

Evolving NEJLT: A proposal

Leon Derczynski
ITU Copenhagen, 16 August 2019
ld@itu.dk

Summary

I propose updating NEJLT to match our field's needs. This will be done by making the Northern European topic secondary; by accepting reviews from *CL conferences into the process; by building an editorial board containing high-profile community members; and by offering a serious NLP journal venue without the embargo of TACL or costly burdens of Elsevier et al.

NEJLT in its current form hasn't worked, so radical and comprehensive changes are one reasonable next step. This document proposes these, aiming to make NEJLT best serve the field of modern NLP.

Motivation

NLP needs a shift in how it publishes. The current situation is that most publications come through conferences, which experience extreme and acute load with little room for scholarly dialogue or for correcting mistakes - and that journals are either highly exclusive, or slow. In the meantime, we still suffer all the problems outlined in Ken Church's "Reviewing the Reviewers" published in CL in 2005, fourteen years ago. I would like to create a journal that is rapid, that is well-tuned to modern publishing in NLP, and is capable of reviewing a diverse range of papers well.

Having journals more available means reducing load on all of us in the community; it serves to soak up the good papers that are otherwise recycled between conference venues. Offering review and publication out of the annual conference cycles also benefits authors, allowing them to publish their work without having to wait - especially important for PhD students - while the extra time means work can be really finished, instead of rushed - important for reviewers and readers. Finally, providing new NLP venues as journals instead of conferences reduces travel, with all the benefits that that brings (cost, time, social, environmental).

NEJLT is a well-indexed journal, which takes time to achieve; at the same time, it is free from the potential constraints of a larger organisation (e.g. ACL) which also having *some* organisation behind it, NEALT, which eases some administration (unlike e.g. ICCL). This makes NEJLT a good frame for a new journal.

I've wanted to put together a journal for our community for some time, and being program chair of COLING last year really crystallised many concepts about how this could become reality in a way that's good for the community. Here I present a concrete and broad proposal for a journal structure, that I have the energy and motivation to carry through.

Scope

The description of the journal's topics will change. The main focus will be identical to that of the ACL conference, with work on Northern European languages added at the bottom. Thus, the journal title becomes simply a regionally-based journal of general scope, rather than a global journal of region-relevant scope (cf. e.g. the APA, which handles all kinds of psychology despite being the American Psychological Association). The journal has been named well, in that we are

afforded this shift without renaming. The Northern European character will be retained by location of publishing offices, and by minimums in editorial board composition (e.g. the journal cannot institute changes which mean that both fewer than 25% and fewer than 10 members of the board have a Northern European affiliation).

The journal will accept full papers and letters.

Marketing

The journal needs to be visible and look serious. This will take some time to build; luckily, we are already indexed, and there are already many serious, visible people behind the journal (i.e. NEALT).

I propose building a standing editorial board of ~20, including high-profile community members. My COLING network and regional network will help with this. To help attract the likely overburdened visible community members, we will set constraints on how much work an editor will be allocated per year. If we find we need to grow our board due to the per-member load constraints, we will probably be busy and visible enough that getting more good editors will be easier.

I see getting high-profile global names as editors as a core part of the journal strategy. Using editors' networks to spread the news and solicit new submissions is a vital part of the plan.

We need a new site, with a redesign and new content. This requires some budget from NEALT.

When things are in place, the journal can be marketed through various channels: the ACL; SIG mailing lists; ELRA; other community lists; by having a maintained social media presence; and by conference sponsorships (budget is required for this, from e.g. NEALT).

High-profile submissions should also be solicited. This could be tough. An ACL endorsement may help, but this might be tough without publication volume.

I expect it to take until 2020Q2 to get ready in stealth mode, while we build a good setup with good people, and then launch a few weeks before ACL reviews are published. There is a chance that this date could be early for ACL2020. First impressions count, so a later launch is preferable to one where things are not yet in place.

Paper format

The journal will accept full papers (from 4000 words to unlimited length) and letters (up to ~1000 words, w/ 2-3 floats).

Papers must be allocated a type, which directly affects how they are reviewed. This explicitly makes space for multiple different kinds of paper. The types considered are:

- Computationally-aided linguistic analysis
- NLP engineering experiment paper
- Reproduction paper
- Resource paper
- Position paper
- Survey Paper

— This inventory can be reviewed periodically. Extended descriptions and review questions for each of these can be found at <https://coling2018.org/paper-types/>.

Papers may include an abstract in a second language after they have been accepted (doing this pre-acceptance means being able to find speakers of that language qualified to review, which is not always possible).

While papers don't have a length limit, they must use their space well, and longer papers must justify their length.

Journal extensions of prior work would not be welcome. They're not a good use of the reviewing pool's time. Rather, work must be new contributions, i.e. no overlapping contributions with other literature (0%), and must not be under review anywhere else for the first three months of review, at the submission's managing editor's discretion. This period is intentionally longer than the 7-week minimum review time to allow for the typical delays, and to make sure authors are serious.

There will be no embargo on rejected papers (cf. TACL, to which ACL/EMNLP/NAACL rejected work may not be submitted for nine months). This gives us an advantage, and helps absorb papers that would otherwise be lightly revised and resubmitted to the next conference, while also giving a route to authors of good but rejected EMNLP papers (EMNLP being last in the cycle).

Letters should make a contribution or talking point; they are somewhere between the CL squib, in that meta-discourse is welcome, and traditional Letters journal articles that make a solid contribution.

Reviewing

Reviews need to be constructive. AEs can help with this. Reviewers can opt in to having their names released with the reviews (though not associated with the particular review they wrote, affording some anonymity). Reviews tend to be better when reviewers know their name will be published with them, but we don't want to push this too much to start with.

A large amount of chairing and reviewing advice was generated in COLING 2018, received with very positive feedback. A small group of reviewers strongly objected to being told that their review quality would be managed, and that they had to be well-behaved; I don't mind leaving those reviewers out of the process. So, these guides and policies would be reused.

I propose fully blind review, and (under the constraints of OJS) having a separate managing editor (ME) and action editor (AE). The ME is aware of the authors' identity, but the AE is not. The AE takes decisions on the manuscript.

We will maintain a large standing pool of reviewers, drawn from the NEALT community and from invitations to those who have reviewed responsively for the field in the past. We expect AEs to be responsive; in return, we limit their load, and will find new AEs if manuscript load increases. The journal should have as much transparency as possible: this means notifying authors when things happen. In customer management, a person who is well-informed will send fewer queries and have a more positive opinion of an organisation, and this human behaviour is likely to extend to this scenario too. Further, because the ME and AE are different people and the authors will only know who the ME is, authors cannot hassle the AE. Rather, this is the ME's job (but one hopes, not required very often at all!).

We give a minimum reviewing time only, of about seven weeks, if everything runs perfectly. The draft process, after receiving a submission, is:

- allocate managing editor (few days)
- ME filters for desk reject, using BMJ procedures as guidelines if needed (few days)
- ME checks paper type and if needed talks with authors about this
- ME waits for an available AE (~few days)
- ME allocates AE and passes on anonymised reviews to AE, if given
- AE accepts (few days)
- AE finds reviewers, keeping the manuscript anonymised (1wk)
- We receive 2 reviews in 1 month (requesting that reviews are done within 2-3 weeks)

- AE maybe finds 1 additional review in 3w, if there were low-quality, late, or non-agreeing reviews
- AE writes meta-review on all reviews available, revealing past reviews to other reviewers (1wk)
- ME contacts authors w. outcome (few days)
- Authors given fixed window to revise (something like 15 weeks for major/minor revisions; no fixed window for accepts)

Further, the journal will allow submission of recent (<1yr) reviews from ACL, EMNLP, NAACL, EACL, AACL, or NeurIPS, in the case of journal articles developed on material rejected from those venues. If authors submit a manuscript with reviews, the original MS and cover letter addressing the reviews must also be supplied, as well as consent for us to verify the reviews with the appropriate conference and for the conference to release the reviews to us. We don't anticipate verifying the reviews - rather, the consent should give a strong enough signal to submitters to not edit their reviews, and means that we are able to check for misuse if we eventually need to. The way the reviews are used is that they may be considered by AEs for the first decision, and are released to reviewers only after the first round of reviewing (details below). This is to avoid biasing our reviewers' initial impression and thus compromising our quality.

Author viewpoint & outcomes

NEJLT is a journal with ACL-quality reviewing and papers, taking COLING2018 style reviewing, where you can send reviews from some top confs along with your revised papers, and that has no ACL embargo.

Papers must reach "minor edits" by the end of the 2nd review cycle.

Possible cycle outcomes are:

- reject (implicit 12-month embargo at NEJLT) - when there are flaws in core/most contribs;
- major edits - flaws in contribs / missing fundamental parts;
- minor edits (<15% of the content needs to change);
- accept / polish+accept (<1% needs to change)

Reviewing takes seven or more weeks, and there are progress updates along the way. We offer full-blind reviewing; nobody associated with acceptance decisions will see an author name.

Implementation

We'll follow the constraints of OJS, reconsidering this if submission volume warrants it (perhaps after one editorial term). That might mean some changes to the proposed procedures, and that's OK. I'd expect about nine months' time required to implement all the changes and to put together to board & marketing, which will also take some budget to achieve - especially a new, serious website.

Publication process

The EiC of NEJLT is also the editor of NEALT, which includes other non-NEJLT work (e.g. NODALIDA proceedings); this may need to be disentangled/refactored, but it makes sense to keep things as they are for now. A good publishing architecture is set up with LiU Press, who I hear are easy to work with and take care of everything. The relationship with LiU Press is mediated through LiU staff; if the LiU connection goes, the relationship with LiU press would have to be renegotiated. Thus, we strongly prefer to keep the LiU connection, perhaps by having a person affiliated with LiU be associated with NEJLT to act as a conduit (we could possibly make this a ~zero-admin position).

Paper type would not be included in the final publication. They are there to encourage good reviewing, rather than to pigeonhole research.

Shifting to biannual issues makes sense for now, while staying at one volume per year.

When the volume of published papers becomes respectable, we would investigate indexing NEJLT in the ACL anthology.

We'll consider adopting / adapting the TACL style files, bearing the proposed "second abstract" feature of NEJLT in mind.

Recognition

Good work should be recognised. Recognising NLP work helps NLP researchers compete. Once volume is there, NEJLT will award the following:

- Annual submission recognition: the best of each paper type, per year
- Periodical submission recognition: the most popular paper every two years (on citation count)
- Reviewer awards: best review per year, and a reward for distinguished contribution (through sustained excellence)

We'll send certificates for each of these.

Further, we're always grateful to reviewers, and so should send them extra thank you notes after three reviews, perhaps with some stickers or some kind of soft marketing material, and acknowledging all reviewers in our annual report (as long as there are more than 15, for privacy).

Policies

We adopt TACL policies in general (and COPE).

We expect the Vancouver authorship policy.

We adopt the ACL arxiv policy in general, though manuscripts published as anonymous preprints (as offered by e.g. OpenReview) are always welcome.