Dear Editor and Reviewers,

We would like to thank you for your valuable feedback, which has enabled us to improve the original version of the manuscript. We have made changes to reflect the proposed suggestions. Below is a list of the received feedback, along with a description of how we have addressed it. In addition, the entire manuscript has been proofread and edited where necessary to improve the presentation.

Sincerely yours,

Giulia Murtas, Veselka Boeva and Elena Tsiporkova

Neuro-LENS: a neuro-symbolic framework integrating incomplete background knowledge and deep learning- Changes applied to address reviewers' comments

Reviewer 1

1. "The computation aspects of the method are unclear: it seems to rely on max likelihood, which can be misleading when not enough data are available"

   We agree with the reviewer that the method is sensitive, as most data-driven methods, to the variability and size of the training dataset. However, an important advantage of our method is its ability to deal with imbalanced datasets, i.e., allowing to bootstrap an initial model from a limited set of available labels, which can be easily upgraded at a later stage when more data becomes available.
   This discussion has been added to the end of Section 4.

2. "Single score computation is limited to binary classes'

   We have expanded upon the explanation in Section 4.1 by describing how scores can be calculated for a multi-class task (see equation (7)).

3. "Second scenario (SCANIA) is difficult to understand and not clear how symbolic approach is applied"

   We agree with the reviewer that, due to its inherent complexity, the second use case (truck failure prediction) was difficult to comprehend. We have made a substantial effort to improve its description. Namely, the entire text presenting the second

scenario has been significantly reworked to improve its clarity and explain how the symbolic approach is applied and integrated with the neural component.

4. “The paper need spolishing (typos, grammar errors, missing references)”

   The entire manuscript has been proofread and edited where necessary to improve the presentation.

5. “Notation to be revised (X used both for features and entities- page 10)”

   The text has been revised to solve the issue of coinciding notations.

6. “Page 17, formula 9: possibility and necessity are used but they both are necessity conditions”

   The formulas were correct, but the accompanying text was not. We fixed the text so that it matched the formulas (formula 9, page 14 in the revised version).

Reviewer 2

1. “Organization of the experimental section”
   We modified the organization of the section as suggested, separating the two use-cases and presenting the problem statement, datasets, method, results, and discussion for each use-case.

2. “Consider dropping truck benchmark altogether”

   The truck failure prediction use-case is essential as it demonstrates that the method can be applied to different types of data (both images and tabular/time series data) and is used to showcase two of the three integration strategies proposed in the paper. We believe that the use-case adds value to the paper and have thus decided to retain it. We think that in the revised version, where the structure has been reorganized and the text has been made clearer, it will be easier to see the added value of presenting this use-case.

3. “Quote performance metric of some baseline methods on the considered benchmarks”

A “Baseline comparison” paragraph was added at the end of the results sections of each use case. Tables 3 and 7 show the results of the baseline comparisons.

4. “Shorten argumentation for modularity on p. 8; the advantages are widely accepted and not controversial”

   The argumentation was shortened to one sentence only.

5. “Consider using "Scene classification" instead of "image scene classification"”

   All occurrences of "image scene classification" were replaced with "scene classification".

6. “The approach presented in the original paper is ~~inserted~~ embedded within a framework”

   The term “embedded” was used instead of “inserted”.

7. “The sentences in bullets on p. 3 should end with periods”

   We added periods at the end of the sentences.

8. “Page 8: $(\forall x \in X)(F(x) \not= \emptyset)$ Unnecessary parentheses”

   The unnecessary parentheses were removed in the specified formulas and in similar ones.

9. “Missing references on page 11”

   We solved the issue with the references to all sections and fixed all missing references.

10. “A deep learning model following one of the integration strategies depicted in ~~Fig. 1~~ Fig.2.”

    Corrected to reference the right figure.

11. “Name or enumerate strategies and use same names throughout the paper”

The naming of the strategies is kept consistent throughout the paper in the revised version.

12. “The section "Neuro-LENS: Neuro-symbolic integration" feels strongly repetitive with respect to the previous section "Neural component. Shorten this part of the paper or move some of the discussed thoughts to Discussion section”

    Shortened the "Neural component" section to avoid repetition with the other sections; reworked the "Neuro-LENS: Neuro-symbolic integration" section to improve clarity and avoid internal repetitions within the section.

13. “Page 13: define tasks clearly before referring to them; task statement for trucks only clear after seeing Table 1”

    A clearer definition of the tasks was added and their description was rewritten.

14. “Mentioning 5 decision classes is inconsistent with Table 1”

    Table 1 refers to the task of detecting abandoned objects. This is now made clearer in the text. The five classes involved in the failure prediction use cases are also explained more clearly.

15. “Figure 4 is quite cryptic”

    We deleted the figure altogether since it was not adding much useful information. A short description of the specification of the trucks was added in the text instead in section 5.2.1.

16. “Discard all considerations concerning data anonymization, not relevant for the approach”

    We discarded the considerations regarding data anonymization and reworked the section describing the SCANIA dataset to make it clearer and more concise.

17. “Page 14 and following section: rewrite text to logically follow the pipeline in Fig. 6“

The text has been written to align with the pipeline shown in Fig. 6 and made more concise to avoid repetitions.

18. “Shorten description of the pipeline- to the point and without repetitions”

    The pipeline description has been shortened and improved by removing all repetitions.

19. “Fig.9 does not seem to convey any useful information (all values are True)”

    The figure was substituted with Table 6, in order to be symmetrical to the first use case of scene classification, where Table 1 shows the inverse and superinverse images of the two considered classes.

20. “Conclusion section should be shortened to convey the essence of results, without repetitions of what already stated earlier”

    The conclusion was shortened to summarize only the results of the experiments.