Learning Analytics for the Academic: An Action Perspective

Alan Dix
(Talis, Birmingham, UK
and University of Birmingham, UK
alan.dix@talis.com)

Justin Leavesley
(Talis, Birmingham, UK
justin.leavsley@talis.com)

Abstract: If learning analytics are to directly benefit students' learning rather than simply inform broad policy decisions, they must be used by academics in the midst of busy and fragmented lives. This paper takes an ecological or action-oriented perspective of the use of learning analytics in higher education, drawing on research sources in psychology, human-computer interaction and visual analytics. We unpack the circumstances during the learning interactions of academics with course materials and students where analytics could trigger or influence action. This leads to a framework based around different academic timescales, and the strategies for synchronising the recognition of need with the potential for execution of teaching and learning interventions.

Keywords: Learning analytics, teaching analytics, action, human–computer interaction, embodiment, learning support
Categories: K.3 COMPUTERS AND EDUCATION, K.3.1 Computer-assisted instruction, H.5.2 User Interfaces, H.1.2 Human information processing

1 Introduction

Learning analytics have obvious benefits to senior management in an increasingly metrics driven world of higher education. However, if they are to directly benefit students' learning rather than simply inform broad policy decisions, they must be used by academics in their on-going relationships with students and management of teaching material. This paper takes an ecological or action-oriented perspective of the use of learning analytics in higher education. We operate from the premise that the benefits of analytics accrue only to the extent to which they inform action; that is we do not simply want academics and students to be able to comprehend available data, but to be able to take effective learning actions based on that information.

We will create a framework of the use of analytics embedded within the complex activities of academic life, drawing on theoretical insights from human–computer interaction, visual analytics, and learning technology; as well as practical experience in the teaching and in the production of software for teaching and learning. Critically we take the limited availability of time as a constraint and the need to close the 'teach learn' loop as a goal.

The resulting framework is intended to aid in the design and development of learning analytics software and user interfaces for the academic. This focus on practical application as well as theoretical foundations is driven from the dual commercial and educational roles of the authors

In the following section we outline the context and background: learning analytics in higher education, the commercial context of this work, an introduction to the embodiment literature, and relevant insights from visual analytics. Section 3 introduces a broad model of the lifecycle of interactions with learning resources as many of the analytics and interventions of interest relate directly to learning resource creation, use and management. Section 4 is the core of the paper applying an action perspective to learning analytics; issues of time, timeliness and timescales will emerge as crucial issues. Finally we will reflect on the broad lessons learnt and also identify gaps and opportunities for learning analytics research.

2 Context and Background

2.1 Drivers of learning analytics in the Higher Education sector

Learning analytics have become a major area of study over recent years. There have been a number of drivers for this.

The first is external; metrics are ubiquitous in academic life. Some of these are required by external monitoring agencies and some compiled by newspapers or other bodies in the creation of league tables (e.g. [THE, 14]). These external metrics are certainly critical financially, whether (in the case of league tables) in affecting student choices, or (in the case of government targets) in determining funding levels.

These external metrics create strong opinions on both sides. For example, in the UK the government Quality Assurance Agency audits universities' own internal 'quality' procedures [QAA, 14], which is often seen as potentially leading to overly bureaucratic mechanism and common denominator education rather than excellence. Such was the feeling over this that in 2001 the London School of Economics and other universities in the UK threatened to break away from QAA [THE, 01]. However, others have argued for greater external measurement to foster better value for money [Palfreyman, 07].

The need for universities to run 'leaner' operations also has an impact internally, with the need for courses and people to justify themselves both academically and financially. It is now common in promotion cases to see numeric data to backup narrative. For research, crude citation counts and H-indices are *de rigueur*. For teaching, the numbers are often less clear but various forms of student satisfaction metric are not uncommon, even if actual learning metrics are all but absent.

Much of this seems far from real learning, however education itself is increasingly required to be more evidence-based. The switch in many countries to using phonics for early reading is a prime example [NRP, 00]. However, it is less clear whether this is possible, let alone appropriate, in higher education.

Metrics to aid direct learning are, in some ways, the poor relation, partly because they are either too late (e.g. end of semester marks), too coarse (e.g. attendance at labs), or simply too uninformative compared with face-to-face feedback. To some extent the rise of MOOCs has changed this, offering the potential for near-real-time, big-data analysis of individual modules [Seaton, et al., 13]. Some of this is due to their volume (enough students to provide meaningful aggregate statistics), but some is simply building on long-standing TEL (technology-enhanced learning) research and the fact that digital delivery can allow finer grained analytics, whether for adapting individual learning or for feedback to the teacher. The lessons from MOOCs are now being re-evaluated to feed back into this broader picture of online learning [Yuan et al., 99].

It is this use of analytics to informal direct teaching and learning that is the focus of this paper, the use of analytics by the teacher to improve on-going and future teaching.

2.2 Development of teaching and learning software

As noted, this work has arisen directly from the practical development needs of specific teaching and learning software, although the issues raised will be common to many other commercial and research platforms. Talis has worked in higher education for many years and includes more than half of UK universities as customers. There are two core products. *Talis Aspire Reading List* (TARL) integrates into Virtual Learning Environments (VLEs) and Library Management Systems (LMS) offering a way to structure and provide resources for individual modules, such as recommended books, articles or web pages to read. *Talis Aspire Digital Content* (TADC) helps academics source portions of copyright material available under different forms of educational copyright agreement, and manages the process, keeping track of whether copying of a single source is within agreed limits, sourcing scans within legal limits and, where needed, automatically interacting with holding institutions such as the British Library.

Originally both these products used off-the-shelf players for content: browser PDF viewers, YouTube/Vimeo videos, etc., as is common in many VLEs. While these are very stable platforms they have disadvantages. First, the variety of different platforms leads to a fragmented learning experience and means that many core aspects of learning: annotation, deep linking, embedded quizzes, are either impossible or dealt with in a variety of different ways outside of the core learning platform.

The second drawback is that because the actual reading or watching of content is outside the learning software, analytics are only available at the level of page views and click-throughs on individual items. This level of analytics is useful; as a practical example, the first author was using TARL as part of a MOOC on Human Computer Interaction in early 2013 [HCIcourse, 13]; the view level analytics showed that while there were many video views, the vast proportion were to the earliest modules. This led him to slow the pace of release of modules, allowing students more time. However, these analytics are still at a relatively coarse level.

To address this a new universal player has been developed and, at the time of writing, at the beginning of trial roll-out. This is capable of viewing a variety of media formats (PDF, video, audio, slides) within a uniform interaction style and providing features such as those mentioned above in a similar way in every format.

As well as allowing improved learning experience during delivery of content, this means that fine-grained analytics will become available. Not only is it possible to see whether someone clicked through or downloaded a PDF, but exactly how much was read, how long was spent in each section, were there were breaks in reading, etc. This could be used at an individual student level, for example during one-to-one mentoring or for automatic personalisation. It could also be used at a class level, if a significant number of students either stop reading at a particular point, or maybe read some parts more slowly, this could be used as a way to hot-spot problematic sections to maybe cover again in revision lectures.

There are clearly many possibilities for using this kind of detailed analytic data, but not all will be equally valuable. We therefore have both a business and educational imperative to understand the way in which this kind of analytics can be used by academics on the ground.

2.3 Action and embodiment

A classic version of cognitive psychology takes cognition, what is in the head, as central: perception's primary role is to inform our internal model of the world, and then our actions in the world are in turn determined by our reasoning over that internal model. This is arguably a strawman, but whether or not true, has led to a number of 'reaction' positions that emphasise the deep entanglement between action and perception.

In psychology Gibson championed an ecological view of perception: we are creatures designed to act in the word, and so we are able to immediately perceive the action potential or 'affordances' of the things around us [Gibson, 79]. Within human computer interaction and user interaction design Norman's seven stage model of interaction has been widely adopted and modified [Norman, 98]; this starts with a goal-determined intention to act followed by multiple stages of elaboration, execution and evaluation. While Norman's goal-centred model is arguably close to the 'strawman' cognitivist approach, he was influenced by Gibson, and does emphasise the tight coupling between action and the perceived effect of our actions.

In philosophy and artificial intelligence notions of embodiment, with various names, trace their roots to Heidegger and Merleau-Ponty, and have become popularised by modern philosophers such as Clark and Gallagher [Heidegger, 27; Merleau-Ponty, 45; Clark, 98; Gallagher, 05]. The details differ, but the general message is that we cannot draw simple bounds for 'mind' or 'self' at the edge of physical bodies. The tools that we use, whether physical tools, such as Heidegger's classic hammer, or cognitive ones, such as adding up using paper and pencil, are as much part of us as are our hands or brain cells. As with Gibson, there is often a focus on the unconsidered or almost semi-automatic response to the world.

Although embodiment philosophy is clearly rooted in physical action, its basic principles have been applied in thinking about digital interactions including the meanings of embodiment for interactions with the web, whether as a massive extension of personal memory or as a matrix for collective intelligence [Carr et al., 09; Halpin et al., 10; Dix, 11].

There are limits certainly to the extreme forms of this position. For example, Gray and Fu showed that in certain circumstances, people do rely on internal models of the world even when it would be more efficient and effective to check what is really happening [Gray and Fu, 01]. However, despite these limits, the broad lesson is that action and perception are inseparable.

2.4 Visual analytics, the big picture

Data analysis is of course a major topic across business and government in general for financial, security, and environmental reasons as well as in education. Making sense of this data typically requires a combination of automatic data processing using statistical or machine learning techniques, combined with effective visualisation to enable human comprehension. The combined area, visual analytics, has had substantial impact in the US with a strong focus on homeland security applications [Thomas et al, 05], and in Europe with a broader agenda [Keim et al. 11].

The core problem of visualisation would normally be seen as one of presenting information in order to allow human understanding and insight. In visual analytics the focus shifts a little more towards the reasoning that occurs due to this, indeed the standard definition is: "the science of analytical reasoning facilitated by visual interactive interfaces" [Thomas et al., 05]. With the sensitisation of the previous section, it is clear there is an element missing.

The process for the creation of the European Roadmap document on visual analytics [Keim et al. 11], started from the existing visual analytics premise, but the focus on 'reasoning' inevitable led to the question "what is the reasoning for?" One portion of the road map focuses on human aspects (albeit still entitled "Perception and Cognition") and a core element of this is the adoption of a broader picture of the role of visual analytics where organisational context and the actions arising from analytics are core (see figure 1).

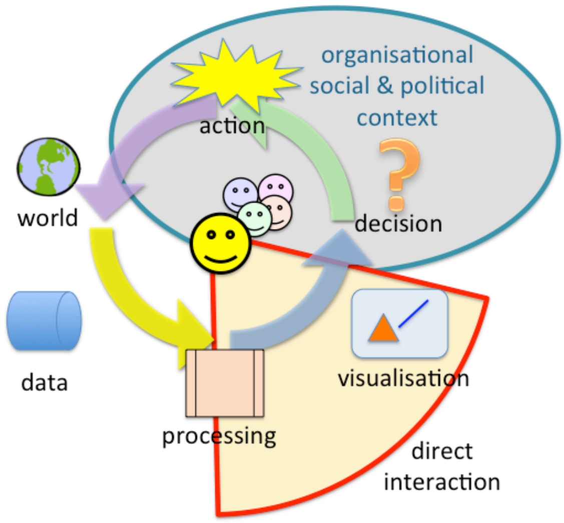

Figure 1: "The human context of visual analytics" (from [Dix et al. 11])

The obvious focus for visualisation and visual analytics is the bottom arc of automatic processing and visualisation of data, but this is not the end of the story. Instead, this direct interaction is seen as informing decisions that lead to actions, which inevitable change the world and create fresh data.

To some extent this is merely making explicit what will happen in practice anyway. However, it leads to core epistemological questions for predictive analytics when the actions taken as a result of prediction may alter the very system being predicted. In addition, very practical insights arise. For example, many systems focus on creating tools to allow an analyst to make sense of the data. Where the decision making process includes others, the analyst's own understanding is not sufficient. In addition it is critical to have tools to present the data in ways that help others, maybe less expert in data analysis, to understand the analyst's insights. That is visualisation is often about persuasion as well as information [Dix, 13].

In an educational context the latter may be critical, for example, if a learning support officer, or librarian notices some issue highlighted by learning analytics and wishes to bring this to an academic's attention, or if an academic wishes to use learning analytics to argue for more resources or to justify marking decisions.

A practical example of the latter occurs in a particular institution where there is a rule that any module where the average mark falls outside 55-65% should be considered for scaling. When this is on the low side a crucial question is whether this is effectively the 'fault' of the module (academic) being too difficult, or whether it is the fault of the students not working hard enough. If students were seen to have never engaged with the course material, this would be an argument against scaling. If however, students were seen to be opening resources, maybe very slowly getting part way through and then stopping, this would be evidence for difficulty.

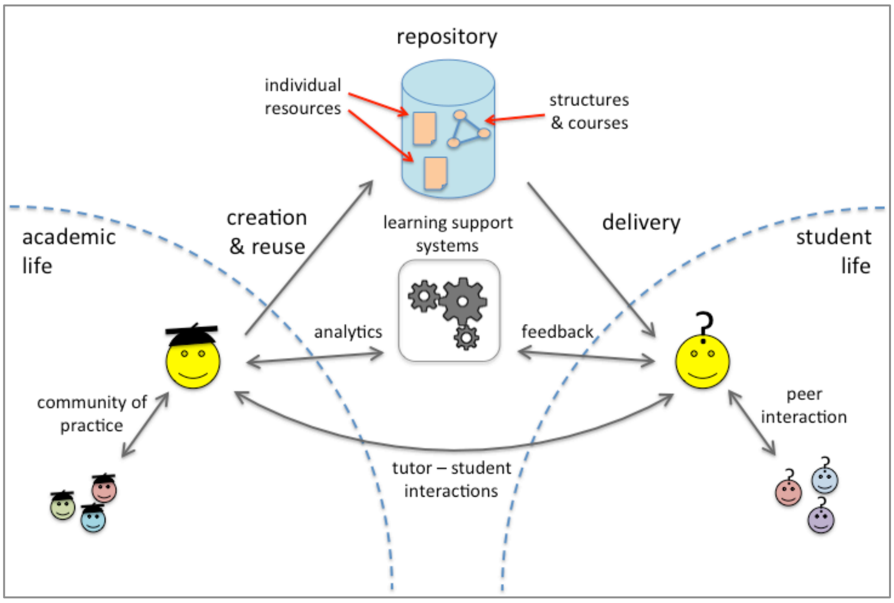

Figure 2: Learning resource lifecycle: actors, agents and events

3 Learning in context

Learning does not happen in isolation at the point when a student listens to a lecture, or reads a textbook, but is set within a broader context of activities, from the inception of a module, to the production of teaching materials, and the award of a qualification. Figure 2 shows a simplified view of a portion of this broad context focused particularly on the role of teaching resources (lecture notes, videos, etc.) as these are at the heart of many approaches to technology enhanced learning.

3.1 The resource lifecycle

We can follow the lifecycle of a resource round this picture:

Creation & reuse – Learning resources are created from scratch, chosen from pre-existing materials, or maybe appropriated (for example, using a Wikipedia page as a resource). Most often this is done in the context of a particular planned course. Even where the resource has been created outwith a particular course (e.g. textbook learning materials), the process of appropriation may involve some form of modification (e.g. editing slides) or annotation.

Repository **–** This includes atomic resources (a video, document, MCQ), but also structures (reading list, course outline, full curriculum). There has been significant R&D focused on standards and frameworks for this, especially focused on 'reusable' resources (SCORM [ADL, 14], Learning Objects [IEEE, 02], etc.). Note 'reusable' is in quotes as the level of actual reuse is often very low beyond the individual or small group involved in teaching a particular course [Weller, 07]. These and new standards are also important in learning analytics as data gathered about learning needs to refer to these resources and the data itself must be stored and reused [Cooper, 14].

Delivery – How to get learning resources into students' hands … or at least before their eyes. In some cases this may be more passive (in the sense of the learning system being passive), with students accessing resources by static course structure (e.g. reading lists), or maybe active, with recommendations based on more sophisticated algorithms [Baker and Siemens, 14].

Individual 'consumption' – The way the student accesses, interacts with and ultimately, hopefully, learns from the material. This is often the focus of pedagogic research. Rather like delivery, this may be passive consumption, or maybe more interactive, including real-time feedback on progress, personal learning models (feeding back into delivery), quizzes, suggestions for further leaning. The US Department of Education briefing document of learning analytics includes a detailed (imagined) scenario where students can view star-rated resources, see how they relate to their leaning goals, take diagnostic tests, and receive detailed analysis of how far they have progressed, all summarised in an individual learning dashboard [Bienkowski et al., 12].

Peer interactions – There is clearly a lot of work on peer learning in traditional settings, but intensified recently by the growth of MOOCs. This area includes use of blogs, peer tutoring, peer assessment, and classic group work. The focus is often as much on meta-cognitive and transferable skills as on the particular content. These interactions may be completely separate from the learning support system, or maybe integrated, for example, in students' comments or annotations on resources.

Direct tutor interactions – Traditionally this would have been face-to-face in lectures, tutorials or lab sessions. Now we also have numerous electronic means, but these are perhaps most interesting when closely embedded or connected to learning resources or more generally mediated by the learning systems. Distance learning systems have included various means to facilitate these tutor–student dialogues. MOOCs on the other hand, by their volume, emphasise the need to limit or control the need for individual personal feedback. In between, and useful at all levels of scale, are means to make sure that feedback is of value to as wide a group as possible.

Indirect tutor interaction and updates – Analytics offer just such an opportunity to channel time and effort into areas where there are common problems, potentially leading to augmenting material within the run of a course. While the majority of available analytics are at a module/course level, there has been some early work on using more detailed time-series visualisation of student video-viewing to offer academic (and student) feedback [Chorianopoulos and Giannakos, 13]. One of the main barriers to this so far has been the limited availability of data from standard playing platforms, a situation we are aiming to redress.

Live update enabled by analytics is of great value, not just for the students of the current run of a module. It is typically just after a module is complete when one becomes aware of the weaknesses of one's material, but of course there are other priorities at this point. When one next comes to update the material for a future run of the same course, it is easy to recall any large-scale issues, but smaller details are often forgotten, leading to the same misunderstandings, or problems recurring. In general bringing cost close to benefit is always a strong heuristic in user centred design [Dix et al., 04]; if academics can see what material is problematic and update it while students are still using it, then there is immediate benefit for the academic's effort, as well as the long-term benefit for future students.

3.2 The bigger picture

The use of resources in a particular course is only a small part of the broader life of the academic and the student.

The student life – Learning is set within a pattern of day-to-day routines (or lack of them), juggling multiple modules, and general time management. As well as direct learning students may be interested in progress monitoring (on individual modules and across their learning), and in some subjects the creation of portfolios of work for future employment. Research shows that "the amount of time and effort students put into academic pursuits … are associated with high levels of learning" [Kandiko and Mawer, 13], but time management is a problem for many students coming out of the structured academic school day and often away from home for the first time. Clear metrics may help students self-regulate, for example seeing whether the amount of effort is as expected by the course and/or more or less than peers. Of course, this same information could also help signal an early warning for failing students.

Academic peer interactions – Academics have a reference group both within their institution and beyond it. This community of practice is very rarely addressed during the day-to-day process of teaching, and more often periodically met at conferences or meetings. This may not matter, indeed during the running of a module, there is usually plenty to do! However, academics clearly do find it valuable to talk about what does and does not work, and generally share experiences. In addition, where there are re-used learning materials (e.g. textbook, video), it is extremely rare for the author of the materials to receive feedback of actual use. Subject to appropriate privacy protection, there is great potential in the return of detailed, but anonymised and aggregated analytics, to resource authors.

The academic life – Just like the student the teaching of an individual module is set amongst other teaching responsibilities, administrative functions, research and even (occasionally!) family life. Time is always the critical resource with multiple competing goals. Selecting where to apply ones effort and focus is always hard, and this is a great opportunity for even simple analytics (such as where drop-outs happen in video). In addition there are professional/contextual pressures such as satisfying external metrics and personal development, both clear opportunities for analytics.

Institutional context – The academics and students are of course usually working within an institution with its own pressures, external requirements, and internal processes. At this point we are in the area of large scale data analytics, which is relatively well covered in the existing literature.

3.3 Reflection – the student and academic life

Looking at this picture, there is some work on understanding the broad life of the student, for example the Society for Research into Higher Education have a 'Student Experience Network' [SHRE, 14] and the QAA commissioned a report on 'Student expectations and perceptions of higher education' [Kandiko and Mawer, 13]. However, it is not clear whether this connects to TEL research or influences the way specific learning systems fit into a student's life.

There is even less work on understanding the academic life, maybe because it is just too familiar. There is effort at a more social science level [Fitzgerald et al., 99], but this does not appear to have had an impact on actual TEL research and development.

The lifecycle model above highlights a number of open problems that are beyond the scope of this paper, for example, means to facilitate peer-peer interactions amongst academics. However, for the rest of this paper, we focus on the lower part of figure 2, the interactions between student and academic, and in particular the way this fits into the rhythms and activities of academic life.

This is not because the whole picture is not important. Indeed, at Talis, we have extensive commercial experience of the upper arc of resource creation and delivery. Similarly, the bulk of educational and TEL research and substantial tool support is focused on the direct learning of students, and this is also the area in which academics exercise their own professional judgement. Both continue to be important and are clearly in need of on-going support technology. However, it is precisely because the existing knowledge is strongest in these areas that we wish to focus on the gaps.

4 Analytics for the academic

As noted, there is little in the way of empirical studies of the nature of academic life, so the remaining analysis is based on a combination of theoretical insights outlined in section 2 and the first author's extensive practical experience of academia. To give context, this experience includes 20 years in posts in three Russell Group and 1994

Group universities (the elite UK groups); but also 6 years (including 2 years as Associate Dean) in 'New' (post 1992) universities, which tend to have a higher teaching balance; and also a period of 3 years as chief external examiner (crossing all subjects) for a higher education college during its transition into university status. While wide in terms of experience across the sector and in terms of the kinds of teaching and student body, these are all UK institutions so there is an inevitable a bias towards UK academic practice. However we believe many of the lessons are of broad applicability.

4.1 Analytics and action

We have argued for the necessity of an action-oriented approach to the use of learning analytics. More generally we are seeking to close the 'teach–learn loop', attempting to open more channels between learner and teacher at the most appropriate and timely points, so, in a way, this is also a natural extension to analytics.

This need for actionable analytics has been noted elsewhere. Clow argues for the importance of seeing learning analytics within a cycle of learners–data–metrics–interventions [Clow , 12]. Chow gives short examples of a student dashboard, an academic making contact with an individual student, and an academic reviewing grades and using this to review teaching for the next cohort; but no systematic analysis. Norris et al. also argue the case for 'action analytics' basing their argument on extensive university and college level examples [Norris et al., 08]; however, these are almost solely action at an institutional planning level. They also note that most current use of 'academic analytics' is backwards looking using past student data to feed into decisions for future cohorts, with relatively little, so far, on using student data to influence their own individual or class learning. A rare exception to this is 'Signals' a system developed at Purdue, which gave students a dashboard with red-amber-green traffic lights against different learning areas [Arnold, 10]. This was based on existing data from grading (relatively frequent in US system) and showed improved outcomes.

As a general rule, action happens when there is need and capability, and when the value (at a moment) exceeds the cost [Dix et al., 04]. There also has to be trigger event, something that prompts the action at a particular moment [Dix et al., 04b]. We wish to focus principally here on academics' actions, and so Figure 3 summarises some of the core drivers and capabilities for academic action with learning analytics, which are expanded in the succeeding sections.

4.2 Need/drivers for action

First the academic must recognise that some action is required. This recognition may be driven from learning analytics, whether automatic, when the system notices a trigger (e.g. if resource is marked essential for a stage in the course and few students have read it) or human when analytics is presented in some way and the academic notices an issue; whether for an individual student or the class as a whole.

The recognition may also come externally to the analytics, for example if a student comes to the academic with a problem, or if the academic notes that students appear confused in class; in this case analytics may act to inform the action, rather than trigger it.

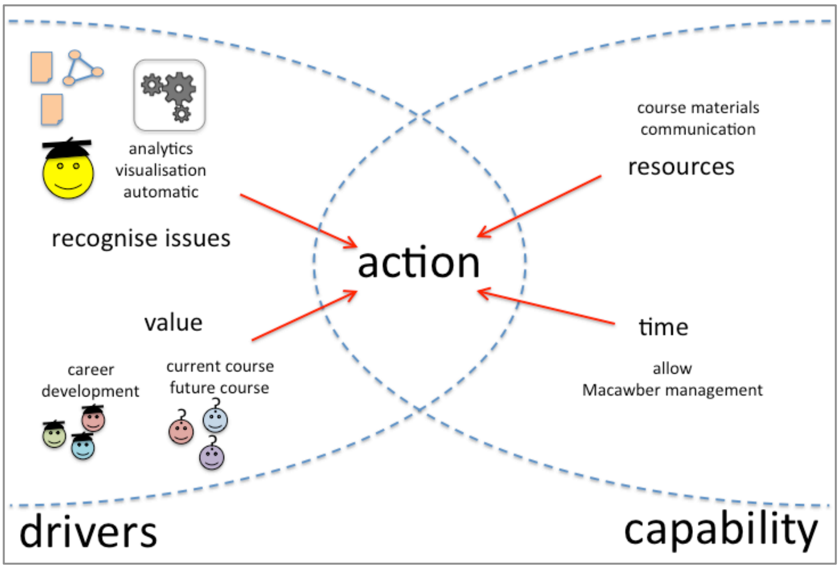

Figure 3: Drivers and capabilities for analytics-driven academic action

However, the need highlighted or informed by learning analytics will not give rise to action unless it is also regarded as giving some sort of value or benefit. This may be of value to particular student, for the course in general, or for broader education, maybe improving the course, or feeding into broader career development. Crucially the value may be in the future, in which case it is typically heavily depreciated.

4.3 Capability for action

The relevant resources need to be at hand. For one-to-one student contact this may mean having easy access to appropriate email addresses, or actually having the student in front of you. The teaching resources connected to the anaytics may also need to be available, either to make sense of a problem, or to be altered and updated. Furthermore, if the action was not triggered by learning analytics directly, relevant analytics need to be available.

For this to be effective tying learning analytics to resources is highly valuable, either making potential resource for action available when analytics demonstrates a need (e.g. easy means to drill through the course details, email groups of students, etc.), and/or to make sure analytics are easily available from resources. For example, when revising materials for a new run of a module, you need to be easily able to see analytics for the previous run of the same module – this is precisely what may be lost or made difficult during roll-over.

Available time and effort are also critical. Academics are always busy, with multiple competing demands. Even if the need for action is recognised, if there is not time at that moment, it may get forgotten. In such cases it is important that the need for action can be easily recorded for later. That is we must support 'Macawber' management – action is not something to do, but something to queue.

The other side of this is that when the academic has time, for example, on a train journey, or a gap between meetings, this is an opportunity for action if it is easy enough to do. That is availability of time is a potential trigger for action.

4.4 Strategies

The above suggests a number of strategies for enabling analytics for action:

analytics in context – making sure analytics are available when already in a teaching context: whether this is a one-to-one meeting with a student, reviewing teaching plans for the coming week, or revising a module at the year end.

automatic triggers – when the system notices a potential problem/issue, sending some sort of notification … and critically making resources easily available from this. The potential problem is this may not correspond to a moment when time is available. Asynchronous notification (e.g. email) can help allow the academic to organise time.

to do list / action queue – this may be an explicit to do list, or maybe just reminder emails in inbox. Crucially different people manage this in different ways, so flexible means to 'put off action' may be needed.

These strategies are not mutually exclusive. For example, the academic might receive periodic email (timed to correspond to likely availability) including traffic light items (automatic trigger), each of which offers the means for immediate action, but also 'ignore', or 'flag for later' (queue). The latter may the mean that there is a red flag marked against the module in a generic teaching dashboard, drilling down through this would then show the flag against a particular unit/week and highlight whatever prompted the trigger (to do list). When the academic is next looking at either general teaching for the week, or the module itself the issue and the available resources to deal with it are at hand (analytics in context).

4.5 Time

Time and timing is obviously critical, no matter how important something is, if the prompts for action do not come at the right time then nothing will happen.

Academic life has multiple rhythms, and effective design for action needs to accommodate these:

Days and hours – Email is often dealt with first thing in the morning while still in bed, or in the evening, as days are too full with people-centric activity; this is prime time for notification and a potential time to look at a 'quick things to do' academic dashboard.

During lectures or other teaching the academic may notice something, maybe a mistake or something that needs updating. Ideally this should be captured at the time, some sort of (virtual) big red button to press saying 'action needed here!' This might need triage soon after while memories fresh (a good 'five minute' action) to add a bit more context or even do the fix.

Lab or seminar style sessions allow more freedom for action in the moment. Libbrecht et al. present an imagined scenario, suggesting how their existing tools could be used in a university maths class to monitor progress in real time and then swop to whole class explanations of problems that emerge when appropriate [Libbrecht et al., 13].

During a one-to-one session with a student, or in some sort of meeting, the academic may need to call up analytics – for the former this would tend to be drill-down details, for the latter more aggregate statistics and visualisations.

Finally there are those odd few minutes in a day when the academic gets a gap – again a single 'what do I need to do' dashboard would be ideal for these potentially dead times.

Week – Academic patterns vary tremendously, especially in terms of weekly routines (if any). One prime time when analytics may be valuable to prompt or inform action is when (on Sunday evening or Monday morning) the academic is looking forward to the week ahead, preparing for classes, etc.

During the week specific activities arise that may also benefit from analytics, notably where there is some sort of exercise/assignment maybe when preparing it or assessing the results, or when looking for potential problem areas that might benefit from a quick reprise at the next lecture or in seminar group.

Months / mid-semester – Some institutions may have formal interim points for reporting on progress part way through courses, monthly staff meetings, others may have more sporadic mid-semester meetings. Irrespective of formal meetings, there are points when the academic reflects on the broad progress of a module.

This is the point at which automated analyses that identify potential student problems might be useful, as there is likely to be sufficient data to have some feel of overall trend for the class as a whole or individual students, but also time to do something about it. For example, one study showed that it was possible to predict 80% of K12 dropouts using a clustering algorithm based on comparison between individual students grading patterns and historical data [Bowers, 10].

However, while this is a potentially useful time frame for analytics triggered intervention, it is also one that may simply be lost in the business of teaching. It is most likely that insights of this sort will be seen during weekly or before-class triage, but need to be raised in a pattern that makes it easy for the academic to schedule action as these are likely to be relatively substantial activities.

End of semester, term or year –Towards the end of the semester or academic year (or sooner depending on internal processes), exams need to be prepared. An awareness of class problems may prompt the academic to timetable additional revision lessons, or maybe even drop questions if a topic is deemed to have been in some way problematic.

We have already seen a real example of a time in exam boards when analytics could be valuable, and in addition there may be times where certain analytics, for example, engagement in online discussion, may be included in assessment criteria. For example, the first author has used usage statistics for group project blogs to help adjudicate when members of a student group differ as to their assessment of individual contributions to group outputs.

Most departments have some sort of end of year teaching review. This often makes use of student questionnaire feedback as well as detailed assessment statistics. However, it rare to have any sort of analytics focused on the process of learning, largely because this has not hitherto been available. Note that here the focus is on modules as a whole, and particularly where some action may be necessary for next run of the module.

Start of or preparing for year – During the preparation for new run of an existing module quite detailed analytics on the previous run are potentially of great value in order to make adjustments. The end-of-year teaching review may highlight high-level points, but finer grained adjustments will tend to be the remit of the individual academic.

Most education systems include some sort of rollover, copying data from old modules to new versions, typically removing old student data to make a fresh start. However, this is potentially problematic as detailed analytics, for example, the patterns of student reading of resource, are likely to be lost. Ideally these should be retained in some sort of educational data warehouse, but this may then cause problems with privacy or data protection unless suitably anonymised or reduced to summary statistics [Slade and Prinsloo, 13].

When preparing new modules or updating old ones, ideally guidance should be available on how to easily add information that would make it easier to generate analytics or alerts. For example, academics might record whether a PDF is essential or extension reading, and the expected engagement time for a resource. Of course, any additional information of this sort needs to be very easy to add and its future value very clear.

Years – Some actions happen more intermittently over the years. This includes major course reviews, and the planning of whole new programmes of study. Obviously where these are brand new, there is little in the way of existing analytics, but new courses typically share modules with existing courses.

For example, an institution may be planning a new MA programme for interactive art hoping to recruit arts-oriented students who wish to employ digital technology. The course development team is considering using an existing Arduino programming module as an introductory programming module for the MA. If the module has previously been offered to a combination of arts and computing students, it would be very useful to drill into previous years' data to see if there were parts that the arts students found particularly difficult with a view to either putting these off to a later module or allocating additional teaching time. Note that individual student data is not needed for this, but aggregate data must to be available at the level of student characteristics (arts vs. computing).

Another process, with either annual or more intermittent timescale is professional development and appraisal. As noted earlier, hard numbers are increasingly importent in making promotion cases, especially as the university committees who consider promotion cases deal with every academic discipline and may be highly political. More qualitative or narrative arguments, unless they are backed by evidence, are likely to be ignored.

5 Conclusions

We have seen how an action oriented approach forces us to consider learning analytics not simply as the collection and visualisation of data, but rather as the collection and visualisation of data *for the purpose of action in a context*.

In some ways this is a truism of design, but easy to forget in the throws of creating the next new thing. Certainly the near universal lack of use or reluctant use of VLEs by academics [Lingard, 07] tells us that there is a gap between the ideals of creating rich learning environments, and the practicalities of academic life.

Our aim has been to map out the territory to clarify goals and issues in our own design of learning analytics systems for academics, and offer this to others. Many implications of this will only unfold in the context of particular design features and options. That is, our goal has been to enumerate and elucidate the dimensions, issues and criteria that need to be considered in design, but not to create specific design options. However, even in the creation of this more theoretical account, we have seen several practical insights.

Perhaps most critical is the importance of matching the point at which analytics (whether automatically or through visualisation) identifies a need for action, and the time when that action is possible. Having some sort way to say "not now but later" and not lose track is critical.

Another example is the importance of preserving appropriate analytics, including quite detailed viewing behaviours, through rollover. Although, once identified, this seems 'obvious', it is not the norm in many university systems.

The most frequent vision for academic analytics is the centralised dashboard, but we have seen through examples and analysis that analytics need to be available in a number of ways depending on the trigger for action: (i) student focused when there are individual meetings or emails, at which time it must be possible to drill down to individual student data; (ii) module focused when preparing for class or updating material, when it is critical that analytics are embedded in the academic's view of a module's resources; (iii) triggered by an automatic notification, in which case it is critical that actions can be scheduled for later (Macawber management) and that the notice links to both broader analytics and the modules/resources/students which gave rise to the alert; and (iv) when the academic has spare time, at which point a standard dashboard is indeed a suitable choice, although it is critical that data and alerts relevant to long-timescale activities (e.g. course updates) can be hidden or de-emphasised when necessary.

Most important however is that we have created a broad framework and vocabulary, which can frame and guide our own and, we hope, others' on-going work.

6 Future Work

The framework and analysis above are currently being used to inform the design process for the incorporation of learning analytics into Talis products. However, the issues are ones which are, or will be, common to many other educational services and platforms, and we hope this work will contribute to these also. As well as contributing to improved teaching and learning, we hope that the availability of rich analytics generated by our own and similar products will be a resource for researchers.

The analysis has highlighted gaps in current understanding of the context for learning analytics in particular and TEL in general; most critically the need for far more extensive empirical work on the nature of academic life in order to inform the design of future learning-support systems.

Acknowledgements

We thank all our colleagues, who design and produce the systems we have described, and to all the students on the 2013 open online HCI course, from whom we learnt such a lot.

References

[ADL, 14] SCORM, Advanced Distributed Learning Initiative, US Govt., accessed 29/4/2014, http://www.adlnet.gov/scorm/

[Arnold, 10] Arnold, K.: Signals: Applying academic analytics, Educause Quarterly, 33(1), 2010, http://www.educause.edu/ero/article/signals-applying-academic-analytics

[Baker and Siemens, 14] Baker, R., Siemens, G.: Educational Data Mining and Learning Analytics. In Cambridge Handbook of the Learning Sciences: 2nd Edition, 2014, http://www.columbia.edu/~rsb2162/BakerSiemensHandbook2013.pdf

[Bienkowski et al., 12] Bienkowski, M., Feng, M., MeansB.: Enhancing Teaching and Learning Through Educational Data Mining and Learning Analytics: An Issue Brief, U.S. Department of Education, 2012, http://www.ed.gov/edblogs/technology/files/2012/03/edm-la-brief.pdf

[Bowers, 10] Bowers, A.: Analyzing the Longitudinal K-12 Grading Histories of Entire Cohorts of Students: Grades, Data Driven Decision Making, Dropping out and Hierarchical Cluster Analysis, Practical Assessment, Research & Evaluation, 15(7) May 2010, http://eric.ed.gov/?id=EJ933686

[Carr et al., 09] Carr, L. and Harnad, S.: Offloading Cognition onto the Web, Proceedings of WebSci'09, 2009.

[Chorianopoulos and Giannakos, 13] Chorianopoulos, K., Giannakos, M.:, Merging Learner Performance with Browsing Behavior in Video Lectures, in WAVe 2013 at LAK 2013 (http://ceur-ws.org/Vol-983/)

[Clark, 98] Clark, A.: Being There: Putting Brain, Body and the World Together Again, MIT Press, 1998.

[Clow , 12] Clow, D.: The learning analytics cycle: closing the loop effectively", in LAK12: 2nd International Conference on Learning Analytics & Knowledge, 2012, Vancouver, BC. http://oro.open.ac.uk/34330/1/LAK12-DougClow-personalcopy.pdf

[Cooper, 14] Cooper, A.: Learning Analytics Interoperability – The Big Picture in Brief, Cetis, Bolton, 2014, http://laceproject.eu/publications/briefing-01.pdf

[Dix et al., 04] Dix, A., Finlay, J.,Abowd, G., Beale, R.: Human-Computer Interaction, third edition, Prentice Hall, 2004, ISBN 0-13-239864-8.

[Dix et al., 04b] Dix, A., Ramduny-Ellis, D., Wilkinson, J.: Trigger Analysis - understanding broken tasks, Chapter 19 in The Handbook of Task Analysis for Human-Computer Interaction. D. Diaper & N. Stanton (eds.). Lawrence Erlbaum Associates, 2004, pp.381-400

[Dix, 11] Dix, A.: A shifting boundary: the dynamics of internal cognition and the web as external representation, In Proceedings of the 3rd International Web Science Conference (WebSci '11). ACM, 2011, DOI:10.1145/2527031.2527056

[Dix et al., 11] Dix, A., Pohl, M., Ellis, G.: Chapter 7: Perception and Cognitive Aspects, In, [Keim et al. 2011], Eurographics Association, 2011, pp. 109–130.

[Dix, 13] Dix, A.: Introduction to Information Visualisation, In Information Retrieval Meets Information Visualization, LNCS 7757, 2013, pp 1-27

[Fitzgerald et al., 12] (2012). Fitzgerald, T., White, J., Gunter, H.(ed.): Hard Labour? Academic Work and the Changing Landscape of Higher Education (International Perspectives on Higher Education Research, Volume 7), Emerald Group Publishing Limited, 2012, http://www.emeraldinsight.com/books.htm?issn=1479-3628&volume=7

[Gallagher, 05] Gallagher, S.: How the Body Shapes the Mind, Oxford University Press, 2005

[Gibson, 79] Gibson, J.: The Ecological Approach to Visual Perception, Lawrence Erlbaum, 1979.

[Gray et al., 01] Gray, W., Fu, W.: Ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head., In Proc. of CHI '01. ACM Press, 2001, pp.112–119, DOI=10.1145/365024.365061

[Halpin et al., 10] Halpin, H., Clark, A., Wheeler, M.: Towards a Philosophy of the Web: Representation, Enaction, Collective Intelligence, In Proc. of the WebSci10, 2010

[HCIcourse, 13] HCIcourse.com, Open Online Course on Human Computer Interaction, 2013 (now also available at interaction-design.org)

[Heidegger, 1927] Heidegger, M.. Sein und Zeit, 1927. (English translation: Being and Time, Harper, 2008)

[IEEE, 02] IEEE Standard for Learning Object Metadata, IEEE 1484.12.1-2002, 2002, http://ltsc.ieee.org/wg12/20020612-Final-LOM-Draft.html

[Kandiko and Mawer, 13] Kandiko, C., Mawer, M.: Student Expectations and Perceptions of Higher Education, King's Learning Institute, King's College London, 2013, https://www.kcl.ac.uk/study/learningteaching/kli/research/student-experience/student-expectations-perceptions-HE.aspx

[Keim et al. 11] Keim, D., Kohlhammer, J., Ellis, E., Mansmann, F. (eds): Mastering the Information Age: Solving Problems with Visual Analytics, Eurographics Association, 2011, ISBN 978-3-905673-77-7.

[Libbrecht et al., 13] Libbrecht, P., Kortenkamp, U., Rebholz, S., Müller, W.: Tales of a Companion Teacher Analytics, IWTA, 2013, http://ceur-ws.org/Vol-985/

[Lingard, 07] Lingard, M.: Why don't all lecturers make use of VLEs?: what can the so-called "laggards" tell us?, Masters thesis, Institute of Education, University of London, 2007, http://eprints.lse.ac.uk/28529/

[Merleau-Ponty, 1945] Merleau-Ponty, M.: Phénomènologie de la Perception, 1945. (English translation: Phenomenology of Perception, Routledge, 1958)

[Norman, 98] Norman, D.: The Design of Everyday Things, MIT Press, 1998.

[Norris et al., 08] Norris, D., Baer, L., Leonard, J., Pugliese, L., Lefrere, P.: Action analytics: Measuring and Improving Performance That Matters in Higher Education, Educause Review, 43(1), 2008, 42-67. http://www.educause.edu/ero/article/action-analytics-measuring-and-improving-performance-matters-higher-education

[NRP, 00] Teaching Children To Read: An Evidence-Based Assessment of the Scientific Research Literature on Reading and Its Implications for Reading Instruction. National Reading Panel, U.S. Department of Health and Human Services, NIH Pub. No. 00-4769, April 2000, http://www.nichd.nih.gov/publications/pubs/nrp/pages/smallbook.aspx

[Palfreyman, 07] Palfreyman , D.: Markets, models and metrics in higher education, Perspectives: Policy and Practice in Higher Education, Vol. 11, Iss. 3, 2007, http://oxcheps.new.ox.ac.uk/MainSite%20pages/Resources/OxCHEPS_OP29iv.pdf

[QAA, 14] Quality Assurance Agency, accessed 28/4/2014, http://www.qaa.ac.uk/

[Seaton, et al., 13] Seaton, D., Bergner, Y., Chuang, I., Mitros, P., Pritchard, D.: Towards Real-Time Analytics in MOOCs, International Workshop on Teaching Analytics (IWTA) 2013, http://ceur-ws.org/Vol-985/

[SHRE, 14] SRHE (Society for Research into Higher Education), Student Experience Network, accessed 25/4/2014, http://www.srhe.ac.uk/networks/student_experience.asp

[Slade and Prinsloo, 13] Slade, S., Prinsloo, P.: Learning analytics: ethical issues and dilemmas, American Behavioral Scientist, 57(10), 2013, pp. 1509–1528, http://oro.open.ac.uk/36594/

[THE, 01] LSE leads revolt against QAA, Time Higher Education Supplement, 23 March, 2001, http://www.timeshighereducation.co.uk/news/lse-leads-revolt-against-qaa/158362.article

[THE, 14] Times Higher Education World University Rankings 2013-2014, http://www.timeshighereducation.co.uk/world-university-rankings/

[Thomas et al, 05] Thomas, J., Cook, K. (ed.): Illuminating the Path: The R&D Agenda for Visual Analytics, 2005. National Visualization and Analytics Center.

[Weller, 07] Weller, M.: Learning objects, learning design, and adoption through succession. Journal of Computing in Higher Education, 19(1), 2007, pp. 26–47.

[Yuan et al., 99] Yuan, L., Powell, S., Olivier, B.: Beyond MOOCs: Sustainable Online Learning in Institutions, CETIS, 2014, http://publications.cetis.ac.uk/2014/898