

D2.1: User Challenge 1
Novel information technology for the early detection
of infectious disease events.

Grand Challenge: the use of modern information and communication technology systems to gather and interpret timely and relevant data and deliver it to those managing an outbreak.

User Challenge 1 - Champion

Professor Roy Anderson, Chief Scientific Adviser, Ministry of Defence

User Challenge 1 - Working Group Members

Professor Will Stewart (technical lead)	Independent consultant
Professor Tony Barnett	London School of Economics
Professor Chris Bartlett	University College London
Professor Christl Donnelly	Imperial College London
Professor Dave de Roure	University of Southampton
Dr Colin Upstill	University of Southampton IT Innovation Centre
Dr Penny Wilson	Independent consultant
Professor Mark Woolhouse (Foresight Science Co-ordination Group)	University of Edinburgh

Introduction/context

Major disease epidemics represent a serious threat to humans, animals and plants, and the need for accurate, timely and relevant information about them is vital. Thus the huge advances in IT and communications over the last 20 years, and the even larger changes currently in progress, have the potential to greatly enhance quality of life. Indeed IT and communications could have huge economic and public health benefits throughout the developed and less-developed world but the gap between what might be done and what has been done is still large and a good deal of effort may be needed to bridge it.

To summarise some issues and trends discussed in more detail later:

The need for information

This need is widely recognised (see for example Butler, Nature 440(2006), 6-7), and illustrations are easy to find:

9/11

The problems created by inadequate (even just delayed) data provision to emergency workers were well illustrated in the aftermath of the September 11 attack on the World Trade Center. In the first four days of the rescue efforts, no coordinated mapping information was available for emergency workers who were forced to draw their own maps or use shopping guides to identify hazards.

2001 foot-and-mouth disease epidemic

One of the key findings of the Royal Society's report on infectious diseases of livestock (Royal Society, 2002) was that: 'Quantitative modelling is one of the essential tools both for developing strategies in preparation for an outbreak and evaluating the effectiveness of control policies during an outbreak.'

Data gathering and distribution infrastructure trends [not disease specific]

Both in the developed and in the less-developed world this capability is improving rapidly. Consider for example the mobile phone. More than 82% of UK adults now possess one (and >90% of children over 10), and mobile phone ownership is rapidly spreading in Africa. By its nature, a mobile handset's approximate location is known to the system, and emerging handsets include GPS to make very precise location possible. Thus population movements can be tracked in real time with obvious potential benefits in tracking a current or potential outbreak, and vulnerable individuals could also be directly contacted.

The cost of this technology, as with most ICT, continues to fall steeply - and this does help the less-developed as well as the developed world - note for example Motorola's project to produce a \$30 handset, or TTP (Cambridge)'s recently-announced intention to license a \$20 phone (FT February 7 2006). Both efforts are specifically targeted at users in less-developed regions.

Sources of information


Disease reporting will be improving fast, see for example UC3 and UC4, but the 9/11 example and the mobile phone data neatly illustrate that helpful information does not come only from traditional or specifically disease-related sources. Other recent examples have included the tracking of dollar bills (by serial number through online bill-tracking websites) as a way of monitoring population movements (Brockman et al Nature 439, 462-465 -26 January 2006), and the use of weather data to inform prediction models of malaria outbreaks (Thomson et al, Nature p 576-9 Vol 439 2 February 2006) because of the weather's impact on the breeding mosquito.

The potential use of information

This is best seen from examples - appended to this report are brief summaries of the use and significance of key information. The examples appended here are:

Challenges of influenza pandemic surveillance (Neil Ferguson and Christl Donnelly)

Expected pattern of spread of an uncontrolled epidemic. Annotation is time sequence in days.





The UK epidemic of foot-and-mouth disease in 2001 (Mark Woolhouse)

Geographic distribution of all 2026 farms reported to be infected with foot-and-mouth disease (FMD) in Britain during the 2001 epidemic

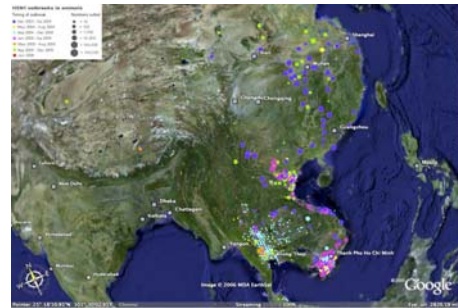


The UK epidemic of BSE (Christl Donnelly)



Data challenges

Among other things, these examples illustrate the need to blend data from different sources, often with maps and geographical information. This is becoming easier - see for example the avian flu data superposed ('mashed up') with maps from Google Earth in Nature, Vol 439 5 January 2006 pp 6-7. However, the problems of sorting data and processing it remain considerable and much work is required.



Document Summary - Key Issues

The Group believes that the benefits of meeting this challenge are huge, but that clearing the biggest hurdles requires attention to both technical and social concerns.

- The actual collection processes may be driven by other developments and unrelated priorities but the problem of extracting and identifying useful information from huge complex systems including both people and machines will be considerable. Ideally these processes would involve uniform international standards but a more practical scenario assumes specific local (probably human-aided) automatic interpretation.
- Various sources, including remote monitoring, weather and mobile phone tracking data can be assembled into very large databases containing some relevant and much irrelevant information. This needs to be filtered and combined with epidemiological models to give timely, accurate response options. Our discussions emphasised the heterogeneous nature of the various database structures and many different representations of data. We also expect data to vary in its reliability and that its provenance may be uncertain. We noted the GEIS news crawler, as an example of effective early warning on doubtful data.
- There are significant public acceptability and privacy issues that may arise.

In parallel with, and intimately connected with technology developments, is the social problem of obtaining agreement for the widespread international management of information for the public good. This involves cultural and other concerns, some of which (for example privacy) are not, yet, high profile but which may become increasingly so in the future. This challenge may provide a useful first example within which the ethical issues can be explored. In any case, the time taken to address the social issues may exceed that needed for the technology to develop - and the need to address them may therefore be urgent.

This document is structured in a standardised manner for clarity. Readers wanting a summary should initially read sections 7 (suggested actions) and 8 (roadmap), each agreed in open meeting (and if possible the examples in the Appendix).

- Section 1 - Lists some of the changes expected in IT that will be significant.
- Section 2 - Discusses the issues and priorities for interfacing information with epidemic management, including modelling issues.
- Section 3 - Explores the complex 'system' (including people) issues involved and the linked issues of who gets any benefits.
- Section 4 - Discusses the financial and international costs and benefits, and impediments.
- Section 5 - Specifically lists likely impediments to deployment.
- Section 6 - Discusses social equity issues.
- Section 7 - Lists the key immediate actions that the group recommends be taken.
- Section 8 - Shows a road map for developments.
- Appendix - describes the issues involved in two specific examples - influenza and foot and mouth disease.

1 Key future capabilities for user challenge systems

The key issue here is the improved acquisition, organisation and management of information. This is very much the core of the ICT (information and communications) revolution, and is an intrinsically global process. ICT available in Africa may be less extensive but is of the same generation as ICT in the developed world – as in the case of relative mobile phone penetration. The changes expected are so large and extensive that they are difficult to summarise, but relevant developments are:

Data organisation and structuring

Data needs to be labelled and classified data to facilitate efficient access. This needs to be at least partly automatic for practical reasons, but the result must be data structured by relevance rather than just by text content.

Semantics Semantics gives meaning to data. The envisioned Semantic Web will give meaning in a manner understandable by machines, to the content of documents on the web. This will involve machine-readable metadata and ontologies.

Semantic interoperability is key to this (agreement on meaning, to achieve mutual understanding), as is the bridging of semantic gaps (matching understanding across domains). Consider, for example, the gap between the low-level information that can be extracted from colour-based segmentation of satellite images and the application-domain semantics of failing crops or animal migration.

Unstructured data Analysis of web pages, blogs, wikis, news feeds, etc may be at least as important as the analysis of structured data from sensors, medical records, satellite and CCTV images. There is a wealth of valuable

information to be gleaned from analysing (possibly a mix of structured and) unstructured multimedia data. Mark-up languages such as XML enable self-describing, self-indexing data. Unstructured text and natural speech pose particular challenges but provide great opportunities.

Data communications and organisation

Data sources will become (arguably have become) pervasive and control and organisation are widely dispersed.

Collaborative means for personal communication Disruptive technologies such as blogs and wikis spread rapidly. Using these new technologies, people share a wide range of personal information including details on health and medical concerns. Such publicly accessible information is becoming universally available. The dynamic social networks by which it spreads can be analysed and may reveal the emergence of new diseases much more rapidly than any other method.

Pervasive data sources Distributed computing devices in the physical world, such as personal devices, wearable computers, devices embedded in everyday objects and sensors in the environment will generate ever-increasing volumes of data. The autonomous, reconfigurable and/or dynamic nature of these sources pose significant problems regarding security, trust and provenance.

Challenges from deployment

These developments will pose many challenges, not all technological. Many of these will directly, but not uniquely, apply to UC1-addressed systems.

Privacy, security, trust, provenance We note the paramount importance of ethics, sovereignty, confidentiality (commercial and otherwise), privacy, data protection, freedom of information, and other rights and obligations when detecting, collecting and processing data. These are underpinned by security (an attribute of systems and procedures that minimises or manages the risk of undesirable system behaviour, even in the presence of malicious, untrusted parties), trust (confident reliance regarding behaviour) and provenance (the credibility and reliability of data).

Distributed and unstructured data increase the challenges. What matters is who can access what and what they can do with it in some defined context, not who they are or what their role is or where they are or where the data is. We need dynamic semantic security and context-based access control.

Methodology and process How do we deal with large-scale distributed systems (of data sources and computing devices)? Established processes are based on single repositories. This is a straitjacket. Process always lags technological capability.

Business understanding is as important for managing disease data as for any other data analysis. Garbage in, garbage out.

2 The role of the systems in managing future risks

Introduction

The problems created by inadequate (even just delayed) data provision to emergency workers were well illustrated in the aftermath of the September 11 attack on the World Trade Center. In the first four days of the rescue efforts, no coordinated mapping information was available for emergency workers who were forced to draw their own maps or use shopping guides to identify hazards. Once GIS equipment and trained operators were on the ground (this having been delayed by the grounding of flights in the days after the attack), the management teams provided detailed maps of the locations of command centres, first aid centres, food sources and even fires (including those detected by heat sensing cameras).

Of course, future risks could be much more widespread than the deadly, but geographically localised, impact of the September 11 attacks. Thus, the needs for data (and data processing) could be much greater.

Data coordination

Volumes of data in and of themselves are not directly useful in managing future risk. To be most helpful, data from various sources needs to be documented, validated, meaningfully linked and filtered.

Database documentation is critical to the appropriate use and interpretation of any dataset. The documentation should include a description of the variables included and possible relational links between these variables as well as a description of the data collection process. Ideally the documentation should also include contact details for someone able to answer questions about the database. It is critical that all users consult the relevant documentation for each database.

Data validation is a systematic process for reviewing a body of data against a set of criteria to provide assurance that the data are adequate for an intended use. Thus the extent of data validation may vary depending on the demands to be placed on the data. Data validation can include: format consolidation; assessment of the level of missing or invalid responses; consistency checks (both internally in a single dataset and between datasets); removal of obsolete data; auditing; and 'deduplication'. (The process of deduplication involves the identification of sets of records that refer to the same entity, despite possible inconsistencies in the data within the different records.)

The coordination and linking of data from different sources can be hastened or hindered by the interoperability (the ability of a system/product to work with other systems/products without special efforts) of each body of data. Even when the software and hardware are interoperable, specific manipulations of datasets may be required to link data from different sources optimally. For example, before two geographically referenced datasets can be integrated, they may need to undergo projection conversion. A projection creates a two-dimensional representation of locations from the earth's three-dimensional surface. Sometimes a particular projection is more appropriate for certain

uses. For example, a projection that accurately represents the shapes of the continents will distort their relative sizes.

The solution to inadequate data (limited data availability and/or the poor quality of available data) could involve greater investment in data collection systems, standardisation of data collection methods, and/or further modelling and analysis of the available data to impute data values as need.

At the other end of the spectrum is the issue of data overload. The available data may be so extensive that the challenge of data filtering is formidable. Data filtering can be critical to allowing a user to find the needed piece of information within a very large body of data whilst potentially drastically increasing the speed at which the user processes the data. Data mining and knowledge discovery in databases is a rapidly growing area of research, bringing together techniques in pattern recognition, data visualisation, machine learning, and artificial intelligence.

Timeliness

The timeliness of information being provided to users and risk managers depends on multiple factors. These include:

- the frequency at which the raw data is collected (e.g. continuous, hourly, weekly, annually or one time only);
- the time taken for data validation and, if necessary, data filtering, modelling and/or analysis;
- the availability of computers or other equipment necessary for the users to receive the data.

Clearly direct real-time user access to databases as they are compiled gives users access to the most up-to-date information. However, if data validation has not taken place, the most up-to-date information may be less useful than the most up-to-date validated information. Furthermore, users may be inadequately prepared to process the raw data in order to inform their decision-making.

Clearly, the decision on how frequently particular data should be collected will depend on the specifics of the system under study. Consideration should be given to the stability of the system (how quickly do changes happen), any limitation on the technology of data collection, the speed of data validation/processing and the need for general trends or specific detail.

The availability of suitable equipment which would receive and interpret the information depends on the level of preparedness within a particular group. Clearly the need for specialist equipment to receive information will reduce the number of users or risk managers who can receive information directly. Thus, as technologies develop, the process of information provision should be kept under review. For example, if data can be transmitted to a common hand-held device, this may be the best medium for communication, though of course consideration should be given to a back-up plan in the event that mobile phone systems were unavailable or overwhelmed by demand.

Links to international processes

In theory, the technological issues involved in coordinating data internationally are generally the same as those that arise within a particular country. An obvious complication arises from differences in language and alphabets. Similarly, attention should be paid to the possibility of different datasets using different units for the same variable (for example, pounds and kilograms). Of course anyone linking datasets should always be alert to this possibility even within a single region.

Data availability may be limited, or access may be more difficult, in a particular setting. Clearly the availability of paper, rather than electronic, records will be particularly challenging. However, ongoing improvements in scanning and character recognition technology will improve the accessibility of such data.

3 The costs and benefits of selected future capabilities and their robustness to future uncertainty

In what follows, the emphasis throughout is on the system.

A technical capability such as a networked system for capture, analysis and modelling of data for early detection of infectious disease events does not stand outside of its social, economic and political environments. The potentially available hard and soft technologies are likely to include:

- Acquisition of individual data from devices akin to mobile telephones
- Logistics tracking
- Integrated models
- Knowledge construction
- Response to users
- Global systems

This is not simply a question of technology. 'Architecture' is fundamental and assessment of potential costs and benefits can only be considered in relation to this. The term 'architecture' describes the organisational, administrative and budgetary structures which enable the capability to come into and remain in existence - in other words to be 'sustainable'. Such arrangements must also take into account issues regarding the robustness of the technology, its security, processes (and their regulation) through which medication becomes available and the development of new information sources and feeds.

We can approach these issues by addressing the following questions:

- 1 What architectures are possible?
- 2 Who are the risk managers?
- 3 What is the public good and what is a public good?
- 4 How can we think about the costs and benefits of such systems?
- 5 What are the equity issues?

1 Possible architectures

Most plant, animal and human health related services have until now been delivered via hierarchically organised systems. This organisational model has its roots in the inflexible and hard-wired social technologies and machines of the 19th and early 20th centuries. Such social technologies have typically been deployed via states or large commercial organisations - the latter sometimes referred to as 'Fordism' after the centralised and cumbersome forms of industrial organisation developed by Henry Ford, among others, in the early twentieth century. The essence of newly developing and foreseeable information technologies is their current and potential disseminated distribution and use. This fits very well with a number of trends in social technology, in particular in relation to the role of the state. We can identify the following broad possibilities or scenarios:

- (i) A liberal-democratic model under which state power will become less easy to exercise in the face of demands for citizen 'choice' and regional and/or local communal control. There would be increased demands that allocation problems be determined via market mechanisms.
- (ii) A liberal-democratic model where state power increases in the face of more disruptive and non-specific and often diffuse threats. In addition, control of technologies, information and resources becomes more centralised despite countervailing pressures against this tendency, these trends resulting in tendencies to disorder and opposition.
- (iii) A Chinese model (which may be restricted to that country alone but is also to a degree apparent in some outliers such as Singapore) where centralised political control remains with the state while allocation problems are resolved via markets which remain subject to a high degree of political regulation. 'Civil society' organisations are subject to firm state regulation or are few in number.
- (iv) A failing state model - probably true of some areas of Africa but not exclusively so - where state structures are ineffective at the most minimal level of security. Most initiatives are very local and technologically limited and variable with an uneven and unpredictable topography of power over some regions of the world.

The key questions arising from these scenarios is what conceivable local and international arrangements could be made through which the available technologies might be managed and made effective for use by 'risk managers'.

Risk Managers

Part of the problem lies in the two words 'risk' and 'managers'. Under each of these scenarios, risks will vary because of the architecture. In addition, the potential for risk management will also be differentially located. With regard to risk, the situation can be summed up as follows:

Scenario (a): Widespread and variously articulated demands for 'individual choice', wide distribution of technology via provision through a global market,

provision of diagnostic and treatment services through a market. Result is difficulty of knowing risk because individual knowledge of health state may be available through communications tracking but emphasis on 'individual choice' could stand in the way of such use at an aggregate level because of legislation and power of civil lobbying groups. Thus, it might be possible for individual health status to be known by individuals but for both individual and group risks to be unknown and subject to wide margins of error and misinterpretation and/or even misuse. Combined with proliferation of information sources, the resulting cacophony of information might mean that risk management by state authorities could become very difficult indeed and impossible for individuals. Given also the specific pathogen + epidemic characteristics, there could be dangerous delays at each stage of the detection and identification chain meaning that any response would be held up.

Scenario (b): This scenario already assumes elevated levels of fear of threats with or without enemies and a base assumption must be a political topography characterised by uneven support for central state power in the face of these threats. There may even be opposition to increased state power. It is in this situation that the best collection and collation of information is likely to be possible and also most likely that the risk managers would be state personnel. But it is also possible that opposition to state access to data (as for example shown with anti-measles vaccination and more general anti-vaccination campaigns) could result in critical levels of politicisation of the issues over a very short time period in less than optimal circumstances for effective public explanation and action.

Scenario (c): In this scenario the party apparatus remains potent but there is a huge range of consumer goods available. Declining levels of trust between government and people resulting from increased inequality, as well as perceived irrelevance of government to much of daily life, could mean that in a critical disease situation the technologies could be deployed relatively effectively but large sections of the population would either not have purchased the technology or remain invisible to it for other reasons. This could mean that a very partial picture emerged over a large area of the globe. In this situation, risk management could be carried out effectively but might be less effective.

Scenario (d): The failing state model is instructive because with no effective state power, individuals and communities make myriad choices. The result is a range of possibilities as between some communities and regions with good data and risk management in place and others where this is impossible. We could consider the different situations as between Somalia and Somaliland and within those two countries as possible models for this unevenness.

But the real issue is of course the global nature of the process for data collection, knowledge production and risk management. And it is the potential problems of interaction and relationships between and not only within these scenario types that pose the key architectural problem.

On immediate inspection it would seem likely that:

- (i) National architecture will be extremely variable - even in settled, prosperous countries with a tradition of 'liberal democracy'.
- (ii) The technical processes by means of which budgets and technologies are deployed, information collected and collated and risk managed will only work in relation to specific local circumstances.
- (iii) International cooperation beyond the local and national will require careful preparation over a very long term.
- (iv) Such preparation will only be effective if populations understand the collection, collation and use of the information as both a public good and in the public good (these terms are discussed in the next section). This is a challenge for politicians and for their education and raises questions about whether such consensus can be achieved internationally in the near future.

There are certainly circumstances where it will be more or less difficult to establish a centralised foundation for these tasks, where appreciation of risk could be very difficult because of the foundations, and where risk perception and management might be very hard to sustain or realise. In regions with a more centralised political tradition and growing inequality, market solutions to allocation could well produce another kind of unevenness along lines of income, wealth and spatial inequality.

The overall conclusion is that in most foreseeable circumstances, the technology which is likely to be available is also likely to be very hard to use for purposes other than individual information. Under many foreseeable political arrangements it will be very hard to use the technology and the information for collective purposes without some legal provision at the national level. At the regional and global level it may be even more difficult. Devolution of control over technology, political mood and normal administrative difficulties combine to severely limit the deployment of technology potential for collective purposes.

If this is the case, then there is much to be done now by way of education and preparation of populations and in the building of diplomatic linkages.

Public good and public goods

The term 'public good' has to be defined in more detail in order to overcome the basic problem resulting from the misunderstanding of the notion of a public good.

In everyday language we speak in quite vague terms of *the* public good. Economists speak of a public good which has a rather different meaning. The more general term *the* public good refers to the motivating or legitimating reason or goal behind a course of action taken on behalf of 'the public' - or a course which may be claimed to be such.

In contrast, the economist's use of the term a public good means something quite specific: a good that has the following characteristics:

- (i) It is non-rivalrous - consumption of its benefits fails to exhibit consumption scarcity; once it has been produced, everyone can benefit from it without diminishing others' enjoyment.
- (ii) It is non-excludable - once it has been created, it is very difficult to impossible to prevent access to the good.
- (iii) It is non-divisible - its benefits cannot be divided.

A traffic light is an example of such a public good. Another (albeit slightly less clear) example is a surveillance system for detection, identification and response to infectious diseases of humans, animals and plants.

The idea of a public good is intuitively simple but politically debated. Whether or not any provision can be or ought to be seen as a public good or whether it should or can be provided via a market mechanism is the stuff of political debate. Availability of individually operable devices may privatise information and even response via market availability of medications, and some of this outside of national controls via, for example, mechanisms such as the internet.

Here lies the problem: how to increase the potential for individualisation of information in the face of a growing requirement that such information be aggregated for the public good so as to constitute a system of information to form a public good.

Costs and benefits

Put very simply, the benefits of these technologies may be thought of in relation to the following dimensions:

- 1 Who are the beneficiaries?
- 2 Over what time period do we consider benefits to accrue?

Who are the beneficiaries?

It is usual in health economic thinking about the first type of question to consider only returns to individuals. If this were the case, we would move to the conclusion that the benefit being individual, the cost should be borne by the individual, either through a direct market transaction or via an insurance-based mechanism. However, it is more realistic to recognise that there is a circle of beneficiaries to most individual health outcomes. The prime beneficiary may be the individual, but there are also benefits to their children, for example, where the avoidance of morbidity and/or mortality protects streams of benefits to dependents of various kinds. But the picture is more complicated as for example avoiding creating orphans benefits the wider society and saves costs which might otherwise accrue. The picture is further complicated by the related issue, that costs (and therefore benefits - which are cost avoided) do not only consist of money-quantifiable costs.

There are also less tangible but very important costs and benefits around processes of unpaid (and often non-quantifiable) work which contributes to social reproduction. Domestic labour is one example of this in rich societies where much socially productive work is performed via state or state-facilitated provision (for example the NHS, pension schemes, unemployment benefits,

child benefit, child care). In less-developed societies or less-developed parts of rich societies where people cannot command such services and goods via the market, real unpaid domestic labour is of great importance. Recent work by Nancy Folbre (Nancy Folbre, *The Invisible Heart: economics and family values*, New York, The New Press, 2001, plus extensive work on 'unpaid labour' for UN agencies) and others suggests that the addition of this work to standard systems of national accounts increases overall GDP by around 17 per cent.

The point of this discussion is that in economic terms beneficiaries can be any of these circles, and who is considered a beneficiary depends on political decisions and mood. If we work within a very constrained perspective, then the individual is the beneficiary and should bear the cost. If we work in the wider perspective, then the entire society benefits and costs should be allocated in such a way as to recognise this. This is a continuum. At one end lies individual market transactions, at the other lies collective provision with cost defrayed in a manner which recognises collective benefits. It is at this point that we return to the issue of *the* public good and a public good. If the devices and their attendant systems are considered to be in the public good, because they protect the ability of a society to cohere and continue in existence while protecting each individual by means of collective action (an assumption which presumably underlies much public health theory from early times), then we confront the question of how to pay for the provision and the relation between that and the precise nature of the provision.

If we ask people to pay individually, then they will only see the individual good side, rather like the identity card issue today. If we say that devices and systems are to be provided or subsidised by the state then they may not see their value at all because it is costless (except through the tax system) or they may be persuaded that there is a public good aspect to the system of which devices form a part and that provision of the system and devices has a public good aspect. That is, the system has that aspect, not the individual devices themselves.

Over what time period do we consider benefits to accrue?

With regard to the second type of question, returns to an investment in business and engineering are not typically considered over a very long term. A typical assumption on cost benefit analysis is an internal rate of return in a period of around 20 years.

However, the benefits of a system associated with any public health issue should probably be considered in the much longer term. What this term might be is hard to say, but a reasonable approach might look to the three-generation unit of many human societies - grandparents, parents and children - plus one generation, as the time span. Assuming each reproductive generation is around 30 years, then the stream of benefits to investment might be considered over 120 years.

4 - The factors influencing the development of the future capabilities

This paper identifies enablers for and barriers to developing analytical and strategic systems to provide timely information to users and risk managers through the use of ubiquitous, continuous and rapid automated capture and transmission of heterogeneous data to networked data repositories.

The following issues are considered:

- i) Who pays?
- ii) International co-operation
- iii) Impediments to deployment
- iv) Commercial motives
- v) Demands of individuals versus public good.

Who pays?

Demand for data pertinent to public health, animal health and plant health issues is likely to come from a variety of sources. Note that the kind of data at issue is not simply data on cases of disease or infection but also underlying demographics (including movements) and distributions of risk factors:

- i) Commercial interests of industries directly affected by a disease outbreak, e.g. in the livestock and agriculture sectors.
- ii) Commercial interests of companies designing, manufacturing and delivering drugs, vaccines and other products related to disease treatment or prevention.
- iii) Retailers, e.g. of non-prescription drugs, drugs for use in animals, or treatments for plants (including consumer markets for products to do with domestic pets or garden plants).
- iv) Government departments, e.g. HPA, DEFRA, MoD and DfiD in the UK, and their equivalents elsewhere.
- v) The EU, possibly through Framework Program 7.
- vi) International agencies, e.g. WHO, OIE, FAO and possibly including UN agencies the UNDP and the World Bank.

International co-operation

Surveillance for infectious diseases is, in many cases, a global concern. Among human diseases this was illustrated by the experience of SARS in 2003 and is manifest now as in fears over a possible influenza pandemic. In both of these cases, countries not immediately affected (including the UK) need to respond to a threat emerging elsewhere in the world. The US is an instructive case study. In recent years the US has suffered a considerable public health and economic burden from diseases arriving from outside its borders, notably HIV/AIDS and West Nile virus. Historically, similar events have stimulated major investments. In the 1940s the CDC was set up to counter malaria, a disease imported in (probably) the previous century.

For animal diseases, the principle of global surveillance is encoded in the OIE list of notifiable livestock diseases, many of them included because of their 'trans-boundary' importance and their possible implications for international

trade. The list includes not only infectious diseases of livestock but also of poultry, rabbits, bees, fish, molluscs and crustaceans.

For plants, international monitoring is more focused on crop pests than infectious diseases, under the auspices of the FAO. There are no international reporting requirements as there are for livestock diseases.

Conversely, some of the world's most important diseases – notably HIV/AIDS, TB and malaria – are now the subject of international initiatives sponsored and funded by national governments, international agencies and charitable foundations. The UN Millennium Goals have the target of reducing the burden of infectious diseases (Goal 6). All of these initiatives generate a requirement for baseline data on disease burdens and monitoring disease trends.

5 The barriers and enablers to implementing the key capabilities

The tracking and early detection of infectious disease events will allow much more effective tracking and control of infectious diseases through populations and across territories. Rapid analysis and sharing of data between public health authorities will permit much better international coordination of public health strategies. However, some ethical challenges do exist:

- This sort of public health data will not normally provide individual level patient information, but may raise patient concerns beyond the point at which public health services are able to manage patient care, or beyond the point at which individuals are able to pay for testing, diagnosis and treatment. This problem may be most acute where no effective treatment exists or is affordable.
- Population screening has the potential to allow the identification of a particular infectious disease with a particular population group, thus creating or reinforcing stigmatisation of that group and resultant social or official discrimination.
- Screening networks will involve sampling, and there may be incentives to avoid or subvert sampling strategies, especially where there is some commercial value in doing so (this may apply in particular to animal, plant or produce screening, but could also apply to human migration).
- Some schemes to make screening attractive may involve trade-offs, similar to those with mobile telephones or supermarket loyalty cards, where an incentive is built in so as to make access by third parties to personal health information attractive to data subjects. This may be combined with compulsory membership, or opt-out membership. Such approaches to screening may lead to an improper or ill-advised weakening of personal privacy protection, with unintended consequences for liberty, privacy or security, on the part of individuals. In part the ethical status of such schemes will depend on the precise balance of concerns with individual privacy and with 'good citizenship' or solidarity expressed in society at any given time.

- An assumption about the value of the technology is that it will facilitate international cooperation. However, in cases where two or more nations are in conflict with each other, militarily or otherwise, such cooperation may not be forthcoming. Indeed, failure to cooperate could be a *causus belli* [motive for conflict] between nations, particularly where the infectious disease is serious and virulent. Further, in many situations of conflict or rapid movement of populations, rapid screening may be a practical impossibility at just the time when it could be most useful.

6 Issues of equity

It is intrinsic to a system of this kind that more benefits accrue to people collectively than to isolated individual users, since the value of any information gathered rises very rapidly as the information becomes more complete. This means that a significant element of central support will be necessary to provide the infrastructure. This does not imply complete state support, just enough to motivate each user. For global long-term issues such as these, this argument applies not only between individuals but also between countries and between generations. In this respect it resembles vaccination.

The question of provision of such a system over a long time period points us to the issue of equity as between generations. This has already been discussed briefly. It also points us to equity considerations between contemporary members of a society. The essence of a system like this is that each individual needs to see good reason to:

- (a) use the technology or participate in its use;
- (b) bear part of the cost to the extent that they feel that their contribution is equal to some perceived utility to themselves and to those who they value;
- (c) recognise that the system provides a public good - in other words, that they can see that provision meets the criteria described above as closely as possible.

It is likely that - as with vaccination - it is only where provision is seen to treat all potential cost bearers and beneficiaries as equal that the system can be sustained. However, with regard to costs and benefits in a political world, equity is not at all a simple arithmetical process of division. To treat unequals equally is as inequitable as to treat equals unequally.

On the basis of this axiom, it is probably necessary to ensure that the public good is best served by ensuring that the DIM system is funded in ways which ensure that all perceive themselves to gain from the provision. In other words, just as in the case of vaccination, the system has to be funded, not all are equally able to fund it, thus unequal contributions are probably the only way of ensuring that a public good is constructed for the public good.

The conclusion is that:

- a) The system within which these technologies are used is likely to be best deployed via collective provision for reasons both of being able to capture the data and use it all and also if people are going to be

prepared to share the individual results and follow a collectively responsible response which protects the public good. In other words the more the system takes on the characteristics of a public good, the more likely it will function for the public good. How this works in relation to the different politico-social architectures outlined above is hard to determine, as also is the political diplomacy of making this happen internationally.

- b) The time period over which the benefits must be seen to accrue assumes inter-generational responsibility. Thus for the costs to be defrayed collectively, the political presentation of the system must exhibit this to potential current beneficiaries.
- c) It is not easy to see how these problems can be solved internationally within existing international architecture such as the UN. Thus there is a diplomatic agenda arising from this discussion and it has to do with global public health systems.

7 Suggested actions

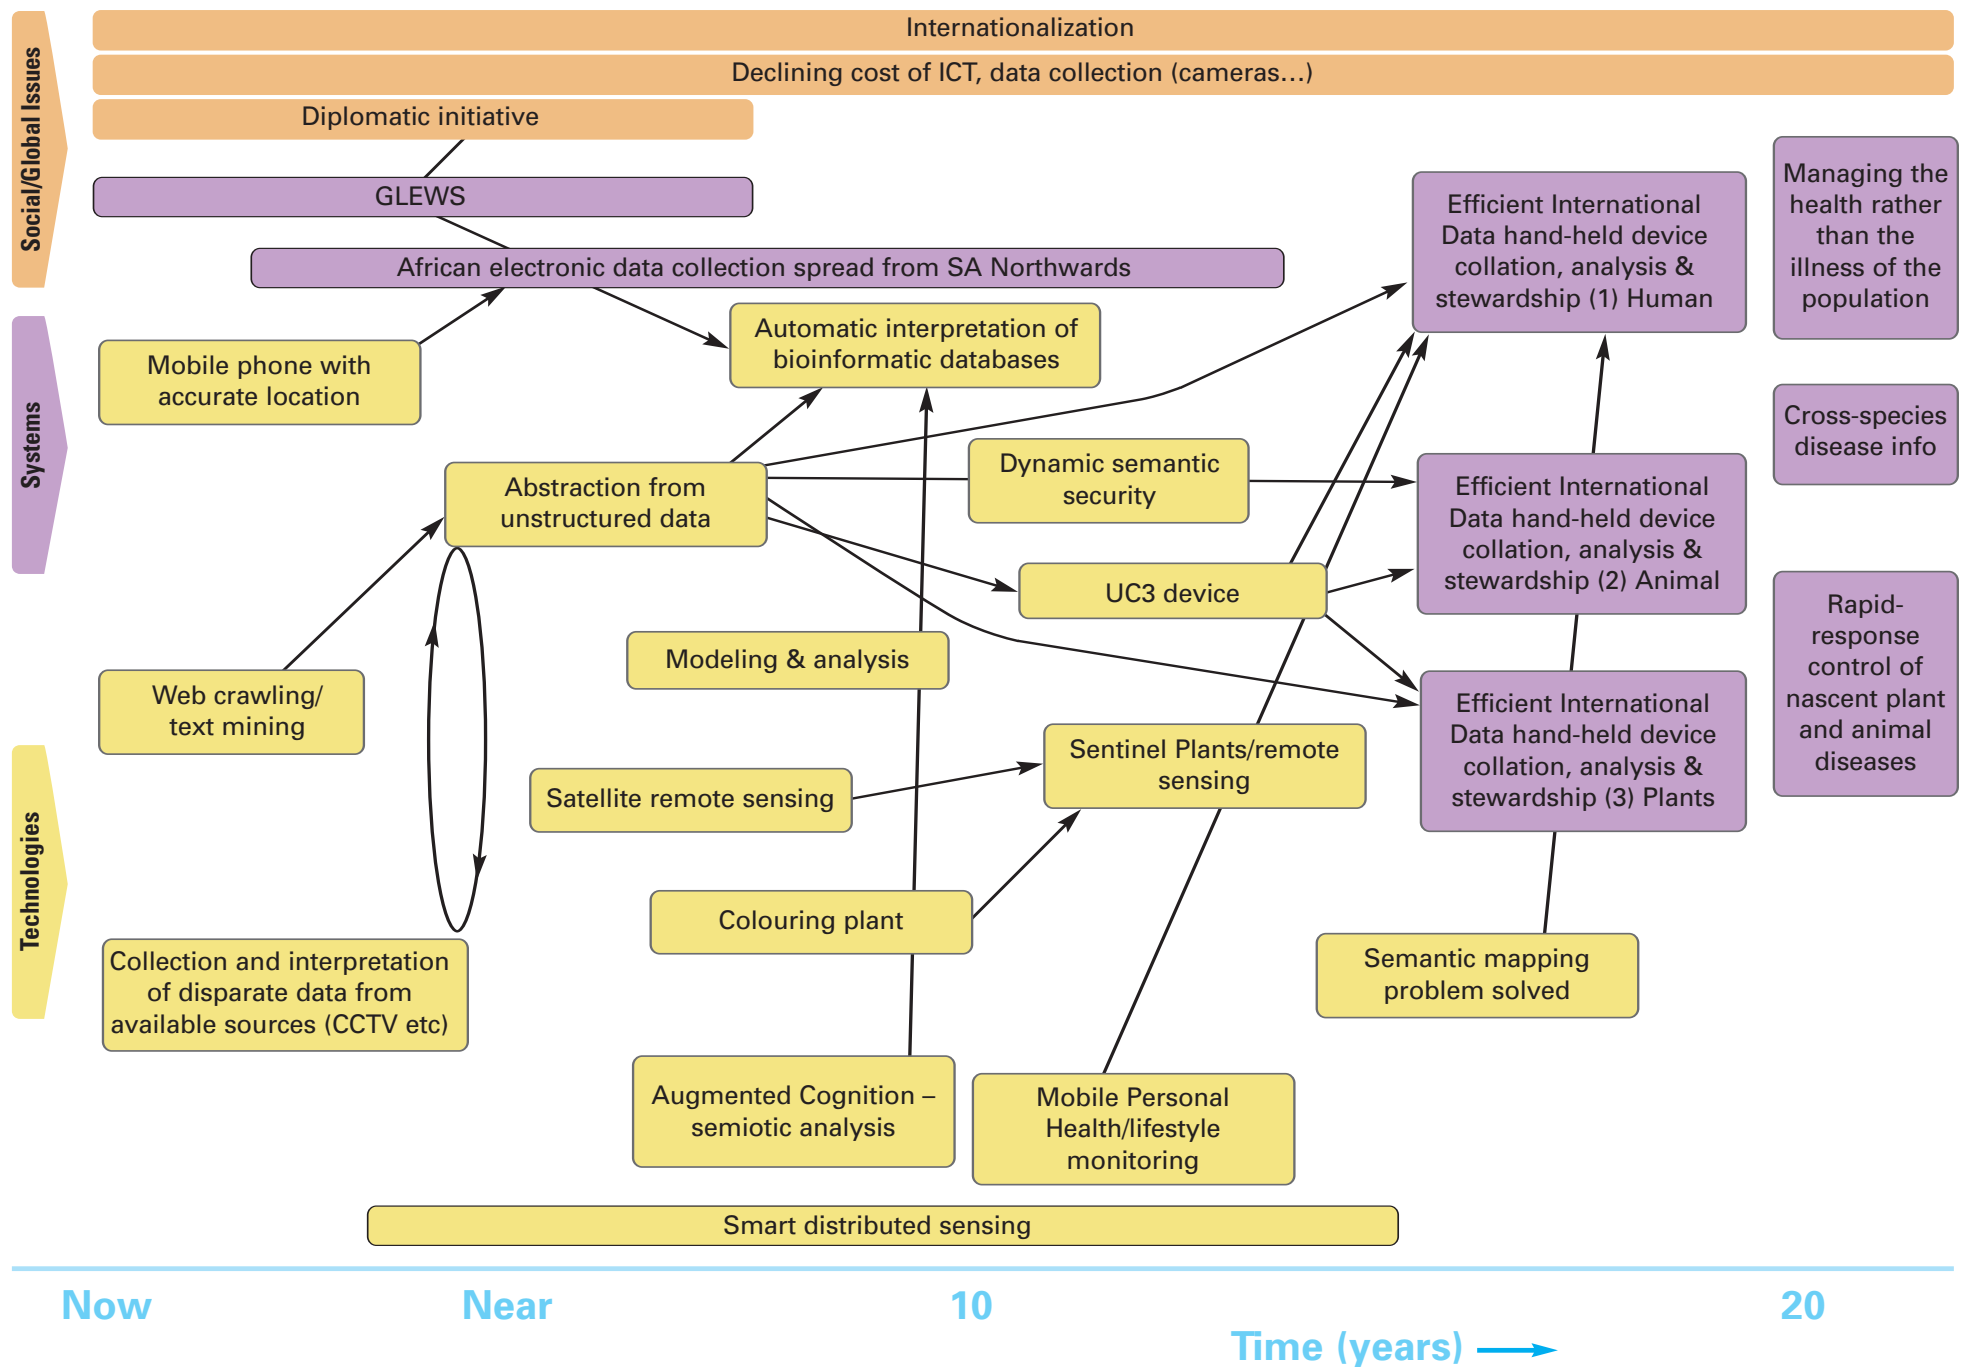
1 UK initially, with international dimension developing:

- a) Assemble key stakeholders and interested parties, including phone operators, academics and medical authorities, to discuss a mechanism for the effective use of personal information for the public and individual good through some trusted mechanism. The problem is to realise the undoubted public good to be obtained, for example in disease control, whilst providing sufficient privacy and personal control over personal data. We recognise that this is a much broader issue than UC1 but suggest that UC1 provides a clear example of the benefit that might be obtained. The mechanism is to be discussed by the group but could involve a trusted gateway that enforced information exchange rules based on the context. Sources for relevant information would be very diverse, from smart cards through loyalty card data to phones.
 - b) Pursue international discussions (for example through OIE, FAO, WHO...) on how this could be extended. Include links to operations such as the US BioSense, a data collection and analysis system of US health data. The whole process would necessarily be complex internationally.
- 2 We suggest a UK pan-research council research (e.g. EPSRC, ESRC, MRC, BBSRC...) initiative on the ITC issues that must be solved to address the above, with the involvement of the wider ITC community and other agencies as appropriate. It is urgent to define the research objectives here. For example, at the least techniques are needed to make extraction of useful information from large volumes of disparately structured data; and this is not easy. There will also be a need for an XML system to define the protocol for information exchange. It will also be desirable to experiment with real systems, for example one could use existing technology to 'trace' the transmission of a pathogen (such as human papilloma virus for example –or a non-sexually transmitted pathogen might be easier to start with) to demonstrate how the process

would work - and to find out hitches now rather than in the middle of an emerging epidemic crisis.

- 3 A key constraint will be the supply of skilled staff able to address these inherently interdisciplinary problems. We suggest that the UK should urgently consider how to address this need.

UC1 Roadmap



UC1 roadmap glossary

(Terms that are not defined are deemed to be self-explanatory)

Abstraction from unstructured data

This milestone is concerned with the ability to find patterns or key elements in unstructured data. Structured data are records in databases or cells in spreadsheets. Unstructured data are emails, reports, PowerPoint presentations, voicemail, photographs, web pages, etc.”

African electronic data collection spread, from South Africa northwards

Availability and collection of data in Africa will vary by region. There can be little doubt that the South African cell phone system as it currently exists means that in that country and adjacent countries in the region as far north as Zambia and Malawi, data collection will be possible – although following a gradient of reducing availability as one travels northwards. In other regions, use of cell phones may be non-existent or intermittent – with only limited hope for improvement. The Democratic Republic of Congo, southern Sudan and Angola probably fall into this category. In Uganda, Kenya and some other countries coverage is good, while in others, it is restricted mainly to the urban areas, for example, in Nigeria where satellite phones are used for local calls by prosperous individuals.

Augmented cognition – semiotic analysis

Computerised data mining, theme identification and analysis will be an invaluable tool in the process of gathering information from diverse sources of variable dependability. However, this cannot be a purely mechanical process. Human interpretation and monitoring will be essential if these data are to be properly interpreted and given meaning. Without this, there will be serious risks of, at best, mere waste and, at worst, missed significant messages.

Automatic interpretation of bioinformatic databases

As described under ‘Augmented cognition – semiotic analysis’, the effective interpretation of data will require human input. This milestone represents the point where bioinformatic data can be analysed automatically without continuous human involvement in interpretation. The word ‘continuous’ is of significance, and quality assurance schemes should be applied to the systems to ensure that high-quality, reliable, information is provided to the user.

Collection and interpretation of disparate data from available sources (CCTV etc.)

Analysing existing data sources for useful tracking or biomedical or behavioural characteristics (face recognition, gait recognition, etc).

Colouring plant

Genetically modified plants that change colour in response to the recognition of volatile organic compounds (VOCs) indicative of disease.

Cross-species disease information

Informatics, modelling and global networks could be used to predict which pathogens, and under what conditions, would be likely make a species jump. 'Hotspots' could be monitored by UC2.

Declining cost of ICT, data collection (cameras etc.)

The story of ICT emphasising that it is the steeply declining cost (as well as rising performance) that is the key driver; arguably the key driver for the modern world.

Diplomatic initiative

The clearly necessary but very challenging drive to obtain widespread agreement in pursuit of a collective benefit.

Dynamic semantic security

This milestone represents the ability to provide active and immediate security to databases and systems through the use of semantics. Security policies are semantically aware (based on the meaning of documents or messages) and dynamic (sensitive to the context in which access is requested)

Efficient international data hand-held device collation, analysis and stewardship (1) Human

Efficient international data hand-held device collation, analysis and stewardship (2) Animal

Efficient international data hand-held device collation, analysis and stewardship (3) Plant

These each define an ideal scenario objective.

GLEWS

Global Early Warning System for Transboundary Animal Diseases and Zoonoses. This is a joint initiative between the World Health Organization, the Food and Agriculture Organization and the World Organization for Animal Health (OIE).

Internationalisation

The drive to a connected world, in every sense.

Managing the health rather than the illness of the population

An ideal – the between this and existing techniques is that health is monitored at all times, with 'illness' defined in terms of departure from health rather than necessarily as a known ill condition. This would also enable the mapping of vulnerability among currently 'well' populations.

Mobile personal health/lifestyle monitoring

A developing mobile-phone-based service uses portable monitoring and location information to keep track of the user's health and lifestyle – this is aimed at personal support, initially of diabetics and other special groups but probably ultimately of all users. It will enable 'health'- as opposed to 'illness'-based services that may be vital for UC1.

Mobile phone with accurate location

While mobile phones inherently provide some location information (obviously needed to route incoming calls/data), the next generation will provide much more precise locations, probably via global positioning systems (GPS). This is driven by commercial opportunities but is very helpful for this User Challenge.

Modelling and analysis

Mathematical representations of relationships between variables, when based on well-estimated parameters and validated (i.e. based on fitting available data and, where possible, by making model predictions and then comparing data to the predictions), can provide valuable insights into the behaviour of systems.

Rapid-response control of nascent plant and animal diseases

An ideal scenario in which departures from normal would be picked up very early, enabling response at a more manageable stage.

Satellite remote sensing

As Earth observation becomes more finely tuned, it will be possible to observe individual plants, animals and humans. The data generated could inform UC1 networks and systems. For example, modelling could be used to predict environmental change and identify 'hotspots' for disease outbreaks.

Semantic mapping problem solved

Semantics is the study of meaning (which can be linguistic or non-linguistic). Considering the richness and diversity of different languages and their dialects, coupled with the use of language by varied and changing cultures, this is not trivial. It is necessary not only to recognise words but to appreciate their context. Semantics gives meaning to data. An ontology is the definition of a domain-specific vocabulary. Semantic mapping accounts for the mismatches between different ontologies.

Sentinel plants/remote sensing

A combination of satellite remote sensing and colouring plants. Sentinel plants, genetically modified to change colour in response to volatile organic compounds (VOCs) indicative of disease, could be placed at specific sites to detect pathogens before mainstream crops become diseased.

UC3 device

The hand-held devices described by UC3 fall into two basic categories:

- 1 simple devices such as urinary dipstick etc.
- 2 sophisticated devices linked to global network(s) – where data can be collected, interpreted, clustered and trended etc.

The second category is represented on the UC1 roadmap.


Web crawling/text mining

Automatically indexing the web and identifying meaningful content in free text. Text mining is a young interdisciplinary field which draws on information retrieval, data mining, machine learning, statistics and computational linguistic

Appendix - Key Examples

1: Challenges of influenza pandemic surveillance

Neil Ferguson and Christl Donnelly



Expected pattern of spread of an uncontrolled epidemic. Time sequence (30-240 days) of an epidemic in Thailand, showing spreading in a single simulation of an epidemic with $R_0 = 1.5$. Red indicates presence of infected individuals, green the density of people who have recovered from infection or died.

Neil M. Ferguson et al, Nature Vol 437 p 209, 8 September 2005

Surveillance and data gathering for influenza pandemics falls into two phases:

- a) Cluster detection. The initial cluster of cases associated with the emergence of a new pandemic virus needs to be identified as soon as possible, wherever the virus emerges in the world. Optimally, the aim would be to pick up a cluster before there were 50 cases.
- b) Tracking a pandemic. Once the spread of a pandemic virus has been detected, the extent of the spread needs to be monitored, both within each affected country and globally.

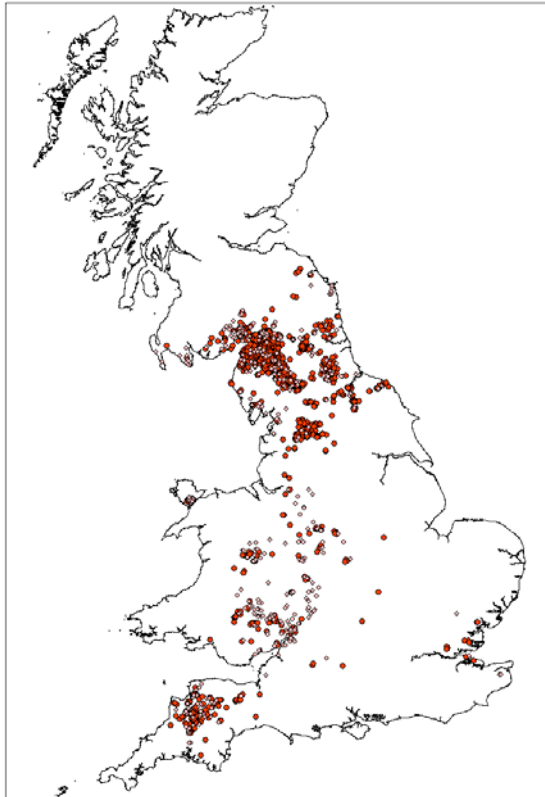
Cluster detection will require systems to detect abnormal outbreaks of respiratory disease in areas of the world with very undeveloped healthcare and/or surveillance systems. In many cases the key countries of concern may have a limited history of openness with the rest of the world about internal problems, and the political structures within such societies may not be optimal for the very rapid decision-making needed in a crisis situation. Novel data mining technologies may be key here – both in terms of the first country affected realising there is a problem (by detecting abnormal patterns of health care demand in a locality) and in terms of global alert. It may be that the rest of world finds out there is a problem before the source country admits to it, via one or more non-traditional data sources: e.g. weblogs, abnormal fluctuations in population travel patterns, covert intelligence gathering.

Pandemic tracking requires near real-time information on the occurrence of disease in a population, and national/global integration of data streams from multiple areas and sources. Most healthcare systems, even those in the developed world, do not have the infrastructure to provide such real-time incidence data, particularly when under the stress a pandemic is likely to impose on medical services. Early in the pandemic it may be possible to track

individual cases (in which case very detailed case data will be invaluable in refining response strategies). Once incidence levels are high, novel approaches to data collection (e.g. web- or telephone-based disease reporting by patients, use of proxy measures such as absenteeism rates, transport system passenger numbers, etc.) will be needed, in addition to major improvements in networking primary care and hospital-based patient database systems.

2: The UK epidemic of foot-and-mouth disease in 2001: data considerations

Mark Woolhouse



Geographic distribution of all 2026 farms reported to be infected with foot-and-mouth disease (FMD) in Britain during the 2001 epidemic. Of these, 1320 were confirmed positive for FMD by laboratory testing (red circles). For a variety of reasons, no laboratory confirmation was obtained for the remaining 706 cases (pink circles), and some of these farms may not, in fact, have had foot-and-mouth disease.

One of the key findings of the Royal Society's report on infectious diseases of livestock (Royal Society, 2002) was that:

'Quantitative modelling is one of the essential tools both for developing strategies in preparation for an outbreak and evaluating the effectiveness of control policies during an outbreak.'

Informed by experience of the 2001 epidemic of foot-and-mouth disease (FMD), the report recognised that developing and using such models was entirely dependent on appropriate and timely input data. These data would cover, as a minimum:

- i) the demography of the livestock population, detailing the location of livestock farms, the numbers and locations of animals on the farms, history of animal movements on and off farms;
- ii) the spread of disease and the measures implemented to contain the spread of disease.

During the 2001 epidemic, well-recognised problems were encountered with the availability, accuracy and accessibility of these data, which exacerbated the challenge of developing and using computer models to assist policy-makers. These problems need to be, and are being, addressed if data from detection, identification and monitoring activities are to contribute to the effective management of infectious disease problems in real time, not just for FMD but for a range of fast moving epidemic diseases of animals, plants and people.

Availability

In 2001, the UK did not maintain databases designed to record the current whereabouts and recent movements of livestock, information which would have been invaluable in helping to contain the FMD epidemic, especially in the early stages (the cattle tracing system (CTS) was already in place but it was not sufficiently up-to-date and accurate for disease control needs). Since 2001, progress has been made in this area. DEFRA has been developing a system known as RADAR (Rapid Analysis and Detection of Animal-related Risks) which will bring together key surveillance information in a structured and consistent way (see www.defra.gov.uk/animalh/diseases/vetsurveillance/radar/project.htm). RADAR and similar initiatives will benefit greatly from improvements in areas such as mapping technologies, smart tags, and the design of reporting systems. However, technology is not the only consideration. A prerequisite for a system such as RADAR to be effective is agreement between data collectors/owners, data users and other stakeholders on the specification of the data needed for predefined purposes.

Accuracy

The 2001 FMD epidemic revealed a number of inaccuracies in both disease data and demographic data. Most striking were inconsistencies in the diagnosis of FMD on individual farms, raising the possibility that some reported cases were not FMD at all (see Figure). At best, this kind of uncertainty decreases public confidence in the way the epidemic is being managed. At worst, it means that limited resources for disease control are sometimes being deployed inappropriately.

Demographic data was also often inaccurate (at least in part because they were not originally collected for disease control purposes – see above). Problems included discrepancies in the geo-referencing of farms and in numbers and types of livestock present.

Quality control is an essential aspect of data management and will become ever more important as the volume of data collected grows. However, the more accurate and precise the data, the more costly they will be to collect. Beyond a certain, definable level it may not be cost-effective to invest in further improving data quality.

Accessibility

During the 2001 FMD epidemic, case data were made available to researchers very promptly, but there were some delays in the provision of

demographic data and some difficulties in collating good information on other aspects of the epidemic, notably the pre-emptive culling effort.

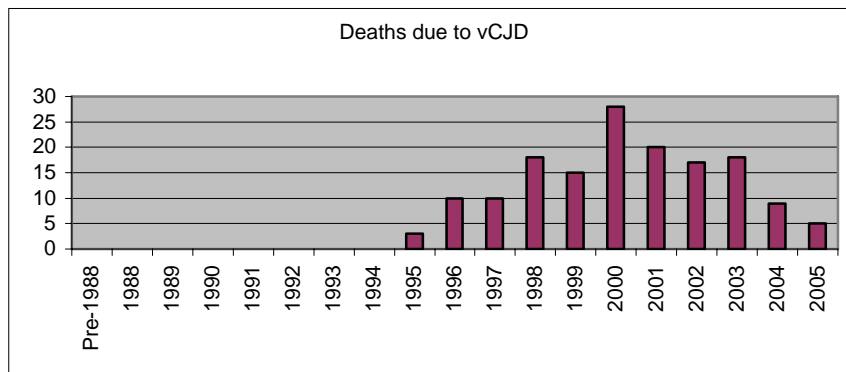
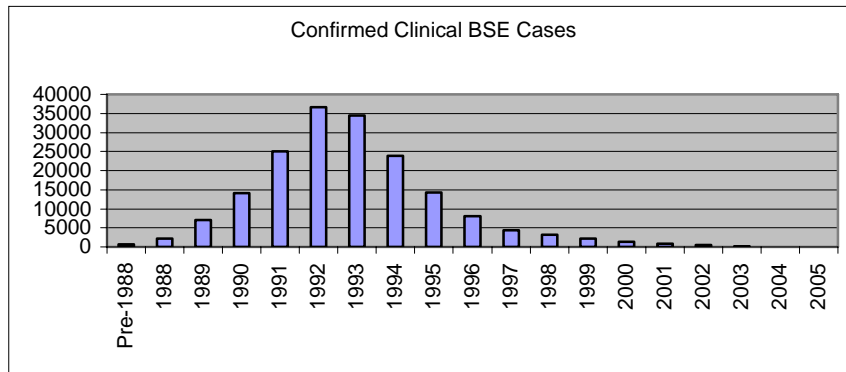
There are two major issues concerning data accessibility. First, there is the practical, technical challenge of transferring data from the field to data users in real time. In many cases – such as outbreaks of FMD or avian influenza in livestock, pandemic influenza in humans, or possible bioterrorism attacks, such transfers must be achieved over time scales of just a few hours. Technological advances that could contribute to advances in this area include computerised read-outs of diagnostic devices, wireless connections to the internet and generic data processing software.

A second set of issues is not technological but cultural. Data sharing raises questions to do with ownership of data and rights of access to data. Traditionally, individuals who have invested effort in acquiring data are often reluctant to grant access to others even, in some cases, when data collection is supported by public funds. Here, a change is required where the objective is seen not just as data collection but also appropriate (and, where possible, pre-committed) data communication. Achievement of this objective must be both recognised and rewarded. A related governance issue which has frequently been cited as putting restrictions on data sharing, is the protection of individual privacy. A proper balance must be found between individuals' rights of this kind and the need for the comprehensive and detailed knowledge about the populations – animal, plant or human – at risk of epidemic diseases that is needed to deploy maximally effective surveillance, prevention and control measures.

If there were an outbreak of FMD in the UK in the near future, there is little doubt that both disease data and demography data would be more readily available, more accurate and more accessible than in 2001. Moreover, the lessons learnt from the experiences of 2001 are more widely applicable and are being used to inform contingency planning for other livestock diseases, as well as human diseases such as influenza. Technological advances seem certain to make possible ever more sophisticated data collection, management and processing systems. Turning this potential into practice will require political will and substantial investment, but would significantly improve our capacity to respond to animal health and public health emergencies.

3 Experience from the BSE outbreak

Christl Donnelly



Following the identification in 1996 of new variant Creutzfeldt-Jakob disease (vCJD), the human disease related to bovine spongiform encephalopathy (BSE) in cattle, a selective culling scheme was undertaken to reduce future BSE incidence by targeting herds previously affected by BSE. Despite this measure, many countries banned the import of British cattle and beef. This ban was lifted more quickly for cattle in Northern Ireland than the rest of the UK because a detailed cattle tracing system had been operational in Northern Ireland for a number of years. A nationwide cattle tracing system was subsequently introduced, tracing individual cattle through their movements between herds. However, the 2001 epidemic of foot-and-mouth disease (FMD) demonstrated the limitations of this system for delivering accurate up-to-date information when it was urgently needed. Current concerns regarding the potential spread of avian influenza showed remaining holes in the agricultural system with efforts being quickly undertaken to record the locations and other details of all those keeping poultry. Clearly an integrated geographically referenced multi-species agricultural database is required to facilitate the timely and efficient implementation of control measures to fight zoonotic and agricultural diseases.

All the reports and papers produced within the Foresight project 'Infectious Diseases: preparing for the future,' may be downloaded from the Foresight website (www.foresight.gov.uk). Requests for hard copies may also be made through this website.

First published April 2006. Department of Trade and Industry. www.dti.gov.uk

© Crown copyright