Optical Alignment using Reinforcement Learning

Marko von der Leyen[1], Ramy Aboushelbaya[1], Ben Spiers[1,*], Henry Davies[2], Abigail James[1], Jordan Lee[1], Heath Martin[1], Robert Paddock[1], Robin Timmis[1], Christian Schroeder de Witt[3], Peter Norreys[1,4]

[1] *Department of Physics, University of Oxford, Oxford, UK,*
[2] *St. Peter's College, University of Oxford, Oxford, UK,*
[3] *Department of Engineering, University of Oxford, Oxford, UK,*
[4] *John Adams Institute, Oxford, UK,*
[*] *now at Rutherford Appleton Laboratory, STFC, Didcot, UK*

marko.vonderleyen@physics.ox.ac.uk

Optical alignment is omnipresent in high-power laser experiments – this holds for both diagnostics and driving beam lines. There are different problems that are associated with alignment. First, when using external laser facilities, a significant part of the allocated time may be used just for setting up the beam line and adjusting the alignment. Second, after perfect alignment has been reached, different factors may degrade the optimal state. These are, for example, unintentional human interaction (i.e. bumping a mirror), drift of the laser beam within the backend, tilting of mirrors after pumping to vacuum, etc. Even degrading optical coatings may change optimal alignment (where the figure of merit can be chosen, for example, to be the highest intensity of the pulse at the focal spot, or largest fringe spacing within interferometers).

Linear control systems may offer a potential route to having a universal alignment system. However, it has vast limitations when a system is not perfectly understood. For example, having an off-axis parabolic mirror as a final focusing optic (which is a common approach in laser-driven wakefield acceleration) is notoriously difficult to align; yet, it is mathematically a solvable problem. Once, scratches or other influences change the perfect conditions, this changes and the lack of characterisation means that there is no accurate description possible (and therefore linear control systems will fail to find the optimal state). In contrast, reinforcement learning (RL) algorithms do not need to 'understand' the physical description of a model to improve a chosen outcome. RL algorithms do not even need to have access to information such as what the action is they are taking at each learning step. Therefore, using transfer learning [1], optical systems can be controlled optimally without human interference. Furthermore, external influences may be countered automatically and in a live manner.

Here, we present results showing that RL agents are able to optimise alignment for systems and components such as MZIs and OAPs without knowledge of the physical description of the system. The system has been trained without simulation data but directly from the laboratory set-up showing its universal applicability.

Acknowledgments

This work has been supported through UKRI-STFC grants ST/X005518/1 and ST/V001655/1.

References

[1] A. Tirinzoni et al., arXiv:2007.00722 (2020)