



## Experimentation with a Probable Human-Shaped Object

**Abstract** This chapter sets out the algorithms that will form the focus for this book and their human and non-human associations. The chapter focuses on one particular algorithmic system developed for public transport security and explores the ways in which the system provided a basis for experimenting with what computer scientists termed human-shaped objects. In contrast to much of the social science literature on algorithms that suggests the algorithm itself is more or less fixed or inscrutable, this chapter will instead set out one basis for ethnographically studying the algorithm up-close and in detail. Placing algorithms under scrutiny opens up the opportunity for studying their instability and the ceaseless experimentation to which they are subjected. An important basis for this experimentation, the chapter will suggest, is elegance. The chapter will suggest that elegance opens up a distinct way to conceive of the experimental prospects of algorithms under development and their ways of composing humans.

**Keywords** Experimentation · Human-shaped objects · Elegance · Objects

### OPENING

How can we get close to the everyday life of an algorithm? Building on the Introduction to this book, how can we make sense of the ways algorithms participate in everyday life, compose the everyday or are

continually involved in the becoming of the everyday? In this chapter, I will focus on the question of how an algorithm can (at least begin to) participate in everyday life. I will set out one particular project that provided a means to get close to the everyday life of an algorithm (or what might be more appropriately termed an ‘algorithmic system’, as we will go on to see). I will then investigate one focus for algorithmic experimentation in this project: efforts to identify humans. As the project developed, the notion of human-shaped objects became more and more apparent in project documents, meetings and demonstrations. We will look into what counts as an anticipated human-shaped object, and we will see how our algorithm struggles to grasp such objects. I will also suggest that the human-shaped object becomes itself a matter for experimentation. This is, I will contend, an important aspect of the everyday life of the algorithm: that it becomes entangled with the incredibly banal everyday life of the human (addressing questions such as what is the shape of a human). I will also suggest that the systems with which algorithms participate in construing effects are not stable and certainly not opaque within these project settings. Instead, algorithms and their system are continually inspected and tested, changed and further developed in order to try and grasp the human-shaped object. Within this experimental setting, we can hence note that the algorithms and their system do not operate entirely in line with expectations or within parameters that instantly make sense. Much of the everyday life of the algorithm is thus made up of attempts to get algorithms, their system, computer scientists or others to adequately account for what it is that they have produced. The chapter begins with a consideration of experimentation.

### WHAT IS EXPERIMENTATION?

The tradition of studying the experiment in science and technology studies (STS) has been focused around the space of the laboratory (Latour and Woolgar 1979), forms of expertise (Collins and Evans 2007) and devices (Latour 1987) that render the laboratory a centre of, for example, calculation. The laboratory becomes the space into which the outside world is drawn in order to approximate its conditions for an experiment. Or alternatively, the conditions of the laboratory are extended into the world beyond the laboratory in order to take the laboratory experiment to the world (Latour 1993). The experiment as such, then, becomes a replicable phenomenon through which some feature of

the world is proclaimed. And we see some parallels drawn with economic experiments that similarly seek to draw up variables to manage and control, a world to be drawn into the economic laboratory or a set of conditions to be extended beyond the laboratory (Muniesa and Callon 2007). The economic experiment, like the laboratory experiment, is as much about demonstration, as it is about discovery (Guala 2008).

In Chapter 5, we will see that in the later stages of the everyday life of our algorithm, these concerns for control and demonstration came to the fore—particularly when research funders wanted to see results. But for now, our algorithm—the abandoned luggage algorithm from the Introduction—sits somewhat meekly and unknowing in an office. It is waiting, but it does not know for what it waits: not for an experiment in the closely controlled sense of the term, not for a pristine laboratory space and not for a set of controlled variables, even human or luggage variables, through which it can demonstrate its capacity to grasp the world around it. To begin with it awaits experimentation.

Experimentation sits somewhat apart from experiments. In place of controls or neatly defined space come proposals, ideas, efforts to try things and see what happens. Experimentation can be as much a part of qualitative social science as it can be a part of algorithmic computer science. In the social sciences, experimentation has been used as an impetus by Marres (2013) to experimentalise political ontology and by Johansson and Metzger (2016) to experimentalise the organisation of objects. What these works point towards is the fundamental focus for experimentation: that the nature of just about anything can be rendered at stake within the experimental realm. Scholars have also begun to conceive of experimentalising economic phenomena (Wherry 2014; Muniesa 2016a, b). This is interesting for drawing our attention towards the ways in which what might otherwise be quite controlled, laboratory-like settings can be opened up for new lines of thought through experimentation. These works draw on a patchy history of social science experimentation that has tended to raise insights and ethical concerns in equal measure. One historical route for the development of such experimentation has been Garfinkel's (1963, 1967) breach experiments. Here, the aim was to disrupt—or breach—taken-for-granted features of everyday life in everyday settings in order to make those features available for analysis. But unlike the laboratory tradition, the breaches for Garfinkel were broadly experimental in the sense of providing some preliminary trials and findings to be further worked on. They were heuristic devices, aiding the

sluggish imagination that provoked new thoughts and new lines of enquiry. Our algorithm is awaiting such provocation. But it is also awaiting experimentation that opens up questions of very fundamental features of everyday life, such as what is a human and how ought we to know. And it awaits experimentation that opens up what might otherwise be a pristine laboratory space to new questions, new forms of liveliness.

Experimentation began from the outset of the project in which the algorithm was a participant. We will begin here with the initial development of the project in order to provide a prior step to experimentation. Although the experimentation was more open than a controlled laboratory experiment, it was not free from any constraints or expectations. The experimentation had a broad purpose that was successively set and narrowed as the experimentation proceeded. What was being experimented upon and what was anticipated as the outcome of experimentation was the product of successive rounds of experimentation. To make sense of these expectations, we need to see how the project produced its algorithms in the first place.

### THE ALGORITHMIC PROJECT

The project upon which this book is based began with an e-mail invitation: Would I be interested in participating in a project that involved the development of a new ‘algorithmic’, ‘smart’ and ‘ethical’ video-based surveillance system? The project coordinator informed me that the project would involve a large technology firm (TechFirm<sup>1</sup>), two large transport firms where the technology would be tested and developed (SkyPort, which owns and operates two large European city airports, and StateTrack, a large European state railway) and two teams of computer scientists (from University 1, UK, and University 2, Poland) and that the project would be managed by a consultancy firm (Consultor, Spain). I was being invited to oversee the ethics of the technology under development and to provide an (at least partially) independent ethical assessment. The project would involve developing a system that would use algorithms to select security-relevant images from the CCTV systems of the airport and train station. It would use this ability to demarcate relevance as a basis for introducing a new, ethical, algorithmic system.

The coordinator suggested the project would provide a location for experimentation with three ethical aims: that algorithms could be used to

reduce the scope of data made visible within a video-based surveillance system by only showing ‘relevant’ images; that algorithms could be used to automatically delete the vast majority (perhaps 95%) of surveillance data that was not relevant; and that no new algorithms or surveillance networks would need to be developed to do so. These aims had been developed by the coordinator into an initial project bid. The coordinator hoped the ‘ethical’ qualities of the project were clear in the way the aims were positioned in the bid as a response to issues raised in popular and academic discussions about, for example, privacy, surveillance and excessive visibility (Lyon 2001; Davies 1996; Norris and Armstrong 1999; Bennett 2005; Van der Ploeg 2003; Taylor 2010; McCahill 2002) and concerns raised with algorithmic surveillance (Introna and Woods 2004). In particular, the project bid set out to engage with contemporary concerns regarding data retention and deletion, as very little data would be kept (assuming the technology worked).

The proposal was a success, and the project was granted €2.8m (about \$3.1m in mid-2015) under the European Union’s 7th Framework Programme. A means to fulfil the promises of ethical algorithms committed to the project bid would now have to be found. This set the scene for early rounds of experimentation.

## ESTABLISHING THE GROUNDS FOR EXPERIMENTATION AND THE MISSING ALGORITHMS

The basis for initial experimentation within the project was a series of meetings between the project participants. Although there were already some expectations set in place by the funding proposal and its success, the means to achieve these expectations and the precise configuration of expectations had not yet been met. I now found myself sat in these meetings as an ethnographer, watching computer scientists talking to each other mostly about system architectures, media proxies, the flow of digital data—but not so much about algorithms.

Attaining a position as an ethnographer on this project had been the result of some pre-project negotiation. Following the project coordinator’s request that I carry out an ethical review of the project under development, I had suggested that it might be interesting, perhaps vital, to carry out an assessment of the system under development. Drawing on recent work on ethics and design and ethics in practice,<sup>2</sup> I suggested that what got to count as the ethics of an algorithm might not be easy

to anticipate at the start of a project or from a position entirely on the outside. It might make more sense to work with the project team, studying the development of the technology and feeding in ethical questions and prompts over the three years of the project. Although this seemed to be an interesting proposition for project participants, questions were immediately raised regarding my ability to be both in the project (on the inside) and offer an ethical assessment (something deemed to be required from the outside). I suggested that during the course of the project I could use the emerging ethnography of system development to present the developing algorithm to various audiences who might feedback interesting and challenging questions, I could set up meetings with individuals who might provide feedback on the developing technology and I would put together an ethics board. The latter would be outsiders to the project, enabling me to move between an inside and outside position, working with, for example, the computer scientists in the project at some points and with the ethics board members at other moments. As we will see in Chapter 3, this ethical assessment formed one part of a series of questions regarding accountability that were never singularly resolved in the project. However, for now, at the outset, my role as ethnographer was more or less accepted, if not yet defined.

But what of the algorithms? In these early project meetings when I was still developing a sense of what the project was, what the technology might be, and what the challenges of my participation might involve, algorithms still retained their mystery. In line with academic writing on algorithms that emphasises their opacity (see Introduction to this book), at this moment the nature and potential of algorithms in the project was unclear. Occasionally during these meetings, algorithms were mentioned, and most project participants seemed to agree with the computer scientists from University 1 and University 2 that the ‘algorithms’ were sets of IF-THEN rules, already complete with associated software/code that could be ‘dropped into’ the system. The system seemed to be the thing that needed to be developed, not the algorithms. As we will see, this notion of ‘dropping in’ an algorithm turned out to be a wildly speculative and over-optimistic assessment of the role and ability of algorithms, but for now in project meetings, the system was key.

Establishing the precise set-up for the algorithmic system under development involved the computer scientists and transport firms involved in the project (the airport and train operator) discussing first steps in technology development. Although this could be described as a negotiation,

it mostly seemed to involve the computer scientists proposing system components and then later an order of system components (more or less setting out how data would flow through the system and how each component could talk to each other) and then the transport firms would respond. There was never an occasion where the transport firms would make the first proposal. This seemed to be a result of the meetings being framed as technical discussions primarily, rather than being focused on, for example, usability. It was also the case that with more than a decade of experience in developing these systems, University 1 and University 2 computer scientists could talk with a fluency, eloquence and technical mastery that no one else could match. When the computer scientists made a proposal, it was up to the transport firms to accept or not the proposal and then it was down to the computer scientists to make any necessary adjustments.

But working together in this kind of complex multiparty, international project was not entirely straightforward. The experience of the computer scientists in developing these kinds of systems was thus a welcome contribution to the project in itself. It established a way of working that others could fit in with. Its absence might have meant a significant number of meetings to decide on ways of meeting. The meetings became framed as technical matters in which the computer scientists would lead partly because of the lack of any alternative way to frame meetings that anyone put forward. The ethnographer certainly didn't propose to have meetings framed around ethnography (at least not yet, see Chapters 3 and 5).

The meetings worked as follows. The participants would be gathered around a semicircle of tables or, on one occasion, an oval table, with a screen and projector at one end. Onto the screen would be projected a technical matter under discussion—often the system architecture. This set out the distinct components of the system under development and the role such components might play. Discussions in meetings then focused on the implications of setting up the system in one way or another, along with discussions of individual components of the system and, in early meetings, new components that might be needed. One computer scientist would sit at a laptop linked to the projector with the system architecture on their screen and make adjustments as discussions continued so that meeting participants could further discuss the emerging system. As they made changes on their laptop, these were projected onto a large screen for meeting participants to discuss. Sometimes two or more computer scientists would gather round the laptop to further

discuss a point of detail or refine precisely what it was that had just been proposed in the meeting and what this might look like for the system. A typical architecture from one of the later meetings looked like this as shown in Fig. 2.1.

By this point in the project, it had been agreed that the existing surveillance cameras in transport hubs operated by SkyPort and StateTrack would feed into the system (these are represented by the black camera-shaped objects on the left). After much discussion, it had been agreed that the digital data from these cameras would need to feed into a media proxy. Both sets of computer scientists were disappointed to learn that the transport hubs to be included in the project had a range of equipment. Some cameras were old, some new, some high definition and some not, and each came with a variable frame rate (the number of images a second that would flow into the algorithmic system). The media proxy was required to smooth out the inconsistencies in this flow of data in order that the next component in the system architecture would then be able to read the data. Inconsistent or messy data would prove troublesome throughout the project, but in these meetings, it was assumed that the media proxy would work as anticipated.

After some discussion, it was agreed that the media proxy would deliver its pristine data to two further system components. These comprised the Event Detection system and the Route Reconstruction

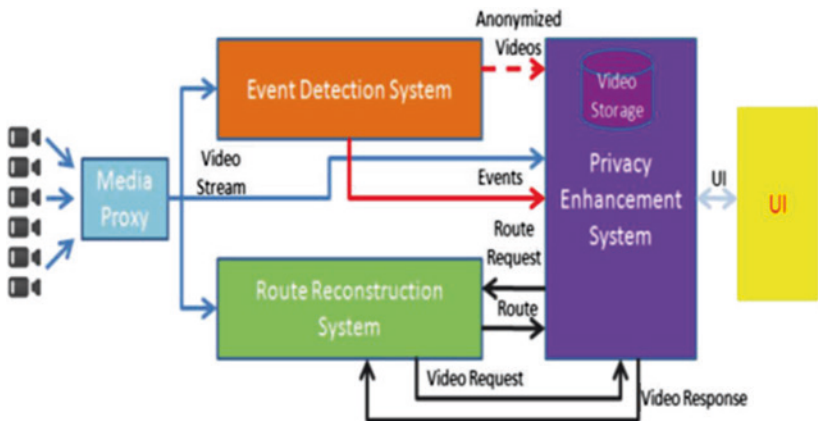


Fig. 2.1 System architecture



system. The Event Detection system was where the algorithms (including the abandoned luggage algorithm of the Introduction to this book) would sit. The idea was that these algorithms would sift through terabytes of digital video data and use IF-THEN rules to select out those events that security operatives in a train station or airport would need to see. In discussions between the computer scientists and transport firms, it was agreed that abandoned luggage, people moving the wrong way (counter-flow) and people moving into forbidden areas (such as the train track in train stations or closed offices in airports) would be a useful basis for experimentation in the project. These would later become the basis for algorithmically experimenting with the basic idea that relevant images could be detected within flows of digital video data. For now, it was still assumed that algorithms could simply be dropped into this Event Detection component of the architecture. Relevant images would then be passed to the User Interface (UI) with all data deemed irrelevant passed to the Privacy Enhancement System. This was put forward as a key means to achieve the ethical aims of the project. It was suggested that only a small percentage of video data was relevant within an airport or train station, that only a small percentage of data needed to be seen and that the rest of the data could be stored briefly in the Privacy Enhancement System before being securely deleted. It later transpired that detecting relevant images, getting the algorithms to work and securely deleting data were all major technical challenges. But for now, in these early project meetings, it was assumed that the system would work as expected.

The Route Reconstruction component was a later addition. This followed on from discussions between the transport firms and the computer scientists, in which it became clear that having an image of, for example, an abandoned item of luggage on its own was not particularly useful in security terms. Transport operatives wanted to know who had left the luggage, where they had come from and where they went next. The theory behind the Route Reconstruction system (although see Chapter 3 for an analysis of this) was that it would be possible to use probabilistic means to trace the history around an event detected by an algorithm. The UI would then give operatives the option to see, for example, how an item of luggage had been abandoned, by whom, with whom they were walking and so on. This would mean that the Privacy Enhancement System would need to store data for as long as these reconstructions were required. It was assumed that most would be performed within

24 hours of an incident. Any data deemed relevant and any reconstructions viewed by operatives would be moved out of the auto-deletion feature of the Privacy Enhancement System and kept (in Video Storage). According to the computer scientists, this should still mean that around 95% of data was deleted and that the ethical aims to see less and store less data would be achieved.

The meetings were discursive fora where the computer scientists took the lead in making proposals and other participants, mostly the transport firms, offered their response. The overall effect was that the algorithmic system began to emerge and take shape, at least on a computer screen and as a projection. The components that would need to be developed were discussed, the future role of algorithms in Event Detection was more or less set, and a specific shape was given to the project's ethical proposals. A technical means was proposed for limiting the range of data that would be seen and the amount of data that would be stored. As we will see, producing a UI and Route Reconstruction system (Chapter 3) and deleting data (Chapter 4) were problematic throughout the life of the project. However, for now we will retain our focus on experimenting with the human-shaped object.

### ELEGANCE AND THE HUMAN-SHAPED OBJECT

With the system architecture agreed at least in a preliminary form, the computer scientists could get on with the task of making the system work. Key to having a working system was to ensure that the data flowing from surveillance cameras through the Media Proxy could be understood by the Event Detection component. If events could not be detected, there would be no system. Figuring out ways to detect abandoned luggage, moving the wrong way and moving into a forbidden space were crucial. Central to Event Detection was the human-shaped object. As digital video was streamed through the system, the algorithms would need to be able to pick out human-shaped objects first, and then the actions in which they were engaged second. Relevant actions for this experimental stage of system development would be a human-shaped object moving the wrong way, a human-shaped object moving into a forbidden space and a human-shaped object abandoning its luggage.

How could this human-shaped object be given a definition that made sense for operationalisation within the system? The algorithms for Event Detection used in video analytic systems are a designed product. They take effort and work and thought and often an amount of reworking.

The algorithms and their associated code for this project built on the decade of work carried out, respectively, by University 1 and University 2. As these long histories of development had been carried out by various colleagues within these Universities over time, tinkering with the algorithms was not straightforward. When computer scientists initially talked of ‘dropping in’ algorithms into the system, this was partly in the hope of avoiding many of the difficulties of tinkering, experimenting and tinkering again with algorithms that were only partially known to the computer scientists. As we saw in the Introduction with the abandoned luggage algorithm, the algorithm establishes a set of rules which are designed to contribute to demarcating relevant from irrelevant video data. In this way, such rules could be noted as a means to discern people and things that could be ignored and people and things that might need further scrutiny. If such a focus could hold together, the algorithms could indeed be dropped in. However, in practice, what constituted a human-shaped object was a matter of ongoing work.

Let’s return to the subject of the Introduction to then explore experimentation with human-shaped objects. As a reminder, these are the IF-THEN rules for the abandoned luggage algorithm (Fig. 2.2):

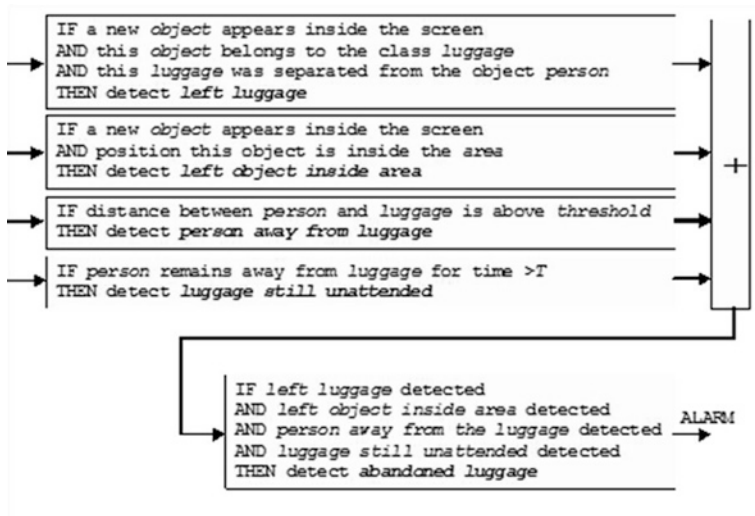


Fig. 2.2 Abandoned luggage algorithm

As I noted in the Introduction, what seems most apparent in these rules is the IF-THEN structure. At its simplest, the ‘IF’ acts as a condition and the ‘THEN’ acts as a consequence. In this particular algorithm, the IF-THEN rules were designed to operate in the following way. IF an object was detected within a stream of digital video data fed into the system from a particular area (notably a train station operated by StateTrack or an airport operated by SkyPort), THEN the object could be tentatively allocated the category of potentially relevant. IF that same object was deemed to be in the class of objects ‘human-shaped’, THEN that object could be tentatively allocated the category of a potentially relevant human-shaped object. IF that same human-shaped object was separate from a luggage-shaped object, THEN it could maintain its position as potentially relevant. IF that same human-shaped object and luggage-shaped object were set apart beyond a specific distance threshold set by the system (say 2 or 10 metres) and the same objects were set apart beyond the temporal threshold set by the system (say 30 seconds or 1 minute)—that is, if the person and luggage were sufficiently far apart for sufficiently long—THEN an alert could be sent to surveillance operatives. The alert would then mean that the package of data relevant to the alert would be sent to the UI and operatives could then click on the data, watch the video of abandoned luggage and offer a relevant response (see Chapter 3). What is important for now is how these putative objects were given shape and divided into relevant and irrelevant entities.

If this structuring and division of various entities (humans, luggage, time, space, relevance and irrelevance) occurred straightforwardly and endured, it might be tempting to argue that this is where the power of algorithms is located or made apparent. A straightforward short cut would be to argue that the algorithm structures the social world and through this kind of statement we could then find what Lash (2007) refers to as the powerful algorithm and what Beer (2009) suggests are algorithms’ ability to shape the social world. We could argue that the outputs of the system demonstrated an asymmetrical distribution of the ability to cause an effect—that it is through the algorithm that these divisions between relevant and irrelevant data can be discerned. However, such a short cut requires quite a jump from the algorithmic IF-THEN rules to their consequences. It ignores the everyday life in which the algorithm must be a successful participant for these kinds of effects to be brought about. If instead we pay attention to the everyday work required

for algorithmic conditions and consequences to be achieved, what we find is not that the algorithm structures the social world. Instead, experimentation takes place (and fails, takes place again, things are reworked and then sometimes fail again or work to a small extent) to constitute the conditions required for the system to participate in the production of effects or the system gets partially redrawn to fit new versions of the conditions. This continual experimentation, rewriting and efforts to achieve conditions and consequences are not only central to the work of computer scientists but also crucial to the life of the algorithm. It is where the distinction between relevant and irrelevant data is continuously in the process of being made for the system. It is where the human-shaped object is drawn up and pursued. It is where the nature of things is made at stake.

The experimental basis designed to enable the algorithm to participate in the everyday life of the airport and train station had, what was for the ethnographer, a peculiar organising principle. The computer scientists of University 1 and University 2 talked of ‘elegance’ during the meetings around system architecture, huddled around the laptop on which they made updates to the system and in the subsequent human-shaped object experimentation that we will now consider. This seemed like an odd term to me in a series of meetings that mostly involved quite specialist, technical language. Elegance seemed to come from a different field—perhaps fashion or furniture design. What could it mean for the computer scientists to talk of elegance, or rather how was the term elegance given meaning by the practical work of the computer scientists?

Ian Glynn (2010) captures something of what elegance can mean in his study of experiments and mathematical proofs. Glynn suggests elegance can be found in scientific and mathematical solutions which combine concision, persuasion and satisfaction. As I followed the experimentation of the computer scientists, this approach to elegance seemed useful for making sense of the ways they discussed system architecture. The composition of the different system components, their location in relation to each other within the system architecture and how they might talk to each other was each discernible as a discussion focused on what might be elegant. However, this came to the fore even more strongly with the human-shaped object. What would count as concise, persuasive and satisfying as a human-shaped object seems a useful way to group together much of the discussion that took place.

For the IF conditions of the Event Detection algorithms to be achieved required coordinated work to bring together everyday competences (among surveillance operatives and computer scientists), the creation of new entities (including lines of code), the further development of components (from algorithmic rules to new forms of classification) and the development of a sense of what the everyday life was in which the algorithms would participate (in the train station or airport). It also required consideration of the changes that might come about in that everyday life. Elegance could be noted as the basis for this coordinated work in the following way. The first point of contention was the technical basis for developing a means to classify putative objects. Readers will recall that first identifying a putative object is important within the stream of digital video data in order that other data can be ignored. What might count as a human-shaped object or a luggage-shaped object as a precise focus for classification was vital. However, what might count as a concise means to achieve this classification was an important but slightly different objective. As project meetings progressed, it became clear that the amount of processing power required to sift through all the data produced in a train station or airport in real time and classify all human-shaped objects precisely would be significant. Face recognition, iris recognition and gait recognition (based on how people walked) were all ruled out as insufficiently concise. These approaches may have been persuasive as a means to identify specific individuals in particular spaces, but their reliability depended on having people stand still in controlled spaces and have their features read by the system for a few seconds. This would not be very satisfying for passengers or airports whose business models depended on the rapid movement of passengers towards shops (Woolgar and Neyland 2013).

How then to be concise and satisfying and persuasive in classifying human-shaped objects? As Bowker and Star (2000) suggest, classification systems are always incomplete. This incompleteness ensures an ambiguity between the focus of algorithmic classification (the putative human-shaped object) and the entity being classified (the possible human going about their business). Concision requires various efforts to classify to a degree that is satisfying and persuasive in relation to the needs of the system and the audiences for the algorithmic system. The system needs to do enough (be satisfying and persuasive), but no more than enough (be concise) as doing more than enough would require more processing of data. In the process of experimenting with human-shaped objects

in this project, various more or less concise ways to classify were drawn up and considered or abandoned either because they would require too much processing power (probably quite persuasive but not concise) or were too inaccurate (quite concise, but produced results that were not at all persuasive). At particular moments, (not very serious) consideration was even given to change the everyday life into which the algorithms would enter in order to make classification a more straightforward matter. For example, computer scientists joked about changing the airport architecture to suit the system, including putting in higher ceilings, consistent lighting and flooring, and narrow spaces to channel the flow of people. These were a joke in the sense that they could never be accommodated within the project budget. Elegance had practical and financial constraints.

A first move in classifying objects was to utilise a standard practice in video analytics: background subtraction. This method for identifying moving objects was somewhat time-consuming and processor intensive, and so not particularly elegant. But these efforts could be ‘front-loaded’ prior to any active work completed by the system. ‘Front-loading’ in this instance meant that a great deal of work would be done to produce an extensive map of the fixed attributes of the setting (airport or train station) prior to attempts at classification work. Mapping the fixed attributes would not then need to be repeated unless changes were made to the setting (such changes included in this project a change to a shop-front and a change to the layout of the airport security entry point). Producing the map provided a basis to inform the algorithmic system what to ignore, helping to demarcate relevance and irrelevance in an initial manner. Fixed attributes were thus nominally collated as non-suspicious and irrelevant in ways that people and luggage, for example, could not be, as these latter objects could not become part of the map of attributes (the maps were produced based on an empty airport and train station). Having a fixed map then formed the background from which other entities could be noted. Any thing that the system detected that was not part of the map would be given an initial putative identity as requiring further classification.

The basis for demarcating potentially relevant objects depended to some degree, then, on computer scientists and their understanding of spaces such as airports, maps that might be programmed to ignore for a time certain classes of objects as fixed attributes, and classification systems that might then also—if successful—provide a hesitant basis for

selecting out potentially relevant objects. It is clear that anything that algorithms were able to ‘do’ was situated in continual reconfigurations of the entities involved—making sense of the everyday life of the algorithm was thus central.

Mapping for background subtraction was only a starting point. Objects noted as non-background entities then needed to be further classified. Although background subtraction was a standard technique in video analytics, further and more precise object classification was the subject of some discussion in the project. Eventually, two means were set for classifying those entities noted as distinct from the background that might turn out to be human-shaped objects, and these became the focus for more intense experimentation in the project. The first of these involved bounding boxes, and the second involved a more precise pixel-based classification. Both approaches relied on the same initial parameterisation of putative objects. To parameterise potential objects, models had to be computationally designed. This involved experimenting with establishing edges around what a human-shaped object was likely to be (in terms of height, width and so on). Other models then had to be built to parameterise other objects, such as luggage, cleaners’ trolleys, signposts and other non-permanent attributes of the settings under surveillance. The models relied on 200-point vector analysis to set in place what made up the edges of the object under consideration and then to which model those edges suggested the object belonged. This was elegant insofar as it would produce rapid, real-time classifications because it was concise, using only a minimal amount of processing power and being achievable very quickly. Parameterisation was presented by the computer scientists as a form of classification that the developing algorithmic system could manage while the system also carried out its other tasks. In this way, parameterisation would act as an initial but indefinite basis for object classification that could be confirmed by other system processes and even later by surveillance operatives when shown images of, for example, an apparently suspicious item of luggage. However, these parameterisations could only be adjudged as satisfactory and persuasive when they were put to use in the airport and train station. There were just too many possible mitigating issues to predict how an initial experimentation with parameterisation would turn out in practice. Initial parameterisation did at least allow the computer scientists to gain some confidence that their putative classifications could be achieved within the




bounds of processing possibility and could be achieved by making selections of relevance from streams of digital video data.

Once parameterised as putative human-shaped or luggage-shaped objects, the action states of these objects also required classification, for example, as moving or not moving. This involved object tracking to ascertain the action state of the objects being classified. To achieve the conditions established in the algorithmic IF-THEN rules, the system had to identify, for example, that a putative item of luggage demarcated as potentially relevant, based on a designed model used to initiate parameterisation, was no longer moving and that a human-shaped object that was derived from a similar process, had left this luggage, had moved at least a certain distance from the luggage and for a certain time. In order to track objects that had been given an initial and hesitant classification, human-shaped objects and luggage-shaped objects would be given a bounding box. This was a digitally imposed stream of metadata that would create a box around the object according to its already established edges (Fig. 2.3).

The box would then be given a metadata identity according to its dimensions, location within the airport or train station (e.g. which

**Fig. 2.3** An anonymous human-shaped bounding box




camera it appeared on) and its direction and velocity. For the Event Detection algorithms of moving into a forbidden space or moving in the wrong direction (e.g. going back through airport security or going the wrong way through an entry or exit door in a rush hour train station), these bounding boxes were a concise form of identification. They enabled human-shaped objects to be selected with what might be a reasonable accuracy and consistency and without using too much processing effort. They were elegant, even if visually they looked a bit ugly and did little to match the actual shape of a human beyond their basic dimensions.

However, for abandoned luggage, something slightly different was required. In experimentation, in order to successfully and consistently demarcate human-shaped objects and luggage-shaped objects and their separation, a more precisely delimited boundary needed to be drawn around the putative objects. This required the creation of a pixel mask that enabled the algorithmic system to make a more precise sense of the human- and luggage-shaped objects, when and if they separated (Fig. 2.4).

This more closely cropped means to parameterise and classify objects could then be used to issue alerts within the system. IF a human-shaped object split from a luggage-shaped object, IF the human-shaped object continued to move, IF the luggage-shaped object remained stationary, IF the luggage-shaped object and human-shaped object were

**Fig. 2.4** A close-cropped pixelated parameter for human- and luggage-shaped object



over a certain distance apart and IF the human-shaped object and luggage-shaped object stayed apart for a certain amount of time, THEN this would achieve the conditions under which the algorithmic system could issue an alert to operatives. As the following figure shows, once a close-cropped image of what could be classified as a luggage-shaped object was deemed by the system to have lingered beyond a defined time and distance from its human-shaped object, then it would be highlighted in red and sent to operatives for confirmation and, potentially, further action (Fig. 2.5).

In place of the concise elegance of the imprecise bounding box, this more precise pixel cropped form of parameterisation was computationally more demanding, requiring a little more time and more processing power. However, it was key to maintaining a set of results in these initial experimentations that satisfied the needs of the system as agreed with SkyPort and StateTrack and could be persuasive to all project partners. That is, it produced results that suggested the project was feasible and ought to continue (although as we will see in Chapter 5, problems with classification continued to be difficult to resolve). The bounding boxes lacked the precision to give effect to the IF-THEN rules of the abandoned luggage algorithm.

The human-shaped object was thus accomplished in two forms—as a bounding box and a more closely cropped image. The bounding boxes although somewhat crude were central to the next stages in algorithmic experimentation in Route Reconstruction and the issuing of alerts to operatives (see Chapter 3) and deletion (see Chapter 4). For now, our algorithm could be satisfied that it had been able to participate in at least a modified, initial and hesitant, experimental form of everyday life. It had

**Fig. 2.5** An item of abandoned luggage



not succeeded entirely in meeting all the goals of the project yet, it had struggled to initially produce a set of results that could elegantly capture sufficient information to accurately and consistently identify abandoned luggage and had to be changed (to a pixel-based process), and it was reliant on digital maps and background subtraction, but it had nonetheless started to get into the action.

## CONCLUSION

In this chapter, I have started to build a sense of the everyday life in which our algorithm was becoming a participant. In experimental spaces, our algorithm was starting to make a particular sense of what it means to be human and what it means to be luggage. The IF-THEN rules and the development of associated software/code, the building of a system architecture and set of components provided the grounds for rendering things like humans and luggage at stake. To return to Pollner's (1974) work (set out in the Introduction), the algorithm was starting to set out the basis for delimiting everyday life. The algorithm was beginning to insist that the stream of digital video data flowing through the system acted on behalf of an account as either luggage-shaped or human-shaped or background to be ignored. In addressing the question how do algorithms participate in everyday life, we have started to see in this chapter that they participate through technical and experimental means. Tinkering with ways to frame the human-shaped object, decide on what might count as elegant through concision, satisfaction and persuasion, are all important ways to answer this question. But we can also see that this participation is hesitant. The bounding box is quite elegant for two of the system's algorithmic processes (moving the wrong way and moving into a forbidden area) but not particularly persuasive or satisfactory for its third process (identifying abandoned luggage). And thus far, all we have seen is some initial experimentation, mostly involving the human-shaped objects of project participants. This experimentation is yet to fully escape the protected conditions of experimentation. As we will see in Chapters 4 and 5, moving into real time and real space, many of these issues in relation to algorithmic participation in everyday life have to be reopened.

It is in subsequent chapters that we will start to look into how the algorithm becomes the everyday and how algorithms can even compose the everyday. For now, these questions have been expressed in limited

ways, for example when the computer scientists joked about how they would like to change the airport architecture to match the needs of the system. In subsequent chapters, as questions continue regarding the ability of the algorithm to effectively participate in everyday life, these questions resurface. In the next chapter, we will look at how the algorithmic system could become accountable. This will pick up on themes mentioned in the Introduction on transparency and accountability and will explore in greater detail the means through which the everyday life of the algorithm could be made at stake. As the project upon which this book is based was funded in order to produce a more ethical algorithmic system, these questions of accountability were vital.

## NOTES

1. The names have been made anonymous.
2. Design and ethics have a recent history (Buscher et al. 2013; Introna and Nissenbaum 2000; Suchman 2011). However, Anderson and Sharrock (2013) warn against assuming that design decisions set in place an algorithmic ethic.

## REFERENCES

- Anderson, B., & Sharrock, W. (2013). *Ethical Algorithms*. Available from: <http://www.sharrockandanderson.co.uk/wp-content/uploads/2013/03/Ethical-Algorithms.pdf>.
- Beer, D. (2009). Power Through the Algorithm? Participatory Web Cultures and the Technological Unconscious. *New Media & Society*, 11(6), 985–1002.
- Bennett, C. (2005). What Happens When You Book an Airline Ticket (Revisited): The Collection and Processing of Passenger Data Post 9/11. In M. Salter & E. Zureik (Eds.), *Global Surveillance and Policing* (pp. 113–138). Devon: Willan.
- Bowker, G., & Star, S. L. (2000). *Sorting Things Out*. Cambridge, MA: MIT Press.
- Buscher, M., Wood, L., & Perng, S.-Y. (2013, May). Privacy, Security, Liberty. In T. Comes, F. Fiedrich, S. Fortier, & J. Geldermann (Eds.), *Proceedings of the 10th International ISCRAM Conference*, Baden-Baden, Germany. Available from: <http://www.iscramlive.org/portal/iscram2013proceedings>. Accessed October 24, 2014.
- Collins, H., & Evans, R. (2007). *Rethinking Expertise*. Chicago, IL: University of Chicago Press.

- Davies, S. (1996). *Big Brother—Britain's Web of Surveillance and the New Technological Order*. London: Pan Books.
- Garfinkel, H. (1963). A Conception of and Experiments With “Trust” as a Condition of Concerted Stable Actions. In O. J. Harvey (Ed.), *The Production of Reality: Essays and Readings on Social Interaction* (pp. 187–238). New York, USA: The Ronald Press Company.
- Garfinkel, H. (1967). *Studies in Ethnomethodology*. Englewood Cliffs, NJ: Prentice-Hall.
- Glynn, I. (2010). *Elegance in Science*. Oxford: Oxford University Press.
- Guala, F. (2008). How to Do Things with Experimental Economics. In D. MacKenzie, F. Muniesa, & L. Siu (Eds.), *Do Economists Make Markets? On the Performativity of Economics* (pp. 128–162). Princeton, NJ: Princeton University Press.
- Introna, L., & Nissenbaum, H. (2000). Shaping the Web: Why the Politics of Search Engines Matters. *The Information Society*, 16(3), 1–17.
- Introna, L., & Woods, D. (2004). Picturing Algorithmic Surveillance: The Politics of Facial Recognition Systems. *Surveillance and Society*, 2(2/3), 177–198.
- Johansson, N., & Metzger, J. (2016). Experimentalizing the Organization of Objects: Re-enacting Mines and Landfills. *Organization*, 23(6), 840–863.
- Lash, S. (2007). Power After Hegemony. *Theory, Culture and Society*, 24(3), 55–78.
- Latour, B. (1987). *Science in Action*. Cambridge, MA: Harvard University Press.
- Latour, B. (1993). *The Pasteurization of France*. Cambridge: Harvard University Press.
- Latour, B., & Woolgar, S. (1979). *Laboratory Life*. London: Sage.
- Lyon, D. (2001). *Surveillance Society: Monitoring Everyday Life*. Buckingham: Open University Press.
- Marres, N. (2013). Why Political Ontology Must Be Experimentalized: On Eco-Show Homes as Devices of Participation. *Social Studies of Science*, 43(3), 417–443.
- McCahill, M. (2002). *The Surveillance Web: The Rise of Visual Surveillance in an English City*. Devon: Willan.
- Muniesa, F. (2016a). Setting the Habit of Capitalization: The Pedagogy of Earning Power at the Harvard Business School, 1920–1940. *Historical Social Research*, 41(2), 196–217.
- Muniesa, F. (2016b) The Problem with Economics: Naturalism, Critique and Performativity. In I. Boldyrev & E. Svetlova (Eds.), *Enacting Dismal Science: New Perspectives on the Performativity of Economics*. London, UK: Palgrave Macmillan.
- Muniesa, F., & Callon, M. (2007). Economic Experiments and the Construction of Markets. In D. MacKenzie, F. Muniesa, & L. Siu (Eds.), *Do Economists Make Markets?* (pp. 163–188). Oxford: Princeton University Press.

- Norris, C., & Armstrong, G. (1999). *The Maximum Surveillance Society: The Rise of CCTV*. London: Berg.
- Pollner, M. (1974). *Mundane Reason*. Cambridge: Cambridge University Press.
- Suchman, L. (2011). Subject Objects. *Feminist Theory*, 12(2), 119–145.
- Taylor, E. (2010). I Spy with My Little Eye. *Sociological Review*, 58(3), 381–405.
- Van der Ploeg, I. (2003). Biometrics and Privacy: A Note on the Politics of Theorizing Technology. *Information, Communication, Society*, 6(1), 85–104.
- Wherry, F. (2014). Analyzing the Culture of Markets. *Theory and Society*, 43(3–4), 421–436.
- Woolgar, S., & Neyland, D. (2013). *Mundane Governance*. Oxford: Oxford University Press.

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

