Bus, bike and random journeys

Crowdsourcing aid distribution in Ivory Coast

Delivering supplies in poor rural areas is difficult and expensive. But people travel; and statistics can piggyback aid supplies on to the network of everyday journeys. **James McInerney**, **Alex Rogers** and **Nicholas R. Jennings** explore an imaginative solution to getting aid to the countryside.

You do not need an expensive dedicated delivery service if your parcels can hitch informal rides. But you have to be able to predict where they will end up...

In many developing countries, half the population lives in rural locations. This is the case in Ivory Coast, Ghana, Liberia, Nigeria; and West Africa is far from alone in providing examples. Outside the cities, access to school materials, medical supplies, mosquito nets, and clothing is restricted[1]. Distribution to such places typically requires direct road transport, which is time-consuming and requires bulk volume to be cost effective. In response to these limitations, alternative, and imaginative, methods of aid distribution have emerged in recent years. For example, Pack For a Purpose (`http://www.packforapurpose.org`) is a non-profit organisation that asks tourists who already have a trip planned for one of 47 developing countries to bring small items (e.g., pencils, deflated soccer balls, stethoscopes) in their spare luggage capacity. Another scheme is Pelican Post (`http://www.pelican-post.org`), which asks donors to send books by post to developing countries. These are promising schemes. However, they fail during periods of conflict, (such as during the post-electoral violence in Ivory Coast in 2011) and they rely on direct outsider support. Arguably it is preferable to empower local populations to be part of their own solutions wherever possible.

Recently we proposed a new distribution method that uses the natural movements of local people to distribute physical packages from one location to another[2]. We considered the possibility of opportunistically using the pre-existing travel routines of a set of local participants by asking them to pick up a package from one exchange point (at a place that they normally visit, at a time that they normally visit it) and then drop it off at another exchange point (such as a lockbox or village store) that is also part of their regular itinerary. By chaining together the mobility of several participants we may cover a large area, possibly a whole country, without having to deploy more expensive and time-consuming infrastructure – and without having to make any extra journeys or use any extra fuel.

For example, if we wish to deliver a package of mosquito nets from the largest city, Abidjan, to a rural village in the west of Ivory Coast, we may first ask Ibrahim, who lives in Abidjan, but often visits his sister in Gagnoa (260 km to the west) on weekends, to pick up the package near his house and drop it off near his sister's house in Gagnoa, when he is there anyway. We may then ask another participant, Phillipe, who lives in Gagnoa, but who works in Taï national park on weekdays (driving past the village each day without realising) to drop the package off at the village on his way to work, taking it a further 90 km. In this way, the participants do not have to significantly change their schedules or travel long distances that they would not have otherwise travelled. The journey of the package in this example delivery is illustrated in Figure 1. Participants in other deliveries might travel by private car, or by bus or motorbike, or

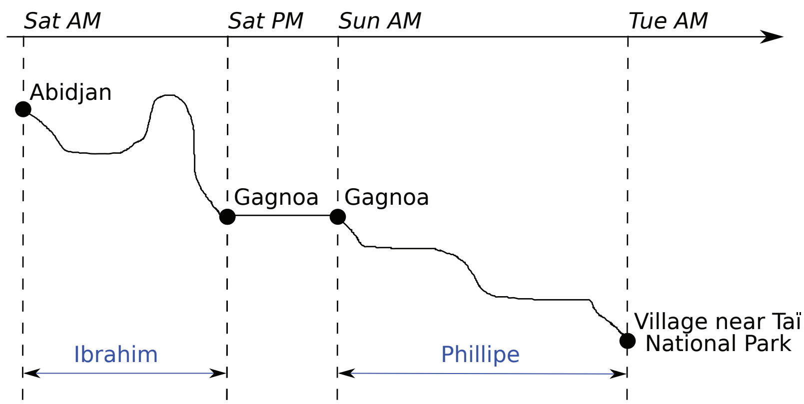

Figure 1. Example trajectory of a physical package across Ivory Coast using crowdsourcing

even, in the shorter last stages of a delivery, by bicycle or on foot.

While potentially appealing, this vision of crowdsourced delivery faces social issues related to trust (e.g., theft or loss) and incentivisation; we briefly discuss these at the end of this article. It faces significant technical barriers as well. Specifically, how should we select the task assignments – the couriers on whom we are piggybacking – to minimise the length of time the delivery will take? Optimising routes is a recurring problem in computer science, and a variety of algorithms have been invented to do it efficiently, without requiring unfeasible amounts of computation time and storage (Dijkstra's algorithm and A* search are two that will be known to specialists). The twist in crowdsourcing settings comes from the unavoidable fact that we are relying on humans to perform tasks, and we can never be really sure how they will behave. Ibrahim may be ill and make no journey, or his sister may be ill so that he makes an extra one, or there may be a family wedding that takes him somewhere else entirely. In package routeing, we face uncertainty about when a participant will choose to travel to the next stage of the package's route. Fortunately, there exists a suite of tools in statistical machine learning and artificial intelligence that allow us to automatically make a robust set of sequential decisions that minimises the delivery delay under human location uncertainty. This is despite the fact that we only have messy and incomplete data about the participants' locations.

Learning mobility

The first component in our system is a probabilistic model of human location behaviour. It allows us to predict, given what we know about an individual's past locations, where is he or she expected to be at a given time in the future. It is probabilistic in the sense that each prediction is represented as a probability distribution over all possible locations, indicating how likely the individual is to be at that location for the given time (which sums to 1, of course, because the person must be somewhere, and they cannot be in more than one place at a time). Before deciding what form the model should take, we need to know what the data looks like. How can we tell where people are now, and where they have been in the past? The answer lies in a phenomenon that has transformed rural African society over the past decade: the mobile phone.

In Britain and the United States smartphones come equipped with GPS. Most mobile phones in Africa do not have such fine-grained global positioning system sensors, so we must make do with cell tower data, which is less precise. It tells us which tower was in contact with which phone, and when. Cell tower data consists of a set of observations (i, x_n, t_n) indicating that participant i was observed near cell tower x_n at date and time t_n. There are three main factors that influence the design of the model:

1. *Cell allocation noise.* The cell tower observations provide discrete measurements on the individual's likely location. However, "likely" is key here: there may be a choice of several towers that the phone can connect to (especially in urban environments) at any single location. At busy times the network operators might not allocate a phone to the tower that is nearest to it. Knowing which tower a phone is connected to does not tell us for certain which area the phone is in. This uncertainty is decided by outside factors (the network operators), and we treat it as noise. Our approach needs to isolate the human presence information in the cell tower allocation to phones and ignore other factors. This implies the need to infer the locations, each of which may be statistically associated with several cell towers.


Minibuses link rural towns, for people – and for aid? iStockphoto/Thinkstock

2. *Sporadic observations*. Since the cell tower only records data when a phone call or text is made (about seven times a day on average) we need a method that can fill in (extrapolate from other observations) large periods of no observability.
3. *Short duration*. We are not guaranteed a long history of data for all individuals. This, combined with the fact that each day may have only a few (or zero) observations, makes learning challenging. Overfitting is a danger when the training data (perhaps just a few weeks of observations) contains characteristics that do not generalise to the rest of the individual's behaviour (i.e., beyond a few weeks). In other words, our data about Ibrahim may come only from June. Our deduction that he also visits his sister weekly in the other months of the year may be quite erroneous.

These considerations suggest the use of the Bayesian framework, which allows us to assume the existence of latent variables that abstract away from the variability of cell allocation (factor 1), and make custom assumptions about the smoothness of location (factor 2). Furthermore, Bayesian non-parametric methods can provide us with powerful guards against overfitting (factor 3).

In more detail, we assume the existence of latent discrete locations that are associated with each observation (i, x_n , t_n), and correspond to places in individual i's routine life (e.g., home, work). These locations are latent (i.e., hidden) in the sense that it is not possible to directly observe a person visiting them, but their existence is implied by the patterns of cell tower visits in space and time. For example, if there are two cell towers near my home, on some occasions my phone might be assigned to one, while on other occasions it might be assigned to the other. Although these assignments are random, over a long enough time we can make an educated guess about the existence of a single place of interest (i.e., my house) at that location. To do this, we can use mixture modelling to infer both the nature and number of latent locations from the data (using a Bayesian non-parametric approach called a Dirichlet process).

To address the problem of filling in large periods of missing data, we assume that behaviour is periodic. Specifically, our model assumes both weekly and daily periodicities in behaviour. The other motivation for using a periodic mobility model is that it allows predictions for arbitrary future time points, enabling optimisation to be done several days ahead. For example, on Tuesday at 2 p.m. in 6 months' time we predict that Ibrahim's movements will be similar to his travel on other Tuesdays, and on other days at 2 p.m.

Now that the model has been specified, it is possible to learn its parameters from any given set of observations. Ibrahim's past travel routines have given him his own unique set of parameters in the model. Once this is done, with the parameters in hand, we are free to disregard the observed data because everything we care to know about an individual's past behaviour is captured in those parameters.

But how do we know the model's predictions will be any good? This is an empirical question about the quality of our assumptions.

Ibrahim's past movements let us predict his future travel plans

George Box, the acclaimed English statistician who died earlier this year, once said "all models are wrong, but some are useful"[3]. One way of assessing the usefulness of our model is to check how much probability mass it assigns to future locations that were subsequently visited by the person – to check our predictions for Ibrahim against the places that he does actually come to visit. We did this for rather more individuals than just Ibrahim. We used a data set from the Orange phone network, describing the cell tower assignments of 50 000 individuals in Ivory Coast. Data from Ivory Coast is particularly interesting for our purposes because the country has faced significant humanitarian and infrastructure crises in recent years, in part due to the post-electoral violence. We computed the parameters of all 50 000 individuals, but held back one observation per person to be used exclusively in testing the predictions (otherwise we would be testing with data that is, in some way, already represented by the parameters, making the reported performance artificially high). Comparing against the next best approach, we found that our model assigned 2.4 times as much probability mass to future locations on average.

So we know that our predictions are pretty good. Next we must work out how we can use them to make optimal decisions about the package route and task assignments.

Optimisation

The optimisation problem is challenging because decisions made in the present affect what decisions can be made in the future. For example, sending the package to the west of the country limits the pool of participants to whom we can next assign tasks. A principled way of solving such sequential decision-making problems using existing methodologies exists in the form of the Markov decision process (MDP).

An MDP describes what happens when an agent – a person, or a robot, or a piece of software – performs an action without knowing exactly what its effect will be. The uncertainty surrounding the effects is represented by a probability distribution that describes the next state of the agent after it performs

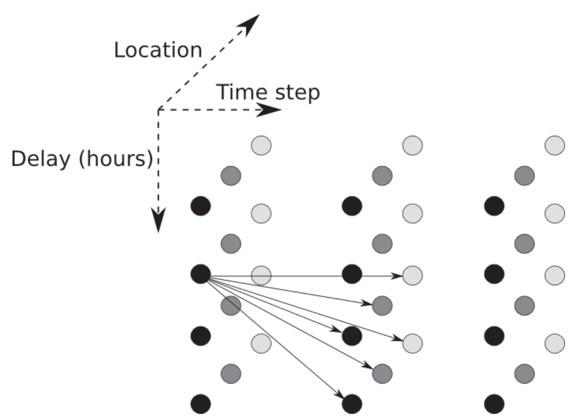

Figure 2. A subsample of the states of our Markov decision process, illustrating the random transition after a single action. Each row represents a different delay, each column represents a different time step, and each shade of colour represents a different location

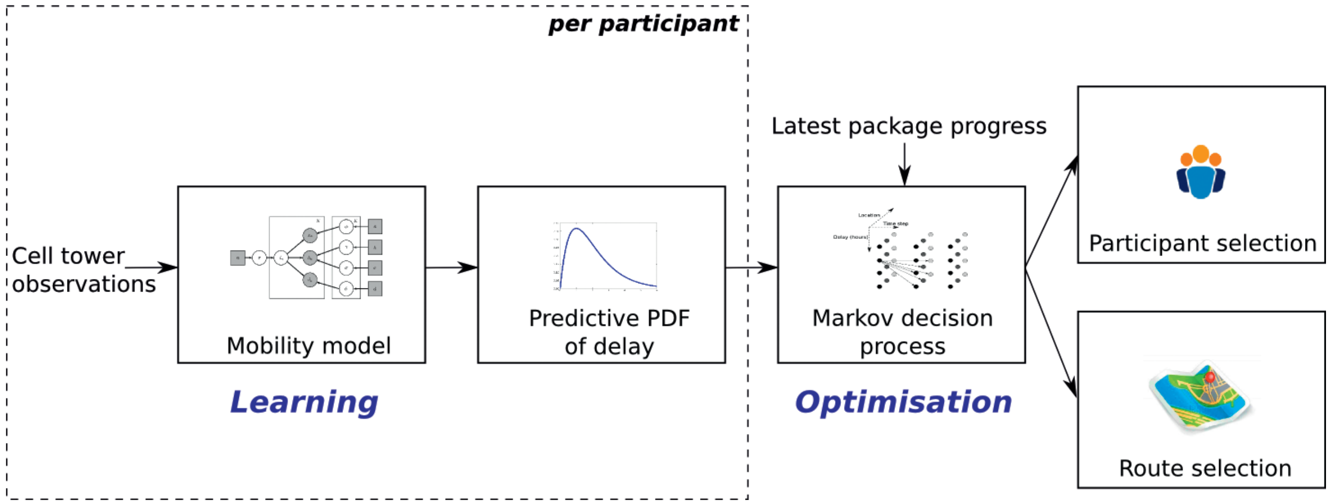

Figure 3. Complete crowdsourcing task assignment system for package delivery under human location uncertainty

an action at its current state. The exact interpretation of states and actions depends on the scenario being modelled. For example, the state of a robot on the surface of Mars could represent its position, while actions could represent the activation of various motors on the robot. The motors might move the robot 3 inches to the north – or might result in its falling down a small cliff. In our case, the set of states represents the joint combination of delay and location in Ivory Coast (see Figure 2), while the set of actions represents the assignment of delivery tasks to participants.

One attractive feature of MDPs is that they come with a set of established methods (e.g., value iteration, policy iteration) for finding the optimal policy that specifies the best action to perform at any given state. Optimal in what respect? To answer this, we need to assign a measure of desirability, or utility, to each state. For the robot, we are happy if it moves to the north, much less happy if it falls down a cliff. In our scenario, the utility is simply the delay the package has experienced so far, though more elaborate representations of utility are certainly possible (e.g., a quadratic function of delay that penalises higher delays much more than lower amounts of delay, or a cost associated with involving a new person that represents the risk that they will lose or steal the package).

In theory, all we have to do is run policy iteration on our MDP and we will have the best task assignment for each state of the package. But there is a catch. Consider again our representation of each state as the combination of a location and possible delay. A moment's thought will make clear that there is no limit to the amount of delay the package may experience. Delays of a week, a month, or even a year between steps in the route, though increasingly unlikely, are not ruled out by our mobility model or the scenario. This presents a major problem in searching the space of optimal decisions. This situation is fairly common for real-world problems: while it is easy to represent a scenario as an MDP, unless you have some clever formulation of states you will sometimes find it computationally unfeasible to identify an optimal policy in practice. Is such a formulation possible here?

To see how we might represent states more compactly, we need to go back to the mobility model. The function of interest is the probability density function describing how long it will take participant i to bring the package from one location to the next, assuming that he was assigned the task after the beginning of the package route. The periodic nature of the model means that this function is also periodic. This is a good thing, because it means there are only a limited number of values the function can take (assuming discrete delays of fixed time blocks, e.g., hours or half days).

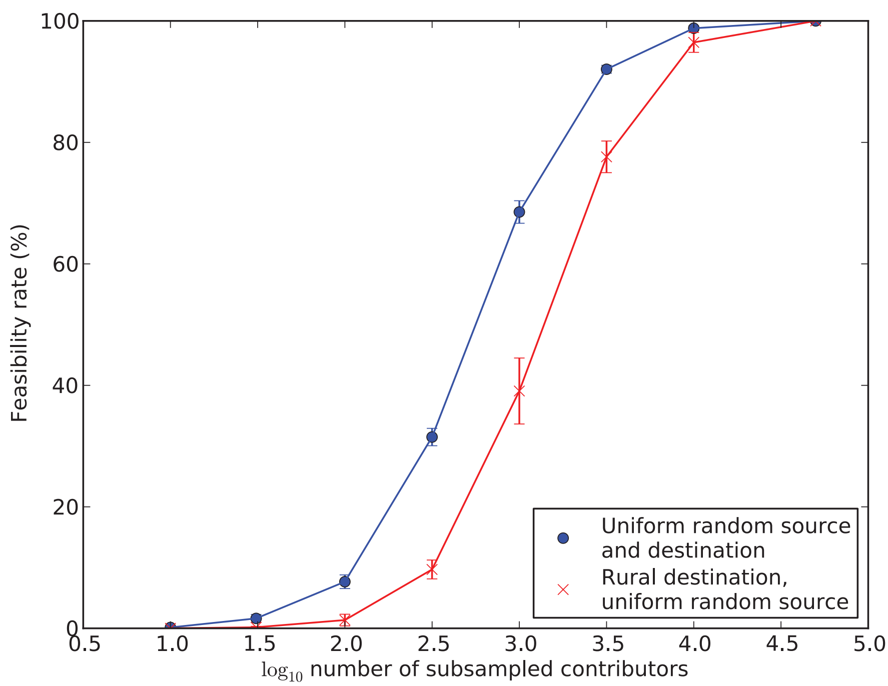

Figure 4. A plot of the percentage of randomly sampled (source, destination) delivery problems that had a solution path of any size, against the $\log_{10}$ size of the number of potential contributors

Formulated in this way, solving the MDP is now a tractable endeavour.

Putting everything together, the whole system for learning and optimisation is laid out in Figure 3. In learning, the mobility patterns of each individual are extracted. From this, it is possible to define the delay probability density function describing the transition probabilities in the MDP. Using the calculated optimal policy of the MDP, the next action (i.e., the participant to ask and the route the package should take) is decided by the package's current position.

So we do not plan out the package's whole route and couriers at the start. Far too many things could go wrong. Instead we wait to see where and when it ends up after the first stage of its journey before deciding who the next courier should be and where he or she should take it. When the second courier has deposited the parcel, we consider what the next destination should be, and who the third courier should be, and so on. That method is far less prone to breakdown.

Simulated deliveries

Before rushing out to deploy our system in the real world it makes sense to ask a few questions. We want to check if anything about the mobility of the participants rules out the feasibility of crowdsourcing package delivery. To do this, we used data about the locations of real people (using the same Orange data set as before) but simulated thousands of delivery problems to be solved with our framework. Our evaluation comprised four key criteria: (1) the number of participants required for acceptable geographical coverage; (2) the number of participants required in any specific delivery (since longer chains imply greater risk of loss and theft); (3) the feasibility of delivering to rural locations, which is expected to be much harder than urban delivery; and (4) the time required for each delivery.

Criterion 1: Number of participants required

Figure 4 shows the percentage of location pairs that were feasible (i.e., that had any path between the source and destination locations). The blue line shows the feasibility for uniform random source and destination locations. We see that the geographical coverage is very poor when there are fewer than $10^{2.5}$ participants. In other words, 300 signed-up couriers will be inadequate. The critical range is around 10^3 – a thousand couriers – when feasibility surges with each new participant. The heavy tail in human location behaviour is one explanation for this effect, where individuals visit many locations a few times (and a few locations many times) in their daily life mobility[4]. Therefore, an acceptable geographic coverage, trading off against recruitment/administration costs, appears to be around $10^{3.5}$ participants. That is, with 3000 participants we can get packages to around 80% of the country.

Criterion 2: Number of participants required for any given delivery problem

Figure 5 shows the number of participants required for the simulated delivery problems we considered. Since unfeasible paths cannot be included when plotting (because they have unspecified numbers of contributors), the number of contributors required for specific paths initially increases with the size of the participant subset, as more paths are made feasible. However, once path feasibility (indicated in Figure 4) goes beyond 20%, the trend is as expected; having a wider pool of participants allows more efficient (i.e., shorter length) paths to be discovered. Combining this with Figure 4, we find that four couriers are enough to get packages to any of our feasible destinations.

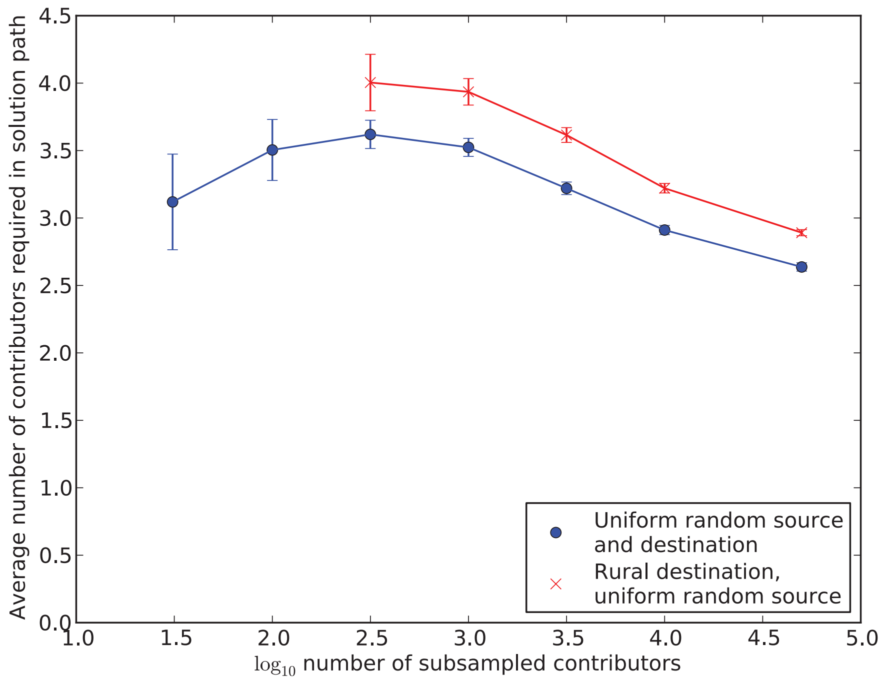

Figure 5. A plot of the average number of contributors required for each specific delivery problem (drawn from the much larger pool of potential contributors) against the $\log_{10}$ size of the potential contributors pool. Note that a majority of rural destinations are infeasible for pool sizes of less than $10_{2.5}$, therefore we are unable to plot the line below this range

Criterion 3: Rural distribution

So far, we have only considered uniformly sampled source and destination test points, which favours urban locations (since there are greater numbers of cell towers in urban areas). We now consider a criterion for rural feasibility, by sampling a set of delivery problems where the destinations are only rural (keeping source locations uniformly sampled, as before). We ran the same analyses as for criteria 1 and 2 with rural destinations, yielding the red lines in Figures 4 and 5. This indicates that restricting the destinations to be rural certainly makes the delivery problem more challenging, but it is still feasible.

Criterion 4: Time required

Now that we know that all three feasibility criteria are met, we consider the problem of learning how to plan the minimisation of delay in delivering the packets.

To estimate the time it would take to make deliveries, we considered the worst-case scenario. If the delays for this hardest scenario are acceptable, then that is encouraging news. Firstly, we considered deliveries to only rural destinations, as delivery is much less of a problem in urban areas. Secondly, given the fact that any recruitment campaign would cost time and money, we considered a small participant pool of 3500 (the smallest number of potential participants we could use while still keeping 80% of feasible routes, according to Figure 4). Under these conditions, the average time required for delivery was 30.0 days, which is 81.3% faster than using the naïve method of finding the route with the least number of contributors (which took 161 days, on average). So we get a big improvement from learning and optimising, though you would not want to send anything urgently in this manner (at least to rural destinations with a small number of participants). So our method is slow – but, on the other hand, it is cheap. Given the low cost of using mobility opportunistically, perhaps a new model of delivery may emerge in which items can be sent continuously, much like the network of blood vessels in the body, as opposed to sending a bulk delivery of items using conventional truck delivery. Mosquito nets could be dispatched from the capital every day; and only halfway along their journey would a final decision be made about which rural village was to receive them.

But can it really work?

What other factors could stop our solution working in practice? To perform routeing under uncertainty, we assumed that the participants would follow their normal mobility patterns when delivering packages, even if these patterns are noisy. Clearly, additional factors could introduce further delay, including disruptions to transport (vehicles can break down, roads can be washed away) and short-term disruptions arising from participants' circumstances (they might be too busy, they might take sick leave). In practical terms, most of the impact of these disruptions could be absorbed by an appropriate task assignment procedure. Specifically, after obtaining a policy from our learning and optimisation approach, the system could ask the selected participants, via automated phone text, whether they are actually willing and able to do the task. We might also offer them incentives (e.g., phone credit that can be simply and immediately added to their account, or free delivery credits for them to use in future) to encourage them to help, provided the package arrives at its final destination. The size of such incentives could vary depending on whether the delivery is for development or for commerce. In general, better pre-task communication would allow participants facing disruptions to be filtered out, limiting the introduction of unexpected delay into the route. On the other hand, some disruptions may not be known at the time of task acceptance, or some participants may simply not be honest about them. Investigating how to update an existing optimal policy with new predictions, and how people respond to incentives in this scenario, are therefore important questions for future research.

Finally, in the worst case (from a routeing perspective), participants may lose or steal packages. A certain amount of loss and theft is assumed even with standard delivery, and is borne as the risk of doing business, or addressed with insurance. In the crowdsourcing setting, this can be taken into account by assigning a cost to each participant (either with a fixed value, or derived from a participant-specific trust evaluation framework. Couriers who have performed well for long periods would be trusted more than untried newcomers). In whatever way the cost of trust is calculated, once obtained, it can be incorporated into the Markov decision process as an added cost in the standard way. Interestingly, keeping unreliable people in the system might be helpful on occasion, specifically, for cases when the high-risk person offers a much faster route to the destination than would otherwise be possible (as long as the package is not irreplaceable, and the sender understands the risks).

It remains to be confirmed how crucial these issues would prove in practice. If they are not crucial, or if they can be mitigated along the lines we suggest, then our system has promise in being cheaper and greener than the conventional alternatives of package delivery. Furthermore, in altruistic applications such as development of poorer countries, getting more citizens involved has the potential to create a more cooperative and inclusive society. This scenario is only one idea in the much larger ORCHID project (`http://www.orchid.ac.uk`), which aims to establish the new science of human–agent interaction. As people and agents form cooperative groups, or collectives, it is important to address challenges such as such as processing uncertain human behaviour, designing incentives, and dealing with the trustworthiness of participants. We only addressed the first; but the benefit to society will be the ability to get more value from the contributions of humans and agents in domains as diverse as citizen science, disaster response, and energy management systems.

On your way, could you carry a small parcel for us? iStockphoto/Thinkstock

References

1. Central Intelligence Agency (2008) *The World Factbook*. Washington, DC: CIA.

2. McInerney, J., Rogers, A. and Jennings, N. R. (2013) Learning periodic human behaviour models from sparse data for crowdsourcing aid delivery in developing countries. In Conference on Uncertainty in Artificial Intelligence (UAI), 2013.

3. Box, G. E. P. and Draper, N. R. (1987) *Empirical Model-Building and Response Surfaces*. New York: John Wiley & Sons.

4. Gonzalez, M., Hidalgo, C. and Barabasi, A. (2008) Understanding individual human mobility patterns. *Nature*, **453**(7196), 779–782.

James McInerney is a PhD candidate researching human location behaviour prediction at the University of Southampton, where Alex Rogers and Nicholas R. Jennings are Professors of Computer Science. Nicholas R. Jennings is a Chief Scientific Advisor to the UK Government and an internationally-recognised authority in the areas of agent-based computing and intelligent systems.