euromicro conference on software engineering and advanced applications (SEAA). IEEE, pp 50–59
- Arulkumaran K, Deisenroth MP, Brundage M, Bharath AA (2017) Deep reinforcement learning: A brief survey. *IEEE Signal Process Mag* 34(6):26–38
- Ayman A, Aziz A, Alipour A, Laszka A (2019) Smart contract development in practice: trends, issues, and discussions on Stack Overflow. [arXiv:1905.08833](https://arxiv.org/abs/1905.08833)
- Bagherzadeh M, Kahani N, Briand L (2021) Reinforcement learning for test case prioritization. *IEEE Trans Software Eng* 48(8):2836–2856
- Bangash AA, Sahar H, Chowdhury S, Wong AW, Hindle A, Ali K (2019) What do developers know about machine learning: a study of ml discussions on stackoverflow. In: 2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR). IEEE, pp 260–264
- Beyer S, Macho C, Di Penta M, Pinzger M (2020) What kind of questions do developers ask on Stack Overflow? a comparison of automated approaches to classify posts into question categories. *Empir Softw Eng* 25:2258–2301
- Brockman G, Cheung V, Pettersson L, Schneider J, Schulman J, Tang J, Zaremba W (2016) Openai gym
- Brockman G, Cheung V, Pettersson L, Schneider J, Schulman J, Tang J, Zaremba W (2016) Openai gym. [arXiv:1606.01540](https://arxiv.org/abs/1606.01540)
- Bruer JJ, Tropp JA, Cevher V, Becker SR (2015) Designing statistical estimators that balance sample size, risk, and computational cost. *IEEE J Selc Topics Signal Process* 9(4):612–624
- Cao Y, Chen L, Ma W, Li Y, Zhou Y, Wang L (2022) Towards better dependency management: A first look at dependency smells in python projects. *IEEE Transactions on Software Engineering*
- Castro PS, Moitra S, Gelada C, Kumar S, Bellemare MG (2018) Dopamine: A research framework for deep reinforcement learning. [arXiv:1812.06110](https://arxiv.org/abs/1812.06110)
- Chen Z, Cao Y, Liu Y, Wang H, Xie T, Liu X (2020) A comprehensive study on challenges in deploying deep learning based software. In: Proceedings of the 28th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering, ESEC/FSE 2020. Association for Computing Machinery, New York, NY, USA, p 750–762. <https://doi.org/10.1145/3368089.3409759>
- de Souza Nascimento E, Ahmed I, Oliveira E, Palheta MP, Steinmacher I, Conte T (2019) Understanding development process of machine learning systems: Challenges and solutions. In: 2019 acm/ieee international symposium on empirical software engineering and measurement (esem). IEEE, pp. 1–6
- Decan A, Mens T, Grosjean P (2019) An empirical comparison of dependency network evolution in seven software packaging ecosystems. *Empir Softw Eng* 24:381–416
- DeCarlo M (2018) Scientific inquiry in social work. *Open Social Work Education*
- Du W, Ding S (2021) A survey on multi-agent deep reinforcement learning: from the perspective of challenges and applications. *Artif Intell Rev* 54:3215–3238
- Dulac-Arnold G, Levine N, Mankowitz DJ, Li J, Paduraru C, Gowal S, Hester T (2021) Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Mach Learn* 110(9):2419–2468
- Frost J (2019) Introduction to Statistics: An Intuitive Guide for Analyzing Data and (2019)

- Garg D, Goel P, Kandaswamy G, Ganatra A, Kotecha K (2019) A roadmap to deep learning: a state-of-the-art step towards machine learning. In: *Advanced Informatics for Computing Research: Second International Conference, ICAICR 2018, Shimla, India, July 14–15, 2018, Revised Selected Papers, Part I 2*. Springer, pp 160–170
- Gibbons JD (1993) *Nonparametric measures of association*. 91. Sage
- GoogleSheet (2020) Google: How to use google sheets. <https://support.google.com/docs/answer/6000292?hl=en&co=GENIE.Platform%3DDesktop>. Accessed: 2023-02-01
- Hamidi A, Antoniol G, Khomh F, Di Penta M, Hamidi M (2021) Towards understanding developers' machine-learning challenges: A multi-language study on Stack Overflow. In: *2021 IEEE 21st International Working Conference on Source Code Analysis and Manipulation (SCAM)*. IEEE, pp 58–69
- Han J, Deng S, Lo D, Zhi C, Yin J, Xia X (2020) An empirical study of the dependency networks of deep learning libraries. In: *2020 IEEE International conference on software maintenance and evolution (ICSME)*. IEEE, pp 868–878
- Haque MU, Iwaya LH, Babar MA (2020) Challenges in docker development: A large-scale study using Stack Overflow. In: *Proceedings of the 14th ACM/IEEE International symposium on empirical software engineering and measurement (ESEM)*, pp 1–11
- Heo K, Lee W, Pashakhanloo P, Naik M (2018) Effective program debloating via reinforcement learning. In: *Proceedings of the 2018 ACM SIGSAC Conference on computer and communications security*, pp 380–394
- He J, Xin Z, Xu B, Zhang T, Kim K, Yang Z, Thung F, Irsan I, Lo D (2023) Representation learning for stack overflow posts: How far are we? [arXiv:2303.06853](https://arxiv.org/abs/2303.06853)
- Hove SE, Anda B (2005) Experiences from conducting semi-structured interviews in empirical software engineering research. In: *11th IEEE International Software Metrics Symposium (METRICS'05)*. IEEE, pp 10–pp
- Huang K, Chen B, Wu S, Cao J, Ma L, Peng X (2022) Demystifying dependency bugs in deep learning stack. [arXiv:2207.10347](https://arxiv.org/abs/2207.10347)
- Humbatova N, Jahangirova G, Bavota G, Riccio V, Stocco A, Tonella P (2020) Taxonomy of real faults in deep learning systems. In: *Proceedings of the ACM/IEEE 42nd International conference on software engineering*, pp 1110–1121
- Islam MJ, Nguyen G, Pan R, Rajan H (2019) A comprehensive study on deep learning bug characteristics. In: *Proceedings of the 2019 27th ACM Joint meeting on european software engineering conference and symposium on the foundations of software engineering*, pp 510–520
- Islam MJ, Pan R, Nguyen G, Rajan H (2020) Repairing deep neural networks: Fix patterns and challenges. In: *Proceedings of the ACM/IEEE 42nd International conference on software engineering, ICSE '20*. Association for Computing Machinery, New York, NY, USA, p 1135–1146. <https://doi.org/10.1145/3377811.3380378>
- Jelodar H, Wang Y, Yuan C, Feng X, Jiang X, Li Y, Zhao L (2019) Latent dirichlet allocation (lda) and topic modeling: models, applications, a survey. *Multimedia Tools Appl* 78:15169–15211
- Kiran BR, Sobh I, Talpaert V, Mannion P, Al Sallab AA, Yogamani S, Pérez P (2021) Deep reinforcement learning for autonomous driving: A survey. *IEEE Trans Intell Transp Syst* 23(6):4909–4926
- Kuhnle A, Schaarschmidt M, Fricke K (2017) Tensorforce: a tensorflow library for applied reinforcement learning. Web page. <https://github.com/tensorforce/tensorforce>
- Li Y (2017) Deep reinforcement learning: An overview. [arXiv:1701.07274](https://arxiv.org/abs/1701.07274)
- Liang E, Liaw R, Nishihara R, Moritz P, Fox R, Goldberg K, Gonzalez J, Jordan M, Stoica I (2018) Rllib: Abstractions for distributed reinforcement learning. In: *International conference on machine learning*. PMLR, pp 3053–3062
- Lorenz U (2022) *Reinforcement Learning From Scratch: Understanding Current Approaches - with Examples in Java and Greenfoot*. Springer International Publishing. <https://books.google.ca/books?id=bV2YEAAAQBAJ>
- Lune H, Berg BL (2017) *Qualitative research methods for the social sciences*. Pearson
- Lwakatere LE, Raj A, Bosch J, Olsson HH, Crnkovic I (2019) A taxonomy of software engineering challenges for machine learning systems: An empirical investigation. In: *Agile processes in software engineering and extreme programming: 20th international conference, XP 2019, Montréal, QC, Canada, May 21–25, 2019, Proceedings 20*. Springer International Publishing, pp 227–243
- McDonald N, Schoenebeck S, Forte A (2019) Reliability and inter-rater reliability in qualitative research: Norms and guidelines for csw and hci practice. *Proceedings of the ACM on human-computer interaction* 3(CSCW), 1–23
- Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G et al (2015) Human-level control through deep reinforcement learning. *Nat* 518(7540):529–533

- Mojica-Hanke A, Bayona A, Linares-Vásquez M, Herbold S, González FA (2023) What are the machine learning best practices reported by practitioners on stack exchange? [arXiv:2301.10516](https://arxiv.org/abs/2301.10516)
- Moritz P, Nishihara R, Wang S, Tumanov A, Liaw R, Liang E, Elibol M, Yang Z, Paul W, Jordan MI et al (2018) Ray: A distributed framework for emerging {AI} applications. In: 13th USENIX symposium on operating systems design and implementation (OSDI 18), pp. 561–577
- Morovati MM, Nikanjam A, Khomh F, Jiang ZM (2023) Bugs in machine learning-based systems: a faultload benchmark. *Empir Softw Eng* 28(3):62
- Morovati, M.M., Tambon, F., Taraghi, M., Nikanjam, A., Khomh, F.: Paper replication package. https://github.com/drlchallenges/drl_challenges. Accessed 2023-02-01
- Morovati MM, Nikanjam A, Tambon F, Khomh F, Jiang ZM (2024) Bug characterization in machine learning-based systems. *Empir Softw Eng* 29(1):14
- Nekkanti H, Reddy S (2016) Surveys in software engineering: a systematic literature review and interview study. Ph.D. thesis, MSc thesis, Blekinge Institute of Technology, Sweden
- Nguyen HA, Nguyen TT, Wilson G Jr, Nguyen AT, Kim M, Nguyen TN (2010) A graph-based approach to api usage adaptation. *ACM Sigplan Notices* 45(10):302–321
- Nguyen TT, Nguyen ND, Nahavandi S (2020) Deep reinforcement learning for multiagent systems: a review of challenges, solutions, and applications. *IEEE Transactions on Cybernetics* 50(9):3826–3839
- Nikanjam A, Braiek HB, Morovati MM, Khomh F (2021) Automatic fault detection for deep learning programs using graph transformations. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 31(1):1–27
- Nikanjam A, Morovati MM, Khomh F, Ben Braiek H (2022) Faults in deep reinforcement learning programs: a taxonomy and a detection approach. *Autom Softw Eng* 29(1):1–32
- Openja M, Majidi F, Khomh F, Chembakottu B, Li H (2022) Studying the practices of deploying machine learning projects on docker. In: Proceedings of the 26th International conference on evaluation and assessment in software engineering, pp 190–200
- Panerati J, Zheng H, Zhou S, Xu J, Prorok A, Schoellig AP (2021) Learning to fly—a gym environment with pybullet physics for reinforcement learning of multi-agent quadcopter control. In: 2021 IEEE/RSJ International conference on intelligent robots and systems (IROS). IEEE, pp 7512–7519
- Panzer M, Bender B (2022) Deep reinforcement learning in production systems: a systematic literature review. *Int J Prod Res* 60(13):4316–4341
- Peruma A, Simmons S, AlOmar EA, Newman CD, Mkaouer MW, Ouni A (2022) How do i refactor this? an empirical study on refactoring trends and topics in Stack Overflow. *Empir Softw Eng* 27(1):11
- Plappert M (2016) keras-rl. <https://github.com/keras-rl/keras-rl>
- Pumperla M, Oakes E, Liaw R (2023) Learning ray: flexible distributed python for machine learning. O'Reilly media. <https://books.google.ca/books?id=vKj0zgEACAAJ>
- Qualtrics (2023) Qualtrics official website. <https://www.qualtrics.com/>. Accessed 04 Jan 2023
- Rao Q, Frtunik J (2018) Deep learning for self-driving cars: chances and challenges. In: Proceedings of the 1st international workshop on software engineering for AI in autonomous systems, pp 35–38
- Reda D, Tao T, van de Panne M (2020) Learning to locomote: understanding how environment design matters for deep reinforcement learning. In: Proceedings of the 13th ACM SIGGRAPH conference on motion, interaction and games, pp 1–10
- Schaarschmidt M, Kuhnle A, Ellis B, Fricke K, Gessert F, Yoneki E (2018) LIFT: reinforcement learning in computer systems by learning from demonstrations. <http://arxiv.org/abs/1808.07903>
- Schoop E, Huang F, Hartmann B (2021) Umlaut: debugging deep learning programs using program structure and model behavior. In: Proceedings of the 2021 CHI conference on human factors in computing systems, pp 1–16
- Shah MB, Rahman MM, Khomh F (2024) Towards enhancing the reproducibility of deep learning bugs: an empirical study. [arXiv:2401.03069](https://arxiv.org/abs/2401.03069)
- StackExchange (2022) Stack exchange data dump. <https://archive.org/details/stackexchange>. Accessed 02 Jan 2023
- Tambon F, Majdinasab V, Nikanjam A, Khomh F, Antoniol G (2023) Mutation testing of deep reinforcement learning based on real faults. In: 2023 IEEE Conference on software testing, verification and validation (ICST). IEEE, pp 188–198
- Vijayaraghavan G, Kaner C (2003) Bug taxonomies: use them to generate better tests. *Star East* 2003:1–40
- Wan Y, Zhao Z, Yang M, Xu G, Ying H, Wu J, Yu PS (2018) Improving automatic source code summarization via deep reinforcement learning. In: Proceedings of the 33rd ACM/IEEE international conference on automated software engineering, pp 397–407
- Xiao Y, Liu J, Wu J, Ansari N (2021) Leveraging deep reinforcement learning for traffic engineering: a survey. *IEEE Commun Surv Tutor* 23(4):2064–2097

- Xie D, Li Y, Kim M, Pham HV, Tan L, Zhang X, Godfrey MW (2022) Docter: documentation-guided fuzzing for testing deep learning api functions. In: Proceedings of the 31st ACM SIGSOFT international symposium on software testing and analysis, pp 176–188
- Yahmed AH, Abbassi AA, Nikanjam A, Li H, Khomh F (2023) Deploying deep reinforcement learning systems: a taxonomy of challenges. [arXiv:2308.12438](https://arxiv.org/abs/2308.12438)
- Yu C, Liu J, Nemati S, Yin G (2021) Reinforcement learning in healthcare: a survey. *ACM Comput Surv (CSUR)* 55(1):1–36
- Zahedi M, Rajapakse RN, Babar MA (2020) Mining questions asked about continuous software engineering: a case study of Stack Overflow. In: Proceedings of the evaluation and assessment in software engineering, pp 41–50
- Zhang S, Ding Z, Dong H (2020) *Deep Reinforcement Learning: Fundamentals*. Springer, Research and Applications
- Zhang T, Gao C, Ma L, Lyu M, Kim M (2019) An empirical study of common challenges in developing deep learning applications. In: 2019 IEEE 30th International Symposium on Software Reliability Engineering (ISSRE). IEEE, pp 104–115
- Zhang X, Zhai J, Ma S, Shen C (2021) Autotrainer: an automatic dnn training problem detection and repair system. In: 2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE). IEEE, pp 359–371
- Zhu W, Zhang H, Hassan AE, Godfrey MW (2022) An empirical study of question discussions on stack overflow. *Empir Softw Eng* 27(6):1–25
- Zolfagharian A, Abdellatif M, Briand L, Bagherzadeh M et al (2022) Search-based testing approach for deep reinforcement learning agents. [arXiv:2206.07813](https://arxiv.org/abs/2206.07813)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.