

AREAS	CONTEXT	ETHIC	REALISATION PRINCIPLES	EVIDENCE - FOR MUSIO
Social Impact	Benefits	AI and robotics shall be seen as delivering good. Doing good is one of the five key ethical principles of the EU guidelines for ethical AI. Good includes commercial prosperity.	1 Deployment of AI and robotics shall be shown to improve the well-being of employees and/or the general public, such as improved safety, working conditions, job satisfaction.	One of the core Musio products 'tagging' has been used to replace the monotonous work involved in manual data entry. At the start of the business co-founder Savage manually tagged 2,000 songs, it took 2 weeks and she described it as 'mind numbingly boring and unpleasant'. Our goal is to replace tasks people won't do or can't because of the sheer volume.
			2 Additional to 1. (or instead of), deployment of AI and robotics shall be supported by a business case that demonstrates it improves competitiveness and is not just 'AI for the sake of AI'.	Our commercial success with clients such as Hipgnosis, prove we are adding valuable AI assistance making tasks more sustainable and achievable. Unlocking their revenue potential. <a href="https://blog.musio.com/2021/07/21/hipgnosisuses-musio-to-boost-searchability-and-make-returns-on-investment/">https://blog.musio.com/2021/07/21/hipgnosisuses-musio-to-boost-searchability-and-make-returns-on-investment/</a> Our use of AI is aimed to solve problems that are difficult and cannot be effectively solved using traditional algorithms.
	Human impact	AI systems should be used to enhance positive social change and enhance sustainability.	3 For any deployments, it shall be clear where the human boundary/interface/ interaction is with the AI/Analytics/Robotics system; and any negative/positive impact on human factors and/or human behaviours is fully understood and mitigated where necessary.	We are very open in the explanation of both our Tagging and Search APIs where the documentation is available online: <a href="https://docs.musio.com/tag/">https://docs.musio.com/tag/</a> List of outputs clearly demonstrates what the AI is capable of. The scope of the AI is clearly limited to the classifications done in the specified methods. The implementation needs to be done by companies and their tech teams. We also provide accuracy scores as a % in our products to allow humans to overwrite and build hierarchy systems on top of our data.
			4 Early analysis, in conjunction with human resources and employees (or their representatives), shall be undertaken to identify potential job role changes or potential human resource impacts and the opportunities for retraining or redeployment.	Our CEO (Hazel Savage) started a company with two goals in mind 'not destroy the industry that has given her, her whole career' and to build something that fits within the existing eco-system. At 3.5 years old we have never replaced a human being's job, our tools are designed and built to augment human capabilities. Understanding the music industry is a key part to being a custodian of a Music Tech company. Our CEO has 15 years experience in this one industry.
			5 Potential for upskilling opportunities or redeployment shall be explored with human resources and employees (or their representatives) when any impact on affected employees is established, to ensure that the organisation has the key capabilities needed to secure emerging opportunities in AI and robotics.	In addition to point 4 Musio has also created new jobs in the Music Industry. As the first VC funded Music Tech company in SG, prior to the start of Musio there were very few roles in Music in Singapore. Musio has created an internal 'Music Team' creating at any one time up to 4 full time roles in the team, we have also trained 9 internships in Music Tech through this team.
			6 Analysis shall be undertaken to assess the impact of the deployment on the supply chain - particularly assessing the likelihood for the technology to have a negative impact on the sustainability of any elements of the supply chain. The same assessment should be afforded to customers as appropriate.	Musio holds no external supply chain. Our product is 100% digital, cloud hosted and global. Musio replaces the solution of simply having 'no data' therefore there is no disruption to an existing process. It could be considered that the release of Music is 'supply chain' based. In which case our adding data to this process has a positive impact in regards to discoverability for 3rd parties.
			7 Where there is potential for negative impact on the sustainability of the supply chain, this shall be discussed with the supply chain partner as soon as possible to give them maximum opportunity to adapt to remain sustainable. This same opportunity should be afforded to customers as appropriate.	In regards to sustainability, Musio only has AWS and GCP as external partners and we rely on those multi billion dollar companies to provide an effective and sustainable service. In fact it is more sustainable for us to use the cloud and not build our own data centre, with more wastage.
	Communication	Knowledge of the human interactions with AI should be provided by key stakeholders.	8 Frequent communication and discussion should be had with all key stakeholders - in particular employees and employee representatives - through a variety of channels.	Musio values open and transparent communication. We do daily Tech meetings (stand ups), we have weekly sales/tech crossover founder meetings and monthly we have an 'all hands' presentation by the founders, and we also have monthly Team feedback sessions using MIRO whiteboard software so we can work virtually and also capture all team members' thoughts.
	Loss of skills	AI systems should be used to enhance positive social change and enhance sustainability.	9 Analysis shall be undertaken as to whether any loss/reduction of skills (which result/cannot be avoided) needs to be sustained, for the good of the business, and how this would be addressed.	There is an argument that AI assisted songwriting destroys music skills. Musio does not operate in the generative AI world, we are only descriptive AI. The ability to describe music is a human trait that we have trained and AI to do. The reality is a human tagging 1 track would do a better job, but tagging 2 million tracks a day is only possible with AI.
	Accuracy/Trust	Safety	AI systems should be safe and secure throughout their operational lifetime. This should be verified where applicable and feasible.	10 A formal risk analysis shall be undertaken with specific emphasis given to identifying and mitigating any hazards to human safety.
11 The algorithms shall be assessed for any bias or discrimination impact and their provenance shall be clearly stated to enable any future Root Cause Analysis or troubleshooting (Note, for complex systems, it may be difficult to assess the risk of bias. A new bias assessment template has been created as part of an ecosystem of AI Assurance tools at [link]).				All classifiers are documented and their behaviour is well understood, but source code and models are kept completely confidential to all but the two co-founders as this is the core value in our startup.
Transparency and traceability		AI systems must provide for transparency and traceability of their design, inputs and outputs.	12 To enable the power of data to be unlocked, all training data shall be good quality and representative and its provenance shall be clearly stated to enable any future Root Cause Analysis or troubleshooting.	Making sure that our AI was trained on only legally acquired data sets has been baked into our business model from day 1. When we started Musio in 2018 we did 3 small partnerships giving us legal access to songs in return for use/ access to our product/tags. These included the FMA (Free Music Archive) and two small publishing companies (one in SG one in the UK). All came directly from Hazel's existing industry relationships. Once we had built the core models and MVP we had enough to prove the technology worked.
			13 The hierarchy of decision making shall be clearly stated regarding human v AI.	The Musio AI does not make any impactful decisions. Suggestions are generated by Playlisting, Tagging, and Search AI; however any final decisions and actions are left to our customers. We also build our products such as search.musio.com with clarity in the UX so the user understands the search commands they are creating and the song recommendations they are receiving. Nothing goes to an end user unless approved/selected by the customer. This ensures beginning to end clarity on AI vs human decision making.
			14 It shall be clear what the insight (forecast/decision making etc.) improvement is compared with a human - forecast improvement and actual.	Our improvements in % terms can be viewed on % levels for our tagging product. One is the % time saved using automation. This can be in the realms of thousands of % quicker. It would take a lifetime for a human to manually assess and tag 2m songs, our AI can do it in one day. The second part is the % accuracy. A human with low volume (1 track) can be 100% accurate. With 1000's of tracks we have seen that accuracy with humans drop to 50-75%. With AI the accuracy is consistently around 90% within the defined scope.
Bias		AI systems must be free from bias or prejudice.	15 It shall be clearly stated how any training data sets have been assured to have no unintentional or unethical biases, noting that, for example, if an AI sub-system is being used to detect anomalies, the training set may need a deliberate bias to ensure sufficient amounts of anomalies occur at different rates.	The music industry has biases very specific to themselves. For a long time artists were judged on physical appearance (not something related to our AI), it is also common for more male than female artists to be successful. Our AI does not have access to any visual cues which eliminates many of these biases. Most of our classifiers do not attempt to judge quality but rather absolute features such as Key and Tempo. We put a lot of effort into designing our datasets to prevent unintentional biases and "random" outputs (which can easily occur for poorly specified problems).
			16 A monitor shall be deployed in the system - this is a sense check of the results comparing actual outputs with likely output ranges for the system in question.	We run manual QA checks on all new/updated classifiers we release to ensure that the output lives up to our quality standards and that the models don't generate incorrect/random classification above the specified thresholds. These checks include both predefined sample sets as well as exploratory tests aimed at finding new faults.
			17 A continuous automated monitor shall be deployed in the system to continuously test the system by using existing test/synthesised data, which already has known and approved outputs.	We have existing simple tests that ensure that the output is consistent and does not change over time. We are in the process of adding more granular automated tests to catch issues, downtime, and flag anomalies on a set of test data.
			18 An independent check shall be deployed in the system - assessment of a proportion of the same data using a completely independent assessment mechanism which is already approved. This is a validation check and could be carried out by a human.	We continuously run manual QA checks on samples from customer deliveries often consisting of new unseen materials. These checks do allow us to catch anomalies and errors that we might have overlooked during the build process.
Validity and reliability		For AI to succeed it must be trusted.	19 A process comprehensiveness check shall be deployed in the system - have the right number of assessments taken place?	We have automated end to end tests which cover the full products including all sub-components.
	20 A faultless transmission of data shall be deployed in the system - use of Cyclic Redundancy Checks (or equivalent) where appropriate.		Our AI works on audio files which gives it a low risk of producing harmful outcomes based on corrupted data. We still have basic checks for invalid/corrupted files submitted to our systems.	
	21 The sparseness of the training set of data and its impact on the validity of the output needs to be clearly stated and justified.		Our datasets have been built with the intent of maximising the diversity of content included to eliminate the case of random outputs outside it's domain of knowledge.	
Governance	Data protection	For AI to succeed it must be trusted.	22 It shall be stated whether there is, or will be, any Personal data or not.	As a company Musio does not hold or interact with personal data in regards to AI. With our tagging product, even the original track, fingerprint and data are not stored after the customer has generated and saved the data. Our Search product does host a copy of a client's music database. Music is not considered to be personal information belonging to a company, it is an asset. For any customer data held by other parts of our systems (non AI related) we do comply with GDPR.
			23 The legitimate purpose for using the Personal data shall be declared and confirmation provided that this has been agreed with the person or employee representative where it refers to an employee.	Not relevant due to no personal data being held or part of our business operations.
			24 The architecture of the system shall protect the data from unwanted access without permission - complying with the principle of 'privacy by design and by default'.	Not relevant due to no personal data being held or part of our business operations.
			25 The architecture of any data storage system should have the facility to, on demand, identify an individual's personal data and update, amend or remove every trace in line with privacy requirements and individuals' rights.	Not relevant due to no personal data being held or part of our business operations.
			26 No Personal data shall be sent outside of the relevant, legal zone (e.g. European Economic Area, US).	Not relevant due to no personal data being held or part of our business operations.
	Export control	For AI to succeed it must be trusted.	27 The data flows (including access/reading of data) shall be described to, discussed with and approved by an Export Control manager to assure compliance with Export Control regulations.	For the export of data. The only way to access an export of data is to be a customer using our API solutions. However to generate the data the customer must first upload their audio. By uploading the audio they are agreeing to use our product (alongside our formal contract) and only then is data generated for export. Customers are then at liberty to copy and store the exported data as they see fit. Musio considers this data to belong to the customer. An owner of a catalogue, we consider already owning the audio to then own the data created using the audio.
	Confidential information	For AI to succeed it must be trusted.	28 All confidential information shall be declared to, discussed with and the architectural protections approved by an IT security expert.	Currently as a seed stage funded company (1.2m SGD) who is chasing commercial viability and product-market fit. We do not have the critical funds available to hire a 3rd party security officer to review our product. Furthermore, we do have IP concerns, given that our product is completely unique and highly valuable, and we do not have the finance to defend a legal attack. Therefore we have decided whilst under 5 years old and up until a series B, we will be responsible for and manage our own cyber security to the best of our ability under the stewardship of our CTO and co-founder Aron.
	Cyber security	For AI to succeed it must be trusted.	29 All confidential information shall be declared to, discussed with and the architectural protections approved by an IT security expert.	As above, 28.
	Accountability	Mechanisms should be put in place to ensure responsibility and accountability for AI systems and their outcomes.	30 Ultimate accountability for the outcomes of the AI system needs to be clearly stated with a business owner clearly identified.	Ultimate accountability for the AI and the Musio company data, falls to the co-founders both actively still with the business, known as Hazel Savage and Aron Pettersson. Our seed funded agreement lays out our responsibilities and legal ownership of operating within the law. We also take a personal approach to this responsibility and believe in each other and our ability to build a safe, fair and legal product and company.
	Responsibility for decisions	Mechanisms should be put in place to ensure responsibility and accountability for AI systems and their outcomes.	31 Algorithmic accountability should fall jointly on the developer and tester, or the DevOps team. They shall clearly state how they have assured confidence in the performance of their individual aspects of the AI system.	As above ultimate accountability for the Tech falls to the two co-founders. There is a wider dev ops and tech team, but based on the size of the company and the relatively flat structure. Responsibility can only lie with the founders.
Risks from re-use/ transfer across processes	For AI to succeed it must be trusted.	32 Transferring knowledge between AI systems should be risk assessed using a formal tool/method to determine where and how the system might fail. Any serious events and their causes must be identified along with the method to detect such events. - which shall be formally reviewed before proceeding.	Any transfer of data is done within a single data center. The integrity of this data center is managed by a 3rd party (GCP/AWS). We trust that these organisations are able to maintain a safe internal environment. However we still run our own tests to ensure the integrity of data within the system. This is all managed by our CTO and co-founder Aron, and subject to the usual QA and accuracy testing as all other products.	