Miscellaneous News and Thoughts

July 9, 2019 SKELJ Uncategorized 0

Just a few random bits:

1. Deep Learning IP Management / Watermarking

Developers have expressed interest in watermarking classifiers (See e.g., this UCSD paper). However, when developers say things like:

> *"The resulting models are therefore considered **to be the IP** of the model builder and need to be protected to preserve the owner's competitive advantage "*

it can have a somewhat cringe-inducing effect on the poor IP attorney reading the paper. How is a trained classifier "IP"? Let's be clear, I'm not talking about all that Monkey-Selfie-Deep-Dreaming-AI-Artist stuff. That's about protection in the output of the classifier. Here, we're instead discussing protection of the classifier itself. In the realm of watermark / IP interplay, at a high-level, let me tentatively posit that:

Relatively clear remedies: contract < patent < trade secret

Wild card: copyright

in terms of leverage effectiveness. That is, having detected unauthorized third party possession of your classifier via watermark (a nontrivial assumption – I'm not even sure the UCSD authors

RECENT UPDATES

DMCA Study

Skelley Snapshots

CPC Based Deal Search (i.e., not a COVID post)

NAVIGATION

Main

Updates

LawMux Knowledgebase

Software

FAQ

Contact

Subscribe

consider transfer learning), the legal vehicles for effecting a remedy should be approached in the preferential order indicated above.

Straightforward Remedies – Contract / Patent / Trade Secret

Trade secret (assuming you've properly maintained the classifier as such) is superior to patents because you have your choice of going into state (state trade secret statute) or federal court (under the DTSA). Trade secret also avoids the claim construction / design around risks of patents, prosecution costs, as well as the inevitable IPR. Both patents and trade secret provide vehicles for suing downstream users, although the DTSA has a knowledge requirement – see the definition of misappropriation in 18 USC 1839:

> *the term "misappropriation" means . . . acquisition of a trade secret of another by a person* ***who knows or has reason to know*** *that the trade secret was acquired by improper means*

Of course, patents have the benefit of protecting the classifier even once it's public and even if it's "innocently" acquired.

Contract will get you into state court, but typically not federal (though given the glacial backlog, you may prefer state in any event). Contract is so much lower than patents and trade secrets, though, because of the privity requirement. If thief A absconds with your classifier, who gives it to B, who gives it to C, you find the watermark in C, and you choose to sue C, you only have privity with A, so there's no breach of contract action against C. "But, but, at least in California, there's an ongoing affirmative duty to return stolen property, isn't there?" Yeah, but: 1) under the terms of the contract were the classifier weights "stolen?" (i.e., lacked title); and 2), even if they were "stolen", how are such weights "property" if neither patents, copyright, nor trade secret applies (maybe you'll get lucky and they passed around a CD/USB and you can recover that, but who does that anymore)?

Ok, those are straightforward. So why have I carved out copyright?

Wildcard Remedies – Copyright

If you assert that the weighted classifier is **copyrighted** your opponent will proceed to open a can of legal insanity by arguing that

there was insufficient authorial creativity on your part in the creation of the classifier to warrant copyright protection.

"How can that be?" you say. "I can receive a copyright in my object code after I run it through a compiler, right?" Yes, yes you can:

> *"A computer program, whether in object code or source code, is a "literary work" and is protected from unauthorized copying, whether from its object or source code version" Apple v. Franklin, 714 F.2d 1240 1983*

"So why can't I similarly receive a copyright in my classifier after I run it through my training regime?"

Because (argues the defense) a compiler (relatively) mechanically translates your (creative) source code into the (derivative) object code, whereas training a, say, deep learning classifier performs a stochastic (and, my, how they'll emphasize *that* word) gradient descent down the cost landscape based on the (likely randomly) selected training inputs. So is there *some* authorial participation in the classifier creation? Sure, you chose the loss function, training inputs, etc. Are those selections enough to warrant copyright in the resulting classifier?

Meh. Maybe?

Seriously, it'll depend on the specifics of the situation and the training system. If this really interests you, I can go collect some caselaw and try to predict which direction a court or the copyright office would come out on this (I know there's certainly some software-generated typography caselaw on point). But to a first order, given the ambiguity, I'd look to the other remedies first before hanging your hat here. Consequently, if you're going to start inserting watermarks as a control mechanism, I'd keep those remedies and the related detection scenarios in mind.

(Usual academic caveat applies – as mentioned in that last paragraph, you'll need to consider your specific situation, so certainly don't take any of this as legal advice – it's just the [probable] high-level state of the law)

2. WIPO BRIP

IP news is weird. If you go onto WIPO's front page you won't find *anything* about the new BRIP reporting system (it seems to be buried down here – basically it's a consolidated database of URLs presenting advertisements next to undesirable content [don't get me started on manga piracy] so that advertisers can withdraw their ads if they wish), however, TorrentFreak and Stanford have already written (July 7 and July 8, respectively) about it.

Why WIPO doesn't celebrate it and why TF and Stanford seem so lukewarm befuddles me (they're both [relatively] balanced, mentioning the pros and cons, but the problem addressed seems rather straightforward).

> *"WIPO provides no process by means of which a website operator can contest inclusion in the BRIP Database."*

Well, sure, I don't provide a process by means of which website operators can contest my including them on my links page. Fortuitously, other legal remedies are available if it's an issue (think, tortious interference with business relations).

Like TAG, I'll (tentatively) suggest that BRIP seems like a responsible development, depending on how it's used. Yes, if you're a search company monetizing based on ad revenue this may seem tedious, but, silver lining, if your customers are willing to pay a premium to control where their ads appear – potential cha-ching.

3. OSS Software

I put some barebones legal software online. I've been working / playing with pyTorch and Cuda lately (yes, I'm dancing all over the development chain), so I'll probably add some basic scripts once I find myself reusing something repeatedly.

PREVIOUS ARTICLE **NEXT ARTICLE**