Course Proposal Form

Please see Page 2 for instructions on which parts of this form to complete, whom to consult with to avoid unnecessary effort, and where to send the completed form.

Proposer(s): **Date:**

Cover page: Basic permanent course information

Unless otherwise noted, items in this section are entered into EUCLID and **cannot** be changed without creating an entirely new course.

Course Name	Foundations of Natural Language Processing
Course Acronym *(used by the School only, e.g., for the Sortable Course List)*	FNLP
Course Level If the course is **only** available to MSc students, then it must be classed as Postgraduate. All other courses, regardless of level, are Undergraduate.	_X_ Undergraduate __ Postgraduate
Normal Year Taken	__UG1 __UG2 _X_ UG3 __UG4 __UG5 __MSc
Also available in years *[This can be changed later if need be.]*	__UG1 __UG2 _X_ UG3 _X_ UG4 _X_ UG5 _X_ MSc
SCQF Credit Level Level 8 should normally be used for pre-honours courses. Level 10 should normally be used for optional UG3 courses (so UG4 students may also take them) and for courses aimed mainly at UG4 students. Level 11 should be used for courses aimed mainly at MSc students, whether or not UG4 students can also take them.	__7 __8 __9 _X_ 10 __11
SCQF Credit Points	__10 __X20 __40 __60 __80 __Other:
Delivery Location	_X_ Campus __On-line Distance Learning
Course Type	_X_ Standard (default) __Dissertation __Online Distance Learning __Other (specify: Placement, Student Led Individually Created Course, Year Abroad)
Marking Scheme By default, courses use a numerical marking scheme. If you wish to use a grade-only marking scheme, your course proposal below should justify this.	_X_ Standard (numerical) __Letter grade only

Guidance for remaining sections:

For an initial course proposal, please complete the **cover page and Section 1 (Case for Support)**, which asks you to describe the need for this course and to provide an overview of the course design, including the learning outcomes. **Please discuss your plans as early as possible with the head of Curriculum Review to avoid unnecessary effort.**

Send the form with these sections completed to the BoS Academic Secretary and head of Curriculum Review (listed on the BoS page) to obtain their comments before filling out the remainder of the form.

If a full proposal is invited, please complete the remaining sections and send to **iss-bos@inf.ed.ac.uk**.

2. Student-facing course description and additional feedback and assessment information.
This section provides most of the information students see in the DRPS entry for this course, as well as related details for BoS consideration.

3. Further information for BoS consideration: sample materials.

4. Additional Course Details required for DRPS. *[Administrative information such as delivery timing and prerequisites.]*

5. Placement in degree programme tables. *[Required for all level 9-11 courses; used to determine where the course will be added to existing degree programme tables.]*

6. Comments from colleagues. *[All course proposal should be sent to relevant colleagues in the area as well as to the appropriate year organizer and BoS Academic Secretary for comment in good time before the BoS meeting. Use this section to indicate what feedback has been solicited and received.]*

Colour coding and item-by-item guidance:

Guidance is provided in italics for each item. Please also refer to the guidance for new course proposals at http://www.inf.ed.ac.uk/student-services/committees/board-of-studies/course-proposal-guidelines. Examples of previous course proposal submissions are available on the past meetings page http://web.inf.ed.ac.uk/infweb/admin/committees/bos/meetings-directory but note that the proposal form was updated in Jan 2019.

Sections in gold are for student view and are required before a course can be entered into DRPS. You must complete these sections even if your course has already been approved based on other documentation.
Sections in orange are for School use but are still required for all courses (even those that have already been approved based on other documentation).
Section in gray are for consideration by the Board of Studies. They are normally required for all new course proposals but may be omitted in some circumstances (e.g., for invited course proposals) if you obtain permission in advance.

1. Case for support

This section is for consideration by the Board of Studies. The final two boxes (Learning Outcomes, Graduate Attributes) will also go into the student-facing course description.

Overall contribution to teaching portfolio and relation to existing curriculum

Please explain (a) what motivates the course proposal (e.g. a previous course having become outdated/inappropriate, an emergent or maturing research area or new research activity in the School, offerings of our competitors) and (b) how it relates to existing courses and degree programmes (including any prerequisite courses). Every new course should make an important contribution to the delivery of our Degree Programmes.

This proposed course implements a part of the curriculum review of our UGT teaching. As set out in that review, it replaces FNLP (currently 10 points) and the NLP parts of Inf2A (which ceases to exist from 2019/20) with a 20 point course entitled FNLP.

The content of FNLP and ANLP look similar. However, they differ in emphasis and delivery. On emphasis, FNLP focuses more on teaching computer scientists some linguistics and empirical methods in NLP (corpus annotation, metrics of evaluation etc), while ANLP focuses more on how established ML methods apply to NLP. The delivery of FNLP and ANLP are quite different for two reasons. First, ANLP is a level 11 course while the proposed FNLP course is a level 10 course---for example, ANLP incorporates a coursework that is entirely open ended and this isn't appropriate for the 3rd years that would be taking FNLP. Secondly, the students that take ANLP have extremely different backgrounds, making their shared knowledge and skill sets almost zero. It would be highly undesirable for our 3rd years to be a part of this very large heterogeneous student body, and furthermore by keeping these courses separate FNLP is able to cover more content (at an introductory level) than ANLP currently does. Finally, ANLP is delivered in the first semester, but AI UG 3rd year students need a 20 point AI course in the second semester, and the proposed FNLP course fulfils that requirement.

Target audience and expected demand

Describe the type of student the course would appeal to in terms of background, level of ability, and interests, and the expected class size for the course based on anticipated demand. A good justification would include some evidence, e.g. by referring to projects in an area, class sizes in similar courses, employer demand for the skills taught in the course, etc

This course is mainly aimed at 3rd year students on all Informatics degrees, but it's also open to UG4 and MSc students. The current 10 point FNLP course has had 140-160 students in each year for the last 3 years. I would expect that a 20 point FNLP course will recruit fewer students, but it will still be very popular because NLP is popular (as evidenced by the heavy demand for UG4 projects) and FNLP is also a pre-requisite to NLU+.

Anticipated Resource Requirements

Estimate how much lecturing, tutoring, exam preparation and marking effort will be needed in steady state, and any additional resources needed to set the course up initially. Provide estimates relative to class size where applicable and discuss how support staff will be recruited and supervised, if the class is likely to be very large. Please mention any scaling limits due to equipment or space. If equipment is required, say how it will be procured and maintained.]

30 lectures, 5 tutorials, 5 labs, 2 pieces of coursework, both involving programming, 1 exam (plus potentially a resit).

The NLP labs and assignments are built on NLTK. NLTK provides labelled corpora plus pre-processing tools (e.g., tokenisation, dictionaries of closed class words etc), which are essential for implementing non-trivial tasks in NLP (e.g., learning statistical parsers, or implementing a word sense disambiguation model). But NLTK is difficult to use and maintain, and so TA time on this course is essential and probably more time is needed than the School standard for a 20 point.

I would expect that we need 100 hours of TA time, plus tutors and demonstrators for the 5 tutorials and 5 labs. Taking a very rough estimate that this course would attract 110 students (so fewer than the current 10 point version of FNLP, but reasons for lower recruitement are detailed below), we would need tutors for around 10 tutorial groups and 4 lab demonstrators. We would also need markers: at 30 minutes per student per assignment, we would need 110 hours of marking time (and so 5 markers doing 25 hours each, taking into account the time needed to acquaint yourself with the assignment and course content). I would anticipate that the lecturers would both set and mark the exams, which is the current arrangement for FNLP.

Quotas, special arrangements or unusual characteristics

Please specify if this course requires any special arrangements such as quotas or other registration arrangements; is a collaboration with another school or institution, or has other atypical characteristics that may affect finances or student registration. Further justification/information may be requested for such courses.

None

Narrative description of the course aims and structure

Please describe the main goals of the course and how the course design will allow students to achieve those goals. This section should be consistent with the student-facing information provided below, but should provide additional information to help colleagues at BoS understand the vision and structure of the course. This description may refer to the learning outcomes and graduate attributes (next two boxes) and should explain how activities such as tutorials, labs, or in-lecture activities will support them, and how the proposed assessments will assess them.

For courses that are important pre-requisites for other courses, this section may also provide content/syllabus information which is too detailed for the student-facing description, such as a lecture-by-lecture syllabus.

This course covers some of the linguistic and algorithmic foundations of natural language processing (NLP). It applies some of the computational techniques from Algorithms and Data Structures and statistical methods from Data Analysis and Inferential Thinking (or Inf2B in the year 2020/21) to learning models of linguistic phenomena from naturally occurring linguistic training data. The course is strongly empirical, using corpus data to illustrate both core linguistic concepts and algorithms, including language modeling, morphology, part of speech tagging, syntactic processing, the syntax-semantics interface, and aspects of both semantic and pragmatic processing. Linguistic and algorithmic content will be interleaved throughout the course.

Topics to be covered include:

1. Lexicon and lexical processing:

* morphology
* language modeling
* Hidden Markov Models and associated algorithms
* part of speech tagging (e.g., for a language other than English) to illustrate HMMs
* smoothing
* Text classification

2. Syntax and syntactic processing:
* The Chomsky hierarchy
* syntactic concepts: constituency (and tests for it), subcategorization, bounded and unbounded dependencies, feature representations
* Context-free grammars
* lexicalized grammar formalisms (e.g., dependency grammar)
* Chart parsing and dependency parsing (eg, shift-reduce parsing)
* treebanks: lexicalized grammars and corpus annotation
* statistical parsing

3. Semantics and semantic processing:
* Word senses: regular polysemy and the structured lexicon; distributional models; word embeddings (including biases found)
* compositionality, constructing a formal semantic representation from a (disambiguated) sentential syntactic analysis.
* predicate argument structure
* word sense disambiguation
* semantic role labelling
* pragmatic phenomena in discourse and dialogue, including anaphora, presuppositions, implicatures and coherence relations.
* Labelled corpora addressing word senses (e.g., Brown), semantic roles (e.g., Propbank, SemCor), discourse information (e.g., PDTB, STAC, RST Treebank).

4. Data and evaluation (interspersed throughout other topics):
* Cross-linguistic similarities and differences
* Commonly used datasets
* Annotation methods and issues (e.g., crowdsourcing, inter-annotator agreement)
* Evaluation methods and issues (e.g., standard metrics, baselines)
* Effects of biases in data

The tutorials and labs will occur in alternate weeks. Each of these lasts 1 hour. However, the labs are designed so that if the student doesn't finish the exercises in the lab itself, they can be completed in the student's own time on any DICE machine.

Summary of Intended Learning Outcomes (MAXIMUM OF 5)

List the learning outcomes of the course. These must be assessable (i.e., observable), so must specify what the student should be able to do concretely, not simply what they should "understand". Use concrete verbs that indicate (a) what type of assessment would be appropriate, and (b) what level of knowledge/thinking is expected (from recall to analysis to novel creation). ***Example verbs:*** *define, explain, implement, compare, justify. Assessments (described later) should be tied to the learning outcomes.*

Outcomes should typically focus more on the types of thinking/skills developed than on the detailed course content, and the level of thinking should be appropriate to the level of the course: outcomes for a Level 11 course should include more

higher-level thinking skills than for a Level 8 course. Further guidance on writing learning outcomes can be found at https://www.ncl.ac.uk/ltds/assets/documents/res-writinglearningoutcomes.pdf

On completion of this course, the student will be able to:

1) Identify and analyze examples of ambiguity in natural language---ambiguity in part-of-speech, word sense, syntax, semantics and pragmatics. Explain how ambiguity presents a problem for computational analysis and NLP applications and some of the ways it can be addressed (see (2) to (5)).

2) Describe and apply standard sequence models (e.g., HMMs), classification models (e.g., Naïve Bayes, MaxEnt); parsing algorithms (e.g., statistical chart parsing and dependency parsing) for processing language at different levels (e.g. morphology, syntax and semantics), and simulate each algorithm on `toy' linguistic examples step-by-step with pen and paper.

3) Explain and provide examples of how sparse data can be a problem for machine learning in NLP; describe and apply methods for addressing the sparse data problem.

4) Given an appropriate NLP problem, students should also be able to identify suitable evaluation measures for testing solutions to the problem, explain the role of annotated corpora in developing those solutions, and assess and justify which sequence of algorithms are most appropriate for solving the problem, based on an understanding of the algorithms in (2) and (3).

5) Implement parts of the NLP pipeline with the help of appropriate support code and/or tools. Evaluate and interpret the results of implemented methods on natural language data sets.

Graduate Attributes, Personal & Professional Skills

List the personal attributes and generic transferrable skills this course will help develop. Examples include

Cognitive skills: *problem-solving, critical/analytical thinking, handling ambiguity*

Responsibility, autonomy, effectiveness: *independent learning, self-awareness and reflection, creativity, decision-making, leadership, organization and time management, flexibility and change management, ethical/social/professional awareness and responsibility, entrepreneurship*

Communication: *interpersonal/teamwork skills, verbal and/or written communication, cross-cultural or cross-disciplinary communication*

Cognitive skills: critical thinking (via tutorials, labs and assessed work), detecting and handling ambiguity (via the study of linguistic ambiguity in this course).
Responsibility, autonomy and effectiveness: self-awareness and reflection (via acquisition of the skill of *perceiving* linguistic ambiguity that in normal human language processing people don't perceive), independent learning (via the labs, required reading and preparation for tutorials), exploration and testing of evidence towards (or against) a hypothesis (via the labs and tutorials), time management (via coursework).
Communication: written communication.

2. Student-facing course description and additional feedback and assessment information

Except where noted, all fields are required and will go into the DRPS entry for the course (for use by students). Important: any text in DRPS is effectively a contract with students, so should not include details that are likely to change from year to year.

Summary Description *Provide a brief official description of the course, around 100 words. This should be worded in a student-friendly way, it is the part of the descriptor a student is most likely to read. If this course replaces another course, please say so in this summary.*	This course replaces FNLP (10 points). This course covers some of the linguistic and algorithmic foundations of natural language processing. It builds on algorithms and models from DAIT (or Inf2B in the year 2020/21) and Inf2-IADS, and applies these to NLP problems. It also equips students for more advanced NLP courses in year 4. The course is strongly empirical, using corpus data to illustrate both core linguistic concepts and algorithms, including language modeling, part of speech tagging, syntactic processing, the syntax-semantics interface, and aspects of semantic processing. The theoretical study of linguistic concepts and the application of algorithms to corpora in the empirical analysis of those concepts will be interleaved throughout the course.
Keywords *Give a list of searchable keywords.*	Natural language, corpus-based methods, machine learning
Course Description *A more detailed student-facing description of the course, which should normally include (a) a more in-depth academic description of the learning aims, nature and context of the course, (b) a rough outline of the content or syllabus, often as bullet points, and (c) a description of how the course will be taught, how students are expected to engage with their learning and how they will be expected to evidence and demonstrate their achievement of the intended learning outcomes.]*	This course covers some of the linguistic and algorithmic foundations of natural language processing (NLP). It applies some of the computational techniques from Algorithms and Data Structures and statistical methods from Data Analysis and Inferential Thinking (or Inf2B in the year 2020/21) to learning models of linguistic phenomena from naturally occurring linguistic training data. The course is strongly empirical, using corpus data to illustrate both core linguistic concepts and algorithms, including language modeling, morphology, part of speech tagging, syntactic processing, the syntax-semantics interface, and aspects of both semantic and pragmatic processing. Linguistic and algorithmic content will be interleaved throughout the course. Topics to be covered include: 1. Lexicon and lexical processing: * morphology * language modeling * Hidden Markov Models and associated algorithms * part of speech tagging (e.g., for a language other than English) to illustrate HMMs * smoothing

	* Text classification 2. Syntax and syntactic processing: * The Chomsky hierarchy * syntactic concepts: constituency (and tests for it), subcategorization, bounded and unbounded dependencies, feature representations * Context-free grammars * lexicalized grammar formalisms (e.g., dependency grammar) * Chart parsing and dependency parsing (eg, shift-reduce parsing) * treebanks: lexicalized grammars and corpus annotation * statistical parsing 3. Semantics and semantic processing: * Word senses: regular polysemy and the structured lexicon; distributional models; word embeddings (including biases found) * compositionality, constructing a formal semantic representation from a (disambiguated) sentential syntactic analysis. * predicate argument structure * word sense disambiguation * semantic role labelling * pragmatic phenomena in discourse and dialogue, including anaphora, presuppositions, implicatures and coherence relations. * Labelled corpora addressing word senses (e.g., Brown), semantic roles (e.g., Propbank, SemCor), discourse information (e.g., PDTB, STAC, RST Treebank). 4. Data and evaluation (interspersed throughout other topics): * Cross-linguistic similarities and differences * Commonly used datasets * Annotation methods and issues (e.g., crowdsourcing, inter-annotator agreement) * Evaluation methods and issues (e.g., standard metrics, baselines) * Effects of biases in data
Assessment Weightings: *These should correspond approximately to the proportion of learning outcomes that each component assesses. More than 30% coursework requires specific justification.* *The expectation for a 10pt course is 20% coursework with the equivalent of one 15-20hr assessed assignment (but possibly split into smaller pieces). See 'components of assessment' below.*	Written Exam ___70__% Practical Exam _____% (*for courses with programming exams)* Coursework ___30__%

Further Assessment Information *Provide any further information that should go on DRPS for students. E.g., if the assessment includes required group work or if students must pass some individual component of assessment as well as the course overall.*	Tutorials and labs will both consist of exercises, from which the students will receive formative feedback from the tutors and demonstrators.
Components of assessment and time spent on assignments (for BoS only) *If not already included in the course narrative description, please describe the type of assessments (oral presentation, report, programming, etc) and **how each component of assessment will assess the intended learning outcomes**. Where coursework involves group work, it is important to remember that every student has to be assessed individually for their contribution to any jointly produced piece of work.* *Also estimate **how many hours** students will spend on assignments. Please see the School policy on Workload and Assessment, which states that students should not be expected to spend more than 6-7 hrs/wk per 10 credits, including contact hours.* *Note that it often desirable to include formative assignments which are not formally assessed but submitted for feedback, often in combination with peer assessment.*	2 pieces of coursework, worth 15% each, both involving implementing algorithms for processing corpus data that will have been taught via the lectures, tutorials and labs. These two pieces of coursework will in fact build on 2 of the lab exercises. We anticipate that these will take around 15 hours each to complete. I anticipate 3 lectures per week for 10 weeks, and the 5 tutorials and 5 labs in alternate weeks. The student will need to spend about 1-2 hours preparing for each tutorial; the tutorials each consist of exercises to be prepared in advance. More able students may complete the lab exercises within the 1 hour lab session, but they are designed to be doable in one's own time as well. The exam will assess LOs 1, 4, and aspects of 2 (describe/hand-simulate) and 3 (explain/examples). Coursework assignments will assess LO 5, and remaining aspects of LOs 2 and 3.
Feedback Information *Provide a high-level description of how and what type of feedback will be provided to students, for inclusion in DRPS.*	Tutorial exercises will be pen and paper (e.g., using an algorithm to analyze a toy example step by step). Students will prepare answers in advance of the tutorial, and present their analyses and get feedback on it during the tutorial. Labs will consist of doing a small amount of programming, implementing algorithms taught in the lectures, running it on corpora and evaluating the results. Coursework will involve more extensive implementation of the algorithms and models taught in lectures. Feedback will be a raw grade plus qualitative feedback. The exam will be pen and paper, and raw marks will be given.

Additional Feedback Information (for BoS use only)

If not already included in the course narrative, provide further details on planned feedback arrangements. This includes how course feedback is solicited from the class and responded to, as well as what feedback students will get (either on work that contributes to their final mark, or not).

The University is committed to a baseline of principles regarding feedback that we have to implement at every level, and the School encourages submission of at least one piece of written work for formative feedback.

In general, formative feedback:

- *Should say how students can improve.*
- *Need not be on individual work (e.g., consider a lecture or document summarizing common issues.)*
- *Can include oral feedback during labs/tutorials*
- *Can include feedback from peers*
- *Clickers/TopHat/equivalents can provide in-class feedback for both students and lecturer.*
- *Is returned in time for other forms of assessment to which it relates, to allow feedforward.*

Students are expected to attend the lab session, although the lab exercises will be designed to be doable on any DICE machine at any time. The students will get model solutions the week after the lab takes place, so as to compare their work to the model answers. Similarly, model answers to the tutorial exercises will be released the week after the tutorial takes place.

Breakdown of Learning and Teaching Activities

State how many hours students spend on each part of the course. The total should be 10 x course credits, but please also see the School policy on Workload and Assessment which states that students should not be expected to spend more than 6-7 hrs/wk per 10 credits, including contact hours.

Assume 10 weeks of lectures slots and 10 weeks of tutorials, but these need not all be used. As a guideline, a 10-pt course typically has 18-20 lecture hours, but should have only around 15 lectures of examinable material; the rest should be used for guest lectures, revision sessions, introductions to assignments, etc.

Contact hours

Hours	Type
30	Lecture Hours
5	Seminar/Tutorial Hours
0	Dissertation Project Supervision Hours
5	Supervised practical/Workshop/Studio hours
0	Feedback/Feedforward hours
2	Summative assessment hours
	Revision Session Hours

Non-contact hours

Hours	Type
158	Directed Learning & Independent Learning hours

Total hours: 200

Reading List/Learning Resources

You are encouraged to create resource lists using LEGANTO

REQUIRED: Dan Jurafsky and James Martin *Speech and Language Processing* (3rd edition online, and 2009 2nd edition for chapters that aren't yet updated in 3rd edition).

RECOMMENDED: Bird, S., E. Klein and E. Loper, *Natural Language Processing with Python*, (2009) O'Reilly Media.

3. Further information for BoS consideration: sample materials

A full proposal for a new course must include examples of exercises and assessment. Please provide these below, along with publicity information if the course is to be advertised outwith the School.

Course information and publicity *The course web page (typically the Learn landing page) will be linked from the Sortable Course List, and information such as timetables and assignment deadlines must be made available prior to the start of the academic year. Please specify here if any additional info/publicity is needed for your course: typically only if it is aimed largely at non-SoI students.*	
Sample tutorial/lab sheet questions *Provide a list of tutorial questions and answers and/or samples of lab sheets. These need not be fully fleshed out but should indicate what sort of exercises will be provided to help students learn the material.*	See Inf2A tutorials 5, 7 and 8. Eg: http://www.inf.ed.ac.uk/teaching/courses/inf2a/tutorials/2018_inf2a_T05_exercises.pdf
Sample assessment materials *If the course is primarily assessed by **exam**, provide a sample exam question with model answers. Any non-standard exam format must be justified. The online list of past exam papers gives an idea of typical and alternative exam formats: http://www.inf.ed.ac.uk/teaching/exam_papers/.* *If the course is largely or primarily assessed by **coursework**, provide a sketch of a possible assignment with an estimate of effort against each sub-task and a description of marking criteria.*	See existing FNLP coursework and past exams. Eg: https://exampapers.ed.ac.uk/record/101479/1/2018735_INFR09028.pdf The questions on the exam for the 20 point version of FNLP will be similar. I would also like to keep the FNLP exam rubric. It has been really successful in testing the student's knowledge across the FNLP syllabus, and I feel this is important for a Foundations course that is then utilized in more advanced courses (e.g., NLU+).
Any other relevant materials *Include anything else that is relevant, possibly in the form of links. If you do not want to specify a set of concrete readings for the official course descriptor, please list examples here.*	

4. Additional Course Details for DRPS

Except where otherwise noted, these fields are required for entry into EUCLID and will be visible to students in the DRPS entry.

Planned Academic Year of Delivery *(The first year you anticipate the course running, e.g. AY 2019-20)*	2020-2021
Course Organiser *(By default, the course proposer)*	Alex Lascarides
Intended Delivery Period	___Semester1 __x_Semester 2 ___Full Year ___Summer ___Other (please specify):
Timetable considerations/conflicts *For School use. Please specify any constraints to be considered (e.g. overlap of popular combinations, other specialism courses, external courses etc). Include whether the semester delivery is constrained or could be flexible.*	Avoid clashes with other Level 9/10 courses, especially compulsory courses and other AI courses. Delivery must be in S2.
Is this course available to visiting students?	__x__Yes (default) ___No **If no, please provide a justification here:**
Required pre-requisite courses *Use sparingly: these are enforced in PATH and can only be waived by approval from the School's Curriculum Approval Officer. Note that cross-year required pre-requisites may prevent MSc students from registering; consider using recommended pre-requisites or "other requirements" instead.*	___No __x_Yes (please specify full course name(s) and code(s)): Inf2 Introduction to Algorithms and Data Structures (Inf2-IADS); Inf2B (INFR08028) for academic year 2020/21 only, and Data Analysis and Inferential Thinking (DAIT) for academic years 2021/22 onwards.
Recommended pre-requisite courses	__x_No ___Yes (please specify full course name(s) and code(s)):
Required co-requisite courses *Specify any courses that must be taken in parallel with the existing course. Note that this leads to a timetabling constraint that should be mentioned elsewhere in the proposal.*	__x_No ___Yes (please specify full course name(s) and code(s)):

Prohibited Combinations *Specify any courses that may not be taken in combination with the proposed course].*	___No __x_Yes (please specify full course name(s) and code(s)): Accelerated Natural Language Processing (ANLP) (INFR11125), Foundations of Natural Language Processing (INFR09028)
Other Requirements/Additional Information *This information is often used by MSc students and students from other Schools to see if they have appropriate background without having done our School's courses. So please avoid course titles, instead list specific knowledge and skills (such as mathematical concepts, programming ability or specific languages, etc).* *Also list any other constraints on registration, for example: "Only available to 4th Year Informatics students including those on joint degrees." or "This course is open to all Informatics students including those on joint degrees, and to students in the School of Mathematics. Other external students whose DPT does not list this course should seek permission from the course organiser."*	___No __X_Yes (please specify): Open to MSc students, so long as they have not taken ANLP, and they have the following expertise: Understanding of basic probability; e.g., Bayes Rule Familiar with basic computational processes: e.g., recursion, dynamic programming Able to code in Python. Basic knowledge of linguistic categories: e.g., Noun, Verb. Familiar with first order logic.
Visiting Student Pre-requisites	___Same as "other requirements" ___Different than "other requirements" (please specify): Understanding of basic probability; e.g., Bayes Rule Familiar with basic computational processes: e.g., recursion, dynamic programming Able to code in Python. Basic knowledge of linguistic categories: e.g., Noun, Verb. Familiar with first order logic.

Commented [GS1]: This is so ug4 students don't take it twice!

5. Placement in degree programme tables: for level 9-11 courses only

This section is for consideration by the Board of Studies and will be used later by ITO to determine where the course will be added to existing degree programme tables.

Is this course restricted to students on a specific degree? *E.g., some courses are only available to students on a specific CDT or MSc.*	__x__No ___Yes (please specify and provide justification):
Is this course compulsory for students on any degree(s)?	__x__No ___Yes (please specify and provide justification):
Any issues for part-time students? *Normally, part-time students have access to the same courses as full-time students on the equivalent degree. If you anticipate any problems with this, please specify here.*	

For optional courses:
If this course is available but non-compulsory for students on various degrees (most courses), please fill in this section. The choices here determine where the course appears in degree programme tables (DPTs) and the 2-3 character tags are displayed in the Informatics sortable course list.

Should this course be tagged as 'ML' (machine learning foundations and methods)? *Courses with the ML tag are typically very high-demand and most degrees limit the number of ML credits. If your course might appeal to a similar audience but draw off students from these large courses, please select 'no' and choose one of the tags below.*	__x__No __Yes
If you chose 'no', please choose at least one of the following tags… *Ideally, select exactly one, unless there is a good argument for more than one. These three are used in various combinations for many of our degrees.*	__**FSS** (CS foundations, systems, and software) __x__**AIA** (artificial intelligence applications and paradigms) __x__**COG** (cognitive science: including HCI and NLP courses, but not most other AI courses. Please restrict to courses most relevant to natural cognition.)
…and also tick if any of the following tags or categories apply. *Do not tick any of these if you selected 'ML' already.*	__**NS** (natural systems: e.g., computation by or about biological or social systems. Many COG courses are also NS. This tag is mainly relevant for MSc in Informatics.) __**SE** (software engineering: including courses that are highly relevant to SE degrees. All SE courses should also be FSS. This tag is mainly relevant for UG SE degrees.) __Databases and data management systems (used for Data Science MSc and MSc(R)) __x__Unstructured data and applications (used for Data Science MSc and MSc(R)) __Level 11 Security courses (used for Security MSc)

If you are not sure which tags are most appropriate or have other questions about this section, please note any comments/issues here.	

6. Comments from colleagues

All course proposal should be sent to relevant colleagues in the area as well as to the appropriate year organizer and BoS Academic Secretary for comment in good time before the BoS meeting. Please indicate here what feedback has been solicited and received.

Additional Comments *Summarise any comments received from relevant individuals prior to proposing the course. If you have not discussed this proposal with others please note this.*	This proposal has drafted via consultation with Sharon Goldwater, Shay Cohen, Adam Lopez, Mirella Lapata, Frank Keller, Ivan Titov and Henry Thompson. The current draft takes all of their feedback into account.
Year Organiser Comments *Year Organisers are responsible for maintaining the official Year Guides for every year of study, which, among other things, provide guidance on available course choices and specialist areas. The Year Organisers of all years for which the course will be offered should be consulted on the appropriateness and relevance on the course. Issues to consider here include balance of course offerings across semesters, subject areas, and credit levels, timetabling implications, fit into the administrative structures used in delivering that year.]*	
BoS Academic Secretary Comments *Proposals must be checked by the Secretary of the Board of Studies prior to discussion at the actual Board meeting. This is a placeholder for their comments, mainly on the formal quality of the content provided above.*	