We're failing to bring intelligent robots into the real world, here's why.

Popular science abstract from PhD thesis by Nicolai Anton Lynnerup: *Intelligent Robotic Learning Research – Bridging the Gap between Academia and Industry*,

Thinking of science as a broad term we think about collectively advancing our understanding of the world around us. We think collaboration across institutions towards a higher goal. A goal of maximizing our knowledge such that we can limit the fears of the unknowns. Sadly, competition has become a greater part of the academic world and neither universities nor publishers of scientific papers have mitigated the risks that competition inherently has in science. It has become a competition of which scientist (or institution) can publish the most papers and who can receive the most citations. This is a direct result of measuring research outcome by key performance indicators (KPIs), hence scoring research by numeric values. Such scoring as the H-index which supposedly should measure both author-level productivity and the scientific impact of the publications. If you want to be a scientist and do research, your work will – in today's standards – be measured by such obscure values. The result: the *publish or perish* aphorism, describes the pressure academic researchers face to either publish their academic work or find another career path. It sounds harsh and it is. It is well-known that competition is healthy in "controlled" doses, but the level of competition within science – and specifically within computational sciences such as artificial intelligence, machine learning, etc. – has led to poor reporting standards of the results generated. This means that we no longer live up to the requirements of a scientific scholarship, in which reproducibility is a *minimum necessary condition* for a scientific finding to be *believable* and *informative*. The ability of a researcher to duplicate the results of a prior study using the same materials as the original investigator is one of the main pillars of science.

Why is an industrial PhD addressing these academic issues? Because the issues reach far beyond the academic world. It impairs the credibility of research such that results is not trusted nor usable in industrial settings. Technological advances in industry are often directly derived from academic proof-of-concept research which is then impossible.

This thesis considers how the field of reinforcement learning – a learning paradigm for agents interacting with the real world – can be utilised in the manufacturing industry to develop new flexible robotic systems that can be easily reconfigured to manufacture new lines of products and personalise products for the end-user. We find that the current state of technological uptake by the industry is impaired by the poor reporting standards of academic findings and thus seek to find a common ground so researchers can stand on top of each other's shoulders and not on each other's feet. We present a set of guidelines and intuitions that combined with practical considerations form a complete development pipeline that can be used to ensure that all details of an academic scholarship (within machine learning) are well-documented. Furthermore, we show how to perform proper statistical evaluation of learning systems especially in cases where data is expensive – or even impossible – to obtain. We also provide the reader with guidelines for managing the expectations of being an AI researcher *in the wild*, dealing with all the hype that the technologies within artificial intelligence have sparked. Following our guide will significantly improve the quality of research within computational sciences such as artificial intelligence and will allow for fast continued progress in this rapidly evolving field.