The Microscopic Human Factors methodology for modelling cognition in crowds and swarm systems*

Colin Marc Henein†
Institute of Cognitive Science, and
Complex and Adaptive Systems Group
Carleton University, Ottawa, Canada

Tony White
School of Computer Science, and
Complex and Adaptive Systems Group
Carleton University, Ottawa, Canada

Abstract

Like in other species there has been an interest in modelling behaviours of human aggregates: crowds. Despite the important influence of cognition on individual human behaviour, it has been a challenge to integrate human factors into swarm-principled crowd models (termed: microscopic crowd models).

To improve this situation we offer the Microscopic Human Factors methodology, which (a) guides modelling of cognitive factors, and (b) increases the transparency of resulting models to non-modelling crowd researchers. Through (a) resulting models become models of people, with better relevance to real crowds. Through (b) resulting models can better contribute to the larger field of inquiry into crowd behaviour. The methodology proposed can also assist in modelling collective behaviour in other species, particularly those species with significant cognitive components in their social behaviour.

The methodology is general, calling for three complete descriptions of the collective behaviour regardless of the implementation formalism to be used: The *specification* delineates the crowd behaviour of interest. The *reduction* identifies the abstract, local behaviour of individuals that underlies the specified collective effect. Formal modelling and implementation details are isolated to the *implementation*. By considering the methodology's layers in reverse order it is possible to make arguments concerning the validity of resulting models.

Keywords: agent-based modelling methodology; cognitive modelling; agent cognition; crowd dynamics; validation; individual-based simulation; multi-agent systems

*This document may be cited as follows: CM. Henein and T. White. The Microscopic Human Factors methodology for modelling cognition in crowds and swarm systems. Technical Report TR-2010-13, Carleton University School of Computer Science, 2010.

†Corresponding author: Institute of Cognitive Science, Carleton University. DT2201-1125 Colonel By Drive. Ottawa, Canada, K1S 5B6. Phone: +1 613 520-2368. Email: Colin_Henein –at– carleton.ca

1 Introduction

When studying swarm intelligence, we often look to the mechanisms at work in social animals to understand how simple behaviours by many of these individuals can produce emergent effects in aggregate. Just as the mechanisms and behaviour of insect and other animal aggregates has been of interest, so has the study of swarms of humans: crowds. Just as in swarm research more generally, the study of crowd dynamics seeks to explain how the behaviour of individuals affects overall crowd movement. Evacuation dynamics, for example, is a sub-field of considerable interest that is concerned with simulation of large crowds motivated to exit quickly (e.g. from a building).

A key to understanding the emergent behaviours of aggregations of a social species is determining species-specific capacities and behaviours that individuals possess, which they collectively contribute to the emergent behaviour [1]. In ants, for example, the pheromone trail is widely acknowledged as a key instrument for stigmergic communication (e.g. [2]). But what of the key species-specific capabilities of humans in aggregate? People have various abilities in the physical domain including some long-range, high resolution sensory capabilities and various physical mechanisms for locomotion and for physically manipulating the environment and neighbouring persons. It is the human capacity for advanced abstract reasoning and communication, however, that immediately sets people in crowds apart from individuals in aggregates of other species. In short, it is to human cognition that we must look in understanding the species-specific determinants of behaviour in crowd dynamics.

In spite of the obvious importance of cognition in understanding individual human behaviour, incorporating cognitive factors into models of crowd behaviour has proven to be a challenge: the history of crowd modelling is not flush with cognitive simulation [3,4]. Jonathan Sime suggested two possible reasons that human behaviour may be omitted from models [5]. First, the popular view that humans panic in crowds, that panic trumps cognition, that crowds are thus inherently non-cognitive and consequently that modelling cognition in crowds is futile. Second, that the extreme computational expense of cognitive simulation makes it practically impossible to model human cognition. On the first point the literature is clear that incidences of panic in crowd disasters are rare, and that people generally take actions that make sense in the context of the information available to them at the time [4–14]. The second point, however, as a practical objection to modelling truly human aspects of behaviour in crowds, has not yet been well addressed; this is particularly true in the area known as *microscopic crowd modelling*[1], in which familiar swarm principles (large numbers of simple and self-organising individuals interacting socially and with the environment to produce complex emergent collective behaviour) are adopted with a view to understanding crowd structure and behaviour.

Despite the difficulties involved in modelling cognition within crowd simulations, it is instructive to note Sime's advice to avoid conceptualisation of people as non-thinking objects like *ball-bearings* [15]. Taking this view, we see the inadequacy of modelling crowds through analogies to fluid dynamics and particle systems like seeds in a hopper. Crowd models that do not simulate human behaviour open themselves to the criticism that they are not models of people; they thus have limited applicability, at best, to human crowds. This type of criticism can be levelled against any swarm model of an aggregate behaviour of an intelligent animal that omits relevant cognitive factors in its elaboration.

The purpose of this article is to describe our methodology to facilitate simulation of cognition-driven human behaviours in microscopic crowd models, which also improves transparency in model design (extending the reach of model results). This mechanism is also applicable to cognitive modelling in non-human aggregates. We begin by first examining exactly what it is about human cognition that is missing from models. We then describe the Microscopic Human Factors methodology that allows us to integrate human behaviour into these models. Although we have previously developed complex models guided by this technique (e.g. [4, 16–20]), we here illustrate it using a very simple example that highlights the structure of the approach. We shall then discuss the validity of models built according to the methodology.

[1]In the crowd literature, the term *microscopic* does not imply physical scale, but rather the level of detail in the simulation. It identifies models that consider behaviour at the level of individuals, as opposed to a monolithic technique. The latter approach considers only aggregate properties like crowd size and density in predicting behaviour.

2 Abstract modelling of cognition

Our methodology targets the integration of *relevant* human factors into crowd models of the swarm type; as we have discussed, cognition, as a defining human characteristic, is of particular interest for modelling. But what is it about cognition that is relevant and thus must be incorporated? The key to obtaining tractability in simulation of large numbers of individuals is in realising that we cannot model everything about cognitive actors, and so we must narrow the scope of the methodology to what is most important.

2.1 Macroscopic cognitive simulation

One way to approach the problem would be to simulate a complete cognition for agents in the form of a cognitive architecture (e.g. ACT-R [21]). (Such an architecture aims to create a complete simulation of cognition, including sensation, perception, memory, processing, and physical constraints on action — even emotion in some cases.) Although this approach might improve relevance to people, cognitive architectures present two major problems for a swarm intelligence approach:

1. they are often heavyweight algorithms, which limits their practicality in swarm-principled simulations with many individuals (although see [22] for a counter-example)

2. they negate one of the key benefits of swarm models: the establishment of causal connections between complex collective behaviour and its origins in the simple, local behaviour of individuals

The latter objection to swarm-intelligent application of cognitive architectures is the more important one. Given that we look to microscopic models to understand how aggregate behaviours emerge from individual ones, it is critical that the implementation of individuals be easy to understand. By introducing a cognitive architecture, however, simple rules that guide individual behaviour give way to complex cognitive simulations that tend to be analytically opaque (cognitive architectures are emergent systems themselves). This can lead to a problem (we term it the *scale mismatch problem*) in which a macroscopic cognitive layer overpowers the microscopic crowd model; the latter becomes little more than an arena, relegated to simulation of the physics and movement of agents, and is divorced from the real action unfolding (potentially even on a different timescale) within the cognitive simulator. This disconnection obscures the causal chain and it becomes difficult to explain the complex collective behaviour under investigation by reference to processes within the individual.

Although cognitive architectures have many benefits, incorporating them at the heart of a microscopic model would be a challenge. Luckily, it is not necessary to import a complete cognitive architecture in order to model cognitive factors in swarm models.

2.2 Behavioural repertoire

Sime has argued that people interact with the information available in the spaces through which they move [5]. This fits well with a modern definition of cognition; Dawson has argued that the hallmark of cognition is information processing [23], and the central assumption of cognitive science is that "the human mind is a complex system that receives, stores, retrieves, transforms and transmits information" [24]. Although modelling a complete cognitive architecture is not desirable, these observations suggest that a focus on information processing would help to resolve Sime's concerns about modelling people as non-thinking objects like ball-bearings; ultimately this would improve the models' relevance to people.

Any model is, by definition, an abstraction that discards detail through simplification in favour of explanatory power. Because of their focus on emergence in understanding collective effects, current microscopic crowd models (like many swarm models) already abstract and simplify behaviour at the level of the individual; they do so by emphasising *effects* rather than *mechanisms* at the individual level. Rather than spending computational resources on perfect physical realism, the models simplify the world to allow many individuals and a large environment to be simulated. Representation of physical space, for example, may be stylised into a grid while proximity to other agents (in reality a visual process) may be represented by stigmergic pheromone mechanisms (see [25] for a crowd model which adopts these principles). This abstraction process is what allows a model to use simple rules at the individual level to produce complex emergent effects at the aggregate level.

In order to avoid the scale mismatch problem, our integration of human behaviour must follow similar design criteria to the existing abstractions in microscopic models. With this in mind, when considering the lack of cognitive influence in crowd models it is not *only* internal processing that is lacking from the models: Human behaviour implies not only a mechanism for processing information but also a rich behavioural repertoire available for deployment by this cognitive mechanism. When we call microscopic crowd models into question because they do not do what people do, we are bemoaning the lack of this behavioural repertoire.

As swarm intelligence and microscopic crowd modellers, we are less concerned with processes internal to an individual and more concerned with the behavioural *actions* of individuals — as it is these that collectively produce emergent effects. Our methodology for integrating human behaviours, then, recognises the swarm focus on modelling effects rather than mechanisms at the individual level. In a swarm context, cognition is important to behavioural simulation *in its role as support for behavioural repertoire*; it is, then, secondary to behavioural repertoire. In other words, in microscopic models that are already strong abstractions of the world, it is the absence of behavioural repertoire — more than the specific cognitive process that underlies it — that calls into question applicability to human crowds.

On this view, we can take Sime's ball-bearing complaints into account by extending the behavioural repertoire of individuals. Our approach — as opposed to representing formal models of cognition through a cognitive architecture — is to implement cognitive behaviour and behavioural actions at the same level of abstraction at which the crowd models currently simulate movement. Consequently, our methodology expresses cognitive characteristics as simple local rules, fusing them with the simple local rules of the existing crowd model. We call these simple local rules Microscopic Human Factors. The methodology is characterised by simplified information processing that triggers behaviour and supports its unfolding over time where required. This simplified view of the human behavioural repertoire reduces the complexity of the cognitive simulation required; this brings the human factors modelling into scale with the existing aspects of the microscopic crowd simulation. This technique ensures that the causal connection between emergent crowd behaviours and local, individual rules is not broken.

In short, by reducing the fidelity of the cognitive simulation (instead focusing on behaviour) we adopt the same strategy that microscopic crowd modellers have used to simulate physical aspects of the simulation. This reduced fidelity is warranted because our goal is not to explain intra-agent processing but rather to induce emergent effects by grouping many agents together in a swarm-intelligent system.

We now turn to an explanation of the Microscopic Human Factors methodology for adding human behaviour to a microscopic model.

3 The Microscopic Human Factors methodology

The purpose of the methodology is to provide a framework for describing human behaviour that:

- guides modelling of *cognitive factors*: simple rules integrated tightly with swarm-principled systems such as a microscopic crowd models; and
- improves model transparency, facilitating input and feedback from domain experts in the modelling process and discussion.

The methodology is general and does not pre-suppose any particular implementation formalism. The methodology does not place limitations on the types of behaviours that can be modelled, but rather, according to swarm principles, enforces the perspective of the individual in explaining the mechanisms of crowd effects. Although the methodology is presented with crowds in mind, as discussed it is equally applicable to modelling of aggregations of non-human animals, particularly intelligent ones.

The methodology requires three distinct descriptions of a behaviour of interest at three different levels (see figure 1). Each level, or layer, is a complete description of the behaviour at a different level of abstraction. Each of these three distinct descriptions is an important part of the methodology, guiding the modeller to consider and disclose their position on what crowds do, how this crowd behaviour originates in individual behaviour and how the individual behaviour can be simulated in a specific implementation formalism. Without presenting all three levels, a modeller leaves unstated an important part of his or her theoretical position, and weakens the connection between his or her work and the real world.

We shall consider each level, its purpose in the methodology, and utility of each for localising and answering objections to the resultant model. To illustrate the three descriptions we shall refer to a trivial example human behaviour[2]: that of selecting a line at a supermarket (figure 2).

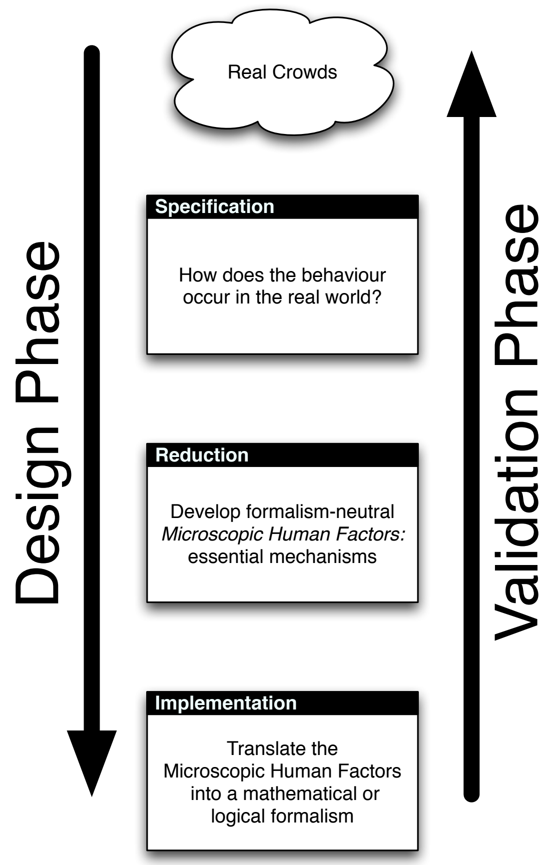

Figure 1: The three levels of the *Microscopic Human Factors* methodology. Each level is a complete description of the behaviour, inspired by knowledge or hypotheses about real crowds. During the design phase the layers are created in order of increasing abstraction (from the behaviour of real crowds) and specificity (in mechanism). During the validation phase the validation of underlying layers supports the layers above.

Specification. The first step in the methodology is to develop a clear statement of how a crowd behaviour, or aspect of crowd dynamics, occurs in the world. This may be a description of a well-understood behaviour, or a hypothesised behaviour for study. Its purpose is to develop a theoretical *specification* — a description at the aggregate level — for the behaviour proposed.

As real cognitive systems are highly interconnected and complex, no model can account for all of crowd behaviour in one simulation. It is important to be clear about what is being modelled, and the specification places a boundary around the simulation of interest. Specifying this boundary is critical because it amounts to a declaration that the specified aggregate behaviour is atomic, and that it can usefully be modelled in isolation. The specification provides both a starting point for modelling, and a desired end point for simulation.

At the specification level, discussion is at the level of real-world crowd behaviour and of the information system embedded in the surroundings; there is no need to discuss agents or modelling practicalities. As a theoretical contribution to the literature on crowd dynamics, this is a very important aspect of the specification. Researchers differ on exactly what is happening in crowds, and these discussions are usefully localised to the specification. Disagreement at this level suggests uncertainty about how people behave in crowds and why; by providing an implemented model to test a position clearly stated in the specification, a swarm-principled crowd

[2]For the purposes of this article, we have selected this simple behaviour in order to focus clearly on the methodology. For several explicit applications of the methodology to real-world problems, the reader is referred to references 4 and 20.

In this trivial example of the *Microscopic Human Factors* methodology, the task is to model selection of a line at the supermarket.

Specification. The specification requires a description of the way the behaviour works in the real world. We hypothesise that people make several types of observations and judgements when selecting a checkout line, in the interest of spending the least time in line. People evaluate the number of shoppers, the contents of others' carts and make judgements about the probable efficiency of the cashiers. We suppose that people prefer the express line to the regular line.

Reduction. Although there are several interesting phenomena at work in the specification, we must extract the essentials of this behaviour, and present it in an individual-focused, formalism-neutral manner for modelling. Although there are several judgements and decisions being made in the specification, we have chosen to reduce these to time as the essential prediction made by individuals. In other words, our model of line selection reduces, in a microscopic context, to an individual's time predictions. Our reduction assigns an approximation of the total number of items (across carts) in a line to be the key predictor of time. It is reasonable, however, to assume that the number of carts will also be considered, as will the type of line (express or not); these factors will be modelled as increases or decreases in each line's time estimation, respectively. The judgements about the cashiers, although interesting, are determined to be non-essential in a person's judgement of time, and are left to future work.

Implementation. At the implementation level we introduce the notation and formalism appropriate to the microscopic model of interest. In this trivial example let us simply assume that the model provides a set of lines $L = \{l_0, l_1, l_2, \ldots, l_n\}$, a determination, $\mathrm{e}(a)$ of whether each line is an express line (with $\mathrm{e}(a) = 1$ if line l_a is an express line, else 0), a set of carts $C_a = \left\{c_{a_1}, c_{a_2}, c_{a_3}, \ldots, c_{a_m}\right\}$ in each line l_a, and a weight of each cart C_{a_b} given by $\mathrm{w}(a, b)$.

The implementation of our line selection algorithm is as follows. First, each line is assigned a unitless time prediction t_a (shown below) then the agent selects the line with the lowest predicted time. The score is determined by summing the weight of the carts in line, penalising each cart with 10 units of weight, and dividing the weight by 2 for express lines.

$$t_a = \frac{1}{1 + \mathrm{e}(a)} \sum_{b < |C_a|} (10 + \mathrm{w}(a, b))$$

Figure 2: A trivial example of the Microscopic Human Factors methodology.

model can add great value to this kind of academic debate.

In our trivial supermarket example of the methodology in action, figure 2, the specification is of the hypothetical type. Note that the description is given in terms of the behaviour of real shoppers. There is no mention of agents, nor of modelling formalisms or constraints. We simply describe the behaviour in question and how we believe it works.

Reduction. The second step in the methodology involves *reducing* the behaviour described in the specification to simple local rules. It is these local rules that we term *Microscopic Human Factors*. This stage is the most important step in the methodology, fulfilling both a theoretical purpose and a practical one.

Practically speaking, it is at this level that an abstraction of the behaviour is developed, describing the *essential* elements of behavioural repertoire that underlie the collective effects of interest. The modeller details the information that is used by individuals to trigger or carry out a behaviour, the local sources of this information and how it is processed. The modeller also covers physical constraints, situational factors and physical abilities of the individuals and how the behaviour is expressed over time. In short, the modeller provides a complete mechanism for the behaviour — and, as swarm models are based on individual rules, the description at this level is from the perspective of the individual within the crowd.

The reduction also amounts to a theoretical position: the enunciation of a mechanism sufficient to reproduce the behaviour of individual people, who collectively will produce the aggregate behaviour described in the specification. This discussion concerns the manner in which a particular crowd behaviour occurs as an interaction between individual people and the space around them. This is an important contribution of the model because

— as compared to aggregate crowd-level behaviour — there is much that researchers do not understand about the behaviour of individuals in crowds.

It is not the data produced by running a model scenario that is its ultimate contribution to our understanding of crowd dynamics, nor the techniques used in its implementation; rather, it is the reduction, which explains how individual behaviour can underlie that of the crowd. When viewed in this light, it is natural to emphasise the importance of the reduction in clearly enunciating the individual crowd member's contribution to crowd behaviour.

The points of the last paragraph argue that the language of the reduction should be one of people in crowds, and not one of modelling mechanism. It being understood that crowd dynamics is a broad interdisciplinary pursuit, crowd modellers benefit from input from psychologists, sociologists, historians, physiologists, physicists, and many other disciplinarians besides. These other disciplines are not interested in models for models' sake, but rather in understanding the mechanisms of crowd dynamics and in understanding how to increase safety in crowded conditions. As a theoretical and mechanistic contribution, the model's reduction is the point of contact with these other disciplines; it must thus be presented in a formalism-neutral and accessible style that is easily understood by non-modellers. This allows a conversation to develop between modellers and other disciplines concerning the mechanisms of crowd behaviour in which models form only a part of the discussion, ultimately benefiting both the modelling of crowds, and study of their properties more generally. Introducing issues of implementation detail would obscure the theoretical contribution just described.

The reduction level exists, in part, to require a *separation* between simplifications introduced for parsimonious understanding of crowd behaviour and simplifications introduced for reasons of modelling practicality; this is another reason to require formalism-neutral language in the reduction. The reduction must be a defensible abstraction from a behavioural perspective in order to ensure that simulation results are ultimately applicable to human crowds. It is not that modelling simplifications are inadmissible, but they are not hypotheses about crowds and must be introduced for what they are: implementation constraints. By making the reduction level explicit, the implementation formalism that most naturally captures the Microscopic Human Factors can be preferred.

In summary, models that do not make the reduction explicit leave a vital step of the modelling process to the imagination of the reader: Without the reduction we do not know whether an implementation detail is supposed to be a statement about real crowds, or whether it is simply an artefact of the formalism selected. Moreover, a reader who is not versed in the specific implementation formalism may perform this imagination incorrectly. If modellers wish their model results to influence the wider, non-modelling audience then a reduction must be provided to ensure — at the point of contact with this wider community — that the mechanism of the model is clear.

The points of the preceding paragraphs are demonstrated in the trivial example of figure 2. First, although the reduction is a complete description of the behaviour, we have performed an abstraction, focusing on *time* as the essential representation made of the lines by the agent. This is a theoretical position as to how the information processing in the line selection behaviour essentially works. Consequently, we provide for a correspondence between the information available to the person (load in the carts, number of carts, line type) and the estimated time. We have identified an aspect of the specification (cashier judgements) as interesting, but non-essential. Second, the reduction is formalism-neutral and individual-centred; this means that the reduction has not pre-supposed what kind of model should be used, but simply focuses on what we believe to be true from the individual shopper's perspective. (The reduction is equally well suited, for example, to implementation in a numerical simulation or in a multi-agent system, but only from an individual's point of view.) The reduction shows what we think is important about the specification from an individual perspective without confusing matters by what it is practical to model using any given approach. Third, the reduction keeps the discussion about the essential mechanisms in crowds accessible to domain experts in supermarket planning, even those who may not be familiar with the formalism selected for the implementation.

Implementation. What remains after the two levels just described is the level of *implementation*. The purpose of the third level is to harmonise the reduction, a functional description of the behaviour from an individual perspective, with an implemented microscopic crowd model. The implementation brings with it not only the advantage that the reduction can be simulated, but also that the requirement to implement is itself a test of the completeness of the specification and reduction.

This level involves a translation of the rules into a specific mathematical/logical formalism. In the case of a new model, this may involve also the selection of an adequate formalism. It is at this stage that the behaviours must be implemented in such a way as to interface closely (i.e. at a similar scale) with the existing rules of the model. From a modelling perspective, an emphasis at this level on scale match yields intertwining physical and behavioural rules in which neither overshadows the other; this yields a model that retains its ability to connect individual action to emergent collective effects.

Some formalisms may be more suited to a natural expression of the reduction than others. Implementation, however, is an area that is often poorly understood by non-modellers, leaving discussion of the model's implementation primarily to modelling researchers. A debate about the appropriateness of applying a particular formalism (with its particular mathematical rules and logic) to the implementation of a fixed specification and reduction can be productive. For this reason, isolating formal details at the level of the implementation is vital; it allows separation of objection to a model on the grounds of implementation from objections that the model takes an inadequate view of crowd mechanisms. In the limit, the implementation details can be altered to make the formalism a satisfactory expression of the reduction.

The simple implementation in figure 2 is sufficient to illustrate two important aspects of models built with our methodology. First, the line selection is done directly using the information represented in the model concerning lines, carts, etc. We do not need to import a separate line-selecting cognitive architecture to sit alongside an existing physical cart-pushing model. Second, the implementation makes a simplification for reasons of modelling practicality (i.e. it does not provide for an exact translation of the reduction into the implementation); while the reduction expresses time in terms of number of items in carts, the implementation works in terms of cart weights. Assuming this is a significant distinction, the modeller would need to argue for the appropriateness of this alteration in the behaviour (perhaps by drawing a parallel between weights and item counts in grocery store cart studies).

Exploring the mechanism. Swarm modelling has more to offer than simple reproduction of emergent collective effects; additional benefits are obtained if a *correspondence* is made between abstract rules at the reduction level and the formal structures that implement these rules at the implementation level. In this case, the model provides a locus for experimentation with the mechanisms anticipated by the reduction — allowing the modeller to modulate or even deactivate some aspects of the simulation. This experimentation can lead to an understanding of how and why the rules provide for emergent crowd level effects as observed in the simulation, and how the new rules interact with previous behaviours of the model. By extension, this provides an opportunity to determine how varying the rules and the situation can give rise to different crowd behaviours, and hence some idea of the range of potential crowd behaviour.

4 Model validation

Validation refers to the process of determining the relationship between a model and the world. In this section we explore challenges in validating complex models, what these models can tell us about the world, and how the Microscopic Human Factors methodology helps to increase model validity.

4.1 Validation of complex models

The study of complex phenomena, such as aggregate behaviour, challenges standard methods of investigation. A crowd, for example, is a massively complex system in which large numbers of cognitive actors interact locally, generally with incomplete information and no central coordination. This swarm system cannot be studied with familiar laboratory methods — like dividing the problem into parts and determining, for each part, the effect of a clear independent variable (e.g. signal intensity) on a simple dependent variable (e.g. reaction time). The category difference here (and the methodological difficulty) occurs because it is precisely the complexity in the situation that is of interest.

The complexity also foils the standard model validation approach, in which we simulate a real-world scenario and the model's output is compared back to the real data. As discussed in section 2.2, the Microscopic Human Factors methodology trades fidelity in individual simulation for tractability and transparency; thus we do not expect that data from organisms in the wild will accord quantitatively with Microscopic Human Factor models. Furthermore, an objection on quantitative grounds to an abstract model built on the Microscopic Human Factors methodology would miss the point of such a simulation: to study aggregate behaviour in abstract situations, to provide qualitative evidence that collective behaviours arise from specific individual behaviours, and to demonstrate mechanisms that can account for these behaviours. Even in a laboratory environment, experimental participants cannot be sufficiently controlled to isolate the effect of a behaviour of interest in the same way that a Microscopic Human Factors model can. Direct experimentation, in short, is not a match for the qualitative, abstract focus of models that aim to draw general conclusions about the contribution of individual behaviours to aggregate behaviours.

The fact that direct experimentation is not a good match for abstract modelling does not mean that such a model has no validity. We take validity to mean relevance to the model's prototype: the real system in the real world (e.g. real crowds). Microscopic crowd models have much to offer in the department of relevance to the real world, and as we shall now see, the methodology supports these arguments of validity.

4.2 Validation with the methodology

One important benefit to the three-level methodology proposed is that, by isolating three different views of the model, multiple points are introduced at which the validity of the resultant model can be naturally considered. In section 3 we presented the levels in order of increasing abstractness. Once the simulation has been made, however, the order of levels can be reversed for the purpose of considering model validity. As we move upwards through the three layers we can use an implemented model that has adopted the Microscopic Human Factors methodology to answer different validation questions.

Implementation. At the level of the implementation we determine whether the Microscopic Human Factors have been appropriately implemented. This includes simple checking, through experimentation, that the implemented model embodies the Microscopic Human Factors described at the reduction level.

A valid implementation also entails selection of an appropriate implementation formalism. An appropriate formalism does not deviate from the reduction by introducing excessive simplifications for reasons of implementation constraint. (Certain reductions may be very difficult to implement, for example, in a purely numerical simulation. This is because information processing in crowds often entails decision-making based on what has happened earlier, suggesting a requirement for memory and logical operators.)

An appropriate implementation formalism also helps to avoid an excessively simplified view of individuals within the aggregate. To return to a crowd example discussed in the introduction, establishing a correspondence between people in crowds and seeds in a hopper may lead to an over-emphasis on physical simulation. The reduction's focus on individuals as the active part of the simulation is key to preventing this kind of over-simplification; it remains, however, to ensure the over-simplification does not recur at the implementation level. In other words, once we note at the level of reduction that individuals are active, the selection of an implementation with static ball-bearing-like agents seems forced.

Reduction. The purpose of this level is to develop simple local rules that provide an abstract mechanism for the behaviour as described in the specification. The validation question at this level, then, concerns whether the simple local rules actually implement the behaviour described at the specification level. If the reduction is valid, then tests run with the local rules are essentially tests of the theory proposed at the specification level. Although there may be multiple ways to reduce the specification to simple local rules, as discussed in section 2.2 we are less concerned with the internal mechanism of agents and more concerned with behavioural repertoire; validity at this level, then, is present if the behaviours generated by the Microscopic Human Factors are sufficient to implement our theory of how crowds work in the world (the specification).

Our trivial example (figure 2) illustrates another point relating to the validity of the reduction. In many cases the behaviour of the specification will be richer than we wish to model in a particular abstraction, resulting in a gap between a richly detailed specification and an abstract reduction. (In the example it was judgements about cashier efficiency that were seen to be interesting but extraneous to the purpose at hand.) This gap is intrinsic to abstract modelling (as suggested in section 2.2) and goes to the heart of the general challenges in validating abstract models (described in section 4.1); we must create abstractions to limit the scope of a phenomenon and achieve understanding, but these abstractions must be defended. The art of abstract modelling is in selecting the right abstraction, a process that must be confirmed as described in the preceding paragraph. Remaining gaps in the reduction usefully identify areas of interesting future work. In the limit, these pieces of future work can be modelled, and their effects on the basic behaviour confirmed.

Specification. The purpose of the specification level is to describe a theory or hypothesis explaining a particular aspect of crowd behaviour that unfolds in the real world. Thus, validation of the model at this level involves determining the extent to which the model is a representation of the behaviour of real crowds. Bearing in mind the discussion of validity of complex models (section 4.1), let us consider how Microscopic Human Factors models are relevant to real crowds.

The validation task, to be carried out through experimentation upon the model, is to show that the specification is *sufficient* to explain the real world behaviour of interest. Sufficiency is stressed here, because the specification is subject to the principles of multiple realisability: There may be multiple versions of the specification that can produce the behaviour. To discuss how a model that expresses sufficiency can tell us about the real world, let us consider two cases for application of the methodology. (Here we assume that, through experimentation on the model, the sufficiency of the specification has already been shown.)

Table 1: Summary of the levels of the Microscopic Human Factors methodology, their content, purpose and value in validation

Level	Design Goal	Vocabulary Used	Validation Goal
Specification	Theory about world	People in crowds	Theory sufficient to account for real-world behaviour
Reduction	Essential local rules	Individuals in crowds	Rules sufficiently capture essence of specification
Implementation	Tight model integration	Formalism-specific	Implementation corresponds to rules & formalism is appropriate

In the first case we use the methodology to study a behaviour that is well understood. In this case, the specification is motivated by — or derives directly from — existing knowledge. Here the modeller can make an argument that the specification is more than sufficient to account for the behaviour because the existing knowledge provides a reason to prefer this particular specification over other possible mechanisms. The model can then be used to draw conclusions about how the real world works, supported both by the methodology (that makes the claims about the world explicit and examines their effects) and the literature (or other reasons that cause us to believe these claims are well-supported).

In the second case, the methodology is used to study a behaviour that is suspected, or to investigate 'what-if' scenarios. Here, the model's contributions are more theoretical in nature. As noted above, microscopic crowd models are generally focused on the qualitative study of emergent effects; they are one tool to help in understanding human behaviour in crowds. Like in much scientific endeavour, particularly in the social sciences, we look for corroborating evidence from multiple disciplines and approaches to answer complex questions; individual methods, such as models built with the Microscopic Human Factors methodology, generally contribute by narrowing the field of inquiry rather than through definitive answers. In this respect, the methodology proposed, with its emphasis on local mechanisms, can contribute by demonstrating the sufficiency of individual-based human-centred hypotheses for crowd effects.

In both cases, the model adds to our knowledge of crowd behaviour by testing the sufficiency and consistency of a theory to reproduce crowd effects. This allows us to make in-principle arguments about how crowds can and cannot work. Where we wish further certainty, we can use the narrowed focus of the model to produce hypotheses that can be subsequently tested in real crowds, or — if the hypotheses cannot be practically and/or ethically tested — make philosophical arguments in support of their applicability to real crowds. Additional support for the theory may derive from case studies of natural experiments, and the results of other disciplines.

Multiple realisations of a particular complex behaviour may be possible; in the case of crowd models, however, those that take account of Sime's concerns about modelling *people* rather than ball-bearings are to be preferred. The advantage of the Microscopic Human Factors methodology is that it provides an explicit mechanism for ensuring human behaviour is integrated at the heart of a microscopic crowd model by representing it through simple local rules. This, then, is a key accomplishment of the methodology: By focusing on local, individual behaviour we can take account of Sime's concerns, integrate behaviours rooted in cognition, stop modelling ball-bearings and thus improve the relevance of microscopic models to people. With this methodology, and the simulations derived from it, we are in a position to better understand what emergent behaviours are likely, and to study how the behaviours of individual people interact in a large-scale movement simulation.

5 Conclusion

The Microscopic Human Factors approach has two main aims:

- as current models have had difficulty incorporating human behaviour in crowds, to guide integration of human factors into microscopic models.
- as crowd dynamics is an interdisciplinary research area, to increase clarity in discourse between modellers and non-modellers (a key to realising the potential contribution of modelling results to the research area more generally).

In modelling a crowd behaviour, the methodology requires three descriptions in increasing levels of abstraction (see table 1). At the specification level, a clear description of how the targeted crowd behaviour occurs in the world is presented. At the reduction level, the specification of how the world works is reduced to a series of abstract individual-centred rules that provide for the behaviour. All modelling and implementation details are relegated to the implementation level, at which the reduced rules are translated into a specific formalism for simulation.

As regards the first aim, having rejected the cognitive architecture approach in the swarm-principled context, the methodology is centred around increasing behavioural repertoire of agents as an interaction between the environment and the individual's social and physical context. We focus on generating the individual behaviours that underlie crowd effects, with low fidelity in simulating processes internal to the individual. This mirrors the approach that swarm-principled microscopic crowd models take in abstracting the physical aspects of crowd simulation as well. The similarity in abstractions produces a scale match that allows cognitive and physical aspects of simulation to be tightly integrated in the reduction and the implementation of the model. Ultimately, this approach allows the modeller to increase the relevance of modelled individuals to real people, and hence the applicability of the resulting model to real crowds.

As regards the second aim, the methodology provides benefits to modellers and non-modellers alike. Due to enhanced plausibility of model results, modellers benefit when their model specifications are informed by (and accord with) our knowledge of crowd behaviour. When a model's explanation of crowd behaviour is made explicit in a formalism-free and individual-focused reduction, crowd experts outside the modelling community can understand the applicability of model results to their own research. The goal is to create a positive feedback loop in which models and other research methods together improve our understanding of crowd dynamics. With benefits to be realised on both sides, the Microscopic Human Factors methodology is designed to support this feedback loop.

Crowds are aggregations of humans with complex emergent behaviour underlain by individual behaviours and interactions. As such, the methodology proposed in this article is also well suited to studying aggregations of other species. This is particularly the case for social species with non-trivial cognitive abilities, in which understanding individual behaviour through direct experimentation may be challenging. With people, as with other animals, the Microscopic Human Factors methodology helps to forge connections between swarm-principled models and the world; these connections will support our understanding of both natural and artificial swarm intelligence.

References

[1] J. A. Schellinck, *A General Perception Based Framework for Modelling Animal Aggregation*. PhD thesis, Carleton University, Ottawa, 2008.

[2] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From natural to artificial systems*. New York: Oxford University Press, 1991.

[3] C. M. Henein and T. White, "The microscopic model and the panicking ball-bearing," in *Pedestrian and Evacuation Dynamics 2008* (W. W. F. Klingsch, C. Rogsch, A. Schadschneider, and M. Schreckenberg, eds.), pp. 569–576, Berlin: Springer, 2010.

[4] C. M. Henein, *Crowds are made of people: Human factors in microscopic crowd models*. PhD thesis, Carleton University, Canada, 2008.

[5] J. D. Sime, "Crowd psychology and engineering," *Safety Science*, vol. 21, pp. 1–14, 1995.

[6] E. Quarantelli and R. R. Dynes, "When disaster strikes," *Psychology today*, pp. 66–70, Feb 1972.

[7] P. J. Keating, "The myth of panic," *Fire Journal*, vol. 76, no. 3, pp. 57–61, 1982.

[8] J. D. Sime, "Affiliative behaviour during escape to building exits," *Journal of environmental psychology*, vol. 3, pp. 21–41, 1983.

[9] N. R. Johnson, "Panic and the breakdown of social order: popular myth, social theory, empirical evidence," *Sociological focus*, vol. 20, pp. 171–183, 1987.

[10] N. R. Johnson, "Panic at the "Who Concert Stampede": An empirical assessment," *Social Problems*, vol. 34, pp. 362–373, 1987.

[11] J. Chertkoff and R. Kushigian, *Don't panic: The psychology of emergency egress and ingress*. London: Praeger, 2001.

[12] W. E. Feinberg and N. R. Johnson, "Primary group size and fatality risk in a fire disaster," in *Human behaviour in fire (2nd international symposium)*, (London), pp. 11–22, Interscience, 2001.

[13] L. Clarke, "Panic: Myth or reality?," *Contexts*, vol. 1, pp. 21–26, 2001.

[14] J. D. Averill, D. Mileti, R. Peacock, E. D. Kuligowski, N. Groner, G. Proulx, P. Reneke, and H. Nelson, *Occupant behavior, egress and emergency communications.* NIST NCSTAR 1-7, Washington: National institute of standards and technology, 2005.

[15] J. D. Sime, "Designing for people or ball-bearings?," *Design Studies*, vol. 6, pp. 163–168, 1985.

[16] C. M. Henein and T. White, "Agent-based modelling of forces in crowds," in *Multi-agent and multi-agent-based simulation* (P. Davidsson, B. Logan, and K. Takadama, eds.), Lecture Notes in Computer Science v. 3415, (New York), pp. 173–184, Springer, 2005.

[17] C. M. Henein and T. White, "Information in crowds: The swarm information model," in *Cellular automata: 7th international conference on cellular automata for research and industry (ACRI 2006)* (S. El Yacoubi, B. Chopard, and S. Bandini, eds.), Lecture Notes in Computer Science v. 4173, (New York), pp. 703–706, Springer, 2006.

[18] C. M. Henein and T. White, "Macroscopic effects of microscopic forces between agents in crowd models," *Physica A*, vol. 373, pp. 694–712, 2007.

[19] C. M. Henein and T. White, "Front-to-back communication in a microscopic crowd model," in *Pedestrian and Evacuation Dynamics 2008* (W. W. F. Klingsch, C. Rogsch, A. Schadschneider, and M. Schreckenberg, eds.), pp. 321–334, Berlin: Springer, 2010.

[20] C. M. Henein and T. White, "Microscopic information processing and communication in crowd dynamics," *Physica A*, in press, 2010. DOI: 10.1016/j.physa.2010.05.045.

[21] J. R. Anderson and C. Lebiere, *The atomic components of thought*. Mahwah, NJ: Erlbaum, 1998.

[22] T. Stewart and R. West, "Python ACT-R: A new implementation and a new syntax," in *Proceedings of the 12th annual ACT-R workshop*, (Trieste, Italy), 2005.

[23] M. R. W. Dawson, *Understanding Cognitive Science*. Oxford: Blackwell, 1998.

[24] N. Stillings, M. H. Feinstein, J. L. Garfield, E. L. Rissland, D. A. Rosenbaum, S. E. Weisler, and L. Baker-Ward, *Cognitive Science: An introduction*. Cambridge: MIT Press, 1987.

[25] C. Burstedde, K. Klauck, A. Schadschneider, and J. Zittartz, "Simulation of pedestrian dynamics using a two-dimensional cellular automaton," *Physica A*, vol. 295, pp. 507–525, 2001.