The Peace Innovation Lab (PIL) was incubated by Stanford University's Persuasive Tech Lab starting in 2008. Formally launched in 2010 by Mark Nelson, Margarita Quihuis, and team, the PIL has a goal that can colloquially be summarized as: to identify the means by which technology can help improve the quality of human interactions. There is no question that technology has improved the "quantity" of human interactions, but whether that has resulted in "better" human interactions is debatable. For example, the #1 application of Artificial Intelligence remains the algorithms used by ecommerce websites to market products, i.e. one of the most powerful technologies of the 21st century is mostly used to make you buy things that you don't need.

In reality, today technology is giving us powerful tools to observe, model, understand and influence the behavior of people. This is the first time in history that we can utilize so many sensors to measure the various aspects of "peace".

The PIL believes that all technologies are "persuasive": any new technology will influence the behavior of its users. It becomes a question not of whether the technology is "persuasive" but of what kind of "persuasion" it should be used for.The unintended consequences of a technology can be devastating, and difficult to roll back. There is no "undo" button in society. Vice versa, some technologies can be powerful media to produce peace. The PIL studies what kind of technology is best for producing peace in a specific environment.

The word "Peace" actually refers to a broad spectrum of human interactions. At the most basic level, "peace" is simply the absence of conflict. In this sense, two machines that interact without hurting each other are at peace—as are two machines not interacting at all. But as you move up the ladder of positive engagement, "peace" becomes an increasing degree of collaboration, with increasing degrees of benefits for the collaborating parties. When the whole becomes greater than the sum of its parts, the collaboration gets particularly interesting, it becomes a form of symbiosis. At that point the "absence of conflict" is more likely to become permanent rather than episodic. Vice versa, war is often the effect of non-mutually beneficial trade.

One key concept on the way to collaboration is "trust". Therefore the question "can technology improve peace" can often be translated, at least initially, into: "can technology improve trust"? The so-called "sharing economy" is largely based on trust algorithms designed to enable the collaboration between parties that never met before. In fact, most ecommerce transactions involve some kind of trust algorithm - at the simpler level, the customer reviews of a service. The designers of trust algorithms (such as Colin Rule, who was Director of Online Dispute Resolution at eBay and PayPal) are de facto "peace engineers".

Mark Nelson's past career as a humanitarian aid worker made him realize that sometimes "aid", no matter how well-intentioned, does not target the root cause of the problem, and sometimes may even lead to negative consequences. The first-order impact of an action on a society (e.g. distribution of food) is easy to predict. The second, third and following orders are not so easy to predict without a scientific model of the society. For example, distributing free food may derail the local agriculture and cause longer-term famine. The PIL aims for a quantitative approach to "peace". The first step is to model

mathematically the human interactions taking place in the society. The second step is to devise actions that will improve the “positive” human engagement in that society. The third step is to measure the results.

The first case study of the PIL was revealing. The boom of Somali piracy that took place in the 2000s had to do with economic conditions created by the collapse of the national government. Fishermen and farmers were displaced by an unlawful use of their waters, lands and resources. Local fishermen started militia coast guards, some of whom then became pirates.A large share of the trade between Asia and Europe flows through the strait between Somalia and Yemen. The economic incentive to become a pirate was obvious. There were other indirect incentives that worked to prevent a solution to the problem: for example, mercenaries hired to defend the cargoes had a vested interest that piracy continued; and insurance companies profited from higher premiums. The lesson learned from Somalia was that economics is one of the strongest sources of conflict, and that measuring trade activity in a region provides a good approximation of human interaction. PIL’s analysis of the problem was relatively simple: the world needs to make it more lucrative for people to invest in peaceful economic activities rather than piracy, for example by instituting a prize in Somalia for the best ideas on how to solve the problem of piracy. Local people tend to come up with the best solutions, if only they are motivated to do so. Pirates themselves might compete for that prize.

The PIL, however, also learned a negative lesson: peace technology requires a large sample of cases. There weren’t enough conflicts like the one in Somalia for analysts to perform comparative analysis and derive generalizable results. The PIL then turned to the institutions that have a lot of data about conflicts: cities. As it turned out, many military strategists such as David Kilcullen (“Out of the Mountains: The Coming Age of the Urban Guerrilla”, 2013) think that most human violence happens in cities. (According to the Bureau of Justice Statistics’ National Crime Victimization Survey of 2015: "the rates of violent crime and property crime were higher in urban areas than the rates in suburban and rural areas"). The PIL was also concerned that persuasive technologies can have unintended consequences: it therefore makes sense to do many small randomized controlled field tests in cities across the world before rolling interventions out globally.

The first PIL city lab was in collaboration with Paris, followed by Buenos Aires, Santa Cruz in Bolivia, Berlin, two in Denmark, etc. The focus on cities helped the PIL realize that steps required to a proper analysis of the problem, for example the need to define what is peace data; the need to understand how people use technology; the need to understand what kind of boundaries are dividing groups (humans seem to be genetically programmed to divide the world into tribes); etc. The city lab deploys sensors in the city to measure behavior; it then uses a combination of behavior design and persuasive technology to design interactions that can increase positive engagement. It aims for a cybernetic loop that progressively designs augmenting technologies that increase mutually beneficial behavior. By their collaborative nature, cities should be natural “peace engines”.

Since then, the PIL has started developing quantitative, predictive, computational methods to sense engagement levels and interaction quality across group boundaries. The ultimate goal is to devise systematic, modular, scalable methodologies to shape behavior that is profitable to both sides of

potential conflicts, enticing global capital markets to reallocate assets towards the solutions that the PIL's network of labs discovers and develops.

The design loop invented by the PIL involves a new form of design thinking that is cooperative instead of competitive. It can be seen as an evolution of the "design science" coined in the 1950s by Buckminster Fuller (who gathered experts in different disciplines to solve complex problems) and of the subsequent Scandinavian design school (that gathered ordinary people to participate in designing the solution). The PIL wants individuals of the community to analyze how they could be valuable to each other and, in a sense, to "design themselves", a veritable case of self-ethnography. Open, crowd-sourced innovation can lead to a spiral (a positive feedback loop) of positive human engagement. Design thinking is usually competitive and therefore not scalable. When design thinking is cooperative, it becomes scalable. More importantly, the user retains the agency and the control over the process. The safest way to deploy persuasive technologies is to let the community control their deployment.

Much of what has happened over the last century in Physics and Computer Science was due to a multidisciplinary "Big Science" approach. Ernest Lawrence at UC Berkeley in nuclear physics and Vannevar Bush at the MIT in the field of computation can be credited as the originators of the "Big Science" approach to solve very complex quasi-intractable problems. Later, the Macy Conferences on Cybernetics, held from 1946 until 1953, lay the foundations for the discussion on the interaction of technology and society that is still raging today. Those conferences featured scholars from such diverse disciplines as computer science (John von Neumann), mathematics (Walter Pitts), philosophy (Norbert Wiener), physiology (Arturo Rosenblueth), Ralph Gerard (neurophysiology), (neuroscience (Warren McCulloch), anthropology (Gregory Bateson), psychology (Lawrence Kubie), sociology (Lawrence Frank), etc. The PIL approach to peace innovation, which is in fact based on a cybernetic loop, is similarly unbounded, and elicits contributions from sociologists, psychologists, economists, game designers, historians, computer scientists, artists, etc.

In addition to the multi-disciplinary approach, the PIL fosters by its own nature a geographically-distributed paradigm because of the need to collect data and run the same experimental protocol in multiple locations in order to infer general rules.

Interaction design is not a perfect science, but technology can increasingly provide means to measure, compute, implement and test, i.e. the fundamental steps of any scientific experiment. While we are a long way from having a science of peace, we feel confident that the tools are out there, and that enough theory has been accumulated in multiple disciplines to help frame the process of peace innovation in scientific terms.