Trust and the Web – A Decline or a Revival?

Piotr L. Cofta

University of Science and Technology, Bydgoszcz, Poland
E-mail: piotr.cofta@utp.edu.pl

Abstract

The Web was conceived as the 'better place' where trust could have flourished, but this is unlikely to be true anymore. There are two research areas that should be at the forefront of addressing this problem: Web Science and computational trust. This paper concentrates on what, if any, is the role of research in bringing trust into the Web, as the topic of trust currently takes a backseat. Amid the stable progress of research in trust in social psychology, economy, and politics, the bleak future of computational trust and trust on Web deserves an analysis. The paper analyses and reflects on the history of research in trust to conclude that this branch of research withered away before reaching maturity of a discipline. From there, it explores what it takes to accelerate and stabilise the development of research, and what challenges should be tackled first.

Keywords: trust on Web, computational trust, Web Science, research discipline.

Introduction

Golbeck [18] notes that "The success of the web is based on its open, unmanaged nature; at the same time, that allows for a wide range of perspectives and intentions". Taking further into account that there is no central authority, no enforcement and no standards, it is equally easy to install a malevolent router and to spread fake news as to observe ethical standards and play by the rules.

Still, the Web is central to our lives: we entrust it with our plans and opinions, we seek answers to burning questions, we use it to conduct financial transactions, we utilise it to stay in touch with friends and family. As the share of our lives that we spend 'on the Web' grows, the question whether the Web can be trusted becomes increasingly important.

Research should be at the forefront of our activities: paving the ways, anticipating problems, delivering solutions and defusing potential show-stoppers. Considering that the Web has been with us at least since 1989 [6], the question of trust should have been solved many times over, and we should be enjoying a fully trustworthy web.

This is not the case. Overly optimistic assumptions about the human nature embedded into blueprints of the early web did not meet the reality. The ability to contact everybody everywhere, the ability to reach any piece of information in an instant did not make us significantly better. If anything, the web allowed us to uncover the side of the human nature that we would rather forget of.

It would not be a problem had trust research been vigorously studied the area and delivered improvements and solutions. But instead interest in trust and the web significantly waned over the years, as if the research community do not believe anymore that trust is worth their time.

This paper formulates the thesis that there is a decline in research about trust that may affect trust on the Web, to the detriment to us all. It uses the mixed methodological approach, reflected in its structure. It starts with stating and scoping the problem. From there it uses the historical literature review, highlighting what is emphatically called the 'golden years' of trust research. Statistical bibliographic analysis is used to demonstrate the recent decline in general research interest, despite occasional valuable contributions. At this junction, the question of the usefulness of trust is answered through a series of examples. Then the generalisation of the problem leads to the listing of potential challenges that may explain the decline. Each challenge is supplemented by an indication of how to overcome it. propositions of how to improve the situation lead to conclusions that close the paper.

Scope

This paper focuses on research in computational trust and Web Science, as they should be at the forefront of studying, defending and re-introducing trust to the Web. The construct of trust is native to psychological and sociological research, as it mostly relates to relationships between people within the society.

However, it is researched across many disciplines, including information technology, and each one has its own definition and methodology when it comes to researching trust.

Research in trust tends to recognise the definition that relates to our decisions and actions: that people often behave above and beyond what the pure risk-based utilitarian approach should suggest. This surplus is called 'trustworthiness' and it can be acknowledged by actions where people make themselves dependent above what their utilitarian risk-based approach suggests. Those actions are called 'trust'.

It is widely acknowledged that trust and trustworthiness deliver benefits that the calculative, risk- based (hence "trustless") approach cannot. While the valuation of those benefits vary, the reasons for such benefits lie e.g. in simplifying decision making [26], minimising social friction by decreasing transaction costs and increasing the speed of transactions [16], in enabling various interactions that otherwise could not have a place (such as innovation and knowledge sharing) and in other similar effects [40].

The scope of the research are of 'trust and the Web' has not yet been clearly defined, so what is outlined below is probably the first attempt of it. What truly differentiates this research area is the scale. Contrasting with the neurology of trust, psychology of trust, or even with organisational trust, 'trust on the Web' has to work on the Web scale to be of any use (see e.g. [20]).

Trust, in order to be processed at the web scale, has to be expressed numerically in a format that can be processed by algorithms. Thus, when trust meets the web, the question of computational forms of trust emerges [28]. As trust is a complex psychological [8], simplifying proxies are used, such as reputation, recommendation or feedback.

Consequently, the following principles can be formulated to define the area of 'trust on the web':

1. Research interest is in understanding visible indications of trust and trustworthiness between the users of the Web, as mediated through technology, so the reasoning about trust can be automated.
2. The web changes our ways of dealing with trust (a quantitative change eventually became a qualitative change), so that rules about small-scale face-to-face trust may not be relevant. Specifically, the Web allows for one-off, remote, anonymous (or pseudonymous) interactions.
3. In order to be processed on a web scale, trust has to be expressed in numerical form, and subject to algorithmic processing.

There are some lines of research that are already considered to be in scope of 'trust and the web'. They may seem to be diverse, yet all of them satisfy those principles, as listed below.

1. Trust in technology (e.g. WebTrust, [http://www.webtrust.org]), i.e. collection of means that assure us that the web does what it is supposed to do, thus not altering the conveyed messages at the technical level.

2. Trust management, i.e. the collection of tools and methods to capture, represent, process, and use trust at the Web scale, i.e. WS-Trust [http://docs.oasis-open.org/ws-sx/ws- trust/v1.4/ws-trust.html].

3. Reasoning about trust by people through the Web, i.e. technical means available to participants to affect and judge trustworthiness of others, e.g. reputation-based systems [18].

4. Application of human ways of dealing with trust in technology to provide novel resolution to some of technical problems of the Web, e.g. trustworthy service composition [9].

5. Ability of the Web to discern between trustworthy and untrustworthy information in an automated manner, e.g. as in Semantic Web [33].

Research in 'Trust on the Web' and its Decline

It is the author's belief that the research in trust is currently in decline. It can be said that researchers simply exhausted its interest in trust and moved their resources elsewhere. There can be even some cynical views that research interest followed grants into some alternative green pastures. If this is the case, however, then the problem of trust on the Web should have been solved, and it is apparently not.

Golden Years in Trust Research

Trust has been a fundamental question for a long time, despite a disagreement whether trust is morally right or wrong. Proponents of trust applauded it as a social virtue and point that without exercising trust, trust will never be reciprocated (e.g. Aristotle, Hume, Kant). Opponents tend to think of trust as a vice and demonstrate the irrationality of trusting in light of visible high level of betrayal (e.g. Plato, Machiavelli, Hobbes). Whatever was the proposition, the importance of trust (and to certain extent the importance of control) to human life is visible everywhere.

While there is no exact definition of 'golden years' in trust research, the author loosely defines them as ranging from 1979 to 1996. They were

defined by the almost continuous inflow of original, bold ideas about trust, its specificity and its role. Several research disciplines focused on trust in their own specific ways, revealing the wealth of approaches and perspectives. As it is hard to acknowledge all contributions, it is worth listing at least some books that made an impact, in the chronological order of their English publications.

Luhmann's books [26] (originally published in German in 1973) and [27] (likewise, in 1984) approached trust from the perspective of social systems theory. He introduced the notion of 'trust as a way to reduce complexity' and the concept of 'double contingency' as a way to establish trust. Trust, as he says, is able to overcome our limitations of bounded rationality and reduce the otherwise unbearable complexity of all foreseeable futures. It becomes psychological and social necessity that can be rationally analysed, weighted and explained. Indeed, "but a complete absence of trust would prevent him [a person] from getting up in the morning" [26].

Deutsch [13] study trust as a way to avoid, manage and resolve conflicts and continue cooperation. Working on the basis of the Prisoner's Dilemma [4] (thus a game-theoretic, experimental approach), his work extended into exploring human motivations and behaviours that accompany the game. It stresses differences between the mental state and the actual behaviour, thus allowing for the rich situational model that covers also pathological situations.

Barber [5] well reflects the spirit of the wave of research by saying "Today nearly everyone seems to be talking about 'trust'.". [ibid., p. 1]. His objective was to make trust into a precise, well defined sociological concept, that is to develop the formal model of trust. His approach is based on exploring various situations of trust using the language of expectations and their fulfilment. Using this language, trust is the generalised expectancy of good behaviour of a trustee over whom there is no control.

Fukuyama's [16] demonstration that there is a link between the level of social trust, welfare and development has created interest in considering trust as a form of a 'social capital' and has led research into closer understanding of the process of creation and distribution of such capital. It has been widely accepted that trust (or, to be more specific, trustworthiness that leads to a vested trust) decreases transactional cost, thus leading to greater efficiency in economical relationships.

Kramer and Tyler [24] edited, and contributed to, an overview of research in trust in organisational sciences. The discovery that trust increases efficiency as well as effectiveness of work, specifically for knowledge workers, made organisations interested in measuring and increasing thrust at the workplace.

The volume combines reports on various aspects and form of research into the current overview of how profound the question of trust became.

Hardin [21] brought the thinking about trust in a systematic way into political sciences. He viewed trust as an 'encapsulated interest', where one can be trusted if one's best interests encapsulates the best interest of a trustor, thus stabilising trust and trustworthiness. He also explored the development of trust as a two-way relationship, differences between trust and trustworthiness, failures in reasoning about trust, as well as distrust.

The Second Wave

The second wave of research started just before 2000 and continues up to our days. It focused on analysing and exploring the variety of approaches generated during the first wave, in order to attempt a synthesis of a kind. One can see it as a natural reaction to the wealth of data and to the disparity of perspectives of the first wave.

Sztompka synthesised a social theory of trust, understood as a fundamental component of human actions. His intention was to "disentangle the problematics of trust" [40] through the clarification of the notion of trust that leads to a theoretical model of the emergence of trust cultures, tested against actual historical events.

Uslaner [43] reverts back to the early question about the ethics of trust and the impact of trust (or the lack of it) on societies. Using mostly quantitative research and the vast amount of available data, he constructed a compelling story about the way trust brings benefits to the society, across economics, politics and everyday life.

Mollering [31] structured an extensive overview of the research from the area into a three-factor model of reason (i.e. trust as a rational choice), routine (i.e. trust as an everyday normality) and reflexivity (where trust develops between active actors).

Cofta [10] studies the relationship between trust and control, seeing them as both competing and complementing. The proposed model of a decision to trust incorporates both elements and is applied specifically to situations there modern technology is a visible component.

The recent outcrop of excellent works from organisational, political and social sciences [36, 44, 49] continue this trend by providing an extensive overview of existing trends in research. It is only unfortunate that there is no similar collaborative effort in the area of computational trust and the Web.

Trust and Technology

When it comes to research, the relationship between trust and technology, specifically trust on the web, is more convoluted. Five lines of research mentioned earlier developed at different pace, drawing from different neighbouring disciplines.

1. Trust in technology is a direct descendant of security research. Trust and security have a long relationship, all the way back to 'Rainbow Books' [14] where the trusted execution environment was defined. Since then, the focus of information security has been to use trust sparingly to make the computing environment do what it was supposed to, despite attacks or faults – to make technology trustworthy. This has led to several developments, such as cryptographic key management, trusted computing [34], and others.

   The development of the Web challenged several established principles, as the open architecture of the Web did not fit the concept of physically protected data centres with defined ownership, services, rules and access rights.

   Still, the concept of trustworthy Web (i.e. the Web that technically can be trusted) follows the same mindset of creating and observing security profiles, applying specific controls and seeking compliance (e.g. http://www.webtrust.org/). The only interesting disparity is that the vagueness of the word 'trust' prevents parties to agree on the exact specification of those profiles.

2. Trust management developed from the more mature information security research, once it became clear that an appropriate architecture of trust is a precondition of a good security. Or, as it has been often said 'encryption is easy, key management is hard'. The particular focus of this area is probably on identification, authentication and authorisation. This is because the Web requires authentication that is much more fluid, standardised and easy to use than the more traditional information system.

   Outcomes such as the WS-Trust protocol [http://docs.oasis-open.org/ws-sx/ws-trust/v1.4/ws- trust.html] as well as Shibboleth [https://www.shibb
   oleth.net/] or single sign-on schemes such as OpenId [https://openid.net/] established themselves on the market.

   While addressing the important issue of trust management on the Web, the mindset visible through this research is derived mostly from information security, and focus on formats, interconnectivity and ease of use.

3. The very different approach, developed much later, was to observe, influence and benefit from trust that exists between people that can be observed through technology. For that approach to work it was necessary to wait for technologies that let deal with trust on a massive scale in an automated way. That is, to wait for the emergence of methods to socially interact through the Internet, mostly social networks, Internet shopping sites etc.

   While first attempts to work with trust were driven by intuition and practice, research quickly followed. The foundation of this approach was the notion of reputation (used as a proxy of trustworthiness), centrally computed from feedback and other relevant information. Books such as [19] provide good account of this approach. More recently, there is a growing interest in blockchain as an alternative way of assuring, verifying and reason about trust.

   This line of research follows very closely principles of the 'trust on web': it is concerned with human actions on the Web, it processes trust algorithmically and its interest is in large-scale schemes characteristic to the web.

4. The notion of reputation eventually seeped back into technology itself and was applied to technical objects and relationships between them, creating a third line of research. Service composition [9] is the primary example, where service providers are afforded trustworthiness and service users are free to choose the most trustworthy ones. In a similar way, autonomous agents [46] can choose whether to collaborate with more or less trustworthy ones. Other examples may include trust-based routing or trust-based access control.

   This probably begged the question whether the concept of trust, taken from social psychology, can be applied to inanimate objects without alteration. While the intentional stance [12] may warrant such an approach, the question whether 'trust' is a valid metaphor persists.

5. Web Science and trust require a special mention here. While the natural position of the Web Science is to observe the human society through the Web (hence using the second approach), the veracity of such observation is contingent on the Web being trustworthy (hence the first approach has to be considered). However, the trustworthiness of the Web is often contingent on trustworthiness of people, making the problem convoluted.

This is particularly visible in the architecture of its flagship Semantic Web [33], where trust resides at the topmost level, achieved by combining of reputation,

provenance and formal reasoning about information. Semantic Web is an example of the final approach: to make the Web trustworthy at the information level, i.e. to construct the web in a way that it will automatically prioritize information that is trustworthy.

What differentiates Web Science is also the fact that this is the first attempt to consciously include trust on the Web into a research project, from the very onset. While this approach did not generate a large research activity, it at least formulated some challenging questions.

The Recent Decline

This paper discusses mostly the problem of the intersection of trust and technology, as this is where trust on the Web, as well as trust in the Web, comes into play. It is the observation of the author that research in this area is in decline. That is, while e.g. research in organisational trust became part and parcel of corporate practices (at least for some corporations) and generated a large number of papers, research in trust and technology, specifically in 'trust on the Web' gradual diminishes.

The reasoning presented here is based on bibliometric analysis of programmes and proceedings of three series of conferences during the last 13 years (2006–2018). Trust, being fairly interdisciplinary, is in scope of several conferences, but those three were selected because not only they all have 'trust' in their names, but also because they focus on various aspects of trust and technology.

Trust is seldom the sole scope of any journal or a conference, probably with an exception of a series of past iTrust conferences (2003–2007) and the Journal of Trust Research [https://www.tandfonline.com/loi/rjtr20] (from 2011). Otherwise, trust tends to be clustered together with topics such as 'security' and 'privacy'.

Looking at Google Scholar, there are currently about 300,000 research papers available with 'privacy', 'security' or 'trust' as keywords. Unfortunately, the use of the sole keyword 'trust' is often misleading, as it is also used in legal, finances or in management sciences to indicate unrelated concepts. Therefore, instead of resolving the true affiliation of those papers, the author resolved to specifically look at conferences that have 'trust' in their primary scope (which usually means also in the title of the conference) and that were run long enough to provide a valuable insight.

This paper picked three series of conferences, of different origin, scope and size that can be considered a reasonable representation of the research

area of trust. The choice is partially opportunistic, as on occasions the author contributed to at least some of them.

- Privacy, Security and Trust (PTS) is a relatively small conference (with about 60 papers presented in the best year). Its objective, expressed by the title of the conference is to seek interdisciplinary perspective on three areas: privacy, security and trust.
- IFIP Trust Management (IFIPTM) is also a small conferences (about 35 papers at its peak). It is a successor of the iTrust conference series and retained its specialisation in trust, gradually moving from various aspects of trust towards greater focus on trust management, built on computational trust.
- IEEE TrustCom is comparatively large conference (190 papers at its peak). It originated from the SDDS workshop that focused on security in communication and eventually embraced trust mostly from the perspective of trusted platform, combining it with a strong focus on computational trust.

This paper analyses proceedings from years 2006–2018, for all those conferences. Some conferences started earlier, but 2006 is the first year for which the full programme and proceedings of each conference were available. They were analysed in order to determine the total number of papers as well as the number of papers related to widely understood notion of trust. The author used his own expertise in determining the actual allocation of each paper, as the allocation to particular session has been sometimes misleading.

The outcome of the analysis is best presented on the following diagram (Figure 1). The solid line represents the fraction of papers devoted to the subject of trust presented across all conferences. Starting from about 50% in 2006, currently papers about trust represent slightly more than 10% of all papers presented (at conferences that have 'trust' in their names). This is due to two factors: the significant growth in the number of papers unrelated to trust (mostly related to security) and the slow decline of the absolute number of papers about trust.

The dashed line indicates a linear regression, showing that the fraction of papers decline by about 2.5% every year. Extrapolating, trust will not be a visible subject on trust-related conferences at all by 2024 – not a great prospect for a subject that is considered to be a fundamental one.

By looking at other conferences that try to incorporate trust into their scope (e.g. TrustBUS), one can see that the problem is not limited to those few and that it is unlikely the Programme's Committee fault. It looks as if a
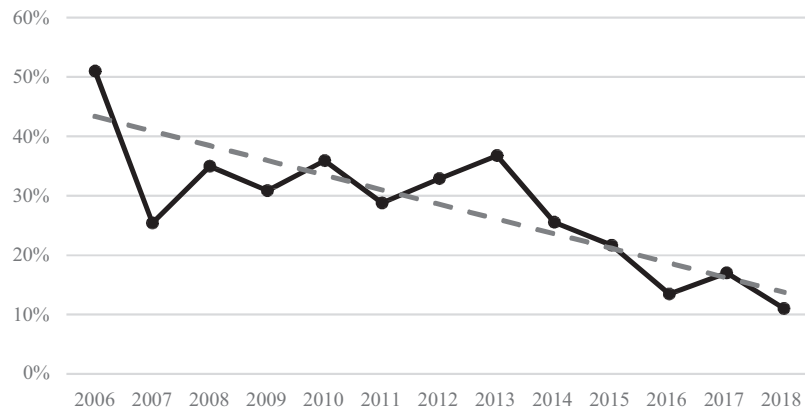
Figure 1   Fraction of papers devoted to trust across all three conferences.

research interest in trust wanes and withers. The question that this paper wants to discuss is 'why'.

Bright Spots

Despite all the above complaints about the decline, there are some bright spots, i.e. papers that reaffirm that trust research can bring interesting developments. Last year (2018) has brought a handful of those, some of which are mentioned below.

Aref and Tran [1] work contribute to the reasoning about computational trust. The key research problem is to develop a reasoning method that is similar enough to human reasoning but that can be expressed in an algorithmic form. Their use of fuzzy logic and Q-learning makes computational trust closer to the way people think about trust.

Information and trust in information is a particularly acute domain, as fake news undermined trust in information available on the Internet. Atele-Williams and Marsh [3] extends the notion of computational trust into a domain of information, with a formulation of computational rules for trust in information.

Trust is an entangled problem that does not yield itself to reasoning that is hierarchical or linear. In fact, logical loops and endless dependencies can make the chain of reasoning practically infinite. Of last year's notable contributions, there was the Tagliaferri and Aldini [41] paper on the logical reasoning about trust, that managed to address the problem of potentially endless recursion in our calculations.

Finally, the paper from Ashtiani and Azgomi [2] brought an unconventional view on trust. Using a similarity to quantum computing, the paper proposed that every situation can be described by the superposition of of trust and distrust. Such superposition cannot be tested without deciding on trust as even by observing the situation (or thinking of it) the decision resolves to trust or distrust.

Are We done with Trust?

Before engaging in any discussion about the decline and the revival of a discipline, it is worth pausing and reflecting whether such a revival is indeed required. The discipline may slow down its progression because it may no longer be of a primary interest to the society or because it solved all of its burning questions.

As for not being the primary interest, trust research is in a precarious position: there is a demand for trust coming from various corners of the Web, yet there is a growing understanding that trust is getting harder to come by. Not only deficiencies in information security make participants weary, but also permissiveness regarding falsification of information [30, 42] make people suspicious. Eventually, it may have a double negative effect: those who are disenchanted will stop participating, thus making further falsification easier.

The second cause of disinterest in a discipline may be its maturity where all burning questions are solved. Again, it is unlikely the case. Had it been an established methodology to create trust wherever it is needed, by any means needed, the answer could have been positive. As it is now, both reputation-based systems and blockchain (to name a few) caught the research community somehow by a surprise. Looking, for example, at the number of papers about blockchain one can see an exponential growth but only after the theory of blockchain was converted into an interesting implementation of a Bitcoin network.

What Makes a Discipline Decline

There is not that many research on research itself, specifically on the decline of the whole discipline. There is often a perception, that the discipline – once created – cannot be rid of, even if it becomes irrelevant. To establish at least some traces of evidence, one has to look at what constitute the normal development of a discipline, and then see how this development is not matching trust research.

Luhmann [27] treats science (and its disciplines) as autopoietic systems that are concerned with differentiation between 'true' and 'false', holding the monopoly on it. Being autopoietic, those systems preserve their structure while renewing themselves with new information, retaining meanings that could have been chosen almost at random early during the formation of a system.

The retention of meanings is done by interpreting incoming information in light of existing meanings. It is characteristic to autopoietic systems that gradually they develop an operational closure – the set of meanings that they operate within the system makes it near impossible to be penetrated from the outside. It is specifically true as the systems develops selective blindness when it reflects on itself. That is the science cannot tell whether the science itself is 'true' or not.

Kuhn [25] describes what is considered a normal development of a scientific disciplines. The evolution starts with a random collection of "mere facts", interpreted in incoherent ways. Through a brief competition between paradigms, a dominant one emerges and is accepted by the research community. As alternative paradigms fade away, the leading one becomes supported by various research structures such as journals, conferences or societies. Eventually the paradigm becomes formalised and taught, creating natural entry barriers to alternative views. It is now only a scientific revolution that can alter the course of research.

Fleck [15] provides interesting observations on how supposed objectivity of science (claimed e.g. by Popper [35]) is undermined by the desire of researchers to conform with the prevailing thought style (proto-ideas), so that science becomes subject to a self-reinforcing mutual influence to the point where it resists everything that contradicts it.

Those few opinions about the development of science have a lot in common. The normal formation of a research discipline starts from humble, almost random research activity. Several paradigms are formed, possibly in line with different proto-ideas. Then one paradigm becomes recognised and the discipline starts both formalising and closing itself, so novel ideas are rejected. It can remain in this state practically indefinitely, with only a drastic change being able to alter or remove it.

If we are to diagnose the discipline of trust research using this description, it will be a diagnose of a premature decline. 'Golden years' bear uncanny resemblance to the random development phase where several paradigms compete, but no winning paradigm emerged, despite a second wave of overviews and synthesis. Naturally, no self-sustained formalisation and closure followed – there is no generally accepted way of thinking and researching trust.

In Search for a Paradigm

The discipline of trust research has its challenges, discussed in the next section. All those challenges, however, are solvable. For the mature discipline they could have been an inspiration. For the research in trust they may be obstacles that are hard to overcome. Yet they can be overcome.

The lack of a shared paradigm is a different problem. It is as troubling as it is critical to the discipline. Without a single, acknowledged paradigm the trust research cannot become a self-sustained discipline, continuing as a patchwork of incompatible research encroaching from and into other disciplines.

There are two cases that illustrate how the lack of a paradigm affects the discipline. The first one is the definition of trust – or rather the multiplicity of it. McKnight and Chervany [29] found no less that 17 semantically different meanings of trust, each one leading to a definition that is at least partially incompatible with others. Consequently, 'my trust' is usually different from 'your trust', and almost every research paper (this one including) has to contain a definition of trust.

The other example is taken specifically from computational trust. The simple question of how to express trust in a numerical form has a large number of answers: from binary to continuous, from one dimensional to three dimensional, from objective to subjective, from fixed numbers to temporal logic (see e.g. [10] for an overview). Results brought in by different researchers cannot be compared or consolidated, thus fragmenting the discipline.

It is worth contrasting this situation with the discipline of information security. While years ago there was a lot of discussion on what and how it should focus on, eventually the CIA (confidentiality, integrity, availability) paradigm became widely accepted, together with probabilistic, risk-based measurements. Once this has been settled, the discipline could have been formalised, taught and organised (see e.g. [22]). Researchers could have concentrated on solving particular problems, knowing that they contribute to the greater whole. As a result, the discipline flourishes, being currently the dominant one even at trust-focused conferences.

The author has of course his preferred paradigm, but it is not a place to discuss it. Instead, it is worth identifying what the paradigm should define. Kuhn [25] uses the word 'paradigm' in two different meanings: as a shared mindset of a research community and as a leading example that everybody willingly adhere to. Either way, it defines the way of perceiving the world, not just details of some models.

Paradigm (after [25]) should define as follows (examples are provided here for clarification, and do not constitute any proposal).

- What is to be observed and scrutinized. That is, for example, trust research should be only concerned with external signals between people, and with their actions, not e.g. with the neurology of a human brain.
- What kind of questions should be asked. While there are no forbidden questions, there may be ones that are irrelevant to the discipline. It may be thus irrelevant to ask a question about different kinds of trust, but it may be relevant to ask questions about the amount of trust.
- What is the primary theory and predictions are made by it. For example, if we accept as a primary theory that trust reduces complexity, then the theory can predict that higher trust is associated with lower complexity.
- How an experiment is to be conducted and research interpreted. For example, the discipline can assume that the 'Game of Trust' [10] or the 'Prisoner's Dilemma' is a valid method to determine the extent of trust without further investigation into subjective psychological processes.

Paradigm should also maintain certain social recognition in a form of professional organisations, leaders, journals, conferences, education materials and media coverage, not to mention governmental support and sources of funding.

Challenges and How to Overcome Them

As stated above, the paradigm is a major challenge that prevents trust research from becoming an established discipline. Still, there are several other challengers worth mentioning, even if only to illustrate that the area of trust is changing and that research should adopt itself to those changes.

Monopoly and Trust

The web has changed significantly since its early days, and one of the most striking change is the monopolisation of its key functions. That is, despite technology that allows everybody to develop and to contribute to the Web, the control over its key content is restricted to a handful of companies.

Monopoly of any kind invalidates the kind of trust that is usually the research subject. Such trust requires a choice: at least a choice of trusting or not trusting. Even better, the choice to trust someone or someone else.

The same way the reputation requires at least two entities to choose from, so that the reputation of the monopolist is irrelevant. It is not that it cannot be computed – it is simply irrelevant.

Still, people behave as if they trust monopolists, despite occasional disappointments. It is because under monopolistic circumstances, the nature of our trust changes. It is possible that instead of a familiar trust-as-a-choice, people resort to the lesser familiar one: ontological trust [17]. This is the trust that people put in what is believed to be the natural order of things. This allows people to trust the world, or rather trust their understanding of the world.

From the perspective of the research in trust on the Web, this introduces the disconcert. The research that concentrates on trust-as-a-choice grows, but it becomes increasingly irrelevant. There are good theories that explain why a user should trust one site more than the other one, but there are no theories explaining how the user can trust the only site available. Reflecting on the Semantic Web, the challenge now may be to have a scheme that allows to calculate how trustworthy is the sole, obscure source of information.

So, the challenge is to address the Web as it is, not as it has been imagined. This challenge is intriguing, because in doing it it may be possible to discover a paradigm that covers all kind of situations, those with a choice and those without it.

The Manichean Evil

The research in trust assumed (for whatever reason) that the silent majority of people on the Web is trustworthy, rational, informed and willing to participate. That is, that the crowd is wise, as epitomised in [39]. In this situation, provisions were made to contain those few untrustworthy ones who (intentionally or not) tamper with the system. For as long as the majority was there, corrective actions worked. Augustine's evil of imperfection was contained and hopefully eradicated.

This situation has changed, and it caught trust researchers somehow unawares. It has been painfully demonstrated [42] that the resourceful few can overpower and overcome trust-based system to the extent that the crowd was no longer wise enough to contain the damage. The Manichean, intentional evil trumped the Augustine's one.

Certainly, there was always some research on possible attacks on trust-based systems (see e.g. [10] for an overview), touching mostly reputation-based schemes [19]. However, it was usually more important to focus on intricacies of calculating reputation rather than considering bleak scenarios where the majority of users are deceived.

This is in a stark contrast with e.g. information security where a thorough analysis of the feasibility of various attacks is a standard part of any viable research. It is as if the mindset of a trust researcher excluded the possibility of large-scale attacks. In fact, there are scantily few research works on distrust, so trust-based systems can neither recognise nor deal with it. When it comes to computational trust, there is even no standardised way of encoding distrust, and some algorithms do not accept its existence.

Considering the amount of trust-based systems that are in constant use, attacks made by distrustful agents will only increase, both in their ferocity and frequency. They can be contained if those attackers and their methods are better understood. Therefore it is urgent to study a dark side of the human nature as well as to improve the resistance of algorithms to various attacks. If nothing else, it is a research opportunity that should not go amiss.

Rediscovery of Trustless Cooperation

Trustless cooperation seems to be almost an oxymoron. From early research on trust (e.g. [4]) cooperation was always assumed to be conditional on trust, to the extent that the presence of cooperation was accepted as a sure sign of the existence of trust, and the extent of cooperation was adopted as a measure of trust [24].

The existence of trustless cooperation puzzled researchers. If trust is not needed to facilitate the cooperation within the society, then possibly research in trust is not needed as well. Cooperation that does not need trust requires a driver that is at least as strong as trust.

It was found in self-interest, namely the common greed. The primary current example of trustless cooperation is Bitcoin [32], the system with its origin shrouded in some secrecy. While Bitcoin is technically a creative combination of distributed ledger, proof-of-work consensus, cryptography and gossip networking, its true strength lies in harnessing greed into supporting a common cause. Unfortunately, the common cause is to let the greed flourish at the expense of our shared environment, as Bitcoin miners consume now more energy than Ireland [https://digiconomist.net/bitcoin-energy-consumption].

The use of self-interest, specifically of the desire to generate private wealth, to derive some common goods is not new, as it is a foundation of a free market [38]. Thus it is not the surprise that it works, and it is not surprising that its work is akin to the early capitalistic enterprises: without any regard to the common goods.

It would be now in place to note that the famous Smith's work is this second book on the subject. The first one [37] discussed, among other topics, 'social transactions' that extend beyond purely commercial ones to include gift exchange, favours, sympathy and passion. There is a parallel with modern works [16] where such 'social transactions' facilitate the creation of mutual trust that supports the perpetuation of the free market. If this is truly the case, then Bitcoin is slated for failure, probably through abandonment, unless it can incorporate some elements of social transactions into itself.

There is a question whether trustless cooperation is entirely trustless. Indeed, it is hard to eliminate trust entirely from any human endeavour, least from the collaborative creation of wealth, even illusory one. Thus certainly there is some trust required to run the Bitcoin, placed e.g. in public repositories that hold the code, the Internet and its ability to interconnect nodes or in the invariability of a human nature.

However, no trust is required in actions of others who participate in the network. There is no Prisoner's Dilemma [4] to solve, no concern for the wellbeing of others (or of the environment) to demonstrate, no trustworthiness to be reciprocated with trust through social transactions: one can entirely self-indulge in one's self-interest. In this sense it is trustless indeed.

Before closing the subject, it is necessary to differentiate between permissionless, open schemes such as Bitcoin and permissioned schemes such as Ripple [https://ripple.com/] or the blockchain technology itself (e.g. HyperLedger [https://www.hyperledger.org/]). Permissioned schemes vest trust in at least some of participants of the scheme, expecting them to cooperate for the benefit of the scheme. Such trust is usually expressed in a form of centrally-managed credentials. Blockchain technology itself is trust-neutral, as it can be used to construct any scheme, from private ones, permissioned via some central authority to permissionless, open ones.

Blockchain as such attracted some interest of the research community. After all, blockchain is an implementation of yet another source of trust: distributed ledger that facilitates accountability, thus making people behave in a trustworthy manner under the threat of social ostracism. Still, while Google Scholar reports exponential growth in research about blockchain as such, research about trust and blockchain is almost invisible.

The Entanglement of Trust

Trust is a hard research topic, because it is an entangled one [31]: trust is contingent on trust, and any change on trust reverberates throughout the system, in

a way that is sometimes. Thus it is hard to isolate the research problem enough to make a satisfactory, reproducible results. Following Luhmann's statement that "the only society is the world society" [27], one may realise that the only system of trust worth studying is the system that includes the world society.

Trust is a social construct with a very simple objective: to make the complexity of the world manageable [26]. There is no requirement that everybody should use the same formula, or that everybody should trust the same way. The famous juxtaposition of two statements: "I trust you because of your good reputation" and "I trust you despite your bad reputation" [23] makes it clear that no single formula can cover all eventualities.

Resolving the question of trust means solving the entangled Byzantine problem [7]. The 'normal' Byzantine problem is complex enough, as we have to assume that, while exchanging messages, we cannot trust some or all elements of our immediate neighbourhood. But a entangled version is more complex still: we cannot even trust the meaning of the message and we cannot fully trust ourselves in figuring out whether to trust or not. We cannot even trust the word 'trust', least the meaning of.

Let's say we have a reputation system and we believe that it predicts the extent of trustworthiness. Reputation is based on opinions, and opinions are trustworthy or not. So we have a second-level reputation system that calculates reputation of opinion-makers. In principle, we should also have a third level system for the reputation of reputation of opinions, and probably all the way down to 'six degrees of separation' where we calculate reputations of all statements all over the world, including our own belief in a reputation of reputation of reputation about ourselves – with the obvious conflict of interest that it brings. And that's just for the start.

This inability to isolate chunks of the problem runs against the grain of specialisation, and leads to fragmented, often disjoint outcomes. Trust, like security, is a whole-system property [7], so even systems constructed from individually trustworthy components may not bring expected level of trust or trustworthiness.

It is the author's belief that this disillusionment can be overcome by introducing the consolidating definition of trust, in the same way information security came up with the triad of confidentiality, integrity and availability. Such definition could deliver the framework that spans the world society while allowing for well-defined, localised innovations. It is also the author's belief that such a framework may arrive from the area of social constructivism, currently neglected by trust research.

What Should be Done

The question whether we have to do anything about trust is intentionally provocative. It has been already established earlier in this paper, that trust is required for the existing Web to operate and that the problem of trust has not been solved to our satisfaction.

Still, it is a valid question: is it possible to have a Web without trust? In light of current interest in trustless cooperation, this is possible. By analogy, it is sufficient that those who invest in the web will have the ultimate say about every detail of its operation and content. The dominant player will be able to rewrite the Web and its content as it pleases. It will be of course a different Web, probably less conducive to innovations and more conducive to censorship, with 'Net neutrality [https://en.wikipedia.org/wiki/Net_neutrality] effectively removed, but a Web nevertheless. Incidentally, this is in some ways an attractive proposition for some, as it may also address some challenges of security, privacy (or the lack thereof), copyright etc.

As the web will be run by few monopolist, trust in the current form will not be required. Some forms of ontological trust will be needed to trust monopolists. Looking at five research areas discuses in this paper. (1) Trust in technology will be replaced by tight security with appropriate controls. (2) Trust management will not be needed beyond agreed access and monitoring controls. (3) Reputation-based systems will become proprietary and users will be judged by monopolists. (4) Human trust to deal with technology will practically stop. (5) Monopolists will decide what will be presented on the Web to be trusted.

Assuming that this is not the desired shape of the Web, here are some action items worth consideration.

Paradigm and Challenges

Some guiding principles of 'trust on the Web' have been already listed: human activities as visible on the Web, large scale operation and computational forms of trust. The question remains whether those three are sufficient as a paradigm of a discipline, whether they both separate it from others and are of sufficient explanatory power.

It is the author's belief that they are not. That without stating clearly what is trust 'on the Web' and how we approach it, we will restrict the research to reactive interpretations of what we see on the Web, without being able to formulate predictive theories. For that, the joint theory of the Web and the society (see e.g. [11] for an attempt) may be of some use.

Challenges, those formulated here, and other that will come, are always welcome. They can be used to sharpen the research mindset and improve the research toolkit. They can be also used to make research more relevant to current situations. Further, it is likely that they will demonstrate weakness of existing approaches and they will stimulate the development of new ones.

Schemes That are Not There

It is both interesting and intriguing that there are schemes for human trust on the Web to be harvested, processed and used, but there are no (or few) schemes that allows for the creation (the 'emergence') of trust. It is as if the Web consumes trust created elsewhere, not contributing to the growth of it. This exploitative approach may not be sustainable.

When the emergence of trust is discussed, there are some basic models that can be applied: Axelrod's game theories [4], Hardin's encapsulation of interests [21] and Luhmann's double contingency [27]. Trustless cooperation, discussed earlier, lies within the domain of calculative behaviour and is not discussed here.

Of those three, game-theoretic approach has been accepted in social scenes, specifically in economy (see e.g. [47]) as model of explaining trust of 'homo economicus'. While such trust may border on purely calculative approach, it is used to explain the emergence of a behaviour that is not justified by short-term rationality. With the emergence of Internet shopping, the model has been adopted for at least some of Web-based activities and then expanded into social networks [45], Web services collaboration [48] and beyond. By all means, to date this is the most successful model explaining he emergence of trust.

Contrasting, the author is not aware of any scheme that specifically model or implement Hardin's encapsulation of interest or Luhmann's double contingency. Those schemes are not there, are not formalised into any computational form and are not implemented. If any, they are the potential research areas that may redefine what we think about trust on the Web, by making the Web into trust-building too.

Resources

Research requires resources. It is an undeniable fact that throwing money at the discipline makes researchers flock to it and the discipline flourish. Figure 2. illustrates this case with the number of papers on trust published in a given
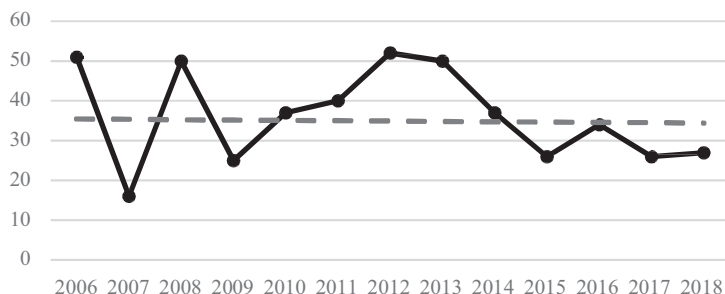
Figure 2   Absolute number of papers on trust, all three conferences.

year. Amid the general decline, years 2010–2015 clearly show the gradual increase, followed by the gradual decrease, in contributions. These years coincide with the period of time where EU run the Seventh Framework Programme (FP7), 2007–2013 (https://ec.europa.eu/research/fp7/index_en.cfm), specifically while taking into account the lead time for a research work to be published.

There is nothing wrong with it. Contrary, it shows that the research in trust can be revived with another round of funding. However, it also shows that funding is not enough – FP7 did not became a catalyst that converted disparate research into a self-sustained discipline.

If there are funds available, this time they may be channelled towards new challenges, but on condition that a shared, interdisciplinary paradigm for trust is used by all participants. It does not have to be the final one, or even the best one, as the evolution of science will take care of improving it. It only has to be accepted by the majority.

Conclusions

Research in trust, specifically trust on the Web is important. However, its situation is troubling. It never became a discipline in its own rights. Instead, it suffers a decline that cannot be resolved simply by an enthusiasm of a few or even by increased funding. Research in trust on the Web has its challenges to overcome, but above all it requires a paradigm that will allow to mature it into a discipline.

By using a variety of methods, the paper outlined and analysed the history of research efforts related to trust, with special focus on computational trust and trust on the Web. It demonstrated how, after relatively promising decades, the decline has set, threatening trust research with irrelevance and disappearance.

The survival, and the revival of trust research requires the conversion from disparate efforts to a research discipline. For that, the unifying paradigm has to be found and promoted, and challenges addressed. The paper withholds from presenting yet another candidate for a paradigm, instead outlining what should such a paradigm cover.

References

[1] Aref A., Tran T. (2018) A hybrid trust model using reinforcement learning and fuzzy logic. Computational Intelligence. 2018;34:515–541. https://doi.org/10.1111/coin.12155

[2] Ashtiani M., Agzomi M. A.: A novel trust evolution algorithm based on a quantum-like model of computational trust. Cogn. Tech Work (2018). https://doi.org/10.1007/s10111-018-0496-9.

[3] Atele-Williams T., Marsh S. (2018). Towards a Computational Model of Information Trust. 124–136. 10.1007/978-3-319-95276-5_9.

[4] Axlerod R. (1997) The Complexity of Cooperation. Agent-Based Models of Competition and Collaboration. Princeton University Press. ISBN0-691-01567-8

[5] Barber B. (1983) The Logic and Limits of Trust. Rutgerts University Press.

[6] Berners-Lee T., et al. (2006) Creating a Science of the Web. Science, 313.5788 (2006): 769–771.

[7] Bishop M. (2005) Introduction to Computer Security. Addison-Wesley. ISBN 0-321-24744-2

[8] Castelfranchi C., Falcone R. (2000) Trust is much more than subjective probability: mental components and sources of trust. In: Proceedings of the 33rd Annual Hawaii International Conference on System Sciences. ISBN: 0-7695-0493-0. DOI: 10.1109/HICSS.2000.926815

[9] Chang, E., Dillon, T., Hussain F. K. (2006) Trust and Reputation for Service-Oriented Environments. John Wiley and Sons. ISBN: 978-0-470-01547-6

[10] Cofta, P. (2007) Trust, Complexity and Control: Confidence in a Convergent World. John Wiley & Sons, Ltd. ISBN: 9780470061305. DOI: 10.1002/9780470517857.

[11] Cofta, P. (2013) The Foundations of a Trustworthy Web. Now Publishers. Foundations and Trends in Web Science: Vol. 3: No. 3–4, pp. 137–385. http://dx.doi.org/10.1561/1800000020

[12] Dennett D.C. (1989) The Intentional Stance. Bradford Books. ISBN: 9780262540537

[13] Deutsch M. (1973) The Resolution of Conflict: Constructive and Destructive Processes. Yale University.

[14] DoD (Depatment of Defence) (1983) Trusted Computer System Evaluation Criteria (TCSEC) 5200.28-STD. Available: http://csrc.nist.gov/publications/secpubs/rainbow/std001.txt

[15] Fleck L. (1979) Genesis and Development of a Scientific Fact. University of Chicago Press.

[16] Fukuyama, F. (1996) Trust: The Social Virtues and the Creation of Prosperity, Touchstone Books, ISBN: 0684825252

[17] Giddens, A. (1991) Modernity and Self-Identity. Self and Society in the Late Modern Age. Cambridge Polity Press. ISBN: 0-7456-0932-5

[18] Golbeck J. (2006) Trust on the World Wide Web: A Survey. In: Foudations and trends in web Science 1:2. ISSN: 1555-077X

[19] Golbeck J. (Eds.) (2008) Computing with Social Trust. Springer. ISBN 978-1-84800-355-2. DOI: 10.1007/978-1-84800-356-9

[20] Hall W., Hendler J., Staab S. (2016) A Manifesto for Web Science @10. Available: http://www.webscience.org/wp-content/uploads/sites/117/2016/12/WebSci-manifesto- v19.pdf

[21] Hardin, R. (2002) Trust and trustworthiness. Russel Sage Foundation.

[22] Harris S., Maymi F. (2016) CISSP All-in-One Exam Guide, Seventh Edition. McGraw-Hill Education. ISBN: 978-0071849272.

[23] Jøsang A. (2007) Trust and Reputation Systems. In: Aldini A., Gorrieri R. (eds) Foundations of Security Analysis and Design IV. FOSAD 2007, FOSAD 2006. Lecture Notes in Computer Science, vol 4677. Springer, Berlin, Heidelberg. ISBN: 978-3-540-74809-0. DOI: 10.1007/978-3-540-74810- 6_8

[24] Kramer, R. M., Tyler T. R. (1996) Trust in organizations. Frontiers of Theory and Research. SAGE Publications. ISBN: 0-8039-5740-8

[25] Kuhn, T. S. (1962) The Structure of Scientific Revolutions. University of Chicago Press.

[26] Luhmann, N. (1979) Trust and Power. John Wiley and Sons.

[27] Luhmann, N. (1995) Social Systems. Stanford University Press.

[28] Marsh S. (1994) Formalizing Trust as a Computational Concept. PhD thesis, University of Stirling, Department of Computer Science and Mathematics.

[29] McKnight, D. H., Chervany, N. L. (1996) The Meanings of Trust. Technical Report MISRC 9604, University of Minnesota MIS Research

Center Working Paper series, WP 96-04. Available: http://www.misrc.umn.edu/workingpapers/fullpapers/1996/9604_040100.pdf

[30] McLaughlin T. (2018) How Facebook's Rise Fueled Chaos and Confusion in Myanmar. Wired. Available : https://www.wired.com/story/how-facebooks-rise-fueled-chaos-and-confusion-in-myanmar/

[31] Möllering, G. (2006) Trust: Reason, Routine, Reflexivity. Elsevier Ltd. ISBN: 978-0-08-044855-8.

[32] Nakamoto S. (2008) Bitcoin: A Peer-to-Peer Electronic Cash System. Available: https://bitcoin.org/bitcoin.pdf

[33] O'Hara K., Hall W. (2008) "Trust on the Web: Some Web Science Research Challenges". In: Julià MINGUILLÓN (coord.). "Web Science". UOC Papers. Iss. 7. UOC. Available: http//www.uoc.edu/uocpapers/7/dt/eng/ohara_hall.pdf. ISSN 1885-1541

[34] Pearson S., et al: Trusted Computing Platforms: TCPA Technology In Context. Prentice-Hall, 2002.

[35] Popper K. R. (1959) The Logic of Scientific Discovery. Hutchinson Press.

[36] Searle R., Nienaber A.-M., Sitkin S. (2018) The Routledge Companion to Trust. Routledge. Taylor and Francis Group. ISBN: 9781138817593

[37] Smith A. (1759) The Theory of Moral Sentiments. Millar, Lodon and Kincaid, Edinburgh.

[38] Smith A. (1776) An Inquiry into the Nature and Causes of the Wealth of Nations. Strahan and Cadell, London.

[39] Surowiecki J. M. (2004) The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations. Doubleday; Anchor. ISBN: 978-0-385-50386-0.

[40] Sztompka, P. (1999) Trust. A Sociological Theory. Cambridge University Press. ISBN 978-0-521-59850-7

[41] Tagliaferri M., Aldini A. (2018) From Knowledge to Trust: A Logical Framework for Pre-trust Computations. In: Gal-Oz N., Lewis P. (eds) Trust Management XII. IFIPTM 2018. IFIP Advances in Information and Communication Technology, vol 528. Springer, Cham. DOI: https://doi.org/10.1007/978-3-319-95276-5_8

[42] Tandoc E.C. Jr., Lim Z. W., Ling R. (2018) Defining "Fake News", Digital Journalism, 6:2, 137-153, DOI: 10.1080/21670811.2017.1360143

[43] Uslaner E. M. (2002) The Moral Foundations of Trust. Cambridge University Press. ISBN: 0 521011035

[44] Usalner E. M. (eds.) (2018) Oxford Handbook of Social and Political Trust. Oxford University Press. ISBN: 9780190274801

[45] Wang, Y., Cai, Z., Yin, G. et al. (2016) A game theory-based trust measurement model for social networks. Computational Social Networks 3: 2. DOI: 10.1186/s40649-016-0027-x

[46] Wierzbicki, A. (2010) Trust and Fairness on Open, Distributed Systems. Studies in Computational Intelligence. vol. 298. Springer. ISBN: 978-3-642-13450-0. DOI: 10.1007/978-3-642-13451-7.

[47] Witteloostuijn van, A. (2003) A Game-Theoretic Framework of Trust. International Studies of Management and Organization. Volume 33, 2003 – Issue 3

[48] Yahyaoui H. (2012) A trust-based game theoretical model for Web services collaboration. Knowledge-Based Systems Volume 27, March 2012, Pages 162-169. DOI: 10.1016/j.knosys.2011.10.014

[49] Zmerli S., van der Meer T. W. G. (eds.) (2017) Handbook on Political Trust. Edward Elgar Publishing. ISBN: 978 1 78254 510 1

Biography

Piotr L. Cofta received his MSc and PhD in Computer Science from Gdansk University of Technology (Poland) and his DSc (habil.) from the Polish-Japanese Academy of Information Technology. His research and industrial career spans organisations such Nokia Research, MIT Media Lab, British Telecom etc. Currently he is a professor of Computer Science at the University of Science and Technology, Bydgoszcz, Poland. His primary research interest is in trust and technology, specifically in computational trust, blockchain and the IoT. He published four books, and authored numerous papers and patents. He is a Senior Member of the IEEE.