A. Appendix A

This section contains the larger versions of all the figures in the paper for enhanced clarity.

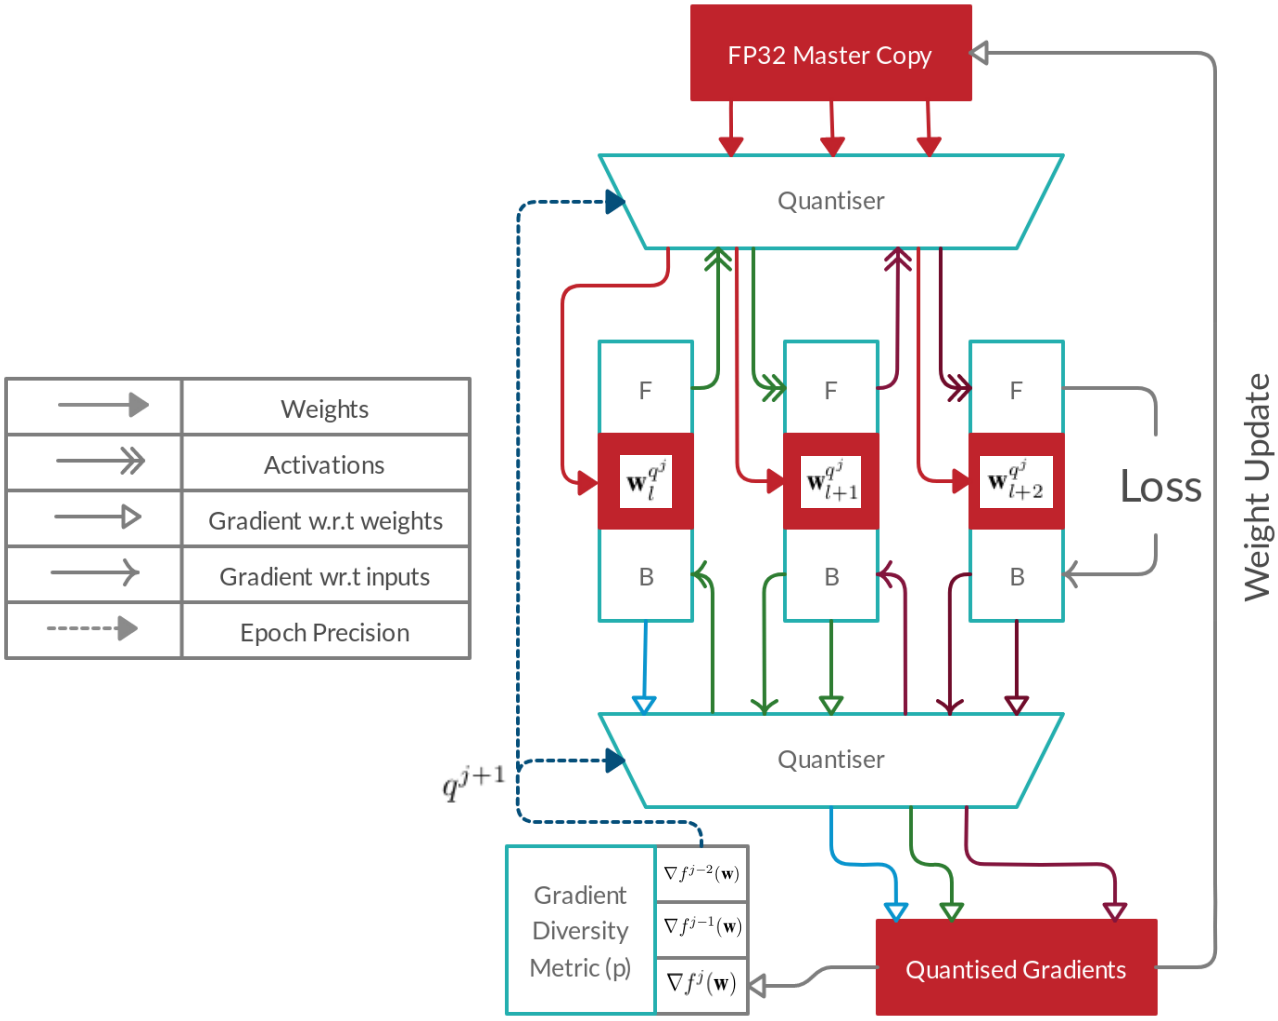

Figure 5. Precision-switching training scheme of MuPPET.

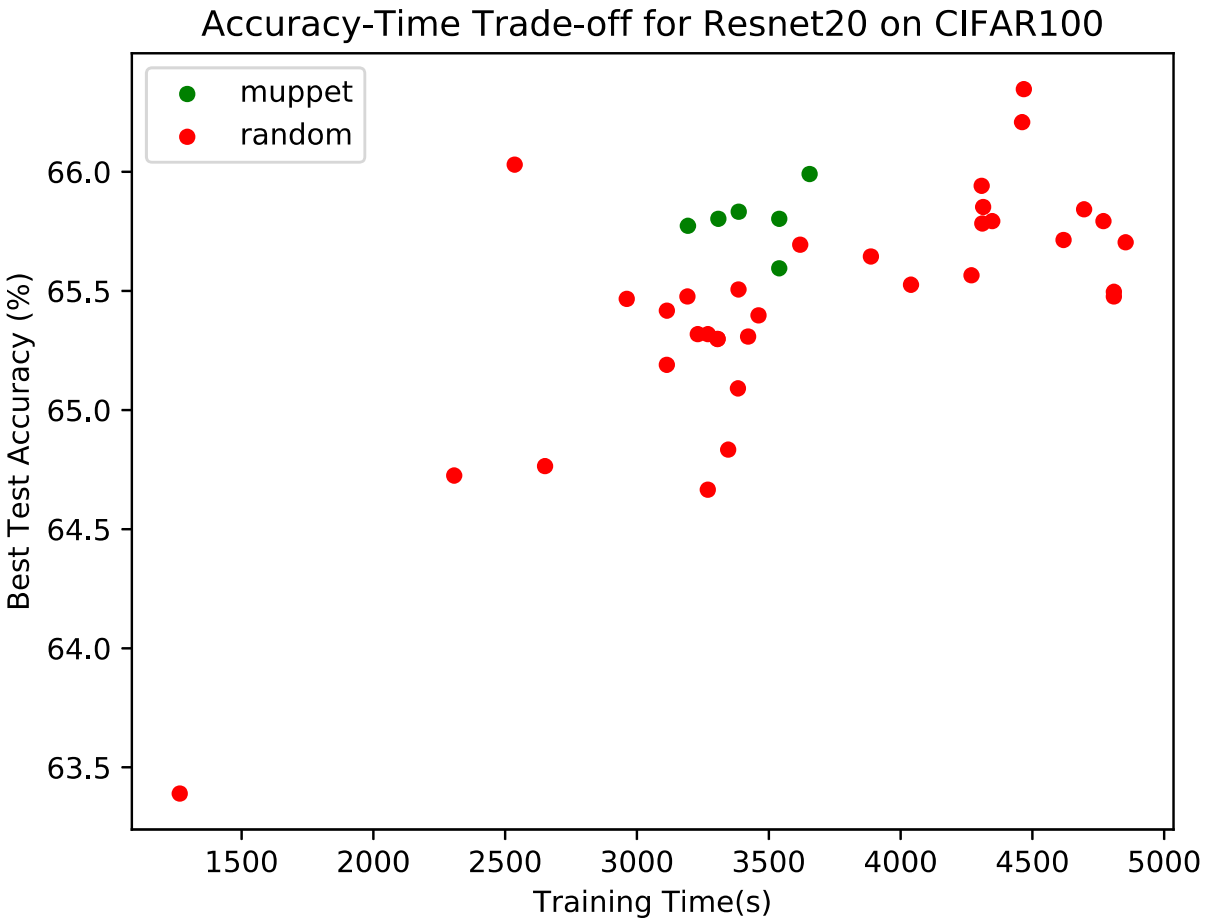

Figure 6. Accuracy vs time tradeoff for ResNet20 MuPPET runs on CIFAR100 - Final Test Accuracy

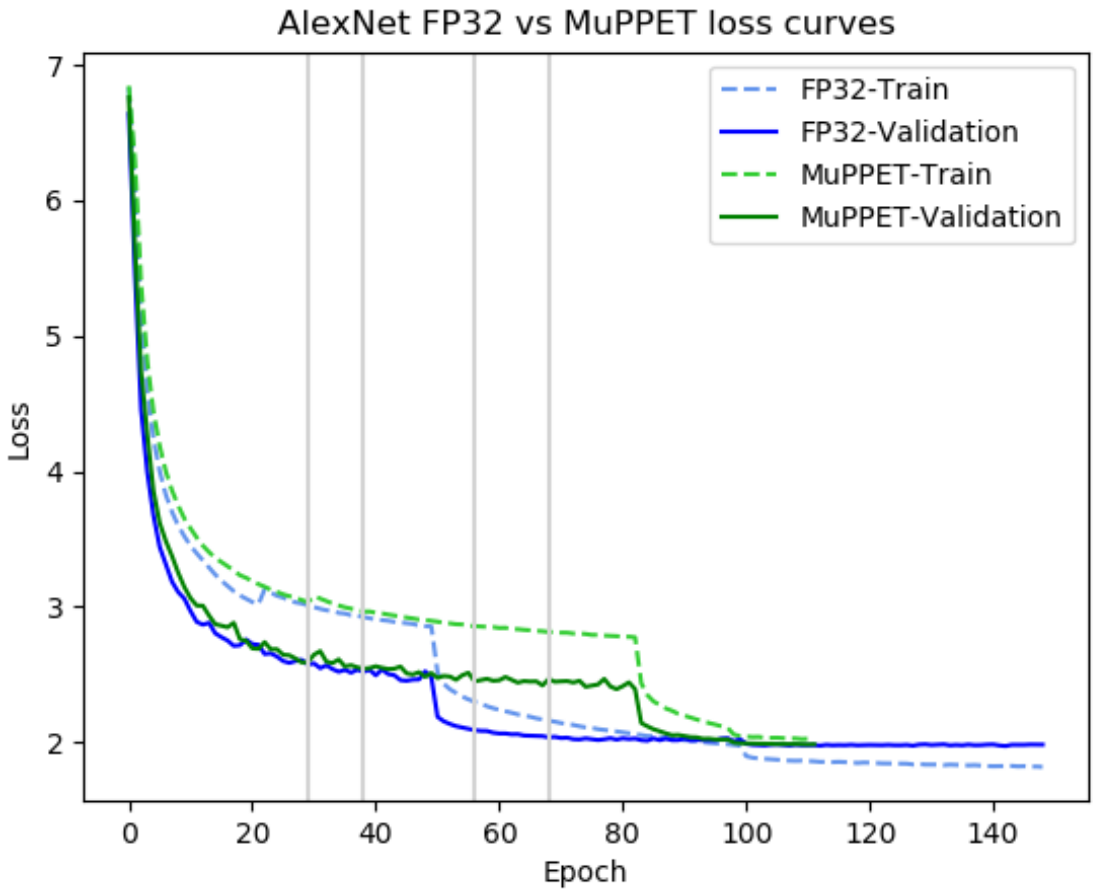

Figure 7. AlexNet training and validation loss values for FP32 and MuPPET on ImageNet.

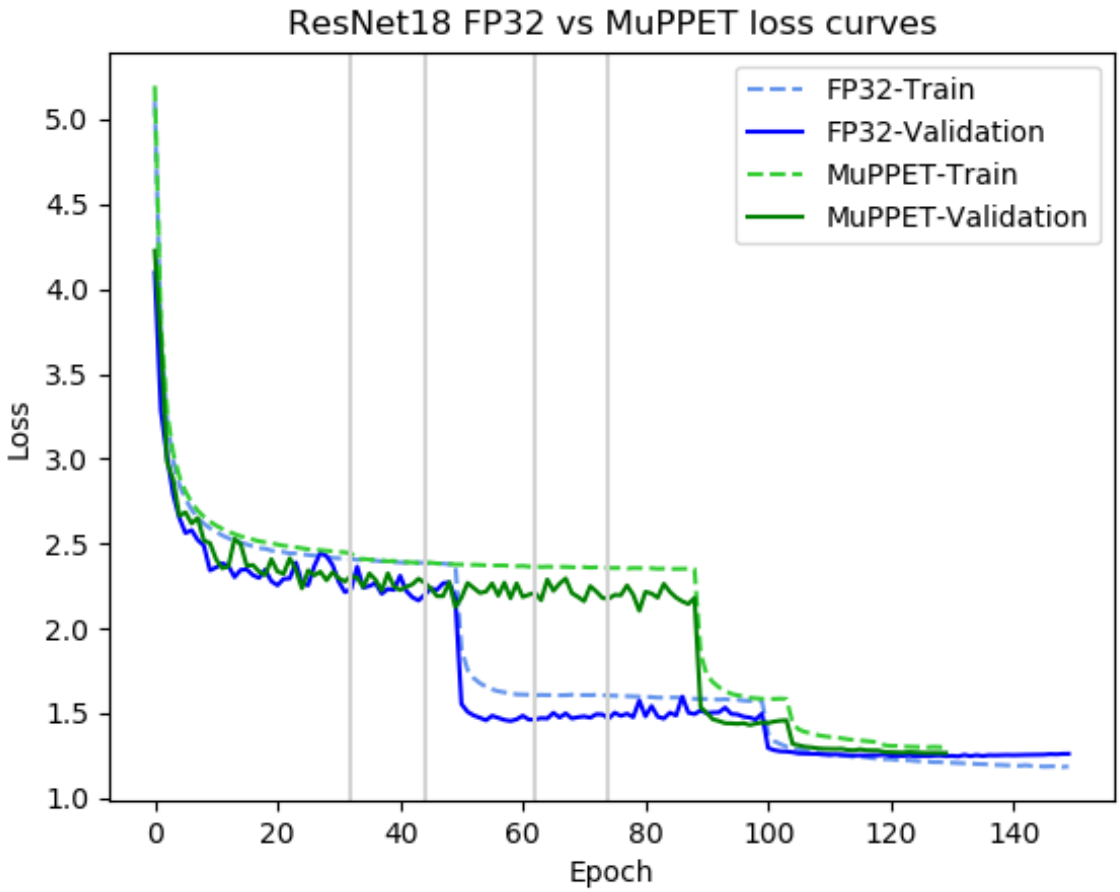

Figure 8. ResNet18 training and validation loss values for FP32 and MuPPET on ImageNet.

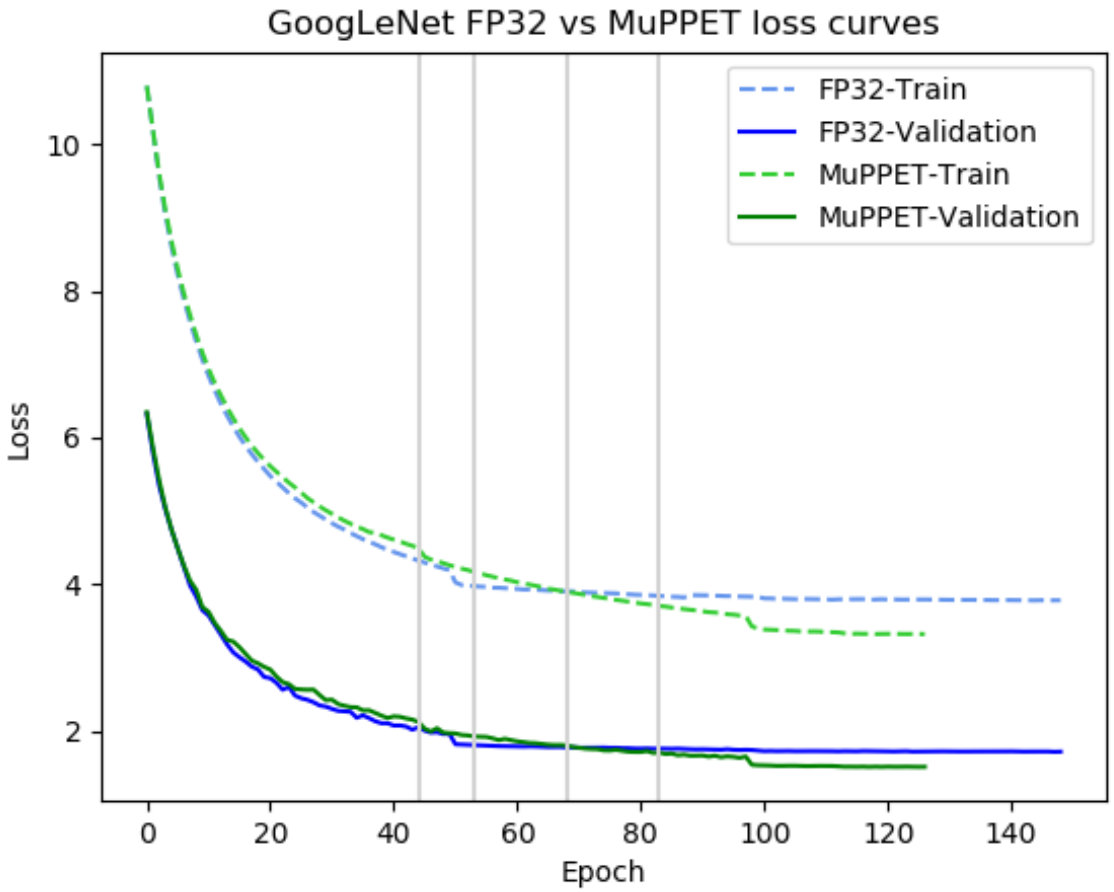

Figure 9. GoogLeNet training and validation loss values for FP32 and MuPPET on ImageNet.

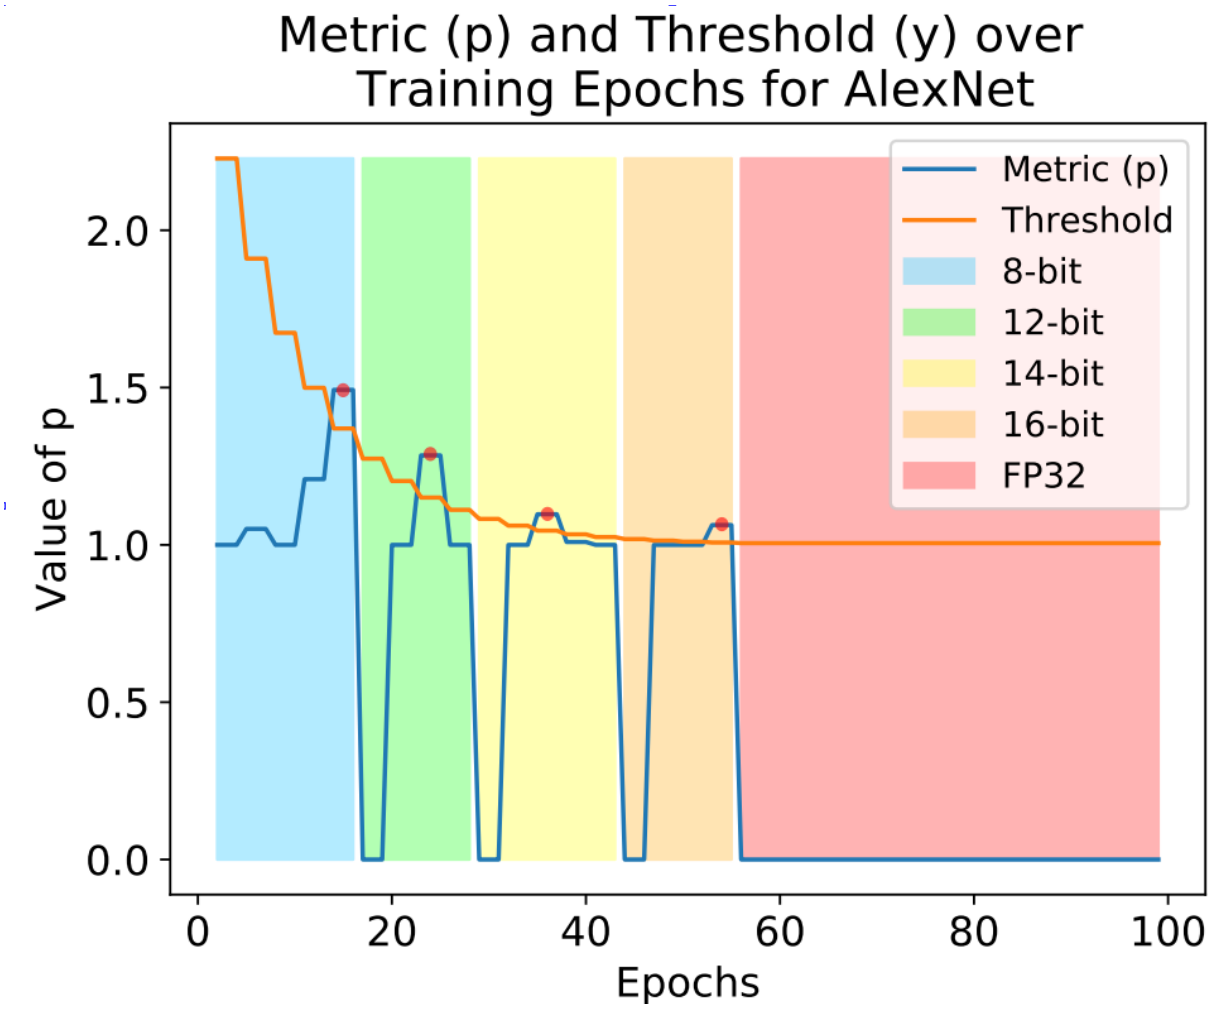

Figure 10. AlexNet CIFAR10 - Demonstration of the generalisability of p over networks, datasets and epochs.

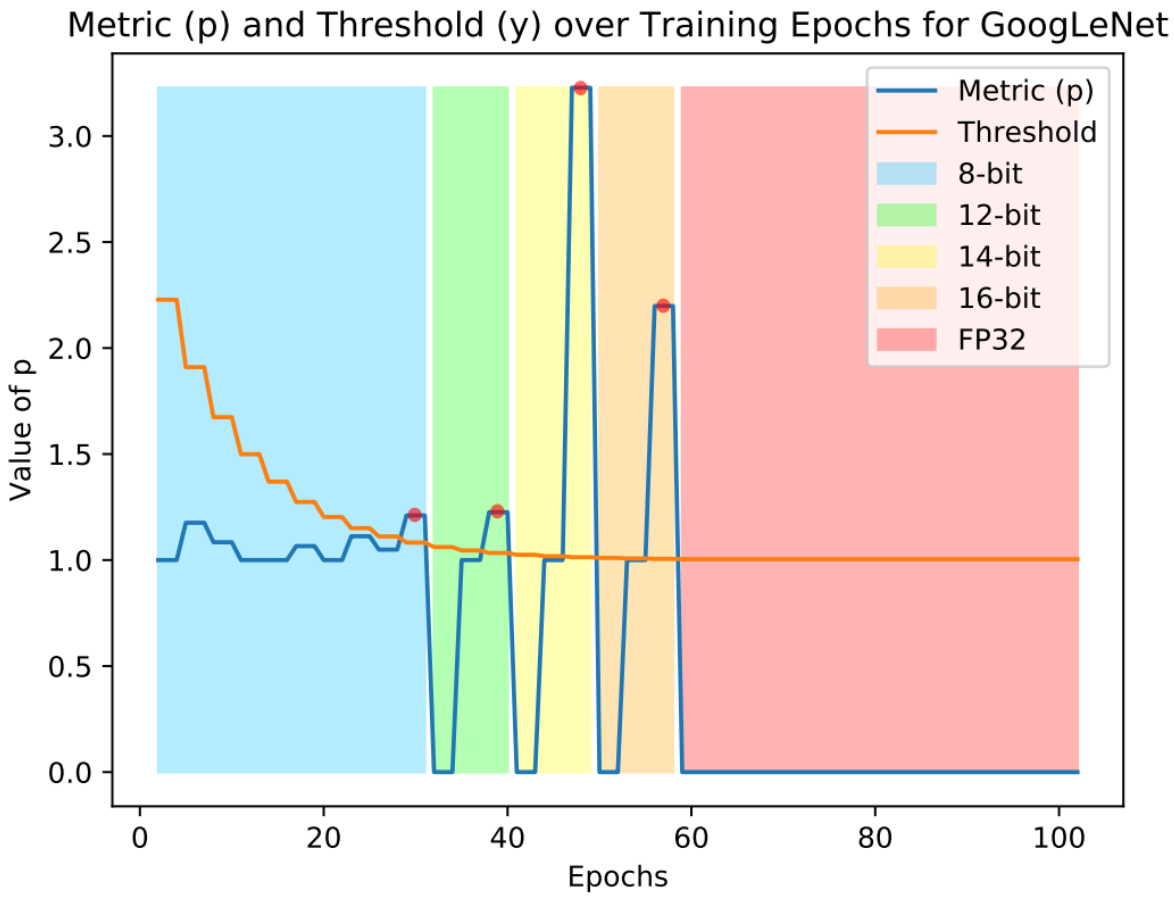

Figure 11. GoogLeNet CIFAR10 - Demonstration of the generalisability of p over networks, datasets and epochs.

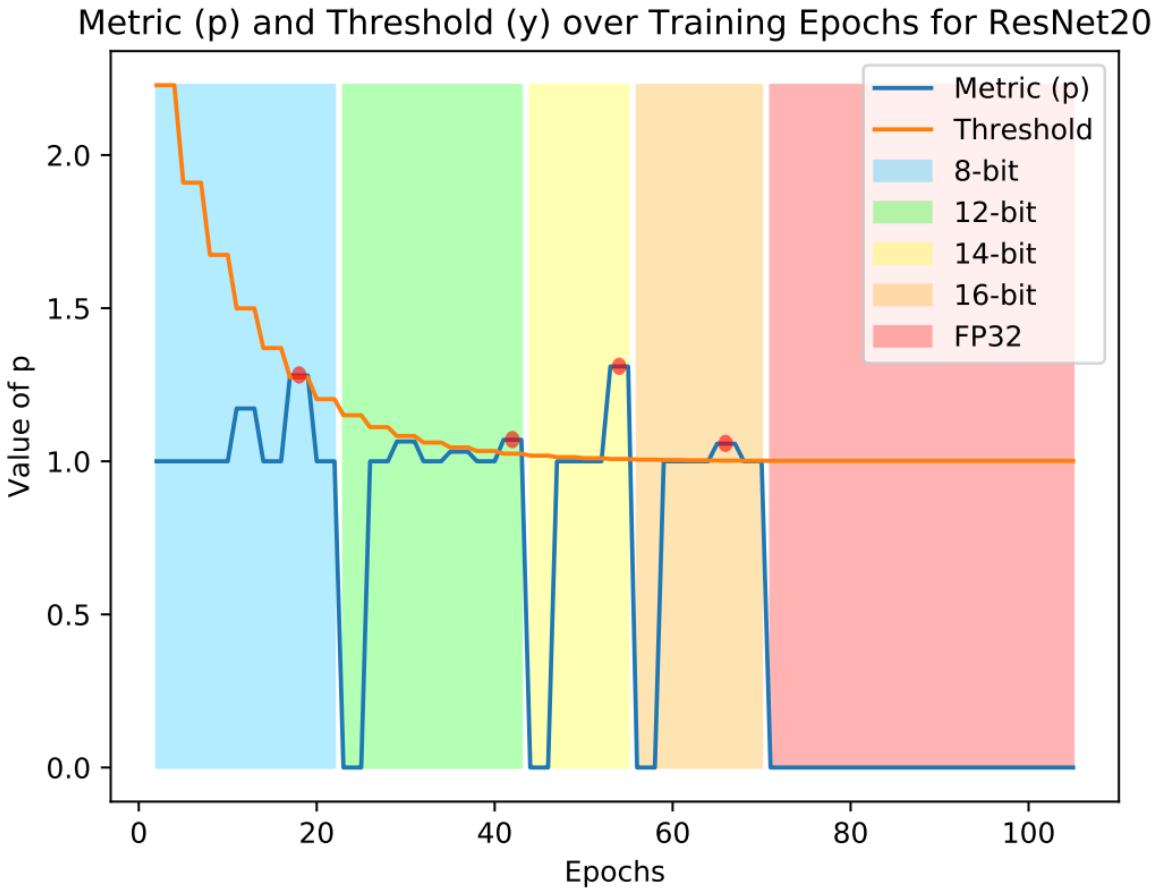

Figure 12. ResNet20 CIFAR10 - Demonstration of the generalisability of p over networks, datasets and epochs.

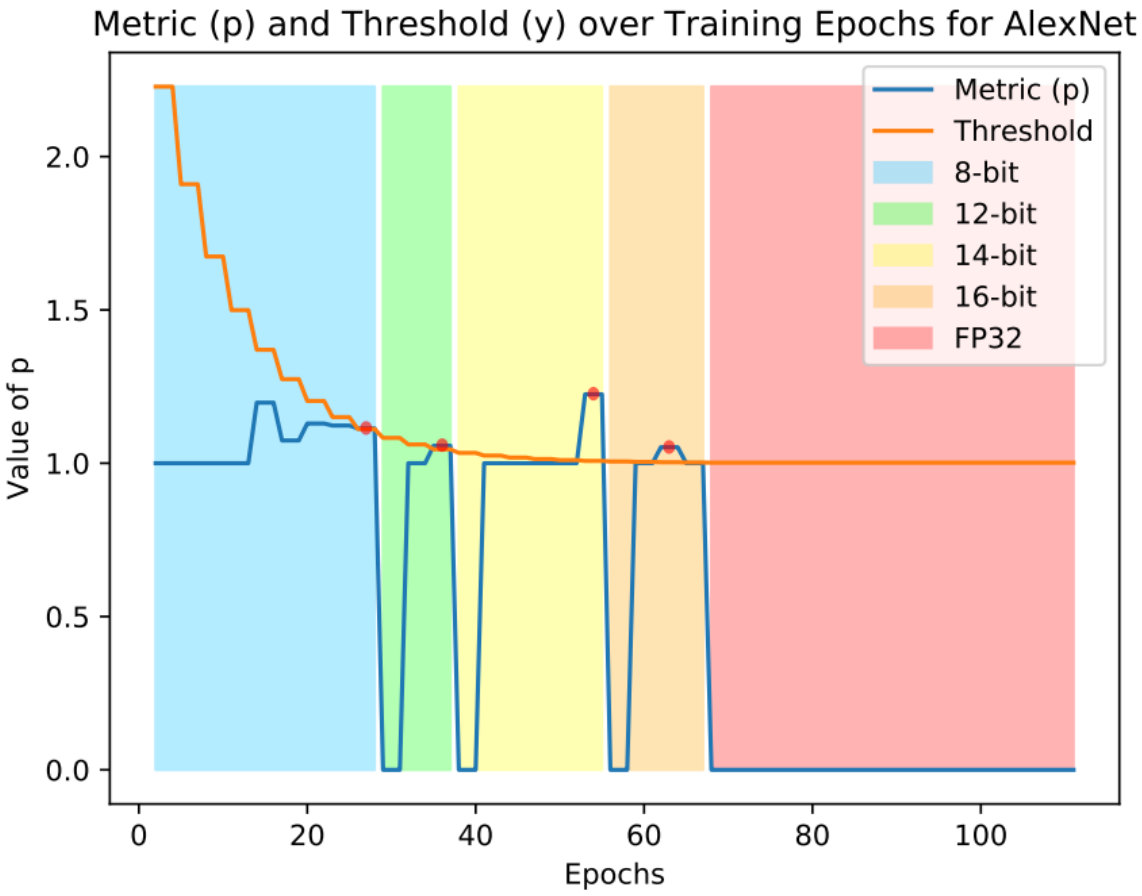

Figure 13. AlexNet ImageNet - Demonstration of the generalisability of p over networks, datasets and epochs.

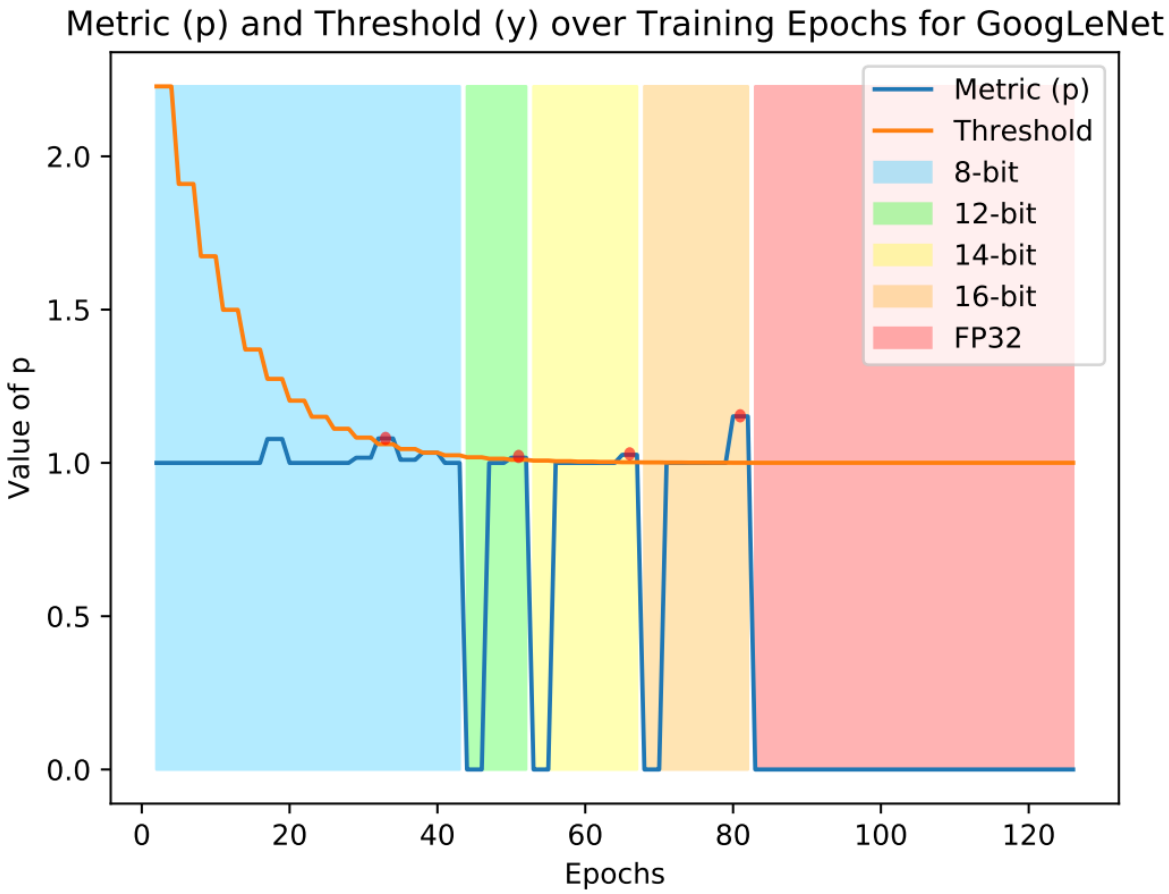

Figure 14. GoogLeNet ImageNet - Demonstration of the generalisability of p over networks, datasets and epochs.

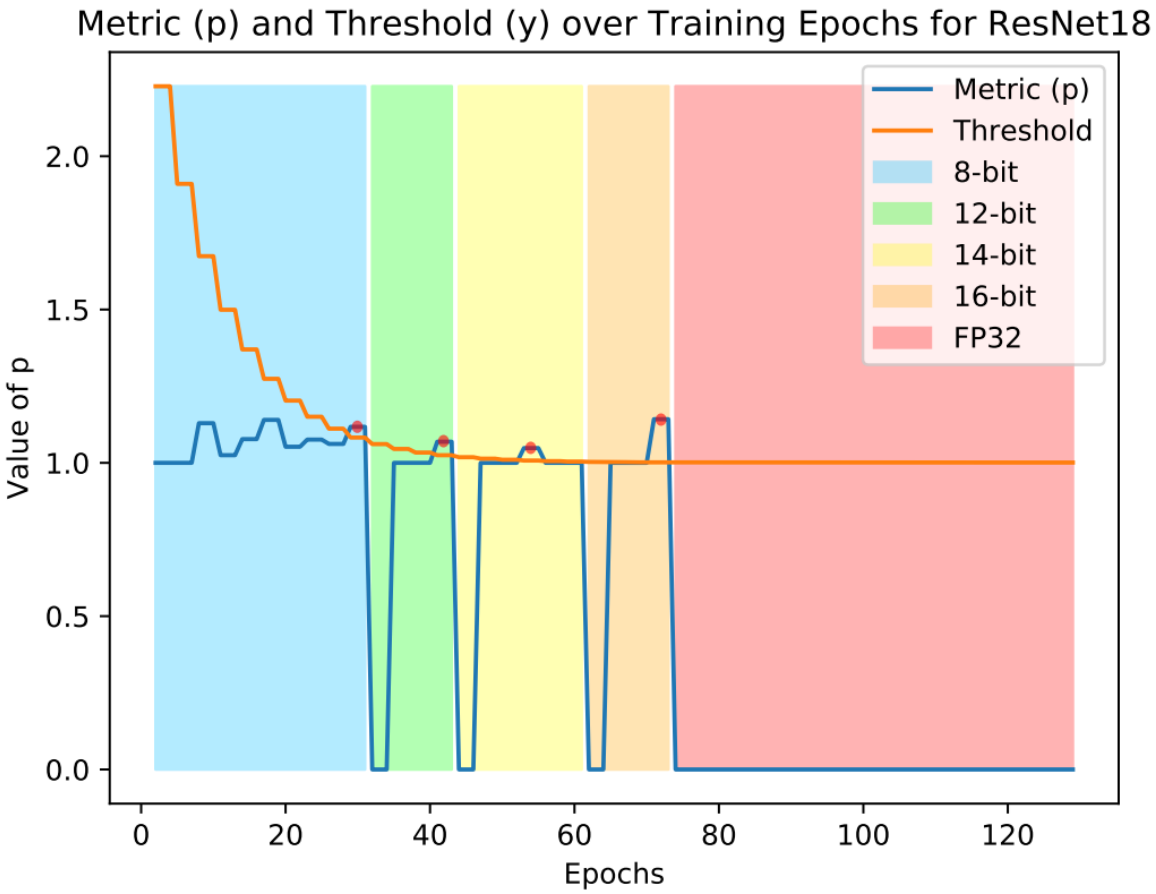

Figure 15. ResNet18 ImageNet - Demonstration of the generalisability of p over networks, datasets and epochs.