

1 Appendix 1




Fig. 1 Each pair of plots represents one run that did not need human help. The lines show the metrics of success in the validation stages. The number of the run is displayed at the left of the plots. The horizontal dotted lines mark the maximum number of packages, while the vertical ones mark the best stage during validation.


2 *Hu-bot cooperation*

Fig. 2 Each row (made of four plots) represents one run that improved only in *efficiency* with human help. The line plots show the metrics of success in the validation stages throughout the learning processes with and without human help: note that naturally, the lines overlap until the stage when the human help starts. The boxplots compare the tests of the best controllers found by each experiment.




Fig. 3 The runs that did not improve with human help. Run 6 actually became a little worse in efficiency with human help. Explanations of the previous figure apply.