



BLACK HAT AI TRACK SUBMISSIONS: OBSERVATIONS AND FEEDBACK

NATHAN HAMIEL (Black Hat Review Board Member)

OCTOBER 4, 2023

Reflecting on the submissions for the AI, ML, and Data Science track for Black Hat conferences for the past couple of years, I wanted to take some time to document a few observations and share some general feedback while my thoughts are still fresh. I hope this information better prepares people for submissions and helps them make the best use of their time with the highest chance of success.

There's always the chance that a great presentation falls through the cracks due to a poor submission. This post aims to help set people on the right track. I also hope this post gives people a bit more confidence to submit, even if they are new to Black Hat or the AI topic. Make our job even harder by submitting great proposals.

Note: I'm not asking for people to provide a 50-page response (this wouldn't be helpful either). I'm hoping people make their content more valuable by using the space available to cover the most important aspects of their submission.

Why Now?

Although we've had this track for a few years now, many of the submissions have been by practitioners working in the space with some academic background, but this year was different. With the massive hype around AI centered on Large Language Models (LLMs), there was an influx of submissions, including submissions by new presenters and people new to the topic. This was great to see. However, many of these submissions fell into a few traps. In this post, I'll highlight these traps by calling out some of my observations and providing some general feedback to help people avoid these pitfalls in the future.

The Primary AI Track

Observation: Many talks selected AI as the primary track, but they were a better fit for another track. In addition, many talks mentioned "AI," but the content had little to do with AI.

You can find the track description for the AI track [here](#).

It's always apparent when a submitter hasn't read the description. I think there's a lot of assumptions. Since Black Hat is a security conference and not an AI conference, the content and description have to be a bit broad, so it can get confusing.

Let me summarize: if your talk is primarily about a problem and you use some machine learning method in your approach, that is NOT a fit for the AI Track as the primary track for the submission. For example, if your talk is about reverse engineering a specific piece of malware and you happen to use ML to assist in that, that would be a better fit for the Reverse Engineering or Malware track as the primary, depending on the content.



If your talk is about using AI tools and approaches to assist in reverse engineering, that would be a good fit. Remember that the AI, ML, or Data Science aspect needs to be the key focus of the submission if you select this track as the primary track.

Black Hat Focus and Attendee Value

I spent an awful lot of time talking with attendees at Black Hat USA this year, asking them questions about the AI track. I asked what they thought of the content and what content they'd like to see. Many people were new to the topic and just trying to figure out where they stood and what they needed to know. This makes sense with all of the hype. However, the overwhelming consensus of people I talked to just wanted something they could use, basically asking for actionable content.

This actionable sentiment makes sense because Black Hat is an applied security conference. We've taken some things in the past that have been more theoretical and academic, but for the most part, the content needs to be useful for attendees immediately.

Actionable doesn't mean that all presentations need a tool or code release; they need content that attendees can use. So, to start with, ask yourself two fundamental questions.

- What do you expect attendees to do after your presentation is over?
- How will attendees use or apply the content and concepts you cover?

Your presentation and the content you cover should serve to answer these two questions.

Actionable on the AI Track

Observation: Submissions often weren't actionable or didn't have an actionable takeaway for attendees.

So, how do you make your content actionable on the AI track? It's pretty easy to determine by answering the two questions posed in the previous section.

- What do you expect attendees to do after your presentation is over?

If your answer to this question is to read my paper, spend months researching, and then publish your own paper with slightly better results, it won't be a good fit for the track. If the answer is understanding the approach I took to solving this problem and allowing them to adapt the code and content to their own environments, then that's a good fit. This means your content has to generalize to the audience or at least a particular segment of the audience.

It doesn't have to be as straightforward as it sounds, though. Many talks on reverse engineering a specific software aren't about the specific software being reversed. It's about the story and the approach. You can give people ideas about how to modify your approach to fit something new. Sylvain Pelissier's Practical Bruteforce of AES-1024 Military Grade Encryption talk is a good example of this. It had a bit of everything, a funny hook, a real-world story, Sylvain's thought process and approach to the problem, as well as perspectives from the affected company. There were multiple takeaways here that attendees could consider



when approaching their own research and product development. I chose this example because I had knowledge of the research from the beginning.

Observation: Submissions appeared to lack enough detail to reproduce the content submitted.

In order to succeed in creating actionable content, you have to provide enough information to make your work reproducible, and you have to provide enough information to bootstrap this effort when necessary. Think about this: if attendees can't reproduce your efforts, they are almost starting from scratch. This isn't helpful. If you can't share enough detail due to confidentiality or intellectual property issues, then you should reconsider submitting to Black Hat because your content appears more like a sales pitch than a value add for attendees.

Now, this level of detail doesn't mean you have to release a tool. It could be an approach or even a glimpse of something that attendees need to prepare for. This could be a roadmap or approach as well as a set of selected techniques and why you chose them. Even if your content is experimental, you must give attendees an idea of where to go next.

Academia vs Industry

Academia and industry are often confronted with different realities and different sets of problems. Both are useful and necessary but still different. Take adversarial attacks against specific image systems and object detectors.

Academia has spent much time ideating new attacks and defenses for these systems. This is great, but industry hasn't cared much because it doesn't impact most of them.

There is certainly some overlap between the two, and a silver lining here is that something not quite fit at an academic conference may be perfect for the practitioners at Black Hat and vice versa. If you are an academic and unsure if the content is a good fit, err on the side of submitting.

Generic Use of "AI" and Simple Overviews

I'm not going to spend much time on these topics because the issues should be self-evident, but since many submissions fell into this area, it's worth addressing.

Observation: Submissions peppered with the term "AI" without any mention of the actual approach.

Quite a few submissions fell into the following category: "We used "AI" for some task." This statement is then followed by a hundred mentions of the term AI. That's not helpful. Which method and approach did you use? If it's about solving the problem and not the approach, then it's a better fit for another track, not the AI track.

Observation: Far too many submissions were a simple overview or involved an uninteresting use case or approach



Simple overviews are not a good fit for Black Hat. There are some exceptions for extremely cutting-edge topics, but when a topic has been covered at length at other venues, it's a good indicator that it's probably not a good fit for Black Hat. This doesn't mean it's not a great talk or subject. Just know when your talk would better fit a regional security event or a blog post.

When it comes to use cases, remember that the audience is filled predominantly with security professionals. So, ensure your use case and content apply to them. Refer back to the actionable section and evaluate actions to ensure they align with expectations for security professionals.

Success and Benchmarking Criteria

Observation: Submissions often didn't contain any success or benchmarking criteria.

If you apply machine learning or deep learning to an approach, specify your success and benchmarking criteria. If you don't, how are reviewers supposed to evaluate your approach? This is critical in understanding whether your approach was successful or not and determining how successful the approach is in light of other approaches.

Far too many submissions fell into the bucket of "We used LLMs for 'X.'" Well, that's great, but did it work? How well did it work? How did using an LLM for this task compare to more traditional approaches? You can see where this is headed.

I was honestly a bit shocked by the lack of this basic information, which was a bit perplexing since it's critical to demonstrating the effectiveness of the approach, even to yourself, while experimenting. The assumption is that you didn't pay any attention to this and were only focused on making something work without regard to effectiveness.

Hype Is What Hype Is

Observation: LLMs were shoehorned into every use case.

With the level of hype around LLMs, it was inevitable that they would be shoehorned into every use case. This was even in cases where the problem itself wasn't interesting or in cases where we already had solid solutions for the problem.

I think of this as experimentation and the natural result of a new technology's introduction. Whenever a new technology comes along, people play around with it, try applying it to different use cases, and see what works. Nothing is wrong with this, but it's time to get real when submitting to a conference.

This is where you need to refer to the previous section on success and benchmarking criteria to demonstrate the value of your submission. It's okay to have a failed experiment or even subpar performance as long as there are takeaways and potential directions for others. Having a lessons-learned style of presentation can be helpful in certain circumstances. Just keep in mind, however, this is very situational.



If you are solving an already solved problem, you better bring it in some way and justify it with examples and success/failure criteria. Using a new technology to solve uninteresting or unimportant problems is also not a good recipe for success. Not every fun project makes a good conference submission.

Submission Issues

Of course, every year, there is no shortage of regular old submission issues unrelated to AI. These are the easy things to avoid, yet people often don't do them. I've got some updates to previous submission guidance I've given, and this isn't the place for that, but I want to hit a couple of highlights for quick reference.

- What's unique about your talk?

Ensure you've covered a unique angle or perspective your talk brings in the submission.

- Would you sit through your own talk?

This is a question almost nobody asks themselves, but it's enlightening on multiple levels.

- Think hard about your takeaways

Your takeaways are the reasons people would attend your talk. Every reviewer has takeaways in the back of their mind when reviewing your submissions. Ensure these are covered in your submission, either spelled out in the appropriate section or painfully obvious from the submission.

- Fill out the form completely

Yes, this actually has to be said. You'd be surprised at the number of people who submit incomplete proposals every single year.

- Get feedback

Find someone who will give you honest feedback and share the submission with them ahead of time. Feedback is the best way to anticipate potential questions and ensure the concepts you think are clear are actually communicated clearly.

- Preemptively answer questions

You can find some of these questions when you ask for feedback, but put your reviewer cap on. Pretend you are reviewing your submission and see if any obvious questions emerge. Your submission should answer more questions than it poses.

Don't Do This

Speaking of questions, don't ask a series of questions in your submission. This isn't a movie trailer; asking questions isn't an opportunity to build suspense with reviewers. I don't know if this is some new trend, but a few submissions did this, and it's not a recipe for success.



I noticed a few submission bodies and outlines were peppered with questions. Examples such as, “Did our approach work?” “Is it possible to implement our approach in production?” You get the point. It’s one thing to have these questions in the abstract since that’s public and will be displayed on the website. It’s another thing to put it in the submission body where reviewers are trying to evaluate the validity of your submission.

Conclusion

My hope is that people find this post helpful and it points people in the right direction. Preparing a submission for a conference can be daunting, but with a bit of preparation and feedback, your submission will have a better chance of getting selected. I’m looking forward to reviewing your submission.

Article used with permission from: <https://perilous.tech/2023/10/04/black-hat-ai-track-submissions-observations-and-feedback/>