

Week 15 - Monday

COMP 4290

Last time

- What did we talk about last time?
- Ethical case studies

Questions?

Project 3

Assignment 5

Ahmed Mohamed Presents

More Case Studies of Ethics

Case V: Proprietary Resources

- Suzie owns a copy of Photoshop which she bought legitimately
- As you know, the software is copyrighted, and the documentation contains a license agreement that the software is for the purchaser only
- Suzie invites Luis to look at the software to see if it will fit his needs
- Luis examines the software on Suzie's computer and likes it
- He wants to try it in a longer test

Different outcomes

- What are the ethical issues in each of the following separate scenarios:
 - Suzie offers to copy the software for Luis to use
 - Suzie copies the software for Luis to use, and he uses it for some period of time
 - Suzie copies the software for Luis to use, and he uses it for some period of time and then buys a copy
 - Suzie copies the software for Luis to try overnight with the understanding that he must bring it back tomorrow without copying it, and he complies
 - Suzie does the same, but Luis makes a copy anyway
 - Suzie does the same, Luis makes a copy, but he eventually buys a copy
 - Suzie does the same, but Luis never returns the copy

Case VI: Fraud

- Alicia is a programmer for a corporation
- Her supervisor Ed tells her to write a program that allows people to edit company accounting information directly
- Alicia knows that programs that can edit company accounts usually have several steps with checks in them
- This program would allow anyone to change the books without a trace
- Alicia mentions these issues to Ed
- Ed says that her job is to write the software she's told to write
- He says that this software can be used to correct mistakes made

Analysis

- Is a programmer responsible for the programs he or she writes?
- Is a programmer an employee who follows orders unthinkingly?
- What degree of personal risk is an employee obliged to accept for opposing an improper action?
- Would a program like the one here ever be justified? When?
- How could a program like this one be controlled?
- Would the ethical issues be changed if Alicia wrote this program on her own?

Case VII: Accuracy of Information

- Emma is a researcher who is analyzing the nutritional content of a cereal called Raw Bits
- She gets a statistical programmer Paul to analyze the data
- His analysis shows that Raw Bits is not nutritious and may be harmful
- He suggests that another set of correlations could show Raw Bits in a more favorable light
 - He claims he could argue any side of any issue with statistics

Analysis

- Is it ethical for Paul to suggest analyzing data to support two different conclusions?
- Is Paul obligated to present both positive and negative analyses?
Is he responsible for their use?
- Is it ethical for Emma to accept positive or negative conclusions if she doesn't understand the statistics?
- She suspects that the company will
 - Get a new researcher if she sends them only the negative results
 - Publicize only the positive results if she sends them both
- What course of action should she take?

Case VIII: Ethics of Hacking or Cracking

- Goli is an independently wealthy computer security specialist
 - She works only for fun
- She attacks commercial products for vulnerabilities and is good at finding them
- She probes systems on the Internet and, when she finds vulnerabilities, she contacts the owners of the sites to offer her services to fix them
- She loves good pastry and plants programs that slow the performance of web sites of bakeries that don't use enough butter in their pastries

Analysis

- Is it ethical for Goli to probe for vulnerabilities in systems?
- What if her probing sometimes causes failures or performance problems?
- How much and to whom should she report the vulnerabilities she finds?
- What if she damaged websites based on an issue more serious than butter?
 - What if she only damaged websites for companies with records of human rights abuses?

Codes of Ethics

10 Commandments of Computer Ethics

I like these because they are short and clear:

1. Thou shalt not use a computer to harm other people.
2. Thou shalt not interfere with other people's computer work.
3. Thou shalt not snoop around in other people's computer files.
4. Thou shalt not use a computer to steal.
5. Thou shalt not use a computer to bear false witness.
6. Thou shalt not copy or use proprietary software for which you have not paid.
7. Thou shalt not use other people's computer resources without authorization or proper compensation.
8. Thou shalt not appropriate other people's intellectual output.
9. Thou shalt think about the social consequences of the program you are writing or the system you are designing.
10. Thou shalt always use a computer in ways that insure consideration and respect for your fellow human beings.

ACM

- The Association for Computing Machinery (ACM) is the most important professional association for computer scientists
 - Followed by the IEEE Computer Society
- The ACM gives out the Turing Award every year
 - The closest thing to a Nobel Prize for computer science
- It's a bigger deal for academics than it is for most professionals
- The ACM holds conferences and publishes journals in many different areas of computer science
- Another thing it does is publish the a code of ethics, which is common for professional associations

ACM Code of Ethics Section 1

- General ethical principles:
 1. Contribute to society and to human well-being, acknowledging that all people are stakeholders in computing.
 2. Avoid harm.
 3. Be honest and trustworthy.
 4. Be fair and take action not to discriminate.
 5. Respect the work required to produce new ideas, inventions, creative works, and computing artifacts.
 6. Respect privacy.
 7. Honor confidentiality.

ACM Code of Ethics Section 2

- Professional responsibilities:
 1. Strive to achieve high quality in both the processes and products of professional work.
 2. Maintain high standards of professional competence, conduct, and ethical practice.
 3. Know and respect existing rules pertaining to professional work.
 4. Accept and provide appropriate professional review.
 5. Give comprehensive and thorough evaluations of computer systems and their impacts, including analysis of possible risks.
 6. Perform work only in areas of competence.
 7. Foster public awareness and understanding of computing, related technologies, and their consequences.
 8. Access computing and communication resources only when authorized or when compelled by the public good.
 9. Design and implement systems that are robustly and usably secure.

ACM Code of Ethics Section 3

- Professional leadership principles:
 1. Ensure that the public good is the central concern during all professional computing work.
 2. Articulate, encourage acceptance of, and evaluate fulfillment of social responsibilities by members of the organization or group.
 3. Manage personnel and resources to enhance the quality of working life.
 4. Articulate, apply, and support policies and processes that reflect the principles of the Code.
 5. Create opportunities for members of the organization or group to grow as professionals.
 6. Use care when modifying or retiring systems.
 7. Recognize and take special care of systems that become integrated into the infrastructure of society.

ACM Code of Ethics Section 4

- Compliance with the Code:
 1. Uphold, promote, and respect the principles of the Code.
 2. Treat violations of the Code as inconsistent with membership in the ACM.

- Full text of the ACM Code of Ethics available here:
 - <https://www.acm.org/code-of-ethics>

AI and Cybersecurity

Ethics and AI

- Julia Bossman lists nine questions to consider when developing AI:
 1. **Unemployment:** What happens to the human workforce as AI replaces human jobs?
 2. **Inequality:** New applications of AI will create wealth. How do we distribute that wealth fairly and equitably among the human population?
 3. **Humanity:** How will AI affect the way we behave and interact?
 4. **Errors:** How do we guard against mistakes? How do we assess blame and make restitution?
 5. **Bias:** How do we prevent human biases from entering the decision-making processes of our automated systems?
 6. **Security:** How do we prevent adversaries from taking over, using, or manipulating our automated systems?
 7. **Unintended consequences:** How do we keep automated systems from making and enacting decisions that will harm people in unintended ways?
 8. **Control:** How do we stay in control of automated systems that continually assess and change their environments?
 9. **Legal status:** What is the legal status of a decision-making automated system? Does it deserve any special treatment? Are there proper and improper ways to handle such automated systems?

Using AI for cybersecurity

- Modern cybersecurity often deals with vast quantities of data
- Using AI and machine learning techniques can make it easier to deal with data:
 - Gathering and analyzing data
 - Summarizing data
 - Extracting insights and patterns from data
 - Automating security management models based on dynamic data analysis
 - Targeting security alerts based on patterns of behavior, minimizing false security alerts
 - Optimizing resources to meet security goals
- Both defenders and attackers can use these tools

Kinds of AI that can be useful

- Supervised learning
 - Use labeled sets of data (like "normal use" and "attacks") to train
 - Then, the tool can identify new data
 - Applications: Detecting DoS and phishing
- Unsupervised learning
 - Uses raw, unstructured data
 - The tool looks for patterns
 - Applications: Detecting deviation from normal use patterns
- Natural language processing (NLP):
 - Interprets normal human language (which LLMs do much better than previous forms of NLP)
 - Syntactic analysis: break sentences down into subjects, verbs, etc.
 - Semantic analysis: extract context and meaning from text
 - Sentiment analysis: detect emotions, positivity, or negativity
 - Application: Detecting malicious domain names and phishing emails
- **Bias is always a risk with AI because training data contains bias**

Dangers of AI

- AI can be used for malicious purposes
 - Make attacks easier to execute
 - Increase volumes of attacks
 - Adapt existing attacks to new contexts
- As AI is used for more applications, the dangers grow
 - AI might have bad training data or it could be manipulated
 - It's really hard to know whether the problem is intentional or not
 - Modern AI decisions are often not explainable
- Who is liable when AI does bad things?

Upcoming

Next time...

- Review up to Exam 2
- Anu Regmi presents

Reminders

- **Finish Assignment 5**
 - **Due tonight by midnight!**
- Office hours canceled tomorrow
- Work on Project 3
 - Try to attack the other projects
- Study for final: 12:30 - 2:30 p.m., Wednesday, 12/10/2025