

European AI Alliance Input for the First Workshop of the AI HLEG

<i>Session 1: Trusted - i.e. responsible and safe - AI</i>	1
<i>Session 2A: Transparency and Accountability</i>	5
<i>Session 2B: Industry & Ecosystems Uptake of AI</i>	10
<i>Session 3A: Use Cases for the Guidelines</i>	12
<i>Session 3B: AI Enablers & Infrastructure</i>	15

Session 1: Trusted - i.e. responsible and safe - AI

AI Infrastructure and Enablers: What does it mean to ensure trusted – i.e. responsible and safe – AI? Which tools can be implemented to achieve such trust? What does it mean to have trusted AI which supports and enables European businesses to be successful in applying AI? How can the above tools be implemented to achieve such businesses' success?

<https://ec.europa.eu/futurium/en/european-ai-alliance/input-requested-first-ai-hleg-workshop-trusted-ie-responsible-and-safe-ai>

What is urgently needed is to develop appropriate indicators - words like 'trust' have an intuitive meaning for all of ,but something more concrete is also needed.
Approaches such as HRIA can be a useful start (human right impact assessment).
'accountability' (interconnected to 'responsibility') needs to be transcribed at some point into more concrete measure(ment)s.

I.e. Accountable to whom; based on what - applying what standards? etc

Differentiation when it comes to specific applications: algorithmic software vs. artificial intelligence, automated vs. intelligent robotic operation

Interpretability (or explainable) AI: when designing such systems, it is important that the people involved at all levels (developers, stakeholders, users, policy makers) are satisfied with their understanding of how the solution works.

This is complex because different people have different needs for information:

- users are given explanations in natural language or visual cues, that give information about the why of the prediction outcome.
- Developers want to know more. The explanations given to the user won't help the developers correct the machine learning model, or understand when the model will misbehave, or what exact data records are causing the model to underperform.

There are lots of unsolved problems like:

- There is no formal definition of interpretability or explanation
- There is no objective measure of how interpretable a model is
- There are no experiments about how much time it takes for a human to understand an explanation

- How to audit explanations in a scalable way?
- How to design UX that facilitates the understanding of an explainable system?

We can ensure the safety by keeping humans responsible. If any AI driven device (robot or other) injures a human, the the owner and producer of that device share the responsibility.

An AI just another human product and humans must remain responsible for what they are doing.

Recent surveys show, that the vast majority of consumers perceives more risks than benefits when companies or government agencies make decisions automatically through algorithms.

Consumers want information on the systems decision making and that the systems are controlled by an independent body.

Making AI trustworthy: Transparency, Audits, Accountability

Trust and acceptance could be promoted by two factors:

A) when automated decisions become transparent and explainable

b) when there is a proper control system in place, that makes sure that the decisions on consumers are lawful and ethically sound.

Therefore it should be considered to establish an independent control system that is able to review and audit socially relevant AI/ADM processes.

This audit should test whether the system conforms with legal laws: anti-discrimination law, unfair competition, data protection and it should analyse individual and social impact of AI.

In order to facilitate an audit in the first place we should consider establishing Standards for transparency-by-design and accountability-by-design.

In order to establish trust, we need to ensure that these Black-Box Systems are independently controlled and audited, so that they adhere to the law.

Value-compliance : The set of principles underlying the functioning of any socially impactful AI/ADS should incorporate and embody fairness and justice.

More generally, we must be mindful about the values reflected in the design and operation of AI/ADS, about what we choose to optimize and automate. We should also acknowledge that apparently neutral data reflect historical or systemic inequality and discrimination.

Opt-out : possibility to be excluded from any data operations and automated decision-making systems and processes, explicitly stated, and implemented as a straightforward,

easy, trouble-free, and inexpensive procedure.

Challenge and reparations : automated decision can be challenged; the operator has the obligation to respond to the concerns of people who feel they have been incorrectly or unfairly assessed, and to rectify any harmful outcomes.

There are also two more complex components of trust, namely transparency and accountability;

With respect to responsible and safe AI, it may also be worthwhile to talk about IT security aspects, namely whether AI aspects can be hacked and misused in novel ways.

A trusted AI system means that its outcomes are at least equivalent to the outcome achieved by a human expert or group of experts doing the same task in similar conditions. AI system can be untrustworthy in case of uncertainties, uncertain measurements, and disturbances. So methods that model and analyze various forms of uncertainty (e.g. reachability analysis) can be used to evaluate the trust of the AI system. Evaluation can be performed against some defined targets. Similar criteria can be used for autonomous systems in all modes of transportation (self-driving cars, automated ships, drones, etc.)

SAFE AI: Safe for who? for humanity? for Europeans? for the environment? for businesses?

"harmless" ? how best to ensure that this rule is embedded into the very core and code of each AI, and how to ensure that no loophole exist that could allow an AI to indirectly harm humanity (e.g. by disrupting the food chain, ...)

"guarantee general humanity safety"? to what extent will it keep us safe and what is the "threat" threshold that govern its decision of whether someone could be considered as "unsafe" to humanity and taken care of / removed ?

A global effort should be made to develop a set of "AI principles" that will act as a core set of values embedded in each and every AI project from the start. One way to make that happen could be to limit the proliferation of AI libraries, and embed in each library these "AI principles" without allowing individual developers to edit them.

Each new library would then have to be audited by a body of trusted experts first and potentially by a narrow "beneficial and trusted" AI with the sole purpose of ensuring that the AI principles are embedded in the library.

- Define "EU trusted AI application"

- Develop and publish a catalogue of criteria for it

- Implement a body that evaluates AI applications under these criteria

- Offer European and other businesses to undergo this evaluation and apply for the label "EU trusted AI application"

- Every effort should be made to eliminate illegal and unfair bias in designing and developing AI decision systems.
- some element of 'bias' is inevitable: algorithms designed for example to detect smugglers could be said to be 'biased' in favour of detecting smuggling. But this is

reasonable.

- Algorithms in general are not unfairly biased. Training is the main problem. The careful use of training data is crucial in ensuring that illegal or unfair bias does not enter the AI system.
- Not easy. More work needs to be done in this area.
- Careful scrutiny of new products and road maps should be the norm.
- promotion of 'AI ethics' modules in training courses for AI engineers.
- *maximum useful transparency in explaining how AI works.*
- *tech sector should accelerate its efforts to provide ways in which AI machines can record and explain their decision making processes.*

two distinct aspect of explainability:

- One is explaining how AI recognises patterns (particularly images).
- The second is about how AI reaches its recommendations/decisions on human situations. Why did it turn down my mortgage application?

- Explaining the role correlation plays in AI decision making and clarifying how this is not the same as causation.
- Explaining that AI decision making is probabilistic, not deterministic.
- Explaining, if possible, the key correlations/criteria used in reaching an AI decision
- Explaining what would need to be different for the AI result to change (eg if your salary was x dollars higher you would have got the loan)
- Explaining the use of filters/criteria (see above) to guard against unacceptable bias.

We need to protect customers: those who pay by money, those who pay by thier personal data or time, as well as those, who have no other choice than deal with AI because of dominance of the provider or mandatory use. All this without restricting reasearch and development.

Regulators can help the commercial sphere to deal with this responsibility by stating standards for training and for testing AI models, where appropriate. The standards can be differentiated by intended usage of AI model: more strict for healthcare, less strict for shops or social media, more strict when it comes to personality traits, less stict when it comes to wheather prediction, same as with other products. Having standards is on one hand an obligation, but on the other hand it is kind of excuse - the standards set baseline, what is considered sufficient care. Without a standard, any care might be not enough. Statistics, ML and AI are powerfull tools and we should take care to not abuse them. It is right to make research compliant with applicable ethical and other standards. The real thread is in inappropriate use or abuse of the results. Humankind has already experience with abuse of statistics of observable human properties to personality, intelligence, somebody's options, somebody's life. We should take care to not repeat such mistakes.

Session 2A: Transparency and Accountability

Transparency and Accountability: **What do the concepts transparency and accountability mean in an AI-context? How can their concrete and practical operationalization be ensured?**

<https://ec.europa.eu/futurium/en/european-ai-alliance/input-requested-first-ai-hleg-workshop-transparency-and-accountability>

Transparency is characterized by the visibility or accessibility of information especially concerning all the processes involved in the implementation of the AI solution.

This means that information about the following should be made publicly available:

- How the data was collected
- How the data is cleaned (the preprocessing step before feeding the data to a machine learning model) and encoded
- What kind of model is used (type, parameters, model architecture, hyperparameters)
- What kind of bias exploration has been done on the data
- What kind of measurements are done to make prevent the deployed AI solution from misbehaving (for example detect model drifting, or change in the distribution of the training data and the production data)
- What measures ensure that the AI solution is not leaking private information from the users, used at the training step (some people have already mentioned differential privacy)

There is no notion of making this explainable to the public. But the point of transparency is that experts can audit the whole process, and in the best case be able to reproduce the complete pipeline.

Accountability refers to the capability of being explained or held answerable.

A necessary requirement to implement audits for accountability is

reproducibility. **Transparency** requires also using transparent, explainable models when possible.

We should keep the amount of regulation under control. Too much regulation might limit further development in the area. Nearly all critical areas are already identified in existing legislation and it is more about its proper application.

Despite we have to prevent (or reasonably reduce) the wrongdoing, we should not make entering the market too difficult by extensive regulation.

There are sensitive areas today, where the requirements are very specific, like medicine, and so should be with AI driven systems. There are areas where we have more freedom, because they are not so critical, like retail sales, where we should keep enough space for creativity.

Review of relevant ADM/AI processes and case-specific measures

1. Inspection and assessment of legal conformity, individual and corporate impact
2. Determining the Relevance of ADM/AI Processes

3. Determining the appropriateness of case-specific measures
 - a. Creating transparency for consumers and the public
 - b. Adaptation of the ADM/AI process: if ADM/AI processes do not comply with legal requirements, these components must be modified or withdrawn from circulation.
 - c. Ban as a last resort: A prohibition or legal prohibition of the use of certain ADM/AI processes can be a justified last resort in certain cases.

General Requirements for ADM/AI Processes

4. Ensure traceability: Traceability-by-design / explainability-by-Design / transparency by-Design
5. Ensuring Falsification
6. Create possibilities for challenging a decision
7. Introducing information rights, labelling and publication obligations
8. Adapting liability: For the reform of the Product Liability Directive, liability for algorithms in the sense of genuine strict liability when used as intended by the consumer, independent of a fault, is an option. It should be sufficient for the provider's liability if an algorithm causes damage when used as intended.

Use of AI risk assessment tools (RAT) by the judiciary in EU Member States

inclusion in the decision-making process: public authorities, civil society organizations, academics, and citizens, in general, should be able to actively participate in the decision-making process about this architecture.

inclusion creates **accountability** since when the meetings are not secret their discussions are open to scrutiny, **transparency** and boosts **trust**, since when you can actively participate, pose questions, express opinions, and receive clarifications, you feel involved and have confidence in the decisions agreed upon.

Since justice is the symbol of fairness, accountability, and transparency, the AI RAT used for its service should be enforcing these principles as well.

Attention needs to be paid: **classification of the risk categories** can influence the decision-making process of the judiciary and lead to biased choices. Furthermore, the potential of an AI RAT to perpetuate **bias** might exist. Quality of the data is of utmost importance (prevent bias). Transparency and fairness

Fundamental rights shall be centrally reflected in the future guidelines of this group.

Ethics in AI need to be formulated based on flexibility and contingency.

any kind of EU-wide ethical guidelines, whilst firmly grounded in the Charter of Human Rights, must be continually re-evaluated against the backdrop of new social and technical developments, and learnings from history.

AI practitioners and regulators should be required to regularly be updated about **new research** into data epistemology, technology and inequality, bias including critical social research into AI and technology.

Considerations around AI and ethics should also firmly be grounded in discussion around **data** and acknowledge that many of the AI advancement happening in the US are fueled by the strong drive to **commercialise data**.

More transparency about how **standards for measuring success in training algorithms** are set, and what kinds of datasets they build on. Transparency on how **AI training data is being gathered, classified, deployed and by whom, to what ends**.

Transparency

Disclosure to the public and regulators: The operator should provide a definition of AI/ADS understandable by the general public.

The operator should disclose their current use of AI/ADS, accompanied by any related self-assessments and outside review processes and their respective results;

Use of an AI/ADS should be clearly indicated on any that communicates the decision to the affected person / institution / group.

AI/ADS description should allow the assessment of: Data (sources, quality, bias, prohibited info - race, ethnicity, sexual preferences, etc., or proxies) ; Correctness / appropriateness of mathematical / statistical approach ; Correct understanding of the subject matter ; Proper usage, i.e., in the contexts and for the purposes the system has been designed for.

Understandable design: The AI/ADS should be specifically designed to be reviewed.

Explainability: The AI/ADS should be accompanied by a natural language explanation of the decision input, process, and output.

Accountability : Responsibility for the consequences of using AI/ADS be assigned to the operator. One key obligation of the operator should be to rectify any harmful outcomes.

- **Human-in-the-loop** : the decision to deploy an AI/ADS, as well as any automated decision of such a system, should link back to a specific human decision-maker.
- **Monitoring and Auditing** : the AI/ADS should be regularly audited; the operator should develop meaningful review processes, subject to independent control, that ensure continuous validation and testing, in order to discover, measure, and track impacts over time.
- **Algorithmic accountability**

There are two approaches possible (at least):

1. explanation individual decision in terms of weights and functions of neural network/SVM. Each result can be tracked back to see, how each of the inputs contributed to the decision. This can be done for each individual decision.

2. description of the training process and the measures that have been taken to prevent biases and ensure fairness.

As with other regulated processes, the users of the system must know, what are their systems doing and to which extent they can rely on them. Each system must be tested and yes, lot of work is ahead of us to develop testing methods for AI. It took some time for legacy systems as well.

Using "black box" excuses is not acceptable, because this can be very easily abused for crime. Would you put your money, health or life into hands of a "black box"?

Proper application of existing legislation, ethical and moral standards on AI:

- Humans must remain responsible for the actions and decisions of AI, and must be able to explain the actions/decisions of AI.
- Compliance of AI with current legislation must be evaluated in broader context.
- Producer's responsibility to ensure that algorithms, training and validation sets are prepared with proper care to ensure compliance with law and regulations, including prevention of any biases or discrimination prohibited by law and to prove this
- No right of AI to defend itself by injuring a human

Labeling AI:

- Clear notification that a product is driven by AI
- Proper explanation of the objectives, reliability, limits, biases, etc. in non-technical language.
- Prevent misleading tools by simulating or copying of appearance or behavior of humans or animals
- Prevent misleading of humans, animals or tools in belief that a human, intentionally or accidentally, being considered AI driven device.
- Proper education of general public

Prevention of abuse where AI creates new areas with potentially huge hard-to-assess damages: AI is entering into communication with humans (and probably other areas) in way unseen before and we should find ways how to regulate these.

Traceability of AI actions:

- Producers and users of AI should be obliged to monitor activities of their system and their compliance with law, regulation and ethical standards
- Saving and preserving of design and making it available in case of investigation
- Saving and keeping logs of decisions and their reasoning
- Ability to explain the outputs, including recommendations, decisions and actions

The requirement of explainability is also very interesting in the medical sector. If the interpretation of this requirement is strict, that excludes the use of some powerful AI methods from medical applications.

My example: We participate in the [Digi-NewB project](#) in which we develop decision support for the early detection of sepsis in preterm infants. If the algorithm we use has to fully explain its decision, that may exclude the use of more accurate algorithms for the sake of explainability.

If anyone knows a position paper type of document or similar from the EU Commission about this, I would be happy to hear about it. The FDA which regulates medical devices in the USA has already recognized the power of AI and [begins to approve AI-based systems as medical devices](#). The manufacturers applying AI in their medical devices need to prove, of course, that their methods work well before they get the FDA approval. If Europe is restricted by a strict interpretation of the requirement of explainability, the European patients do not get as accurate diagnoses as the patients in the USA. Even the uncertainty about the situation is a hindrance to system development.

Explainability of AI: A very important factor in limiting Black Box

AI solutions. Consider the case of Mount Sinai Hospital that employed the AI solution Deep Patient to predict cases of schizophrenia, which is otherwise a notoriously difficult

thing to do for doctors. While Deep Patient could indeed do it more accurately, the problem was that doctors had no clue why/how, and had to blindly trust the AI. The idea here is that if an AI solution makes a prediction or decision, it should also be able to explain the rationale behind it. Sooner or later, we will have ethical laws in place to ensure this. For now, we can at least have a rating of the degree of transparency an AI solution has, based on how well users can see/understand the reasoning behind its predictions.

Journalism has lost enough trust with the general population already, and the answer to that can only be more transparency. And if we expect that transparency of journalists, we should also expect transparency of the algorithms they use. If we ever get to the point where "AI" influences editorial decisions like agenda setting or presentation (headlines, teasing, etc.), we have to be able to explain that decision to the readers - which is why algorithmic accountability is a huge topic.

Transparency

There is a needs for:

- a system to provide rough overview in a **natural language explanation** of the decision making process for any specific instance, of the hypotheses underlying the general decision making process, and these explanations have to be sufficiently concise to be understandable to a member of the general public without AI background.
- host of statistical methods to gather insight into a system's decision making process. These kinds of statistical explanations are likely not helpful to the general public, but are key to enable **expert auditing**.
- transparency on **data** the system has been trained on and which kind of bias may be contained within them. Ideally, the data itself should be available for further scrutiny, albeit in anonymized form such that privacy interests of citizens are taken into account. Further, the training process of any system should be recorded to be available for later auditing.

Accountability:

- **Legal liability** for a system's actions and decisions **should always rest with humans**, and should not be pushed off to the unpredictability of the system itself.
- In line with the GDPR, AI System should provide a rough summary in natural language for the subjects of the decisions themselves, as well as more **detailed information** which can be used by experts **as basis for potential legal measures**.
- **Regular audits** of large-scale AI systems as well as continuous validation and testing measures, subject to independent control. AI developers and/or users could be required to change the system.

Session 2B: Industry & Ecosystems Uptake of AI

AI Infrastructure and Enablers: **How can AI's uptake be accelerated across the entire European industry? Which measures can we take to ensure a level playing field in AI's uptake across all sectors, companies of all sizes, and all regions?**

<https://ec.europa.eu/futurium/en/european-ai-alliance/input-requested-first-ai-hleg-workshop-industry-ecosystems-uptake-ai>

I would like to propose a desirable structural AI Certification-chain whose objective is that to guarantee a basilar security and reliability of the product as this hit the market.

I envision the following hierarchical structure:

- EU Certification body: to charter operational institutes (Data agencies and Application Software).
- Data agencies (private or public), possibly segmented by field of interest (e.g. medicine, home care, school learning, autonomous driving, etc.). To maintain/mine huge quantity of data on a specific domain; certified to train AI applications.
- AI original application software producers (programmers).
- AI educational institutes and schools to educate students and benefit the future of Industry.

In the absence of a similar structure we run the risk to let super-powerful multinationals supplant designated EU Institutions as legitimate authorities, what has already happened with Google-Play Store!

It is crucial for EU to have common strategic plan. EC should encourage governments of EU countries to act in increasing uptake of AI. What we need is a combination of education, research, entrepreneurship and funding to build a sustainable and competitive AI ecosystem. All these ingredients are available in EU, they just need more focus to different aspect of AI.

AI is about autonomous and adaptive software systems capable of learning, reasoning and making decisions by themselves.

If such systems would exist, and may be adopted at a mass level within EU industry space, consider first what could be the possible socio-economical and ethical implications. Once (as a society and decision makers) we have a solid understanding of implications of such mass deployment, we can then discuss about an accelerated uptake of the AI technologies where appropriate.

Create a virtual experience platform from the moment of generating knowledge to using AI's systems. Virtual experience is designed for conduct experiments, analyze systems and applications, to identify flaws, negative consequences, and timely actions to adjust. This will economy material resources, as well as reduce the negative consequences. Thus, we will reduce risks in real situations and show opportunities.

- **Training:** not only from a technical expertise point of view, but also making business leaders understand the capabilities and limitations of AI to correctly manage expectations.
- **Legal security:** one of the main AI adoption stoppers might be lack of clarity regarding legal implications linked to the use of AI.

- **AI democratization:** empower API & cloud based AI to avoid expertise or organizational adoption barriers.

Expertise is key: Make available free digital courses and provide a platform to practically experiment with AI and exchange ideas (fora).

- Make data available across sectors, companies and regions.
- Make NLP and speech engines available for European languages.
- Fund free education programs, and consider offering a financial bonus upon successful completion.

Concentrate economic effort on most advanced AI takers/users (Countries) on condition of exchanging know-how and investment in a less advanced sphere.

EU should set an appropriate educational budget to the benefit of less fortunate member countries.

Session 3A: Use Cases for the Guidelines

Use-Cases for the Guidelines: **Which Use-Cases are interesting/useful to explain, challenge and test the Guidelines?**

<https://ec.europa.eu/futurium/en/european-ai-alliance/input-requested-first-ai-hleg-workshop-use-cases-guidelines>

As soon as AI/ADM processes are based on personal data, privacy issues play an intrinsic role in these processes and their potential impact. Many of the risks in connection with AI/ADM processes mentioned below would increase if the level of data protection were lowered.

It is often postulated that data protection would prevent AI applications. It is overlooked that the GDPR already provides for far-reaching exceptions that also allow AI applications for personal data.

Data protection only plays a role at all when personal data are processed. However, this is not the case with many AI applications.

Recommendation systems and decision assistants: Preserving the decision-making autonomy of individuals

Examples: Smart Digital Assistants for Purchasing, media selection, health consulting, financial consulting

Possible approaches:

Labelling obligation for AI systems that interact with consumers, or make decisions over consumers, establishment of a control system. This could entail transparency towards consumers and experts (audit team), meaningful explanation of decision logic including criteria and database.

Risk scoring and categorization of consumers: Determination of access to services and markets

Examples: Behaviour-based insurance tariffs, credit scoring, automatic job-candidate pre-selection, trustworthiness scores.

Possible Approaches:

Duty to labelling, when decisions on consumers are based on scores. Transparency of the AI/ADM processes for consumers and experts (audit) with respect to decision logic, criteria and database, establishment of an independent control system, liability in case of damage. Possibilities to intervene in the case of unjustified decisions.

Smart Home: Ensuring liability

Examples: Smart entertainment products, home automation, smart meters, ...

Possible Approaches:

Liability in the event of damage caused by ADM and AI, or a low security Device in the

network: Product liability regardless of fault lies with the manufacturer/supplier. Also for damages caused by connected devices. Result: The consumer has a contact person who is liable. The distribution of responsibility between the companies involved (within the internal relationship) is not the responsibility of the consumer (a task he is not up to anyway). Market mechanisms will lead to companies adopting high standards for security, not allowing their devices to connect to low-security devices.

Personalised prices: Power and knowledge asymmetries to the detriment of consumers

Examples: Finely tuned price differentiation in online trading through collected user data. With the amount of information available about individual consumers, companies can manage an increasingly finely structured price differentiation (in real time).

Possible Approaches:

An ethical discussion and decision are needed as to whether personalized prices should be banned in some social and economic sectors/contexts. The exploitation of health emergencies in terms of price is particularly questionable from an ethical point of view. E.g. by making health products (like special mattresses for people with back pain) more expensive for consumers who urgently need them.

Data protection-conscious consumers must not be denied access to goods and services beyond what is legally required if they refuse to consent to data processing.

• Health and Transport.

- Privacy of citizens for an open data platform use case
- autonomous vehicle use case
- **An accountability challenge:** How to govern the automated decision making of virtual agents or AI engines?

- Guidelines should likely provide advice regarding the norms within which systems scoring and representing human beings should operate, both in terms of development and in terms of application.
- Risk scoring (credit defaults, car crashes, or financial risks)
- Many on-line companies, e.g. in advertisement, are interested in representing their users according to certain groups, such as psychological profiles. Grouping profiles can mean that private information is inferred without direct consent, groupings could reinforce stereotypes, and actions based on groupings could be discriminatory
- Search engines: adaptations to existing data can reinforce and even exacerbate existing biases and stereotypes.
- Public institutions may be increasingly inclined to guide decisions by AI systems in order to save resources: e.g. predictive policing. How to avoid over and under policing depending on the population and/or incomes.

Automation of airborne and ground automated systems supporting respectively pilots' and air traffic controllers' activities.

Airborne automation example: TCAS

We analysed the Airbus Auto Pilot/Flight Director (AP/FD) TCAS mode as an onboard example to classify the level of automation. The AP/FD TCAS mode enhances the existing TCAS functionality by implementing a TCAS vertical guidance feature into the

Auto Flight computer. Since this innovation represents an increase in the level of automation, it is interesting to compare it with the traditional TCAS functionality in the light of LOAT.

Ground automation example: ERATO

An example of ground automation is the ERATO Filtering Function. ERATO stands for En Route Air Traffic Organizer and represents a toolset aimed at supporting the controller in ensuring separation among aircraft in a concerned airspace sector. More specifically the ERATO Filtering Function helps the controller in identifying all the potential intruders of a given flight in the medium-short term (i.e. up to 20 minutes to the predicted conflict).

From "AI"-powered Content Management Systems like [Forbes' Bertie](#) to tools that use "AI" to create automated transcripts (e.g. [Trint](#)), I do see a lot of opportunities there, that could make journalists work easier, faster and more efficient.

AI in agriculture: satellites looking over our fields, 3 times per day the milk robots are analysing every little parameter of what the cow is producing and we have thousands of sensors in, on or next to the combines. ... Agriculture 4.0, where every cow and plant gets its "personal" treatment, is coming closer. We already have a number of applications based on AI in agriculture, but there is much more to expect.

Session 3B: AI Enablers & Infrastructure

AI Infrastructure and Enablers: What are the most important building blocks at EU level to enable European businesses to successfully apply AI? How can such building blocks be ensured, and how can investments therein be stimulated?

<https://ec.europa.eu/futurium/en/european-ai-alliance/input-requested-first-ai-hleg-workshop-ai-infrastructure-and-enablers>

Because of cultural, economic and institutional boundaries at national level, Europe has a more fragmented market than for example the US. In the digital sector, where economies of scale are quite important, the companies have a smaller basis to start with. This puts the EU's companies at a disadvantage compared to their competitors. With initiatives, such as "Building a European Data Economy", we can overcome these barriers and create a level playing field. Also for certain types of infrastructure, such cross-border cooperation is essential. For example, the developing of Copernicus, would not make sense for each MS separately.

Basically, those companies which own training data in Europe often lack resources in R&D (mostly it's still R these days) to come up with their own chatbots. Therefore, a harbour for training data would be priority. This could trigger a large set of hackers and enthusiasts which could apply existing research from good papers in Github to novel domains. Those models could be reused by startups or dedicated groups in companies, without being faced by the technical hurdle of training models on their own. Those companies will provide fancy interfaces to their existing customers to enrich their services and keep their customers happy.

The more data you have, the better your product; the better your product, the more data you can collect; the more data you can collect, the more talent you attract; the more talent you attract, the better your product.

Europe has no choice but to rely on its domestic market by offering :

- an innovative compromise that would fight against our propensity to always choose free and ease of use rather than privacy
- another model reconciling both
- to a certain extent to allow users to monetize their data while keeping control of them

European start-ups could therefore rely on a "European architecture" and thus offer services by focusing on the BtoB market of AI, which remains the Achilles heel of digital giants, or even by designing overlayers built on the "digital" back of the GAFA by already focusing on the technologies of the day after tomorrow.

What if, for example, a European startup offered you a system of recommendations that allowed you to increase the value of your data? This is the scenario of Marc Elsberg's book Zero, which raises immense ethical questions of course.

- DataOps tools & methodology
- API-based AI consumption model
- Functional monitoring
- Explainable AI solutions

For machine learning, EU can consider It infrastructure providing: BigData platform, ML toolkit, API management tool, Generic Portal, GUI platform to experiment.

For conversational AI, EU can consider : NLP infrastructure and text2Speech and speech2text in all EU languages and IT infrastructure providing : A platform, that lists the utterances (intents & data entities) used in different languages and use cases, API to interface with Enterprises, Generic Portal, GUI platform to experiment.

should we measure if those systems are behaving as expected optimizing a specific function, we could find that overall those systems are behaving as expected, improving some aspect of our lives, even though some specific individual system is not performing as expected.

If behaviour requirements (responsibilities, liability, etc) are set for each specific instance of a system, in some cases, we could lose the benefit of the overall systems.

New kind of regulatory environment could be beneficial to address this issue; it could be an enabler to accelerate development and application of AI to some areas.