# Privacy Preservation for Visual Data in Smart Cities

**Name**: Lucas Maris
**Laboratory' s name**: Ubiquitous Computing Systems Laboratory
**Supervisor' s name**: Keiichi Yasumoto

**Abstract** (should be within 1st page)

The realisation of smart cities rests on the successful deployment of large-scale sensing, where collected data is meant and expected to benefit the city's inhabitants. With a better view of pedestrian flows, cities could better fulfil their responsibilities of providing, designing and maintaining appropriate infrastructure, efficient transportation, or attractive tourism facilities. Video cameras are already prevalent in cities and could thus serve as very powerful and versatile sensors, but the rich nature of the information they capture raises substantial privacy concerns. Finding a balance between data utility for smart city tasks and privacy for individuals remains a significant challenge. This thesis explores how and to what extent pedestrian image data can be protected, and aims to address the core problem of protecting sensitive data without compromising the effectiveness of smart city tasks such as cross-camera person re-identification and demographic predictions. The proposed solution extends the notion of differential privacy to the image domain through a novel ε-IDP mechanism, which leverages pixelisation and colour quantisation to reduce the dimensionality of image data prior to noise addition and is shown to sensibly reduce the amount of required noise to achieve differential privacy. The effect of ε-IDP is extensively evaluated on common vision tasks and workable privacy-utility trade-offs are identified, at which performances remain reasonable while achieving better privacy guarantees than existing methods, and at which over half of the respondents of a user survey feel that privacy is sufficiently protected. Two relaxations of this mechanism, (ε,δ)-IDP and SegCAM-IDP, respectively based on failure probability and body part importance, are also introduced and their privacy-utility balance is compared to the original mechanism, highlighting the first may be of use when utility is to be prioritised over privacy, and the latter when some utility tasks are more important than others. To further characterise privacy, this thesis also formulates image domain equivalents of the k-anonymity and l-diversity privacy metrics, which use trained attribute classifiers as a proxy for the database tables required by these methods. Their behaviour confirms the proposed privacy mechanisms significantly reduce the privacy leakage from image data (sensibly better privacy metrics in all scenarios) while maintaining sufficient utility for practical vision-based applications (generally better performance than existing privacy-preserving methods).