Extracting Black-box Deep Learning Models via Software Based Power Consumption Measurements

DESIGN DOCUMENT

May22-8

Client/Advisor: Prof. Berk Gulmezoglu

Team Members/Roles:

Noah George	- Data Collection/Facilitator
David Swarts	- Data Collection/Testing
Mike Mazur	- Data Collection/Report Manager
Austin Van Brogen	- Neural Networks/Scribe
Danielle Rodriguez	- Neural Networks
Long Ly	- Neural Networks/Client Communication

sdmay22-08@iastate.edu
http://sdmay22-08.sd.ece.iastate.edu/

Revised: 11/29/21

Executive Summary

Development Standards & Practices Used

Agile development

Summary of Requirements

- Determine the accuracy and limitations of power side channel attacks
- Distinguish between different Neural Networks based on:
    - Layer depth
    - Hyperparameters
- Find a regression model between total execution time of the network and total # of layers/depth
- Quantify power measurements into meaningful training data points for RNN
- Use software-based programs (Bash scripting) to collect power consumption data
- Utilize python and Tensorflow libraries to train deep learning models

Applicable Courses from Iowa State University Curriculum

- Com S 474: Introduction to Machine Learning
- Com S 573: Machine Learning
- CprE 482X: Hardware design for Machine Learning
- EE 425: Machine learning: A signal Processing Perspective
- EE 526: Deep Learning: Theory and Practice
- Cpre 381, operating systems
- CprE 581: Computer Systems Architecture

New Skills/Knowledge acquired that was not taught in courses

Bash Scripting, Creating and Training Neural Networks, TensorFlow, Keras Library

Table of Contents

List of figures/tables/symbols/definitions

Side-channel Attack - In cyber security, this type of attack is based on any attacks that obtain information from the system implementation of a computer rather than its software implementation (i.e. through algorithm, buffers, applications etc.)

Neural Network - A neural network is made up of layers of simulated neurons. Each neuron has an edge weight that represents it's strength of connection.

Hyperparameters - Parameters that determine a network's structure

Team

1.1 Team Members

- Austen Van Brogen
- Noah George
- Michael Mazur
- Long Ly
- David Swarts
- Danielle Rodriguez

1.2 Required Skill Sets for Your Project

- Programming Skills (Java/Python/tensorflow/Matlab)
- Knowledge of AI/machine learning algorithms
- Project planning skills
- Knowledge of computer hardware (GPU/CPU)
- Knowledge of circuits systems
- Knowledge of power consumption in the system under different parameters

1.3 Skill Sets covered by the Team

- **Austen Van Brogen:** Programming Skills (C, Java, Python, Javascript, Assembly), Project planning skills, knowledge of computer hardware, decent hardware knowledge, **SE**
- **Noah George:** Programming(Java, C), Embedded systems, **EE**
- **Long Ly:** programming (C, Java, Ruby, Python, x86), proficient knowledge of circuits
- **Michael Mazur:** Programming skills (C, Java, Python, Assembly), hardware skills/knowledge, **CprE**
- **David Swarts:** Programming Skills (C, java, matlab), knowledge of computer hardware(GPU/CPU), knowledge of Circuits systems, **EE**
- **Danielle Rodriguez:** Programming Skills (C, java, python, javascript), Embedded Systems

1.4 Project Management Style Adopted by the team

- Agile

1.5 Initial Project Management Roles

- **Austen Van Brogen** - Meeting Scribe
- **Noah George** - Facilitator
- **Long Ly** - Client Interaction
- **Michael Mazur** - Report Manager
- **David Swarts** - Test Engineer
- **Danielle Rodriguez** - Test Engineer

1 Introduction

1.6 Problem Statement

In current day settings, deep learning is being successfully used in many different areas, e.g. computer vision, natural language processing and business intelligence. Deep learning architectures are also deployed for automating critical decision making in security applications and/or malware and intrusion detection. However, while these models carry extreme profitability, these deep machine learning neural networks require time, money and significant effort to set up. From collecting massive amounts of data samples to fine-tuning the network for performance, these massive networks are intellectual property for companies. As a result, these neural networks algorithms are typically presented to users in a black-box model. This type of commercial model does not reveal any information to the service users other than the output prediction of the input through the network.

The objective of this project, then, is to extract a black-box deep neural network model and infer information about the algorithm. We will then try to rebuild a substitute neural network model with functionality close to the target model. It can be shown that the neural network model is susceptible to certain timing side-channel attacks. From these attacks, a low-resolution channel can be seen for adversaries to infer depth (# of layers in the neural network). Other side-channel exploits, either with power, memory accesses, and/or cache attacks, can give fine-grained information about the target model during execution. The main problem, then, is to infer target neural network attributes through the usage of power sensor based (available in Intel and NVIDIA processors) side channels with the minimum number of network queries. Given enough attributes, such as layer's depth, the neural network parameters and hyperparameters, and extracted training data, it's possible to predict a new substitute architecture, closely related to the target architecture, through a Recurrent Neural Network (RNN) controller.

1.7 Requirements & Constraints

Requirements:

- Determine the accuracy and limitations of power side channel attacks
- Distinguish between different Neural Networks based on:
    - Layer depth
    - Hyperparameters
- Find a regression model between total execution time of the network and total # of layers/depth
- Quantify power measurements into meaningful training data points for RNN
- Use software-based programs (Bash scripting) to collect power consumption data
- Utilize python and Tensorflow libraries to train deep learning models

Constraints:

- Time
    - 20 Weeks of 6 people committing 10 hrs/week
- The ability to gather power usage will be limited by how quickly the system being targeted updates power usage data through its various programs
- Software Privilege. Individuals with higher privilege can have access to tools for the system that give more information about the power usage, so how much power usage data can be gathered will be limited by how much privilege the attacker can gain
- Available Hardware
    - RTX 3090

1.8 Engineering Standards

Systems and hardware are run on adequate and consistent power settings to prevent failure and measurement errors

Tensor Flow (https://www.tensorflow.org/guide) library to train deep learning model based on target model

1.9 Intended Users and Uses

Users

1. Security firms
    a. AMD
    b. Intel
2. Firms who own proprietary deep learning models
    a. Healthcare infrastructure
    b. Bank systems

    c. Google(DeepMind)
3. Researchers
    a. Professors
4. Cloud Service Providers
    a. Microsoft Azure
    b. Amazon Web Services

Uses

1. Security for any circuits not explicitly designed against power attacks
2. Analysis for hardware and machine learning algorithm vulnerabilities
3. Mount security and privacy attacks against the target model
4. Violate users' confidentiality and privacy of their input data
    a. eg. Hospitals
5. Testing new systems/hardware
    a. An automatic tool that reveals vulnerabilities in new hardware
    b. Checking if new hardware has fixed previous vulnerabilities

2 Project Plan

2.1 Project Management/Tracking Procedures

We will use an Agile approach to manage our project. This is a research project, so goals and objectives will change as we progress through the project. We will want to be able to adjust our expectations in the future. Agile is beneficial when the full project requirements are not completely understood.

What will your group use to track progress throughout the course of this and the next semester. This could include Git, Github, Trello, Slack or any other tools helpful in project management.

We will use the Git repos scrum board to track what everyone is doing weekly and to plan what to continue working on.

2.2 Task Decomposition

We're going to split up into two teams. One team will handle the collection of power consumption data, while the other team will train the oracle neural network model.

Power data collection team goal: collect power information through side-channels from the GPU

Task decomp:

1. Gain access to the GPU Machine and learn how to connect to it
2. Explore what sensors/interfaces/type of softwares are available on the GPU

3. Create automatic tool to collect power consumption measurements
4. Use analysis tools to draw meaningful conclusions to modify oracle model

Neural network team goal: Create and train a neural network that can function similarly to the black-box model

Task decomp:

1. Research concepts related to neural network & deep learning
2. Make neural networks of different types
3. Train the oracle model using Tensorflow library
4. Recover target model architecture through oracle model

Once these two teams have completed their tasks, the team goal will change to trying to create a substitute architecture of the original neural network.

1. Determine characteristics of neural networks based on power data recovered from side-channels
2. Recover substitute neural network architecture
3. Determine time needed for successful attacks on different networks

2.3 Project Proposed Milestones, Metrics, and Evaluation Criteria

Important Milestones:

- Getting access to GPU machine
- Creating a type of each neural network
- Training a neural network
- Infer neural network depth based on power data from side-channels
- Identify neural network types using power data from side-channels
- Use power data to create substitute neural networks

We can measure success with these metrics:

- Identify model within 20 neural network queries
- Identify model within 20 seconds
- Identify neural network model layers and hyperparameters is > 90% accuracy *
- Reconstructed training data is > 50% accuracy

* percent accuracy will vary depending on the layer of the model, some layers may not achieve 90% accuracy

2.4 Project Timeline/Schedule

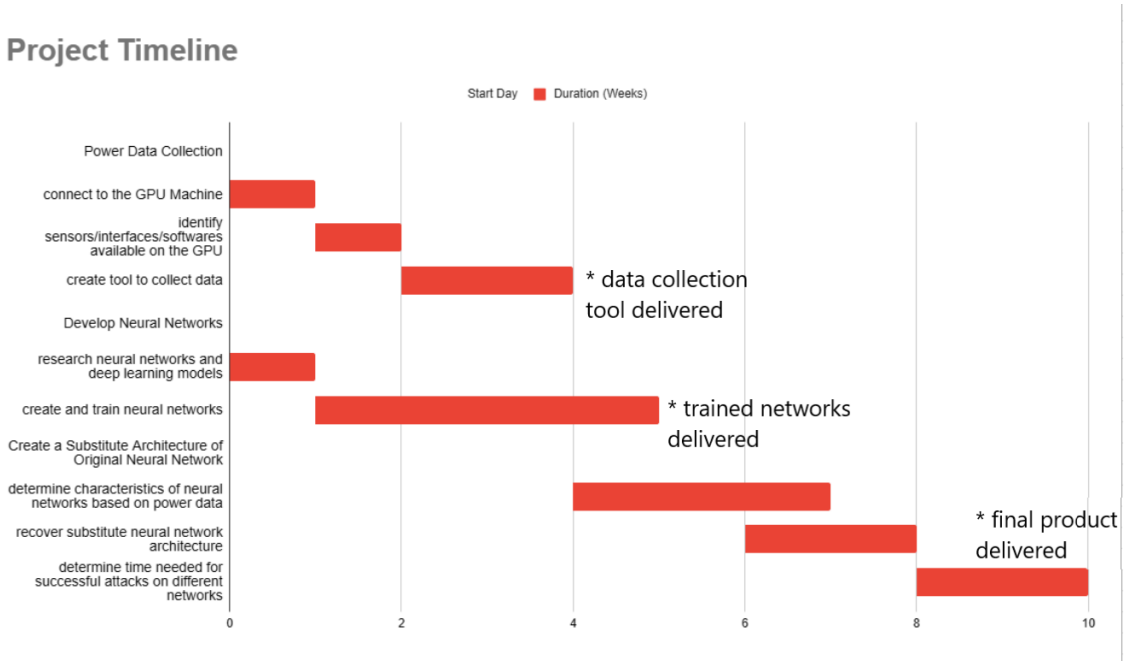

2.5 Risks And Risk Management/Mitigation

Risk: The GPU temperature sensor is not suitable(Too slow or inaccurate/ Does not exist) - 0.4 chance

Mitigation plan: use the sensor on the CPU (Intel RAPL interface)

Agile projects can associate risks and risk mitigation with each sprint

2.6 Personnel Effort Requirements

Task name	Person Hours
Gain access to the GPU Machine and learn how to connect to it	51
Explore what sensors are available on the GPU	51
Create tool to collect data	60
Make neural networks of different types	90
Determine characteristics of neural networks based on power data	90
Recover substitute neural network architecture	120

Determine time needed for successful attacks on different networks	120
	~500

2.7 Other Resource Requirements

This project will require a GPU equipped machine for training the neural network and for power data collection.

3 Design

3.1 Design Context

3.1.1 Broader Context

This project is to determine how much information can be gathered from a side channel attack. In an increasingly digital world, security is getting more and more important. From companies to governments to private citizens, many parties have private data that can be stolen. This project is for anyone who wants to keep their data secret.

Area	**Description**	**Examples**
Public health, safety, and welfare	Improved digital security protects companies and people's sensitive data	Any company or organization that needs to keep information a secret. This project specifically pertains to institutions with machine learning algorithms.
Global, cultural, and social	Large systems which implement deep learning and handle sensitive data value and prioritize security. Enhance awareness of the time of data that can be inferred using side channel attacks will help to improve security systems	Development or operation of the solution would violate a profession's code of ethics, implementation of the solution would require an undesired change in community practices

Environmental	This project will use electricity, which will be generated in part by fossil fuels.	The electricity from almost any location will be partially produced from fossil fuels.
Economic	This project could help save money by preventing algorithms from being stolen.	This project will help to discover the possibilities of side channel attacks so researchers and cyber security firms will be able to develop strategies to counteract them.

3.1.2 User Needs

Companies need a way to gauge the security risks when determining strategies to keep their proprietary algorithms from being stolen and used by other parties.

Governments need to assess the vulnerability of their systems from potential bad actors.

Researchers need to know how effective a power side channel attack is to learn about what's possible when it comes to side channel attacks.

3.1.3 Prior Work/Solutions

This project is based on a research paper[1] that used the Intel RAPL (Running Average Power Limit) interface to steal data and cryptographic keys. We want to take this idea and apply it to neural networks running on a GPU.

[2] shows how a neural network can be stolen and reconstructed by a side-channel attack.

3.1.4 Technical Complexity

The project requires a knowledge of computer science and machine learning. This project also requires us to be proficient electrical, computer, and software engineers to be able to create both machine learning algorithms, a way to measure power consumption, and understanding of how computers function to determine where to measure power from.

This problem requires us to develop machine learning algorithms, develop a precise power measurement tool, and to create a way to steal information from a GPU based on power measurements.

3.2 Design Exploration

3.2.1 Design Decisions

This project may be research based, but there are still several design decisions we need to make. They include:

- We will need to decide which Machine Learning algorithms to analyze
- We will need to determine which methods to use for gathering power consumption data
- We will need to decide how much data to gather before analyzing
- We will need to decide whether to use Matlab or Python for analyzing the data

3.2.2 Ideation

Brainstorming, Reflection on previous design considerations, Discovery with client, Reflection on past experience, lotus blossom

3.2.3 Decision-Making and Trade-Off

The criteria for our design decision and making is very much led by our clientele, Dr. Gulmezoglu. Also due to the research-like scope of this project, many softwares used is to be in accordance with those available to the mass public (i.e. Python & bash scripting). While limitations are to be expected with coding languages chosen, like Python, it is the best under the circumstances to be used with an NVIDIA 3090 GPU and the TensorFlow library. As the semester and project progresses, we will hopefully be able to clarify these different design trade-offs.

3.3 Proposed Design

3.3.1 Design Visual and Description

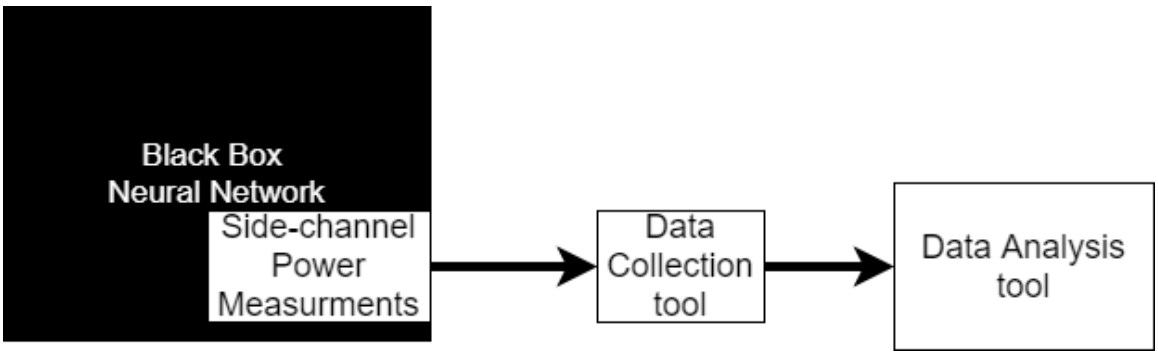

We will be running a neural network on a GPU machine without knowing the architecture. We will create a data collection tool with software to collect power measurements from the system. Then, we will create a data analysis tool to try to identify and recreate the black box neural network.

3.3.2 Functionality

In the real world, this research is intended to be used by owners of black box neural networks to protect their secret designs. If we can prove that neural networks can be stolen this way, these owners will want to take steps to prevent access to these power measurements.

It is also possible that this technique will be used by someone to steal black box neural networks.

3.3.3 Areas of Concern and Development

One of our biggest concerns is that the date we collect will not be accurate enough, or that we will not be able to collect the data fast enough for it to be useful. If we determine that we can't do anything with the data we collect, we will have to change tracks and look at other places we can collect power data, like from the CPU using the Intel RAPL.

3.4 Technology Considerations

In terms of technological softwares needed, most of the neural network model training will be accomplished through the Python platform running Tensorflow and the TensorFlow built-in version of Keras. Depending on the complexity of the model, later deep learning TensorFlow libraries might be added but for now Keras will suffice. On the power collection side of things, we will utilize a bare-bone bash-scripting language structure to automate the collection of power draw on the 3090 GTX.

3.5 Design Analysis

We believe, with the current state of our project, that this idea is still worth exploring. We have not yet begun to create neural networks, so there is still work for us to complete.

3.6 Design Plan

Our design plan will have three main modules (Data Analysis, Data Collection, Neural Network). We will run a Neural network on a GPU and record data from it such as how much power it's using. The Data Analysis module will use the parameters recorded with the Data Collection tool to try to figure out as much about the neural network being ran on the GPU as possible.

4 Testing

4.1 Unit Testing

Our program will consist of three main components: the power measurement consumption unit, the neural network training unit, and the analysis tool. For each of these units we will test function return values and execution time to achieve the timing goals described in the requirements.

Power Measurement Consumption Unit-

Since this unit's main function is to collect power consumption data, we will test the access to power consumption data from different user spaces (i.e. root vs non-root access) and from different platforms (i.e. sensors, GPU, etc.) Additionally, we will test the collected data for accuracy.

Training Unit-

In this unit, we will build and train our neural networks using python's TensorFlow library. We will test these networks the performance and accuracy of these networks with varying hyperparameters and layers.

Analysis Tool-

Our analysis tool will take the power consumption data as input, perform analysis, then output information about the neural networks. We will test the tool's ability to handle varying input lengths and check for loss or corruption of data. Additionally, we will test the output of the tool to check for its accuracy in determining the hyperparameters, number of layers, and layer types of the networks with respect to the different inputs.

Python provides multiple libraries for unit testing and special extensions for creating mock objects and testing neural networks built with Tensorflow

-**pytest:** unit testing framework provided which allows for testing APIs, UIs, and databases

-**unittest**: unit testing framework provided by Python library

-**tf.test**: extension of unit test which contains assertions tailored to Tensorflow code

4.2 Interface Testing

The project will have two main interfaces in order to allow the project to properly work. The first one will be an interface to show the output of the neural network layers as a diagram, and the second one will be an interface that will allow us to test our neural network properly.

For the first interface, we need something that easily displays the number of layers in the black-box in order to make it easier for the user to to understand the inner workings of the black box. This will require a custom interface that takes the output from the black box and makes a nice little diagram to show what our neural network believes the inside of the black box looks like.

The second interface will be to make an interface that allows us to test if inputs given to an interface match expected outputs. This interface also ties into the first one as it checks to see if the diagram produced using the prior interface is correct to the expected output.

For both of these interfaces we can use Mockito as it specializes in testing interfaces.

4.3 Integration Testing

We have to Integrate power measurements from the processor, to our analysis tool, and Connect that with the hyper parameters from the neural network we're training. We want to be able to combine all of our code into one program and to do so, we can use Python.

We can use Python to interface our power analysis tool with the processor running our neural network. The power measurements tool can be implemented in Python. We can use machine

learning libraries like the Tensorflow library for Python to interface our data recording tool and the neural network itself, so everything can be integrated seamlessly.

We will test the integration by checking if the data from our power measurements reflects the data in the analysis tool. Then we will check if the measurements match our model for the hyper parameters of the neural network being run.

4.4 System Testing

To fully test the research system, we will have to utilize power consumption and training tools unit tests alongside integration testing. These unit tests will help to verify the stability of the training model and power consumption measurements being run on the RTX 3090. Depending on the level of complexity of the model, interface testing will be limited to the softwares, and testing libraries, they were created upon. As such, interface testing in this sense will be fully defined later in the design flow. In turn, while our research project does not require too many interfaces, we will have to utilize the above integration testing strategy to verify the stability of the system. Our system will be mainly our power consumption measurement unit running in parallel with the oracle model. The training model will then learn in sync with the power measurement output.

4.5 Regression Testing

Necessary Features:

- Data collection tool
- AIs
- Power Sensors
- GPU

Ways to mitigate breaking old functionality:

- Keep old tests to see if a new implementation break anything old
- CI/CD to ensure if new code breaks the system it won't be merged into the official code
- Requirements driven testing
- Back up the AIs in case new training corrupts them

4.6 Acceptance Testing

Because this is a research project our only requirement is that we come to a conclusion that we can support with evidence. Therefore, we will have to create enough meaningful data that our statistical error is sufficiently small.

Our initial targets are to:
-Identify model within 20 queries
-Identify model within 20 seconds
-Identified model is > 90% accuracy *
-Reconstruct image training data with > 50% accuracy

We will maintain contact with our client to ensure these goals are met, and in the case we can't meet these expectations, create updated goals. We have no subjective goals for this project because it is a software-based research project.

4.7 Results

From our testing, we hope to be able to make our project more efficient, reliable, and able to comply with the requirements to the best of its ability. Testing is important in order to make sure that mistakes are ironed out and that bugs don't become features. They will help improve our ability to detect errors in our system and overall allow us to make the project overall better.

5 Implementation

Our plan is to create two separate teams to work on the Neural Networking and the Data analysis parts of the project. Each team will consist of 3 members and will focus on only those parts for the rest of the semester. The overall design of the project only consists of three major components (see 3.3.1). According to the professor, this should not take too long and the two teams should be able to finish their respective components rather quickly.

So far, our data collection team has created a tool to collect data and we have collected preliminary measurements of a neural network processing training data. There is a clear difference between running the network at different layer amounts, so our project looks promising.

6 Professionalism

6.1 Areas of Responsibility

Area of responsibility	Definition	NSPE Canon	IEEE Code of ethics	Own words
Work Competence	Perform work of high quality, integrity, timeliness, and professional competence	Perform services only in areas of their competence; Avoid deceptive acts	Maintain and improve our technical competence and to undertake technological tasks for others only if qualified by training or experience	This pledges to stay up to date with technological progress, as well as to only work on what we are qualified to work on. This differs from NSPE by adding that an engineer not only

				needs to be competent, but also improve.
Financial Responsibility	Deliver products and services of realizable value and at reasonable costs	Act for each employer or client as faithful agents or trustees.	Reject bribery in all its forms; Avoid real or perceived conflicts of interest whenever possible	The IEEE code does not directly state that products should be financially valuable, but its does say to avoid bribery and other conflicts of interest.
Communication Honesty	Report work truthfully, without deception, and understandable to stakeholders.	Issue public statements only in an objective and truthful manner; Avoid deceptive acts	Seek, accept, and offer honest criticism of technical work, acknowledge and correct errors, be honest and realistic in stating claims or estimates based on available data, and credit properly the contributions of others;	Be honest in criticism, statistics and data, and give credit where credit is due. The IEEE standard is very similar to the NSPE canon here, although it doesn't limit honesty to public statements.
Health Safety and Well-Being	Minimize risks to safety, health, and well-being of stakeholders.	Hold paramount the safety, health, and welfare of the public.	I.1. to hold paramount the safety, health, and welfare of the public, to strive to comply with ethical design and sustainable development practices, to protect the privacy of others, and to disclose promptly factors that might endanger the public or the environment;	Your work should hold the safety of others in high regard. Don't work on anything that could endanger the public, and make sure to protect the privacy of others.
Property Ownership	Respect property, ideas, and information of clients and others.	Act for each employer or client as faithful agents or trustees.	5. to seek, accept, and offer honest criticism of technical work, to acknowledge and correct errors, to be honest and realistic in stating claims or estimates based on available data, and to credit properly the contributions of others;	Don't try to take other's work and pass it off as your own.
Sustainability	Protect environment and natural resources locally and globally	N/A	I.1. to hold paramount the safety, health, and welfare of the public, to strive to comply with ethical design and sustainable	You should make sure your work doesn't harm the environment.

			development practices, to protect the privacy of others, and to disclose promptly factors that might endanger the public or the environment;	
Social Responsibility	Produce products and services that benefit society and communities.	Conduct themselves honorably, responsibly, ethically, and lawfully so as to enhance the honor, reputation, and usefulness of the profession.	(But also I.1, above) I.2. to improve the understanding by individuals and society of the capabilities and societal implications of conventional and emerging technologies, including intelligent systems;	Make sure to maintain people's privacy. Make sure to think about the societal implications of your work. Help society learn about emerging systems.

6.2 Project Specific Professional Responsibility Areas

Work Competence: This applies to our project because if the team members can't perform work of high quality, then the results of the power analysis won't be accurate and the entire project will be pointless.Our team has been performing medium-high for this because while none of us had previous experience with power analysis or machine learning, we've been doing good research to improve our abilities and have met with our client several times to ensure we understand exactly what is needed.

Financial Responsibility: This field applies to our project because, while we won't be making anything physical, we will still need some hardware in order to build our software product. We will be using hardware that the university already owns, so we won't be purchasing any new material. The financial cost will be the value of the time we use on the machine, which is pretty negligible.

Communication Honesty: This field applies to the project because we have a client that we need to communicate with and deliver the product to. Our performance here has been high because we've been completely open with communication with them and have made sure to go to them when we aren't sure what sort of actions need to be taken for the product.

Health, Safety, and Well-being: This field doesn't have much application to our product as it's a pure software product so there's no threat of any physical harm.The biggest concern here would be that it could be used to expose data vulnerabilities of companies that a malicious actor might wish

to take advantage of, which could be quite harmful. Our performance has been high here, as we haven't exposed any data vulnerabilities and no one has been harmed.

Property Ownership: This area of responsibility is applicable to our project due to the research-like scope of this particular subject. Many of our provided research and development will be based on source codes or example models that might have been under the ownership of another creator. Simultaneously, many of our documents/reference papers relating to our project are under the publication of other researchers and research institutions. Our team is performing high in this area as we are currently making sure to cite sources and keep notes of outside links & references.

Sustainability: In the area of sustainability, our team is performing quite high due to our low resource usage of GPU machines in the lab environment. This will change more over time throughout the second semester when we are starting the development process. However, outside of the GPU machines, the area of sustainability does not have many applications due to our area of research and the project's topic at large.

Social Responsibility: In terms of social responsibility, our team has been performing quite well with keeping the confidentiality of our clientele/advisor and project designs from outside groups. Because of the power-side channel attack implications of this project, we will have to be more careful during development to not violate any vital data within the GPU's neural network/machine.

6.3 Most Applicable Professional Responsibility Area

Communication Honesty

Our group will not keep things from one another and work together without holding information back. We need to be able to trust one another to get our respective jobs done. One way that we have been able to demonstrate our group's commitment to communication honesty is holding integrity between what we say we can get done by some checkpoint and what is actually delivered. Life happens and sometimes we will not be able to achieve everything we set out to when we made the plan with our other teammates. The important thing is to communicate with the team about your situation so that they can adjust their plans accordingly.

Communication is not just necessary for keeping a happy healthy team environment, it is also crucial for delivering a reliable product. When individuals within a team work on separate components of a shared project, it is important to be transparent about how your component works so that the other team members can ensure that the pieces will fit together smoothly.

7 Closing Material

7.1 Discussion

We have identified the nvidia-smi tool where power information is available to measure. Using this interface, we have created a tool that can collect power data from the GPU.

7.2 Conclusion

Our goal is to collect and analyze power data from the GPU with the goal of using it to differentiate between different types of neural networks. So far, we have created a tool that can collect power information. We have not yet created any neural networks to collect data from.

In the coming semester, we plan to begin to create these neural networks to collect data from. Then, we will be able to analyze this data and draw better conclusions.

7.3 References

[1] M. Lipp et al., "PLATYPUS: Software-based Power Side-Channel Attacks on x86," 2021 IEEE Symposium on Security and Privacy (SP), 2021, pp. 355-371, doi: 10.1109/SP40001.2021.00063.F

[2] Duddu, Vasisht et al.,"Stealing Neural Networks via Timing Side Channels." ArXiv abs/1812.11720 (2018)

7.4 Appendices

NVIDIA-SMI Documentation:

https://developer.nvidia.com/nvidia-system-management-interface

7.4.1 Team Contract

Team Members:

1) Michael Mazur
2) Austen Van Brogen
3) Noah George
4) Long Ly
5) David Swarts
6) Danielle Rodriguez

Team Procedures

1. Day, time, and location (face-to-face or virtual) for regular team meetings:
    - Team meeting: Thursdays after TA meeting
    - TA meetings: Thursday 2 pm (virtual)
    - Client meetings: Friday 2:30 pm (virtual)
2. Preferred method of communication updates, reminders, issues, and scheduling (e.g., e-mail, phone, app, face-to-face):
    - Discord for virtual meetings and Parks library for face-to-face meetings.
3. Decision-making policy (e.g., consensus, majority vote):
    - Majority vote at team meetings
4. Procedures for record keeping (i.e., who will keep meeting minutes, how will minutes be shared/archived):
    - Austen Van Brogen will record the duration of meetings and record basic stuff in a shared document.

Participation Expectations

1. Expected individual attendance, punctuality, and participation at all team meetings:
    - Team members are expected to show up to 80% of meetings and be no more than 15 minutes late. If a team member cannot attend a meeting, they should give ample notice time.
2. Expected level of responsibility for fulfilling team assignments, timelines, and deadlines:
    - Team members are expected to fulfill their assignments before a deadline, however, if there is an issue with meeting the deadline, they should notify the rest of the team beforehand.
3. Expected level of communication with other team members:
    - Team members are expected to communicate with each other on Discord in a timely manner. Typically members are expected to respond within a 24 hour time period.
4. Expected level of commitment to team decisions and tasks:
    - Team members are expected to equally contribute to decisions and tasks. If a team member is not contributing fairly the team will meet to discuss repercussions.

Leadership

1. Leadership roles for each team member (e.g., team organization, client interaction, individual component design, testing, etc.):
    - Long Ly: Client Interaction
    - Noah George : Facilitator
    - Michael Mazur: Report Manager
    - Austen Van Brogen**:** Meeting Scribe
    - David Swarts: Test Engineer
2. Strategies for supporting and guiding the work of all team members:
    - Weekly meetings to share progress and ask questions
3. Strategies for recognizing the contributions of all team members:
    - Round of applause at weekly meetings.

Collaboration and Inclusion

1. Describe the skills, expertise, and unique perspectives each team member brings to the team.
    - **Michael Mazur:** Programming skills (C, Java, Python, Assembly), hardware skills/knowledge, CprE
    - **Noah George:** Programming(Java, C, Python), Embedded systems, EE

- **Austen Van Brogen:** Programming Skills (C, Java, Python, Javascript, Assembly), Project planning skills, knowledge of computer hardware, decent hardware knowledge, SE
- **Long Ly:** programming (C, Java, Ruby, Python, x86), knowledge of circuits, knowledge of embedded systems programming
- **David Swarts:** Programming Skills (C, java, matlab), knowledge of computer hardware(GPU/CPU), knowledge of Circuits systems, EE
- **Danielle Rodriguez:** Programming Skills (C, java, python, javascript), Embedded Systems

2. Strategies for encouraging and support contributions and ideas from all team members:
    - Consistently inform members of responsibilities.
    - Ask for input
    - Respectfully listen and engage in their ideas during meetings
    - Constructive criticism
3. Procedures for identifying and resolving collaboration or inclusion issues (e.g., how will a team member inform the team that the team environment is obstructing their opportunity or ability to contribute?)
    - If a member feels that they are being obstructed, they should voice their concerns in discord or at a group meeting. If the obstruction prevails, then the member should contact the TA/professor.

Goal-Setting, Planning, and Execution

1. Team goals for this semester:
    - Create an achievable plan of action for the upcoming semester
    - Structure project's timeline and deadlines
    - Have proficient knowledge of the project at hand
    - Understand client's expectations and requirements
    - Create design plans tailored to the client's idea
    - Obtain an initial design structure or prototype
    - Obtain clear knowledge of team's available resources and availability
2. Strategies for planning and assigning individual and team work:
    - We will be using git as our main software for version control and task management. Tasks will be posted to the git repo's scrum board. Members will be assigned tickets/issues with deadlines to be met. Issues assigned are expected to be completed or near completion on the day of the team meeting. Any problems and questions on implementation should be addressed with the team as soon as possible.
3. Strategies for keeping on task:
    - Visit scrum board daily
    - Visit main communication software with team daily
    - Communicate deadlines and arising problems
    - Communicate team meetings and availability

- Don’t procrastinate!

Consequences for Not Adhering to Team Contract

1. How will you handle infractions of any of the obligations of this team contract?
    - Depending on the severity of the infarction, the team will either discuss the matter with the member under question or take the matter to a TA/ professor.

2. What will your team do if the infractions continue?
    - The team will take the matter to the TA/professor and ask for the member to be taken out of the group.

1) _____ DATE _____

2) _____ DATE _____

3) _____ DATE _____

4) _____ DATE _____

5) _____ DATE _____

6) _____ DATE _____

7) _____ DATE _____

8) _____ DATE _____