



FORHUMANITY

AN INTRODUCTION

SYSTEMIC SOCIETAL IMPACT ANALYSIS (SSIA)

April 2022
RYAN CARRIER

<https://forhumanity.center/>

This article is designed to consider, track and measure the importance (risk and/or potential negative impact) of socio-technical systems to individuals, communities, nation-states or society-at-large in order to signal shifting levels of risk that likely require risk reassessment

The introduction of the EU Artificial Intelligence Act (proposed) drew attention to Artificial Intelligence and Autonomous Systems deemed to be High Risk. Found in Annex III of the proposed act - the Act states the following:

“This list of high-risk AI systems in Annex III contains a limited number of AI systems whose risks have already materialized or are likely to materialize sooner. To ensure that the regulation can be adjusted to emerging uses and applications of AI, the Commission may expand the list of high-risk AI systems used within certain predefined areas, by applying a set of criteria and risk assessment methodology.”

High risk, as referred to in the EU and UK’s General Data Protection Regulation, directly refers to the risk to the rights and freedoms of natural persons, specifically, but limited to data protection and privacy. The guidance rightly goes on to list a wide variety of additional potential rights and freedoms at risk.

The clarity of scope is strong. Our query lies in the measurement of thresholds - when boundaries are crossed and risks to rights and freedom may manifest - how are these determined and navigated? When are risks high - what are the levels, metrics, thresholds? When do those risks change - becoming something in need of mitigation?

SSIA is designed to consider, track and measure the importance (risk and/or potential negative impact) of socio-technical systems to individuals, communities, nation-states or society-at-large in order to signal shifting levels of risk that likely require risk reassessment



► Why is SSIA necessary?

History tells us that barring catastrophe, without accountability, without signals insisting that we change course - the status quo will prevail. This thought process aggregates up to the corporate level when a collection of individuals make the same assessment about risk and subsequently manifest a higher risk tolerance and appetite than might be prudent for the entity. In response to individual levels of risk tolerance and appetite, we can examine measures taken to overcome this bias. Consider the annoying seat belt beep that reminds you to put it on. Engineers and psychologists tested and proved time and again that these types of nudges are very beneficial to risk management and moreover without them, we are greatly tempted to accept the risk, even when that choice is imprudent given the associated cost or effort of risk mitigation (clicking in that seat belt). The Systemic Societal Impact Analysis is the annoying seat belt, for the entity - a reminder that says "risk present, maybe you should reassess, reassess, reassess, reassess" - until we finally do.

As a species we tend to be risk-takers and frequently underestimate risk and overestimate our risk appetite and risk tolerance. 3 reasons You underestimate Risk - by Srini Pillay, collects a series of studies that document this phenomenon - something he calls "Risk Blindness". Risk Blindness is omnipresent in AI, algorithmic and autonomous systems, highlighted by a weak risk management culture in technology companies that pales in comparison to the embedded risk management cultures (even if imperfect) in healthcare and finance.

This undervaluation of risk, aggregated by the sum of all individuals presents a blind spot of risk management at the systemic societal level.

The Systemic Societal Impact Analysis is a critical part of ForHumanity's comprehensive Risk Management Framework, designed to monitor and quantify "a creeping risk that builds over time, unnoticed, under appreciated, potentially insidiously, until it is too late" - that unfortunate moment when monumental risk occurs linked to AI, algorithmic and autonomous systems.

► What is Systemic Risk?

Borrowing from our Algorithm Risk Assessment (ARA) work - ForHumanity has determined that risk management is optimized when we deploy a pool of Diverse Inputs and Multi Stakeholder Feedback to identify a wide array of risk inputs from the system.

Systemic Risk will be identified when a wide array of assessors spot the same risk input combined with sufficient severity and likelihood to create a meaningful risk.

Some examples might include damage to people's health including mental health, damage to safety, security of a community, country, damage to democratic systems, significant impact on the environment (part of ForHumanity's definition of Ethical), destabilization of financial markets or other collective mechanisms which benefit people at the individual, community, nation-state or societal level systemically.

Considerations - amount of impacted people

The sheer number of impacted data subjects or stakeholders affects the overall assessment of risk. A system or process that impacts 5 to 10 people is entirely different from a system that impacts 100 million people. The reader should understand that is not a reflection of the value or importance of an individual. A system that risks the lives of even 5 to 10 people, remains high risk. However, if the severity of risk is the same for a 5-to-10-person system versus a system with 100 million people, the 100-million-person system is riskier inherently. View it from this standpoint, how much remediation is needed to fix things? If the individual impact is the same, the sheer magnitude makes the risk greater.

Considerations - scope of impact


Talking about systemic impact, we are also accounting for all elements of society and/or the environment in which we live. Therefore, it is necessary to include risk impacts on societal constructs (laws, community procedures and moral frameworks) and importance itself. We might evaluate an AI, algorithmic or autonomous system as low risk to specific individuals or protected category persons, but if our soil is poisoned, or the air is polluted, or energy consumption is excessive then that specific risk must be included in this analysis - a cost to all.

Systematic Societal Impact Analysis overcomes a cognitive dissonance between individual risk assessment and the collective. In the collective, we find the importance of the SSI and what it measures.

When we consider our own individual risk or cost, it might seem meaningless, but when considered as a collective, the sum of miniscule risks can amount to a meaningful and damaging impact.

► What is Systemic Impact?

Systemic Impact refers to the entire opportunity set of engaged entities or individuals and the combined ecosystem under which they collectively function. It is in the importance of that collectivity where we find Systemic Impact, a wide array of meaningful impact across any of the four Spheres of Influence, a simple example would be use of the term "Kleenex" when we mean "facial tissues", it is obvious that the brand Kleenex has become synonymous with tissue. That is Systemic Impact. When we "google" something and mean "use a search engine" or we "fedex something" and mean "send it overnight or by the quickest means possible". We have experienced forms of Systemic Societal Impact. We will discuss in more detail below the differing measures and manifestations of Systemic Societal Impact.



► But why would we care about Systemic Societal Impacts?

Issues appear from Systemic Impact impacts when the Importance, Saturation, Authority and Dependency exerted by entities rises with insufficient risk mitigations. How many people realize that when you search now on Google, paid ads occupy the top spot? For example, searching “AI Ethics” yields an entire page of Google Ads, and searching on Duck Duck Go yields some useful genuine links. The point is that as of Jan 2022 according to statcounter Global Stats - Google held just short of 90% of the search engine market in the United States. You see what Google wants you to see. That is Systemic Societal Impact. When people talk about negative influences on Elections in the United States, they are referring to Facebook and the way information is disseminated to individuals as described concisely by Strategy Lab marketing out of Canada. Our point in highlighting Systemic Societal Impact is to argue that when entities reach certain thresholds it puts our basic rights and freedoms at risk.

This paper aims to provide some guidance - beginning to understand the intrinsic risk associated with socio-technical systems or services and to identify metrics and thresholds where risk reassessment and further mitigations are likely to be necessary because Systemic Societal Impact is increasing. However, before we get into these metrics there is one other foundational principle required - perspective.



► Spheres of influence

Who is impacted by these socio-technical systems? What is their individual perception of the system and how does it vary from one person or group to another? How does each group weigh and consider these Systemic Societal Impacts? The SSI process identifies perspective separated into four distinct Spheres of Influence:

1. Individual - allows us to examine the risk impact to any single individual or group of unaffiliated individuals (an assessment of a collection of individuals)
2. Community - collection of people, usually centered around common characteristics such as location, school, hobbies, professions, beliefs etc. (an assessment of a collection of communities)
3. Nation-State - often defined by both boundaries and common laws (an assessment of country-level adoption)
4. Humanity - characterizing impact to the species, especially in the areas of inalienable human rights and our shared resources beyond the state-level, such as the environment

“ForHumanity advocates Systemic Societal Impact Analysis as a consistent process to understand the evolving impact across each of the spheres of influence..”



FOR HUMANITY



Individuals



Community



Nation State



Humanity

When we assess the systemic risk of a system, it must be grounded in “who” is at-risk. The SSI examines four specific people groups that might be subject to unique or disproportionate risk at any point in time. Because each of these groups are different manifestations of humanity, it is critical that the risk-assessment process always consider all four Spheres of Influence. It is easy to imagine a system that is apparently minimally-risky to humanity broadly, but of the utmost importance to an individual - for example, a sentencing algorithm for a criminal offense. This example represents a life-changing impact to an individual, but beyond activists and watchdog groups, has minimal relevance to most members of society to the degree that they are likely unaware that such a system might exist. This is not a judgment, but instead an examination of direct impact to each of our four Spheres of Influence.

Spheres of Influence are the specific “who” questions we want to independently apply the four metrics of the SSI criteria upon. The result could be a wide variety of systemic impacts, where some Spheres of Influence are largely risk-free while others are critically impacted.

Metrics of Systemic Societal Impact

We can now return to the interaction between these Spheres of Influence and the specific ways in which AI, algorithmic and autonomous systems impact them.

The paper identifies four critical areas of measurement for these metrics:

1. Importance - the magnitude of impact measured at the individual or people group level in regard to survival, health and well-being
2. Saturation - the market share, and dominance of a system associated with a specific Sphere of Influence
3. Authority - the amount of influence the system wields in regard to values, trust, and decision-making for a specific Sphere of Influence
4. Dependency - the economic and personal total cost of withdrawal from a system weighed against the economic and personal need for accessing the system

When Facebook was first built, it was a niche social system, in a single university. The impact of which would barely register as meaningful to any Sphere of Influence. However, Facebook, like most systems, was not static and changed over time. First, it transitioned from meaningless to the students at the University to meaningful. Then Facebook expanded to become meaningful at multiple Universities. Let's pause there for a moment.

Was Facebook meaningful, at that moment, to society-at-large? No. Was Facebook meaningful to families or different generations? No. But it was very impactful to college students - first.

That point is important to understand. It highlights that societal impact is relative to specific Spheres of Influence, often manifesting in different ways at any point in time. Each of our four metrics are relative to the Sphere of Influence and relative to individual assessors.

Returning to our analogy, in early 2005, Facebook had no power over the non-college age communities, or the nation-state or society at-large, in fact it was only just launching a high school version. However, it might have been reasonable to assess that Facebook might have had a huge systemic impact on college-age students. Systematic Societal Impact measures risk and negative impact on each Sphere of Influence and those measurements should be assessed regularly and submitted for Risk Treatment as needed.

Importance

This metric reflects on the area of our lives that is impacted by the system. If the system impacts our source of humor, laughs and entertainment that is less risky than the system we use for information about how we manage our money, manage our health, or vote. A suggested scale of the metric is listed below

1. Casual and Fun - hobbies
2. Knowledge and Facts - the size of something, the temperature at which water boils
3. Consumer interactions and managing daily life
4. Relational and Societal views - political
5. Survivorship needs - financial, health, job security
6. Core Beliefs - family relations, belief systems

When assessing the risks from a system, assessors must take care to understand the influence the system may have subconsciously on users. Presumably, users and multi-stakeholders can explain and understand the conscious impact of a system, therefore this section is about issues of nudging, self-affirmation impact from likes, dislikes, thumbs up, responding with comments, creating sub interest groups as a result of following specific post/ person/ entity, creating multiple related events, accumulation of followers, etc. This assessment is likely provided by domain experts with the detachment to discern the hierarchical need of users in the context of the system beyond their conscious assessments. Especially in light of psychological techniques that technology companies, social media entities and even hackers are deploying, the risk impact of subconscious manipulation must be weighed to determine if the risk from the system might be higher than perceived. This manipulation can be weighed based on some of the following factors::

1. High increase of creation/ adaptability of specific influence groups
2. Highly specific content increase which has an incremental count of followers
3. Incremental increase in focus/need/desire designed to trigger or motivate emotional response/defense/rebuttal regarding beliefs on a specific content/ beliefs which could trigger a movement, violent behavior, force for change in a community, state, nation.
4. High response in terms of followers, likes, dislikes, views, comments etc.
5. High acceptability and usage of a system in a specific community/ group

FOR HUMANITY

The application of this risk assessment should be applied to the measurement of Importance and may increase by one or two steps the overall assessment - based on the amount of subconscious manipulation. The EU Artificial Intelligence Act puts heavy restrictions on these kinds of techniques; therefore, a company should take a proactive, honest and self-aware approach to evaluating these risks. Failure to recognize them and manage these risks will open an entity up to regulatory enforcement.

Scaling Importance and scaling applicability to the other metrics

All scales in a Systematic Societal Impact Analysis can be relative. Assessors can choose more or fewer thresholds, however the point remains the same, "Importance" varies from meaningless to meaningful and as a system rises in the "Importance" metric, there are thresholds crossed and reassessment must occur.

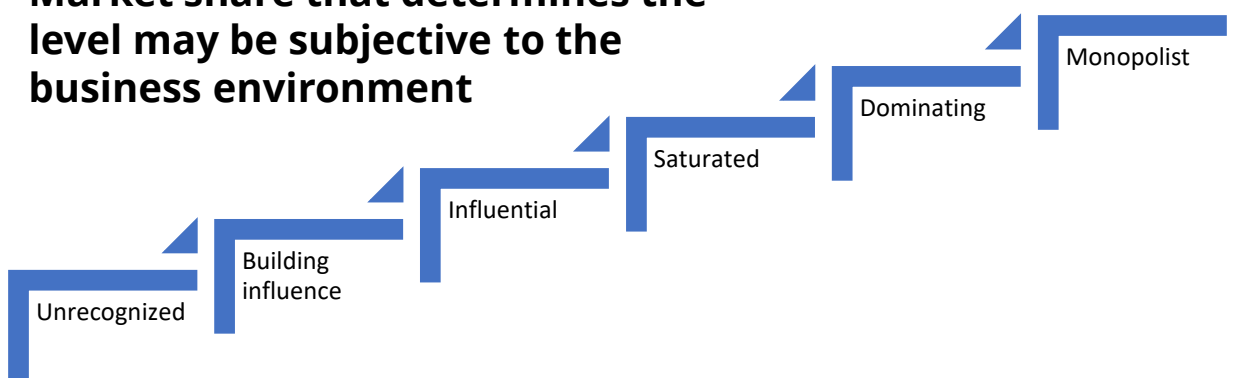
The Importance metric assessment challenge lies in the difficulty to discreetly and accurately quantify systems. Using Facebook again, there are biases in assessment, for example, an embarrassment factor that will lead people to downplay their informational impact from that social media platform.

The typical response will be to downplay a system's importance, or even to fully realize a system's influence. Managing these response biases requires that the assessor and assessment process carefully craft their assessment tools to extract an accurate and unbiased impression of the Importance the system has on an affected individual. Anonymity may be a critical tool for assessment.

Saturation

The SSI requires analysis of how deeply the system or service has penetrated any Sphere of Influence. Therefore, Independent Audit of AI Systems suggests predetermined thresholds to help risk managers understand when the impact of a system has reached a more meaningful or impactful level. For now, a simple example should suffice - a system that is used by 5000 people has different societal Saturation than a system with 2 billion users. Thus, it can be reasonably concluded that somewhere between 5,000 and 2 billion is one threshold (or more) that gets crossed requiring a reassessment of risk. We lay out a range of thresholds that can be measured quantitatively to denote when a risk reassessment must be engaged.

Market share that determines the level may be subjective to the business environment





Crossing Saturation thresholds in any Sphere of Influence should force assessors to reassess risk for that Sphere of Influence. As Saturation increases, Systemic Societal impact increases for the same reason that economists are concerned on behalf of consumers in the face of oligopolies or monopolies - the entities wield too much power, too much control and consumers (or in this case the Spheres of Influence) become beholden to the whims and wishes of the monopolist. Price setting, information control, policy influence, political control and reductions to rights and freedoms occur in the face of rising Saturation.

A critical issue with Saturation is that it is often an implied goal (measure of success) in capitalism - increase market share being that battle cry. That logic results in risks that have proven to require regulatory oversight if not outright antitrust countermeasures. From a human perspective, the greater the Saturation the greater the negative impact from errors, omissions, bias, economic and ethical choices made by the entity.

This paper is not about advising corporations, however, on this point, entities should mitigate at all costs excessive Saturation, either through reduced barriers to entry or avoidance of outright expansion. These measures will result in risk mitigation that benefits humans.

Authority

Authority pertains to the weight and importance a Sphere of Influence places on the flow of information coming from the system. Stock Exchanges have an enormous Authority regarding the price discovery of a stock, in fact it is often the ONLY source of price discovery, and thus the risk of errors, omissions, failures, misinformation is enormous. Amazon on product pricing, Google on "internet searches" rendering information as a primary source and the risk impact from wrong or biased information is high. This system requires an awareness at the individual level that is perceptual and unlikely to be well quantified by assessors. This measure of authority is domain-dependent and comparisons across domains are only relationally relevant.

FOR HUMANITY

Moreover, as they grew, gaining Authority, these entities needed to take proactive steps to mitigate risk resulting from their positions of Authority.

1. Not Trusted
2. Skeptical but an input to information gathering
3. One of many sources (with this source having high reliability)
4. A key resource for information receipt
5. A primary source of information

As a system climbs this scale of authority, the Spiderman principle applies “With great power comes great responsibility”. This process requires an entity to simultaneously engage in self-awareness and external feedback as to their Authority status as determined by Diverse Inputs and Multi Stakeholder Feedback. Each time the collective assessment increases to a higher threshold, the entity should reassess risk and potential negative impact. Illustrating the risks of increasing Authority regarding information could mean higher standards, increased fact-checking, obligations to provide two sides of a debatable issue (or fair consideration of options), avoidance of hyperbole and robust, transparent distinctions between fact and opinion.

Dependency

Systems, where users may opt out, avoid or otherwise completely eliminate their need for a system are inherently less risky than systems to which they are beholden, such as government services or at their place of employment, where opting out is not a free choice without cost. Dependency has an impact on the riskiness of the system.

The measurement of Dependency however is a one-sided evaluation - leaving “this” system will cost me how much? Does it cost me my job, my social networks, forms of ex-communication, and/or reduced access to goods and services? Or worse, do I have no choice at all, such as Border Control, Social welfare, Government surveillance, where Dependency equals the inability to enter a country, receive services or even leave home. The measurement of Dependency captures the perspective of the Sphere of Influence and the exit costs associated with an opt-out choice. The higher the exit costs, measured comprehensively, then the greater the Dependency and greater the risk inherent in the system.

Exploring this further - opting out of a social network. Clearly, I am able to opt out, but does that hinder my ability to effectively engage with my friends who primarily use that system as their form of communication, maybe the social network was a place to conduct business. Opting out has costs and if those costs are too high, then do we really have a choice? So when we examine this scale, we want to recognize the exit cost to determine if risk to people has increased as a result of Dependency on the system.

1. Full freedom - no Dependency
2. Freedom to opt-out - limited discomfort
3. May opt out - meaningful social cost
4. May opt out - meaningful economic or reputational harm
5. May opt out - hinders rights, freedoms and survivability
6. No Freedom - complete Dependency

This analysis is important to ensure that Dependency on the system is captured and recorded as perceived by each Sphere of Influence. Any system with a high degree of Dependency, with a high exit cost to the user is inherently riskier - one reason why governmental systems pose a high risk.

► Governance, Oversight, and Accountability applied to the SSI?

Independent Audit of AI Systems requires documentary evidence to satisfy all elements of certification criteria. When engaging in systemic risk evaluation, the entity is expected to apply their Code of Ethics and Risk Management Framework upon the process. Accountability for conducting the Systemic Societal Impact Analysis rests with the Algorithm Risk Committee and the Ethics Committee - requiring objective application and assessment.

In addition, this risk assessment establishes ongoing thresholds that must be monitored and maintained. When thresholds are crossed/breached the Algorithmic Risk Committee should initiate a re-assessment of risk. This re-assessment must create a cascading effect including the identification of new mitigations and their subsequent deployment. It is a complete reset on the risk assessment feedback loop for Risk Inputs and subsequent Risk Treatment.


“Systemic Societal Impact, also should feed into the Risk Management process regarding Risk Appetite and Risk Tolerance”

► Systemic Societal Impact and Risk Tolerance/Risk Appetite

As more and more thresholds are crossed by an entity, increasing their Systemic Societal Impact, this should also feed into the Risk Management process regarding Risk Appetite and Risk Tolerance. There is a further, all too common risk embedded in a rise in Systemic Societal Impact - the commensurate rise in Risk Appetite and Risk Tolerance of an entity. Call it “hubris”, attribute it to “wealth”, confirmation bias on the “rightness” of the entity’s choices or the simple financial ability to absorb risk, this parallelism is always negative to humans. Again, the Spiderman Principle applies, “with great power comes great responsibility” - this is the only path towards sustainability.

Entities, despite apparent public reverence and self aggrandizement, are at their core societal constructs allowed to exist to produce benefit for humans. As the concentration of that benefit narrows, societal tolerance wanes. Entities must be held accountable for a failure to reduce Risk Tolerance and Risk Appetite as they increase in their Systemic Societal Impact.

This is a self-awareness examination, one that will likely increase the need for Risk Assessment and Risk Treatment while simultaneously suggesting a dynamic modulation in Risk Appetite and Risk Tolerance in the context of a robust, evolving Code of Ethics reflective of the organization’s true reason for being, granted by people for the benefit of people.



It is precisely this type of assessment that entities tend to avoid because the pressures of capitalism demand growth in earnings and profitability, precisely at the moment when entities should responsibly elevate and celebrate their users, clients and prospects from whom trust, and commerce have been granted.

ForHumanity and Independent Audit of AI Systems calls on entities that are uncovering meaningful systemic societal impact to focus on sustainability - their purpose in society beyond the dogma of straight capitalism and strive for a balance that allows for power, combined with social responsibility.

The SSI operated correctly and objectively, like all self-assessments, can provide early warning signs that the entity is moving into elevated levels of risk with the potential to create significant negative impact across any and all Spheres of Influence. ForHumanity suggests that these entities accept their responsibility attributable to Importance, Authority, Dependency and Saturation and act prudently, otherwise humans will be at-risk and other remedies will have to be applied.



FORHUMANITY



Images: Freepik