

Dealing with the “Fake News” Problem

Submission

to the

Joint Standing Committee on Electoral Matters

17th January, 2019

Dr Carlo Kopp, Dr Kevin B. Korb, Dr Bruce I. Mills

Executive Summary

- The potential damage produced by “*fake news*” can be significant, and it should be treated as a genuine threat to the proper functioning of democratic societies;
- The problem of social media users irresponsibly propagating falsehoods, without knowing nor caring that these are falsehoods, and not making any effort to determine veracity or bias, needs to be dealt with;
- It does not take much “*fake news*”, under some circumstances, to wholly disrupt consensus forming in a group. If allowed to propagate unchecked, very little “*fake news*” can do a huge amount of damage;
- Regulatory or other mechanisms that might be introduced to disrupt, interdict or remove “*fake news*” from social media will confront serious challenges in robustly identifying what is or is not ‘*fake news*’;
- The “*fake news*” problem in mass media is similar in many ways to the quality assurance problems observed in manufacturing industries decades ago. Mandatory process based quality assurance standards for mass media could address much of this problem, as quality assurance problems in other sectors have been resolved;
- A model that should be explored is the exploitation of media organisations with track records of bias free and high integrity news reporting to provide fact checking services to social media organisations;
- The commonly proposed model of “*inoculation*”, in which users are taught to think critically and test evidence for veracity and bias, may be difficult and expensive to implement, and lazy social media users may simply not bother;
- The long-term solution is a robust educational system that rigorously teaches critical thinking from the very outset, starting with primary schools. Students will need to learn the ability to test source credibility, identify reasoning errors, and develop respect for tested fact over subjective opinion. This would solve much of the “*fake news*” problem as it would destroy most of the ‘market’ for “*fake news*”;
- A demerit point system uniformly applied across social media platforms, where users are penalized for habitually propagating “*fake news*” is an alternative to “*inoculation*”, but confronts the same problems as fact checking – who determines what is or is not “*fake news*”, and is this free of errors and bias?
- There may not exist any “*silver bullet*” solution to this problem, and comprehensive strategies similar to those required to defeat biological contagions may be needed. These will involve interdicting “*fake news*” production and distribution, and “*inoculation*” of social media users.
- There is genuine potential for abuse, and potential for free speech to be impaired, where social media are subjected to censorship mechanisms, so any regulatory models considered will have to be carefully designed and tested to prevent improper exploitation by any parties;

Introduction

Australia is not immune to the global pandemic of fake news, deception, cyber attacks, social media botnet attacks, and other digital platform hosted mischief, that has engulfed most Western democracies.

We are university computer science researchers with an interest in the problem of “fake news” and potential measures to deal with this pandemic in the digital media. We have recently published research that exposes both the sensitivity of a population to the effects of “fake news”, and the sensitivity of “fake news” to costs incurred, leading us to provide evidence to the UK Parliament’s DCMS Select Committee’s current *Inquiry on Disinformation and ‘Fake News’*.

We hope that that through shining some further light on this perverse problem we can usefully aid the Committee in forming its understanding of the problem Australia confronts.

In this submission we will aim to address a range of related problems and two of the stated *Terms of Reference*:

- ways to address the spread of deliberately false news online during elections; and
- measures to improve the media literacy of Australian voters.

To date Australia has not been subjected to the kind of focused high intensity multi-pronged state-sponsored information attack campaign that the UK was subjected to during the Brexit referendum, the US was subjected to during the 2016 presidential election, and Ukraine has been subjected to continuously since the covert invasion of Crimea by Russia. However, Australia has been exposed to large volumes of “fake news”, and this appears to be most frequently the result of Australian mass media and social media users redistributing deceptive “fake news” content that was being widely distributed in the global mass and social media.

There is no reason to believe that Australia would cope any better than the UK, US or even Ukraine were it subjected to the same style of state-sponsored information attack campaign. There is no evidence to support the proposition that Australian government and private sector cyber security is any better than that in the UK, US or even Ukraine, and no evidence that Australian social media users, public, and mass media are any less gullible than their peers in nations previously subjected to such attacks.

As we noted to the UK DCMS Select Committee last year: *“There can be little doubt that the dissemination of disinformation and misinformation through social media and mass media is a major risk to the functioning of parliamentary democracies. Existing legal frameworks, and arguably, widely accepted community standards of behavior in social media, are clearly not capable of coping with this problem. The impacts go well*

beyond what choices voters make during elections and referendums, as a population that is immersed in a deluge of misinformation and intentional disinformation will become confused or develop false beliefs about the world they live in. The term “alternative reality” could be applied. While the Orwellian model is frequently discussed in this context, the prospect of communities in a perpetual and universal state of confusion and chaos is a no less undesirable outcome. Universal false beliefs and frequently widespread confusion are a feature of totalitarian states past and present, and such outcomes are clearly incompatible with a functioning democratic system.”

The problems we now see with gullible social media users, public, and mass media falling for naïve or trivial deceptions will be immensely exacerbated as “deep fakes” become exploited for mischief and state sponsored propaganda.

We agree with Chesney and Citron’s observations: *“Harmful lies are nothing new. But the ability to distort reality has taken an exponential leap forward with “deep fake” technology. This capability makes it possible to create audio and video of real people saying and doing things they never said or did. Machine learning techniques are escalating the technology’s sophistication, making deep fakes ever more realistic and increasingly resistant to detection. Deep-fake technology has characteristics that enable rapid and widespread diffusion, putting it into the hands of both sophisticated and unsophisticated actors. While deep-fake technology will bring with it certain benefits, it also will introduce many harms. The marketplace of ideas already suffers from truth decay as our networked information environment interacts in toxic ways with our cognitive biases. Deep fakes will exacerbate this problem significantly. Individuals and businesses will face novel forms of exploitation, intimidation, and personal sabotage. The risks to our democracy and to national security are profound as well.”*

The deep fake problem was predicted by one the authors of this submission two decades ago, in an Information Warfare [conference paper](#). [The warning was simply ignored.](#)

Understanding the “Fake News” Problem

Most of the academic research in the area of “fake news” has been narrowly focused on specific problems, and empirical in nature.

Much excellent work has been done on studying the mechanisms of how “fake news” is propagated and diffused in social media and mass media. No less important empirical work has been done mining social media data and metadata, to establish the scale of the problem, identify perpetrators, and in some instances, explore the demographics of “fake news” redistribution.

Our research, published in late 2018, intentionally explored more fundamental mechanisms. We aimed to answer two key questions:

- A. How sensitive is consensus forming (i.e. cooperation) in a population to the presence of “*fake news*”?; and
- B. How do costs incurred by producers and propagators of “*fake news*” impact the effects of “*fake news*” in the population?

We detail the research and computer simulation methods employed, involving artificial agents who survive or die out depending on their behaviour and environment, in publications listed in the *References* at the end of this submission.

In the real world, the costs incurred by producers and propagators of “*fake news*” can be external costs (typically monetary), such as fines, penalties, exclusions, expenditures in creating and distributing fakes, or internal costs (i.e. social costs) to individuals, such as feelings of loss or embarrassment due to being ridiculed or shamed by peers.

An unexpected finding of our research was that even a remarkably small percentage of deceivers in a population, in our simulations less than 1%, could catastrophically disrupt cooperative behaviours in the simulated population.

In the extreme case of cost-free deceptions, where “*fake news*” producers in the simulated population are unhindered in any way, cooperative behaviours vanished altogether. Put differently, almost everybody ended up deceiving everybody else.

Only under conditions where the cost of deceptions was larger than zero, did cooperative behaviour survive in the population. Where costs to deceivers were very high, cooperative behaviours actually thrived at the expense of deceptive behaviours.

Another important finding was that for all simulations, the ability of deceiving players to survive depended very strongly on the cost of deceptions. If the cost of deception was high enough, deceivers simply could not survive in the population.

Applying these results to social media, two important messages are:

- increasing the cost of generating and repeating “*fake news*” is a priority for maintaining a healthy information environment;
- allowing even a small amount of “*fake news*” to flourish cost-free can be catastrophic.

These results are reflected in what has already been empirically observed in social media, despite the fact that our simulation involved a population of very simple software robots playing a trivial social game, and employing trivial deceptions.

What we found most remarkable is how closely such a simple simulation was able to capture behaviors seen in the ostensibly more complex real world of social media.

What are the practical implications of these research results for dealing with the real world of “*fake news*” distribution in social and mass media?

The first and arguably most important implication is that very little “*fake news*” is required to create a lot of mayhem in a population, and prevent consensus forming that is critical to public debates. The common misconception is that a lot of “*fake news*” is required to produce a lot of mayhem, but the opposite is demonstrably true. It doesn’t take a lot of poison to ruin a well.

What matters is that whether victims of “*fake news*” deceptions are confused, or end up believing in falsehoods, their ability to reach consensus will have been disrupted.

In the context of politics, voters who are confused and / or believe in falsehoods are very likely to be in a state of anxiety, and therefore not likely to make well considered and informed choices at the ballot box.

Our modeling of the “*fake news*” problem was very specifically focused on small populations of “*influencers*”, who actively debate issues and are mutually engaged in a discussion. Where influencers cannot come to a consensus, in turn followers (i.e. in a political or other wider community debate much of the general public) cannot align themselves to a consensus position.

This to date poorly understood problem is one of the major reasons why the global pandemic of “*fake news*” has been and continues to be so destructive to democratic societies.

This is also the reason why genuinely effective measures to deal with this problem *must be found and deployed at the earliest opportunity*.

The second result of broad interest from our simulations is that attaching a high cost to the production, but especially the distribution and re-distribution, of “*fake news*” may prove to be the most effective tool available to defeat the “*fake news*” pandemic.

A key consideration is how to best disrupt or interdict the mass distribution and re-distribution of ‘fake news’.

Over a decade ago Information Warfare research found that *proxy delivery* was a major “*force multiplier*” in the distribution of toxic propaganda – the case study was that mass media distributing violent media content produced by terrorists were directly acting as proxies for the terrorists producing the propaganda, whether they knew it or not. This problem was also directly identified in the UK DCMS *Disinformation and “fake news” inquiry*, and elsewhere.

It is a sad truth that social media users who share “*fake news*” are likewise acting as proxies for the producers of “*fake news*”, multiplying the reach and footprint of the “*fake news*”, especially if these social media users have large networks of followers.

In contemporary public discourse, such social media users are typically cast as victims of “*fake news*”, which they often are, but every time such social media users share

“fake news” they also become active participants in the *“fake news”* producer’s deception, or put bluntly *“proxy deceivers”*.

An open question is to what extent can social media users be expected to take responsibility for what they share, given that far too often they will have little or no understanding of what they are sharing, or its veracity?

This is without doubt a genuinely perverse problem.

Why is this so? Psychologists have identified the *Dunning-Kruger Effect*, whereby the ignorant become overly confident in their belief that they understand what they do not understand. The less such individuals understand, the more confident they become that they have mastery of the subject. The famous *“Darwin Awards”* provide a wealth of sorry examples.

Social media users who are the most poorly equipped to identify misinformation or disinformation are thus the most likely to confidently believe it and in turn confidently propagate it to other users as factual content.

The commonly propounded view that attaching a cost to the distribution of *“fake news”* in social media is straightforward has no basis in fact. The opposite is clearly true.

One approach to putting a cost to the distribution of *“fake news”* is the informal *“outing”* of habitual posters of *“fake news”*. This accords well with the evolutionary psychology of cheater detection.

Another approach often proposed is for social media organisations to be more proactive, and set up *“fake news”* detection units to identify and flag *“fake news”* posts accordingly.

This approach will be always be expensive in personnel costs, as at this time Artificial Intelligence is not up to this task, contra claims by Facebook and others.

But both of these approaches must confront the frequently challenging problem of determining exactly what is or is not *“fake news”*.

As many have noted repeatedly, unpalatable facts or truths are too often falsely labeled as *“fake news”*.

Popularity is not a reliable guide to the truth, despite the popularity of the *“argumentum ad populum”* logical fallacy!

This problem reflects the well documented propensity to seek conformity with a majority view, even if the majority view is obviously and completely false. The extensive literature on Janis’ *Groupthink* and the *Pluralistic Ignorance* problem is rich in examples. The widespread use of Bots in social media is specifically intended to exploit this behavior by creating the illusion of popularity where there is little or none.

Fact checking is often propounded as the panacea solution to “fake news” and related deceptions.

The reliability and objectivity of fact checkers can vary widely, and can sometimes be very poor.

Unfortunately, ground truths are frequently obscured by political, ideological or cultural bias. More frequently, fact checker limitations in understanding result in defective fact checking that simply adds to the confusion produced by the “fake news” being fact-checked.

The existence of multiple Internet fact checking websites that rate or rank the objectivity or biases in other fact checking websites speaks for itself!

There is some evidence that complaint handling channels in some social media platforms have been prone to bias, reflecting much the same problem as seen with fact checking.

The recent *New York Times* report by Max Fisher ([*Inside Facebook's Secret Rulebook for Global Political Speech - The New York Times*](#)) shows many of the difficulties involved when social media platform operators attempt censorship of user postings. We note that many of these problems are not new and were encountered over a decade before the advent of web based platforms, in user debates happening on *Usenet* email based discussion groups. The problem with trolls and their toxic behaviours was well known a quarter of a century ago.

A model that should be explored is the exploitation of media organisations with track records of bias free and high integrity news reporting to provide fact checking services to social media organisations. We note that this approach is not foolproof but may work better than social media organisations attempting to recruit and train novices as fact checkers. There is also the unavoidable challenge of who determines which media organisations have the best track records in bias free and high integrity reporting.

There are precedents for large scale audits of digital systems. The global “Y2K” effort to audit software for software bugs in date processing that would produce failures when the calendar rolled over from 1999 to 2000 was mostly successful. Many bugs were found, and only a handful of systems actually failed when “Y2K” arrived. The “Y2K” effort was successful due to the imperative of avoiding failures that might cost money or even lives. This imperative was widely acknowledged and accepted, globally. It does not appear that any such consensus as yet exists on the genuine dangers presented by the pandemic of “fake news” and other deceptions in the information sphere.

This lack of consensus is amply demonstrated by the recently published survey of “*Government Responses To Malicious Use Of Social Media*” by Bradshaw, Neudert and Howard, that demonstrates widely varying responses to this problem, refer below.

Observations on Measures to Improve the Media Literacy of Australian Voters

The reason why “*fake news*” exists, and existed prior to the advent of the Internet, is the simple reality that many people are gullible. A population of gullible Internet users creates opportunities for that gullibility to be exploited to advantage – whether commercial, ideological, or political. The cynical view would be simply that Barnum’s famous quip on suckers still applies.

The most interesting recent benchmark on the gullibility of younger audiences is the widely reported Stanford University study by Wineburg et al (2016), refer below. The results should not come as a surprise. In two of the three areas assessed no less than 80% of the students evaluated produced deficient assessments of the integrity of online content.

A more recent study by Guess et al (2019) found that older social media users were more prone to propagate “*fake news*” reports, although the study excluded a number of common categories that might be counted as “*fake news*” in other analyses.

Improving the media literacy of Internet users is often referred to as the “*inoculation*” model, in which users are taught to think critically and test evidence for veracity and bias.

Implementing “*inoculation*” may be difficult and expensive, and lazy social media users may simply not bother.

This is because Internet users have to first be convinced that becoming inoculated against “*fake news*” and other Internet deceptions matters and that their time and effort is being properly invested. Internet users not caring about the veracity or integrity of digital content they are redistributing is demonstrably a major part of the “*fake news*” problem.

Anecdotally, Internet users who have been *scammed*, *phished*, *spear-phished*, *hacked* or otherwise subjected to embarrassment or material losses appear more often cautious about believing what they find in their inbox or news feed.

Until an Internet user is convinced that it is worth their time and effort to learn, it is unlikely that any amount of expensive “*push*” advertising and promotion of “*inoculation*” will be taken seriously, let alone acted upon.

We can speculate that a common belief in users who have not been impacted is the idea that “*it cannot happen to me*”, or “*I am not important enough to be a target*”.

Yet experience in the UK, US, Germany, Baltic States and Ukraine indicates otherwise. So-called “*ordinary*” Internet or social media users have genuine value to an attacker, especially if the agenda is propaganda distribution and/or political manipulation.

If they are able to cast a vote or influence friends or family who vote, they have value as a target. If they share the deceptive content, their value increases in proportion to their sharing activity.

Any effort to advertise the importance of “*inoculation*” must first and foremost be convincing, and convey the message that every voter can be improperly exploited if they choose not to be careful.

Teaching Internet and social media users to think critically and test evidence will present its own challenges. As anybody with teaching experience in higher education will observe, teaching critical thinking and skepticism about the veracity of some stated factoid is often difficult. While some students are naturally inclined to think in this manner, a majority are not and have to be taught in an almost “*ab initio*” manner. The Stanford University study shows this very clearly.

It is worth observing that multiple industry surveys aimed at determining what employers want in tertiary graduates rank critical thinking skills very highly. One of the authors of this submission conducted such a survey in 2010, across Australian and overseas IT industry organisations, and discovered that all surveyed organisations wanted critical thinking skills, and many answered “yes” to the survey question whether such skills should be explicitly examinable in tertiary courses. Put bluntly, gullible employees are a liability for their employers.

An experimental effort currently under way in the Czech Republic is the “[FakeScope](#)” game, developed by students, in which players aim to determine what is or is not “*fake news*” in the game. This approach exploits competitive behavior and aims to engage the players so they become invested in catching the “*fake news*” producer.

Game based teaching approaches may be less effective with mature age Internet users and effort will need to be invested in finding strategies that are effective, and time efficient.

The long-term solution is of course a robust educational system that rigorously teaches critical thinking from the very outset, starting with primary schools. Students will need to learn the ability to test source credibility, identify reasoning errors, and develop respect for tested fact over subjective opinion. This would solve much of the “*fake news*” problem as it would destroy most of the ‘market’ for “*fake news*”.

As recent experience with benchmarking primary and secondary education learning outcomes has shown, this might require major changes in teaching strategies.

Recent observations in [The Australian by Dr Clio Cresswell at the University of Sydney](#), that many tertiary students are challenged in traditional reasoning methods, accord closely with the recent teaching experiences of the authors of this submission.

Cosmetic changes to primary and secondary education will not produce the intended changes in learning outcomes.

Improving the quality of education and explicitly including critical thinking skills as a assessed learning outcome does not address the near term problem we observe, and will require increased investment in education regardless of how it is implemented.

Considering Measures to Improve the “Fake News” Literacy of Australian Media

The advent of the Internet has produced serious disruptions in the mass media as a sector, globally, and Australia is no exception.

The traditional mass media model was one in which media organisations were mostly vertically integrated, performing in-house content production and content distribution to content consumers. In the globalized digital world, most media organisations in Australia perform very little substantial media content production, being focused primarily on reusing third party content and packaging it for their target audience. This model is very common across the English-speaking world, with a small number of large US and UK media organisations producing a large fraction of original content, that is syndicated out across smaller media platforms, or often simply copied, emulated or plagiarized by the latter.

The result of this model is that fewer and fewer journalists over time develop the traditional experience we still expect of journalists, in conducting research, testing sources for veracity and integrity, and ensuring that what is written correctly reflects facts. Far too often the skills exercised are “copy and paste” from other sources, with little or no effort invested in determining whether the story is true or accurate. This is the reason why, globally, mass media are frequently complicit in the distribution of “fake news”, and foreign state sponsored propaganda.

Some practices in some parts of the media exacerbate this problem. One is the notion that articles should be “balanced”, to present both sides of a story. When one side of the story is “fake news” or state sponsored propaganda, the journalist and media organization become complicit in the deception involved, unless they explicitly identify the deceptive nature of the content in question. Far too often the journalist or media organisation treat the “fake news” or state sponsored propaganda no differently than a legitimate source, which to many less informed readers legitimizes or validates the fake as correct and true.

This problem is compounded by partisanship, or to employ the recently coined US term, “hyper-partisanship” in some parts of the media. Partisanship in media often reflects the biases and expectations of the target audience of that media organization. Commercial imperatives will drive media organisations to maximize audience sizes, and this creates pressure to appeal to the prevalent biases, especially political biases, in that audience.

Unfortunately, where there is a partisan agenda, accuracy can often be impaired to produce a narrative that is coherent with the partisan agenda, rather than hard facts.

A recent critique by [Hewitt in the Washington Post](#), dealing with partisan reporting, presents a good case study, which is sufficiently distant from the Australian debate on media partisanship to avoid our domestic partisan biases as an audience.

Hewitt argues thus: *“The second, subsidiary message was for the media: The quick, and often negative, interpretation of every Trump statement has real-world consequences in the conduct of U.S. foreign policy. Regarding coverage of Trump’s Syria policy, a senior administration official told reporters accompanying Bolton on his trip that media coverage has exacerbated the problem of understanding the United States’ Syria policy. Hair-trigger “hot takes” can be totally wrong but still widely read and believed. Turkish President Recep Tayyip Erdogan follows the American media; he declined to meet with Bolton. But we have no way of knowing what message he was acting on.”*

This instance shows that partisanship at the expense of careful and thoughtful analysis produced what was in effect *“fake news”* that was most likely swallowed by a foreign leader, who himself was eager to find reasons to criticize the US president.

Another interesting example, involving state sponsored propaganda rather than partisan bias induced filtering, has been the adoption of Russian propaganda tropes by many Western media organisations reporting on the conflict between Russia and Ukraine. Most Western media organisations, and many Ukrainian media organisations, continue to label the Russian mercenary and regular military forces deployed in Eastern Ukraine as *“rebels”* or *“separatists”*, reflecting Russian propaganda disseminated widely in global media during the Russian invasion of Ukraine in 2014. That five years later even *“reputable”* media organisations in the West continue to use deceptive labels coined by Russia’s very skilled propagandists in 2014 illustrates a fundamental and pervasive breakdown in media professionalism.

Another related example involved *New York Times* journalists Broad and Sanger [speculating that North Korea acquired its ballistic missile rocket engine technology from Ukraine](#), who built the rocket motors in question during the Cold War, despite the publicly known fact that Russia had hundreds of these rocket engines in mothballs, and Ukraine only a very small number in museum exhibits. What followed was a major propaganda offensive by the Russians to exploit this at Ukraine’s expense, and [very loud public complaints of bias by the Ukrainians](#). The motives of the journalists remain to be explained, the misleading report was never retracted.

Many other examples exist. The result of poor journalistic practices, laziness or unreasonable time pressures to publish, a lack of deeper understanding of the subject matter, and frequently partisanship, is that far too frequently we see mass media producing and/or distributing what amounts to *“fake news”*. Much of the exodus of traditional media content consumers away from established media platforms, and the often cited lack of public trust in the media, can be explained by these common systemic problems.

The *“fake news”* problem observed in mass media is similar in many ways to the quality assurance problems that bedeviled manufacturing industries decades ago.

There is a potential payoff in establishing process based quality assurance standards for mass media, and mandating compliance with such standards.

Fundamentally, too many contemporary media organization staff suffer the very same gullibility problems widely observed in social media users.

References:

1. Kopp C, Korb KB, Mills BI "Information-theoretic models of deception: Modelling cooperation and diffusion in populations exposed to 'fake news'". PLoS ONE 13(11), November, 2018: e0207383, URI: <http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0207383>
2. Kopp, C, Korb, K.B, "We made deceptive robots to see why fake news spreads, and found its weakness", The Conversation, November, 2018, URI: <https://theconversation.com/we-made-deceptive-robots-to-see-why-fake-news-spreads-and-found-a-weakness-104776>
3. Kopp, C, Korb, K.B, Mills, B.I., "Understanding the Inner Workings of "Fake News"", Science Trends, November, 2018, URI: <https://sciencetrends.com/understanding-the-inner-workings-of-fake-news/>
4. Kopp, C., "Understanding the Deception Pandemic", Presentation Slides, Australian Skeptics Seminar, 16th July, 2018, Melbourne, Australia, URI: <http://users.monash.edu/~ckopp/Presentations/Understanding-The-Deception-Pandemic-2018-B.pdf>
5. Carlo Kopp, Kevin B. Korb, Bruce I. Mills, Written Evidence, "Inquiry on Disinformation and 'fake news'", House of Commons Digital, Culture, Media and Sport Committee, 12 December 2018, URI: <http://data.parliament.uk/writtenevidence/committeeevidence.svc/evidencedocument/digital-culture-media-and-sport-committee/disinformation-and-fake-news/written/93672.html>
6. Wineburg, Sam and McGrew, Sarah and Breakstone, Joel and Ortega, Teresa. (2016). "Evaluating Information: The Cornerstone of Civic Online Reasoning." Stanford Digital Repository. Available at: <http://purl.stanford.edu/fv751yt5934>
7. Samantha Bradshaw, Lisa-Maria Neudert, Philip N. Howard, "COUNTERING THE MALICIOUS USE OF SOCIAL MEDIA: GOVERNMENT RESPONSES TO MALICIOUS USE OF SOCIAL MEDIA", NATO STRATCOM COE, Riga, November 2018, 11b Kalciema Iela, Riga LV1048, Latvia, URI: <https://www.stratcomcoe.org/download/file/fid/79655>
8. Kopp C., "Considerations on deception techniques used in political and product marketing", Proceedings of the 7th Australian Information Warfare and Security Conference, 4 December 2006 to 5 December 2006, School of Computer Information Science, Edith Cowan University, Perth WA Australia, pp. 62-71. URI: <http://users.monash.edu/~ckopp/InfoWar/Lectures/Deception-IWC7-2006-BA.pdf>
9. Kopp C., "Classical Deception Techniques and Perception Management vs. the Four Strategies of Information Warfare", in G Pye and M Warren (eds), Conference Proceedings of the 6th Australian Information Warfare & Security Conference (IWAR 2005), Geelong, VIC, Australia, School of Information Systems, Deakin University, Geelong, VIC, Australia, ISBN: 1 74156 028 4, pp 81-89. URI: <http://users.monash.edu/~ckopp/InfoWar/Lectures/Deception-IWC6-05.pdf>
10. Kopp C., "The Analysis of Compound Information Warfare Strategies", in G Pye and M Warren (eds), Conference Proceedings of the 6th Australian Information Warfare & Security Conference (IWAR 2005), Geelong, VIC, Australia, School of Information Systems, Deakin University, Geelong, VIC, Australia, ISBN: 1 74156 028 4, pp 90-97. URI: <http://users.monash.edu/~ckopp/InfoWar/Lectures/Method-IWC6-05.pdf>
11. Kopp C., "Shannon, Hypergames and Information Warfare", in W Hutchinson (ed), Proceedings of the 4th Australian Information Warfare & Security Conference 2003 (IWAR 2003). Perth WA Australia, 28 - 29 November 2003, Edith Cowan University, Churchlands WA Australia, ISBN: 0-7298-0524-7. URI: http://users.monash.edu/~ckopp/InfoWar/Lectures/_JIW-2002-1-CK.pdf
12. Kopp C. and Mills B.I., "Information Warfare and Evolution", in W Hutchinson (ed), Proceedings of the 3rd Australian Information Warfare & Security Conference 2002 (IWAR 2002). Perth WA Australia, 28 - 29 November 2002, Edith Cowan University, Churchlands WA Australia, ISBN: 0-

7298-0524-7, pp 352-360. URI: http://users.monash.edu/~ckopp/InfoWar/Lectures/_JIW-2002-2-CK-BIM.pdf

End of Submission