

The Knowledge Engineering Review, Vol. 30:3, 237–264. © Cambridge University Press, 2014. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/3.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

doi:10.1017/S026988891400023X

First published online 19 November 2014

Self-organising management of user-generated data and knowledge

SAM MACBETH and JEREMY V. PITT

Intelligent Systems and Networks Group, Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2BT, UK;

e-mail: samuel.macbeth06@imperial.ac.uk; j.pitt@imperial.ac.uk

Abstract

The proliferation of sensor networks, mobile and pervasive computing has provided the technological push for a new class of participatory-sensing applications, based on sensing and aggregating user-generated content, and transforming it into knowledge. However, given the power and value of both the raw data and the derived knowledge, to ensure that the generators are commensurate beneficiaries, we advocate an open approach to the data and intellectual property rights by treating user-generated content, as well as derived information and knowledge, as a common-pool resource. In this paper, we undertake an extensive review of experimental, commercial and social participatory sensory applications, from which we identify that a decentralised, community-oriented governance model is required to support this approach. Furthermore, we show that Ostrom's institutional analysis and development framework, in conjunction with a framework for self-organising electronic institutions, can be used to give both an architecture and algorithmic base for the requisite governance model, in terms of operational and collective-choice rules specified in computational logic. This provides, we believe, the foundations for engineering knowledge commons for the next generation of participatory-sensing applications, in which the data generators are also the primary beneficiaries.

1 Introduction

Participatory sensing (Burke *et al.*, 2006) is the process of leveraging user devices which are capable of various sensor measurements, to gather data in a bottom-up fashion and gain knowledge from the analysis of this data. It has already been applied in many varying domains, from traffic and transportation (Mathur *et al.*, 2010; Costa *et al.*, 2012), to environmental conditions (Mendez *et al.*, 2011; Hasenfratz *et al.*, 2012), product pricing (Deng & Cox, 2009) and behavioural information (Miluzzo *et al.*, 2008).

In all of these applications, individual users or devices are gathering data which is then aggregated by a third party. Having collected this data, it is primarily this third party who reaps the benefits from the analysis of the data. The value of data gathered through these means has been estimated in the billions of dollars (Manyika *et al.*, 2011). While some organisations provide some return to contributors, usually in services, the equitability of this arrangement is debatable (van Dijck, 2009). This has led to a call for the empowerment of users such that they can achieve a fair exchange for their data (Buckingham Shum *et al.*, 2012).

A key issue is the provision of tasks for participatory sensing. This is the process by which an entity determines a set of parameters to sense, and builds community and infrastructure for this purpose. This is often a top-down process: of the applications we have mentioned the majority are centralised. The result is

that a single entity holds all of the collected data (often taking ownership or property rights of the data provided by others (O'Hara & Shadbolt, 2010)), sets the policies regarding access to data and knowledge derived from the data, and controls how these policies are changed.

This approach is in direct contrast to how some of the most successful resources for data and knowledge have developed on the Internet. The Open Data Institute has shown the benefits of open access to government data and that significant additional value can be generated by allowing others open and permissive access (Shadbolt *et al.*, 2012). Wikipedia¹ is an example where individuals contributing knowledge have created an encyclopaedia at a scale which would be infeasible to do in a traditional top-down manner. Its continued success is down to its decentralised governance which allows its users to direct the trajectory of the site's policy (Famiglietti, 2011).

To understand the implications of how these organisational structures affect user participation and the associated benefits from the data and knowledge generated, we look at how social organisations have developed around knowledge repositories. Hess and Ostrom (2003) argue that information can be seen as a common-pool resource (CPR) and thus analysed using the considerable existing literature on the commons (Hess & Ostrom, 2006). Some initial investigations have been applying this approach to user-generated content (Pitt, 2012b) which we take further here.

From her significant field work on physical commons, Ostrom (1990: 42) outlined 'the problem of supplying a new set of institutions'. This is the problem that, in order to form an organisation, someone must first provide an initial set of rules by which the organisation and its members are governed. This is a difficult task as there are many stakeholders to satisfy under changeable conditions. Through survey of both successful and unsuccessful institutions Ostrom extracted a set of principles which were more prevalent in successful cases.

In participatory sensing we can empower users by providing them with the ability to supply institutions, and the knowledge to assess the effect of the rules and organisational structure governing sensing applications. We review the literature analysing knowledge commons to derive an analytic framework for institutional design of participatory-sensing applications construed as a provision and appropriation system. We can then derive guidelines for first, how a self-organising participatory-sensing application can function as a provision and appropriation system using a knowledge commons, and second, how such a system can be supplied.

This paper is organised as follows: in Section 2 we introduce the theory behind information and knowledge as a commons, and how participatory sensing can be seen as such a commons. We then review participatory-sensing applications in Section 3. This is followed by an analysis of participatory sensing as a knowledge commons according to Ostrom's institutional analysis and development (IAD) framework in Section 4. This analysis allows us to then create a formal representation of a system for self-organising management of a participatory-sensing application as an information and knowledge commons and subsequently evaluate this system. Our evaluation shows that such institutions could enable autonomous, heterogeneous agents to manage information and knowledge commons used for participatory-sensing applications while maintaining important criteria such as sustainability, participation standards and equity.

We finish, in Section 5, with some conclusions derived from this critical review and analysis of some representative participatory applications, in particular that:

- Data clouds in open participatory-sensing applications can be construed as information and knowledge commons and thus characterised by provision and appropriation actions.
- A system for access control (i.e. provision and appropriation) in participatory-sensing applications can be designed according to Ostrom's institutional design principles for self-governing institutions and formally specified in an action language.

This provides the foundations for engineering knowledge commons for the next generation of participatory-sensing applications, in which the data generators are also the primary beneficiaries.

¹ <http://www.wikipedia.org>

2 Data and knowledge as a commons

In this section we address how participatory sensing can be seen as an information and knowledge commons. We give some background on the commons and introduce Ostrom's work on managing CPRs. We describe how the characteristics of information and knowledge as resources mean that repositories of them can also be viewed as commons, and therefore managed in the same way.

The commons is a term generally used to describe shared resource systems. They are often also called common-pool resource systems. These systems are characterised as systems where the ownership of the resource system (land, air, etc.) and resource units (trees, radio frequencies, etc.) are shared, public property or not covered by any property rights; and it is difficult to exclude access to the resource and resource units to others. Hardin (1968) theorised that these properties would inevitably lead to overuse and depletion of the resource, a 'Tragedy of the Commons', and that enclosure (via privatisation or centralisation) of the resource system was the only solution.

2.1 Governing the commons

Having done extensive field work analysing both cases where this tragedy had been overcome and where it had not, Ostrom (1990) contested that it was in fact possible to sustainably manage these resources as a commons. The key to this success was how the commons were governed and how the individuals involved in the use and maintenance of the commons were engaged in this governance. Ostrom identified eight principles which, when all satisfied, increase the chances of a successful commons (Ostrom, 1990: 91–101):

1. Clearly defined boundaries: 'Individuals or households who have rights to withdraw resource units from the CPR must be clearly defined, as must the boundaries of the CPR itself'.
2. Congruence between appropriation and provision rules and local conditions: 'Appropriation rules restricting time, place, technology, and/or quantity of resource units are related to local conditions and to provision rules requiring labor, materials, and/or money'.
3. Collective-choice arrangements: 'Most individuals affected by the operational rules can participate in modifying the operational rules'.
4. Monitoring: 'Monitors, who actively audit CPR conditions and appropriator behaviour, are accountable to the appropriators or are the appropriators'.
5. Graduated sanctions: 'Appropriators who violate operational rules are likely to be assessed [sic] graduated sanctions (depending on the seriousness and context of the offence) by other appropriators, by officials accountable to these appropriators, or by both'.
6. Conflict-resolution mechanisms: 'Appropriators and their officials have rapid access to low-cost local arenas to resolve conflicts among appropriators or between appropriators and officials'.
7. Minimal recognition of rights to organise: 'The rights of appropriators to devise their own institutions are not challenged by external governmental authorities'.
8. Nested enterprises: 'Appropriation, provision, monitoring, enforcement, conflict resolution, and governance activities are organised in multiple layers of nested enterprises'.

Ostrom later formalised the analysis methodology used as the IAD framework (Ostrom, 2005). This framework examines the relevant factors in an institution which manages a commons, allowing the evaluation of institutions with respect to Ostrom's eight principles, and with respect to other institutions which have been analysed in the same way. The framework can also be used for the design of institutions, and we will use the framework in this way later in this article.

Since Ostrom's (1990) book many new areas have been recognised as commons. It has been deemed applicable to many shared resource problems: from transport to radio spectrum and the Internet (Hess, 2000). Many have emerged from technological innovations which have either created a new virtual resource, or transformed the properties of a resource such that it is feasible to treat it as a commons. One example is knowledge (Hess & Ostrom, 2006), with open-access publishing as a specific case study.

2.2 Data, information and knowledge commons

The shift to treating knowledge as a commons has largely been enabled by the digitalisation of information which has separated the knowledge from the physical format of publication (Hess & Ostrom, 2003). In this section we consider data, information and knowledge as resources, and how they are managed. For this discussion we require a definition of these terms within this context. Most definitions give a hierarchy where data is just raw values, information is data organised in context, and knowledge is an understanding of meaning of information and how to use it (Machlup, 1983; Davenport & Prusak, 2000; Shadbolt, 2013). For example the tuple (51,0) would be data; adding the context that this tuple corresponds to latitude and longitude transforms it into information; and knowledge would be the ability to provide directions to travel to this location along public roads. It is generally thought that with current technology the creation of knowledge requires human intervention (Davenport & Prusak, 2000), but data and information can be generated independently by devices. In this work we will use these definitions for data, information and knowledge.

Hess and Ostrom (2003) assessed how information and knowledge can be a commons by using a classification of *exclusion* and *subtractability*. Exclusion is a measure of how easy it is to exclude individuals from the benefits of the good, either through physical or legal means. Subtractability (also called rivalry) is a measure of the extent to which the benefits consumed by one individual subtract from the benefits available to others. Combined these give a two-dimensional classification of goods (see Table 1) giving four types of goods: public goods, CPRs, toll or club goods, and private goods.

The shift that has occurred in information and knowledge due to technology is that its distribution has become non-rivalrous (Hess & Ostrom, 2003; Bollier, 2007). Previously, information and knowledge would have to be put into tangible form, for example, a book, to allow others to use it. In this form it has high subtractability: by owning a book you prevent someone else owning that copy of the book, and the supply of books is limited by the publisher. This meant that information and knowledge goods were either common-pool or private goods depending on whether they were distributed by a library (difficult exclusion) or sold (easy exclusion). In digital form, however, these goods can be freely copied for negligible cost. They are infinite and undepletable (Bollier, 2007). Therefore, these goods would count as either public goods or club goods.

In the context of information and knowledge as goods, copyright law enables exclusion to be controlled by the information or knowledge creator. This allows a book to simultaneously be a private good—when purchased by an individual—and part of a CPR—when made available by a library. Digital information has gained the same flexibility, both through the protection of the law, but also through technological means via access-control and digital rights management (Lessig, 2004). Wikipedia provides a public knowledge resource, having chosen to not have any exclusion, while Journals use pay-walls to exclude access for those without subscriptions.

Lessig (2004) also shows that copyright law has also been used to artificially change the subtractability of digitalised information and knowledge. For example, some publishers have imposed e-book lending

Table 1 Types of goods

Exclusion	Subtractability	
	Low	High
Difficult	Public goods Useful knowledge Sunsets	Common-pool resources Libraries Irrigation systems
Easy	Toll or club goods Journal subscriptions Day-care centres	Private goods Personal computers Doughnuts

Hess and Ostrom (2006).

limits on libraries, preventing too many simultaneous withdrawals. This effectively prevents this knowledge from becoming a public good, ensuring that library lending remains a CPR.

Therefore, we have a resource which is able to be cast into any of the four goods categories. While information and knowledge should naturally be a public good, technological and legal control allows the properties to be changed in order to create a more beneficial good for the creator. Copyright law generally advocates that creative work should eventually become public goods by falling into the public domain, and that the protected state of work is purely an incentive mechanism (Samuelson, 2006). The argument for this is that, because information and knowledge also aid in the creation of new information and knowledge, having a large public domain will have a cumulative effect and increase overall knowledge as a result. We reach a dilemma where the knowledge creators wish to exert their control to achieve a short-term benefit which causes a sub-optimal outcome for all the users of their knowledge, and also their long-term productivity.

The premise of the information and knowledge commons is that a method of collective action as seen in past commons literature can be used here in order to tap into the benefits of making information and knowledge more available (Hess & Ostrom, 2006; Fuster Morell, 2010; Shadbolt, 2013). Ostrom and Hess (2007) analysed the knowledge commons in the same way as physical commons: using the IAD framework.

The creation of an institution in order to manage knowledge and information in practical scenarios is certainly non-trivial. There is a problem of supply (Ostrom, 1990: 42)—that someone must provide the initial organisational structure and institutional rules—and, as we will see in the next section, there is a fine line between the emergence of centralised or collective governance. Rules which satisfy Ostrom's eight principles are often very specific to the resource characteristics so we do not have a pre-existing 'library' of rules available when looking to 'supply' an institution with which to manage a new resource. The difficulty of implementing effective collective governance structures is arguably a contributing factor to why in many cases initiators decide to follow Hardin's advice—privatisation seems to be the only way.

2.3 Application to participatory sensing

Ostrom provides an approach for preventing the tragedy of the commons in social systems using rules interpreted and acted upon by human actors. In order to use this approach for the kind of large-scale multi-user content-creation applications seen in participatory sensing we need to reconcile this information and knowledge with that of the information and knowledge commons. Furthermore, as we are working in a purely digital domain with high rates of interaction, we should put our institutional rules into digital form too. This will enable both human and computational participants to interact with the rules as well as the capability to handle a greater speed and quantity of actions than a social system could.

Participatory sensing deals with the creation of sensor networks from user devices (Burke *et al.*, 2006). It tasks large numbers of independent individuals to contribute information to a central pool. This information is then aggregated and analysed to generate new, valuable information and knowledge.

We can view this pool of information which is generated, as a commons. It is a collection of information contributed by a large number of independent actors. The value of the pool is cumulative, in that additional contributions to the pool increase its value. Thus, this CPR has the properties required to be considered an information commons (Bollier, 2007).

Furthermore, we can consider the knowledge required to aggregate and analyse the information pool gathered by sensing. This can also be seen as a resource which could be shared with others to reduce the data processing burden on an individual, or just to share knowledge for others to improve upon. While not currently common in participatory-sensing applications, some platforms do accommodate contributions of analytical algorithms (Kansal *et al.*, 2007). Thus, if knowledge is contributed to the participatory-sensing commons then we can also consider it a knowledge commons.

Therefore, the information and knowledge collected and generated by participatory-sensing applications can be seen as an information and knowledge commons. This commons can be managed by an institution of human actors. However, in many cases it can be desirable to be able to transfer this task to computational agents. For example, in situations where decisions must be made quickly, frequently and

repeatedly, a computational agent is likely to be much better suited to this role—Wikipedia allow the use of *bots*² for many such tasks such as detecting and reversing malicious edits and checking for copyright violations. In socio-technical systems there are both human and computational actors, so there needs to be some kind of common interface within which they can interact. In addition, there are purely computational domains where we may benefit from these institutions which aim to balance the needs of multiple, conflicting stakeholder positions. For these domains we use an electronic institution, which can be considered to be a computational analogue of Ostrom's informal definition of an institution as "a set of working rules that are used to determine who is eligible to make decisions in some arena, what actions are allowed or constrained, ... [and] contain prescriptions that forbid, permit, or require some action or outcome" (Ostrom, 1990: 51). The institution is intended to enable agents to interact in order to achieve both individual and collective goals by regulating and constraining their behaviour, voluntarily, according to a set of conventional rules.

In this context, Ostrom's work dealing with social systems has been translated to electronic institutions, principally by formalising the definition in computational logic and, critically, extending it with the use of institutionalised power (Jones & Sergot, 1996) for the formal representation of permitted, forbidden and obliged actions. Effectively, we are axiomatising Ostrom's principles for computational agents (Pitt *et al.*, 2012) such that agents can understand institutional rules and choose whether or not to abide by them. These formalisations are a form of algorithmic governance, translating the written rules and protocols from social systems into axioms and algorithms which agents can understand and a computer can execute.

Therefore, by using electronic institutions for algorithmic governance according to Ostrom's principles for enduring institutions, we can create sustainable organisations and communities around participatory sensing.

3 Participatory sensing and the knowledge commons

Having outlined the theoretical case for self-organising governance of participatory sensing as a knowledge commons, this section presents a review of existing participatory-sensing applications in order to assess the current approach to governance models. We outline a set of criteria which we use to evaluate a representative sample of both experimental applications from the literature and commercial applications leveraging participatory sensing in their product. We also look at two socio-technical systems dealing with information and knowledge as an example of effective self-organising management of such a resource.

There are two other relevant surveys on participatory-sensing applications. Christin *et al.* (2011) surveys privacy in 30 sensing applications. The survey identifies privacy threats by looking at what is sensed and the granularity of the sensing in each application. Tilak (2013) performs a simple survey of the domains of several applications and the hardware and software tools used for sensing. This survey differs from these, in that our focus is on the organisational structures and where benefits are accrued in each application.

3.1 Evaluative criteria for participatory sensing

Our survey criteria assess three distinct points in the sensing process: what information flows into the system and how; how information is managed inside the system; and what benefits are generated by the system and who can access them. The aim is to determine the flow of value (in the form of information and knowledge) into and out of the system, and if and how users are incentivised to contribute to and sustain the system.

We classify information flowing into the application in two ways. First, by the source of the information. It can be people centric, collecting information about the user, or environment centric, capturing information about the participants' surroundings (Kanhere, 2013). The latter is public information, any other user could gather the same data by being in the same physical location (or equivalent). Second, we look at how the data is gathered. This can be explicitly contributed or implicitly gathered (Shadbolt, 2013).

² <http://en.wikipedia.org/wiki/Wikipedia:Bots>

Most participatory sensing is explicit—the user knows that they are contributing information. However, increasingly information is being gathered implicitly, often without user knowledge.

The management of information is dictated by the organisational structure of the system, both technical (the architecture) and social (governance). A sensing system's architecture and governance combine to specify where the power and assets lie in a system. A centralised governance will mean that a small number of actors have control over the system and the rules all users must adhere to. Less-centralised governance methods will spread this power around and possibly require consensus for rule changes and/or elected positions. A centralised architecture is one where all data is aggregated under one entity's control. While this is often a practical solution to data aggregation, it gives this individual power through control of the core assets of the system. Effective governance is required to limit these powers. The architecture can also be decentralised, for example, using peer-to-peer (p2p) technologies. This can have a limiting effect on the governance of a system as there is no single point where control can be applied.

We classify the benefits of a participatory-sensing application on two factors: who they are beneficial for and who can access them. An output benefit from an application can either be *personal*—tailored to an individual based on the information they have contributed—or *community*—aggregated information useful to everyone on the given task. Access to benefits can be *public* or *private*, which denote whether some kind of membership is required to receive benefits.

In total we have six criteria:

- information source: *environment* or *people*;
- gathering method: *explicit* or *implicit*;
- governance: *centralised* or *community*;
- architecture: *centralised server* or *p2p*;
- benefits: *personal* and/or *community*; and
- access: *private* and/or *public*.

3.2 Review of participatory-sensing applications

We took a sample of experimental applications from the literature covering multiple domains including environment monitoring, traffic management and price comparison. The reviewed applications are as follows:

- The *Pothole Patrol* system (Eriksson *et al.*, 2008) uses vibration and GPS data from participating vehicles to detect road surface conditions.
- *CenceMe* (Miluzzo *et al.*, 2008) is a mobile phone application which attempts to infer the current context of the user from the device's sensors as well as neighbouring devices' sensors.
- The *LiveCompare* (Deng & Cox, 2009) phone application users contribute pictures of grocery products' price tags which are analysed to create a price matching service.
- *VTrack* (Thiagarajan *et al.*, 2009) collects location information from smartphones in order to provide travel time estimation for drivers.
- In *Parknet* (Mathur *et al.*, 2010) vehicles monitor for available road-side parking spaces which is then aggregated to provide real-time parking availability over a city-wide area to various third parties.
- *VibN* (Miluzzo *et al.*, 2011) gathers smartphone sensor data to identify events in a city in real time.
- *P-Sense* (Mendez *et al.*, 2011) measures air pollution from smartphones to generate pollution maps of the local area.
- *Cloud2Bubble* (Costa *et al.*, 2012) collects data on ambient conditions on public transport from smartphones and suggests to users changes that they can make to their journey to give themselves a better quality of experience.
- *Gas-Mobile* (Hasenfratz *et al.*, 2012) uses low-cost sensors connected to smartphones to measure air pollution levels.

Each application was evaluated according to our given criteria. The results are shown in the 'Experimental applications' section of Table 2.

Table 2 Survey of participatory-sensing applications

Application	Source	Method	Governance	Architecture	Benefits	Access
Experimental applications						
Pothole Patrol (Eriksson <i>et al.</i> , 2008)	Environment	Explicit	—	Centralised	Community and personal	—
CenceMe (Miluzzo <i>et al.</i> , 2008)	People	Explicit	—	Centralised	Personal	Private (app)
LiveCompare (Deng & Cox, 2009)	Environment	Explicit	—	Centralised	Community	Private
VTrack (Thiagarajan <i>et al.</i> , 2009)	Environment	Explicit	—	Centralised	Personal	Private
Parknet (Mathur <i>et al.</i> , 2010)	Environment	Explicit	—	—	Community	—
VibN (Miluzzo <i>et al.</i> , 2011)	Environment	Explicit	—	Centralised	Community	—
P-Sense (Mendez <i>et al.</i> , 2011)	Environment	Explicit	—	Hybrid	Community	—
Cloud2Bubble (Costa <i>et al.</i> , 2012)	Environment	Implicit	—	Centralised	Personal	Private (app)
Gas-Mobile (Hasenfratz <i>et al.</i> , 2012)	Environment	Explicit	—	p2p	Community	—
Commercial applications						
Waze	Environment	Implicit and explicit	Centralised	Centralised	Community	Public and private (app)
OpenSignal	Environment	Explicit	Centralised	Centralised	Community	Public
Social systems						
Wikipedia	N/A	Explicit	Community	Centralised	Community	Public
F/OSS	N/A	Explicit	Community	Centralised and p2p	Community and personal	Public and private

p2p, peer-to-peer; F/OSS, free/open-source software.

When a classification cannot be made on a criterion due to lack of available information we mark it as not specified (denoted by —).

The reviewed papers all follow a similar methodology. Each one identifies a domain to which participatory sensing can be applied, determines a method for gathering appropriate data from users, develops an algorithm for producing useful analytics from the data set and deploys a small test application to verify the method. Under this method there is no need to consider how to incentivise participation in the application via fair governance and access to benefits. We made the following observations from our survey:

- Data gathering is almost always explicit. Users are recruited specifically for the purpose of the sensing campaign.
- Governance is not considered in any case. In deployed applications it is assumed that the institution has complete control over the application with no oversight.
- Architectures are mostly centralised. This is the simplest way to deploy such a system when scaling is not a requirement. In some cases the authors considered scaling, and thus architectures which utilised some p2p technologies were used (Mendez *et al.*, 2011; Hasenfratz *et al.*, 2012).
- Access to the system's benefits is only specified in four cases, and this access is always private, requiring certain membership or ownership of an app. As the authors' aims were generally to show that the proposed algorithms generate correct inferences from the sensors then access to this data for users was not considered in many cases.

To see how participatory sensing is deployed in real-world scenarios we extended the survey to two commercial applications, Waze³ and OpenSignal⁴. Waze is a smartphone application which provides car navigation information to the user as well as alerts about hazards which may affect users' journeys. Information is sent back while the app is in use to track the speed of certain routes and traffic levels. Users may also actively report events which affect their progress, such as accidents or flooding. Waze uses this information to first build a map of possible driving routes and detect new roads, and second to provide optimal routing to users with accurate arrival predictions. OpenSignal aggregates information on mobile phone signal quality from user devices in order to generate maps of the quality of service available under each provider across the world.

The survey of these applications is shown in Table 2 ('Commercial applications' section). In both cases a central governance is used. The companies control any changes to policies governing the data they hold and access to it. The applications differ in how data are collected and benefits accessed. OpenSignal is explicit about data gathering—the sole purpose of their app is the gathering of signal strength data. Their aggregated data set is then made open to the public via their websites and publications. Therefore, this application does not provide incentives for users to participate in the gathering effort beyond some personal stats on their own gathered data. Waze, on the other hand, hides its data gathering behind the app's functionality as a satellite navigation tool. While the app is performing navigation for the user, data is fed back to Waze in order to improve the navigation performance. Through its design the app implicitly enforces a reciprocal relationship where information on traffic and road conditions is exchanged for optimised travel directions. The app encourages additional explicit contributions in the form of tagging roadwork locations, etc., which is then provided to other users to alert them about hazards or to improve routing. While access to the real-time travel map is made publicly available through the Waze website, it does not provide the navigation available in the app. Therefore, the application sustains itself simply by providing accurate navigation via the app which in turn causes users to implicitly contribute information to improve the service.

From this review we draw two main conclusions. First, that the current work on participatory sensing focuses on the provision of sensing campaigns and algorithms for the processing of data for those campaigns. Second, where governance and infrastructure has been provisioned for large-scale, long running participatory-sensing deployments it tends to be centralised. The fact that centralised governance is provisioned should not be a surprise. For the organisations involved it satisfies a self-interested strategy,

³ <http://www.waze.com>

⁴ <http://opensignal.com>

which is also an effective one in terms of return on investment and/or profit. This governance choice does not seem to hinder the operation of the system as users can be incentivised to contribute.

3.3 Extension to knowledge commons

To assess possible alternatives to centralised governance of participatory-sensing systems we extend our survey to a particular kind of social system which is accessed and controlled via the Internet. These systems involve large numbers of participants and contributors, who access and build on the information and knowledge in a shared resource. They can be seen as knowledge commons. We look at two examples: Wikipedia⁵ and Free/open-source software⁶ (FOSS). Wikipedia is an example of a very successful deployment of community governance which has created a knowledge commons which is arguably richer than what could be produced through other means. In FOSS we have a general framework which has enabled anyone in the world to create and contribute to software which can compete with commercial products, and also give free and libre⁷ access to this software. We can test both of these systems against the same criteria we have used previously, and the results are again in Table 2 ('Social systems' section).

Under our survey criteria, Wikipedia is classified as gathering information explicitly. Contributions are new articles, modifications to existing ones or moderations of others' submissions, and therefore users generating this content are aware that they are provisioning it to Wikipedia. The result of these contributions is a knowledge resource which is available to anyone, and no specific rewards are reserved to encourage editing or moderation actions. While its architecture is centralised under one website, Wikipedia has a complex community governance which allows policies on the site to change, provided consensus is behind the change. This governance has developed over time to become more open in response to criticism and concerns over possible involvement from commercial organisations (Fuster Morell, 2011). The need for dynamic governance is in part driven by the permissive licence given to content on the site, which allows anyone to publish the content elsewhere, and has in the past led to the forking of whole sections to separate sites (Famiglietti, 2011).

Like Wikipedia, contributions to FOSS are explicit. Contributions take the form of code commits, code review, bug reporting, mailing-list discussion and general project administration. There are some cases of open-source software which gather anonymised usage information, a process which would count as implicit information gathering, however, currently this is quite rare (although increasingly common in commercial software). The governance of FOSS projects tends to be informal and dynamic (Schweik, 2007). In addition, as the FOSS community has matured, tools for both the administration of source code (code repositories, bug trackers, etc.) and the administration of project governance (mailing lists, discussion forums, etc.) have been developed. The availability of these tools is generally free and commoditised to such an extent that creating the infrastructure for a new FOSS project can be done at zero monetary cost⁸.

For both Wikipedia and FOSS, work analyses have been done with respect to Ostrom's framework. Forte *et al.* (2009) describe how Wikipedia's governance conforms to Ostrom's eight principles. Schweik (2007) uses the IAD analysis on FOSS. The innovation of Copyleft (Stallman, 1999), which allows the creation of code licensed under this regime to be considered as a public good, is credited as being a key rule-in-use in FOSS by legally binding works into the commons and keeping them there. It offers a middle ground between the copyright extremes of 'all rights reserved' and public domain. Schweik (2007) also advocates that the FOSS collaborative paradigm can be applied to any collaboration around intellectual property. These two examples show that successful systems have been built to develop and manage shared resources of information and knowledge using decentralised, community governance. Like in physical

⁵ <http://www.wikipedia.org>

⁶ Software which is free *and* open source, as advocated by the Free Software Foundation (<http://www.fsf.org>), amongst others.

⁷ In the context of open-source software these terms are used to describe two different meanings of 'free': 'for zero price' and 'with little or no restriction' (libre).

⁸ Sites such as <http://sourceforge.net> and <https://github.com> provide a full suite of tools for free to projects licensed under FOSS licenses.

commons the governance system can be shown to conform to Ostrom's principles, and further analysis can identify some particular features which contribute to their success.

3.4 Summary

From our survey of participatory-sensing applications we have identified several issues. First, a lack of governance consideration. This led to a review of the governance that does exist in commercial applications. Our observation is that the high cost of provisioning a participatory-sensing application under the current paradigm contributed to a centralised governance which seeks to achieve a return on investment. Finally, we looked at social systems for managing knowledge commons which have been shown to conform to Ostrom's design principles to sustain and develop a rich commons of knowledge, as well as a framework which enables very cheap and easy provision of infrastructure and governance for information and knowledge commons. This is the motivation for specifying a system which uses a collective governance model for participatory sensing, as presented in the next section.

4 A commons for participatory sensing

We have seen how information and knowledge can be treated as a commons, and some examples where collective governance has been supplied for this purpose. In this section we propose an architecture for provision and appropriation of information and knowledge for participatory sensing and analyse it, using the IAD framework (Ostrom, 2005), as an information and knowledge commons. We then specify a formal system for participatory sensing using this architecture, show how it can accommodate Ostrom's principles and evaluate the expected outcomes of such an institution according to the evaluative criteria of the IAD. Finally, we consider how formal system can provide a library of rules which can be instantiated in order to 'supply' an institution.

In order to define a participatory-sensing application, we take a set of four general user roles, as defined by Burke *et al.* (2006), and add a fifth role. These roles are:

- *initiators*, who initiate the application and form an organisation around it;
- *gatherers*, users who participate in the information gathering and provision it;
- *evaluators*, who verify and classify received information;
- *analysts*, who process the information to create conclusions on the data, often in the form of new information and/or knowledge; and
- *consumers* who demand the derived, or second-order, information and knowledge.

We consider a 'role' in this context to be an institutionally assigned label to denote what is expected and/or permitted for a user to do. Note that user roles are not mutually exclusive and a user may occupy many roles simultaneously within one institution. For example, in a large proportion of cases *gatherers* are also *consumers*, and in fact their compensation for their gathering efforts is the right to consume. Equally, *initiators* often are *evaluators* and *analysts* too. Therefore, the model allows appropriation of knowledge by a user occupying the role of consumer, if that user also occupies the role of *analyst*. The formulation of role in this section specifically allows agents to occupy multiple roles in the same institution, and indeed different roles in different institutions.

We consider this participatory-sensing application in the form of a provision and appropriation system, where the resources being provisioned and appropriated are information and knowledge. Figure 1 illustrates such a system and how each role interacts with the resource. Four boxes represent user roles interacting with the resource. Some roles will require additional agent internals, such as *gatherers* requiring appropriate sensors and *analysts* requiring knowledge to aggregate sensor information. Arrows represent movement of information and knowledge; dotted arrows represent optional actions.

The IAD framework can be used to analyse such a system, develop it and evaluate it. The IAD is split into nine areas of analysis (Figure 2). The left side of the framework looks at what the resource and the community using it is like, and rules which have been created for resource and institution management. The middle section deals with where interactions occur and who these interactions are between. The right-hand side

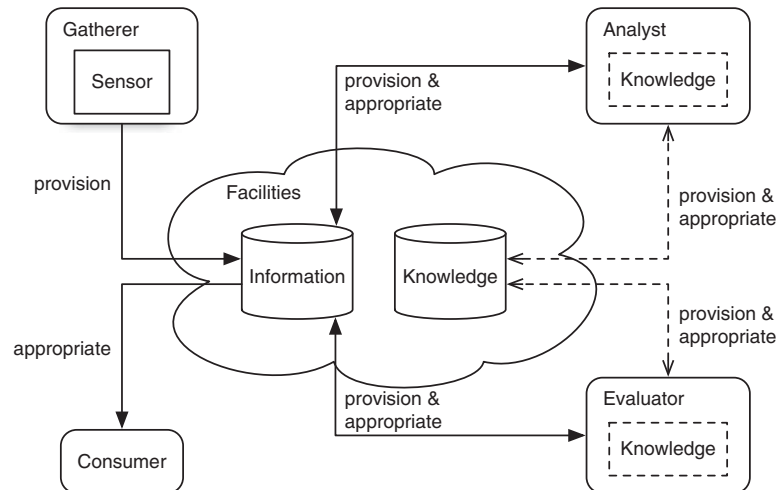


Figure 1 Participatory sensing as a provision and appropriation system. Dotted arrows represent optional actions for that role

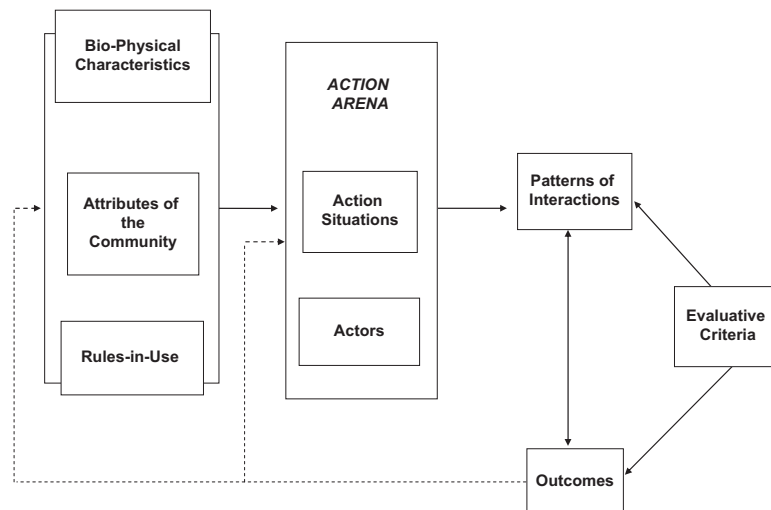


Figure 2 Institutional analysis and development framework (Ostrom & Hess, 2007: 44)

observes what the outcomes are and evaluates them. Our use of the IAD in the section is split into three parts. First, we provide an analysis of a participatory-sensing system using the left-hand side of the framework, looking at resource and community characteristics and rules-in-use. Second, we propose a new set of rules-in-use to address each of Ostrom's principles for sustainable institutions. Finally, we use the observational and evaluative aspects on the right of the IAD framework to evaluate our proposed system. The final part of this section addresses the issue of supply of such an institution.

4.1 IAD analysis

We now present an analysis of participatory sensing as a provision and appropriation system. Following the IAD framework we discuss the characteristics of the information and knowledge gathered through sensing as a resource (bio-physical characteristics), the community of individual actors involved in the sensing process (attributes of the community), and how institutional rules are, or could be, created (rules-in-use).

4.1.1 Bio-physical characteristics

Typically this area is concerned with the physical attributes of the resource, in terms of *flow* and *facility*. This distinction separates the resource units (*flow*) from where they are stored and generated (*facility*). In order to deal with the complexity of information and knowledge Hess and Ostrom (2003) put forward an alternative method, considering *ideas*, *artefacts* and *facilities*. In this framework, *ideas* are data, information and knowledge in intangible form; *artefacts* are the expression of ideas into tangible, observable form; *facilities* store artefacts and make them available. As intangible objects, *ideas*, by definition, cannot be represented within a computational system—they must be expressed into files, databases or algorithms to be used. In this form they are then *artefacts*. Therefore, we treat *ideas* as a resource flow external to the system, and instead manage the tangible derivative works from them: *artefacts*.

In participatory sensing, the *artefacts* are raw sensor data, contextualised sensor information, algorithms for the analysis of data and rich information generated from these algorithms. Some of these artefacts will be protected by copyright law, and this will affect what control the organisation is able to have over them. The status of data and information (as we have defined them) with regards to copyright is currently ambiguous, traditionally it was not copyrightable, however, databases are copyrightable (Miller *et al.*, 2008). This copyright protection is important as it affects how easily the organisation can protect its assets. Bad protection limits the ability of the organisation to prevent *forks* of the information. This will affect the excludability of the artefacts, as, while one may be able to exclude access to one database, if the information can be freely copied elsewhere then this control is lost. It has become prevalent in the age of digital information that when individuals feel there will be no or little chance of punishment for sharing information, they do so on a large scale (Lessig, 2004: 62–66).

The *facilities* constitute how the data, information and knowledge are stored. A suitable facility depends on the properties of the resource (e.g. where this storage is physically located, i.e. under whose control it is) and who is willing to underwrite its costs. In existing participatory-sensing applications the initiator, evaluators and/or analysts provide this infrastructure and pay for it. Alternatives would be distributed or p2p databases where a set of individuals cooperatively provide infrastructure and bear that cost (perhaps being compensated by those who cannot contribute). A p2p system could be one where each individual looks after their own data and responds to requests for it (cf. Global Sensor Network (GSN) (Aberer *et al.*, 2006) and Open Mustard Seed (OMS)⁹), or a robust replicated system where every user has a copy of the whole data set (cf. decentralised version-control system Git¹⁰ and crypto-currency Bitcoin¹¹) or anywhere between these extremes. We may also have different facilities for different artefact types within a single organisation if the artefact quantity and transitivity differ.

We can furthermore consider computational resources as facilities. Evaluators and analysts use their knowledge (i.e. algorithms) with the information stored in the facility to generate new information. This process requires computational resources which are provided by the evaluator or analyst themselves. Alternatively, participants could provision resources or share cost for this computation. The commodification of computational resources makes this very easy to achieve in practice.

4.1.2 Attributes of the community

To identify the community for the knowledge commons, Ostrom and Hess (2007) begin by assessing information users, information providers and information policymakers. *Users* appropriate information, *providers* constitute both providers of artefacts and facilities, and *policymakers* are those who partake in the organisational governance. Each of these groups contain multiple sub-groups, each with different interests and agendas.

The *users* in participatory sensing are those appropriating information in order to apply knowledge to it (*evaluators* and *analysts*) and those appropriating this derived data (*consumers*). As in many cases consumers are also providers as they have a reciprocal relationship with the analysts.

Providers constitute the users providing sensor information (*providers*), users providing both information derived from knowledge and/or the knowledge itself (*evaluators* and *analysts*), and users

⁹ <http://idcubed.org/open-platform/platform/>

¹⁰ <http://git-scm.com>

¹¹ <http://bitcoin.org>

providing infrastructure for *facilities*. This is a diverse group which is likely to involve almost all users in the application. However, it is also one of the most critical groups, as a lack of provision will be a key factor in the success or failure of the organisation. Big data has shown us that the value of information is additive, often exponentially so, and gaining traction—a critical mass of providers—is important. We have seen in Section 3 some methods for incentivising contribution.

Finally, *policymakers* constitute a more diverse and abstract group of community decisionmakers. Any *user* or *provider* can be a *policymaker*, but equally, membership of those groups is not a requirement. Currently, many participatory-sensing applications have a very small number of predefined *policymakers*, and this is in direct violation of Ostrom's third principle—that those affected by the rules can participate in their modification. Therefore, we would like to see larger groups of decisionmakers like in Wikipedia and FOSS. This does not mean that all *policymakers* will have equal power (in terms of capability to enact change), however, policy-making communities can be nested, as we see in social communities and the systems reviewed in Section 3.

4.1.3 Rules-in-use

The rules-in-use dictate what users must, must not or may do in an organisation. The IAD breaks these rules down into three levels: operational, collective-choice and constitutional. Operational rules deal with day-to-day operations regarding the resource, who can provision and appropriate what and when. Collective-choice rules determine how operational rules can be changed, and constitutional rules determine who can participate in collective-choice decisions and how collective-choice rules can be changed.

For a knowledge commons Ostrom and Hess (2007) advocate the creation of rules by allocating bundles of property rights, a method which has also been adopted for the Creative Commons licenses¹². Seven rights were identified (Schlager & Ostrom, 1992):

- Access: the right to enter an area and enjoy non-subtractive benefits. Typically, this right would be for physical access to the property, though this concept can be extended to information commons. In participatory sensing, access would grant the ability to browse or search the repository, but not to extract anything. This can be likened to browsing a library without the right to loan anything.
- Contribution: the right to contribute artefacts to the repository, or right of provision. This can be discriminated by the type of information/knowledge being contributed.
- Extraction: the right to obtain artefacts, or right of appropriation. As with contribution this could be discriminated by information/knowledge type, for example, only certain experts may be allowed to extract sensor information, but everyone can extract the information generated by these experts.
- Removal: the right to remove one's artefacts from the resource.
- Management/participation: 'The right to regulate internal use patterns and transform the resource by making improvements' (Ostrom & Hess, 2007: 52). The first part of this right is a collective-choice right, but the second part is also applied at the operational level. Management of the resource could involve pruning, aggregating or compressing information in order to keep facility costs down.
- Exclusion: 'The right to determine who will have access, contribution, extraction and removal rights and how those rights may be transferred' (Ostrom & Hess, 2007: 53). This is a collective-choice right which controls the level of access for users of the resource.
- Alienation: 'The right to sell or lease extraction, management/participation, and exclusion rights' (Ostrom & Hess, 2007: 53). In this definition the sale of a right prohibits the seller from utilising that right once it is transferred. Therefore, it is not a right which is particularly applicable to information commons.

With these rights we can make rules to describe the operational and most of the collective-choice levels of many organisations managing information and knowledge. Rights give users the institutional *power*, *permission* and/or *obligation* to take actions within the context of a knowledge commons. We can formalise

¹² <http://creativecommons.org>

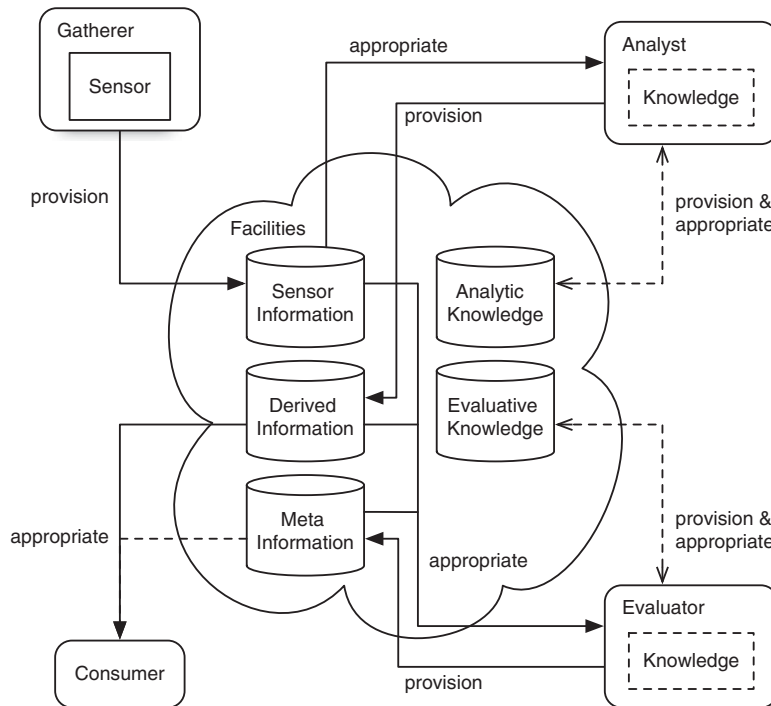


Figure 3 Participatory sensing as a provision and appropriation system with multiple types of information and knowledge. Raw information from sensors is provisioned then knowledge is used to generate multiple different types of information from this. Dotted arrows represent optional actions for that role

these relationships to write rules which are machine readable, such that agents can understand when they are permitted or obliged to perform actions, and when someone has broken the rules (Artikis, 2011).

4.2 Formal characterisation

Having analysed participatory sensing as a knowledge commons, we can now begin to formally characterise such a system. Pitt *et al.* (2012) used the event calculus (EC) (Kowalski & Sergot, 1986) and institutionalised power (Jones & Sergot, 1996) to formally characterise a resource-allocation system and address six of Ostrom's (1990) eight principles. We follow the same methodology, applying it instead to a provision and appropriation system with multiple resource types.

Figure 3 illustrates the participatory-sensing provision and appropriation system with multiple levels of information and knowledge. The same actions, *provision* and *appropriate*, can be used for several different resource types. In addition to those shown in the diagram, facilities and institutions can be provisioned. We assume that we are able to distinguish between each of these resource types and thus tailor rules. In our representation we use simple predicates to make these distinctions.

4.2.1 The EC and institutionalised power

For this formalisation we wish to use a language which is able to represent and reason about action, agency, social constraints and change. We continue with the EC, as we have extended the specification from Pitt *et al.* (2012). In addition, its clarity of exposition and executable specification are useful for this task.

The EC is a logical formalism for representing and reasoning about actions or events and their effects, based on a many-sorted first-order predicate calculus. An EC specification consists of an *action description* which includes axioms that define: action occurrences, using happensAt predicates; the effects of actions, using initiates and terminates predicates; and the values of *fluents*, using initially and holdsAt predicates. Table 3 summarises EC predicates which we use in our specification. Variables start with an uppercase letter

Table 3 Main predicates of the event calculus

Predicate	Meaning
<i>Act</i> happensAt <i>T</i>	Action <i>Act</i> occurs at time <i>T</i>
Initially <i>F = V</i>	The value of fluent <i>F</i> is <i>V</i> at time 0
<i>F = V</i> holdsAt <i>T</i>	The value of fluent <i>F</i> is <i>V</i> at time <i>T</i>
<i>Act</i> initiates <i>F = V</i> at <i>T</i>	The occurrence of action <i>Act</i> at time <i>T</i> initiates a period of time for which the value of fluent <i>F</i> is <i>V</i>
<i>Act</i> terminates <i>F = V</i> at <i>T</i>	The occurrence of action <i>Act</i> at time <i>T</i> terminates a period of time for which the value of fluent <i>F</i> is <i>V</i>

Table 4 Agent actions

Action	Description
<i>provision(A, I, Obj)</i>	Agent <i>A</i> provisions object <i>Obj</i> to institution <i>I</i>
<i>appropriate(A, I, Obj)</i>	Agent <i>A</i> appropriates object <i>Obj</i> from institution <i>I</i>
<i>apply(A, I, Role)</i>	Agent <i>A</i> applies for the role <i>Role</i> in institution <i>I</i>
<i>assign(G, A, I, Role)</i>	Gatekeeper agent <i>G</i> assigns the role of <i>Role</i> in institution <i>I</i> to agent <i>A</i>
<i>report(M, A, I, Reason)</i>	Monitor agent <i>M</i> reports agent <i>A</i> in institution <i>I</i> for the reason given by <i>Reason</i>
<i>appeal(A, I, S)</i>	Agent <i>A</i> appeals the sanction level <i>S</i> in institution <i>I</i>
<i>uphold(C, A, I, S)</i>	Head agent <i>C</i> upholds the appeal by agent <i>A</i> for the sanction level <i>S</i> in institution <i>I</i>

and are assumed to be universally quantified. Note that the underscore ‘_’ denotes the anonymous variable which can stand for (be unified with) any value. Predicates, function symbols and constants start with a lowercase letter. *Fluents* are properties which can have different values at different points in time.

In order to represent and reason about permissions and access control within an institution we need a formalisation of *institutionalised power* (Jones & Sergot, 1996). This term refers to the feature of institutions whereby designated agents, often acting in specific roles, are empowered, obliged or permitted to perform certain actions which in turn may modify certain institutional facts. In EC we formalise these powers as fluents which indicate whether an agent has or had a specific power at a time.

4.2.2 Addressing Ostrom’s principles

Pitt *et al.* (2012) demonstrated that six of Ostrom’s principles could be axiomatised for a resource-allocation problem. There are several differences between this problem and our provision and appropriation system. Therefore, we follow the same methodology but modify the axioms where appropriate. In resource-allocation systems there is a physical resource which is accessed through appropriations, while in our case the agents interact with provisions as well as appropriations. In addition, resource allocation deals with scarce, highly excludable resources, meaning that one agent’s appropriation excludes another agent from doing the same. In the case of information and knowledge, a key feature which we identified in Section 2, is the lack of scarcity and excludability.

The formal characterisation consists of a set of actions which agents can perform in the action arena of the participatory-sensing application, a set of fluents which describe the state of the system at discrete time points, and rules which describe how the agents’ actions affect the fluents. Through this characterisation we enumerate how certain rules can satisfy some of Ostrom’s principles, and may lead to certain outcomes. The rules are sourced from both social systems that we have reviewed and from technical solutions which are available for the virtual domain. Table 4 lists agent actions, Table 5 lists fluents and Table 6 lists the predicates and function symbols.

Table 5 Fluents

Fluent	Values	Description
<i>role_of</i> (<i>A</i> , <i>I</i> , <i>Role</i>)	<i>Boolean</i>	<i>true</i> iff agent <i>A</i> has the role of <i>Role</i> in institution <i>I</i>
<i>provided</i> (<i>A</i> , <i>I</i> , <i>Obj</i>)	<i>Boolean</i>	<i>true</i> iff agent <i>A</i> successfully provisioned <i>Obj</i> to institution <i>I</i>
<i>appropriated</i> (<i>A</i> , <i>I</i> , <i>Obj</i>)	<i>Boolean</i>	<i>true</i> iff agent <i>A</i> successfully appropriated <i>Obj</i> from institution <i>I</i>
<i>applied</i> (<i>A</i> , <i>I</i> , <i>Role</i>)	<i>Boolean</i>	<i>true</i> iff agent <i>A</i> successfully made an application for the role <i>Role</i> in institution <i>I</i>
<i>acMethod</i> (<i>I</i> , <i>Role</i>)	{ <i>none</i> , <i>attribute</i> , <i>discretionary</i> , <i>vote</i> }	Specifies the access-control method for the role <i>Role</i> in institution <i>I</i>
<i>appLimit</i> (<i>I</i> , <i>Role</i> , <i>Rtype</i>)	(<i>Integer</i> , <i>time</i>)	The current rate limit on appropriations of the resource type <i>Rtype</i> for the role <i>Role</i> in institution <i>I</i> . Value is a two-tuple of number of appropriations over a discrete time period
<i>review_score</i> (<i>A</i> , <i>I</i>)	[0...1]	A rating of the agent <i>A</i> 's provisions to the institution <i>I</i>
<i>score_threshold</i> (<i>I</i>)	[0...1]	The threshold of provision quality required by the institution <i>I</i>
<i>appealed</i> (<i>A</i> , <i>I</i> , <i>S</i>)	<i>Boolean</i>	<i>true</i> iff agent <i>A</i> has made an appeal over the sanction level <i>S</i> in institution <i>I</i>
<i>sanction_level</i> (<i>A</i> , <i>I</i>)	<i>Integer</i>	The current sanction level of agent <i>A</i> in institution <i>I</i>
<i>offences</i> (<i>A</i> , <i>I</i>)	<i>Integer</i>	The number of offences committed by agent <i>A</i> in institution <i>I</i>
<i>adrMethod</i> (<i>I</i>)	{ <i>arbitration</i> , <i>mediation</i> , <i>negotiation</i> , ...}	Specifies the dispute resolution method used in institution <i>I</i>
<i>licenceReq</i> (<i>I</i>)	{ <i>none</i> , <i>copyleft</i> , ...}	Specifies the required licence for artefacts in institution <i>I</i>
<i>licence</i> (<i>Obj</i>)	{ <i>none</i> , <i>copyleft</i> , ...}	Gives the current licence type of artefact <i>Obj</i>
pow (<i>A</i> , <i>Action</i>)	<i>Boolean</i>	Agent <i>A</i> has the institutionalised power to perform action <i>Action</i>
per (<i>A</i> , <i>Action</i>)	<i>Boolean</i>	Agent <i>A</i> has the institutionalised permission to perform action <i>Action</i>
obl (<i>A</i> , <i>Action</i>)	<i>Boolean</i>	Agent <i>A</i> is obliged to perform action <i>Action</i>

Table 6 Predicate/function symbols

Predicate/function	Values/range	Description
<i>type</i> (<i>Obj</i>)	{ <i>sensor_information</i> , <i>derived_information</i> , <i>meta_information</i> , <i>analytic_knowledge</i> , <i>evaluative_knowledge</i> }	Determines the artefact type of <i>Obj</i>
<i>countAppropriations</i> (<i>A</i> , <i>I</i> , <i>R</i> , <i>T1</i> , <i>T2</i>)	<i>Integer</i>	Counts the number of appropriations by agent <i>A</i> of resources of type <i>R</i> in institution <i>I</i> between times <i>T1</i> and <i>T2</i>
<i>countProvisions</i> (<i>A</i> , <i>I</i> , <i>R</i> , <i>T1</i> , <i>T2</i>)	<i>Integer</i>	Counts the number of provisions by agent <i>A</i> of resources of type <i>R</i> in institution <i>I</i> between times <i>T1</i> and <i>T2</i>
<i>derivedFrom</i> (<i>S</i> , <i>Alg</i> , <i>O</i>)	<i>Boolean</i>	<i>true</i> iff artefact <i>O</i> was derived from artefact <i>S</i> using <i>Alg</i>

Principle 1: clearly defined boundaries

Defining boundaries for a digital resource is much easier than with physical resources. First, the resource facility is not a pre-existing physical area, it is a virtual portal where the information and knowledge are stored. Access to the resource is much easier to control due to the availability of authentication mechanisms (e.g. public key infrastructure, identity management, etc.) which allow the verification of users accessing the resource. Taking open-source software as an example, the facilities (the version-control

system, bug/issue tracker and mailing list) have fine-grained access control, preventing unauthorised access to the resource.

This access control can be implemented in a distributed system using a role-based system. In addition to the operational user roles for participatory sensing we add roles for institutional tasks, in this case an agent with the role of *gatekeeper*. This role empowers this agent to assign roles according to the specified access-control method for the institution. An agent who applies for a role can be subsequently assigned to the role by the *gatekeeper*, provided the conditions of entry for the chosen access-control method are satisfied.

$apply(A, I, Role)$ initiates $applied(A, I, Role) = true$ at $T \leftarrow$
 $\mathbf{pow}(A, apply(A, I, Role)) = true$ holdsAt T
 $\mathbf{pow}(A, applied(A, I, Role) = true$ holdsAt $T \leftarrow$
 $role_of(A, I, Role) = false$ holdsAt T
 $assign(G, A, I, Role)$ initiates $role_of(A, I, Role) = true$ at $T \leftarrow$
 $\mathbf{pow}(G, assign(G, A, I, Role)) = true$ holdsAt T
 $\mathbf{pow}(G, assigned(A, I, Role) = true$ holdsAt $T \leftarrow$
 $applied(A, I, Role) = true$ holdsAt $T \wedge$
 $acMethod(I, Role) = discretionary$ holdsAt $T \wedge$
 $role_of(G, I, gatekeeper) = true$ holdsAt T

The axioms above give the example of discretionary assignment, where the gatekeeper can decide whether to assign a role or not. An agent A performing the *apply* action initiates a period of time during which the fluent $applied(A, I, Role)$ is *true* if it is empowered to perform that action; it is empowered to perform that action if the agent does not already occupy this role. Similarly, an agent G can make the institutional fact (fluent) true that an agent A is assigned to a role $Role$ if it is empowered to perform the *assign* action. This is the case if agent A has applied for the role, agent G occupies the role of gatekeeper and the access-control method is discretionary.

Note that we could also use attribute-based access control, where the gatekeeper may only assign the role if the applicant satisfies certain conditions. We could allow agents to vote on new applicants, or, if we want more open access to certain roles, we can compel the gatekeeper to assign a role to all applicants:

$\mathbf{obl}(G, assign(G, A, I, Role)) = true$ holdsAt $T \leftarrow$
 $Role = gatherer \wedge$
 $applied(A, I, Role) = true$ holdsAt $T \wedge$
 $acMethod(I, Role) = none$ holdsAt $T \wedge$
 $role_of(G, I, gatekeeper) = true$ holdsAt T

Using roles we can define who will be *empowered*, that is, have the right, to provision and appropriate certain resources, according to Figure 3.

$\mathbf{pow}(A, provision(A, I, Obj)) = true$ holdsAt $T \leftarrow$
 $role_of(A, I, gatherer) = true$ holdsAt $T \wedge$
 $type(Obj) = sensor_information$
 $\mathbf{pow}(A, provision(A, I, Obj)) = true$ holdsAt $T \leftarrow$
 $role_of(A, I, analyst) = true$ holdsAt $T \wedge$
 $type(Obj) = derived_information$
 $\mathbf{pow}(A, provision(A, I, Obj)) = true$ holdsAt $T \leftarrow$
 $role_of(A, I, evaluator) = true$ holdsAt $T \wedge$
 $type(Obj) = meta_information$

The axioms for appropriation have the same structure, simply changing the *provision* action to *appropriation* and the role and resource type where applicable. Figure 3 specifies which roles can appropriate each resource type.

Principle 2: congruence between appropriation and provision rules and local conditions

The appropriation and provision rules must be relevant for the local conditions, and therefore are often quite application specific. These rules cover time, place, technology and quantity of resource units appropriated or provisioned. The principle is typically violated when the rules cause over-regulation of an abundant resource or under-regulation of a scarce one.

In participatory sensing, scarcity may arise in access to the resource, for example, through an excessive quantity of requests going to the facility. One method of dealing with such congestion would be to rate limit requests. If we take an example of a consumer appropriating road traffic information from the facility, an incorrectly implemented rate limit could severely limit the usefulness of the resource. As the nature of resource usage in this case is bursty (large number of requests in one go when route planning, then very few for a period of time afterwards), a quota system of x requests per minute is extremely more restrictive than $60 \times 24 \times x$ requests per day, despite corresponding to the same overall quota.

Rate limits can be expressed as follows:

$$\begin{aligned} \text{pow}(A, \text{appropriate}(A, I, \text{Obj})) = \text{true holdsAt } T \leftarrow \\ \text{type}(\text{Obj}) = \text{Rtype} \wedge \\ \text{role_of}(A, I, \text{Role}) = \text{true holdsAt } T \wedge \\ \text{appLimit}(I, \text{Role}, \text{Rtype}) = (\text{Limit}, \text{Period}) \text{ holdsAt } T \wedge \\ \text{countAppropriations}(A, I, \text{Rtype}, T - \text{Period}, T) < \text{Limit} \end{aligned}$$

In Pitt *et al.* (2012) this principle is addressed by looking at when it is appropriate to demand resources. Here, the patterns of interaction are likely to be more free form. However, in some applications we may want to force a more symmetric relationship to compel agents to contribute to the resource:

$$\begin{aligned} \text{pow}(A, \text{appropriate}(A, I, \text{Obj})) = \text{true holdsAt } T \leftarrow \\ \text{countProvisions}(A, I, \text{Rtype}, 0, T) = P \wedge \\ \text{countAppropriations}(A, I, \text{Rtype}, 0, T) \leq P \end{aligned}$$

Principle 3: collective-choice arrangements

‘Most individuals affected by the operational rules can participate in modifying the operational rules’—this is certainly not the case in current participatory-sensing applications. To achieve this, a framework is required to allow users to propose and agree on new rules and rule changes in the context of the application. Many such frameworks exist for human organisations, and some have been formalised to be usable for agent systems (Artikis *et al.*, 2005), including robust voting mechanisms (Pitt *et al.*, 2006).

Pitt *et al.* (2012) formulate a voting mechanism for collective-choice arrangements through effective enfranchisement of members of the institution. This enfranchisement is achieved by empowering agents to vote and granting an entitlement associated with this right, ensuring a correct result is declared from a ballot. This entitlement is implemented as an obligation for a particular agent to declare the result of a vote according to an agreed-upon winner-determination method. This general voting system is fit for purpose across many domains, indeed it is just a formalisation of protocols used in many social systems (Pitt *et al.*, 2006), and so we can adopt it as-is for this domain.

An issue remains, though, of how or whether agents can make reasoned decisions about rules, let alone create new rules, without human intervention. This problem has been explored with the implementation of the game of nomic played by agents (Holland *et al.*, 2013). Normative synthesis (Morales *et al.*, 2013) offers an approach to the generation of rules. We can use the concepts of learning, self-simulation and self-awareness to provide some reasoning about the consequences of rules.

Principle 4: monitoring

Monitoring can be implemented, provided sufficient auditing capabilities are available (Pitt *et al.*, 2012). Wikipedia sets a good example here: all users are able to monitor each other through the edit history, which gives a fine-grained breakdown of what has been contributed and by whom. However, this level of transparency may seem too much for some, as full visibility of users actions will come with privacy concerns. In technical systems, this auditing can be done via a log of institutional actions performed by agents. This creates a narrative, like the ones which can be processed with the EC, and the process of monitoring is simply looking for actions which are performed when the actor did not have the permission to perform that action at the time.

Effective monitoring can occur in several different ways. First, some systems can be designed such that monitoring is a side-effect arising from normal use. Ostrom (1990) observed this in fisheries which allocated areas of a lake to different individuals on a rota. If an individual, upon arriving in their allotted zone, found someone else using it they would obviously have detected the violation of the rules and be able

to report it. Second, individuals can be incentivised to monitor for violations. Unlike in the first case this involves some additional effort on the individual's part. This incentivisation can be in the form of a reward for finding violations or payment for performing audits. Finally, the institution can arrange for an independent third party to monitor the system, and pay for this by collecting contributions from members.

In the participatory-sensing scenario our access-control procedures provide protection over access to the resource, therefore monitoring is more concerned with the content of provisions: the quality of information and knowledge provided, and whether a user is permitted to provision a particular artefact. These requirements touch on different concepts and thus need to be handled differently.

In the case of the quality of provisioned information and knowledge, the actions of *evaluators* can be a form of monitoring. Meta information provisioned by *evaluators* can provide indications of the trustworthiness and accuracy of sensor information, derived information and even other meta information (see Figure 3). This process is akin to review, rating and trust systems seen in e-commerce and other online applications. Provided that there are *evaluators* provisioning this kind of information in the institution, this is a form of side-effect monitoring. Incentives can also be created to encourage more provision of meta information.

We may assign a *monitor* role to an agent who is required to monitor for low-quality provisions. To do this, they can simply look for information provisioned by *evaluators* that suggests an agent's data is below some quality threshold. If this is the case, it will trigger a report of that agent (indeed, the monitor may even be obliged to report the offence):

$report(M, A, I, O)$ initiates $monitored_offence(A, I) = O$ at $T \leftarrow$
 $pow(M, report(M, A, I, O)) = true$ holdsAt T
 $pow(M, report(M, _ , I, _)) = true$ holdsAt $T \leftarrow$
 $role_of(M, I, monitor) = true$ holdsAt T
 $obl(M, report(M, A, I, baddata)) = true$ holdsAt $T \leftarrow$
 $review_score(A, I) = S$ holdsAt $T \wedge$
 $score_threshold(I) = Thres$ holdsAt $T \wedge$
 $S < Thres$

The other concern, provision of artefacts when the individual does not have the permission to, is more difficult to monitor. There are two possible malicious actions here which we may want to monitor: an agent provisioning an artefact which has already been provisioned previously; and an agent provisioning an artefact which was generated by another, but has not been provisioned to this institution. In both cases the agent acts maliciously in order to gain benefits of provision without having to generate an artefact legitimately.

However, any open system facing the problem of non-compliant behaviour has to consider the related problems: how are these actions detectable, and if so, is it even worthwhile to monitor? It has been shown (e.g. Balke *et al.*, 2013) that the cost of monitoring can outweigh the benefits of preventing or punishing non-compliant behaviour. This can be particularly acute in any system with endogenous resources, where the cost of monitoring has to be 'paid for' from the same CPR that is itself being monitored.

Principle 5: graduated sanctions

In order to have graduated sanctions, we require multiple levels of punishment. First, as we have control over the system's boundaries through access control, we can implement temporary and permanent banishment as possible sanctions. Second, we can revoke certain provision and/or appropriation rights. Finally, if there exists some form of micro-payment system, we can levy fines against members. Therefore, given that we have several forms of sanction which differ in both context and severity, we should be able to implement graduated sanctions.

Pitt *et al.* (2012) use a sanction level which increases on each new offence for that agent. At higher sanction levels a nominated agent (in this case a role of *head* assumes this responsibility) is empowered (and may be obliged) to impose stronger sanctions. This method assumes equal weight to each offence, though we may define some offences which can cause multiple level increases.

Principle 6: conflict-resolution mechanisms

In Ostrom's work this principle was mainly used to deal with conflict due to ambiguities arising from the way that rules were written. Rules in a computational system (if written correctly) should not have any ambiguities, so theoretically there should be no need for this principle.

However, as we have collective-choice arrangements, we have a system which has a rule which says that the rules should be changeable. Suber (1990) hypothesised that such a system will inevitably lead to paradox or gaps in the rule set. Given this, the question is whether to give agents recourse to retrospectively claim for losses due to such issues through conflict resolution.

In Pitt *et al.* (2012) alternative dispute resolution (ADR) is used as a fast, cheap and effective conflict-resolution mechanism. A simple appeals procedure is given, allowing one to appeal sanctions. A more complete specification of an ADR protocol is given in Pitt (2012b). We will simply give an overview of the semantics of the appeal process, and the method of deciding whether an appeal is valid and should be upheld. How one argues one's case is beyond the scope of this work, but there are several approaches to choose from.

For example, the agent occupying the requisite role (*head*) is empowered to sanction another agent if an offence has been reported (see Principle 4), and this action increases the sanction level against this agent.

$sanction(H, I, A, O)$ initiates $sanction_level(A, I) = S1$ at $T \leftarrow$
 $\mathbf{pow}(H, sanction(H, I, A, O)) = true$ holdsAt $T \wedge$
 $sanction_level(A, I) = S$ holdsAt $T \wedge$
 $S1 = S + 1$
 $\mathbf{pow}(H, sanction(H, I, A, O)) = true$ holdsAt $T \leftarrow$
 $role_of(H, I, head) = true$ holdsAt $T \wedge$
 $monitored_offence(A, I) = O$ holdsAt T

An agent may appeal against its sanction level:

$appeal(A, I, S)$ initiates $appealed(A, I, S) = true$ at $T \leftarrow$
 $\mathbf{pow}(A, appeal(A, I, S)) = true$ holdsAt T
 $\mathbf{pow}(A, appeal(A, I, S)) = true$ holdsAt $T \leftarrow$
 $role_of(A, I, _)$ holdsAt $T \wedge$
 $sanction_level(A, I) = S$ holdsAt T

The *head* may then lower the sanction level if the appeal is successful. If the ADR method for the institution is arbitration, then the *head* agent can decide the outcome of the appeal process itself.

$uphold(H, A, I, S)$ initiates $sanction_level(A, I) = S1$ at $T \leftarrow$
 $\mathbf{pow}(H, uphold(H, A, I, S)) = true$ holdsAt $T \wedge$
 $sanction_level(A, I) = S$ holdsAt $T \wedge$
 $S1 = S - 1$
 $\mathbf{pow}(H, uphold(H, A, I, S)) = true$ holdsAt $T \leftarrow$
 $role_of(H, I, head) = true$ holdsAt $T \wedge$
 $appealed(A, I, S) = true$ holdsAt $T \wedge$
 $adrMethod(I) = arbitration$ holdsAt T

An appeal which is upheld could also initiate review and retraction of the monitored offence (axioms not specified here).

Principle 7: minimal recognition of rights to organise

This principle states that external entities should not interfere with the rule-making capabilities of the organisation. Heller (1998) wrote that over-regulation leads to under-use of a CPR, which he called a 'Tragedy of the Anti-Commons'. The use of copyright law has the capability of subverting the internal rules of an information commons.

If users attempt to assert their intellectual property rights in participatory-sensing applications they could gain more rights than the organisation initially permitted them with, or gain the power to prevent others performing actions. For example, an analyst may determine that information he generates and provisions is his intellectual property. If an evaluator then wants to appropriate this information and generate some meta information about it, this could be seen as a derivative work, and therefore require the permission of the analyst. Scenarios like this need to be prevented as they stifle productivity (evaluators will be cautious about what information they appropriate for fear of infringement) and undermine the authority of the organisation.

Luckily, there has already been significant success in dealing with these issues. Licenses have been used to ensure that certain rights are released when information is provisioned. In open-source software Copyleft (Stallman, 1999) is an example which uses copyright law to its advantage to provide additional benefits in return. The Creative Commons¹³ licenses simplify the process of inferring what rights one has regarding specific content. Care must be taken to choose appropriate licenses for the different forms of information and knowledge generated from the sensing application.

In order to protect itself, an institution may require that contributions are licensed in a particular way:

$$\begin{aligned} \text{per}(A, \text{provision}(A, I, \text{Obj}) \text{ holdsAt } T \leftarrow \\ \text{licenceReq}(I) = L \text{ holdsAt } T \wedge \\ \text{licence}(\text{Obj}) = L \text{ holdsAt } T \end{aligned}$$

We can also formalise aspects of some open-source licences. For example, the right to redistribute; the right to read source code (which we interpret as an obligation to provision any knowledge which Copylefted information is derived with); and the obligation to provision any derivative works back to the institution(s) it originated from.

$$\begin{aligned} \text{provision}(A, I, \text{Obj}) \text{ initiates } \text{obl}(A, \text{provision}(A, I, \text{Alg})) \text{ at } T \leftarrow \\ \text{derivedFrom}(_, \text{Alg}, \text{Obj}) = \text{true} \wedge \\ \text{licence}(\text{Obj}) = \text{copyleft} \text{ holdsAt } T \wedge \\ \text{provisioned}(_, I, \text{Alg}) = \text{false} \text{ holdsAt } T \\ \text{per}(A, \text{provision}(A, _, \text{Obj}) \text{ holdsAt } T \leftarrow \\ \text{licence}(\text{Obj}) = \text{copyleft} \text{ holdsAt } T \\ \text{provision}(A, I_2, \text{Obj}) \text{ initiates } \text{obl}(A, \text{provision}(A, I, \text{Obj})) \text{ at } T \leftarrow \\ \text{derivedFrom}(\text{Source}, _, \text{Obj}) = \text{true} \wedge \\ \text{licence}(\text{Source}) = \text{copyleft} \text{ holdsAt } T \wedge \\ \text{provisioned}(_, I, \text{Source}) = \text{true} \text{ holdsAt } T \wedge \\ I \neq I_2 \end{aligned}$$

Note that the presence of a derivative work can only be detected once it is provisioned elsewhere. In addition, what constitutes *derivedFrom* can be difficult to pinpoint in many cases (Lessig, 2004).

4.3 Evaluation

We now move on to the evaluative aspects of the IAD. We perform this evaluation with respect to current participatory-sensing applications and our proposed self-organising approach based on Ostrom's design principles. We first look at possible outcomes arising from applications, then determine a set of criteria to evaluate them with.

4.3.1 Outcomes

We have talked of some current outcomes seen in participatory-sensing applications, namely enclosure and inequity. This analysis should help us understand why we see these, but we will also look at other potential outcomes we may observe from different organisational approaches.

- Enclosure: what we see in often participatory-sensing and other applications based around information collection is the enclosure of the database. The data collector wants to keep the information to try and derive as much value from it themselves as possible. This is first inefficient; we see that, when data is made open much more total value is derived than any individual could achieve by themselves (Shadbolt *et al.*, 2012). Second, it deprives the original information provider, the user, from their own property.
- Access: the opposite to the negative outcome of enclosure is access. This has a beneficial effect of having an information and knowledge resource where there was none before. The benefits of access increase the more open the information is.

¹³ <http://creativecommons.org>

- Equity: equity and inequity are outcomes based on the fairness of the allocation of the resource, in terms of information and knowledge availability. Fairness can be evaluated in several different ways, many of which are subjective. We can use a more objective measure by assessing whether each user is compensated according to the value of their input to the system.
- Pollution: the insertion of incorrect information can affect the quality of the resource. This information could be contributed accidentally, for example, a gatherer, whose sensor is faulty, provisioning information or maliciously, for example, doctoring information to gain a competitive edge (Oreskes & Conway, 2010). While the *evaluator's* role can be partly to validate information and guard against pollution, this can be very difficult to practice, as has been seen from the manipulation of Facebook likes, product reviews and Google PageRank. Pollution can be countered by trust and reputation systems, or redundancy in the information-gathering process.
- Degradation and depletion: these outcomes are caused by either the removal of information or insufficient provision of new information. These relate to the endurance of a system. They will occur either when the system does not react to changing external or internal conditions, or if the system evolves into a 'bad' state, causing participation to be no longer worthwhile for some agents. Lack of reactivity to changing conditions is likely to be caused by a lack of flexibility in the institution, either by not having mechanisms in place for reacting, or such mechanisms being blocked by entrenched interests or inappropriate decision-making processes (North, 1990). Similarly, 'bad' states can be caused by malicious groups of agents manipulating the institution for short-term benefit.

4.3.2 Evaluative criteria

Finally, to assess the efficacy of the organisational structure on the outcomes of the system, we specify a set of evaluative criteria. We evaluate outcomes and participant's interactions with measurable values, allowing objective comparison of systems and organisations. We take general criteria applied to knowledge commons, as proposed by Ostrom and Hess (2007), describe how we can measure them, and then discuss the effect we expect our framework will have on the values of the criteria. These criteria are: increase in knowledge, sustainability and preservation, participation standards, economic efficiency and equity through fiscal equivalence.

Certainly a successful participatory-sensing application will lead us to have increased knowledge, in terms of more information derived from a large pool of sensor information. However, we may also see an improvement in the knowledge being applied to the information by analysts and evaluators. In academia it is generally accepted that open collaboration can improve the quality of knowledge generated. Allowing participants to share their knowledge and then allow others to modify and improve it is an efficient way of increasing knowledge, and one which could be leveraged in participatory sensing. Therefore, the increase in knowledge can be measured as the quantity and quality of information generated by the application.

Quantity of provision is achieved through high participation standards and incentivisation. In our survey we saw several methods of incentivising provision, and our self-organising approach makes several of these methods available as rules-in-use. Quality is maintained through effective monitoring. We have discussed how monitoring can be achieved in our system and be responsive to increases in malicious actions using Principles 4 and 5. The differentiating factor between our proposed framework and other participatory-sensing applications is the pool of knowledge available in parallel with information. This permissive sharing of knowledge, allowing for iterative improvements by many disparate contributors, is an effective method for the increase of knowledge, as demonstrated by academia, FOSS and Wikipedia.

Sustainability and preservation of the organisation can be measured in terms of how long the system endures before participation ceases or falls below a certain threshold (where the system is not longer economically viable). To achieve sustainability it must be possible to meet the needs of the majority of participants over the long term, and also to react to changes in conditions. For example, combating a rise in malicious actors before the information pool becomes over-polluted.

Linked to sustainability are participation standards. This relates to the number of participants in the system (and in each user role), as well as the level of contribution from each of those participants. Low levels of participation or asymmetric participation (i.e. appropriating more than one provides) may be

interpreted as free-riding, depending on the system's rules. Free-riding is a frequent negative behaviour in the commons and is likely to affect participation standards if not controlled.

Through collective choice, a self-organising institution should be more sustainable than dictatorial institutions, as each agent is properly enfranchised and is able to address any issues they have with the current rules. Equally effective self-organisation should imply good reactivity to changing conditions. However, we must consider the possible implications of inappropriate collective-choice procedures which can cause deadlocks. Voting mechanisms can be manipulated, or unrealistic win criteria may be required for the passing of motions (Pitt *et al.*, 2011). In social systems we can identify these issues and attempt to work around them, but in computational systems collective choice will simply grind to a halt. Therefore, we must be careful in the supply of institutions that these issues are considered.

Participation standards can be maintained through both enfranchisement and incentivisation. Participatory sensing also has a promotional aspect to it, which helps to gain and maintain participation. For example, OpenSignal largely relies on media publicity to attract users. This is an aspect where open access to information helps. Users are more likely to engage with an organisation if they initially get something for free. This gives them an idea of the product and its value, and may then lead to a more active role in the institution. Both Wikipedia and FOSS projects work based on this concept, however, they only need a very small proportion of appropriators to provision to the knowledge pool.

It is also possible to measure the economic efficiency of a participatory-sensing application. We can count the utility generated by the consumers through the use of information appropriated and compare this to the total cost of maintaining the resource. This cost contains facility costs and the individual costs incurred by participants in information gathering and processing.

If we compare centralised with self-organising approaches according to this metric, it is difficult to determine which will perform better. A self-organising institution requires additional facilities for the execution of the governance layer of the application. Our position is, however, that a self-organising institution will be able to generate more value under the same conditions, due to the outcomes we have discussed, than a dictatorial one. In order to be more efficient overall, this increase in generated value must outweigh the additional operating costs.

Finally, we can assess equity through fiscal equivalence. This is a measure of the ratio of effort put in versus benefit extracted from a system. An observation from other knowledge commons is that consumers often benefit for little or no effort (e.g. Wikipedia, open-access publishing, FOSS). However, this benefit is not at a cost to the knowledge providers—in these examples consumption by others usually benefits the providers, helping them achieve fiscal equivalence—but is a form of redistributive equity. In participatory sensing there is not such a strong benefit to information providers from its consumption so we focus more on fiscal equivalence and equality.

4.4 Supply of a participatory-sensing commons

The final question we must answer is how we can enable the supply of an institution as we have described for participatory sensing. Our requirements are a platform which can be easily deployed and is capable of providing a technical means to administer the rules-in-use needed to satisfy Ostrom's principles, and therefore create a sustainable institution around the participatory-sensing application. We review the suitability of existing platforms for use with our framework, and advocate one which we believe fits our criteria best.

There exist several platforms specifically for the deployment of participatory sensing. CarTel (Hull *et al.*, 2006) is a centralised system designed to accommodate heterogeneous sensor data and provides a Web-based portal for data access. SenseWeb (Kansal *et al.*, 2007) provides a shared resource for concurrent sensing applications with a centralised coordinator directing information acquisition and aggregation. It enables users demanding information to provide incentives for those who are able to provide it. The GSN (Aberer *et al.*, 2006) is a p2p network of sensor nodes which can advertise and respond to queries for data streams across the network. G-Sense (Perez *et al.*, 2010) is a platform which uses a hybrid architecture to achieve scalability, using local centralised hubs which collaborate with a global p2p network. Demirbas *et al.* (2010) leverage Twitter as infrastructure for information provision using encoded tweets.

In general, similar to what we found in our review in Section 3, these platforms are dealing with a particular technical challenge of participatory sensing. Therefore, these platforms do not specify how rules for membership, provision and appropriation quantities, monitoring and accountability can be implemented. In p2p platforms such as GSN, individuals may be able to control access to their own data by selectively honouring requests, however, there is no governance oversight to keep consistency and fairness. The use of Twitter as an independent third-party infrastructure provider provides a more transparent operation, but sacrifices all control over access to the resource.

OMS¹⁴ is an open-source project which aims to give users control over their data in the cloud through a technical architecture based on trust and the formation of self-organising groups. A key feature is a rule engine governing access to data on a per-user basis. Such a platform could enable the formation of institutions for participatory sensing with transparency over information access and use, but also with overrides for users. The OMS architecture is p2p, enabling sharing of facility costs by design, and offering scalability and data privacy. For these reasons we see OMS as an implementation avenue for our framework.

5 Summary and conclusions

This paper has articulated the usefulness of the participatory-sensing paradigm and the shortfalls arising from the lack of consideration of the impact of governance in current applications in this area. Our approach aims to find a method of democratising the large-scale aggregation of user-generated information, with participatory sensing as an example of this process. As these systems are digital, and often online, using electronic institutions and algorithmic governance enables us to bring both the enforcement of rules, as well as governance determining what the rules are, into one place. It also enables autonomous agents to interact with the governance layer along with human actors. In future work, we envision enabling this interaction through serious games and gamification, with an emphasis on the visualisation of the 'state' of the common pool and the rule (see Bourazeri *et al.*, 2012 for an example using SmartMeters with energy as the common pool), and through design contractualism, whereby design decisions and governance models are manifested both in the code and in the interface (Pitt, 2012a).

In summary, the two main contributions are, namely:

- Data clouds in open participatory-sensing applications can be construed as information and knowledge commons and thus characterised by provision and appropriation actions.
- A system for access control (i.e. provision and appropriation) in participatory-sensing applications can be designed according to Ostrom's institutional design principles for self-governing institutions and formally specified in an action language.

We presented in Section 2 the literature on how information and knowledge can be seen as a commons, and, being a system with the purpose of gathering information in order to generate new knowledge, that participatory sensing enables the creation of information and knowledge commons. Following a review of participatory-sensing applications we were able to characterise a generic participatory-sensing application as a provision and appropriation system in Section 4. We then formally defined a framework for the management of an information and knowledge commons in participatory sensing.

We have two primary motivations for why such a framework is needed in this domain. First, through our review of participatory sensing in Section 3, we identified a lack of governance supplied for applications. Given that participatory sensing is a knowledge commons, appropriate governance, as stressed by Ostrom's (1990) work, is important for its successful management. The success of certain knowledge commons such as Wikipedia and FOSS, which have been retrospectively shown to conform to the principles which Ostrom proposed, further reinforces this point. Thus, as information and knowledge is gathered more and more on the Internet, it is important that governance is supplied in order to fairly protect the interests of all involved parties and stakeholders.

¹⁴ <http://idhypercubed.org>

Second, the digitisation of information and knowledge has caused a shift in its properties as a good (Hess & Ostrom, 2003). However, legal interpretations have not yet been fully updated to this new reality, leading to exploitative and restrictive consequences (Lessig, 2004). We concluded that this gives the knowledge creator complete control over what class of good their knowledge falls in to, and leads to a dilemma where knowledge is enclosed for short-term benefits and a negative long-term outcome.

With the presentation of this framework we are taking the first steps to the supply of institutional governance for managing information and knowledge commons in participatory-sensing applications. What remains to do is a quantitative comparison of the different rule permutations which can be instantiated on top of our framework and an objective comparison between centralised and community governance. Our framework is flexible enough to be able to represent a centralised governance as well as the self-organising approach which is required for Ostrom's principles. It is the aim of our future work to answer these questions.

In conclusion, we have shown that consideration of governance and supply thereof is important in order to achieve the potential of the participatory-sensing paradigm. By drawing on the management of the commons, we can create efficient and empowering institutions around information and knowledge resources, which can exploit the power of open data. In addition, given increasing concerns over data acquisition and data privacy in the digital age, our general representation of a system for the provision and appropriation of information could have applications beyond those which we have discussed here.

Acknowledgements

The first author is supported by EPSRC Studentship Grant No. EP/P505550/1.

References

- Aberer, K., Hauswirth, M. & Salehi, A. 2006. A middleware for fast and flexible sensor network deployment. In *Proceedings of the 32nd International Conference on Very Large Data Bases*, 1199–1202.
- Artikis, A. 2011. Dynamic specification of open agent systems. *Journal of Logic and Computation* **22**(6), 1301–1334.
- Artikis, A., Kamara, L., Pitt, J. & Sergot, M. 2005. A protocol for resource sharing in norm-governed ad hoc networks. In *Declarative Agent Languages and Technologies II*, Leite, J., Omicini, A., Torroni, P. & Yolum, P. (eds). Lecture Notes in Computer Science **3476**, 221–238, Springer.
- Balke, T., de Vos, M. & Padget, J. 2013. I-ABM: combining institutional frameworks and agent-based modelling for the design of enforcement policies. *Artificial Intelligence and Law* **21**(4), 371–398.
- Bollier, D. 2007. The growth of the commons paradigm. In *Understanding Knowledge as a Commons: From Theory to Practice*, Hess, C. & Ostrom, E. (eds). MIT Press, 27–40.
- Bourazeri, A., Pitt, J., Almajano, P., Rodriguez, I. & López-Sánchez, M. 2012. Meet the meter: visualising smartgrids using self-organising electronic institutions and serious games. In *Sixth IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops (SASOW)*, 145–150.
- Buckingham Shum, S., Aberer, K., Schmidt, A., Bishop, S., Lukowicz, P., Anderson, S., Charalabidis, Y., Domingue, J., Freitas, S., Dunwell, I., Edmonds, B., Grey, F., Haklay, M., Jelasity, M., Karpištsenko, A., Kohlhammer, J., Lewis, J., Pitt, J., Sumner, R. & Helbing, D. 2012. Towards a global participatory platform. *The European Physical Journal Special Topics* **214**(1), 109–152.
- Burke, J. A., Estrin, D., Hansen, M., Parker, A., Ramanathan, N., Reddy, S. & Srivastava, M. B. 2006. Participatory sensing. In *International Workshop on World-Sensor-Web*, 117–134.
- Christin, D., Reinhardt, A., Kanhere, S. S. & Hollick, M. 2011. A survey on privacy in mobile participatory sensing applications. *Journal of Systems and Software* **84**(11), 1928–1946.
- Costa, P. M., Pitt, J., Falcão e Cunha, J. & Galvão, T. 2012. Cloud2Bubble: enhancing quality of experience in mobile cloud computing settings. In *Proceedings of the Third ACM Workshop on Mobile Cloud Computing and Services*, 45–52.
- Davenport, T. & Prusak, L. 2000. *Working Knowledge: How Organizations Manage What They Know*, Harvard Business School Press.
- Demirbas, M., Ali Bayir, M., Akcora, C. G., Yilmaz, Y. S. & Ferhatosmanoglu, H. 2010. Crowd-sourced sensing and collaboration using Twitter. In *2010 IEEE International Symposium on a World of Wireless Mobile and Multi-media Networks (WoWMoM)*, 1–9.
- Deng, L. & Cox, L. P. 2009. LiveCompare: grocery bargain hunting through participatory sensing. In *Proceedings of the 10th Workshop on Mobile Computing Systems and Applications – HotMobile '09*, ACM Press, 1–6.
- Eriksson, J., Girod, L., Hull, B., Newton, R., Madden, S. & Balakrishnan, H. 2008. The pothole patrol: using a mobile sensor network for road surface monitoring. In *Proceedings of the 6th International Conference on Mobile Systems, Applications, and Services (MobiSys 2008)*, Breckenridge, CO, USA, June 17–20, 2008; pp. 29–39.

- Famiglietti, A. 2011. The right to fork: a historical survey of de/centralization in Wikipedia. In *Critical Point of View: A Wikipedia Reader*, Lovink, G. & Tkacz, N. (eds). Institute of Network Cultures, 296–308.
- Forte, A., Larco, V. & Bruckman, A. 2009. Decentralization in Wikipedia governance. *Journal of Management Information Systems* **26**(1), 49–72.
- Fuster Morell, M. 2010. *Governance of Online Creation Communities: Provision of Infrastructure for the Building of Digital Commons*. European University Institute, PhD thesis.
- Fuster Morell, M. 2011. The Wikimedia foundation and the governance of Wikipedia's infrastructure: historical trajectories and its hybrid character. In *Critical Point of View: A Wikipedia Reader*, Lovink, G. & Tkacz, N. (eds). Institute of Network Cultures, 325–341.
- Hardin, G. 1968. The tragedy of the commons. *Science* **162**(3859), 1243–1248.
- Hasenfratz, D., Saukh, O., Sturzenegger, S. & Thiele, L. 2012. Participatory air pollution monitoring using smartphones. In *Mobile Sensing: From Smartphones and Wearables to Big Data*. http://research.microsoft.com/en-us/um/beijing/events/ms_ipsn12/papers/ms_ipsn-hasenfratz.pdf, 5pp.
- Heller, M. A. 1998. The tragedy of the anticommons: property in the transition from marx to markets. *Harvard Law Review* **111**(3), 621–688.
- Hess, C. 2000. Is there anything new under the sun?: a discussion and survey of studies on new commons and the Internet. In *Constituting the Commons: Crafting Sustainable Commons in the New Millennium, the Eighth Biennial Conference of the International Association for the Study of Common Property*. <http://hdl.handle.net/10535/384>, 27pp.
- Hess, C. & Ostrom, E. 2003. Ideas, artifacts, and facilities: information as a common-pool resource. *Law and Contemporary Problems* **111**(1998), 111–145.
- Hess, C. & Ostrom, E. 2006. *Understanding Knowledge as a Commons*, MIT Press.
- Holland, S., Pitt, J., Sanderson, D. & Busquets, D. 2013. Reasoning and reflection in the game of nomic: self-organising self-aware agents with mutable rule-sets. In *7th IEEE International Conference on Self-Adaptation and Self-Organizing Systems Workshops, SASOW*, 101–106.
- Hull, B., Bychkovsky, V., Zhang, Y., Chen, K., Goraczko, M., Miu, A., Shih, E., Balakrishnan, H. & Madden, S. 2006. CarTel: a distributed mobile sensor computing system. In *4th ACM SenSys*. pp. 125–138.
- Jones, A. J. I. & Sergot, M. 1996. A formal characterisation of institutionalised power. *Logic Journal of IGPL* **4**(3), 427–443.
- Kanhere, S. 2013. Participatory sensing: crowdsourcing data from mobile smartphones in urban spaces. In *Distributed Computing and Internet Technology, 9th International Conference, (ICDCIT) 2013*. Bhubaneswar, India, February 5–8, pp. 19–26.
- Kansal, A., Nath, S., Liu, J. & Zhao, F. 2007. SenseWeb: an infrastructure for shared sensing. *Multimedia, IEEE* **14**(4), 8–13.
- Kowalski, R. & Sergot, M. 1986. A logic-based calculus of events. *New Generation Computing* **4**(1), 67–95.
- Lessig, L. 2004. *Free Culture: How Big Media Uses Technology and the Law to Lock Down Culture and Control Creativity*, Penguin Press.
- Machlup, F. 1983. Semantic quirks in studies of information. In *The Study of Information: Interdisciplinary Messages*, Machlup, F. & Mansfield, U. (eds). Wiley, New York, pp. 641–671.
- Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C. & Byers, A. H. 2011. *Big Data: The Next Frontier for Innovation, Competition, and Productivity*. Technical report, June, McKinsey Global Institute.
- Mathur, S., Jin, T., Kasturirangan, N., Chandrasekaran, J., Xue, W., Gruteser, M. & Trappe, W. 2010. Parknet. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services – MobiSys '10*, ACM Press, 123.
- Mendez, D., Perez, A. J. & Labrador, M. A. 2011. P-Sense: a participatory sensing system for air pollution monitoring and control. In *2011 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, 344–347.
- Miller, P., Styles, R. & Heath, T. 2008. Open data commons, a license for open data. In *Proceedings of the 1st Workshop About Linked Data on the Web (LDOW2008)*. <http://ceur-ws.org/Vol-369/paper08.pdf>, 5pp.
- Miluzzo, E., Lane, N. D., Fodor, K., Peterson, R., Lu, H., Musolesi, M., Eisenman, S. B., Zheng, X. & Campbell, A. T. 2008. Sensing meets mobile social networks. In *Proceedings of the 6th International Conference on Embedded Networked Sensor Systems, SenSys 2008*, ACM Press, 337–350.
- Miluzzo, E., Papandrea, M., Lane, N. D., Sarroff, A. M., Giordano, S. & Campbell, A. T. 2011. Tapping into the vibe of the city using VibN, a continuous sensing application for smartphones. In *Proceedings of 1st International Symposium on from Digital Footprints to Social and Community Intelligence*, ACM, 13–18.
- Morales, J., Lopez-Sanchez, M., Rodriguez-Aguilar, J. A., Wooldridge, M. & Vasconcelos, W. 2013. Automated synthesis of normative systems. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems, AAMAS '13*, International Foundation for Autonomous Agents and Multiagent Systems, 483–490.

- North, D. C. 1990. *Institutions, Institutional Change and Economic Performance*, Cambridge University Press.
- O'Hara, K. & Shadbolt, N. 2010. Privacy on the data web. *Communications of the ACM* **53**(3), 39–41.
- Oreskes, N. & Conway, E. M. 2010. *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*, Bloomsbury Press.
- Ostrom, E. 1990. *Governing the Commons*, Cambridge University Press.
- Ostrom, E. 2005. *Understanding Institutional Diversity*, Princeton University Press.
- Ostrom, E. & Hess, C. 2007. A framework for analyzing the knowledge commons. In *Understanding Knowledge as a Commons: From Theory to Practice*, Hess, C. & Ostrom, E. (eds). MIT Press, 41–81.
- Perez, A., Labrador, M. & Barbeau, S. 2010. G-Sense: a scalable architecture for global sensing and monitoring. *IEEE Network* **24**(4), 57–64.
- Pitt, J. 2012a. Design contractualism for pervasive/affective computing. *IEEE Technology and Society Magazine* **31**(4), 22–29.
- Pitt, J. 2012b. Managing user-generated content as a knowledge commons. In *Logic Programs, Norms and Action - Essays in Honor of Marek J. Sergot on the Occasion of His 60th Birthday*, Artikis, A., Craven, R., Kesim, N., Sadighi, B. & Stathis, K. (eds). Springer, 401–424.
- Pitt, J., Kamara, L., Sergot, M. & Artikis, A. 2006. Voting in multi-agent systems. *The Computer Journal* **49**(2), 156–170.
- Pitt, J., Schaumeier, J. & Artikis, A. 2012. Axiomatization of socio-economic principles for self-organizing institutions. *ACM Transactions on Autonomous and Adaptive Systems* **7**(4), 1–39.
- Pitt, J., Schaumeier, J. & D'Agostino, M. 2011. Modelling resource allocation in open embedded systems: some logical and epistemological issues. *Logic and Philosophy of Science* **9**(1), 1–15.
- Samuelson, P. 2006. Challenges in mapping the public domain. In *The Future of the Public Domain: Identifying the Commons in Information Law*, Guibault, L. & Hugenholtz, B. (eds). Kluwer Law International, 7–25.
- Schlager, E. & Ostrom, E. 1992. Property-rights regimes and natural resources: a conceptual analysis. *Land Economics* **68**(3), 249–262.
- Schweik, C. M. 2007. Free/open-source software as a framework for establishing commons in science. In *Understanding Knowledge as a Commons: From Theory to Practice*, Hess, C. & Ostrom, E. (eds). MIT Press, 277–310.
- Shadbolt, N. 2013. Midata: towards a personal information revolution. In *Digital Enlightenment Forum Yearbook 2013: The Value of Personal Data*, Hildebrandt, M., O'Hara, K. & Waidner, M. (eds). IOS Press, 202–224.
- Shadbolt, N., O'Hara, K., Berners-Lee, T., Gibbins, N., Glaser, H., Hall, W. & Schraefel, M. 2012. Linked open government data: lessons from data.gov.uk. *IEEE Intelligent Systems* **27**(3), 16–24.
- Stallman, R. 1999. The GNU operating system and the free software movement. In *Open Sources: Voices from the Open Source Revolution*, DiBona, C., Ockman, S. & Stone, M. (eds). O'Reilly Media, 53–70.
- Suber, P. 1990. *The Paradox of Self-Amendment: A Study of Law, Logic, Omnipotence, and Change*, Peter Lang.
- Thiagarajan, A., Ravindranath, L., LaCurts, K., Madden, S., Balakrishnan, H., Toledo, S. & Eriksson, J. 2009. VTrack: accurate, energy-aware road traffic delay estimation using mobile phones. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems — SenSys '09*, ACM Press, 85–98.
- Tilak, S. 2013. Real-world deployments of participatory sensing applications: current trends and future directions. *ISRN Sensor Networks* **2013**, Article ID 583165, 8pp. <http://www.hindawi.com/journals/isrn/2013/583165/cta/>.
- van Dijck, J. 2009. Users like you? Theorizing agency in user-generated content. *Media, Culture & Society* **31**(1), 41–58.