

Article

Volunteers in the Smart City: Comparison of Contribution Strategies on Human-Centered Measures

Stefano Bennati ^{1,*}, Ivana Dusparic ², Rhythima Shinde ³ and Catholijn M. Jonker ^{3,4}

¹ Chair of Computational Social Science, ETH Zürich, Clausiusstrasse 50, 8092 Zürich, Switzerland

² School of Computer Science and Statistics, Trinity College Dublin, Dublin 2, Ireland; duspari@scss.tcd.ie

³ Interactive Intelligence Group, TU Delft, Mekelweg 4, 2628 Delft, The Netherlands; shinde@ifu.baug.ethz.ch (R.S.); C.M.Jonker@tudelft.nl (C.M.J.)

⁴ LIACS, Leiden University, Niels-Bohr-Weg 1, 2333 CA Leiden, The Netherlands

* Correspondence: sbennati@ethz.ch

Received: 28 September 2018; Accepted: 26 October 2018; Published: 31 October 2018

Abstract: Provision of smart city services often relies on users contribution, e.g., of data, which can be costly for the users in terms of privacy. Privacy risks, as well as unfair distribution of benefits to the users, should be minimized as they undermine user participation, which is crucial for the success of smart city applications. This paper investigates privacy, fairness, and social welfare in smart city applications by means of computer simulations grounded on real-world data, i.e., smart meter readings and participatory sensing. We generalize the use of public good theory as a model for resource management in smart city applications, by proposing a design principle that is applicable across application scenarios, where provision of a service depends on user contributions. We verify its applicability by showing its implementation in two scenarios: smart grid and traffic congestion information system. Following this design principle, we evaluate different classes of algorithms for resource management, with respect to human-centered measures, i.e., privacy, fairness and social welfare, and identify algorithm-specific trade-offs that are scenario independent. These results could be of interest to smart city application designers to choose a suitable algorithm given a scenario-specific set of requirements, and to users to choose a service based on an algorithm that matches their privacy preferences.

Keywords: participatory sensing; smart cities; public good; privacy; fairness

1. Introduction

Recent years have seen a substantial increase in active user participation in the smart city, and, with it, an increase in the resources contributed by the citizens. Sensor data is one type of user-contributed resource that is at the base of many smart city applications [1]. Collecting data in a smart city allows the applications to predict the needs of the citizens [2], thus enabling the creation of more advanced and more efficient services with a high potential for innovation [3].

User participation is crucial for the success of some smart city applications, but it entails costs that disincentivize users. For example, transmitting privacy-sensitive data in a participatory sensing scenario increases the risk of disclosure and misuse of private information, e.g., discrimination. Similarly, increasing energy availability in the smart grid scenario, e.g., by postponing the use of appliances, comes with a risk of disproportionately low access to the resource and unfair treatment [4]. In order for the smart city applications to be successful, these costs and risks must be reduced.

Examples of existing solutions for reducing contribution costs are privacy-enhancing technologies [5] that reduce the risks associated with disclosing information [6], and fair resource-management technologies [7], which improve the perceived fairness of the system. However,

comparing different implementations of smart city algorithms is difficult because they are tied to specific assumptions about the scenario and the resource.

The first contribution of this paper is to introduce a design principle that enables such a comparison by means of a scenario-independent model of user contributions to a smart city application, which is based on the theory of public goods and voluntary contribution games [8]. This principle is based on the concept of contribution strategy: the orchestration of demand and supply of some common resource, i.e., the smart city service, is driven by independent contributions of individuals in a population. This approach is independent of the characteristics of the resource; hence it can be combined with existing mechanisms that work at the content level, e.g., privacy-preserving algorithms. This also makes data sharing across multiple smart city applications easier, thanks to the uniform way of dealing with different kinds of data.

Following this design principle, a simulation framework is developed (available on Github at [9]) with the goal of understanding user participation in smart city services by addressing the research question: how do different contribution strategies compare in terms of privacy, fairness and social welfare?

The second contribution of this paper is to verify the applicability of the design principle in two different examples of smart city application scenarios, i.e., traffic congestion information and charging of electric vehicles, using real-world datasets. Similar validation in different scenarios can provide useful insights for the design of other resource-based smart city applications.

The third contribution is to compare different contribution strategies—centralized optimization and two flavors of localized learning, with and without contextual awareness—in terms of trade-offs between system efficiency, user privacy or resource distribution fairness, and to identify scenario-independent trade-offs that characterize a specific algorithm.

Specifically, localized solutions, which work only on local data, are found to deliver higher privacy and equality than a centralized solution. Centralized optimization offers instead higher efficiency thanks to the global knowledge of data. Localized algorithms with and without contextual knowledge are evaluated, and those considering the current context in the decision-making are found to deliver higher fairness.

This work can be of interest to designers of smart city applications and services that look for a guideline for the choice of algorithm, given specific scenario priorities in terms of different measures. Another target audience is that of citizens that are concerned with the risk of abuse of their data, and particularly privacy-aware prosumers. The results help quantify different service implementations along measures such as privacy and fairness; thus help the citizens choose a service provider that best matches their preferences, reducing the concerns about the system and fostering user participation.

The rest of the paper is organized as follows: Section 2 introduces the proposed design principle and describes how it can be applied to two smart city scenarios, Section 3 describes the contribution strategies and the measures used to evaluate them, Section 4 describes and discusses the main results obtained from simulating voluntary contributions in different scenarios, Section 5 situates the work in the literature, and Section 6 concludes the paper discussing possible avenues for future work and giving general recommendations about the choice of contribution strategies.

2. Model

This section introduces the design principle, based on the theory of public goods and Voluntary Contribution Games [8], and it describes how it can be applied to the smart city scenarios of participatory sensing, with the example of traffic congestion [10], and smart grids, with the example of electric vehicle charging [11]. Voluntary contribution games involve the need of coordination between users to provide some public good, e.g., financing a playground with private donations. The contribution of this paper is to present Voluntary contribution games as a design principle for modeling generic smart city services, and not limited to modeling only a specific scenario, e.g., smart grids [12]. The rest of this section presents the details of our proposed design principle.

Table 1 illustrates the notation of the Voluntary contribution games model; symbols are listed in order of appearance.

Table 1. Mathematical notation, in order of appearance.

Math Symbol	Description
Time-Independent Variables	
S	The service provider
$V = \{1, \dots, n\}$	The set of n users
$T \in \mathbb{N}_{>0}$	The number of rounds in the simulation
Q	The total quality of service after T rounds
$\mathcal{A} = \{D, C\}$	The action set
Functions	
$f : V \rightarrow \mathbb{R}$	The quality function
$S : \mathbb{R} \times \mathbb{R} \rightarrow \mathcal{R}$	The success function
$G : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$	The payoff function for successful rounds
$B : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$	The payoff function for unsuccessful rounds
Time-Dependent Variables	
r_i	The resource produced by user i
$v_i \in \mathbb{R}$	The value associated to disclosing r_i
$c_i \in \mathbb{R}$	The cost associated to transmitting r_i
$p_i \in \mathbb{R}$	The privacy leaked when disclosing r_i
$a_i \in \mathcal{A}$	The action of user i
$A_+ \subseteq V$	The set of contributors
$q \in \mathbb{R}$	The service quality
$\tau \in \mathbb{R}$	The quality requirement
$G \in \mathbb{R}$	The payoff for a successful round
$B \in \mathbb{R}$	The payoff for an unsuccessful round
U_i	The utility of user i
Scenario: Smart Grid	
$\pi_i \in \mathbb{R}$	The energy production of household i
$\beta_i \in \mathbb{R}$	The baseline consumption of household i
$\sigma \in \mathbb{R}$	The energy surplus

$V = \{1, \dots, n\}$ is the set of users and S is a service provider. In each round $t \leq T$, each user produces a resource r_i with value v_i . Users perform an action $a_i \in \mathcal{A} = \{C, D\}$: C corresponds to contributing the resource and D to opt out from contribution. The set of contributors is denoted as $A_+ = \{i \in V : a_i = C\}$.

Contribution is costly; thus, contributors pay a cost c_i that depends on the characteristics of resource and communication medium. Contribution might also entail a privacy cost p_i , which models the risks of revealing private information to third parties.

The service provider determines the service quality $q = f(A_+) = \sum_{i \in A_+} v_i$ based on all contributions received. A quality requirement τ , either global or per user, is generated at every timestep. Not all users are required to contribute in order to satisfy the requirement $\tau \leq \sum_i v_i$. A round is successful, i.e., the service can be provided, if the quality is higher than the threshold: the success function is defined as $S(q, \tau) = \{G \text{ if } q \geq \tau \text{ else } B\}$. Each user gets the same positive payoff $G = G(\tau, q)$ from accessing the service, including those who did not contribute. If the quality threshold is not met, the service cannot be provided and every agent receives a large negative payoff $B(\tau, q) > 0$. Given that payoffs are distributed equally, there is no incentive for users to contribute unilaterally: the public goods theory predicts the existence of an equilibrium where nobody contributes.

The individual goal of the users is to maximize their individual utility over T rounds (see Table 2):

$$U_i = \begin{cases} G(\tau, q) & \text{if } q \geq \tau \\ -B(\tau, q) & \text{if } q < \tau \end{cases} - \begin{cases} c_i & \text{if } a_i = C, \\ 0 & \text{otherwise.} \end{cases}$$

The social goal, i.e., the goal of the service provider, is to maximize the total quality of the service over T rounds $Q = \max \sum_{t=1}^T f(A_+)$.

The model relies on the following assumptions:

- Users contribute their resources to a central entity, which uses them to provide a service.
- Contributions cannot be doctored, e.g., to reduce the cost.
- Users are not allowed to communicate with each other; this is generally the case for simple devices such as smart meters.
- The system implements a privacy-preserving resource-management algorithm that optimizes the use of the resource.

Table 2. Definition of utilities for agent i . Let $q_{-i} = \sum_{j \in A_+ \setminus \{i\}} v_j = q - v_i$ be the total contributions, excluding agent i , and τ be the global requirement. This game qualifies as a threshold public goods game if $G(\tau, q) + B(\tau, q) > c_i$, which is always verified for large enough values of G or B .

q_{-i}	$q_{-i} + v_i < \tau$	$\tau - v_i \leq q_{-i} < \tau$	$\tau \leq q_{-i}$
Outcome	Failure	Depends on a_i	Success
$U_i(a_i = C)$	$-B(\tau^t, q^t) - c_i^t$	$G(\tau^t, q^t) - c_i^t$	$G(\tau^t, q^t) - c_i^t$
$U_i(a_i = D)$	$-B(\tau^t, q^t)$	$-B(\tau^t, q^t)$	$G(\tau^t, q^t)$

2.1. Implementation of the Model in Real-Life Scenarios

This section discusses the generality of the proposed model by describing how real-life data-based smart city applications can be modelled using our proposed design principle. We introduce smart grid (electric vehicle charging) and traffic congestion information (based on data contributed by participatory sensing) scenarios. We describe the data used by the model and illustrate examples of a public good, contribution, cost and value in each scenario.

2.1.1. Smart Grid: EV Charging

In the smart grid scenario, we assume a neighborhood composed of households connected to each other via the infrastructure of the smart grid. Households consume a variable quantity of energy, depending on the appliances in use, and produce a variable quantity of renewable energy, depending on weather conditions. Electric vehicles (EVs), associated with households connected to the smart grid, are required to be periodically recharged. Demand-side management mechanism [13] is required to schedule and postpone appropriately EV charging, based on current demand and renewable energy production, i.e., energy availability.

The *public good* consists of the total energy surplus $\sigma = \pi - \beta$, which is available to all households for charging the EVs. Its *value* is computed from the current production of renewable energy $\pi = \sum_i \pi_i$, and the current network load $\beta = \sum_i \beta_i$ [14]. We assume that the energy surplus is lower than the total demand, i.e., at any given time $\pi < \beta$; hence, a black-out is inevitable if all households decide to use the surplus simultaneously. *Contribution* to the public good is then defined as opting out from consumption, i.e., renouncing to charge the EV of a charge v_i that might depend on contextual variables such as the current charge level, the availability of a charging station, and the current energy surplus. It is assumed that the values are privacy-sensitive as they might be used to infer habits, e.g., the work schedule, of the owners. Contribution entails a comfort cost that is assumed to be proportional to the corresponding value, as the utility of an EV depends on it being charged. *Values* v_i ,

which corresponds to the charge that the EV can accumulate during a time period, are generated uniformly at random between 1 and $n2$, while the corresponding costs c_i are determined by sampling a normal distribution centered around the value v_i . The threshold τ is dynamically computed at every timestep from the dataset to be inversely proportional to the difference between production and consumption: $\tau = 1 - (\pi_i - \beta_i)$.

The baseload level β —non reschedulable load for each household—determines the aggregated consumption of the network and is obtained from Irish Smart Meter trial data [15]. On a single household level, on average, it ranges from 0.8 kW, during the night, to 3 kW, at evening peak time. The energy production level π is based on Irish wind production data, obtained from Irish electricity grid operator EirGrid [16], which represents an estimate of the total electrical output of all wind farms on the system. Individual contribution values and costs are randomly sampled from a uniform distribution that ranges from 1 to a maximum value that is determined by the simulations' parameters.

2.1.2. Participatory Sensing: Traffic Congestion Information

In this scenario, we assume a service provider which optimizes the congestion of the whole road network [17] and provide live traffic information to the users. Traffic statistics are considered a *public good* as their creation requires that enough users *contribute* their mobility traces to the service provider [18]. The value v_i of individual measurements reflects the novelty of the information, which might depend on the actions of other users [19], e.g., duplicated information due to local correlation in measurements. In this specific case, measurements cannot be linked with one another as the dataset lacks GPS coordinates; therefore, the *value* of novelty is approximated with the change of speed, as a sudden change in speed is considered more informative than keeping a constant speed.

In this scenario, contributing entails a privacy *cost* that comes from revealing privacy-sensitive information such as an individual's location [20] or destination, as this information can reveal users' habits, even if the locations are obfuscated [21,22]. Another source of private information is the travel speed, as it might reveal violation of the road laws. Depending on the contribution strategy, private information might be disclosed even by non-contributors, e.g., computing the optimal allocation of contributors might require all users to reveal their location.

Real mobility traces of cars are provided by the 2011 Atlanta Regional Travel Survey by the U.S. Department of Energy, National Renewable Energy Laboratory (NREL) [23], containing travel speed of 25,797 participants, encoded second-by-second GPS information (speed) over the course of three days. Trips with less than 1000 data points are removed from the dataset, which results in a total of 663 people in 377 households. The values and cost are generated from the data and rescaled within the range from 1 to a maximum value that is determined by the simulation's parameters. The distribution of frequencies of value/cost pairs in the NREL dataset is shown in Figure A3 in the Appendix.

Contributing data comes at a transmission cost c_i —that might depend on the characteristics of the message, e.g., size, and of the medium, e.g., congestion, power consumption—but also, most relevantly to our scenario, privacy cost p_i —that might depend on disclosing private information, such as location and speed [21]. Therefore, *costs* are determined as the distance to the only points of interests known in the data: the origin and destination of a trip. The cost c_i is the highest at either the source or the destination and decreases as the distance from these points increases (see Figure 1). The *public good* is created if the sum of individual contributions is higher than a certain threshold τ , which is set to 80% of the size of the population, which means that the service is successfully generated if at least 80% of the users contribute a value of 1.

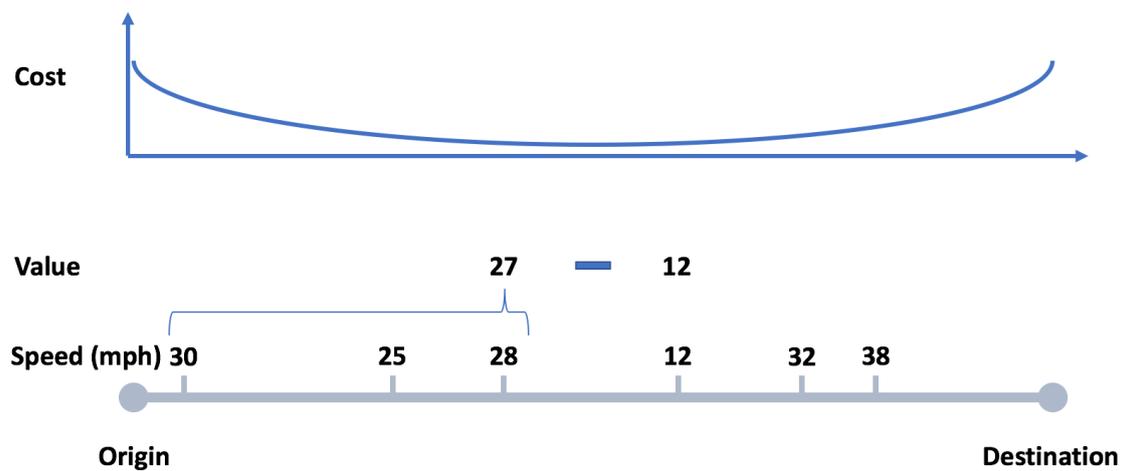


Figure 1. Definition of costs and values in the traffic data. Costs are defined as the distance from the source and the destination, with the assumption that knowing these location conveys privacy-sensitive information. The costs are maximal at the source and at the destination, and decrease in between. Values are defined as the difference between the current speed and the average past speed, with the assumption that a sudden change of speed conveys information about traffic congestion.

3. Evaluation of Contribution Strategies

Following the implementation of the two smart-city scenarios which rely on user data contributions described in the previous section, we now present three commonly used algorithms in smart city applications, which we implement in both of these scenarios, in order to evaluate them with respect to human-centred measures. We first describe in detail the implementation of the algorithms on which contribution strategies are based, and then describe the measures on which the contribution strategies are evaluated.

Broadly, algorithms for optimisation in smart cities can be classified into centralized and distributed algorithms. Centralized contribution strategies translate in the real world into optimization algorithms that are implemented by a service provider, e.g., load balancing in smart grids, where contributions are decided at the central level and users have no decisional power. In contrast, distributed strategies give the choice to the users allowing them to act independently of others and reflecting their personal privacy concerns. These distributed strategies model human decision-making by evaluating the trade-off between the benefits from accessing the service and an estimation of the privacy cost for the user. In reality, human decision is more complex than this, e.g., the value of privacy expressed by people does not necessarily reflect in their actions, the so-called privacy paradox [24]. The purpose of these contribution strategies is not to model accurately human behavior, but instead to evaluate whether distributed decision-making, possibly performed by the users themselves, that optimizes for the individual utility would compare against centralized decision making that optimizes for the global utility.

3.1. Algorithms

This section describes the contribution strategy algorithms chosen for the analysis in our framework. The criteria for the choice of algorithms is their diffusion and application to smart city scenarios. Enough background and details on the algorithms are given to provide the basis for understanding the specificity of our implementation, which are then described together with their parameters and the evaluation metrics. The contribution of this paper is to compare these algorithms on several smart city application scenarios and verify that trade-offs between evaluation measures depend on the algorithms and are independent of the application scenario. Our algorithm choice favored well-established and general-purpose algorithms as opposed to algorithms with state-of-the-art

efficiency, as the goal is to highlight trade-offs between measures over different scenarios. The review and comparison of scenario-specific algorithms are not aligned with the goals of this paper and are therefore out of scope.

3.1.1. Centralized Algorithms: Optimisation

Centralized or top-down algorithms rely on a central optimizer that satisfies the public good while minimizing the cost of contribution. This problem can be modeled as the well-known Knapsack problem: minimize $\sum_{i=1}^n c_i * a_i$ subject to $\sum_{i=1}^n v_i * a_i > \tau$, where the weight of items is given by the contribution value and the value of each item is defined as the inverse of the cost. We chose a customized “fully polynomial time approximation scheme” (FPTAS) that reaches the knapsack constraints from above, instead of from below, such that the threshold can be met.

3.1.2. Localized Algorithms

Decentralized algorithms distribute decision-making at the local level and allow communication between agents for coordination [25] or learning [26]. In this paper, we focus on *localized* algorithms, a type of decentralized algorithms that operate only on local knowledge, without assuming the availability of specialized hardware for communication [27].

Aspiration Learning

Aspiration learning is a learning algorithm that is specifically tailored for coordination games [28]. Agents contribution does not consider the current context, e.g., current value or cost; it is instead based on the agent’s *aspiration value*, which determines how satisfied the agent is with the status quo, given its previous experiences in the game.

At every turn, the agent updates its aspiration level $\rho_i(t)$ and chooses an action $\alpha(t)$. The action for the current turn will be the same action executed in the past turn if the current reward exceeds the aspiration level; otherwise, a random action will be selected with probability

$$p_i(t) = 1 - \max(h, 1 + c * [u_i(\alpha(t)) - \rho_i(t)]),$$

where $u_i(\alpha(t))$ is the reward obtained for executing action $\alpha(t)$. The aspiration level is updated according to the following formula:

$$\rho_i(t + 1) = \rho_i(t) + \epsilon[u_i(\alpha(t)) - \rho_i(t)] + r_i(t),$$

where $r_i(t)$ is a noise component and the value of $\rho_i(t + 1)$ is bound between $\bar{\rho}$ and $\underline{\rho}$.

The algorithm is implemented following the details contained in the original paper, with the following choice of parameters: $\lambda = 10^{-4}$; $c = 0.05$; $h = 0.01$; $\zeta = 0.05$; $\epsilon = 10^{-4}$; $\bar{\rho} = 1.1$; $\underline{\rho} = -1.1$.

Q-Learning

Q-Learning is a model-free unsupervised reinforcement learning approach that considers both the history and the current context in the decision [29]. The reward in our implementation is defined as $R_q = S(q, \tau) - c_i \times a_i$, i.e., a component related to the success of the public good, driven by the collective action, and a negative component driven by the individual cost and individual action. The implementation of Q-Learning relies on TensorFlow, the schema of the network is detailed in Figure 2. Q-Learning is influenced by the choice of parameters α , the learning rate, and γ , the learning discount. In our simulations, γ is set to 0, as the future states are independent from the choice of action: the scenario can be classified as contextual bandits, where the reward of an action depends on the current state, but the chosen action does not reflect the next state. A parameter sweep on the value of α showed no impact of the learning rate on the performance of the algorithm; hence, the default value of $\alpha = 0.001$ has been chosen for the experiments.

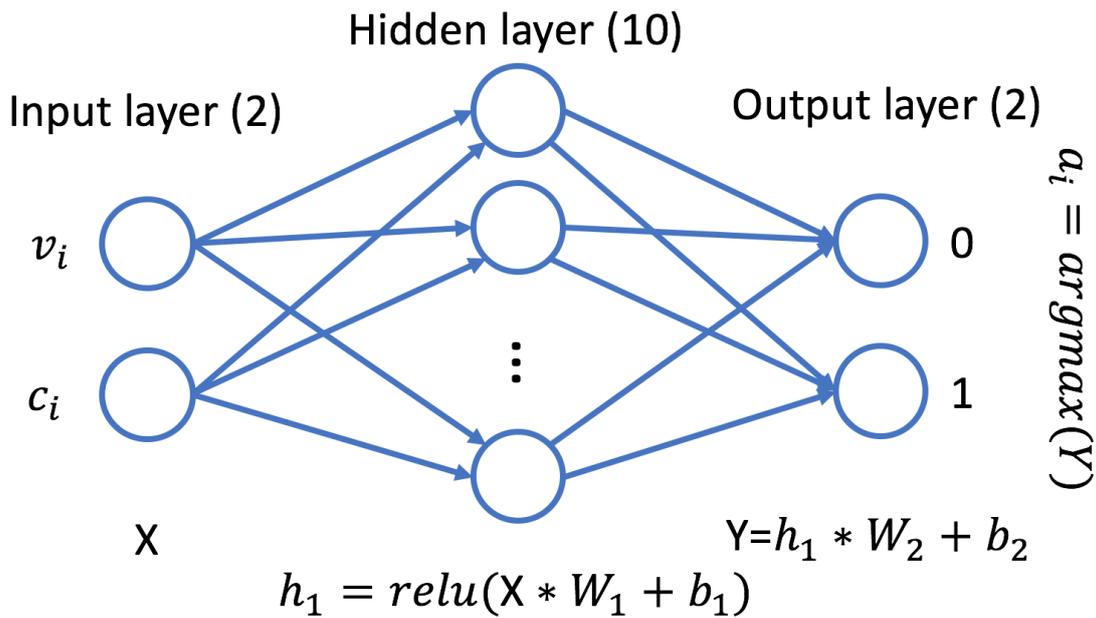


Figure 2. Diagram of the neural network on which the Q-Learning mechanism is based. The inputs are the current value and cost of contribution and the outputs are the weights given to the possible actions. The chosen action is determined by applying argmax to the output of the network.

A disadvantage of reinforcement learning is its sensitivity to initial conditions—for example, a multi-agent learning process might converge to an inefficient equilibrium where nobody contributes. In order to make this outcome less likely, agents are pre-trained to prefer contribution in order to bias the initial exploration period. Pre-training is a reasonable solution as it can be performed during device manufacturing and its effect on the behavior of agents fades off quickly as agents start learning.

3.2. Measures

The performance of different contribution strategies is quantified with the following measures (see Table 3):

- (a) Success rate: The fraction of the threshold that has been covered by contributions, or 1 if the total contribution is higher than the threshold.
- (b) Efficiency: The ratio between the requirement and the sum of contributions. Efficiency is 1 if the sum of contributions is equivalent to the threshold, e.g., all agents contribute 1, it is lower than 1 if the sum of contributions is larger than the threshold.
- (c) Social welfare: The average reward, sum of a constant benefit and a constant negative cost (only for contributors).
- (d) Privacy: the fraction of agents that did not disclose private information during the current timestep.
- (e) Fairness: The Gini coefficient represents the fairness of the current round of contributions (0 is total equality). It is computed for each timestep, as the Gini coefficient of the values that agents contributed in that timestep, and aggregated across time and repetitions.
- (f) Fairness over time: It represents the fairness of the history of contributions. It measures the Gini coefficient computed across agents on the individual histories of contribution, from the start of the simulation to the current timestep.

The definition of privacy can be made scenario-specific by adopting an appropriate privacy measure, e.g., K-anonymity, differential privacy [30].

Table 3. Measures used during evaluation.

	Measure	Definition
(a)	Success	$\Sigma^t = \min(1, q^t / \tau^t)$
(b)	Efficiency	$E^t = \tau^t / q^t$ if $\tau^t \leq q^t$ else 0
(c)	Social welfare	$W^t = \frac{1}{T} \sum_{i=1}^T U_i^t$
(d)	Privacy	$P^t = 1 - A_+^t / n$
(e)	Fairness	$F^t = \frac{1}{n} \left(n + 1 - 2 \left(\frac{\sum_{i=1}^n (n+1-i)y_i}{\sum_{i=1}^n y_i} \right) \right)$ where $y_i = v_i^t$
(f)	Fairness, over time	$G = \frac{1}{n} \left(n + 1 - 2 \left(\frac{\sum_{i=1}^n (n+1-i)y_i}{\sum_{i=1}^n y_i} \right) \right)$ where $y_i = \sum_{t=1}^T v_i^t$

3.3. Model Design

This section describes the design of the components in the simulation framework and their interaction.

Data Generation

The data generation function contains the logic to generate a new data point for each user and to compute the respective contribution values, contribution costs and common good threshold. The generation function depends on the scenario that is modeled, and might either rely on existing data or generate artificial data randomly. This function is called at the beginning of each timestep in order to provide a new value and cost to each agent, as well as system-wide parameters such as the common good threshold.

Decision Function

The decision function maps an input, reflecting the current state of the system, to one of many available actions. The decision function can consider past experience, current context or answer reactively to the current input, depending on the algorithm that is implemented. Some algorithms can improve their performance by learning from feedback about the effect of the action on the environment.

Evaluation Function

The evaluation function classifies the current state of the system with respect to a number of measures, which are used to evaluate the behavior of the system. It does so by accessing the state of the system and of all individual agents; this global knowledge of the system allows it to compare the performance obtained empirically with the maximum theoretical performance.

Reward Function

The reward function evaluates the effect that the action of an individual agent had on the state of the system and computes a proportionate reward, which might be used to condition the future action of that agent such that they align to some global goal.

Supervisor

The supervisor is the main component of the model, which is in charge of coordinating each timestep by updating the state of the system and evaluating the outcome of the set of actions being selected by the population. The supervisor implements the communication system for agents, on which measurements, actions and feedback are exchanged. Another important task of the supervisor is that of evaluating and logging the state of the system for successive analysis. In case of a centralized configuration, the supervisor is also in charge to instruct agents on the actions they have to perform.

Agent

The agent class implements generic functions that are required for the working of the simulation environment. The agent class is modular as it can support different logic of operation, e.g., different ways of generating data and different algorithms for decision-making. Agents obtain a new set of input values at the beginning of each timestep, from the respective data generation function. The agent's own decision function selects one of the available actions for execution, based on the current input and the state of the system: in case of a localized configuration, the agents choose autonomously what action to perform, given the current input values and, possibly, the past experience of the agent. Experience is updated and accumulated by the feedback obtained from the Supervisor. In case of a centralized configuration, actions are decided at the central level and executed by the local decision function.

3.4. Parameters of the Model

Table 4 presents an overview of the parameters of the simulation environment. Parameters that have been kept constant across experiments are represented by a number, while parameters that varied across or within experiments, e.g., parameter sweep, are represented as tuples.

Table 4. List of parameters in the simulation environment. Tuples represent the values of parameters that have been tested during the parameter sweep.

Parameter	Values
Simulation length	5000
Repetitions	20
Population size	(5,7,15,20,30,50)
Scenario: Smart grid	
Minimum value	1
Maximum value	(2,3,4,6,8,10)
Scenario: Participatory sensing	
Minimum value	1
Maximum value	(3,4,6,8,10)
Public Good Threshold	0.8
Q-Learning	
Alpha	0.001
Gamma	0.0
Dropout probability	0.8
Hidden layer size	10
Aspiration Learning	
λ	10^{-4}
c	0.05
h	0.01
ζ	0.05
ϵ	10^{-4}
$\bar{\rho}$	1.1
$\underline{\rho}$	-1.1

4. Results and Analysis

This section presents the results of computational experiments. Trade-offs are evaluated between the three contribution strategies presented in Section 3.1 and two baselines: "full" where all users contribute, and "random" where users have 50% probability of contributing at each round (50% chance of contribution does not imply 50% chance of success because each contribution is on average greater than 1.).

All experiments are performed in a Python simulation framework and run on the computing cluster of ETH Zü. Results presented in this paper represent the average state of 20 simulations after 5000 timesteps, and error bars represent the confidence intervals at 95%. The choice of cost, value and public good values is described in Section 2.1.

Full results are shown in Figure 3 and present the comparisons of the contribution strategies by the six measures discussed in Section 3.2. The plots show aggregated results across a range of population sizes, from 5 to 50 users. The decision to aggregate the visualization is motivated by the independence of the results on the size of the population, i.e., the number of sensors, caused by the choice of the public good threshold to be proportional to the size of the population. For completeness, Figures A1 and A2 in the Appendix illustrate the (lack of) variation of these results with respect to the population size and action-state space size. If the threshold would be constant, an increase in the number of sensors, i.e., potential contributors, would make the creation of the public good easier and hence affect all measures.

Success Rate

Concerning success rate (Figure 3a), full contribution and centralized optimization always succeed, i.e., are always able to provide the services, while Q-Learning does not guarantee success and fails in about 2–3% of the cases.

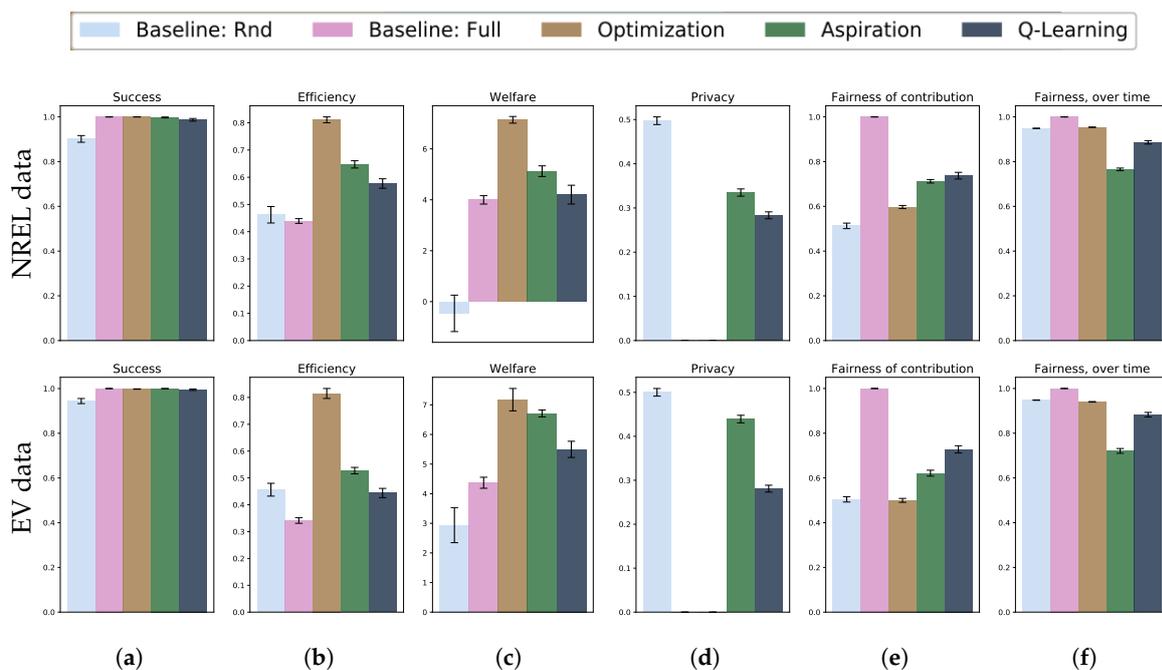


Figure 3. Comparison of contribution strategies. The plots show trade-offs between centralized optimization, aspiration learning and Q-Learning in terms of success rate (a), efficiency (b), social welfare (c), privacy (d) and fairness (e,f). Optimization offers the highest success rate and efficiency, while aspiration learning and Q-Learning offer higher privacy. Q-Learning offers higher fairness than aspiration learning on both measures.

Efficiency

Concerning efficiency (Figure 3b), centralized optimization is the most efficient solution, as it finds the subset of users whose contribution satisfies the requirement at the lowest cost (although the chosen approximation algorithm does not guarantee to find the global minimum). Optimization is not successful if the total requirement is larger than the sum of contributions from all agents, but the experiments are generated to be successful if all users contribute. Efficiency measures how close the total contribution approaches the needs of the system; hence, the baseline in which all agents

contribute will have the lowest possible efficiency. Efficiency increases with the size of the population as a higher number of possible solutions—combinations of individual contributions—makes it more likely to find an efficient solution.

Social Welfare

Similarly, optimization scores the highest value of social welfare (see Figure 3c), while localized strategies reach a welfare around 30% lower. Social welfare is the difference between the rewards from the public good and the costs of contributions, so a negative value indicates that costs are higher than gains. Differently from the previous result, the performance of aspiration learning is equivalent to that of Q-Learning, as opposed to that of optimization.

Privacy

Concerning privacy (Figure 3d), neither centralized optimization nor full contribution grant any privacy to the users, as they require full knowledge about the state of the system. Conversely, localized contribution strategies allow a fraction of the user to keep their data private. This fraction increases with the population size for aspiration learning, while Q-Learning trades lower privacy off for higher fairness. Random contribution offers the highest privacy—around 50% of users (as each user has a 50% chance of contributing)—at the expenses of other measures, e.g., success rate and efficiency.

Fairness

Fairness is measured in two ways: “fairness of contributions” compares the actions of all agents at the current time t (Figure 3e), while “fairness of contribution over time” considers the histories of contribution up to time t (Figure 3f). Full contribution requires all users to contribute; thus, it trades perfect fairness off for other measures such as efficiency. Optimization offers low fairness because it considers only the current state, and users in certain states, e.g., with high values, are more likely to contribute than others.

Fairness over Time

Conversely, optimization offers high fairness over time because agents are randomly assigned to states; hence, the chance of being in any state is over time the same. This result might not hold if states are not randomly assigned, e.g., some users are more likely to obtain high values/costs than others. Aspiration learning bases decisions only on the history of decisions. This leads to higher fairness, as contributions are independent of the state, and to lower fairness over time, caused by individual differences in training that accumulate over time. These values decrease with the population size, while other contribution strategies are not affected by this parameter. Q-Learning scores high values in both measures as it considers both the current context and the history of actions. This allows agents to learn similar behaviors by interacting with one another.

In conclusion, the simulation results highlight a set of trade-offs between the measures of efficiency, privacy and fairness, which are summarized in Table 5. These trade-offs are consistent across scenarios. This suggests that they depend on the characteristics of the algorithms; hence, these results can be the basis for recommending to system designers appropriate contributions strategies for a given cyberphysical system and application domain. Specifically, centralized optimization assures the success of the service and provides high efficiency; it is hence appropriate for mission-critical services for which computation and network constraints are not an issue, e.g., measuring the current load on the smart grid to prevent outages. Localized strategies trade efficiency and reliability off for higher privacy and fairness; therefore, they are best suited for applications where privacy concerns might reduce user adoption, e.g., participatory sensing. Among them, Q-Learning offers the highest fairness; hence, it is ideal for applications where fair access to the service is desirable, e.g., charging of electric vehicles.

Table 5. Summary of advantages and disadvantages of contribution strategies.

Type	Algorithm	Pros	Cons
Baseline	Full	Success	Efficiency
Baseline	Random	Privacy	Success
Centralized	Knapsack	Success, Efficiency	Privacy, Fairness
Localized	Aspiration	Privacy	Fairness over time
Localized	Q-Learning	Fairness (over time)	Efficiency

5. Related Work

The concept of smart city is broad and difficult to define [31]. Many application scenarios can be placed under the umbrella of smart cities [1]. In this work, we focus on a few popular application scenarios: participatory sensing [32,33], with the examples of traffic control [34,35] and traffic congestion maps [36], and smart grid [37], with the example of charging Electric Vehicles (EVs) [38]. The choice of these applications is motivated by the interest shown by the smart city and privacy communities: Privacy is an important component of the smart city [5] that is particularly well studied in the literature, with numerous privacy-preserving solutions being developed for the application scenarios of traffic [39], participatory sensing [20], and the smart grid [40]. Regarding methodology, a variety of solutions have been applied to these application scenarios, ranging from optimization [41,42] to machine learning [43,44].

This work quantifies trade-offs between privacy and fairness. Privacy is a well-studied topic in smart grids [6,45,46] and in participatory sensing [20,47,48], with location privacy being of special interest [49–51]. Fairness has been investigated in demand response programs of smart grids [52,53] and in participatory sensing [54], and only recently fairness has been studied independently of the application scenario [7,55,56]. However, only limited literature considers trade-offs between fairness and privacy [57].

These application scenarios have some similarities, e.g., routing protocols are used in both smart grids and traffic control [58]. The most relevant similarity between application scenarios is that both deal with producing a service from user-contributed data. This makes them suitable to be modeled by a Voluntary Contribution Game (VCG) [8], which has been successfully applied to smart grids [12].

In our modeling, we assume that a demand-response algorithm which optimized the use of available energy is in place [59]. This allows us to abstract away details about the smart grid network and focus on individual production and consumption. We also require as the demand-side management algorithm to be privacy preserving, e.g., [60,61]; otherwise, successive privacy considerations would not be meaningful. We have the same requirement for a traffic control algorithm [35] that is privacy preserving [62].

This work analyzes different centralized and decentralized control algorithms. Centralized algorithms are preferable over decentralized ones when it comes to accuracy, as they can take decisions considering the global state of the system. Nevertheless, there are drawbacks that make decentralized algorithm preferable in some particular conditions. The most obvious limitation for centralized algorithms is the intractability of nondeterministic polynomial time (NP) complete problems, e.g., the charging-scheduling problem [63]. Another limitation relates to the communication footprint, which is especially a problem for large wireless sensor networks, as the traffic coming from numerous sensors might saturate the communication medium [64]. A more appropriate solution for this scenario is a decentralized approach that relies on short range communication, or localized algorithms that allow no intra-agent communication and therefore have constant communication footprint. Another advantage of a distributed/localized algorithm over centralized optimization is the increased fault tolerance due to the absence of a single point of failure, i.e., the central optimizer, as well as an increased scalability to large population sizes. Other non-technical constraints might play a role in making decentralized more desirable, e.g., privacy [4] or fairness [53]. Finally, localized algorithms empower

the users to choose autonomously about their own actions and to keep ownership of data. A more in-depth comparison of centralized and decentralized algorithms is provided in [26].

Reinforcement learning (RL) algorithms are frequently used in situations where both traditional AI-based approaches, such as planning, or supervising learning are not practical or scalable [65], and where a model of the environment is costly, or even impossible, to obtain. RL has been applied extensively to both smart city applications under consideration, e.g., Ref. [66] uses RL to reduce energy costs, and [67] uses RL in participatory sensing. The main advantage of using RL over other decentralized optimization approaches, e.g., evolutionary computing or Monte Carlo tree search, as evaluated in [26], is that it operates by selecting an action based only on the current state, i.e., current set of conditions, rather than pre-calculating longer-term (e.g., daily) schedules/actions; therefore, it does not need a prediction of the future conditions, and it does not need to re-calculate the schedules if underlying conditions change. Another advantage is that a learning algorithm obtains rewards by interacting with the environment, so it can work in environments in which the relation between state, actions and rewards is not known a priori—for example, because it depends on the collective action of a population. In addition, as Q-learning is a widely established technique, numerous extensions exist that enable its implementation in numerous varieties based on domain requirements, e.g., multi-goal implementations using W-Learning [43,44] and collaborative implementations using Distributed W-Learning [68].

Table 6 compares the current work with the most relevant literature along three dimensions: the kind of target measure, the organization of the system and the application scenario. The table highlights the uniqueness of the current work, as no other work generalizes over application scenarios, while considering trade-offs between privacy and fairness, as well as trade-offs between different system organizations, which is the main contribution of this paper.

Table 6. Comparison of related literature.

Paper	Measure		Organization			Application Scenario	
	Privacy	Fairness	Centralized	Decentralized	Smart Grid	Traffic Management	Participatory Sensing
Current work	✓	✓	✓	✓	✓	✓	✓
[39]	✓	✗	✗	✓	✗	✓	✗
[67]	✓	✗	✗	✓	✗	✗	✓
[62]	✓	✗	✓	✗	✗	✓	✗
[6,40,45,46,61]	✓	✗	✓	✓	✓	✗	✗
[20,47–51]	✓	✗	✓	✓	✗	✗	✓
[52,53]	✗	✓	✓	✗	✓	✗	✗
[54]	✗	✓	✓	✗	✗	✗	✓
[55]	✓	✓	✗	✓	✗	✗	✓
[7]	✗	✓	✗	✓	✗	✓	✓
[56]	✗	✓	✓	✓	✓	✓	✓
[57]	✓	✓	✓	✗	✓	✗	✗
[4]	✓	✗	✗	✓	✓	✓	✓

6. Conclusions and Future Work

The goal of this paper is to provide a tool for understanding user participation in established smart city services. In order to do this, we first proposed a scenario-independent design principle based on voluntary user contribution and a simulation framework for smart city applications. The applicability of this framework was verified using real-world data from the two application scenarios of traffic congestion monitoring and electric vehicle charging. Secondly, we used this framework to measure the effects, along various dimensions, of citizens' participation.

Voluntary contributions empower users to control the ownership of their resource, e.g., by contributing data towards a service, independently of the type of resource and its use. Different contribution strategies produce trade-offs along measures such as efficiency, privacy and fairness, which we quantified in both implemented scenarios.

Results suggest that such trade-offs depend on characteristics of the algorithms and not on characteristics of the scenarios. Therefore, they can be used as implementation recommendations to service providers and system designers about the choice of contribution strategies, regardless of the specific smart city domain. Specifically, centralized optimization is found to offer the highest efficiency of contributions and to guarantee the success of the application. It is hence recommendable for mission-critical services that require high availability, e.g., load management on a smart grid. Localized strategies are, on the other hand, less efficient and might cause the system to fail with low probability, but compensate for this by improving on privacy and fairness of contribution. Localized strategies are hence more suitable for privacy-sensitive applications that require user adoption, such as participatory sensing, or applications where fairness in redistribution is crucial, e.g., charging of electric vehicles.

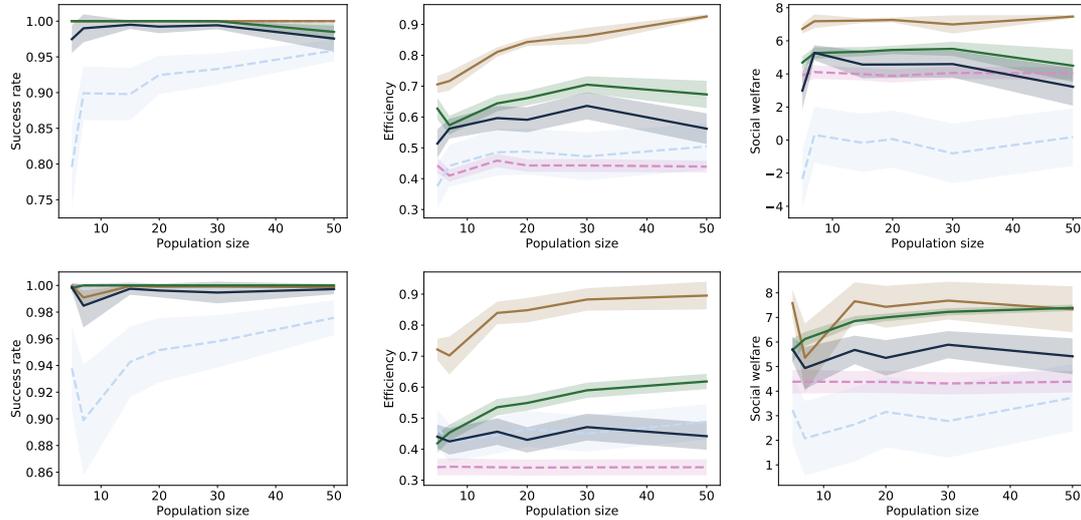
Modeling of scenarios with different cost and value characteristics is left to future work—for example, “negative” contribution in traffic congestion, where users contribute to the public good by choosing a longer route over the shortest but congested route [69]. Incentive mechanisms help increase user participation [70] but require quantifying privacy [24], which could also be addressed in future work. The current work considers only localized algorithms, i.e., it assumed no communication between user devices. Relaxing this limitation is a worthy avenue for future work. Communication between user devices would introduce privacy concerns due to data exchange but would allow analysis of new classes of algorithms, such as decentralized optimization. Finally, it would be beneficial to verify the proposed design principle based on public good theory in other smart city domains, in order to further evaluate its generality and suggest as a complete blueprint for smart city applications. Other potential application areas could, for example, include pollution or noise monitoring and real-time public transport information, or any other similar data-rich application relying on user contributions in which system performance, fairness and user privacy are of concern.

Author Contributions: The authors contributed to the research as follows: S.B., I.D. and C.M.J. contributed to the conceptualization and methodology, S.B. and R.S. contributed to software creation and validation. S.B. and I.D. contributed to data curation and visualization, S.B., I.D. and C.M.J. contributed to writing the manuscript.

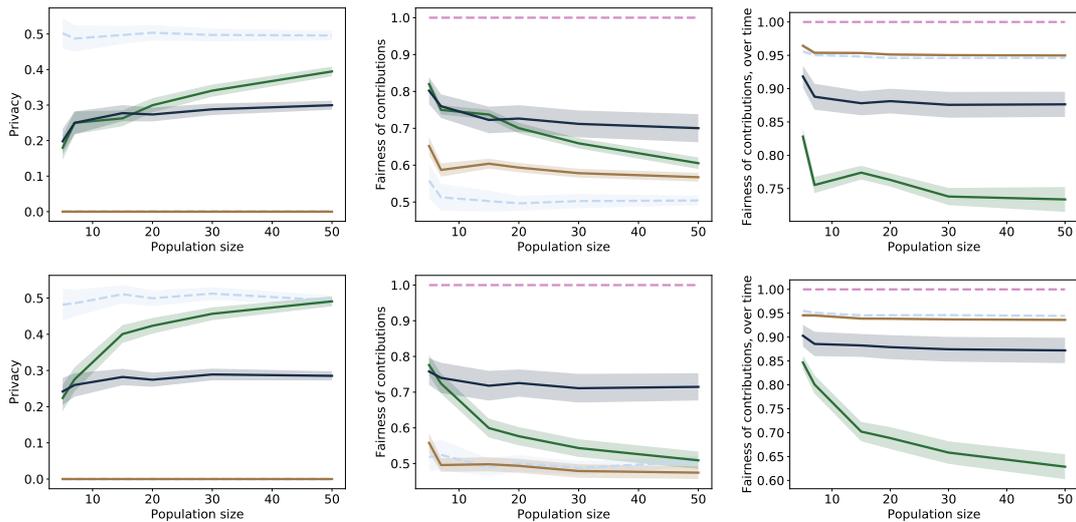
Funding: Stefano Bennati acknowledges support by the European Commission through the ERC Advanced Investigator Grant ‘Momentum’ (Grant No. 324247).

Conflicts of Interest: The authors declare no conflict of interest. The founding sponsors had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Appendix

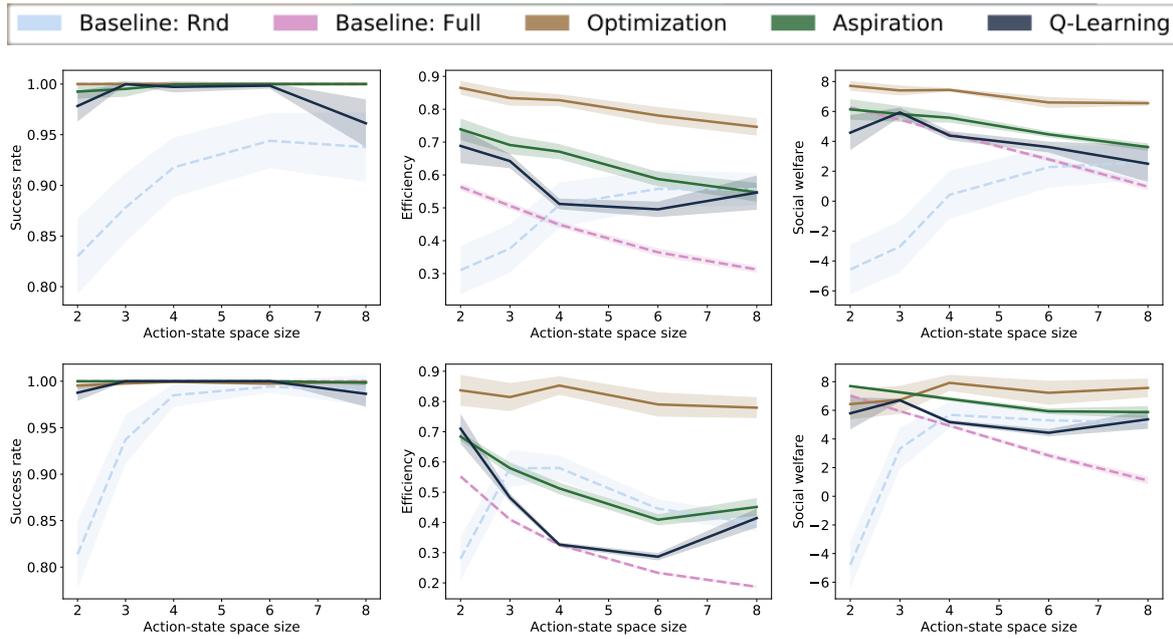


(a) Ratio of success in service creation; (b) Efficiency of total contributions; (c) Social welfare of the population;

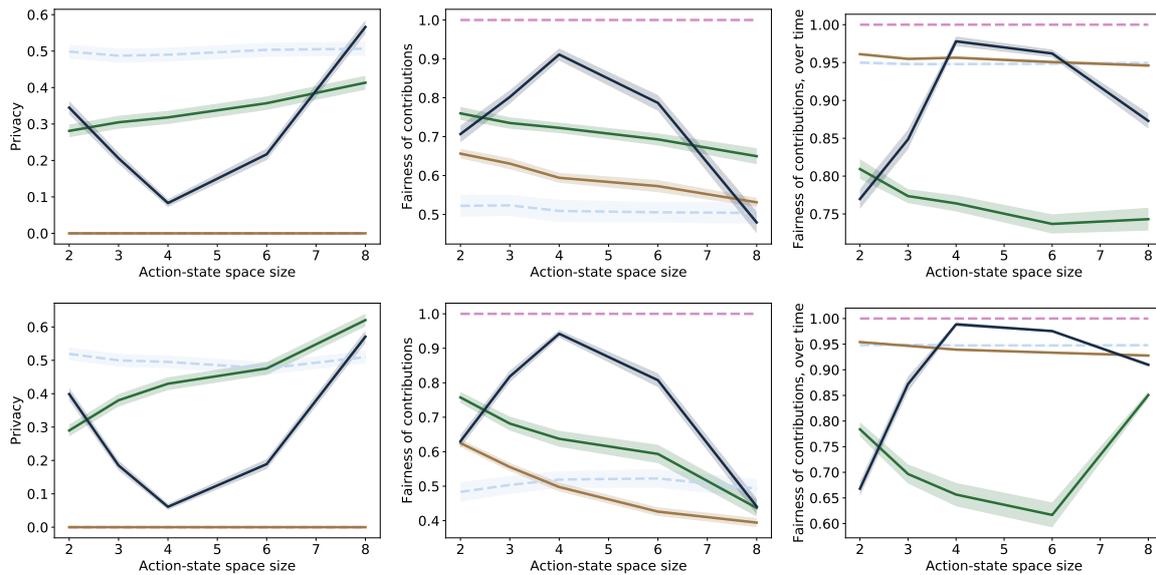


(d) Average privacy of users; (e) Average fairness of contributions; (f) Average fairness over time.

Figure A1. Dependence of results on the population size (x -axis). Dashed lines represent baselines, solid lines represent contribution strategies. The plots show that most measures are independent of the population size. This result is driven by the assumption that the public good threshold is proportional to the size of the population. If this threshold were constant, successfully generating the public good would become easier with an increase in the number of agents/contributors. Aspiration learning is the algorithm with the highest dependency on population size, as a larger population size leads to higher efficiency and privacy, as well as lower fairness. This result can depend on the lack of contextual information in aspiration learning, in contrast to Q-Learning, which could lead to reduced performance for smaller and less flexible populations.



(a) Ratio of success in service creation; (b) Efficiency of total contributions; (c) Social welfare of the population;



(d) Average privacy of users; (e) Average fairness of contributions; (f) Average fairness over time.

Figure A2. Dependence of results on the range of values and costs (x -axis). Dashed lines represent baselines, solid lines represent contribution strategies. The plots show a trend for which all contribution strategies tend to perform worse for larger action-state spaces. Q-Learning is the measure that is most affected by the size of the action-state, in particular, it can be seen that its performance varies significantly on most measures for intermediate sizes. This result is driven by the fixed capacity of the neural network, which makes it more suitable for a certain degree of complexity in the input; the performance of a Q-Learning contribution strategy could be adapted to a given environment by changing the capacity of the network to a more suitable value.

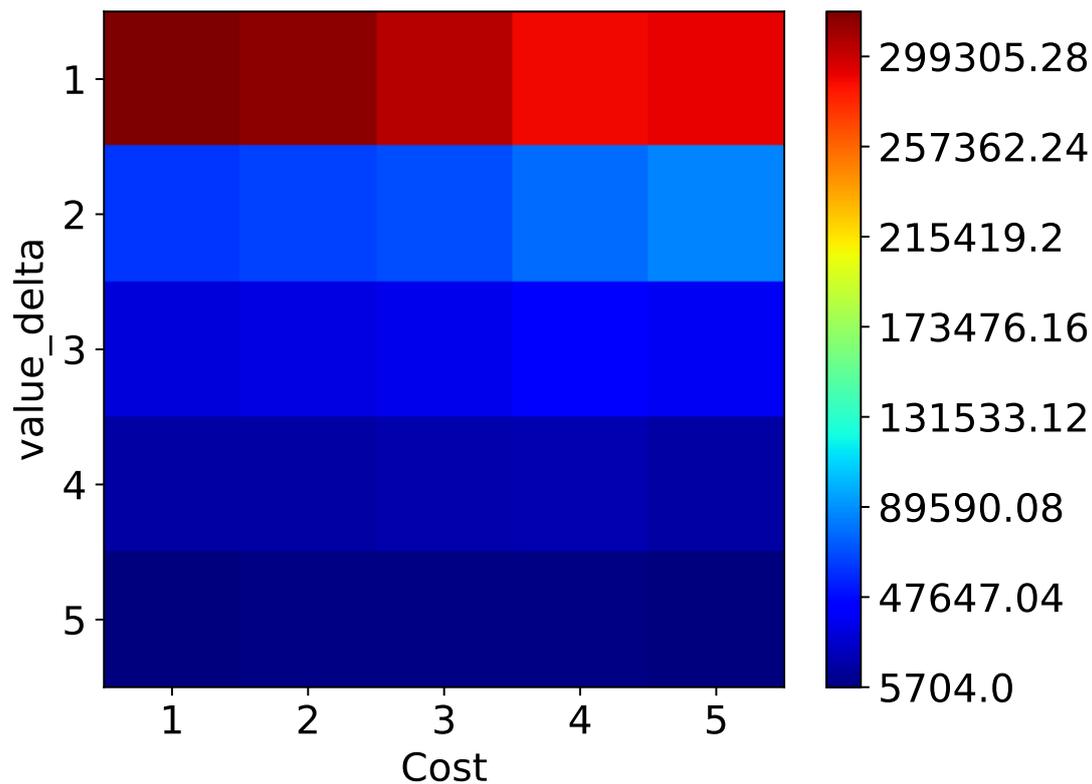


Figure A3. Average frequency of value/cost combinations in the NREL dataset. Data points with a combination of low value and high cost are more frequent in the dataset than data points with high value and high cost. Values are normalized by the maximum speed of the trip and binarized in N integer values.

References

- Gaur, A.; Scotney, B.; Parr, G.; McClean, S. Smart city architecture and its applications based on IoT. *Procedia Comput. Sci.* **2015**, *52*, 1089–1094. [CrossRef]
- Al Nuaimi, E.; Al Neyadi, H.; Mohamed, N.; Al-Jaroodi, J. Applications of big data to smart cities. *J. Int. Serv. Appl.* **2015**, *6*, 25. [CrossRef]
- Hashem, I.A.T.; Chang, V.; Anuar, N.B.; Adewole, K.; Yaqoob, I.; Gani, A.; Ahmed, E.; Chiroma, H. The role of big data in smart city. *Int. J. Inf. Manag.* **2016**, *36*, 748–758. [CrossRef]
- Bennati, S.; Pournaras, E. Privacy-enhancing Aggregation of Internet of Things Data via Sensors Grouping. *Sustain. Cities Soc.* **2018**, *39*, 387–400 [CrossRef]
- Eckhoff, D.; Wagner, I. Privacy in the Smart City—Applications, Technologies, Challenges and Solutions. *IEEE Commun. Surv. Tutor.* **2017**, *20*, 489–516 [CrossRef]
- Finster, S.; Baumgart, I. Privacy-Aware Smart Metering: A Survey. *IEEE Commun. Surv. Tutor.* **2015**, *17*, 1088–1101. [CrossRef]
- Tham, C.K.; Luo, T. Fairness and social welfare in service allocation schemes for participatory sensing. *Comput. Netw.* **2014**, *73*, 58–71. [CrossRef]
- Bagnoli, M.; Lipman, B.L. Provision of Public Goods: Fully Implementing the Core through Private Contributions. *Rev. Econ. Stud.* **1989**, *56*, 583–601. [CrossRef]
- Bennati, S. The Simulation Framework. 2018. Available online: <https://github.com/bennati/EnergyVCG> (accessed on 26 October 2018).
- Duan, L.; Kubo, T.; Sugiyama, K.; Huang, J.; Hasegawa, T.; Walrand, J. Motivating smartphone collaboration in data acquisition and distributed computing. *IEEE Trans. Mob. Comput.* **2014**, *13*, 2320–2333. [CrossRef]
- Wang, X.; Shao, C.; Wang, X.; Du, C. Survey of electric vehicle charging load and dispatch control strategies. *Proc. CSEE* **2013**, *33*, 1–10.

12. Wolsink, M. The research agenda on social acceptance of distributed generation in smart grids: Renewable as common pool resources. *Renew. Sustain. Energy Rev.* **2012**, *16*, 822–835. [[CrossRef](#)]
13. Gellings, C.; Chamberlin, J. *Demand-Side Management: Concepts and Methods*; OSTI: Oak Ridge, TN, USA, 1987.
14. Saha, A.; Kuzlu, M.; Khamphanchai, W.; Pipattanasomporn, M.; Rahman, S.; Elma, O.; Selamogullari, U.S.; Uzunoglu, M.; Yagcitekcin, B. A home energy management algorithm in a smart house integrated with renewable energy. *IEEE Innov. Smart Grid Technol. Conf. Eur.* **2015**. [[CrossRef](#)]
15. Eirgrid Group. System Demand Data [Dataset]. 2018. Available online: <http://smartgriddashboard.eirgrid.com/#all/demand> (accessed on 26 October 2018).
16. Eirgrid Group. Wind Generation Data [Dataset]. 2018. Available online: <http://smartgriddashboard.eirgrid.com/#all/wind> (accessed on 26 October 2018).
17. Scellato, S.; Fortuna, L.; Frasca, M.; Gómez-Gardeñes, J.; Latora, V. Traffic optimization in transport networks based on local routing. *Eur. Phys. J. B* **2010**, *73*, 303–308. [[CrossRef](#)]
18. Brown, J.W.S.; Ohrimenko, O.; Tamassia, R. Haze: Privacy-preserving real-time traffic statistics. In Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems—SIGSPATIAL'13, New York, NY, USA, 5–8 November 2013; pp. 530–533.
19. Vuran, M.; Akan, Ö.; Akyildiz, I. Spatio-temporal correlation: Theory and applications for wireless sensor networks. *Comput. Netw.* **2004**, *45*, 245–259. [[CrossRef](#)]
20. Christin, D. Privacy in mobile participatory sensing: Current trends and future challenges. *J. Syst. Softw.* **2015**, 1–12. [[CrossRef](#)]
21. Tsoukaneri, G.; Theodorakopoulos, G.; Leather, H.; Marina, M.K. On the inference of user paths from anonymized mobility data. In Proceedings of the 2016 IEEE European Symposium on Security and Privacy, Saarbrücken, Germany, 21–24 March 2016; pp. 199–213.
22. De Montjoye, Y.A.; Hidalgo, C.A.; Verleysen, M.; Blondel, V.D. Unique in the Crowd: The privacy bounds of human mobility. *Sci. Rep.* **2013**, *3*. [[CrossRef](#)] [[PubMed](#)]
23. NREL [Dataset]. Transportation Secure Data Center. 2015. Available online: www.nrel.gov/tsdc (accessed on 26 October 2018).
24. Norberg, P.A.; Horne, D.R.; Horne, D.A. The privacy paradox: Personal information disclosure intentions versus behaviors. *J. Consum. Aff.* **2007**, *41*, 100–126. [[CrossRef](#)]
25. Kennedy, J.; Eberhart, R. Particle Swarm Optimization. *IEEE Int. Conf. Neural Netw.* **1995**, *4*, 1942–1948.
26. Dusparic, I.; Taylor, A.; Marinescu, A.; Golpayegani, F.; Clarke, S. Residential demand response: Experimental evaluation and comparison of self-organizing techniques. *Renew. Sustain. Energy Rev.* **2017**, *80*, 1528–1536. [[CrossRef](#)]
27. Fadlullah, Z.M.; Fouda, M.M.; Kato, N.; Takeuchi, A.; Iwasaki, N.; Nozaki, Y. Toward intelligent machine-to-machine communications in smart grid. *IEEE Commun. Mag.* **2011**, *49*, 60–65. [[CrossRef](#)]
28. Chasparis, G.C.; Shamma, J.S.; Arapostathis, A. Aspiration learning in coordination games. *Decis. Control* **2010**, 5756–5761. [[CrossRef](#)]
29. Sutton, R.S.; Barto, A.G. *Reinforcement Learning: An Introduction*; MIT press: Cambridge, UK, 1998; Volume 1.
30. Damiani, M.L. Location privacy models in mobile applications: Conceptual view and research directions. *Geoinformatica* **2014**, *18*, 819–842. [[CrossRef](#)]
31. Cocchia, A. Smart and digital city: A systematic literature review. In *Smart city*; Springer: Berlin, Germany, 2014; pp. 13–43.
32. Burke, J.A.; Estrin, D.; Hansen, M.; Parker, A.; Ramanathan, N.; Reddy, S.; Srivastava, M.B. Participatory Sensing. Center for Embedded Network Sensing. 2006. Available online: <https://escholarship.org/uc/item/19h777qd> (accessed on 26 October 2018).
33. Aggarwal, C.C.; Abdelzaher, T. Social sensing. In *Managing and Mining Sensor Data*; Springer: Berlin, Germany, 2013; pp. 237–297.
34. Li, L.; Wen, D.; Yao, D. A survey of traffic control with vehicular communications. *IEEE Trans. Intell. Transp. Syst.* **2014**, *15*, 425–432. [[CrossRef](#)]
35. Baskar, L.D.; De Schutter, B.; Hellendoorn, J.; Papp, Z. Traffic control and intelligent vehicle highway systems: A survey. *IET Intell. Transp. Syst.* **2011**, *5*, 38–52. [[CrossRef](#)]
36. Taylor, M.A.; Woolley, J.E.; Zito, R. Integration of the global positioning system and geographical information systems for traffic congestion studies. *Transp. Res. Part C Emerg. Technol.* **2000**, *8*, 257–285. [[CrossRef](#)]

37. Fang, X.; Misra, S.; Xue, G.; Yang, D. Smart grid—The new and improved power grid: A survey. *IEEE Commun. Surv. Tutor.* **2012**, *14*, 944–980. [[CrossRef](#)]
38. Mwasilu, F.; Justo, J.J.; Kim, E.K.; Do, T.D.; Jung, J.W. Electric vehicles and smart grid interaction: A review on vehicle to grid and renewable energy sources integration. *Renew. Sustain. Energy Rev.* **2014**, *34*, 501–516. [[CrossRef](#)]
39. Freudiger, J.; Raya, M.; Félegyházi, M.; Papadimitratos, P.; Hubaux, J.P. Mix-zones for location privacy in vehicular networks. In *ACM Workshop on Wireless Networking for Intelligent Transportation Systems (WiN-ITS)*; ACM: Vancouver, BC, Canada, 2007.
40. McDaniel, P.; McLaughlin, S. Security and privacy challenges in the smart grid. *IEEE Secur. Priv.* **2009**, *7*, 75–77. [[CrossRef](#)]
41. Papageorgiou, M.; Diakaki, C.; Dinopoulou, V.; Kotsialos, A.; Wang, Y. Review of road traffic control strategies. *Proc. IEEE* **2003**, *91*, 2043–2067. [[CrossRef](#)]
42. Vardakas, J.S.; Zorba, N.; Verikoukis, C.V. A survey on demand response programs in smart grids: Pricing methods and optimization algorithms. *IEEE Commun. Surv. Tutor.* **2015**, *17*, 152–178. [[CrossRef](#)]
43. Dusparic, I.; Monteil, J.; Cahill, V. Towards autonomic urban traffic control with collaborative multi-policy reinforcement learning. In *Proceedings of the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, Rio de Janeiro, Brazil, 1–4 November 2016; pp. 2065–2070. [[CrossRef](#)]
44. Dusparic, I.; Harris, C.; Marinescu, A.; Cahill, V.; Clarke, S. Multi-agent residential demand response based on load forecasting. *IEEE Conf. Technol. Sustain.* **2013**, 90–96. [[CrossRef](#)]
45. Erkin, Z.; Troncoso-pastoriza, J.R.; Legendijk, R.L.I.L.; Perez-Gonzalez, F.; Pérez-gonzález, F.; Erkin, Z.; Troncoso-pastoriza, J.R.; Systems, S.M. Privacy-preserving data aggregation in smart metering systems: An overview. *IEEE Signal Process. Mag.* **2013**, *30*, 75–86. [[CrossRef](#)]
46. Zeadally, S.; Pathan, A.S.K.; Alcaraz, C.; Badra, M. Towards Privacy Protection in Smart Grid. *Wirel. Pers. Commun.* **2013**, *73*, 23–50. [[CrossRef](#)]
47. Giannetsos, T.; Gisdakis, S.; Papadimitratos, P. Trustworthy people-centric sensing: Privacy, security and user incentives road-map. In *Proceedings of the 2014 13th Annual Mediterranean Ad Hoc Networking Workshop, MED-HOC-NET*, Piran, Slovenia, 2–4 June 2014; pp. 39–46.
48. Li, N.; Zhang, N.; Das, S.K.; Thuraisingham, B. Privacy preservation in wireless sensor networks: A state-of-the-art survey. *Ad Hoc Netw.* **2009**, *7*, 1501–1514. [[CrossRef](#)]
49. Krumm, J. A survey of computational location privacy. *Pers. Ubiquitous Comput.* **2009**, *13*, 391–399. [[CrossRef](#)]
50. Hoh, B.; Gruteser, M.; Xiong, H. Enhancing Security and Privacy in Traffic-Monitoring Systems. *IEEE Pervasive Comput.* **2006**, *5*, 38–46.
51. Hubaux, J.P.; Čapkun, S.; Luo, J. The security and privacy of smart vehicles. *IEEE Secur. Priv.* **2004**, *2*, 49–55. [[CrossRef](#)]
52. Vuppala, S.K.; Padmanabh, K.; Bose, S.K.; Paul, S. Incorporating fairness within demand response programs in smart grid. In *Proceedings of the IEEE PES Innovative Smart Grid Technologies (ISGT)*, Manchester, UK, 5–8 December 2011; pp. 1–9.
53. Baharlouei, Z.; Hashemi, M.; Narimani, H.; Mohsenian-Rad, H. Achieving optimality and fairness in autonomous demand response: Benchmarks and billing mechanisms. *IEEE Trans. Smart Grid* **2013**, *4*, 968–975. [[CrossRef](#)]
54. Peng, J.; Zhu, Y.; Zhao, Q.; Zhu, H.; Cao, J.; Xue, G.; Li, B. Fair Energy-Efficient Sensing Task Allocation in Participatory Sensing with Smartphones. *Comput. J.* **2017**, *60*, 850–865. [[CrossRef](#)]
55. Holzbauer, B.O.; Szymanski, B.K.; Bulut, E. Socially-aware market mechanism for participatory sensing. In *Proceedings of The First ACM International Workshop On Mission-Oriented Wireless Sensor Networking*, New York, NY, USA, 5–7 November 2012; p. 9.
56. Jabbari, S.; Joseph, M.; Kearns, M.; Morgenstern, J.; Roth, A. Fairness in Reinforcement Learning. 2016. Available online: <http://xxx.lanl.gov/abs/1611.03071> (accessed on 26 October 2018).
57. Baharlouei, Z.; Hashemi, M. Efficiency-fairness trade-off in privacy-preserving autonomous demand side management. *IEEE Trans. Smart Grid* **2014**, *5*, 799–808. [[CrossRef](#)]
58. Saputro, N.; Akkaya, K.; Uludag, S. A survey of routing protocols for smart grid communications. *Comput. Netw.* **2012**, *56*, 2742–2771. [[CrossRef](#)]

59. Siano, P. Demand response and smart grids, a survey. *Renew. Sustain. Energy Rev.* **2014**, *30*, 461–478. [[CrossRef](#)]
60. Wang, Z.; Yang, K.; Wang, X. Privacy-preserving energy scheduling in microgrid systems. *IEEE Trans. Smart Grid* **2013**, *4*, 1810–1820. [[CrossRef](#)]
61. Gong, Y.; Cai, Y.; Guo, Y.; Fang, Y. A Privacy-Preserving Scheme for Incentive-Based Demand Response in the Smart Grid. *IEEE Trans. Smart Grid* **2016**, *7*, 1304–1313. [[CrossRef](#)]
62. Work, D.B.; Tossavainen, O.P.; Blandin, S.; Bayen, A.M.; Iwuchukwu, T.; Tracton, K. An ensemble Kalman filtering approach to highway traffic estimation using GPS enabled mobile devices. In Proceedings of the 47th IEEE Conference on Decision and Control, Cancun, Mexico, 9–11 December 2008; pp. 5062–5068.
63. Zhu, M.; Liu, X.Y.; Kong, L.; Shen, R.; Shu, W.; Wu, M.Y. The charging-scheduling problem for electric vehicle networks. In Proceedings of the IEEE Wireless Communications and Networking Conference, Istanbul, Turkey, 6–9 April 2014; pp. 3178–3183.
64. Rawat, P.; Singh, K.D.; Chaouchi, H.; Bonnin, J.M. Wireless sensor networks: A survey on recent developments and potential synergies. *J. Supercomput.* **2014**, *68*, 1–48. [[CrossRef](#)]
65. Moriarty, D.E.; Schultz, A.C.; Grefenstette, J.J. Evolutionary algorithms for reinforcement learning. *J. Artif. Intell. Res.* **1999**, *11*, 241–276. [[CrossRef](#)]
66. O’Neill, D.; Levorato, M.; Goldsmith, A.; Mitra, U. Residential demand response using reinforcement learning. In Proceedings of the 2010 First IEEE International Conference on Smart Grid Communications (SmartGridComm), Gaithersburg, MD, USA, 4–6 October 2010; pp. 409–414.
67. Bennati, S.; Jonker, C.M. PriMaL: A Privacy-Preserving Machine Learning Method for Event Detection in Distributed Sensor Networks. 2017. Available online: <http://xxx.lanl.gov/abs/1703.07150> (accessed on 26 October 2018).
68. Dusparic, I.; Cahill, V. Autonomic multi-policy optimization in pervasive systems. *ACM Trans. Auton. Adapt. Syst.* **2012**, *7*, 1–25. [[CrossRef](#)]
69. Huberman, B.A. Social Dilemmas and Internet Congestion. *Science* **1997**, *277*, 535–537. [[CrossRef](#)]
70. Radanovic, G.; Faltings, B. Incentive Schemes Participatory Sensing. In Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, Istanbul, Turkey, 4–8 May 2015; pp. 1081–1089.



© 2018 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).