



Peer Community Journal

Section: Genomics

RESEARCH ARTICLE

Published
2022-02-04

Cite as

Agnès Baud, Mariène Wan,
Danielle Nouaud, Nicolas
Francillonne, Dominique
Anxolabéhère and Hadi
Quesneville (2022) *Traces of
transposable elements in
genome dark matter co-opted
by flowering gene regulation
networks*, Peer Community
Journal, 2: e14.

Correspondence

hadi.quesneville@inrae.fr

Peer-review

Peer reviewed and
recommended by
PCI Genomics,

[https://doi.org/10.24072/pci.
genomics.100004](https://doi.org/10.24072/pci.genomics.100004)



This article is licensed
under the Creative Commons
Attribution 4.0 License.

Traces of transposable elements in genome dark matter co-opted by flowering gene regulation networks

Agnès Baud¹, Mariène Wan¹, Danielle Nouaud²,
Nicolas Francillonne^{id,3}, Dominique
Anxolabéhère^{2,4}, and Hadi Quesneville^{id,3,1}

Volume 2 (2022), article e14

<https://doi.org/10.24072/pcjournal.68>

Abstract

Transposable elements (TEs) are mobile, repetitive DNA sequences that make the largest contribution to genome bulk. They thus contribute to the so-called “dark matter of the genome”, the part of the genome in which nothing is immediately recognizable as biologically functional. We developed a new method, based on *k-mers*, to identify degenerate TE sequences. With this new algorithm, we detect up to 10% of the *A. thaliana* genome as derived from as yet unidentified TEs, bringing the proportion of the genome known to be derived from TEs up to 50%. A significant proportion of these sequences overlapped conserved non-coding sequences identified in crucifers and rosids, and transcription factor binding sites. They are overrepresented in some gene regulation networks, such as the flowering gene network, suggesting a functional role for these sequences that have been conserved for more than 100 million years, since the spread of flowering plants in the Cretaceous.

¹Université Paris-Saclay, INRAE, URGI, 78026, Versailles, France, ²IJM, Institut Jacques Monod, CNRS UMR 7592, Université Paris-Diderot, Paris, France, ³Université Paris-Saclay, INRAE, Bioinformatics facility, 78026, Versailles, France, ⁴Université Paris-Sorbonne, Paris, France



Introduction

Transposable elements (TEs) are mobile repetitive DNA sequences that make a major contribution to the bulk of the genome in many organisms (Feschotte et al., 2002; Kidwell, 2002; Schnable et al., 2009; Huang et al., 2017; Wicker et al., 2018). They can represent up to 85% of some genomes as in wheat and maize.

TEs invade genomes, through their ability to amplify. However, they are also controlled by their host, through multiple pathways involving RNAi machinery. They invade genomes in a recurrent manner, through bursts of transposition that are rapidly halted by host defense mechanisms. TEs remain quiescent in the genome for long periods of time, until they are reactivated by events such as genomic shocks. TE sequences also accumulate mutations, which may inactivate the sequence by rendering it too degenerate to be functional. The TE sequence thus gradually “blends into” the background genome sequence until it is no longer recognizable. It then contributes to the so-called “dark matter of the genome”, the part of the genome in which nothing is immediately recognizable as biologically functional.

Little is known about the evolution and impact of TE sequences over long periods of time. We explored this question, by developing an innovative repeat annotation approach, which we call *cross-species TE annotation* because it uses closely related species to enhance detection sensitivity for ancient, degenerate repeated sequences (Maumus & Quesneville, 2014). We analyzed the genome of several relatives of *A. thaliana* that diverged about 5-40 million years (My) ago (Huang et al., 2016). We generated a library of consensus repeat sequences that we appended to the *A. thaliana* TE reference library, to compile a “Brassicaceae” library. This compiled TE library was used to annotate the *A. thaliana* Col-0 genome to explore the effects of the long-term presence of TEs on genome evolution. Our Brassicaceae TE annotation, excluding annotations overlapping CDS, covers more than 31.8 Mb (26.7%) of the *A. thaliana* genome, and is highly sensitive for the detection of degenerate TE sequences, as it identified one third more TEs than the current official annotation (Buisine et al., 2008). The detection of many TE copies in *A. thaliana* on the basis of consensus sequences built from sequences in related species provides evidence in support of these *A. thaliana* repeats originating from the common ancestors of these species.

However, our ability to recognize the part of the dark matter derived from TEs remains limited by the sensitivity of current alignment algorithms. We present here a new tool that we developed to improve this strategy. Our new algorithm can find older and more degenerate TE sequences. Indeed, with this tool, we were able to detect up to 10% more of the *A. thaliana* genome as material derived from as yet unidentified TEs. By combining several strategies and tools, we were able to bring the proportion of the genome of this species known to be derived from TEs up to 50%. Interestingly, the new sequences detected were generally very short and located in the 500 bp immediately upstream from genes. Their epigenetic status and nucleotide composition attest to their origination from an ancient TE. Moreover, long-term conservation in orthologous positions and overlap with experimentally identified transcription factor binding sites (TFBS), suggest that they have been co-opted for new functional roles. Interestingly, these sequences were found to be overrepresented in the 5' sequences of flowering genes. A significant proportion of these TEs overlap with TFBSs able to bind transcription factors (TFs) known to be involved in flowering. Their overlaps with conserved non-coding sequences (CNS) suggest a long-term impact of TEs on flowering, since the initial global spread of flowering plants in the Cretaceous period.

Materials and Methods

Genome sequences

Genome sequences were obtained from the following sources: *A. thaliana* ecotype Col-0 (TAIR10 release) (<http://www.phytozome.com/arabidopsis.php>); *A. lyrata* (v1.0, <http://www.phytozome.com/alyrata.php>); *C. rubella* (initial release, <http://www.phytozome.com/capsella.php>); *A. alpina* (preliminary release, courtesy of Eva-Maria Willing, George Coupland, and Korbinian Schneeberger); *Schrenkiella parvulum* (formerly *Thellungiella parvula*; v2.0, <http://thellungiella.org/data/>); and *B. rapa* (v1.2, <http://www.phytozome.com/napacabbage.php>).

Genome annotation

TAIR10 gene and TE annotations were retrieved from the URGI website (https://urgi.versailles.inra.fr/gb2/gbrowse/tairv10_pub_TEs/).

The *Arabidopsis thaliana* “Brassicaceae” TE annotation was developed in a previous published study (Maumus & Quesneville, 2014). Briefly, in this previous work, for all the genomes from the five *Arabidopsis thaliana* ecotypes that have been assembled (Col-0, Ler-1, Kro-0, Bur-0 and C24), *Arabidopsis lyrata*, *Capsella rubella*, *Arabis alpina*, *Brassica rapa*, *Thellungiella salsuginea*, and *Schrenkiella parvula*, the TEdenovo pipeline from the REPET package (v2.0) (Quesneville et al., 2005; Flutre et al., 2011; Hoede et al., 2014) was used with default parameters and combining the similarity and structural branches. Consensus sequences derived from the structural branches, which use LTR Harvest, were retained only when they presented pfam domains typical of LTR retrotransposons. Consensus sequences were classified with REPET, by checking for characteristic structural features and similarities to known TEs from Repbase Update (17.01) (Bao et al., 2015), and by scanning against the Pfam library (26.0) (El-Gebali et al., 2018) with HMMER3 (Eddy, 2011). All the consensus repeat sequences generated were compiled into a “Brassicaceae” library, which we used to annotate the Col-0 genome with TEannot from the REPET package and default settings.

Brassicaceae TE copies

In order to annotate TE copies in each Brassicaceae genomes, we performed a new REPET analysis taking advantage of its improvements and of some new genome assembly. Hence, for all genomes from *Arabidopsis thaliana* Col-0 ecotype, *Arabidopsis lyrata*, *Capsella rubella*, *Arabis alpina*, *Brassica rapa* and *Schrenkiella parvula*, we used REPET package v2.5 with its two pipelines, TEdenovo and TEannot. We used the similarity branch of TEdenovo with default parameters on each genome, followed by TEannot with default parameters (sensitivity 2). From this first annotation, we selected consensus sequences containing at least one full-length copy (*i.e.* aligned over more than 95% of the length of the consensus sequence), on which we performed a second run of TEannot. This procedure has been shown to improve the quality of annotation (Jamilloux et al., 2017). Copies from the consensus annotated as 'PotentialHostGene' were removed.

Prediction accuracy

True positives (TP) are defined as predicted TE nucleotides that truly belong to a TE copy. False positives (FP) are the predicted TE nucleotides that do not really belong to a TE copy. True negatives (TN) are the nucleotides correctly predicted not to belong to a TE copy (correct rejection), and false negatives (FN) are the true TE copy nucleotides missed by the TE prediction process.

Sensitivity, the true positive rate, given by the formula $TP/(TP+FN)$, is obtained by calculating the fraction of nucleotides in the predicted TE overlapping with the TE reference annotation.

Specificity, also referred to as the true negative rate, is less straightforward to calculate. It can be calculated according to the formula $TN/(TN+FP)$, but TN and FP are difficult to determine for TEs, as they can only be known if we are sure that we have identified all the TE copies in the genome, which does not really seem possible. However, as a first approximation, we can consider that genes are not TEs, and are not derived from TEs, and use this information to obtain more accurate estimates for TN and FP. This is obviously an approximation as TEs are known to be sometimes part of genes. Hopefully this is rare compared to other regions of the genome, in particular if we consider the coding sequence (CDS) as it discards introns as well as 5' and 3' UTR where TEs can be found frequent. In this context, FP are predicted TE nucleotides that overlap a gene CDS annotation, and TN are CDS regions not predicted to be TEs.

Accuracy, given by $ACC=(TP+TN)/(TP+TN+FN+FP)$, is the rate of correct predictions.

Epigenetic data

We used a small-RNA map from Lister et al. (2008) (Lister et al., 2008) corresponding to dataset GSM277608 from the Gene Expression Omnibus (GEO) database (<http://www.ncbi.nlm.nih.gov/geo/>). The original coordinates were projected onto the TAIR10 assembly. The occurrences of multiply mapping reads were distributed evenly between genomic copies. This small-RNA dataset is derived from inflorescences of plants grown at 23°C with a 16 hours / 8hours dark cycle. Small RNAs from 15 to 35 nucleotides were extracted from bulk RNA extraction.

We used the 10 chromatin mark maps (H3K18ac, H3K27me1, H3K27me3, H3K36me2, H3K36me3, H3K4me2, H3K4me3, H3K9ac, H3K9me2 et H3) from Luo *et al.* (Luo et al., 2013).

Reads overlapping an annotation were counted with the CompareOverlapping.py script (option -O) of the S-Mart package (Zytnicki & Quesneville, 2011).

We normalized counts by calculating the ratio of the mean number of reads overlapping an annotation to the number of overlapping reads from the input.

Hierarchical clustering was performed on the epigenetic marks, based on the normalized ratio, with the *seaborn* python library, the correlation metric and a standard-scale normalization for each mark.

TFBS and CNS data

We use TFBS data compiled by Heyndrickx *et al.* (Heyndrickx et al., 2014) from ChIP-seq experiments, which we downloaded from http://bioinformatics.psb.ugent.be/cig_data/RegNet/.

CNSs data were retrieved from the work of Van de Velde *et al.* (de Velde et al., 2014) , from http://bioinformatics.psb.ugent.be/cig_data/Ath_CNS/Ath_CNS.php, and that of Haudry *et al.* (Haudry et al., 2013) from http://mustang.biol.mcgill.ca:8885/download/A.thaliana/AT_CNS.bed.

Analysis of binding motifs

We searched for binding motifs with the MEME suite server (Bailey et al., 2009) from <http://meme-suite.org>. We used MEME-CHIP (Machanick & Bailey, 2011) and JASPAR2018 CORE non-redundant databases.

Analysis of orthologous genes

We used OrthoMCL (Li et al., 2003) version 2.0 to identify genes orthologous between *A. thaliana*, *A. lyrata*, *C. rubella*, and *S. parvulum*. From the 21689 groups of orthologs obtained, we retained only 6921 for which four genes were identified, one from each species, to limit the detection of false-positive paralogs by this method.

Statistical analysis

We used the python libraries *pynum*, *scipy* for statistics, *matplotlib* for graphics and *panda* for data manipulation. *Jupyter notebooks* were used to monitor the analysis.

Sequence and coordinate manipulation

We obtained random sequences with *shuffle* from the SQUID 1.9g package (Eddy) and *revseq* from the Emboss 6.1.0 (Rice et al., 2000) package.

Genome coordinates were manipulated with the S-Mart package (Zytnicki & Quesneville, 2011). In particular, we used *modifyGenomicCoordinates* (version 1.0.1) and *CompareOverlapping* (version 1.0.4) to extend coordinates within the 5' region of genes, and find overlaps, respectively.

Results

Duster: a new approach for analyzing old degenerate transposable elements

Following their divergence from a common ancestor, repeat families have different destinies in different genomes. A repeat family may stop multiplying in one species, but may continue to multiply in another closely related species. The burst of transposition in an autonomous repeat family is a highly selective process: only the copies that have accumulated limited numbers of mutations remain functional and are able to transpose during the burst. This selective burst of transposition drives multiplication of the best conserved copies, *i.e.* those most similar to the ancestral sequence. Therefore, the TE families that remain active in some genomes should more similar to the ancestral sequence for longer than the decaying pool of related sequences in other genomes. Consequently, a repeat copy from one species may be considered to be relatively old if it closely resembles a sequence obtained from another species.

We previously showed (Maumus & Quesneville, 2014) that identifying TEs in a species by comparison with reference sequences found in the studied species but also in closely related species leads to the detection of older TE copies than searches exclusively with the reference sequence from the study species. Indeed, this approach detects old TE sequences that would not otherwise be recognized.

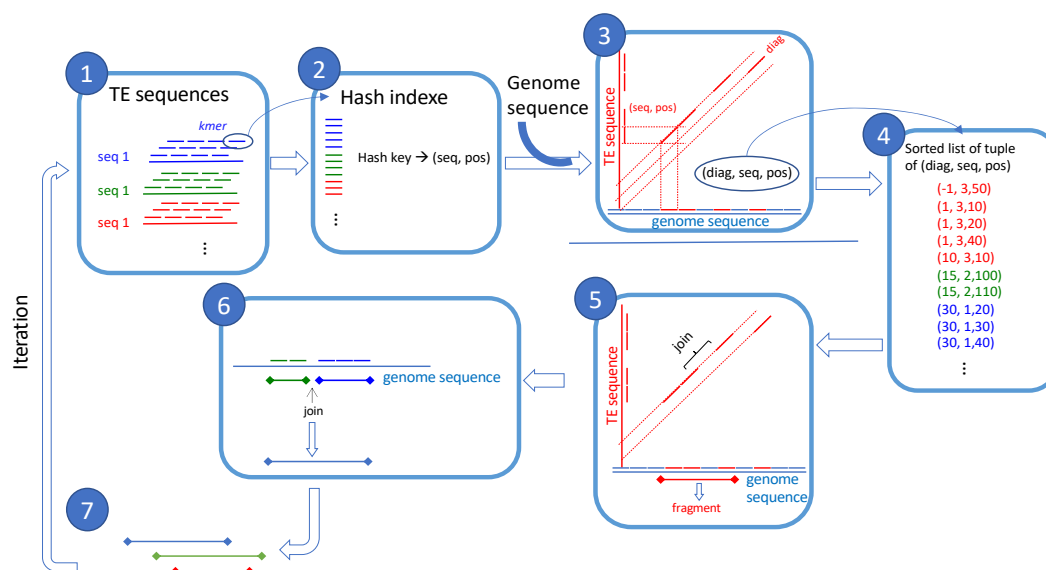


Figure 1: Duster algorithm overview. 1- TE sequences are cut in overlapping kmer of length w . 2- Build a hash table from kmer hash value to store their positions and sequence numbers. 3- Search kmer on genome sequence that match a TE sequence kmer, record diagonal (difference of start sequence position) with TE sequence number and position on TE sequence. 4- Sort the list of tuple by TE sequence numbers, diagonals, and positions. 5- Find consecutive kmer matches in the sorted list and join if they are on the same TE sequence, same diagonal, and with position distant by less than d kmer. Obtain fragments. 6- Join consecutive fragments on the genome sequence if distant by less than f nucleotides. 7- Take the fragment sequences as a new TE sequence set for another algorithm iteration.

We developed a program called Duster that compares a genome sequence, here considered as a query sequence, to a large number of TE sequences, *i.e.* a sequence library. Its algorithm used k -mers of size w (parameter $-w$) to search for similar sequences without the need to generate nucleotide alignments (figure 1). Hashed k -mer values can be used to speed up the search. Sensitivity is achieved by allowing one mismatch in k -mers every n consecutive nucleotides (parameter $-k$). Details of the algorithm are provided below, but it can be summarized as comparing k -mers between the genome shifting on the genome by few nucleotides (parameter $-S$), and each sequence from the library shifting by 1 nucleotide, and reporting matches when at least two k -mers are found on the same alignment diagonal (*i.e.* the differences between the coordinates in the query and the sequence library are identical) with a maximal distance of d k -mers (parameter $-d$). The region bounded by the two-extreme k -mer position are reported as matching. Two matching regions on the genome separated by less than x k -mers are merged (parameter $-f$). At the end of this first pass, the region identified on the genome can be used as a new sequence library for a new search (the $-n$ parameter). This procedure is repeated until genome coverage increases by less than 1% if $-n$ is set to 0.

The Duster algorithm

Definition and notations: We consider the problem of searching for occurrences of part of a query sequence Q within a library of subject sequences $L = \{S_1, S_2, \dots, S_d\}$. Each sequence in L is labeled with an index value s . We use the term k -mer to denote a contiguous sequence of DNA bases that is w bases long. The offset of a k -mer within a sequence S is the position of its first base with respect to the first base of S . We use the letter j to denote offsets and the notation S_j to denote the k -mer of S that has an offset j . The position within L of each occurrence of each k -mer may then be described by an (s, j) pair.

k -mer hash function: We store k -mer as a *hash-value* to speed up comparisons and to reduce memory requirement. The hash-value is obtained from the k -mer nucleotide sequence, by encoding each of the four possible nucleotides as two binary digits, as follows: $f(A \text{ or any value different from } C, G, T) = 00_2$, $f(C) =$

01₂, f(G) = 10₂, f(T) = 11₂. With this encoding system, any k -mer $K = b_1b_2 \dots b_k$ can be represented in a unique manner by a $2k$ bit integer

$$(1) \quad E(K) = \sum_{i=1}^k 4^{i-1} f(b_i)$$

Constructing the hash table: The first stage of the algorithm is the conversion of L into a hash table. The hash table is stored in memory as two data structures, a vector of positions P and a vector V of pointers into P . There are 4^k pointers in V , one for each of the 4^k possible k -mers. The pointer at position $E(K)$ of V points to the entry of P describing the position of the first occurrence of the k -mer in the library L . We can obtain the positions of all occurrences of K in L by moving P along the sequence from this point until it reaches the location indicated by the pointer at position $E(K)+1$ of V . The hash table is constructed by making two passes through L . On each pass, we consider only non-overlapping k -mers in S . With the first pass, we count all nonoverlapping occurrences in L of each possible k -mer. We then use these data to calculate the pointer positions for V and allocate the correct amount of memory for P . We ignore all k -mers with a frequency of occurrence exceeding a cutoff threshold. This has the advantage of both reducing the size of the hash table and effectively filtering out spurious matches from low-complexity DNA sequences (see below). We make a second pass through the data, using V to place the k -mer positions in P at the correct positions. During these two passes, any ambiguous or unrecognized characters, such as “N”, are translated into “A”. Stretches of unrecognized characters are thus translated into k -mers containing only “A”s. This is generally the commonest k -mer in the genome and will therefore be excluded from matching by the cutoff threshold.

Sequence search: We use the hash table to search for occurrences of a query sequence Q within L . We proceed along Q , from base 0 to base $n-w$, where n is the length of Q . At any base position t , we obtain the list of r positions of the occurrences of the k -mer Q_t . These positions are given by $E(Q_t)$ in V . We take this list of positions:

$$(2) \quad (s_1, j_1), (s_2, j_2), \dots, (s_r, j_r)$$

and use it to compute a list of hits:

$$(3) \quad H_1=(s_1, j_1-t, j_1), H_2=(s_2, j_2-t, j_2), \dots, H_r=(s_r, j_r-t, j_r)$$

We add these hits to a cumulative list M of hits that grows as t moves from 0 to $n-w$. Depending on the sensitivity required, we shift t by a few nucleotides. The elements of a hit are referred to, from left to right, as the *index*, *shift*, and *offset*. At the end of the sequence Q , M is sorted, first by *index* and then by *shift*. The last step is to scan M , looking for consecutive hits for which the *index* and *shift* are identical. Two hits are considered consecutive if they are separated by less than d k -mers. Note that such hits correspond to a succession of w bases consistent with an alignment between Q and a particular sequence S from L . If we sort these hits by *offset*, we obtain regions of exact matching between the two sequences. We retain only matching regions with at least two consecutive k -mers separated by fewer than d non-overlapping k -mers. We can create gapped alignments including mismatches, by joining exact matching regions if they are sufficiently close on the genome sequence (parameter $-f$). The method described above finds matches only on the direct DNA strand. For the identification of matches on the reverse strand, the search is simply repeated on the sequence complementary to Q .

Improving sensitivity: In the algorithm described above, hits are obtained only if two sequences contain identical k -mers, limiting the sensitivity of the search. Indeed, if two sequences have diverged considerably, it is likely to be difficult to find exactly matching k -mers. We therefore introduced some flexibility into the k -mer matching procedure, by modulating the hashing function. When calculating the k -mer hash value, we skipped one in every n nucleotides. Hence, two k -mers with different bases may have the same hash values if the difference concerns the skipped nucleotides. If they have the same hash value, they are considered to match according to the algorithm.

Filtering options: k -mers can be removed from the initial hash table if they fail to fulfill certain criteria. First, we remove all k -mers containing only ‘A’, as such k -mers are overrepresented in most genomes. Note that in masked sequences, ‘N’ and ‘X’ are also converted to ‘A’ in our hashing function, automatically

removing masked stretches of a sequence. For each *k-mer*, we also calculate the ratio of observed to expected occurrences, and remove all occurrences with a ratio below 1, which can be explained by chance alone. The expectation is calculated with a probabilistic background, modeled by a Markov chain in which order is a parameter (order 1 by default). Some *k-mers* can be removed, according to a threshold set by the user. Thus, *k-mers* for which the ratio of observed to expected occurrence is below a given threshold (1 in the example above) can be removed. This ratio shows to what extent the observed number of occurrences diverges from expectations. Overrepresented *k-mers* are removed as they are considered not to be specific enough. We remove those with occurrences above a given percentile (with 100% as the default setting). *K-mers* can also be removed on the basis of high entropy and low diversity, calculated as the number of different *k-mers* of the size of the background Markov chain order, divided by the maximum number of *k-mers* of this size.

Iterative search: The sequences of the library *L* may be considered as a sample of sequences to be searched for the sequence *Q*. Consequently, they may not contain all the *k-mers* required for a sensitive detection of *Q* in *L*. We thus implemented an iterative search to enrich the *k-mer* repertoire with the sequence *Q*. We reran the algorithm described above, replacing *L* with the regions identified in *Q*, and filtering out sequences below a threshold length. This procedure can be repeated until *Q* coverage changes by less than 1% between successive iterations.

Assessment of the performance of Duster

We assessed the performance of Duster, by calculating its prediction accuracy. This accuracy (ACC) was obtained by calculating the sensitivity (Sn) and specificity (Sp) of the predicted TE annotation by comparing the prediction with a reference annotation at nucleotide level (see Materials and Methods). We used the official annotation for *A. thaliana* from TAIR as the reference here. ACC takes both Sn and Sp into account, providing a convenient aggregate value. We therefore decided to maximize this value in our benchmark tests, for which the sequence library used was that of TE sequences from the TAIR annotation.

We chose the parameter set that gave the best results for Duster in our hands empirically, by optimizing annotation accuracy with TE copy sequences from other Brassicaceae species (data not shown). We used this parameter to compare Duster performances, benchmarking with tools implementing other algorithms that could be used for similar analyses. For this comparison, we chose BLAST (Altschul et al., 1997) and MegaBLAST (Morgulis et al., 2008), two widely used sequence comparison algorithms. Neither of these tools was designed to be run on a long genomic sequence. We therefore ran them through Blaster (Quesneville et al., 2003), which pre- and post-processes the input sequence and output results, respectively, to facilitate their use.

Table 1 shows the results obtained with Duster, BLAST and MegaBLAST. Duster was run twice, varying the distance required between two *k-mers* (*-d* parameter), using the values of 0 or 5, and the position shift on the genomic sequence (*-S* parameter), using values of 15 (size of the *k-mer*) or 7 (*k-mers* overlapping by 7 bp). The parameter (*-d* 0; *-S* 15) is, by definition, less sensitive than (*-d* 5; *-S* 7). Note that parameter *-f* is the maximum sequence distance between two consecutive fragments on the genome sequence to be joined, and *-n* is the number of iterations. Here, the two sets of parameters are set to consider only sequences of more than 100 bp in length, with only one iteration. We also chose a *k-mer* length of 15 (parameter *-w*) and a potential nucleotide mismatch every four nucleotides (parameter *-k*).

We show that Duster outperforms the other tools in terms of speed, taking only five to seven minutes, versus 38 minutes at best with MegaBLAST run in parallel on four threads. Sensitivity was higher for Duster and BLAST, at 0.99. Duster had a lower specificity, but a higher coverage, suggesting that our tool detects many more previously unidentified potential TEs. As a different way of assessing the false-positive rate, we ran Duster on a shuffled genome sequence respecting dinucleotide composition, and a reversed but noncomplementary sequence. Coverage remained below 0.001 for the shuffled sequence, and below 0.01 for the reversed sequence.

Based on the way we compute false positives, this suggests that many genes have regions derived from old TEs not detected with other tools. As this is one purpose of our new tool, we consider this to be a positive result, particularly given that it is not biologically inconsistent.

Table 1: Comparison of tool performances with TE sequences from the official TAIR A. thaliana TE annotation. *computed on a Linux workstation with Intel Xeon® CPU E3-1270 v3 @ 3.50 GHz x 8 and 15.5 Go of RAM.

Tool	Parameters	Cov.	Sn	Sp	Time*
Duster	-w15; -k4; -d0; -f100; -S15; -n1	0.27 (shuff.: 0% rev.:0.1%)	0.99	0.91	5.0 m
Duster	-w15; -k4; -d5; -f100; -S7; -n1	0.31 (shuff.: 0.1% rev.:1%)	0.99	0.87	7.0 m
Blaster/MegaBLAST	-S2 ; -L200	0.20	0.96	0.99	1.18 h
Blaster/MegaBLAST-4Threads	-S2 ; -L200	0.20	0.96	0.99	38 m
Blaster/BLAST	-S2	0.23	0.99	0.98	17 h
Blaster/BLAST-2Threads	-S2 ; -L200	0.22	0.97	0.98	8.8 h
Blaster/BLAST-4Threads	-S2 ; -L200	0.22	0.97	0.98	6.15 h

Transposable elements account for up to 50% of the A. thaliana genome

Assuming that Duster would be able to detect interesting new TE sequences in the *A. thaliana* genome, we ran an analysis with all the Brassicaceae TE copies from *Arabidopsis thaliana*, *Arabidopsis lyrata*, *Capsella rubella*, *Schrenkiella parvulum*, *Arabis alpina*, and *Brassica rapa* that we had previously annotated with the REPET package (see Material and Methods). We used the parameter setting with $-d$ 5 and $-S$ 7, but changed $-n$ to 0, allowing iteration of the algorithm until it reached a genome coverage difference between two successive iterations of less than 1%.

The TAIR10, Brassicaceae and Duster TE annotations together accounted for 49.75% of the genome sequence. This figure is 29.72% higher than that for the TAIR10 TE reference annotation (20.03%), and 10.60% higher than that for the Brassicaceae TE annotation (39.15%).

Structural properties of Duster-specific copies

We characterized the new set of repeats identified by Duster, by using the annotations to extract copies that did not overlap with any Gene, TAIR10 TE, Brassicaceae, or *A. thaliana* REPET annotations (see Materials and Methods). We identified 19608 TE copies that were Duster-specific. We did the same for the TAIR10 and Brassicaceae annotations, thereby obtaining 177 TAIR10-specific and 5139 Brassicaceae-specific copies, by removing any copies with no overlap to another annotation.

We characterized these copies by comparing their length, chromosome distribution, and position relative to genes (figure 2). Duster-specific copies appeared to be significantly shorter than Brassicaceae-specific, TAIR10-specific, and TAIR10 copies (Figure 2A, chi-squared p -value respectively 3.09×10^{-192} , 2.70×10^{-8} , $<10^{-293}$). Figure 2B shows the distance to the closest 5' or 3' TE copy for each annotated gene. Duster-specific copies are more abundant close to genes than other copies (all chi-squared p -value $<10^{-293}$, versus Brassicaceae-specific, TAIR10-specific, and TAIR10 copies). Similarly, Brassicaceae-specific copies were more abundant than TAIR10-specific copies. They were more frequently found upstream from genes (Figure 1B, chi-squared p -value $<10^{-293}$), as were Brassicaceae-specific, and TAIR10 TE copies (all chi-squared p -value $<10^{-293}$). Figure 2C shows the distribution of TE copies over the chromosomes. It shows that Duster TE copies, and, to a lesser extent, Brassicaceae TE copies, follow the chromosomal distribution of genes (see the right panel of figure 2C), whereas TAIR10 TEs follow the opposite pattern. Duster and Brassicaceae TEs have a different chromosomal distribution from the annotated TEs from TAIR10.

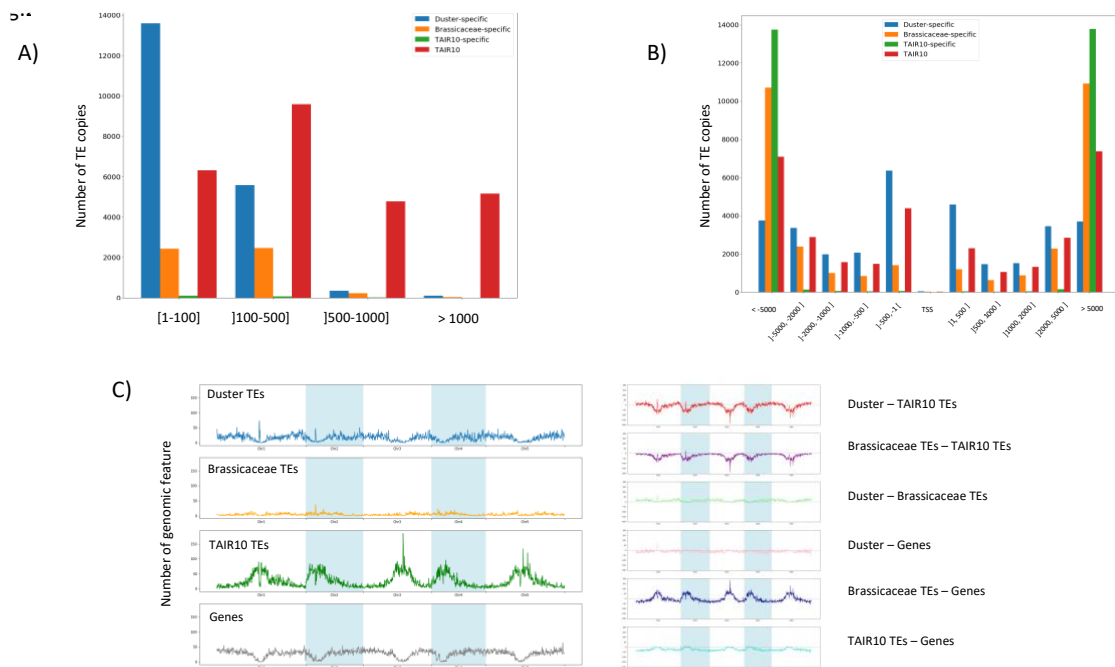


Figure 2: Structural characteristics of Duster-specific, Brassicaceae-specific, TAIR10-specific, and TAIR10TE copies. (A) TE length distribution, (B) distance to the closest 5', or 3' TE copy for each annotated gene, (C) TE copy distribution on the chromosomes. Left panel: density plot, for 100 kb windows overlapping by 10 kb. Right panel: density differences, in 100 kb bins.

Finally, we investigated the nucleotide composition of the sequences, including dinucleotides. The counts are presented as a radar plot in Figure 3A. The profile is similar for all TE copies (Duster-specific, Brassicaceae-specific, TAIR10-specific, and TAIR10 TEs). Interestingly, TAIR10-specific copies had the strongest bias towards TT, AA, AT, and TA dinucleotides, followed by Duster-specific copies. These biases, also shared by other TE copies but to a lesser extent, were thought to be a consequence of the process by which methylated cytosine is deaminated. The greater “A-T” richness of TAIR10-specific and Duster-specific copies compared to TAIR10 TE copies may indicate that they have undergone a mutation over a longer period and are therefore more ancient.

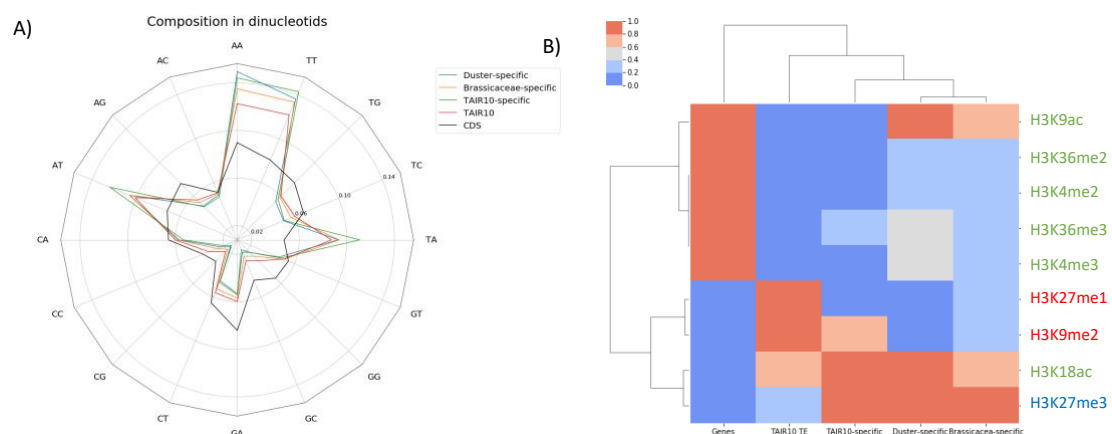


Figure 3: Composition of Duster-specific, Brassicaceae-specific, and TAIR10-specific copies. (A) Radar plot of the dinucleotide composition of the sequences. (B) Hierarchical clustering of TEs and genes with respect to heterochromatin marks (H3K27me1 and H3K9me2) and euchromatin marks (H3K18ac, H3K27me3, H3K36me2, H3K36me3, H3K4me2, H3K4me3, and H3K27me3).

H3K9ac). Histone marks colors denote the the type of chromatine. Green marks characterize transcriptionally active genes. Blue marks pinpoint genes transcriptionally repressed by transcriptional inhibitors. Red marks identify closed chromatin, also called heterochromatin.

Epigenetic profiles

We investigated the epigenetic status of the identified TE copies, considering small RNAs, and chromatin marks. The small RNAs which sizes range from 15 to 35 nucleotids were taken from Lister et al. (Lister et al., 2008), for which mapped data were available. There were 4.17%, 20.14%, 16.95%, and 60.44% of matching TE copies from the Duster-specific, Brassicaceae-specific, TAIR10-specific, and TAIR10 TE datasets, respectively, in the intersection between this dataset and our annotations, indicating a low targeting by small RNA of the method-specific annotation compared to TAIR10 annotated TEs.

We analyzed nine epigenetic marks from Luo et al. (Luo et al., 2013), also available as mapped data. The hierarchical clustering algorithm identified distinctly different profiles for genes and for TAIR10 TEs (figure 3B). TAIR10 TEs were enriched in heterochromatin marks H3K27me1 and H3K9me2, and genes with euchromatin marks H3K36me2, H3K36me3, H3K4me2, H3K4me3, and H3K9ac as expected.

The clustering algorithm associated Duster-specific, Brassicaceae-specific, and TAIR10-specific copies with the TAIR10 TE profile, indicating that their profiles were more similar to a typical TE profile than to a gene profile. However, Brassicaceae-specific, TAIR10-specific, and Duster-specific marks copies that had very similar profiles which differ from TEs. Their copies appeared to have very few heterochromatic, however the euchromatin marks H3K27me3 and H3K18Ac are predominant for method-specific TEs. Interestingly, H3K27me3 is known to be a repressive mark preferentially associated with genes expressed at low levels or in a tissue-specific manner (Turck et al., 2007; Zhang et al., 2007; Oh et al., 2008; Jacob et al., 2010) and H3K18ac is usually associated with promoters (Weiner et al., 2015).

TE conservation in flowering plants

We investigated the conservation of TE copies by searching for overlaps with known conserved non-coding sequences (CNSs) identified in previous studies. We compared the TE copies with CNSs identified in crucifers (Haudry et al., 2013) and rosids (de Velde et al., 2014). For both datasets, a substantial proportion of the TE copies overlapped with these CNSs (5.32 to 19.8%, see table 2). Some TE sequences overlapped with CNSs conserved in 12 rosid species: *Arabidopsis thaliana*, *Carica papaya*, *Glycine max*, *Malus domestica*, *Populustrichocarpa*, *Fragaria vesca*, *Medicago truncatula*, *Lotus japonicus*, *Theobroma cacao*, *Ricinus communis*, *Manihot esculenta*, and *Vitis vinifera*. Fossil rosids dating back to the Cretaceous period, estimated by a molecular clock between 125 and 99.6 million years ago (Davies et al., 2004; Magallón & Castillo, 2009; Hohmann et al., 2015), have been found. Our findings therefore reveal a remarkable conservation of 1521 and 1213 TE insertions identified by the Duster and Brassicaceae methods, respectively, over more than 100 million years, twice as many as can be detected with the traditional annotation approach available for the TAIR10 TE annotation. We also show here that the Duster approach can detect more TEs overlapping with CNSs than the Brassicaceae method.

The 65 rosid CNSs from the 12 rosid species associated with Duster-specific copies in the 500 bp upstream of genes included 58 Duster-specific sequences. A MEME-ChIP (Bailey et al., 2009) analysis of these sequences identified a significant 15-nucleotide TTTT(TTTT(G/T)TTT(G/T)(G/C) motif (E-value 3.4×10^{-6}) in 27 sites. This motif matched MA1281.1 (AT5G02460), MA1274.1 (OBP3), MA1278.1 (OBP1), MA1268.1 (AT1G69570), MA1267.1 (AT5G66940), MA1272.1 (AT2G28810), MA1371.1 (IDD4), MA1279.1 (COG1), MA1156.1 (JKD), MA1374.1 (IDD7), MA1160.1 (AT1G14580), MA1158.1 (MGP), MA1157.1 (NUC), MA1125.1 (ZNF384), MA1159.1 (SGR5), MA1277.1 (Adof1), and MA1370.1 (IDD5) with all q-values $< 5 \times 10^{-2}$, described in the JASPAR database [36] as C2H2 zinc finger factors of the Dof-type. Interestingly, some of these motifs were found to be related to: (i) AT3G55370 (OBP3), which is known to encode a nuclear DOF domain-containing TF expressed primarily in roots that is responsive to salicylic acid in leaves and petals, (ii) AT1G69570 (CDF5), which is involved in flower development and photo-periodism, (iii) AT1G29160 (COG1), which acts as a negative regulator in phytochrome-mediated light responses, (iv) AT3G50410 (OBP1), which acts as a positive regulation of transcription and play an important roles in plant growth and development and (v) AT5G02460.1, which is probably involved in early processes in vascular development.

Table 2: Number of CNSs and percentage overlap with TE annotations

	Crucifer CNSs	Rosid CNSs	Rosid CNSs – score 11
All Duster	14774 14.1%	20758 19.8%	1521 1.45%
All Brassicaceae	7798 9.39%	13702 16.5%	1213 1.46%
All TAIR10 TEs	1583 5.32%	4934 16.6%	527 1.77%
Duster-specific	3659 18.7%	2580 13.2%	194 0.99%
Brassicaceae-specific	1848 36.0%	2535 49.3%	246 4.79%
TAIR10-specific	27 15.2%	27 15.2%	3 1.69%
5' Duster-specific	745 15.5%	968 20.2%	65 1.36%
5' Brassicaceae-specific	270 30.6%	452 51.3%	26 2.95%
5' TAIR10-specific	2 6.67%	4 13.3%	0 0.0%

We then looked in detail at the conservation of Duster-specific, Brassicaceae-specific, and TAIR10-specific copies in the Brassicaceae. We considered only regions close to orthologous genes found with OrthoMCL (see Materials and Methods). We focused on *A. thaliana*, *A. lyrata*, *C. rubella*, and *S. parvulum*, as these species have divergence times of 5 to 40 My. Overall, we found 6265 groups of four orthologous genes (one gene for each species) containing a TE copy annotated (available in one of the TAIR10 TEs, Brassicaceae, or Duster datasets) within the gene extended by 500 bp upstream (29% of *A. thaliana* genes). Regions encompassing both the orthologous gene and its upstream sequence were pairwise aligned with the cognate *A. thaliana* region. Within this 6265 regions, we found 2457, 353, and 11 TE copies from respectively Duster-specific, Brassicaceae-specific, and TAIR10-specific datasets. We considered a TE copy to be present if more than 50% of the *A. thaliana* annotated TE copy nucleotides were identical in the pairwise alignment. The TEs were oldest in the Duster-specific set, followed by the Brassicaceae-specific set, as shown by the height of the “111” bars of the histogram, which corresponds to the presence of a TE at orthologous positions in all four species (Figure 4). Interestingly, the “000” bars was also quite high. This bar corresponds to TEs found only in *A. thaliana*, but which belonged to method-specific sets and therefore escaped the reference *A. thaliana* TE annotation or TE detection by the simple REPET *de novo* procedure limited to *A. thaliana*. They were therefore detectable only with TEs found in other species. These copies may result from horizontal transfer from these other species, or may simply have been identified in other genomes because they are better conserved in those genomes or have enough copies to built reliable consensus. Consequently, we probably identified here TEs that had a poor success of invasion following the horizontal transfer as there are too few copies in each family, to characterise them properly. This result illustrates the utility of our cross-species TE annotation approach and the greater efficiency of Duster than of the REPET annotation procedure.




Figure 4: Numbers of sequences conserved in orthologous positions between species. The three-digit code indicates the species in which the sequences were present. A “1” indicates presence and a “0” absence, in *A. lyrata* (position 1), *C. rubella* (position 2), and *S. parvulum* (position 3).

Contribution of TEs to the architecture of gene regulatory networks

We investigated the functional role of these TE sequences, which may have been co-opted for some regulatory purpose. We chose to study two gene regulatory networks (GRNs) in which TEs might be suspected to play a role. The genes controlling flower development in *Arabidopsis thaliana* are good candidates, as some alleles have been reported to be controlled by a TE sequence in *A. thaliana*: the FLOWERING WAGENINGEN (FWA) locus (Kinoshita et al., 2007), and FLOWERING LOCUS C (FLC) (Liu et al., 2004; Strange et al., 2011).

We considered the genes reported by Chen et al. 2018 (Chen et al., 2018) in their paper describing the architecture of GRNs controlling flower development in *Arabidopsis thaliana*. We searched the 500 bp immediately upstream from these genes for the presence of Duster-specific, Brassicaceae-specific, and TAIR10-specific TEs.

An enrichment in Duster-specific regions was observed in the 5' sequences of flowering genes: 33.1% of these sequences contained such regions, versus only 17.5% for all genes (chi-squared p -value= 2.4×10^{-7} , table 3). Brassicaceae-specific regions and specific regions identified by TAIR10 displayed no particular enrichment, with 4.46% versus 3.21% (chi-squared p -value=0.37), and 0% vs 0.11% (chi-squared p -value=0.68), respectively of 5' sequences containing such regions.

Table 3: Method-specific TEs overlap with the 500 bp immediately upstream from genes

	Duster-specific	Brassicaceae-specific	TAIR10-specific
All genes	17.5%	3.21%	0.11%
All flowering genes	33.1%	4.46%	0.0%
Stress genes	16.5%	4.71%	0.0%

We further explored the overrepresentation of Duster-specific and Brassicaceae-specific TEs in GRNs, by focusing on stress GRNs genes, which are also thought to be linked to TEs, as reported by several studies suggesting that transposition events may be triggered during plant stress responses including salt (Naito et al., 2009), wounding (Mhiri et al., 1997), bacteria (Grandbastien et al., 2005), and viruses (Buchmann et al., 2009). We focused on the genes expressed in various stress conditions described by Barah et al. (Barah et

al., 2016). We searched in the 500 bp immediately upstream from these genes for Duster-specific, Brassicaceae-specific, and TAIR10-specific TEs (table 3). We found no enrichment of these upstream regions in Duster-specific TE copies (chi-squared p -value=0.73) or TAIR10-specific TE copies (chi-squared p -value=0.67). However, we found an enrichment for Brassicaceae-specific TEs (chi-squared p -value=5.1 x 10⁻⁷).

TEs and transcription factor binding sites

The conservation and overrepresentation described above suggest a probably functional role for these TEs. We then investigated their ability to regulate gene expression, by assessing their ability to bind TFs. We investigated the co-occurrence of the transcription factor binding sites (TFBSs) identified with 27 TFs in ChIP-seq experiments by Heyndrick et al. (Heyndrickx et al., 2014) and the various TE annotations studied here.

We found that TFBSs were more frequently present in Duster regions (29.0%) than in Brassicaceae and TAIR10 regions (19.5%, and 14.6%, respectively, chi-squared p -values all < 10⁻³⁰⁰; Table 4). This pattern was even more marked for the analysis of method-specific regions: 53.9%, 29.7% and 24.9% of these regions, respectively, overlapped with TFBSs (chi-squared p -values all < 10⁻³⁰⁰). This trend was even stronger for analyses limited to the 500 bp immediately upstream from genes (49.3% for Duster-specific and 38.2% for Brassicaceae-specific TEs, chi-squared p -value < 10⁻³⁰⁰, TAIR10-specific TEs being untestable due to the low counts). Interestingly, these regions contained 567 Duster-specific and 48 Brassicaceae-specific TFBS regions, associated with more than seven TFs, and referred to hereafter as *hot TFBSs*. The identification of these regions suggests that there may be a hub of target genes involved in the important function of crosstalk between different processes (Barabási & Oltvai, 2004).

Table 4: Method-specific TEs overlap with TFBS identified by ChIPseq. Occurrences are given between parenthesis.

	Hot TFBSs	TFBSs
All Duster	1.10% (1155)	29.0% (30400)
All Brassicaceae	0.51% (426)	19.5% (16206)
All TAIR10 TEs	0.40% (120)	14.6% (4331)
Duster-specific	2.89% (567)	53.9% (10575)
Brassicaceae-specific	0.93% (48)	29.7% (1528)
TAIR10-specific	1.13% (2)	24.9% (44)
5' Duster-specific	2.38% (114)	49.3% (2365)
5' Brassicaceae-specific	0.79% (7)	38.2% (337)
5' TAIR10-specific	3.33% (1)	6.67% (2)

An analysis of the 5' sequences of flowering genes identified 6 key TFs known to be involved in flowering that could bind to both Duster-specific and Brassicaceae-specific regions (AGL-15, AP1, PI, AP3, SEP3, SOC1). Few sites for TFs involved in circadian rhythm and light response (PIF3, PRR5, PRR7) and development (GL3) were identified. Most were found in Duster-specific regions, with very few in Brassicaceae-specific regions, and none in TAIR10-specific regions. Interestingly, some Duster-specific regions were associated with several TFBSs.

We found 1757 and 1009 Duster-specific regions overlapping with crucifer and rosid CNS, respectively, and a TFBS. We found that 84 of these regions were highly conserved, as they overlapped with CNSs present in the 12 rosid species used for their identification, suggesting a presence in the common ancestors of the rosids more than 100 My ago. We also found that 9 of these highly conserved Duster-specific regions overlapped with a *hot TFBS*, suggesting the presence of a highly conserved hub of target genes involved in crosstalk between different processes. The top five highly conserved TFBSs from Duster-specific regions were AGL-15, AP1, SEP3, PRR5, and PIF4 (31, 23, 20, 14, 14 occurrences, respectively), all but one of which are directly involved in flowering process, the exception being PRR5, which is more closely related to circadian rhythms and light responses.

The CNSs associated with Duster-specific copies in the 500 bp upstream of the gene present in the 12 rosid species included 58 Duster-specific sequences: 42 target genes of floral regulators according to Chen

et al. (Chen et al., 2018), 16 for which Duster-specific regions were colocalized with a highly conserved CNS and a TFBS (Table 5).

Table 5: Sixteen floral regulators genes from Chen et al. (Chen et al., 2018) with a Duster-specific region in the 500 bp immediately upstream, colocalizing with highly conserved CNSs and TFBSs

Gene	TFBS	Profil from MEME-suite
AT1G20930	AGL-15	SOC1
AT1G49760	AGL-15	
AT1G75110	GL1	
AT2G05920	AGL-15; PIF4	SOC1
AT2G30400	AGL-15	
AT2G30970; AT2G30960	AGL-15; AP1; SEP3	
AT2G33750	SEP3	SOC1
AT2G41370	AGL-15; AP1; AP3; PI	
AT3G02040	AP3; AGL-15; PI	SOC1
AT3G14172	AGL-15	
AT3G19170	AGL-15; PIF4; PIF3; PI; AP3	SOC1
AT4G00810	PIF	SOC1
AT4G37260	SEP3	SOC1
AT4G39980	SEP3, AP2, AGL-15	
AT5G03680	AP1; SEP3; FLM; PIF4	
AT5G13990	PIF4	

Using the *MEME Suite*, we identified 7 of these genes as containing a sequence corresponding to a binding motif of the SOC1 TF, a MADS box factor active in flowering time control which may integrate signals from the photoperiod, vernalization and autonomous floral induction pathways. Thirteen of downstream genes were controlled by the floral regulation motifs of one or several type II TF-MADSs (AG, AP1, AP2, AP3, BLR, ETT, FLM, JAG, LFY, PI, RGA, SEP3, SVP).

Discussion

The need for a new dedicated repeat detection algorithm

RepeatMasker (Smit et al.), Censor (Jurka et al., 1996; Kohany et al., 2006), and Blaster [19] are the tools most frequently used to annotate TE sequences in genomes. All these tools combine BLAST (or BLAST-like with seed and extend algorithms) calls with pre- and post-processing for the analysis of genomic sequences. They are all, therefore, subject to the intrinsic limitations of these algorithms, including a reliance on seeds to find alignments. These seeds in BLAST are *k*-mers with a default size of 11 nucleotides. BLAST requires two *k*-mers on the same diagonal (i.e. alignment without gaps) to proceed with the analysis of an alignment to assess its relevance. An alignment score threshold determined with a probabilistic model is used for this assessment of relevance. These features may account for the lower sensitivity for this method than for Duster.

First, two *k*-mers are required to initiate an alignment. With the default BLAST parameters, this requires an exact match of at least 22 nucleotides between two sequences. This requirement can be decreased, as seed length is a parameter of BLAST that can be set, and it is decreased to 14 nucleotides for some implementations (seed size of 7 with WU-BLAST), but it still needs an exact match. For Duster, we allow mismatches in the *k*-mers, and the two *k*-mers may overlap. With the settings used for this analysis, we required a match of 21 nucleotides, but with some mismatches allowed.

Second, in the statistical test based on an alignment score threshold, even if the required exact match of 22 nucleotides is found, a gap-free alignment is produced for testing with the probabilistic model. The result depends on sequence length and on a model that is mathematically sophisticated, but too simple biologically, in that it considers successive nucleotides to be independent and equally probable. We now know that neither of these assumptions hold true for real sequences. Consequently, the model is of debatable value and may reject some alignments differently according to sequence lengths. In Duster, we retain all regions that match two *k*-mers, and the empirically chosen parameters yielded very few false positives (0.001).

We can see here that BLAST is not the most appropriate algorithm for finding small degenerate TEs. It was developed for a different purpose: identification of the best match within databases to a sequence

used as a query. Its use to identify TEs constitutes a major deviation from its initial purpose, for which it performs well.

Duster was designed for the express purpose of finding old and degenerate TE copies. In addition to having a different *k*-mer strategy, it is essentially an alignment-free algorithm. BLAST attempts many alignments before reporting a match. Duster does not really require an alignment, just boundary coordinates, accounting for the greater speed of this algorithm. Boundaries may be considered imprecise as they are based on *k*-mers and their precision is therefore limited by *k*-mer size and the shift in the coordinates of the *k*-mer on the genomic sequence. With the parameters used here, the precision is about seven nucleotides. We think that this is sufficient for the identification of regions, and it may not be appropriate to aim for greater precision in the identification of very old, degenerate TE copies.

Improving sensitivity is always at the cost of specificity which is particularly difficult to evaluate in the context of an uncomplete TEs identification. Many TEs remain to be discovered even in a well studied genome such as *A. thaliana*. False positive rate is consequently difficult to measure. We proposed here a proxy for this assessment using CDS (see Materials and Methods) allowing to compare Duster with BLAST on the same basis. The good sensitivity performance of Duster is partly due to the chosen balance with specificity. However, sensitivity has also a cost in term of speed, and we think that Duster algorithm improve mainly on this feature and open the route for future new more efficient algorithms. Improving the sensitivity for BLAST algorithm to reach this of Duster, would have made this work more difficult or even impossible for computing reasons.

The work presented here highlights the utility of specifically developed tools for addressing certain difficult biological questions. It highlights the need for a new generation of sequence-finding tools, tailored to the particular biological question posed and perhaps replacing BLAST with more adapted algorithms.

Long-term impact of TE copies

TEs are important sources of variation on which selection can operate in the evolution of species. Many examples of TEs generating new phenotypes have been reported in plants (Lisch, 2013). This phenomenon is known as “domestication” when the TE sequences are retrained in new genes, or “co-option” when TE insertions affect existing genes. TE sequences that become functional in the host are conserved by selection, which can be recognized over long periods of time. Other TE copies devoid of function in the host are progressively removed from the genome through the accumulation of point mutations and deletions. However, gene regulatory regions accumulate point mutations and deletions at a slower rate than other non-coding regions, because of their function. Consequently, TE insertions in these regions, even if neutral, may be difficult to remove once established.

TFs control the transpositional activity of TEs by binding to them, but they have also been shown to bind TEs in regions not supposed to be transcribed. Are they remnant of old TFBS or illegitimate because of a particular base composition of the regions? Whatever is the answer, the corresponding TFBS are, therefore, widespread throughout the genome. In some cases, TEs from the same family may be inserted close to several genes. This may lead to nearby genes being regulated by the same TF, potentially leading to their evolution into gene regulation networks. Genome-wide assessment revealed that hundreds of TEs have been co-opted into the regulatory regions of mammalian genes (Bejerano et al., 2006; Lowe et al., 2007). TEs have also been involved in both the creation of new regulatory networks (Cordaux et al., 2006; Wang et al., 2007) and the rewiring of preexisting ones (Ackerman et al., 2002). Such networks are observed, for example, for the DAYSLEEPER gene in *A. thaliana* (Bundock & Hooykaas, 2005). This gene has features in common with hAT transposases, suggesting that it may have been domesticated as a new TF. In order to operate, transposases are able to recognize TE DNA motifs close to the sequence boundaries in order to clive DNA and initiate the transposition process. This suggests that this domesticated transposase may have conserved this property, as TF also bind specifically DNA motifs. This binding may be functional as it leads to the domestication of the transposase. Overall, this suggests that TE sequences, dispersed throughout the genome, are targeted by DAYSLEEPER, and may regulate host genes. Another interesting example is provided by the retrotransposon ONSSEN in *Arabidopsis* (Ito et al., 2011). Thieme *et al.* (Thieme et al., 2017) showed that, following the heat stress-dependent mobilization of ONSSEN, the progenies of treated plants contain up to 75 new ONSSEN insertions. Progenies with additional ONSSEN copies display broad environment-dependent phenotypic diversity. This finding suggests that some of the new TE insertions affect the expression of genes in a temperature-dependent manner. It also suggests that TE sequences may

have contributed to individual local adaptation through the mutations they induce during bursts of transposition. Some of these bursts of transposition may result from activation by environmental stresses, promoting environment-sensitive phenotypes.

However, little is known about the long-term impact of TEs. Our CNS data suggest that some of the TEs identified were inserted in their current positions more than 100 My ago, during the Cretaceous period. The most important evolutionary event during this period was, perhaps, the spread of flowering plants (Angiosperms) to colonize the entire planet. Flowering plants were particularly successful at colonizing new areas and replacing the older established flora, with which they competed. TEs undoubtedly played an important role in this process. Some of the TE insertions we detected may, indeed, have played this role. Duster-specific copies appear to be old, degenerate, short, and surprisingly close to genes, lying in the 5' flanking sequences known to correspond to gene regulatory regions. Their maintenance specifically in these zones suggests that they supply the host with a function, probably in the regulation of the neighboring gene.

The Duster-specific TEs identified here may have played an important role in building new pathways allowing flowering plants to adapt to their environment. Indeed, we found that Duster-specific TE copies were overrepresented in the 5' regions of genes of the GRN for flowering (table 3). A significant proportion of these copies overlapped with TFBSs known to bind TFs involved in the control of flowering. Moreover, the histone H3K27me3 mark was identified predominantly in method-specific TEs (see figure 3). This histone mark has been reported to be associated with genes expressed at low levels of in a tissue-specific manner (Jacob et al., 2010), such as those involved in flower development.

Our results suggest a possible link between the success of flowering plants during the Cretaceous period and the co-option of TEs in the flowering GRN. However, further analyses are required to demonstrate a causal role. This study is a first step in this direction, identifying previously unknown candidate TEs.

Flowering has been studied in considerable detail, generating a wealth of data. The data used here are, therefore, clearly biased towards flowering. However, other impacts on other GRNs may subsequently be discovered with our ancient TE annotation, as and when new data become available.

Interestingly, our results suggest that identifying very old TE copies could facilitate the identification of TE-based regulatory modules selected a long time ago. They support the detection of TFBS in ChipSeq experiment, but also suggest a TE-based origin for many TFBS.

Conclusions

In this study, we investigated the contribution of TEs to the bulk genome of *Arabidopsis* over a timescale that remains inaccessible to other approaches, through the use of a new tool that we developed, called Duster. Duster uses a new efficient algorithm, which identified an additional 10% of nucleotides as belonging to TEs. We have, thus, dug deeper into the dark matter than previous studies, leading to the recognition of old, degenerate TE sequences undetectable with other methodologies.

This study delivers a key result, improving our understanding of plant evolution and plant adaptation, by providing clues for identifying ancient TE remnants in gene regulatory regions underlying potential regulation modules. Some of the TE copies identified here may have been selected a long time ago, to drive adaption to changing environments.

Acknowledgment

We thank KlaasVandepoele for reading a first version of this manuscript and drawing our attention to interesting TFBS and CNS data that enabled us to improve the manuscript. We also thank Michaël Alaux and Johann Confais for their comments on the manuscript. This work was performed with the facilities of the Plant Bioinformatics Facilities (<https://doi.org/10.15454/1.5572414581735654E12>). A previous version of this article has been peer-reviewed and recommended by Peer Community In Genomics (<https://doi.org/10.24072/pci.genomics.100004>).

Author contributions

HQ: conceived the study, developed Duster, ran the software. AB: identified method-specific sets and described their general features. MW: performed conservation analysis on the Brassicaceae. DA, DN, HQ: performed CNS, TFBS and profile analyses. DA and HQ: drafted the manuscript. All the authors read and commented on the manuscript.

Materials & Correspondence

All requests for materials and information should be addressed to Hadi Quesneville at hadi.quesneville@inrae.fr.

Code availability

The Duster code is available on github, and is distributed as part of the TEfinder package at <https://archive.softwareheritage.org/swh:1:dir:e00658f617274bc59c721de13863ca6c4ce47d8d>.

Data availability

The Brassicaceae TE copies from *Arabidopsis thaliana* Col-0 ecotype, *Arabidopsis lyrata*, *Capsella rubella*, *Arabis alpina*, *Brassica rapa* and *Schrenkiella parvula*, used for Duster annotation, and TAIR10 TE, Duster, and Brassicaceae annotations can be downloaded from <https://doi.org/10.15454/VWTJYS>. The data are also available in a JBrowse genome browser at https://urgi.versailles.inra.fr/jbrowse/gmod_jbrowse/?data=myData%2FAtha.

Conflict of interest disclosure

The authors of this article declare that they have no financial conflict of interest with the content of this article. Hadi Quesneville is one of the PCI Genomics recommenders.

References

- Ackerman H, Udalova I, Hull J, Kwiatkowski D (2002) Evolution of a polymorphic regulatory element in interferon-gamma through transposition and mutation. *Molecular Biology and Evolution*, 19, 884–890. <https://doi.org/10.1093/oxfordjournals.molbev.a004145>
- Altschul SF, Madden TL, Schäffer AA, Zhang J, Zhang Z, Miller W, Lipman DJ (1997) Gapped BLAST and PSI-BLAST: a new generation of protein database search programs. *Nucleic Acids Research*, 25, 3389–3402. <https://dx.doi.org/10.1093%2Fnar%2F25.17.3389>
- Bailey TL, Boden M, Buske FA, Frith M, Grant CE, Clementi L, Ren J, Li WW, Noble WS (2009) MEME Suite: tools for motif discovery and searching. *Nucleic Acids Research*, 37, W202–W208. <https://doi.org/10.1093/nar/gkp335>
- Bao W, Kojima KK, Kohany O (2015) Repbase Update, a database of repetitive elements in eukaryotic genomes. *Mobile DNA*, 6, 11. <https://doi.org/10.1186/s13100-015-0041-9>
- Barabási A-L, Oltvai ZN (2004) Network biology: understanding the cell's functional organization. *Nature Reviews. Genetics*, 5, 101–113. <https://doi.org/10.1038/nrg1272>
- Barah P, B N MN, Jayavelu ND, Sowdhamini R, Shameer K, Bones AM (2016) Transcriptional regulatory networks in *Arabidopsis thaliana* during single and combined stresses. *Nucleic Acids Research*, 44, 3147–3164. <https://doi.org/10.1093/nar/gkv1463>
- Bejerano G, Lowe CB, Ahituv N, King B, Siepel A, Salama SR, Rubin EM, Kent WJ, Haussler D (2006) A distal enhancer and an ultraconserved exon are derived from a novel retroposon. *Nature*, 441, 87–90. <https://doi.org/10.1038/nature04696>

- Buchmann RC, Asad S, Wolf JN, Mohannath G, Bisaro DM (2009) Geminivirus AL2 and L2 proteins suppress transcriptional gene silencing and cause genome-wide reductions in cytosine methylation. *Journal of Virology*, 83, 5005–5013. <https://doi.org/10.1128/JVI.01771-08>
- Buisine N, Quesneville H, Colot V (2008) Improved detection and annotation of transposable elements in sequenced genomes using multiple reference sequence sets. *Genomics*, 91, 467–475. <https://doi.org/10.1016/j.ygeno.2008.01.005>
- Bundock P, Hooykaas P (2005) An Arabidopsis hAT-like transposase is essential for plant development. *Nature*, 436, 282–284. <https://doi.org/10.1038/nature03667>
- Chen D, Yan W, Fu L-Y, Kaufmann K (2018) Architecture of gene regulatory networks controlling flower development in Arabidopsis thaliana. *Nature Communications*, 9, 4534. <https://doi.org/10.1038/s41467-018-06772-3>
- Cordaux R, Udit S, Batzer MA, Feschotte C (2006) Birth of a chimeric primate gene by capture of the transposase gene from a mobile element. *Proceedings of the National Academy of Sciences of the United States of America*, 103, 8101–8106. <https://doi.org/10.1073/pnas.0601161103>
- Davies TJ, Barraclough TG, Chase MW, Soltis PS, Soltis DE, Savolainen V (2004) Darwin's abominable mystery: Insights from a supertree of the angiosperms. *Proceedings of the National Academy of Sciences*, 101, 1904–1909. <https://doi.org/10.1073/pnas.0308127100>
- Eddy SR (2011) Accelerated Profile HMM Searches. *PLOS Computational Biology*, 7, e1002195. <https://doi.org/10.1371/journal.pcbi.1002195>
- Eddy SR SQUID - library of functions for biological sequence analysis Copyright (C) 1992-2002 Washington University School of Medicine. <http://eddylib.org/software.html>
- El-Gebali S, Mistry J, Bateman A, Eddy SR, Luciani A, Potter SC, Qureshi M, Richardson LJ, Salazar GA, Smart A, Sonnhammer ELL, Hirsh L, Paladin L, Piovesan D, Tosatto SCE, Finn RD (2018) The Pfam protein families database in 2019. *Nucleic Acids Research*, 47, D427–D432. <https://doi.org/10.1093/nar/gky995>
- Feschotte C, Jiang N, Wessler SR (2002) Plant transposable elements: where genetics meets genomics. *Nature Reviews Genetics*, 3, 329–341. <https://doi.org/10.1038/nrg793>
- Flutre T, Duprat E, Feuillet C, Quesneville H (2011) Considering transposable element diversification in de novo annotation approaches. *PloS One*, 6, e16526. <https://doi.org/10.1371/journal.pone.0016526>
- Grandbastien M-A, Audeon C, Bonnivard E, Casacuberta JM, Chalhoub B, Costa A-PP, Le QH, Melayah D, Petit M, Poncet C, Tam SM, Van Sluys M-A, Mhiri C (2005) Stress activation and genomic impact of Tnt1 retrotransposons in Solanaceae. *Cytogenetic and Genome Research*, 110, 229–241. <https://doi.org/10.1159/000084957>
- Haudry A, Platts AE, Vello E, Hoen DR, Leclercq M, Williamson RJ, Forczek E, Joly-Lopez Z, Steffen JG, Hazzouri KM, Dewar K, Stinchcombe JR, Schoen DJ, Wang X, Schmutz J, Town CD, Edger PP, Pires JC, Schumaker KS, Jarvis DE, Mandáková T, Lysak MA, van den Bergh E, Schranz ME, Harrison PM, Moses AM, Bureau TE, Wright SI, Blanchette M (2013) An atlas of over 90,000 conserved noncoding sequences provides insight into crucifer regulatory regions. *Nature Genetics*, 45, 891–898. <https://doi.org/10.1038/ng.2684>
- Heyndrickx KS, de Velde JV, Wang C, Weigel D, Vandepoele K (2014) A Functional and Evolutionary Perspective on Transcription Factor Binding in Arabidopsis thaliana. *The Plant Cell*, 26, 3894–3910. <https://doi.org/10.1105/tpc.114.130591>
- Hoede C, Arnoux S, Moisset M, Chaumier T, Inizan O, Jamilloux V, Quesneville H (2014) PASTEC: an automatic transposable element classification tool. *PloS One*, 9, e91929. <https://doi.org/10.1371/journal.pone.0091929>
- Hohmann N, Wolf EM, Lysak MA, Koch MA (2015) A Time-Calibrated Road Map of Brassicaceae Species Radiation and Evolutionary History. *The Plant Cell*, tpc.15.00482. <https://doi.org/10.1105/tpc.15.00482>
- Huang P, Studer AJ, Schnable JC, Kellogg EA, Brutnell TP (2017) Cross species selection scans identify components of C4 photosynthesis in the grasses. *Journal of Experimental Botany*, 68, 127–135. <https://doi.org/10.1093/jxb/erw256>
- Huang C-H, Sun R, Hu Y, Zeng L, Zhang N, Cai L, Zhang Q, Koch MA, Al-Shehbaz I, Edger PP, Pires JC, Tan D-Y, Zhong Y, Ma H (2016) Resolution of Brassicaceae Phylogeny Using Nuclear Genes Uncovers Nested Radiations and Supports Convergent Morphological Evolution. *Molecular Biology and Evolution*, 33, 394–412. <https://doi.org/10.1093/molbev/msv226>

- Ito H, Gaubert H, Bucher E, Mirouze M, Vaillant I, Paszkowski J (2011) An siRNA pathway prevents transgenerational retrotransposition in plants subjected to stress. *Nature*, 472, 115–119. <https://doi.org/10.1038/nature09861>
- Jacob Y, Stroud H, Leblanc C, Feng S, Zhuo L, Caro E, Hassel C, Gutierrez C, Michaels SD, Jacobsen SE (2010) Regulation of heterochromatic DNA replication by histone H3 lysine 27 methyltransferases. *Nature*, 466, 987–991. <https://doi.org/10.1038/nature09290>
- Jamilloux V, Daron J, Choulet F, Quesneville H (2017) De Novo Annotation of Transposable Elements: Tackling the Fat Genome Issue. *Proceedings of the IEEE*, 105, 474–481. <https://doi.org/10.1109/JPROC.2016.2590833>
- Jurka J, Klonowski P, Dagman V, Pelton P (1996) CENSOR—a program for identification and elimination of repetitive elements from DNA sequences. *Computers & Chemistry*, 20, 119–121. [https://doi.org/10.1016/S0097-8485\(96\)80013-1](https://doi.org/10.1016/S0097-8485(96)80013-1)
- Kidwell MG (2002) Transposable elements and the evolution of genome size in eukaryotes. *Genetica*, 115, 49–63. <https://doi.org/10.1023/A:1016072014259>
- Kinoshita Y, Saze H, Kinoshita T, Miura A, Soppe WJJ, Koornneef M, Kakutani T (2007) Control of FWA gene silencing in *Arabidopsis thaliana* by SINE-related direct repeats. *The Plant Journal: For Cell and Molecular Biology*, 49, 38–45. <https://doi.org/10.1111/j.1365-313X.2006.02936.x>
- Kohany O, Gentles AJ, Hankus L, Jurka J (2006) Annotation, submission and screening of repetitive elements in Repbase: RepbaseSubmitter and Censor. *BMC bioinformatics*, 7, 474. <https://doi.org/10.1186/1471-2105-7-474>
- Li L, Stoeckert CJ, Roos DS (2003) OrthoMCL: identification of ortholog groups for eukaryotic genomes. *Genome Research*, 13, 2178–2189. <https://doi.org/10.1101/gr.1224503>
- Lisch D (2013) How important are transposons for plant evolution? *Nature Reviews. Genetics*, 14, 49–61. <https://doi.org/10.1038/nrg3374>
- Lister R, O'Malley RC, Tonti-Filippini J, Gregory BD, Berry CC, Millar AH, Ecker JR (2008) Highly integrated single-base resolution maps of the epigenome in *Arabidopsis*. *Cell*, 133, 523–536. <https://doi.org/10.1016/j.cell.2008.03.029>
- Liu J, He Y, Amasino R, Chen X (2004) siRNAs targeting an intronic transposon in the regulation of natural flowering behavior in *Arabidopsis*. *Genes & Development*, 18, 2873–2878. <https://doi.org/10.1101/gad.1217304>
- Lowe CB, Bejerano G, Haussler D (2007) Thousands of human mobile element fragments undergo strong purifying selection near developmental genes. *Proceedings of the National Academy of Sciences of the United States of America*, 104, 8005–8010. <https://doi.org/10.1073/pnas.0611223104>
- Luo C, Sidote DJ, Zhang Y, Kerstetter RA, Michael TP, Lam E (2013) Integrative analysis of chromatin states in *Arabidopsis* identified potential regulatory mechanisms for natural antisense transcript production. *The Plant Journal: For Cell and Molecular Biology*, 73, 77–90. <https://doi.org/10.1111/tpj.12017>
- Machanic P, Bailey TL (2011) MEME-ChIP: motif analysis of large DNA datasets. *Bioinformatics*, 27, 1696–1697. <https://doi.org/10.1093/bioinformatics/btr189>
- Magallón S, Castillo A (2009) Angiosperm diversification through time. *American Journal of Botany*, 96, 349–365. <https://doi.org/10.3732/aib.0800060>
- Maumus F, Quesneville H (2014) Ancestral repeats have shaped epigenome and genome composition for millions of years in *Arabidopsis thaliana*. *Nature Communications*, 5, 4104. <https://doi.org/10.1038/ncomms5104>
- Mhiri C, Morel J-B, Vernhettes S, Casacuberta JM, Lucas H, Grandbastien M-A (1997) The promoter of the tobacco Tnt1 retrotransposon is induced by wounding and by abiotic stress. *Plant Molecular Biology*, 33, 257–266. <https://doi.org/10.1023/A:1005727132202>
- Morgulis A, Coulouris G, Raytselis Y, Madden TL, Agarwala R, Schäffer AA (2008) Database indexing for production MegaBLAST searches. *Bioinformatics (Oxford, England)*, 24, 1757–1764. <https://doi.org/10.1093/bioinformatics/btn322>
- Naito K, Zhang F, Tsukiyama T, Saito H, Hancock CN, Richardson AO, Okumoto Y, Tanisaka T, Wessler SR (2009) Unexpected consequences of a sudden and massive transposon amplification on rice gene expression. *Nature*, 461, 1130–1134. <https://doi.org/10.1038/nature08479>
- Oh S, Park S, Nocker S van (2008) Genic and Global Functions for Paf1C in Chromatin Modification and Gene Expression in *Arabidopsis*. *PLOS Genetics*, 4, e1000077. <https://doi.org/10.1371/journal.pgen.1000077>

- Quesneville H, Bergman CM, Andrieu O, Autard D, Nouaud D, Ashburner M, Anxolabéhère D (2005) Combined evidence annotation of transposable elements in genome sequences. *PLoS computational biology*, 1, 166–175. <https://doi.org/10.1371/journal.pcbi.0010022>
- Quesneville H, Nouaud D, Anxolabéhère D (2003) Detection of new transposable element families in *Drosophila melanogaster* and *Anopheles gambiae* genomes. *Journal of Molecular Evolution*, 57, S50–59. <https://doi.org/10.1007/s00239-003-0007-2>
- Rice P, Longden I, Bleasby A (2000) EMBOS: the European Molecular Biology Open Software Suite. *Trends in genetics*, 16, 276–277. [https://doi.org/10.1016/s0168-9525\(00\)02024-2](https://doi.org/10.1016/s0168-9525(00)02024-2)
- Schnable PS, Ware D, Fulton RS, Stein JC, Wei F, Pasternak S, Liang C, Zhang J, Fulton L, Graves TA, Minx P, Reily AD, Courtney L, Kruchowski SS, Tomlinson C, Strong C, Delehaunty K, Fronick C, Courtney B, Rock SM, Belter E, Du F, Kim K, Abbott RM, Cotton M, Levy A, Marchetto P, Ochoa K, Jackson SM, Gillam B, Chen W, Yan L, Higginbotham J, Cardenas M, Waligorski J, Applebaum E, Phelps L, Falcone J, Kanchi K, Thane T, Scimone A, Thane N, Henke J, Wang T, Ruppert J, Shah N, Rotter K, Hodges J, Ingenthron E, Cordes M, Kohlberg S, Sgro J, Delgado B, Mead K, Chinwalla A, Leonard S, Crouse K, Collura K, Kudrna D, Currie J, He R, Angelova A, Rajasekar S, Mueller T, Lomeli R, Scara G, Ko A, Delaney K, Wissotski M, Lopez G, Campos D, Braidotti M, Ashley E, Golser W, Kim H, Lee S, Lin J, Dujmic Z, Kim W, Talag J, Zuccolo A, Fan C, Sebastian A, Kramer M, Spiegel L, Nascimento L, Zutavern T, Miller B, Ambroise C, Muller S, Spooner W, Narechania A, Ren L, Wei S, Kumari S, Faga B, Levy MJ, McMahan L, Van Buren P, Vaughn MW, Ying K, Yeh C-T, Emrich SJ, Jia Y, Kalyanaraman A, Hsia A-P, Barbazuk WB, Baucom RS, Brutnell TP, Carpita NC, Chaparro C, Chia J-M, Deragon J-M, Estill JC, Fu Y, Jeddelloh JA, Han Y, Lee H, Li P, Lisch DR, Liu S, Liu Z, Nagel DH, McCann MC, SanMiguel P, Myers AM, Nettleton D, Nguyen J, Penning BW, Ponnala L, Schneider KL, Schwartz DC, Sharma A, Soderlund C, Springer NM, Sun Q, Wang H, Waterman M, Westerman R, Wolfgruber TK, Yang L, Yu Y, Zhang L, Zhou S, Zhu Q, Bennetzen JL, Dawe RK, Jiang J, Jiang N, Presting GG, Wessler SR, Aluru S, Martienssen RA, Clifton SW, McCombie WR, Wing RA, Wilson RK (2009) The B73 maize genome: complexity, diversity, and dynamics. *Science (New York, N.Y.)*, 326, 1112–1115. <https://doi.org/10.1126/science.1178534>
- Smit A, Hubley R, Green P RepeatMasker Open-4.0. 2013–2015. <http://www.repeatmasker.org>
- Strange A, Li P, Lister C, Anderson J, Warthmann N, Shindo C, Irwin J, Nordborg M, Dean C (2011) Major-Effect Alleles at Relatively Few Loci Underlie Distinct Vernalization and Flowering Variation in *Arabidopsis* Accessions. *PLOS ONE*, 6, e19949. <https://doi.org/10.1371/journal.pone.0019949>
- Thieme M, Lanciano S, Balzergue S, Daccord N, Mirouze M, Bucher E (2017) Inhibition of RNA polymerase II allows controlled mobilisation of retrotransposons for plant breeding. *Genome Biology*, 18, 134. <https://doi.org/10.1186/s13059-017-1265-4>
- Turck F, Roudier F, Farrona S, Martin-Magniette M-L, Guillaume E, Buisine N, Gagnot S, Martienssen RA, Coupland G, Colot V (2007) *Arabidopsis* TFL2/LHP1 Specifically Associates with Genes Marked by Trimethylation of Histone H3 Lysine 27. *PLOS Genetics*, 3, e86. <https://doi.org/10.1371/journal.pgen.0030086>
- de Velde JV, Heyndrickx KS, Vandepoele K (2014) Inference of Transcriptional Networks in *Arabidopsis* through Conserved Noncoding Sequence Analysis. *The Plant Cell*, 26, 2729–2745. <https://doi.org/10.1105/tpc.114.127001>
- Wang J, Keightley PD, Halligan DL (2007) Effect of divergence time and recombination rate on molecular evolution of *Drosophila* INE-1 transposable elements and other candidates for neutrally evolving sites. *Journal of Molecular Evolution*, 65, 627–639. <https://doi.org/10.1007/s00239-007-9028-6>
- Weiner A, Hsieh T-HS, Appleboim A, Chen HV, Rahat A, Amit I, Rando OJ, Friedman N (2015) High-Resolution Chromatin Dynamics during a Yeast Stress Response. *Molecular Cell*, 58, 371–386. <https://doi.org/10.1016/j.molcel.2015.02.002>
- Wicker T, Gundlach H, Spannagl M, Uauy C, Borrill P, Ramírez-González RH, De Oliveira R, Mayer KFX, Paux E, Choulet F, International Wheat Genome Sequencing Consortium (2018) Impact of transposable elements on genome structure and evolution in bread wheat. *Genome Biology*, 19, 103. <https://doi.org/10.1186/s13059-018-1479-0>
- Zhang X, Clarenz O, Cokus S, Bernatavichute YV, Pellegrini M, Goodrich J, Jacobsen SE (2007) Whole-Genome Analysis of Histone H3 Lysine 27 Trimethylation in *Arabidopsis*. *PLoS Biology*, 5, e129. <https://doi.org/10.1371/journal.pbio.0050129>

Zytnicki M, Quesneville H (2011) S-MART, a software toolbox to aid RNA-Seq data analysis. PLoS One, 6, e25988. <https://doi.org/10.1371/journal.pone.0025988>