Today we'll look at layer 8+ issues through the lens of standoff biometry. In particular, we'll focus on face detection and recognition from video

Small mention of interesting things

- Summarize some part of what Dan Geer said in Identity as Privacy

Small mention of interesting things

- Signatures and Diffie
- “What we have is a very weak biometric mechanism [i.e., the handwritten signature --DS] that works quite well in practice---except that it’s choked by procedural rules that vary by country and by application.” (p. 263)
  - https://www.washingtonpost.com/politics/does-your-vote-count-in-florida-it-might-depend-on-your-signature/2018/11/14/7625251a-e762-11e8-a939-9469f1166f9d_story.html

The Seven-Layer OSI Model classifies protocols by function

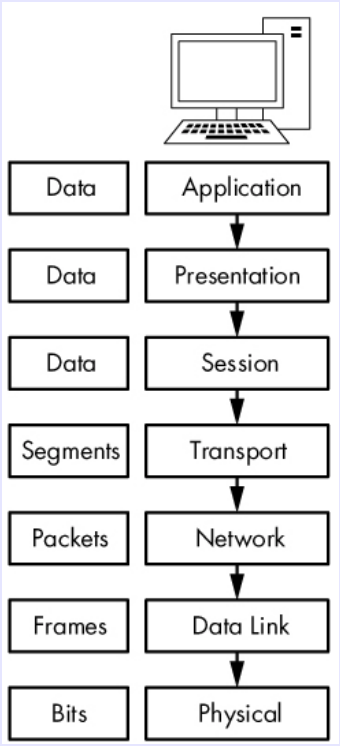

Application layer (layer 7) The topmost layer of the OSI model provides a means for users to access network resources. This is the only layer typically seen by end users, as it provides the interface that is the base for all of their network activities.

Presentation layer (layer 6) This layer transforms the data it receives into a format that can be read by the application layer. The data encoding and decoding done here depends on the application layer protocol that is sending or receiving the data. The presentation layer also handles several forms of encryption and decryption used to secure data.

Session layer (layer 5) This layer manages the dialogue, or session, between two computers. It establishes, manages, and terminates this connection among all communicating devices. The session layer is also responsible for establishing whether a connection is duplex (two-way) or half-duplex (one-way) and for gracefully closing a connection between hosts rather than dropping it abruptly.

Transport layer (layer 4) The primary purpose of the transport layer is to provide reliable data transport services to lower layers. Through flow control, segmentation/desegmentation, and error control, the transport layer makes sure data gets from point to point error-free. Because ensuring reliable data transportation can be extremely cumbersome, the OSI model devotes an entire layer to it. The transport layer utilizes both connection-oriented and connectionless protocols. Certain firewalls and proxy servers operate at this layer.

Network layer (layer 3) This layer, one of the most complex of the OSI layers, is responsible for routing data between physical networks. It sees to the logical addressing of network hosts (for example, through an IP address). It also handles splitting data streams into smaller fragments and, in some cases, error detection. Routers operate at this layer.

Data link layer (layer 2) This layer provides a means of transporting data across a physical network. Its primary purpose is to provide an addressing scheme that can be used to identify physical devices (for example, MAC addresses). Bridges and switches are physical devices that operate at the data link layer.

Physical layer (layer 1) The layer at the bottom of the OSI model is the physical medium through which network data is transferred. This layer defines the physical and electrical nature of all hardware used, including voltages, hubs, network adapters, repeaters, and cabling specifications. The physical layer establishes and terminates connections, provides a means of sharing communication resources, and converts signals from digital to analog and vice versa.

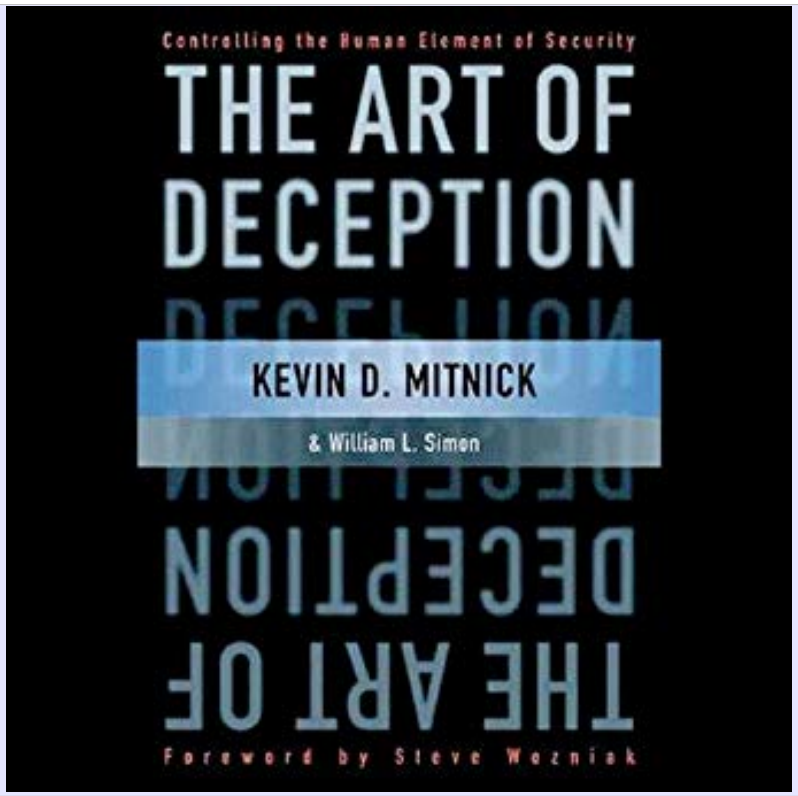

- one way to attack privacy is to attack the people who are involved in protecting it.
- “amateurs target systems, professionals target people”

- Examples
  - Big one: Social Engineering
  - The Art of Deception (story time: The embarassed security guard. From Tucson!)
  - Remember: this is told from a security perspective. Think about “virtual attrition” WRT trust as a privacy issue. What is needed to trust with people like this guy around?
- Contract law
- Compulsion

- Incentives, costs, benefits, ..

Example Technologies at layer 8

- PGP web of trust
- Ceremonies for verifying keys in instant messengers
- shoulder surfing countermeasures
- reputation systems
- **biometrics countermeasures**

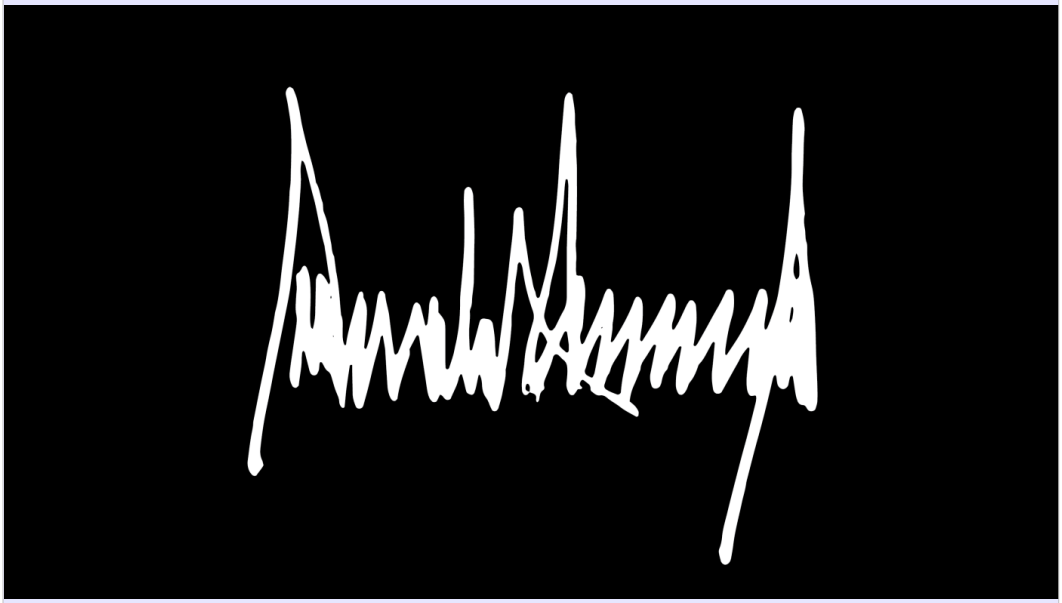

Biometrics before computers

- handwritten signatures
- facial features
- fingerprints
- hand geometry
- seals
- pronunciation

- ... mostly not rigorous; imprecise

Biometrics with computers

"Technical progress in image acquisition guarantees observability now; technical progress in standoff biometrics guarantees identifiability real soon now. Venture capitalists regularly hear new ideas in standoff biometry, and Moore's law is its friend ...

Geer is using 'observability' and 'identifiability' advisedly here: what do they mean (think back to the anonymity reading by Danezis and Diaz)

“We will soon live in a society where identity is not an assertion (“Call me Dan”) but rather an observable (“Sensors say that’s Dan”). Your breath, the microwave emissions of your beating heart, the idiosyncratic anomalies of how you type, talk, and walk say who you are.”

Dan Geer, Identity as Privacy, in 2013

“Moore’s law is its friend ... . We will soon live in a society where identity is not an assertion (“Call me Dan”) but rather an observable (“Sensors say that’s Dan”).”

What does this mean, and why might it be important?

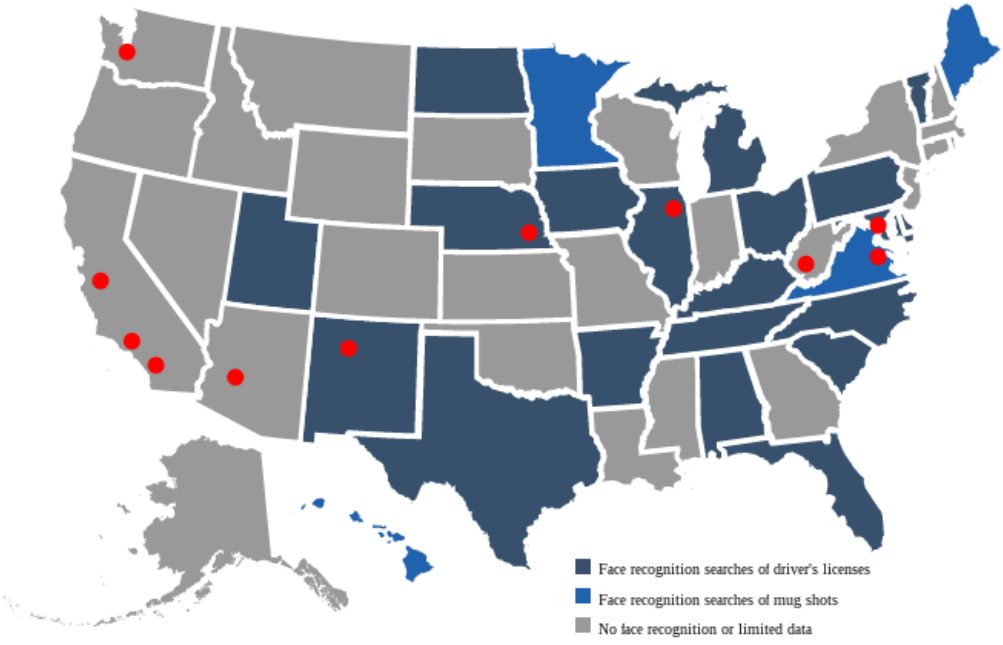

Figure 1: Law Enforcement Face Recognition Use By Jurisdiction

Click on state or city for more information

- The Police Department for the city of San Diego uses facial-recognition technology to attempt real-time identification of every face that passes in front of a camera
- The FBI's Next-Generation Identification (NGI) searches 16 states' driver's license databases, American passport photos, and photos from visa applications, amounting to nearly 412 million records

  https://www.perpetuallineup.org
-

FindfFace was a consumer app doing something similar to NGI, in a way: it searched the russian OSN VKontakt by your face, in real time. It caused a lot of hand-wringing

“In identity verification, the subject looks straight at the camera under controlled lighting conditions, and their face is compared with the one on file. A related but harder problem is found in forensics, where we may be trying to establish whether a suspect’s face fits a low-quality recording on a security video. The hardest of all is surveillance, where the goal may be to scan a moving crowd of people at an airport and try to pick out anyone who is on a list of perhaps a few hundred known suspects.” (Biometrics, 265)

Google's photo service is a good example of the commercial opportunities afforded by good facial recognition technology: it organizes photo albums by recognizing who is photographed. The chief benefit to Google derives from linking this biometric data to data from its more than 60 services: knowing what a customer is doing, and with whom they are doing it, is important to predicting and shaping future purchasing behavior. For an empirical study of the privacy harms of Google sharing across its services, see James C. Cooper, "Anonymity, Autonomy, and the Collection of Personal Data: Measuring the Privacy Impact of Google's 2012 Privacy Policy Change," 2017, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2909148.

All this for 10% decrement in performance?
https://link.springer.com/chapter/10.1007/978-3-319-41501-7_38

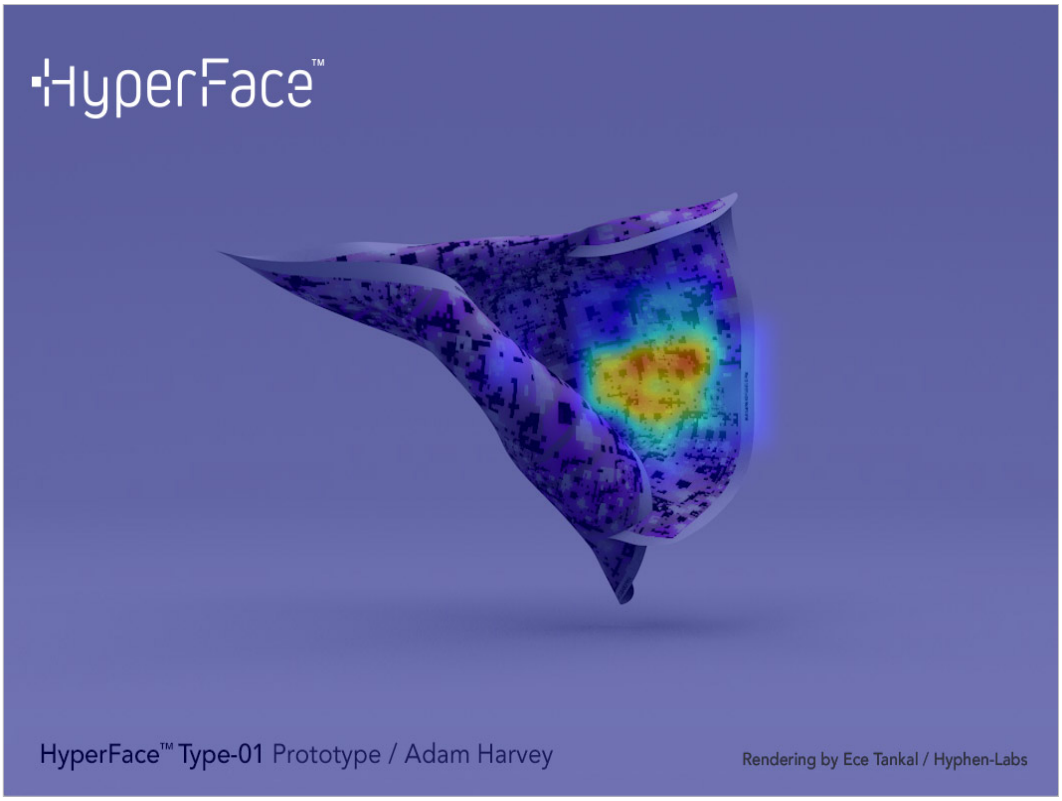
HyperFace™
HyperFace™ Type-01 Prototype / Adam Harvey
Rendering by Ece Tankal / Hyphen-Labs


TIME
Tech
World
Business
Tech
Health
Science
Entertainment
Leery of Facial Recognition? These Glasses Might Help
BBC
NEWS TECHNOLOGY
'Privacy visor blocks facial recognition software'
NBC NEWS TECHNOLOGY
LED-powered 'privacy visor' thwarts facial recognition
>Click News
Losing Face
取り上げられたメディアは
TIME, BBC, NBC, ABC,
NY Times, Spiegel,
ACM Tech Newsなど
海外300以上。
海外でも待ち望まれています！

Figure 1: Can you see the face in each of these images? ***Facebook's automatic face detector can detect and localize all six faces shown above****, even when we try to hide the face by (A) adding occluding noise, (B) adding distortion, (C) blurring the face region, or (D) altering image lighting. Deliberate countermeasures such as (E) wearing a "privacy visor" with infrared LEDs [24] or (F) wearing "Dazzle"-style makeup [9] are not always effective—sometimes Facebook can see through these disguises too. Facebook may not be able to* recognize *these faces, but if Facebook can* detect *them, it may prompt friends to tag these hard faces and reveal their identity.*

Wilbur et al., "Can we still avoid automatic face detection?"

Notice the remark about Facebook prompting users. This is akin to Anderson's story about handwriting recognition in a mixed-initiative architecture.

Figure 8: All 30 missed faces in the COFW training set.

Tricking state of the art face detection

The color of each keypoint in Fig. 8 shows face detection probability among images where that keypoint is occluded. From this figure, we can see that Facebook's face detector will only find 60% of faces that have the nose tip occluded. This is likely because objects that cover the nose tip often occlude many more points as well. Further, covering the mouth region lowers detection probability more than covering the eye regions. These results, and those in 5.1, seem to indicate that covering the nose and mouth regions may be a reasonable way to hide from some face detectors, but more exploration is necessary.

Disguised Face Identification

Figure 1: The figure (left) illustrates the 14 facial key-points annotated for both the introduced datasets. The description of the facial points is as: Eyes region (cyan): P1- left eyebrow outer corner, P2- left eyebrow inner corner, P3- right eyebrow inner corner, P4- right eyebrow outer corner, P5- left eye outer corner, P6- left eye center, P7- left eye inner corner, P8- right eye inner corner, P9- right eye center, P10- right eye outer corner; Nose region (yellow): P11- nose; Lip region (green) P12- lip left corner, P13- lip centre, P14- lip right corner. Few key points have been shown on the right.

Figure 2: The illustration shows samples images with different disguises from both the Simple and Complex face disguise (FG) datasets. As seen from the image, the samples from the complex background dataset have a relatively complicated background as opposed to the simple dataset.

data. However, such datasets are not available (small:

On arxiv since late August, to be published at ICCVW 2017. It is unclear how much performance can be improved.

…Researchers with the University of Cambridge in the UK, the National Institute of Technology and Indian Institute of Science have developed a deep learning approach to solving the problem of ‘Disguised Facial Identification’, aka, how to identify people at protests who have covered their faces.

…The approach relies on the creation of two new datasets, both of which contain 2,000 images each, and which label the 14 key points essential for facial identification on each person’s face. A simple variant of the dataset has simple backgrounds, while the harder version has noisy, more complex backgrounds. Both datasets appear to consist of portrait-style photographs, and feature male and female subjects aged between 18 and 30, wearing a variety of disguises, including: ‘(i) sun-glasses (ii) cap/hat (iii) scarf (iv) beard (v) glasses and cap (vi) glasses and scarf (vii) glasses and beard (viii) cap and scarf (ix) cap and beard (x) cap, glasses, and scarf.’

…The results: The resulting Disguised Face Identification (DFI) framework can identify a person wearing a cap, face-covering scarf, and glasses, about 55% of the time in the simple dataset, and 43% of the time in the complex one. So don’t put down that protest wear just yet – the technology has a ways to go. In the long run, perhaps this will increase the likelihood of people using rigid masks – like the V for Vendetta one adopted by anonymous – instead of soft ones like scarves, balaclavas, and so on. I also think that the datasets and underlying machine learning techniques will need to get dramatically better and larger for this sort of approach to be tractable and practical – especially when dealing with diverse groups of protesters.

Disguised Face Identification

55% accuracy on easier dataset

Figure 1: The figure (left) illustrates the 14 facial key-points annotated for both the introduced datasets. The description of the facial points is as: Eyes region (cyan): P1- left eyebrow outer corner, P2- left eyebrow inner corner, P3- right eyebrow inner corner, P4- right eyebrow outer corner, P5- left eye outer corner, P6- left eye center, P7- left eye inner corner, P8- right eye inner corner, P9- right eye center, P10- right eye outer corner; Nose region (yellow): P11- nose; Lip region (green) P12- lip left corner, P13- lip centre, P14- lip right corner. Few key points have been shown on the right.

Figure 2: The illustration shows samples images with different disguises from both the Simple and Complex face disguise (FG) datasets. As seen from the image, the samples from the complex background dataset have a relatively complicated background as opposed to the simple dataset.

data. However, such datasets are not available (small:

On arxiv since late August, to be published at ICCVW 2017. It is unclear how much performance can be improved.

…Researchers with the University of Cambridge in the UK, the National Institute of Technology and Indian Institute of Science have developed a deep learning approach to solving the problem of ‘Disguised Facial Identification’, aka, how to identify people at protests who have covered their faces.

…The approach relies on the creation of two new datasets, both of which contain 2,000 images each, and which label the 14 key points essential for facial identification on each person’s face. A simple variant of the dataset has simple backgrounds, while the harder version has noisy, more complex backgrounds. Both datasets appear to consist of portrait-style photographs, and feature male and female subjects aged between 18 and 30, wearing a variety of disguises, including: ‘(i) sun-glasses (ii) cap/hat (iii) scarf (iv) beard (v) glasses and cap (vi) glasses and scarf (vii) glasses and beard (viii) cap and scarf (ix) cap and beard (x) cap, glasses, and scarf.’

…The results: The resulting Disguised Face Identification (DFI) framework can identify a person wearing a cap, face-covering scarf, and glasses, about 55% of the time in the simple dataset, and 43% of the time in the complex one. So don’t put down that protest wear just yet – the technology has a ways to go. In the long run, perhaps this will increase the likelihood of people using rigid masks – like the V for Vendetta one adopted by anonymous – instead of soft ones like scarves, balaclavas, and so on. I also think that the datasets and underlying machine learning techniques will need to get dramatically better and larger for this sort of approach to be tractable and practical – especially when dealing with diverse groups of protesters.

Obfuscation

- “If privacy both as impossible-to-observe and impossible-to-identify is dead, then what might be an alternative? If you’re an optimist or an apparatchik, your answer will tend toward rules of procedure administered by a government you trust or control. If you’re a pessimist or a hacker/maker, your answer will tend toward the operational, and your definition of a state of privacy will be mine: the effective capacity to misrepresent yourself.”

  Geer, ‘Identity as Privacy’

Glamouflage

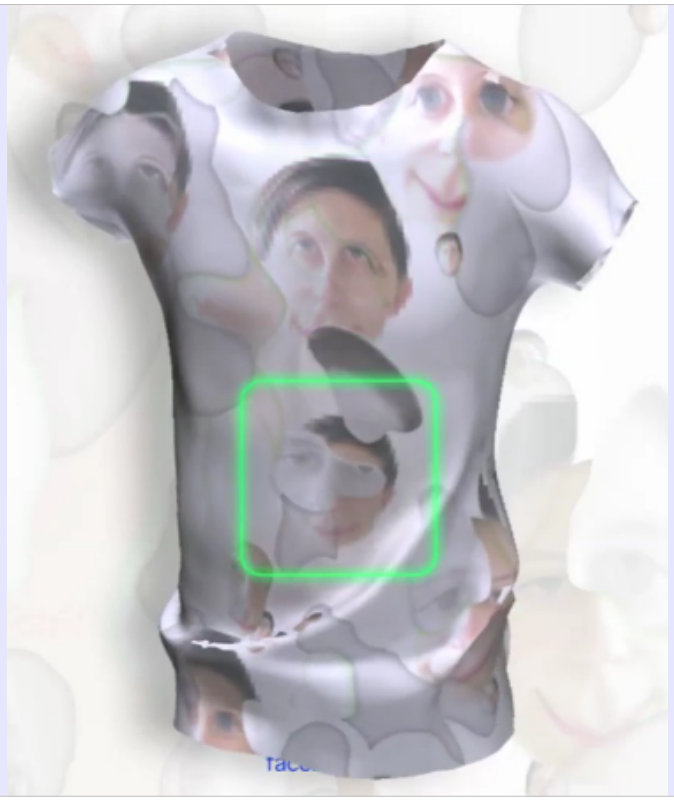

Another approach: make more video collectors trustworthy

- Personalized Privacy Assistants
  - A registry of IoT devices that respects users requests about how their data will be handled
  - https://youtu.be/j5btHZKgwaI
- A good idea for those willing to participate in the registry (which is probably a lot of device owners, and participation could be made compulsory in some contexts)
- No good if the device owner spurns the registry, or if they try to use it to mislead

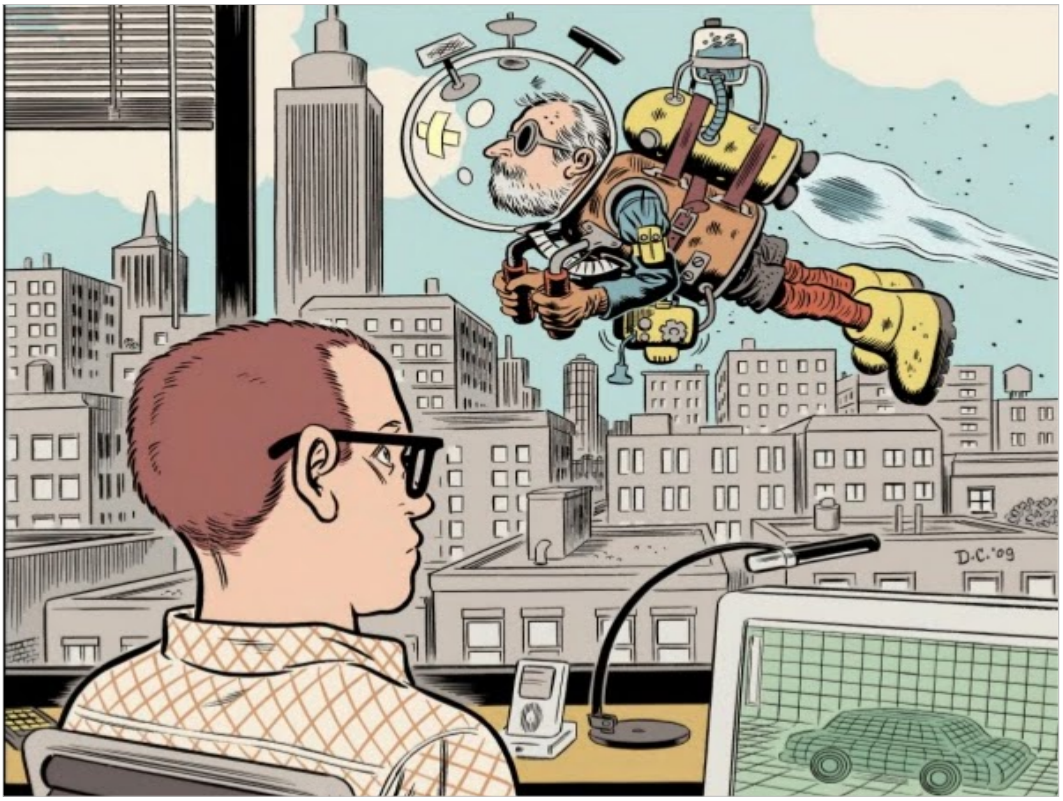
D.C.'09

Small mention of interesting things

- In the VM, the command line uses ksh with vim bindings. (Notice you'll need bash for conda, if you use it).
- demonstration (petlib)