Dionysios Georgiadis
Martin Raubal
25 November 2020

An interdisciplinary review on weak signal detection

#3

Table of Contents

Abstract

When managing systems that are embedded in complex environments, being able to anticipate and navigate discontinuities can make the difference between perishing and thriving. To this end, one may actively scan the environment for weak signals: seemingly minor irregularities that may foretell impending discontinuities. This approach has been studied across literature, from the detection of ball bearing faults in rotating machinery, to the rupture of fuel tanks, prediction of riots and political uprising, and financial crises. Despite the universality of this challenge, each literature has focused on idiosyncratic definitions, tools and solutions. As a result, there is no unifying framework for exploiting weak signals, and interdisciplinary convergent research is impeded. Seeking to remedy this issue, we take a holistic view on how weak signals are exploited across literature and propose a unified conceptual framework. Based on this framework, we identify research directions that may enable social sciences to draw methodology from mathematically rigorous fields. This approach stands to drastically improve weak signal exploitation in social systems by reducing human involvement and lowering the detecting threshold for weak signals.

Introduction

From the cheerful acceptance of the Trojan horse in antiquity, to the Mexican gulf oil spill in 2010 (Schoemaker et al. 2013), decision makers have been known to ignore subtle hints of impeding disasters, leaving observers wondering how the clear signs could have been missed. Management science dubs these signs weak signals, seemingly insignificant pieces of information that are hard to detect and to interpret — but catalytic to our understanding (Ansoff 1975). Organisations that systematically scan their environment for irregularities (Daft et al. 1988) may be able to detect weak signals and harness their predictive power — enabling decision makers to anticipate, avert or even take advantage of catastrophes. The concept of weak signals has also been employed in the literature on futures, or futurology (Sardar 2010), with the aim of detecting emerging trends in society as a whole (Rossel 2012).

Beyond social sciences (managerial sciences and futures), research on engineering and signal processing has been focusing on a similar problem: analysing data to uncover features that are both vitally important and obscure (Wang et al. 1999; Wang and He 2003; Qiu et al. 2006). The detection of weak signals is routine in multiple applications, ranging from GPS (Global Positioning System) positioning (Glennon and Dempster 2004), to the early detection of percolating microcracks in fuel tanks (Gould and Eldredge 1977), and the detection of minor defects in electrical motors (Thomson and Fenger 2001). Such applications are typically enabled by non-trivial data analytics methods, which rely heavily on concepts from nonlinear dynamics (Anifrani et al. 1995; Gould and Eldredge 1977; Harmer et al. 2002). Such methods are capable of automatically detecting and/or harnessing weak signals.

The automated systematic approach to harnessing weak signals that we encounter in mathematically rigorous fields contrasts the approaches found in management and futures — where the detection and interpretation of weak signals rely on expert opinion and ad-hoc analysis (Mendonça et al. 2004). Contemplating this contrast, one may wonder whether weak signal detection and interpretation in social sciences could benefit from the methodological arsenal of mathematically rigorous fields. However, so far, this question has not been systematically pursued. In fact, there is no framework unifying the phenomenology of weak signals across both social and engineering disciplines. As a result, individual literatures have developed their own idiosyncratic tools, concepts, and methods — resulting in a disconnect that obscures the potential for convergent interdisciplinary research.

The present report is a first foray to bridge this disconnect, by drawing from multiple literature to present a unified framework for systematically detecting and exploiting weak signals. We argue that such a framework can be built using concepts from certain works from complex systems and nonlinear dynamics that investigate: i) the causes and mechanisms underlying abrupt transitions in dynamical systems, and ii) the precursors of such events. The mathematical abstractness of the nonlinear dynamics literature has prevented the development of an idiosyncratic framework; and consequently, the concepts of nonlinear dynamics can be translated in both managerial and engineering sciences. A body of work has been consolidating complex systems and nonlinear dynamics with social and/or organisational sciences (Kaivo-oja 2012; Todnem By 2005; Somit and Peterson 1992; Timmermans 2008), while many engineering weak signal applications rely on nonlinear dynamics and complex systems (Sornette 2002; Gould and Eldredge 1977). However, few works have focused on using methodology from nonlinear dynamics and complex systems, to aid weak signal detection and interpretation in managerial sciences (Durlauf 1999; Goldstein et al. 2010).

As an example of how nonlinear systems may offer a unified framework for harnessing weak signals, we may consider a specific vibrational pattern that the nonlinear dynamics literature has identified as a precursor to a particular type of discontinuous transition (Saleur and Sornette 1996). Engineering works have empirically linked this vibrational pattern with the impeding breakdown of electromechanical equipment (Anifrani et al. 1995). More interestingly, another string of works from econophysics has focused on using the same oscillatory pattern to predict major changes in a social system: financial asset prices driven by trader behaviour (Sornette et al. 1996). The example demonstrates how nonlinear dynamical system theory can contribute to harnessing weak signals in engineering as well as social sciences.

In the present report we categorise literature on weak signals in four broad literary groups and review each group in an individual section. The sections are: Organisational Science (section 2); Futures & Strategic Foresight (section 3); Signal Processing & Engineering (section 4); and Nonlinear Dynamics (section 5). Each of these sections summarises the main research questions, definitions, insights, and challenges found within the respective literary group. Notably, two prominent sub-groups were also identified: weak signal detection for GPS applications, and stochastic resonance for weak signal detection. These two clusters are discussed in sections 4 and 5, respectively.

Section 6 elaborates on the connections between the aforementioned four literary strands and presents a list of the overarching challenges. Doing so naturally produces a set of unified definitions for weak signals and other associated concepts.

In the final section, we discuss implications of the proposed framework, by examining the reviewed literature under the lens of the newly proposed definitions and discussing how new interdisciplinary research avenues may be identified.

Throughout the following pages, we will be using the term ‘environmental scanning’ to refer to the task of finding weak signals. This term is borrowed from organisational sciences, the domain of managerial sciences that arguably has the most pedagogical terminology for weak signals and their associated concepts. The term ‘wildcard information’ is borrowed from futures literature to indicate crucial information that can be revealed via successfully detecting and interpreting weak signals.

Organisational science

Organisational science focuses on ‘the examination of how individuals construct organizational structures, processes, and practices and how these, in turn, shape social relations and create institutions that ultimately influence people’ (Clegg and Bailey 2007, pp xliii). Here, we first discuss how weak signals are defined in this field, and then we examine two different approaches to harnessing weak signals.

In organisational science, the topic of weak signals was first discussed in Ansoff's seminal work, where he argued that traditional forms of forecasting that rely on historical data were incapable of predicting strategic surprises, described as 'sudden, urgent, unfamiliar changes in a firm's perspective which threaten either a major profit reversal or the loss of a major opportunity' (Ansoff 1975, pp 22). Ansoff's underlying assumption is that strategic surprises do not emerge without warning. Based on this assumption, which has been explored in subsequent works (Kaivo-oja 2012; Seidl 2004), managers may have the opportunity to scan their environment for such warnings, and anticipate strategic surprise. In this context, weak signals are defined as 'first symptoms of strategic discontinuities; they are symptoms of possible change in the future' (Holopainen and Toivonen 2012, pp 200.

Ansoff's fundamental assumption is that strategic surprises do not emerge without warning, and in the following sections we will see that nonlinear dynamics is in partial agreement: certain dynamical systems produce precursors before they undergo a change of regime. Certain works have focused on exploring and formalising this connection from a theoretical standpoint, by using a constructivist approach to map phenomenology of firm dynamics to statistical mechanics (Kaivo-oja 2012; Meyer et al. 2005). Other works have developed methodology based on formal methods from nonlinear systems in order to tackle practical weak signal-related problems in organisations. The latter category primarily draws from statistical mechanics (Levy 2008) or graph theory. For example, a network science based method to detect lone wolf terrorists through social media activity has been presented in (Brynielsson et al. 2013). A more general approach to network science methods for weak signal detection in social networks can be found in (Mossel et al. 2015).

Ansoff considered firms, and outlined the challenges of detecting and aggregating environmental information to extract the predictive power of weak signals. He also emphasised how acting on weak signals fundamentally differs from acting on strong signals (Ansoff 1980). This framework has been expanded in subsequent works, which more thoroughly discuss the detection and aggregation of environmental information (see Welz et al. 2012; Ilmola and Kuusi 2006; Narchal et al. 1987; Daft et al. 1988; Schoemaker and Day 2009; and Amanatidou and Guy 2008). The result is the de facto framework or the study of weak signals across all managerial and social sciences. Arguably, it also is the most intuitive and pedagogical exposition of the weak signal phenomenology — and that is why we chose to use most of its terminology for this report.

To enable organisations to harness the predictive power of weak signals, Ansoff suggested the formation of strategic early warning systems (SEWS) — top-down structures that empower executives (Reinhardt 1984). Remarkably however, a second approach has been put forward from organisational science, based on field observations from high reliability organisations (HROs). HROs have been able to effectively deal with strategic surprise by interpreting weak signals, without a SEWS. Instead of a top-down system, HROs rely on a highly effective form of distributed cognition, which allows for the detection and collective interpretation of weak signals.

Strategic early warning systems

The concept of environmental scanning was discussed by Aguilar (1967), as a methodology to search a company's environment for all relevant strategic information. Ansoff addressed this challenge in the context of weak signals using the notion of filters: conditions that potential weak signals have to adhere to in order to reach the strategic planning process (Ansoff 1975). Since then, a large body of works has developed, exploring the use of textual analysis for the detection of weak signals in social media. Prominent examples can be found in: predicting rioting through Twitter activity using natural language processing and classification methods (Charitonidis et al. 2015, 2017). Gaining business insights by detecting weak signals from online sources has also been explored by applying traditional natural language processing (NLP) methods on tweets (Pépin et al. 2017; Kayser and Bierwisch 2016), news (Yoon 2012; Thorleuchter and Van den Poel 2013), and mixed textual sources (see Mühlroth and Grottke 2018; Song, 2012 for a thorough review of techniques and applications of weak signal mining from online sources).

Although the organisational science literature has explored the challenge of environmental scanning, much less emphasis has been given to a closely related question: how to define the scope of this environmental search. Should we consider only media directly relating to the company (customers, employees, etc.) or set a broader scope including competitors, legislators, news outlets, etc.? This question has been more extensively explored in the futures literature.

In Ansoff's framework, the weak signals acquired by scanning one's environment need to be meaningfully aggregated to reveal useful information. Ansoff suggested that when aggregating weak signals, companies may rely on expert judgment and past experience. The resulting interpretation will be gradual: 'when a threat/opportunity first appears on the horizon, we must be prepared for very vague information, which will progressively develop and improve with time. This progression may be characterized by successive states of knowledge' (Ansoff 1975, pp 24). In stark contrast, strong signals have a clear interpretation and impact as soon as they appear. A number of methods have been suggested for aggregating weak signals, relying either on the judgement of experts (Rossel 2012), the DELPHI method (Nowack et al. 2011), the judgment of an expert aided by computer software (Rouiah and Ould-Ali 2002; Lesca and Lesca 2011; Pépin et al. 2017), or prediction markets (Graefe et al. 2010). Remarkably, human involvement is a vital part in all of this literary strand. A review of methods using automated analytics to aid experts in the interpretation of weak signals can be found in (Eckhoff et al. 2014).

The final challenge considered by organisational science in the context of weak signals is organisational flexibility: a set of principles and practices that can enable an organisation to adapt to environmental discontinuities (Bessant et al. 2005; Mendonça et al. 2004). However, the focus of this report is on the detection and aggregation of weak signals, and not on organisational strategy — therefore this particular topic will not be discussed here. If the reader is interested in this topic, more can be found in Aaker and Mascarenhas (1984) and references therein.

Organisational mindfulness and sensemaking

Interestingly, another branch of organisational sciences has provided an entirely different take on detecting and aggregating weak signals. By empirically studying the practices of high reliability organisations (nuclear plants operators, aircraft carrier crews, etc.), Weick was able to derive five organising principles that enabled groups to avert catastrophes by effectively filtering, aggregating, and reacting to information through a distributed cognitive process dubbed 'sensemaking' (Weick 1987). This empirical work gave rise to literature that studies how organisations can survive or even thrive in the presence of strategic surprise (Lengnick-Hall and Beck 2016). In contrast to the previously discussed methods, this approach does not consider a SEWS system specifically built for exploiting weak signals. Instead, sensemaking is a distributed cognitive process, embedded within the organisation (Weick 1995).

For effective sensemaking to take place, individuals at the front lines of the organisation must be free of cognitive biases and actively look out for crucial information (Weick 1995). This information must then be communicated to those who are best suited for interpretation (Weick 1995). Again, this information-routing process must be free of cognitive biases; for example, experts should be given crucial technical information first, instead of supervisors (Weick 1988). Over time, individuals learn who is best suited for which kind of crucial information, and the organisation develops a property dubbed 'transactional memory'. This process allows for the collective intelligence of the organisation to be tapped, and for anomalies to be timely detected and interpreted. In this way, sensemaking can enable organisations to deal with strategic discontinuities, and thrive in turbulent environments (as demonstrated by studies in healthcare (Thomas et al. 1993), marketing (Neill et al. 2007), and humanitarian aid (Stephenson 2005). Therefore, in contrast to the top-down Ansoff-ian early warning systems, sensemaking can be understood as a bottom-up approach to tackling the challenges of weak signals.

An example of the value of transactive memory can be found in the development of the open-source operating system Gentoo, which depends on a network of collaborating developers. Once a software issue appears, the respective expert in the network is informed and mobilised to settle it. Initially, the process of routing the issue to the expert occurs in a distributed fashion: the issue is passed from one developer to the next until the expert is found. In this way, transactive memory is formed: the individuals within the organisation slowly learn who is an expert on what, and become adept at this process. Eventually, a single specialist may be tasked exclusively with routing issues, in order to accelerate this process. This approach results in a boost in terms of productivity but robs individuals of their transactive memory. The importance of transactive memory becomes apparent when the routing specialist leaves the organisation and performance plummets (Zanetti et al. 2013).

Effectively detecting and routing crucial information within the organisation can be a challenge, and to aid individuals, the concept of mindfulness has been proposed. Mindfulness has been described as the psychological process of purposely bringing one's attention to experiences occurring in the present moment without judgment (Baer 2003). Empirical and theoretical works have demonstrated how mindfulness can enable sensemaking and empower organisations to deal with weak signals through the emergence of collective cognition. An example on how the lack of mindfulness could result in catastrophes can be found in the Bhopal disaster in India, an event where a series of mishaps led to the explosion of a pesticide factory and to the deaths of thousands (Broughton 2005). According to Weick the root cause of the disaster was lack of mindfulness: ‘The breakdowns included a low standard of plausibility, minimal doubt, infrequent updating of both mental models and current hunches, and mindless action’ (Weick 2010, pp 532). Broadly speaking, mindfulness can be understood as an inquisitive, highly perceptive mental state.

It should be noted that certain SEWS rely on either using panels of experts to collectively interpret weak signals, or collecting information from crowds. Therefore, these systems rely on distributed cognitive processes, which means that the boundary between the sensemaking and SEWS approaches is not a clear line.

Futures & strategic foresight

Futures literature, or futurology, is a field of sociology that dwells on postulating possible, probable, and preferable futures for society as a whole (Sardar 2010). Strategic foresight has a similar scope, but focuses on producing planning-oriented, actionable insights (Vecchiato 2012). Both futures and strategic foresight dwell on the practices and methods that enable the early detection of emerging discontinuities in social systems. They also focus on interpreting the consequences of discontinuities from the perspective of various decision makers (firms, governments, policy makers, etc.). The ultimate goal is to inform future courses of action in order to ensure the long-term survival and success of societal structures.

Both these fields have a large body of work dedicated to weak signals, and rely on the framework of Ansoff. Some works have suggested that a slightly different weak signals framework needs to be developed specifically for futures and strategic foresight (see (Holopainen and Toivonen 2012, and references therein for a review of such works). However, the suggested alterations to Ansoff’s framework are either minor or entirely semantic — and thus of no interest to the current report.

Peripheral vision, information overload and analytics

A prominent challenge in identifying weak signals is determining the scope for environmental scanning (Hiltunen 2008). This challenge was not explored in depth in Ansoff’s work, which vaguely suggested focusing on organisations and individuals intimate to the company (Ansoff 1975, 1980). Works from the futures literature argue that by doing this, the scope of the search can become narrow, resulting in missing weak signals — a condition dubbed lack of ‘peripheral vision’ (Day and Schoemaker 2008, pp 43). The same problem — lack of peripheral vision — can result when introducing overly strict filters for weak signals (Ilmola and Kuusi 2006). On the other hand, selecting an overly broad scope of search can result in information overload, fatiguing the decision maker and potentially increasing the difficulty of sensemaking (Schoemaker et al. 2013). In futures and strategic foresight, determining the scope of search for weak signals relies on empirically derived rules, a summary of which can be found in Hiltunen (2008, 2007). Additionally, in many cases the scope is overly narrow because of limited data accessibility. The recent boom of social media has helped in solving this problem: social media are now among the most prominent avenues for monitoring and observing social systems.

In the case of textual data from social media, machine-learning methods have been considered to limit the scope of the search for weak signals (Sriram et al. 2010). These methods rely on a combination of natural language processing, manual labelling, and supervised classification algorithms to classify textual content in various categories, such as relevant/irrelevant, warning/emotional, support/complaint, etc. Such applications can be found in Imran et al. (2013), in filtering disaster-related tweets; and Culotta (2010), who focuses on epidemic-related tweets. The detection of small, localised events from searching the entire throughput of Twitter for an entire city are considered in Suma et al. (2017). A combination of natural language processing and machine learning is used for the detection of weak signals in generic online textual data sources in Thorleuchter and Van den Poel (2013). Atefeh and Khreich (2015) review numerous works on event detection by monitoring Twitter. Additionally, Sándor (2009) considers natural language processing in monitoring textual streams and automatically detecting topics where discourse is emerging. The same challenge in identifying relevant social media content has been explored for non-textual information, but to a far lesser extent. Non-textual data has been considered primarily with images. Here, Alam et al. (2018) explore the classification of disaster-related images in tweets, and Gupta et al. (2013) propose an algorithm for automatically identifying fake images tweeted during Hurricane Sandy in the United States. Most of the afore-cited methods employ the following approaches: i) first, all possible social media content is accumulated, ii) then a fraction of the accumulated content is manually labelled (e.g., as relevant/irrelevant), and iii) finally a supervised classification algorithm is trained based on the manual labels.

Besides social media platforms, other potential sources for weak signals in futures and strategic foresight studies are: patents; news articles; individuals (policy makers, expert scientists, governmental employees, journalists, CEOs); crowdsourcing; online communities: and groups of amateurs with an apparent talent for anticipating developments in specific fields, dubbed 'expert amateurs' by Hiltunen (2010, pp 108).

After the content produced by online media has been classified as relevant or irrelevant, weak signal detection and interpretation follow. For the detection and interpretation of weak signals, methodological works from this literary stream use identical approaches to the ones from organisational sciences: they rely on expert judgement and methods such as DELPHI (see previous section for citations). That being said, works from futures and strategic foresight consider weak signals on a broad social scale, instead of a company-wide scale.

Signal processing & engineering

Uncovering information buried under excessive noise is the main challenge in numerous engineering applications (Wang et al. 2013), e.g., detecting objects on radar, performing early diagnosis of upcoming errors in electromechanical equipment, or identifying minor artefacts in medical imaging data. Works found here essentially consider mathematical denoising techniques that are applied to quantitative data. These techniques are routinely used in engineering and signal processing, resulting in practical benefits such as reduced downtime, aversion of catastrophes, and optimised operations.

Generally, these methods are designed to detect a specific predefined signal obscured by noise, and therefore, a description of a weak signal must be available for the methods to be used. Examples of weak signal descriptions include specific vibrational patterns or nonlinear trends. These descriptions or signatures are studied in a second smaller body of work, which uses either experimental or formal methods to discover and describe indicators of crucial events (e.g., specific vibrational patterns that allow to foresee the breakdown of an electric motor). These signature indicators are referred to as 'weak signals' in this literary group. The successful automation of weak signal detection and interpretation in this literary strand stands in stark contrast with efforts from the social sciences, which all rely on expert judgement and manual work.

Dealing with low signal-to-noise ratio

Generalised methods that can be used to detect specific weak signals across different domains have been developed. The following is a non-exhaustive list of the most prominent of these methods:

- One example of such a method makes use of **stochastic resonance**, a phenomenon that originates from biology (Benzi et al., 1981) and will be explored in the next section. In short, some nonlinear, multistable systems may become more sensitive to low amplitude signals in the presence of noise, and this property has been used to detect weak signals, as demonstrated in Hari et al. (2012), Lei et al. (2013), Wiesenfeld and Moss (1995), and Asdi and Tewfik (1995). Applications include non-invasive gearbox diagnostics (Lei et al. 2013), fault diagnosis of rotating machinery (Li 2005), and sensitive magnetometers (Wiesenfeld and Moss 1995). A general method for stochastic resonance based weak signal detection is presented in Asdi and Tewfik (1995). Additionally, reviews on specific nonlinear oscillators that rely on stochastic resonance for weak signal detection can be found in Wang et al. (1999), and Jalilvand and Fotoohabadi (2011).
- **Fast Fourier Transforms** (FFT) have been successfully used for: the detection of weak signals in detecting sustained tachycardia from electrocardiogram data (Cain et al. 1984), and fault diagnosis in arc fault detection (Wang et al. 2014).
- **Wavelet analysis** is 'a mathematical technique which can decompose a signal into multiple lower resolution levels by controlling the scaling and shifting factors of a single […] function' (Liu and Menzel 2016, subsection 3.1 par. 1). The efficacy of wavelet analysis for weak signals detection has been demonstrated for predictive diagnostics in motor bearings (Qiu et al. 2006; Wang et al. 2015), and radar (Ehara et al. 1994) applications.
- **Blind source separation** is an unsupervised learning technique that allows the decomposition of a single mixed signal into its constituents through a Hebbian-inspired (adaptive) rule (Comon and Jutten 2010). This method has been applied to predict the onset of Alzheimer's disease from electroencephalogram data (Cichocki et al. 2005; Solé-Casals and Vialatte 2015).
- Machine learning can also be used in weak signal detection. A **manifold learning method** was used in combination with **wavelet analysis** for the detection of ball bearing faults in rotating machinery (Wang et al. 2015) A recent review of works using unsupervised learning for the early prediction of component failure can be found in

(Amruthnath and Gupta 2018). The challenges of using machine learning for predicting weak signals indicative of cybersecurity threats in IoT (Internet of Things) applications have been summarised in Abeshu and Chilamkurti (2018). Weak signals in cosmological data have been captured using deep learning techniques (George and Huerta 2018; Haiman 2018). Generally, a major challenge of using machine-learning methods for weak signal detection is sample imbalance, as we typically have more samples of non-weak-signals.

- Finally, a number of works from this literary group dwells on weak signal detection for **satellite positioning systems** applications, such as GPS. These works all have a very clear, common goal and share methodological similarities. All approaches developed in this literature are specific to GPS signals and have not found broad application in other domains. These methods will be discussed at the end of the current section.

In part, the success of weak signal applications in engineering seems to stem from the luxury of experimentation. One can typically obtain more data on an engineering application by running lab experiments. Data helps, either in identifying clear signatures of important events, or by giving a better description of the engineered system in normal operation. Social sciences do not have the same luxury as experiments are harder or even impossible to run. However, the data explosion driven by social hyperconnectivity has helped ameliorate this issue.

Satellite positioning systems

Broadly speaking, the receiver of a satellite positioning system must collect and then convert signals from satellites into measurements of position, velocity and time. This information is acquired from at least four satellites, in order to determine the user position. The operation of a receiver involves three major tasks: detection of satellite signals, tracking of the signals, and using the obtained information to estimate the user position. Satellite positioning systems have been used for many years in numerous fields (including air, naval and automotive navigation), but their use in complex environments (e.g. cities, indoors, dense forests, inside tunnels) is still limited. This is because satellite positioning signals are very faint, as the power of the satellite emitter is spread over a large fraction of the globe. Additional problems arise as well: the signal is subject to electromagnetic noise, the signal may be reflected and reach the receiver in more than one path (each of different length), and atmospheric conditions may influence the speed of the signal. Various studies demonstrate the significant degradation of satellite signals when subjected to this type of environment (Melgard et al. 1994; MacGougan et al. 2001; Seco-Granados et al. 2012).

Methods developed to address satellite positioning signal detection can broadly be classified into three categories: the use of complementary radio frequency signals, the use of advanced signal processing techniques, and the use of external sensors (Andrianarison 2018). Most of the proposed methods in the recent literature are intended for software-defined radio (SDR) receivers, due to their flexibility for research and development. SDRs are well suited for research, since all-software signal processing algorithms (whether acquisition or tracking) can easily be controlled, and there is no need to change hardware when adding or modifying algorithms.

All of the following techniques can increase the sensitivity of the receiver to satellite signals at the expense of either: computational power, energy and time (Andrianarison 2018). That is to say that the presented techniques can improve the sensitivity of the receiver by either increasing power consumption, increasing the time it takes to acquire the signal, or utilising more hardware. Many applications impose limits on power, computational resources, and time. For example, a mobile phone user has access to limited computing power, and would rather increase battery life than increase the accuracy of GPS by a few metres. Therefore, works on weak signal detection for satellite positioning systems focus not only on detecting the weakest possible signal, but also on making the detection process as computationally efficient and fast as possible.

Temporal integration methods

Coherent and non-coherent integration are temporal averaging techniques, which exploit the repetition of the satellite signal in time. By taking temporal averages, temporally uncorrelated noise is reduced. The two methods differ in that coherent integration averages the raw signal, while non-coherent integration averages the squared signal. Generally, longer integration windows increase the sensitivity of the receiver, at the cost of longer acquisition time (more details can be found in DiEsposti (2001)). However, the structure of satellite signals imposes an upper limit on the coherent integration window. Non-coherent integration is designed to aid coherent integration, by expanding the maximum window of the latter. In this way, non-coherent integration increases the sensitivity of the receiver when the coherent integration reaches its limit (Borio et al. 2009). While longer non-coherent integration windows allow to increase the coherent window further, it also introduces a type of distortion known as ‘squaring loss’. Thus, a compromise between the two effects must be found (Borio et al. 2009). The temporal limits of integration windows are the result of the structure of the signal broadcasted by satellites, and newer systems (like the Galileo) have changed their signal in a way that allows for longer integration windows.

Assisted positioning

Increasing the efficacy of satellite positioning systems is also possible by using an external reference station to provide additional information to the receiver. This family of techniques is called assisted GPS, and was first introduced in Taylor and Sennott (1984). The external reference station may broadcast information on the position and speed of satellites to aid the receiver. In order to infer the position and speed of a satellite, the unaided receiver needs to scan the entire space for [position x speed] combinations. In contrast, an aided receiver can a priori reduce this search space based on the information provided by the reference station (Van Diggelen 2009). The reduced search space allows for more effective usage of the receiver’s computational power.

A pedagogical overview of additional methods on weak signal acquisition for satellite positioning systems can be found in Andrianarison (2018).

Nonlinear dynamics

In a broad sense, nonlinear dynamics study the behaviour of systems that respond to inputs in a disproportionate manner. There are two literary strands from nonlinear dynamics related to the topic of weak signals. While both strands take a theoretical approach and rely on formal methods, they focus on different aspects around weak signals and consider the concept differently.

Generating mechanisms & precursors

The first strand considers two major challenges: i) uncovering the origins of discontinuities in dynamic systems, and ii) determining whether such events are associated with specific precursors. These precursors can serve as weak signals that foretell of the impending discontinuity (see a more in-depth explanation in Huang et al. (1998)). Indeed, after such precursors are identified within the nonlinear dynamics literature they may be exported towards engineering to serve as signatures of weak signals (Gould and Eldredge 1977; Sornette 2002; Anifrani et al. 1995; Sornette et al. 1996).

Self-organised criticality

Regarding the first challenge (generating mechanisms of discontinuous events), let us consider an example from evolutionary biology: based on fossil records the evolution of species was found to be described by long periods of non-adaptation (stasis), separated by short bursts of massive cascades of adaptation (cladogenesis) (Gould and Eldredge 1977). The origins of this non-trivial property of the evolutionary process puzzled evolutionary biologists for decades, until eventually an explanation was offered by the theory of self-organised criticality, or SOC (Bak and Sneppen 1993). SOC is a property of certain systems to spontaneously 'tune' themselves towards a critical state. Here, critical implies a state characterised by divergent correlation length or susceptibility (Bak et al. 1988). More intuitively, a divergent correlation length means that seemingly local changes can have far-reaching consequences, while divergent implies that small perturbations may result in system-wide events.

While the concept of SOC originated from the statistical mechanics literature, it has been identified as a generating mechanism of discontinuous change across disparate systems, including: geophysics (Sornette and Sornette 1989), financial markets (Stauffer and Sornette 1999), sociology (Weisbuch et al. 2001), environmental sciences (Pueyo 2007), and neuronal dynamics (Arcangelis et al. 2006; Georgiadis and Sornette 2019; Hesse and Gross 2014).

Models exhibiting SOC can provide a formal explanation of the origins of weak signals in certain systems. System-wide events birthed in systems exhibiting SOC may be perceived as environmental discontinuities by observers. In certain models, such events may exhibit precursory behaviour. In such cases, precursors can serve as weak signals, making the identification of precursors and the development of methods for their detection a topic of study (Huang et al. 1998; Amundsen et al. 2012).

Phase transitions

Phase transitions are phenomena describing physical systems undergoing discontinuous change (most notably, change between states of matter). While originating from physics, phase transitions have been successful as analogues in other fields as well. Notably, phase transitions have been put forward as models of discontinuous change in complex systems such as social networks (Fronczak et al. 2007), finance (Levy 2008), cortical tissue (Beggs and Plenz 2003), and technical infrastructures (Carreras et al. 2004), among others.

To provide an informal explanation of a phase transition we may consider a system associated with quantity measuring order (for example, the density of liquid), and a control parameter that we are allowed to vary slowly (such as temperature). Assume that we slowly increase the control parameter, until it reaches a characteristic critical value — at which we will observe a discontinuous change in either the order parameter or one of its derivatives (with respect to the control parameter).

Certain types of phase transitions are associated with specific precursors (for example, the divergence of the order parameter as we approach the critical control parameter value). When using phase transitions as metaphors for complex systems, we may attempt to predict discontinuities by scanning the environment for the respective precursors. Such precursors may be spatiotemporal — such as the percolation a network extracted from the system's correlation matrix (Rodríguez-Méndez et al. 2016); or temporal — such as a super-exponential trend leading to singularities in finite-time (Yan et al. 2010).

An in-depth explanation of phase transitions and their applications across domains is out of scope for this report. The interested reader may see cited works for more information.

Bifurcation theory

Consider a dynamical system undergoing a gradual change of some of its parameters. In certain systems this gradual change may cause a sudden shift of its dynamical qualities: the stability of equilibria might change (e.g. from stable to unstable), the equilibria or limit cycles may vanish or spontaneously appear, and limit cycles may become equilibria or vice versa (Strogatz et al. 1994). Bifurcation theory looks at such behaviour through the use of abstract analytical models, and has identified a number of generating mechanisms for discontinuous dynamics (Iooss and Adelmeyer 1998). The parsimonious models studied in bifurcation theory provide us with formal examples of how a seemingly 'weak' signal can have drastic effects on systemic behaviour. These examples can be used as powerful analogies to numerous applications, and thus bifurcation theory is the theoretical basis for many weak signal detection methods found in engineering, such as nonlinear oscillators and stochastic resonance. A review of different bifurcation types is out of scope for this report, and the interested reader may consult Iooss and Adelmeyer (1998) and references therein. For a more pedagogical introduction we recommend Strogatz et al. (1994).

The perspective of risk management

Besides exploring the mechanisms of discontinuous change, works in nonlinear dynamics focus on investigating the origins of precursory signals, and developing tools for their early detection. Whether a specific discontinuous event is associated with a precursor or not can be a major debate in risk management (Sornette 2009) and generally depends on the mechanism giving birth to the discontinuity. For example, in some cases of SOC large-scale events do not produce precursors, and therefore cannot be anticipated — while discontinuities associated with second-order phase transitions generally produce precursors.

Discontinuities that cannot be anticipated due to lack of precursors, may give rise to Black Swan events. In such cases, scanning one's environment for weak signals is pointless — it is better to prepare for adversity rather than try to predict it (Taleb 2007). In contrast, discontinuities with the potential to be predicted (that is, accompanied by precursors) have been dubbed ‘Dragon Kings’ and have been identified in numerous domains including financial markets, nuclear accidents (Wheatley et al. 2017), financial flash crashes (Johnson et al. 2012), and electricity prices (Janczura and Weron 2012). In such cases, investment in detecting precursors that allows the prediction of a discontinuity may be an optimal strategy.

Generally, what makes Dragon King events predictable is that they are produced by dynamics different from their less impactful counterparts (Sornette and Ouillon 2012). Thus, Dragon Kings are typically encountered as outliers of some distribution, and statistical tests to identify such outliers have been developed (Sornette and Ouillon 2012; Pisarenko and Sornette 2012). Using such tests may provide empirical evidence on whether the outlier events are: i) unpredictable extremely low probability events (Black Swans), or ii) generated by a fundamentally different mechanism and thus potentially predicable (Dragon Kings). Another approach to establish the predictability of a particular type of discontinuity is via assembling a bottom-up model of the system in question and running simulations. The results can then be explored for the detection of specific signatures prior to extreme events (for example using superposed epoch analysis). Typical precursors to Dragon Kings include increasing temporal autocorrelation in the system dynamics, increasing spatial correlation, and diverging variance (Sornette and Ouillon 2012).

Stochastic resonance

A second strand of works from nonlinear dynamics looks at weak signals in another context, by exploring how some input of small magnitude (referred to as a weak signal) can cause macroscopic change in a large system. Therefore, here the term ‘weak signal’ indicates the disproportionality between cause and effect.

Stochastic resonance is a prominent example of a mechanism that allows weak input to cause large-scale change (Benzi et al. 1981). Multistable systems that are exposed to noise tend to amplify the effect of a weak signal for a finite level of noise, as indicated by the term ‘stochastic resonance’. A pedagogical and thorough explanation of stochastic resonance can be found in McDonnell and Abbott (2009).

For an intuitive, informal description of stochastic resonance we can think of a multistable dynamical system, with two stable states A and B. These two points lie far away from one another (that is, they describe very different states of the system). Assume that A has a lower potential than B, and that the system starts off resting at A. Next, this system is subjected to a weak input, which alters the potentials of the system’s equilibria. Now, state B has a marginally lower potential than A. However, A is a stable equilibrium, and the system cannot escape the local region of stability surrounding A. Assume that we now apply noise to the system. The noise may cause the system to escape the local area of stability near point A, and jump to point B. This will be perceived by us as a radical change in the system’s state. This mechanism can be used to amplify the weak input provided to the system, aiding in weak signal detection.

A more tangible example can be found in the study of the Earth’s climate. The slight wobble in the Earth’s trajectory around the sun causes a change in the Earth’s received solar energy by

0.1%. Nevertheless, this seemingly minuscule signal has been proposed as a potential driver of ice ages (Wiesenfeld and Moss 1995).

Stochastic resonance has also been observed in biology (Hänggi 2002) — for example, to increase the sensitivity of sensory systems. Additionally, stochastic resonance has been used in signal processing and engineering to design sensors capable of detecting signals in low signal-to-noise environments. See in the respective section (Dealing with low signal-to-noise ratio) examples of uses of stochastic resonance in engineering and signal processing.

Unified framework

In this section we argue that the four literary strands from the previous section describe the same general challenge, and that differences between the four frameworks are the result of each strand focusing on a different aspect of this common challenge. We present a unifying framework that consolidates the approaches of all four literary strands and discuss how the proposed unified framework can account for the differences in the definition of weak signals across literatures. We then mention interdisciplinary works on weak signals that focus on the import/export of methods between the four identified literary strands, and illustrate how the proposed unified framework can also account for the import/export relationships between the four literary strands.

Universal challenges in harnessing weak signals

Let us briefly restate the concept of weak signals, as defined in each of the four literary strands. Futures, and organisational sciences utilise approximately the same definition — perceiving a weak signal as information that is hard to spot and to interpret, but stands to offer vital predictive power. Engineering and signal processing consider a weak signal as a signal in a low signal-to-noise ratio environment. Finally, nonlinear dynamics systems describe a weak signal as a signal of relatively low amplitude that can alter the macroscopic behaviour of some large system. We propose the following unifying definition:

Weak Signal: Information that is i) non-trivial to obtain and/or interpret, but ii) may offer mission-critical predictive power if obtained and properly interpreted.

According to this definition, what constitutes a weak signal depends on the goals or mission of an actor: hard-to-spot signals with far-reaching consequences that are not relevant to that mission are not weak signals from the perspective of the agent.

As previously stated, the identification of signatures is a major focus of nonlinear dynamics — with some examples being critical slowing down, increasing spatial or temporal correlations, and diverging variance. We define a weak signal signature as:

Weak Signal Signature: A quantifiable property of signals that is characteristic of weak signals.

The quantifiability mentioned in the definition above directly implies that weak signal signature can be used to aid in the search of weak signals in data.

Aided by these definitions we may identify the fundamental challenges shared across literature.

Challenge 1 - Allocating attention

According to futures, the first step to harnessing weak signals is determining the potential sources to be scanned for them. We will refer to this challenge as attention allocation. As existing futures literature has shown, setting too broad a scope can cause information overload, making the search unnecessarily hard, while setting a narrow scope can result in missing weak signals. This challenge is dealt with very differently across the literature. For example, signal processing and engineering merely focus on all measurements coming from the system of interest (radar, motor, MRI sensors, or otherwise). However, in futures the scope is determined based on the opinion of experts, and the available data sources. Generally, it appears that the less coupled a system is with its environment, the easier it is to determine the scope of the search. Finally, works from organisational science and nonlinear dynamics can help us uncover new relevant sources that should be included in our scope. For example, the different colours of sparks in asynchronous electric motors have been empirically associated with different malfunctions. This nontrivial connection reveals an additional feature of interest (spark colour). Thus, the colour of sparks (a feature that appears unimportant) should be included in the scope of weak signal search.

Challenge 2 - Detecting weak signals

All literary strands agree on weak signals being non-trivial to detect. Signal processing and engineering have developed mathematical methods for detecting a specific weak signal signature by denoising numerical data, including spectral methods, filtering, and chaos oscillators (for more details on the methods see the section on signal processing methods). Research on satellite positioning systems has developed temporal averaging methods (see coherent integration and non-coherent integration in section 4.2) that allow the detection of signals below the thermal noise threshold. Similarly, works from nonlinear dynamics that focus on the sensory reception of biological organisms have described how the phenomenon of stochastic resonance can be used to detect signals buried under excessive noise (see section 5.2). Furthermore, futures and organisational sciences consider media produced by social systems and rely on semi-automated methods for weak signal detection. Such methods typically utilise textual analysis, and present the results to appointed experts for manual assessment via the DELPHI method, or majority voting (for details, see section 3). Finally, the sensemaking literature from organisational sciences recommends the adoption of specific practices that enable organisations to identify weak signals in a bottom-up fashion — including deference to expertise (that is assigning tasks based on expertise and not on rank or status); preoccupation with failure (actively considering how processes and systems could fail and taking proactive measures); and reluctance to simplify (preventing the simplification of complex problems by actively challenging long held beliefs). Further details can be found in section 2.3. Before we move on to the next challenge, it should be emphasised that the detection of weak signals can be greatly aided if we have available a description of the weak signal at hand. Having such a description transforms the problem of weak signal detection from scanning for any sort of irregularity to scanning for a specific pattern. We will discuss the challenge of obtaining such descriptions in the following paragraphs (Challenge 5).

Challenge 3 - Interpreting weak signals

Weak signals need to be interpreted to reveal the respective wildcard information. The interpretation tasks are handled differently based on: i) the number of historical occurrences of a specific weak signal and ii) whether the given system is changing in time or not. Signal processing and engineering focus on detecting specific weak signal signatures that have already been associated with specific outcomes (e.g., a specific vibrational pattern of an electric motor may imply a specific defect in the motor's rotor). In such cases, the weak signal interpretation task has been dealt with beforehand: either through formal methods or through experimentation where a given signature has been associated with a given wildcard. In contrast to engineered systems, social systems change in time, and tend to be harder to model mathematically. Consequently, the task of interpretation cannot be resolved beforehand. Instead futures and organisational sciences rely on a combination of i) past incidents and ii) expert judgement to interpret weak signals. The sensemaking literature suggests another alternative: adopting certain practices that allow the collective knowledge of the organization to be utilised in the interpretation of a weak signal.

Challenge 4 - Exploiting wildcards

In managerial literature, once the wildcard information has been obtained, a decision maker needs to determine how it will affect decision making. Here, the adaptability of the organisation managed by the decision maker is crucial. Methods and practices that increase adaptability are a topic of study. However, this current report focuses on the detection and interpretation of weak signals and will not explore this challenge.

Challenge 5 - Discovering weak signal signatures

The search and interpretation of weak signals (Challenges 1, 2 and 3) can be aided if we possess some description of such signals. For the case of quantitative data, literature on nonlinear dynamics has provided a number of such descriptions in the form of specific patterns in space or time (critical slowing down, diverging variability, etc.). These descriptions can be formulated mathematically as temporal or spatial trends, making the search of weak signals a trend-fitting task — for which advanced mathematical tools have been developed.[1] Additionally, identifying signatures aids in interpreting weak signals, since certain signatures have been associated with specific outcomes, e.g., diverging susceptibility is linked with a second-order phase transition. In contrast, social science literature has not identified such signatures. Instead, works from social sciences identify either clusters or anomalies in the data as potential weak signals, and provide them to an expert for evaluation. Additionally, there is no clear link between weak signals and specific outcomes. The interpretation of the former is left entirely up to the judgement of the expert.

Challenge 6 - Understanding the origins of discontinuities

Nonlinear dynamics has identified a number of mechanisms that generate discontinuities, and for some of them certain precursors have been described. Therefore, understanding the mechanisms generating the said discontinuities may allow us to draw from the nonlinear dynamics literature and to identify precursors that may serve as signatures for weak signals (see paragraph for Challenge 4).

[1] In many applications, spatial trends may be strong signals (very obvious, with clear interpretation). However, spatial trends can also constitute weak signals (nontrivial to detect and to interpret). For example, an increasing number of parked cars at retail stores (viewed by satellite) is predictive of increased revenues for the supermarket chain (Melgard et al. 1994)

An intuitive overview of challenges

In this paragraph we provide an intuitive, abstracted hypothetical scenario in order to unify the previously mentioned six challenges. The purpose of this example is to aid the reader in digesting the previous paragraph, by providing a high-level description of the challenges in a concrete setting. Let us assume that we are quality managers in a haystack factory. We have been tasked with: i) cleaning a number of haystacks from all foreign objects (such as nails, needles, screws, and other objects), and ii) understanding what is the origin of these foreign objects so that we can take preventative measures in the future.

1. **Challenge 1:** Determining where to search for the foreign objects. Should we only search some of the haystacks, or perhaps only the tops of haystacks?
2. **Challenge 2:** Determining the search method. For example, are we to sift through the hay manually, stalk by stalk, or can we use machinery to accelerate the process?
3. **Challenge 3:** Relying on our findings in order to understand the origins of the foreign particles. For example, assume that we found a few screws in the same haystack, which happens to be right next to an old conveyor belt. In that case, we may assume that the screws came loose from the conveyor belt and found their way to the haystacks.
4. **Challenge 4:** Relying on our findings, it is reasonable to assume that the conveyor belt is missing some screws and that it is likely to fail soon. The high probability of the failure of the conveyor belt is our wildcard information, which we may now exploit by thoroughly maintaining the conveyor belt.
5. **Challenge 5:** Determining a common feature of all foreign objects. In our example, we may notice that all foreign objects are made of metal. This observation now allows us to vastly simplify Challenge 2 by using a metal detector.
6. **Challenge 6:** Finally, understanding the origins of the weak signals (in this case screws) may help us govern systems more efficiently. Here, we may move hay further away from mechanical equipment with small parts, and keep this in mind when designing future haystack factories.

A depiction of the identified challenges and their relation to the four literary strands is given in Figure 1. Note the Challenge 4 is omitted, as it is not related to the detection of weak signals and not considered in this report. The figure depicts the five remaining challenges in purple, with arrows connecting challenges that are coupled. Literary strands are depicted in green, and connected to challenges on which they focus. Finally, red arrows indicate strong export of concepts and tools from one literary strand to another.

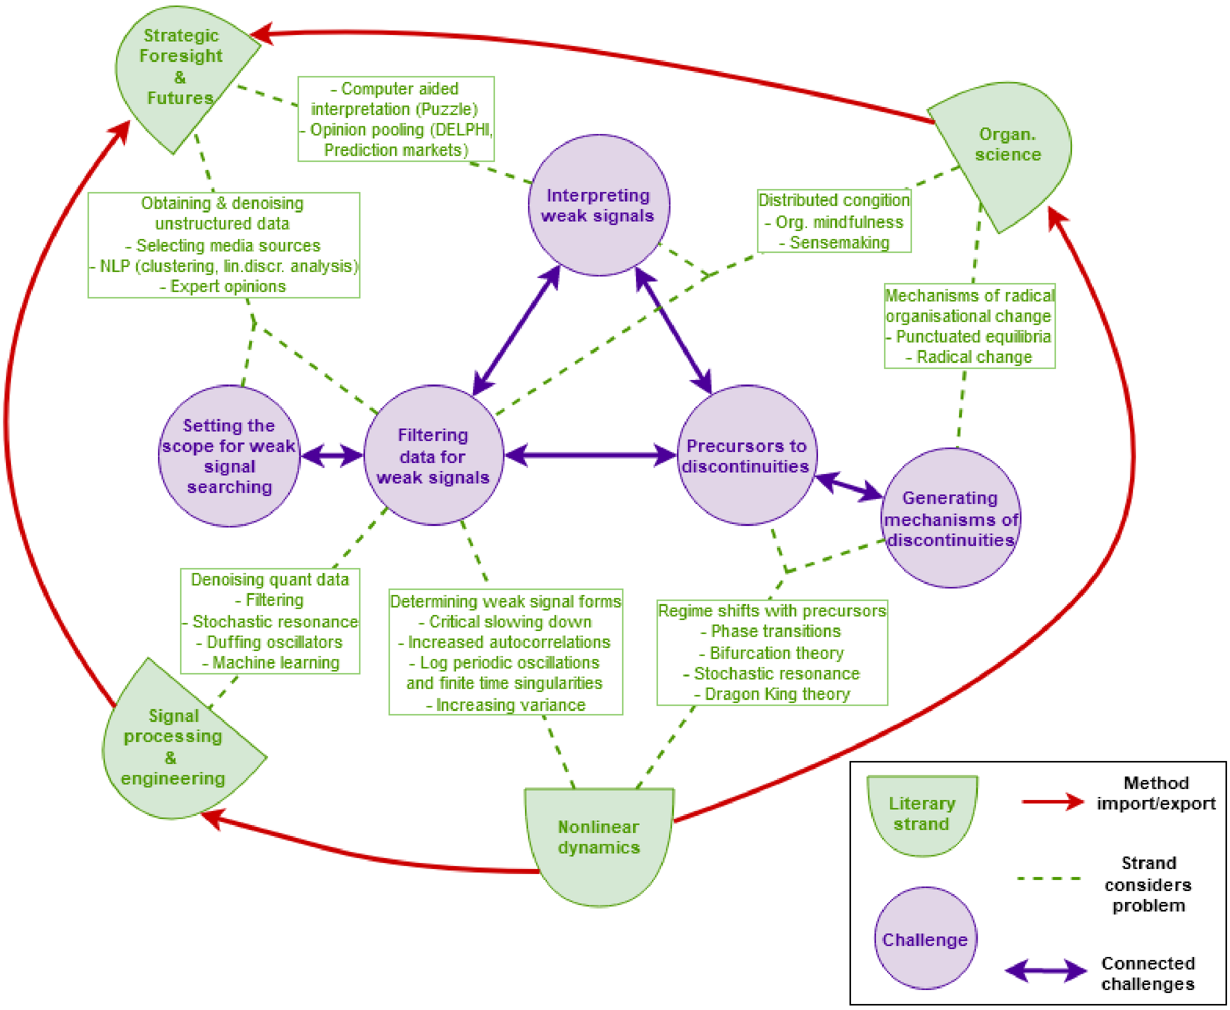

Figure 1: The fundamental challenges in harnessing weak signals (purple network), and their relationship with the four literary strands (green network). Red arrows represent flow of methodology between literatures. The text written over the green dashed lines is a non-exhaustive list of keywords and topics linking literary strands with challenges.

Interpreting literature diversity

In this section we present how the framework (i.e., challenges and definitions) presented in the previous section can be used to consolidate all reviewed works. We argue that all four literary strands have been dealing with the same phenomenology and track the origins of all their differences to domain-specific issues. Specifically, all works are in pursuit of the same objective (harnessing weak signals), which is captured by the challenges and definitions presented in the previous section. However, in each strand this objective is impeded by a different subset of the six previously described challenges, and, rather idiosyncratically, the framework produced by each strand focuses on its respective subset. Thus, we claim that the variations between the four literary strands are the result of missing the bigger picture by overemphasising on domain-specific challenges.

Arguably, the most prominent differentiation between the four frameworks is that social and managerial sciences emphasise weak signal interpretation, while engineering, signal processing, and nonlinear dynamics do not. We propose that this difference is the result of weak signals being harder to interpret in social systems than in engineered ones. Engineering and nonlinear dynamics often have the luxury of precise models and experimentation, both of which can be used to aid in the interpretation of weak signals. In contrast, social systems are less accurately described by formal models and harder to experiment with, making weak signal interpretation more challenging.

There are also differences to be found between the concept of weak signals as defined in nonlinear dynamics and signal processing and engineering. Nonlinear dynamics emphasise the disproportionality between a weak signal and its outcome, while engineering does not. This difference stems from the fact that most engineering works are focused on methodological contributions, and as a result do not explore the mechanisms of weak signal generation. Instead, works from engineering rely on importing methods from nonlinear dynamics when needed.

Automation of weak signal applications

Here we propose an intuitive categorisation of weak signal applications, that can be used to: i) justify the flow of import/export works between the four literary strands; ii) help us identify future research avenues; and, more importantly, iii) help us determine the methodological tools used to automatically tackle the challenges of a weak signal application. Building on the previously identified challenges, we propose four levels of automation maturity for any weak signal challenge.

- **Applications on the first level of maturity.** Applications on this level need experts to tackle all challenges to harnessing weak signals: setting the weak signal scope, detecting weak signals, and interpreting weak signals. Works on futures are at this maturity level, and focus on how to best leverage the knowledge of experts, providing them with computer aid, pooling multiple experts (DELPHI), using criteria to detect ‘hidden’ experts, or using tools to catalyse sensemaking and organisational mindfulness. An additional obstacle to graduating to this level of maturity is system observability: the state of certain systems may not be directly measurable. For example, the perception of the public for some societal issue may only be inferred via studies, or sentiment analysis of textual content. Social media platforms such as Twitter have helped in increasing the observability of social systems.
- **Works on the second level of maturity** have determined the potential sources for weak signals (see Challenge 1 in the previous section). If we think of a weak signal as a ’needle in a haystack’, the second level of maturity refers to weak signal applications where we have clearly defined the boundaries of the haystack. In certain applications, this task can be accomplished via feature selection: skimming through available historical data and determining which information is relevant to our application. Feature selection is a mature field in predictive modelling and machine learning, and a wealth of methodological tools are available. Graduating to this maturity level can be a challenge because of i) limited data availability, i.e., one may not have access to all relevant historical data, or ii) because the potential sources for weak signals may change in time and thus one has to continuously redefine the scope for the weak signal search. The latter reason is why most works from social and organisational sciences are unable to graduate to this level. For example, in order to ensure that the scope is up to date, many

firms include blindspot analysis in their business cycle. Blindspot analysis is a laborious manual process involving expert judgment.

- **The third level of maturity** includes works that have obtained weak signal signatures (i.e., an accurate, machine-understandable description of weak signals; see Challenge 2 in the previous section). Such a description allows one to algorithmically detect weak signals in a specific dataset. Denoising is typically still an issue, as signatures may be obscured by noise. However, having a machine-understandable definition of a weak signal allows for the development of advanced denoising tools, and the automated identification of weak signals. Thinking in terms of the 'needle in a haystack', works on this level of maturity have accurately defined what a needle looks like. Works may fail to graduate to this maturity level because of complex or idiomatic weak signals that require ad-hoc descriptions or denoising tools. For example, textual analysis of social media streams may be hindered by false news, spelling errors, use of emoticons, irony, and other obstacles. To identify and remove such obstacles, purpose-specific methods need to be developed.
- **A fourth level of maturity** includes works that have clearly associated weak signal signatures with specific wildcards (see Challenge 3 in the previous section). A typical example comes from engineering, where the log-periodic power law pattern (a particular weak signal signature) has been established to be indicative of imminent material rupture (the wildcard). Associating signatures with specific wildcards can be achieved i) via formal methods, if the underlying system can be modelled accurately; ii) experimentally, with simulation of laboratory experiments; or iii) empirically, by investigating historical data. At this level of maturity, the task of weak signal-based prediction can become automated, as demonstrated in numerous engineering applications. Graduating to this maturity level can be a major challenge — or arguably impossible — for systems with time-varying behaviour. Instead, once this signature is detected, one must rely on expert judgement to infer the associated wildcard. The majority of works from signal processing and engineering are in this category.

Notably, the nonlinear dynamics stream of works is not explicitly mentioned in any specific maturity level. This is because works from nonlinear dynamics can be found at all maturity levels, focusing on making applications graduate to a higher maturity level through the use of formal methods or simulation. Specifically, nonlinear dynamics may enable an application to graduate from the second level of maturity to the third by identifying signatures (see Challenge 5 in previous section) or from the third level to the fourth by associating signatures with wildcards.
The presented levels of maturity can also justify the import/export patterns between literary strands. As mentioned, nonlinear dynamics exports to works at the second and third levels of maturity. Additionally, it is the view of the authors that methods and concepts are exported from literatures of high maturity towards lower maturity literatures. For example, works from engineering and signal processing that are at the highest maturity level do not import any methods. Futures — found at the first level of maturity — imports from organisational science works that are found on the second level of maturity.

Discussion

We have identified four literary strands that are all preoccupied — in varying degrees — with a similar goal: utilising hard-to-spot signals in order to drastically enhance system performance. Each literary strand is faced by different challenges, and in response has developed tools and definitions that focus around these idiosyncratic challenges:

- Strategic foresight deals with systems of non-stationary scope — and thus focuses on allocating attention and leveraging expert opinion.
- Engineering and signal processing primarily deal with closed systems described by precise analytical models — and thus focus on detecting signals obscured by excessive noise through advanced analytics.
- Weak signal works from organisational science dwell on the endogenous dynamics of complex systems, which lack precise models — and rely on functionalist approaches and empirical studies.
- Works from nonlinear systems that focus on weak signals deal with systems characterised by well-understood but non-trivial dynamics — and rely on formal methods to unravel the mechanisms that may generate weak signals in this context.

In spite of sharing a similar goal, works from different literatures may employ significantly different approaches in harnessing the predictive power of weak signals: social sciences focus on levering and enabling the intuition of experts and prescribing rules that enable groups to cope with weak signals; engineering and signal processing rely on experimentation and advanced analytical methods to capture weak signal signatures; and finally, nonlinear dynamics utilise parsimonious mathematical models as analogues in order to identify weak signal signatures, or to produce novel weak signal detection methods.

The methodological differences between these literatures have resulted in the false impression of a divide, impeding interdisciplinary work. As a remedy, we propose a holistic framework for weak signal detection (consisting of a set of definitions and a set of challenges), which can be used to unify weak signal detection across all considered literary groups. Based on this framework we identify four levels of maturity for weak signal applications: from applications where each weak signal requires ad-hoc analysis, to systems that identify, interpret, and exploit weak signals fully automatically.

The value of the proposed framework lies in its potential to reveal future research directions for radical interdisciplinary research, which may help weak signal applications to graduate to higher maturity levels. This is done by exporting methods from highly mature literatures to less mature ones — the most prominent example being the use of methods from engineering and signal processing for social science and management applications. Given the non-stationary and hard-to-predict behaviour typically encountered in social systems, we expect that expert judgement will continue to be a driver behind weak signal detection and interpretation.

Social sciences may benefit by importing methodology (from engineering, signal processing, and nonlinear dynamics) in order to augment expert judgement. Given the recent boom in social data analytics, it is almost cliché to broadly point out the potential of quantitative methods for social systems. Indeed, the pitfalls of drawing naive analogies between analytical and social sciences have been discussed in organisation science studies (Oswick et al. 2011). However, identifying research topics via the proposed framework avoids this pitfall by revealing connections between specific underlying challenges, which allows us to identify targeted interdisciplinary projects.

As engineering systems increase in scale and become more intertwined with human action, they also become increasingly complex. This trend stands to exacerbate discontinuous events. Also, precise models of engineered systems give way to the unpredictability of human action, and the complexity of non-trivial endogenous dynamics. Thus, coping with discontinuities via utilising weak signals will require the development of methods that do not rely on precise models and are in this sense non-traditional in engineering. For the development of such non-traditional tools, engineering may benefit from importing works from organisational science, which has been studying weak signal behaviour solely within non-precisely understood systems. Additionally, to capture more accurately the non-trivial endogenous dynamics characterising large-scale engineered systems, one may rely on works from statistical mechanics, where tools for describing the emergent behaviour of large-scale systems have been developed.

A prime example can be found in software development: continuously developed and distributed software infrastructure can give rise to unforeseen incidents (such as cascading outages of services). A number of methods aimed at preventing such incidents are already well established in the industry, such as routinely subjecting individual components to thorough tests. However, to deal with the increasing complexity of large, distributed software infrastructure, novel methods are being developed that do not rely on testing for precise behaviours of components or strict assumptions — such as system-agnostic methods for automated stress testing in production conditions (Basiri 2016; Rosenthal et al. 2020).

References

Aaker, David A. and Briance Mascarenhas (1984). ‘The need for strategic flexibility’, *Journal of Business Strategy*, 5(2): 74–82.

Abeshu, Abebe and Naveen Chilamkurti (2018). ‘Deep learning: The frontier for distributed attack detection in fog-to-things computing’, *IEEE Communications Magazine*, 56 (2): 169–75.

Aguilar, Francis Joseph (1967). *Scanning the Business Environment*. New York: Macmillan.

Alam, Firoj, Ferda Ofli, and Muhammad Imran (2018). ‘CrisisMMD: Multimodal Twitter datasets from natural disasters’, paper presented at the 12th International AAAI Conference on Web and Social Media, Palo Alto.

Amanatidou, Effie and Ken Guy (2008). ‘Interpreting foresight process impacts: Steps towards the development of a framework conceptualising the dynamics of “foresight systems”’, *Technological Forecasting and Social Change*, 75(4): 539–57.

Amundsen, David Skålid, Julien Scheibert, Kjetil Thøgersen, Jørgen Trømborg, and Anders Malthe-Sørenssen (2012). ‘1D model of precursors to frictional stick-slip motion allowing for robust comparison with experiments,’ *Tribology Letters*, *45*: 357–69.

Amruthnath, Nagdev and Tarun Gupta (2018). ‘A research study on unsupervised machine learning algorithms for early fault detection in predictive maintenance’, in *2018 5th International Conference on Industrial Engineering and Applications*, 355–61. IEEE.

Andrianarison, Maherizo (2018). ‘New methods and architectures for high sensitivity hybrid GNSS receivers in challenging environments’. PhD dissertation. Institut Supérieur de l’Aéronautique et de l’Espace.

Anifrani, Jean-Charles, Christian Marcel Le Floc’h, Didier Sornette, and Bernard Souillard (1995). ‘Universal log-periodic correction to renormalization group scaling for rupture stress prediction from acoustic emissions’, *Journal de Physique I*, 5(6): 631–38.

Ansoff, H. Igor (1975). ‘Managing strategic surprise by response to weak signals’, *California Management Review*, 18(2): 21–33.

Ansoff, H. Igor (1980). ‘Strategic issue management’, *Strategic Management Journal*, 1(2): 131–48.

Arcangelis, Lucilla de, Carla Perrone-Capano, and Hans J. Herrmann (2006). ‘Self-organized criticality model for brain plasticity’, *Physical Review Letters*, 96(2): 028107.

Asdi, A. S. and Ahmed H. Tewfik. (1995). ‘Detection of weak signals using adaptive stochastic resonance’, in *1995 International Conference on Acoustics, Speech, and Signal Processing*, *Vol. 2*, 1332–35. IEEE.

Atefeh, Farzindar and Wael Khreich (2015). ‘A survey of techniques for event detection in Twitter’, *Computational Intelligence*, 31(1): 132–64.

Baer, Ruth A. (2003). ‘Mindfulness training as a clinical intervention: A conceptual and empirical review’, *Clinical Psychology: Science and Practice*, 10(2): 125–43.

Bak, Per and Kim Sneppen (1993). 'Punctuated equilibrium and criticality in a simple model of evolution', *Physical Review Letters*, 71(24): 4083.

Bak, Per, Chao Tang, and Kurt Wiesenfeld (1988). 'Self-organized criticality', *Physical Review A,* 38(1): 364.

Basiri, Ali, Niosha Behnam, Ruud de Rooij, Lorin Hochstein, Luke Kosewski, Justin Reynolds, and Casey Rosenthal (2016). 'Chaos engineering', *IEEE Software*, 33(3): 35–41.

Beggs, John M. and Dietmar Plenz (2003). 'Neuronal avalanches in neocortical circuits', *Journal of Neuroscience,* 23 (35): 11167–77.

Benzi, Roberto, Alfonso Sutera, and Angelo Vulpiani (1981). 'The mechanism of stochastic resonance', *Journal of Physics A: Mathematical and General*, 14 (11): L453–57.

Bessant, John, Richard Lamming, Hannah Noke, and Wendy Phillips (2005). 'Managing innovation beyond the steady state', *Technovation*, 25(12): 1366–76.

Borio, Daniele, Cillian O'Driscoll, and Gerard Lachapelle (2009). 'Coherent, noncoherent, and differentially coherent combining techniques for acquisition of new composite GNSS signals', *IEEE Transactions on Aerospace and Electronic Systems*, 45(3): 1227–40.

Broughton, Edward (2005). 'The Bhopal disaster and its aftermath: A review', *Environmental Health*, 4(1): 6.

Brynielsson, Joel, Andreas Horndahl, Fredrik Johansson, Lisa Kaati, Christian Mårtenson, and Pontus Svenson (2013). 'Harvesting and analysis of weak signals for detecting lone wolf terrorists', *Security Informatics*, 2(1): 1–15.

Cain, Michael E., H. Dieter Ambos, Francis X. Witkowski, and Burton E. Sobel (1984). 'Fast-Fourier transform analysis of signal-averaged electrocardiograms for identification of patients prone to sustained ventricular tachycardia', *Circulation*, 69(4): 711–20.

Carreras, Benjamin A., David E. Newman, Ian Dobson, and A. Bruce Poole (2004). 'Evidence for self-organized criticality in a time series of electric power system blackouts', *IEEE Transactions on Circuits and Systems I: Regular Papers*, 51(9): 1733–40.

Charitonidis, Christos, Awais Rashid, and Paul J. Taylor (2015). 'Weak signals as predictors of real-world phenomena in social media', in *2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 864–71. IEEE.

-------- (2017). 'Predicting collective action from micro-blog data', in *Prediction and Inference from Social Networks and Social Media*, eds. Jalal Kawash, Nitin Agarwal, and Tansel Özyer, 141–70. Cham: Springer.

Cichocki, Andrzej, Sergei L. Shishkin, Toshimitsu Musha, Zbigniew Leonowicz, Takashi Asada, and Takayoshi Kurachi (2005). 'EEG filtering based on blind source separation (BSS) for early detection of Alzheimer's disease', *Clinical Neurophysiology*, 116(3): 729–37.

Clegg, Stewart and James R Bailey (2007). *International Encyclopedia of Organization Studies*. Thousand Oaks, CA: Sage Publications.

Comon, Pierre and Christian Jutten (2010). *Handbook of Blind Source Separation: Independent Component Analysis and Applications*. Oxford: Academic Press.

Culotta, Aron (2010). 'Towards detecting influenza epidemics by analyzing Twitter messages', in *Proceedings of the First Workshop on Social Media Analytics*, 115–22. New York: ACM.

Daft, Richard L., Juhani Sormunen, and Don Parks (1988). 'Chief executive scanning, environmental characteristics, and company performance: An empirical study', *Strategic Management Journal*, 9(2): 123–39.

Day, George S. and Paul J. H. Schoemaker (2008). 'Are you a "vigilant leader?"', *MIT Sloan Management Review*, 49(3): 43.

DiEsposti, Raymond (2001). 'GPS PRN code signal processing and receiver design for simultaneous all-in-view coherent signal acquisition and navigation solution determination', *Proceedings of the 2007 National Technical Meeting of the Institute of Navigation*, 91–103. Fairfax: ION.

Durlauf, Steven N. (1999). 'How can statistical mechanics contribute to social science?', *Proceedings of the National Academy of Sciences*, 96(19): 10582–84.

Eckhoff, Robert, Mark Markus, Markus Lassnig, and Sandra Schon (2014). 'Detecting weak signals with technologies overview of current technology-enhanced approaches for the detection of weak signals', *International Journal of Trends in Economics Management & Technology (IJTEMT)*, 3(5).

Ehara, Naoki, Iwao Sasase, and Shinsaku Mori (1994). 'Weak radar signal detection based on wavelet transform', *Electronics and Communications in Japan (Fundamental Electronic Science)*, 77(8): 105–14.

Liu, Zhiyong, and Menzel, Lucas (2016). 'Identifying long-term variations in vegetation and climatic variables and their scale-dependent relationships: A case study in Southwest Germany', *Global and Planetary Change*, *147*, 54-66.

Fronczak, Piotr, Agata Fronczak, and Janusz A Hołyst (2007). 'Phase transitions in social networks', *European Physical Journal B*, 59(1): 133–39.

George, Daniel and Eliu Antonio Huerta (2018). 'Deep learning for real-time gravitational wave detection and parameter estimation: Results with advanced LIGO data', *Physics Letters B*, 778: 64–70.

Georgiadis, Dionysios and Didier Sornette (2019). 'Pattern phase diagram of spiking neurons on spatial networks', *Physical Review E*, 99(4): 042410.

Glennon, Eamonn P. and Andrew G. Dempster (2004). 'A review of GPS cross correlation mitigation techniques', paper presented at *The 2004 International Symposium on GNSS/GPS*, Sydney.

Goldstein, Jeffrey, James K. Hazy, and Joyce Silberstang (2010). 'A complexity science model of social innovation in social enterprise', *Journal of Social Entrepreneurship*, 1(1): 101–25.

Gould, Stephen Jay and Niles Eldredge. (1977). 'Punctuated equilibria: The tempo and mode of evolution reconsidered', *Paleobiology*, 3(2): 115–51.

Graefe, Andreas, Stefan Luckner, and Christof Weinhardt (2010). 'Prediction markets for foresight', *Futures*, 42(4): 394–404.

Gupta, Aditi, Hemank Lamba, Ponnurangam Kumaraguru, and Anupam Joshi (2013). ‘Faking Sandy: Characterizing and identifying fake images on Twitter during Hurricane Sandy’, in *Proceedings of the 22nd International Conference on World Wide Web*, 729–36. New York: ACM.

Haiman, Zoltán (2018). ‘Machine learning uncovers new cosmological information’, *arXiv Preprint arXiv:1810.12220*.

Hari, V. N., G. V. Anand, A. Benjamin Premkumar, and A. S. Madhukumar (2012). ‘Design and performance analysis of a signal detector based on suprathreshold stochastic resonance’, *Signal Processing*, 92(7): 1745–57.

Harmer, Gregory Peter, Bruce R. Davis, and Derek Abbott (2002). ‘A review of stochastic resonance: Circuits and measurement’, *IEEE Transactions on Instrumentation and Measurement*, 51(2): 299–309.

Hänggi, Peter (2002). ‘Stochastic resonance in biology: How noise can enhance detection of weak signals and help improve biological information processing’, *ChemPhysChem*, 3(3): 285–90.

Hesse, Janina and Thilo Gross (2014). ‘Self-organized criticality as a fundamental property of neural systems’, *Frontiers in Systems Neuroscience*, 8: 166.

Hiltunen, Elina (2007). ‘Where do future-oriented people find weak signals?’, *Future Takes*, 6(2).

Hiltunen, Elina (2008). ‘Good sources of weak signals: A global study of where futurists look for weak signals’, *Journal of Futures Studies*, 12(4): 21–44.

Hiltunen, Elina (2010). *Weak signals in organizational futures learning*. Helsinki School of Economics.

Holopainen, Mari and Marja Toivonen (2012). ‘Weak signals: Ansoff today’, *Futures*, 44(3): 198–205.

Huang, Y., Hubert M. Saleur, Charles Sammis, and Didier Sornette (1998). ‘Precursors, aftershocks, criticality and self-organized criticality’, *Europhysics Letters*, 41(1): 43–48.

Ilmola, Leena and Osmo Kuusi (2006). ‘Filters of weak signals hinder foresight: Monitoring weak signals efficiently in corporate decision-making’, *Futures*, 38(8): 908–24.

Imran, Muhammad, Shady Elbassuoni, Carlos Castillo, Fernando Diaz, and Patrick Meier (2013). ‘Practical extraction of disaster-relevant information from social media’, in *Proceedings of the 22nd International Conference on World Wide Web*, 1021–24. New York: ACM.

Iooss, Gérard and Moritz Adelmeyer (eds) (1998). *Topics in Bifurcation Theory and Applications*, 2nd edition. Singapore: World Scientific.

Jalilvand, Abolfazl and Hadi Fotoohabadi (2011). ‘The application of duffing oscillator in weak signal detection’, *ECTI Transactions on Electrical Engineering, Electronics, and Communications*, 9(1): 1–6.

Janczura, Joanna and Rafal Weron (2012). ‘Black Swans or Dragon-Kings? A simple test for deviations from the power law’, *European Physical Journal Special Topics*, 205(1): 79–93.

Johnson, Neil, Guannan Zhao, Eric Hunsader, Jing Meng, Amith Ravindar, Spencer Carran, and Brian Tivnan (2012). ‘Financial Black Swans driven by ultrafast machine ecology’, *arXiv Preprint arXiv:1202.1448*.

Kaivo-oja, Jari (2012). ‘Weak signals analysis, knowledge management theory and systemic socio-cultural transitions’, *Futures*, 44(3): 206–17.

Katona, Zsolt, Marcus Painter, Panos N Patatoukas, and Jean Zeng (2018). ‘On the capital market consequences of alternative data: Evidence from outer space’, in *9th Miami Behavioral Finance Conference*, Miami.

Kayser, Victoria and Antje Bierwisch (2016). ‘Using Twitter for foresight: An opportunity?’, *Futures*, 84: 50–63.

Lei, Yaguo, Dong Han, Jing Lin, and Zhengjia He (2013). ‘Planetary gearbox fault diagnosis using an adaptive stochastic resonance method’, *Mechanical Systems and Signal Processing*, 38(1): 113–24.

Lengnick-Hall, Cynthia A. and Tammy E Beck (2016). ‘Resilience capacity and strategic agility: Prerequisites for thriving in a dynamic environment’, in *Resilience Engineering Perspectives, Vol. 2*, eds. Christopher P. Nemeth and Erik Hollnagel, 61–92. Boca Raton: CRC Press.

Lesca, Humbert and Nicolas Lesca (eds) (2011). *Weak Signals for Strategic Intelligence: Anticipation Tool for Managers.* London: ISTE and Wiley.

Levy, Moshe (2008). ‘Stock market crashes as social phase transitions’, *Journal of Economic Dynamics and Control*, 32(1): 137–55.

Li, Chongsheng (2005). ‘Study of weak signal detection based on second FFT and chaotic oscillator’, *Nature and Science*, 3(2): 59–64.

MacGougan, Glenn, Gerard Lachapelle, Rakesh Nayak, and Alexander Wang (2001). ‘Overview of GNSS signal degradation phenomena’, in *Proceedings of the International Symposium Kinematic Systems and Geodesy, Geomatics and Navigation*. Calgary: University of Calgary.

McDonnell, Mark D. and Derek Abbott (2009). ‘What is stochastic resonance? Definitions, misconceptions, debates, and its relevance to biology’, *PLoS Computational Biology*, 5(5): e1000348.

Melgard, Tor E., Gerard Lachapelle, and Hazen Gehue (1994). ‘GPS signal availability in an urban area-receiver performance analysis’, in *Proceedings of 1994 IEEE Position, Location and Navigation Symposium — PLANS ’94*, 487–93. IEEE.

Mendonça, Sandro, Miguel Pina e Cunha, Jari Kaivo-oja, and Frank Ruff (2004). ‘Wild cards, weak signals and organisational improvisation’, *Futures*, 36(2): 201–18.

Meyer, Alan D., Vibha Gaba, and Kenneth A. Colwell (2005). ‘Organizing far from equilibrium: Nonlinear change in organizational fields’, *Organization Science*, 16(5): 456–73.

Mossel, Elchanan, Allan Sly, and Omer Tamuz (2015). ‘Strategic learning and the topology of social networks’, *Econometrica*, 83(5): 1755–94.

Mühlroth, Christian and Michael Grottke (2018). ‘A systematic literature review of mining weak signals and trends for corporate foresight’, *Journal of Business Economics*, 88(5): 643–87.

Narchal, R. M., K. Kittappa, and P. Bhattacharya (1987). ‘An environmental scanning system for business planning’, *Long Range Planning*, 20(6): 96–105.

Neill, Stern, Daryl McKee, and Gregory M. Rose (2007). ‘Developing the organization’s sensemaking capability: Precursor to an adaptive strategic marketing response’, *Industrial Marketing Management*, 36(6): 731–44.

Nowack, Martin, Jan Endrikat, and Edeltraud Guenther (2011). ‘Review of Delphi-based scenario studies: Quality and design considerations’, *Technological Forecasting and Social Change*, 78(9): 1603–15.

Oswick, Cliff, Peter Fleming, and Gerard Hanlon (2011). ‘From borrowing to blending: Rethinking the processes of organizational theory building’, *Academy of Management Review*, 36(April): 318–37.

Pépin, Lambert, Pascale Kuntz, Julien Blanchard, Fabrice Guillet, and Philippe Suignard (2017). ‘Visual analytics for exploring topic long-term evolution and detecting weak signals in company targeted tweets’, *Computers & Industrial Engineering*, 112: 450–58.

Pisarenko, Vladilen and Didier Sornette (2012). ‘Robust statistical tests of Dragon-Kings beyond power law distributions’, *European Physical Journal Special Topics*, 205(1): 95–115.

Pueyo, Salvador (2007). ‘Self-organised criticality and the response of wildland fires to climate change’, *Climatic Change*, 82(1–2): 131–61.

Qiu, Hai, Jay Lee, Jing Lin, and Gang Yu. (2006). ‘Wavelet filter-based weak signature detection method and its application on rolling element bearing prognostics’, *Journal of Sound and Vibration*, 289(4–5): 1066–90.

Reinhardt, W. A. (1984). ‘An early warning system for strategic planning’, *Long Range Planning*, 17(5): 25–34.

Rodríguez-Méndez, Víctor, Víctor M. Eguíluz, Emilio Hernández-García, and José J. Ramasco (2016). ‘Percolation-based precursors of transitions in extended systems’, *Scientific Reports,* 6: 29552.

Rosenthal, Casey, Lorin Hochstein, Aaron Blohowiak, Nora Jones, and Ali Basiri (2020). *Chaos Engineering*. Sabastapool: O'Reilly Media.

Rossel, Pierre (2012). ‘Early detection, warnings, weak signals and seeds of change: A turbulent domain of futures studies’, *Futures*, 44(3): 229–39.

Rouiah, Kamel and Samia Ould-Ali (2002). ‘PUZZLE: A concept and prototype for linking business intelligence to business strategy’, *Journal of Strategic Information Systems*, 11(2): 133–52.

Saleur, Hubert M. and Didier Sornette (1996). ‘Complex exponents and log-periodic corrections in frustrated systems’, *Journal de Physique I*, 6(3): 327–55.

Sardar, Ziauddin (2010). ‘The namesake: Futures; futures studies; futurology; futuristic; foresight — What’s in a name?’, *Futures*, 42(3): 177–84.

Sándor, Ágnes (2009). 'Automatic detection of discourse indicating emerging risk', *Critical Approaches to Discourse Analysis Across Disciplines. Risk as Discourse — Discourse as Risk: Interdisciplinary Perspectives*, 4.

Schoemaker, Paul J. H. and George S. Day (2009). 'How to make sense of weak signals', *Leading Organizations: Perspectives for a New Era*, 37.

Schoemaker, Paul J. H., George S. Day, and Scott A. Snyder (2013). 'Integrating organizational networks, weak signals, strategic radars and scenario planning', *Technological Forecasting and Social Change*, 80(4): 815–24.

Seco-Granados, Gonzalo, Jose Lopez-Salcedo, David Jimenez-Banos, and Gustavo Lopez-Risueno (2012). 'Challenges in indoor global navigation satellite systems: Unveiling its core features in signal processing', *IEEE Signal Processing Magazine*, 29(2): 108–31.

Seidl, David (2004). 'The concept of weak signals revisited: A re-description from a constructivist perspective', in *Managing the Future: Foresight in the Knowledge Economy*, eds. Haridimos Tsoukas and Jill Shepherd, 153–70. Oxford: Blackwell.

Solé-Casals, Jordi and François-Benoît Vialatte (2015). 'Towards semi-automatic artifact rejection for the improvement of Alzheimer's disease screening from EEG signals', *Sensors*, 15(8): 17963–76.

Somit, Albert and Steven A. Peterson (1992). *The Dynamics of Evolution: The Punctuated Equilibrium Debate in the Natural and Social Sciences*. Ithaca: Cornell University Press.

Song, Bihao (2012). 'Weak signal detection on Twitter datasets: A non-accumulated approach for non-famous events'. Master's thesis. TU Delft.

Sornette, Anne and Didier Sornette (1989). 'Self-organized criticality and earthquakes', *Europhysics Letters*, 9(3): 197.

Sornette, Didier (2002). 'Predictability of catastrophic events: Material rupture, earthquakes, turbulence, financial crashes, and human birth', *Proceedings of the National Academy of Sciences*, 99(suppl 1): 2522–29.

Sornette, Didier (2009). 'Dragon-Kings, Black Swans and the prediction of crises', *arXiv Preprint arXiv:0907.4290*.

Sornette, Didier, Anders Johansen, and Jean-Philippe Bouchaud (1996). 'Stock market crashes, precursors and replicas', *Journal de Physique I*, 6(1): 167–75.

Sornette, Didier and Guy Ouillon (2012). 'Dragon-Kings: Mechanisms, statistical methods and empirical evidence', *European Physical Journal Special Topics*, 205(1): 1–26.

Sriram, Bharath, Dave Fuhry, Engin Demir, Hakan Ferhatosmanoglu, and Murat Demirbas (2010). 'Short text classification in Twitter to improve information filtering', in *Proceedings of the 33rd International SCM SIGIR Conference on Research and Development in Information Retrieval*, 841–42. ACM.

Stauffer, Dietrich and Didier Sornette (1999). 'Self-organized percolation model for stock market fluctuations', *Physica A: Statistical Mechanics and Its Applications*, 271(3–4): 496–506.

Stephenson, Max, Jr. (2005). ‘Making humanitarian relief networks more effective: Operational coordination, trust and sense making’, *Disasters*, 29(4): 337–50.

Strogatz, Steven H., Mark Friedman, A. John Mallinckrodt, Susan McKay (1994). ‘Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering,’ *Computers in Physics*, 8(5): 532.

Suma, Sugimiyanto, Rashid Mehmood, and Aiiad Albeshri (2017). ‘Automatic event detection in smart cities using Big Data Analytics’, in *International Conference on Smart Cities, Infrastructure, Technologies and Applications*, eds. Rashid Mehmood, Budhendra Bhaduri, Iyad Katib, Imrich Chlamta, Vol. 224, 111–22. Cham: Springer.

Taleb, Nassim Nicholas (2007). ‘Black Swans and the domains of statistics’, *The American Statistician*, 61(3): 198–200.

Taylor, Ralph E. and James W. Sennott (1984). *Navigation System and Method* (US4445118A). Google Patents. https://patents.google.com/patent/US4445118A/en

Thomas, James B., Shawn M. Clark, and Dennis A. Gioia (1993). ‘Strategic sensemaking and organizational performance: Linkages among scanning, interpretation, action, and outcomes’, *Academy of Management Journal*, 36(2): 239–70.

Thomson, William T. and Mark Fenger (2001). ‘Current signature analysis to detect induction motor faults’, *IEEE Industry Applications Magazine*, 7(4): 26–34.

Thorleuchter, Dirk and Dirk Van den Poel (2013). ‘Weak signal identification with semantic web mining’, *Expert Systems with Applications*, 40(12): 4978–85.

Timmermans, Jos (2008). ‘Punctuated equilibrium in a non-linear system of action’, *Computational and Mathematical Organization Theory*, 14(4): 350.

Todnem By, Rune (2005). ‘Organisational change management: A critical review’, *Journal of Change Management*, 5(4): 369–80.

Van Diggelen, Frank (2009). *A-GPS: Assisted GPS, GNSS, and SBAS.* Boston and London: Artech House.

Vecchiato, Riccardo (2012). ‘Environmental uncertainty, foresight and strategic decision making: An integrated study’, *Technological Forecasting and Social Change*, 79(3): 436–47.

Wang, Guanyu, Dajun Chen, Jianya Lin, and Xing Chen (1999). ‘The application of chaotic oscillators to weak signal detection’, *IEEE Transactions on Industrial Electronics*, 46(2): 440–44.

Wang, Guanyu and Sailing He (2003). ‘A quantitative study on detection and estimation of weak signals by using chaotic duffing oscillators’, *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 50(7): 945–53.

Wang, Jie, Lu Yang, Long Gao, and Qiang Miao (2013). ‘Current progress on weak signal detection’, in *2013 International Conference on Quality, Reliability, Risk, Maintenance, and Safety Engineering (Qr2mse)*, 1812–18. IEEE.

Wang, Yi, Guanghua Xu, Lin Liang, and Kuosheng Jiang (2015). ‘Detection of weak transient signals based on wavelet packet transform and manifold learning for rolling element bearing fault diagnosis’, *Mechanical Systems and Signal Processing*, 54: 259–76.

Wang, Zhan, Stephen McConnell, Robert S. Balog, and Jay Johnson (2014). ‘Arc fault signal detection — Fourier transformation vs. wavelet decomposition techniques using synthesized data’, in *2014 IEEE 40th Photovoltaic Specialist Conference*, 3239–44. IEEE.

Weick, Karl E. (1987). ‘Organizational culture as a source of high reliability’, *California Management Review*, 29(2): 112–27.

Weick, Karl E. (1988). ‘Enacted sensemaking in crisis situations’, *Journal of Management Studies*, 25(4): 305–17.

Weick, Karl E. (1995). *Sensemaking in Organizations*, *Vol. 3*. Thousand Oaks, CA: Sage.

Weick, Karl E. (2010). ‘Reflections on enacted sensemaking in the Bhopal disaster’, *Journal of Management Studies*, 47(3): 537–50.

Weisbuch, Gérard, Sorin Solomon, and Dietrich Stauffer (2001). ‘Social percolators and self organized criticality’, in *Economics with Heterogeneous Interacting Agents*, eds. Alan Kirman and Jean-Benoît Zimmermann, 43–55. Cham: Springer.

Welz, Kirill, Leo Brecht, Anja Pengl, Julian V Kauffeldt, and Daniel RA Schallmo (2012). ‘Weak signals detection: Criteria for social media monitoring tools’, in *International Society for Professional Innovation Management Innovation Symposium*, 1. ISPIM.

Wheatley, Spencer, Benjamin Sovacool, and Didier Sornette (2017). ‘Of disasters and Dragon Kings: A statistical analysis of nuclear power incidents and accidents’, *Risk Analysis*, 37(1): 99–115.

Wiesenfeld, Kurt and Frank Moss (1995). ‘Stochastic resonance and the benefits of noise: From Ice Ages to grayfish and squids’, *Nature*, 373(6509): 33.

Yan, Wanfeng, Ryan Woodard, and Didier Sornette (2010). ‘Diagnosis and prediction of tipping points in financial markets: Crashes and rebounds’, *Physics Procedia*, 3: 1641–57.

Yoon, Janghyeok (2012). ‘Detecting weak signals for long-term business opportunities using text mining of web news’, *Expert Systems with Applications*, 39(16): 12543–50.

Zanetti, Marcelo Serrano, Ingo Scholtes, Claudio Juan Tessone, and Frank Schweitzer (2013). ‘The rise and fall of a central contributor: Dynamics of social organization and performance in the Gentoo community’, in *2013 6th International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE)*, 49–56. IEEE.

Acknowledgment

This work is an outcome of the Future Resilient Systems project at the Singapore-ETH Centre (SEC) supported by the National Research Foundation, Prime Minister's Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme.

Contact

Singapore-ETH Centre
Future Resilient Systems
1 Create Way, #06-01 CREATE Tower
138602 Singapore

www.frs.ethz.ch

Publisher: Singapore-ETH Centre