ASSURING AUTONOMY
INTERNATIONAL PROGRAMME

DEMONSTRATOR PROJECT
FINAL REPORT

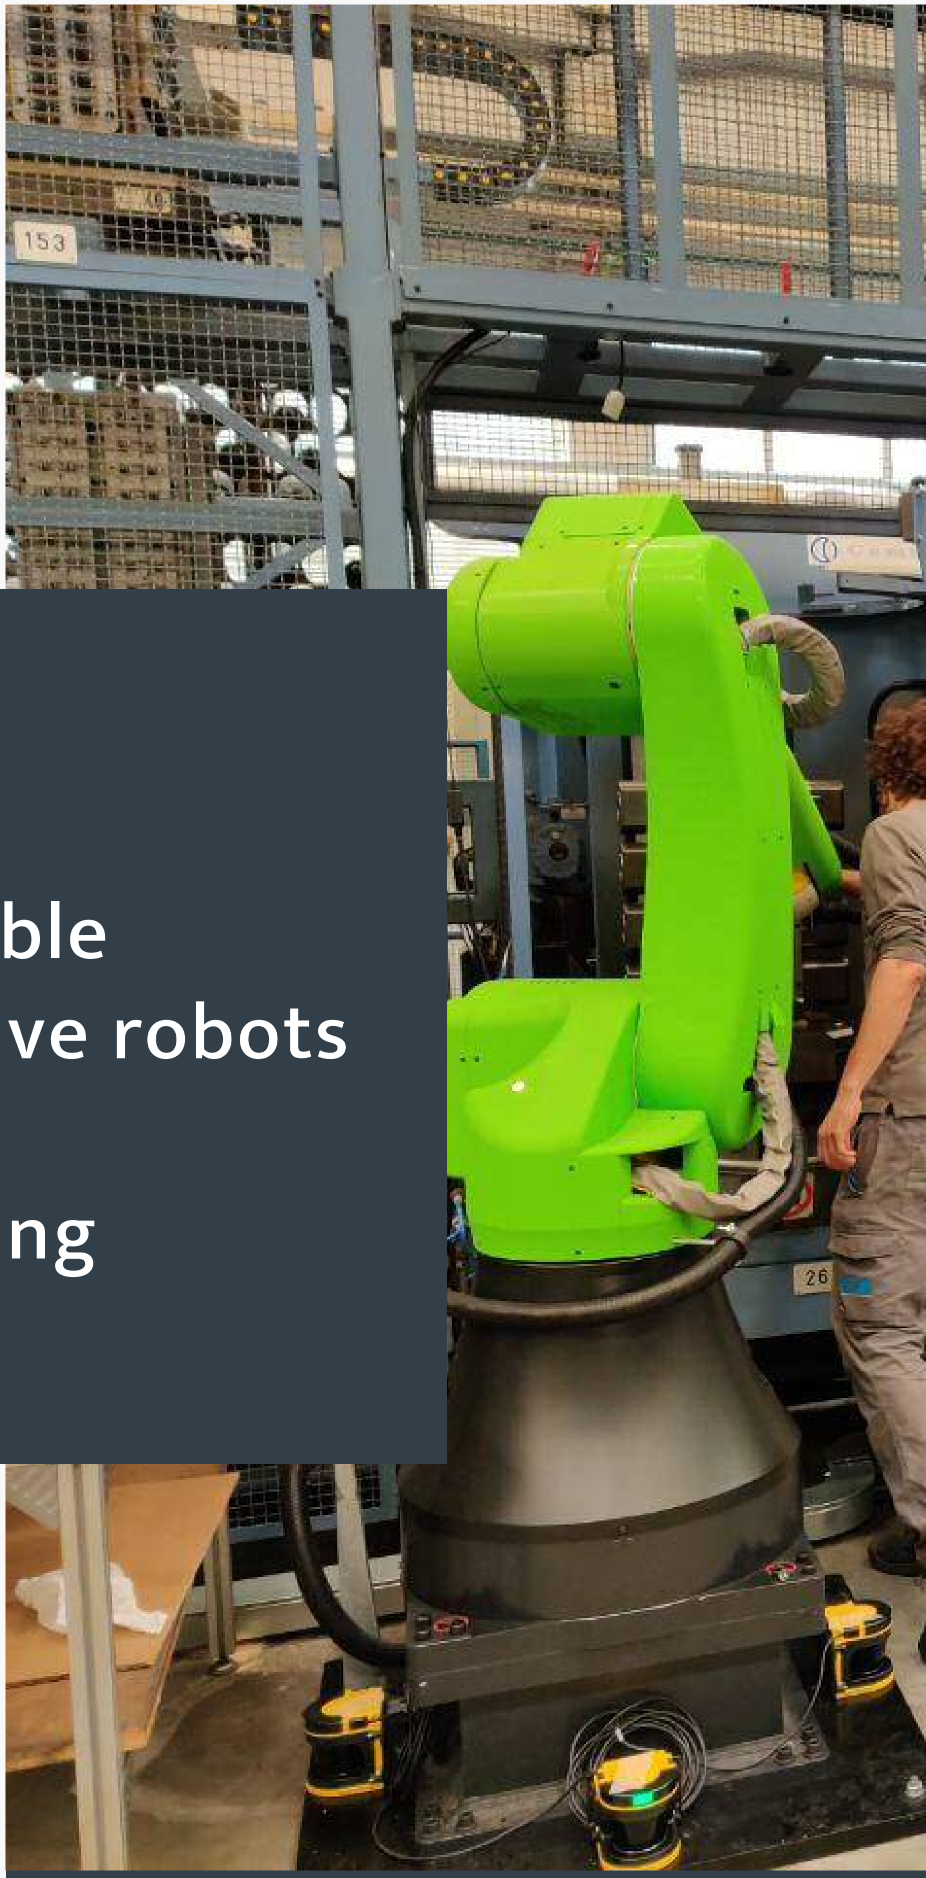

RECOLL
Safety of REconfigurable COLLaborative robots for flexible manufacturing systems

APRIL 2020

Assuring Autonomy International Programme

RECOLL Final Report

2) OUTCOME SUMMARY

Flexible machining systems are key production assets in situations requiring reconfiguration under flexible production pressure and failures. The human-centred manufacturing paradigm is transforming the role of operators by increasing the number of supervisory added-value tasks on larger portions of systems with the assistance of collaborative robots. Dynamical task allocation and layout configuration require an on-the-fly safety assessment of the resulting application. Collaborative robotics can transform the role of operators of flexible machining systems, allowing them to take on more value-added tasks, leaving to robots repetitive or assistive tasks within the same shared workspace. At the same time, operator safety during hybrid tasks dealing with dynamic allocation or layout changes has to be assured.

The project has been dedicated to trace and study the safety-related human-robot behaviour (movements, layout occupation, voluntary/accidental contacts, near misses, etc.) in a prototype machining production setup. The novelty of the methodology investigated consisted of the usage of the tracing interacting human behaviour in dynamic collaboration with robots and production systems with the aim to assess the standard industrial practice of the risk estimation data. The project has been therefore dedicated to combine such findings into the robot-machines co-planning in order to deliver production schedules that optimize the expected human-robot interaction conditions (safety-oriented task allocation). The quantitative evaluation of safety risk assessment according to current standards has been investigated, and, in particular, possible differences within the number of hazardous situations occurring during cooperative work between robot and humans have been analysed.

The RECOLL demonstrator is a collaborative Load/Unload station of a Flexible Manufacturing System (FMS), equipped with two different camera sensors to allow a complete human motion tracking.

Figure 1: Front view of setup. Load/unload station used for experiments is the left one, the closed station still dedicated to manual pallet assembly. Pallets and fixtures are visible inside the automatic storage around the two load/unload stations. (Courtesy of CEMBRE S.p.A., Italy)

3) INDUSTRIAL PRACTICE AND STATE OF THE ART

Modelling the interaction between human and robot is the base for the deployment of risk analysis and assessment, which are mandatory for the real installation of an industrial robotized application. The modelling consists in the definition of all the possible cooperation behaviour of the human and how to configure the correct robot behaviour according to the functionality of the task. Such interaction model has been used for the selection of the safety sensors and for the configuration of the safety options of the robot controller. An example of human-robot interaction modelling is reported in the table of *Figure 2*. Currently, the preliminary safety risk assessment of a collaborative robotic cell requires a model of the expected tasks to be executed and of the volumes occupied by both the robot and the human. This model is generated *a priori*, but an analysis during the actual task execution is useful to evaluate its accuracy and to understand if all the possible risks and hazardous situations have been considered.

situation	region	Path/locations	Intended use	Foreseeable error/misuse	analysis
...					
R loading the part, while H is in manual station.	c/d (R) – s (H)	2-3	H access from 2-3-4 for fixing a part, while R is leaving 3-5, past e.	H anticipates access to the pallet while R is still departing	Regular, periodic sequencing along the path 1-2-3-4, identify alterations in time periods.
H leaves the station while R is completing load/unload cycle	p/e/q	1-3-8(7)	• H go to 1, retrieve information about the status of the working cycle • H goes to 3 for directly checking the status of the machine.		How often H changes the approach pattern, and principal directions w.r.t. the current state of R
Competing access to 4, as long as both R and H need to load and operate on the pallet	a/b	3-4	R in force limitation, mostly holding or moving slowly, H with limited available space	H in unexpected positions for accessing 4, wrong gestures, procedures not followed	How long the synchronization takes, how stable is the cycle period, how much the rate of near misses changes

Figure 2: Example of task model for the risk assessment: all the actions, occupancy areas, paths and misuses are represented

For this purpose, human and robot activities should be monitored, during the whole operational cycles, both spatially and temporally, this was done [1] in terms of robot task recognition. However, data on industrial robot-related fatalities indicate that even in traditional applications of industrial robots, safety is not a solved problem – especially for the variety of operational phases in which the human operator is, by necessity, physically close to a mechanical arm or vehicle. Studies [2] show that many robot accidents do not occur under normal operating conditions, but during programming, program touch-up, refinement, maintenance, repair, testing, setup or adjustment. During many of these operations the operator, programmer or corrective maintenance worker may temporarily be within the robot's working envelope, where unintended operation could result in injuries. One important aspect of designing robots for safety is the ability to quantitatively assess the risk of injuries in accidents, so as to be able to compare different design, control solutions and to optimize them.

The embedding of the interaction model within the safety rules that run on the robot controller is described by the standard in *4* different interaction modalities:

- *Safety-rated monitored stop*: When the robot system is in the collaborative workspace, the safety-rated monitored function is active and the robot motion is stopped, the operator is permitted to enter the collaborative workspace.
- *Hand guiding*: Operator uses a hand-operated device to transmit motion commands to the robot system. The robot shall utilize a safety-rated monitored speed function and stop function.

- *Speed and separation monitoring*: During robot motion, the robot system never gets closer to the operator than the protective separation distance (otherwise the robot stops). When the operator moves away from the robot system, the robot itself can resume motion automatically. Speed and separation monitoring shall apply with all persons within the collaborative workspace. Possibilities for robot control system can be: a) speed reduction with safety-rated monitoring stop or b) execution of alternative path.
- *Power and force limiting*: Physical contact between robot system (including the work piece) and an operator can occur (intentionally or unintentionally). Contact events between the collaborative robot and body party of the operator could come about a number of ways (e.g. intended or incidental contact; failure modes). Types of contact between moving parts of the robot system and areas on a person's body are categorized: quasi-static contact or transient contact.

Two important aspects in safety of robot manipulators have to be taken into account: (i) unexpected collisions between robot and human operator, (ii) physical interactions which are expected according to the task model. In the first case, the human-robot interaction may involve any part of the manipulator structure and operator body; furthermore, it may occur at any time during the execution of the planned trajectory. In literature [3] [4] several indices of impact severity are proposed, which can be mapped (through extensive experimental campaigns and statistical correlations) to the probability of causing a certain level of injury. Some examples are: the *Gadd's severity index* (GSI), the *3ms criterion*, the *viscous injury response* (VC) and the *thoracic trauma index* (TTI). The most widely used index in the automotive industry is presently the so-called *head injury criterion* (HIC). These metrics have constituted a useful basis to begin the development and evaluation of safe robotic concepts. The conditions under which these indices were formulated are quite different from those actually encountered in robotics. Experimental campaigns [5] measuring the effects of impacts of a robot arm using standard crash-test facilities have indicated that classical severity indices used in the automobile industry cannot be transferred without corrections to robotics field. A different set of safety methods and strategies is needed in cases where successful task completion requires people and robots to collaborate intimately, for example in intelligent assist devices (IADs), human extenders and collaborative robots (cobots). Collaborative robot purpose is to relieve humans from fatigue, stress and injuries in manipulating heavy and/or awkward parts. Cobots presume a division of control between human and robot, with a robot perhaps supporting a payload and allowing a human to guide it, subject to constraint surfaces or virtual walls.

[1] D. J. Rude, S. Adams, and P. A. Beling, "Task recognition from joint tracking data in an operational manufacturing cell," Journal of Intelligent Manufacturing, vol. 29, no. 6, pp. 1203–1217, 2018.

[2] A. Bicchi, M. A. Peshkin, and J. E. Colgate, "Safety for physical human– robot interaction," *Springer handbook of robotics*, pp. 1335–1348, 2008.

[3] C. W. Gadd, "Use of a weighted-impulse criterion for estimating injury hazard," SAE Technical Paper, Tech. Rep., 1966.

[4] J. A. Newman and N. Shewchenko, "A proposed new biomechanical head injury assessment function-the maximum power index," SAE Technical Paper, Tech. Rep., 2000.

[5] S. Haddadin, A. Albu-Schaffer, and G. Hirzinger, "Safety evaluation of¨ physical human-robot interaction via crash-testing." in *Robotics: Science and Systems*, vol. 3, 2007, pp. 217–224.

4) METHODOLOGY

In RECOLL, the production follows different rules with respect to a standard manufacturing plant: the sequence of pieces to be worked is not known in advance and is determined by the main supervisor, called jFMX. Plants are often customized for different needs, and the variability of piece types, fixtures and pallets is really high. In this scenario, human presence is particularly useful with respect to standard robotized process, since the high variability reduces dramatically the possibility fully automatized processes.

In *Figure 3* below, the layout of the fenceless cell is displayed. To have a fully working system, the design of the safety rules embedded in the robot controller must therefore guarantee that the human operator may access freely the collaborative cell all the time, also when his presence is not strictly necessary, as the task schedule. Therefore, the main objective of RECOLL has been to define a protocol that provides quantitative information on all the aforementioned aspects, monitoring and analysing human behaviour.

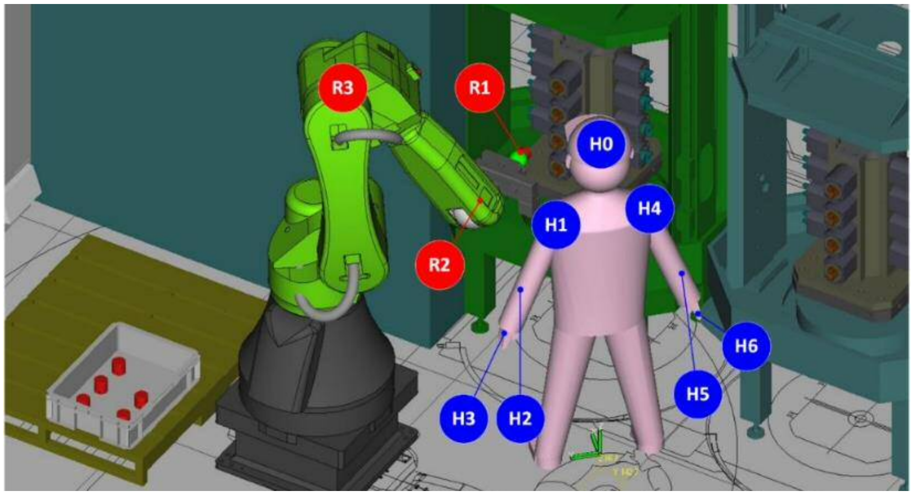

Figure 3: Points of interest of the human subject and of the robot, monitored during task execution

A direct interaction between human and robots in the same workspace is expected and the problem of identifying, quantifying and limiting dangerous and hazardous situations is crucial. A hazard can derive from the intended use in which robot and human are intentionally very close each other in the same shared space and accidental access to the cell or a misuse (wrong execution of a task). The number of occurrences of hazards deriving from the intended use defines the baseline of expected risk level, while the ones of the other two types are related to variable exposure of users to risks. To achieve this goal, motion tracking sensors can be used to gather, process, analyse and extract information from paradigmatic collaborative tasks. This provides a description of safety-related human-robot behaviour in terms of movements, layout occupation, contacts and near misses, in a prototype machining production setup.

In the *Figure 3* above, the Point of Interest (POI) of the human operator and the robot that have been used in the analysis are depicted, and the tracking of the POI is the proxy used to measure the human behaviour during the application. Data structure such as the one described in *Figure 4* must be acquired during all the process so that any possible correlations between the parameters in the table can be studied.

The cell was equipped with a high-payload robot as the Fanuc CR35. The chosen safety sensor was a laser scanner, in which were implemented the safety areas shown in *Figure 5b*.

t	POI	Intended task	Intended region	Actual region	Distance from R1	...	Distance from R1	Distance from H0	...	Distance from H6	Safety state	
	R1	[assign from list in 3.3 List of robot sub-tasks]		n/a (always as intended)	n/a						[assign from 4.1 List of robot system safety states]	
	R2											
	R3											
	H0	[assign from list in in 3.4 List of human sub-tasks]	[can be **multiple**, visited during the allocated task]	[**single** region, at the current t]				n/a				
	...											
	...											
	H6											

Figure 4: General structure of the data frame acquired during the monitoring phase

5) MOTION TRACKER

In order to track all the movements during task execution, two camera sensors were used; *Figure 3* shows the map of the Points of Interest (POI), whose position and relative distances were recorded.

- The main sensor (Microsoft-Kinect One®) was placed in location "***S***" (Figure 5*a*) in order to track large motions and detect any access to the shared workspace by the human operator. The SDK provided by the sensor producer allows easily the robust re-construction of the human skeleton, granting the possibility to stream the human joints position as sensor output, making easy the tracking of the POI.
- A Time of Flight (ToF) camera was placed in region "*a*" (Figure 5*a*), to track with higher precision, the small movements of the hands and head of the human operator. The point cloud in output from the TOF has been spatially filtered, to isolate the operator's hands/head from the surrounding environment. The centre of mass of each blob so identified is then computed and considered as POI.

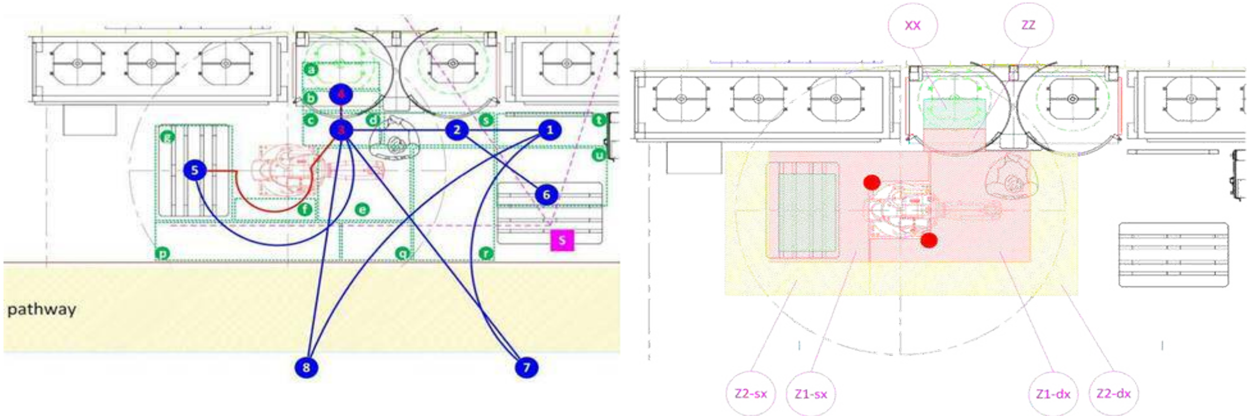

Figure 5: (a) Layout used to represent regions of occupancy (in green) with the correspondent labels, expected paths by human (in blue) and robot (in red); (b) Safety areas implemented through the laser scanner.

Each POI, identified from the main or the secondary sensor, is continuously monitored during task execution. For each instant of time, the software tool creates, visualizes and saves the data structure shown in the table of *Figure 4*; in particular, human operator POIs were monitored in terms of:

- Position with respect to the world frame
- Relative distances, which were chosen to be the elbow, shoulder and end-effector joints
- Sub-task in execution
- Regions of occupancy: the position of each POI was mapped inside a set of region labels

Along the task execution, both the robot and the operator are expected to occupy one of the labelled regions used for partitioning the workspace, and a set of intended regions of occupancy can be defined *a priori*, for each action to be executed. These intended regions are then recorded and put in comparison with the actual regions of occupancy for each time instant.

A difference between the two for a significant number of samples is the estimator of how well the preliminary model described in the previous paragraph fits the actual case-study, and consequently, it is a first hint of hazardous situations not expected and not included in the model.

Significant hazards are also related to the relative distances between POIs shorter than the minimum safety threshold corresponding to a potential contact. This threshold, established during the risk assessment stage as *0.5m*, defines the safety-stop area (in red in *Figure* 5*b*) in which the robot stops its motion when human presence is detected; the only exceptions to this rule occurred during collaborative tasks which required closest interaction. The combination of these two measures allows understanding when it is expected that human and robot POIs have small distances between each other and distinguishes it from cases in which this happened as a consequence of an unintended use.

6) DEMONSTRATOR DESCRIPTION

The RECOLL demonstrator scenario was the testing of two different working cycles carried out at a Load and Unload Station (LUS) of an FMS, where pieces are mounted on a fixture carried by a pallet. The main significant differences between these two cycles were the number and the dimension of the mechanical parts to be loaded and unloaded on the fixture and the different types of cooperative work between the human operator and the robot. In the first cycle, the parts to be machined were *8* for each fixture, each one with small dimension and relatively negligible weight. In each fixture different part types can coexist, depending on the position in which they are fixed; therefore the operators experience is crucial to avoid possible mistakes on that. At the beginning of the cycle the finished pieces, which have already been machined, are unloaded from the fixture and placed in a box in region "*g*", with respect to the layout shown in *Figure 5*. After that the unloading phase is finished and the raw materials can be loaded into the fixture. Each piece is recognized by the robots' own vision system, consisting in a camera placed in the end effector, so that the piece can be grasped and correctly placed in position. Since the friction between the piece and the fixture is really high, the robot needs the help of the human operator in both the loading and the unloading stage. In the second cycle, a single part of *15kg* has to be placed into its fixture, which is composed of various subparts to be screwed and unscrewed, in different positions; human operator's dexterity is required for this purpose while the robot carries the load. Also in this case, the process starts by unloading the finished piece and putting it in a box in region "*g*", then the raw piece is loaded into the pallet region in "*a*".

The cooperative work phases in the two cycles are different each other:

- In the first cycle, the expected physical interaction between the human and the robot is very low, and for a limited amount of time, an alternate and repetitive access to the cell is expected. The first type was executed *8 times*; the nominal time required for each repetition to complete is *18 minutes* for a total of *126* recording minutes.
- In the second case, the considerable weight of the load to be machined demands the human operator and the robot to work together simultaneously, resulting in a continuous physical interaction, but being a single part, with less repetitive actions. The second type was executed *22 times*, *5 minutes* each, for a total of *110* recording minutes.

The time of execution of each cycle is expressed considering a correct execution of the whole sub-task list, according to the intended use. Misuses during the execution can lead to huge differences in time execution from one repetition to another: when an unexpected action occurs, the robot goes in error state and suspend the execution of any movement. Depending on the moment in which this happens, there can be a considerable difference in the time required to restore the correct state of the cell. In general, the errors were managed distinguishing two different situations:

- The error state caused can be restored by the human operator without affecting the following tasks to be executed: in this case the robot waits for the human to confirm that the situation has been managed and then continues with the following actions.
- The error state inhibits the robot and/or the human from performing any further sub-task: in this case the whole cell must be reset to the initial state and the cycle must be restarted.

The second case is the most dangerous one in terms of safety for the operator (i.e. the robot could be stopped while carrying a load in an area potentially occupied by the human) and also for affecting the productivity of the whole plant. The data has been analysed considering three statistical values:

- Probability P_r that the actual region of occupancy is different from the intended, i.e., it represents how well the human and robot actions are represented in the preliminary model.
- Probability P_d that for each POI, the minimum distance is less than the safety threshold of *0.5m,* i.e., it is a proxy of how much the human and the robot have a close physical interaction.
- Combined risk probability P_{rd}: the aforementioned relative distance is below the safety threshold. This is a proxy for the performed interaction and identifies any possible residual hazardous situations not included in the interaction model.

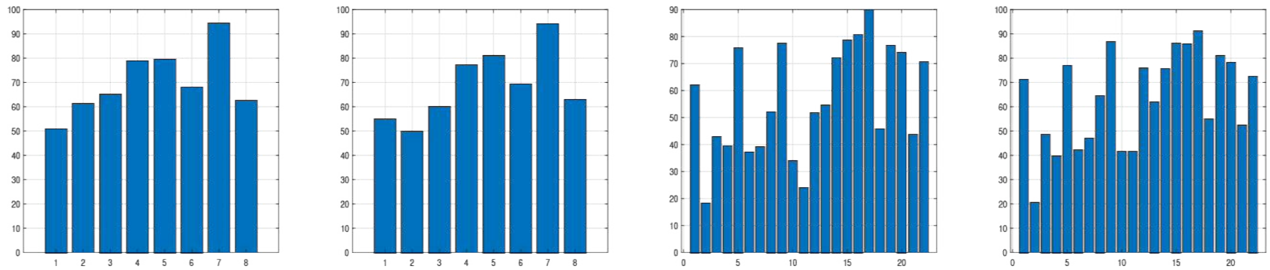

(a) (b) (c) (d)

Figure 6: Percentage that the head of the operator was actually in the intended region: Graphs 6a and 6c are referred to the first and the second task cycle from the main sensor; 6b and 6d are the first and the second cycle from the TOF camera

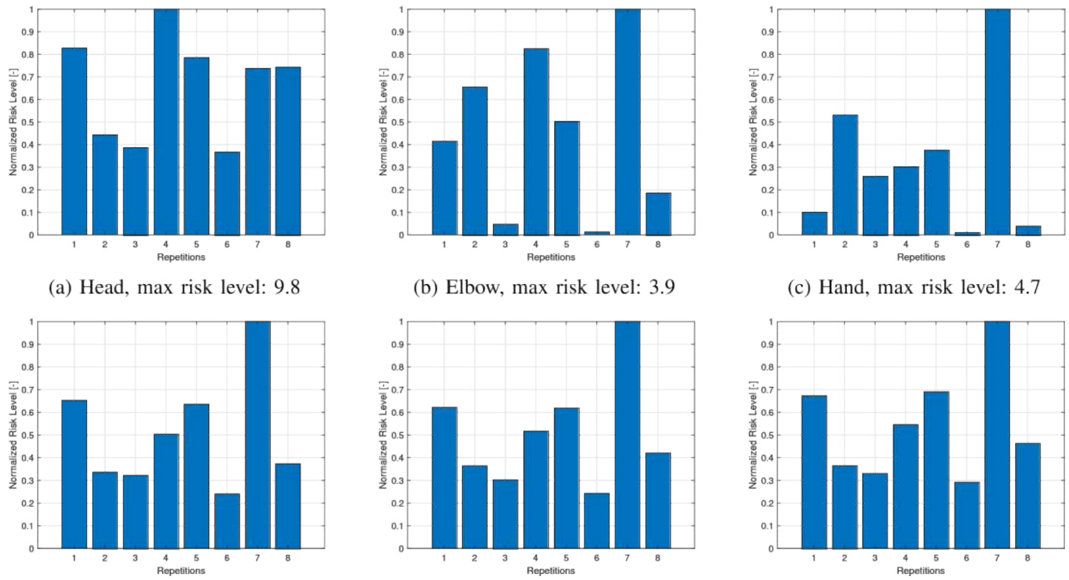

(a) Head, max risk level: 9.8 (b) Elbow, max risk level: 3.9 (c) Hand, max risk level: 4.7

(d) Head, max risk level: 13.5 (e) Elbow, max risk level: 14.5 (f) Hand, max risk level: 13.7

Figure 7: Normalized probability of risk in tasks of the first type. Graphs 7a to 7c are from the main sensor, 7d to 7f from TOF camera

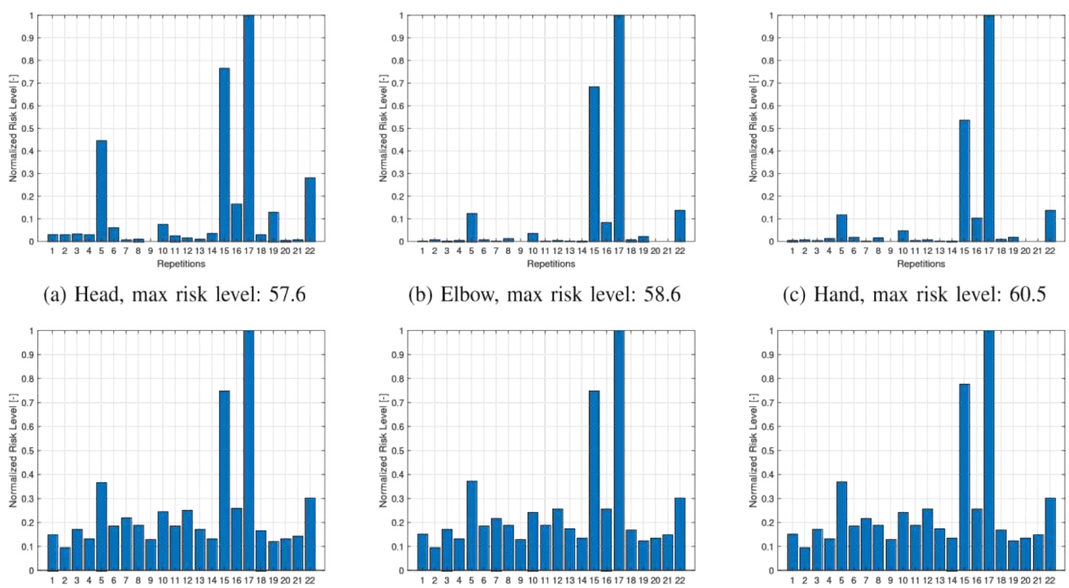

(a) Head, max risk level: 57.6 (b) Elbow, max risk level: 58.6 (c) Hand, max risk level: 60.5

(d) Head, max risk level: 66.1 (e) Elbow, max risk level: 66.3 (f) Hand, max risk level: 66.5

Figure 8: Normalized probability of risk in tasks of the second type. Graphs 8a to 8c are taken from the main sensor; 8d to 8f from TOF camera actual region occupied by the human is different from the intended one.

During the tests, the human operator could freely access to the cell and go away, also when it was not involved in tasks strictly necessary to the robot to continue its action. This provided the human operator with flexibility in movement and time scheduling, but, on the other hand, complicated a lot the possibility to model in advance the human position at any moment, increasing the region error P_r. This is shown in the bar graphs in *Figure 6*, where P_r is represented respectively in the first task type in *6a* and *6b*, and in the second type in *6c* and *6d*. Looking at these values, it seems clear that it is not possible to establish a regular pattern in data recorded trials which followed the intended use, and consequently it is not even possible to identify any outlier. However, a minor part of these unexpected accesses situations resulted in potential hazards and contacts: this can be seen when looking at the graphs in *Figure 7 and Figure 8*, where P_{rd} of the human operator's head, elbow and hand were recorded respectively from the main sensor in graphs *(a)* to *(c)*, and from the TOF camera in graphs *(d)* to *(f).* Here, values in tasks-cycles of both types follow a regular patter during nominal execution, so that it is possible to quantitatively define an "intended use".

Once the mean value for the probability has been established, it is possible to identify a threshold, and the outliers as the tasks when the probability of risk overcome the threshold. For excluding any possible error of the sensors, only outliers coming from both of the two cameras were identified as such. In addition, two operations were done in a preliminary stage:

- Filtering of data acquired in time instants in which the human was not in the scene.
- Filtering of data in which the relative distance was under the safety threshold for less than a window of *4* samples (if less, samples were supposed to be a result of a sensor fault condition).

In the first analysis performed, the values of probability P_{rd} were considered in general for each repetition of the experiment, considering the whole task cycle. *Figure 7* and *Figure 8* show the results for the *6* measured POIs in both the two task cycles. As a preliminary consideration, if looking at the values in general, in experiments of the second task-type (*Figure 8*), in which the low execution time allowed a higher number of experiments, a clear difference between repetitions within the mean value and outliers is noticeable. On the other hand, in tasks of the first type (*Figure 7*) the higher execution time and the huge number of different sub-tasks allowed fewer experiments and less repetition to be analysed.

In the first type of task cycle, data coming from the TOF camera (*Figure 7d*, *e* and *f*), showed a better and more regular behaviour: having a single outlier identified as trial number *7*, which found correspondence also in the values of *Elbow* and *Hand* recorded by the main sensor (*Figure 7b* and c). In graphs relative to the second type of working cycles, represented in *Figure 8*, it is evident that both sensors identified as outliers' the experiment repetitions number *5, 15* and *17*.

This means that in all of the four mentioned experiments when the human subject accessed the cell at an unexpected moment, it resulted in a contact situation excluded from the preliminary model.

Although there is a mean correspondence between the outliers in the task cycles, not clear information on what happened during the considered experiments can be extracted. Indeed, each cycles is composed of different sub-tasks, and embedding everything in a unique value results in a poor approximation.

All the data coming from the two task types have been divided into single human and robot sub-tasks: each complete task cycle was composed of a sub-set of *10* elementary actions. Any interaction occurred outside one of these sub-tasks (situations of fault condition of the robot, resulted from errors in the execution of one or more actions), was modelled within an additional action called *"others"* (see *Figure 9* and *Figure 10).*

Within the data of each cycle, each sub-task can occupy a variable percentage of samples; this depends on:

- The percentage of time in which the human operator performs its action.
- The percentage of time in which the human operator is in the scene during an autonomous robot action, also if his presence is not expected by the task model.

Action with a percentage below the 2% are not relevant and were not considered.

In *Figure 9*, results show that in the majority of the repetitions the sub-tasks with the higher risk probability are the collaborative actions, in which an interaction is expected. Specifically, looking at the probability levels of the action "*insert piece*" during repetitions in which the execution followed the intended use, is it possible to see that the operator may not be in the position predicted by the task model (so the actual regions of occupancy may be different from the intended ones).

Except that for the above sub-task, in all the other cases the probability level P_{rd} remained below the outlier limit, represented by the horizontal line in all graphs, for all the sub-tasks for each POI.

The only repetition in which this limit was overcome according to all the POI-measures (apart from *9a*), was number *7*, and the corresponding actions causing this were an error state ("*others*") and a sub-task in which human presence was not expected like "*unload*".

The same considerations can be done for the second task cycle if looking at *Figure 10*: in repetitions within the mean value, the higher risk probability was corresponding to the collaborative actions "*screw piece*" and "*unscrew piece*"; while in the three outliers (repetitions number *5, 15* and *17*) the high risk probability was monitored during an error state or during actions "*pre-unload*" and "*predispose*", which are meant to be executed autonomously by the robot.

Finally, the ROS monitoring software was used to reproduce the experiments which resulted to be outliers, with the intention of verifying that these values resulted from an actual wrong execution of the task and not from a sensor fault condition or a data processing error. *Figure 11* and *Figure 12* show a correct and incorrect execution of the same sub-task respectively, as it can be seen in the graphical user interface that was specifically developed. The complete structure of data of *Figure 4* is recorded in the upper part of the interface, while robot and human movements are visualized in a virtual model of the collaborative cell in the bottom of the screen. Reaching this point, further consideration can be done: in all the aforementioned cases, the sub-tasks identified as outliers, occurred just before or after a collaborative action. This happened for different reasons:

- Human operator accidentally gave confirmation to have executed an action which still had to be performed, letting start the robot action too early (repetition n. *7* in the first task cycle).
- Human operator started his action in advance, when the robot did not finish the previous sub-task and was still moving, which was the case of repetitions number *5* and *15* of the second task cycle.
- Human operator didn't follow the expected paths (see *Figure 5a*) to exit from the shared workspace after performing his action, which was the case of repetition number *17* of the second task cycle.

The errors detected in the aforementioned outliers resulted from a robot fault caused by these errors.

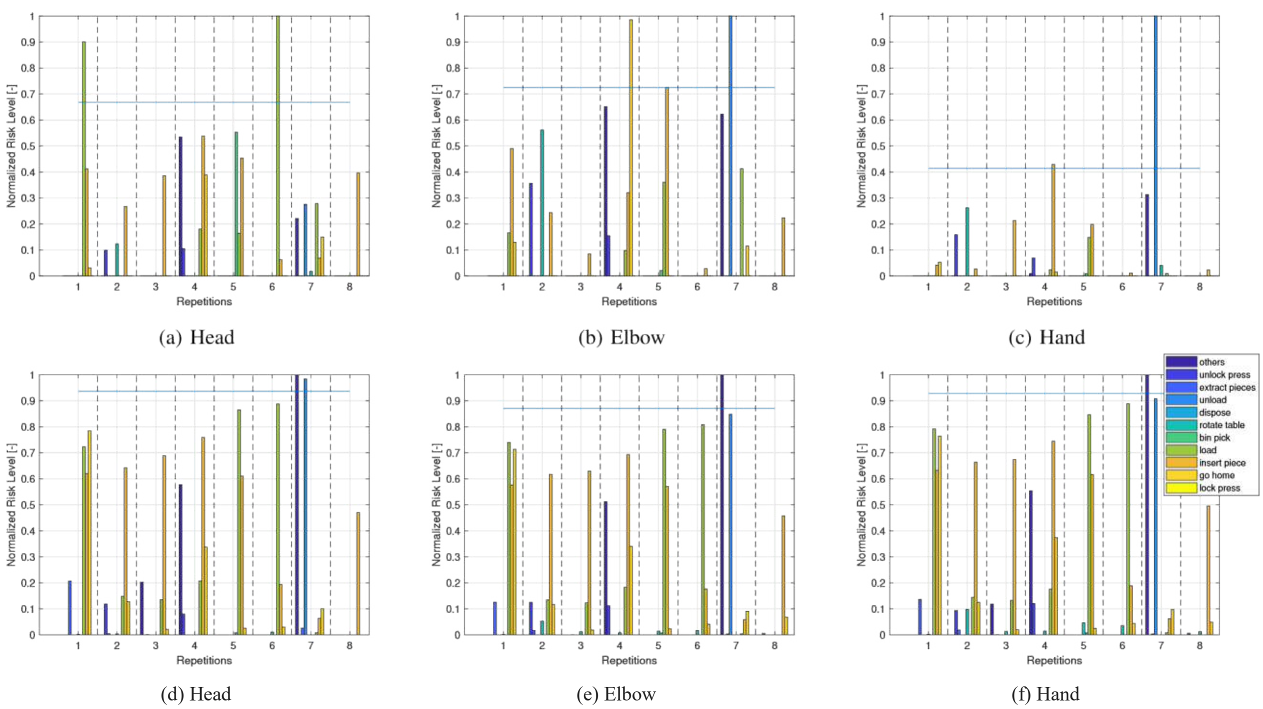

(a) Head (b) Elbow (c) Hand

(d) Head (e) Elbow (f) Hand

Figure 9: Normalized Probability of risk for all the sub-task which were monitored during the execution of task cycles of the first type. Graphs 9a to 9c are from the main sensor, 9d to 9f from TOF camera. The horizontal line represents the outlier limit

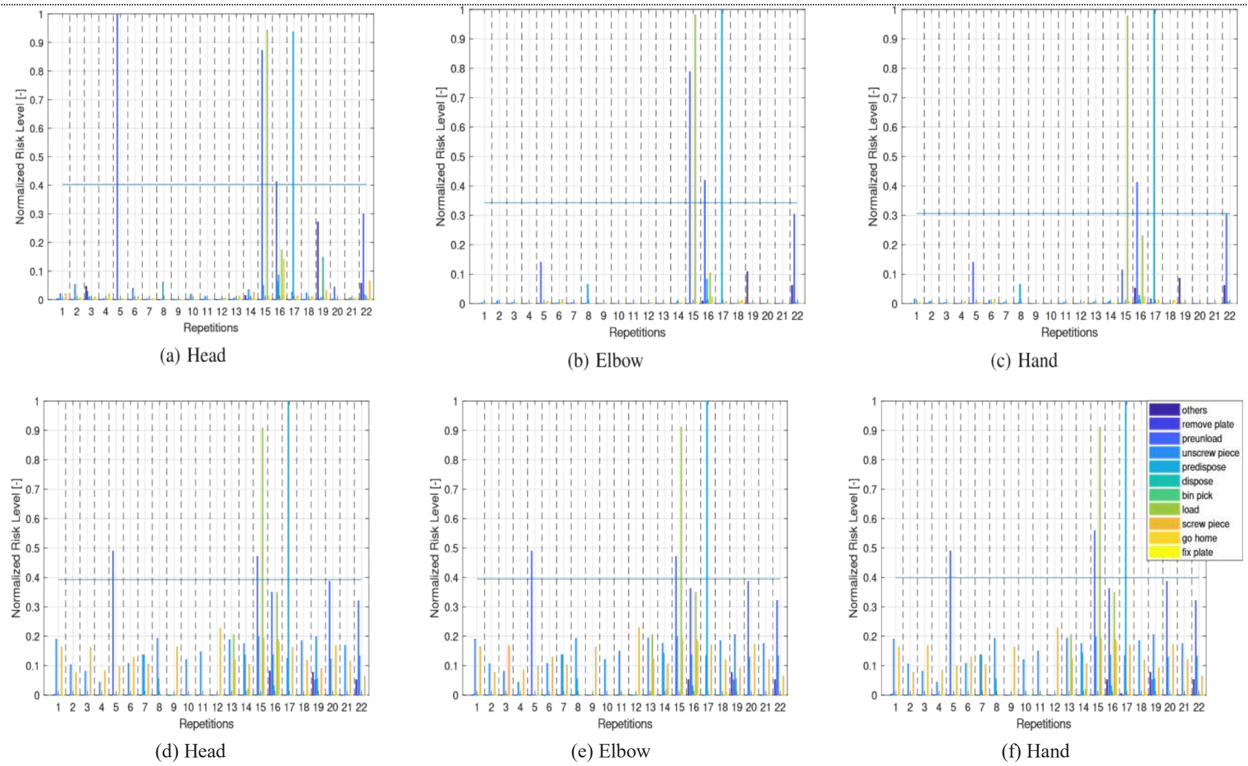

Figure 10: Normalized Probability of risk for all the sub-task which were monitored during the execution of task cycles of the second type. Graphs 10a to 10c are taken from the main sensor, 10d to 10f by TOF camera; the horizontal line is the outlier threshold

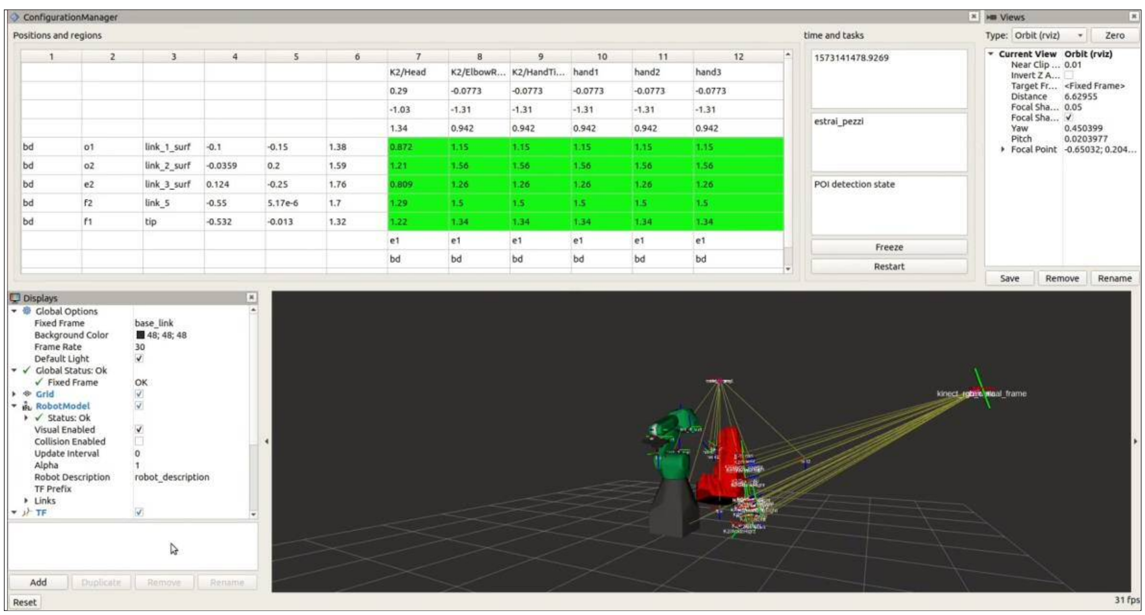

Figure 11: Task executed correctly, human accessed the pallet area when is expected to do it

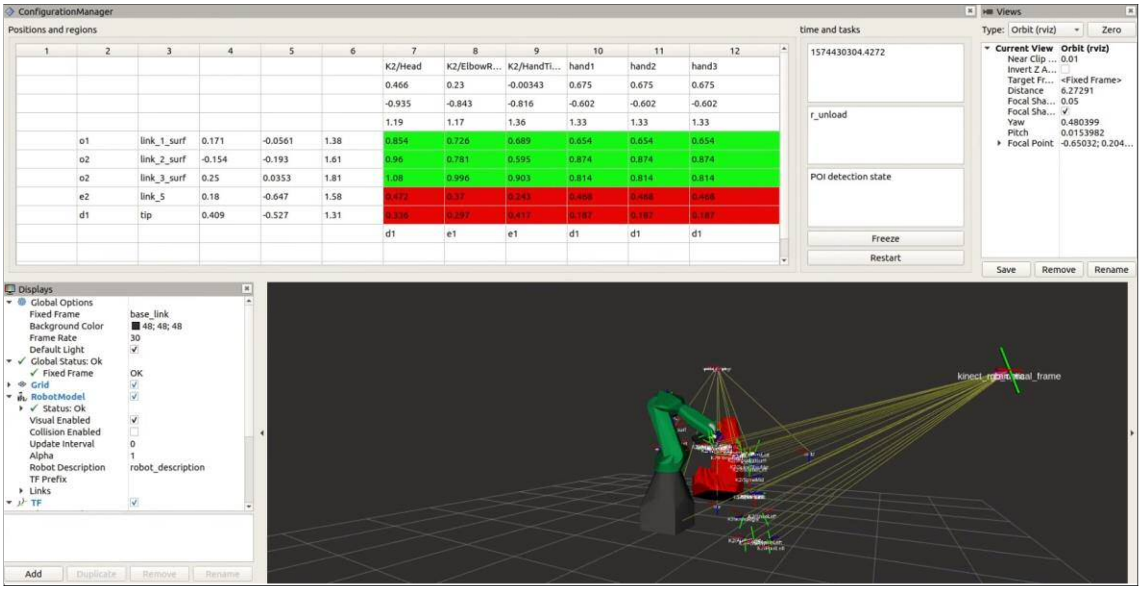

Figure 12: Error situation, human accessed the pallet area during a robot action, exposing himself to risk

7) OUTCOMES

The performed analysis showed that the system was capable of detecting unexpected possible hazards and collisions, and of distinguish them from regular patterns of physical interactions included in the model.

The safety rules were defined using a traditional and conservative approach; despite that, not all hazards were completely excluded. Oppositely, the most dangerous situations and hazards were observed during and after an error state of the robot as a result of the conservative safety rules which were applied. Most of those errors came after non-dangerous interactions and accesses to the shared workspace. For this reason, the traditional approach, although limiting the productivity of the collaborative cell did not fit the hybrid cooperative working scenarios that were proposed in this work.

The data analysis confirmed these considerations, giving us quantitative values of the occurrence probability of these hazards that were not included in the preliminary model and in safety-rules. Furthermore, results showed that most of these unexpected risks occurred during robot autonomous tasks, and not in actions in which the robot was supposed to be collaborative. A more dynamic approach in the definition of safety areas and safety rules that take into account human actual behaviour would be useful in future to overcome the limits that this method revelled.

The integration of RAS and human operators within flexible manufacturing centres entails many effects in the production experience, notably on efficiency and safety.

Finally, the main outcomes of the project are:

Outcome #1: A generalizable and well-documented approach to obtaining evidence for the safety of collaborative robots deployed in flexible manufacturing systems. The data-driven observations and records of human-robot interaction have provided figures about the concentration of added-value tasks for humans, error-reduction figures, workspace occupancy and human patterns figures, distribution of potentially critical actions (e.g. near misses).

Outcome #2: The identification of metrics to evaluate the correctness of the intended usages allows estimating the value of operator activity. Indeed, the application of the methodology allows considering what is the number of tasks that the human and the robot can safely perform in collaborative mode, weighted with the relative added-value of each task.

Outcome #3: Thanks to the outcomes coming from the assessment proposed, an efficient task planning and scheduling software tool may deployed, allowing likely an improvement of the throughput of 10-20%, while the errors should decrease of 50%.

ASSURING
AUTONOMY
INTERNATIONAL PROGRAMME