Preface

Cities are realizing that opening access to their many data sources and using semantic models to provide a holistic view of this heterogeneous data can unleash economic growth, optimize their operational and strategic goals while addressing computational sustainability issues. We call the cities committed to a semantic infrastructure as a way to integrate, analyze and standardize access to their open data, Semantic Cities.

In a semantic city, available resources are harnessed safely, sustainably and efficiently to achieve positive, measurable economic and societal outcomes. Enabling city information as a utility, through a robust (expressive, dynamic, scalable) and (critically) a sustainable technology and socially synergistic ecosystem could drive significant benefits and opportunities. Data (and then information and knowledge) from people, systems and things is the single most scalable resource available to city stakeholders to reach the objective of semantic cities.

Two major trends are supporting semantic cities — open data and the semantic web. *Open data* is the idea that data should be accessible from everyone to use and republish as they wish, without restrictions from copyright, patents, or other mechanisms of control. A number of cities (including, London, Chicago, Washington D.C., and Dublin) have made their data publicly available, leveraging the technologies and principles from open (linked) data and the semantic web, interconnecting heterogeneous data. These technologies, principles and good practices are maturing and are becoming a perfect play-field for research-grade, scalable, and robust AI techniques.

This workshop aims to bring clarity and foster the communication among AI researchers, domain experts and city and local government officials beyond open data to models, standards, and reasoning for cities.

- Mark Fox, Freddy Lecue, Sheila McIlraith, Biplav Srivastava, Rosario Uceda-Sosa