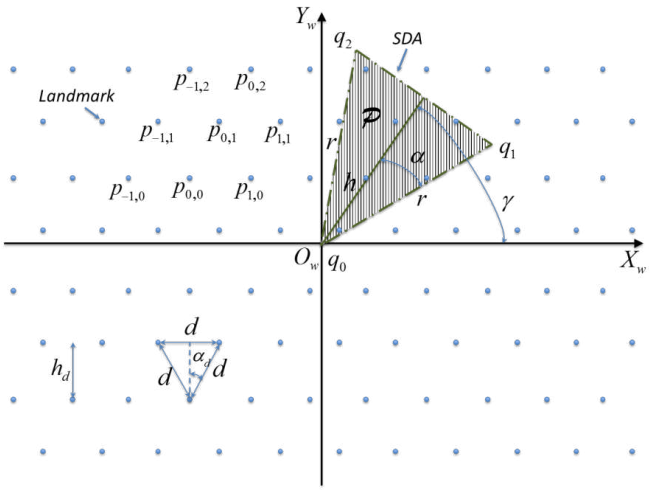

Figure 30: Lattice points representing the possible landmark positions in a wide-open, unbounded room. The FoV of the camera is represented by the shadowed region inside polytope $\mathcal{P}$.

The placement optimization problem can now be formulated as the problem of finding the maximum distance between landmarks, i.e. $max(d)$, such that for any position and orientation of $\mathcal{P}$, at least one lattice point lies inside the polytope area. This is equivalent to finding whether the polytope contains a lattice point for any possible position $\mathcal{P}$. Due to the symmetry and periodicity of the triangular lattice, in order to generate all the possible lattice positions, it is sufficient to move the reference point $p_{0,0}$ inside the rectangle $\mathcal{R}_d$, and change the polytope orientation angle γ in $[0, \pi/2]$, as depicted in Figure 31.

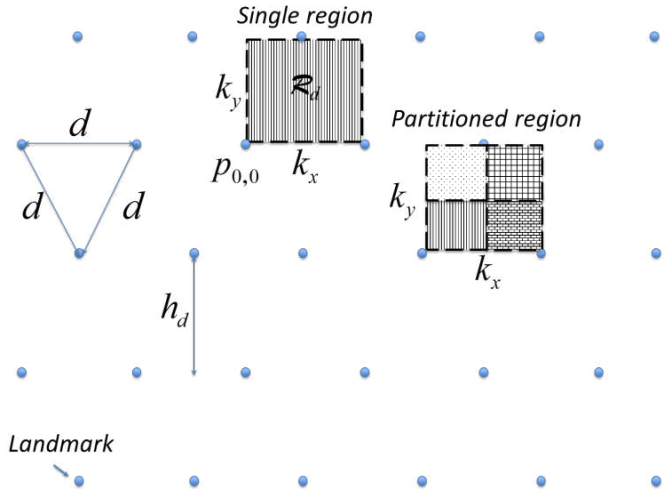

Figure 31: Possible landmark positions in the rectangle $\mathcal{R}_d$ and corresponding partitions.

Let $b = 2r \cdot \sin(\alpha)$ be the FoV maximum width, i.e. the length of the base of $\mathcal{P}$. In order to solve the optimization problem, two geometrical constraints must be satisfied, i.e. $d \leq r$ and $d \leq b$. Indeed, with reference to Figure 31, if these constraints are not met, at least one triple of values k_x, k_y and γ exists such that the sensor cannot detect any landmark (e.g. $\mathcal{P}_1$ and $\mathcal{P}_2$ in Figure 31). Consider now the polytope labelled as $\mathcal{P}_3$ in Figure 32. In this case, even if the base of the polytope is larger than the distance between landmarks (i.e. $b > d$), the first available lattice point lies in in the grey rectangular area adjacent to $\mathcal{P}_3$. This scenario suggests that we can possibly use a sensor with a larger FoV.

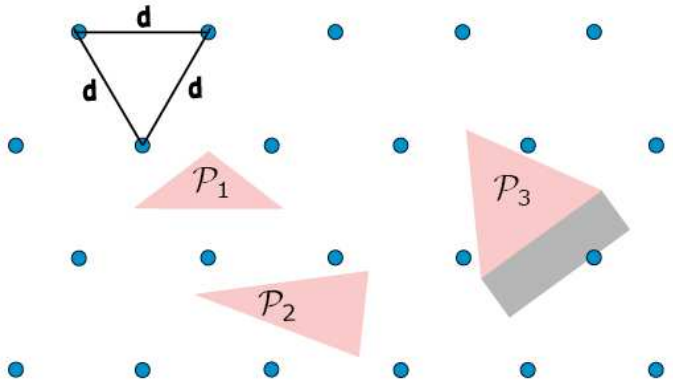

Figure 32: Pictorial examples of missed landmark detection, when the geometrical constraints $d \leq r$ (e.g. $\mathcal{P}_1$) and $d \leq b$ (e.g. $\mathcal{P}_2$) are not satisfied.

With reference to Figure 33, let us consider a camera with a triangular FoV $\mathcal{P}^{\mathrm{v}}$ included into $\mathcal{P}$, i.e. with $r^v \leq r$, $d \leq r^v$ and $d \leq b^v = 2r^v \sin(\alpha)$ Let be $\mathcal{R}^v$ the rectangle, with base b^v and height $H = h - h^v = (r - r^v)\cos(\alpha)$. Since $\{\mathcal{P}^{\mathrm{v}} \cup \mathcal{R}^v\} \subseteq \mathcal{P}$, $\{\mathcal{P}^{\mathrm{v}} \cup \mathcal{R}^v\}$ can be regarded as an inner approximation of polytope $\mathcal{P}$, which becomes increasingly accurate as $r^v \to r$.
From this perspective, the optimal placement problem can be rephrased as follows: maximize the landmark distance d such that $\forall p_{0,0} \in \mathcal{R}_d$ and $\forall \gamma \in [0, \pi/2)$ there exists at least one $p_{i,j} \in \{\mathcal{P}^{\mathrm{v}} \cup \mathcal{R}^v\}, i, j \in$ □.
To find an analytical solution to this problem, first of all notice that $\max d = \min(r, b_v)$ as $b^v \leq b$. Given that α, b and r are known parameters of the sensor while $b^v$is unknown, using simple geometric arguments, it can be shown that $b^v = 2\tan(\alpha)(r \cdot cos(\alpha) - H)$.

Therefore, in order to maximize b^v it is sufficient to minimize H. To this purpose, we should find the smallest rectangle $\mathcal{R}^v$ ensuring existence of at least one lattice point $p_{i,j}$ inside $\{\mathcal{P}^{\mathrm{v}} \cup \mathcal{R}^v\}$. If a landmark lies in $\mathcal{P}^{\mathrm{v}}$ the problem is straightforwardly solved. If instead no landmarks lie in $\mathcal{P}^{\mathrm{v}}$, then a landmark $p_{i,j}$ should be inside $\mathcal{R}^v$. This results in the fact that $\mathcal{R}_d \subseteq \mathcal{R}^v$. Since $d \leq b^v$, the minimum value of H such that $\mathcal{R}_d \subseteq \mathcal{R}^v$ is

$$H = h_d = d\cos\frac{\pi}{6} = \frac{d\sqrt{3}}{2} \tag{5.14}$$

Thus, the optimal distance between landmarks becomes

$$d^{\dagger} = b^v = r\frac{2\sin\alpha}{1+\sqrt{3}\tan\alpha} \tag{5.15}$$

Notice that $0 < d^{\dagger} < r$ because $\alpha \in [0, \pi/2]$.

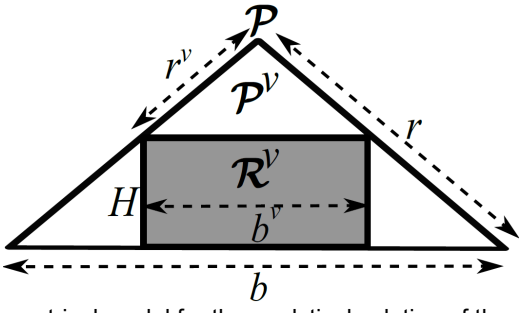

Figure 33: Approximate geometrical model for the analytical solution of the optimal placement problem.

5.2.1 Simulation results and analysis

In order to validate the correctness of the optimal placement strategy based on the analytical criterion described above, several simulations have been performed and compared with the results obtained with a numerical approach. In the latter case, $\mathcal{R}_d$ was partitioned into a settable

number of smaller sub-regions (i.e. by bisecting k_x and k_y iteratively and changing γ by small steps within each sub-region). Further details on this approach are reported in [63]. Ultimately, we checked whether analytical and numerical results converge to the same solution.
Figure 34 shows the optimal landmark distances normalized by r as a function of the sensor angular semi-range $\alpha \in (0, \pi/2)$. In the numerical case, the results refer to three different partitions of R_d (i.e. 1, 8 and 64 sub-regions). From Figure 34 it is quite evident that as the number of partitions grows, the result approaches to the optimal one and with smaller fluctuations due to the finer granularity of the regions explored. Ultimately, the numerical solutions with 64 sub-regions are hardly distinguishable from the optimal solution found analytically using (5.15).

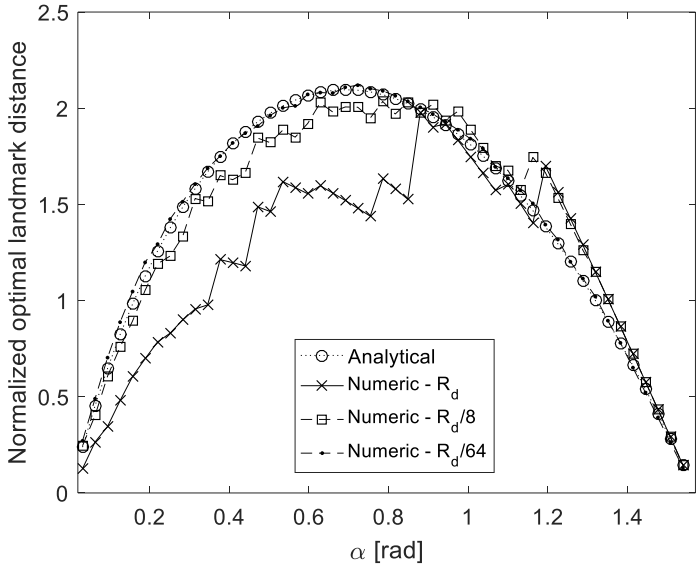

Figure 34: Optimal landmark distances (normalized by r) computed numerically and analytically as a function of the sensor angular semi-range of the camera α.

In order to evaluate more clearly to what extent the estimated landmark distance values are close to the optimal ones, further Monte Carlo simulations (10^5 runs) have been performed by randomly changing both camera position and orientation for some values of α. In each case, the actual distance d between landmarks was set purposely larger than $d^\dagger$ by a factor δ. Table 17 reports both the optimal values of $d^\dagger/r$ and the minimum values of factor δ for which, given $d = \delta \cdot d^\dagger$, there exists at least one configuration of the sensor (out of 10^5) where no landmark is detected. The results in Table 17 show that for small values of α, $d^\dagger$ is just slightly different from the real optimal value.

α [rad]	0.2	0.4	0.6	0.8	1.0	1.2	1.4
$\frac{d^\dagger}{r}$	0.32	0.46	0.53	0.53	0.46	0.35	0.18
δ	1.08	1.02	1.02	1.02	1.02	1.01	1.01

Table 17: Optimal landmark distances (computed analytically and normalized by *r*) and minimum values of factor δ for which no landmark is detected in at least one out of 10^5 randomly generated positions and orientations of the camera.

5.3 Application of Optimal Placement for Collaborative Localisation in a Public Scenario

In order to evaluate the impact of optimal landmark placement on localisation accuracy in a real case study, the results of some Monte Carlo simulations in two different indoor environments

are reported in the following, i.e. a large wide-open room without any obstacle, and a more realistic scenario based on the map of the Department of Information Engineering and Computer Science (DISI) of the University of Trento. In both environments, the *FriWalk* trajectories were generated using settings similar to those described in Section 5.1.3. By applying the optimal placement strategy described in Section 5.2.1, we found that the optimal distance between landmarks is $d^{\dagger} \approx 2$ m.

In the ideal case of a wide-open room without obstacles where the triangular lattice is deployed on the floor with distances between adjacent landmarks given by $d^{\dagger}$, the front camera of the *FriWalk* is always able to detect at least one landmark, regardless of camera's position and orientation. To verify this, 200 random-walk trajectories of 180 s each were generated within a 10×10∼m wide-open room. The results obtained in this case are comparable with those reported in Tab. 17. Indeed, by increasing the distance between adjacent landmarks by just a few percents with respect to the optimal value, it may happen that no landmark is detected.

The trajectories in an even more realistic environment (i.e. the DISI premises) were generated using the so-called Social Force Model (SFM). This provides realistic human-like paths along with collision avoidance mechanisms [68].

Figure 34 shows the DISI map along with six trajectories. While initial position and final destinations are generated randomly, the main difference with respect to the case of wide-open room is that now the optimal landmarks placement is affected by the presence of obstacles and walls. Anytime a point of the lattice is located inside a wall or an obstacle reported in the DISI map, obviously the corresponding landmark is not available in practice.

This problem can be partially addressed by shifting all QR codes by a fixed amount until the number of those falling inside walls or obstacles is minimum. However, while in the wide-open room case, anytime the camera detects a landmark, any one of its six neighbours can be detected next (as depicted in Figure 35 on the left), the presence of walls and obstacles can make the transition from one landmark to another impossible or much less likely, e.g. because the trajectory is constrained by a wall, as shown in Fig. 34 on the right. This kind of situations may considerably affect RMS estimation errors even if the landmarks layout is as close as possible to the optimal one.

This is due to the fact that when obstacles are present, the probability of detecting a landmark even in the case of optimal deployment decreases from 96% when r = 1 m to r = 87% when r = 8 m. As a result, the RMS positioning errors of state variables x_k, y_k and θ_k may increase by between 20% and 30% with respect to the results shown in Figure 29 when $d = d^{\dagger} = 2$ m.

5.3.1 Optimal placement with trajectory observations

An overview of the approach used in this section to overcome some of the limitations highlighted previously is the following. We start from a collection of data on the trajectories followed by the agents in the environment. Using this information, we derive a stochastic abstraction that reproduces the trajectories with a reasonable approximation. A stochastic abstraction is, in this context, a discrete–time Markov chain, in which states are associated with the "typical" locations visited by the target, and transitions describe the motion between two locations resulting from the observed trajectories and are annotated with a probability accounting for the frequency with which the transition is observed. The use of this type of stochastic abstractions is frequent in the literature [74], [75], [76]. In our case, a transition is also associated with a level of uncertainty, which accounts for the loss of localisation accuracy incurred in moving between start and end locations without using landmark information.

To describe a useful abstraction, we make use of the QR landmarks deployment described in Section 5.2. Indeed, with such a choice we have that a landmark will be always in sight of the *FriWalk*, apart from a limited set of locations depending on the particular geometry of the environment. Hence, the locations among which construct the Markov chain abstractions are given by the QR positions determined in Section 5.2. In practice, to solve this problem it is possible to cast the problem of finding the minimum number of sensors and their optimal position into the framework of mixed binary linear programming (MBLP). An important

theoretical result of this approach is that the problem can be addressed with a MBLP of limited size even in the presence of cyclic paths between the different locations.

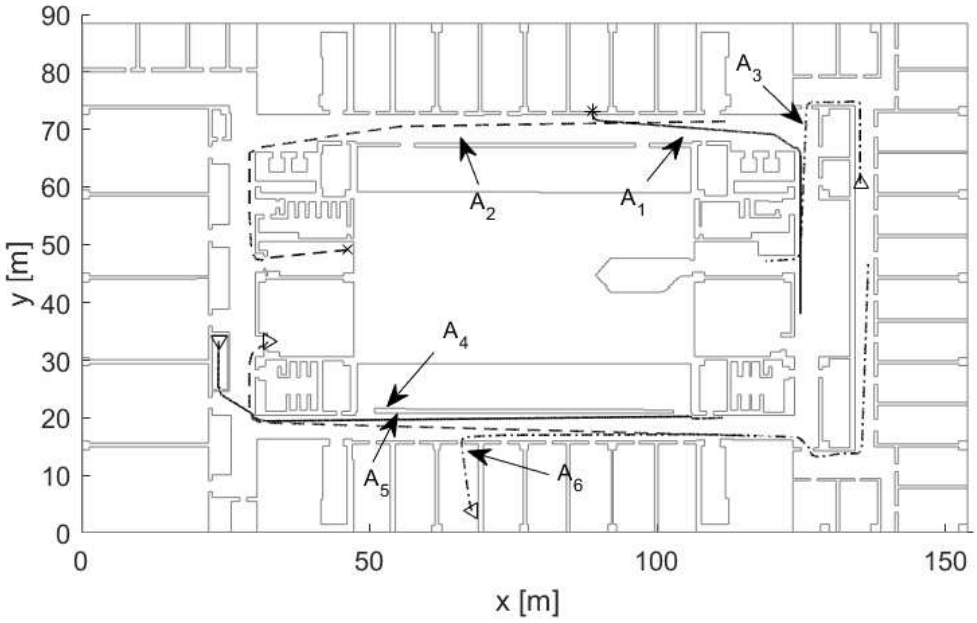

Figure 34: Six examples of agents' trajectories in the premises of the Department of Information Engineering and Computer Science of the University of Trento.

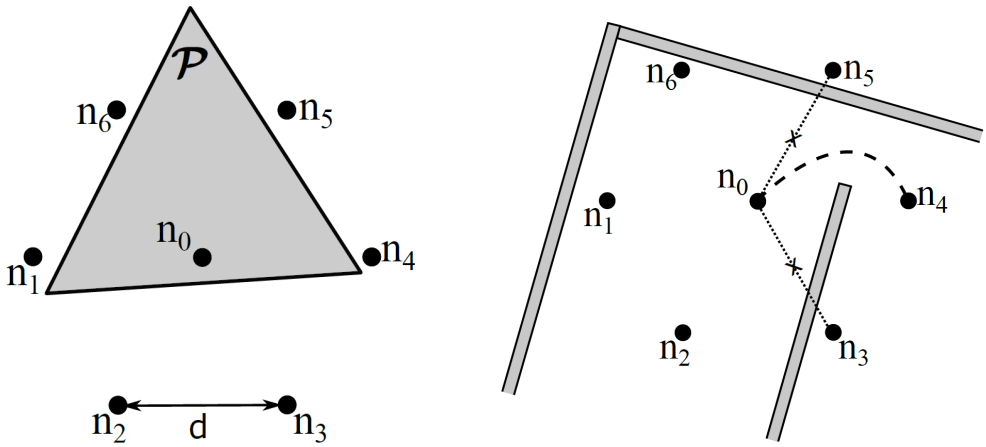

Figure 35: Qualitative relationship between the FoV of the robot's camera, the position of a detected landmark and its six neighbours over a portion of a triangular lattice in the case of a wide-open environment (on the left) and a room with obstacles, e.g. walls (on the right).

The core of the proposed abstraction relies on an oriented graph including three sets of nodes, i.e. the source nodes *S*, the intermediate nodes *V* and the destination nodes *D*. Each of these nodes set maps on a possible QR position as determined in Section 5.2. The general properties of such sets of nodes are shortly summarized below (see [77] for details):

- The source and destination nodes can have just outgoing or incoming edges, respectively;
- Position and heading at both source and destination nodes are assumed to be known or

measured with negligible uncertainty;

- At time t_0, the robot is assumed to be located in any one of the source nodes with a given probability;
- Every destination node is reachable from at least one source node;
- Each intermediate node represents the position of one of the possible QR landmarks to be deployed in the environment. Therefore, a binary variable is associated with each node. This variable will be set equal to 1 if a QR has to be placed in it or 0 otherwise;
- A node will be referred to as a *successor* of another node v_i for if the probability that it would be visited immediately after v_i (namely in one step) is larger than zero.

As a consequence of the assumptions above, the resulting system can be regarded as a discrete–time absorbing Markov chain. Even though the step duration in this case is not constant (as it depends on the actual path of the robot) this is not an issue for the solution of the optimization problem considered. Note that the set of graph edges joining the various nodes includes all the possible transitions with probability larger than 0. In the destination nodes instead just self-transitions are possible (i.e. their probability is equal to 1), since they are absorbing states. An example of the resulting graph is given in Figure 36.

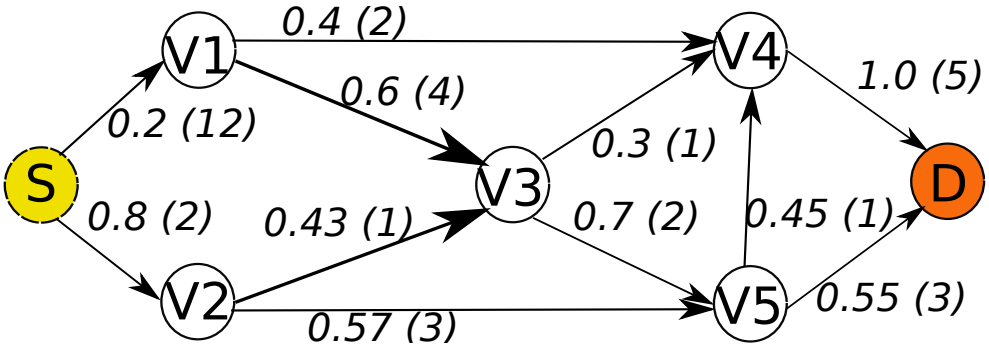

Figure 36: Graph example.

We assume that the motion of the robot is based on direct observations, if available, or synthesized by a random walk model and it is constrained only by the walls or fixed obstacles of the room. Moreover, a final important remark about the proposed abstraction concerns with the positioning uncertainty due to dead reckoning. In particular, the positioning uncertainty accumulated over the path between successor nodes can be described by a random variable. In theory, this variable is a continuous random variable. However, for the sake of simplicity and without loss of generality, we can assume that it is multiple of a "base" uncertainty quantity (for instance due to the minimum resolution of the measurement system).

As an example, the scenario considered in the testing phase to prove the solution effectiveness represents the 200 m^2 exposition area "Salone Donatello" of the Bargello National Museum in Florence, Italy. As reported in Figure 37, our model considers the room entrances and exits as source and destination nodes and, additionally, the set of POIs like paintings and sculptures (marked with pointers). The room is split into regular squares on a grid, representing intermediate nodes. Each square covers approximately 3 m^2 and there are N = 50 available squares on which the sensor can be placed. The range and the shape of the area of detection has been set in order to satisfy these constraints:

- For every position of the robot in this environment at least one node is in the detection range of the vision system, as the result of the placement of Section 5.2;
- The Robot is assumed to move freely in the space, as hypothesized previously;
- If the robot moves on a diagonal trajectory from a node to another one of the grid, the sensing system firstly detects a sensor on the same row or column.

The deployed sensing uncertainties and the ego-motion uncertainty are supposed to be modelled as in [78]. The Markov chain describing the motion of the vehicle inside the environment has been generated taking into account the attractiveness of the destination nodes

as well as the presence of the POI. More in depth, each POI has been modelled as an attractive potential for the generation of the node sequences by modifying the probabilities of the Markov model transition matrix, as exemplified in Figure 38. This definition of the attractive potential fields has a close resemblance to standard robot motion planning algorithms. Even if simulated data are used for this example, mainly due to the impossibility of accessing real data from the museum, the model here synthesized is in perfect accordance with widely used
stochastic motion models derived from direct observations, such as the Social Force Model widely used in the ACANTO project.

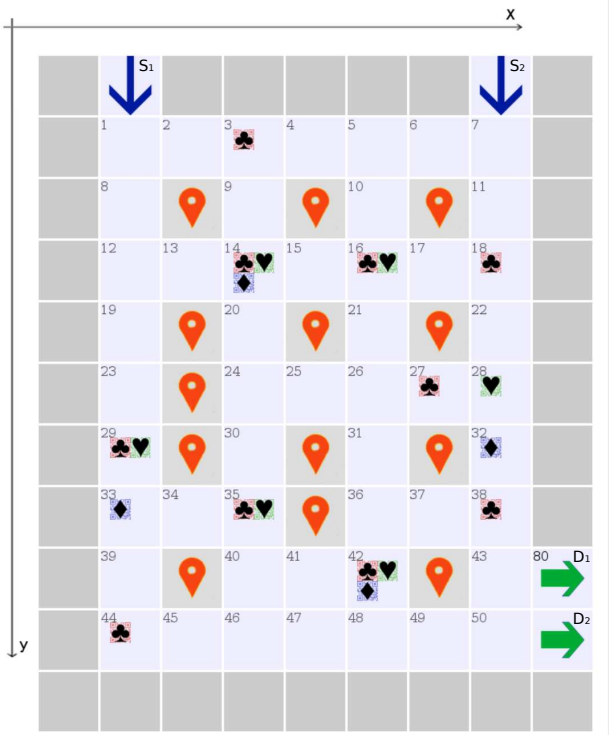

Figure 37: Map of the "Salone Donatello" of the Bargello National Museum use case with optimal sensors placement for increasing target uncertainty of H = 5 (clubs), H = 10 (hearts) and H = 15 (diamonds).

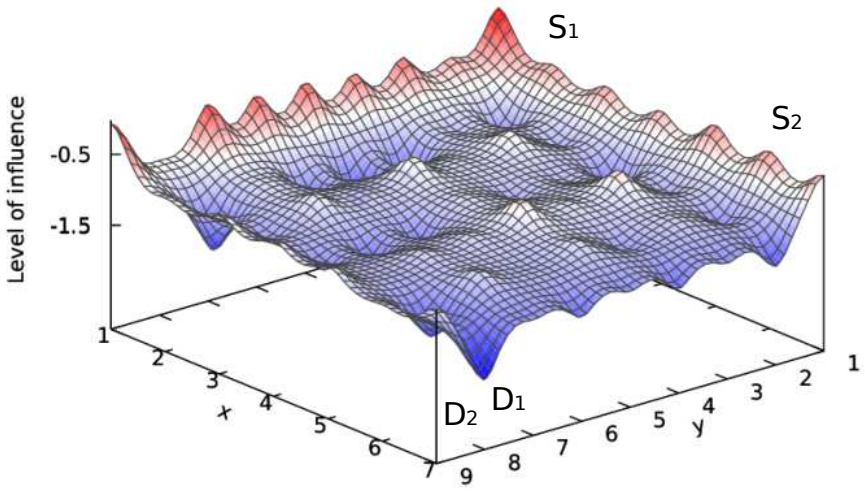

Figure 38: Points of interest influence on the Markov chain transition matrix of the use case of Figure 37.

The optimization algorithm determines the placement of the QR codes assuming a 95 % confidence and different values for the maximum uncertainty *H*. All the computations has been performed on a desktop equipped with a consumer quad core CPU and 16 GByte of RAM, while the total amount of trajectories, evaluated by a front-end coded on purpose in C++. The computation time of this front-end tool and the GLPK optimizer software for the described example are reported in Table 18, while the optimal position of the sensors for each target uncertainty *H* are depicted with different shapes in Figure 37.

$\alpha\ [rad]$	**0.2**	**0.4**	**0.6**	**0.8**	**1.0**	**1.2**	**1.4**
$\frac{d^{\dagger}}{r}$	0.32	0.46	0.53	0.53	0.46	0.35	0.18
δ	1.08	1.02	1.02	1.02	1.02	1.01	1.01

Table 18: Performance analysis of the "Salone Donatello" of the Bargello national museum use case.

The number of deployed sensors lowers down from 10 to 4, as expected, when *H* increases from 5 to 15 (i.e., a lower accuracy is requested). However, both the front-end and the optimizer computation times increase with lower requests in term of uncertainty (being the front-end the more computational intensive). This is due by the larger number of paths to be checked. It is worthwhile to point out that a time limit can be imposed to the software that in these cases provides a suboptimal but feasible solution. For instance, for $H = 15$, the optimizer reports a suboptimal solution with 5 sensors in less than 10 minutes, one third of the time reported for the optimal 4 sensors. Finally, the optimal solutions have been fed to an evaluation tool that verifies the constraints and provides the cumulative distribution functions of the positioning uncertainty after optimal placement for each choice of *H*. The solution is computed for a maximum value of $k = 22$ steps, which means that over the 88 % of the trajectories reached a destination node. This value can be considered as an index of the likelihood of the results.

Chapter 6
Relation with others WPs

We have presented a set of algorithms and methods to support our aim of user-centric sensing in a non-invasive way. Due to our systematic bottom-up approach of measuring biologic and biometric signals, derivation of higher level information has been simplified. Furthermore by establishing ZMQ as a flexible and standardized framework for communication, information exchanges between the various project partners and, hence, to other work packages has been eased.

From the clinical use case triggered by workshops held at the University Hospital of Getafe, the project consortium has identified a plurality of tangible use cases resulting in automated activity analysis for various diagnostic procedures and exercises. The rationale behind automation is a higher degree of quantification and hence repeatability as opposed to currently manual measurements taken, as well as the opportunity to execute these tests without the presence of highly qualified medical professionals up to even unsupervised execution of certain exercises. We present a fully automated solution for the prominent SPPB-test by the contributions of FORTH, SIEMENS and UNITN. In this respect, WP3 does directly support the goals of the clinical validation in WP8.

The user state model is meant to support activity analysis (WP4 and WP5) and hence the public scenario by transforming the raw sensor values into higher level semantic level (i.e. vigilance, activity index, stress) and can thus give an estimate on the overall well being on both the emotional as well as physiological condition. Exploiting this kind of information for reactive activity planning (i.e. during activity execution) as well as retrospective activity evaluation based on statistical data is left open for the future.

Localization is an underlying enabler for all the activities, clinical tests and navigation tasks that are foreseen in the project. Indeed, all the *FriWalk* features are feasible only once its location is known in the environment. From this perspective, we can safely state that localization is one of the basic pillar for the other tasks in WP3, as well as for WP4, WP5 and WP8. When collaboration comes into play, also a tight link is enforced with WP7 for what concerns the communication requirements.

Bibliography

[1] Guralnik, J. M., Ferrucci, L., Pieper, C. F., Leveille, S. G., Markides, K. S., Ostir, G. V., Studenski, S., Berkman, L. F. and Wallace, R. B. (2000), 'Lower extremity functionand subsequent disability: consistency across studies, predictive models, and value of gait speed alone compared with the short physical performance battery', Journals of Gerontology Series A: Biological Sciences and Medical Sciences, 55: M221–M231.

[2] Schöps, T., Engel, J., Cremers, D., (2000), 'Semi-Dense Visual Odometry for AR on a Smartphone', In International Symposium on Mixed and Augmented Reality (September 2014).

[3] Lucey P., Cohn J.F., Kanade T., Saragih J., Ambadar Z., Matthews I. (2010) 'The Extended Cohn-Kande Dataset (CK+): A complete dataset for action unit and emotion-specified expression'. In: CVPRW. pp 94–101.

[4] Lucey P., Cohn J.F., Prkachin K.M., Solomon P.E., Matthews I. (2011) "PAINFUL DATA: The UNBC-McMaster Shoulder Pain Expression Archive Database" In: FG. pp 57–64.

[5] Paltoglou G, Thelwall M (2013) 'Seeing stars of valence and arousal in blog posts'. TAFFC 4:1, pp 116–123.

[6] Prkachin K.K.M, Solomon P.P.E. (2008) 'The structure, reliability and validity of pain expression: Evidence from patients with shoulder pain'. Pain 139:2, pp 267–274.

[7] Prkachin K.M., Hughes E., Schultz I., Joy P., Hunt D. (2002) 'Real-time assessment of pain behaviour during clinical assessment of low back pain patients'. Pain 95:1, pp 23–30.

[8] Russell J.A. (1983), "Pancultural aspects of the human conceptual organization of emotions." J Pers Soc Psychol 45:6, pp 1281–1288.

[9] Scherer KR (2005) "What are emotions? And how can they be measured?" Soc Sci Inf 44:4, pp 695–729.

[10] Williams A.C. de C. (2002) 'Facial expression of pain: an evolutionary account'. Behav Brain Sci 25:4, pp 475–480.

[11] Yin L., Wei X., Sun Y., Wang J., Rosato M.J. (2006) "A 3D Facial Expression Database For Facial Behaviour Research" FG pp 211–216.

[12] Asthana A, Zafeiriou S, Cheng S, Pantic M (2013) 'Robust Discriminative Response Map Fitting with Constrained Local Models'. In: CVPR. pp 3444–3451.

[13] Balakrishnan G, Durand F, Guttag J (2013) "Detecting Pulse from Head Motions in Video," In: CVPR. pp 3430–3437.

[14] Boyd S, Parikh N, E Chu BP, Eckstein J (2011) "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," Found Trends Mach Learn 3:1, pp 1–122.

[15] De Haan G, Jeanne V (2013) "Robust pulse rate from chrominance-based rPPG," IEEE Trans Biomed Eng 60:10, pp 2878–2886.

[16] Hassner T, Harel S, Paz E, Enbar R (2015) "Effective face frontalization in unconstrained images," In: CVPR. pp 4295–4304.

[17] Jeni LA, Cohn JF, Kanade T (2015) "Dense 3D face alignment from 2D videos in real-time," In: FG. pp 1–8.

[18] Li X, Chen J, Zhao G, Pietikäinen M (2014), "Remote heart rate measurement from face videos under realistic situations," In: CVPR. pp 4264–4271.

[19] Poh M, McDuff D, Picard R (2010), "Non-contact, automated cardiac pulse measurements using video imaging and blind source separation.'. Opt Express 18:10, pp 10762–10774.

[20] Poh MZ, McDuff DJ, Picard RW (2011), "Advancements in noncontact, multiparameter physiological measurements using a webcam'. IEEE Trans Biomed Eng 58:1, pp 7–11.

[21] Tulyakov S., Sebe N. (2016) "Regressing a 3D face shape from a single image," In: ICCV. pp 3748–

3755.

[22] Verkruysse W., Svaasand L.O., Nelson J.S. (2008), "Remote plethysmographic imaging using ambient light," Opt Express 16:26, pp 63–86.

[23] Wang W., Cui Z., Yan Y., Feng J., Yan S., Shu X., Sebe N. (2016), "Recurrent Face Aging", In: CVPR. pp 2378–2386.

[24] Welch P.D. (1967), "The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms," IEEE Trans Audio Electroacoust 15:2, pp 70–73.

[25] Wu H.-Y., Rubinstein M., Shih E., Guttag J., Durand F., Freeman W. (2012) "Eulerian video magnification for revealing subtle changes in the world,". ACM Trans Graph 31:4, pp 1–8.

[26] Xiong X., De La Torre F., Xuehan X., De La Torre F. (2013) "Supervised descent method and its applications to face alignment" In: CVPR. pp 532–539.

[27] Zhang Z., Girard J.M, Wu Y., Zhang X., Liu P., Ciftci U., Canavan S., Reale M., Horowitz A., Yang H., Cohn J.F., Ji Q., Yin L. (2016), "Multimodal Spontaneous Emotion Corpus for Human Behaviour Analysis," In: CVPR. pp 3438–3446.

[28] A. Makris and A. Argyros, "Model-based 3D Hand Tracking with on-line Shape Adaptation," 2015, p. 77.1-77.12.

[29] J. Myers, 2003. Exercise and cardiovascular health. Circulation 107.1, e2-e5.

[30] D.R. Bassett et al., 2008. Walking, Cycling, and Obesity Rates in Europe, North America, and Australia. Journal of Physical Activity and Health, 5, 795-814.

[31] A.A. Hakim et al., 1998. Effects of Walking on Mortality among Nonsmoking retired men. The New England Journal of Medicine.

[33] J.J. Pilcher et al., 1996. Sleep quality versus sleep quantity: Relationships between sleep and measures of health, well-being and sleepiness in college students. Journal of Psychosomatic Research, 42(6), 583-596.

[34] R.S. Lazarus et al., 1962. A laboratory study of psychological stress produced by a motion picture film. Psychological Monographs: General and Applied, 76(34), 1-35.

[35] K.A. Muscatell and N.I. Eisenberger, 2012. A Social Neuroscience Perspective on Stress and Health. Social and Personality Psychology Compass, 6/12, 890-904.

[36] E. Kristal-Boneh et al., 2000. The association of resting heart rate with cardiovascular, cancer and all-cause mortality. European Heart Journal, 21, 116- 124.

[37] P. Greenland et al., 1999. Resting Heart Rate is a Risk Factor for Cardiovascular and Noncardiovascular Mortality. American Journal of Epidemiology, 149(9), 853- 862.

[39] P. Palatini and S. Julius, 2004. Elevated heart rate: a major risk factor for cardiovascular disease. Clinical and experimental hypertension, 26(7-8), 637-644.

[40] M. Partinen et al., 1982. Sleep disorders in relation to coronary heart disease. Acta Medica Scandinavica, 211.S660, 69-83.

[42] K.R. Fox, 1999. The influence of physical activity on mental well-being. Public health nutrition. 2.3a, 411-418.

[43] P. Hassmen et al., 2000. Physical exercise and psychological well-being: a population study in Finland. Preventive medicine. 30.1, 17-25.

[44] D. Michel and A. Argyros. Apparatuses, methods and systems for recovering a 3-dimensional skeletal model of the human body, March 2016.

[45] D.J. Terry et al., 1993. Effects of work stress on psychological well-being and job satisfaction: The stress-buffering role of social support. Australian Journal of Psychology. 45.3, 168-175.

[46] D. Michel, K. Panagiotakis, and A. Argyros. Tracking the articulated motion of the human-body with two rgbd cameras. Machine Vision Applications.

[47] A Shapiro, A Feng, R Wang, H Li, M Bolas, G Medioni, E Suma. Rapid avatar capture and simulation using commodity depth sensors, 27th Conference on Computer Animation and Social Agents, Houston, TX, May, 2014

[48] A. W. Feng, D. Casas, A. Shapiro. Avatar reshaping and automatic rigging using a deformable model, ACM SIGGRAPH Conference on Motion in Games, Paris, France, November, 2015

[49] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, Yaser Sheikh. Convolutional pose machines, CVPR 2016

[50] OpenNI. OpenNI User Guide. OpenNI organization, November 2010

[51] F. Ofli, R. Chaudhry, G. Kurillo, R. Vidal, and R. Bajcsy. Berkeley MHAD: A Comprehensive Multimodal Human Action Database. In IEEE Workshop on Applications on Computer Vision (WACV), 2013.

[52] H. Hamer, K. Schindler, E. Koller-Meier, and L. Van Gool. Tracking a Hand Manipulating an Object. In IEEE ICCV, 2009.

[53] S. Frattasi, M. Monti, and R. Prasad, "The synergic localisation system (SLS)," in Proc. Vehicular Technology Conference (VTC), vol. 4, Dallas, TX, USA, Sep. 2005, pp. 2753–2755.

[54] R. Buehrer, T. Jia, and B. Thompson, "Cooperative indoor position location using the parallel projection method," in Proc. International Conference on Indoor Positioning and Indoor Navigation (IPIN), Zurich, Switzerland, Sep. 2010, pp. 1–10.

[55] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probabilistic approach to collaborative multi-robot localisation," Autonomous Robots, vol. 8, no. 3, pp. 325–344, 2000.

[56] T. Bailey, M. Bryson, H. Mu, J. Vial, L. McCalman, and H. Durrant-Whyte, "Decentralised cooperative localisation for heterogeneous teams of mobile robots," in Proc. IEEE International Conference on Robotics and Automation (ICRA), Shanghai, China, May 2011, pp. 2859–2865.

[57] A. Prorok and A. Martinoli, "A reciprocal sampling algorithm for lightweight distributed multi-robot localisation," in Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), San Francisco, CA, USA, Sep. 2011, pp. 3241–3247.

[58] D. Taniuchi, X. Liu, D. Nakai, and T. Maekawa, "Spring model based collaborative indoor position estimation with neighbor mobile devices," IEEE Journal of Selected Topics in Signal Processing, vol. 9, no. 2, pp. 268–277, Mar. 2015.

[59] Wang, Xiaoqin, Y. Ahmet Şekercioğlu, and Tom Drummond, "Vision-based cooperative pose estimation for localisation in multi-robot systems equipped with RGB-D cameras," Robotics 4.1 (2014): 1-22.

[60] S. Panzieri, F. Pascucci, L. Sciavicco, and R. Setola, "Distributed cooperative localisation," Journal of Information Technology Research, vol. 6, no. 3, pp. 49–67, Jul. 2013.

[61] D. Fontanelli, D. Macii, P. Nazemzadeh and L. Palopoli, "Collaborative localisation of robotic wheeled walkers using interlaced Extended Kalman Filters," Proc. IEEE International Instrumentation and Measurement Technology Conference Proceedings, Taipei, Taiwan, May 2016, pp. 1-6.

[62] S. S. Kia, S. Rounds and S. Martinez, “Cooperative Localisation for Mobile Agents: A Recursive Decentralized Algorithm Based on Kalman-Filter Decoupling,” in IEEE Control Systems, vol. 36, no. 2, pp. 86-101, Apr. 2016.

[63] P. Nazemzadeh, D. Fontanelli, and D. Macii, “Optimal Placement of Landmarks for Indoor Localisation using Sensors with a Limited Range,” in Proc. Int. conference on Indoor Positioning and Indoor Navigation (IPIN), Alacalà de Henares, Madrid, Spain, Oct. 2016.

[64] Y. Ma, S. Soatto, J. Kosecka, and S. S. Sastry, “An invitation to 3-D vision: from images to geometric models,” Springer Science & Business Media, 2012, vol. 26.

[65] A. A. Khaliq, F. Pecora, and A. Saffiotti, “Inexpensive, reliable and localisation-free navigation using an RFID floor,” in European Conf. on Mobile Robots (ECMR), Lincoln, United Kingdom, Sep. 2015, pp. 1–7.

[66] Y.-C. Wang, C.-C. Hu, and Y.-C. Tseng, “Efficient placement and dispatch of sensors in a wireless sensor network,” IEEE Trans. Mobile Comput., vol. 7, no. 2, pp. 262–274, Feb. 2008.

[67] L. H. Erickson and S. M. LaValle, “An art gallery approach to ensuring that landmarks are distinguishable,” in Robotics:Science and Systems, Los Angeles, CA, USA, Jun. 2011.

[68] D. Helbing and P. Moln´ar, “Social force model for pedestrian dynamics,” Phys. Rev. E Stat. Nonlin. Soft., vol. 51, pp. 4282–4286, May 1995.

[69] T. Baltrušaitis, P. Robinson, L.P. Morency, “Openface: an open source facial behaviour analysis toolkit”, 2016 IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 1-10, 2016, IEEE.

[70] Tassi, Patricia, and Alain Muzet. "Sleep inertia." Sleep medicine reviews 4.4 (2000): 341-353.

[71] Emery, Charles F., et al. "Psychological and cognitive outcomes of a randomized trial of exercise among patients with chronic obstructive pulmonary disease." Health Psychology 17.3 (1998): 232.

[72] Walmsley, D. Jim, and Gareth J. Lewis. "The pace of pedestrian flows in cities." Environment and Behaviour 21.2 (1989): 123-150.

[73] Norris, Richard, Douglas Carroll, and Raymond Cochrane. "The effects of physical activity and exercise training on psychological stress and well-being in an adolescent population." Journal of psychosomatic research 36.1 (1992): 55-65.

[74] M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun, “Learning motion patterns of people for compliant robot motion,” The International Journal of Robotics Research, vol. 24, no. 1, pp. 31–48, 2005.

[75] T. Sasaki, D. Brscic, and H. Hashimoto, “Human-Observation-Based Extraction of Path Patterns for Mobile Robot Navigation,” IEEE Transactions on Industrial Electronics, vol. 57, no. 4, pp. 1401–1410, April 2010.

[76] H. Kretzschmar, M. Kuderer, and W. Burgard, “Learning to predict trajectories of cooperatively navigating agents,” in 2014 IEEE International Conference on Robotics and Automation (ICRA), May 2014, pp. 4015–4020.

[77] F. Zenatti, D. Fontanelli, L. Palopoli, D. Macii, P. Nazemzadeh, “Optimal placement of passive sensors for robot localisation”, 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 4586--4593, 2016, IEEE

[78] P. Nazemzadeh, F. Moro, D. Fontanelli, D. Macii, and L. Palopoli, “Indoor Positioning of a Robotic Walking Assistant for Large Public Environments,” IEEE Trans. on Instrumentation and Measurement, vol. 64, no. 11, pp. 2965–2976, Nov 2015

[79] Breiman, Leo. "Random forests." Machine learning, 45.1, pp. 5-32, 2001

[80] E. Foxlin, “Pedestrian tracking with shoe-mounted inertial sensors,” IEEE Comput. Graph. Appl., vol. 25, no. 6, pp. 38–46, Nov./Dec. 2005

[81] Wallen, Matthew P., et al. "Accuracy of heart rate watches: Implications for weight management." PloS one 11.5 (2016): e0154420.

[82] Wang, Robert, et al. "Accuracy of Wrist-Worn Heart Rate Monitors." Jama cardiology (2016).