In deep learning, summarization models are traditionally trained using a maximum likelihood objective with reference summaries. Another line of work explores self-supervised approaches that do not require and are not limited by references. In this thesis, we opt for the latter approach. Our main contributions include the design of a novel dense reward function for summarization and its application for fine-tuning a sequence-to-sequence model via reinforcement learning. We build the whole training pipeline in a modular fashion, separately evaluating and tuning a supervised pre-training module, the reinforcement learning algorithm, and the reward function. After connecting all these components together, we also tune our self-learning approach as a whole. We evaluate the final checkpoints using 12 automatic and 3 manual metrics, revealing an improvement in reference-free metrics in nearly all cases.