

InteNSE'23: The 1st International Workshop on Interpretability and Robustness in Neural Software Engineering

May 14th, 2023

[Co-located with ICSE'23, Melbourne, Australia]

Call for Papers

Important Dates: AoE (UTC-12h)

Important Links:

Submission Deadline: January 13th

Notification to Authors: February 23rd

Camera-Ready Deadline: March 17th

Submission: <https://intense-icse2023.hotcrp.com>

Website: <https://intense23.github.io>

Twitter: <https://twitter.com/intenseWorkshop>

Why InteNSE?

Code is the most recent modality of interest in the Machine Learning (ML) domain. Learning-based techniques have been shown to improve and revolutionize software engineering and analysis tasks, including code completion and synthesis, code captioning and documentation, code search and clone detection, recovering variable names and types, and software testing, verification, and debugging. Recent studies have even shown the great potential of neural models of code for decomposing the natural language reasoning problem into programmatic steps and solving it by executing the generated programs. However, most of these techniques consider ML as closed-box, i.e., only considering the final performance of the developed models to evaluate their effectiveness. Without a systematic and rigorous approach to interpreting neural models concerning code modality, one cannot validate their actual performance, generalizability, decision certainty, if/how they are learning what they are supposed to do, and why they make particular decisions.

InteNSE is an interdisciplinary workshop for research at the intersection of ML and Software Engineering (SE). The InteNSE workshop aims to promote the required **awareness** and **knowledge** necessary for both research communities to interpret neural models of code and design robust neural models for software analysis and engineering. The InteNSE program will combine academia and industry in a quest for well-founded practical solutions. We invite international experts in academia and industry with both ML and SE backgrounds to (1) discuss their research, (2) evaluate prior ML4Code research, and (3) identify the roadmap for the research domain.

Workshop Components

InteNSE consists of the following four components to enable the **knowledge** and **awareness** concerning the interpretability and robustness of ML4Code.

Keynotes. We have two awesome keynote speakers, *Dr. Michael Pradel (University of Stuttgart)* and *Dr. Kla Tantithamthavorn (Monash University)*, who will discuss their insights on ML4Code robustness and interpretability. The details of the talks will appear in the program.

Papers. We accept three types of papers:

1. *Full Workshop Papers:* 4-6 pages papers related to the workshop's topics of interest that are (1) presenting novel research ideas and preliminary results or (2) interpreting or assessing the robustness of existing ML4Code models.

2. *Posters*: 2 pages papers related to the workshop’s topics of interest that are (1) proposing a statement of vision or position or (2) presenting novel ideas without any preliminary results. For the second category of papers, authors should clearly indicate their evaluation plan.
3. *Journal First Papers*: The papers accepted to appear at IEEE Software XAI4SE special issue are welcome to present their papers at the workshop. The authors of journal first papers at other top software engineering venues, including TSE, TOSEM, EMSE, and JSS are also welcome to represent their published work at InteNSE. Journal first papers do not appear in the workshop proceeding.

Papers with negative results are also welcome. Examples of negative results are: interpretability techniques that did not reveal meaningful knowledge or adversarial attacks against which the model was robust.

Hackathon. InteNSE offers a unique Hackathon-style experience to in-person participants interested in exploring the interpretability and robustness of neural models of code. Workshop organizers will set up the tool set and framework required for the hackathon, including ready-to-use Jupyter notebooks. The participants will (1) assess *how* and *why* the closed-box neural models of code make specific decisions and (2) evaluate the robustness of such models against adversarial inputs.

Participation in the Hackathon and accessing the tool is only available for **in-person** registered attendees. Participants are required to have their own laptops to use the tools. A separate call for participation in Hackathon will be announced later.

Expert Panel. The expert panel will comprise people from academia and industry with a research background in neural software analysis. The main theme of the discussion would be the state of the ML4Code in academia and industry, the importance of interpretability and explainability of such models to move the research frontiers, and potential risks concerning the robustness of neural models of code. The list of panelists and details about the discussions will appear in the program.

At the end of the panel, we will identify open challenges and a promising roadmap in advancing the role of interpretability and robustness in neural software engineering research.

Topics of Interest

We welcome research related to different aspects of software/code, including *code completion and synthesis, program analysis, software testing and debugging, formal verification and proof synthesis, neurosymbolic programming, and prompting*. Specifically, we are interested in both **theoretical** and **empirical** papers that explore one or more of the following perspectives related to ML4Code:

- *Interpretability*
 - Why interpret neural models of code?
 - How to interpret neural models of code?
 - What are the limitations of ML4Code models?
 - How to leverage interpretability for improving neural models of code?
 - Do the neural models perceive the code the same way humans do?
- *Robustness*
 - What are the consequences of brittle neural models of code?
 - How to assess the robustness of neural models of code?
 - How to quantify the robustness of neural models of code?
 - How to develop/train robust neural models of code?
 - What are the impacts of robustness on other requirements (e.g., generalization)?
- *Applications* Explaining and root causing the following concerns in ML4Code

- False positives and false negatives that result in performance degradation
- Biases in the model or data that questions the fairness of the model
- Model uncertainty, out-of-distribution, and out-of-sample detection
- Security-critical applications (e.g., addressing adversarial examples and data reconstruction).
- Other applications for neural software engineering/analysis.

If your submission does not match this category, but you believe it would be of interest to the workshop, don't hesitate to get in touch with the general chair Reyhaneh Jabbarvand (reyhaneh@illinois.edu).

Submission Guidelines

Submissions must conform to the ICSE 2023 Proceedings Formatting Guidelines (The title in 24pt font and full text in 10pt type. Latex users must use `\documentclass[10pt,conference]{IEEEtran}` without including the `compsoc` or `compsocconf` options). The page limit is strict and **purchases of additional pages in the proceedings is not allowed**. The official publication date is the date the proceedings are made available in the ACM or IEEE Digital Libraries. This date may be up to two weeks prior to the first day of ICSE 2023. The official publication date affects the deadline for any patent filings related to published work. InteNSE will employ a double-blind review process. No submission may reveal its authors' identities. The authors must make every effort to honor the double-blind review process. In particular, the authors' names must be omitted from the submission, and references to their prior work should be in the third person. The Workshop will follow the ACM SIGSOFT rules on Conflicts of Interest and Confidentiality of Submissions. https://www.sigsoft.org/policies/pgmcommittee.html#con_int

Participation

Participation in the workshop is open to anyone interested in the topic. This includes graduate and undergraduate students as well as faculty and researchers at academic institutions. We also welcome participants from the industry. Workshop attendees are expected to have basic familiarity with general ML concepts.

Organization

Organizing Committee

Reyhaneh Jabbarvand (University of Illinois at Urbana-Champaign)
 Saeid Tizpaz-Niari (University of Texas at El Paso)
 Earl T. Barr (University College London)
 Satish Chandra (Google)

Program Committee

TBA