



Unveiling the commercial value of the responsible use of AI: A collaboration between Loomi and Digital Catapult

February 2021

Machine Intelligence Garage cohort company Loomi explores the responsible use of AI and the impact it can have on its business.





Contents

3 Foreword by Luciano Floridi

4 Introduction

5 About AI Ethics Deep Dives

7 Digital Catapult's AI Ethics Framework

8 About Loomi

9 Ethics roadmap and milestones

First milestone: recommendations for the subsequent month

Second milestone: recommendations for the next quarter

Third milestone: recommendations before business model change or soft launch

Fourth milestone: recommendations for every six months

Fifth milestone: recommendations to consider within a year

11 Discussion and interpretation of outcomes

Loomi's enhancement of transparent communication

How Loomi enhanced trustworthiness, and how the team continue to make responsible data decisions

How Loomi developed clarity around benefits

How Loomi ensures long term acknowledgement and management of risk

How Loomi has integrated product design decisions to enhance trustworthiness

Collaborative expectations between Digital Catapult and startups

16 Feedback from the Loomi team

17 Final remarks

18 Endnotes



Foreword

By Professor Luciano Floridi

(University of Oxford, Chair of the Machine Intelligence Garage Ethics Committee)

There is an evident need to apply and refine the responsible use of artificial intelligence (AI) in practice. With the risk of a growing and individuals becoming increasingly cynical about the value of digital technologies, yet increasingly relying on them, it is essential to ensure that AI systems are designed and re-designed with ethics in mind. This will provide truly beneficial products and services, with no counter effects, thus supporting strong, sustainable businesses, while fostering a fair society. This case study highlights how the practical application of ethics can provide innovators with the tools they need to make thoughtful and mindful decisions relating to the design, development, and deployment of their AI products.

This case study is an exercise in realism. It is important to recognise that there are challenges to implementing ethics, not least the tension between short-term costs and potential longer-term benefits. However, this case study also shows why time and energy spent now can reap large dividends in the future, while any delay in engaging with ethical opportunities and risks is short-sighted.

The vectors that propel history into the future are countless and interact with each other in very complex ways. However, some of these vectors and their interactions are up to us to determine and change. And insofar as it is possible to shape and influence the tech trends of the future to be pro-ethical, startups are arguably well-placed to lead this, because it is easier to embed good practices from the beginning than to retrofit them later. Start ups are constrained in terms of resource, so the path is unclear, therefore the ethics journey outlined in this paper is to be interpreted as tentative: Digital Catapult and Loomi are testing and piloting ethics in practice. It is evident that ethics has shaped and enabled a team to create a valuable AI product, whilst enjoying positive commercial outcomes. This is the story of only one startup, but we hope that it is also the first chapter of many more thoughtfully designed and conscientious AI products. This case study shows how much can be achieved through good will, intelligence, and cooperation. The future of AI is clearly open to many more successes like this.



Introduction

In recent years AI ethics has become a major buzzword. As of 2020 there were over 160 published guidelines on building ethical AI. These frameworks, guides, and methodologies have been central in moving us closer to understanding of what responsible use of AI is, but more needs to be done to help businesses take a practical approach to integrating ethics in AI products.

Digital Catapult, the UK's leading independent research and innovation centre for advanced digital technologies, is testing methods to achieve responsible AI in practice. Digital Catapult has appointed an independent, expert committee of AI ethics specialists, who have helped us devise and implement ethics support for machine learning startups.

This report showcases the collaborative work between Loomi, an AI startup, and Digital Catapult and its Ethics Committee. The Loomi team focuses on machine learning to offer virtual personal assistants and other highly specialised language processing Services. It highlights the value that ethical approaches offer businesses developing AI products and demonstrates how long term commitment to ethical processes or methodologies can help AI companies to achieve positive commercial outcomes. The CEO and founder of the startup believe the significant benefits of using AI responsibly through this process were to:

- attract investors
- increase the company's client base
- enhance the user experience design of an AI platform
- increase the predictive accuracy of AI algorithms
- enhance explainability and transparency of inputs and consequent outputs of AI algorithms

By sharing a company's navigation through responsible AI use, Digital Catapult hopes to inspire other organisations to be candid in sharing learning curves, near misses and achievements within this space. These learnings will accelerate practical AI ethics, and enable companies to translate values and principles into action and practice. While this is only the beginning, and these reflections are specific to one UK based startup, it is hoped that this will serve to illuminate the importance and value of embedding responsible usage into the early stages of AI design.

Digital Catapult's AI ethics service for startups, is part of a portfolio of activities aimed at leveraging the unique capability and appetite of the UK AI ecosystem¹ to grow responsible AI adoption. Other components include our industry working group for responsible AI adoption², research and development of tools for responsible AI³, and research to set the theoretical foundations of this work in collaboration with our academic partners⁴.



About AI Ethics Deep Dives

Deep Dives are available to Digital Catapult's cohort of startups as part of the Machine Intelligence Garage,⁵ a machine learning acceleration programme for early-stage machine learning startups.

The AI Ethics Deep Dive process spans approximately 12 months, led by the Ethics Committee, who have an extensive range of data ethics expertise. The Ethics Committee is chaired by Luciano Floridi, Professor of

Philosophy and Ethics of Information at the University of Oxford. Luciano is also a Turing Fellow and Chair of the Data Ethics Group for the Alan Turing Institute. The full list of the Ethics Committee Members at the time of publishing can be found [here](#).⁶

The Deep Dives are flexibly designed, and tailored to meet the needs of startups.

A typical timeline of a Deep Dive may look like the following.

First 1-3 months:

Ethics roadmap is defined in collaboration between the ethics committee and the startup. Several half day workshops.

Months 3-10:

Additional support and advice in actioning ethics recommendations and milestones. Quarterly meetings.

Last 10-12 months:

Conclusion, next steps, and case study.



About AI Ethics Deep Dives

The AI Ethics Committee provides startups with the following.

- **One-on-one bespoke company ethics consultations.** Ethics Committee members will learn about the startup in advance, preparing company-specific questions and areas of ethical interest to discuss during the consultation. The consultations are very collaborative, informal and encouraging.
- **AI Ethics Sprints, in the form of half-day workshops.** The workshops are tailored to the specific needs of startups, informed by the initial kick-off consultations.
- **Bespoke company interviews and surveys.** Gathering more information from various employees and functions, to advise the approach with the best knowledge and information possible.
- **Recommendations for, tools and techniques** that will be relevant at multiple stages of development.⁷

- **Responsible adoption of AI roadmaps for startups.**

After diving deep into the specific startup's stage of development, company culture, technical and business needs, the expert will develop a defined ethics roadmap for the company. Each stage of the roadmap involves a series of milestones for progressing ethical development. Each stage involves a consultation with members of the Ethics Committee who recommend tangible and actionable milestones that could positively impact their products and services, company culture, and business outcomes.

Running the Deep Dive over a year ensures that the participants are conducting a meaningful and sustained review of every aspect of their business, including benefits, risks, responsible data use, trust, diversity and inclusion, transparency and business model. Implementing and reaping the benefits of the roadmap recommendations does not happen overnight. A year-long engagement enables the companies to consult with the Ethics Committee on implementing their recommendations, and adapt them incrementally.



Digital Catapult's AI Ethics Framework

Digital Catapult's initiatives in practical AI ethics, including the Deep Dives, are all based on its AI Ethics Framework - the first iterative and practical guide to be designed and developed to meet the needs of early-stage startups developing machine learning solutions across the UK.⁸ The AI Ethics Framework was developed in collaboration with the Ethics Committee, and translates high-level principles into practical questions that illuminate how they are relevant to business, people and technology decisions. This is the approach later taken by Europe's High Level Expert Group for Trustworthy AI.⁹ This translation process is a critical first step in putting principles into practice. The Framework has been designed to ensure that companies are encouraged to have dynamic and iterative conversations about using AI responsibly, with a view to implementing any processes or changes necessary. Each of its core principles has an associated set of questions to facilitate a reflective, consultative and practical approach.

The AI Ethics Framework was developed in collaboration with Digital Catapult's Ethics Committee, and is updated bi-annually as issues become pertinent and new challenges arise. The Ethics Framework provides a set of considerations that enable startups to adopt a constructive ideal of responsibility - doing good where possible, instead of a constrained ideal, which solely focuses on following legal rules. The Framework asks questions that will provoke thoughtful discussion and the considerate development and deployment of machine learning technologies. It provides companies with a process that will facilitate conscientious, practical decision-making.

Principles and example questions within the framework are:

Be clear about the benefits of the product or service

- How can the products or services be monitored and tested to ensure they meet these goals, purposes and intended applications?

Know and manage your risks

- Who or what might be at risk from the intended and non-intended applications of your product/service?
- Consider all potential groups at risk - whether individual users, groups, society as a whole, or the environment.

Use data responsibly

- Are the data uses proportionate to the problem being addressed?
- Can individuals remove themselves from the dataset? Can they also remove themselves from any resulting models?

Be worthy of trust

- Within the company, are there sufficient processes and tools built-in to ensure meaningful transparency, auditability, reliability and suitability of the product output?

Diversity, equality and inclusion

- Are there processes in place to establish whether the product or service might have a negative impact on the rights and liberties of individuals or groups?

Transparent communication

- Is there a communication strategy and process if something goes wrong?

Consider your business model

- What happens if the company is acquired? What happens to the data and software, for example.



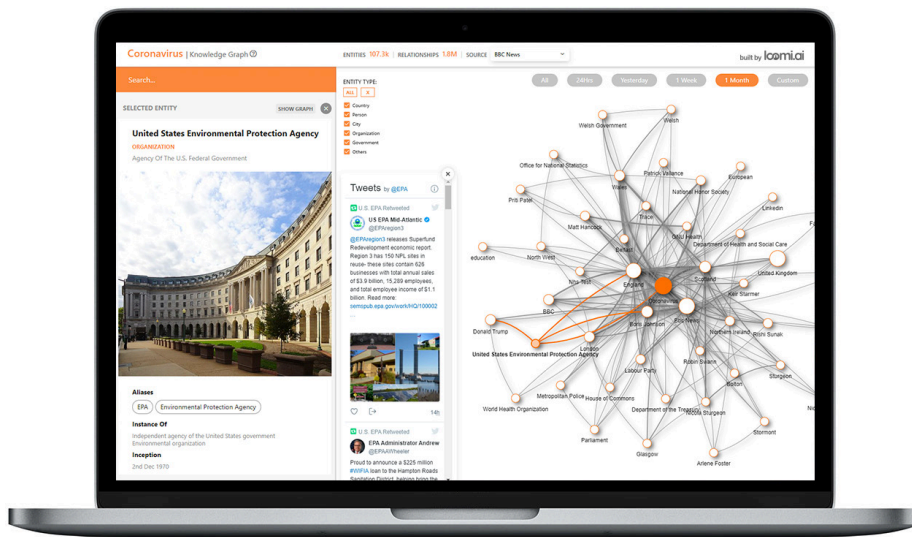
About Loomi

Loomi automates data management and offers virtual personal assistants by utilising a diverse set of highly specialised language processing services to both a B2B and B2C clientele. From helping media and financial services companies automate data management by creating custom knowledge graph ontologies used to improve accuracy in content categorisation; tagging and building highly contextual recommendation engines, to the creation of custom artificial intelligence personal assistants based on these ontologies. These AI-powered assistants seamlessly orchestrate the flow of all information, from email to news, into one platform. Aiming to operate in a similar way to the most qualified personal assistants, the Loomi platform learns what is important to its users, prioritises this information and ignores other content or messaging - saving users time from sifting through extensive content. Loomi consensually takes information from various sources and presents the relevant information to its users in a filtered and prioritised way, using three core function areas: to do, to know, and to learn. Loomi believes this will help users with information overload, contribute to increased productivity and decrease stress. This tool

might be especially useful for other startups and SMEs, who often have to deliver projects under tight timelines using small teams. A great example of what Loomi can do is demonstrated in its freely accessible coronavirus research tool, available [here](#).¹⁰ More information is available on Loomi's [website](#).¹¹

The Loomi team consists of approximately 10 members and it has a number of corporate clients from media to financial services.

Loomi is committed to ethics for a number of reasons; there is a desire to be socially responsible with advanced technology, and to 'do good where possible', as well as understanding that using AI responsibly may be a conduit to enabling its product to be more successful. An AI platform, such as Loomi's, requires an abundance of trust from its users - from the data privacy of its users to ensuring the product creates reliable and accurate outcomes. Without the public perception and reality of both, the success of a product may be limited. For these combinations of motives, Loomi undertook the Deep Dive programme with Digital Catapult.





Ethics roadmap and milestones

This section will provide the details of Loomi’s ethics roadmap. The roadmap was developed by the Ethics Advisor leading the Deep Dive. The roadmap was created through thoughtful consultation with Loomi, and following the broad principles outlined in Digital Catapult’s Ethics Framework.

The consultation highlighted a few areas of focus: enhance transparency of communication; enhance trustworthiness and continue to make responsible data decisions; develop clarity around the product’s benefits; acknowledge and manage long term risk; integrate AI ethics into the product design process.

The milestones below address these focus areas, suggesting actions that are aligned with Loomi’s product development roadmap at the time.

FIRST MILESTONE: RECOMMENDATIONS FOR THE SUBSEQUENT MONTH

- Develop clear policies and documentation relating to user information before the beta goes live.
- Consider how best to communicate with users and others about the way data is used in the product, and how the Loomi system works. External review for understandability for anyone without intimate knowledge of the product.
- Ensure terminology and features are meaningfully explained. Semantics that come up often, such as

“metadata”, and view features such as “my graph vs whole graph”, or indicators such as “similar people”, “similar attributes” which appear on Loomi’s platform, need to have explanations.

- Support team to have training in information security, data protection, plus training and support about what to do if they encounter information that affects them or may be criminal or concerning in nature. It should be clear to users, even at beta stage, who is supporting them, what personal information access they have in what situations, and how to get support or ask a question.

SECOND MILESTONE: RECOMMENDATIONS FOR THE NEXT QUARTER

- Explore potential consequences of Loomi product use for different stakeholders. This may include utilising methodologies such as Doteveryone’s Consequence Scanning¹² tool.
- Consider how to engage beta users to get meaningful feedback relating to questions of user diversity, information presentational risks.
- Ensure the proper documentation of datasets, in its motivation, composition, collection process, recommended uses. Facilitate better understanding between developers and users, and look to prioritise transparency and accountability. One possibility is using the “datasheets for datasets” methodologies.¹³



Continued: Ethics roadmap and milestones

THIRD MILESTONE: RECOMMENDATIONS BEFORE BUSINESS MODEL CHANGE OR SOFT LAUNCH

- Run a workshop with external stakeholders and social science expertise to explore the different information, control and trust dynamics for the planned market and business model. This workshop should include potential effects on different stakeholder groups, including those who do not use the product, and also scenarios where the product. Scenarios would include a product becoming very successful and dominating a market or being less successful with a risk of being discontinued.
- Scenario plan stress exits and business failure: how this affects the product, users/customers, access to datasets.
- Back up marketing claims with evidence: For example, if the product claims to benefit users by reducing online addiction, how is this evidence validated?

FOURTH MILESTONE: RECOMMENDATIONS FOR EVERY SIX MONTHS

- Alongside planned security reviews, Loomi is advised to undertake an ethics review approximately every six months. If significant product changes are planned, a review should be aligned with the development planning stage.

FIFTH MILESTONE: RECOMMENDATIONS TO CONSIDER WITHIN A YEAR

- Consider engaging with academic researchers to explore the consequences of the product further. For example, the benefits and risks of products and evaluating potential bias.



Discussion and interpretation of outcomes

This section will discuss the outcomes and interpret the impact of the ethics roadmap. Each key area of outcome and discussion will be signposted by the ethical principle outlined above in Digital Catapult's Ethics Framework. For clarity, the principle in each subtitle is italicised. The outcomes are broadly chronological and in line with the above roadmap; this section navigates the results of following the ethics recommendations.

The Loomi team states that a huge challenge, especially in the early stages, can be prioritisation - ethics alone will demand significant attention. It is essential to combine the most important ethical considerations with the product roadmap, this ensures the risk register is enhanced with the most pertinent ethical challenges that may arise and team members are methodological in how to approach them. Loomi highlighted that the Deep Dive allowed them to consider in detail how to action Digital Catapult's Ethics Framework to enable concrete changes. Loomi discusses the importance of having the opportunity to institutionalise ethics within its business, and having a clear direction in how to convert ethical ideas or principles into reality, through design and product decisions.

LOOMI'S ENHANCEMENT OF TRANSPARENT COMMUNICATION

As a result of the advice from the first milestone in the roadmap, Loomi wanted to be more transparent in communications, presenting the benefits and risks of its product clearly, to ensure users can be aptly informed. Consequently, Loomi changed the entire way it communicates the product. Initially, the focus centred around ethics and responsible development. Now Loomi is a very conscientious communicator, especially in simplifying and utilising non-technical terms. This challenge of communication is not to be overlooked. Digital Catapult has worked with many highly technical teams. It is commonplace for language or explanations that appear obvious to like-minded technical teams, to be opaque for a non-technical audience. This communication barrier can have an impact on a company's ability to raise investment, win new clients, and provide a high quality user experience. The Loomi team quickly understood its product needed to be communicated clearly, in plain English.

Clear communication was also critical to ensure transparency and trust for Loomi's client base. Loomi added certain design features to enhance this: this. As an example, for each key recommendation or piece of information, the user can click on a "why am I seeing this?" functionality, which highlights the relevant inputs and outputs as an explanation. This re-framing of the communication around the product enabled Loomi to emphasise customer-centricity over product-centricity. The results of this were evident in the initial development stages after the Deep Dive, as Loomi rapidly expanded its client base. Its number of users grew quickly, as Loomi experienced an increased 45% conversion rate of leads visiting the website to sign up as customers. In this one small example, the need to use AI responsibly



Continued: Discussion and interpretation of outcomes

is clear - not only does it encourage startups to innovate responsibly, it can also create positive business outcomes. By focusing on the ethical principle of transparency in communication, Loomi were able to attract customers with enhanced communication, and increase its client base rapidly. Of course, it is difficult to attribute the expansion in the client base solely to enhanced communication, the Loomi team highlight that many other variables remained constant at this time. This communication strategy was the biggest change to its business in this period, suggesting a correlated success.

HOW LOOMI ENHANCED TRUSTWORTHINESS, AND HOW THE TEAM CONTINUE TO MAKE RESPONSIBLE DATA DECISIONS

In light of this customer expansion, Loomi also acquired more data. The team understood the importance of using data responsibly, and even though GDPR is an essential starting point, they felt privacy features could go above and beyond the minimum set by the legislation. As previously explored in AI ethics literature, GDPR does not cover all data management questions, often leaving a lot of 'grey areas'.¹⁴ To fortify its privacy and security features, Loomi created different layers of permissions in the database, only granting access to selected individuals at specific data points. Loomi's security processes ensure that raw email text is never stored within the system, to enhance privacy. Instead, the pertinent information is extracted, with the full, raw email discarded from its internal database. Loomi also ensures that information linked to individual users is encrypted.

With much more data and now a growing client base, client trust and responsible data decisions remain paramount. These principles and values have intrinsic appeal, as well as being commercially important: if

customers feel they cannot trust the platform, for example due to inaccuracies, there can be high churn - especially in competitive arenas, such as virtual personal assistants.

The roadmap for the responsible use of AI encouraged the startup to make tangible actions to enhance trustworthiness. The need for high accuracy became evident, as any misclassification and miscategorisation of information could diminish trust in the product and service, and potentially face mistrust from users. Loomi looked to enhance user trust by prioritising the product's high accuracy rates, improving entity matching, categorisation, and tagging by providing custom ontologies and taxonomies. The Loomi team eventually increased its model's predictive accuracy rates from 60-70% to 98%, sometimes even higher for specific domains rich in public information. As part of developing this new stack, Loomi progressed from a simple taxonomy to building a more sophisticated, underlying ontology, as a general reference point for organising knowledge. In brief, while taxonomies identify hierarchical relationships within a category, ontologies offer a higher level of sophistication into a larger structure, including associated words and categories.¹⁵ The benefits of using ontologies for virtual assistants include enhancing inference capabilities, and expanding contextualised knowledge, such as answering questions that have not been pre-programmed into the model. This creates a more powerful product because it can: 'understand' how each element relates to one another, capture intent in natural language and is also more scalable and cost effective.¹⁶ The desire to enhance trust with its users and as a brand, gave Loomi the impetus to approach its technical development in this manner, demonstrating where ethics considerations are aligned with business and product quality.



Continued: Discussion and interpretation of outcomes

With the ethics recommendations to look deeper into its data sources, a few issues became apparent. In the early stages of machine learning development, it is imperative to be aware of potential bias within data: bias could impact on the quality of service for certain groups of society, and directly result in discriminatory treatment.¹⁷ Whilst the expectation was to find bias, as the prototype was trained using the Enron emails dataset, the prompt and guidance in examining data sources in more detail accurately identified how the dataset was biased. The types of users and speech within the Enron emails were biased towards specific corporate language, including jargon and colloquialisms, which are only relevant within financial services. The Enron dataset includes emails from about 150 individuals. However, the Loomi team quickly understood that they needed more data, especially types of writing and speech that were contextually different - more informal, personal, and not commercially driven. At this stage in the roadmap Loomi discovered that the way the data was structured had limitations, especially in a B2B context. Consequently, it is useful to understand how looking at ethics in the early stages of machine learning development can advance a product or service. Through close examination of the potential biases present in original data sources, Loomi improved the accuracy and efficacy of its product simultaneously.

The Loomi team highlights conversations with investors gaining traction at this stage, and the founder and CEO reports the team's ability to handle difficult questions around data management, privacy, ethics and compliance to be a major selling point. Loomi's founder and CEO highlights that being part of an ethics Deep Dive, with third parties examining the intricacies of the company's machine learning development continues to place them at a unique advantage.

HOW LOOMI DEVELOPED CLARITY AROUND BENEFITS

The ethics roadmap encouraged Loomi to think further around the benefits they were presenting to users. Well-intentioned assumptions around the efficacy of a product do not suffice, so the milestones on the ethics roadmap ensured Loomi always considered how its product benefited users and looked to find evidence to support these assumptions. As an early-stage startup, it is not easy to collate a large enough sample to estimate causality and impact. This has continually pushed forward the need to collect more data, and ensure more predictive accuracy. Claims on small samples of B2C efficacy is difficult, but is much more achievable in its B2B application - as there are clear metrics of performance, with a much larger sample size. Indeed, when Loomi pivoted to a larger B2B clientbase, it was able to satisfy the criteria to confidently attribute causality between variables and ultimately support claims with evidence.



Continued: Discussion and interpretation of outcomes

HOW LOOMI ENSURES LONG TERM

ACKNOWLEDGEMENT AND MANAGEMENT OF RISK

The ethics roadmap encouraged Loomi to consider implementing a bi-annual ethics review. In working with a number of machine learning startups, programmers and computer scientists are often in search of a finite checklist or tickbox, to measure if their company is ethical. The reality is that checklists of this sort do not instil a commitment to ethical evaluation within a company, which will lead to a long-term, concrete business responsibility. The search for companies to be branded as ethical or fair is perhaps less relevant or useful. Rather, objectives should centre around forming design and business choices that consistently make fairer or more ethical products and services.

This makes periodic responsible AI reviews such as these important. Loomi found that the Deep Dive instilled a level of critical analysis to day-to-day data science and product design, which has inspired the Loomi team to be conscientious innovators. A bi-annual procedure such as this, ensures a more formal, company-wide, structured review continues to be a priority and ingrained into company culture. Loomi has raised the importance of this, consistently being aware of new or changed risks, and managing them better - to the benefit of its users and company.

HOW LOOMI HAS INTEGRATED PRODUCT DESIGN DECISIONS TO ENHANCE TRUSTWORTHINESS

Alongside Digital Catapult, Loomi took the initiative to work with additional third parties. It collaborated with researchers at Brunel University, taking on an additional ethics-based project. The final objective was to utilise design thinking to foster communication to non-experts about complex algorithms in AI, ultimately improving users' trust and enabling a large range of users to make the most of the service. As a result of implementing design thinking and human-centred design techniques, Loomi ensures its graphic representation of inputs and outputs clearly denotes how it is structured, highlighting the user's individual, personalised data, and where the user can toggle and change preferences on the platform.

As an extension of this, many design choices were undertaken to be worthy of trust from its users. By way of illustration, one design choice includes what users consistently see within its platform. It is common currency for many platforms to create "filter bubbles", whereby a user's beliefs are reinforced by only consistently seeing viewpoints that align. This is based on an assumption of what users want to see and is how many social media algorithms are engineered - with a view to retain users on a platform, for as long as possible. The user has limited access to opposing viewpoints, creating feedback loops and echo

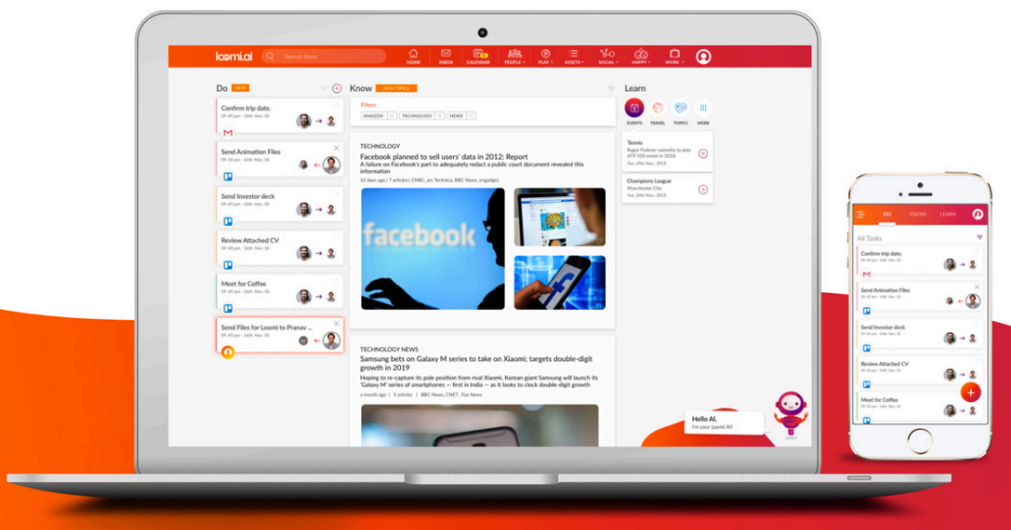


Continued: Discussion and interpretation of outcomes

chambers. In social media teams, the programmers (known as addition engineers)¹⁸ are tasked to keep the user engaged for as many seconds, minutes or hours as possible. Conversely, Loomi made design decisions that turned this trend on its head, mitigating against users spending too much time checking emails, or reading irrelevant news.¹⁹ Loomi gives back this agency to the user, so the design of the product does not rely on the assumption that users only want to see what they have historically been interested in. Instead, Loomi offers a full range of opinions, so users have to modify their preferences explicitly, should they not wish to see contrasting or opposing beliefs. Ultimately, Loomi's objective is not to keep users engaged on its platform for as long as possible, rather to optimise time spent by providing the relevant, pertinent information, so the user can streamline their day as much as possible.

COLLABORATIVE EXPECTATIONS BETWEEN DIGITAL CATAPULT AND STARTUPS

In this report, in partnership with Loomi, Digital Catapult hopes to be as candid as possible, to best advise inspired companies and startups to deal with the challenges of using AI responsibly. Given Loomi's status as a high growth startup, with a rapidly expanding client base, some milestones may still be in the stages of exploration. This is the reality of being a startup, with smaller teams: oftentimes teams are working well above capacity to get product to soft launch. For this reason, Digital Catapult's roadmaps are recommendations rather than specific actions, as startups may not have capacity to satisfy each item at every predefined stage. This will be part of the process and journey for many startups. Ethics roadmaps may be adapted to reflect this within the year of engagement.





Feedback from the Loomi team

Loomi's biggest takeaways from this Deep Dive were the approach and mindset that was learned throughout the process. Loomi stated how ethics became ingrained into its company culture; every time a new challenge crops up, the need to discuss ethics emerges. Loomi illustrates the importance of considering ethics right from the beginning of the machine learning development journey, as this directly impacts the ease with which it was able to onboard the process. The CEO and founder of Loomi, Al Ramich, stressed the positive impact that DotEveryone's Consequence Scanning Tool²⁰ has had on the company's approach, strategy, and responsible decision making. The methodology can be considered as one of the suite of tools in an ethical agile machine learning sprint that the Ethics Committee might recommend to startups in the process.

Al also highlighted that he did not expect the extent to which this ethics dive not only made the team more conscientious, but also improved their ability to grow their client base and to enhance their product offerings. He identified the aspects that the Deep Dive helped with most to be product development and tackling investor questions.

From a business perspective, Al also indicated that although the immediate benefits of engaging in ethics are somewhat invisible, it is very much worth investing the time in ethics for the mid to long-term impact and success it can bring to a product or service.

"10/10 in terms of benefit to our company, 10/10 in terms of service"

**Al Ramich,
Founder and CEO of Loomi**



Final remarks

This report provides an insight into the importance of using AI responsibly for AI companies. The Deep Dive offers companies guidance to critically evaluate the design, development and deployment of AI solutions. Whilst it is hoped that companies will be interested in aiming to do 'good' in society, this report illustrates how ethics also serves a utilitarian, commercially-driven function for scaling AI companies. As this case study demonstrates, it is evident that strong ethics can contribute to an expansion in customer base, enhanced predictive accuracy and valuable user-centric design choices.

Digital Catapult hopes to inspire future startups to engage with AI ethics, and realise the potential it can bring to their AI products.

Startups interested in the Deep Dive process, should apply to join Machine Intelligence Garage, Digital Catapult's machine learning acceleration programme, in which Deep Dives are a core offering for selected companies. More information can be found [here](#), or

reach out directly at:

appliedAIethics@digicatapult.org.uk



Endnotes

¹ <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf>

² Add reference to adoption report

³ <https://github.com/digicatapult/dc-federated>

⁴ For example: <https://link.springer.com/article/10.1007/s11948-019-00165-5>

⁵ For more information around eligibility and how to apply, see <https://www.digicatapult.org.uk/for-startups/acceleration-programmes/machine-intelligence-garage>

⁶ <https://www.digicatapult.org.uk/for-startups/other-programmes/ai-ethics-committee>

⁷ This might range from tools such as Doteveryone's consequence scanning tool, to Aequitas: A Bias and Fairness Toolkit. For a deeper look into existing Applied AI Ethics tools, see Digital Catapult's typology.

⁸ For more information about our Applied AI ethics activities in general, please see our previous publication: https://www.digicatapult.org.uk/news-and-insights/publication/lesson-in-practical-ai-ethics?utm_source=comms_pack&utm_medium=social&utm_campaign=ai_ethics_paper

⁹ <https://ec.europa.eu/digital-single-market/en/high-level-expert-group-artificial-intelligence>

¹⁰ <https://coronavirus.loomi.ai>

¹¹ <https://www.loomi.ai/>

¹² https://doteveryone.org.uk/wp-content/uploads/2019/04/Consequence-Scanning_Agile-Event-Manual_TechTransformed_Doteveryone.pdf

¹³ <https://arxiv.org/pdf/1803.09010.pdf%E2%80%8B>

¹⁴ See: <https://www.law.ox.ac.uk/business-law-blog/blog/2018/10/right-reasonable-inferences-re-thinking-data-protection-law-age-big>

¹⁵ For more information on ontologies and taxonomies, see <https://www.earley.com/blog/what-difference-between-taxonomy-and-ontology-it-matter-complexity>

¹⁶ <https://www.earley.com/blog/what-difference-between-taxonomy-and-ontology-it-matter-complexity>

¹⁷ For example, see Amazon's infamous discriminatory hiring tool: <https://www.bbc.com/news/technology-45809919>

¹⁸ <https://www.washingtonpost.com/news/monkey-cage/wp/2018/08/06/its-no-accident-that-facebook-is-so-addictive/>

¹⁹ "Irrelevant" defined and inputted into the platform by each individual user.

²⁰ <https://www.doteveryone.org.uk/project/consequence-scanning/>



CATAPULT
Digital



Digital Catapult
101 Euston Road
London NW1 2RA

0300 1233 101

www.migarage.ai

www.digicatapult.org.uk