



D6.4 First exploitation strategy

Submission date: 31/03/2025

Due date: 31/03/2025

DOCUMENT SUMMARY INFORMATION

Grant Agreement No	101120763	Acronym	TANGO
Full Title	It takes two to tango: a synergistic approach to human-machine decision-making		
Start Date	01/10/2023	Duration	48 months
Project type	Research and Innovation Action (RIA)	Funding programme	Horizon Europe
Deliverable	D6.4: First exploitation strategy		
Work Package	WP6 – Maximising impact		
Type	R	Dissemination Level	PU
Lead Beneficiary	EITD		
Authors	Annalisa Andaloro EITD		
Co-authors			
Reviewers	Claudia Esposito (CARR), Maria Carolan (CARR)		



**Funded by
the European Union**

DISCLAIMER

Funded by the European Union under Horizon Europe Programme - Grant Agreement 101120763 - TANGO.

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Health and Digital Executive Agency (HaDEA). Neither the European Union nor the granting authority can be held responsible for them.

While the information contained in the documents is believed to be accurate, it is provided “as is” and the authors(s) or any other participant in the TANGO consortium make no warranty of any kind with regard to this material including, but not limited to the implied warranties of merchantability and fitness for a particular purpose. Neither the TANGO Consortium nor any of its members, their officers, employees, or agents shall be responsible or liable in negligence with respect to any inaccuracy or omission herein. Without derogating from the generality of the foregoing neither the TANGO Consortium nor any of its members, their officers, employees, or agents shall be liable for any direct or indirect or consequential loss or damage caused by or arising from any information advice or inaccuracy or omission herein. The reader uses the information at his/her sole risk and liability.

COPYRIGHT

© Copyright in this document remains vested in the contributing project partners.

This document contains original unpublished work except where clearly indicated otherwise.

Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation, or both. Reproduction is authorised, provided the source is acknowledged.

DOCUMENT HISTORY

Version	Date	Changes	Contributor(s)
V0.1	28/02/2025	First release	Annalisa Andalaro
V0.2	14/03/2025	First draft internally reviewed	Claudia Esposito (CARR), Maria Carolan (CARR)
V1	28/03/2025	Final release	Annalisa Andalaro

PROJECT PARTNERS

Partner	Country	Short name
UNIVERSITA DEGLI STUDI DI TRENTO	IT	UNITN
TECHNISCHE UNIVERSITAT DARMSTADT	DE	TUDa
UNIVERSITA DI PISA	IT	UNIFI
CONSIGLIO NAZIONALE DELLE RICERCHE	IT	CNR
SCUOLA NORMALE SUPERIORE	IT	SNS
UNIVERSITE PARIS CITE	FR	UPC
FONDAZIONE BRUNO KESSLER	IT	FBK
CARR COMMUNICATIONS LIMITED	IE	CARR
ISTRAZIVACKO-RAZVOJNI INSTITUT ZA VESTACKU INTELIGENCIJU SRBIJE	RS	IVI
SURGICAL SCIENCE SWEDEN AB	SE	SuS
UNIVERSITATSKLINIKUM HEIDELBERG	DE	UKHD
CENTRE FOR EUROPEAN POLICY STUDIES	BE	CEPS
BCAM - BASQUE CENTER FOR APPLIED MATHEMATICS	ES	BCAM
U-HOPPER SRL	IT	UH
INTESA SANPAOLO SPA	IT	ISP
EIT DIGITAL	BE	EITD
SHARE FONDACIJA	RS	SHARE
A11-INITIATIVE FOR ECONOMIC AND SOCIAL RIGHTS	RS	A11
MINISTARSTVO ZA BRIGU O PORODICI I DEMOGRAFIJU	RS	MFWD
AZIENDA PROVINCIALE PER I SERVIZI SANITARI	IT	APSS
SWANSEA UNIVERSITY	UK	UoS
THE UNIVERSITY OF WARWICK	UK	UoW



**Funded by
the European Union**

LIST OF ACRONYMS

Acronym	Definition
DoA	Description of Action
EC	European Commission
HaDEA	European Health and Digital Executive Agency
WP	Work Package
ER	Exploitable Result
KER	Key Exploitable Result
TRL	Technology Readiness Level

EXECUTIVE SUMMARY

The First Exploitation Strategy of TANGO project is meant to provide both general, high-level guidance to address the exploitation of the whole set of Key Exploitable Results of the project and guarantee on point guidance to technology transfer of single Key Exploitable Results.

D6.4 aims at ensuring the alignment of all exploitation actions to the overarching project aim, identifying specific opportunities for Key Exploitable Results which have a potential for market transfer as well defining an overall strategy to ensure long-term sustainability of results after project conclusion. Both commercial and non-commercial exploitation are addressed in this document, depending on the specific needs and ambitions of each result.

After a short introduction about the purpose of this document, content is organized in two main sections:

- The first illustrates the general approach adopted to structure the Exploitation Strategy
- The second clarifies the details of development and the specific approach adopted for each of the Key Exploitable Results,

The overall and result specific exploitation strategies are subject to updates as soon as project results are further developed and the case studies of the project will provide all validation elements supporting TANGO hybrid human-machine decision-making processes.

The First Exploitation Strategy is delivered at M18 of project duration while all project results are still far from being finalized. For this reason, D6.7 – Second Exploitation Strategy has been planned for M36 when all project results, including those related to case study applications will be close to an ending in terms of expected development within the course of the TANGO project.

D6.4 has been prepared by EITD as the main beneficiary, collecting relevant contributions from all WP Leaders, as well as all consortium members who have expressed their positions and opinions concerning each Key Exploitable Result identified in the Grant Agreement.

TABLE OF CONTENTS

1 Introduction	9
1.1 Purpose of the document	9
1.2 Relation to other project work	9
1.3 Structure of the document	9
1.4 Glossary adopted in this document and external references	10
2 Exploitation Strategy approach	10
2.1 Exploitation Strategy set-up activities	10
2.2 Exploitation strategy roadmap and activities until M18	11
2.2.1 Consolidation of the list of Key Exploitable Results	11
2.2.2 Definition of a preliminary Project Exploitation Strategy	11
2.2.3 Brainstorming on exploitation for each KER	12
2.2.4 Definition of a preliminary exploitation plan for each KER (ongoing until M36)	12
3 Key Exploitable Results and their preliminary exploitation plans	13
3.1 KER 1 and KER 2: Methods and Algorithms human-machine mutual understanding and hybrid Decision Making; TANGO Software ecosystem	13
3.1.1 Original Exploitation Strategy for KER 1 and KER 2 as marked in the GA	13
3.1.2 Status update on exploitation strategy at M18	14
3.2 KER 3: TANGO Personal Assistant for health and wellbeing	15
3.2.1 Original Exploitation Strategy for KER 3 as marked in the GA	15
3.2.2 Status update on exploitation strategy at M18	15
3.3 KER 4: TANGO AI system for medical decision making	16
3.3.1 Original Exploitation Strategy for KER 3 as marked in the GA	16
3.3.2 Status update on exploitation strategy at M18	16
3.4 KER 5: TANGO AI system for credit lending and financial applications	16
3.4.1 Original Exploitation Strategy for KER 3 as marked in the GA	16
3.4.2 Status update on exploitation strategy at M18	17
3.5 KER 6: TANGO AI system for social policy making	17
3.5.1 Original Exploitation Strategy for KER 3 as marked in the GA	17
3.5.2 Status update on exploitation strategy at M18	17
4 Next steps	18
5 Conclusions	19

LIST OF FIGURES

Figure 1: Example Figure 1	10
Figure 2: Example Figure 2	11

LIST OF TABLES

Table 1: Example Table	8
------------------------	---

1 Introduction

1.1 Purpose of the document

This deliverable includes the first release of the Exploitation Strategy for TANGO project and its Key Exploitable Results. The aim of D6.4 First Exploitation Strategy, in alignment with the project's Description of Action (DoA), is to ensure alignment of all exploitation actions to the overarching project aim, identifying specific opportunities for Key Exploitable Results which have a potential for market transfer as well defining an overall strategy to ensure long-term sustainability of results after project conclusion. Whenever possible, the roadmap to market topic has been addressed, although the Technology Readiness Level of project results is oftentimes too low to allow for straightforward market transfer by the time the project ends.

Both commercial and non-commercial exploitation are addressed in this document, depending on the specific needs and ambitions of each result.

D6.4 is a public deliverable so it is also intended to be made available to the general public interested in the TANGO project as well as to other EU funded projects involving technology and market transfer of research results within the domain of hybrid human-machine decision making systems backed by Artificial Intelligence.

1.2 Relation to other project work

Due to the scope of D6.4 First Exploitation Strategy and the objective of tracing a roadmap for possible market transfer of research results, the work in this deliverable is related to all tasks and deliverables developed within WP5 TANGO case studies but also connects to other WPs in the project, in particular WP2 Algorithms for Synergistic Human Machine Learning, WP3 Algorithms for Hybrid Decision Making and WP4 The TANGO software ecosystem on the technical side.

D6.4 is part of WP6 on the Communication, Dissemination and Exploitation and, as such, Exploitation Strategy is designed to make the most of synergies with the Communication and Dissemination activities included in T6.1 and T6.2.

D6.4 is a direct output of T6.3 Exploitation and IP Management which lasts until M48 of the project and supports the development of exploitation activities for Key Exploitable Results, as well as any IP management activity. D6.4 is connected directly to T5.1, T5.2, T5.3 and T5.4 as it will set the approach to ensure long-term sustainability of results developed within TANGO case studies, as well as to WP2, WP3 and WP4 as it will set the approach to exploitation of the whole TANGO API and software ecosystem, which are joint results generated by the above mentioned WPs.

D6.4 will provide valuable input for D6.7 Second Exploitation Strategy.

1.3 Structure of the document

This document is presented in two major sections:

- Section 2
- Section 3

Section 2 is focusing on setting the high-level exploitation approach and strategy for the whole TANGO project in order to ensure its long-term sustainability. It provides an overview of activities planned and describes the process followed to define the list of activities.

Section 3 describes TANGO's Key Exploitable Results technical details and development expectations, setting up a preliminary exploitation plan for each result.

1.4 Glossary adopted in this document and external references

- Exploitable Result: in the context of European Project, a Project Result is labelled as exploitable when it shows potential for commercial development through technology and market transfer activities.
- Key Exploitable Result (KER) is a particularly promising Exploitable Results on which the Exploitation activities of a project are focusing. Such results can be identified already at proposal stage, but can undergo changes and updates during the course of project duration.
- Technology Readiness Level (TRL) is the level of deployment of a specific project result. In the case of TANGO, expected TRL at project end is 5 out of 9, where 9 stands for complete market transfer.

2 Intellectual property management approach

Intellectual property (IP) generated within the course of project activities will be managed according to the general rules set in the project Consortium Agreement and in accordance with the regulations of Horizon Europe.

EITD is leading a specific task T6.3 on Exploitation and IP management, and, as task Leader, is receiving updates on IP developments during the exploitation workshops.

At current status, no major IP generation has been registered, as most of the project results are still under development. However, a shared Excel table is available to all partners to record ownership for the Key Exploitable Results. As part of the exploitation activities, IP matters are regularly discussed.

By the end of the project, consolidated lists of ownership for results generated during project duration will be produced and approved by all partners, to set the ground for fruitful exploitation of such results beyond project end.

Additional updated on IP management and generated IP will be included in the Second Exploitation Strategy, expected for M36.

3 Exploitation Strategy approach

3.1 Exploitation Strategy set-up activities

EITD worked on the Exploitation Strategy for TANGO, setting up a proposal that has been presented officially to the consortium during M12, in a dedicated Exploitation Workshop (held via online meeting).

After collection of feedback from all partners, a consolidated approach was shared during M13 in a live update session, in-person during the second GA meeting of the consortium (Paris, October 2024). During the session, follow-up activities have been defined in order to achieve the definition of an Exploitation Roadmap and subsequent action points for the following months.

The main pillars of Exploitation Strategy definition proposed for the project and stated in the DoA are:

- Exploitation Strategy for TANGO project as a whole, core methods and algorithms
- Exploitation of case-study related results

These two pillars are supported by the creation of synergies with technical and non-technical community of stakeholders (link to T6.2 Clustering, networking and knowledge transfer led by UNITN) which will enable greater awareness about the project topics, closer contact with target groups potentially interested in exploitation and preliminary validation of Key Exploitable Results in terms of perspective adoption, both academic, commercial and non-commercial.

Based on the pillars described above, and taking into account that the expected TRL at project end for TANGO decision making support system components is 5, an exploitation strategy roadmap has been defined and is described in the following section.

3.2 Exploitation strategy roadmap and activities until M18

The Exploitation Strategy Roadmap includes the following activities, whose status at M18 is marked in brackets:

- Consolidation of the list of Key Exploitable Results (completed)
- Definition of a preliminary Project Exploitation Strategy (completed)
 - Definition of owners for each KER (completed)
 - Definition of consortium members with an interest to directly follow exploitation for each KER (completed)
- Brainstorming on exploitation for each KER (ongoing until M36)
- Definition of a preliminary exploitation plan for each KER (ongoing until M36)
- *Consolidation of the complete exploitation plan for each KER and for the project as a whole (not started yet, expected for M36, to be included in D6.7 Second Exploitation Strategy)*

3.2.1 Consolidation of the list of Key Exploitable Results

In order to consolidate the list of Key Exploitable Result, EITD organized an online First Exploitation Workshop to provide a refresher about exploitation obligations under Horizon Europe to all consortium members and providing the list of Key Exploitable Results included in the DoA.

The list of Key Exploitable Results is reported here below

KER 1 Methods and Algorithms human-machine mutual understanding and hybrid decision making

KER 2 TANGO Software Ecosystem

KER 3 TANGO Personal Assistant for health and wellbeing

KER 4 TANGO AI system for medical decision making

KER 5 TANGO AI system for credit lending and financial applications

KER 6 TANGO AI system for social policy making

During the First Exploitation workshop the list of Key Exploitable Results was discussed and the consortium members reached an agreement that no further results would have been added to the list for the moment being. The consortium was also made aware that new Key Exploitable Results can be added to the list at any point in time during project duration, provided the owners give due notice to EITD who is the main beneficiary for T6.3 Exploitation and IP management.

3.2.2 Definition of a preliminary Project Exploitation Strategy

EITD led a discussion on how to approach the preparation of the exploitation strategy for the project with all consortium partners, and agreed on the final version with the project coordinator. The discussion started during the First Exploitation Workshop and was finalized during the Project Plenary Meeting held in Paris in October 2024, where all project partners were attending.

The preliminary Exploitation Strategy is based on two levels.

The first level deals with defining, for each Key Exploitable Result, clear ownership of the result and a mapping of all consortium partners potentially interested in contributing to exploitation activities during project duration and beyond. This mapping activity was carried out through a call to action from EITD as T6.3 leader to all consortium partners, where each was asked to mark on a spreadsheet whether they were owners of a specific result and/or willing to contribute to exploitation activities. The spreadsheet is archived in the project shared repository so partners can access it at all times and check status. It is a living mapping, which can be updated in case new contributions to a result arise, or more results need to be tracked. Mapping ownership and interest to exploit from the early stage of the project is of the utmost importance to avoid conflicts regarding IP management during the course of project, and address any issues in due time so to avoid delays in scientific publications or hindering dissemination and communication opportunities, in case any IP items is addressed for protection.

The second level is described in the following sub-section and concerns a deep-dive on each result, to highlight exploitation potential.

3.2.3 Brainstorming on exploitation for each KER

EITD led the brainstorming on KERs both during a collective workshop where all partners were invited to participate (Second Exploitation Workshop, held in January 2025) and in smaller focus groups where results owners mapped as per 2.2.2 participated to further discuss vertically on their specific result (focus group activities further described in 2.2.4).

The collective session was useful to map general high-level needs from project partners concerning exploitation, provide guidance on exploitation obligations within the framework of Horizon Europe and highlight potential synergies among target groups referring to each specific result. In addition, this session pinpointed the relevance of community creation around the project as an enabler to promote single and combined project results.

During these discussions, cross-contamination of exploitation ideas highlighted the opportunity to transfer hybrid human-machine decision making systems across industries, transferring results of validation in specific case-studies developed within WP5 from the original sector to other, offering increased replicability.

3.2.4 Definition of a preliminary exploitation plan for each KER (ongoing until M36)

EITD organized follow-up sessions to the collective brainstorming developed during the Second Exploitation Workshop adopting the following structure:

- KER 1 and 2 – addressed during the same follow-up session, due to strong correlation among the status of development and potential for exploitation
- KER 3-4-5-6 - addressed during single follow-up sessions, in order to deep-dive on specificities of each result and focus on potential market transfer within the same industry, whenever possible

These sessions have highlighted some cornerstones for the preliminary exploitation plan, such as

- confirming or enlarging the target groups
- revising the originally planned means of exploitation
- double checking feasibility for commercial exploitation.

As the expected TRL of results by project end is 5 out of 9, the meaning for commercial exploitation needs to be put into context, most likely leading to preliminary expressions of interests, further collaborative industrial research and development

Further details on each KER will be provided in section 3.

4 Key Exploitable Results and their preliminary exploitation plans

This section describes each Key Exploitable Result in detail and provides insights on their exploitation potential, considering the status of development until M18. TANGO will deliver an updated version of the Exploitation Strategy in D6.7 Second Exploitation Strategy at M36.

As a general consideration, KER 1 and 2 are described collectively while KER 3, 4, 5 and 6 are addressed separately, as they stem from the four case studies of the project. For these latter, business case exercises will take place in the following months, as soon as the status of result development and evidence collection provide sufficient elements to start with the business model canvas, value proposition and value mapping of the results themselves. Business cases for KERs from 3 to 6 will be included in D6.7 Second Exploitation Strategy at M36.

Key Exploitable results identified and confirmed up to M18 for TANGO are the following:

KER 1 Methods and Algorithms human-machine mutual understanding and hybrid decision making

KER 2 TANGO Software Ecosystem

KER 3 TANGO Personal Assistant for health and wellbeing

KER 4 TANGO AI system for medical decision making

KER 5 TANGO AI system for credit lending and financial applications

KER 6 TANGO AI system for social policy making

4.1 KER 1 and KER 2: Methods and Algorithms human-machine mutual understanding and hybrid Decision Making; TANGO Software ecosystem

KER 1 and 2 are addressed in the same sub-section as all the exploitation effort concerning these two KERs has been running in parallel until M18.

Ownership and interest to exploit for KER 1 and KER 2 have been clarified, and owners on both sides agreed on continuing the joint exploitation strategy set-up effort in order to optimize synergies in the cases of joint exploitation. The expectation is that KER 1 and KER 2 will also have stand-alone exploitation paths, as some components of the software ecosystem have a strong market potential as single products.

4.1.1 Original Exploitation Strategy for KER 1 and KER 2 as marked in the GA

Here below two tables summarize the original expectation set for the exploitation strategy concerning KER 1 and KER 2. It is clearly seen that KER 1 aimed at a scientific and non-commercial exploitation for the immediate future, while KER 2 aimed at a commercial exploitation, although part of KER 2 (TANGO library) is still intended for open source dissemination. Commercial exploitation is meant to focus exclusively on TANGO API and a business case is to be developed within the course of the project by partners who expressed an interest in commercial exploitation of KER 2.

Exploitable Result	Methods and algorithms human-machine mutual understanding and hybrid DM (R3-R4)
Description	The research produced by the consortium partners will lead to valuable IP, open-source code, and several new follow-up research lines in human-centric AI. For this KER, exploitation is mostly in terms of further research.
Target Audience	Cognitive science and digital technology research communities
Means of Exploitation	Approach: Open science - open access for publications describing the methods and algorithms; open source for the implementation of said methods and algorithms.
Exploitable Result	TANGO software ecosystem (R5-R6)
Description	Core TANGO technology (library and API) will be leveraged by the industrial partners to strengthen their competitive positioning and offer new products/services in the HDSS area to the market. For this KER, exploitation is of commercial nature.
Target Audience	Healthcare industry companies building digital personal assistants for health and wellbeing and medical decision support systems; Financial services providers (incl. retail banks)
Means of Exploitation	While the TANGO library will be open sourced, the TANGO API will not. It will be offered as a SaaS to facilitate the development of new products and services in the area of human-AI systems. Market analysis will be carried out and a business case developed.

4.1.2 Status update on exploitation strategy at M18

As far as KER 1 is concerned, non-commercial exploitation options have been listed during the Second Exploitation Workshop and are summarized below

- Publication of a white paper clarifying status of development of the methods and algorithms for hybrid decision-making and describing potential industrial and civil applications, targeting general public and the scientific community at large
- Publication of a position paper on the same topics, targeting policy makers
- Creation of training modules to promote adoption and further development of the methods and algorithms, targeting e.g. PhD candidates, early stage researchers, R&D staff in public organization and corporates or SMEs
- Community building on the topic of interactive explainability and its translation from theory to practice, animated by a set of subsequent and regular recurring events to progress with the methodology, the development of the algorithm and its potential adoption, targeting experts in the field of AI for data analysis and decision making, human cognition experts and others.
- Search for additional stream of funding to allow further research development

An option to single out specific elements of the methods and algorithms was discussed, but no strategy has currently been established to exploit stand-alone elements of the algorithms. Follow-up discussions will further explore this topic in the months to come.

Specific elements to be singled out of KER 1 and KER 2 for potential stand-alone exploitation are well summarized in the project technical report for M18, and a list of scientific contributions connected to their respective peer reviewed papers and source-code is reported below for two different WPs (also available in the technical report).

Table 1 reports contributions to code under WP2, and Table 2 reports contributions to code under WP3.

Table 1: WP2 contributions with links to pdf and code

Paper	pdf	code
[Banerjee et al., 2024] Debodeep Banerjee, Stefano Teso, Burcu Sayin, and Andrea Passerini. Learning To Guide Human Decision Makers With Vision-Language Models. arXiv:2403.16501.	pdf	code
[Barrainkua et al., 2024] Ainhize Barrainkua, Paula Gordaliza, Jose A Lozano, Novi Quadrianto. Uncertainty Matters: Stable Conclusions under Unstable Assessment of Fairness Results. AISTATS, 2024.	pdf	code
[Bortolotti et al., 2024] A neuro-symbolic benchmark suite for concept quality and reasoning shortcuts. NeurIPS Dataset & Benchmarks Track, 2024.	pdf	code
[Busch et al., 2024] Busch, F. P., Kamath, R., Mitchell, R., Stammer, W., Kersting, K., & Mundt, M. (2024). Where is the truth? the risk of getting confounded in a continual world. arXiv preprint arXiv:2402.06434.	pdf	code
[Calanzone et al., 2025] Diego Calanzone, Stefano Teso, Antonio Vergari.	pdf	code

Paper	pdf	code
Logically Consistent Language Models via Neuro-Symbolic Integration. ICLR 2025.		
[Cappuccio et al 2024] Cappuccio E., Kathirgamanathan B., Andrineko G., Andrineko N., Rinzivillo S. <i>Integrating Human Knowledge for Trustworthy and Explainable AI</i> . Submitted to Machine Learning Journal .	–	–
[Delfosse et al., 2024] Quentin Delfosse, Sebastian Sztwiertnia, Mark Rothermel, Wolfgang Stammer, Kristian Kersting. Interpretable concept bottlenecks to align reinforcement learning agents. Accepted at NeurIPS, 2024.	pdf	code
[DTW+25] A. Dix, T. Turchi, B. Wilson, A. Monreale, M. Roach, <i>Talking Back – human input and explanations to interactive AI systems</i> . Proceedings of the 2nd International Workshop on Adaptive XAI @ The 30th Annual ACM Conference on Intelligent User Interfaces, 2025.	pdf	–
[Kehrenberg et al., 2024] Thomas Kehrenberg, Myles Bartlett, Viktoriia Sharmanska, Novi Quadrianto. Addressing Attribute Bias with Adversarial Support-Matching. TMLR, 2024.	pdf	code
[Lenders et al., 2024] Lenders, D., Pugnana, A., Pellungrini, R., Calders, T., Giannotti, F., & Pedreschi, D. (2024). Interpretable Fair Abstention Classifier. Accepted to ECML PKDD 2024	pdf	code
[Longa et al., 2025] Antonio Longa, Steve Azzolin, Gabriele Santin, Giulia Cencetti, Pietro Liò, Bruno Lepri, and Andrea Passerini. Explaining the Explainers in Graph Neural Network: A Comparative Study. ACM Computing Surveys, 57(5), pp. 1-37, 2025.	pdf	–
[Marconato et al., 2024] E. Marconato, S. Bortolotti, E. van Krieken, A. Vergari, A. Passerini, and S. Teso. BEARS Make Neuro-Symbolic Models Aware of their Reasoning Shortcuts. Proceedings of UAI, 2024.	pdf	code
[Marconato et al., 2025] Marconato, E., Lachapelle, S., Weichwald, S., & Gresele, L. <i>All or None: Identifiable Linear Properties of Next-Token Predictors in Language Modeling</i> . In The 28th International Conference on Artificial Intelligence and Statistics.	pdf	–
[Mazzoni et al., 2024] Federico Mazzoni, Riccardo Guidotti, Alessio Malizia. A Frank System for Co-Evolutionary Hybrid Decision-Making. To be presented at IDA 2024.	pdf	code
[Morettin et al., 2025] Paolo Morettin, Andrea Passerini, Roberto Sebastiani. <i>Probabilistic ML Verification via Weighted Model Integration</i> . Manuscript in preparation, 2025.	pdf	–
[Ng and Quadrianto, 2024] Yeat Jeng Ng, Novi Quadrianto. Assessing and Enhancing Robustness of Active Learning Strategies to Spurious Bias. Under Review .	–	–
[Ng et al., 2024] Yeat Jeng Ng, Viktoriia Sharmanska, Thomas Kehrenberg, Anastasia Pentina, Novi Quadrianto. Disagreement-based Active Learning for Robustness Against Subpopulation Shifts. HLDM Workshop at ECML 2024.	pdf	code
[Palomba et al., 2025] Palomba, F., Pugnana, A., Alvarez, J., Ruggieri, S.,	pdf	code

Paper	pdf	code
(2025): <i>A Causal Framework for Evaluating Deferring Systems</i> . Accepted at AISTATS 2025.		
[PPG24] Punzi, C., Pellungrini, R. Giannotti, F (2024). L2IoRe: a method for explaining the reject option. Accepted at DS2024.	pdf	–
[Pugnana et al., 2025] Pugnana, A., Massidda, R., Giannini, F., Barbiero, P., Espinosa Zarlenga, M., Pellungrini, R., Dominici, G., Giannotti, F., Bacciu, D. <i>Deferring Concept Bottleneck Models: Learning to Defer Interventions to Inaccurate Experts</i> . Under Review	–	–
[Punzi et al., 2024] Punzi, C., Pellungrini, R., Setzu, M., Giannotti, F., Pedreschi, D. <i>AI, Meet Human: Learning Paradigms for Hybrid Decision Making Systems</i> . arXiv preprint arXiv:2402.06287.	pdf	–
[Stammer et al., 2024] Wolfgang Stammer, Antonia Wüst, David Steinmann, Kristian Kersting. <i>Neural Concept Binder</i> . NeurIPS, 2024	pdf	code
[Steinmann et al., 2024] David Steinmann, Wolfgang Stammer, Felix Friedrich, Kristian Kersting. Learning to Intervene on Concept Bottlenecks. Proceedings of ICML, 2024.	pdf	code
[Willig et al., 2024] Moritz Willig, Tim Tobiasch, Florian Peter Busch, Jonas Seng, Devendra Singh Dhami, Kristian Kersting. <i>Systems with Switching Causal Relations: A Meta-Causal Perspective</i> . CLR Workshop @ NeurIPS, 2024.	pdf	–
[Wuest et al., 2024b] Antonia Wüst, Tim Tobiasch, Lukas Helff, Devendra Singh Dhami, Constantin A. Rothkopf, Kristian Kersting. <i>Bongard in Wonderland: Visual Puzzles that Still Make AI Go Mad?</i> . Sys2-Reasoning Workshop @ NeurIPS, 2024.	pdf	–
[Wust et al., 2024a] Antonia Wüst, Wolfgang Stammer, Quentin Delfosse, Devendra Singh Dhami, Kristian Kersting. Pix2Code: Learning to Compose Neural Visual Concepts as Programs. Proceedings of UAI, 2024.	pdf	code

Table 2: WP3 contributions with links to pdf and code

Paper	pdf	code
[BNL+24] Marco Bronzini, Carlo Nicolini, Bruno Lepri, Jacopo Staiano, Andrea Passerini. Unveiling LLMs: The Evolution of Latent Representations in a Dynamic Knowledge Graph. Proceedings of COLM, 2024.	pdf	code
[BP25] Cesare Barbera, Andrea Passerini. <i>Interpretable and Efficient Counterfactual Generation for Real-Time User Interaction</i> . Under Review at TMLR .	–	–
[BRR+2025] Raul Barbosa, Salvatore Rinzivillo, Jacques Robin, Andrea Beretta and Henrique Madeira. <i>A Software Architecture for Verifiable and Explainable Classification</i> . Under review at Machine Learning Journal	–	–

Paper	pdf	code
[CFL+2024] Cappuccio, E., Fadda, D., Lanzilotti, R., and Rinzivillo, S. (2024). <i>Fiper: a visual-based explanation combining rules and feature importance</i> . CoRR,abs/2404.16903.	pdf	
[CFL+2025] Eleonora Cappuccio, Daniele Fadda, Rosa Lanzilotti, Salvatore Rinzivillo. <i>TRACE - A Tool for the Visualization of Rule-Based Explanation</i> . Under Review at IEEE Transactions on Visualization and Computer Graphics .	–	–
[CHN+24] Luca Corbucci, Mikko A. Heikkilä, David Solans Noguero, Anna Monreale, Nicolas Kourtellis: <i>PUFFLE: Balancing Privacy, Utility, and Fairness in Federated Learning</i> . ECAI 2024: 1639-1647.	pdf	code
[CMP24] Luca Corbucci, Anna Monreale, Roberto Pellungrini: <i>Enhancing Privacy and Utility in Federated Learning: A Hybrid P2P and Server-Based Approach with Differential Privacy Protection</i> . SECUREPT 2024: 592-602	pdf	code
[DTL+25] De Toni, Giovanni, Stefano Teso, Bruno Lepri, Andrea Passerini. <i>Time Can Invalidate Algorithmic Recourse</i> . Under submission (2025).	pdf	–
[DVT+24] Giovanni De Toni, Paolo Viappiani, Stefano Teso, Bruno Lepri, Andrea Passerini. <i>Personalized Algorithmic Recourse with Preference Elicitation</i> . Transactions on Machine Learning Research, 2024.	pdf	code
[DWR+24] Alan Dix, Ben Wilson, Matt Roach, Tommaso Turchi, and Alessio Malizia. <i>Epistemic Interaction – tuning interfaces to provide information for AI support</i> . SYNERGY Workshop on Designing and Building Hybrid Human–AI Systems, 2024	pdf	–
[EDL+24] Seyedehdelaram Esfahani, Giovanni De Toni, Bruno Lepri, Andrea Passerini, Katya Tentori, Massimo Zancanaro. <i>Exploiting Preference Elicitation in Interactive and User-centric Algorithmic Recourse: An Initial Exploration</i> . Proceedings of UMAP, 2024.	pdf	–
[GMB+2025] Claudio Giovannoni, Carlo Metta, Andrea Berti, Sara Colantonio, Anna Monreale, Francesca Pratesi, Salvatore Rinzivillo: <i>Integrating Multimodal Learning and Explainable AI for Enhanced and Interpretable Prostate Lesion Classification</i> . Under Review at Machine Learning Journal .	–	code
[HNS+2024] Rami Haffar, Francesca Naretto, David Sánchez, Anna Monreale, Josep Domingo-Ferrer: <i>GLOR-FLEX: Local to Global Rule-Based EXplanations for Federated Learning</i> . FUZZ 2024: 1-9	pdf	code
[LFL24] Veronica Lachi, Francesco Ferrini, Antonio Longa, Bruno Lepri, Andrea Passerini. <i>A Simple and Expressive Graph Neural Network Based Method for Structural Link Representation</i> . In Proceedings of the GRaM Workshop @ ICML 2024.	pdf	code
[MCR2025] Carlo Metta, Eleonora Cappuccio, Salvatore Rinzivillo. <i>Interactive Visual Exploration of Latent Spaces for Explainable AI: Bridging Concepts and Features</i> . To be presented at the 2nd Adaptive XAI (AXAI) Workshop , IUI 2025, Cagliari, Italy. 2025.	–	code
[NFN+2025] Gergely Dániel Németh, Eros Fanì, Yeat Jeng Ng, Barbara	–	–

Paper	pdf	code
Caputo, Miguel Angel Lozano, Nuria M Oliver, Novi Quadrianto: <i>Attention to Diversity Is All You Need: Addressing Spurious Correlations and Data Heterogeneity in Federated Learning</i> . Under Review .		
[NMG2025] Francesca Naretto, Anna Monreale, Fosca Giannotti: <i>Evaluating the Privacy Exposure of Interpretable Global and Local Explainers</i> . Trans. Data Priv. 18(2): 67-93 (2025)	pdf	code
[NRS+2025] Ben Wilson, Chiara Natali, Matt Roach, Darren Scott, Alma Rahat, David Rawlinson, Federico Cabitza. <i>Dimensions of Human-Machine Combination: Prompting the Development of Deployable Intelligent Decision Systems for Situated Clinical Contexts</i> . Accepted and to be appeared at eCSCW 2025 .	–	–
[PGM+25] Andrea Passerini, Aryo Gema, Pasquale Minervini, Burcu Sayin, Katya Tentori, <i>Fostering effective hybrid human-LLM reasoning and decision making</i> . Front. Artif. Intell., research topic on “Hybrid Human Artificial Intelligence: Augmenting Human Intelligence with AI”, 2025.	pdf	–
[RMM+24] Russo, F.M., Metta, C., Monreale, A., Rinzivillo, S., Pinelli, F. (2025). Explainable AI in Time-Sensitive Scenarios: Prefetched Offline Explanation Model. In: Pedreschi, D., Monreale, A., Guidotti, R., Pellungrini, R., Naretto, F. (eds) Discovery Science. DS 2024. Lecture Notes in Computer Science, vol 15244. Springer, Cham. https://doi.org/10.1007/978-3-031-78980-9_11 .	pdf	code
[SGM+2024] Francesco Spinnato, Riccardo Guidotti, Anna Monreale, Mirco Nanni. <i>Fast, Interpretable, and Deterministic Time Series Classification With a Bag-of-Receptive-Fields</i> . IEEE Access 12: 137893-137912 (2024)	pdf	code
[SKL+2025] Javier Sanguino Baustiste, Thomas Kehrenberg, Jose A. Lozano, Novi Quadrianto. <i>Performative Prediction Made Practical</i> . Under review at ICML 2025 .	–	–
[SMS+24] B. Sayin, P. Minervini, J. Staiano and A. Passerini. <i>Can LLMs Correct Physicians, Yet? Investigating Effective Interaction Methods in the Medical Domain</i> . Proceedings of the 6th Clinical Natural Language Processing Workshop @NAACL 2024.	pdf	code
[SSS+25] B. Sayin, I.B. Schlicht, J. Staiano, P. Minervini, S. Allievi and A., Passerini. <i>MedSyn: Enhancing Diagnostics with Human-AI Collaboration</i> . Under Review at HHAI 2025 .	–	–
[TMT25] Turchi, T., Mikhaylova, D., Troccoli, M., Malizia, A., Cimino, M. G. C. A., Galatolo, F., La Mantia, G., Campisi, G., & Di Fede, O.. <i>The Reality Check: Ecological Validity Challenges in AI-Assisted Oral Cancer Diagnosis</i> . Submitted to HHAI 2025: Hybrid Human AI Systems for the Social Good .	–	–
[TPD+24] Samuele Tonati, Marzio Di Vece, Roberto Pellungrini, Fosca Giannotti. <i>Ensemble Counterfactual Explanations for Churn Analysis</i> . DS 2024: 332-347	pdf	code
[VFC-2025] Matyas Vincze, Laura Ferrarotti, Leonardo Lucio Custode, Bruno Lepri, Giovanni Iacca, <i>SMOSE: Sparse Mixture of Shallow Experts for Interpretable Reinforcement Learning in Continuous Control Tasks</i> . Accepted at the 39th annual AAAI conference on artificial intelligence, 2025	pdf	–

Paper	pdf	code
[WSR24] Ben Wilson, Darren Scott, Matt Roach, Emily Nielsen, Berndt Müller. <i>Designing for Situated AI-Human Decision Making: Lessons Learned from a Primary Care Deployment</i> . SYNERGY Workshop on Designing and Building Hybrid Human–AI Systems, 2024	pdf	–
[ZLA24] Olga Zaghen, Antonio Longa, Steve Azzolin, Lev Telyatnikov, Andrea Passerini, Pietro Liò. <i>Sheaf Diffusion Goes Nonlinear: Enhancing GNNs with Adaptive Sheaf Laplacians</i> . In Proceedings of the GRaM Workshop @ ICML 2024.	pdf	–
[ZPR2024] M. K. Zuziak, R. Pellungrini and S. Rinzivillo. <i>One-Shot Clustering for Federated Learning</i> . 2024 IEEE International Conference on Big Data (BigData), Washington, DC, USA, 2024, pp. 8108-8117, doi: 10.1109/BigData62323.2024.10825763	pdf	code

As far as KER 2 is concerned, commercial exploitation still looks like the most viable option, as anticipated during the proposal preparation phase. However, the TANGO Library which is part of KER 2 will be made available through Open Source licensing (e.g. Apache 2.0) to allow reuse in both research and commercial settings.

As exploitation for KER 2 will be a combination of commercial and non-commercial, further discussion will be taking place among owners and partners interested in exploitation follow-ups to further progress on the following points:

- validate target groups identified during the proposal phase (e.g. healthcare industry companies and financial service providers)
- evaluate whether additional target groups will be added to maximize adoption of the TANGO software ecosystem within the industry
- deep-dive on specific target groups developing vertical business cases taking into account specific contexts and features of the industrial sectors potentially involved, this point being linked to the four KERs deriving from the case study applications

On a side note, a preliminary discussion on the healthcare use case (surgical simulator) was initiated during the Second Exploitation Workshop, although no major decision has been taken so far, as the results are still at a very early development stage.

4.2 KER 3: TANGO Personal Assistant for health and wellbeing

4.2.1 Original Exploitation Strategy for KER 3 as marked in the GA

Here below a table summarizes the original expectation set for the exploitation strategy concerning KER 3. For this result, the exploitation path is open to both developing further research to replicate the adoption in additional sectors and scenarios, as well as to commercial exploitation through licensing mechanisms.

Exploitable Result	TANGO personal assistant for health and wellbeing (R7)
Description	The personal assistant developed to support women in the perinatal period will help them monitor and improve their mental and health state, identifying subjects at risk. Feasibility analysis will show applicability to real-life scenarios and potential applicability to other scenarios of hybrid DM for health and wellbeing. Furthermore adoption can be obtained in near medical fields such as neonatology and paediatrics.
Target Audience	Research institutes and healthcare providers, insurances and public institutes caring for vulnerable groups
Means of Exploitation	Potential licensing for the development of digital health solutions with self-assessment and screening functionalities. Market analysis will be carried out and a business case developed.

4.2.2 Status update on exploitation strategy at M18

KER 3 exploitation strategy is based on engagement actions with the relevant target group community to raise awareness about the potential of the application and maximize replication in other hospitals. Community workshops will be organized as soon as the result of the case study will be available, involving medical staff from an extended number of hospitals, to present the latest advancements and explore early adoption possibilities within the community.

KER 3 is a software delivered to the general public through a mobile application, which was partially already available to the public before the start of the TANGO project. However, within TANGO an extensive set of features is being developed and will become integral part of the result.

On a side note, although not initially foreseen in the development plan for TANGO, the owners are considering some additional features to improve the exploitation potential of the result, such as

- integrating a chatbot based on LLMs in the application, to improve user experience
- making more languages available, to improve usability across geographies and social categories

4.3 KER 4: TANGO AI system for medical decision making

4.3.1 Original Exploitation Strategy for KER 3 as marked in the GA

Here below a table summarizes the original expectation set for the exploitation strategy concerning KER 4. For this result, the exploitation path is commercial with a focus on medical school and healthcare providers. The simulator can be licensed to interested parties based on their interest and capability to implement it.

Exploitable Result	TANGO AI system for medical decision making (R7)
Description	The HDSS developed to support surgical teams in intraoperative DM will support the training of surgical teams to improve coordination and prevent unfavourable outcomes. Empirical activities on the augmented-reality equipped operating room at UNITN will show applicability to real-life scenarios. Potential applicability to other scenarios of team DM for health and industry will be investigated.
Target Audience	Medical schools, healthcare providers.
Means of Exploitation	Licensing for the usage of the TANGO-powered HDSS-enabled augmented reality simulator. Market analysis will be carried out and a business case developed.

4.3.2 Status update on exploitation strategy at M18

KER 4 exploitation strategy is grounded on the licensing of the simulator product to other interested parties, which is the result element more prone to direct market transfer. Further discussion will be needed to verify

whether owners of case study results have an interest to develop a commercial product to be licensed, working on market analysis and building a proper business case for the product.

To date, preliminary discussions with other clinical hospital universities have been initiated to understand whether there is potential for replication outside of the original application field, leveraging the optimized and verified hybrid decision making system. Owners of the result will be responsible for carrying out such discussion protecting the sensitivity and confidentiality of information shared, and eventually an NDA form would be used in case relevant technical information needs being shared outside of the consortium to support potential exploitation of KER 4.

4.4 KER 5: TANGO AI system for credit lending and financial applications

4.4.1 Original Exploitation Strategy for KER 3 as marked in the GA

Here below a table summarizes the original expectation set for the exploitation strategy concerning KER 5. For this result, the exploitation path is clearly of commercial nature, although there is also room for further research activities on algorithmic recourse.

Exploitable Result	TANGO AI system for credit lending and financial applications (R6-R7)
Description	The HDSS developed for the credit lending case study will be used to improve transparency, fairness and efficiency in the credit lending process within ISP. The results of its use will potentially boost its adoption by other financial institutions in Europe and worldwide, as well as its adaptation to other financial DM scenarios.
Target Audience	Financial service providers, including retail banks
Means of Exploitation	Initially, this will be exploited by ISP for innovating their own processes. Extension to other institutions could be based on a licensing scheme. Market analysis will be carried out and a business case developed.

4.4.2 Status update on exploitation strategy at M18

KER 5 combines algorithmic recourse, fairness and explainability in a single decision making support system for the credit lending industry. At current status, exploitation strategy for the result is grounded on the following elements:

- no confirmed interest to commercially exploit the result within the timeframe of TANGO project
- non-commercial research oriented exploitation seems to be the focus for result's owners
- further investigation will explore potential replication in other sectors than credit lending, such as human resources management, which is quite promising.

The lack of interest for commercial exploitation is also triggered by the lack of actual results given the preliminary stage of development of the system. For this reason, commercial exploitation will be further investigated at a later point in time during project development.

The exploitation potential outside of the credit lending sector will be supported by community workshop to validate interest and collect feedback from the relevant industries.

4.5 KER 6: TANGO AI system for social policy making

4.5.1 Original Exploitation Strategy for KER 3 as marked in the GA

Here below a table summarizes the original expectation set for the exploitation strategy concerning KER 5. For this result, the exploitation path foreseen is non-commercial but still aiming at maximizing the replication potential across several departments of the government of Serbia, as well as across geographies.

Exploitable Result	TANGO AI system for social policy making (R6-R7)
Description	The HDSS developed for the social policy making case study will be adopted by the Government of Serbia for improving the transparency, fairness, and effectiveness of its welfare and targeted social policies. The results will potentially boost its adoption by other national and local administrations (for example, several local administrations in Italy could be interested in adopting the developed framework).
Target Audience	Governments, public administrations at the local level, EU institutions, NGOs.
Means of Exploitation	Initially, this will be exploited by the Government of Serbia at various levels and across its departments. The solution will be made available to other governments at no cost to facilitate its reuse across diverse scenarios and contexts.

4.5.2 Status update on exploitation strategy at M18

Exploitation strategy for KER 6 is grounded on seamless integration of the TANGO decision making support system within the platform that the Serbian government is currently using to implement welfare policy decisions. The first deployment of TANGO will focus on maternity leave policy support measures for female researchers. Upon successful completion of this case, further cases will be analyzed.

The expectation is that decision makers at the policy level will be able to constantly evaluate the impact of their policy actions, at first through the E-science portal, which is hosting data related to research staff employed in Serbia, then in other contexts within the same country. Integration of the TANGO algorithms and decision support system will be performed in line with applicable Serbian laws, also accounting for European regulations with regards to personal data processing and protection (GDPR).

Community workshops are the foundation of the long-term exploitation strategy as, they will enable

- Collection of geographic specific needs with regards to implementation
- Discussion on personal data access issues across geographies, being data the enabler for the tool
- Awareness raising on the result to foster adoption in different countries

Given the current political situation in Serbia, the development of these results are expected to be impacted by some delay, as of February 2025. However, no major impact on exploitation strategy are foreseen for the moment being.

5 Next steps

As the First Exploitation Strategy for TANGO is released at M18, T6.3 Exploitation and IP management will carry on its activities until the end of the project at M48, with an additional deliverable D6.7 Second Exploitation Strategy foreseen at M36.

To meet project's exploitation objectives, the high-level next steps are envisaged as follows:

- EITD will keep on leading and animating the discussion on identified KERs, providing IP support as needed
- KERs owners and partners interested in exploitation will regularly update on any new emerging issues for existing KERs
- New KERs will be collected, in case they emerge during the course of project, and developed following the same exploitation strategy approach applied for the existing ones
- All consortium members will support exploitation linking project KERs to other projects, leveraging dissemination opportunities and their professional networks, tracking their contributions for reporting purposes
- Community workshops will support exploitation activities for all KERs
- Business cases will be developed for KERs 3 to 6
- EITD will release the final and consolidate version of the Exploitation Strategy at M36
- EITD will follow-up on IP management issues and support exploitation activities until project end

6 Conclusions

TANGO project exploitation strategy is preliminary defined in this document, released at M18. The following release of the exploitation strategy for TANGO is foreseen at M36, D6.7 Second Exploitation Strategy.

TANGO project will develop six Key Exploitable Results, each of them following a custom exploitation strategy.

The exploitation strategy described in this report is the result of collective brainstorming, focus groups and interviews with project beneficiaries who are either owner of the specific result, or expressed an interest to support exploitation activities.

KER 1 and 2 are high-level methodological and software tools, not linked to a specific application. For these two results, non-commercial and commercial exploitation paths will be explored.

KER 3 to 6 are case-study specific results, following development of the four case studies of the project. For each of these results, a business case will be created in the months to come and also non-commercial exploitation opportunities will be evaluated.

In the next months, further effort will be dedicated to more detailed elaboration of the overall and specific exploitation strategy, with specific action points assigned both to the task leader EITD (coordination and general IP support) and to project partners (single KER exploitation strategy follow-up actions, such as community workshops, clustering activities with other projects, elaboration of business cases).



**Funded by
the European Union**