NLPCC-ICCPOL 2016 Shared Task:
Stance Detection in Chinese Microblogs

This task is designed to evaluate stance detection techniques for Weibo texts. The stance detection task aims to automatically determine whether the author of a Weibo text is in favor of the given target, against the given target, or neither. Noted that the given target may not be present in the Weibo text. This means that stance detection is different from traditional target/aspect sentiment analysis.

Consider the target-Weibo pair:

Target: 俄罗斯在叙利亚的空袭行动

Weibo: 9 月 30 日开始至今，俄在叙利亚发起空隙，共死亡 1331 人，其中 403 人是一般的民众, 其中的三分之一是无辜平民陪葬。

It is observed that the author of the Weibo text is against the given target. The aim of stance detection is to evaluate the techniques of detecting the stance of the writer of each Weibo text. Normally, a system needs to identify relevant information that may not be present in the focus text. For example, if one emphasizes the death of civilians, then he or she is likely against the air strikes by Russia. Thus, we provide a domain corpus pertaining to each of the targets, from which systems can gather helpful information for stance detection.

There are two tasks in this evaluation. **Task A is a mandatory task and each participant must submit the results for this task. However, Task B is an optional task and the participants can decide whether to submit the results for this task by themselves.**

Task A (supervised/semi-supervised framework): This task aims to detect stance towards five targets such as "俄罗斯在叙利亚的空袭行动" and "《太阳的后裔》热播". A total of 3,000 labeled instances for all targets will be provided as training data as well as a large amount of unlabeled data.

Task B (unsupervised framework): This task aims to detect stance towards another two targets. No training data will be provided. Instead, a large set of Weibo texts associated with each target, without any stance annotation, will be provided.

The possible stance labels are described as follows:

FAVOR: The author is in favor of the target (e.g., directly or indirectly by supporting someone/something, by opposing or criticizing someone/something opposed to the target, or by echoing the stance of somebody else).

AGAINST: The author is against the target (e.g., directly or indirectly by opposing or criticizing someone/something, by supporting someone/something opposed to the target, or by echoing the stance of somebody else).

NONE: none of the above.

Submission Format: The test data file has the same format as the training file, except for the class label which will be "UNKNOWN" for all instances. The participating system is required to replace "UNKNOWN" with its predicted class to create the submission file.

The format of training data file is:

<ID> <tab> <Target> <tab> <Weibo> <tab> <Stance>

where

<ID> is an internal identification number;

<Target> is the given target;

<Tweet> is a Weibo text;

<Stance> is the stance label.

Evaluation: The macro-average of F-score(FAVOR) and F-

score(AGAINST) is used as the bottom-line evaluation metric.

Each participant is allowed to submit only one running result per task.