Assessing the Quality of Computer Science Courses Using Taxonomy Tools

Neville Ivor Williams, B.Sc.

A thesis submitted in total fulfilment of the requirements
for the degree of Doctor of Philosophy

School of Computer Science, Engineering and Mathematics
Faculty of Science and Engineering
Flinders University, Adelaide, South Australia

30 June 2015

Contents

List of Figures

List of Tables

Abstract

In the current era of higher education, where there are pressures to both attract and retain students in degree courses, two significant factors that influence student choices are the reputation of the university and the quality of the courses offered. The research undertaken in relation to this thesis is the latter factor dealing with course quality, and in particular the development of new approaches and an innovative metric that may be used to provide indicative guidance about the expected learning rigour to which students will be exposed in the selected course.

While traditional approaches to assessing course quality have focussed on the examination of student assessments, assignments, examination outcomes, project work and interviews with staff and students, little has been done to examine the learning demands placed upon the students. In this thesis, the specifications for the subjects in courses have been scrutinised using the SOLO Taxonomy, and quantified by a method previously described by Brabrand and Dahl (2007) to generate scores for the various subjects in a course. By aggregation according to the course rules it has been possible to develop learning rigour profiles for each year level in the course and an overall course profile which highlights the different types of learning expectations for the course. In addition to the overall course profile a numeric value labelled the C-Index has been calculated and it has been proposed that this value should be interpreted as an indicator of the level of learning rigour expectation for the course.

With the detailed level of analysis that occurred in constructing a course profile, a composite view of the subjects in the course allowed for further grouping and analysis to take place. When the subject data was compared within year levels it was clear that some subjects appeared to place much higher learning demands on students than others. Although being outside the scope of this research to determine whether such demands were reasonable or not, the analysis has been able to identify where potential problems may exist in courses when either the demands are too great or are not sufficiently strong

for the year level concerned. The methodologies used in this research are proposed as being beneficial tools for university curriculum groups to assist in monitoring the internal quality control aspects of the courses for which they are responsible.

The context of this research has been in the domain of Information Technology and Computer Science, where course quality and accreditation are important matters. The techniques proposed will provide additional tools to accreditation and benchmarking teams by providing course profile information that may be used to support the observations they make about the accompanying course materials.

The outcomes of this research include the creation of a new metric labelled the C-Index; the description of a methodology to construct a course profile; a proposed method to identify "subjects of interest" within a degree program; and the documentation of an approach that may be used for course benchmarking either within a particular university or across universities.

Declaration

I certify that this thesis does not incorporate without acknowledgement any material previously submitted for a degree or diploma in any university; and that to the best of my knowledge and belief it does not contain any material previously published or written by another person except where due reference is made in the text.

Neville Ivor Williams ____________________ Date ____________________

Publications

During the course of this research a number of relevant articles have been submitted and published. The list of publications is:

1. Williams, N. I., 2013, *A Quality Framework to Measure Course Objectives*, Proceedings of the 4th International Conference on e-Learning, SDIWC, 14-25, Ostrava, Czech Republic.

2. Williams, N. I., 2013, *Constructing a Course Profile by Measuring Course Objectives*, International Journal of Digital Information and Wireless Communications, 3(4), 16-28.

3. Williams, N. I., 2014, *Course Quality Starts with Knowing Its C-Index*, Issues in Informing Science and Information Technology, 11, 225-237.
   http://iisit.org/Vol11/IISITv11p225-237Williams0453.pdf

Chapter 1

Introduction

1.1 Motivation

In the higher education sector a continuing issue remains at the forefront of the teaching and learning agenda, and that is the quality of the degree programs offered. There are many initiatives undertaken to investigate the quality of teaching, the quality of assessment, the quality of graduates, and other output evaluations. Oliver et al. (2007) highlighted the undertakings at Curtin University in mapping curricula as part of the increased level of demand for accountability in the university sector and the potential benefit in improved teaching and learning. In the Faculty of Science and Engineering at Flinders University "... a strategic decision to appoint a dedicated Quality Assurance (QA) Coordinator to assist in its preparations for the AUQA audit" was undertaken in 2005 (Smith and Martin, 2006). The University of Sydney noted that it was preferable to focus on 'quality enhancement' rather than 'quality assurance' (McLean and Sachs, 2006). Keane and Labhrainn (2005) identified that substantial efforts in the development of systems for evaluating teaching and course quality in higher education were well established in the US, UK, and Australia, and was proposing the more systematic introduction of such systems into the Irish University system. Most efforts are post-event or post-process activities that have an important role in attempting to validate and maintain institutional quality standards. Applications of this approach are particularly evident at times of course accreditations or during benchmarking processes when various forms of documentation are provided as evidence of effective quality teaching, learning and assessment processes being in place.

When these evaluation and audit processes occur, a significant element that is examined is the documentation associated with a complete program (a degree, or a course) and its component elements (the individual subjects, or topics, or courses). In particular, the aims and objectives are reviewed, and then the outputs and deliverables associated with the member items are examined and evaluated. The parenthesised terms above are listed to show the variability in terminology usage across the education sector where for example the term *course* may mean either a whole degree program in one institution or a semester (subject) of study in another.

The main focus of this research is the initial part of this documentation, namely the aims and objectives of individual component elements, and to propose that a profile can be established for each of the subjects in a degree program, and by extension therefore to arrive at an overall course profile that may be used as an indicator of course intent. When implemented, this profile can be used to provide key stakeholders with a predictive capacity that presently does not exist. For example, students could compare courses in a quantitative manner to supplement their qualitative decision making on course selection. University administrations could compare courses within their institution to confirm consistency or identify inconsistency between department offerings. External course reviewers and evaluators could establish baseline expectations for the conduct of their reviews and audits. Being an input-side or pre-process activity, there is an inherent value in such a profile being created.

While at first appearing to be either confronting, or perhaps an impossible dream, it should be pointed out that many other areas of endeavour have metrics that are used to provide initial expectations for evaluators on which to base their judgements. Simple examples include the ‘degree of difficulty’ factor used in judging some Olympic events such as diving, gymnastics, dance and similar. In the University research sector there exists the h-index, an index that attempts to measure both the productivity and impact of the published work of a researcher (Hirsch, 2005). Tofallis (2012) discussed the issue of attempting to modify the approach to determining university rankings in the UK, and noted that "the 'league tables' ... are here to stay". One of the conclusions presented by Oliver et al. (2004) after using the Bloom Taxonomy to analyse a number of Computer Science subjects was that "...there may also be potential for using a Bloom analysis in order to standardise results across a range of courses in a similar fashion to diving competitions ...". In the health sector in Australia the case-mix approach identifies a 'standard' time in hospital for various medical procedures, and in financial accounting there exists the 'standard cost' for production of component items in manufacturing processes. Why then should it not be possible to establish a baseline value that may be used as an indicator to the educational potential of course-work studies? As will be shown in the remainder of this thesis, a proposed course profile indicator is feasible.

1.2 Research Question

The fundamental question of this research is whether taxonomic tools are able to be applied to the course objectives for university level Information Technology and Computer Science courses to provide an indicator of course quality.

There are supporting questions that must be answered in order to arrive at a definitive and supported result in answer to the primary question. In particular,

- Which taxonomy tools are appropriate for Information Technology and Computer Science courses?
- Is there a suitable metric that is able to be derived using the taxonomy tools?
- Does the tool metric provide a useful measure for assessing the learning rigour of a course?
- Do similar courses return a similar result using the metric?

1.3 Contribution

There are several significant contributions that emerge in this thesis, which has examined the documentation of degree courses in the area of Computer Science, Information Technology, and Software Engineering.

An underlying theme which has driven this is to consider the nature of technological computing-oriented degree courses as having a faint alignment with software products. Stemming from this is the idea that a student's education is largely a software process that generates a suitable outcome, a graduate, at the end of the process. In the tertiary sector a great deal of effort is applied to the quality control stages at various output levels, but there is little formal assessment done to evaluate the input specifications to that process. An assertion that is made in this thesis is that the learning objectives for the subjects in a degree course are equivalent to the specifications for a software product, and may therefore be validated by a process that is separate from the output testing which we know as subject assessment that confirms whether a student has met the subject specifications or not. It is this external verification of the input-side specification through the assessment of subject learning objectives that is a new approach in the determination of a suitable metric for degree courses.

Following on from previous research in the area, the major contribution of this thesis is the proposal of a course metric for degree courses, designated as the C-Index. There is a formal definition given for both the C-Index, a metric to describe an evaluation of the learning level expectation for a degree course, and the p-index, a companion metric to describe the associated learning level expectations for the various year levels within the degree course. The thesis chapters 4 and 5 provide detailed information about the theoretical determination of these metrics and the supporting experimental evidence.

While the C-Index metric determination is conceptually straight-forward, the obvious question that arises is whether it can support widespread usage. The extended experiment described in Chapter 7 addresses this question and demonstrates the potential for the metric to be applied in benchmarking exercises when reviewing degree courses in the Computer Science and Information Technology domains. Future research options are to explore its applicability beyond these domain areas.

1.4 Thesis Structure

This thesis is organised in the following manner:

Chapter 1 introduces the topic area and describes the nature of the research problem being addressed. An integral part of this introduction is to provide the motivation for the research and to describe the primary research question and the supporting questions that are necessary components leading to a qualified answer to the fundamental hypothesis of the primary research question.

Chapter 2 presents a review of the literature within the area of relevant educational theory, learning taxonomies, quality concepts and applications in the Computer Science domain.

Chapter 3 describes the methodology used in this research and the rationale behind the approach taken. It is an important chapter as it sets out the experimental design philosophy and highlights the underlying structures of the models to be used in the experimental components of the research.

Chapter 4 describes the initial calculation techniques explored using both the SOLO Taxonomy and the revised Bloom Taxonomy on a known course where access to course and subject coordinators was readily available. Although based on the methodology of previous researchers in this area, the study was conducted to confirm the relevance to the Australian context and to extend the methods used to create a course metric. With the in-depth analysis of the individual subjects in a course a substantial amount of data became available and the resultant aggregation of subject scores into frequency tables, which are labelled SOLO Distributions (Brabrand and Dahl, 2007), enabled a course profile to be constructed showing the relevant proportions of learning levels specified in the degree. The course profile was most informative when presented graphically, showing the shift of learning demands through each of the years of study and culminating in an overall profile view. As highlighted in the discussion section of this chapter the results obtained using the SOLO scale and the Bloom scale were very similar.

Chapter 5 builds on the calculation methods described in Chapter 4 and proposes the formalised mathematical structure for the creation of a Course Index, named as the C-Index. The formal methods also describe the statistic p-index which is the nominated term for an individual index score for a year level of a degree course. To support the creation of the p-index and C-Index metrics, a more detailed study across several degree courses in a known environment was undertaken and there are comparative results shown and discussed in the latter part of this chapter.

Further analysis of the course data has been shown to be a potentially useful tool to compare the expected learning demands of subjects within a course, and may therefore be considered as a procedure to assist with the internal quality control of the subject descriptors. While there is an expected amount of variation in learning demands from one subject to another there is also an implicit expectation that the learning demands for subjects in a particular year level should be approximately similar. Yorke (2003) makes a similar assertion stating " . . . most programmes in higher education are based on a set of general assumptions in which the subject discipline, rather than student development, is dominant." Accordingly, Chapter 6 explores several approaches to improve the identification of potential areas where the expression of the expected learning demands of subjects may need to be reviewed. This chapter discusses the application of those techniques to several courses and demonstrates the outcomes in both tabular and graphical formats.

A second consequence of close analysis of the subjects in a course was found to be that courses could be compared in a quantitative manner, either internally within a single institution, or externally across several institutions. A broader study across several university courses was undertaken, and is described in Chapter 7. To control the experiment and validate the metric as a comparative measure, courses from Australian Universities in the areas of Computer Science, Information Technology, or Software Engineering were selected where the details of the course schedule and rules, plus the behavioural objectives or learning outcomes, were available on the University's public web pages. The investigation across several Australian Universities provides a potential application for the methodology as a benchmarking tool. In the realm of benchmarking, it is clearly more valuable if the tool is applicable beyond the Australian boundaries, but this was not incorporated in this thesis.

The overall project and research is reviewed in Chapter 8 where the methodology is critiqued, and the results are analysed and discussed in detail.

Finally, Chapter 9 presents the conclusions and proposes several areas for future research activity.

Chapter 2

Related Literature

2.1 Overview

As indicated in the title of this thesis, this research brings together a number of different elements from the domains of Learning Theory, Educational Learning Taxonomies, the discipline of Computer Science and the concepts of quality in a higher education environment. In particular, the literature covered looked at Constructivist Theory and the works of John Biggs, Bloom's Taxonomy in its more recent revised form, the SOLO Taxonomy, and a number of representative works demonstrating the application of these theories and models to the teaching and learning of Computer Science. The discipline area of Computer Science in the context of this thesis should be interpreted as an overarching term that includes the more specific sub-disciplines of Computer Science, Information Technology, and Software Engineering.

2.2 Relevant Learning Theory

2.2.1 Constructivist Theory

A popular branch of more contemporary learning theory is known as the constructivist approach, which has the following attributes (Schmidt and Winterhalter, 2004):

- Learning is the construction and refinement of knowledge structures in learners' minds;
- The construction process depends mainly on the personal effort and engagement of the learner;
- Knowledge cannot be transferred or trained, but must be built in each individual learner; and
- Learning should be self-determined and situated in real-life situations.

Liu (2003) cites several authors in describing constructivist theory and proposes that "...most constructivists agree on these four essential characteristics" which influence learning as being:

- Learners construct their own learning;
- New learning depends on current understanding;
- Learning is facilitated by social interaction; and
- Meaningful learning occurs within authentic learning tasks.

It can be seen that these two views are very similar and emphasise the basis of the theory as considering that learning is a cumulative process that builds upon previous experience and knowledge. The secondary part of the theory is that the success of the approach is dependent on the engagement of the learner and the recognition of the learning activities being appropriately relevant to the learner. The applications of contructivist approaches are seen in the stage-based models and problem-based learning (Wilson, 1996; Jonassen, 1999).

It is clear that under this model, the common teacher driven approach is inappropriate, and, as a consequence, the function of the teacher becomes more about being a resource that the student can use in their own synthesis of suitable knowledge structures. Adopting this approach has a clear implication of workload pressures on teachers wishing to be able to facilitate student learning when they have substantial numbers of students in their classes. Wang (2011) identified that "...learning outcomes now represent the guiding principles in curriculum design." Wang particularly stressed the importance of carefully designing the intended learning outcomes for a course prior to implementing Outcomes Based Education (OBE) successfully.

2.2.2 The Revised Bloom Taxonomy

Beyond the mere philosophy of teaching and learning, a number of studies have been undertaken to explore the variety of approaches to the implementation of teaching, or the practice of teaching. One of the key platforms that gained a great deal of support was the taxonomy of educational objectives proposed by Bloom (1956), which subsequently became widely referred to as "Bloom's Taxonomy". The underlying basis of Bloom's ideas were to create a framework for classifying the statements of what was expected for students to learn through the teaching process. While the original publication of Bloom's work dates back to the 1950s, further discussion and analysis has taken place over many years, and has been updated to now incorporate amended aspects in what is described as the Revised Bloom Taxonomy (Anderson and Krathwohl, 2001; Krathwohl, 2002). In essence, the revised taxonomy has expanded the Knowledge dimension of the original taxonomy and has become represented as a two-dimensional matrix mapping the Knowledge dimension against the Cognitive dimension as shown in Table 2.1 (Krathwohl, 2002). Use of this tabular form allowed the analysis of the objectives of a unit or course of study, and in particular, enabled an indication of the extent to which more complex types of knowledge and cognitive processes were involved.

Table 2.1: Revised Bloom Taxonomy Matrix

Knowledge Dimension	**Cognitive Dimension**					
	Remember	Understand	Apply	Analyse	Evaluate	Create
Factual Knowledge						
Conceptual Knowledge						
Procedural Knowledge						
Metacognitive Knowledge						

In the accompanying table (Table 2.1), the terms in the cognitive dimension are self-explanatory, and similarly, the first three terms in the knowledge dimension are equally self-explanatory. However, the fourth term, "Metacognitive Knowledge" requires further explanation. In a related work, Pintrich (2002) discusses the importance of metacognitive knowledge and expands on the three distinct types proposed by Anderson and Krathwohl. Specifically, Pintrich considers the first type, "Strategic Knowledge", as incorporating the knowledge of strategies for learning, thinking and problem solving in the domain area. The second type, "Knowledge about cognitive tasks", includes the ability to discern more about the nature of the problems to be solved and to begin to know about the "what" and "how" of different strategies as well as "when" and "why" the strategies may be appropriate. The third type, "Self-Knowledge", includes understanding about one's own strengths and weaknesses with respect to learning.

It was found that this tabular form was able to be applied across a range of granularities, from the fine-grained analysis of a module in a larger teaching program, to broader analyses of subject objectives (Meerbaum-Salant et al., 2010; Fuller et al., 2007; Thompson et al., 2008; Johnson and Fuller, 2006). The application of the revised Bloom Taxonomy matrix involves the examination of learning objectives and classifying them into the appropriate cells of the matrix.

Several studies have investigated the suitability of the Bloom Taxonomy in the field of Computer Science (Scott, 2003; Oliver et al., 2004; Whalley et al., 2006; He and Brandt, 2007; Gluga et al., 2012a), and most appear to examine the various micro-level aspects of individual subject components such as the practical tests, assignment work, and examinations. That knowledge about the use of taxonomies is relevant to lecturers in the field of Computer Science is emphasised in the study by Gluga et al. (2012a) in which the project developed a training package in using the Bloom Taxonomy for the teachers of a programming fundamentals subject. The broad aim of this project was to provide a stronger appreciation of learning competence progression in programming subjects and to then see a closer link between future teaching activities and the achievement of the expressed learning objectives.

The analysis undertaken by Oliver et al. (2004) described the determination of a 'Bloom Rating' using a scale of 1 to 6 corresponding to the cognitive levels of the Bloom Taxonomy for parts of the assessment instruments in several subjects that were considered, and concluded that there were observable differences in the two different subject streams reviewed. In the first stream there were three subjects from a programming stream that were examined, and in the second stream there were three subjects from a data communications and networking stream that were examined. It was highlighted in Sitthiworachart (2004) that the Bloom levels 1 to 3 were considered as *surface learning*, and the levels 4 to 6 were viewed as *deep learning*.

2.2.3 The SOLO Taxonomy

A strong proponent of the constructivist approach was John Biggs, who coined the phrase 'constructive alignment' (Biggs and Tang, 2007), and describes it as "... we start with the outcomes we intend students to learn, and align teaching and assessment to those outcomes. The outcome statements contain a learning activity, a verb, that students need to perform to best achieve the outcome ...". The constructionist part flows from the general philosophy that learning is built upon the activities that students carry out, with learning resulting from what they do, and is not about what teachers do (Biggs, 2011). The natural extension of this idea is that the teaching process is merely the catalyst to learning.

Some of the key ideas are summarised on Biggs' personal web-page (Biggs, 2011), particularly:

- Constructive alignment is an example of outcomes-based education (OBE). His version is concerned with only improving teaching and learning and as such has been successfully implemented in universities all over the world.
- Constructive alignment can be used for individual courses, for degree programmes, and at the institutional level, for aligning all teaching to graduate attributes.
- The SOLO Taxonomy (Structure of the Observed Learning Outcome) helps to map levels of understanding that can be built into the intended learning outcomes and to create the assessment criteria or rubrics.

A detailed discussion of the SOLO Taxonomy, which "provides a measure of the quality of assimilation in terms of structural complexity" and leads to the ability to "assess student work in terms of its quality . . . " (Biggs, 2011), is given in several Biggs publications (Biggs and Collis, 1982; Biggs, 1979, 1999; Biggs and Tang, 2007), but a succinct description is available in Biggs (1979), which outlines the 5-level taxonomy as:

- Level 1 – Pre-Structural; The response has no logical relationship to the display, being based on inability to comprehend, tautology or idiosyncratic relevance.
- Level 2 – Uni-Structural; The response contains one relevant item from the display, but misses others that might modify or contradict the response. There is a rapid closure that oversimplifies the issue.
- Level 3 – Multi-Structural; The response contains several relevant items, but only those that are consistent with the chosen conclusion are stated. Closure is selective and premature.
- Level 4 – Relational; Most or all of the relevant data are used, and conflicts resolved by the use of a relating concept that applies to the given context of the display, which leads to a firm conclusion.
- Level 5 – Extended Abstract; The context is seen only as one instance of a general case. Questioning of basic assumptions, counter examples and new data are often given that did not form part of the original display. Consequently a firm closure is often seen to be inappropriate.

In this same publication, Biggs recognises that the SOLO Taxonomy is functionally close to the Bloom Taxonomy, and also highlights that it has been applied across a wide range of subject areas.

The application of the SOLO Taxonomy to the assessment of learning outcomes (objectives) involves the review of the objectives in terms of the functionality expected at the various levels. In particular, there are typical verbs associated with each level that are likely to appear in statements of learning objectives. Important features associated with the SOLO Taxonomy as presented by Biggs and Collis (1982) are the notions that the SOLO Taxonomy is hierarchical and student learning tends to be progressive from the more quantitative outcomes associated with Levels 2 and 3, through to becoming more qualitative at the higher Levels 4 and 5. Additionally the terms *surface learning*, at the quantitative stage, and *deep learning*, at the qualitative stage are highlighted.

Slack et al. (2003) examined the relevance of the SOLO Taxonomy and the attributes of students in gaining deep learning in their subject area, observing that "…students who are personally involved in learning from real life situations are the ones who are most likely to experience deep learning" and further cited other research (McAllister et al (1997) in Slack et al. (2003)) identifying that "…the deep learning approaches were in stark contrast to the surface learning approaches exhibited by students who sought only to memorise and reproduce information or skills."

Killen (2005) distinguishes between *deep knowledge* and *deep understanding*, explaining that *deep knowledge* is considered as a characteristic of the content that students are studying, and that *deep understanding* is something that develops in the mind of the learner as they learn about *deep knowledge*. Killen further suggests that attaining *deep understanding* would correspond to being classified at the highest levels on both the SOLO Taxonomy and the revised Bloom Taxonomy, thus making it equivalent to the more generally used *deep learning* expressed in both taxonomies.

2.3 Applications of Taxonomies in Learning Research

While there are many studies in the education sector that have explored the use of Bloom's Taxonomy (original or revised) and the SOLO Taxonomy in various ways, in the context of this thesis the major focus is on the Science-oriented domain areas, and particularly the Computer Science – Information Technology domain at University level.

2.3.1 The Brabrand-Dahl Study

One large-scale study conducted in Denmark across the courses offered at the University of Aarhus and the University of Southern Denmark investigated the stated course objectives of all the science course subjects at the two universities using SOLO Taxonomy classifications (Brabrand and Dahl, 2007).

This study attempted to provide a quantitative value conversion from the qualitative base of the taxonomy structure and considered some 550 syllabi from the science faculties at the two universities. The approach in this study listed a number of typical verbs associated with the SOLO Taxonomy, adopting the Biggs and Collis proposition that levels 2 and 3 provided mostly quantitative outcomes and levels 4 and 5 were more qualitative in nature, as shown in Table 2.2. The mapping of each learning objective statement to a value was then given by the level number that the verb(s) in the objective most closely matched.

Table 2.2: Prototypical Verbs According to the SOLO Taxonomy (Brabrand and Dahl, 2007)

Quantitative		**Qualitative**	
SOLO 2 Uni-structural	**SOLO 3** Multi-structural	**SOLO 4** Relational	**SOLO 5** Extended Abstract
Paraphrase	Combine	Analyse	Theorize
Define	Classify	Compare	Generalize
Identify	Structure	Contrast	Hypothesize
Count	Describe	Integrate	Predict
Name	Enumerate	Relate	Judge
Recite	List	Explain causes	Reflect
Follow (simple) instructions	Do algorithm Apply method	Apply Theory (to its domain)	Transfer Theory (to new domain)

While the initial intention of using the SOLO Taxonomy is to classify learning objectives into the appropriate SOLO categories, the work undertaken by Brabrand and Dahl enabled a relative measure of competencies to be established across the courses in the science faculties in the universities in the study. The body of evidence in the Brabrand and Dahl work has established a method to create a quantitative measure based on the statements of learning objectives.

The method used by Brabrand and Dahl in the examination of syllabi was to count the frequencies of the verbs used in the learning objectives for the subjects and apply an average to the subject. It was further enhanced by using what is described as a 'double-weight averaging scheme', which meant that compound statements of learning objectives such as "identify ... and compare ..." would result in an averaging for that single objective of (S2 + S4)/2. In this approach, the values 2 to 5 were applied to the learning objectives based on their verb classification. The outcome of this method is to create a singular value for each subject syllabus objective within the range 2 to 5, and ultimately generate a single value for each subject. The SOLO-1 (Pre-Structural) level was omitted as this is the ab-initio or naïve state which would not appear as part of any learning objective as all teaching and learning activities would be targeted to levels 2 and above. As described by the authors, there is an underlying assumption that the distance between each SOLO level is equal to enable the values 2 to 5 to be used in this manner. The term for this metric given by Brabrand and Dahl is "**SOLO Average**". The equal distance assumption is supported by a similar approach with the Oliver study (Oliver et al., 2004) to analyse the cognitive difficulty of two streams of ICT subjects in which the values 1 to 6 were used to correspond to the various levels of the Bloom Taxonomy for assignment and examination tasks in each of three subjects from a programming stream and each of three subjects from a data communications and networks stream. In that study the relative weights of the elements comprising the assessments were used to calculate an overall score which the authors labelled as the "Bloom Rating".

In applying this approach, it becomes clear that the double-weight averaging scheme will return a singular result in the range 2 to 5 for each subject. When examining various syllabus statements it is also apparent that it is necessary to take an average of the scores as individual subjects may have few or many learning objective statements. Therefore the mean is a simple but effective method to arrive at a standardised score for a subject.

A by-product of applying the double-weight averaging scheme to the assessment of behavioural objectives occurred as Brabrand and Dahl counted the instances of each of the descriptors in the various SOLO categories. In their analysis they examined the frequencies of the descriptors in each of the SOLO-2 to SOLO-5 categories which they described as the "**SOLO Distribution**". Over the course of the study they were able to compare the proportions of the various SOLO categories across domain areas.

In the paper discussion the method used to determine the SOLO Distribution could have been either of two possibilities. The first was to proportionalise the frequency count within each of the learning outcome statements. That is the sum of the frequencies equalled the number of learning outcome statements, or each learning outcome statement contributed a value of 1 in the determination of the SOLO Distribution. This meant that where learning outcomes were expressed in compound statement terms, the aggregations of the SOLO scores may have resulted in fractional values. The distribution percentage was then calculated relative to the number of learning outcome statements. This approach was used in the analysis of an individual subject (Figure 10 in Brabrand and Dahl (2007)). The underlying rationale for this choice was that each learning outcome for a subject should have equal value rather than being distorted by individual counts that could occur with the compound statements. For the purposes of this research the author has chosen to label this as the 'proportional' approach.

The second possibility was to simply count the number of occurrences of the relevant verbs that were given in the stated learning outcomes. Depending on how the learning outcome statements were worded, complex and compound statements might contribute to giving a higher number of SOLO verbs for a subject than would be the case with the proportional approach. Those subjects with many learning outcome statements would also contribute more raw count frequencies than those with few statements unless standardised by an appropriate method. It was noted in a subsequent paper by Brabrand and Dahl that using the raw counting method may have the effect of lowering the overall SOLO score for a subject when the learning outcome statement contained a list of several content elements at the same SOLO level (Brabrand and Dahl, 2009). For the purposes of this research the author has chosen to label this as the 'simplex' approach.

The use of the SOLO classifications is quite simple at the conceptual level, and it also has an implied equality of learning competencies within each level. Hence, any learning activity that is classified at a particular SOLO level may be thought of as being educationally equivalent to every other learning activity at that level.

Interesting observations and conclusions reported in the Brabrand-Dahl study include:

- The terms 'surface understanding (or surface learning)' and 'deep understanding (or deep learning)' are easy to define in conjunction with the SOLO Taxonomy – surface learning implying that the student is confined to action at the lower SOLO levels (2-3), whereas deep learning implies that the student can act at any SOLO level. Students producing a high-level response at SOLO 4-5 are deemed to have a deep understanding of the subject matter;

- The contributing elements in the calculation of the SOLO average were able to be used across individual subjects and whole courses, giving rise to a SOLO Distribution;
- There were notable differences between the SOLO Distributions for Computer Science, Natural Science and Mathematics, with Computer Science scoring a higher number of the upper level SOLO verbs than the other two domain areas.

2.3.2 Learning Taxonomies in Computer Science

More specific studies concentrating on the application of the SOLO Taxonomy to the Computer Science domain have been undertaken in works such as that of Sheard et al. (2008) "Going SOLO to Assess Novice Programmers", and Lister et al. (2006) "Not Seeing the Forest for the Trees: Novice Programmers and the SOLO Taxonomy".

The Lister study claimed to apply the SOLO Taxonomy to "the study of how novice programmers manifest their understanding of code . . . ", which had not previously been done. Important features of this paper were the interpretation of the five SOLO levels as they related to the programming context in Computer Science; the observation that level 5, the extended abstract response, was unlikely to be observed in the focus group; and that the SOLO taxonomy was a useful organising framework for comparing work relevant to the testing of novice programmers.

The Sheard study built on the previous Lister study and aimed to address some of the deficiencies or inconclusive outcomes of that work. The more interesting outcomes from this study were that some support for the Lister assertion that a better SOLO reading performance should produce a better code writing result was shown, and that a higher level of SOLO responses was obtained from postgraduate students than undergraduate students in an introductory programming unit (asserted as being a result of postgraduate students having developed higher level thinking skills during their undergraduate degree).

The nature of learning progression was a secondary part of the Brabrand-Dahl study and is discussed comprehensively in "Using the SOLO taxonomy to analyze competence progression in science curricula" (Brabrand and Dahl, 2009). The findings from this study supported the use of the SOLO taxonomy as an analysis tool in a number of ways, particularly:

- that "the use of the SOLO Taxonomy showed that competency progression in terms of SOLO does indeed exist, . . . , from undergraduate to graduate level";
- that "the SOLO Taxonomy has 'been proven' to be a good tool for analyzing competence progression";
- that "not all verbs have a fixed SOLO-level and that some are connected with the faculty in question"; and
- that "the use of the SOLO language might hopefully result in more clear explanations to respond to student questions about the relevance of topic matter, and result in fewer non-operational and ambiguous learning objectives such as 'understanding'."

2.3.3 The ACM-IEEE Computer Science Curricula 2013 Report

The ACM and IEEE-Computer Society have cooperated over many years to propose appropriate suggestions about curriculum content in courses related to Computer Science, Computer Engineering, Information Systems, Information Technology and Software Engineering (ACM/IEEE, 2013). The Computer Science Curricula 2013 (CS2013) document incorporates guidelines to address a "redefined body of knowledge" which was "the result of rethinking the essentials necessary for a Computer Science curriculum."

Interesting elements that appear in that report include the observations of:

- *Terminology* – adoption of the term "course" to mean an institutionally recognised unit of study, while noting that some institutions may also use other terminologies such as "module" or "paper";

- *Bloom's Taxonomy* – acknowledgement of reference to Bloom's Taxonomy in the development of the three-level mastery classification system. While Bloom's levels were not chosen to be applied directly, the mastery classifications reflect aspects of learning classifications that have some similarity to the Bloom levels. The three mastery levels were described as
  - Familiarity – implying student understanding of a concept at a basic awareness level;
  - Usage – implying the ability to use or apply a concept in a concrete manner; and
  - Assessment – implying that the student is able to consider a concept from multiple viewpoints and select an appropriate approach from understood alternatives.
- *Learning Outcomes* – are not of equal size and do not have a uniform mapping to curriculum hours. Recognition was made that the proposed learning outcomes for the body of knowledge courses may not exactly match those used by institutions.
- An updated list of Knowledge Areas to reflect changes in the discipline since the previous revision in 2008.

While the CS2013 document provides a number of suggestions about content and possible approaches in the various knowledge areas described in their Body of Knowledge, it remains the domain of individual institutions to design their own degree programs which may or may not align closely with those recommendations and suggestions. Equally, the suggested three-level mastery classification scheme is not currently in common usage within educational institutions, and will therefore not be covered further.

2.3.4 The CDIO Approach

CDIO (Conceive - Design - Implement - Operate) is an initiative aimed at improving undergraduate education in engineering internationally. Originating at MIT (USA) and extending to other universities, it is underpinned by four tenets (Berggren et al., 2003):

- Curriculum reform to ensure that students have opportunities to develop the knowledge, skills and attitudes to conceive and design complex systems and products;
- Improved level of teaching and learning necessary for deep understanding of technical information and skills;
- Experiential learning environments provided by laboratories and workshops; and
- Effective assessment methods to determine quality and improve the learning process.

CDIO is more about a holistic approach to the teaching and learning in engineering programs and is not an educational taxonomy as such. Accordingly it does not really fit within the scope of this thesis.

2.3.5 Alternative Taxonomies

The application of the Revised Bloom Taxonomy and the SOLO Taxonomy to the Computer Science discipline is not fully endorsed by researchers in the field of Computer Science education and learning. When there have been difficulties in fitting the peculiarities of the discipline, attempts have been made to either modify one or other of the taxonomies or create a blended form that more closely meets their viewpoint. An example of the hybrid approach was proposed in Meerbaum-Salant et al. (2010), where their conclusion was that "the combined taxonomy captured the cognitive characteristics of CS practice". Notwithstanding the claim to there potentially being merit in the combined taxonomy approach, there were some anomolies discovered, with a possible reason being given that the taxonomy which required discrete classification categories may not be entirely suitable for some aspects where a continuum of development across cognitive categories would have been more appropriate.

In another case, it was suggested that there were problems in applying the Bloom taxonomy to the Computer Science discipline, referring to works of Lahtinen (2007) that describe the poor fit of cognitive tasks under the Bloom taxonomy framework for beginning programming students (Fuller et al., 2007). In response to the difficulties, the researchers proposed a revised taxonomy for application in the Computer Science discipline, with a focus directed more towards programming related subject areas. There do appear to be some strong arguments proposed for the case that neither the Bloom nor the SOLO taxonomy provide a good fit for programming related tasks. However, programming is just one aspect of Computer Science and Information Technology courses overall, and the nature of this research has not addressed the appropriateness of either taxonomy at the low-level detail of individual learning elements within subjects. Therefore, even though the criticism of the Bloom and SOLO taxonomies for a number of programming related tasks may be valid, it does not exclude their use in this research.

A different approach was taken in Bower (2008), where the objective was to create a taxonomy of the task types found in computing disciplines. Again oriented more to the programming types of tasks, the taxonomy presents a hierarchical list of task types which were claimed to more closely match the types of learning activities expected of students in the discipline area.

Although these various alternative taxonomy approaches were interesting they each tended to focus on the detail level learning activities associated with the development of programming language competency and problem solving skills. There are important messages in these research initiatives and most propose further study to extend their findings, but they are beyond the scope of this thesis.

2.4 Quality in the Higher Education Sector

When considering educational programs at any level there is always an underlying question as to whether the program is a good one or not. This is especially the case in the higher education sector where universities are being asked to be more accountable for their operations. The two terms of quality and accountability are often used in a shared close relationship sense as the improvement in one tends to drive improvement in the other. Liu (2009) stated that "...accountability is needed in higher education for the same reasons that it is needed in K-12 education and in any other area of education. Because a good education has become a pathway to opportunities and success, stakeholders deserve to know whether institutions have done their best to maximize student learning and have effectively utilized public resources." Although Liu was discussing the US Voluntary System of Accountability (VSA) for publicly funded institutions, the sentiments expressed are another indicator of the greater focus being given to higher education quality and accountability. In part this drive towards better educational programs in the higher education sector has spawned the creation of a number of journals devoted to quality matters, or as stated in Harvey and Williams (2010), "...it began as a result of the Quality in Higher Education project, a funded project to explore the meanings of quality in the early 1990s. In addition, the journal has provided a professional publication for the International Network of Quality Assurance Agencies in Higher Education."

There has been much debate over many years about the various interpretations of the term 'quality' and what that means in the higher education sector context. One of the more significant papers in this area was that of Melrose (1998) which proposed three fundamental paradigms of curriculum evaluation, labelling them as functional, transactional, and critical. It was further proposed in Harvey and Williams' discussion of this work (Harvey and Williams, 2010) that "any model or tool for curriculum evaluation ... developed by an institute or group of staff has an underlying philosophy ... that can be matched to one or more of these paradigms."

The functional paradigm proferred by Melrose (1998) has been given an alternate label of 'technical' and has the key attribute of attempting to measure program outcomes against pre-stated goals in order to comply with governmental or institutional objectives. Melrose proposed that typically the evaluators working within this paradigm did so as an independent expert whether working alone or within a team employed for the evaluative task. Accordingly behavioural objective goal attainment models are representative exemplars for this paradigm.

The transactional paradigm also has an alternate label of 'naturalistic', having a key attribute of being based on "liberal humanism and subjectivist ethics" ((House, 1978) cited in Melrose (1998)). Typical of this approach is the use of focus groups and interviews to gather data to be interpreted by collaborative groups of relevant stakeholders such as educators and students who plan and implement the evaluation process. It was considered that this approach allowed for emergent change to planned programs as they were underway.

The third of the paradigms, the critical paradigm, was considered to be 'emancipatory', as it enabled learning communities to be self-evaluating, with critical reflection, and ultimately empowered to set their own standards.

Although the functional paradigm may at first appear as the most attractive to university administrations, Melrose has argued that this is appropriate for a compliance approach and the application of interpreting quality as 'fitness for purpose'. However it was also pointed out that the transactional paradigm was a better classification where there was input from the target audience regarding the program development, and indeed may have a stronger link to the critical paradigm when quality is viewed as a transformation leading to the notion of empowerment of the learner.

Other attempts to define quality in the higher education sector have been made, but Tam (2001) has made a clear case to recognise various interpretations of quality depending on context and viewpoint. In particular, Tam addresses the different roles and usages of quality related terms including *quality control*, *quality assurance*, *quality audit*, *quality assessment*, and *indicator systems*. Many of these are based on the loose association with the manufacturing-production approach, examining inputs and outputs, and proceeds to identify that the associated indicators do not give rise to being able to comment on the student experience within the higher education realm. Perhaps the most significant point arising from the discussion is that there are many different component parts that contribute to the notion of quality in higher education, with no singular best approach that covers all viewpoints. More simply, different stakeholders place higher value on different indicators that best suit their interest.

Jordens and Zepke (2009) proposed that a curriculum be considered as a network comprising a number of elements and described it as "...the sum of learning experiences in a unit of study (module, paper, course or programme) and encompasses discipline knowledge, teaching and learning activities and the learning environment in a structure that facilitates the desired learning in the individual student." Their subsequent deduction was that the quality of a curriculum is determined by how well it achieves the purpose of facilitating desired learning in a student or group of students. From this, the focus of the measurement of quality moved towards the quality assurance processes, citing external quality assurance agencies in higher education with a general aim of ensuring that "curricula meet similar quality standards." They also noted that even though other authors have suggested that education needs to take a transformative view of quality, the contemporary quality assurance approaches appear to neglect that view. Effectively, when agencies, either internal or external, adopt a standards-based approach they are really demonstrating a compliance mode typical of the functional paradigm.

There is a clear imperative that perceptions of quality are important in the higher education sector as pointed out by Tofallis (2012) when he discussed the issue of attempting to modify the approach to determining university rankings in the UK, and noted that "the 'league tables' ... are here to stay." Accordingly, the manner in which the quality of institutions or the courses they offer is significant as it will ultimately affect the number of students attending or wishing to attend particular universities.

2.4.1 Quality Management

The concept of quality, or more particularly quality management, has taken on different interpretations depending on the context and domain area in which it appears. It has perhaps gained its greatest common usage in matters relating to manufacturing and production industries and is often attributed to the work of Walter Shewhart in introducing scientific method in improving the work process during the 1920s (Zairi, 2013). His work using statistical process control, and the subsequent efforts of Deming, Juran and Crosby have led to the application of quality principles to management processes and the term "Total Quality Management" has arisen to become a paradigm followed by many enterprises (Zairi, 2013).

There is no particular reason why universities cannot adopt some or many of the quality management principles in the education process. It has been suggested that there are five underlying quality management principles that need to be embraced (Kuei and Lu, 2013), which are:

- Facilitating increased awareness of quality and market signals;
- Enabling conditions for quality;
- Adopting a systems approach;
- Achieving greater communication and alignment between cross-organisational units; and
- Examining for congruence with quality objectives.

One of the key elements that becomes clear when one delves further into the quality management paradigm is the philosophy of continuous improvement and building quality into the product. When applied to an educational program, the opportunities begin at the course design and specification stage. Hence if one adopts the view that the learning objectives or learning outcomes of the subjects in a course of study may be considered as the specifications for the subjects in a degree program, then the level of learning rigour that the course proposes may be gauged.

What has been shown in this chapter is that the use of learning taxonomies has been investigated in many different ways across many segments of the Computer Science and Information Technology fields. While a large proportion of these investigations appears to have concentrated on the relevance of particular taxonomies to the teaching and learning of programming within Computer Science and Information Technology at University, especially at the introductory or first-year level, sometimes resulting in attempts to create more tailored versions or hybrid approaches, it is clear that the use of taxonomic tools in the research of learning and teaching in Computer Science and Information Technology courses is essential. It is equally important to recognise the nature of analysis and evaluation criteria being used, and to understand the quality paradigm within which the research was undertaken. As will be explained in the next chapter, the methodology for this research was based on applying taxonomic evaluation techniques across a number of University level courses in the Computer Science and Information Technology field, with the result that a new course profile metric was able to be determined.

Chapter 3

Methodology

The preceding chapter has discussed the importance of using educational taxonomies in the research of teaching and learning in University courses generally, and in the fields of Computer Science and Information Technology in particular. The strong body of evidence from previous research has contributed to the creation of new applications of techniques and the formulation of a new evaluation metric that enables course profiles for degree programs to be determined.

This chapter describes the methodology of the approach undertaken in this research. It contains the broad design criteria and a detailed explanation of the elements critical to the significant parts of the research. In relation to the broad aim of the thesis as contributing to part of the quality picture of degree courses, the approach is best classified as falling within the functional paradigm.

3.1 Methodology Design Overview

In order to answer the key research questions of this thesis it was necessary to consider an approach that would enable the analysis of the statements of individual subject learning objectives or learning outcomes in a straight-forward manner. The data to be assembled was extracted from University public documents containing course and subject descriptions from several universities where degrees in Information Technology, Computer

Science, or Software Engineering were offered. To achieve a valid research outcome, the descriptors given by the universities was taken from the published web pages, and then submitted to the analysis process for evaluation and subsequent analysis.

In the preliminary study stage, the data collection took place in a known environment where the researcher had ready access to the subject coordinators to corroborate the evaluation of behavioural objective statements. Once the interpretive learning had been completed in this known environment, the broader research phase to examine the course descriptions from other universities was able to be conducted.

Following the data collection, and using similar approaches to the Danish study, the data were analysed across whole degree programs. The construction of a course metric was devised, and trialled on the 'known' data. The proof of concept and formalisation of appropriate metrics were completed, and then the formulaic model was tested.

While the established metrics were interesting, they would only become valuable if the techniques were applicable across a wider range of data sets. Accordingly additional data sets were constructed from the web pages of the selected Universities. Following the gathering of a comprehensive set of data, the relevant data analyses were undertaken, and conclusions about the proposed metrics were made.

3.2 Design Elements

The initial data collection involved editing published course materials to extract the various subject statements of learning objectives, along with an organisation chart to describe the structure of the degree programs considered. A set of procedures was established to enable the conversion of the objective statement collection into a database structure that allowed for the classification of the statements by either or both the revised Bloom and SOLO taxonomies.

It was found that the structure of degree programs have varying degrees of complexity from university to university but generally fall into different classes of subjects within the relevant degree program. In each year level, there is some mix of compulsory subjects, usually labelled as **CORE** subjects, choices from a specific restricted list, often labelled as **SELECTIVE** subjects, and other more general subjects from basically an open list that are labelled **ELECTIVE** subjects. In order to profile the degree programs effectively, it was necessary to construct a mapping algorithm and establish some rules to be applied in the quantification. Accordingly the approach taken was:

- all of the core and selective subjects were analysed individually;
- the contribution of the selective subjects was limited to the extent of selectives required in the degree program, and a weighted average score was included in the overall degree score calculation;
- as the general elective subjects were too numerous to uniquely identify, a weighted average of the core subject score for that year level was used. The rationale for this decision was that the elective subjects should be approximately equal in terms of workload and cognitive skills that were required of the core subjects in the discipline area at each year level of the degree program.

Looking at degree structure rules, typical examples may include statements such as "...five of the following subjects from list A, plus two of those in list B, plus one general elective ..." in each year level of the degree program.

3.3 Experimental Design

3.3.1 Stage One

The initial stage required a comprehensive review of relevant literature in the areas of learning theory, educational taxonomies, and applications of these to the field of Computer Science and Information Technology teaching and learning at University level. Although there are many research studies that have been conducted over the last fifty years, it was decided that only those that had relative recency should be used, given the dynamic nature of the discipline area. The exception to this underlying premise was in the case of any seminal works that were found, and which clearly needed to be mentioned because of their significance to the research study of this thesis.

3.3.2 Stage Two

Having considered the previous work of Brabrand and Dahl (Chapter 2, Section 2.3.1), a modified form of this study was undertaken to validate the approach in the Australian environment. The similarity included using the same approach to determine subject scores under the SOLO taxonomy to enable a consistency for comparison purposes, and the extension was to repeat the experiment using the Revised Bloom taxonomy as well. The use of both taxonomies enabled a comparative parallel experiment to occur on a known data set base. The major differences from the Brabrand-Dahl study were that the score distributions across the individual subjects enabled profiles to be created for year levels, and overall degree programs, which was only briefly mentioned in the Danish study, but were considered to be important in the Australian context.

The outcome of this important stage was the establishment of the quantitative measures used in conjunction with course rules for a degree program, and the documentation of the techniques used for the year-level and overall degree profile creation. The full discussion of these elements are found in Chapter 4.

3.3.3 Stage Three

The experimental techniques used in the previous stage were reviewed and analysed in order to create a formal specification for the metric calculation. Once specified formally, the method was tested by applying the method to several known courses. In particular the degrees of Bachelor of Information Technology, the Bachelor of Computer Science, and the Bachelor of Engineering (Software) at Flinders University were used to validate the mathematics of the method and the resultant C-Index scores.

The analysis of the results of this stage of the experiment effectively constitutes a proof of concept for the theory proposed, and is discussed in detail in Chapter 5.

3.3.4 Stage Four

With the accumulated data from the previous stage, opportunities for additional meaningful analysis were pursued. Key motivating questions at this point were to consider whether there were interesting observations able to be made from the data analysis and to determine ways of interpreting those results.

3.3.5 Stage Five

Further studies of comparable degree programs in Computer Science and Information Technology were examined to determine the effectiveness of the metric and distribution results as a baseline comparison tool. Under the parameters outlined earlier, the degree programs that were eligible for consideration were those which had a clearly defined set of course structure rules and availability of the individual subject learning objectives on the public pages of the university web-site. While a number of Australian University sites were reviewed, the selected courses for comparison were the Bachelor of Information Technology at Swinburne University, the Bachelor of Information Technology at the University of Queensland, and the Bachelor of Information Technology at the University of Newcastle.

3.3.6 Stage Six

Chapter 8 is devoted to an analysis of the experimental data and discussion of different aspects of the results.

The final part of the experiment, discussed in Chapter 9, was to reflect on the research, documenting the discussions about the experiment and drawing appropriate conclusions from the data.

Chapter 4

From Taxonomy to a Metric

This chapter discusses the quantification of course objectives using the Bloom and SOLO taxonomies, and describes the method used to determine a metric representing the combined value for the topics in a degree course. [1]

The innovative application of previously described methods has given rise to numerical values that assist in the evaluation of degree course specifications at several levels of detail. Initially the analysis of individual subjects that comprise a degree course provides a baseline set of scores. By aggregating those individual scores by year level, a set of year-level scores can be generated. Finally the aggregation of the year-level scores produces an overall score, or, as proposed in this thesis, a new value labelled the C-Index.

4.1 Context Introduction

For the purposes of this discussion the focus will be constrained to the undergraduate degree programs of the higher education sector, and particularly to a single degree to

[1] A substantial part of the material in this chapter (and relevant parts of Chapters 1 & 2) was presented in the paper "A Quality Framework to Measure Course Objectives" at the Fourth International Conference on E-Learning, Ostrava, Czech Republic, July 2013, and also in the journal paper "Constructing a Course Profile by Measuring Course Objectives", International Journal of Digital Information and Wireless Communications (IJDIWC), 2013, Vol 3. Issue 4, pp 16-28.

validate the approach. The initial goal was to confirm the validity of the Brabrand-Dahl study as described in Section 2.3.1 in the Australian context using a known degree program to provide a baseline measure, and then explore additional aspects that arise from the data analysis. An undergraduate degree program is normally named a Bachelor of . . . , and the specific example used in this chapter is a Bachelor of Information Technology. In describing the Bachelor of Information Technology, it may be referred to as either the Bachelor of Information Technology **degree** or the Bachelor of Information Technology **course**.

4.2 The Higher Education Sector

Structurally, a degree program comprises a number of specified studies that must be undertaken in an acceptable combination to satisfy the requirements of the particular degree. Typically the studies are organised on a semester basis, and, depending on the institution concerned, the studies may have the same weighting value in each semester, or there may be differences. For each subject, there is a set of aims and objectives that are intended to provide information about the content of the subject and the skills and knowledge that a student should attain. To clarify the use of terms in this thesis, a brief set of interpretation definitions and equivalences is given in Table 4.1 below:

Table 4.1: Terminology Interpretations

Term	**Meaning**	**Alternative**
Course	A complete degree program	Degree, award
Course Rule	Specification for the combination of subjects to be completed in order to satisfactorily complete the course	Degree Regulations, Schedule of Study
Subject	A prescribed study program in a specific discipline area, typically over one semester	Topic, Course
Unit Value	The effective weight of the subject in the student load, typically expressed as a fraction of a full-time year	Course credits, Credit Points, Units
Learning Objective	A student learning objective written in behavioural terms	Learning Outcome

4.3 Relationship between Elements

The teaching component of the higher education enterprise may be viewed as a composite set of the elements just discussed and arranged in an hierarchial order as shown in Table 4.2.

Table 4.2: Degree Hierarchy Structure

Degree Programs	Course 1	Subject 1	Learning Objective 1
			Learning Objective 2
			...
			Learning Objective m_1
		Subject 2	
		...	
		Subject n_1	
	Course 2	Subject 1	
		Subject 2	
		...	
		Subject n_2	
	...		

At the subject level, the subject specification may be thought of as the set of the learning objectives for that subject. Typically these are expressed in behavioural terms and are therefore usually prefaced with a statement such as or similar to "*On successful completion of this subject the student will be able to ...*".

In practice, each degree/course has its own course aims and objectives, which are presumably addressed by one or more of the individual subject learning objectives. These overarching aims and objectives are intended to convey a sense of the overall graduate attributes that should be realised in the successful students, and provide some thematic relevance or intent across the subjects in the course.

University standards require that each subject has an approved assessment and examination scheme, and a fundamental principle of university teaching is that the assessment plan tests the achievement of the subject learning objectives. On the assumption that this principle is valid and applied in every case, it is reasonable to assume that any student who has received a passing grade has met the subject specifications. Of course the reality is that the assessment of students is not quite so simplistic otherwise there would exist just Pass and Fail as the two possible outcomes for students. What students, educators, and potential employers wish to see is a qualifier on the level of pass attained, so we have grading systems that extend beyond the simple Pass/Fail criteria and include additional classifications such as Credit, Distinction, and High Distinction. Some systems allocate grades in the range A to E, or A to F, with similar interpretations being applied to the final grade. Rather than being purely indicators of success, these categories generally show some form of performance index, and may include other factors such as the way in which students have applied themselves to the subject at hand. Typically those students who engage well with the subject will achieve higher grades than those students who minimise their efforts to satisfy the subject requirements. The relative performance of students is used by universities world-wide and accumulated into a statistic known as GPA (Grade Point Average). This statistic is then used for subsequent admissions to other courses or for the award of scholarships.

An interesting question that arises is this:
Is it feasible to construct a meaningful a-priori profile of a degree course based on subject learning objectives?

4.4 Determination of an individual subject profile

In order to achieve a satisfactory course profile, it is necessary to examine the individual subjects that make up the course, and then aggregate the individual subject assessments to create an overall view.

Fortunately there have been several studies undertaken in the field of learning objectives, and two in particular deal with the development of taxonomies for learning objectives in an attempt to provide qualitative approaches to the examination of learning objectives. As discussed in Chapter 2 Section 2.2.2, the Revised Bloom Taxonomy has been a major research platform in various branches of study with respect to student learning. It was found that the tabular form of the Revised Bloom Taxonomy was able to be applied across a range of granularities, from the fine-grained analysis of a module in a larger teaching program, to broader analyses of subject objectives. The application of the Revised Bloom

Taxonomy matrix involves the examination of learning objectives and classifying them into the appropriate cells of the matrix as shown in Table 4.3. Although the classification process is somewhat subjective in nature, by following a consistent approach a workable set of data was able to be obtained.

Table 4.3: Revised Bloom Taxonomy Matrix (also shown as Table 2.1)

Knowledge	**Cognitive Dimension**					
Dimension	Remember	Understand	Apply	Analyse	Evaluate	Create
Factual Knowledge						
Conceptual Knowledge						
Procedural Knowledge						
Metacognitive Knowledge						

The second significant model is that proposed by Biggs in the form of the SOLO Taxonomy (Structure of Observed Learning Outcome), which is described as a "means of classifying learning outcomes in terms of their complexity" and leading to the ability to "assess student work in terms of its quality . . . " (Biggs, 2011). Earlier publications from Biggs (1979), which refers to an even earlier study by Collis and Biggs, outlines the 5 level structure of the SOLO Taxonomy and discusses the intent and interpretation of each of the 5 levels, also explained in detail in Chapter 2, Section 2.2.3:

1. Pre-Structural
2. Uni-Structural
3. Multi-Structural
4. Relational
5. Extended Abstract

The application of the SOLO Taxonomy to the assessment of learning outcomes (objectives) involves the review of the objectives in terms of the functionality expected at the various levels. In particular, there are typical verbs associated with each level that are likely to appear in statements of learning objectives.

The Brabrand-Dahl study discussed in Chapter 2, Section 2.3.1 described a consistent method to provide a quantitative value conversion from the qualitative base of the taxonomy structure for the 550 syllabi from the science faculties at two universities in Denmark (Brabrand and Dahl, 2007). The analysis of syllabi learning objectives in this study resulted in a list of typical verbs associated with the SOLO Taxonomy, and identified levels 2 and 3 as providing mostly quantitative outcomes and levels 4 and 5 as being more qualitative in nature, as shown in Table 4.4 (and also Table 2.2). The mapping of learning objective statement to a value was then given by the level number that the verb(s) in the objective most closely matched.

Table 4.4: Prototypical Verbs According to the SOLO Taxonomy (Brabrand and Dahl, 2007)

Quantitative		**Qualitative**	
SOLO 2 Uni-structural	**SOLO 3** Multi-structural	**SOLO 4** Relational	**SOLO 5** Extended Abstract
Paraphrase	Combine	Analyse	Theorize
Define	Classify	Compare	Generalize
Identify	Structure	Contrast	Hypothesize
Count	Describe	Integrate	Predict
Name	Enumerate	Relate	Judge
Recite	List	Explain causes	Reflect
Follow (simple)	Do algorithm	Apply Theory	Transfer Theory
instructions	Apply method	(to its domain)	(to new domain)

The method employed by Brabrand and Dahl enabled a relative measure of competencies to be established across the courses in the science faculties in the universities concerned based on the stated learning objectives. The implied equivalence of learning activities within the same SOLO level, and the equal distance assumption between levels were significant pragmatic necessities to generate the type of result gained in this research.

This is somewhat different from the revised Bloom Taxonomy which differentiates knowledge types within cognitive levels, but an interesting question arises to consider whether similar approaches can be used with the revised Bloom Taxonomy as a classification and metric determination tool. Under the equal distance assumption proposed in Brabrand and Dahl, the cognitive levels within any one knowledge dimension should change by an equal amount. Similarly, a constant distance value between knowledge dimension levels should apply within any one cognitive dimension. Accordingly, using an integral unit value a score value table can be constructed as in Table 4.5. However, this in itself creates an issue – should the scoring table be constructed using column-first preference as shown in this example or should it be constructed using row-first preference, or is there some other pattern of scoring that should be used in this circumstance?

Table 4.5: Revised Bloom Ranking Schedule

Knowledge Dimension	**Cognitive Dimension**					
	Remember	Understand	Apply	Analyse	Evaluate	Create
Factual Knowledge	1	5	9	13	17	21
Conceptual Knowledge	2	6	10	14	18	22
Procedural Knowledge	3	7	11	15	19	23
Metacognitive Knowledge	4	8	12	16	20	24

Given that the revised Bloom Taxonomy matrix contains 24 cells, the resultant scale will be in the range 1 to 24. In pure numeric terms the scores obtained using this scale will be vastly different from those using the SOLO scale where the range is between 2 and 5. However it is worth examining whether meaningful outcomes are obtained using the two taxonomies.

Given that behavioural objectives are written as statements of intended student behaviours and learning outcomes, which is about the cognitive skills rather than the subject content, and acknowledging that subject content should become more in-depth as a student progresses through their studies, it is reasonable to remove the depth of knowledge factor in determining a profile that examines the cognate skills. Since the knowledge dimension addresses the nature of content within a subject, the comparison is not really comparing like with like by ranking against the SOLO scores. Therefore, to be more reflective of a properly constructed test to compare similar items, namely the cognate skills specified by learning objectives, an adjusted scale based purely on the cognitive dimension by collapsing the knowledge dimension to a single integral value resulted in a scoring range between 1 and 6, where 1 was assigned to Remember and 6 was assigned to Create. This is consistent with the approach adopted in Oliver et al. (2004) where the Bloom scale was used to examine several programming subjects within an IT degree at an Australian regional University, also using the integral values 1 to 6 in order to arrive at an overall subject score.

With two possible measuring instruments available, the question of how to determine an individual subject metric must now be answered. When reviewing syllabus learning objectives it becomes clear that many are framed in compound terms that is to "do x and do y", or to "understand x, y, and z". The evaluation of compound objective statements can be resolved by one of three methods, namely:

i. to expand the compound statements into multiple simple statements, which in many instances would create a much longer list of objectives. The potential problem with this approach is that an objective of single intent but expressed in compound form would provide a doubling or tripling of scores, thus inflating the score value of the objective.
ii. to evaluate the compound statement and average the individual parts that would be the simple statements under the expansion approach. In this method, the inflationary problem of the first method is overcome and it gives a score within the scaling range for the specific objective. This is the method that was adopted by Brabrand and Dahl.
iii. to use the maximum classification value obtained by inspecting the statement of the learning objective. While simplistic in nature, this method tends to err on the side of generosity when evaluating compound objective statements.

For consistency and comparison purposes it was decided in this research to adopt the same approach as Brabrand and Dahl and use the 'double-weight averaging scheme'.

4.5 Methodology Used

In this study the syllabi for a degree in Information Technology were examined and rated in conjunction with the individual subject coordinators using both the SOLO Taxonomy and the revised Bloom Taxonomy scales. Consistent with the Brabrand-Dahl method, a score in the range 2 to 5 was assigned for each objective under the SOLO ranking, and a score in the range 1 to 6 was assigned for each objective under the revised Bloom ranking. The average score for each objective was calculated when the objective was expressed in compound terms, and then an overall average was calculated for each subject. The relative weight of the subject is given in terms of its unit value, so this weighting was applied to the subject score to arrive at the year level aggregate.

The scores obtained were then grouped by the year level of the course to consider whether there were year level differences, and finally a score for the degree program was calculated.

In the particular degree program examined, there were three classes of subjects, the **Core** subjects which were compulsory, **Selective** subjects where students have a narrow choice from a limited list of subjects, and **Elective** subjects where students may choose from a broad range of subjects. A total of 20 syllabi statements were examined in this degree course to provide the data for the core and selective subjects. To effectively deal with the mix of subject types, the following rules were applied:

a) The compulsory core subjects were evaluated as distinct entries;
b) The selective subjects were evaluated individually but the number of required selective subjects were included as cumulative average values. That is, where the course rule made a statement such as “include 2 of the following 5 subjects ...”, then the average score for the 5 subjects was calculated and weighted by the specified number of selectives required;
c) The elective subjects needed for each year level were included as the average of the core subjects for that year level.

4.5.1 Worked Example

To demonstrate the application of this approach, the Bachelor of Information Technology degree at Flinders University was studied. This is a normal 3-year degree program, typical of many similar degrees offered at universities in Australia and elsewhere. The degree structure requires a total of 108 units, based on 36 units per year.

For this degree, the course rule specifies (Flinders University, 2012c):

Core - Year 1 topics

```
36 units comprising:
 COMP1001  Fundamentals of Computing  (4.5 units)
 COMP1101  Information and Communications Technology 1A  (4.5 units)
 COMP1102  Computer Programming 1  (4.5 units)
 COMP1111  Information Technology Applications*  (4.5 units)
 COMP1401  Professional Skills in Computing**  (4.5 units)
 STAT1412  Data Analysis Laboratory*  (4.5 units)

Plus 9 units of elective topics*** from across the University where entry
requirements are met.
```

Core - Year 2 topics

```
36 units comprising:
 COMP2731  Software Engineering 1  (4.5 units)
 COMP2741  Application Development  (4.5 units)
 COMP2761  Database and Conceptual Modelling  (4.5 units)
 COMP2772  Web-Based Systems Development  (4.5 units)
 ENGR2792  Software Engineering 2  (4.5 units)

Plus 9 units of elective topics from across the University where entry
requirements are met.

Plus one of:

 BUSN3027  E-Business  (4.5 units)
 COMP2762  Operating Systems  (4.5 units)
```

Core - Year 3 topics

```
36 units comprising:
 COMP3721  Enterprise Information Security  (4.5 units)
 COMP3732  Enterprise Systems  (4.5 units)
 COMP3751  Interactive Computer Systems  (4.5 units)
 COMP3771  Advanced Database  (4.5 units)
 ENGR3704  Project Management for Engineering and Science  (4.5 units)

Plus 4.5 units of elective topics from across the University where entry
requirements are met.

Plus either:

 COMP3782  Information Technology Project  (4.5 units) AND
 CSEM upper-level topic# (4.5 units)

OR

 COMP3792  Information Technology Practicum  (9 units)
```

Note: *the variously asterisked subjects (topics) have supplemental information about substitution with alternative similar topics if approved by the course coordinator.*

As can be seen in this example, the year 1 studies comprise core topics and elective topics, years 2 and 3 comprise core topics, selective topics, and elective topics. In order to construct an overall degree profile, the individual subjects (topics) need to be examined and classified separately. In the example below, the behavioural objectives or learning outcomes are all effectively compound statements other than the last.

When looking at individual topics, in this case COMP1001, the stated learning outcomes are:

At the completion of the topic, students are expected to be able to:

1. *Be familiar with the fundamentals, nature and limitations of computation*

2. *Be familiar with standard representations of data and the translation to and from standard forms*

3. *Be aware of the structure of information systems and their use*

4. *Understand the social and ethical implications of the application of information systems*

5. *Construct simple imperative programs*

Discussion with the subject coordinator to confirm the interpretation of these objectives resulted in the following assessments on the SOLO and Bloom scales as shown in Table 4.6. In order to conform to the double-weight averaging scheme, each of the component parts of an objective was discussed and analysed, then listed in a schedule under each of the SOLO and Bloom scales. The integer value associated with the level for each of the component parts was summed, and then averaged to give a score for that objective. The overall score for the subject (*topic* in the language of Flinders University) was calculated by averaging the individual objective scores.

Table 4.6: COMP1001 Classification

Objective	SOLO Analysis	SOLO Score	Bloom Analysis	Bloom Score	Bloom Adjusted	Adjusted Score
1	S3, S3, S4	3.3	B5, B6	5.5	B2, B3	2.5
2	S3, S3, S3	3.0	B5, B6	5.5	B2, B3	2.5
3	S3, S3	3.0	B5, B6	5.5	B2	2.0
4	S3, S3	3.0	B9, B10	9.5	B2, B4, B5	3.7
5	S4	4.0	B23	23.0	B6	6.0
Average		**3.26**		**9.80**		**3.34**

In the Danish study, a substantial effort of rewriting the intended learning outcomes had been undertaken prior to the study being done (Brabrand and Dahl, 2007). This was not the case in the test environment and therefore some interpretive efforts were made to translate the statements into what they would most likely become in the language of the SOLO Taxonomy. The SOLO Distribution for the subject in this example becomes $(0.67 + 1 + 1 + 1)/5 = 73\%$ of S3 and $(0.33 + 1)/5 = 27\%$ of S4 using the proportional distribution approach identified in Section 2.3.1, or $9/11 = 82\%$ of S3, and $2/11 = 18\%$

of S4 using the raw count or simplex approach also described in the same section. While these values appear to differ significantly at the individual subject level, the research looked at the whole degree course, and it was decided to investigate whether there was a substantial difference overall between the simplex method and the proportional method in determining the SOLO Distribution.

This process was repeated for each of the core and selective subjects in the course.

4.6 Results

In the Flinders University environment, a full-time student would normally undertake 36 units of study per year. This gives rise to a weighting factor for individual subjects (topics) as a proportion of a full-time load. Hence 4.5 units has a weighting of 0.125, 9 units has a weighting of 0.25, 18 units has a weighting of 0.5, and other weightings can be calculated similarly as a fraction of the 36 units.

Following the approach described in the previous section, the accumulated data for the 20 specific core and selective subjects of the degree are given in the following table (Table 4.7).

The rows labelled "Core Topic Average" are derived from the Core Topics, and this value is used in the relevant year level elective (highlighted with *).

With further consideration of each year level, a weighted score under each of the three taxonomy scales was determined and the following summary (Table 4.8) of classifications for the subjects was obtained.

4.7 Discussion

In evaluating the behavioural objectives for the subjects in this degree there were several interesting points that were revealed. These are discussed as separate items below.

Table 4.7: BInfoTech Analysis

Subject Code	Weight	SOLO Score	Revised Bloom	Adjusted Bloom
COMP1001	0.125	3.26	9.80	3.34
COMP1101	0.125	3.20	10.03	2.90
COMP1102	0.125	3.50	10.07	4.00
COMP1111	0.125	3.67	10.75	3.56
COMP1401	0.125	3.13	11.67	3.98
STAT1412	0.125	3.67	12.00	3.67
Core Topic Average		3.43	10.72	3.57
Elective Yr 1*	0.250	3.43	10.72	3.57
COMP2731	0.125	3.25	9.75	3.25
COMP2741	0.125	3.63	9.375	3.13
COMP2761	0.125	3.42	10.06	3.75
COMP2772	0.125	3.92	12.50	4.50
ENGR2792	0.125	3.50	11.21	3.71
Core Topic Average		3.54	10.58	3.67
BUSN3027	0.125	3.78	11.78	3.44
COMP2762	0.125	3.63	11.00	3.25
Selective Average	0.125	3.71	11.39	3.35
Elective Yr 2*	0.250	3.54	10.58	3.67
COMP3721	0.125	4.00	11.62	4.36
COMP3732	0.125	4.13	10.13	4.25
COMP3751	0.125	3.70	14.17	4.17
COMP3771	0.125	3.39	15.25	3.67
ENGR3704	0.125	3.92	9.86	3.92
Core Topic Average		3.83	12.20	4.07
COMP3782	0.125	3.76	11.13	3.79
Upper level topic	0.125	3.83	12.20	4.07
COMP3792	0.250	4.08	9.50	4.00
Selective Average	0.250	3.94	10.58	3.93
Elective Yr 3*	0.125	3.83	12.20	4.07

Table 4.8: SOLO vs Bloom Scores

	Weighted SOLO Total	Weighted Bloom Total	Adjusted Bloom Scores
First Year	3.43	10.72	3.57
Second Year	3.56	10.68	3.63
Third Year	3.85	11.80	4.04
Degree Total	10.84	33.20	11.24
Degree Average	3.62	11.07	3.75

Appropriateness of a Taxonomic approach. Two distinct taxonomic approaches have been used in this study, namely the revised Bloom Taxonomy and the SOLO Taxonomy, as vehicles to investigate the learning outcomes or objectives of the subjects in a particular course. A similar question has been raised with respect to the use of Bloom's Taxonomy in the Computer Science domain (Johnson and Fuller, 2006). In that work, Johnson and Fuller proposed a slightly different structure to cater for "the idea that application is the aim of computer science teaching". No firm resolution was given, but the issue of whether the Bloom Taxonomy is appropriately suitable in the Computer Science arena was raised, and has been discussed in Chapter 2, Section 2.3.5. Other works have proposed that the revised Bloom Taxonomy was useful in Computer Science teaching, particularly where multiple staff members were involved in the subject (Thompson et al., 2008).

Further developments and research into a computer-science specific learning taxonomy have been undertaken by various researchers with a newer model addressing the perceived deficiencies in both the Bloom and SOLO taxonomies (Fuller et al., 2007; Bower, 2008). These research activities in concert with the Brabrand and Dahl efforts highlight and support that a taxonomic approach is relevant, even though the taxonomic tools currently available may not yet be the best fit, or may need some refinement for domain areas such as Computer Science. The experience gained in this study suggests that it is more an issue of interpretation of the standard descriptors used in the classifications rather than changing the classification framework to suit the domain, otherwise one spawns a whole new set of taxonomies for various discipline domains, each of which then need to be validated.

Current written form of the statements of behavioural objectives. The standard and consistency of the current behavioural objective statements was quite variable for this course. A significant number were quite vague and therefore difficult to classify appropriately. However, the vaguely expressed objectives were more easily classified using the Bloom Taxonomy than with the SOLO Taxonomy. The challenge for educational institutions is to ensure that the stated learning objectives accurately reflect what is being taught, what is being expected of students, and subsequently what is proposed to be learned in the subjects of the course.

Language-rich bias. The subjects which have a stronger focus on language elements such as report writing and critiquing of subject materials tend to score more highly in both taxonomies. Some subject areas such as computer programming may involve quite complex levels of problem solving and formulation of creative approaches to resolve issues, but these elements were not explicitly stated in the subject learning objectives. Discussions with the subject coordinators highlighted that their impressions of some of the tasks required of the students involved the higher order taxonomy classifications, yet the subject learning objectives did not adequately express this.

Interpretation opportunities. The two dimensional nature of the Bloom taxonomy makes subsequent investigation of comparative subsets somewhat more difficult computationally. On the other hand, the SOLO approach allows for more internal analysis to be undertaken with relative ease, as can be seen in the simple scatter-plot of the subject scores (Figure 4.1).

The SOLO Distribution for the course was calculated using both the *proportional* method and the *simplex* method to determine whether there was either a close similarity or a substantial difference. The proportional method results are shown in Table 4.9 and Figure 4.2 while the simplex method results are shown in Table 4.10 and Figure 4.3. Comparing the two sets of data, there was a small difference in some of the component values, but the overall view is very similar.

The nature of the simplex method (see Section 2.3.1) is that the simple counting method

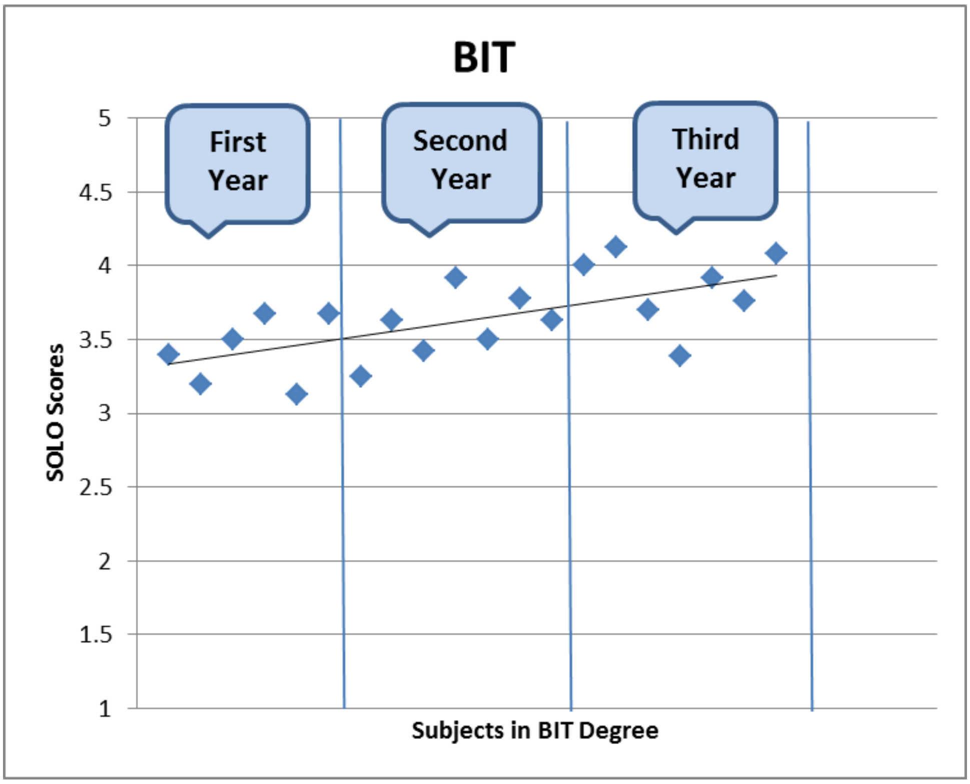

Figure 4.1: Scatter Plot of SOLO Levels in the Information Technology Degree

generates a higher number of SOLO scores when the learning outcome is expressed as a compound statement. The proportional method (also in Section 2.3.1) is a little more computationally complex, but it does more accurately reflect the assumption stated in the later paper by Brabrand and Dahl (2009) that each learning outcome competence should have a proportionate weight in the subject. When the data accumulated by the simplex method is weighted according to the subject contribution to the course, the overall impact may not be significantly different from the proportional method. As there is currently no agreed 'best' way to determine the SOLO Distribution, and the goal of this research is to review overall patterns, the simplex method was decided to be used in subsequent analyses.

Table 4.9: SOLO Summary of Subjects by Year Level (Proportional)

	Solo2	Solo3	Solo4	Solo5
First Year	5%	53%	39%	4%
Second Year	5%	39%	47%	8%
Third Year	3%	25%	49%	11%
Overall	4%	39%	49%	8%

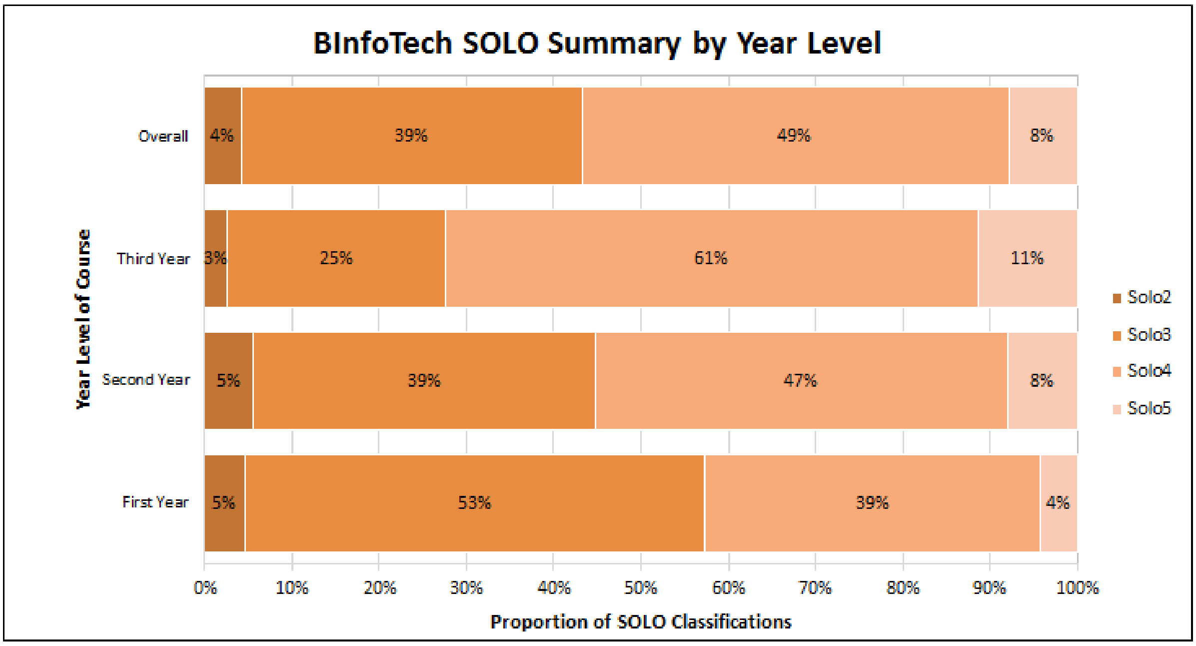

Figure 4.2: Relative SOLO Levels in the Information Technology Degree (Proportional)

Table 4.10: SOLO Summary of Subjects by Year Level (Simplex)

	Solo2	Solo3	Solo4	Solo5
First Year	5%	49%	39%	7%
Second Year	7%	41%	42%	10%
Third Year	4%	25%	55%	16%
Overall	5%	36%	47%	12%

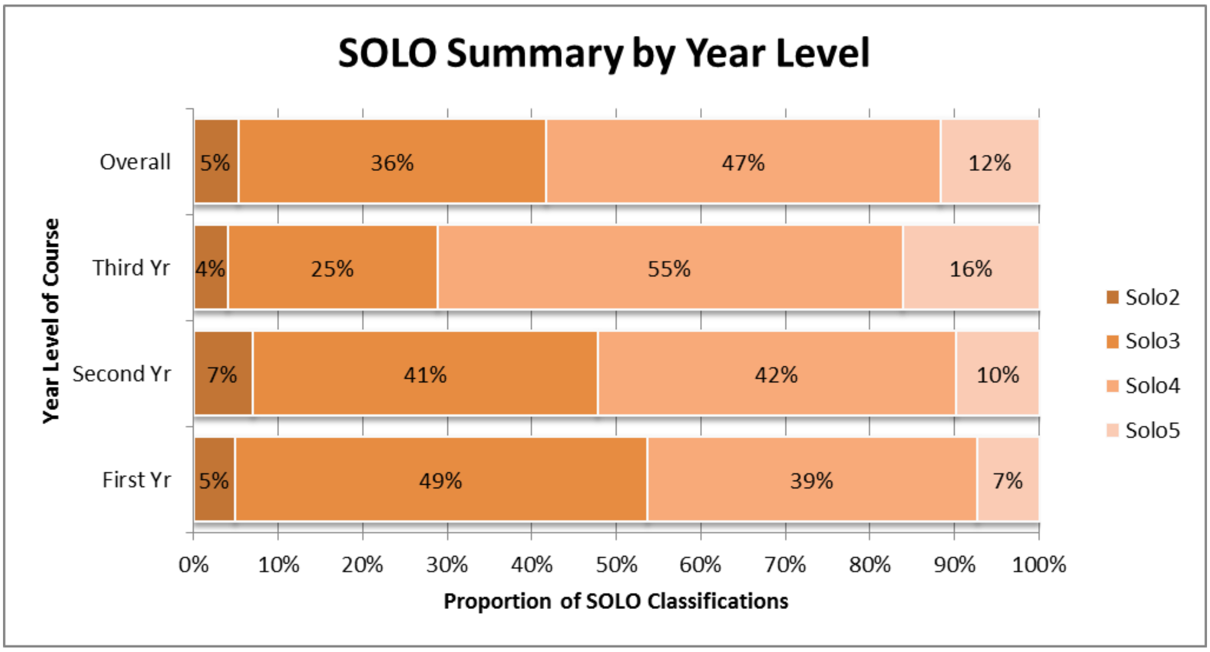

Figure 4.3: Relative SOLO Levels in the Information Technology Degree (Simplex)

Using the adjusted Bloom scale to focus only on the cognitive dimension, a comparable set of data was obtained with the equivalent statistic listed as the Adjusted Bloom Score in Table 4.8, and the detailed breakdown is shown in Table 4.11 with the associated graphic in Figure 4.4. It would appear that the adjusted one-dimensional Bloom approach could be applied as readily as the SOLO approach.

Meaningful result? The process applied in this research project has demonstrated that a statistic can be determined for a particular course of study. At this point of the research the meaningfulness of that statistic is yet to be determined, either with the SOLO Taxonomy or the Bloom Taxonomy. Subsequent work is required to provide comparative data and overall calibration for this metric. What has been revealed is that the closer analysis of the subject behavioural objectives for this degree across year levels does match

Table 4.11: Bloom Summary of Subjects by Year Level

	Bloom1	Bloom2	Bloom3	Bloom4	Bloom5	Bloom6
First Year	0%	21%	31%	22%	16%	10%
Second Year	6%	14%	26%	20%	22%	12%
Third Year	0%	11%	18%	37%	20%	13%
Overall	2%	15%	25%	26%	19%	12%

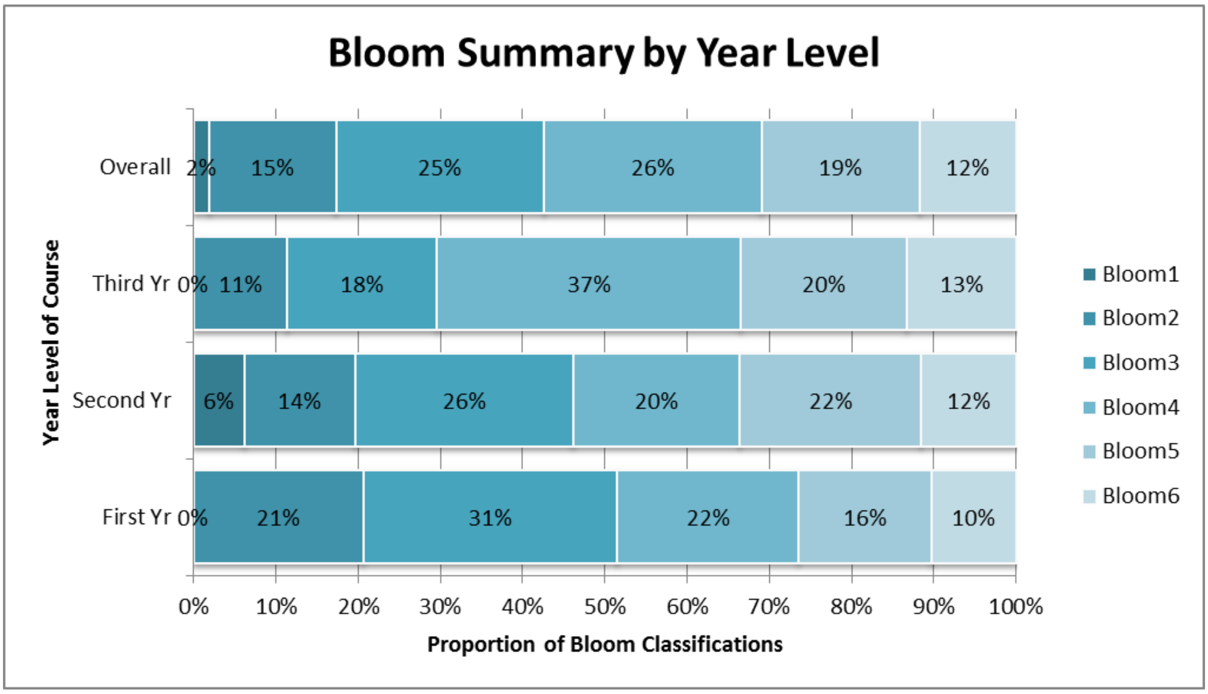

Figure 4.4: Relative Bloom Levels in the Information Technology Degree

the naïve expectations - namely that as one progresses through the degree studies from first year to second year to third year there is a shift of emphasis from the lower more functional or quantitative SOLO levels to the more sophisticated qualitative levels. The data in Table 4.8 demonstrates an increasing "SOLO Average" through the year levels, and provides a total of 10.84 for the course, or an average of 3.62 if one wanted to arrive at a single indicator figure within the scaling range. This finding is consistent with the findings of a separate study by Brabrand and Dahl (2009) that explored the use of the SOLO Taxonomy to examine competence progression from undergraduate to graduate level studies. An almost identical result was obtained when using the Bloom Taxonomy, adjusted to consider only the cognitive elements. The fact of being able to establish a metric suggests that there is an opportunity to further develop a set of expanded tools that may be useful in the quality and benchmarking domain for degree courses.

4.8 Conclusions

Traditionally it has been the case that teachers, academics and educators generally have rejected the notions of measurement and accountability in respect of the teaching process, even though they subject their students to exactly those elements. Many previous attempts at measurement of the education sector have been derived from administrators attempting to apply accounting principles which overlook many of the peculiarities of the education sector and invariably fail or invoke feelings of angst and hostility towards their implementation. This research has introduced a concept for a metric that is systemic in nature, measuring attributes of the 'system' via the individual subjects that comprise a course of study, and ultimately generates a measure for the overall course of study. Being a pre-activity indicator it is independent of the approach taken by the teaching team and the peculiarities of the particular cohort of students. Individual academics have control over the attributes being measured in that they are the ones who write the behavioural objectives for their subjects and therefore contribute to the specifications for the subjects under their control as they have always done. How this metric is used within academia and the meanings and interpretations associated with it remain to be seen in future works.

One of the major findings in this part of the research was that the standard of written behavioural objectives in the course examined was somewhat inconsistent. Some of the subjects had well-formed statements and made it clear about what was intended in the subject. Others were vague and provided minimal useful information about the subject content or intended student expectations. It was noted in the work of Gluga et al. (2012b) that a specific training program had been created for computer science educators on applying Bloom to the classification of programming questions. From an institutional perspective, a recommendation would be to tighten the statements of behavioural objectives to improve the subject specifications. With better and more consistent statements of objectives the key stakeholders who make use of those subject specifications will be better informed, and more reliable data based on those stated objectives may be obtained.

This research has demonstrated that it is feasible to construct a course profile for a degree using either the SOLO Taxonomy or the amended Bloom Taxonomy to evaluate the subject learning objectives for the course. Although the numeric values given in Table 4.8 are potentially useful indicators, the distribution of expected learning activity across year levels has proven to be much more interesting and informative when displayed either in tabular form (Table 4.10, Table 4.11), or graphically as in Figure 4.3 and Figure 4.4. When used in conjunction with other examination tools and inspection of output artefacts, the profile of expected learning activities in the course may be a valuable instrument that finds application in course comparisons, benchmarking, and the evaluation of course quality.

The language-rich subjects tended to score higher in the methodology used in this research. Although this may be a slight impediment to the technically oriented courses, the overall influence of the language-rich subjects tends to be overshadowed by the inherent ratio of technical to less/non-technical subjects in structuring technically oriented degree programs.

The next stage associated with this research is to expand the data sets involved and make decisions about the relative ease of assessing courses on the basis of a taxonomic analysis of subject learning objectives, to investigate the ways to better interpret the results obtained, and to assess the usefulness of the approach in course benchmarking.

Chapter 5

Formalising the Metric

The approach given in the previous chapter established a method for documenting a course profile for a university degree course. The formalisation of the technique and application of that method to other degree courses in conjunction with the comparative analysis aspects between three known sets of data is the subject of this chapter. [1]

The establishment of a standardised technique is a critical element to the widespread application and use of a statistical measure. Accordingly, it is necessary to express the calculation of this metric in a formal mathematical manner, as described in the subsequent sections.

5.1 Defining the C-Index

It is proposed that the C-Index is a measure of the cognitive learning level prescribed by a course of study based on the stated learning objectives. In this research it has been based on the SOLO taxonomy of learning objectives, but could be equally well determined from an analysis based on the revised Bloom taxonomy, as the similarity between the two taxonomies has been demonstrated in the earlier work in Chapter 4. Although the choice

[1] A substantial amount of the material in this chapter (and relevant parts of Chapters 1, 2 & 4) was presented in the paper "Course Quality Starts with Knowing Its C-Index" at the InSITE 2014 Conference, University of Wollongong, New South Wales, Australia, July 2014.

between using the SOLO taxonomy or the revised Bloom taxonomy within this definition appears to be arbitrary, there appears to be more support for using the SOLO taxonomy in this manner, partly because the classification of learning objectives and tasks is more clearly visible within the SOLO framework.

It has previously been shown that a statistical measure labelled as the SOLO score can be determined for a particular subject (Brabrand and Dahl, 2007). By extension across a complete degree program, as described in Chapter 4, the weighted mean of the SOLO scores will generate a standardised score for the individual year levels within a course, and the accumulation of these weighted means gives rise to the C-Index for the course. In order to accommodate different course lengths – some three-year degrees, some four-year degrees, and perhaps longer – it is necessary to standardise the overall value obtained. The most straight-forward approach then is to use a measure of central tendency; calculating the mean of the various year-level scores determines the overall C-Index for a course.

Given that the C-Index is a statistic based on the cognitive learning levels in a degree course, the specific interpretation should be deemed to be a Course-Index.

Formally, the C-Index is expressed as:

$$C = \frac{1}{y} \sum_{i=1}^{y} \sum_{j=1}^{n_i} w_{ij} S_{ij} \quad \text{(5.1)}$$

where

S_{ij} = SOLO score for subject j in year i of the course;

w_{ij} = the weight of the subject expressed as the fraction of a full-time year of study;

n_i = number of subjects included in the full-time year of study for year i;

y = year levels in the course of study.

Further explained, the statistic is calculated by determining the SOLO score for each subject prescribed in the course schedule for the particular degree program, and then aggregated according to the subject weight in the given year of the course. The resultant figure will be a weighted average SOLO score or p-index for the particular year level of the course. The term p-index is chosen to represent a partial assessment of the overall course. By repeating over each of the other year levels, a set of p-index values is obtained and then the mean of these year-level indices returns the course C-Index.

Hence, an alternative expression for the C-Index could be given as:

$$C = \frac{1}{y}\sum_{i=1}^{y} p_i \qquad (5.2)$$

where

p$_i$ = year-level p-index = $\sum_{j=1}^{n_i} w_{ij}S_{ij}$, for the subjects in that year-level.

There are of course other definitions of C-Index, one being in the realm of association theory (Garcia, 2012), and another being a software tool for indexing books, journals and other textual material (Indexing Research, 2012). However, this being a quite different domain area, the author sees no real terminology conflict.

5.2 Applying the C-Index

The data for this part of the study was obtained in the same manner as the initial data, but restricted to just applying the SOLO scale rather than both the SOLO and Bloom scales, and assembled in the same way using the same assumptions with respect to the core, selective, and elective subjects. In order to make a controlled comparison, three degree courses from the same School at Flinders University were chosen, and their relevant course rules were applied accordingly. The specific courses were

1. Bachelor of Information Technology (Flinders University, 2012c)

2. Bachelor of Computer Science (Flinders University, 2012a)

3. Bachelor of Engineering (Software) (Flinders University, 2012b)

There was some overlap of data between the Bachelor of Computer Science and the Bachelor of Information Technology degree, with both of these particular degrees having some subjects in common, and to a lesser extent the Bachelor of Engineering (Software). The Bachelor of Engineering (Software) was included for two reasons – the first being that it is related to the other two degrees, but secondly it is a four-year degree program to compare with three-year degree programs. There were 20 subjects that were analysed in the Bachelor of Information Technology, 28 subjects in the Bachelor of Computer

Science, and 44 subjects in the Bachelor of Engineering (Software). In the Bachelor of Information Technology there were a relatively small number of selective and elective subjects, which meant that only 20 subjects were evaluated. With the greater number of selective subjects in the Bachelor of Computer Science, a larger number of subjects were required to be evaluated. The reason for there being a substantially higher number of subjects evaluated in the Bachelor of Engineering (Software) is that there are multiple streams from which to choose in that degree. The streams were treated in a similar manner to the selective subjects to arrive at an average score for the streams.

5.2.1 Comparison of Results

Application of the methodology resulted in the following outcomes for the three degrees.

Table 5.1: Course Comparison Scores

Course Year Level	BInfoTech Weighted SOLO Scores	BCompSc Weighted SOLO Scores	BEng(SW) Weighted SOLO Scores
First Year	3.43	3.45	3.55
Second Year	3.56	3.63	3.68
Third Year	3.85	3.77	3.87
Fourth Year	-	-	4.00
Degree Total	10.84	10.85	15.10
C-Index	3.62	3.62	3.78

As can be seen in the accompanying results table (Table 5.1), for each of the degrees there was an increasing progression in the weighted SOLO scores through each year of the study program, which was both expected and reassuring. It was expected in that one would hope that the learning required in each year of a course did become more involved and more demanding and that the statement of learning outcomes accurately reflected this. It was reassuring that the courses examined did display this characteristic.

Interestingly, the BInfoTech and BCompSc C-Index values were the same even though their individual year-level scores were slightly different. A partial explanation for this outcome is that these two courses share a significant number of common subjects in their study program.

The difference in the C-Index for the BEng(SW) degree highlights the impact of a four-year degree compared to a three-year degree, where there is the expectation that the later year subjects will contain more advanced work, and these results support that assertion.

Other statistical measures were explored (see Table 5.2), including the sample standard deviation to consider the spread of the data, and a year-weighted mean. In the accompanying Table 5.2 the value for the year-weighted mean has been calculated using a simple integral value of the year level (1 for first year, 2 for second year, 3 for third year and 4 for fourth year) and then normalised by the sum of the weights. With so few data points the standard deviation is unlikely to reveal any particularly useful information for an individual course at the overall level. Its application is more likely to be relevant when used as a comparative value to compare similar or related courses, or perhaps discipline areas.

Table 5.2: Course Comparison Other Scores

Course Year Level	BInfoTech Weighted SOLO Scores	BCompSc Weighted SOLO Scores	BEng(SW) Weighted SOLO Scores
First Year	3.43	3.45	3.55
Second Year	3.56	3.63	3.68
Third Year	3.85	3.77	3.87
Fourth Year	-	-	4.00
Degree Total	10.84	10.85	15.10
C-Index	3.62	3.62	3.78
Std Deviation	.210	.160	.199
Yr-Weighted Mean	3.69	3.67	3.85

What can be seen from Table 5.2 is that the BInfoTech displays the greatest amount of spread, by virtue of having the largest standard deviation value, and the BCompSc displays the least amount of spread based on the year-level scores.

Under the hypothesis that later year subjects are perceived to be 'more important' than early year subjects, a calculation was performed to determine the weighted mean score based on year level. While this does slightly bias the outcome towards the later year scores, it is debatable as to whether it provides a better view of the course, or whether it is simply an unnecessary complicating factor. It could validly be argued that the material content of later year studies does require higher order cognitive skills at the SOLO 4 and SOLO 5 levels, but this has only become possible because of the formative learning that has occurred in the earlier years of study and the development of learning maturity in the student. As a result the educational rigour may be viewed as being comparable across the year levels of the course as it is more closely matched to what a student can be hoped to achieve at that stage in their educational development. This could be considered as their expected learning potential. Consequently it follows that the year-weighted mean is more likely to be an unnecessary complication on the calculation than being able to provide a more meaningful C-Index value.

Interpreting the C-Index

While the mathematical calculations have returned a quantitative value under the method described, the question that arises is "What does it mean?". When looking at the underlying assumptions upon which the methodology is based, two elements in particular are significant. The first is the *equal distance assumption* proposed in the initial Brabrand-Dahl study (2007), where the progression from one SOLO level to the next was deemed to be of equivalent difficulty. The second is the implied assumption that all of the learning tasks within a particular SOLO level are at least approximately equal.

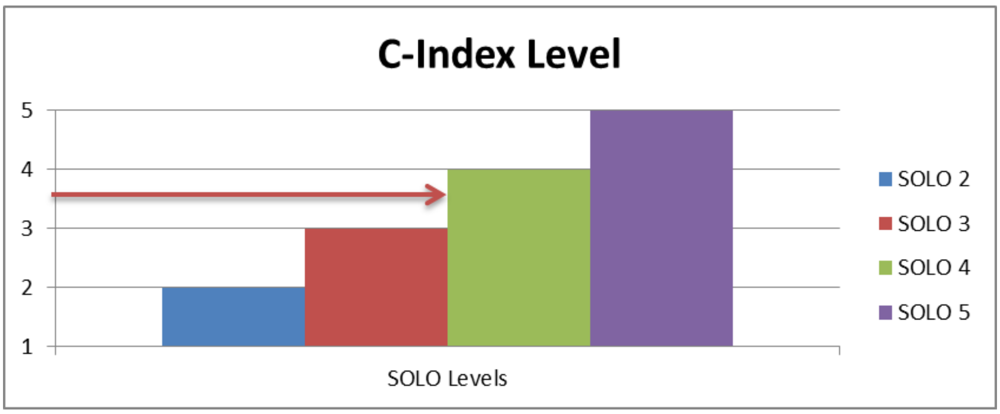

Figure 5.1: C-Index Level of the BInfoTech Degree

Accordingly, the C-Index value of 3.62 for the Bachelor of Information Technology degree in the data set is beyond SOLO Level 3, and is nearing the achievement of SOLO Level

4. The same initial comment applies to the other two degrees in this data set. Knowing that SOLO Level 3 items are classified as multi-structural and includes learning tasks such as classification and application of method, compared with SOLO Level 4 where the tasks are classified as relational and require analysis and application of theory, the scores indicate a substantial learning expectation is required of students in these courses. From the interpretation given by Brabrand and Dahl, this places the overall average of expected learning outcomes well into the qualitative or deep-learning region rather than the more quantitative surface-learning area. Importantly, in the definition of the SOLO Taxonomy levels it is clearly stated that the levels are cumulative (Biggs and Collis, 1982), which means that the learning skills achieved at Level 4 include those from the lower Level 3 (and Level 2).

A corollary of the equal-distance assumption is that the SOLO scales may be considered as a continuous scale that then allows a graphical representation to be given, either at the summary overall level view (see Figure 5.1), or at the more detailed year-level view (see Figure 5.2).

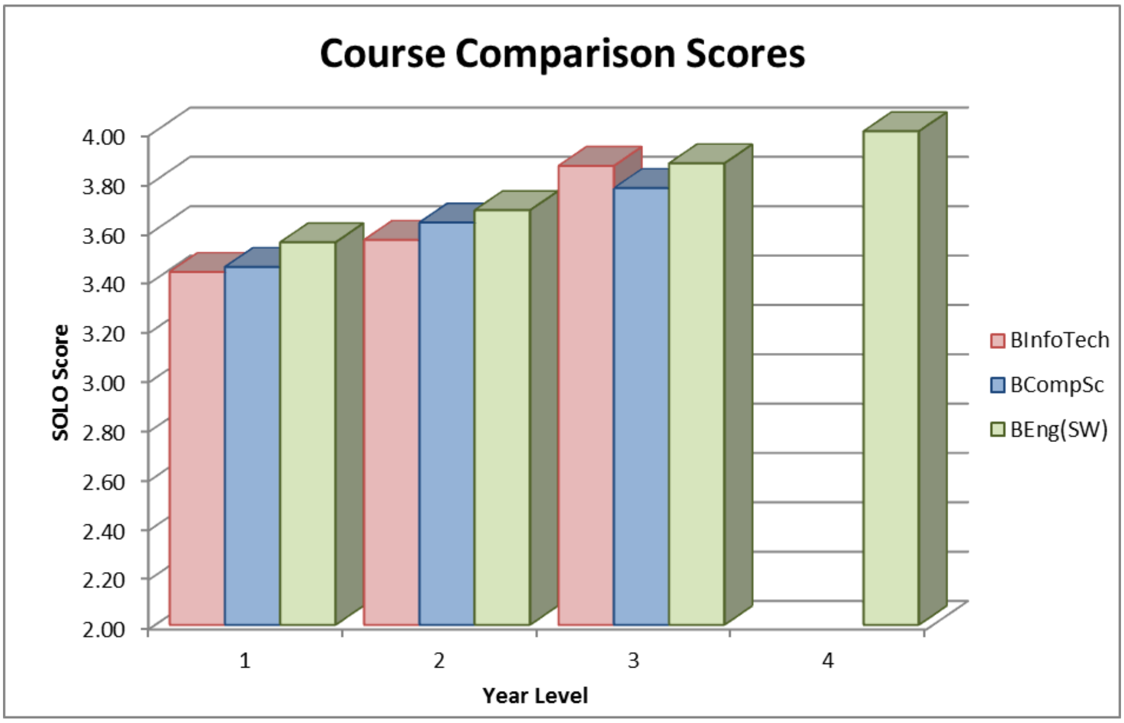

Figure 5.2: Relative Year Level SOLO Scores in the Degree Courses

Another representation for the comparison of degree courses is given in Figure 5.3, which

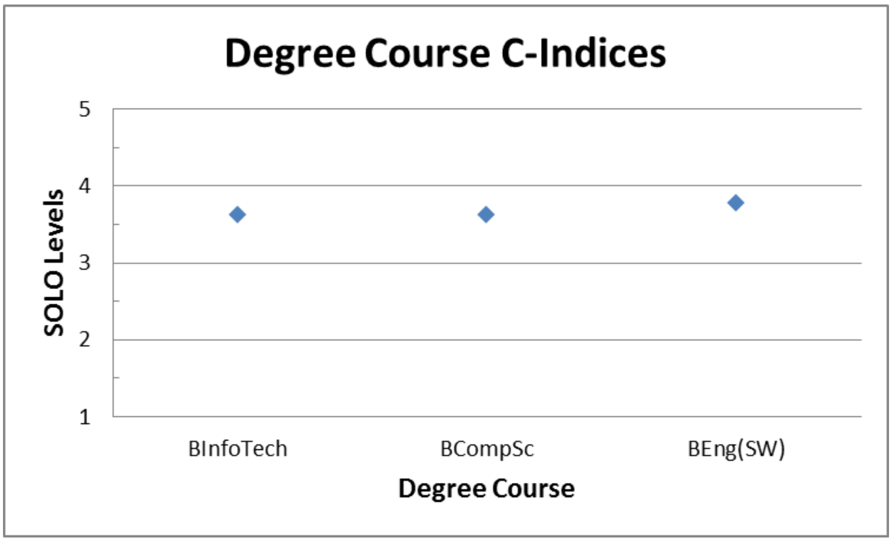

Figure 5.3: Comparative Degree Course Indices

shows a simple plot of the overall C-Index for each distinct degree against the SOLO Score.

When the C-Index is used in conjunction with the SOLO Distribution for a degree (for example the BInfoTech degree discussed in Chapter 4 and displayed in Table 4.10 and Figure 4.3) the value of 3.62 suggests that the course aims to require students to undertake learning near the SOLO Level 4. The Distribution data indicates that approximately 60% of the course is oriented to *deep understanding* (SOLO-4 and SOLO-5) and around 40% of the course is oriented towards *surface learning* (SOLO-2 and SOLO-3) (Brabrand and Dahl, 2007).

This finding is consistent with the findings of the separate study by Brabrand and Dahl (2009) that explored the use of the SOLO taxonomy to examine competence progression from undergraduate to graduate level studies.

The Effect of Averaging

The method used to determine the subject year-level scores and the course C-Index involves a number of repetitions of using averaging techniques. There is some validity in the argument that repeated averaging may throw into question the merit of the statistic obtained as the overall granularity of the data set may become coarser. One of the strong criticisms of the mean as a measure of central tendency is that it is unduly affected by extreme observations (Moore and McCabe, 2003; Selvanathan et al., 2007), but in this case there is no opportunity for extreme values to occur as the scoring range is between 2 and 5. However it is proposed that the resultant value is a guidance number that should be used in conjunction with other factors rather than being taken as a stand-alone value on which to base interpretations and judgements.

Summary

This chapter has contributed two significant elements to the thesis. Firstly the formal statement of the method to calculate the C-Index makes it clear how to combine the analysis of SOLO scores into subject scores, followed by generating a year-level score using the weighting factor of the the contribution of each subject to a student's year of study, and finally averaging the year-level scores over the full degree program to produce the course C-Index. It has been proposed that the C-Index value may be used as an indicator of the level of learning rigour that can be expected for the course in question.

The second contribution was to introduce potential ways to view the initial result of the C-Index calculations as either a simple placement on a linear scale for the overall value, or in a more detailed view of the respective year-level scores, or as a simple scatter-plot when used to compare several degree courses.

In this chapter the data examined was from several degree programs at a single university. The goals were to explore the C-Index calculation on known courses, and to investigate supplementary analytical techniques based on that data. The results obtained have validated the methods used and have clarified the concepts introduced. In the next chapter, the main goals are to explore approaches to using the C-Index calculation methods for application in internal quality control.

Chapter 6

Internal Quality Control

The approach given in the previous chapters established a method for documenting a course profile for a university degree course, and provided a formal mathematical description of the technique used to determine the C-Index for a course. The C-Index calculation involves a number of detailed analyses of individual subjects in order to provide a set of subject SOLO scores which can be aggregated under the course rules to arrive at the final C-Index value. Those intermediate calculations generate a set of data which may potentially enable further analysis to be performed in a meaningful way. Using the three previously examined courses, a comparative analysis of the intermediate data to consider opportunities for application as internal quality control tools is the subject of this chapter.

6.1 C-Index as an Internal Quality Control Tool

It has been proposed that the C-Index value may be used as an indicator of course rigour, or the level of learning expectation within the course. When reviewing the summary values shown in the accompanying table (Table 6.1) it can be seen that the C-Index values for each of the three courses are not too dissimilar, although there appears to be a greater amount of variation in the standard deviations for the three courses. One of the questions that arises is whether each of the courses displays a comparable amount of rigour? A second question that arises is whether the learning expectation at each year level is appropriate for each of the courses? To answer these questions it becomes necessary to drill down to the more detailed levels of the underlying data and investigate the individual subject scores in a comparative manner.

Table 6.1: Course Comparison Scores

Course Year Level	BInfoTech Weighted SOLO Scores	BCompSc Weighted SOLO Scores	BEng(SW) Weighted SOLO Scores
First Year	3.43	3.45	3.55
Second Year	3.56	3.63	3.68
Third Year	3.85	3.77	3.87
Fourth Year	-	-	4.00
Degree Total	10.84	10.85	15.10
C-Index	3.62	3.62	3.78
Std Deviation	.210	.160	.199

6.1.1 Comparing Subject Rigour

According to the glossary of educational terms prepared by the Great Schools Partnership (2014) (available at http://edglossary.org/rigor) "*The term rigor is widely used by educators to describe instruction, schoolwork, learning experiences, and educational expectations that are academically, intellectually, and personally challenging.*" The Southern Cross University (2014) discusses intellectual rigour as part of its set of graduate attributes, emphasising intellectual rigour as " *. . . having clarity in thinking and an ability to think carefully, deeply and with rigour when faced with new knowledge and arguments. . . . It also relates to the ability to analyse and construct knowledge with depth, insight and intellectual maturity.*" If all subjects demanded the same level of rigour, then it would be reasonable to expect each to have the same SOLO score. It is clear in the calculations of SOLO scores that they are not all the same, and therefore one must ask "*how different are they*?" or perhaps more appropriately, "*how close are they*?" in terms of their learning expectations as described by the statements of learning objectives.

When one typically begins looking for statistical support to answer such questions, a usual approach is to make assumptions about the data being normally distributed or approximately normally distributed. In such cases it becomes possible to make use of elements such as the Central Limit Theorem and the properties of normally distributed data including the so-called 'Empirical Rule' or the '68-95-99.7 Rule' where (approximately) 68% of observations lie within one standard deviation, 95% of observations lie within two standard deviations and 99% of observations lie within three standard deviations. (Selvanathan et al., 2007; Moore and McCabe, 2003) However, with small data sets, the approximately normal assumption cannot be made. Selvanathan et al. (2007) suggests that one should work with sample sizes of at least 100 wherever possible in using the Chi-Square test for normality of a data set as smaller data sets usually fail to reject the null hypothesis that the data are normally distributed. Alternatively, using the mean and standard deviation values and not reliant on the assumption of a normal distribution, under Chebyshev's Theorem (Selvanathan et al., 2007), at least 75% of observations should lie within two standard deviations of the mean, and at least 90% of observations should lie within three standard deviations of the mean. Using the data calculated previously and shown (again) in Table 6.1, a re-evaluation of the subject SOLO-scores in the courses based on their closeness to the overall C-Index can be determined. Using the two-standard deviation and three-standard deviation values as control limits above and below the C-Index, a different view of the course data emerges. The results of this analysis is shown for each of the three degrees in Table 6.2.

Table 6.2: SOLO Score Distributions by Course

Subject SOLO Score Range	# of Subjects BInfoTech	# of Subjects BCompSc	# of Subjects BEng(SW)
More than 3 std dev below mean	0	1	0
Between 2 and 3 std dev below mean	1	3	0
Within 2 std dev of mean	17	19	41
Between 2 and 3 std dev above mean	2	2	2
More than 3 std dev above mean	0	3	1

Initial observations suggest that the degrees BInfoTech and BEng(SW) appear to be consistent in terms of their statements of learning objectives, but the BCompSc seems to have too many subjects that fall outside the acceptable limits, and may therefore become subjects of interest. Hence, using the data in the examples in this section, a course review of the BCompSc might choose to look back over the individual subjects and consider the

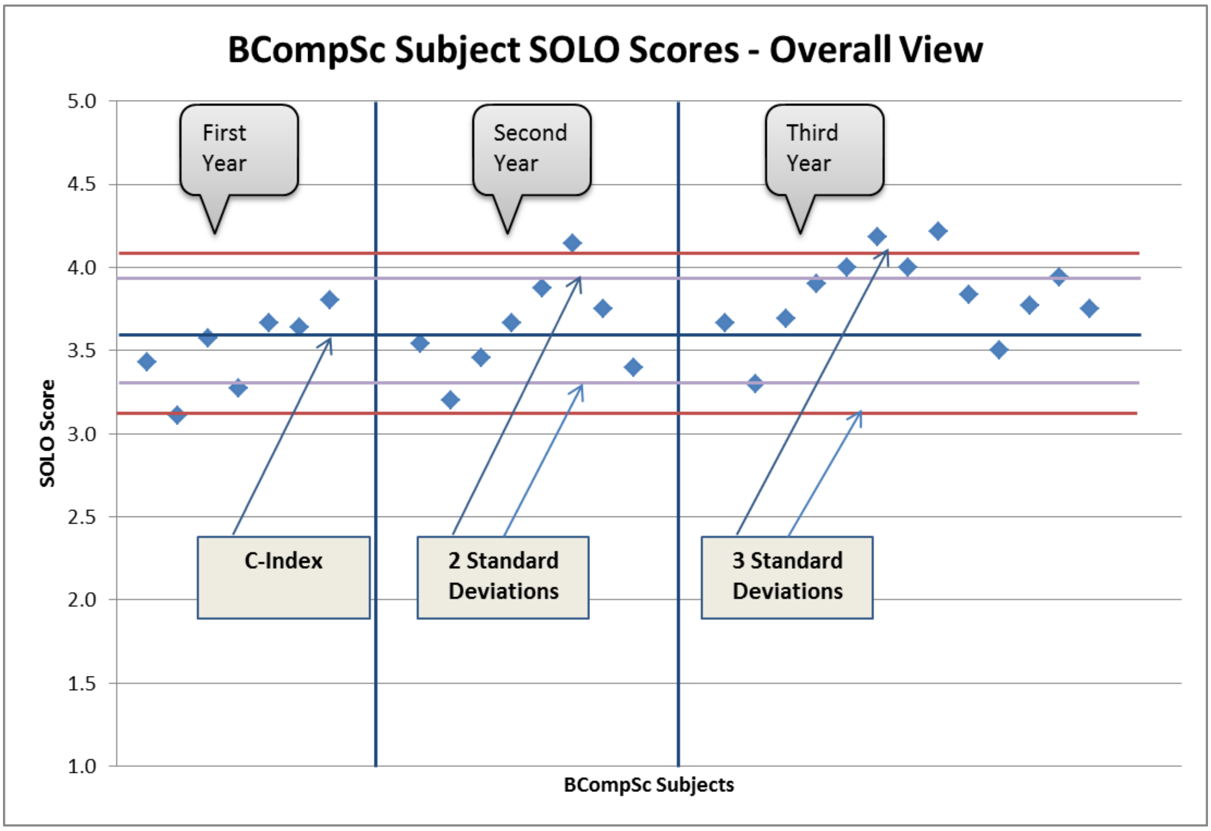

Figure 6.1: BCompSc Subject Analysis – Overall

learning requirement specifications of subjects where the subject SOLO score was outside the range of either two or three standard deviations from the mean - ie. 3.30 to 3.94 for two standard deviations from the mean, or 3.14 to 4.10 for three standard deviations from the mean. A graphical representation of this data is shown in Figure 6.1, where the overall C-Index is plotted and the two and three standard deviation boundaries are highlighted so that it is clear as to which subjects fall outside the relevant boundary limits.

Under this proposed analysis method for the BCompSc degree, there are potentially 4 subjects (14%) that could come under some scrutiny as they fall outside the three standard deviation range, or 9 subjects (32%) that are outside the two standard deviation range. Under Chebyshev's Theorem this observation is abnormal since there should be at most 10% outside the three standard deviation limits and at most 25% outside the two standard deviation limits.

For those which are below the range boundaries the question that arises is whether the stated or expected learning demands should be raised to be more consistent with other subjects in the degree or is it that the statement of learning objectives is inadequately expressed and therefore does not match the level of learning rigour that will be demanded of the students? In either case the observation that the subject SOLO score falls below a boundary level implies that there is some degree of inconsistency when compared with the other subjects in the course. This at least suggests that a review of the subject learning objectives may be needed.

For those above the range boundaries the converse applies, resulting in the question of whether too much is being asked of the students in that degree program at that stage of their learning, or equally the expression of learning expectation may be higher than that being delivered? As with those that fall below the boundary levels, those that lie above the boundary levels are equally identified as being worthy of further review.

Depending on the outcomes to those review questions, it may be that if the statements of learning objectives are deemed to be appropriate then other questions could be framed in terms of the learning support being provided to the students to enable them to better cope with the higher levels of learning expectation. It is also important to identify the year level of the subjects that are flagged as being of interest. In the above example it was noted that four of the five subjects above the upper boundaries were final year subjects, and of those below the range boundaries the lowest scoring subject was a first-year introductory subject and the other three subjects in the two to three standard deviation range included one first-year subject, one second-year subject and one third-year subject.

A different view could be held when looking at the data from a different perspective. For example, the scatter plot chart in Figure 4.1 showed a clear upwards trend for the BInfoTech subjects, and this is similarly the case for the BCompSc in Figure 6.1, so the simple control limits based on the overall course C-Index may not necessarily reveal what was hoped. The underlying proposition of using control limits in the production setting is that the expected output is the same or within a particular tolerance level. Therefore, in the academic setting we would normally expect the SOLO scores to increase from year to year, yet equally have an expectation that the academic rigour in each year level should be approximately similar. Hence it may be better to perform the analysis based on the year-level means, which would give a stepped set of control limits upon which to frame the evaluations. This may be the most appropriate approach given that the method of calculating the C-Index is based on averaging year-level scores.

6.1.2 Examination of Year-Level Scores

To further explore the step-wise approach to examine the subjects in each year level, a scatter-plot was prepared with the individual subject SOLO-scores grouped by year level in the course. Overlaid on the scatter-plot were the year-level means for each year level and control limit boundaries of two standard deviations. In the first part of the analysis the overall course standard deviation was used, but it could also be considered that the year-level standard deviation may be more appropriate for analysing subject data, and this approach has been used in the second part of this analysis.

Analysis with Overall Standard Deviation

The results of the first part of this analysis are shown in the accompanying charts (Figure 6.2, 6.3 and 6.4).

In this view, the BInfoTech subjects appeared to be within the control limits in the first two years of the degree, and there was one subject just under the lower control limit in the third year of the program.

The BCompSc subjects appeared to be quite satisfactory in the first year of the degree, having just one subject slightly outside the control limits above and below; the second year program suggested that there was one subject notably above the two standard deviation control limit and one slightly below; and the third year program had one subject notably below the lower control limit and two subjects slightly above the upper control limit. The earlier observation of 9 subjects being outside the two standard deviation boundaries should now be amended to note that there are more likely to be only 5 subjects from the second and third year programs that may fall into the *subject of interest* category.

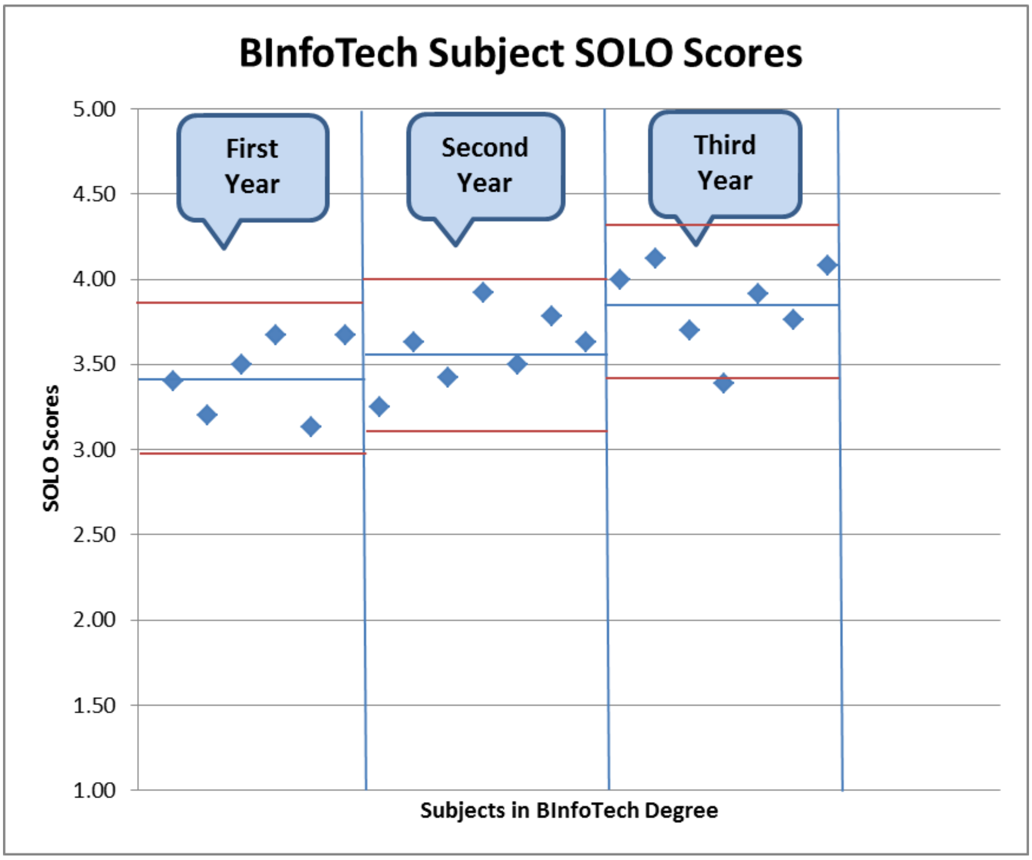

Figure 6.2: BInfoTech Subject Analysis

A similar analysis of the BEng(SW) subjects really only highlighted three subjects that are of interest, namely one in each of the second and third year programs that may be under-specified relative to the other subjects, and one second year subject that is above the upper control limit. There were three other subjects that were only just outside the two standard deviation control limit.

Analysis with Year-Level Standard Deviation

When a further investigation of the data was conducted using the year level segmentation, and applying the year-level standard deviation values to determine the control limits a somewhat different view was obtained. In particular, it was found that there was a difference in the standard deviations of the SOLO Scores in each year of study.

For the BInfoTech, the standard deviations of the SOLO Scores in year one, two, and

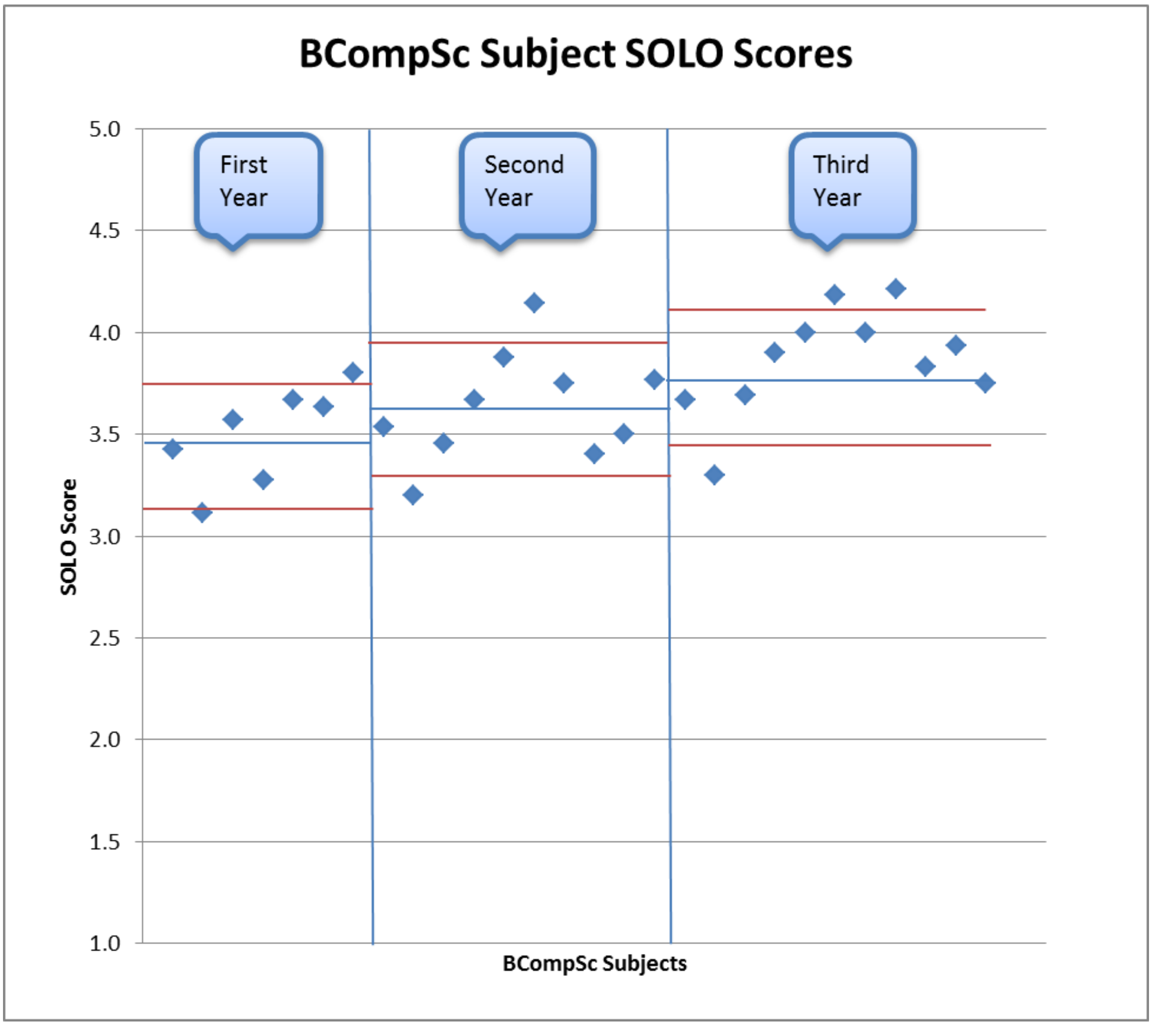

Figure 6.3: BCompSc Subject Analysis

three were 0.210, 0.208, and 0.252 respectively. While not significantly different from the overall standard deviation of 0.210, the increased third-year value changed the outcome to suggest that all subjects were within the two standard deviation control limit, as shown in Figure 6.5.

The greatest difference was seen in the BCompSc data where, compared with the overall standard deviation of 0.16, the standard deviations for year one, two and three were 0.201, 0.297, and 0.257 respectively. When the year-level standard deviations were overlaid on the scatter-plot data (Figure 6.6), it appeared that all subject scores lie within the two standard deviation control limits.

A similar outcome was evident in the analysis of the BEng(SW) data, where the use

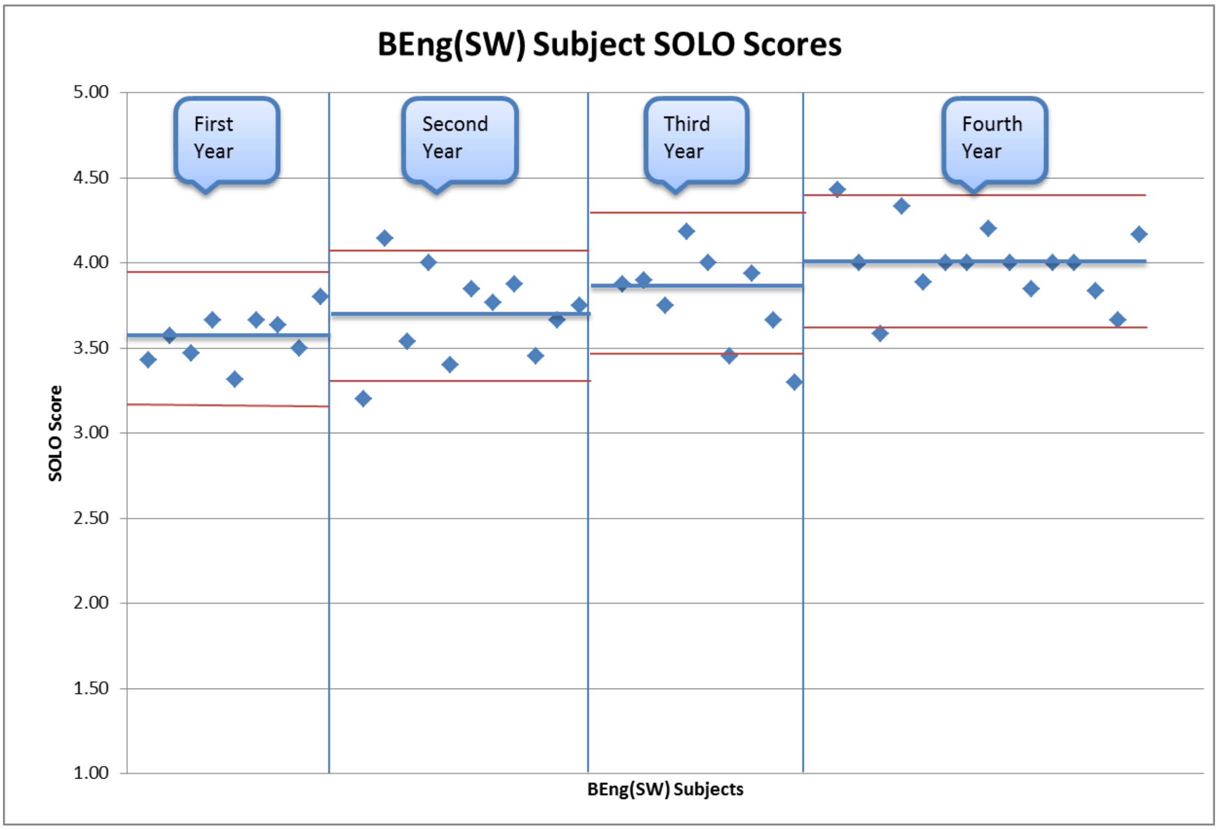

Figure 6.4: BEng(SW) Subject Analysis

of the year-level standard deviations to calculate the 2-standard deviation control limits highlighted just one third year subject that was slightly below the lower control limit (Figure 6.7).

It is beyond the scope of this research to answer those questions relating to the structure of the learning objectives, for that is the task of the course architects and curriculum designers of those subjects. What is shown however, is that this approach may be used as a tool to highlight or flag particular subjects as being worthy of further scrutiny for the reasons mentioned. What has been revealed is that a year-level segmentation of degree study programs is far preferable to considering simply the overall C-Index and placing control limits around that value. By definition, where there is an increasing progression of year-level SOLO Scores – that is, an upward trend line – then applying control limits around that C-Index would naturally expose the lower scoring subjects in the early years and the higher scoring subjects in later years as being outside the control limits.

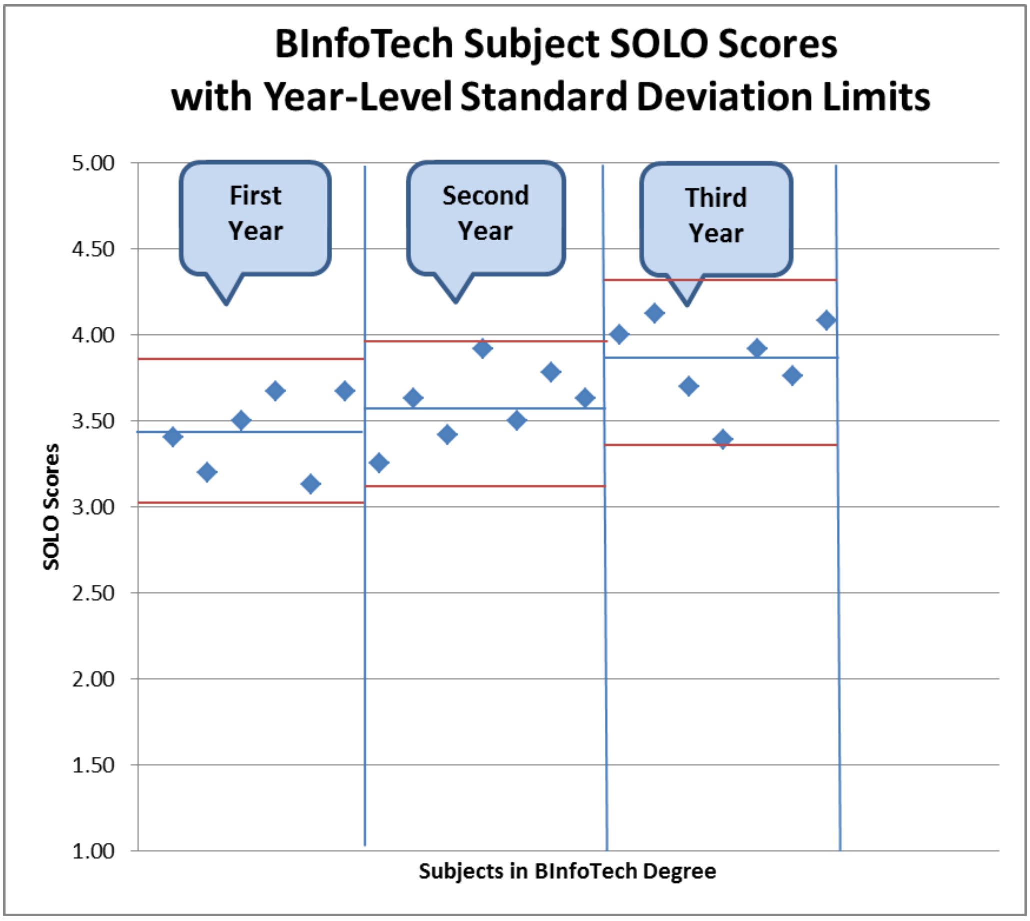

Figure 6.5: BInfoTech Subject Analysis by Year Level

Although the simpler approach is to use the overall course standard deviation as the metric upon which to calculate control limits for each year level, the analysis conducted tends to support the proposition that the year-level analysis is more appropriate. In terms of interpretation, it becomes clear as to which year levels within the degree program have the greatest amount of variation in SOLO Scores, simply by observing the year levels with the highest standard deviations. Again, it is the role of the course architects and curriculum designers to determine the reasonability factors for the subjects in the degree. It may be the role of university academic administrations to consider the establishment of standards and policies that they may impose on various departments to ensure some level of compliance with C-Index metrics and allowable variations.

This can best be demonstrated by reviewing a comparative table (Table 6.3) in which the year-level SOLO Scores and the corresponding year-level standard deviations are

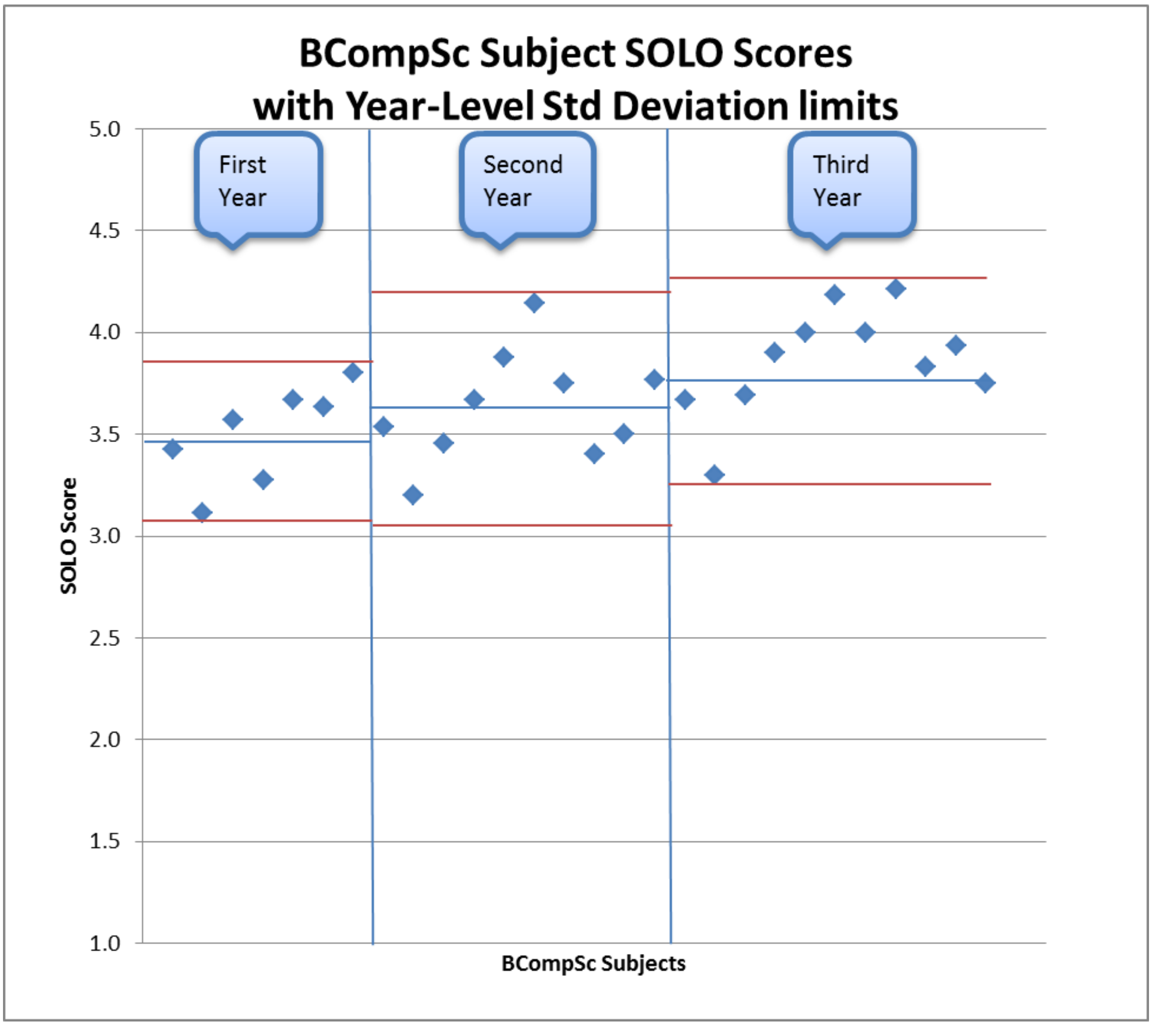

Figure 6.6: BCompSc Subject Analysis by Year Level

seen all together. In addition to highlighing where the greatest variations occur within courses in terms of the proposed learning rigour, the table also shows how an average of the year-level standard deviations would be a more representative value to use as a simple metric than the calculated course standard deviation when that is calculated as the simple standard deviation of the year-level SOLO Scores. This is particularly evident in the BCompSc course where the average over the three year-level standard deviations is 0.252, yet the initial calculated course standard deviation is 0.160. In a practical sense, it would be more appropriate to use this "mean course standard deviation" as the metric on which to calculate control limits in order to highlight the *subjects of interest* that may require review and revision of their stated learning objectives if a single standard deviation value was to be used.

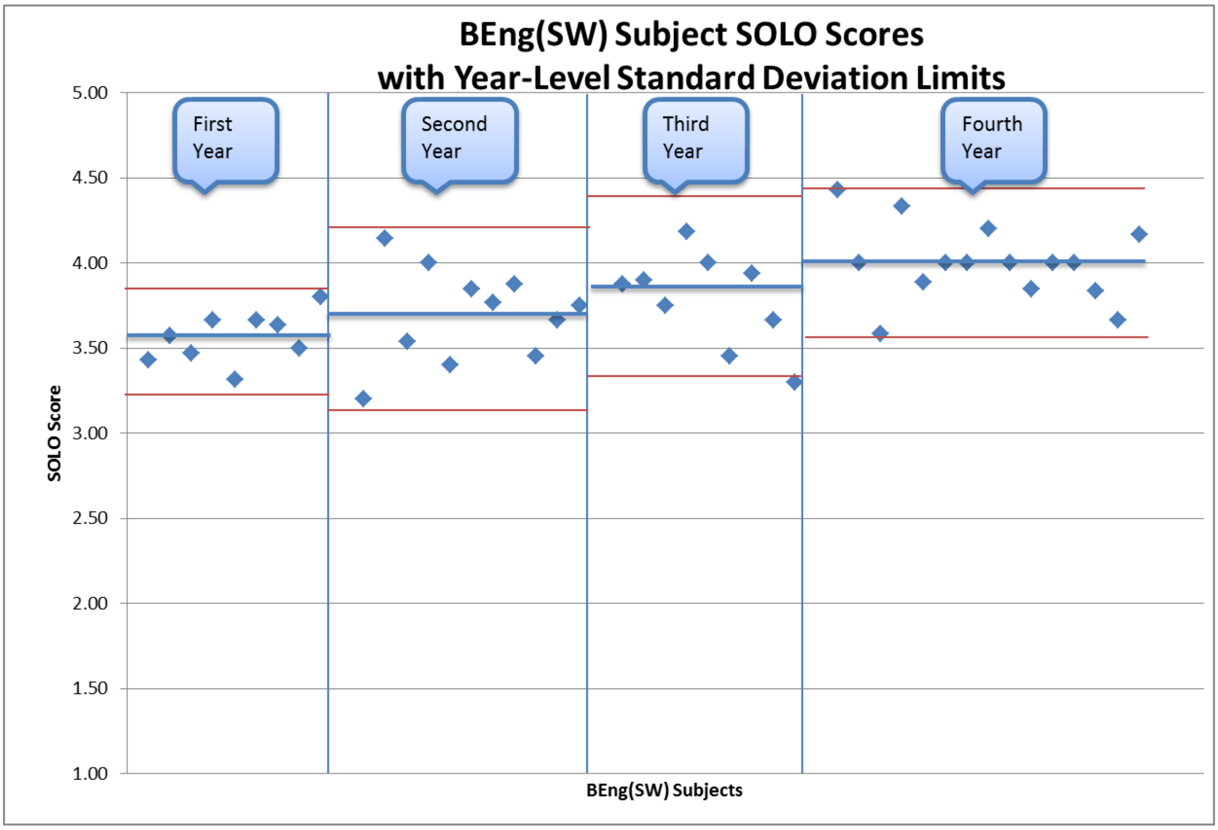

Figure 6.7: BEng(SW) Subject Analysis by Year Level

This chapter has contributed two significant elements to the thesis. Several approaches to the analysis of the underlying SOLO scores were investigated with the overall conclusion being that year-level analysis was the most appropriate method to adopt. The graphical charts based on a simple scatter plot of subject SOLO scores, segmented by course year-level, and overlaid with year-level means and two year-level standard deviations as boundary limits enabled a clear visual representation of where potential *subjects of interest* were in relation to the other subjects in the course. The subjects of interest may be viewed as outliers in the course data set and therefore are highlighted as potentially worthy of further scrutiny with respect to the statements of learning objectives.

It was suggested that it is the role of curriculum assessment groups to determine whether the statements of learning intent are specified as being too simple or too difficult. Importantly it was highlighted that there were several potential standard deviation values that could be used. In the first instance there was the initial course standard deviation which was calculated from the aggregated year-level scores. It was demonstrated that this value

Table 6.3: Year-Level Standard Deviation Distributions by Course

	Course					
	BInfoTech		BCompSc		BEng(SW)	
Year Level	SOLO	Std. Dev.	SOLO	Std. Dev.	SOLO	Std. Dev.
Year 1	3.43	0.194	3.45	0.201	3.55	0.131
Year 2	3.56	0.198	3.63	0.297	3.68	0.278
Year 3	3.85	0.225	3.77	0.257	3.87	0.274
Year 4	–	–	–	–	4.00	0.225
Average	3.62	0.206	3.62	0.252	3.78	0.227
Calculated Course Std. Dev.	0.210		0.160		0.199	

was too imprecise to be used for further decision making. The second approach was to use the year-level standard deviations, on the basis that the year-level segmentation was most likely to be more reflective of the way the course structures were viewed. A third option was proposed as using the mean of the year-level standard deviations as an overall indicator of the subjects' learning rigour variability within the course.

From the point of view of internal quality control, the data in Table 6.3 clearly highlighted the year-levels within courses where the greatest disparity in subject SOLO scores exists, namely those with the highest value in the standard deviation column. In the event of a course review, the second year program of the BCompSc and the second year program of the BEng(SW) should be good starting places to consider the learning rigour comparability.

Another option that would be available to University administrations through their curriculum advisory committees would be to review courses across faculties or perhaps University-wide and establish course specification performance metrics where standardised allowable variations were documented. Since the SOLO Taxonomy approach has a known range of values, it would be quite feasible to consider the subject learning outcomes for courses and create a standard that reflected the allowable variation that had been agreed to or prescribed. In a practical sense such a standard would state something along the lines of "*. . . the year-level standard deviation for the subject learning outcome SOLO Scores in the . . . degree should not exceed 0.25 . . .*" and the curriculum management group would then have a means to examine individual courses for internal consistency, and even to make appropriate comparisons across different courses. Using the C-Index calculations also enables those curriculum committees to compare the learning rigour specified at each year-level of individual courses based on the year-level means, and this too could be used as an indicator of comparability between courses. Ultimately, it would be equally possible to specify a faculty or University-wide standard for the acceptable ranges of year-level means as a measure of course learning rigour. For example, it may be that a University chooses to specify that the first-year mean SOLO Score should lie between 3.40 and 3.60; the second-year mean SOLO Score should lie between 3.55 and 3.75; the third-year mean SOLO Score should lie between 3.70 and 3.90; the fourth-year mean SOLO Score should lie between 3.85 and 4.05; or other range values deemed appropriate. (Note: The values given in the examples in this paragraph are purely arbitrary for demonstration purposes.)

In this chapter the data examined was from several degree programs at a single university. The goals were to explore the C-Index calculation on known courses, and to investigate other analytical techniques based on that data. The results obtained have validated the methods used and have proven the concepts introduced. In the next chapter, the main goals are to extend the data sets to degree programs from other universities in Australia, and test whether the analytical methods devised are applicable in the wider context.

Chapter 7

Benchmarking

In the previous chapter it was shown that comparative profiles for degree courses were able to be prepared using the data from a single institution in a related discipline area. In order for the techniques described to have wider applicability, it was necessary to apply those analytical methods to the courses of other institutions' degree programs.

Accordingly, the aim of this part of the research was to demonstrate that the techniques were robust and applicable to 'any' degree program. It will be shown in the remainder of this chapter that the analytical methods were effective in their application to other courses. The profiling techniques and calculation of C-Index values for degree courses has meant that there is now a quantitative metric that can be used as part of the benchmarking processes that are so important to universities as they strive to maintain quality standards.

7.1 Selection of Courses

For the purposes of this part of the research it was necessary to obtain appropriate data from more than one university. However, in order to minimise the introduction of additional variables in the research, the course data selected was chosen to be from universities offering a degree in Information Technology. To maintain data independence, it was further decided to restrict the course data to those universities where both the course rules and detailed subject descriptions containing the behavioural objectives or learning outcomes were publicly available on the university web site. Given the initial part of the research was conducted in Flinders University in South Australia, it was also decided that the comparisons made should come from Universities in other states in Australia.

Initially those universities in Australia that offered degrees in Information Technology were selected as potential data sources, with their selection being based on whether there was a clear statement of the course structure as well as making available the detailed subject descriptions including the behavioural objectives or learning outcomes. While there were quite a number of universities that could have satisfied this requirement, a small cross-section across several states with universities of different types were decided upon. In Australia there are currently four groups of Universities and another 'group' of those which are not members of one of the other groups (Australian Education Network, 2014). Those groups are the 'Group of Eight' (go8), the 'Australian Technology Network' (ATN), the 'Innovative Research Universities' (IRU), and the 'Regional Universities Network'. Although not specifically relevant, a university comparable with Flinders University (IRU member) was deemed to be necessary as well as at least two others. In particular the following universities and degree courses in Australia were selected:

- Swinburne University of Technology – Bachelor of Information Technology (Swinburne University, 2012). Swinburne University is located in Melbourne in the State of Victoria. (non-member group)
- University of Queensland – Bachelor of Information Technology (University of Queensland, 2014). The University of Queensland is located in Brisbane in the State of Queensland. (go8 member)
- University of Newcastle – Bachelor of Information Technology (University of Newcastle, 2014). The University of Newcastle is located in Newcastle, a regional city on the east coast of New South Wales. (IRU member)

It was found that a fairly common practice among Australian universities was to offer a particular degree program with students being able to choose from one of several streams of interest to construct a degree using their main interest area as a major component of the degree. In an in-depth analytical study of these degree programs, each of the major streams would have been analysed to build a set of alternative pathway scores. However, for the proof-of-concept approach it was decided to review either the generic pathway if available, or a more traditional software engineering or application development stream if it was necessary to focus on a streamed major to comply with the specific degree requirements.

7.2 Data Analysis

7.2.1 Swinburne University BIT

The Bachelor of Information Technology offered by Swinburne University appears to have a slightly unusual structure based on the interpretations gained from their public website information. A highlighted strength is the Industry Based Learning, which is a little confusing in that the detailed specifications for the subject suggest 0 credit points, yet the degree rules indicate the Industry Based Learning counts for 100 credit points. This degree course is a 3-year full-time degree program with traditional semesters 1 and 2, supplemented by a summer semester between years 1 and 2, and again between years 2 and 3. The course schedule requires a total of 400 credit points, with 100 of those allocated as Industry Based Learning and the other 300 being comprised of 24 subjects, each with a unit value of 12.5 credit points. On this basis, the nominal weight for each subject is therefore 0.125, although there is a slightly uneven distribution of when the subjects are actually studied according to the recommended study sequence. Additionally the two blocks of Industry Based Learning were deemed to have a value of 50 credit points each with a corresponding weight of 0.500.

In the calculations made, the list of published ICT electives was averaged by their subject code being in the 1xxx or 2xxx classifications (4 subjects) as electives for the second year program, and the 3xxx or 4xxx electives (20 subjects) were averaged to determine the third year elective scores. Although the course rule does allow another 14 suggested non-

ICT electives or other non-ICT electives subject to other factors, these were not used in the calculations of the elective subject scores. The core and selective subjects comprised 32 distinct subjects to be assessed. The results of the individual subject assessments are given in Appendix B.1 and Appendix B.2.

Overall, an analysis of the Swinburne Bachelor of Information Technology degree returned a C-Index value of 3.49, which is a little less than the score obtained for the Flinders degree. The profile of the degree, based on the SOLO Distribution obtained is given in Figure 7.1, and this shows similar characteristics to other degree courses examined, although it does appear to be somewhat under-stated in the higher order SOLO Level 5 (5% at third year level, compared with 16% in the Flinders degree).

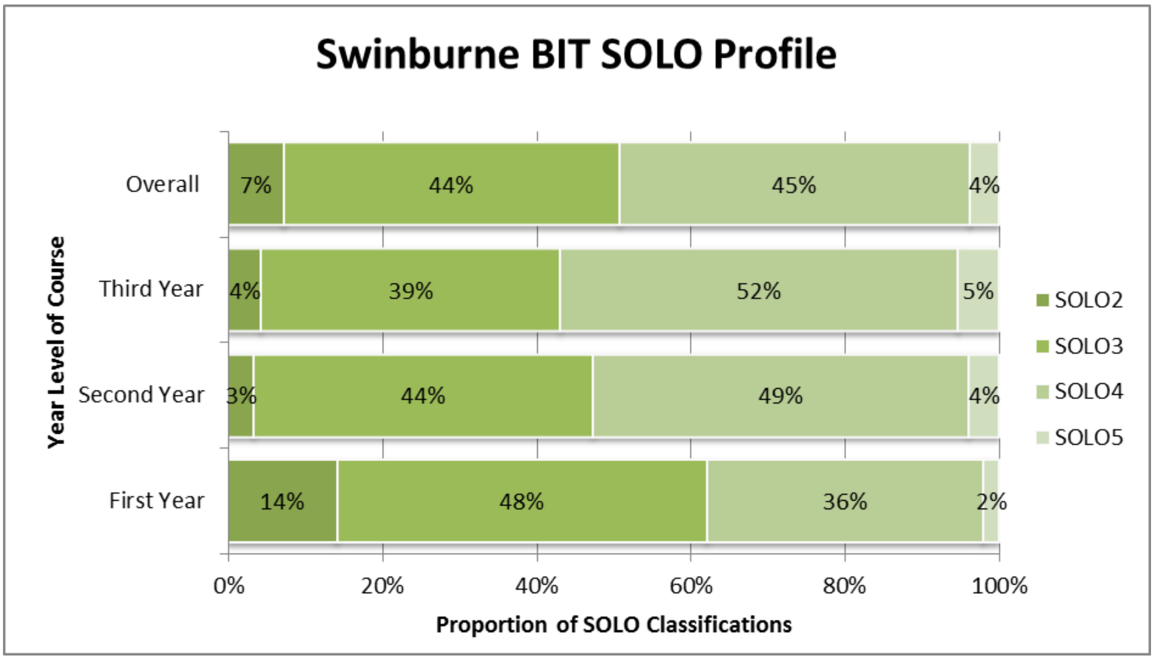

Figure 7.1: Swinburne BIT Analysis Summary

Observations

When a comparable year-based analysis was undertaken using the year-level SOLO Scores and the associated year-level standard deviations as described in the previous chapter, the picture revealed was more interesting (see Table 7.1 and Figure 7.2). In particular, there did not appear to be a clear upward trend of learning expectation across the degree.

While there was an expected and familiar learning rigour jump from first year to second year, there was no corresponding jump in the third year expectations. Rather, it seemed to taper off after second year with there being negligible difference between the second year SOLO Score (3.58) and the third year SOLO Score (3.59).

Using the techniques described in the previous chapter on internal control, a set of control limit boundaries for two and three standard deviations for each year level was prepared, and is shown in Table 7.2. As seen in Figure 7.2 there was one subject in the first year program that was below the two standard deviation lower control limit, one subject in second year that was close to the upper control limit and all other subjects were within the two standard deviation control limit boundaries. The other point of interest was that just one of the subjects listed in the degree had a SOLO Score of 4 or above.

Table 7.1: Swinburne BIT Year Level Summary

Year Level	Mean	Std Dev
First Year	3.30	0.222
Second Year	3.58	0.244
Third Year	3.59	0.125
Overall C-Index	3.49	

Table 7.2: Swinburne BIT Subject Control Limits

Control Limits	First Year	Second Year	Third Year
Year-Level Standard Deviation	0.222	0.244	0.125
3 std dev below mean	2.63	2.85	3.22
2 std dev below mean	2.85	3.09	3.34
Year-Level Score (mean)	3.30	3.58	3.59
2 std dev above mean	3.74	4.07	3.84
3 std dev above mean	3.96	4.32	3.97

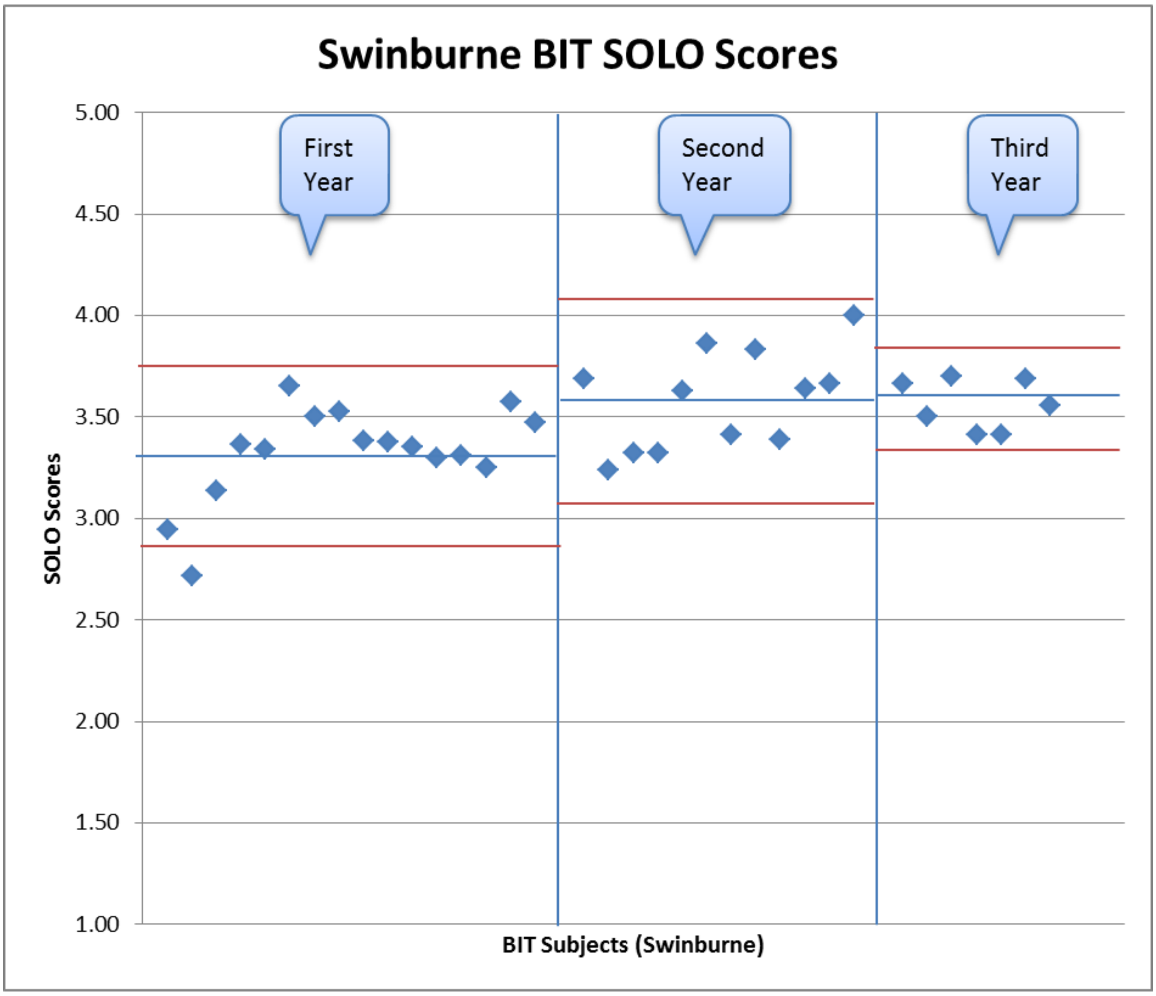

Figure 7.2: Swinburne BIT Subject Analysis by Year Level

7.2.2 University of Queensland BInfTech

The University of Queensland offers several variants of the Bachelor of Information Technology. There is a generic offering with no majors of study, a version with one major of study, another version with two single majors of study, and a version with a dual major of study (University of Queensland, 2014). The majors are appropriate sequences of related subjects that enable students to build particular specialisations into their degree studies such as Computer Systems and Networks, Human-Computer Interaction, Software Design, Games Modelling, Health Informatics, Information Security, and others.

Whilst it would be possible to determine C-Index scores for each of the variant options, for consistency with other University courses, only the variant with no majors has been evaluated in this research.

The unit value system adopted by the University of Queensland assigns two units to a typical semester subject. Hence the normal degree program of 48 units over three years would typically involve a student undertaking 16 units per year, or four subjects per semester. There are several project-based or research-based subjects that may be taken in the final year of the study program depending on student choices and these have a four-unit value. The course information for the degree is quite broad in its specification, giving only general unit requirements for the completion of the degree. At the time of the research being conducted, it appeared that there had been a revision to the course and a consequent renaming of the subjects as the listed recommended study program guide referred to subjects with either a different name or a different subject code. Hence the data gathered refers to the course as it would be from 2015.

The stated requirement for the Bachelor of Information Technology (BInfTech) – no major option – for students commencing in 2015 was shown as:

> BInfTech with no major, #48 comprising -
> a. at least #18 from Part A; and
> b. at least #6 from Part B, with at least #2 from Part B1; and
> c. at least #8 from Part C; and
> d. the balance from electives being courses from Part D or other courses approved by the Executive Dean; with no more than #8 of level 1 courses.

As with other courses examined, the University of Queensland degree has a mix of compulsory subjects (core), selections from a limited range of subjects (selective), and a broader range of elective subjects. In the above schedule of degree requirements, Part A lists the compulsory subjects (10 in total), Part B lists the introductory electives (9 subjects), Part C lists the advanced electives (17 subjects), and Part D lists the other electives (21 subjects). The subjects in Part B and Part C were considered as selective subjects as they were more domain specific, and those in Part D were regarded as recommended electives, but not specifically evaluated as were the other subjects.

With the less specific course directives in the requirements statement, the interpretation of subjects undertaken in each year level has been done on a more generically aggregated basis. As a result the data obtained may not be properly reflective of the true learning expectations for students in this degree course.

Observations

What can be seen is that the University of Queensland course does present a quite similar overall profile to the other courses that were examined in this discipline area, although the first and second year subjects appeared to score a little lower than the other courses observed, which has meant that the profile suggests a not so strong focus on the higher order learning activities in those earlier years. However, this view is countered by the strong impact of the third year subjects which appear to be much more demanding in their proposed learning rigour. The sharp jump can be seen in the profile graphic (Figure 7.3) where the higher order learning expectations (SOLO 4 and SOLO 5) account for 59% in third year compared with 33% in second year, and the marked increase in the year-level score from 3.27 to 3.71 (Table 7.3).

Table 7.3: University of Queensland BInfTech Year Level Summary

Year Level	Mean	Std Dev
First Year	3.18	0.156
Second Year	3.27	0.257
Third Year	3.71	0.278
Overall C-Index	3.39	

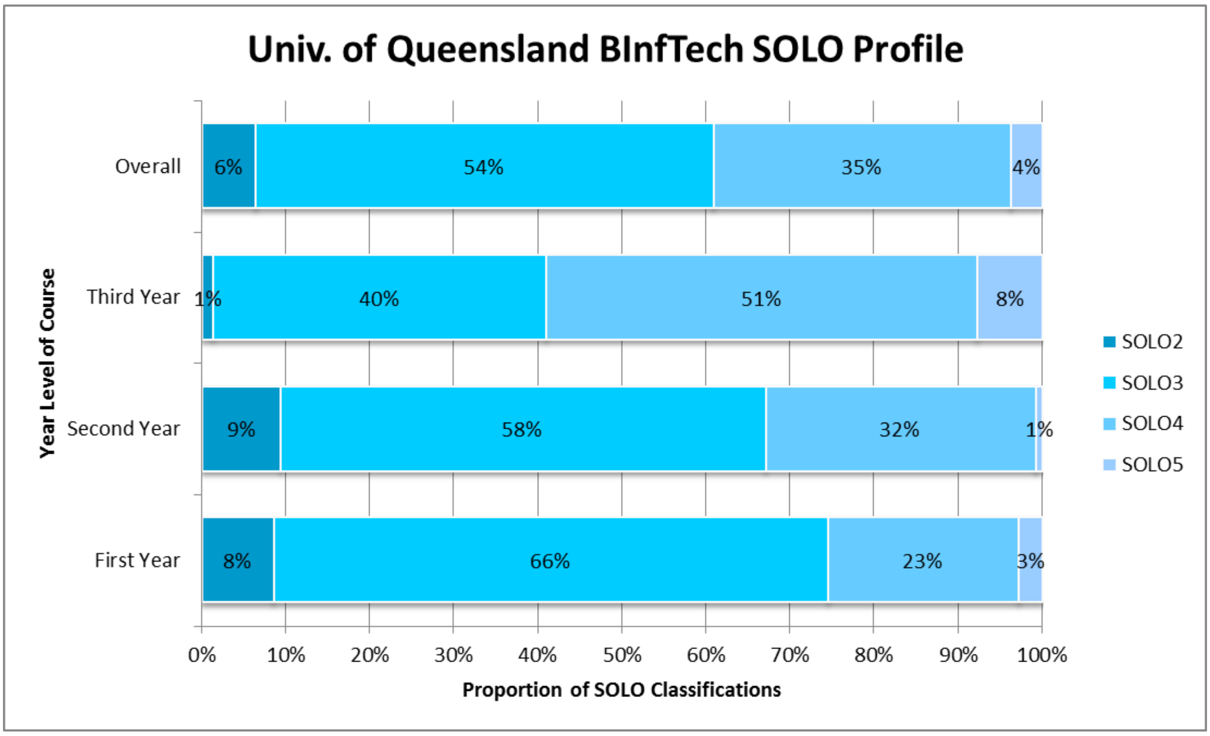

Figure 7.3: University of Queensland BInfTech Analysis Summary

Upon examining the scores of the University of Queensland subjects by year level it was clear that effectively all of the subjects fell within the two standard deviation control limits with only a small number of subjects that might come under consideration as 'subjects of interest' – one in first year, one in second year and two in third year that were around the upper control limits and one third year subject that was close to the lower control limit (Figure 7.4). There was only a small difference between the first year and second year year-level scores, but a much more significant upwards shift in the third year-level score. Potentially this could be interpreted as an indicator to highlight that perhaps the second year subjects could need to be reviewed as a whole in order to provide a smoother progression from first year to second year to third year in this degree.

Table 7.4: University of Queensland BInfTech Subject Control Limits

Control Limits	First Year	Second Year	Third Year
Year-Level Standard Deviation	0.156	0.257	0.278
3 std dev below mean	2.71	2.50	2.88
2 std dev below mean	2.86	2.76	3.16
Year-Level Score (mean)	3.18	3.27	3.71
2 std dev above mean	3.49	3.79	4.27
3 std dev above mean	3.64	4.05	4.55

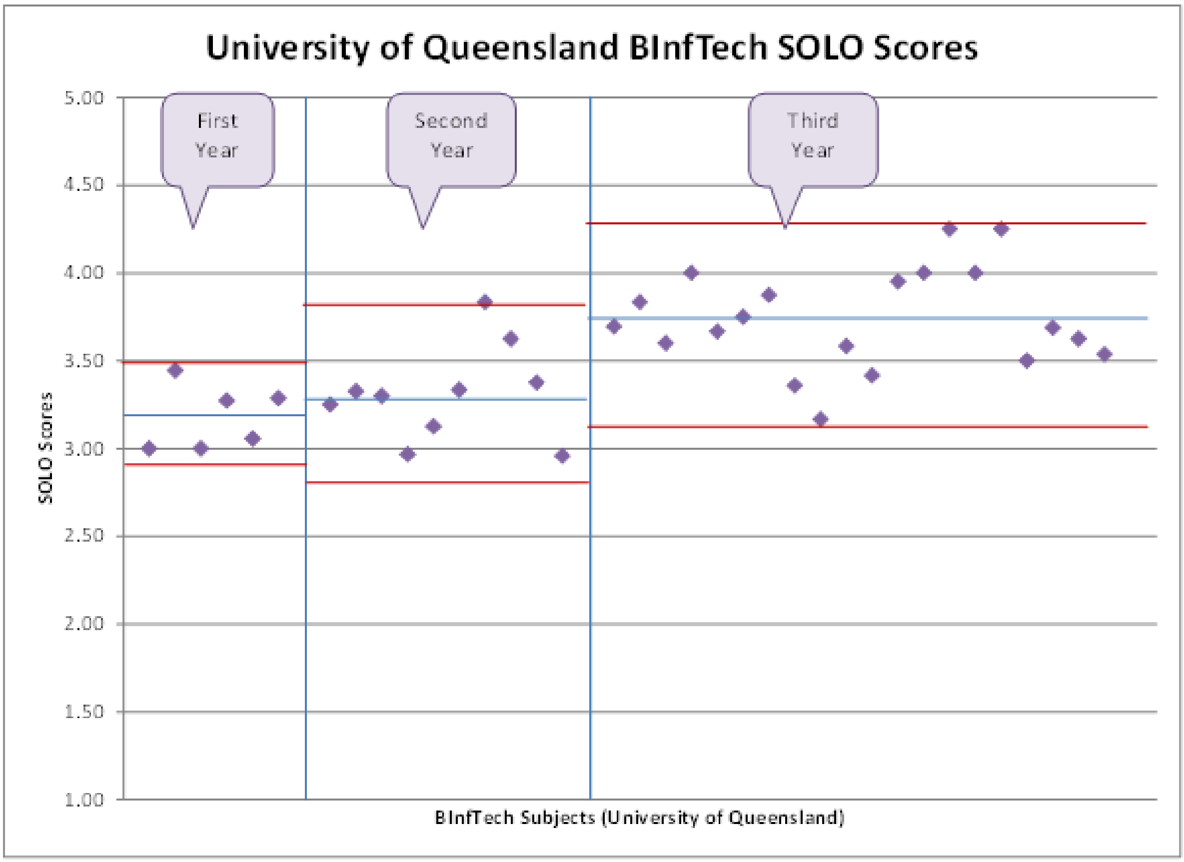

Figure 7.4: University of Queensland BInfTech Subject Analysis by Year Level

7.2.3 University of Newcastle BIT

As with the other universities considered, the University of Newcastle Bachelor of Information Technology has some degree of flexibility in its offering. The course was revised to a new structure and format from the beginning of 2014, and the descriptions, especially in the course specification, contain references to the previous structure. The requirements for the completion of the degree state (University of Newcastle, 2014):

> The 240 units required to complete the degree must include:
> a. All core courses (100 units);
> b. A major sequence (80 units);
> c. Electives (as many units as are required to bring the total units up to 240 units. Some students may wish to apply this to a second major. Students can elect to count core course INFT3970 IT Major Project towards one major only).
>
> Please note:
> No more than 100 units in total at 1000 level will be counted towards the award.
> At least 60 units must be completed at 3000 level.

In the degree program there are four majors defined, namely:

- Data Analytics Major;
- Digital Media and Entertainment Major;
- Enterprise Information Technology Major; and
- Software Development and Applications Major.

and each has a nominated set of compulsory *courses* and a set of directed *courses* which are the preferred electives for students in that major stream to choose. The course rules for the degree are generous in that the directed subjects are not mandated as needing a specific number of them within the major stream, but clearly it would be sensible for students undertaking the relevant major to select subjects from the directed subjects list.

For the purposes of this research only the Software Development and Applications Major was chosen to be examined as this appeared to be closest to a generic Bachelor of Information Technology as seen in other universities.

Each of the subjects (courses in the University of Newcastle jargon) in this degree is based on a value of 10 units other than the major project subject which is 20 units. Over a typical three year program, each year would then require 80 units, which gives a subject equivalent full-time load of 0.125 for a 10-unit subject.

The structure of this degree is quite flexible in that it does allow a reasonable number of elective subjects to be chosen in areas within the degree theme or from outside the degree speciality. As with other degrees investigated the list of elective subjects was too numerous to examine each one individually. It would be possible for an internal curriculum group to identify the most common set of electives normally chosen by students, but that information was not readily available under the approach taken in this research. Accordingly, since there was not a specific constraint on the number of directed subjects at each year level, the calculation method for this degree was to use the selective subjects average, weighted to the balance of subjects needed beyond the core subjects in each year level to determine the year-level scores and ultimately the overall C-Index.

Observations

The results of the analysis for the University of Newcastle Bachelor of Information Technology degree are shown in Table 7.5, and this data shows a good amount of progression from year one to year two to year three in the learning rigour demands, with the overall C-Index for the degree being calculated as 3.49. The corresponding profile for the degree is shown in Figure 7.5, where that learning rigour progression is quite clear. It was noted that there was a strong increase in the higher order SOLO levels over the three year program going from 32% in first year up to 64% in third year.

From the internal quality control perspective, the subjects in the revised course structure all fall within the two standard deviation control limit boundaries, although there appears to be one subject in first year and one subject in third year that can be labelled as subjects of interest as they are close to the lower control limit boundaries, and one subject in third year that is close to the upper control limit boundary (See Figure 7.6).

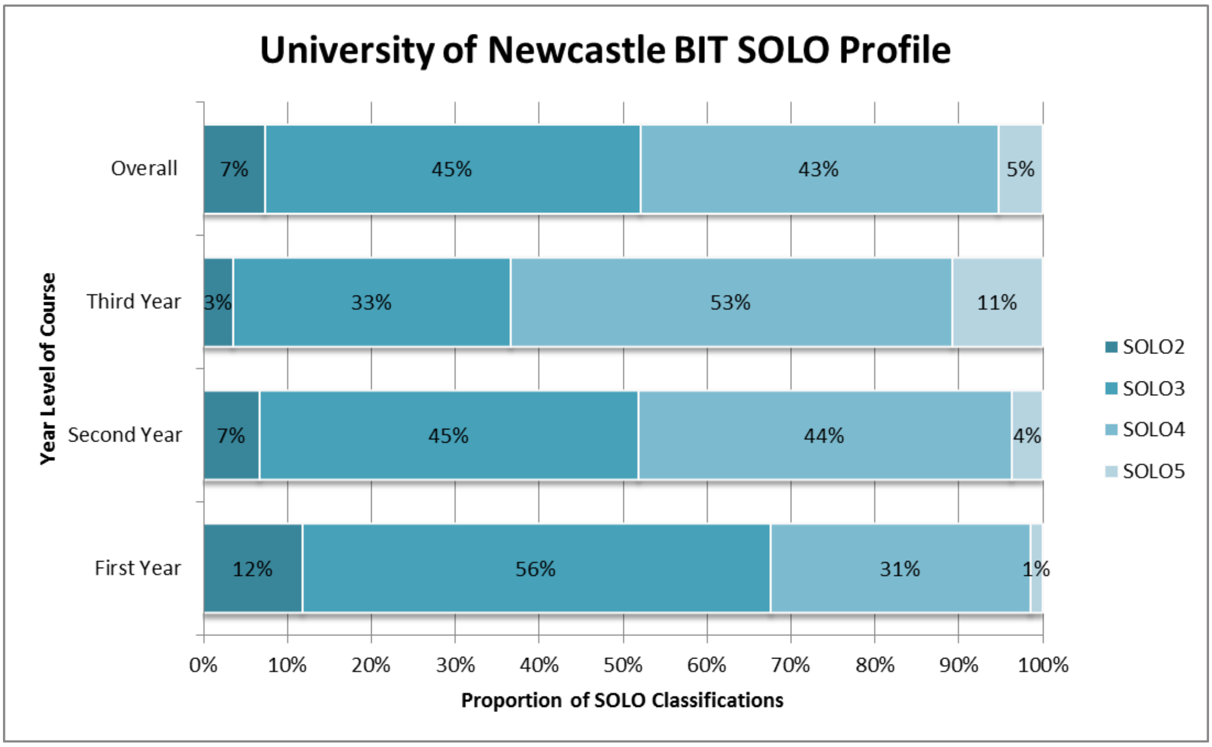

Figure 7.5: University of Newcastle BIT Analysis Summary

Table 7.5: University of Newcastle BIT Year Level Summary

Year Level	Mean	Std Dev
First Year	3.18	0.214
Second Year	3.51	0.183
Third Year	3.79	0.333
Overall C-Index	3.49	

Table 7.6: University of Newcastle BIT Subject Control Limits

Control Limits	First Year	Second Year	Third Year
Year-Level Standard Deviation	0.214	0.183	0.333
3 std dev below mean	2.54	2.96	2.79
2 std dev below mean	2.75	3.15	3.12
Year-Level Score (mean)	3.18	3.51	3.79
2 std dev above mean	3.61	3.88	4.45
3 std dev above mean	3.82	4.06	4.79

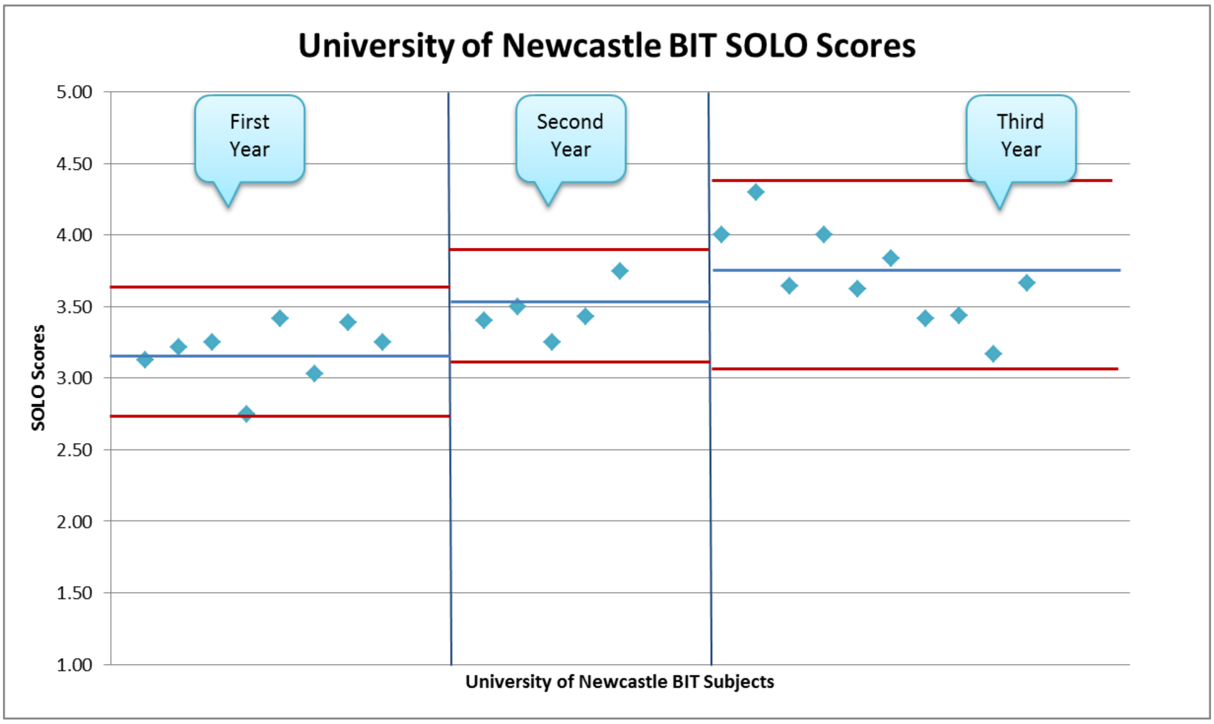

Figure 7.6: University of Newcastle BIT Subject Analysis by Year Level

7.3 Benchmarking Results

There have been two distinct approaches to benchmarking described in this thesis. In the earlier developmental part where several different degree courses from a single university were evaluated, it was shown that the evaluation approach enabled a set of year-level scores to be calculated for the individual degree programs. The approach and those results were given in Chapter 5, with a summarised version shown here in Table 7.7.

Table 7.7: Course Scores – Single University

	BInfoTech		**BCompSc**		**BEng(SW)**	
Course Year Level	SOLO Score	Std Dev	SOLO Score	Std Dev	SOLO Score	Std Dev
First Year	3.43	0.194	3.45	0.201	3.55	0.131
Second Year	3.56	0.198	3.63	0.297	3.68	0.278
Third Year	3.86	0.225	3.77	0.257	3.87	0.274
Fourth Year	-		-		4.00	0.225
C-Index	3.62		3.62		3.78	

From an institutional viewpoint the data can be examined across degree programs to ascertain whether there is a comparable level of learning rigour displayed in each year level of the courses, and equally to determine whether each of the degrees have a comparable C-Index. It was previously discussed that the implied progression from one year level to the next in learning rigour expectations should be seen in increasing year-level scores is evident in this group of degrees that were examined. The other factor which becomes a corollary to that implied progression assumption is that a four-year degree should score higher than a three-year degree, resulting in a higher C-Index, and this was also evident in the data shown.

The underlying questions that curriculum review committees might ask when presented with this data include some of the following:

- Are the C-Index values appropriate for these courses?
- Are the year-level programs comparable across the courses?
- Is the amount of year-level progression in learning rigour expectation suitable in each course?

- Is the standard deviation in expected learning rigour acceptable in each year level in each course?
- Is the relative proportion of low-order and high-order learning appropriate in each course?

The last of these questions requires one to look at the SOLO Distributions for each course at either the overall level or at the detailed year-level. The comparative overall SOLO Distributions for the three Flinders University courses examined, namely the Bachelor of Information Technology, the Bachelor of Computer Science, and the Bachelor of Engineering (Software) are shown in the accompanying Figure 7.7. The detailed profile for the Bachelor of Information Technology was previously given in Chapter 4, Figure 4.3.

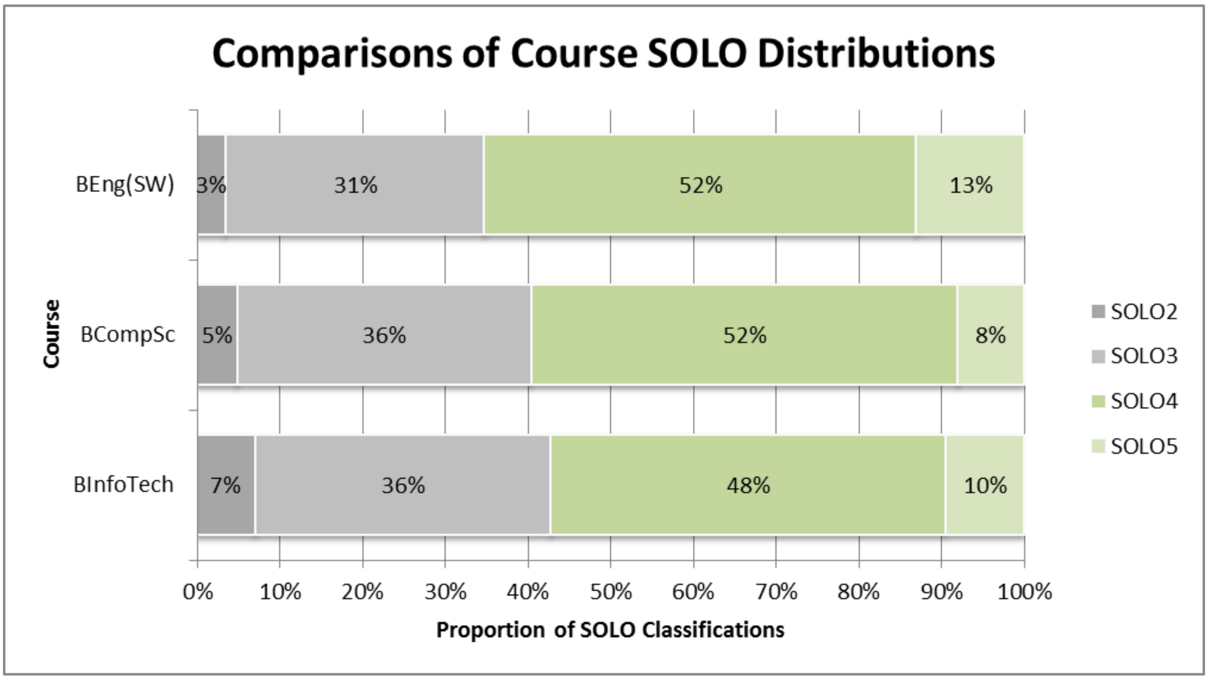

Figure 7.7: Comparisons of Flinders University Courses

The higher order SOLO levels of SOLO-4 and SOLO-5 indicate expectations of deep learning compared with the more superficial learning demands at SOLO-2 and SOLO-3 levels. The courses shown have combined SOLO-4 and SOLO-5 proportions of 58%, 60% and 65% respectively. This demonstrates a level of consistency across those courses, which are from the same general field of study domain.

The second approach to benchmarking occurs when the course evaluation process is applied to comparable courses across different institutions. In the earlier part of this chapter, the Bachelor of Information Technology courses at three different universities were examined with both profile determination and C-Index calculations being performed. The same set of questions is applicable to cross-institutional reviews, and the accompanying table (Table 7.8) shows the summarised data for the different universities whose courses were examined.

Table 7.8: Course Scores – Multiple Universities

	Flinders		**Swinburne**		**Queensland**		**Newcastle**	
Course Year Level	SOLO Score	Std Dev	SOLO Score	Std Dev	SOLO Score	Std Dev	SOLO Score	Std Dev
First Year	3.43	0.194	3.30	0.222	3.18	0.156	3.18	0.214
Second Year	3.56	0.198	3.58	0.244	3.27	0.257	3.51	0.183
Third Year	3.86	0.225	3.59	0.125	3.71	0.278	3.79	0.333
C-Index	3.62		3.49		3.39		3.49	

On the basis of the data presented, there were several potentially interesting observations that could be made. In the first instance, there is the observation that the University of Queensland degree has the lowest C-Index of the four universities examined. Similarly the Swinburne University third-year program scored the lowest of the final year programs. The greatest amount of subject score variation was seen in the third-year program at the University of Newcastle and the least amount was in the third-year program at Swinburne University.

When the overall course profiles for these four courses were compared (see Figure 7.8) it appeared that the University of Queensland course was not as strong in the higher order learning rigour demands as the other three universities. Possible reasons for this outcome were discussed in the earlier section (Section 7.2.2) where that course was discussed in greater detail. The profiles for the University of Newcastle and Swinburne University courses were very close to one another, and the Flinders University course appeared to score a little higher in the overall higher order learning rigour demands.

While it is possible to compare the courses against each other, as has been done in this case, the ultimate benchmarking tool would have established a baseline standard against which the courses would be compared. At this point in time, such a standard does not exist, but it is proposed that the techniques described in this thesis may become a starting point for such course standardisations.

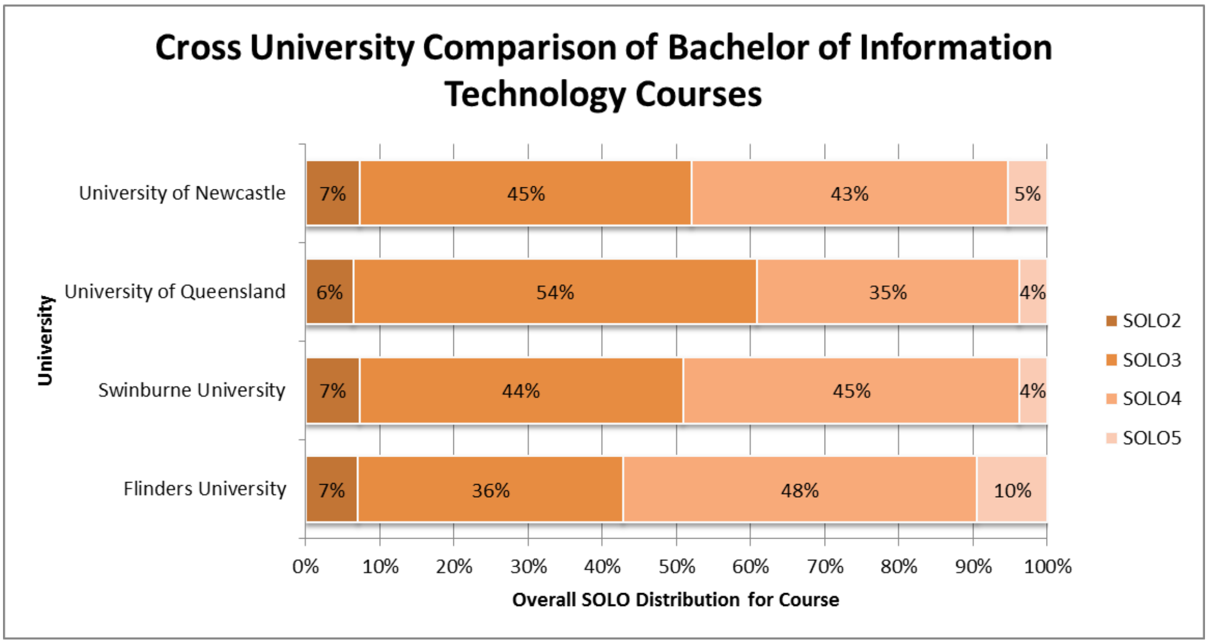

Figure 7.8: Comparisons of Bachelor of Information Technology Courses

This chapter has combined the techniques and approaches described in earlier chapters to demonstrate the applicability of those techniques to the benchmarking of courses across universities. The data collection involved the retrieval of subject learning outcomes for each of the required core subjects and selective subjects in conjunction with the course rules for each of the degree programs at the universities selected. The number of subject evaluations conducted were 31 from Swinburne University, 36 from the University of Queensland, and 23 from the University of Newcastle. For each university the subjects were grouped by year level and identified as core or selective subjects according to the relevant course rules. From these appropriately scored data, year-level scores, SOLO Distributions, C-Index calculations, year-level standard deviations and course profiles were prepared. For analytical and comparison purposes the data were presented in both tabular and graphical formats, with supplemental annotations showing the upper and lower control limit boundaries at the two standard deviation mark based on year-level scores.

In the next chapter a review of the methodology, techniques, approaches, and an analysis of this research will be discussed.

Chapter 8

Analysis and Discussion

This chapter reflects on the approach taken, reviews the results obtained from the various experiments conducted in this research, and discusses those results in relation to the research questions posed initially.

8.1 Methodology Discussion

The methodology that has been used in this research was based on the well-established SOLO Taxonomy to go beyond the initial ideas proposed when that taxonomy was first devised. In the educational theory domain, the use of taxonomies has typically been qualitative in their application to educational programs and their relative components. Research undertaken by Brabrand and Dahl suggested a conversion of the qualitative into the quantitative through the use of the 'double-weight averaging scheme' to arrive at the *SOLO Average* score for an individual subject or group of subjects, and a subsequent *SOLO Distribution* which was derived from counting the number of descriptors in the relevant SOLO categories. An interesting point of difference from that study was that the Danish study had been conducted after a major undertaking to revise the subject statements of learning outcomes was completed, whereas in this research the candidate courses in the Universities considered contained a mix of well-formed statements and older-style statements that did not align as closely with the terminologies of the SOLO

Taxonomy. Where Brabrand and Dahl were able to more straight-forwardly classify the 'intended learning outcomes', in this research a more subjective judgement was required to interpret the estimated likely intention of the stated objectives and learning outcomes.

Accordingly, the first part of the research was done in conjunction with experienced subject coordinators to assist in the interpretation of the behavioural objectives or learning outcomes for the subjects under their control. The outcome from this part of the approach was to gain a level of confidence in being able to reasonably accurately interpret the intent of the stated learning outcomes in a consistent manner within a familiar discipline area. With this knowledge, combined with restricting the data sets to the domain of Information Technology and Computer Science degrees, the subjectivity in classifying all of the behavioural objectives and learning outcomes for the constituent subjects was consistent across the project. Because there was this element of subjectivity involved rather than a purely objective classification method, a valid concern is whether other researchers would obtain results that were quite close to those presented in this thesis. It is certainly possible that personal bias could influence the judgements made about the classification and subsequent scoring of individual subjects when other researchers apply the methodologies described in this thesis. However, it is also a reasonable belief that higher education professionals are both prudent and responsible when making assessments on educational matters. Consequently the approach proposed should be sufficiently robust to accommodate individual differences in the interpretation of stated learning objectives, particularly when those evaluations are undertaken by specialists in the relevant discipline area.

The initial data collection was taken from a known University site (Flinders University), where full and ready access to course and subject information was available, to enable a controlled study to take place to validate the techniques used and effectively confirm the proof of concept for the proposed theory. The expanded study was undertaken to assess the methods against external data sources, using similar degree programs from other Australian Universities. The choice of institution was based on whether there were readily accessible statements about the course rules and the subject content for the selected degrees.

The varying complexity of course rules at different institutions and the difficulty of retrieving some of the subject information effectively resulted in not being able to automate most of the detail level analysis. The statements of behavioural objectives and learning outcomes varied from very brief and vague to being quite verbose and detailed. Accordingly, a simple mapping of these statements into SOLO classifications was not possible, and a more extensive manual process was required to calculate the individual SOLO scores for each objective, and then determine the SOLO Average for each subject.

In the selected domain for this research, namely the fields of Information Technology, Computer Science, and Software Engineering, the course rules tend to be more tightly specified than for other more generic courses such as a Bachelor of Science or a Bachelor of Arts, yet even so they still varied quite a lot between the institutions examined. Equally the subject statements of behavioural objectives or learning outcomes varied considerably from being quite vague to being very specific, and also from simple clear statements to complex compound statements. These factors meant that the interpretation of the statements usually required a judgement to be made about the intent of the statement as much as what was actually stated.

Once the individual subject learning outcomes had been assessed, the determination of the subject SOLO Score and its SOLO Distribution was a relatively simple calculation that was used in the overall aggregation of data for the year level. In the Brabrand and Dahl study the method used to determine the SOLO Distribution was based on a proportional contribution of SOLO scores in each of the intended learning outcomes. This meant that when a single learning outcome was expressed in a compound form with differing SOLO levels, a fractional contribution for each of those SOLO levels was applied to the overall SOLO Distribution for the subject. The alternative approach, labelled the simplex approach (see Section 2.3.1), was to use the raw count of the number of SOLO levels described in the learning outcomes for the subject. After comparing the two approaches it was decided to adopt the simplex approach to determine the SOLO Distributions as it was computationally a little simpler; the differences between the two methods were not substantial at the overall level, although they may be observed to be more significantly different at the individual subject level; and potentially more straight-forward to follow for casual users of the methodology. The subsequent calculation of year-level scores was dependent on the complexity of the course rule for the course, and again required an individually constructed calculation template for each year level. The final aggregations for overall analysis of the course, including the course profile and the internal quality control charts were the most straight-forward segments to produce.

8.2 Results Interpretation

Using the methodology proposed in this thesis, a C-Index value has been obtained for the degree courses reviewed. As has been previously discussed, there are just five levels of the SOLO Taxonomy, and the scoring system provides a value potentially in the range 1 to 5, but is actually only in the range 2 to 5 since all learning activities are aimed at levels 2 or above. Even using decimal values to say 2 decimal places, as has been done in the calculations shown, the resultant value is in a scaling range that is foreign to most readers. For example, is there a significant difference between a degree course which has a C-Index of 3.25 compared with one that has a C-Index of 3.68?

An earlier discussion in Chapter 2, Section 2.2.3 (Biggs and Collis, 1982) pointed out that the SOLO Taxonomy classifications were cumulative in nature, implying achievement at Level 4 also meant that Level 3 and Level 2 skills were mastered. Combined with the equal distance assumption, it became possible to create a linear scale on which the C-Index scores could be plotted. However the limited range of the scale 1 to 5 makes it difficult to gauge differences unless they are quite substantial.

Most educators are familiar with percentage scales, so there appears to be potential merit in converting the C-Index values to a number in the range 1 to 100. One simplistic approach would be simply to adjust the C-Index to a percentage value, based on the maximum score being 5. Hence in the above example, 3.25 would convert to 65.0%, and 3.68 would convert to 73.6%. Academics would readily concede that there is a noticeable difference in these two percentage scores.

However, even this is not a true reflection, as the effective range is just 2 to 5, so a more accurate conversion should highlight the progression from 2 to 3 to 4 to 5 by accentuating the difference regions. Accordingly a proposed scaling formula is (C-Index - 1) * 25, resulting in 3.25 converting to 56.25, and 3.68 converting to 67.0, which both lie in the 1 to 100 range, but support slightly different viewpoints. In particular, if we examine the 'breakpoint scores', the conversion points can be seen in the following table (Table 8.1), which also shows the simple percentage breaks.

Table 8.1: Scaled C-Index Scores

C-Index Score (Breakpoints)	Simple Percentage	Scaled SOLO Score
2	40.0	25.0
3	60.0	50.0
4	80.0	75.0
5	100.0	100.0

If one considers the widely held view (Biggs and Collis, 1982; Brabrand and Dahl, 2007) that *surface learning* occurs at Levels 2 and 3, and *deep learning* occurs at Levels 4 and 5, and that most academics would hope that the degrees in which they teach would aim more towards deep learning, then a scale that proposes an outcome in what the academic world considers to be advanced levels of achievement is more likely to be accepted. Equally, an *ordinary* level of achievement should have a resultant score at around a typical passing grade score. To this end, the mapping algorithm proposed does deliver such scores, with the score of 50 being the differentiation point between the surface learning scores and the deep learning scores.

If one were to adopt the Scaled SOLO Score approach, those degree programs that scored closer to 50 than 75 could be thought of as providing largely superficial coverage of their subject areas, whereas the higher scores promise an intention of demanding deeper coverage of their subject areas. By virtue of considering the underlying premise of this research, namely that the examination of behavioural objectives and/or learning outcomes of the subjects in a degree course, the calculated results present an intention of what is proposed to be covered within the degree, and the depth of learning that should be achieved by a student in that course. What can be concluded from this approach is that a significant difference in the course specification analysis may be used as an indicator of course quality if we associate quality with depth of coverage.

Returning to the data collected in this study and applying the Scaled C-Index Score, the results can be seen in Table 8.2.

Table 8.2: Scaled C-Index Result Scores

Degree Course	C-Index Score	Scaled C-Index
Flinders University BCompSc	3.62	65.50
Flinders University BEng(SW)	3.78	69.50
Flinders University BInfoTech	3.62	65.50
Swinburne University BIT	3.49	62.20
University of Queensland BInfTech	3.39	59.75
University of Newcastle BIT	3.49	62.20

At first glance, one might be tempted to say that the Flinders University BInfoTech offers the most rigorous Information Technology degree (of those degree courses examined) in terms of learning expectation placed upon students. However, the research process has assessed the statements of behavioural objectives or proposed learning outcomes, so a more appropriate interpretation of the data would be that each of the courses examined have a significant proportion of surface learning and make a good attempt at developing deep learning.

8.3 Specific Outcomes and Contributions

There have been several distinct outcomes arising from this research, each of which contributes new knowledge to the discipline areas of Computer Science and Information Technology, and the associated development of educational programs in these areas.

8.3.1 Course Profiling

The application of the quantification of behavioural objectives using the double weighted averaging scheme proposed by Brabrand and Dahl using the SOLO Taxonomy gives rise to a *SOLO Average* for an individual subject of study. When applied over a set of related subjects according to the course rules for a degree program, the methodology described in this thesis then generates a **year-level score** for each of the year levels in the degree, and by aggregating the number of SOLO classifications for each subject in the year level a *SOLO Distribution* can be determined to highlight the proportion of each type of SOLO classification in each year level. The subsequent aggregation across the whole of the degree program then gives the overall **Course Profile** and the **C-Index**.

The terms SOLO Average and SOLO Distribution are inherited from the Brabrand and Dahl study. New terms developed as part of this research are:

1. **Year-Level Score**;
2. **Course Profile**; and
3. **C-Index**;
4. **SOLO Distribution – proportional method**; and
5. **SOLO Distribution – simplex method**.

Although the year-level score is an intermediate item that is used in calculating the course C-Index, in itself it may not at first appear to be particularly important. However it is a good indicator when viewed in relation to the other year-level scores to see whether the progression of learning rigour demands seems reasonable for the particular degree course. It does become significant when comparing one course against another in that it can be noted that year-level scores are or are not approximately equivalent, thus implying that the different courses may be demanding similar levels of learning rigour (or not).

The Course Profile appears to have great merit for tasks such as capturing a quick snapshot of the nature of learning expectation in a course. Another potential application for using the Course Profile would be in the marketing of the degree program to future students, highlighting the pattern of learning expectation as they move from first year through to their final year.

The C-Index is a single value statistic which attempts to assess the overall learning rigour proposed for a degree course. Its value as an indicator is yet to be proven in wider testing but it does have the potential to highlight the extent to which deep learning can be anticipated in the particular degree.

In determining a SOLO Distribution for a course it was found that there were two possible approaches to calculating the contribution of SOLO levels for an individual subject. The proportional method assigned a proportional value of SOLO level for each learning outcome, and then the sum of each of the level scores was divided by the number of learning outcomes and converted to a percentage to give the subject SOLO Distribution. The simplex method aggregated the raw counts of each of the SOLO levels expressed in the learning outcomes, which meant that complex and compound learning outcome statements could over-emphasise some SOLO levels in a subject. When expressed as a percentage value this was observed to be different from the proportional method scores for an individual subject. However, in the overall weighted aggregation, where the 'big picture' view was the major focus, the individual differences became less significant.

8.3.2 Internal Quality Control

The Internal Quality Control contribution has arisen from a more detailed study of the individual subjects and their SOLO Average scores within the degree. It was found that grouping the subjects by their year levels enabled some degree of comparative analysis to be undertaken. In all degree programs examined there was a reasonable amount of variation in the SOLO Average scores within year levels. By using the standard deviation statistic a nominal two standard deviation set of control limits was applied to the year-level score to highlight those subjects that were close to or outside the upper and lower control limits.

The new terms developed for considerations about internal quality control were:

1. **Subjects of Interest**
2. **Control Limit Boundaries**

In this era of standardisation and normalisation, the determination of Subjects of Interest within courses can become most valuable to highlight where learning rigour expectations may be too high or too low for the year of study. It has been pointed out previously that there may well be valid reasons for the individual subjects appearing as Subjects of Interest, but the process should be seen as a great benefit to University Courses and Curricula groups during their course review and development periods. The secondary benefit which comes as a direct result of having performed these analyses is that the same data would allow external reviewers and course benchmarking panels to readily identify focus areas for their purposes.

The Control Limit Boundaries proposition is based on using the SOLO Scores for the set of subjects within each year level of the course and determining the standard deviation of those scores. The Year-Level Score is calculated as the weighted mean of the subject SOLO Scores, and the Control Limit Boundaries are calculated for each year level as two standard deviations above and below the Year-Level Score. There is an underlying assumption that each of the subjects in a particular year level should have a similar level of learning rigour for the internal quality control approach to function as described.

8.3.3 Course Benchmarking

The concept of course benchmarking is common in Australian Universities, particularly in the science, engineering, and technological areas where many of the courses are subject to accreditation by external professional societies to satisfy the academic requirements for membership by the course graduates. Current tools used tend to be based on outputs – examination papers, assignments, projects, and similar. While these are essential elements, to date there has been little on the course specification side that can provide helpful information to benchmarking teams and accreditation panels. The research undertaken in this thesis has considered the preliminary information for the subjects in a course of study, namely the behavioural objectives or learning outcomes that have been specified. The research proposition is that these items express the learning rigour expectation for each of the subjects in a course, and by applying the methodologies described it is possible to arrive at a value to represent that learning rigour expectation, which has been labelled as the C-Index.

The C-Index is not intended to replace the current benchmarking tools, but rather to supplement them. When used in conjunction with the Course Profile information, whether tabular or graphical, benchmarking teams may be able to conduct their benchmarking exercises more efficiently as they would have the course intent information as well as the course output information to review. The positive feature of using pre-delivery material is that the variables of the nature of the student cohort, the factors of the teaching team(s) and other environmental issues are removed. As discussed in the earlier section on this matter, student cohorts change from year to year, teaching groups change, environmental factors including changes to resources such as equipment and software can impact significantly so that overall the course outputs will have some variation from year to year, yet the course specifications remain stable over a period of years and will only change at course review times.

8.4 Limitations

As with any developmental work in a new area there are limitations as to what can be done, and the research in this thesis is no exception. A number of the issues and potential problems have been discussed in earlier sections of this thesis, but a more succinct grouping of those limitations is listed below:

1. The resulting subject SOLO Scores were very much dependent on the quality of the expression of the learning outcomes and behavioural objectives. Where subjects had been relatively recently updated it was obvious that greater attention had been given to attempting to express the learning outcomes in more taxonomy specific language in many cases. There appeared to be evidence of a greater knowledge of appropriate learning theory and educational taxonomy jargon for some of those subjects. Not every subject in every university examined could lay claim to that, but anecdotally at least it appears to strengthen the argument that greater attention to learning outcome descriptions is becoming a higher priority in Australian Universities.

2. Where learning outcomes were expressed in complex compound statements it was necessary to parse the different ideas being defined to best determine the SOLO levels covered in the statement as well as the number of each type. In some cases the outcomes were a similar task level across several sub-topic areas within the subject, and therefore may have received a score of 3 instances at SOLO-2 for example. In other cases there were several distinct outcomes at differing SOLO levels, and may therefore have been scored at 2 instances of SOLO-3 and 1 instance of SOLO-4 for example.

3. A criticism that could be levelled at this research is that of subjectivity and bias, where the researcher has been required to make judgemental decisions about classifications of learning outcome statements. During the initial stages of this research the interpretation of learning outcome statements was undertaken with subject co-ordinators in order to develop some degree of expertise in determining the intent of the statements. Thereafter that expertise was applied to other subjects and assessments were made on the basis of that early learning. The issue of bias was minimised by constraining the research to fields of study with which the researcher was familiar – that is the areas of Information Technology, Computer Science and Software Engineering. In extending this research, it would be preferable to have curriculum specialist teams from within the domain area to evaluate the subject learning outcomes to minimise the subjectivity and bias elements.

4. One of the obvious limitations of using the SOLO scores is the very narrow band of values that can be applied in the coding of the learning outcomes, where the choices are 2, 3, 4, or 5. When combined with the relatively small set of common verbs used to describe learning outcomes, and the possibility that the learning outcomes have been created by less well-informed academics, it means that there is likely to be only a small observable difference in scores for subjects across year levels. It could be possible for example that a specific learning objective appeared to be the same in a first-year subject, a second-year subject, and repeated in a third-year subject, yet the expected depth of learning may be quite different. To clarify this, a sample learning objective of "describe the structure of a relational database . . . ." could

be part of an introductory computing subject, repeated in a subsequent database subject, and repeated again in an advanced database subject, yet the intended learning expectations would be quite different for each of those subjects. Based purely on the verb "describe", each of these would be classified as SOLO-3.

Using this same example, it becomes clear that the authors of the learning outcomes do need to be more aware of how learning outcome statements may impact on the interpretations of these subject specifications. In the event that institutions decide to adopt the approaches described in this thesis, it is recommended that appropriate training be given to those people who are to apply such coding techniques to the statements of learning objectives for the purpose of preparing course profiles and associated analyses.

5. It was highlighted in section 2.3.1 that individual subjects may have few or many learning objective statements. Whether this is a considered decision by the subject administrator as needing to specify many outcomes to cover the subject matter, or a disingenuous attempt as 'needing to put something down' without giving a great deal of thought to what was stated, is an unknown factor. Presumably the curriculum managers in Universities would mediate what was stated prior to widespread publication of those learning objective statements.

   Should the theories proposed in this thesis become adopted by Universities, there is another possibility that subject administrators might choose to manipulate their learning objective statements to return a higher subject SOLO Score, and potentially overstate the level of expected learning demand for the subject. Although this would be feasible for a subject administrator to be "smart" about using the descriptions to boost the SOLO Score, it should ultimately be moderated by the other internal quality control processes the University would have in place that confirmed a strong correlation between what was stated as intended and what was being delivered and/or achieved in the subject.

6. The approaches described in this thesis are really only applicable to courses where the course rules are well defined and have sufficient specificity to allow the associated C-Index calculations to be performed relatively easily. Those courses which offer

many optional pathways make it difficult to reliably determine course profiles and C-Index values. Nevertheless it would be feasible to evaluate 'typical' course selections and highlight the corresponding course profiles and C-Index values for these more broad-based degree courses.

7. In the discussion of the methodology for this research (Section 8.1) it was pointed out that the complexity of course rules meant that a modified template needed to be created for each year level of each course that was examined. The high level of manual intervention needed posed a natural restriction on the breadth of study that was able to be achieved, and further poses a potential restriction on subsequent researchers who wish to expand on this study.

Chapter 9

Conclusions and Future Work

9.1 General Concluding Remarks

The motivation for this research has been to explore the hypothesis that "*taxonomic tools are able to be applied to the course objectives for university level Information Technology and Computer Science courses to provide an indicator of course quality*" (see Section 1.2 and below in Section 9.1.1). It has been shown that the application of a new methodology and formalised approach has enabled the quantification of subject behavioural objectives or learning outcomes in such a way that a profile for a university course of study can be established when using those stated subject objectives as input. In so quantifying the subject objectives and applying the course rules for the degree program an instrument has been created that can contribute to the information base of various stakeholders associated with that degree. This may include university departments, faculties, and central administrations to enable them to assess the projected perception of the levels of study required in their degrees. It may also provide supporting data to external agencies such as accreditation panels and benchmarking teams when considering matters such as course quality and standards. Another stakeholder group who may find such information useful in their decision making is the end-user client – the student and perhaps their parents, along with advisers at schools and universities. The course profiles determined during this research have applications in university marketing and publicity to show the proposed levels of educational rigour for those degree programs that have constructed course profiles using the approach described in this thesis. As highlighted in Chapters 4 and 5, the graphical representation that is derived from the quantification of the course subject objectives can be a useful supplement to the current word-based descriptors of degree courses.

9.1.1 Research Question Outcomes

In order to properly conclude this thesis it is necessary to return to the initial research questions to examine the outcomes of the research in relation to those questions.

The fundamental question of this research is whether taxonomic tools are able to be applied to the course objectives for university level Information Technology and Computer Science courses to provide an indicator of course quality.

There are supporting questions that must be answered in order to arrive at a definitive and supported result in answer to the primary question. In particular,

- *Which taxonomy tools are appropriate for Information Technology and Computer Science courses?*
- *Is there a suitable metric that is able to be derived using the taxonomy tools?*
- *Does the tool metric provide a useful measure for assessing the learning rigour of a course?*
- *Do similar courses return a similar result using the metric?*

The fundamental research question has been answered in the affirmative in the introductory comments for this chapter, and in the discussion chapter in Section 8.3.1 in particular where the resultant outcome of the application of the techniques described have given rise to the new concept of a course profile and the new metric labelled the C-Index.

The first of the supporting questions on the clarification of which taxonomy tools are appropriate is discussed more fully in Section 9.2.1, where it has been stated that either the SOLO Taxonomy or the revised Bloom Taxonomy could be used, although the C-Index results would be different as the two taxonomies do use different scales.

The answer to the second of the supporting questions has been clearly identified throughout the thesis, and the C-Index is a new metric that has been proposed. This is further discussed in Section 9.2.3.

The third of the supporting questions relating to the usefulness of the C-Index has been highlighted in Sections 8.3.2, 8.3.3, 9.2.4, and 9.2.5.

The fourth of the supporting questions concerning the similarity between comparable courses was discussed at various stages, in particular in Sections 8.3.3 and 9.2.5.

9.1.2 Overall Remarks

It has been demonstrated that the methodology does indeed return a degree profile and a corresponding C-Index for the course. The very nature of the C-Index is that it proposes an intended level of educational rigour for the relevant degree. This of course does not prescribe what is actually delivered in the degree program, as that will be determined by the way in which the teaching teams choose to deliver the material and the demands they place upon students in the learning and assessment practices at each stage of the degree. However, as a guidance figure, the C-Index can be used to interpret some of the unstated impressions relating to the degree program. What is not stated, and cannot be determined from the structure of the existing sets of learning outcomes observed, is the level of learning needed in order to pass a particular subject. In the courses examined, there were no indications about which learning outcomes were essential, desirable, or optional. Whether it would be feasible to prescribe the subject Pass/Fail criteria based on the statements of learning outcomes could be the subject of further research.

The degree courses chosen in this research have all been from the Information Technology and Computer Science domain to demonstrate the efficacy of the approach, and that there is some consistency across the domain area. On the basis of the investigations undertaken, it appears that some of the courses are aiming towards the more practical side of the domain area by specifying more strongly the application and implementation considerations while other courses aim to provide a more strategic approach with stronger considerations in attempting to get students to extend their thinking into higher order issues that involve making sound judgements and predictions. A common feature that was evident from the analysis of the SOLO Distributions by year level is that in almost every case there was a significant component of low-level superficial coverage of subject matter in the first year of the course, giving way to increased levels of higher order learning expectations in later years of study.

The issue of subjectivity in the analysis of the behavioural objectives and learning outcomes has been discussed earlier, and this is a valid concern. At some point of course, someone has to make a decision whenever classifications of anything are made. In the instance where some may argue that the decision-making is subjective, others might equally argue that they are 'exercising their professional judgement'. For this particular field of assessing educational statements of learning outcomes, there is less likelihood of indiscriminant subjectivity as the assessors are more likely to be professional people making an informed judgement of the material they are examining. However, by following a structured approach in a consistent manner the inaccuracies that may have arisen through subjective assessments of learning outcome statements have been minimised. For subsequent researchers to extend upon this research approach it would require some initial assessor familiarisation to be undertaken to again ensure a relatively consistent outcome was obtained.

9.2 Specific Outcomes and Contributions

There have been several specific outcomes from this research which can also be viewed as contributions to the area.

9.2.1 The Use of Educational Taxonomies in Computer Science

The use of educational taxonomies is widespread in all fields of study, but there have been some concerns raised about the appropriateness of their use in the field of Computer Science. The arguments for and against were presented in Chapter 2, and the research in this thesis has been based on the affirmative view that educational taxonomies are indeed relevant to the field of Computer Science. The matter of whether a new taxonomy should be created for the field of Computer Science was rejected, at least at the strategic level. The issue of being able to better accentuate the aims of application development and implementation in existing taxonomies was acknowledged. Although this can cause potential under-valued scoring in some of those subject areas, degrees in either Computer Science or Information Technology do comprise more than just application development and implementation, and so the overall impact is not such a major concern. Careful reconstruction of the learning objectives to more accurately reflect the learning skills required of the student, particularly those that would be classified at the higher order taxonomy levels, are needed in a number of the subjects examined in this research.

The two major educational taxonomies in common use are the Bloom Taxonomy (either in its original form or in the newer revised form), and the SOLO Taxonomy. It has been shown within the context of this research that either taxonomy could have been used, but for consistency with other studies the taxonomy of choice was the SOLO Taxonomy.

9.2.2 Validation of Other Studies

While there were many previous studies in the field of Computer Science that explored the use of educational taxonomies in various ways, the one major study that became the platform from which to launch this research was that conducted by Brabrand and Dahl in Denmark, in which they proposed a method to quantitatively establish both a metric and a distribution of results based on the application of numeric values to the subject learning outcomes according to classification against the SOLO Taxonomy. In that work they defined the terms *SOLO Average* and *SOLO Distribution* and introduced a technique called the *double-weight averaging scheme*.

The approach was applied to the subjects of a particular degree in the Australian context to confirm the method described as being both workable and applicable. Following the validation of the approach, the subsequent steps were to shift from a broad discipline oriented focus to a narrower course oriented focus to investigate what the analytical techniques might reveal, and then to extend that analysis across several courses. The outcome from those studies was that a method had been developed to compare degree courses at the specification level for similarity or difference.

Although there was a difference in the focus of the research, the underlying techniques that were adopted were based on the Brabrand and Dahl methods, and there were sufficient similarities in results to validate the methodology and also the observation of progression as year-levels increase.

9.2.3 Creation of a Course Metric

Repeated application of the analytical techniques led to the formulation of a working model that enabled the creation of new tools for making course comparisons. In particular, new concepts that have now been defined are the **p-index**, a year-level weighted SOLO Average score, the **C-Index**, a degree program metric, and the **Course Profile**, a graphical representation of the SOLO Distributions across the various year levels of the course.

As highlighted in the discussions section (Chapter 8), the C-Index represents a quantitative value that can be used as a metric to indicate the overall level of academic rigour that may be expected in a degree program. Its merit is primarily as a guidance metric that proposes to suggest the depth of learning that a student should be prepared for in that degree.

In the development of the working model for this research it became clear that an important outcome of the research was the year-level profile that describes the shift in emphasis of learning requirements as students progress through their studies. In the domain area researched, namely the field of Information Technology and Computer Science, there was a marked shift away from the low-level superficial coverage of material in first year to the higher level quantitative and qualitative requirements in later years. Many of the degree programs examined have specified similar proportions of learning outcomes at Level 3 and Level 4 of the SOLO Taxonomy, and the main point of difference appeared to be in the proportion of Level 5 requirements. The matter of two approaches to determining the profile distribution, namely the proportional method and the simplex method, was discussed, and in this research the simplex method was used to prepare course profiles. Subsequent extensions to this research could focus on exploring those two methods to determine the best approach that becomes most widely applicable and generally accepted.

9.2.4 Course Internal Quality Control

While the course profile, at either the overall level or the more detailed year level, was an important contribution, the additional data analysis on the base-level data demonstrated the variability in subject SOLO Scores. After considering a number of statistical measures and graphical representations of the data, the view that the differences in learning rigour expectations were of particular interest from a strategic management perspective emerged. Looking at the year-level score for a course, it was proposed that this value should be representative of the learning rigour for the subjects in that course at that year level. The corollary to that proposition was that subjects whose SOLO Score was too far away from that representative value, by virtue of being well above or well below, should be candidates for review as they may be expecting too much from students or conversely not expecting enough for that year level.

The distance measure used was the standard deviation value, which is a statistical measure of spread for a set of data. In particular the boundary limits chosen were defined as two standard deviations either side of the year-level score, where the standard deviations were calculated for each year level, and these were labelled as the **Control Limit Boundaries**. Those subjects whose SOLO Score fell outside the Control Limit Boundaries were identified as **Subjects of Interest**, and therefore became potential candidates for review. Equally, subjects that lay very near the Control Limit Boundaries were also recommended to be considered as potential candidates for review.

The information conveyed in the Internal Quality Control graphics should be of great benefit to University Curriculum Committees and Course Standards groups as they move towards meeting the increased standardisation and accountability initiatives in the university sector.

9.2.5 Course Benchmarking

While the initial part of this research was conducted at a single University, and the analyses were done on several different degree courses, the secondary part of the research was to investigate the applicability of the method on a broader scale. Given that the underlying base technique had been adopted from an overseas environment and validated in the local context, it should have been the case that broader application would be successful.

It has been shown in Chapter 7 that the techniques used in this research are applicable to the Australian context. The limiting factors have been found to be in the complexity of course structures, the level of public availability of detailed course information, and the quality of the statements of behavioural objectives and learning outcomes. In some cases the rules for the course structure were not sufficiently clear, and in other cases the underlying details of individual subject content such as the behavioural objectives or learning outcomes were not able to be retrieved unless one was an enrolled student. There were difficulties in interpretation, and therefore classification, arising from either vague or ambiguous statements in the learning outcomes, or mixed statements that confused content matter with learning outcome. Courses which were specific in nature, such as a Bachelor of Computer Science, were sufficiently distinct to allow an analysis to occur, whereas generic degrees such as a Bachelor of Arts or Bachelor of Science would be much more difficult upon which to attempt this type of analysis.

However, despite the difficulties mentioned, for the courses where there was sufficient clarity of course structure and availability of behavioural objectives and learning outcomes, the methodology and techniques have been effective.

With regard to the use of the Course Profile and C-Index as aids to benchmarking of courses, the limited range of the C-Index scores between 2 and 5 makes fine-grained interpretive comparisons somewhat difficult. Accordingly the concept of the **Scaled C-Index** was discussed in Chapter 8 to convert the range into a more easily interpreted range between 1 and 100. Importantly the scaling algorithm was designed to provide a breakpoint at 50 to identify the *surface learning* scores at 50 or below, and the *deep learning* scores at above 50. Hence, the interpretation for benchmarking purposes is that the further past 50 a course scores, then the greater amount of deep learning is being prescribed for the course. There is an implicit suggestion that the more deep learning a course provides then the better the 'quality' of the course. This re-scaling of the C-Index was not tested on academics during this research, but has been put forward as a proposition only. It therefore remains as a research question in any future research arising from this thesis.

9.3 Future Work

This research has provided a valuable foundation for future work in the evaluation of degree programs using a taxonomic analysis of subject behavioural objectives and learning outcomes. In this thesis the particular domain area field of study was that of Computer Science and Information Technology, and the methodology proposed, the techniques used and the model for analysis that was developed have been confirmed at the proof of concept level. Future works in this area that are seen as both viable and important are to extend the study base more widely at both national and international levels, and to investigate the applicability of the techniques to other domain areas.

A useful by-product of this research would be for universities to re-visit their statements of learning outcomes to more closely align the student learning requirements with more precise language that is indicative of the intention of the subjects in question, and which gives a clearer statement in the language of contemporary educational taxonomies. A potential result of such actions would be that a more objective analysis of course information could be undertaken and the information could be presented in both the course profile graphic form and the C-Index metric form, which could then appear in the relevant course marketing materials.

There are clear messages about the reviewing of learning outcome statements for universities wishing to evaluate and improve the overall quality of their courses. In domain areas where the accreditation of courses by external professional bodies and associations is a significant feature of their status, the development of improved learning outcome statements should contribute to a more straight-forward benchmarking process, and the documentation of C-Indices and Course Profiles should facilitate this by providing an indication of what the accreditation assessors can expect from the courses being evaluated. These tools will never be stand-alone elements for course evaluation as they are simply measures of intent, offering promises of the level of educational rigour that should be present in the particular course. What is actually delivered by the educational institution and the learning that is undertaken by students in the course is another aspect that course evaluators need to assess in order to arrive at an overall view of course quality, but having instruments such as the C-Index and Course Profile will enable them to make a judgement about whether what is promised is being delivered.

Appendix A

Flinders University Course Data

The appendix sections contain the summarised data obtained in this research after the individual subject evaluations have been completed. The data shown in this Appendix includes the degree course data for the three courses examined at Flinders University.

1. The Bachelor of Information Technology – BInfoTech
2. The Bachelor of Computer Science – BCompSc
3. The Bachelor of Engineering (Software) – BEng(SW)

A.1 Appendix A1 - Flinders University BInfoTech

The Bachelor of Information Technology data in this appendix contains the subject evaluation data using the SOLO Taxonomy scores, the Revised Bloom Taxonomy scores and the Adjusted Bloom Taxonomy scores as discussed in Chapter 4, and the individual subject weights.

Table A.1: Flinders BInfoTech Data Analysis

Subject Code	Weight	SOLO Score	Revised Bloom	Adjusted Bloom
COMP1001	0.125	3.26	9.80	3.34
COMP1101	0.125	3.20	10.03	2.90
COMP1102	0.125	3.50	10.07	4.00
COMP1111	0.125	3.67	10.75	3.56
COMP1401	0.125	3.13	11.67	3.98
STAT1412	0.125	3.67	12.00	3.67
Core Topic Average		3.43	10.72	3.57
Elective Yr 1*	0.250	3.43	10.72	3.57
COMP2731	0.125	3.25	9.75	3.25
COMP2741	0.125	3.63	9.375	3.13
COMP2761	0.125	3.42	10.06	3.75
COMP2772	0.125	3.92	12.50	4.50
ENGR2792	0.125	3.50	11.21	3.71
Core Topic Average		3.54	10.58	3.67
BUSN3027	0.125	3.78	11.78	3.44
COMP2762	0.125	3.63	11.00	3.25
Selective Average	0.125	3.71	11.39	3.35
Elective Yr 2*	0.250	3.54	10.58	3.67
COMP3721	0.125	4.00	11.62	4.36
COMP3732	0.125	4.13	10.13	4.25
COMP3751	0.125	3.70	14.17	4.17
COMP3771	0.125	3.39	15.25	3.67
ENGR3704	0.125	3.92	9.86	3.92
Core Topic Average		3.83	12.20	4.07
COMP3782	0.125	3.76	11.13	3.79
Upper level topic	0.125	3.83	12.20	4.07
COMP3792	0.250	4.08	9.50	4.00
Selective Average	0.250	3.94	10.58	3.93
Elective Yr 3*	0.125	3.83	12.20	4.07

A.2 Appendix A2 - Flinders University BCompSc

The accompanying table shows the evaluation data for the Bachelor of Computer Science course using the SOLO Taxonomy only. The individual subject weights are also listed.

Table A.2: Flinders BCompSc Data Analysis

Subject Code	Weight	SOLO Score
COMP1001 Fundamentals of Computing	0.125	3.43
COMP1101 Information and Communications Technology 1A	0.125	3.11
COMP1102 Computer Programming 1	0.125	3.57
COMP1401 Professional Skills in Computing	0.125	3.27
MATH1121 Mathematics 1A	0.125	3.67
Core Topic Average		3.41
MATH1122 Mathematics 1B	0.125	3.64
STAT1412 Data Analysis Laboratory	0.125	3.80
Selective Average	0.125	3.72
Elective Yr 1*	0.250	3.41
COMP2711 Computer Programming 2	0.125	3.54
COMP2731 Software Engineering 1	0.125	3.20
COMP2761 Database and Conceptual Modelling	0.125	3.45
COMP2762 Operating Systems	0.125	3.67
COMP2781 Computer Mathematics	0.125	3.88
COMP3712 Computer Programming 3	0.125	4.14
ENGR2782 Computer Networks	0.125	3.75
ENGR2792 Software Engineering 2	0.125	3.40
Core Topic Average		3.63
COMP3751 Interactive Computer Systems	0.125	3.67
COMP3771 Advanced Database	0.125	3.30
COMP3772 Computer Science Project	0.125	3.69
ENGR3704 Project Management for Engineering and Science	0.125	3.90
Core Topic Average		3.64
COMP3721 Enterprise Information Security	0.125	4.00
COMP3722 Theory and Practice of Computation	0.125	4.18
COMP3732 Enterprise Systems	0.125	4.00
COMP3742 Intelligent Systems	0.125	4.21
COMP3752 Computer Game Development	0.125	3.83
ENGR2711 Engingeering Mathematics	0.125	3.50
ENGR2721 Microprocessors	0.125	3.77
ENGR3701 Computer Organisation and Design	0.125	3.94
ENGR3791 Software Engineering 3	0.125	3.75
Selective Average	0.500	3.91

A.3 Appendix A3 - Flinders University BEng(SW)

This appendix shows the results of the evaluation of the subjects in the Bachelor of Engineering (Software) course using the SOLO Taxonomy for scoring.

Table A.3: Flinders BEng(SW) Data Analysis – years 1 and 2

Subject Code	Weight	SOLO Score
COMP1001 Fundamentals of Computing	0.125	3.43
COMP1102 Computer Programming 1	0.125	3.57
ENGR1201 Digital Electronics 1	0.125	3.47
ENGR1202 Analog Electronics 1	0.125	3.67
ENGR1401 Professional Skills for Engineers	0.125	3.32
MATH1121 Mathematics 1A	0.125	3.67
MATH1122 Mathematics 1B	0.125	3.64
Core Topic Average		3.54
PHYS1332 Engineering Physics 1	0.125	3.50
STAT1412 Data Analysis Laboratory	0.125	3.80
Selective Average	0.125	3.65
COMP2731 Software Engineering 1	0.125	3.20
COMP3712 Computer Programming 3	0.125	4.14
ENGR2701 Engineering Programming	0.125	3.54
ENGR2711 Engineering Mathematics	0.125	4.00
ENGR2792 Software Engineering 2	0.125	3.40
Core Topic Average		3.66
Stream 1 – Electronics		
ENGR2712 Electronic Design and Automation	0.125	3.85
ENGR2721 Microprocessors	0.125	3.77
ENGR2722 Signals and Systems	0.125	3.88
Stream 1 Average	0.375	3.83
Stream 2 – Computing		
COMP2761 Database and Conceptual Modelling	0.125	3.45
COMP2762 Operating Systems	0.125	3.67
ENGR2782 Computer Networks	0.125	3.75
Stream 2 Average	0.375	3.62
Stream Average	0.375	3.73

Table A.4: Flinders BEng(SW) Data Analysis – years 3 and 4

Subject Code	Weight	SOLO Score
COMP2781 Computer Mathematics	0.125	3.88
ENGR3704 Project Management for Engineering and Science	0.125	3.90
ENGR3791 Software Engineering 3	0.125	3.75
Core Topic Average		3.89
ENGR3700 Engineering Practicum	0.375	4.18
ENGR3710 International Engineering Practicum	0.375	4.00
Selective Average	0.375	4.09
Stream 1 – Electronics		
COMP2761 Database and Conceptual Modelling	0.125	3.45
ENGR3701 Computer Organisation and Design	0.125	3.94
Stream 1 Average	0.250	3.70
Stream 2 – Computing		
COMP3751 Interactive Computer Systems	0.125	3.67
COMP3771 Advanced Database	0.125	3.30
Stream 2 Average	0.250	3.48
Stream Average	0.250	3.73
ENGR4710A Engineering Project	0.125	4.43
ENGR4710B Engineering Project	0.125	4.00
ENGR4791 Software Engineering 4	0.125	3.58
Core Topic Average		4.00
COMP4701 Advanced Enterprise Security	0.125	4.33
COMP4702 Computer Supported Cooperative Work and Groupware	0.125	3.89
COMP4706 Advanced Conceptual Modelling and Knowledge Engineering	0.125	4.00
COMP4707 Advanced Data Mining	0.125	4.00
COMP4709 Computational Genomics	0.125	4.20
COMP4712 Embodied Conversational Agents	0.125	4.00
COMP4716 Information Retrieval and Text Processing	0.125	3.85
COMP4720 Advanced Studies in Computer Science	0.125	4.00
ENGR4708 Scalable Computing	0.125	4.00
ENGR4720 Advanced Studies in Engineering	0.125	3.83
ENGR4761 Image Processing	0.125	3.67
ENGR4742 Standards, Ethics and Compliance	0.125	4.17
Selective Average	0.500	3.99
Elective Yr 4*	0.125	4.00

Appendix B

Other Australian Universities Course Data

The data shown in this Appendix includes the degree course data for selected Australian Universities offering a Bachelor of Information Technology. The selected universities were:

1. The Bachelor of Information Technology – Swinburne University of Technology
2. The Bachelor of Information Technology – University of Queensland
3. The Bachelor of Information Technology – Newcastle University

B.1 Appendix B1 - Swinburne University BIT

Table B.1: Swinburne BIT Data Analysis – year 1

Subject Code	Weight	SOLO Score
Semester 1		
HBC110N Accounting for Managers	0.125	2.94
HIT1401 Introduction to Business Information Systems	0.125	2.71
HIT1402 Database Analysis and Design	0.125	3.13
HIT1403 ICT Environments	0.125	3.36
Choose one of:		
HIT1301 Algorithmic Problem Solving, or	0.125	3.34
HIT1404 Introduction to Programming in .NET	0.125	3.65
Selective 1 Average	0.125	3.50
Semester 2		
HBSH100 Behaviour and Communication in Organisations	0.125	3.50
HIT2405 Requirements Analysis and Modelling	0.125	3.52
HIT2422 Database Systems	0.125	3.38
HIT2416 Enterprise Systems	0.125	3.38
Choose one of:		
HIT2302 Object-Oriented Programming or	0.125	3.35
HIT2425 Business Systems Programming in .NET	0.125	3.30
Selective 2 Average	0.125	3.33
Summer Semester		
HIT3407 Information Systems Project Management	0.125	3.31
HIT3405 Business Process Modelling	0.125	3.25
Choose one of:		
HIT3408 Information Systems Risk and Security or	0.125	3.57
HIT3413 Business Intelligence	0.125	3.48
Selective 3 Average	0.125	3.52
Core Subject Average		3.25
First Year Score		3.30

Table B.2: Swinburne BIT Data Analysis – years 2 and 3

Subject Code	Weight	SOLO Score
Stage 2		
Semester 1		
HIT2414 Mobile Business and Connectivity	0.125	3.69
HIT3410 Systems Acquisition and Implementation Management	0.125	3.24
Elective	0.125	3.33
Elective	0.125	3.33
And choose one of:		
HIT2037 Software Development in Java or	0.125	3.63
HIT3303 Data Structures and Patterns or	0.125	3.86
HIT3304 Database Programming or	0.125	3.41
HIT3119 Enterprise Java or	0.125	3.83
HIT3421 Database Implementation or	0.125	3.39
HIT3412 Business Information Systems Analysis	0.125	3.64
Selective Average	0.125	3.63
Semester 2		
HIW051 Industry-Based Learning	0.500	3.67
Summer Semester		
HBSH200 Organisation Behaviour	0.125	4.00
Second Year Score		3.58
Stage 3		
Semester 1		
HIW052 Industry-Based Learning	0.500	3.67
Semester 2		
HIT3424 Information Systems Management	0.125	3.50
HIT3044 Professional Issues in Information Technology	0.125	3.70
Elective	0.125	3.41
Elective	0.125	3.41
And choose one of:		
HIT3416 Industry Project (Analytical) or	0.125	3.69
HIT3427 Configuring Business Information Systems Solutions	0.125	3.56
Selective Average		3.62
Core Subject Average		3.62
Third Year Score		3.59

B.2 Appendix B2 - University of Queensland BInfTech

Table B.3: University of Queensland BInfTech Data Analysis

Subject Code	Weight	SOLO Score
First Year Subjects		
CSSE1001 Introduction to Software Engineering I	0.125	3.00
DECO1100 Design Thinking	0.125	3.44
DECO1400 Introduction to Web Design	0.125	3.00
DECO1800 Design Computing Studio I – Interactive Technology	0.125	3.27
INFS1200 Introduction to Information Systems	0.125	3.06
MATH1061 Discrete Mathematics	0.125	3.29
Second Year Subjects		
DECO2800 Design Computing Studio 2 – Testing & Evaluation	0.125	3.25
CSSE2002 Programming in the Large	0.125	3.33
INFS2200 Relational Database Systems	0.125	3.30
COSC2500 Numerical Methods in Computational Science	0.125	2.97
CSSE2010 Introduction to Computer Systems	0.125	3.13
CSSE2310 Computer Systems Principles and Programming	0.125	3.33
DECO2200 Graphic Design	0.125	3.83
DECO2300 Digital Prototyping	0.125	3.63
DECO2500 Human-Computer Interaction	0.125	3.38
SCIE2100 Introduction to Bioinformatics	0.125	2.96
Third Year Subjects		
DECO3800 Design Computing Studio 3 – Proposal	0.125	3.69
DECO3801 Design Computing Studio 3 – Build	0.125	3.83
CSSE3006 Special Projects in Comp. Systems and S/W Eng.	0.250	3.60
COMP3301 Operating Systems Architecture	0.125	4.00
COMP3506 Algorithms & Data Structures	0.125	3.67
COMP3702 Artificial Intelligence	0.125	3.75
COMS3000 Information Security	0.125	3.88
COMS3200 Computer Networks I	0.125	3.36
COSC3000 Visualization, Computer Graphics & Data Analysis	0.125	3.17
COSC3500 High-Performance Computing	0.125	3.58
CSSE3002 The Software Process	0.125	3.42
DECO3500 Social & Mobile Computing	0.125	3.95
DECO3850 Physical Computing & Interaction Design Studio	0.250	4.00
INFS3200 Advanced Database Systems	0.125	4.25
INFS3202 Web Information Systems	0.125	4.00
INFS3204 Service-Oriented Architectures	0.125	4.25
INFS3222 Systems Analysis & Design	0.125	3.50
MATH3201 Scientific computing: adv. techniques and app...	0.125	3.69
MATH3202 Operations Research & Mathematical Planning	0.125	3.63
MATH3302 Coding & Cryptography	0.125	3.54

B.3 Appendix B3 - University of Newcastle BIT

Table B.4: University of Newcastle BIT Data Analysis

Subject Code	Weight	SOLO Score
First Year Subjects		
CORE		
COMP1050 Internet Communications	0.125	3.13
INFT1001 Foundations of Information Technology	0.125	3.21
INFT1004 Introduction to Programming	0.125	3.25
DIRECTED		
INFO1010 Introduction to Information Systems and Technology	0.125	2.75
GENG1003 Introduction to Procedural Programming	0.125	3.42
MATH1510 Discrete Mathematics	0.125	3.03
SENG1110 Introduction to Software Engineering 1	0.125	3.39
SENG1120 Introduction to Software Engineering 2	0.125	3.25
Second Year Subjects		
CORE		
INFT2009 Systems Modelling	0.125	3.40
INFT2031 Systems & Network Administration	0.125	3.50
INFT2040 Database Management Systems	0.125	3.25
DIRECTED		
SENG2050 Introduction to Web Engineering	0.125	3.43
DESN2270 Web Multimedia	0.125	3.75
Third Year Subjects		
CORE		
INFT3100 Project Management	0.125	4.00
INFT3920 Contemporary Issues in Information Technology	0.125	4.30
INFT3970 IT Major Project	0.250	3.64
DIRECTED		
COMP3260 Data Security	0.125	4.00
INFT3007 The Information Resource	0.125	3.63
INFT3940 Information Technology Applications	0.125	3.83
INFT3960 Games Production	0.125	3.42
SENG3130 Software Architecture & Quality Management	0.125	3.44
SENG3300 User Interface Design	0.125	3.17
SENG3400 Network and Distributed Computing	0.125	3.67

Bibliography

ACM/IEEE (2013). *Computer Science Curricula 2013.* ACM IEEE, New York.

Anderson, L. W. and Krathwohl, D. R. (2001). *A Taxonomy for Learning, Teaching and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives.* Addison Wesley Longman, New York, abridged edition.

Australian Education Network (2014). Groupings of Australian Universities. *http://www.australianuniversities.com.au/directory/australian-university-groupings/*, Accessed 27 May, 2014.

Berggren, K.-F., Brodeur, D., Crawley, E. F., Ingemarsson, I., Litant, W. T., Malmqvist, J., and Östlund, S. (2003). CDIO: An international initiative for reforming engineering education. *World Transactions on Engineering and Technology Education*, 2(1):49–52.

Biggs, J. (1979). Individual Differences in Study Processes and the Quality of Learning Outcomes. *Higher Education*, 8(4):381–394.

Biggs, J. (2011). Constructive Alignment. http://www.johnbiggs.com.au/constructive_alignment.html, Accessed: Sept 2011.

Biggs, J. and Tang, C. (2007). *Teaching for Quality Learning at University.* SRHE and Open University Press, Maidenhead, UK, 3rd edition.

Biggs, J. B. (1999). *Teaching for Quality Learning at University.* SRHE and Open University Press, Buckingham, UK.

Biggs, J. B. and Collis, K. F. (1982). *Evaluating the Quality of Learning: The SOLO Taxonomy.* Academic Press, New York.

Bloom, B., Engelhart, M., Furst, E., Hill, W., and Krathwohl, D. (1956). *Taxonomy of educational objectives: the classification of educational goals; Handbook 1: Cognitive Domain*. Longmans Green, New York.

Bower, M. (2008). A Taxonomy of Task Types in Computing. In *ITiCSE'08*, pages 281–285. ACM.

Brabrand, C. and Dahl, B. (2007). Constructive Alignment and the SOLO Taxonomy: A Comparative Study of University Competences in Computer Science vs Mathematics. In Lister, R. and Simon, editors, *Seventh Baltic Sea Conference on Computing Education Research (Koli Calling 2007)*, volume 88. Australian Computer Society CRPIT.

Brabrand, C. and Dahl, B. (2009). Using the SOLO Taxonomy to Analyze Competence Progression of University Science Curricula. *Higher Education*, 58(4):531–549.

Flinders University (2012a). Bachelor of Computer Science Course Rule. *http://www.flinders.edu.au/courses/rules/undergrad/bcsc.cfm*, Accessed Sept. 2012.

Flinders University (2012b). Bachelor of Engineering (Software) Course Rule. *http://www.flinders.edu.au/courses/rules/undergrad/bengs.cfm*, Accessed Sept. 2012.

Flinders University (2012c). Bachelor of Information Technology Course Rule. *http://www.flinders.edu.au/courses/rules/undergrad/bit.cfm*, Accessed Sept. 2012.

Fuller, U., Johnson, C. G., Ahoniemi, T., Cukierman, D., Hernan-Losada, I., Jackova, J., and et al. (2007). Developing a Computer Science-specific Learning Taxonomy. *SIGCSE Bulletin*, 39(4):152–170.

Garcia, E. (2012). C-Indices and Measures of Associations. *http://www.miislita.com/semantics/c-index-2.html*, Accessed 7 Sept. 2012.

Gluga, R., Kay, J., Lister, R., Kleitman, S., and Lever, T. (2012a). Coming to terms with Bloom: an online tutorial for teachers of programming fundamentals. In *Fourteenth Australasian Computing Education Conference (ACE2012)*, volume 123, pages 147–156. ACS.

Gluga, R., Kay, J., Lister, R., Kleitman, S., and Lever, T. (2012b). Over-confidence and confusion in using Bloom for programming fundamentals assessment. In *Proceedings of the 43rd ACM technical symposium on Computer Science Education*, pages 147–152. ACM.

Great Schools Partnership (2014). Glossary of Education Reform. *http://edglossary.org/rigor*, Accessed 22 Sept. 2014.

Harvey, L. and Williams, J. (2010). Fifteen Years of Quality in Higher Education. *Quality in Higher Education*, 16(1):3–36.

He, L. and Brandt, P. (2007). WEAS: A Web-based Educational Assessment System. In *ACM-SE45. Proceedings of the 45th Annual Southeast Regional Conference*, pages 126–131.

Hirsch, J. (2005). An index to quantify an individual's scientific research output. *Proceedings of the National Academy of Sciences*, 102(46):16569–16572.

Indexing Research (2012). Cindex. *http://www.indexres.com/home.php*, Accessed 7 Sept. 2012.

Johnson, C. G. and Fuller, U. (2006). Is Bloom's Taxonomy Appropriate for Computer Science? In *6th Baltic Sea Conference on Computing Education Research*, pages 120–123. University of Eastern Finland.

Jonassen, D. H. (1999). Designing constructivist learning environments. *Instructional design theories and models: A new paradigm of instructional theory*, 2:215–239.

Jordens, J. Z. and Zepke, N. (2009). A Network Approach to Curriculum Quality Assessment. *Quality in Higher Education*, 15(3):279–289.

Keane, E. and Labhrainn, I. (2005). Obtaining student feedback on teaching & course quality. *Briefing paper*, 2:1–19.

Killen, R. (2005). *Programming and Assessment for Quality Teaching and Learning.* Thomson, Southbank, Victoria.

Krathwohl, D. R. (2002). A Revision of Bloom's Taxonomy: An Overview. *Theory into Practice*, 41(4):212–218.

Kuei, C.-h. and Lu, M. H. (2013). Integrating quality management principles into sustainability management. *Total Quality Management & Business Excellence*, 24(1-2):62–78.

Lahtinen, E. (2007). Categorization of Novice Programmers: A Cluster Analysis Study. In *19th Annual Workshop of the Psychology of Programming Interest Group*, pages 32–41. University of Joensuu.

Lister, R., Simon, B., Thompson, E., Whalley, J. L., and Prasad, C. (2006). Not Seeing the Forest for the Trees: Novice Programmers and the SOLO Taxonomy. In *The 11th Annual SIGCSE Conference on Innovation and Technology in Computer Science Education*, pages 118–122.

Liu, O. L. (2009). Measuring Learning Outcomes in Higher Education. *R & D Connections*, 10:RDC10.

Liu, Y. (2003). Improving online interactivity and learning: A constructivist approach. *Academic Exchange Quarterly*, 7(1):174–178.

McLean, B. and Sachs, J. (2006). Implementing and Sustaining Quality Enhancement in a devolved structure. In *Proceedings of the Australian Universities Quality Forum 2006*, pages 101–105.

Meerbaum-Salant, O., Armoni, M., and Ben-Ari, M. (2010). Learning Computer Science Concepts from Scratch. In *Proceedings of the Sixth International Workshop on Computing Education Research*, pages 69–76.

Melrose, M. (1998). Exploring Paradigms of Curriculum Evaluation and Concepts of Quality. *Quality in Higher Education*, 4(1):37–43.

Moore, D. S. and McCabe, G. P. (2003). *Introduction to the Practice of Statistics*. W.H. Freemand and Company, New York, fourth edition.

Oliver, B., Jones, S., Ferns, S., and Tucker, B. (2007). Mapping curricula: ensuring work-ready graduates by mapping course learning outcomes and higher order thinking skillls. In *Australian Technology Network Evaluation and Assessment Conference*. ATN.

Oliver, D., Dobele, T., Greber, M., and Roberts, T. (2004). This Course Has a Bloom Rating of 3.9. In *Sixth Australasian Computing Education Conference*, volume 30, pages 227–231. ACS.

Pintrich, P. R. (2002). The Role of Metacognitive Knowledge in Learning, Teaching, and Assessing. *Theory into Practice*, 41(4):219–225.

Schmidt, A. and Winterhalter, C. (2004). User Context Aware Delivery of e-Learning Material: Approach and Architecture. *Journal of Universal Computer Science*, 10(1):38–46.

Scott, T. (2003). Bloom's Taxonomy Applied to Testing in Computer Science Classes. *Journal of Computing Sciences in Colleges*, 19(1):267–274.

Selvanathan, A., Selvanathan, S., Keller, G., and Warrack, B. (2007). *Australian Business Statistics*. Thomson, South Melbourne, fourth edition.

Sheard, J., Carbone, A., Lister, R., Simon, B., Thompson, E., and Whalley, J. L. (2008). Going SOLO to Assess Novice Programmers. In *ITiCSE'08*, pages 209–213.

Sitthiworachart, J. and Joy, M. (2004). Using Web-based Peer Assessment in Fostering Deep Learning in Computer Programming. In *International Conference on Education and Information Systems*.

Slack, F., Beer, M., Armitt, G., and Green, S. (2003). Assessment and Learning Outcomes: The Evaluation of Deep Learning in an On-line Course. *Journal of Information Technology Education*, 2:305–317.

Smith, P. and Martin, D. (2006). Translating Policy into Quality Outcomes Through a Devolved Management Structure: One Faculty's Experience. In *Proceedings of the Australian Universities Quality Forum 2006*, pages 148–153.

Southern Cross University (2014). Graduate Attribute 1: Intellectual Rigour. *http://scu.edu.au/teachinglearning/index.php/116*, Accessed 22 Sept. 2014.

Swinburne University (2012). Bachelor of Information Technology Course Rule. *http://courses.swinburne.edu.au/courses/Bachelor-of-Information-Technology-I050/local*, Accessed 15 November, 2012.

Tam, M. (2001). Measuring Quality and Performance in Higher Education. *Quality in Higher Education*, 7(1):47–54.

Thompson, E., Luxton-Reilly, A., Whalley, J. L., Hu, M., and Robbins, P. (2008). Bloom's Taxonomy for CS Assessment. In *Tenth Australasian Computing Education Conference*. ACS.

Tofallis, C. (2012). A different approach to university rankings. *Higher Education*, 63(1):1–18.

University of Newcastle (2014). Bachelor of Information Technology Course Rule. *http://www.newcastle.edu.au/degrees/bachelor-of-information-technology/handbook*, Accessed 25 October, 2014.

University of Queensland (2014). Bachelor of Information Technology Course Rule. *http://www.uq.edu.au/study/program_list.html?acad_prog=2230*, Accessed 25 September, 2014.

Wang, L. (2011). Adaptation of outcome-based learning in an undergraduate English education programme. *Research in Higher Education Journal*, 12:1–17.

Whalley, J. L., Lister, R., Thompson, E., Clear, T., Robbins, P., Kumar, A., and Prasad, C. (2006). An Australasian Study of Reading and Comprehension Skills in Novice Programmers, using the Bloom and SOLO Taxonomies. In *Eighth Australasian Computing Education Conference*. ACS.

Wilson, B. G. (1996). *Constructivist learning environments: Case studies in instructional design*. Educational Technology.

Yorke, M. (2003). Formative assessment in higher education: Moves towards theory and the enhancement of pedagogic practice. *Higher Education*, 45(4):477–501.

Zairi, M. (2013). The TQM Legacy - Gurus' contributions and theoretical impact. *The TQM Journal*, 25(6):659–676.