

**Lunch – brief statement (2 ½ minutes) – what will happen in the near future –  
frame up 5 years and 10 years**

Forty years ago Alan Kay said “Technology is anything that wasn't around when you were born.” Similarly, whatever the world of the future is like, that will be the normal for the people born into it.

A sizable part of the population already lives in a mixed reality world. On a very deep level I can barely remember what life was like before the internet. People are blending the virtual and the real, and work and play, without thinking.

The entertainment technology community is part of a larger technology and business ecosystem that is racing to create and perfect, multisensory, interactive, immersive experiences where the virtual elements are indistinguishable from real elements. The tools to create and distribute believable yet false or distorted information in any media are available to anyone on the planet with an internet connection who is willing to put in the time and effort to learn how to do it. And the barriers to using the tools are dropping as the tools and their user interfaces get better.

We now live in a post-evidence world where anything can be faked or distorted. It can take an expert to detect when a digital asset has been altered or falsified.

Researchers and companies are also well on their way to developing the tools to create individualized psychological profiles of every one of us. Built using the backchannel data that we agree to give away in the mostly unread click licenses, as well as data gathered by devices on the people and structures around us, those

profiles are increasingly effective at delivering personalized messages that are emotionally engaging. Messages that resonate with us can press heavily on what we consider to be our rational beliefs, regardless of their ‘truthiness.’

In our discussion of ‘fake news,’ it may make more sense to discuss a means to identify, elevate, and reward the bounded set of data, information, and knowledge that we can verify to be true, reliable, and undistorted, rather than to try to detect and react to an unbounded flow of the false, the distorted, and the fake.

## **Main Session**

### **1) (1 minute) Intro**

I’m Philip Lelyveld. I run the Immersive Media Initiative at the Entertainment Technology Center within the University of Southern California’s School of Cinematic Arts. Prior to joining the Center I spent 10 years leading initiatives related to the future of entertainment for The Walt Disney Company’s corporate strategic planning group. I’m here to discuss my current thoughts on how entertainment technology and the language of storytelling fit into the fake news landscape, and visa versa.

### **2) (5 min) Global trends linked to fake news– Define fake news. Go into the history and current state of the industry.**

Fake News is 1) verifiably false or 2) partially reported information taken out of its original context. I'd like to mention two related terms; disinformation and propaganda. Law Professor Peter Kruger says disinformation is about changing the actual facts to fictional facts; propaganda is about getting you to ignore facts and make emotionally-biased conclusions. Disinformation is actual fake news; propaganda is spin.

I'm not going to discuss propaganda, but I am going to talk about the power of emotion as I discuss how we got here, from the storytelling perspective, in five points.

Point 1 – Historically the creation and distribution tools for storytelling have been developed and refined for the professional market, then simplified and released to the consumer market. Think still cameras and movie cameras, camcorders, but also audio and video editing tools, filters and effects generators, etc. While originally developed to meet the special needs of a niche community, these tools have become general purpose tools useful for all sorts of unintended purposes; some constructive and some destructive to society. In addition to this, with the birth of the Maker Movement, we also have tools being developed at the grass roots level by citizen-inventors and creators that are adopted by the professional community. Innovation is alive and unconstrained.

Point 2 – At the birth of the internet Mitch Kapor said “Architecture is politics.” It was architecturally designed with a political intent; to be immune to any attempt by government or other forces to control it. An intended element of this is online anonymity. The founders wanted to protect free speech and political protest. As a

result, on the internet nobody knows that you are a dog, or a bot, or one person masquerading as a crowd. Associated with this, global online culture as reflected by the Creative Commons and Open Source movements strongly object to any efforts to embed Intellectual Property controls into the infrastructure of the web. The unconstrained flow of data, including valuable IP, helped steer innovation towards the use and monetization of this data. Mashup Culture delivered yet another ecosystem of creative tools and resources to both the professional community and the global public.

Point 3 – Let’s talk about business models and the attention economy. There are only three business models; I pay, you pay, they pay. I’m not paying for you. You aren’t paying when you have been conditioned to expect things for free. They, the advertisers, are willing to pay in exchange for your willingness to give away a rapidly growing amount of your personal data. But they want your attention. A good narrative, whether it is an ad, a fictional story, or news, creates ‘sticky’ memories when it engages you emotionally. Extremes in character, plot, and theme are more memorable and sticky, regardless of their rationality, reasonableness, or truthfulness.

Point 4 – Quality Control has been discounted, or at least redefined, as the narratives we experience have gone real-time and we are presented with a continuous flood of new experiences to choose from. In the 24 hour news cycle we barely notice corrections.

Point 5 – We increasingly live in a constructed reality that blends the virtual and the real. This isn’t new. The telephone is a teleportation device for your voice. It

delivers a disembodied virtual 'you' into someone else's reality, and visa versa. Given enough stimuli, enough multisensory input, the virtual can become indistinguishable from the real. As the tools for creating alternate versions of the world and the people in it become widely available, we are left with Morpheus's line from The Matrix; "Real is simply electrical signals interpreted by your brain."

The open question that I suggest we focus on is, in a world where everything can be faked, and where everyone has the tools to create and distribute 'fake,' how do we identify and protect the real?

**3e) (6 min) Technologies and concepts that change the way people interact with each other and get their information. Emerging tools and their implications.**

Good storytelling draws you into a constructed world. The goal of entertainment technology and the language of storytelling is to construct an experience that will emotionally engage you. The best constructed experiences work to hide the limitations of their techniques so that the experience comes across as believably as possible.

I'm sure you've all seen videos where the voice and image of world a leader or celebrity is manipulated to make it appear that they are making outrageous statements, or scenes where completely artificially constructed life forms interact seamlessly with live actors and the scene around them. Those are the easy lies, the transparent fakery in service of story that we recognize, accept, and enjoy.

Let me go a little deeper.

For over a decade tools like Dramatica have been used by scriptwriters to assess how well they have developed character, plot, and theme.

Market researchers use eye tracking, brainwave monitor, heart rate, and other bio-response data to design more effective ads. The ads are factually correct, but you must also recognize that the sound, image, pacing, and other artistic language elements are studiously designed to illicit a desired emotional impact. The data they use can also be gathered by properly equipped virtual and augmented reality headsets and body computing devices like smart watches and fitness bands.

Affectiva, a company spun out of MIT, is one of a number of companies working to marry computer vision and deep learning to determine emotions from a person's nonverbal cues and facial expressions. The videogame company Flying Mollusk used Affectiva's software to produce a psychological thriller game whose difficulty changes with the player's level of fear.

Marketers hope these techniques will help them understand emotions well enough to tell when a person is open to a transaction. Imagine the potential uses, both good and bad, for all of the personal data you are releasing through the click license in exchange for the promise of a more personalized experience.

Facebook announced last summer that they were working on a brain-computer interface for typing and other interactions. Imagine other uses for the backchannel data; data that is necessarily gathering to make that work properly. Giving up this privacy is already now included in some click licenses. Do we need to start thinking about a right to mental privacy?

Virtual Reality is often called an empathy engine. Because there is no frame, like the rectangle of a movie or TV screen, separating you from the sights and sounds of the experience, you are more likely to be emotionally and viscerally engaged. Fable Studio has developed an AI character named Lucy. Lucy's story conceit is that you are Lucy's imaginary friend. She learns about you, gives you tasks, and stays with you, cross-platform, as you move through your day. Once you have bonded with her, imagine Lucy's power as a virtual trusted friend to influence you. Fable Studio's primary business is storytelling and world building, not AI research. But combine story with AI efforts to profile and predict individual human behavior, and you have a powerful resource being developed with, at least until recently, little discussion of the ethics, risks, and unintended consequences. I am aware of a number of 'ethics in technology' initiatives around the globe, and last month there were news stories about Harvard, MIT, and other schools developing courses on ethics related to computer science and artificial intelligence, but I have not seen any of them gain traction yet. Should we have the right to understand the framework for an AI's personalized decision-making, including how it is filtering data?

To sum up, tools and techniques to engage our attention, and to understand and influence us, individually, on both a conscious and subconscious level, are being invented and refined in offices, test labs, and basements around the world.

Expect to see them as either standalone products and services, or embedded into other products and services – including entertainment experiences and advertising - as they prove their usefulness.

#### **4) (6 min) Suggestions on how to address fake news**

I start from the position of immersive experiences, which include virtual reality, augmented reality, and mixed reality. Those immersive experiences can include enhanced explorations of real places, people, and events around the globe, imaginary places, people, and events, and any type of simulation in-between. Within those immersive experiences you can have social interactions with avatars that are telepresences of real humans (remember when Mark Zuckerberg sent his avatar on a boat tour of hurricane-ravaged Puerto Rico), avatars driven purely by code, or actual human beings as 'avatars' or actors within an immersive experience (Pokemon developer Flint Dille says the best designed in-game avatars are other human beings because they behave very realistically).

These immersive experiences are going to be informed by the internet of things which will feed data either directly into the experience or into the code that is creating the experience.

They will be informed by AI that will create the most realistic elements of the experience for you and will respond in contextually logical ways to your actions.

No new entertainment medium ever replaces an old one. People just rebalance the entertainment mix that they seek out. So we will add immersive experiences to text on physical paper and digital platforms; personal passive and interactive audiovisual experiences on TV, phones, tablets, laptops, game consoles, and other devices; and group and virtual group audiovisual experiences such as movie theatres, plays, opera, concerts, e-sports and sports events in both physical and virtual spaces.



All of these are channels that can be used to delivery data (raw elements), information (meaning), and knowledge (context).

For the most part these channels will deliver anything. The channels don't check for errors or distortions or lies.

There is a continuum between the absolutely true and undistorted, and the absolutely false and distorted. It is straightforward to check whether a sensor on the Internet of Things is sending true data or if it is malfunctioning. It is not nearly as straightforward to tell if a statement of opinion has been sampled and used in a manner that changes its original meaning.

In situations where there is uncertainty it is much easier and less risky to prove that something is not true than to prove something is not false.

So rather than try to stop fake news, I propose that we focus on mechanisms to identify, protect, and elevate what is most probably true and undistorted. To recognize and reward sources for being trustworthy.

One approach would be to tie all posts back to their online identities.

Fundamental to the proposal is a mechanism that securely develops a reputation profile that is securely attached to the source's online identity. The profile would then be securely linked to any and all posts linked to that online identity, should the online identity choose to do so.

We can look to cryptography and the blockchain ledger infrastructure for ideas on how to do this.

By doing that, anyone viewing the posted material would be able to track back to the identity and make an informed determination of whether both the material and the source are credible or not credible, based on their own personal criteria. This approach won't directly take a person out of their personal echo chamber of ideas, but it could give them a sense of the trustworthiness of the sources.

A key aspect of this structure is that first time posters, including fake identities and bots, would have low reliability ratings, and their contribution could be discounted accordingly. Also, contributions without a tie back to a reputation profile would be discounted and flagged as unverified or therefore of unknown reliability.

MetaCert (MetaCert.com) is one company that has implemented a variant of this approach, and released it as a browser add-on. They aggregate data about news sites from trusted fact checkers (ex. Snopes, Politifact). They then crawl the web, find, and label social media accounts owned by those news sites. MetaCert is one company tracking a very limited number of news sites and their posts. I am suggesting expanding that approach to all online media from all sources.

The approach could be developed voluntary, allowing for self-regulation on an industry and societal level. A small group of key infrastructure players and influencers could come together and, assuming their intentions are aligned, they could create a beta version, trial and evolve it, and seed the infrastructure with it. The player's involvement could even be useful for their organization's messaging. They would be working cooperatively to create a social good.

The approach won't cover all situations or completely solve the problem of lies and distortions masquerading as truth. But if it yields results that the market finds appealing, it could redefine the boundaries of trust within an open market.

**5) (5 min) Don't try to assess each EC initiative. Give an opinion on direction from my perspective.** (Tie it back to the IP industry's efforts to monetize IP and combat piracy from Napster to today.)

Innovation is never stifled, it is only redirected by changes in the ecosystem that it operates within. Those changes can include regulation, the bottom-line imperatives of business, and societal pressure.

The authors of 'fake news' are innovators. They will innovate around any technical, legal, or social obstacle you put in place in order to reach and build their audience. As you put up more obstacles they will find more sophisticated ways to bypass them. If you want to counter their effect (and I say counter their effect because stopping them would require fundamental changes to the architecture and philosophical foundation of the internet), then you must understand their tools and techniques and out-play them. You need to identify their drivers and develop counter drivers.

On the technical, tactical, and granular level of the fake news response, as I mentioned before, I suggest a technical approach that elevates what can be objectively graded as trusted or reliable, rather than taking on the gargantuan subjective task of grading for lies and distortions.

A gathering of key players whose interests in elevating reliable information sources are aligned could develop the definitions, screening criteria, policing and enforcement mechanisms, and infrastructure of an alpha or proof-of-concept release. If, through the usual fail fast and fix process, the approach proves out, then it would become the seed solution that others could contribute to and build on. This would be a 'seal of approval' system for good players.

Explaining why this approach was developed (to combat fake news), how it works, and how it protects democratic processes without diminishing online freedoms, could become part of public education, media literacy, and other social actions that raise awareness of both the problem of "fake news" and ongoing work to mitigate its impact.

The bigger question is; how do you identify, control or limit psychological manipulation by outsider agents on a societal scale. One possible approach is to involve psychologists, sociologists, ethicists, etc. in the development of code (an AI) that scours the web looking for evidence of social manipulation related to political issues.

The system would need ongoing human oversight to train and quality control the process. It would also need human intervention when the code returns material that falls on the boundaries of its judgement criteria. Regular human involvement in the training and evolution of the algorithm's judgement framework is key.

In his article *Build a Better Monster: Morality, Machine Learning, and Mass Surveillance*, author Maciej Ceglowski writes "...the algorithms have learned that users interested in politics respond more if they're provoked more, so they

provoke.” When provocation stops working, the innovators will try other strategies.

As the fake news generators change tactics, the detection and response systems need to be guided to respond to the new profile of the attack. Detecting social manipulation efforts by outsiders is an ongoing process, not a closed problem with a static deliverable.

The storytelling arts and sciences have regularly evolved to keep up with the rising sophistication and expectations of our audiences. Similarly, if you are going to detect and respond to the outsider innovators and storytellers propagating targeted psychological manipulation tools including fake news, then you have to put in place technical and social processes that can also evolve to stay one step ahead of them, and that contribute to helping the average citizen make informed decisions based on reliable data.

## **6) (1 min) Concluding remarks. Main take-away.**

Innovation is never stifled, it is only redirected by regulation, the bottom-line imperatives of business, and societal forces. A good story captures and holds your attention. A great story creates “sticky” memories by engaging your emotions. We are well on our way to developing and deploying the tools that will allow anyone to create and distribute personalized ‘sticky’ memories on a global scale. The tools don’t care whether the story is fact or fiction, true or false. It may make more sense to create a system that identifies, elevates, and rewards a bounded set of data, information, and knowledge that we can verify to be true,

reliable, and undistorted, than to try to detect and react to an unbounded flow of false, distorted, and fake content.

(filed under; Event / 180212a EU question thoughts)