

## HABILITATION THESIS REVIEWER'S REPORT

### Masaryk University

#### Applicant

RNDr. Vít Musil, Ph.D.

#### Habilitation thesis

Bridging Continuous Learning and Discrete Optimization in Artificial Intelligence

#### Reviewer

Prof. Nils Jansen

#### Reviewer's home unit, institution

Ruhr-University Bochum

### Review of the habilitation thesis

#### *Summary of the thesis*

This thesis collects eleven previously published articles by the candidate, Vít Musil. The works sit at the intersection of discrete, combinatorial decision-making and continuous, gradient-based learning. The core idea is to connect the two with explicit surrogate gradients instead of ad hoc relaxations, so that discrete solvers and controllers stay exact at execution time while remaining trainable end-to-end. Two streams follow: differentiable programming for combinatorial solvers, and strategy synthesis for patrolling games and long-run stochastic control. The introduction is concise and substantive. It fixes notation and explains clearly why standard automatic differentiation fails on these problems.

#### *Assessment of the quality of the habilitation candidate*

I see it as good practice that the candidate acknowledges using AI tools in preparing the thesis. The candidate's trajectory also is very nice: a PhD in mathematical analysis (Charles University, 2018), postdocs at the Czech Academy of Sciences, the University of Florence, and Masaryk University, plus long-standing international collaborations. The candidate has repeatedly moved into new research environments and made himself useful there. The work is substantial and has been published at premier venues in the field: ICLR (twice, one spotlight), ICML (twice), CVPR (best-paper nomination), ECCV, UAI (twice), IJCAI (twice), and ICAPS. Declared contribution shares (16–33%) reflect substantial work within two coauthor groups, concentrated in method design, proofs, and mathematical framing.

#### *Assessment of the thesis document*

I like the document. It identifies one structural problem shared by both streams: combinatorial  $\arg \min/\arg \max$  operators and long-run strategy-evaluation functionals have a Jacobian that is zero almost everywhere. The exposition of finite-memory strategies in Chapter 3 is precise and self-contained. Generally, everything reads very well and is well structured.

One thing I would have liked to see discussed more: the two streams are framed as parallel instances of relaxing a discrete decision into a continuous one, but they are not quite parallel. The differentiable-programming stream trains against a fixed solver

whose feasible set does not change at deployment. The patrolling stream additionally has to deal with an adversary reacting to the deployed strategy. That is a modeling difference, not just a different application area, and the thesis could discuss this better.

#### *Assessment of the contributions*

Blackbox Backpropagation is the base result: a solver that only exposes a linear objective still gets an exact gradient of a well-defined, piecewise-affine surrogate, from one extra solver call on perturbed costs. CombOptNet extends this to learning the constraint set itself. It also has a nice, genuinely useful trick: reparameterizing half-space constraints so the optimization stays stable. LPGD then unifies Blackbox Backpropagation, implicit differentiation, and proximal/Frank-Wolfe updates under a single Lagrange–Moreau envelope. The applications are also very nice and somewhat novel/refreshing.

Regstar’s dynamic program is, to me, the strongest single result. It evaluates a randomized patrolling strategy’s protection value and its gradient by aggregating exponentially many attack paths, in polynomial time. That is what makes gradient ascent here trustworthy, not just something that happens to work on the benchmarks tried. The follow-up work formalizes “security holes”: the vulnerability created when switching between two individually optimal strategies. The memory-assignment paper turns conflicting attack gradients into a simple, robust rule for where to add controller memory. Of the two IJCAI papers, the first extends the objective language to hitting-time moments and edge frequencies. The second, on mean-payoff scheduling for maintenance, has what I think is the most creative algorithmic idea in this half of the thesis: sampling long executions from a low-memory randomized controller to find good deterministic cycles.

#### **Reviewer’s questions for the habilitation thesis defence** (number of questions up to the reviewer)

- How sensitive are the guarantees to misspecification of the model?
- Is there a natural bridge to the robust-MDP/robust-POMDP literature, where the uncertainty itself is part of the model?
- What are real-world applications of your work?
- What’s next :)?

#### **Conclusion**

The candidate’s profile and the thesis are of high quality, connecting a mathematically rigorous treatment of surrogate gradients to two genuinely different application domains as a coherent research program. I recommend the candidate for habilitation.

The habilitation thesis entitled *Bridging Continuous Learning and Discrete Optimization in Artificial Intelligence* by RNDr. Vít Musil, Ph.D. **fulfils** requirements expected of a habilitation thesis in the field of Informatics.

Date: 21.09.2026