

THE LAW OF KNOWLEDGE DISTILLATION

Jyh-An Lee

The Chinese University of Hong Kong Faculty of Law

First articulated in influential work by Geoffrey Hinton, Oriol Vinyals, and Jeff Dean, knowledge distillation has become a widely adopted technique for training smaller, more efficient “student” models to emulate larger “teacher” models, often achieving performance close to far more computationally intensive systems. Their early results—and subsequent industrial practice—showed that distilling ensembles into a single model can deliver substantial efficiency gains and lower latency with minimal loss of accuracy. As a result, distillation is now a staple for model compression, deployment on resource-constrained devices, and transferring capabilities across domains.

Over the past two years, however, several U.S. AI companies (including OpenAI and Anthropic) and government officials have publicly alleged that certain Chinese AI firms—such as DeepSeek, Moonshot AI, and MiniMax—have engaged in large-scale misappropriation of proprietary model capabilities via knowledge distillation and related model-extraction techniques. Although the technique itself is general-purpose and widely used, the legal stakes turn on how the teacher model was accessed and what contractual or statutory limits govern that access. This paper examines whether, and under what circumstances, knowledge distillation could constitute copyright infringement, trade secret misappropriation, and breach of contract. It also considers how intellectual property law and policy should adapt as distillation becomes more prevalent and assesses the broader implications for the U.S.–China AI competition.